Properties of Tabulated Bidirectional Reflectance Distribution Functions

by

Joel M. DeYoung

B.S., Harding University, 1994

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE STUDIES

(Department of Computer Science)

We accept this thesis as conforming
to the required standard

The University of British Columbia

August 1996

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 acquired that is represented by a table of numbers. The generality of these measured BRDFs is useful for generating realistic images, but the inevitable inaccuracy associated with the data gathering process can lead to a BRDF that is more general than it needs to be, or that lacks certain physical properties.

This thesis proposes measures for several properties of 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.

Contents

List of Tables

List of Figures

Acknowledgements

First of all, I wish to thank Alain Fournier, my research supervisor, for agreeing to supervise me even while on sabbatical. Alain suggested the topic of this thesis and has provided invaluable help with the results. I have learned much from Alain's way of looking at the world, as he is never afraid to hold conventional wisdom up to the scrutiny of common sense.

Paul Lalonde provided a great deal of help during all stages of this work. His patient answers to my many questions have been a great benefit. Paul also served as the student reader for this thesis. John Amanatides of York University was the second reader for this thesis. My thanks go to John for his useful feedback and comments. Kevin Coughlan and Marcelo Walter also deserve my thanks for reading over an almost finished draft. Many other members of the Imager Lab were very helpful when I presented various ideas and problems to them.

Bob Lewis assisted me with issues relating to physical plausibility, and was also extremely handy with several tools used to prepare the thesis document. Jean-Luc Duprat always seemed to have the answers to my many LaTeXquestions. Jim Varah provided help with details of the singular value

decomposition algorithm. Greg Ward of Lawrence Berkeley Labs helped with feedback on both gonioreflectometers and BRDF interpolation.

Thanks are due to my office mates: Sijian Zhang, Joseph Wu, Jean-Luc Duprat, Ronald Beirouti, and Brian Fuller, for indulging me with many interesting conversations on topics that were rarely related to computer graphics. Such sessions provided needed diversions from school work.

I promised the producers of the CBC Stereo program, Realtime, that I would mention their show in my thesis if they would send me a t-shirt. You can hear it Saturday evenings, coast to coast. Check your local listings. They should feel good for helping to clothe a grad student.

I wish to thank my parents for endowing me with a desire to learn, and providing much needed encouragement during every stage of my education.

Finally, thanks to my wife and best friend Michelle, for not only working several jobs to put me through grad school, but also for her enduring support of everything I do. These last two years would scarcely have been possible without her help—and I know they would not have been nearly as fun.

JOEL M. DEYOUNG

The University of British Columbia

August 1996

Chapter 1

Introduction

The computation of reflectance is one of the most fundamental operations in realistic image synthesis. Realistically characterizing how light reflects off surfaces is also one of the most challenging aspects of creating photorealistic computer graphics images.

One way to represent the reflectance of a surface is with a *bidirectional reflectance distribution function*, or BRDF. This thesis is written in the context of this approach. Section 1.1 gives a more precise definition of the BRDF, and Section 1.2 provides a motivation for the results presented here.

1.1 Bidirectional Reflectance Distribution Functions

Consider the geometry in Figure 1.1. 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

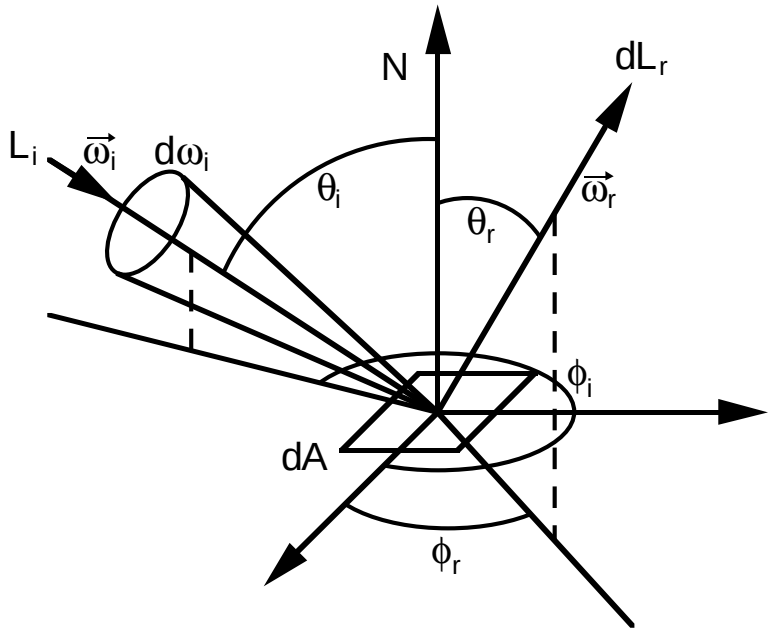

Figure 1.1: Shading geometry

direction $\vec{\omega_r}$. The amount of reflected radiance L_r is proportional to the incident irradiance E:

$$dL_r \propto dE \tag{1.1}$$

The differential incident irradiance dE can be rewritten in terms of the radiance from $\vec{\omega_i}$ and the solid angle $d\omega_i$:

$$dE = L_i \cos\theta_i d\omega_i \tag{1.2}$$

The constant of proportionality in Equation 1.1 is called the bidirectional reflectance distribution function, or BRDF. Using the substitution from Equation 1.2, the BRDF f_r can be expressed as:

$$f_r(\vec{\omega_i} \rightarrow \vec{\omega_r}) = \frac{dL_r}{L_i \cos\theta_i d\omega_i} \tag{1.3}$$

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 as $f_r(\theta_i, \phi_i, \theta_r, \phi_r)$ to express this dependency explicitly. Note that in this form, f_r almost fully characterizes the anisotropic reflection of a surface, with several simplifying assumptions.

The first simplification made in Equation 1.3 is to express the BRDF as independent of wavelength. This is physically incorrect, and means that reflection from surfaces that cause the light to shift in frequency, such as thin film interference, is not included in this model. Furthermore, this definition of the BRDF ignores phenomena caused by subsurface scattering, as well as changes in the phase and polarization of the light. It is also assumed that the time between the arrival of the incident light and the emittance of the reflected light is negligible, even though this is not the case in fluorescent materials.

It also should be noted that the BRDF can only characterize opaque surfaces since its domain is the hemisphere above the surface normal. The bidirectional transmittance distribution function, or BTDF, describes how light travels through a surface. All of the results persented here that operate on BRDFs can be analogously applied to BTDFs.

Integrating ω_i over the hemisphere Ω_N surrounding dA results in the *fundamental reflectance equation*, since it expresses reflected radiance in terms of incident radiance:

$$L_r = \int_{\Omega_N} f_r(\vec{\omega_i} \rightarrow \vec{\omega_r}) L_i \cos\theta_i d\omega_i \tag{1.4}$$

1.2 Properties of BRDFs

The simplest and fastest way to obtain a BRDF that can be used in realistic image synthesis is to derive a closed-form illumination model. Several such models have been widely used for a number of years. However, they often make non-trivial assumptions about the surface reflectance, making them incapable of representing some surfaces.

As an alternative to using a closed-form illumination model, BRDFs can be measured from physical samples. Since they are obtained experimentally from real surfaces, these measured BRDFs 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, it would be useful to examine properties of a BRDF that indicated how much error was introduced by the measurement process, so the BRDF data could 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.

This thesis proposes techniques for measuring four such properties. Two of them, reciprocity and energy conservation, are properties that any realistic BRDF should have—since they are defined by physical principles. Measuring

these represents an attempt to find and eliminate errors introduced into the data. The other two properties, isotropy and separability, are measures of the compressibility of the data, since they are formulated to represent how much error would be introduced by changing the data to possess these properties.

The techniques for measuring these properties, as well as the techniques for transforming the BRDF data so that it possesses them, are tools that are useful as a post-processing step after the data is obtained experimentally. These techniques represent the major contribution of this thesis.

1.3 Outline of the Rest of the Thesis

Chapter 2 summarizes research related to the work in this thesis, including a historical overview of closed-form illumination models, and techniques used to obtain BRDFs experimentally. Chapter 3 surveys several ways that BRDF data has been represented and used for rendering tasks, and suggests several new ideas in this area. Chapters 4 and 5 make up the core results of the thesis. Chapter 4 defines the measuring schemes for the four BRDF properties: reciprocity, energy conservation, isotropy, and separability, while Chapter 5 explains the techniques for transforming BRDF data to possess one or more of these properties. Chapter 6 summarizes the contributions of the thesis and proposes several areas worthy of further research. Appendix A contains an overview of the principles of radiometry that underlie the work in this thesis, including a more thorough explanation of the derivation in Section 1.1 above. Appendix B contains a glossary of notation.

Chapter 2

Related Work

Approaches to the realistic characterization of reflection fall into two categories: modeling, that is deriving a mathematical representation for light reflection; and measuring, which involves using data obtained through empirical measurement from real world surfaces. This chapter gives a historical overview of these two types of techniques.

2.1 Analytical Reflection Models

Early attempts at approximating realistic BRDFs provided closed-form solutions of the reflectance function. For this reason, these approaches are referred to as *analytical* reflection models. One of the earliest models which is still widely used today, was proposed by Bui-Tuong Phong [4][1]. The Phong model

[1] This Vietnamese researcher's name has caused some degree of confusion in the computer graphics community. Bui-Tuong was his hyphenated family-generation name, while Phong was his given name. The fact that his family name was written first on his 1975 illumination model paper caused some to get his family name and given name mixed up.

incorporates three parameters representing three types of reflected light: ambient, diffuse, and specular. His model was a great improvement over previous shading algorithms. However, it was based on *ad hoc* observation of reflectance rather than physical principles.

Blinn extended this model to be more physically accurate [3]. He based his model on work in physics by Torrance and Sparrow [34, 35] that was concerned with reflection from roughened surfaces. Blinn's work represented the first in a long line of physically based approaches to illumination.

Cook and Torrance's physically based illumination model [7] was also based on the work of Torrance and Sparrow, but made use of physics results from Beckmann and Spizzichino [2] as well. This was the first model to include wavelength dependence, which was motivated by the difference in color of the specular highlights on metal surfaces.

All of the above models are isotropic, meaning that the reflection is constant as the surface is rotated about its normal. Very few surfaces in the real world actually exhibit such ideal behavior. Kajiya first proposed an illumination model that allowed for anisotropic reflection [21]. Poulin and Fournier presented an anisotropic reflection model based on a microgeometry of parallel cylinders [30].

Several other extensions have been made to these models, most notably the addition of polarized light, which is included in models by Wolff and Kurlander [40] and He et al. [17, 18]. The latter is currently the most complex analytical model in computer graphics.

2.2 Tabulated BRDFs

Despite the complexity of the most recent analytical models, there are still many situations where an analytical approach cannot accurately model the reflection from a real world surface. One example is reflection from hair. In such cases, it is desirable to compute reflection from empirically measured data. This section presents a survey of techniques used to obtain realistic reflectance data through such measurements.

2.2.1 Gonioreflectometers

A device used for measuring BRDFs is called a *gonioreflectometer*. A simplified version of a gonioreflectometer is shown in Figure 2.1. The sample to be measured is placed in the center of the device. A light source and detector are moved about the hemisphere above the sample, and measurements of the reflectance are taken every few degrees. The output of the device is a table of numbers. For this reason, the data is referred to as a *tabulated* BRDF.

There are several problems with gonioreflectometer devices that restrict their practical use in computer graphics. First of all, they are expensive. Of the relatively few labs in North America that have gonioreflectometers, some will accept a surface and make several measurements for several hundred dollars. A full BRDF measurement can easily cost thousands of dollars. Secondly, the number of moving parts in the device can cause the data to be quite noisy, requiring the use of perfectly dark and clean environments. Lastly, the devices

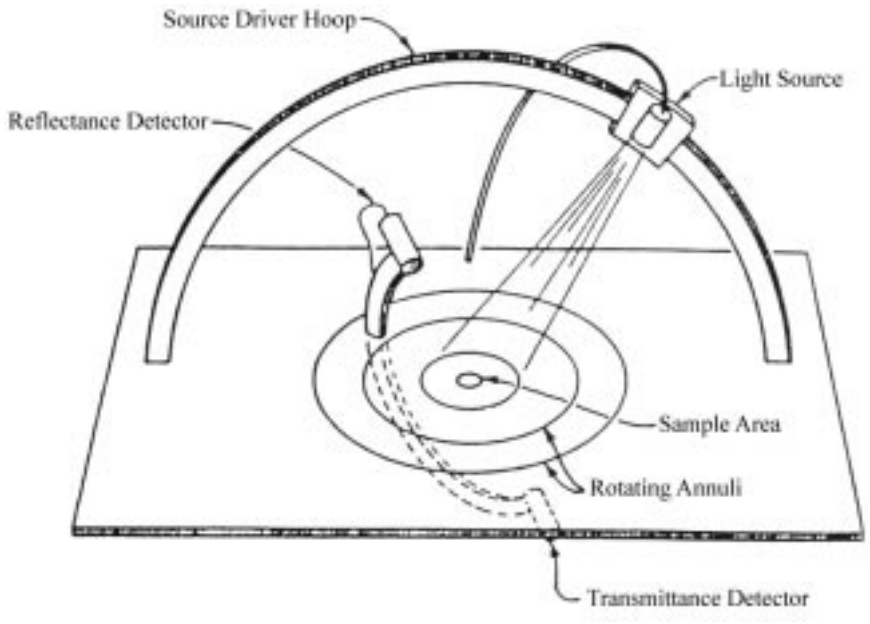

Figure 2.1: A simplified version of a gonioreflectometer (copied from [37])

are usually quite slow, requiring several hours to sample an entire BRDF at even a low resolution.

2.2.2 Ward's Gonioreflectometer

In response to the restrictions imposed by gonioreflectometers, Greg Ward developed a novel implementation of a gonioreflectometer that is relatively fast and inexpensive [37].

Figure 2.2 shows a diagram of what Ward termed an *imaging* gonioreflectometer. The sample to be measured is placed at the edge of a half-silvered hemisphere, facing inward. A collimated light source shines from the outside of the hemisphere onto the sample. A CCD camera with a fisheye lens collects

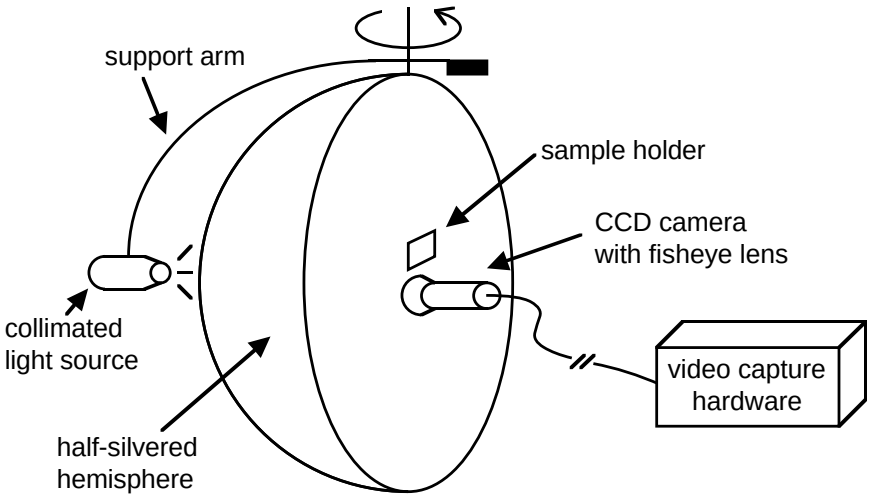

Figure 2.2: The imaging gonioreflectometer constructed by Greg Ward (adapted from [37])

the light reflected from the sample onto the inside of the hemisphere. The direction of incident reflection is changed by moving the light source about the equator of the hemisphere and by rotating the surface sample about its normal.

This device is useful because only two degrees of freedom rely on moving parts. One image collected by the CCD camera represents all the measurements for a given incident direction. This characteristic greatly reduces the amount of time required to collect data for an entire BRDF.

For each pixel collected by the CCD camera, transformations must be performed to account for the imperfections in the hemisphere and the fisheye lens. As a result, the reflected angles in each sample are nonuniform. This is the biggest obstacle to using data from this device directly. Ward fit his data

to the parameters of an anisotropic Gaussian illumination model and then shaded with this model.

2.2.3 Virtual Gonioreflectometers

As an alternative to obtaining BRDF data directly, several researchers have simulated gonioreflectometers in software. This approach is useful since it is inexpensive, and both fast and repeatable.

Cabral et al. [5] proposed a method for computing reflectance functions from bump maps. This work was expanded and presented in detail by Westin et al. [39]. Their approach was to model a small patch of surface with microscopic detail, and then compute the BRDF by casting rays into this microgeometry with a ray tracer. Using this technique the BRDF can be measured to any resolution. The resulting tabulated BRDF data is then used directly for rendering.

The use of virtual gonioreflectometers to generate tabulated BRDFs has been extended by several researchers [15, 16] to include other effects such as wavelength dependence and subsurface scattering.

This approach has several disadvantages. First, it simply reduces the problem to modeling realistic microgeometries. For extremely complex surfaces, this can prove to be quite difficult. Second, most ray tracers ignore interreflection between surface elements of the microgeometry. Third, reflections computed in the virtual gonioreflectometer are dependent on the reflection model used in the ray tracer. Most ray tracers use relatively simple models

of reflection such as Phong shading. Shortcomings of the underlying reflection model may show up in the resulting BRDF.

A virtual gonioreflectometer was used to generate the examples BRDFs used in this thesis. Using *Optik*, a ray tracer developed at the University of Toronto and the University of British Columbia, several rays were cast into the microgeometry for each sample, distributed in a jittered fashion about the microgeometry patch. The value of the BRDF at a given sample was computed by taking the average of all the jittered measurements.

Figure 2.3 shows a simple microgeometry consisting of a sawtooth with alternating red and blue sides. Figure 2.4 shows two plots of the BRDF that result from this microgeometry. The black line indicates the direction of incident light. The colored blobs show the distribution of reflected light where the distance from the origin shows the magnitude of the BRDF. Notice in the plot on the left that most of the reflected light is red and is reflected back toward the direction of incident light. As the incident light direction is moved to the other side in the plot on the right, notice the same effect holds, except the reflected light is mostly blue. This is the behavior that is expected from a surface whose microgeometry is like Figure 2.3. Figure 2.5 shows a plane rendered with the sawtooth BRDF.

A more realistic microgeometry is shown in Figure 2.6. The stochastically perturbed cylinders represent a velvet surface. Figure 2.7 shows the results of applying the velvet BRDF over a surface representing cloth draped over a chair.

Figure 2.3: Sawtooth microgeometry

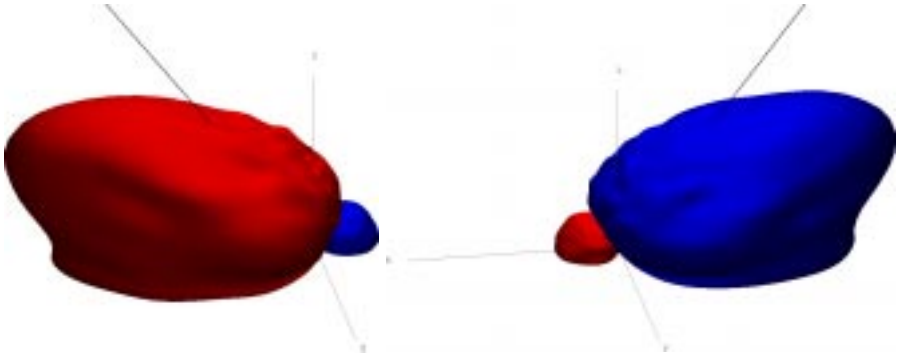

Figure 2.4: Plots of the sawtooth BRDF, left to right: $(\phi_i = 30°, \theta_i = 0°)$, $(\phi_i = 30°, \theta_i = 180°)$

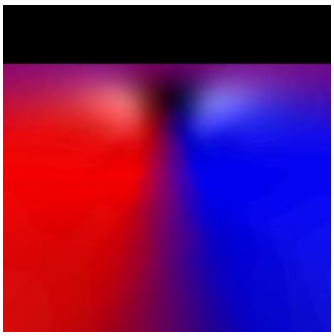

Figure 2.5: A plane rendered with the sawtooth BRDF

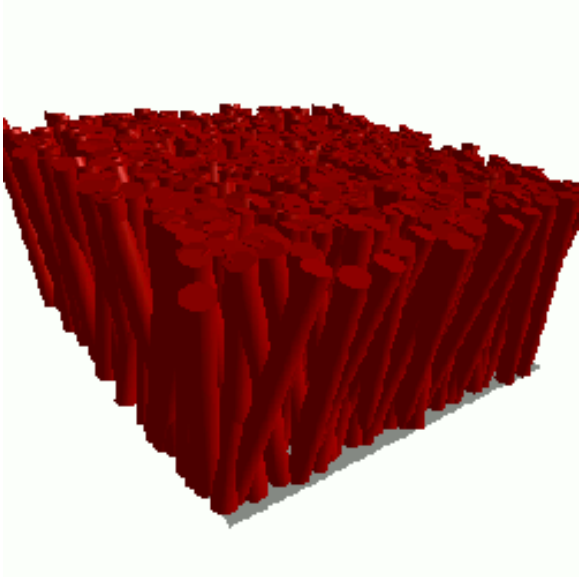

Figure 2.6: Velvet microgeometry

Figure 2.7: Velvet chair

Chapter 3

Representing and Using Tabulated BRDFs

One of the largest hurdles to overcome when using tabulated BRDFs for rendering tasks is choosing a representation of the data that stores the samples in an efficient way, and provides a reasonable way of interpolating between the samples. This chapter gives a cursory overview of several such techniques. Some of these have been used before, while others are merely suggestions of techniques that may be worthwhile of further study.

It is important to note that the purpose of this chapter is to supply the reader with an idea of some of the practical issues involved in using measured BRDF data. No attempt is made to perform a thorough evaluation and comparison of the techniques presented. Such an undertaking is beyond the scope of this thesis, and would require a great deal of further study.

3.1 Tabulated BRDF Samples

When a tabulated BRDF is sampled on a uniform grid, the samples are usually stored lexicographically in the ϕ_i, θ_i, ϕ_r, and θ_r directions. In this case, the number of samples n_{f_r} is the product of the number of samples in each of these directions, which are n_{ϕ_i}, n_{θ_i}, n_{ϕ_r}, and n_{θ_r} respectively.

$$n_{f_r} = n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r} \tag{3.1}$$

In the case of non-uniform samples, each of the n_{f_r} samples is characterized by a 5-tuple, the jth sample being:

$$(\phi_{i_j}, \theta_{i_j}, \phi_{r_j}, \theta_{r_j}, f_r(\phi_{i_j}, \theta_{i_j}, \phi_{r_j}, \theta_{r_j})) \tag{3.2}$$

Suppose the samples are uniform in the incident direction, but are non-uniform in the reflected direction. Recall that this is the case with data returned from Ward's imaging gonioreflectometer [37]. In this case, the samples are referred to as *semi-uniform*. If the situation is reversed, with non-uniform samples in the incident direction and uniform samples in the reflected direction, then semi-uniformity still holds. Under an assumption of reciprocity[1], the incident and reflected directions can be reversed. Semi-uniformity can be exploited in a variety of situations, which will be explained in Chapter 4.

The problem of interpolating between non-uniform or semi-uniform samples is a difficult one. In the general non-uniform case, evaluating the BRDF at an arbitrary position using the samples themselves involves interpolating non-uniform samples in four dimensions. When building other represen-

[1]Reciprocity is covered in Section 4.2.1.

tations of the BRDF from the tabulated data, as presented below, a technique to interpolate between non-uniform samples is also needed, at best in two dimensions. While this problem is discussed somewhat in Section 3.2.1, it demands further investigation.

Several of the techniques presented below consider the problem of defining a continuous surface over the domain of the hemisphere using the BRDF as a distance function. This is referred to as a *BRDF surface*. Approaching the interpolation problem in this way reduces it to a surface representation problem. However, when using this approach, each BRDF surface only represents the BRDF for one incident direction.

3.2 BRDF Representation and Interpolation

3.2.1 Naive Storage

The most naive BRDF representation involves simply storing the samples themselves. In the uniform grid case, they are usually stored in some meaningful lexicographic order. In the non-uniform grid case, the order is arbitrary.

To obtain values of the BRDF between the sample points, an interpolant must be introduced. In the uniform sample case, a quadrilinear interpolation scheme can be used. Using the four sample positions around both the incident and reflected directions results in a 16 sample hypercube that is used in the interpolation. While quadrilinear interpolation is extremely fast, it does not

provide tangential continuity of the BRDF surface. This lack of continuity can sometimes cause visible artifacts in the rendered image.

The non-uniform grid case presents 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.

This non-uniform interpolation technique raises the issue of how to compute distance between sample points: both when choosing the j closest samples, and when computing the distance between a sample point and the point being evaluated. A naive solution would be to simply use the Euclidean distance between the points on the unit hemisphere. More complex approaches may yield better results, but such a determination necessitates further investigation.

3.2.2 Interpolating Spline

Another way of representing the BRDF is by fitting an interpolating spline surface through the sample points. The spline can be computed in a variety of forms, including tangentially continuous and non-tangentially continuous. This representation is chosen to provide smooth interpolation between the sample points, but its main disadvantage is that interpolating splines often

oscillate undesirably. This introduces high frequency errors in the data that can show up as visible artifacts in rendered images.

3.2.3 B-spline

An alternative representation is to use the uniformly sampled data points as the control mesh of a B-spline surface. This has an advantage over the interpolating spline approach in that the values between the sample points are interpolated smoothly. The main disadvantage to this approach is that the BRDF values at a sample point are not the original measured values at that point. However, as long as the sample points are sufficiently dense relative to the curvature of the BRDF surface, the relative error will be negligible.

3.2.4 Hierarchical B-spline

Hierarchical B-splines were first introduced by Forsey and Bartels [10] and have been used effectively as an approach to facial animation. Hierarchical B-splines are B-spline surfaces with multiple levels of refinement. Areas of the surface that are very smooth are represented with a coarse control mesh, while the control mesh is denser in areas of greater detail. This representation of the BRDF surface may be beneficial since BRDF data often has different amounts of detail at different areas of its surface.

Wong proposed a technique using a least squares approach to compute a hierarchical B-spline from a dense set of points on a surface [41]. This technique could be used on a BRDF to generate a hierarchical B-spline representation of

the data, which would not only greatly reduce storage costs, but also provide a smooth interpolant.

3.2.5 Spherical Harmonics

Spherical harmonics have often been employed in representing measured BRDFs in computer graphics [5, 33, 39].

Spherical harmonics are the two dimensional analogue of the Fourier series on a sphere. Any function defined on the sphere f, such as a BRDF surface, can be represented by an infinite series of spherical harmonic basis functions.

$$f(\theta,\phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} C_{l,m} Y_{l,m}(\theta,\phi) \tag{3.3}$$

In the above equation, $Y_{l,m}(\theta,\phi)$ is the spherical harmonic basis function of order l and degree m, and $C_{l,m}$ is some constant. Of course in practice the function must be approximated by a finite number of terms.

The spherical harmonic basis functions are defined as

$$Y_{l,m}(\theta,\phi) = \begin{cases} N_{l,m} P_{l,m}(\cos\theta)\cos(m\phi) & \text{if } m > 0 \\ N_{l,m} P_{l,m}(\cos\theta)/\sqrt{2} & \text{if } m = 0 \\ N_{l,m} P_{l,|m|}(\cos\theta)\sin(|m|\phi) & \text{if } m < 0 \end{cases} \tag{3.4}$$

where $N_{l,m}$ is a normalizing constant defined as

$$N_{l,m} = \sqrt{\frac{(2l+1)(l-|m|)!}{2\pi(l+|m|)!}} \tag{3.5}$$

and $P_{l,m}(t)$ are the *associated Legendre polynomials*, defined by the recurrence

$$\begin{aligned} P_{0,0}(t) &= 1 \\ P_{m,m}(t) &= (1-2m)\sqrt{1-t^2}P_{m-1,m-1}(t) \\ P_{m+1,m}(t) &= t(2m+1)P_{m,m}(t) \\ P_{l,m}(t) &= t\left(\frac{2l-1}{l-m}\right)P_{l-1,m}(t) - \left(\frac{l+m-1}{l-m}\right)P_{l-2,m}(t) \end{aligned} \tag{3.6}$$

Plots of the first few real spherical harmonic basis functions can be found in Figure 3.1.

When fitting a function to a spherical harmonic representation, the constants $C_{l,m}$ must be found. Since the spherical harmonic basis functions are orthogonal, this can be done by taking the inner product of the function f and the appropriate basis function, and integrating that expression over the sphere [24].

$$C_{l,m} = \int_0^{2\pi}\int_0^{\pi} f(\theta,\phi)Y_{l,m}(\theta,\phi)sin\theta d\theta d\phi \tag{3.7}$$

There are several drawbacks to using spherical harmonics to represent BRDFs. First, high frequency effects, which often occur in BRDF data as specular reflection, require a large number of terms to represent accurately. In addition, the symmetry of the spherical harmonic basis functions imposes symmetry in the reconstructed BRDF data, requiring a large number of terms to accurately represent BRDFs with asymmetric features.

3.2.6 Adaptive Geodesic Sphere

Gondek et al. used an adaptive geodesic sphere approach to represent the BRDF [15]. Subdivision of a facet of the geodesic sphere into four subfacets

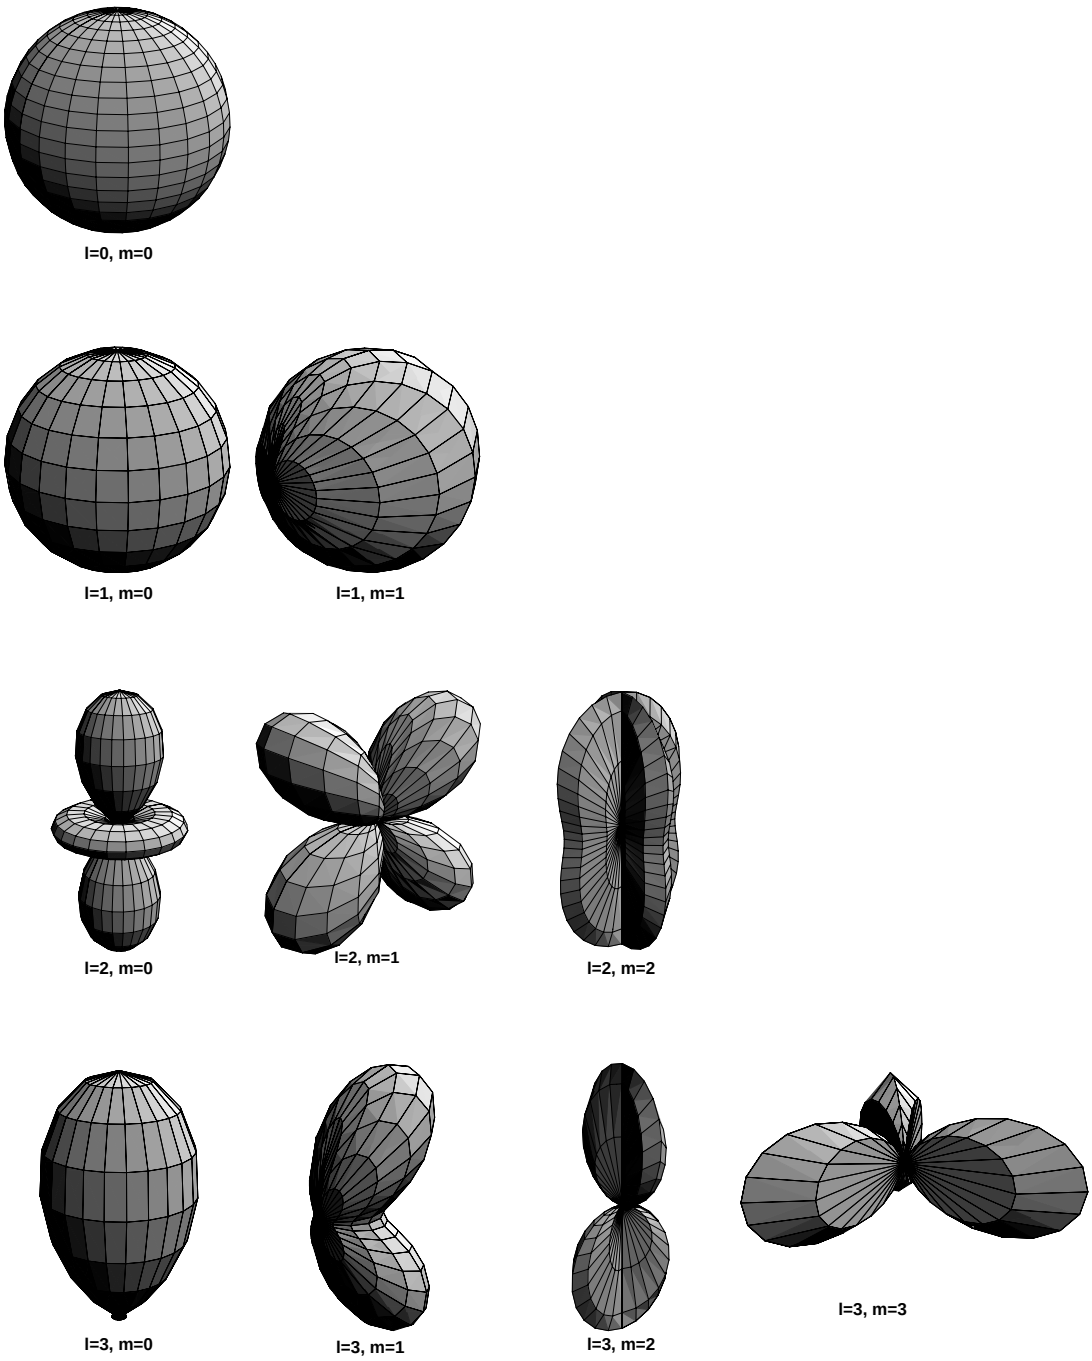

Figure 3.1: Plots of the first few real spherical harmonic basis functions

was made whenever the root mean squared difference of those four samples was above some error tolerance. Bilinear interpolation on the four closest samples was performed to evaluate the BRDF at an arbitrary position.

3.2.7 Wavelets

Wavelets are basis functions that are useful for representing signals that contain a great deal of high frequency information in some areas, but very little in others [25]. This property makes wavelets attractive as a representation for BRDF data.

Wavelets can be formulated to represent functions on the sphere, making them applicable to the BRDF representation problem. Schröder and Sweldens used BRDFs as an example of their technique for representing functions on the sphere with wavelets [32]. Their approach only represented the BRDF for one incident direction, however. Lalonde proposed solutions to the more general problem of representing the entire four-dimensional BRDF with wavelets [22]. Both of these approaches to the representation problem allow for a great deal of compression with relatively little error.

Chapter 4

Measuring BRDF Properties

In this chapter, several schemes for measuring properties of tabulated BRDFs are presented. Section 4.1 begins with a general discussion on measuring BRDF properties. The four measurements, which are reciprocity, energy conservation, isotropy, and separability, are then defined in Section 4.2.

4.1 Characteristics of the Measurements

Reciprocity and energy conservation express how closely the BRDF obeys physical laws. Lewis termed this *physical plausibility* [23]. A BRDF that is not physically plausible will not necessarily produce unrealistic looking images, but in some tasks, such as physically based global illumination, physical plausibility may be a prerequisite to rendering.

In theory, measured BRDFs should always be physically plausible by virtue of the way they are created. However, systematic errors or noise can be

introduced by a gonioreflectometer, causing the resulting BRDF to be implausible. BRDFs obtained with a virtual gonioreflectometer may lack plausibility because of shortcomings in the underlying reflection model used to reflect light into the microgeometry. Measurements of reciprocity and energy conservation and the corresponding transformations explored in Sections 5.1.1 and 5.1.2 are attempts to measure and eliminate the consequences of these shortcomings.

Isotropy and separability can be thought of as measures of compressibility, since the associated transformations detailed in Sections 5.1.3 and 5.1.4 result in considerable savings in storage space. Another way to think of these two measurements is as the amount of relative error introduced by the corresponding transformations in Sections 5.1.3 and 5.1.4.

It should be noted that these measurements do *not* represent an attempt to measure how much effect these BRDF properties will have on generated images, but strictly to measure 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 the extent to which these properties show up in an image. Measuring these effects would necessitate a separate investigation.

For consistency, each of the measurements has been formulated such that a value of zero implies that 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, mean-

ing that measurements from two tabulated BRDFs of different resolutions are comparable.

Each of these measuring techniques works directly on the samples obtained from a gonioreflectometer or equivalent device, as described in Chapter 2. If it is desired that the BRDF be stored using some other representation, these measurements, and their corresponding transformations in Chapter 5, must be performed as an intermediate step before fitting the measured data to the desired representation.

4.2 Properties

4.2.1 Reciprocity

Helmholtz's reciprocity rule states that if the incident and reflected directions are reversed, the value of the BRDF should not change [36]. 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

$$P_r = \sqrt{\frac{\sum_{\phi_i} \sum_{\theta_i} \sum_{\phi_r} \sum_{\theta_r} [f_r(\phi_i, \theta_i, \phi_r, \theta_r) - f_r(\phi_r, \theta_r, \phi_i, \theta_i)]^2}{2 n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r}}} \tag{4.1}$$

Notice that an additional factor of 2 is placed 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 uniform grid, but are not necessarily matched as specified above, then the value of $f_r(\phi_r, \theta_r, \phi_i, \theta_i)$ can be computed using quadrilinear interpolation as described in Section 3.2.1. If the samples are on a non-uniform grid, then the reciprocal can likewise be computed using a non-uniform interpolation method such as in Section 3.2.1.

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

4.2.2 Energy Conservation

The law of energy conservation, which is a direct consequence of the Second Law of Thermodynamics [27], 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

$$\frac{M}{E} = \int_{\Omega_N} f_r(\vec{\omega_i} \rightarrow \vec{\omega_r}) \, (N \cdot \vec{\omega_r}) \, d\omega_r \tag{4.2}$$

For a BRDF f_r to conserve energy, $\frac{M}{E}$ must be less than or equal to one for all possible values of $\vec{\omega_i}$[1].

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 4.2 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 $(N \cdot \vec{\omega_r}) = \cos\theta_r$, the integral

[1]The fraction $\frac{M}{E}$ is sometimes referred to as *albedo*.

can be expressed as

$$\frac{M}{E} = \int_{\phi_r} \int_{\theta_r} f_r(\phi_i, \theta_i, \phi_r, \theta_r) \sin\theta_r \cos\theta_r \; d\theta_r \; d\phi_r \tag{4.3}$$

When the samples are uniform in the θ and ϕ directions, the integral in Equation 4.3 can be approximated with a sum. For consistency, one is subtracted from the resulting integral so that energy conserving behavior will be expressed as zero. This yields

$$, (\phi_i, \theta_i) = \max(0, \sum_{\phi_r} \sum_{\theta_r} [f_r(\phi_i, \theta_i, \phi_r, \theta_r) \sin\theta_r \cos\theta_r \; \Delta\phi_r \Delta\theta_r] - 1) \tag{4.4}$$

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

To determine energy conserving behavior of the entire BRDF, the average of all , values is taken, giving a final measure.

$$P_{ec} = \frac{\sum_{\phi_i} \sum_{\theta_i} , (\phi_i, \theta_i)}{n_{\phi_i} n_{\theta_i}} \tag{4.5}$$

Notice that there is no need to assume that the samples are uniform in the ϕ_i and θ_i directions. This means that this measurement can be used for both uniform *and* semi-uniform sample grids.

When the samples are collected on a non-uniform grid, the integral must be approximated in some other way. 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 4.2 is the "differential-to-finite form factor" between the triangle and the infinitesimal surface element dA. This can be computed using classic methods from the field of radiative transfer theory such

as the *Nusselt analog* [29]. The integral , (ϕ_i, θ_i) is then obtained by summing the terms from all of the triangles.

4.2.3 Isotropy

In general, BRDFs as described in Equation 1.3 are *anisotropic*, meaning that the value of the BRDF can vary 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

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = f_r^{iso}(\theta_i, \phi_r - \phi_i, \theta_r) \tag{4.6}$$

This simplification can reduce storage and computational costs considerably, reducing the total number of samples in Equation 3.1 to

$$n_{f_r} = n_{\theta_i} n_{\phi_r} n_{\theta_r} \tag{4.7}$$

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

$$\mu_i(\theta_i, \phi_+, \theta_r) = \frac{\sum_{\phi_i} f_r(\phi_i, \theta_i, \phi_i + \phi_+, \theta_r)}{n_{\phi_i}} \tag{4.8}$$

Isotropy is related to the average deviation from the mean μ_i, expressed as

$$\sigma_i(\theta_i, \phi_+, \theta_r) = \sqrt{\frac{\sum_{\phi_i} [f_r(\phi_i, \theta_i, \phi_i + \phi_+, \theta_r) - \mu_i(\theta_i, \phi_+, \theta_r)]^2}{n_{\phi_i}}} \tag{4.9}$$

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

To obtain the final measure of isotropy, the average of all σ_i values over all possible values of $(\theta_i, \phi_+, \theta_r)$ is computed.

$$P_i = \frac{\sum_{\theta_i} \sum_{\phi_+} \sum_{\theta_r} \sigma_i(\theta_i, \phi_+, \theta_r)}{n_{\theta_i} n_{\phi_r} n_{\theta_r}} \tag{4.10}$$

In the case of non-uniform samples, 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 method described in Section 3.2.1. The resulting set of BRDF values can be used in the computation of μ_i in Equation 4.8.

4.2.4 Separability

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

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = f_r^{in}(\phi_i, \theta_i) \times f_r^{out}(\phi_r, \theta_r) \tag{4.11}$$

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

$$n_{f_r} = n_{\phi_i} n_{\theta_i} + n_{\phi_r} n_{\theta_r} \tag{4.12}$$

An additional benefit is that radiosity computations, like those described by Cohen and Wallace [6], can be done efficiently with non-Lambertian surfaces [26].

For the discussion of separability, an assumption of the existence of matching sample pairs is made, as in Section 4.2.1. If this is in fact not the case, the missing pairs can be generated using the interpolation techniques

described in Section 3.2.1. Of course, if one assumes 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)[2]. If the matrix A is filled with the BRDF values, the two separated functions f_r^{in} and f_r^{out} can be thought of as two vectors u and v whose outer product results in A. To find u and v, A must be factored. The singular value decomposition (SVD) numerical technique is ideal for this, since it provides a measure of how close to factorable A is, as well as actually determines the most probable values for u and v (see Section 5.1.4)[3].

The singular value decomposition separates A into three matrices.

$$A = UDV^T \tag{4.13}$$

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-zero elements in D specifies the *rank* of A. The rank of A 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

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

[3]For more details on the SVD, the reader is directed to the book by Watkins [38]. Code implementing SVD can be found in Numerical Recipes [31].

separable the function is, it is necessary to define another $Q \times Q$ matrix $D^{(1)}$.

$$D^{(1)} = \begin{bmatrix} d_{11} & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{4.14}$$

$D^{(1)}$ is then substituted for D. Performing the multiplication in Equation 4.15 results in $A^{(1)}$, the matrix of rank one based on A. Notice that the sparse nature of $D^{(1)}$ allows this multiplication to be performed quickly.

$$A^{(1)} = UD^{(1)}V^T \tag{4.15}$$

The magnitude of the difference between $A^{(1)}$ and A gives a measure of separability. This magnitude is measured in a way that is similar to the Frobenius norm of the difference matrix [14], except that each term in the sum is divided by the total number of elements in the matrix. This removes the dependence of P_s on the number of samples in the tabulated BRDF.

$$P_s = \sqrt{\sum_p \sum_q \frac{[A^{(1)}_{pq} - A_{pq}]^2}{n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r}}} \tag{4.16}$$

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 in D acting as weights. Equation 4.16 only deals with expressing A as the product of the first such column and row. Fournier showed in [11] 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 function

products, it is said to be k-separable. The k-separated BRDF is reconstructed as:

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = \sum_{j=1}^{k} [f_r^{in(j)} \times f_r^{out(j)}] \tag{4.17}$$

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

$$D^{(k)} = \begin{bmatrix} d_{11} & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & d_{22} & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & \ddots & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & d_{kk} & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ \vdots & & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \tag{4.18}$$

A re-multiplication is then performed to get the rank k matrix $A^{(k)}$.

$$A^{(k)} = UD^{(k)}V^T \tag{4.19}$$

This provides a measure of the k-separability $P_s^{(k)}$.

$$P_s^{(k)} = \sqrt{\sum_p \sum_q \frac{[A_{pq}^{(k)} - A_{pq}]^2}{n_{\phi_i} n_{\theta_i} n_{\phi_r} n_{\theta_r}}} \tag{4.20}$$

Chapter 5

Transforming BRDFs

It may be desirable to change a tabulated BRDF to possess one or more of the four properties covered in the previous chapter. 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.

It may also be desirable to change the BRDF to be *closer* to possessing a certain property, placing it intuitively in between the original BRDF and one fully possessing the property. The constant δ indicates how close to a given property the BRDF should be. Each section of this chapter first shows how to transform a BRDF to fully possess the property ($\delta = 1$), and then generalizes the technique to allow for more gradual transformations ($0 < \delta < 1$).

If the values of the BRDF 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.

5.1 Transformations

5.1.1 Reciprocity

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

$$f_r^r(\phi_i, \theta_i, \phi_r, \theta_r) = \frac{f_r(\phi_i, \theta_i, \phi_r, \theta_r) + f_r(\phi_r, \theta_r, \phi_i, \theta_i)}{2} \tag{5.1}$$

Making the BRDF *more* reciprocal is achieved by linearly interpolating between a given sample and the average between that sample and its reciprocal. This is expressed as

$$f_r'(\phi_i, \theta_i, \phi_r, \theta_r, \delta) = (1 - \delta) f_r(\phi_i, \theta_i, \phi_r, \theta_r) + \delta f_r^r(\phi_i, \theta_i, \phi_r, \theta_r) \tag{5.2}$$

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 its reciprocal will be the result of non-uniform 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.

5.1.2 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 , (ϕ_i, θ_i) in Equation 4.4. It is necessary to divide each BRDF value by a scalar x that will make , (ϕ_i, θ_i) exactly zero.

$$\begin{gathered}
\sum_{\phi_r} \sum_{\theta_r} [\frac{f_r(\phi_i,\theta_i,\phi_r,\theta_r)}{x} \sin\theta_r \cos\theta_r \ \Delta\phi_r \Delta\theta_r] - 1 = 0 \\
\frac{1}{x} \sum_{\phi_r} \sum_{\theta_r} [f_r(\phi_i, \theta_i, \phi_r, \theta_r) \sin\theta_r \cos\theta_r \ \Delta\phi_r \Delta\theta_r] = 1 \\
\frac{1}{x} = \frac{1}{\sum_{\phi_r} \sum_{\theta_r} [f_r(\phi_i,\theta_i,\phi_r,\theta_r) \sin\theta_r \cos\theta_r \ \Delta\phi_r \Delta\theta_r]} \\
x = , (\phi_i, \theta_i) + 1
\end{gathered} \tag{5.3}$$

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

$$f_r^{ec}(\phi_i, \theta_i, \phi_r, \theta_r) = \frac{f_r(\phi_i, \theta_i, \phi_r, \theta_r)}{, (\phi_i, \theta_i) + 1} \tag{5.4}$$

To transform a tabulated BRDF to be *closer* to conserving energy, , (ϕ_i, θ_i) is simply multiplied by δ.

$$f_r'(\phi_i, \theta_i, \phi_r, \theta_r, \delta) = \frac{f_r(\phi_i, \theta_i, \phi_r, \theta_r)}{\delta, (\phi_i, \theta_i) + 1} \tag{5.5}$$

It could be argued that this technique is somewhat artificial. It only changes those incident directions that violate energy conservation by scaling

the data down 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. This causes differences in the data that may be noticeable in rendered images. This undesirable effect would be particularly striking in an animation where a sequence of images contained shading computations that used a variety of incident directions, each of which was transformed by a different scaling factor.

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 surface in the real world exhibits such ideal behavior. Even the most perfect reflectors ever observed have values of $\frac{M}{E}$ less than one [19].

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 , (ϕ_i, θ_i), changing Equation 5.4 to

$$f_r^{ec}(\phi_i, \theta_i, \phi_r, \theta_r) = \frac{f_r(\phi_i, \theta_i, \phi_r, \theta_r)}{\max_{(\phi_i, \theta_i)} , (\phi_i, \theta_i) + 1} \tag{5.6}$$

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 , (ϕ_i, θ_i) in Equation 4.4 to

$$, (\phi_i, \theta_i) = \max(0, \sum_{\phi_r} \sum_{\theta_r} [f_r(\phi_i, \theta_i, \phi_r, \theta_r) \sin\theta_r \cos\theta_r \ \Delta\phi_r \Delta\theta_r] - \tau) \tag{5.7}$$

The energy conserving BRDF is then computed by changing Equation 5.4 to

$$f_r^{ec}(\phi_i, \theta_i, \phi_r, \theta_r) = \frac{\tau \cdot f_r(\phi_i, \theta_i, \phi_r, \theta_r)}{, (\phi_i, \theta_i) + \tau} \tag{5.8}$$

5.1.3 Isotropy

Transforming a tabulated BRDF to be isotropic is very simple. The average BRDF value μ_i from Equation 4.8 is used for all values with the same ϕ_+. This yields

$$f_r^{iso}(\theta_i, \phi_+, \theta_r) = \mu_i(\theta_i, \phi_+, \theta_r) \tag{5.9}$$

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

$$f_r'(\phi_i, \theta_i, \phi_r, \theta_r, \delta) = (1-\delta) f_r(\phi_i, \theta_i, \phi_r, \theta_r) + \delta \mu_i(\theta_i, \phi_r - \phi_i, \theta_r) \tag{5.10}$$

5.1.4 Separability

Recall from Section 4.2.4 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}[1].

$$f_r^{in}(\phi_i, \theta_i) = d_{11} u_p \tag{5.11}$$

$$f_r^{out}(\phi_r, \theta_r) = v_q \tag{5.12}$$

Finding the jth separated component simply involves using the jth column of U, $u^{(j)}$ and the jth row of V^T, $v^{(j)}$.

$$f_r^{in(j)}(\phi_i, \theta_i) = d_{jj} u_p^{(j)} \tag{5.13}$$

[1] In Equations 5.11 through 5.15, $p = \phi_i n_{\theta_i} + \theta_i$ and $q = \phi_r n_{\theta_r} + \theta_r$.

$$f_r^{out(j)}(\phi_r, \theta_r) = v_q^{(j)} \tag{5.14}$$

This is identical to Fournier's technique of separating a BRDF into a sum of k separated components [11].

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

$$f_r'(\phi_i, \theta_i, \phi_r, \theta_r, \delta) = (1 - \delta)A_{pq} + \delta A_{pq}^{(k)} \tag{5.15}$$

5.2 Experiments and Results

Figure 5.1 shows an example of performing a reciprocity transform on a Phong BRDF where $P_r = 0.117646$. Lewis showed in [23] that Phong shaders are never reciprocal. Performing the transformation resulted in a noticeable change in the BRDF data, but the resulting rendered spheres in Figure 5.2 are virtually identical. Differences in P_r do not seem to affect rendered images a great deal.

Figure 5.3 demonstrates a 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 at the right is the energy conserving version. The reduction in size is more subtle than in Figure 5.4, where the incident polar angle is larger ($\theta_i = (90 - \epsilon)°$, for some small ϵ). This is exactly what is expected since Phong's energy conserving behavior becomes worse at large incident polar angles. Figure 5.5 shows spheres rendered with these two BRDFs. Again, the visual difference

is negligible. However, the underlying difference between the two BRDFs is very important, especially if the BRDFs are used in global illumination computations.

Figure 5.6 shows an example of a BRDF transformed to be isotropic. This anisotropic BRDF that simulates brushed metal was generated by running a microgeometry of parallel cylinders through a virtual gonioreflectometer[2]. 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 5.7 shows plots of several scratched metal BRDFs. The one on the left is the original BRDF used in Figure 5.6. The one on the right is the isotropic version of that BRDF. The middle one is halfway between the other two ($\delta = 0.5$).

Figure 5.8 demonstrates separating a BRDF. The microgeometry in Figure 2.6 was used to obtain a velvet BRDF, which was used to render the top chair. It was then separated ($k = 1$) and used to render the bottom-left chair. The bottom-right chair was rendered with a version made up of a sum of five separated components ($k = 5$). Notice that this version is virtually indistinguishable from the original, yet the BRDF only requires 4.4% of the storage space! Figure 5.9 shows the value of $P_s^{(k)}$ for the velvet BRDF as k increases. Curves for other BRDFs exhibited similar behavior.

[2]The use of parallel cylinders to simulate brushed metal was inspired by the work of Poulin and Fournier [30].

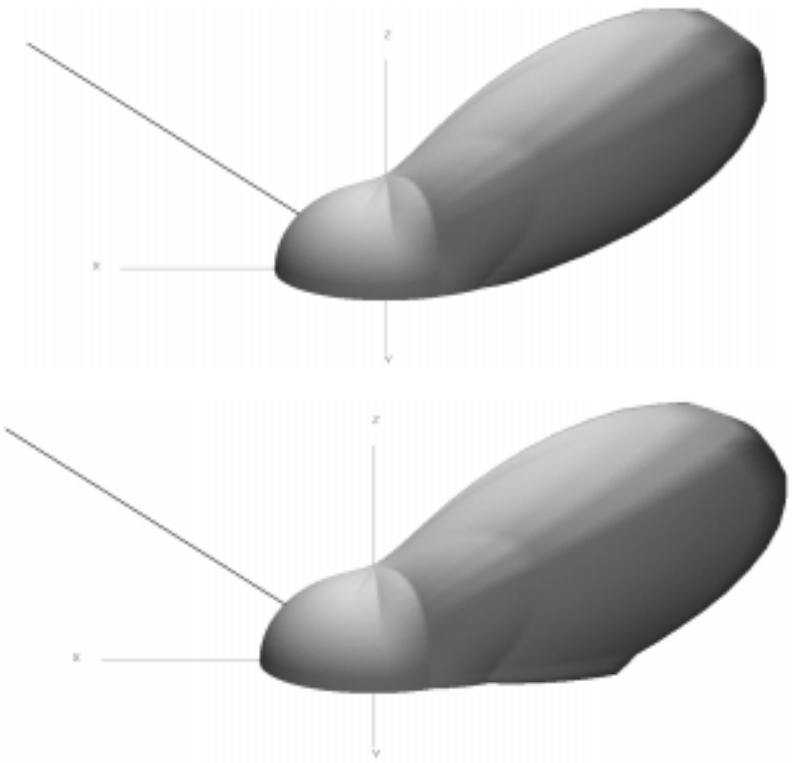

Figure 5.1: Two plots of a Phong BRDF ($\phi_i = 0°, \theta_i = 50°$), before and after a reciprocity transformation

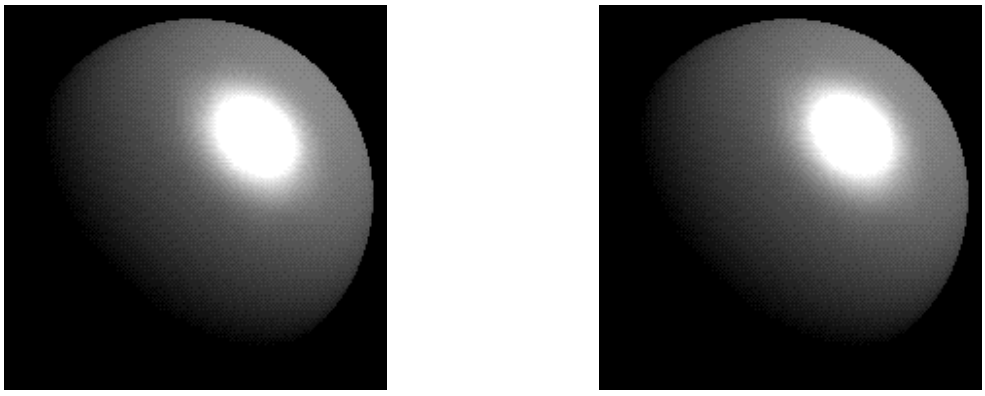

Figure 5.2: Two spheres rendered with the BRDFs from Figure 5.1

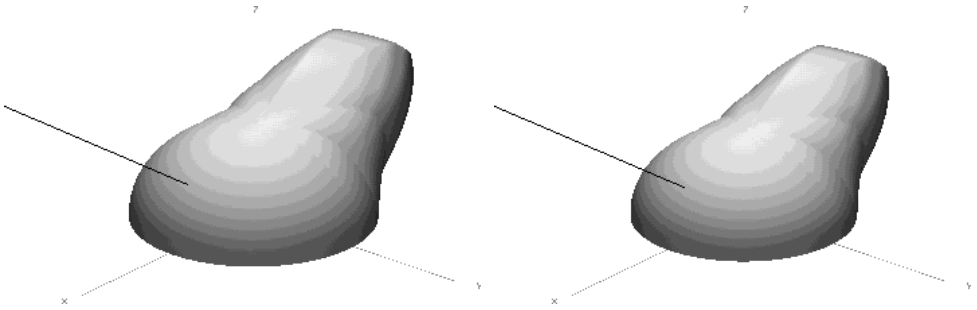

Figure 5.3: Two plots of a Phong BRDF ($\phi_i = 0°, \theta_i = 50°$), before and after an energy conservation transformation

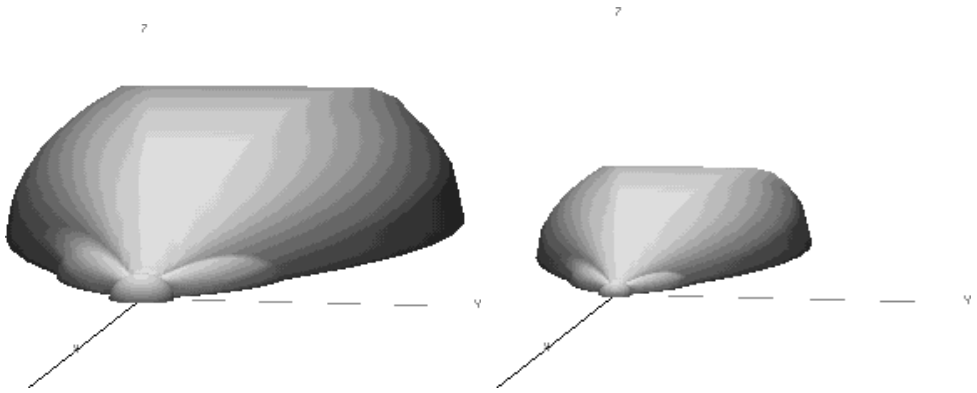

Figure 5.4: Two plots of a Phong BRDF ($\phi_i = 0°, \theta_i = (90 - \epsilon)°$), before and after an energy conservation transformation

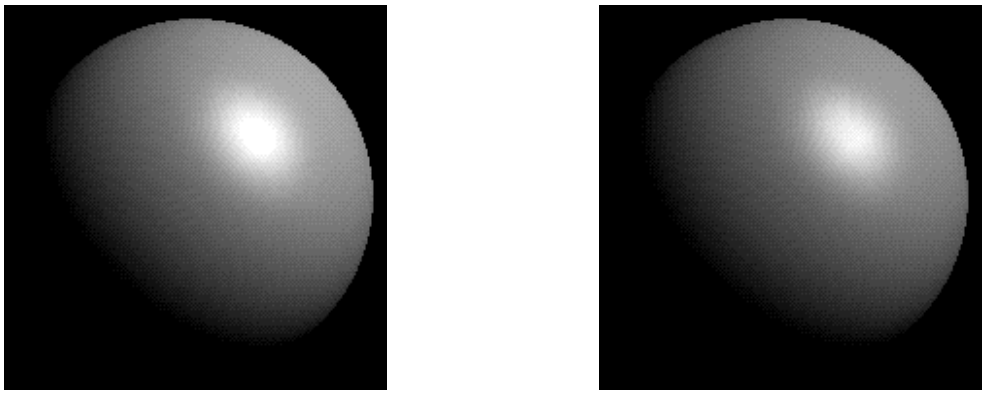

Figure 5.5: Two spheres rendered with the BRDFs from Figures 5.3 and 5.4.

Figure 5.6: Two teapots rendered with an anisotropic BRDF representing brushed metal and the transformed isotropic version

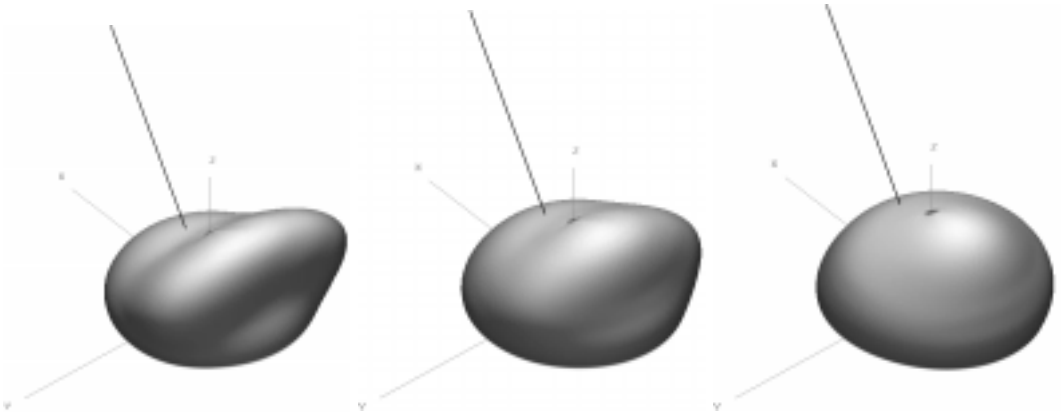

Figure 5.7: Plots of three brushed metal BRDFs ($\phi_i = 0°$, $\theta_i = 20°$) transformed to be isotropic with different δ values

Figure 5.8: Three velvet chairs

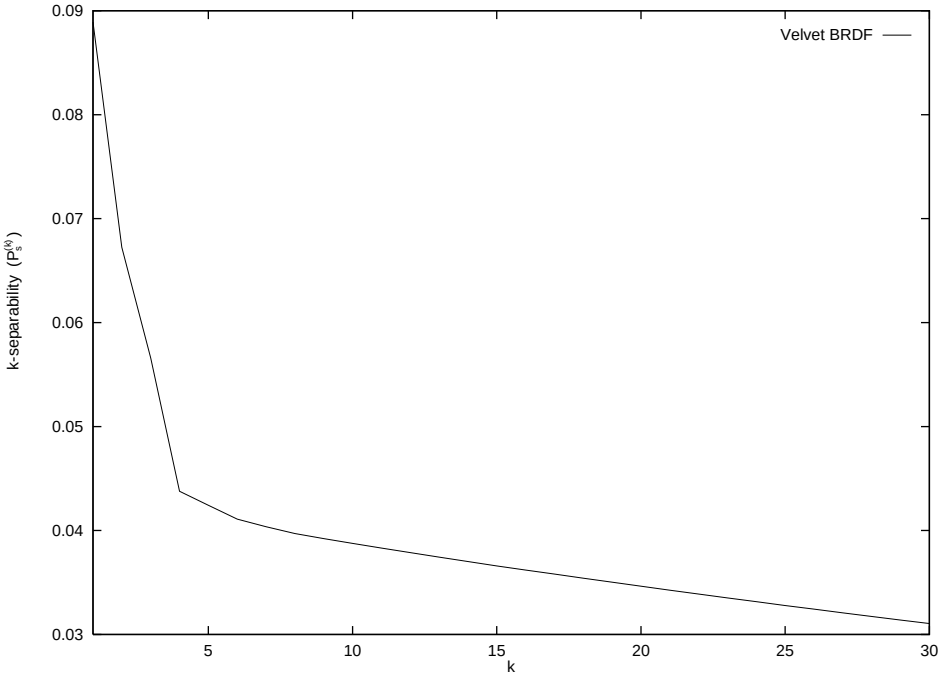

Figure 5.9: Measure of k-separability of the velvet BRDF for $1 \leq k \leq 30$.

Chapter 6

Conclusion

6.1 Contributions of this thesis

This thesis has proposed techniques for measuring four properties of BRDFs: reciprocity, energy conservation, isotropy, and separability. As well, methods for transforming BRDF data to possess those properties have been presented.

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.

6.2 Future Research

While the primary contributions of this thesis are the properties and transformations described in Chapters 4 and 5, several related issues have been touched upon that demand further investigation.

The differences in the techniques for BRDF representation and interpolation presented in Chapter 3 are of interest. Namely, a comprehensive comparison of the different storage schemes should be carried out. Furthermore, the interpolation techniques for non-uniform and semi-uniform samples should be studied further. There is a great deal of literature on non-uniform sampling and interpolation. Research that investigated how work in this field applied to the non-uniform BRDF interpolation problem would be useful.

The property measurements and transformations themselves are rather *ad hoc*. This is especially evident when considering the variety of options presented for transforming a BRDF to conserve energy (Section 5.1.2). It would be interesting to conduct a more rigorous comparison of these techniques to determine which is more useful in different contexts. Furthermore, the possibility of using non-linear interpolation in a partial transformation (where $0 < \delta < 1$) should be examined, since the only techniques presented here linearly interpolate between the original and transformed BRDF values.

Appendix A

Principles of Radiometry

This section overviews the radiometric principles underlying this thesis. Basic concepts and definitions of radiometric units are provided in Section A.1. Important principles related to the reflection of light from surfaces are covered in Section A.2. Section A.3 suggests some additional references.

A.1 Measuring Light

This thesis revolves around modeling light and reflection based on physical principles. Such an undertaking is driven by concepts in the field of radiometry.

Radiometry is itself based on the abstract physical models of transport theory. Transport theory was developed to explain the behavior of particles flowing in three dimensions, and is therefore applicable to many disciplines. Glassner gives an excellent overview of transport theory as it relates to computer graphics [13].

Radiometry can be thought of as a specific application of transport theory, where the particles being studied are photons, and their behavior is constrained by their unique properties.

A.1.1 Radiometry and Photometry

The field of radiometry has in some sense been motivated by the much older field of photometry. Like radiometry, photometry is concerned with the nature of light, but specifically in relation to its effect on the human visual system. Radiometry is the study of light energy and reflection independent of the characteristics of any particular detector.

While it is important to be aware of how the human visual system responds to light, incorporating such factors into an otherwise objective study of light makes things drastically more complex.

A.1.2 Radiometric Units

The most basic radiometric unit is *radiant energy*, denoted Q, which is measured in joules. For a given wavelength λ, each photon can be thought of as having a constant amount of energy, determined by Planck's constant h and the speed of light in a vacuum c.

$$Q = \frac{hc}{\lambda} \qquad [J] \tag{A.1}$$

Considering the amount of energy that flows through a surface per unit time yields a measure of *radiant flux* or *radiant power*. It is denoted by Φ, and

is measured in watts, or joules per second.

$$\Phi = \frac{dQ}{dt} \qquad \left[W \equiv \frac{J}{s}\right] \tag{A.2}$$

To consider how light interacts with surfaces, it is necessary to have a measure of radiant flux per unit area. This measurement is called *radiant flux area density*[1], and is denoted by ξ.

$$\xi = \frac{d\Phi}{dA} \qquad \left[\frac{W}{m^2}\right] \tag{A.3}$$

Measuring radiant flux area density is more interesting and useful when it is clear whether the measure refers to a surface receiving energy or emitting energy. If the energy is arriving at a surface, it is referred to as *irradiance*, and denoted by E.

$$E = \frac{d\Phi}{dA} \qquad \left[\frac{W}{m^2}\right] \tag{A.4}$$

If the energy is leaving a surface, it is referred to as the *exitance* M. In some computer graphics literature, exitance is known as *radiosity* and denoted with B.

$$M = B = \frac{d\Phi}{dA} \qquad \left[\frac{W}{m^2}\right] \tag{A.5}$$

It is sometimes useful to express a ratio of radiant flux to solid angle rather than radiant flux to area, since point light sources, which are common in lighting models, have an area of zero. The ratio of radiant flux to solid angle is known as the *radiant intensity*, and is designated with I.

$$I = \frac{d\Phi}{d\omega} \qquad \left[\frac{W}{sr}\right] \tag{A.6}$$

[1]The ANSI/IES standard [20] does not have a notation for generic radiant flux area density, since irradiance and exitance are used more in practice.

Symbol	Name	Definition	Units
Q	radiant energy	hv	J
Φ	radiant flux (power)	dQ/dt	$J/s \equiv W$
ξ	radiant flux area density	$d\Phi/dA$	W/m^2
E	irradiance	$d\Phi/dA$	W/m^2
M	exitance	$d\Phi/dA$	W/m^2
I	radiant intensity	$d\Phi/d\omega$	W/sr
L	radiance	$d^2\Phi/(d\omega \cdot dA)$	$W/(sr \cdot m^2)$

Table A.1: Summary of radiometric units

The final radiometric measure related to this thesis, and perhaps the most important, is *radiance*. Radiance is a measure of radiant flux per unit projected area per unit solid angle, and is designated with L.

$$L = \frac{d^2\Phi}{dA \cdot d\omega \cdot \cos\theta} \qquad \left[\frac{W}{m^2 \cdot sr}\right] \tag{A.7}$$

where θ is the angle between the normal N of the surface element dA, and the direction of the flux Φ.

Observe that using the definitions above, we can express radiance in terms of radiant intensity, irradiance, or radiant exitance:

$$L = \frac{dI}{dA \cdot \cos\theta} = \frac{dE}{d\omega \cdot \cos\theta} = \frac{dM}{d\omega \cdot \cos\theta} \qquad \left[\frac{W}{m^2 \cdot sr}\right] \tag{A.8}$$

A summary of the units that have been presented are shown in Table A.1.

A.1.3 Foreshortening in Radiance

It is not immediately clear why the cosine term is included in the definition of radiance above. Although the units imply that radiance is only dependent on

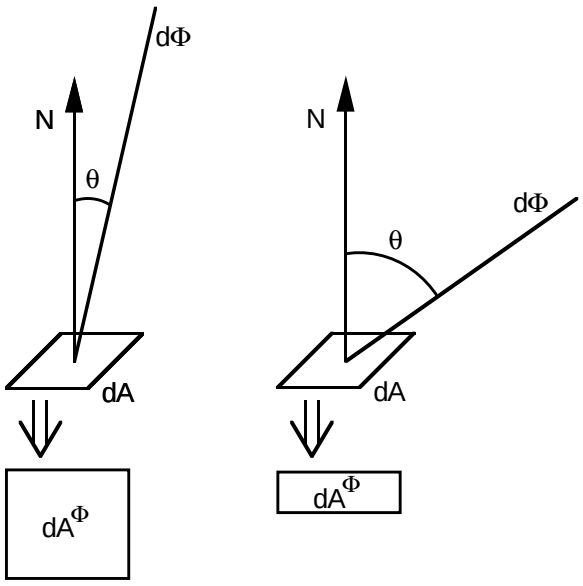

Figure A.1: Radiance demonstrated with projected area

radiant flux, area, and a solid angle, it is also dependent on the direction of the flux (either arriving or leaving) relative to the surface element being irradiated or emitting light. This implies that radiance includes a dimensionless direction as part of its unit.

Consider the following alternate definition of radiance:

$$L = \frac{d^2\Phi}{dA \cdot d\omega \cdot \cos\theta} = \frac{d^2\Phi}{dA^\Phi \cdot d\omega} = \frac{d^2\Phi}{dA \cdot d\omega^N} \qquad \left[\frac{W}{m^2 \cdot sr}\right] \quad \text{(A.9)}$$

The notation dA^Φ refers to the surface element dA projected in the direction of the flux Φ. Likewise, $d\omega^N$ refers to the solid angle $d\omega$ project in the direction of the normal N. Each of these terms takes the place of the cosine term, and simply represent different ways of thinking about the dependence of radiance on a direction vector. Figures A.1 and A.2 demonstrate these two perspectives.

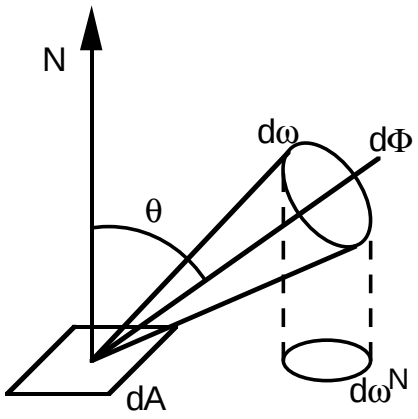

Figure A.2: Radiance demonstrated with projected solid angle

Figure A.1 shows radiant flux arriving from two different directions, one with a larger θ value than the other. The rectangles below each of the diagrams indicate how the surface element would look from the direction of the flux, that is under an orthographic projection in that direction. The size of this projected area is proportional to the cosine of the angle θ. Figure A.2 shows the solid angle $d\omega$ projected onto the plane of the surface element dA, that is in the direction of the normal to dA. Again, the size of this projected area is proportional to the cosine of the angle θ.

Another way to rationalize the cosine term is to think of water flowing through a wire hoop. As the face of the hoop is oriented away from the direction of the flowing water, the amount of water passing through the hoop decreases. This change in flow is proportional to the cosine of the angle between the normal of the hoop face and the direction of the flow of water. The

water can be thought of as the radiant flux Φ and the hoop can be thought of as the surface element dA.

A.2 Reflection from Surfaces

Nicodemus et al. clarify that there is a notable difference between the terms reflection and reflectance [28]. *Reflection* is defined as the process whereby electromagnetic flux, incident on a stationary surface or medium leaves that surface or medium from the incident side without changing frequency. By contrast, *reflectance* is the fraction of the incident flux that is reflected.

This section demonstrates the motivation for categorizing surface reflectance properties with bidirectional reflectance distribution functions. Several other important measures of reflectance are also derived and explained.

A.2.1 The BSSRDF

Consider the geometry of a patch of surface A in Figure A.3. Radiant flux arriving at a differential surface element dA_i at point p from direction $\vec{\omega_i}$ through a differential solid angle $d\omega_i$ results in a certain amount of radiance reflected from a point q in direction $\vec{\omega_r}$. The amount of reflected radiance L_r is proportional to the amount of incident flux Φ_i.

$$dL_r \propto d\Phi_i \tag{A.10}$$

The constant of proportionality in Equation A.10 is called the *bidirectional surface scattering reflectance distribution function*, or BSSRDF, and is

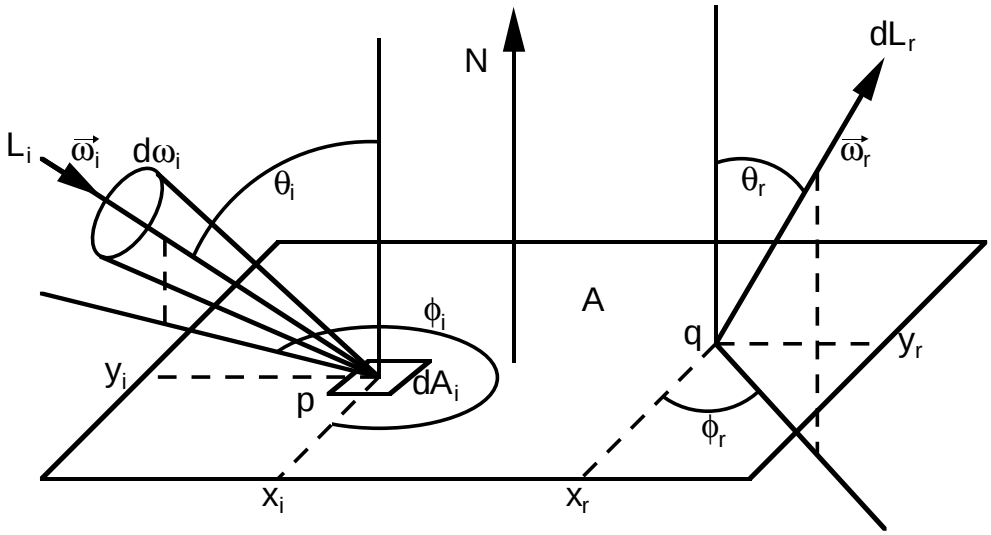

Figure A.3: Geometry of the BSSRDF

denoted S.

$$S = \frac{dL_r}{d\Phi_i} \qquad \left[m^{-2} \cdot sr^{-1}\right] \tag{A.11}$$

Notice that the BSSRDF provides a very general categorization of reflectance, making no simplifying assumptions other than those made by the geometrical ray optics model. Recall that the positions p and q refer to the pairs (x_i, y_i) and (x_r, y_r) respectively, which are relative to some parametric coordinate system on the surface. The directions $\vec{\omega_i}$ and $\vec{\omega_r}$ are represented with the pairs (ϕ_i, θ_i) and (ϕ_r, θ_r) respectively. Thus, the BSSRDF is a function of eight independent variables. It can be written as $S(\phi_i, \theta_i, x_i, y_i, \phi_r, \theta_r, x_r, y_r)$ to express these dependencies explicitly. This complexity makes the BSSRDF very difficult to measure, store, or use in practice.

A.2.2 The BRDF

The generality of the BSSRDF allows for very complex reflectance properties, most notably that of subsurface scattering. If the possibility of subsurface scattering is eliminated, the dependence on the points p and q is removed. This assumption implies that the amount of incident flux is constant over all of the surface patch A. Furthermore, the reflectance properties of the surface do not vary at different positions on the surface. With these assumptions, it is convenient to integrate the BSSRDF over the entire surface patch A. This results in the *bidirectional reflectance distribution function*, or BRDF, which is denoted[2] f_r.

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = \int_A S(\phi_i, \theta_i, x_i, y_i, \phi_r, \theta_r, x_r, y_r)\, dA \qquad \left[sr^{-1}\right] \tag{A.12}$$

It would be convenient to express the BRDF as a radiometric relation between incident radiance L_i and reflected radiance L_r, as in Figure A.4. To this end, S in the above equation is substituted with its definition from Equation A.11.

$$f_r(\phi_i, \theta_i, \phi_r, \theta_r) = \int_A \frac{dL_r}{d\Phi_i}\, dA \tag{A.13}$$

Notice that the assumption of the uniformity of A above implies that the reflected radiance is no longer a function of the positions p and q on the patch.

[2]The r subscript in f_r represents the distance between the points p and q. Historically, the BRDF has been presented as dependent on this distance. However, instead of being listed as an explicit dependency, the r was placed in the subscript. In the field of computer graphics, r is almost always fixed at a value of zero. Consequently, the presentation here has ignored the BRDF's dependency on this variable.

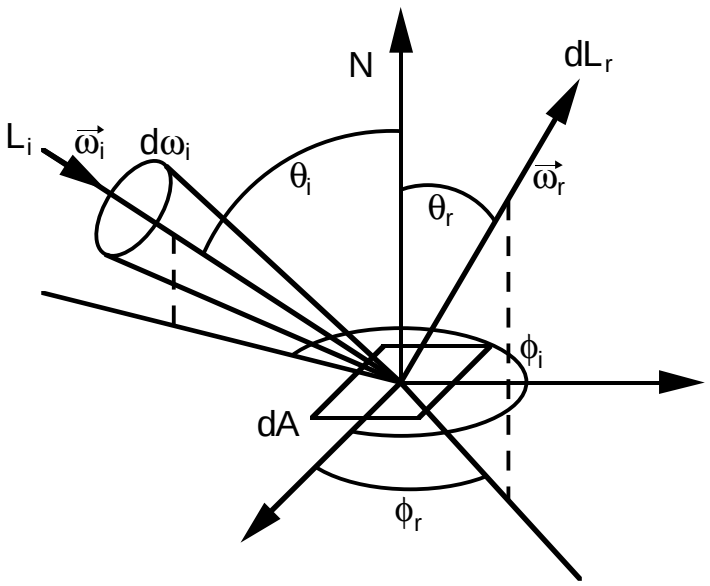

Figure A.4: Geometry of the BRDF

This observation results in the following equality:

$$dL_r = \int_A dL_r \tag{A.14}$$

Thus, the integral in Equation A.13 can be removed. Using the definitions from Section A.1.2 allows the following derivation:

$$\begin{aligned} f_r(\phi_i, \theta_i, \phi_r, \theta_r) &= \int_A \frac{dL_r}{d\Phi_i}\, dA \\ &= \frac{dL_r\ dA}{d\Phi_i} \\ &= \frac{dL_r}{dE} \\ &= \frac{dL_r}{L_i\ \cos\theta_i\ d\omega_i} \end{aligned} \tag{A.15}$$

A.2.3 Other Measures of Reflectance

Values of a BRDF can vary from zero to infinity. However, it is sometimes useful to measure reflectance as a ratio of reflected flux to incident flux. Such

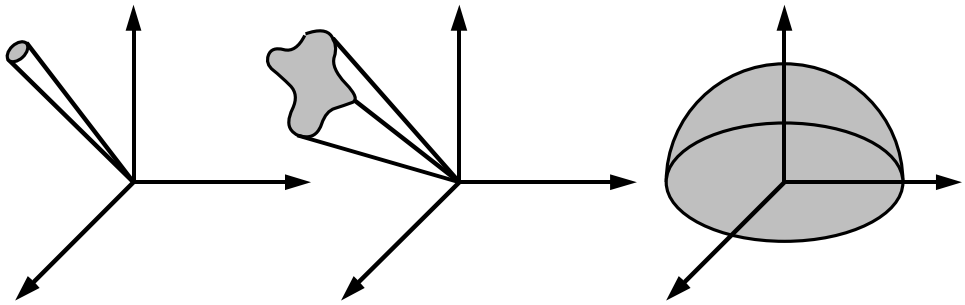

Figure A.5: Left to right: directional, conical, and hemispherical solid angles

a measurement is referred to simply as the *reflectance*, and is denoted ρ.

$$\rho = \frac{d\Phi_r}{d\Phi_i} \qquad \text{[dimensionless]} \qquad \text{(A.16)}$$

Since it is a ratio, the value of ρ varies from zero to one.

Another measurement is the *reflectance factor*. It is denoted R and is defined as the ratio of reflected flux to the reflected flux of a perfectly diffuse reflector (Φ_d). A perfectly diffuse reflector is a hypothetical surface which reflects radiation equally in all directions, and does not absorb any energy[3].

$$R = \frac{d\Phi_r}{d\Phi_d} \qquad \text{[dimensionless]} \qquad \text{(A.17)}$$

Both reflectance and the reflectance factor can be expressed as a function of three things: the incident solid angle, the reflected solid angle, and the BRDF. Depending on the situation, the type of solid angle used can be either differential, finite, or defined over the whole hemisphere. Figure A.5 demonstrates this point. The standard terms for these solid angles are *directional*, *conical*, and *hemispherical*[4]. Since ρ and R are dependent on both incident

[3]For a perfectly diffuse reflector, $f_r = \frac{1}{\pi}$, $\rho = 1$, and $R = 1$.

[4]In [6], Hanrahan uses the more descriptive terms *differential*, *finite*, and *hemispherical*.

Reflected	Incident		
	directional	conical	hemispherical
directional	bidirectional	conical-directional	hemispherical-directional
conical	directional-conical	biconical	hemispherical-conical
hemispherical	directional-hemispherical	conical-hemispherical	bihemispherical

Table A.2: Names of the nine types of reflectance

and reflected solid angles, there are nine different types of reflectance for each, all of which can be described in terms of the BRDF f_r. The names of these are given in Table A.2. The definitions of each of these reflectances and reflectance factors are presented several places in the literature [13, 20, 28], so for simplicity they are not reproduced here.

It is worth noting that the definition of $\frac{M}{E}$ in Equation 4.2 used to measure energy conservation is in fact directional-hemispherical reflectance. Recall that the integral used in the energy conservation measurement varies between zero and one, dependent on the energy reflected over the entire hemisphere due to energy from a single incident direction. A value of zero implies that no energy is reflected but is all absorbed, while a value of one implies that no energy is absorbed but is all reflected. This is precisely the definition of directional-hemispherical reflectance.

A.3 Further Reading

The presentation of the material in this appendix follows the work by Nicodemus et al. [28], as well as the standards proposed by the American National

Standards Institute and the Illuminating Engineering Society of North America [20]. Discussions of the principles of radiometry as they relate to computer graphics can be found in the book by Glassner [13] and the chapter in Cohen and Wallace's book by Hanrahan [6].

Appendix B

Glossary of Notation

In the few cases that a symbol has been used for more than one purpose, the definitions are separated by semicolons.

Table B.1: Glossary of Notation

Symbol	Definition
A	Matrix representing a tabulated BRDF;
	Patch of irradiated surface
$A^{(k)}$	The closest matrix of rank k to A
B	Radiosity (equivalent to radiant exitance M)
c	Speed of light in a vacuum $(299,792,458\ m/s)$
D	Diagonal matrix containing singular values of A
$D^{(k)}$	D with all but the k largest values zeroed
dA	Differential surface element
dA_i	Incident differential surface element

Glossary of Notation (continued)

Symbol	Definition
dA^{Φ}	dA projected in the direction of flux Φ
$d\omega_i$	Differential solid angle in direction of incident light
$d\omega_r$	Differential solid angle in direction of reflected light
$d\omega^N$	Differential solid angle projected in direction of N
E	Irradiance
f_r	Bidirectional reflectance distribution function (BRDF)
f'_r	Partially transformed BRDF (i.e. $\delta < 1$)
f_r^{ec}	Energy conserving BRDF
f_r^{iso}	Isotropic BRDF
f_r^{r}	Reciprocal BRDF
f_r^{in}	Incident portion of a separated BRDF
$f_r^{in(j)}$	The jth component of the incident portion of a k-separated BRDF ($1 \leq j \leq k$)
f_r^{out}	Reflected portion of a separated BRDF
$f_r^{out(j)}$	The jth component of the reflected portion of a k-separated BRDF ($1 \leq j \leq k$)
h	Planck's constant ($6.626176 \times 10^{-34} J \cdot s$)
I	Radiant intensity
L	Radiance
M	Radiant exitance (equivalent to radiosity B)

Glossary of Notation (continued)

Symbol	Definition
N	Vector normal to the surface
n_{f_r}	Total number of samples in a tabulated BRDF
n_{ϕ_i}	Number of ϕ_i samples in a tabulated BRDF
n_{θ_i}	Number of θ_i samples in a tabulated BRDF
n_{ϕ_r}	Number of ϕ_r samples in a tabulated BRDF
n_{θ_r}	Number of θ_r samples in a tabulated BRDF
P_{ec}	Measure of BRDF energy conservation
P_i	Measure of BRDF isotropy
P_r	Measure of BRDF reciprocity
P_s	Measure of BRDF separability
$P_s^{(k)}$	Measure of BRDF k-separability
p	Point of incident flux on a surface element (x_i, y_i);
	$\phi_i n_{\theta_i} + \theta_i$
Q	Radiant energy;
	Dimension of BRDF matrices A, U, V, and D
q	Point of reflected radiance from a surface element (x_r, y_r);
	$\phi_r n_{\theta_r} + \theta_r$
R	Reflectance factor
S	Bidirectional surface scattering reflectance distribution function
	(BSSRDF)

Glossary of Notation (continued)

Symbol	Definition
U	Matrix representing set of f_r^{in}'s
u	Vector representing f_r^{in}
V	Matrix representing set of f_r^{out}'s
v	Vector representing f_r^{out}
,	Energy conservation of a BRDF for a single incident direction
δ	BRDF transformation factor
Δ	Radian distance between samples
θ_i	Incident polar angle
θ_r	Reflected polar angle
λ	Wavelength of light (usually measured in nanometers)
μ_i	Mean BRDF value for a given $(\theta_i, \phi_+, \theta_r)$
ρ	Reflectance
σ_i	Standard deviation of BRDF values for a given $(\theta_i, \phi_+, \theta_r)$
τ	Maximum allowable exitance to incident irradiance ratio
Φ	Radiant flux (radiant power)
Φ_d	Radiant flux from a perfectly diffuse reflector
ϕ_i	Incident azimuth angle
ϕ_r	Reflected azimuth angle
ϕ_+	$\phi_r - \phi_i$
ξ	Radiant flux area density

Glossary of Notation (continued)

Symbol	Definition
Ω_N	Hemisphere in the direction of N
$\vec{\omega_i}$	Incident direction of light (equivalent to (ϕ_i, θ_i))
$\vec{\omega_r}$	Reflected direction of light (equivalent to (ϕ_r, θ_r))

Bibliography

[1] Howard Anton. *Calculus*. John Wiley and Sons, Inc., 1988.

[2] Petr Beckmann and Andre Spizzichino. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon Press, 1963.

[3] James F. Blinn. Models of light reflection for computer synthesized pictures. In *Computer Graphics (SIGGRAPH '77 proceedings)*, pages 192–198, 1977. San Jose, California: July 20–22, 1977.

[4] Bui-Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, June 1975.

[5] Brian Cabral, Nelson Max, and Rebecca Springmeyer. Bidirectional reflection functions from surface bump maps. In *Computer Graphics (SIGGRAPH '87 proceedings)*, pages 273–281, 1987. Anaheim, Calafornia: July 27–31, 1987.

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

[7] Robert L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics*, 1(1):7–24, January 1982.

[8] Joel DeYoung, Paul Lalonde, and Alain Fournier. Acquiring and using realistic reflectance data in computer graphics images. In *Arkansas Computer Conference*, pages 77–82, 1996. Searcy, Arkansas: March 7–8, 1996.

[9] James D. Foley, Andries van Dam, Steven K. Feiner, John F. Hughes, and Richard L. Phillips. *Introduction to Computer Graphics*. Addison-Wesley Publishing Company, 1994.

[10] David Forsey and Richard Bartels. Hierarchical b-spline refinement. In *Computer Graphics (SIGGRAPH '88 proceedings)*, pages 205–212, 1988. Atlanta, Georgia: August 1–5, 1988.

[11] Alain Fournier. Separating reflection functions for linear radiosity. In *Sixth EUROGRAPHICS Workshop on Rendering*, pages 383–392, 1995. Dublin, Ireland: June 12–14, 1995.

[12] Andrew Glassner, editor. *Graphics Gems*. Academic Press, 1990.

[13] Andrew Glassner. *Principles of Digital Image Synthesis*. Morgan Kauffman Publishers, 1995.

[14] Gene H. Golub and Charles F. van Loan. *Matrix Computations*. John Hopkins University Press, 1989.

[15] Jay S. Gondek, Gary W. Meyer, and Jonathan G. Newman. Wavelength dependent reflectance functions. In *Computer Graphics (SIGGRAPH '94 proceedings)*, pages 213–220, 1994. Orlando, Florida: July 24–29, 1994.

[16] Pat Hanrahan and Wolfgang Krueger. Reflection from layered surfaces due to subsurface scattering. In *Computer Graphics (SIGGRAPH '93 proceedings)*, pages 165–174, 1993. Anaheim, California: August 1–6, 1993.

[17] Xiao D. He, Patrick O. Heynen, Richard L. Phillips, Kenneth E. Torrance, David H. Salesin, and Donald P. Greenberg. A fast and accurate light reflection model. In *Computer Graphics (SIGGRAPH '92 proceedings)*, pages 253–254, 1992. Chicago, Illinois: July 26–31, 1992.

[18] Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. A comprehensive physical model for light reflection. In *Computer Graphics (SIGGRAPH '91 proceedings)*, pages 175–186, 1991. Las Vegas, Nevada: July 28–August 2, 1991.

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

[20] IES Nomenclature Committee. American national standard nomenclature and definitions for illuminating engineering. Technical Report ANSI/IES

RP-16-1986, American National Standards Institute and Illuminating Engineering Society of North America, 345 East 47th Street, New York, NY, 10017, USA, June 1986.

[21] Jim Kajiya. Anisotropic reflection models. In *Computer Graphics (SIGGRAPH '85 Proceedings)*, 1985. San Francisco, California: July 22–26, 1985.

[22] Paul Lalonde. *Representation and Uses of Light Distribution Functions*. PhD thesis, Department of Computer Science, University of British Columbia, 1996. Work in progress.

[23] Robert R. Lewis. Making shaders more physically plausible. In *Fourth EUROGRAPHICS Workshop on Rendering*, pages 47–61, 1993. Paris, France: June 14–16, 1993.

[24] Thomas M. MacRobert. *Spherical Harmonics*. Pergamon Press, 1967.

[25] Yves Meyer. *Les Ondelettes: Algorithmes et Applications*. Armand Colin, 1993.

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

[27] Sir Isaac Newton. *Philosophiae Naturalis Principia Mathematica (1686)*. Harvard University Press, 1972.

[28] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometrical considerations and nomenclature for reflectance. Technical Report 160, National Bureau of Standards, October 1977.

[29] Wilhelm Nusselt. Graphische Bestimmung des Winkelverhaltnisses bei der Warmestrahlung. *Verein deutscher Ingenieure Zeitschrift*, 72:673ff, 1928.

[30] Pierre Poulin and Alain Fournier. A model for anisotropic reflection. In *Computer Graphics (SIGGRAPH '90 proceedings)*, pages 273–281, 1990. Dallas, Texas: August 6–10, 1990.

[31] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing.* Cambridge University Press, 1990.

[32] Peter Schröder and Wim Sweldens. Spherical wavelets: efficiently representing functions on the sphere. In *Computer Graphics (SIGGRAPH '95 proceedings)*, pages 161–172, 1995. Los Angeles, California: August 6–11, 1995.

[33] François X. Sillion, James R. Arvo, Stephen H. Westin, and Donald P. Greenberg. A global illumination solution for general reflectance distributions. In *Computer Graphics (SIGGRAPH '91 proceedings)*, pages 187–196, 1991. Las Vegas, Nevada: July 28–August 2, 1991.

[34] Kenneth E. Torrance and Ephraim M. Sparrow. Polarization, directional distribution, and off-specular peak phenomena in light reflected from roughened surfaces. *Journal of the Optical Society of America*, 56(7):916–925, July 1966.

[35] Kenneth E. Torrance and Ephraim M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1114, September 1967.

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

[37] Gregory J. Ward. Measuring and modeling anisotropic reflection. In *Computer Graphics (SIGGRAPH '92 proceedings)*, pages 265–272, 1992. Chicago, Illinois: July 26–31, 1992.

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

[39] 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. Chicago, Illinois: July 26–31, 1992.

[40] Lawrence B. Wolff and David B. Kurlander. Ray tracing with polarization parameters. *IEEE Computer Graphics and Applications*, 10(6):44–55, November 1990.

[41] David Wong. Multiresolution surface construction for hierarchical B-splines. Master's thesis, Department of Computer Science, University of British Columbia, 1995.