A SIGNAL-PROCESSING FRAMEWORK FOR FORWARD AND INVERSE RENDERING

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Ravi Ramamoorthi
August 2002

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Pat Hanrahan
(Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Marc Levoy

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Jitendra Malik
(UC Berkeley)

Approved for the University Committee on Graduate Studies:

Abstract

The study of the computational aspects of reflection, and especially the interaction between reflection and illumination, is of fundamental importance in both computer graphics and vision. In computer graphics, the interaction between the incident illumination and the reflective properties of a surface is a basic building block in most *rendering* algorithms, i.e. methods that create synthetic computer-generated images. In computer vision, we often want to undo the effects of the reflection operator, i.e. to invert the interaction between the surface reflective properties and lighting. In other words, we often want to perform *inverse rendering*—the estimation of material and lighting properties from real photographs. Inverse rendering is also of increasing importance in graphics, where we wish to obtain accurate input illumination and reflectance models for (forward) rendering.

This dissertation describes a new way of looking at reflection on a curved surface, as a special type of convolution of the incident illumination and the reflective properties of the surface (technically, the bi-directional reflectance distribution function or BRDF). The first part of the dissertation is devoted to a theoretical analysis of the reflection operator, leading for the first time to a formalization of these ideas, with the derivation of a convolution theorem in terms of the spherical harmonic coefficients of the lighting and BRDF. This allows us to introduce a signal-processing framework for reflection, wherein the incident lighting is the signal, the BRDF is the filter, and the reflected light is obtained by filtering the input illumination (signal) using the frequency response of the BRDF filter.

The remainder of the dissertation describes applications of the signal-processing framework to forward and inverse rendering problems in computer graphics. First, we address the forward rendering problem, showing how our framework can be used for computing and displaying synthetic images in real-time with natural illumination and physically-based

BRDFs. Next, we extend and apply our framework to inverse rendering. We demonstrate estimation of realistic lighting and reflective properties from photographs, and show how this approach can be used to synthesize very realistic images under novel lighting and viewing conditions.

Acknowledgements

Firstly, I would like to thank my advisor, Pat Hanrahan, for convincing me to come to Stanford, for insightful discussions and technical advice during all these years on the whole spectrum of graphics topics, and for the perhaps more important discussions and advice on following appropriate scientific methodology, and development as a scientist and researcher. It has indeed been an honor and a privilege to work with him through these four years.

I would also like to thank my other committee members, Marc Levoy, Jitendra Malik, Ron Fedkiw, and Bernd Girod, for the advice, support, encouragement and inspiration they have provided over the years. I especially want to thank Marc for the many discussions we had in the early stages of this project.

Over the course of my time here, it has been exciting and energizing to work with and be around such an amazing group of people in the Stanford graphics lab. I want to thank Maneesh Agrawala, Sean Anderson, James Davis, Ziyad Hakura, Olaf Hall-Holt, Greg Humphreys, David Koller, Bill Mark, Steve Marschner, Matt Pharr, Kekoa Proudfoot, Szymon Rusinkiewicz, Li-Yi Wei and many others for being such wonderful friends, colleagues and co-workers. In particular, I want to thank Steve and Szymon for the many hours spent discussing ideas, their patience in reviewing drafts of my papers, and their help with data acquisition for the inverse rendering part of this dissertation.

Over the last four years, the Stanford graphics lab has really been the epicenter for graphics research. I can only consider it my privilege and luck to have been a part of this golden age.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The central problem in computer graphics is creating, or *rendering*, realistic computer-generated images that are indistinguishable from real photographs, a goal referred to as *photorealism*. Applications of *photorealistic rendering* extend from entertainment such as movies and games, to simulation, training and virtual reality applications, to visualization and computer-aided design and modeling. Progress towards photorealism in rendering involves two main aspects. First, one must develop an algorithm for physically accurate light transport simulation. However, the output from the algorithm is only as good as the inputs. Therefore, photorealistic rendering also requires accurate input models for object geometry, lighting and the reflective properties of surfaces.

Over the past two decades, there has been a signficant body of work in computer graphics on accurate light transport algorithms, with increasingly impressive results. One consequence has been the increasing realism of computer-generated special effects in movies, where it is often difficult or impossible to tell real from simulated.

However, in spite of this impressive progress, it is still far from routine to create photorealistic computer-generated images. Firstly, most light transport simulations are slow. Within the context of interactive imagery, such as with hardware rendering, it is very rare to find images rendered with natural illumination or physically accurate reflection functions. This gap between interactivity and realism impedes applications such as virtual lighting or material design for visualization and entertainment applications.

A second difficulty, which is often the limiting problem in realism today, is that of obtaining accurate input models for geometry, illumination and reflective properties. Entertainment applications such as movies often require very laborious fine-tuning of these parameters. One of the best ways of obtaining high-quality input illumination and material data is through measurements of scene attributes from real photographs. The idea of using real data in graphics is one that is beginning to permeate all aspects of the field. Within the context of animation, real motions are often measured using *motion capture* technology, and then retargetted to new characters or situations. For geometry, it is becoming more common to use *range scanning* to measure the 3D shapes of real objects for use in computer graphics. Similarly, measuring rendering attributes—lighting, textures and reflective properties—from real photographs is increasingly important. Since our goal is to invert the traditional rendering process, and estimate illumination and reflectance from real images, we refer to the approach as *inverse rendering*. Whether traditional or image-based rendering algorithms are used, rendered images now use measurements from real objects, and therefore appear very similar to real scenes.

In recent years, there has been significant interest in acquiring material models using inverse rendering. However, most previous work has been conducted in highly controlled lighting conditions, usually by careful active positioning of a single point source. Indeed, complex realistic lighting environments are rarely used in either forward or inverse rendering. While our primary motivation derives from computer graphics, many of the same ideas and observations apply to computer vision and perception. Within these fields, we often seek to perform *inverse rendering* in order to use intrinsic reflection and illumination parameters for modeling and recognition. Although it can be shown that the perception of shape and material properties may be qualitatively different under natural lighting conditions than artificial laboratory settings, most vision algorithms work only under very simple lighting assumptions—usually a single point light source without any shadowing.

It is our thesis that a deeper understanding of the computational nature of reflection and illumination helps to address these difficulties and restrictions in a number of areas in computer graphics and vision. This dissertation is about a new way of looking at reflection on a curved surface, as a special type of convolution of the incident illumination and the reflective properties of the surface. Although this idea has long been known qualitatively, this is the first time this notion of reflection as convolution has been formalized with an analytic convolution formula in the spherical domain. The dissertation includes both a theoretical analysis of reflection in terms of signal-processing, and practical applications of this frequency domain analysis to problems in forward and inverse rendering. In the rest of this chapter, we briefly discuss the main areas represented in the dissertation, summarizing our contributions and giving an outline of the rest of the dissertation. At the end of the chapter, table 1.1 summarizes the notation used in the rest of the dissertation.

1.1 Theoretical analysis of Reflection: Signal Processing

The computational study of the interaction of light with matter is a basic building block in rendering algorithms in computer graphics, as well as of interest in both computer vision and perception. In computer vision, previous theoretical work has mainly focussed on the problem of estimating shape from images, with relatively little work on estimating material properties or lighting. In computer graphics, the theory for forward global illumination calculations, involving all forms of indirect lighting, has been fairly well developed. The foundation for this analysis is the rendering equation [35]. However, there has been relatively little theoretical work on inverse problems or on the simpler reflection equation, which deals with the direct illumination incident on a surface.

It should be noted that there is a significant amount of qualitative knowledge regarding the properties of the reflection operator. For instance, we usually represent reflections from a diffuse surface at low resolutions [59] since the reflection from a matte, or technically *Lambertian*, surface *blurs* the illumination. Similarly, we realize that it is essentially impossible to accurately estimate the lighting from an image of a Lambertian surface; instead, we use mirror surfaces, i.e. gazing spheres.

In this dissertation, we formalize these observations and present a mathematical theory of reflection for general complex lighting environments and arbitrary reflection functions in terms of signal-processing. Specifically, we describe a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. Under these assumptions, we are able to derive an analytic formula for the reflected light field in terms of the spherical harmonic coefficients of the BRDF and the lighting. Our formulation leads to the following theoretical results:

Signal-Processing Framework for Reflection as Convolution: Chapter 2 derives our signal-processing framework for reflection. The reflected light field can therefore be thought of in a precise quantitative way as obtained by convolving the lighting and reflective properties of the surface (technically, the bi-directional reflectance distribution function or BRDF), i.e. by filtering the incident illumination using the BRDF. Mathematically, we are able to express the frequency-space coefficients of the reflected light field as a product of the spherical harmonic coefficients of the illumination and the BRDF. We believe this is a useful way of analyzing many forward and inverse problems. In particular, forward rendering can be viewed as *convolution* and inverse rendering as *deconvolution*.

Analytic frequency-space formulae for common lighting conditions and BRDFs Chapter 3 derives analytic formulae for the spherical harmonic coefficients of many common lighting and BRDF models. Besides being of practical interest, these formulae allow us to reason precisely about forward and inverse rendering in the frequency domain.

Well-posedness and Conditioning of Forward and Inverse Problems: Inverse problems can be ill-posed—there may be several solutions. They are also often numerically ill-conditioned, which may make devising practical algorithms infeasible. From our signal-processing framework, and the analytic formulae derived by us for common lighting and BRDF models, we are able to analyze the well-posedness and conditioning of a number of inverse problems. This analysis is presented in chapter 3, and explains many previous empirical observations, as well as serving as a guideline for future research in inverse rendering. We expect fruitful areas of research to be those problems that are well-conditioned. Additional assumptions will likely be required to address ill-conditioned or ill-posed inverse problems. This analysis is also of interest for forward rendering. An ill-conditioned inverse problem corresponds to a forward problem where the final results are not sensitive to certain components of the initial conditions. This often allows us to approximate the initial conditions—such as the incident illumination—in a principled way, giving rise to more efficient forward rendering algorithms.

Figure 1.1: *A scene rendered in real time using our approach, described in chapter 5. The illumination is measured light in the Grace Cathedral in San Francisco, obtained by photographing a mirror sphere, courtesy of Paul Debevec. The surface reflective properties include a number of measured BRDFs.*

1.2 Forward Rendering

Lighting in most real scenes is complex, coming from a variety of sources including area lights and large continuous lighting distributions like skylight. But current graphics hardware only supports point or directional light sources, and very simple unphysical surface reflection models. One reason is the lack of simple procedural formulas for efficiently computing the reflected light from general lighting distributions. Instead, an integration over the visible hemisphere must be done for each pixel in the final image.

Chapters 4 and 5 apply the signal-processing framework to interactive rendering with natural illumination and physically-based reflection functions. An example image is shown in figure 1.1. That image includes natural illumination and a number of physically-based or measured surface reflective properties.

As is common with interactive rendering, we neglect the effects of global effects like cast shadows (self-shadowing) and interreflections, and restrict ourselves to distant illumination. Since the illumination corresponds to captured light in an environment, such rendering methods are frequently referred to as *environment mapping*.

Chapter 4 demonstrates how our signal-processing framework can be applied to *irradiance environment maps*, corresponding to the reflection from perfectly diffuse or Lambertian surfaces. The key to our approach is the rapid computation of an analytic approximation to the irradiance environment map. For rendering, we demonstrate a simple procedural algorithm that runs at interactive frame rates, and is amenable to hardware implementation. Our novel representation also suggests new approaches to lighting design and image-based rendering.

Chapter 5 introduces a new paradigm for prefiltering and rendering environment mapped images with general isotropic BRDFs. Our approach uses spherical frequency domain methods, based on the earlier theoretical derivations. Our method has many advantages over the angular (spatial) domain approaches:

Theoretical analysis of sampling rates and resolutions: Most previous work has determined reflection map resolutions, or the number of reflection maps required, in an ad-hoc manner. By using our signal-processing framework, we are able to perform error analysis, that allows us to set sampling rates and resolutions accurately.

Efficient representation and rendering with Spherical Harmonic Reflection Maps: We introduce *spherical harmonic reflection maps (SHRMs)* as a compact representation. Instead of a single color, each pixel stores coefficients of a spherical harmonic expansion encoding view-dependence of the reflection map. An important observation that emerges from the theoretical analysis is that for almost all BRDFs, a very low order spherical harmonic expansion suffices. Thus, SHRMs can be evaluated in real-time for rendering. Further, they are significantly more compact and accurate than previous methods.

Fast precomputation (prefiltering): One of the drawbacks of current environment mapping techniques is the significant computational time required for prefiltering, or precomputing reflection maps, which can run into hours, and preclude the use of these approaches

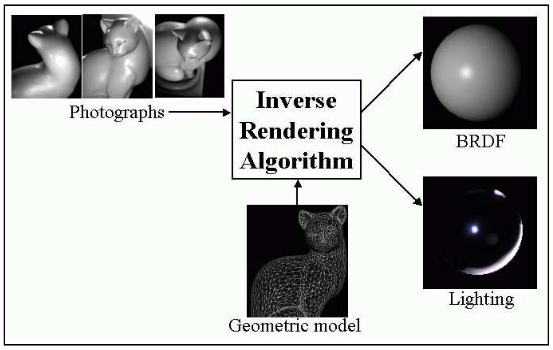

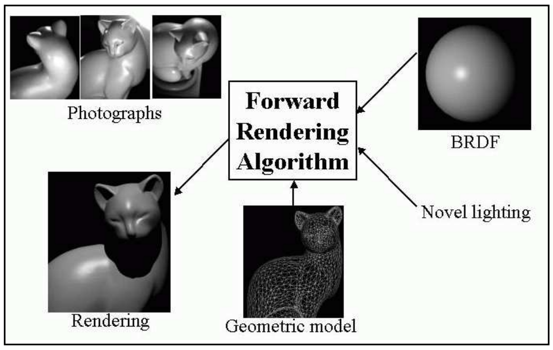

Figure 1.2: *Inverse Rendering. On the left, we show a conceptual description, while the right shows how illumination may be manipulated to create realistic new synthetic images.*

in applications involving lighting and material design, or dynamic lighting. We introduce new prefiltering methods based on spherical harmonic transforms, and show both empirically, and analytically by computational complexity analysis, that our algorithms are orders of magnitude faster than previous work.

1.3 Inverse Rendering

Finally, chapter 6 discusses practical inverse rendering methods for estimating illumination and reflective properties from sequences of images. The general idea in inverse rendering is illustrated in figure 1.2. On the left, we show a conceptual description. The inputs are photographs of the object of interest, and a geometric model. The outputs are the reflective properties, visualized by rendering a sphere with the same BRDF, and the illumination, visualized by showing an image of a chrome-steel or mirror sphere. On the right, we show how new renderings can be created under novel lighting conditions using the measured BRDF from the real object. We simply need to use a standard physically based rendering algorithm, with the input BRDF model that obtained by inverse rendering.

There are a number of applications of this approach. Firstly, it allows us to take input photographs and relight the scene, or include real objects in synthetic scenes with novel illumination and viewing conditions. This has applications in developing new methods for image-processing, and in entertainment and virtual reality applications, besides the usefulness in creating realistic computer-generated images.

The dissertation uses the theoretical framework described earlier, and the formal analysis of the conditioning of inverse problems, to derive new robust practical algorithms for inverse rendering. Our specific practical contributions are:

Complex Illumination: As stated in the introduction, one of our primary motivations is to perform inverse rendering under complex illumination, allowing these methods to be used in arbitrary indoor and outdoor settings. We present a number of improved algorithms for inverse rendering under complex lighting conditions.

New Practical Representations And Algorithms: The theoretical analysis is used to develop a new low-parameter practical representation that simultaneously uses the spatial or angular domain and the frequency domain. Using this representation, we develop a number of new inverse rendering algorithms that use both the spatial and the frequency domain.

Simultaneous Determination Of Lighting And BRDF: In many practical applications, it might be useful to determine reflective properties without knowing the lighting, or to determine both simultanously. We present the first method to determine BRDF parameters of complex geometric models under unknown illumination, simultaneously finding the lighting.

$\vec{L}, \vec{N}, \vec{V}, \vec{R}$	Global incident, normal, viewing, reflected directions
B	Reflected radiance
B_{lp}	Coefficients of Fourier expansion of B in 2D
B_{lmpq}, B_{lmnpq}	Coefficients of isotropic, anisotropic basis-function expansion of B in 3D
L	Incoming radiance
L_l	Coefficients of Fourier expansion of L in 2D
L_{lm}	Coefficients of spherical-harmonic expansion of L in 3D
ρ	Surface BRDF
$\hat{\rho}$	BRDF multiplied by cosine of incident angle
$\hat{\rho}_{lp}$	Coefficients of Fourier expansion of $\hat{\rho}$ in 2D
$\hat{\rho}_{lpq}, \hat{\rho}_{ln;pq}$	Coefficients of isotropic, anisotropic spherical-harmonic expansion of $\hat{\rho}$
θ_i', θ_i	Incident elevation angle in *local*, *global* coordinates
ϕ_i', ϕ_i	Incident azimuthal angle in *local*, *global* coordinates
θ_o', θ_o	Outgoing elevation angle in *local*, *global* coordinates
ϕ_o', ϕ_o	Outgoing azimuthal angle in *local*, *global* coordinates
$\vec{X}$	Surface position
α	Surface normal parameterization—elevation angle
β	Surface normal parameterization—azimuthal angle
γ	Orientation of tangent frame for anisotropic surfaces
R_α	Rotation operator for surface orientation α in 2D
$R_{\alpha,\beta}, R_{\alpha,\beta,\gamma}$	Rotation operator for surface normal (α, β) or tangent frame (α, β, γ) in 3D
$F_k(\theta)$	Fourier basis function (complex exponential)
$F_k^*(\theta)$	Complex Conjugate of Fourier basis function
$Y_{lm}(\theta, \phi)$	Spherical Harmonic
$Y_{lm}^*(\theta, \phi)$	Complex Conjugate of Spherical Harmonic
$f_{lm}(\theta)$	Normalized θ dependence of Y_{lm}
$D_{mn}^l(\alpha, \beta, \gamma)$	Representation matrix of dimension $2l+1$ for rotation group $SO(3)$
$d_{mn}^l(\alpha)$	$D_{mm'}^L$ for y-axis rotations
Λ_l	Normalization constant, $\sqrt{4\pi/(2l+1)}$
I	$\sqrt{-1}$

Table 1.1: *Common notation used throughout the dissertation.*

Chapter 2

Reflection as Convolution

In the introduction to this dissertation, we have discussed a number of problems in computer graphics and vision, where having a deeper theoretical understanding of the reflection operator is important. These include *inverse rendering* problems—determining lighting distributions and bidirectional reflectance distribution functions (BRDFs) from photographs—and *forward rendering* problems such as rendering with environment maps, and image-based rendering.

In computer graphics, the theory for forward global illumination calculations has been fairly well developed, based on Kajiya's rendering equation [35]. However, very little work has gone into addressing the theory of inverse problems, or on studying the theoretical properties of the simpler reflection equation, which deals with the direct illumination incident on a surface. We believe that this lack of a formal mathematical understanding of the properties of the reflection equation is one of the reasons why complex, realistic lighting environments and reflection functions are rarely used either in forward or inverse rendering.

While a formal theoretical basis has hitherto been lacking, reflection is of deep interest in both graphics and vision, and there is a significant body of qualitative and empirical information available. For instance, in their seminal 1984 work on environment map prefiltering and rendering, Miller and Hoffman [59] qualitatively observed that reflection was a convolution of the incident illumination and the reflective properties (BRDF) of the surface. Subsequently, similar qualitative observations have been made by Cabral et al. [7, 8], D'Zmura [17] and others. However, in spite of the long history of these observations, this notion has never previously been formalized.

Another interesting observation concerns the "blurring" that occurs in the reflection from a Lambertian (diffuse) surface. Miller and Hoffman [59] used this idea to represent irradiance maps, proportional to the reflected light from a Lambertian surface, at low resolutions. However, the precise resolution necessary was never formally determined. More recently, within the context of lighting-invariant object recognition, a number of computer vision researchers [18, 30, 91] have observed empirically that the space of images of a diffuse object under all possible lighting conditions is very low-dimensional. Intuitively, we do not see the high frequencies of the environment in reflections from a diffuse surface. However, the precise nature of what is happening computationally has not previously been understood.

The goal of this chapter is to formalize these observations and present a mathematical theory of reflection for general complex lighting environments and arbitrary BRDFs. Specifically, we describe a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. Under these assumptions, we are able to derive an analytic formula for the reflected light field in terms of the spherical harmonic coefficients of the BRDF and the lighting. The reflected light field can therefore be thought of in a precise quantitative way as obtained by convolving the lighting and BRDF, i.e. by filtering the incident illumination using the BRDF. Mathematically, we are able to express the frequency-space coefficients of the reflected light field as a product of the spherical harmonic coefficients of the illumination and the BRDF.

We believe this is a useful way of analyzing many forward and inverse problems. In particular, forward rendering can be viewed as *convolution* and inverse rendering as *deconvolution*. Furthermore, in the next chapter, we are able to derive analytic formulae for the spherical harmonic coefficients of many common BRDF and lighting models. From this formal analysis, we are able to determine precise conditions under which estimation of BRDFs and lighting distributions are well posed and well-conditioned. This analysis also has implications for *forward rendering*—especially the efficient rendering of objects under complex lighting conditions specified by environment maps.

The goal of this and the following chapter are to present a unified, detailed and complete description of the mathematical foundation underlying the rest of the dissertation. We will briefly point out the practical implications of the results derived in this chapter, but will refer the reader to later in the dissertation for implementation details. The rest of this chapter is organized as follows. In section 1, we discuss previous work. Section 2 introduces the reflection equation in 2D and 3D, showing how it can be viewed as a convolution. Section 3 carries out a formal frequency-space analysis of the reflection equation, deriving the frequency space convolution formulae. Section 4 briefly discusses general implications for forward and inverse rendering. Finally, section 5 concludes this chapter and discusses future theoretical work. The next chapter will derive analytic formulae for the spherical harmonic coefficients of many common lighting and BRDF models, applying the results to theoretically analyzing the well-posedness and conditioning of many problems in inverse rendering.

2.1 Previous Work

In this section, we briefly discuss previous work. Since the reflection operator is of fundamental interest in a number of fields, the relevant previous work is fairly diverse. We start out by considering rendering with environment maps, where there is a long history of regarding reflection as a convolution, although this idea has not previously been mathematically formalized. We then describe some relevant work in inverse rendering, one of the main applications of our theory. Finally, we discuss frequency-space methods for reflection, and previous work on a formal theoretical analysis.

Forward Rendering by Environment Mapping: The theoretical analysis in this paper employs essentially the same assumptions typically made in rendering with environment maps, i.e. distant illumination—allowing the lighting to be represented by a single environment map—incident on curved surfaces. Blinn and Newell [5] first used environment maps to efficiently find reflections of distant objects. The technique was generalized by Miller and Hoffman [59] and Greene [22] who precomputed diffuse and specular reflection maps, allowing for images with complex realistic lighting and a combination of Lambertian and Phong BRDFs to be synthesized. Cabral et al. [7] later extended this general method to computing reflections from bump-mapped surfaces, and to computing environment-mapped images with arbitrary BRDFs [8]. It should be noted that both Miller and Hoffman [59], and Cabral et al. [7, 8] qualitatively described the reflection maps as obtained by convolving the lighting with the BRDF. In this paper, we will formalize these ideas, making the notion of convolution precise, and derive analytic formulae.

Inverse Rendering: We now turn our attention to the inverse problem—estimating BRDF and lighting properties from photographs. Inverse rendering is one of the main practical applications of, and original motivation for, our theoretical analysis. Besides being of fundamental interest in computer vision, inverse rendering is important in computer graphics since the realism of images is nowadays often limited by the quality of input models. Inverse rendering yields the promise of providing very accurate input models since these come from measurements of real photographs.

Perhaps the simplest inverse rendering method is the use of a mirror sphere to find the lighting, first introduced by Miller and Hoffman [59]. A more sophisticated *inverse lighting* approach is that of Marschner and Greenberg [54], who try to find the lighting under the assumption of a Lambertian BRDF. D'Zmura [17] proposes estimating spherical harmonic coefficients of the lighting.

Most work in inverse rendering has focused on BRDF [62] estimation. Recently, image-based BRDF measurement methods have been proposed in 2D by Lu et al. [51] and in 3D by Marschner et al. [55]. If the entire BRDF is measured, it may be represented by tabulating its values. An alternative representation is by low-parameter models such as those of Ward [85] or Torrance and Sparrow [84]. Parametric models are often preferred in practice since they are compact, and are simpler to estimate. A number of methods [14, 15, 77, 89] have been proposed to estimate parametric BRDF models, often along with a modulating texture.

However, it should be noted that all of the methods described above use a single point source. One of the main goals of the theoretical analysis in this paper is to enable the use of inverse rendering with complex lighting. Recently, there has been some work in this area [16, 50, 64, 75, 76, 90], although many of those methods are specific to a particular illumination model. Using the theoretical analysis described in this paper, we [73] have presented a general method for complex illumination, that handles the various components of the lighting and BRDF in a principled manner to allow for BRDF estimation under general lighting conditions. Furthermore, we will show that it is possible in theory to separately estimate the lighting and BRDF, up to a global scale factor. We have been able to use these ideas to develop a practical method [73] of *factoring* the light field to simultaneously determine the lighting and BRDF for geometrically complex objects.

Frequency-Space Representations: Since we are going to treat reflection as a convolution and analyze it in frequency-space, we will briefly discuss previous work on frequency-space representations. Since we will be primarily concerned with analyzing quantities like the BRDF and distant lighting which can be parameterized as a function on the unit sphere, the appropriate frequency-space representations are spherical harmonics [32, 34, 52]. The use of spherical harmonics to represent the illumination and BRDF was pioneered by Cabral et al. [7]. D'Zmura [17] analyzed reflection as a linear operator in terms of spherical harmonics, and discussed some resulting ambiguities between reflectance and illumination. We extend his work by explicitly deriving the frequency-space reflection equation (i.e. convolution formula) in this chapter, and by providing quantitative results for various special cases in the next chapter. Our use of spherical harmonics to represent the lighting is similar in some respects to previous methods such as that of Nimeroff et al. [63] that use steerable linear basis functions. Spherical harmonics, as well as the closely related Zernike polynomials, have also been used before in computer graphics for representing BRDFs by a number of other authors [43, 79, 86].

Formal Analysis of Reflection: This paper conducts a formal study of the reflection operator by showing mathematically that it can be described as a convolution, deriving an analytic formula for the resulting convolution equation, and using this result to study the well-posedness and conditioning of several inverse problems. As such, our approach is similar in spirit to mathematical methods used to study inverse problems in other areas of radiative transfer and transport theory such as hydrologic optics [67] and neutron scattering. See McCormick [58] for a review.

Within computer graphics and vision, the closest previous theoretical work lies in the object recognition community, where there has been a significant amount of interest in characterizing the appearance of a surface under all possible illumination conditions, usually under the assumption of Lambertian reflection. For instance, Belhumeur and Kriegman [4] have theoretically described this set of images in terms of an illumination cone, while empirical results have been obtained by Hallinan [30] and Epstein et al. [18]. These results suggest that the space spanned by images of a Lambertian object under all (distant) illumination conditions lies very close to a low-dimensional subspace. We will see that our theoretical analysis will help in explaining these observations, and in extending the predictions to arbitrary reflectance models. In independent work on face recognition, simultaneous with our own, Basri and Jacobs [2] have described Lambertian reflection as a convolution and obtained similar analytic results for that particular case.

This chapter builds on previous theoretical work by us on analyzing planar or flatland light fields [70], on the reflected light field from a Lambertian surface [72], and on the theory for the general 3D case with isotropic BRDFs [73]. The goal of this chapter is to present a unified, complete and detailed account of the theory in the general case. We describe a unified view of the 2D and 3D cases, including general anisotropic BRDFs, a group-theoretic interpretation in terms of generalized convolutions, and the relationship to the theory of Fredholm integral equations of the first kind, which have not been discussed in our earlier papers.

2.2 Reflection Equation

In this section, we introduce the mathematical and physical preliminaries, and derive a version of the reflection equation. In order to derive our analytic formulae, we must analyze the properties of the reflected light field. The light field [20] is a fundamental quantity in light transport and therefore has wide applicability for both forward and inverse problems in a number of fields. A good introduction to the various radiometric quantities derived from light fields is provided by McCluney [56], while Cohen and Wallace [12] introduce many of the terms discussed here with motivation from a graphics perspective. Light fields have been used directly for rendering images from photographs in computer graphics, without considering the underlying geometry [21, 48], or by parameterizing the light field on the object surface [88].

After a discussion of the physical assumptions made, we first introduce the reflection equation for the simpler *flatland* or 2D case, and then generalize the results to 3D. In the next section, we will analyze the reflection equation in frequency-space.

2.2.1 Assumptions

We will assume curved convex homogeneous reflectors in a distant illumination field. Below, we detail each of the assumptions.

Curved Surfaces: We will be concerned with the reflection of a distant illumination field by curved surfaces. Specifically, we are interested in the variation of the reflected light field as a function of surface orientation and exitant direction. Our goal is to analyze this variation in terms of the incident illumination and the surface BRDF. Our theory will be based on the fact that different orientations of a curved surface correspond to different orientations of the upper hemisphere and BRDF. Equivalently, each orientation of the surface corresponds to a different integral over the lighting, and the reflected light field will therefore be a function of surface orientation.

Convex Objects: The assumption of convexity ensures there is no shadowing or interreflection. Therefore, the incident illumination is only because of the distant illumination field. Convexity also allows us to parameterize the object simply by the surface orientation. For isotropic surfaces, the surface orientation is specified uniquely by the normal vector. For anisotropic surfaces, we must also specify the direction of anisotropy, i.e. the orientation of the local tangent frame.

It should be noted that our theory can also be applied to concave objects, simply by using the surface normal (and the local tangent frame for anisotropic surfaces). However, the effects of self-shadowing (cast shadows) and interreflections will not be considered.

Homogeneous Surfaces: We assume untextured surfaces with the same BRDF everywhere.

Distant Illumination: The illumination field will be assumed to be generated by distant sources, allowing us to use the same lighting function anywhere on the object surface. The lighting can therefore be represented by a single environment map indexed by the incident angle.

Discussion: We note that for the most part, our assumptions are very similar to those made in most interactive graphics applications, including environment map rendering algorithms such as those of Miller and Hoffman [59] and Cabral et al. [8]. Our assumptions also accord closely with those usually made in computer vision and inverse rendering. The only significant additional assumption is that of homogeneous surfaces. However, this is not particularly restrictive since spatially varying BRDFs are often approximated in practical graphics or vision applications by using a spatially varying texture that simply modulates one or more components of the BRDF. This can be incorporated into the ensuing theoretical analysis by merely multiplying the reflected light field by a texture dependent on surface position. We believe that our assumptions are a good approximation to many real-world situations, while being simple enough to treat analytically. Furthermore, it is likely that the insights obtained from the analysis in this paper will be applicable even in cases where the assumptions are not exactly satisfied. We will demonstrate in chapter 6 that in practical applications, it is possible to extend methods derived from these assumptions to be applicable in an even more general context.

We now proceed to derive the reflection equation for the 2D and 3D case under the assumptions outlined above. Notation used in chapter 2, and reused throughout the dissertation, is listed in table 1.1. We will use two types of coordinates. Unprimed global coordinates denote angles with respect to a global reference frame. On the other hand, primed local coordinates denote angles with respect to the local reference frame, defined by the local surface normal and a tangent vector. These two coordinate systems are related simply by a rotation.

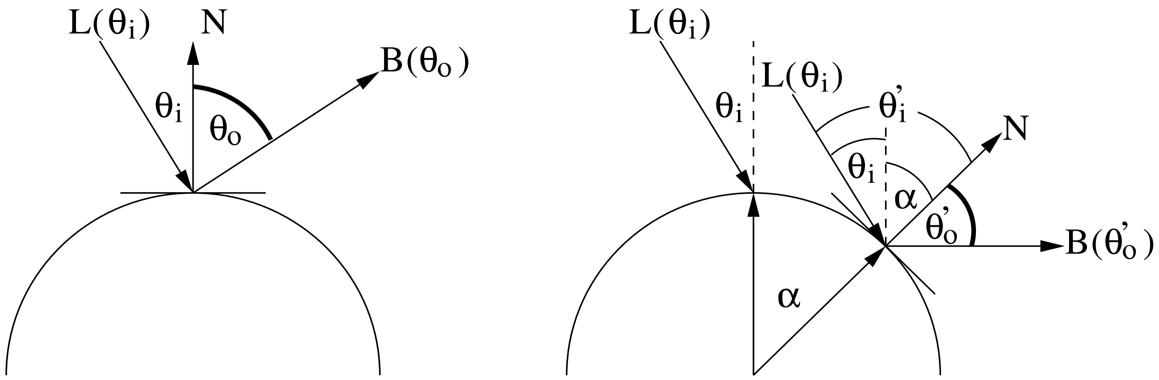

Figure 2.1: *Schematic of reflection in 2D. On the left, we show the situation with respect to one point on the surface (the north pole or* $0°$ *location, where global and local coordinates are the same). The right figure shows the affect of the surface orientation* α*. Different orientations of the surface correspond to rotations of the upper hemisphere and BRDF, with the global incident direction* θ_i *corresponding to a rotation by* α *of the local incident direction* θ_i'*. Note that we also keep the local outgoing angle (between* N *and* B*) fixed between the two figures*

2.2.2 Flatland 2D case

In this subsection, we consider the *flatland* or 2D case, assuming that all measurements and illumination are restricted to a single plane. Considering the 2D case allows us to explain the key concepts clearly, and show how they generalize to 3D. A diagram illustrating the key concepts for the planar case is in figure 2.1.

In local coordinates, we can write the reflection equation as

$$B(\vec{X}, \theta_o') = \int_{-\pi/2}^{\pi/2} L(\vec{X}, \theta_i')\rho(\theta_i', \theta_o') \cos\theta_i' \, d\theta_i'. \tag{2.1}$$

Here, B is the reflected radiance, L is the incident radiance, i.e illumination, and ρ is the BRDF or bi-directional reflectance distribution function of the surface, which in 2D is a function of the local incident and outgoing angles (θ_i', θ_o'). The limits of integration correspond to the *visible half-circle*—the 2D analogue of the upper hemisphere in 3D.

We now make a number of substitutions in equation 2.1, based on our assumptions. First, consider the assumption of a convex surface. This ensures there is no shadowing or interreflection; this fact has implicitly been assumed in equation 2.1. The reflected radiance therefore depends only on the distant illumination field L and the surface BRDF ρ. Next, consider the assumption of distant illumination. This implies that the reflected light field depends directly only on the surface orientation, as described by the surface normal $\vec{N}$, and does not directly depend on the position $\vec{X}$. We may therefore reparameterize the surface by its angular coordinates α, with $\vec{N} = [\sin\alpha, \cos\alpha]$, i.e. $B(\vec{X}, \theta'_o) \rightarrow B(\alpha, \theta'_o)$ and $L(\vec{X}, \theta'_i) \rightarrow L(\alpha, \theta'_i)$. The assumption of distant sources also allows us to represent the incident illumination by a single environment map for all surface positions, i.e. use a single function L regardless of surface position. In other words, the lighting is a function only of the global incident angle, $L(\alpha, \theta'_i) \rightarrow L(\theta_i)$. Finally, we define a transfer function $\hat{\rho} = \rho\cos\theta'_i$ to absorb the cosine term in the integrand. With these modifications, equation 2.1 becomes

$$B(\alpha, \theta'_o) = \int_{-\pi/2}^{\pi/2} L(\theta_i)\hat{\rho}(\theta'_i, \theta'_o)\, d\theta'_i. \tag{2.2}$$

It is important to note that in equation 2.2, we have mixed local (primed) and global (unprimed) coordinates. The lighting is a global function, and is naturally expressed in a global coordinate frame as a function of global angles. On the other hand, the BRDF is naturally expressed as a function of the local incident and reflected angles. When expressed in the local coordinate frame, the BRDF is the same everywhere for a homogeneous surface. Similarly, when expressed in the global coordinate frame, the lighting is the same everywhere, under the assumption of distant illumination. Integration can be conveniently done over either local or global coordinates, but the upper hemisphere is easier to keep track of in local coordinates.

Rotations—Converting between Local and Global coordinates: To do the integral in equation 2.2, we must relate local and global coordinates. One can convert between these by applying a rotation corresponding to the local surface normal α. The *up-vector* in local coordinates, i.e $0'$ is the surface normal. The corresponding global coordinates are clearly α. We define R_α as an operator that rotates θ'_i into global coordinates, and is given in 2D simply by $R_\alpha(\theta'_i) = \alpha + \theta'_i$. To convert from global to local coordinates, we apply the inverse rotation, i.e. $R_{-\alpha}$. To summarize,

$$\begin{aligned} \theta_i &= R_\alpha(\theta'_i) &&= \alpha + \theta'_i \\ \theta'_i &= R_\alpha^{-1}(\theta_i) &&= -\alpha + \theta_i. \end{aligned} \tag{2.3}$$

It should be noted that the signs of the various quantities are taken into account in equation 2.3. Specifically, from the right of figure 2.1, it is clear that $\mid \theta_i' \mid = \mid \theta_i \mid + \mid \alpha \mid$. In our sign convention, α is positive in figure 2.1, while θ_i' and θ_i are negative. Substituting $\mid \theta_i' \mid = -\theta_i'$ and $\mid \theta_i \mid = -\theta_i$, we verify equation 2.3.

With the help of equation 2.3, we can express the incident angle dependence of equation 2.2 in either local coordinates entirely, or global coordinates entirely. It should be noted that we always leave the outgoing angular dependence of the reflected light field in local coordinates in order to match the BRDF transfer function.

$$B(\alpha, \theta_o') = \int_{-\pi/2}^{\pi/2} L\left(R_\alpha(\theta_i')\right) \hat{\rho}\left(\theta_i', \theta_o'\right) \, d\theta_i' \quad (2.4)$$

$$= \int_{-\pi/2+\alpha}^{\pi/2+\alpha} L\left(\theta_i\right) \hat{\rho}\left(R_\alpha^{-1}(\theta_i), \theta_o'\right) \, d\theta_i. \quad (2.5)$$

By plugging in the appropriate relations for the rotation operator from equation 2.3, we can obtain

$$B(\alpha, \theta_o') = \int_{-\pi/2}^{\pi/2} L\left(\alpha + \theta_i'\right) \hat{\rho}\left(\theta_i', \theta_o'\right) \, d\theta_i' \quad (2.6)$$

$$= \int_{-\pi/2+\alpha}^{\pi/2+\alpha} L\left(\theta_i\right) \hat{\rho}\left(-\alpha + \theta_i, \theta_o'\right) \, d\theta_i. \quad (2.7)$$

Interpretation as Convolution: Equations 2.6 and 2.7 (and the equivalent forms in equations 2.4 and 2.5) are *convolutions*. The reflected light field can therefore be described formally as a convolution of the incident illumination and the BRDF transfer function. Equation 2.5 in global coordinates states that the reflected light field at a given surface orientation corresponds to *rotating* the BRDF to that orientation, and then integrating over the upper half-circle. In signal processing terms, the BRDF can be thought of as the filter, while the lighting is the input signal. The reflected light field is obtained by filtering the input signal (i.e. lighting) using the filter derived from the BRDF. Symmetrically, equation 2.4 in local coordinates states that the reflected light field at a given surface orientation may be computed by *rotating* the lighting into the local coordinate system of the BRDF, and then doing the integration over the upper half-circle.

It is important to note that we are fundamentally dealing with rotations, as is brought out by equations 2.4 and 2.5. For the 2D case, rotations are equivalent to translations, and equations 2.6 and 2.7 are the familiar equations for translational convolution. The main difficulty in formally generalizing the convolution interpretation to 3D is that the structure of rotations is more complex. In fact, we will need to consider a generalization of the notion of convolution in order to encompass rotational convolutions.

2.2.3 Generalization to 3D

The flatland development can be extended to 3D. In 3D, we can write the reflection equation, analogous to equation 2.1, as

$$B(\vec{X}, \theta'_o, \phi'_o) = \int_{\Omega'_i} L(\vec{X}, \theta'_i, \phi'_i) \rho(\theta'_i, \phi'_i, \theta'_o, \phi'_o) \cos\theta'_i \, d\omega'_i. \tag{2.8}$$

Note that the integral is now over the 3D upper hemisphere, instead of the 2D half-circle. Also note that we must now also consider the (local) azimuthal angles ϕ'_i and ϕ'_o.

We can make the same substitutions that we did in 2D. We reparameterize the surface position $\vec{X}$ by its angular coordinates (α, β, γ). Here, the surface normal $\vec{N}$ is given by the standard formula $\vec{N} = [\sin\alpha\cos\beta, \sin\alpha\sin\beta, \cos\alpha]$. The third angular parameter γ is important for anisotropic surfaces and controls the rotation of the local tangent-frame about the surface normal. For isotropic surfaces, γ has no physical significance. Figure 2.2 illustrates the rotations corresponding to (α, β, γ). We may think of them as essentially corresponding to the standard Euler-angle rotations about Z, Y and Z by angles α,β and γ. As in 2D, we may now make the substitutions, $B(\vec{X}, \theta'_o, \phi'_o) \to B(\alpha, \beta, \gamma, \theta'_o, \phi'_o)$ and $L(\vec{X}, \theta'_i, \phi'_i) \to L(\theta_i, \phi_i)$, and define a transfer function to absorb the cosine term, $\hat{\rho} = \rho\cos\theta'_i$. We now obtain the 3D equivalent of equation 2.2,

$$B(\alpha, \beta, \gamma, \theta'_o, \phi'_o) = \int_{\Omega'_i} L(\theta_i, \phi_i) \hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) \, d\omega'_i. \tag{2.9}$$

Rotations—Converting between Local and Global coordinates: To do the integral above, we need to apply a rotation to convert between local and global coordinates, just as in 2D. The rotation operator is substantially more complicated in 3D, but the operations

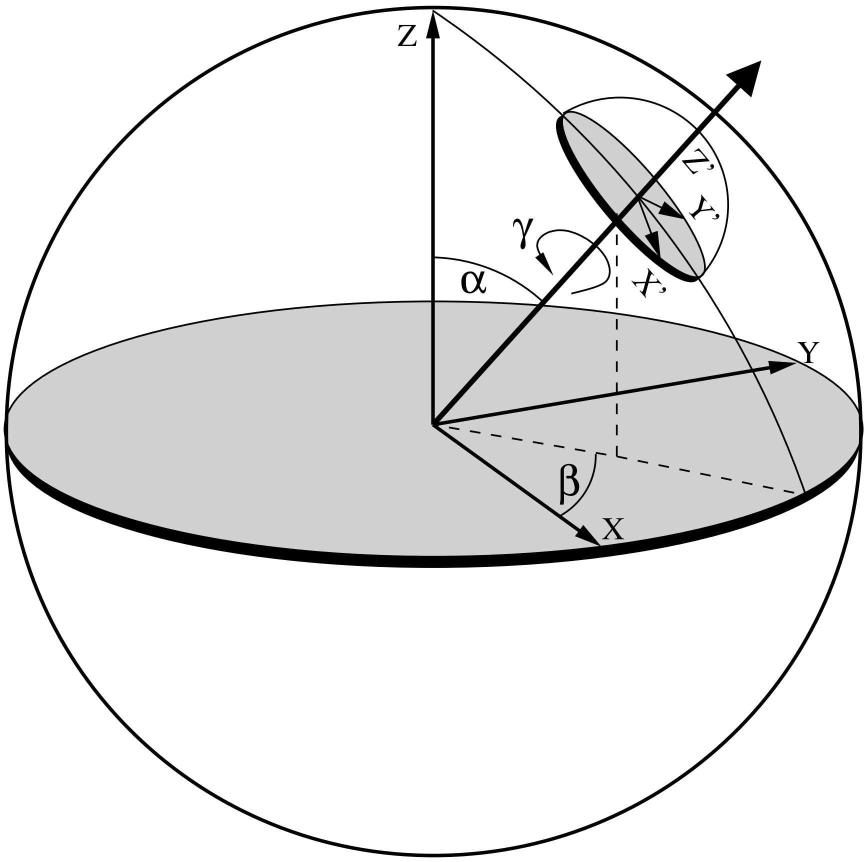

Figure 2.2: *Diagram showing how the rotation corresponding to (α, β, γ) transforms between local (primed) and global (unprimed) coordinates. The net rotation is composed of three independent rotations about Z,Y,and Z, with the angles α, β, and γ corresponding directly to the Euler angles.*

are conceptually very similar to those in flatland. The *north pole* $(0', 0')$ or +Z axis in local coordinates is the surface normal, and the corresponding global coordinates are (α, β). It can be verified that a rotation of the form $R_z(\beta)R_y(\alpha)$ correctly performs this transformation, where the subscript z denotes rotation about the Z axis and the subscript y denotes rotation about the Y axis. For full generality, the rotation between local and global coordinates should also specify the transformation of the local tangent frame, so the general rotation operator is given by $R_{\alpha,\beta,\gamma} = R_z(\beta)R_y(\alpha)R_z(\gamma)$. This is essentially the Euler-angle representation of rotations in 3D. We may now summarize these results, obtaining the 3D equivalent of equation 2.3,

$$\begin{aligned}(\theta_i, \phi_i) &= R_{\alpha,\beta,\gamma}(\theta_i', \phi_i') &&= R_z(\beta)R_y(\alpha)R_z(\gamma)\left\{\theta_i', \phi_i'\right\} \\ (\theta_i', \phi_i') &= R^{-1}_{\alpha,\beta,\gamma}(\theta_i, \phi_i) &&= R_z(-\gamma)R_y(-\alpha)R_z(-\beta)\left\{\theta_i, \phi_i\right\}.\end{aligned} \tag{2.10}$$

It is now straightforward to substitute these results into equation 2.9, transforming the integral either entirely into local coordinates or entirely into global coordinates, and obtaining the 3D analogue of equations 2.4 and 2.5,

$$B(\alpha, \beta, \gamma, \theta_o', \phi_o') = \int_{\Omega_i'} L\left(R_{\alpha,\beta,\gamma}(\theta_i', \phi_i')\right)\hat{\rho}(\theta_i', \phi_i', \theta_o', \phi_o')\, d\omega_i' \tag{2.11}$$

$$= \int_{\Omega_i} L(\theta_i, \phi_i)\hat{\rho}\left(R^{-1}_{\alpha,\beta,\gamma}(\theta_i, \phi_i), \theta_o', \phi_o'\right)\, d\omega_i. \tag{2.12}$$

As we have written them, these equations depend on spherical coordinates. It might clarify matters somewhat to also present an alternate form in terms of rotations and unit vectors in a coordinate-independent way. We simply use R for the rotation, which could be written as a 3×3 rotation matrix, while ω_i and ω_o stand for unit vectors corresponding to the incident and outgoing directions (with primes added for local coordinates). Equations 2.11 and 2.12 may then be written as

$$B(R, \omega_o') = \int_{\Omega_i'} L\left(R\omega_i'\right)\hat{\rho}(\omega_i', \omega_o')\, d\omega_i' \tag{2.13}$$

$$= \int_{\Omega_i} L(\omega_i)\hat{\rho}\left(R^{-1}\omega_i, \omega_o'\right)\, d\omega_i, \tag{2.14}$$

where $R\omega_i'$ and $R^{-1}\omega_i$ are simply matrix-vector multiplications.

Interpretation as Convolution: In the spatial domain, convolution is the result generated when a filter is *translated* over an input signal. However, we can generalize the notion of convolution to other transformations T_a, where T_a is a function of a, and write

$$(f \otimes g)(a) = \int_t f\left(T_a(t)\right) g(t)\, dt. \tag{2.15}$$

When T_a is a translation by a, we obtain the standard expression for spatial convolution. When T_a is a rotation by the angle a, the above formula defines convolution in the angular domain.

Therefore, equations 2.11 and 2.12 (or 2.13 and 2.14) represent rotational convolutions. Equation 2.12 in global coordinates states that the reflected light field at a given surface orientation corresponds to *rotating* the BRDF to that orientation, and then integrating over the upper hemisphere. The BRDF can be thought of as the filter, while the lighting is the input signal. Symmetrically, equation 2.11 in local coordinates states that the reflected light field at a given surface orientation may be computed by *rotating* the lighting into the local coordinate system of the BRDF, and then doing the hemispherical integration. These observations are similar to those we made earlier for the 2D case.

Group-theoretic Interpretation as Generalized Convolution: In fact, it is possible to formally generalize the notion of convolution to groups. Within this context, the standard Fourier convolution formula can be seen as a special case for $SO(2)$, the group of rotations in 2D. More information may be found in books on group representation theory, such as Fulton and Harris [19] (especially note exercise 3.32). One reference that focuses specifically on the rotation group is Chirikjian and Kyatkin [10]. In the general case, we may modify equation 2.15 slightly to write for compact groups,

$$(f \otimes g)(s) = \int_t f(s \circ t) g(t)\, dt, \tag{2.16}$$

where s and t are elements of the group, the integration is over a suitable group measure, and $\circ$ denotes group multiplication.

It is also possible to generalize the Fourier convolution formula in terms of representation matrices of the group in question. In our case, the relations do not exactly satisfy equation 2.16, since we have both rotations (in the rotation group $SO(3)$) and unit vectors. Therefore, for frequency space analysis in the 3D case, we will need both the representation matrices of $SO(3)$, and the associated basis functions for unit vectors on a sphere, which are the spherical harmonics.

2.3 Frequency-Space Analysis

Since the reflection equation can be viewed as a convolution, it is natural to analyze it in frequency-space. We will first consider the 2D reflection equation, which can be analyzed in terms of the familiar Fourier basis functions. We then show how this analysis generalizes to 3D, using the spherical harmonics. Finally, we discuss a number of alternative forms of the reflection equation, and associated convolution formulas, that may be better suited for specific problems.

2.3.1 Fourier Analysis in 2D

We now carry out a Fourier analysis of the 2D reflection equation. We will define the Fourier series of a function f by

$$\begin{aligned} F_k(\theta) &= \frac{1}{\sqrt{2\pi}} e^{Ik\theta} \\ f(\theta) &= \sum_{k=-\infty}^{\infty} f_k F_k(\theta) \\ f_k &= \int_{-\pi}^{\pi} f(\theta) F_k^*(\theta) d\theta. \end{aligned} \tag{2.17}$$

In the last line, the $*$ in the superscript stands for the complex conjugate. For the Fourier basis functions, $F_k^* = F_{-k} = (1/\sqrt{2\pi})\exp(-Ik\theta)$. It should be noted that the relations in equation 2.17 are similar for any orthonormal basis functions F, and we will later be able to use much of the same machinery to define spherical harmonic expansions in 3D.

Decomposition into Fourier Series: We now consider the reflection equation, in the form of equation 2.6. We will expand all quantities in terms of Fourier series.

We start by forming the Fourier expansion of the lighting, L, in global coordinates,

$$L(\theta_i) = \sum_{l=-\infty}^{\infty} L_l F_l(\theta_i). \tag{2.18}$$

To obtain the lighting in local coordinates, we may rotate the above expression,

$$\begin{aligned} L(\theta_i) = L(\alpha + \theta_i') &= \sum_{l=-\infty}^{\infty} L_l F_l(\alpha + \theta_i') \\ &= \sqrt{2\pi} \sum_{l=-\infty}^{\infty} L_l F_l(\alpha) F_l(\theta_i'). \end{aligned} \tag{2.19}$$

The last line follows from the form of the complex exponentials, or in other words, we have $F_l(\alpha + \theta_i') = (1/\sqrt{2\pi}) \exp\left(Il(\alpha + \theta_i')\right)$. This result shows that the effect of rotating the lighting to align it with the local coordinate system is simply to multiply the Fourier frequency coefficients by $\exp(Il\alpha)$.

Since no rotation is applied to B and $\hat{\rho}$, their decomposition into a Fourier series is simple,

$$\begin{aligned} B(\alpha, \theta_o') &= \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} B_{lp} F_l(\alpha) F_p(\theta_o') \\ \hat{\rho}(\theta_i', \theta_o') &= \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} \hat{\rho}_{lp} F_l^*(\theta_i') F_p(\theta_o'). \end{aligned} \tag{2.20}$$

Note that the domain of the basis functions here is $[-\pi, \pi]$, so we develop the series for $\hat{\rho}$ by assuming function values to be 0 outside the range for θ_i' and θ_o' of $[-\frac{\pi}{2}, \frac{\pi}{2}]$. Also, in the expansion for $\hat{\rho}$, the complex conjugate used in the first factor is to somewhat simplify the final result.

Fourier-Space Reflection Equation: We are now ready to write equation 2.6 in terms of Fourier coefficients. For the purposes of summation, we want to avoid confusion of the indices for L and $\hat{\rho}$. For this purpose, we will use the indices L_l and $\hat{\rho}_{l'p}$. We now simply multiply out the expansions for L and $\hat{\rho}$. After taking the summations, and terms not depending on θ_i' outside the integral, equation 2.6 now becomes

$$B(\alpha, \theta_o') = \sqrt{2\pi} \sum_{l=-\infty}^{\infty} \sum_{l'=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} L_l \hat{\rho}_{l'p} F_l(\alpha) F_p(\theta_o') \int_{-\pi}^{\pi} F_{l'}^*(\theta_i') F_l(\theta_i') \, d\theta_i'. \tag{2.21}$$

Note that the limits of the integral are now $[-\pi, \pi]$ and not $[-\frac{\pi}{2}, \frac{\pi}{2}]$. This is because we have already incorporated the fact that the BRDF is nonzero only over the upper half-circle into its Fourier coefficients. Further note that by orthonormality of the Fourier basis, the value of the integrand can be given as

$$\int_{-\pi}^{\pi} F_l'^{*}(\theta_i') F_l(\theta_i')\, d\theta_i' = \delta_{ll'}. \tag{2.22}$$

In other words, we can set $l' = l$ since terms not satisfying this condition vanish. Making this substitution in equation 2.21, we obtain

$$B(\alpha, \theta_o') = \sqrt{2\pi} \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} L_l \hat{\rho}_{lp} F_l(\alpha) F_p(\theta_o'). \tag{2.23}$$

Now, it is a simple matter to equate coefficients in the Fourier expansion of B in order to derive the Fourier-space reflection equation,

$$B_{lp} = \sqrt{2\pi} L_l \hat{\rho}_{lp}. \tag{2.24}$$

This result reiterates once more that the reflection equation can be viewed as a convolution of the incident illumination and BRDF, and becomes a simple product in Fourier space, with an analytic formula being given by equation 2.24.

An alternative form of equation 2.24 that may be more instructive results from holding the local outgoing angle fixed, instead of expanding it also in terms of Fourier coefficients, i.e. replacing the index p by the outgoing angle θ_o',

$$B_l(\theta_o') = \sqrt{2\pi} L_l \hat{\rho}_l(\theta_o'). \tag{2.25}$$

Note that a single value of θ_o' in $B(\alpha, \theta_o')$ corresponds to a slice of the reflected light field, which is *not* the same as a single image from a fixed viewpoint—a single image would instead correspond to fixing the *global* outgoing angle θ_o.

In summary, we have shown that the reflection equation in 2D reduces to the standard convolution formula. Next, we will generalize these results to 3D using spherical harmonic basis functions instead of the complex exponentials.

2.3.2 Spherical Harmonic Analysis in 3D

To extend our frequency-space analysis to 3D, we must consider the structure of rotations and vectors in 3D. In particular, the unit vectors corresponding to incident and reflected directions lie on a sphere of unit magnitude. The appropriate signal-processing tools for the sphere are spherical-harmonics, which are the equivalent for that domain to the Fourier series in 2D (on a circle). These basis functions arise in connection with many physical systems such as those found in quantum mechanics and electrodynamics. A summary of the properties of spherical harmonics can therefore be found in many standard physics textbooks [32, 34, 52].

Although not required for understanding the ensuing derivations, we should point out that our frequency-space analysis is closely related mathematically to the representation theory of the three-dimensional rotation group, $SO(3)$. At the end of the previous section, we already briefly touched on the group-theoretic interpretation of generalized convolution. In the next subsection, we will return to this idea, trying to formally describe the 2D and 3D derivations as special cases of a generalized group-theoretic convolution formula.

Key Properties of Spherical Harmonics: Spherical harmonics are the analogue on the sphere to the Fourier basis on the line or circle. The spherical harmonic Y_{lm} is given by

$$\begin{aligned} N_{lm} &= \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} \\ Y_{lm}(\theta,\phi) &= N_{lm}P_{lm}(\cos\theta)e^{Im\phi}, \end{aligned} \tag{2.26}$$

where N_{lm} is a normalization factor. In the above equation, the azimuthal dependence is expanded in terms of Fourier basis functions. The θ dependence is expanded in terms of the associated Legendre functions P_{lm}. The indices obey $l \geq 0$ and $-l \leq m \leq l$. Thus, there are $2l+1$ basis functions for given order l. Figure 2.3 shows the first 3 orders of spherical harmonics, i.e. the first 9 basis functions corresponding to $l = 0, 1, 2$. They may be written either as trigonometric functions of the spherical coordinates θ and ϕ or as polynomials of the cartesian components x, y and z, with $x^2 + y^2 + z^2 = 1$. In general, a spherical harmonic Y_{lm} is a polynomial of maximum degree l. Another useful relation is that $Y_{l-m} = (-1)^m Y^*_{lm}$. The first 3 orders (we give only terms with $m \geq 0$) are given by the following expressions,

$$\begin{aligned}
Y_{00} &= \sqrt{\frac{1}{4\pi}} \\
Y_{10} &= \sqrt{\frac{3}{4\pi}}\cos\theta &&= \sqrt{\frac{3}{4\pi}}z \\
Y_{11} &= -\sqrt{\frac{3}{8\pi}}\sin\theta e^{I\phi} &&= -\sqrt{\frac{3}{8\pi}}(x+Iy) \\
Y_{20} &= \frac{1}{2}\sqrt{\frac{5}{4\pi}}\left(3\cos^2\theta - 1\right) &&= \frac{1}{2}\sqrt{\frac{5}{4\pi}}\left(3z^2-1\right) \\
Y_{21} &= -\sqrt{\frac{15}{8\pi}}\sin\theta\cos\theta e^{I\phi} &&= -\sqrt{\frac{15}{8\pi}}z(x+Iy) \\
Y_{22} &= \frac{1}{2}\sqrt{\frac{15}{8\pi}}\sin^2\theta e^{2I\phi} &&= \frac{1}{2}\sqrt{\frac{15}{8\pi}}(x+Iy)^2 .
\end{aligned} \tag{2.27}$$

The spherical harmonics form an orthonormal basis in terms of which functions on the sphere can be expanded,

$$\begin{aligned}
f(\theta,\phi) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l} f_{lm} Y_{lm}(\theta,\phi) \\
f_{lm} &= \int_{\phi=0}^{2\pi}\int_{\theta=0}^{\pi} f(\theta,\phi) Y^*_{lm}(\theta,\phi)\sin\theta\, d\theta d\phi.
\end{aligned} \tag{2.28}$$

Note the close parallel with equation 2.17.

The rotation formula for spherical harmonics is

$$Y_{lm}\left(R_{\alpha,\beta,\gamma}(\theta,\phi)\right) = \sum_{m'=-l}^{l} D^l_{mm'}(\alpha,\beta,\gamma) Y_{lm'}(\theta,\phi). \tag{2.29}$$

The important thing to note here is that the m indices are *mixed*—a spherical harmonic after rotation must be expressed as a combination of other spherical harmonics with different m indices. However, the l indices are not mixed; rotations of spherical harmonics with order l are composed entirely of other spherical harmonics with order l. For given order l, D^l

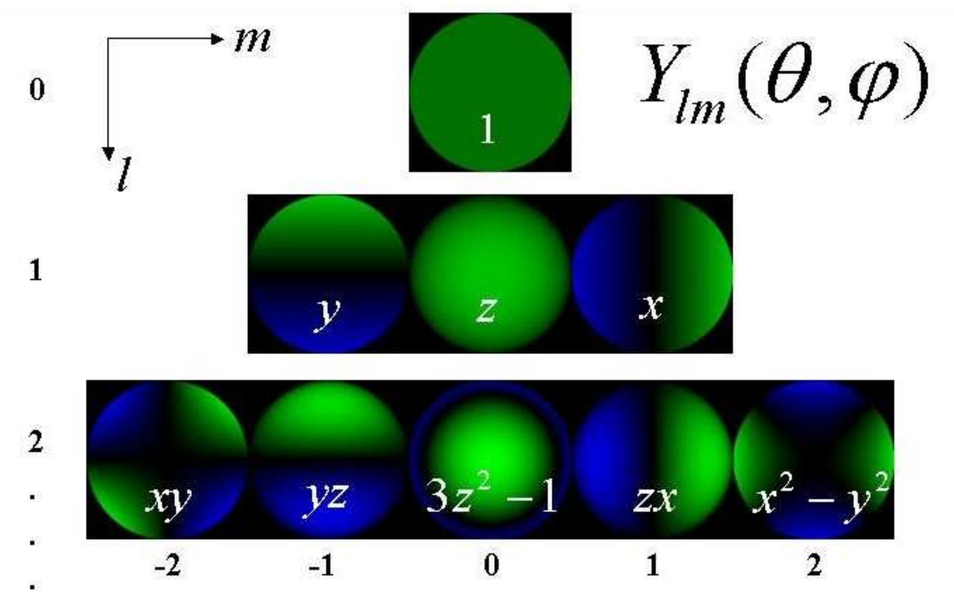

Figure 2.3: *The first 3 orders of real spherical harmonics ($l = 0, 1, 2$) corresponding to a total of 9 basis functions. The spherical harmonics Y_{lm} may be written either as trigonometric functions of the spherical coordinates θ and ϕ or as polynomials of the cartesian components x, y and z, with $x^2 + y^2 + z^2 = 1$. In general, a spherical harmonic Y_{lm} is a polynomial of maximum degree l. In these images, we show only the front the sphere, with green denoting positive values and blue denoting negative values. Also note that these images show the real form of the spherical harmonics. The complex forms are given in equation 2.27.*

is a matrix that tells us how a spherical harmonic transforms under rotation, i.e. how to rewrite a rotated spherical harmonic as a linear combination of all the spherical harmonics of the same order. In terms of group theory, the matrix D^l is the $(2l + 1)$-dimensional representation of the rotation group $SO(3)$. The matrices D^l therefore satisfy the formula,

$$D^l_{mm'}(\alpha, \beta, \gamma) = \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} Y_{lm}(R_{\alpha,\beta,\gamma}(\theta, \phi)) Y^*_{lm'}(\theta, \phi) \sin\theta \, d\theta \, d\phi. \tag{2.30}$$

An analytic form for the matrices D^l can be found in standard references, such as Inui et al. [32]. In particular, since $R_{\alpha,\beta,\gamma} = R_z(\beta)R_y(\alpha)R_z(\gamma)$, the dependence of D^l on β and γ is simple, since rotation of the spherical harmonics about the $z-$axis is straightforward,

$$D^l_{mm'}(\alpha, \beta, \gamma) = d^l_{mm'}(\alpha) e^{Im\beta} e^{Im'\gamma}, \tag{2.31}$$

where d^l is a matrix that defines how a spherical harmonic transforms under rotation about the $y-$axis. For the purposes of the exposition, we will not generally need to be concerned with the precise formula for the matrix d^l, and numerical calculations can compute it using a simplified version of equation 2.30 without the z rotations (i.e. $\beta = \gamma = 0$),

$$d^l_{mm'}(\alpha) = \int_{\phi=0}^{2\pi}\int_{\theta=0}^{\pi} Y_{lm}(R_y(\alpha)(\theta,\phi))Y^*_{lm'}(\theta,\phi)\sin\theta\, d\theta\, d\phi. \tag{2.32}$$

For completeness, we give below the relatively complicated analytic formula, as derived in equation 7.48 of Inui et al. [32],

$$\begin{aligned}
\xi &= \sin^2\frac{\alpha}{2} \\
N(l,m,m') &= (-1)^{m-m'}\sqrt{\frac{(l+m)!}{(l-m)!(l+m')!(l-m')!}} \\
d^l_{mm'}(\alpha) &= N(l,m,m') \times \xi^{-(m-m')/2}(1-\xi)^{-(m+m')/2}\left(\frac{d}{d\xi}\right)^{l-m}\xi^{l-m'}(1-\xi)^{l+m'},
\end{aligned} \tag{2.33}$$

as well as analytic formulae for the the first three representations (i.e. $d^l_{mm'}$ with $l = 0, 1, 2$),

$$\begin{aligned}
d^0(\alpha) &= 1 \\
d^1(\alpha) &= \begin{pmatrix} \cos^2\frac{\alpha}{2} & \frac{\sin\alpha}{\sqrt{2}} & \sin^2\frac{\alpha}{2} \\ -\frac{\sin\alpha}{\sqrt{2}} & \cos\alpha & \frac{\sin\alpha}{\sqrt{2}} \\ \sin^2\frac{\alpha}{2} & -\frac{\sin\alpha}{\sqrt{2}} & \cos^2\frac{\alpha}{2} \end{pmatrix}
\end{aligned}$$

$$d^2(\alpha) = \begin{pmatrix}
\cos^4\frac{\alpha}{2} & 2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & 2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & \sin^4\frac{\alpha}{2} \\
-2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \cos^2\frac{\alpha}{2}(-1+2\cos\alpha) & \sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & (1+2\cos\alpha)\sin^2\frac{\alpha}{2} & 2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} \\
\frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & -\sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \frac{1}{2}(3\cos^2\alpha-1) & \sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha \\
-2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & (1+2\cos\alpha)\sin^2\frac{\alpha}{2} & -\sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \cos^2\frac{\alpha}{2}(-1+2\cos\alpha) & 2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} \\
\sin^4\frac{\alpha}{2} & -2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & -2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \cos^4\frac{\alpha}{2}
\end{pmatrix}. \tag{2.34}$$

To derive some of the quantitative results in section 2.4 and the next chapter, we will require two important properties of the representation matrices D^l, which are derived in appendix A,

$$\begin{aligned}
D^l_{0m'}(\alpha,\beta,0) = \quad & d^l_{0m'}(\alpha) & = \sqrt{\frac{4\pi}{2l+1}}Y^*_{lm'}(\alpha,\pi) \\
D^l_{m0}(\alpha,\beta,\gamma) = \quad & d^l_{m0}(\alpha)e^{Im\beta} & = \sqrt{\frac{4\pi}{2l+1}}Y_{lm}(\alpha,\beta).
\end{aligned} \tag{2.35}$$

Decomposition into Spherical Harmonics: As for the 2D case, we will now expand all the quantities in terms of basis functions. We first expand the lighting in global coordinates,

$$L(\theta_i, \phi_i) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} Y_{lm}(\theta_i, \phi_i). \tag{2.36}$$

To obtain the lighting in local coordinates, we must rotate the above expression, just as we did in 2D. Using equation 2.29, we get,

$$L(\theta_i, \phi_i) = L\left(R_{\alpha,\beta,\gamma}(\theta_i', \phi_i')\right) = \sum_{l=0}^{\infty} \sum_{m=-l}^{+l} \sum_{m'=-l}^{l} L_{lm} D^l_{mm'}(\alpha, \beta, \gamma) Y_{lm'}(\theta_i', \phi_i'). \tag{2.37}$$

We now represent the transfer function $\hat{\rho} = \rho \cos\theta_i'$ in terms of spherical harmonics. As in 2D, we note that $\hat{\rho}$ is nonzero only over the upper hemisphere, i.e. when $\cos\theta_i' > 0$ and $\cos\theta_o' > 0$. Also, as in 2D, we use a complex conjugate for the first factor, to simplify the final results.

$$\hat{\rho}(\theta_i', \phi_i', \theta_o', \phi_o') = \sum_{l=0}^{\infty} \sum_{n=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-p}^{p} \hat{\rho}_{ln,pq} Y^*_{ln}(\theta_i', \phi_i') Y_{pq}(\theta_o', \phi_o') \tag{2.38}$$

Spherical Harmonic Reflection Equation: We can now write down the reflection equation, as given by equation 2.11, in terms of the expansions just defined. As in 2D, we multiply the expansions for the lighting and BRDF. To avoid confusion between the indices in this intermediate step, we will use L_{lm} and $\hat{\rho}_{l'n,pq}$ to obtain

$$\begin{aligned}
B(\alpha, \beta, \gamma, \theta_o', \phi_o') &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{m'=-l}^{l} \sum_{l'=0}^{\infty} \sum_{n=-l'}^{l'} \sum_{p=0}^{\infty} \sum_{q=-p}^{p} L_{lm} \hat{\rho}_{l'n,pq} D^l_{mm'}(\alpha, \beta, \gamma) Y_{pq}(\theta_o', \phi_o') T_{lm'l'n} \\
T_{lm'l'n} &= \int_{\phi_i'=0}^{2\pi} \int_{\theta_i'=0}^{\pi} Y_{lm'}(\theta_i', \phi_i') Y^*_{l'n}(\theta_i', \phi_i') \sin\theta_i' \, d\theta_i' d\phi_i' \\
&= \delta_{ll'} \delta_{m'n}.
\end{aligned} \tag{2.39}$$

The last line follows from orthonormality of the spherical harmonics. Therefore, we may set $l' = l$ and $n = m'$ since terms not satisfying these conditions vanish. We then obtain

$$B(\alpha, \beta, \gamma, \theta_o', \phi_o') = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-p}^{p} L_{lm} \hat{\rho}_{ln,pq} \left(D^l_{mn}(\alpha, \beta, \gamma) Y_{pq}(\theta_o', \phi_o') \right). \tag{2.40}$$

This result suggests that we should expand the reflected light field B in terms of the new basis functions given by $C_{lmnpq} = D^l_{mn}(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o)$. The appearance of the matrix D^l in these basis functions is quite intuitive, coming directly from the rotation formula for spherical harmonics. These basis functions are *mixed* in the sense that they are a product of the matrices D^l and the spherical harmonics Y_{pq}. This can be understood from realizing that the reflected direction is a unit vector described by two parameters (θ'_o, ϕ'_o), while the surface parameterization is really a rotation, described by three parameters (α, β, γ). Finally, we need to consider the normalization of these new basis functions. The spherical harmonics are already orthonormal. The orthogonality relation for the matrices D^l is given in standard texts on group theory (for instance, equation 7.73 of Inui et al. [32]). Specifically,

$$\int_{\gamma=0}^{2\pi} \int_{\beta=0}^{2\pi} \int_{\alpha=0}^{\pi} \left(D^l_{mn}(\alpha, \beta, \gamma) \right)^* \left(D^{l'}_{m'n'}(\alpha, \beta, \gamma) \right) \sin\alpha \, d\alpha \, d\beta \, d\gamma = \frac{8\pi^2}{2l+1} \delta^{ll'} \delta_{mm'} \delta_{nn'}. \tag{2.41}$$

In the equation above, the group-invariant measure $d\mu(g)$ of the rotation group $g = SO(3)$ is $\sin\alpha \, d\alpha \, d\beta \, d\gamma$. The integral of this quantity $\mu(g) = 8\pi^2$, which can be easily verified. Therefore, to obtain an orthonormal basis, we must normalize appropriately. Doing this,

$$\begin{aligned}
C_{lmnpq} &= \sqrt{\frac{2l+1}{8\pi^2}} D^l_{mn}(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o) \\
B &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-p}^{p} B_{lmnpq} C_{lmnpq}(\alpha, \beta, \gamma, \theta'_o, \phi'_o) \\
B_{lmnpq} &= \int_{\phi'_o=0}^{2\pi} \int_{\theta'_o=0}^{\pi} \int_{\gamma=0}^{2\pi} \int_{\beta=0}^{2\pi} \int_{\alpha=0}^{\pi} U(\alpha, \beta, \gamma, \theta'_o, \phi'_o) \sin\alpha \sin\theta'_o \, d\alpha \, d\beta \, d\gamma \, d\theta'_o \, d\phi'_o \\
U(\alpha, \beta, \gamma, \theta'_o, \phi'_o) &= B(\alpha, \beta, \gamma, \theta'_o, \phi'_o) C^*_{lmnpq}(\alpha, \beta, \gamma, \theta'_o, \phi'_o).
\end{aligned} \tag{2.42}$$

Although this appears rather involved, it is a straightforward expansion of the reflected light field in terms of orthonormal basis functions. As written, since we are assuming anisotropic surfaces for full generality, the reflected light field is a function of five variables, as opposed to being a function of only two variables in 2D. We should note that it is generally impractical to have the full range of values for the anisotropic parameter, i.e. the tangent frame rotation, γ for every surface orientation. In fact, γ is often a function of the surface orientation (α, β). However, our goal here is to write the completely general formulae. In the next subsection, we will derive an alternative form for isotropic surfaces which corresponds more closely to observable quantities.

Finally, we can write down the frequency space reflection equation by comparing equations 2.40 and 2.42 and equating coefficients. This result is comparable to its 2D counterpart, given in equation 2.24, and as in 2D, is a convolution. In frequency-space, the reflected light field is obtained simply by multiplying together coefficients of the lighting and BRDF, i.e. by *convolving* the incident illumination with the BRDF,

$$B_{lmnpq} = \sqrt{\frac{8\pi^2}{2l+1}} L_{lm}\hat{\rho}_{ln,pq}. \tag{2.43}$$

As in 2D, an alternative result without expanding the output dependence may be more instructive,

$$B_{lmn}(\theta'_o, \phi'_o) = \sqrt{\frac{8\pi^2}{2l+1}} L_{lm}\hat{\rho}_{ln}(\theta'_o, \phi'_o). \tag{2.44}$$

We reiterate that the fixed *local* outgoing angle in the above equation does *not* correspond to a single image, but to a more general slice of the reflected light field. In a single image, the *local* viewing angle is different for different points in the image, depending on the relative orientation between the surface normal and viewing direction. On the other hand, a single image corresponds to a single *global* viewing direction, and hence a single *global* outgoing angle.

In summary, we have shown that the direct illumination integral, or reflection equation, can be viewed in signal processing terms as a convolution of the incident illumination and BRDF, and have derived analytic formulae. These analytic results quantify the qualitative observations made by many researchers in the past. In 2D, the formulae are in terms of the standard Fourier basis. In 3D, we must instead use spherical harmonics and the representation matrices of the rotation group, deriving a generalized convolution formula. Still, the extension from 2D to 3D is conceptually straightforward, and although the mathematics is significantly more involved, the key idea that the reflected light field can be viewed in a precise quantitative way as a convolution still holds.

2.3.3 Group-theoretic Unified Analysis

While not required for understanding the rest of this chapter, it is insightful to attempt to analyze the 2D and 3D derivations as special cases of a more general convolution formula in terms of the representation theory of compact groups. Our analysis in this subsection will be based on that in Fulton and Harris [19] and Chirikjian and Kyatkin [10].

Convolution can be defined on general compact groups using equation 2.16. To analyze this in the frequency domain, we need a generalization of the Fourier transform. It is possible to define

$$f_l = \int_G f(g) D^l(g)\, dg. \tag{2.45}$$

In this equation, f_l is the generalization of the Fourier transform, corresponding to index l, $f(g)$ is the function defined on the group G of which g is a member, and $D^l(g)$ is the (irreducible) representation matrix labeled with index l, evaluated at the group element g. Here, the group-invariant measure for integration is written dg or $d\mu(g)$.

To obtain some intuition, consider the flatland case where the group corresponds to rotations in 2D, i.e. $G = SO(2)$. The elements g are then simply the angles ϕ, and the representation matrices are all 1-dimensional and correspond to the standard Fourier series, i.e. $D^l = e^{Il\phi}$. Thus, equation 2.45 corresponds directly to the standard Fourier series in 2D. Now, consider the case where the group is that of 3D rotations, i.e. $G = SO(3)$. In this case, D^l corresponds to the $2l+1$-dimensional representation, and is a $(2l+1) \times (2l+1)$ representation matrix. The generalized Fourier transform is therefore *matrix-valued*. For a general compact group, we can generalize the notion of the Fourier transform to a matrix-valued function labeled by indices corresponding to the group representation. This reduces to the standard Fourier series for the 2D or flatland case, since the group representations are all one-dimensional and correspond directly to complex exponentials.

Once we have the generalization of the Fourier transform, one can derive [19] a convolution formula corresponding to equation 2.16,

$$(f \otimes g)_l = f_l \times g_l. \tag{2.46}$$

It should be noted that the multiplication on the right-hand side is now a matrix multiplication, since all coefficients are matrix-valued. In the 2D flatland case, these are just standard Fourier coefficients, so we have a simple scalar multiplication, reducing to the standard Fourier convolution formula. For 3D rotations, the convolution formula should involve matrix multiplication of the generalized Fourier coefficients (obtained by integrating against the representation matrices of $SO(3)$).

However, it is important to note that the 3D case discussed in the previous subsection does not correspond exactly either to equation 2.16 or 2.46. Specifically, all operations are not carried out in the rotation group $SO(3)$. Instead, we have rotations operating on unit vectors. Thus, it is not possible to apply equation 2.46 directly, and a separate convolution formula must be derived, as we have done in the previous subsection.

Note that we use the associated basis functions, i.e. spherical harmonics, and not the group representations directly, as basis functions for the lighting and BRDF, since these quantities are functions of directions or unit vectors, and not rotations. For the reflected light field, which is a function of the rotation applied as well as the outgoing direction (a unit vector), we use mixed basis functions that are a product of group representations of $SO(3)$ and the spherical harmonics. The convolution formula we derive in equation 2.43 is actually simpler than equation 2.46, since it does not require matrix multiplication.

In the remainder of this section, we will derive a number of alternative forms for equation 2.43 that may be more suitable for particular cases. Then, in the next section, we will discuss the implications of equations 2.43 and 2.44 for forward and inverse problems in computer graphics.

2.3.4 Alternative Forms

For the analysis of certain problems, it will be more convenient to rewrite equation 2.43 in a number of different ways. We have already seen one example, of considering the outgoing angle fixed, as shown in equation 2.44. In this subsection, we consider a few more alternative forms.

Isotropic BRDFs

Isotropic BRDFs are those where rotating the local tangent frame makes no difference, i.e. they are functions of only 3 variables, $\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \hat{\rho}(\theta'_i, \theta'_o, \mid \phi'_o - \phi'_i \mid)$. With respect to the reflected light field, the parameter γ, which controls the orientation of the local tangent frame, has no physical significance for isotropic BRDFs.

To consider the simplifications that result from isotropy, we first analyze the BRDF coefficients $\hat{\rho}_{ln,pq}$. In the BRDF expansion of equation 2.38, only terms that satisfy isotropy, i.e. are invariant with respect to adding an angle $\triangle\phi$ to both incident and outgoing azimuthal angles, are nonzero. From the form of the spherical harmonics, this requires that $n = q$. Furthermore, since we are considering BRDFs that depend only on $\mid \phi'_o - \phi'_i \mid$, we should be able to negate both incident and outgoing azimuthal angles without changing the result. This leads to the condition that $\hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q)p(-q)}$. Finally, we define a 3-index BRDF coefficient by

$$\hat{\rho}_{lpq} = \hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q)p(-q)}. \tag{2.47}$$

Note that isotropy reduces the dimensionality of the BRDF from 4D to 3D. This is reflected in the fact that we now have only three independent indices. Furthermore, half the degrees of freedom are constrained since we can negate the azimuthal angle without changing the BRDF.

Next, we remove the dependence of the reflected light field on γ by arbitrarily setting $\gamma = 0$. It can be verified that for isotropic surfaces, γ mathematically just controls the origin or 0-angle for ϕ'_o and can therefore be set arbitrarily. Upon doing this, we can simplify a number of quantities. First, the rotation operator is now given simply by

$$R_{\alpha,\beta} = R_{\alpha,\beta,0} = R_z(\beta)R_y(\alpha). \tag{2.48}$$

Next, the representation matrices can be rewritten as

$$D^l_{mn}(\alpha, \beta) = D^l_{mn}(\alpha, \beta, 0) = d^l_{mn}(\alpha)e^{Im\beta}. \tag{2.49}$$

It should be noted that removing the dependence on γ weakens the orthogonality condition (equation 2.41) on the representation matrices, since we no longer integrate over γ.

The new orthonormality relation for these matrices is given by

$$\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi}\left(D^l_{mn}(\alpha,\beta)\right)^*\left(D^{l'}_{m'n}(\alpha,\beta)\right)\sin\alpha\,d\alpha\,d\beta=\frac{4\pi}{2l+1}\delta^{ll'}\delta_{mm'}. \tag{2.50}$$

In particular, the orthogonality relation for the index n no longer holds, which is why we have used the index n for both D and D' instead of using n and n', as in equation 2.41. The absence of an integral over γ leads to a slight weakening of the orthogonality relation, as well as a somewhat different normalization than in equation 2.41. For the future discussion, it will be convenient to define normalization constants by

$$\Lambda_l=\sqrt{\frac{4\pi}{2l+1}}. \tag{2.51}$$

We now have the tools necessary to rewrite equation 2.43 for isotropic BRDFs. Since we will be using the equations for isotropic BRDFs extensively in the rest of this dissertation, it will be worthwhile to briefly review the representations of the various quantities, specialized to the isotropic case.

First, we define the expansions of the lighting in global coordinates, and the results from rotating this expansion,

$$\begin{aligned}L(\theta_i,\phi_i) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l}L_{lm}Y_{lm}(\theta_i,\phi_i)\\ L(\theta_i,\phi_i)=L\left(R_{\alpha,\beta}(\theta_i',\phi_i')\right) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{+l}\sum_{m'=-l}^{l}L_{lm}D^l_{mm'}(\alpha,\beta)Y_{lm'}(\theta_i',\phi_i').\end{aligned} \tag{2.52}$$

Then, we write the expansion of the isotropic BRDF,

$$\hat{\rho}(\theta_i',\theta_o',\mid\phi_o'-\phi_i'\mid)=\sum_{l=0}^{\infty}\sum_{p=0}^{\infty}\sum_{q=-\min(l,p)}^{\min(l,p)}\hat{\rho}_{lpq}Y^*_{lq}(\theta_i',\phi_i')Y_{pq}(\theta_o',\phi_o'). \tag{2.53}$$

The reflected light field, which is now a 4D function, can be expanded using a product of representation matrices and spherical harmonics,

$$C_{lmpq}(\alpha,\beta,\theta_o',\phi_o') = \Lambda_l^{-1}D^l_{mq}(\alpha,\beta)Y_{pq}(\theta_o',\phi_o')$$

$$\begin{aligned}
B(\alpha,\beta,\theta_o',\phi_o') &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-\min(l,p)}^{\min(l,p)} B_{lmpq}C_{lmpq}(\alpha,\beta,\theta_o',\phi_o') \\
B_{lmpq} &= \int_{\phi_o'=0}^{2\pi}\int_{\theta_o'=0}^{\pi}\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi} U(\alpha,\beta,\theta_o',\phi_o')\sin\alpha\sin\theta_o'\,d\alpha\,d\beta\,d\theta_o'\,d\phi_o' \\
U(\alpha,\beta,\theta_o',\phi_o') &= B(\alpha,\beta,\theta_o',\phi_o')C^*_{lmpq}(\alpha,\beta,\theta_o',\phi_o'). \qquad (2.54)
\end{aligned}$$

It should be noted that the basis functions C_{lmpq} are orthonormal in spite of the weakened orthogonality of the functions D^l_{mq}, as expressed in equation 2.50. Note that the index q in the definition of C_{lmpq} is the same (coupled) for both factors D^l_{mq} and Y_{pq}. This is a consequence of isotropy, and is not true in the anisotropic case. Therefore, although the representation matrices D^l no longer satisfy orthogonality over the index q (corresponding to the index n in equation 2.50), orthogonality over the index q follows from the orthonormality of the spherical harmonics Y_{pq}.

Finally, we can derive an analytic expression (convolution formula) for the reflection equation in terms of these coefficients,

$$B_{lmpq} = \Lambda_l L_{lm}\hat{\rho}_{lpq}. \qquad (2.55)$$

Apart from a slightly different normalization, and the removal of γ and the corresponding index n, this is essentially the same as equation 2.43. We will be using this equation for isotropic BRDFs extensively in the next chapter, where we quantitatively analyze the reflection equation for many special cases of interest.

We may also try to derive an alternative form, analogous to equation 2.44, by holding the outgoing elevation angle θ_o' fixed. Since the isotropic BRDF depends only on $\mid \phi_o' - \phi_i' \mid$, and not directly on ϕ_o', we do not hold ϕ_o' fixed, as we did in equation 2.44. We first define the modified expansions,

$$\begin{aligned}
\hat{\rho}(\theta_i',\theta_o',\mid\phi_o'-\phi_i'\mid) &= \sum_{l=0}^{\infty}\sum_{q=-l}^{l}\hat{\rho}_{lq}(\theta_o')\left(\frac{1}{\sqrt{2\pi}}Y^*_{lq}(\theta_i',\phi_i')\exp(Iq\phi_o')\right) \\
B(\alpha,\beta,\theta_o',\phi_o') &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{q=-l}^{l} B_{lmq}(\theta_o')\left(\frac{1}{\sqrt{2\pi}}\Lambda_l^{-1}D^l_{mq}(\alpha,\beta)\exp(Iq\phi_o')\right). \qquad (2.56)
\end{aligned}$$

Then, we may write down the isotropic convolution formula corresponding to equation 2.44,

$$B_{lmq}(\theta'_o) = \Lambda_l L_{lm} \hat{\rho}_{lq}(\theta'_o). \tag{2.57}$$

Reciprocity Preserving

One of the important properties of physical BRDFs is that they are *reciprocal*, i.e. symmetric with respect to interchange of incident and outgoing angles. However, the transfer function $\hat{\rho} = \rho \cos\theta'_i$ as defined by us, does not preserve this reciprocity of the BRDF. To make the transfer function reciprocal, we should multiply it by $\cos\theta'_o$ also. To preserve correctness, we must then multiply the reflected light field by $\cos\theta'_o$ as well. Specifically, we define

$$\begin{aligned} \tilde{\rho} &= \hat{\rho}\cos\theta'_o = \rho\cos\theta'_i\cos\theta'_o \\ \tilde{B} &= B\cos\theta'_o. \end{aligned} \tag{2.58}$$

With these definitions, all of the derivations presented so far still hold. In particular, the convolution formulas in equation 2.43 and 2.55 hold with the replacements $B \rightarrow \tilde{B}, \hat{\rho} \rightarrow \tilde{\rho}$. For example, equation 2.55 for isotropic BRDFs becomes

$$\tilde{B}_{lmpq} = \Lambda_l L_{lm} \tilde{\rho}_{lpq}. \tag{2.59}$$

The symmetry of the transfer function ensures that its coefficients are unchanged if the indices corresponding to incident and outgoing angles are interchanged, i.e. $\tilde{\rho}_{lpq} = \tilde{\rho}_{plq}$. In the more general anisotropic case, $\tilde{\rho}_{ln,pq} = \tilde{\rho}_{pq,ln}$. We will use the frequency-space reflection formula, as given by equation 2.59, whenever explicitly maintaining the reciprocity of the BRDF is important.

Reparameterization by central BRDF direction

Consider first the special case of *radially symmetric or 1D BRDFs*, where the BRDF consists of a single symmetric lobe of fixed shape, whose orientation depends only on a well-defined central direction $\vec{C}$. In other words, the BRDF is given by a 1D function u as

$\hat{\rho} = u(\vec{C} \cdot \vec{L})$. Examples are Lambertian $\hat{\rho} = \vec{N} \cdot \vec{L}$ and Phong $\hat{\rho} = (\vec{R} \cdot \vec{L})^s$ models. If we reparameterize the BRDF and reflected light field by $\vec{C}$, the BRDF becomes a function of only 1 variable (θ_i' with $\cos\theta_i' = \vec{C} \cdot \vec{L}$) instead of 3. Refer to figure 2.4 for an illustration. Further, the reflected light field can be represented simply by a 2D reflection map $B(\alpha, \beta)$ parameterized by $\vec{C} = (\alpha, \beta)$. In other words, after reparameterization, there is no explicit exitant (outgoing) angular dependence for either the BRDF or reflected light field.

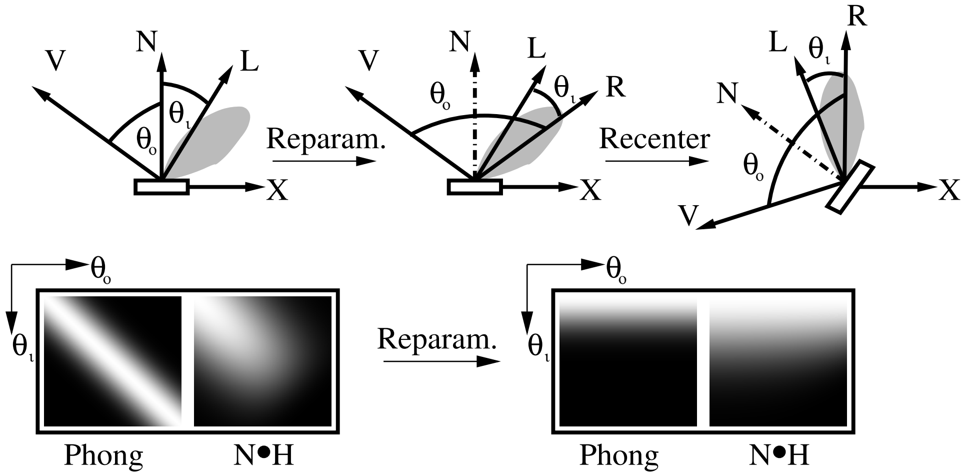

Figure 2.4: *Reparameterization involves recentering about the reflection vector. BRDFs become more compact, and in special cases (Phong) become 1D functions.*

We may now write the BRDF and equations for the reflected light field as

$$\begin{aligned}
\hat{\rho}(\theta_i') &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta_i') \\
\hat{\rho}_l &= 2\pi \int_0^{\pi/2} \hat{\rho}(\theta_i') Y_{l0}(\theta_i') \, \sin\theta_i' \, d\theta_i' \\
B(\alpha, \beta) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{q=-l}^{l} \hat{\rho}_l L_{lm} D_{mq}^l(\alpha, \beta) \int_0^{2\pi} \int_0^{\pi} Y_{lq}(\theta_i', \phi_i') Y_{l0}(\theta_i') \, \sin\theta_i' \, d\theta_i' d\phi_i' \\
&= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \hat{\rho}_l L_{lm} D_{m0}^l(\alpha, \beta).
\end{aligned} \tag{2.60}$$

In computing $\hat{\rho}_l$, we have integrated out the azimuthal dependence, accounting for the factor of 2π. In the last line, we have used orthonormality of the spherical harmonics. Now, we use the second property of the matrices D from equation 2.35, i.e. $D_{m0}^l(\alpha, \beta) =$

$\Lambda_l Y_{lm}(\alpha, \beta)$. Therefore, the reflected light field can be expanded simply in terms of spherical harmonics,

$$B(\alpha, \beta) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm} Y_{lm}(\alpha, \beta). \tag{2.61}$$

The required convolution formula now becomes

$$B_{lm} = \Lambda_l \hat{\rho}_l L_{lm}. \tag{2.62}$$

In the context of Lambertian BRDFs (for which no reparameterization is required), it has been noted by Basri and Jacobs [2] that equation 2.62 is mathematically an instance of the Funk-Hecke theorem (as stated, for instance in Groemer [24], page 98). However, that theorem does not generalize to the other relations previously encountered. With respect to equation 2.55, we have essentially just dropped the indices p and q corresponding to the outgoing angular dependence. It is important to remember that the reflected light field is now expanded in terms of spherical harmonics. B is simply a filtered version of L, with each frequency l being attenuated by a different amount, corresponding to the BRDF transfer function $\hat{\rho}_l$.

For general BRDFs, the radial symmetry property does not hold precisely, so they cannot be reduced exactly to 1D functions, nor can B be written simply as a 2D reflection map. Nevertheless, a reparameterization of the specular BRDF components by the reflection vector (or other central BRDF direction) still yields compact forms. To reparameterize, we simply recenter the BRDF (and the reflection integral) about the reflection vector $\vec{R}$, rather than the surface normal, as shown in figure 2.4. The reflection vector now takes the place of the surface normal, i.e. $\vec{R} = (\alpha, \beta)$, and the dependence on the surface normal becomes indirect (just as the dependence on $\vec{R}$ is indirect in the standard parameterization). The angles θ_i' and θ_o' are now given with respect to $\vec{R}$ by $\cos\theta_i' = \vec{R} \cdot \vec{L}$ and $\cos\theta_o' = \vec{R} \cdot \vec{V}$, with $B(\alpha, \beta, \theta_o', \phi_o')$ a function of $\vec{R} = (\alpha, \beta)$ and $\omega_o = (\theta_o', \phi_o')$. Once we have done this, we can directly apply the general convolution formulas, such as equation 2.55.

This section has presented a frequency-space analysis of the reflection equation. We have shown the simple quantitative form that results from this analysis, as embodied by equations 2.43 and 2.55. The mathematical analysis leading to these results is the main contribution of this chapter, showing quantitatively that reflection can be viewed as a convolution. The next section gives an overview of the implications of these results for forward and inverse problems in rendering. The next chapter will work out a number of special cases of interest.

2.4 Implications

This section discusses the implications of the theoretical analysis developed in the previous section. Our main focus will be on understanding the well-posedness and conditioning of inverse problems, as well as the speedups obtained in forward problems. In this section, we make some general observations. In the second part of the paper, we will quantitatively analyze a number of special cases of interest.

We will deal here exclusively with the 3D case, since that is of greater practical importance. A preliminary analysis for the 2D case can be found in an earlier paper [70]. The quantitative results in 2D and 3D are closely related, although the fact that the 3D treatment is in terms of spherical harmonics, as opposed to the 2D treatment in terms of Fourier series, results in some important differences. For simplicity, we will also restrict the ensuing discussion to the case of isotropic BRDFs. The extension to anisotropic surfaces can be done using the equations derived earlier for the general anisotropic case.

2.4.1 Forward Rendering with Environment Maps

We first consider the problem of rendering with *environment maps*, i.e. general lighting distributions. For the purposes of rendering, it is convenient to explicitly write the formula for the reflected light field as

$$B(\alpha, \beta, \theta_o', \phi_o') = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} L_{lm} \hat{\rho}_{lpq} \left(D^l_{mq}(\alpha, \beta) Y_{pq}(\theta_o', \phi_o') \right). \tag{2.63}$$

If either the lighting or the BRDF is low frequency, the total number of terms in the summation will be relatively small, and it may be possible to use equation 2.63 directly for shading a pixel. In chapter 4, we will demonstrate the practicality of this approach for

Lambertian BRDFs, where we can set $p = q = 0$, and use $l \leq 2$, i.e. only 9 spherical harmonic terms.

In the general case, frequency space analysis allows for setting sampling rates accurately, and enables compact frequency domain representations. Further, just as image convolutions are often computed in the Fourier rather than the spatial domain, computing the reflected light field is more efficient in frequency space, using equation 2.63, rather than in angular space. Chapter 5 describes the practical implementation of these ideas.

2.4.2 Well-posedness and conditioning of Inverse Lighting and BRDF

In this subsection, we briefly discuss how to apply ideas from the theoretical analysis to determine which inverse problems are well-posed, i.e. solvable, versus ill-posed, i.e. unsolvable, and also determine the numerical conditioning properties. At the end of this subsection, we will also relate these results to the general theory of linear integral equations. An important duality should be noted here. Forward problems for which an efficient frequency domain solution is possible, such as those involving diffuse surfaces and/or soft lighting, have corresponding inverse problems that are ill-conditioned. Turned around, ill-conditioned inverse problems allow us to get a very good solution to the forward problem by using very coarse low-frequency approximations of the initial conditions. For instance, Lambertian surfaces act as low-pass filters, the precise form of which we will explore in the next chapter, blurring the illumination. Therefore, high-frequency components of the lighting are not essential to rendering images of diffuse objects, and we can make very coarse low-frequency approximations to the lighting without significantly affecting the final image. This leads to more efficient algorithms for computer graphics, and illustrates one of the benefits in considering a signal-processing view of reflection.

Inverse-BRDF

We first address the question of BRDF estimation. Our goal is to consider this problem under general illumination conditions, and understand when the BRDF can be recovered, i.e. BRDF estimation is well posed, and when the BRDF cannot be recovered, i.e. estimation is ill-posed. We would also like to know when BRDF recovery will be well-conditioned, i.e.

numerically robust. An understanding of these issues is critical in designing BRDF estimation algorithms that work under arbitrary lighting. Otherwise, we may devise algorithms that attempt to estimate BRDF components that cannot be calculated, or whose estimation is ill-conditioned.

For isotropic surfaces, a simple manipulation of equation 2.55 yields

$$\hat{\rho}_{lpq} = \Lambda_l^{-1} \frac{B_{lmpq}}{L_{lm}}. \tag{2.64}$$

In general, BRDF estimation will be well-posed, i.e. unambiguous as long as the denominator on the right-hand side does not vanish. Of course, to be physically accurate, the numerator will also become 0 if the denominator vanishes, so the right-hand side will become indeterminate. From equation 2.64, we see that BRDF estimation is well posed as long as for all l, there exists at least one value of m so that $L_{lm} \neq 0$. In other words, all orders in the spherical harmonic expansion of the lighting should have at least one coefficient with nonzero amplitude. If any order l completely vanishes, the corresponding BRDF coefficients cannot be estimated.

In signal processing terms, if the input signal (lighting) has no amplitude along certain modes of the filter (BRDF), those modes cannot be estimated. BRDF recovery is well conditioned when the spherical harmonic expansion of the lighting does not decay rapidly with increasing frequency, i.e. when the lighting contains high frequencies like directional sources or sharp edges, and is ill-conditioned for soft lighting. Equation 2.64 gives a precise mathematical characterization of the conditions for BRDF estimation to be well-posed and well-conditioned. These results are similar to those obtained by D'Zmura [17] who states that there is an ambiguity regarding the BRDF in case of *inadequate illumination*. In our framework, inadequate illumination corresponds to certain frequencies l of the lighting completely vanishing.

Inverse Lighting

A similar analysis can be done for estimation of the lighting. Manipulation of equation 2.55 yields

$$L_{lm} = \Lambda_l^{-1} \frac{B_{lmpq}}{\hat{\rho}_{lpq}}. \tag{2.65}$$

Inverse lighting will be well-posed so long as the denominator does not vanish for all p, q for some l, i.e. so long as the spherical harmonic expansion of the BRDF transfer function contains all orders. In signal processing terms, when the BRDF filter truncates certain frequencies in the input lighting signal (for instance, if it were a low-pass filter), we cannot determine those frequencies from the output signal. Inverse lighting is well-conditioned when the BRDF has high-frequency content, i.e. its frequency spectrum decays slowly. In physical terms, inverse lighting is well-conditioned when the BRDF contains sharp specularities, the ideal case of which is a mirror surface. On the other hand, inverse lighting from matte or diffuse surfaces is ill-conditioned. Intuitively, highly specular surfaces act as high-pass filters, so the resulting images have most of the high frequency content in the lighting, and the lighting can be estimated. On the other hand, diffuse surfaces act as low-pass filters, *blurring* the illumination and making it difficult or impossible to recover the high frequencies.

Analysis in terms of theory of Fredholm Integral equations

We now briefly put our results on the well-posedness of inverse lighting and BRDF problems into a broader context with respect to the theory of Fredholm integral equations. Inverting the reflection equation to solve for the lighting or BRDF is essentially a Fredholm integral equation of the first kind. By contrast, the (forward) global illumination problem typically considered in rendering is a Fredhom integral equation of the second kind. Fredholm integral equations of the first kind may be written generally as

$$b(s) = \int_t K(s,t) f(t)\, dt, \tag{2.66}$$

where $b(s)$ is the known quantity (observation), $K(s,t)$ is the kernel or operator in the equation, and $f(t)$ is the function we seek to find. To make matters concrete, one may think of f as the incident illumination L, with the kernel K as corresponding to the (rotated) BRDF operator, and $b(s)$ as corresponding to the reflected light field. Here, t would represent the incident direction, and s would represent the surface orientation and outgoing direction.

The theory of linear integral equations, as for instance in Cochran [11], analyzes equation 2.66 based on the structure of the kernel. In particular, assume we may find a basis function expansion of the form

$$\begin{aligned} b(s) &= \sum_{i=1}^{n} b_i u_i(s) \\ K(s,t) &= \sum_{i=1}^{n} K_i u_i(s) v_i^*(t) \\ f(t) &= \sum_{i=1}^{\infty} f_i v_i(t), \end{aligned} \tag{2.67}$$

where each of the sets of functions u_i and v_i (with v_i^* being the complex conjugate) is linearly independent. Here, n is the number of terms in, or rank of the kernel, K. If n is finite, the kernel is referred to as degenerate. It should be noted that if the function sets u and v were orthonormal, then we would have a result of the form $b_i = K_i f_i$. In effect, we have constructed an expansion of the form of equation 2.67 using orthonormal basis functions involving group representations and spherical harmonics, thereby deriving the convolution result.

As long as the kernel has finite rank n, it annihilates some terms in f, (for $i > n$), and the integral equation is therefore ill-posed (has an infinity of solutions). If the kernel has numerically finite rank, the integral equation is ill-conditioned. Our analysis can be seen as trying to understand the rank of the kernel and its degeneracies in terms of signal processing, thereby determining up to what order the function f can be recovered. In the future, it may be possible to directly apply the theory of integral equations to analyze the well-posedness and conditioning of inverse problems for which simple analytic formulae such as our convolution relation are not readily available.

2.4.3 Light Field Factorization

Having analyzed estimation of the BRDF and lighting alone, we now consider the problem of *factorizing* the light field, i.e simultaneously recovering the lighting and BRDF when both are unknown. An analysis of this problem is very important theoretically in understanding the properties of the light field. There is also potential for practical applications in many different areas. Within BRDF estimation, being able to factor the light field allows us to estimate BRDFs under uncontrolled unknown illumination, with the lighting being recovered as part of the algorithm. Similarly, it would be useful to be able to recover the lighting from an object of unknown BRDF. Factorization reveals the structure of the light field, allowing for more intuitive editing operations to be carried out in order to synthesize novel images for computer graphics. Factorization also reduces the dimensionality, and is therefore useful in compressing light fields that are usually very large.

We first note that there is a global scale factor that we cannot recover. Multiplying the lighting everywhere by some constant amount and dividing the BRDF uniformly by the same amount leaves the reflected light field, which is a product of the two, unchanged. Of course, physical considerations bound the scale factor, since the BRDF must remain energy preserving. Nevertheless, within this general constraint, it is not possible to estimate the absolute magnitudes of the lighting and BRDF. However, we will demonstrate that apart from this ambiguity, the light field can indeed be factored, allowing us to simultaneously determine both the lighting and the BRDF.

An important observation concerns the dimensionality of the various components. The isotropic BRDF is defined on a 3D domain, while the lighting is a function of 2D. On the other hand, the reflected light field is defined on a 4D domain. This indicates that there is a great deal of redundancy in the reflected light field. The number of knowns, i.e. coefficients of the reflected light field, is greater than the number of unknowns, i.e. coefficients of the lighting and BRDF. This indicates that factorization should be tractable. Indeed, for fixed order l, we can use known lighting coefficients L_{lm} to find unknown BRDF coefficients $\hat{\rho}_{lpq}$ and vice-versa. In fact, we need only one known nonzero lighting or BRDF coefficient for order l to bootstrap this process, since inverse lighting can use any value of (p, q) and inverse-BRDF computation can use any value of m.

It would appear from equation 2.55 however, that there is an unrecoverable scale factor for each order l, corresponding to the known coefficient we require. In other words, we may multiply the lighting for each order l by some amount (which may be different for different frequencies l) while dividing the BRDF by the same amount. However, there is an important additional physical constraint. The BRDF must be reciprocal, i.e. symmetric with respect to incident and outgoing angles. The corresponding condition in the frequency domain is that the BRDF coefficients must be symmetric with respect to interchange of the indices corresponding to the incident and outgoing directions. To take advantage of this symmetry, we will use the reciprocal form of the frequency-space equations, as defined in equation 2.59.

We now derive an analytic formula for the lighting and BRDF in terms of coefficients of the reflected light field. Since we cannot recover the global scale, we will arbitrarily scale the DC term of the lighting so $L_{00} = \Lambda_0^{-1} = \sqrt{1/(4\pi)}$. Note that this scaling is valid unless the DC term is 0, corresponding to no light—an uninteresting case. Using equations 2.59, 2.64, and 2.65, we obtain

$$\begin{aligned}
L_{00} &= \Lambda_0^{-1} && : \text{Global Scale} \\
\tilde{\rho}_{0p0} &= \tilde{B}_{00p0} && : \text{Equation 2.64 } (l = q = 0) \\
L_{lm} &= \Lambda_l^{-1} \frac{\tilde{B}_{lmpq}}{\tilde{\rho}_{lpq}} && : \text{Equation 2.65} \\
&= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{l00}} && : \text{Set } p = q = 0 \\
&= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{0l0}} && : \text{Reciprocity, } \tilde{\rho}_{0l0} = \tilde{\rho}_{l00} \\
&= \Lambda_l^{-1} \frac{\tilde{B}_{lm00}}{\tilde{B}_{00l0}} && : \text{Plug in from 2}^{\text{nd}} \text{ line} \\
\tilde{\rho}_{lpq} &= \Lambda_l^{-1} \frac{\tilde{B}_{lmpq}}{L_{lm}} && : \text{Equation 2.64} \\
&= \frac{\tilde{B}_{lmpq} \tilde{B}_{00l0}}{\tilde{B}_{lm00}} && : \text{Substitute from above for } L_{lm}.
\end{aligned} \tag{2.68}$$

Note that in the last line, any value of m may be used. If none of the terms above vanishes, this gives an explicit formula for the lighting and BRDF in terms of coefficients of the output light field. Assuming reciprocity of the BRDF is critical. Without it, we would not be able to relate $\tilde{\rho}_{0l0}$ and $\tilde{\rho}_{l00}$ above, and we would need a separate scale factor for each frequency l.

Therefore, up to global scale, **the reflected light field can be factored into the lighting and the BRDF**, provided the appropriate coefficients of the reflected light field do not vanish, i.e. the denominators above are nonzero. If the denominators do vanish, the inverse-lighting or inverse-BRDF problems become ill-posed and consequently, the factorization becomes ill-posed. Note that the above relations are one possible factorization formula. We may still be able to factor the light field even if some of the $\tilde{\rho}_{l00}$ terms vanish in equation 2.68, by using different values of $\tilde{\rho}_{lpq}$ with $p \neq 0$.

Of course, the results will be more and more ill-conditioned, the closer the reflected light field coefficients in the denominators come to 0, and so, in practice, there is a maximum frequency up to which the recovery process will be possible. This maximum frequency will depend on the frequency spectrum of the reflected light field, and hence on the frequency spectra of the lighting and BRDF. When either inverse-lighting or inverse-BRDF computations become ill-conditioned, so will the factorization. Therefore, the factorization will work best for specular BRDFs and high-frequency lighting. In other cases, there will remain some ambiguities, or ill-conditioning.

2.5 Conclusions and Future Work

In this chapter, we have presented a theoretical analysis of the structure of the reflected light field from a convex homogeneous object under a distant illumination field. We have shown that the reflected light field can be formally described as a convolution of the incident illumination and the BRDF, and derived an analytic frequency space convolution formula. This means that reflection can be viewed in signal processing terms as a filtering operation between the lighting and the BRDF to produce the output light field. Furthermore, inverse rendering to estimate the lighting or BRDF from the reflected light field can be understood as deconvolution. This result provides a novel viewpoint for many forward and inverse rendering problems, and allows us to understand the duality between forward and inverse problems, wherein an ill-conditioned inverse problem may lead to an efficient solution to a forward problem. We have also discussed the implications for inverse problems such as lighting recovery, BRDF recovery, light field factorization, and forward rendering problems such as environment map prefiltering and rendering. The next chapter will make these ideas concrete for many special cases, deriving analytic formulae for the frequency spectra of many common BRDF and lighting models. Following that, the rest of this dissertation will develop practical implications of the theoretical analysis from this chapter, showing how frequency domain methods may be used for forward and inverse rendering.

It should be noted that the analysis in this chapter is based on the specific assumptions noted here, and is only one way in which the reflection operator can be analyzed. Other researchers have derived analytic formulae for many useful special cases that go beyond our assumptions. For instance, Soler and Sillion [80] derive a convolution relation for calculating soft shadows assuming planar objects. Arvo [1] and Chen and Arvo [9] develop methods for computing irradiance from planar luminaires including near-field effects which we do not treat here. In the future, it would be interesting to consider perturbative methods that could unify our results with some of these previous analytical derivations.

More generally, we have studied the computational properties of the reflection operator—given a complex illumination field and arbitrary BRDF—in the frequency domain. However, there are many other ways these *computational fundamentals of reflection* can be studied. For instance, it might be worthwhile to consider the differential properties of reflection, and to study perceptual metrics rather than physical ones. Another important area is the formal study of the conditioning of forward and inverse problems, possibly directly from an eigenanalysis of the kernel of the Fredholm integral equation. We believe this formal analysis will be increasingly important in deriving robust and efficient algorithms in the future. While we have made a first step in this direction, other issues such as how our results change when we have only a limited fraction of the reflected light field available, or can move our viewpoint only in a narrow range, need to be studied. In summary, we believe there are a number of domains in graphics and vision that benefit greatly from a fundamental understanding of the reflection operator. We believe the work described in this chapter is a first step in putting an analysis of reflection on a strong mathematical foundation.

Chapter 3

Formulae for Common Lighting and BRDF Models

An analysis of the computational properties of the reflection operator is of interest in both computer graphics and vision, for analyzing forward and inverse problems in rendering. Building on previous qualitative observations by Miller and Hoffman [59], Cabral et al. [7, 8], D'Zmura [17] and others, the previous chapter formalized the notion of reflection as a spherical convolution of the incident illumination and the BRDF.

Specifically, we were able to develop a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. Under these assumptions, we were able to derive an analytic formula for the reflected light field in terms of the product of spherical harmonic coefficients of the BRDF and the lighting. Our formulation allows us to view forward rendering as *convolution* and inverse rendering as *deconvolution*.

In this chapter, we will primarily be concerned with the well-posedness and numerical conditioning of inverse problems. We analytically derive the spherical harmonic coefficients for many common lighting and BRDF models. In this way, we analyze the well-posedness and conditioning of a number of inverse problems, explaining many previous empirical observations. This analysis is also of interest for forward rendering, since an ill-conditioned inverse problem corresponds to a forward problem where the final results are not sensitive to certain components of the initial conditions, allowing for efficient approximations to be made.

The rest of this chapter is organized as follows. In section 1, we briefly summarize the main results from the previous chapter that we will use here. Section 2 is the main part of this chapter, and works out analytic formulae for spherical harmonic coefficients of many common lighting and BRDF models, demonstrating the implications of the theoretical analysis. This section is a more detailed version of the derivations in our previous SIGGRAPH paper [73], and includes verification of the spherical harmonic formulae from first principles, as well as a discussion of light field factorization for the special cases of interest. Finally, section 3 concludes this chapter and discusses future work.

3.1 Background

In this section, we briefly summarize the main theoretical results derived in the previous chapter, introducing the notation and terminology required in the next section. This section may be skipped by readers familiar with the previous chapter of the dissertation. In this chapter, we will only discuss results in 3D, since that is of greater practical importance. For simplicity, we will also restrict ourselves to isotropic BRDFs. A more complete derivation of the convolution formula, along with a number of alternative forms, is found in the previous chapter. Notation used in this and the previous chapter is in table 1.1.

3.1.1 Reflection Equation and Convolution Formula

The assumptions we are making here are convex curved homogeneous isotropic surfaces under a distant illumination field. Under these circumstances, the reflection equation can be written as (c.f. equation 2.12)

$$B(\alpha,\beta,\theta_o',\phi_o') = \int_{\Omega_i'} L\left(R_{\alpha,\beta}(\theta_i',\phi_i')\right) \hat{\rho}(\theta_i',\phi_i',\theta_o',\phi_o')\, d\omega_i'. \tag{3.1}$$

It is possible to derive a frequency-space convolution formula corresponding to equation 3.1. For this purpose, we must expand quantities in terms of spherical harmonics.

Specifically, the illumination can be written as (c.f. equation 2.52)

$$\begin{aligned} L(\theta_i, \phi_i) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} Y_{lm}(\theta_i, \phi_i) \\ L(\theta_i, \phi_i) = L\left(R_{\alpha,\beta}(\theta_i', \phi_i')\right) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{+l} \sum_{m'=-l}^{l} L_{lm} D^l_{mm'}(\alpha, \beta) Y_{lm'}(\theta_i', \phi_i'). \end{aligned} \quad (3.2)$$

Then, we write the expansion of the isotropic BRDF (c.f. equation 2.53),

$$\hat{\rho}(\theta_i', \theta_o', \mid \phi_o' - \phi_i' \mid) = \sum_{l=0}^{\infty} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} \hat{\rho}_{lpq} Y^*_{lq}(\theta_i', \phi_i') Y_{pq}(\theta_o', \phi_o'). \quad (3.3)$$

The reflected light field, which is now a 4D function, can be expanded using a product of representation matrices and spherical harmonics (c.f. equation 2.54),

$$\begin{aligned} C_{lmpq}(\alpha, \beta, \theta_o', \phi_o') &= \Lambda_l^{-1} D^l_{mq}(\alpha, \beta) Y_{pq}(\theta_o', \phi_o') \\ B(\alpha, \beta, \theta_o', \phi_o') &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} B_{lmpq} C_{lmpq}(\alpha, \beta, \theta_o', \phi_o'). \end{aligned} \quad (3.4)$$

Finally, we can derive an analytic expression (convolution formula) for the reflection equation in terms of these coefficients (c.f. equation 2.55).

$$B_{lmpq} = \Lambda_l L_{lm} \hat{\rho}_{lpq}. \quad (3.5)$$

We may also derive an alternative form, holding the outgoing elevation angle θ_o' fixed (c.f. equation 2.57),

$$B_{lmq}(\theta_o') = \Lambda_l L_{lm} \hat{\rho}_{lq}(\theta_o'). \quad (3.6)$$

If we seek to preserve the *reciprocity* of the BRDF, i.e. symmetry with respect to incident and outgoing angles, we may multiply both the transfer function and the reflected light field by $\cos\theta_o'$, defining (c.f. equation 2.58)

$$\begin{aligned} \tilde{\rho} &= \hat{\rho} \cos\theta_o' = \rho \cos\theta_i' \cos\theta_o' \\ \tilde{B} &= B \cos\theta_o'. \end{aligned} \quad (3.7)$$

With these derivations, equation 3.5 becomes (c.f. equation 2.59)

$$\tilde{B}_{lmpq} = \Lambda_l L_{lm} \tilde{\rho}_{lpq}. \tag{3.8}$$

The symmetry of the transfer function ensures that its coefficients are unchanged if the indices corresponding to incident and outgoing angles are interchanged, i.e. $\tilde{\rho}_{lpq} = \tilde{\rho}_{plq}$.

Many models, such as Lambertian and Phong BRDFs are *radially symmetric or 1D BRDFs*, where the BRDF consists of a single symmetric lobe of fixed shape, whose orientation depends only on a well-defined central direction $\vec{C}$. If we reparameterize by $\vec{C}$, the BRDF becomes a function of only 1 variable (θ_i' with $\cos\theta_i' = \vec{C} \cdot \vec{L}$) instead of 3. In this case, we may write the BRDF and equations for the reflected light field as (c.f. equation 2.60)

$$\begin{aligned} \hat{\rho}(\theta_i') &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta_i') \\ B(\alpha,\beta) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm} Y_{lm}(\alpha,\beta). \end{aligned} \tag{3.9}$$

The required convolution formula now becomes (c.f. equation 2.62)

$$B_{lm} = \Lambda_l \hat{\rho}_l L_{lm}. \tag{3.10}$$

3.1.2 Analysis of Inverse Problems

The convolution formula in equation 3.5 (or equation 2.55) can be used to analyze the well-posedness and numerical conditioning of inverse problems. For the inverse-BRDF problem, we manipulate equation 3.5 to yield (c.f. equation 2.64)

$$\hat{\rho}_{lpq} = \Lambda_l^{-1} \frac{B_{lmpq}}{L_{lm}}. \tag{3.11}$$

In general, BRDF estimation will be well-posed, i.e. unambiguous as long as the denominator on the right-hand side does not vanish.

A similar analysis can be done for estimation of the lighting (c.f. equation 2.65),

$$L_{lm} = \Lambda_l^{-1} \frac{B_{lmpq}}{\hat{\rho}_{lpq}}. \tag{3.12}$$

Inverse lighting will be well-posed so long as the denominator does not vanish for all p, q for some l, i.e. so long as the spherical harmonic expansion of the BRDF transfer function contains all orders.

Finally, we can put these results together to derive an analytic formula for *factoring* the reflected light field, i.e. determining both the lighting and BRDF in terms of coefficients of the reflected light field. See equation 2.68 for details. We are able to show that up to global scale, **the reflected light field can be factored into the lighting and the BRDF**, provided the appropriate coefficients of the reflected light field do not vanish.

The next section will derive analytic formulae for the frequency spectrum of common lighting and BRDF models, explaining the implications for the well-posedness and conditioning of inverse problems in terms of the results stated above.

3.2 Derivation of Analytic Formulae

This section discusses the implications of the theoretical analysis developed in the previous section. Our main focus will be on understanding the well-posedness and conditioning of inverse problems. We consider a number of lighting distributions and BRDFs, deriving analytic formulae and approximations for their spherical harmonic spectra. From this analysis, we quantitatively determine the well-posedness and conditioning of inverse problems associated with these illumination conditions and BRDFs. Below, we first consider three lighting conditions—a single directional source, an axially symmetric configuration, and uniform lighting. Then, we consider four BRDF models—a mirror surface, a Lambertian surface, the Phong BRDF, and a microfacet model.

3.2.1 Directional Source

Our first example concerns a single directional source. The lighting is therefore described by a delta function in spherical coordinates. Let (θ_s, ϕ_s) refer to the angular coordinates of the source. Then,

$$\begin{aligned} L(\theta_i, \phi_i) &= \delta\left(\cos\theta_i - \cos\theta_s\right)\delta\left(\phi_i - \phi_s\right) \\ L_{lm} &= \int_0^{2\pi}\int_0^{\pi} \delta\left(\cos\theta_i - \cos\theta_s\right)\delta\left(\phi_i - \phi_s\right) Y_{lm}^*(\theta_i, \phi_i)\, \sin\theta_i\, d\theta_i d\phi_i \\ &= Y_{lm}^*(\theta_s, \phi_s). \end{aligned} \quad (3.13)$$

Note that in the equation above, the delta function has the correct form for spherical coordinates. The same form will be used later to study mirror BRDFs.

It will simplify matters to reorient the coordinate system so that the source is at the north pole or +Z, i.e. $\theta_s = 0$. It is now straightforward to write

$$\begin{aligned} L_{lm} &= Y_{lm}^*(0) \\ &= \Lambda_l^{-1}\delta_{m0} \\ B_{lmpq} &= \delta_{m0}\hat{\rho}_{lpq} \\ \hat{\rho}_{lpq} &= B_{l0pq}. \end{aligned} \quad (3.14)$$

In angular space, a single observation of the reflected light field corresponds to a single BRDF measurement. This property is used in image-based BRDF measurement [51, 55]. We see that in frequency space, there is a similar straightforward relation between BRDF coefficients and reflected light field coefficients. BRDF recovery is well-posed and well-conditioned since we are estimating the BRDF filter from its impulse response.

It is instructive to verify equation 3.14 directly from first principles. We first note that in angular space, the reflected light field for a directional source at the north pole can be written as

$$B(\alpha, \beta, \theta_o', \phi_o') = \hat{\rho}(\alpha, \pi, \theta_o', \phi_o'). \quad (3.15)$$

Note that the incident angle for surface normal (α, β) is given by (α, π). Clearly the elevation angle of incidence must be α, and because of our standard right-handed sign convention, the azimuthal angle is π. We will now show that equation 3.14 is simply a frequency-space version of equation 3.15 by expanding out B and $\hat{\rho}$, using the expressions in equation 3.14. We will need to use the first property of the representation matrices from equation 2.35. The first line below simply derives the form of C_{l0pq}, making use of equation 2.35. In the next line, we expand the left hand side of equation 3.15, $B(\alpha, \beta, \theta'_o, \phi'_o)$ in terms of C_{l0pq}. Note that since the directional source is at the north pole, there is no azimuthal dependence and we can assume that $m = 0$. Finally, we expand the right hand side of equation 3.15, and equate coefficients.

$$\begin{aligned}
C_{l0pq}(\alpha, \beta, \theta'_o, \phi'_o) &= \Lambda_l^{-1} D^l_{0q}(\alpha, \beta) Y_{pq}(\theta'_o, \phi'_o) \\
&= Y^*_{lq}(\alpha, \pi) Y_{pq}(\theta'_o, \phi'_o) \\
B(\alpha, \beta, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} B_{l0pq} Y^*_{lq}(\alpha, \pi) Y_{pq}(\theta'_o, \phi'_o) \\
\hat{\rho}(\alpha, \pi, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} \hat{\rho}_{lpq} Y^*_{lq}(\alpha, \pi) Y_{pq}(\theta'_o, \phi'_o) \\
B_{l0pq} &= \hat{\rho}_{lpq}.
\end{aligned} \tag{3.16}$$

The last line comes from equating coefficients, and this confirms the correctness of equation 3.14, thereby verifying the convolution formula for the special case of a single directional source.

Finally, we may verify the factorization relations of equation 2.68 for the case when both the BRDF and lighting are unknown *a priori*, and the lighting is actually a single directional light source.

$$\begin{aligned}
L_{lm} &= \Lambda_l^{-1} \frac{\tilde{B}_{lm00}}{\tilde{B}_{00l0}} \\
&= \Lambda_l^{-1} \delta_{m0} \left(\frac{\tilde{\rho}_{l00}}{\tilde{\rho}_{0l0}} \right) \\
\tilde{\rho}_{lpq} &= \frac{B_{l0pq} B_{00l0}}{B_{l000}} \\
&= \tilde{\rho}_{lpq} \left(\frac{\tilde{\rho}_{0l0}}{\tilde{\rho}_{l00}} \right).
\end{aligned} \tag{3.17}$$

We see that these relations give the correct answer if the BRDF obeys reciprocity, and provided the appropriate BRDF coefficients do not vanish. If the BRDF coefficients do vanish, the factorization is ill-posed since there is an ambiguity about whether the lighting or BRDF coefficient is 0. This is related to the associativity of convolution.

In summary, we have derived an analytic frequency-space formula, that has been verified from first principles. A directional source corresponds to a delta function, whose frequency spectrum does not decay. Therefore, BRDF estimation is well-posed and well-conditioned for all frequencies. In effect, we are estimating the BRDF filter from its impulse response. This is a frequency-space explanation for the use of directional sources in BRDF measurement, especially in the newer image-based methods for curved surfaces [51, 55].

3.2.2 Axially Symmetric Distribution

We now consider a lighting distribution that is symmetric about some axis. For convenience, we position the coordinate system so that the Z axis is the axis of symmetry. Such a distribution closely approximates the illumination due to skylight on a cloudy day. We should also note that a single directional source, as just discussed, is a trivial example of an axially symmetric lighting distribution. The general property of these configurations is that the lighting coefficients L_{lm} with $m \neq 0$ vanish, since they have azimuthal dependence. The frequency-space reflection formulas now become

$$\begin{aligned} B_{lmpq} &= \delta_{m0}\Lambda_l L_{l0}\hat{\rho}_{lpq} \\ \hat{\rho}_{lpq} &= \Lambda_l^{-1}\frac{B_{l0pq}}{L_{l0}}. \end{aligned} \tag{3.18}$$

It is important to note that the property of axial symmetry is preserved in the reflected light field, since $m = 0$. The remaining properties are very similar to the general case. In particular, BRDF estimation is well conditioned if the frequency spectrum of the illumination does not rapidly decay, i.e. there are sharp variations with respect to the elevation angle (there is no variation with respect to the azimuthal angle).

3.2.3 Uniform Lighting

Our final lighting configuration is that of uniform lighting. This can be considered the canonical opposite to the case of a directional source. Note that a uniform lighting distribution is symmetric about any axis and is therefore a special case of axially symmetric lighting. For uniform lighting, only the DC term of the lighting does not vanish, i.e. we set $L_{00} = \Lambda_0^{-1}$, with other coefficients being zero. The relevant frequency-space reflection formulas become

$$\begin{aligned} L_{lm} &= \delta_{l0}\delta_{m0}\Lambda_0^{-1} \\ B_{lmpq} &= \delta_{l0}\delta_{m0}\delta_{q0}\hat{\rho}_{0p0} \\ \hat{\rho}_{0p0} &= B_{00p0}. \end{aligned} \tag{3.19}$$

Note that $q = 0$ since $l = 0$ and $\mid q \mid \leq l$. We may only find the BRDF coefficients $\hat{\rho}_{0p0}$; the other coefficients cannot be determined. In other words, we can determine only the $0-$order coefficients ($l = 0$). This is because the input signal has no amplitude along higher frequencies, making it impossible to estimate these higher frequencies of the BRDF filter. A subtle point to be noted is that reciprocity (symmetry) of the BRDF can actually be used to double the number of coefficients known, but the problem is still extremely ill-posed.

All this may be somewhat clearer if we use the convolution formula from equations 3.6 or 2.57, where the dependence on the ougoing elevation angle is not expanded into basis functions,

$$B_{lmq}(\theta_o') = \delta_{l0}\delta_{m0}\delta_{q0}\hat{\rho}_{00}(\theta_o'). \tag{3.20}$$

This formulation makes it clear that only the lowest frequency, i.e. DC term of the transfer function contributes to the reflected light field, with B being independent of surface orientation. From observations made under uniform lighting, we can estimate only the DC term of the BRDF transfer function; higher frequencies in the BRDF cannot be estimated. Hence, a mirror surface cannot be distinguished from a Lambertian object under uniform lighting.

We may also verify equation 3.20 by writing it in terms of angular space coordinates.

First, we note the form of the basis functions, noting that $Y_{00} = \Lambda_0^{-1}$ and that $D_{00}^0 = 1$,

$$
\begin{aligned}
C_{000} &= \sqrt{\frac{1}{2\pi}}\Lambda_0^{-1} \\
\hat{\rho}_{00}(\theta_o') &= \sqrt{\frac{1}{2\pi}}\int_0^{2\pi}\int_0^{2\pi}\int_0^{\pi/2} \hat{\rho}(\theta_i', \theta_o', \mid \phi_o' - \phi_i' \mid) \sin\theta_i'\, d\theta_i' d\phi_i' d\phi_o' \\
&= \sqrt{2\pi}\int_0^{2\pi}\int_0^{\pi/2} \hat{\rho}(\theta_i', \theta_o', \mid \phi \mid)\, \sin\theta_i'\, d\theta_i' d\phi. \qquad (3.21)
\end{aligned}
$$

In the last line, we have set $\phi = \phi_o' - \phi_i'$ and integrated over ϕ_o', obtaining a factor of 2π. Substituting in the expansion of equation 3.20, we obtain

$$
\begin{aligned}
B &= B_{000}(\theta_o')C_{000} \\
&= \hat{\rho}_{00}(\theta_o')\sqrt{\frac{1}{2\pi}}\Lambda_0^{-1} \\
&= \Lambda_0^{-1}\int_0^{2\pi}\int_0^{\pi/2} \hat{\rho}(\theta_i', \theta_o', \mid \phi \mid)\, \sin\theta_i'\, d\theta_i' d\phi, \qquad (3.22)
\end{aligned}
$$

which is the expected angular-space result, since the lighting is constant and equal to Λ_0^{-1} everywhere. Thus, we simply integrate over the BRDF.

We next consider factorization of the light field to estimate both the lighting and the BRDF. An examination of the formulas in equation 2.68 shows that we will indeed be able to estimate all the lighting coefficients, provided the BRDF terms $\tilde{\rho}_{0p0}$ do not vanish. We will thus be able to determine that the lighting is uniform, i.e. that all the lighting coefficients L_{lm} vanish unless $l = m = 0$. However, once we do this, the factorization will be ill-posed since BRDF estimation is ill-posed. To summarize, we have derived and verified analytic formulae for the case of uniform lighting. BRDF estimation is extremely ill-posed, and only the lowest-frequency terms of the BRDF can be found. Under uniform lighting, a mirror surface and a Lambertian surface will look identical, and cannot be distinguished. With respect to factoring the light field when the lighting is uniform, we will correctly be able to determine that the illumination is indeed uniform, but the factorization will remain ill-posed since further information regarding the BRDF cannot be obtained. In signal processing terms, the input signal is constant and therefore has no amplitude along nonzero frequencies of the BRDF filter. Therefore, these nonzero frequencies of the BRDF filter cannot be estimated.

We will now derive analytic results for four different BRDF models. We start with the mirror BRDF and Lambertian surface, progressing to the Phong model and the microfacet BRDF.

3.2.4 Mirror BRDF

A mirror BRDF is analogous to the case of a directional source. A physical realization of a mirror BRDF is a gazing sphere, commonly used to recover the lighting. For a mirror surface, the BRDF is a delta function and the coefficients can be written as

$$\begin{aligned}
\hat{\rho} &= \delta\left(\cos\theta_o' - \cos\theta_i'\right)\delta\left(\phi_o' - \phi_i' \pm \pi\right) \\
\hat{\rho}_{lpq} &= \int_0^{2\pi}\int_0^{2\pi}\int_0^{\pi/2}\int_0^{\pi/2} \hat{\rho}(\theta_i', \phi_i', \theta_o', \phi_o') Y_{lq}(\theta_i', \phi_i') Y_{pq}^*(\theta_o', \phi_o') \sin\theta_i' \sin\theta_o' \, d\theta_i' d\theta_o' d\phi_i' d\phi_o' \\
&= \int_0^{2\pi}\int_0^{\pi/2} Y_{pq}^*(\theta_i', \phi_i' \pm \pi) Y_{lq}(\theta_i', \phi_i') \, \sin\theta_i' \, d\theta_i' d\phi_i' \\
&= (-1)^q \delta_{lp}.
\end{aligned} \tag{3.23}$$

The factor of $(-1)^q$ in the last line comes about because the azimuthal angle is phase shifted by π. This factor would not be present for retroreflection. Otherwise, the BRDF coefficients simply express in frequency-space that the incident and outgoing angles are the same, and show that the frequency spectrum of the BRDF does not decay with increasing order.

The reflected light field and BRDF are now related by

$$\begin{aligned}
B_{lmpq} &= \Lambda_l (-1)^q \delta_{lp} L_{lm} \\
L_{lm} &= \Lambda_l^{-1} (-1)^q B_{lmlq}.
\end{aligned} \tag{3.24}$$

Just as the inverse lighting problem from a mirror sphere is easily solved in angular space, it is well-posed and easily solved in frequency space because there is a direct relation between BRDF and reflected light field coefficients. In signal processing terminology, the inverse lighting problem is well-posed and well-conditioned because the frequency spectrum of a delta function remains constant with increasing order. This is a frequency-space explanation for the use of gazing spheres for estimating the lighting.

It is easy to verify equation 3.24 from first principles. We may write down the appropriate equations in angular space and then expand them in terms of spherical harmonics.

$$
\begin{aligned}
B(\alpha, \beta, \theta'_o, \phi'_o) &= L\left(R_{\alpha,\beta}(\theta'_o, \phi'_o \pm \pi)\right) \\
&= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} Y_{lm}\left(R_{\alpha,\beta}(\theta'_o, \phi'_o \pm \pi)\right) \\
&= \Lambda_l (-1)^q \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} \left(\Lambda_l^{-1} D^l_{mq}(\alpha, \beta) Y_{lq}(\theta'_o, \phi'_o)\right),
\end{aligned}
\tag{3.25}
$$

from which equation 3.24 follows immediately.

Factorization: The result for factoring a light field with a mirror BRDF is interesting. We first note that, unlike in our previous examples, there is no need to perform the symmetrizing transformations in equation 3.7 since the transfer function is already symmetric with respect to indices l and p ($\delta_{lp} = \delta_{pl}$). We next note that equation 2.68 seems to indicate that the factorization is very ill-posed. Indeed, both the denominators, B_{00l0} and B_{lm00}, vanish unless $l = 0$. In fact, it is possible to explain the ambiguity by constructing a suitable BRDF,

$$
\begin{aligned}
\hat{\rho}_l &= f(l)(-1)^q \delta_{lp} \\
\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{q=-l}^{l} f(l) Y^*_{lq}(\theta'_i, \phi'_i) Y_{lq}(\theta'_o, \phi'_o),
\end{aligned}
\tag{3.26}
$$

where $f(l)$ is an arbitrary function of frequency. When $f(l) = 1$, we obtain a mirror BRDF. However, for any $f(l)$, we get a valid BRDF that obeys reciprocity. Physically, the BRDF acts like a mirror except that the reflectance is different for different frequencies, i.e. it may pass through high frequencies like a perfect mirror but attenuate low frequencies, while still reflecting them about the mirror direction. It is not clear that this is a realistic BRDF model, but it does not appear to violate any obvious physical principles.

It is now easy to see why factorization does not work. The function $f(l)$ cannot be determined during factorization. Without changing the reflected light field, we could multiply the coefficients L_{lm} of the lighting by $f(l)$, setting the BRDF to a mirror. In other words, there is a separate global scale for each frequency l that we cannot estimate. Reciprocity of the BRDF is not much help here, since only the "diagonal" terms of the frequency spectrum are nonzero.

Note that if the lighting coefficients do not vanish, we will indeed be able to learn that the BRDF has the form in equation 3.26. However, we will not be able to make further progress without additional assumptions about the form of the BRDF, i.e. of the function $f(l)$. In certain applications, we may want to turn this ambiguity to our advantage by selecting $f(l)$ appropriately to give a simple form for the BRDF or the lighting, without affecting the reflected light field. The ambiguity, and its use in simplifying the form for the reflected light field, are common to many reflective BRDFs, and we will encounter this issue again for the Phong and microfacet models.

Reparameterization by Reflection Vector: For reflective BRDFs, it is often convenient to reparameterize by the reflection vector, as discussed in section 2.3.4, or at the end of section 3.1.1. The transfer function can then be written simply as a function of the incident angle (with respect to the reflection vector), and is still a delta function. Since there is no dependence on the outgoing angle after reparameterization, we obtain

$$\begin{aligned} \hat{\rho}(\theta_i') &= \frac{\delta\left(\cos(0) - \cos\theta_i'\right)}{2\pi} \\ \hat{\rho}_l &= Y_{l0}(0) = \Lambda_l^{-1} \\ B_{lm} &= L_{lm}. \end{aligned} \tag{3.27}$$

In the top line, the factor of 2π in the denominator is to normalize with respect to the azimuthal angle. The bottom line follows from the identity for mirror BRDFs that $\Lambda_l \hat{\rho}_l = 1$.

Therefore, we see that after reparameterization by the reflection vector, the BRDF frequency spectrum becomes particularly simple. The reflected light field corresponds directly to the incident lighting. The BRDF filter just passes through the incident illumination, and the reflected light field is therefore just an image of the lighting without filtering or attenuation. Hence, the illumination can be trivially recovered from a mirror sphere.

3.2.5 Lambertian BRDF

For a Lambertian surface, the BRDF is a constant, corresponding to the surface albedo—which for purposes of energy conservation must not be greater than $1/\pi$. For simplicity, we will omit this constant in what follows. The transfer function is a *clamped cosine* since it is equal to the cosine of the incident angle over the upper hemisphere when $\cos\theta_i' > 0$ and is equal to 0 over the lower hemisphere when $\cos\theta_i' < 0$. A graph of this function, along with spherical harmonic approximations to it up to order 4 is in figure 3.1.

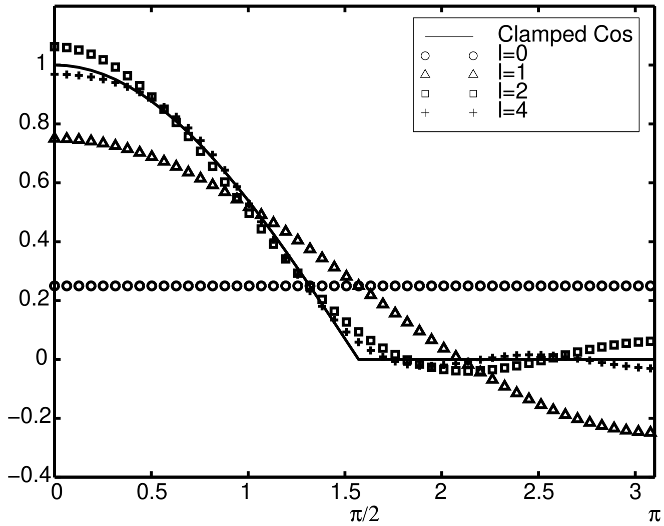

Figure 3.1: *The* clamped cosine *filter corresponding to the Lambertian BRDF and successive approximations obtained by adding more spherical harmonic terms. For $l = 2$, we already get a very good approximation.*

Since the reflected light field is proportional to the incident irradiance and is equal for all outgoing directions, we will drop the outgoing angular dependence and use the form of the convolution formula given in equation 3.10,

$$\begin{aligned} \hat{\rho} = \max(\cos\theta_i', 0) &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta_i') \\ B_{lm} &= \Lambda_l \hat{\rho}_l L_{lm}. \end{aligned} \tag{3.28}$$

It remains to derive the form of the spherical harmonic coefficients $\hat{\rho}_l$. To derive the spherical harmonic coefficients for the Lambertian BRDF, we must represent the transfer function $\hat{\rho}(\theta'_i) = \max(\cos\theta'_i, 0)$ in terms of spherical harmonics.

We will need to use many formulas for representing integrals of spherical harmonics, for which a reference [52] will be useful. The spherical harmonic coefficients are given by

$$\hat{\rho}_l = 2\pi \int_0^{\frac{\pi}{2}} Y_{l0}(\theta'_i) \cos\theta'_i \sin\theta'_i \, d\theta'_i. \tag{3.29}$$

The factor of 2π comes from integrating 1 over the azimuthal dependence. It is important to note that the limits of the integral range from 0 to $\pi/2$ and not π because we are considering only the upper hemisphere. The expression above may be simplified by writing in terms of Legendre polynomials $P(\cos\theta'_i)$. Putting $u = \cos\theta'_i$ in the above integral and noting that $P_1(u) = u$ and that $Y_{l0}(\theta'_i) = \Lambda_l^{-1} P_l(\cos\theta'_i)$, we obtain

$$\hat{\rho}_l = 2\pi\Lambda_l^{-1} \int_0^1 P_l(u) P_1(u) \, du. \tag{3.30}$$

To gain further insight, we need some facts regarding the Legendre polynomials. P_l is odd if l is odd, and even if l is even. The Legendre polynomials are orthogonal over the domain $[-1, 1]$ with the orthogonality relationship being given by

$$\int_{-1}^1 P_a(u) P_b(u) = \frac{2}{2a+1} \delta_{a,b}. \tag{3.31}$$

From this, we can establish some results about equation 3.30. When l is equal to 1, the integral evaluates to half the norm above, i.e. $1/3$. When l is odd but greater than 1, the integral in equation 3.30 vanishes. This is because, for $a = l$ and $b = 1$, we can break the left-hand side of equation 3.31 using the oddness of a and b into two equal integrals from $[-1, 0]$ and $[0, 1]$. Therefore, both of these integrals must vanish, and the latter integral is the right-hand integral in equation 3.30. When l is even, the required formula is given by manipulating equation 20 in chapter 5 of MacRobert[52]. Putting it all together, we obtain,

$$\begin{aligned}
l = 1 \quad \hat{\rho}_l &= \sqrt{\frac{\pi}{3}} \\
l > 1, odd \quad \hat{\rho}_l &= 0 \\
l \;\; even \quad \hat{\rho}_l &= 2\pi\sqrt{\frac{2l+1}{4\pi}}\frac{(-1)^{\frac{l}{2}-1}}{(l+2)(l-1)}\left[\frac{l!}{2^l(\frac{l}{2}!)^2}\right].
\end{aligned} \tag{3.32}$$

This formula is plotted in figure 3.2.

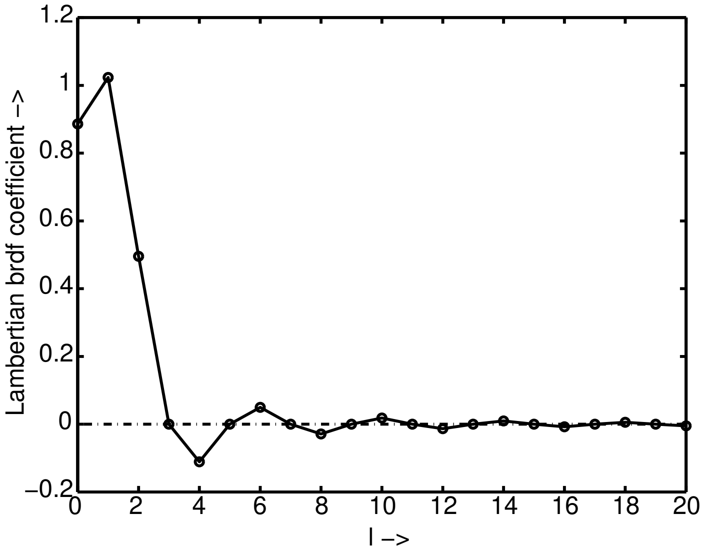

Figure 3.2: *The solid line is a plot of $\hat{\rho}_l$ versus l, as per equation 3.32. It can be seen that odd terms with $l > 1$ have $\hat{\rho}_l = 0$. Also, as l increases, the BRDF coefficients decay rapidly.*

There are two important points to note here. Firstly, the transfer function is identically zero for odd frequencies greater than 1. Secondly, it can be shown by applying Stirling's formula, that the bracketed term falls off asymptotically as $1/\sqrt{l}$, cancelling the square root. Therefore, $\hat{\rho}_l \sim l^{-2}$. The reflected light field therefore falls off as $\Lambda_l \hat{\rho}_l \sim l^{-5/2}$. This rapid falloff means the Lambertian BRDF effectively behaves like a low-pass filter, letting only the lowest frequencies of the lighting through. Numerically, we may write the first few terms for the BRDF filter as

$$
\begin{aligned}
\Lambda_0\hat{\rho}_0 &= 3.142 \\
\Lambda_1\hat{\rho}_1 &= 2.094 \\
\Lambda_2\hat{\rho}_2 &= 0.785 \\
\Lambda_3\hat{\rho}_3 &= 0 \\
\Lambda_4\hat{\rho}_4 &= -0.131 \\
\Lambda_5\hat{\rho}_5 &= 0 \\
\Lambda_6\hat{\rho}_6 &= 0.049.
\end{aligned}
\tag{3.33}
$$

We see that already for $l = 4$, the coefficient is only about 4% of what it is for $l = 0$. In fact, it can be shown that over 99% of the energy of the BRDF filter is captured by $l \leq 2$. By considering the fact that the lighting must remain positive everywhere [2], similar worst-case bounds can be shown for the approximation of the reflected light field by $l \leq 2$. **Therefore, the irradiance, or equivalently, the reflected light field from a Lambertian surface can be well approximated using only the first 9 terms of its spherical harmonic expansion**—1 term with order 0, 3 terms with order 1, and 5 terms with order 2. Note that the single order 0 mode Y_{00} is a constant, the 3 order 1 modes are linear functions of the Cartesian coordinates—in real form, they are simply x, y, and z—while the 5 order 2 modes are quadratic functions of the Cartesian coordinates. Therefore, the irradiance can be well approximated as a quadratic polynomial of the Cartesian coordinates of the surface normal vector.

We first consider illumination estimation, or the inverse lighting problem. The fact that the odd frequencies greater than 1 of the BRDF vanish means the inverse lighting problem is formally ill-posed for a Lambertian surface. The filter zeros the odd frequencies of the input signal, so these terms cannot be estimated from images of a convex Lambertian object. This observation corrects a commonly held notion (see Preisendorfer [67], volume 2, pages 143–151) that radiance and irradiance are equivalent in the sense that irradiance can be formally inverted to recover the radiance. Different radiance distributions can give rise to the same irradiance distribution. For further details, see [72]. Moreover, in practical applications, we can robustly estimate only the first 9 coefficients of the incident illumination, those with $l \leq 2$. Thus, inverse lighting from a Lambertian surface is not just formally ill-posed for odd frequencies, but very ill-conditioned for even frequencies. This result explains the ill-conditioning observed by Marschner and Greenberg [54] in estimating the lighting from a surface assumed to be Lambertian.

The 9 parameter approximation also gives rise to a simple algorithm (described in detail in chapter 6) for estimating the illumination at high angular resolution from surfaces having both diffuse and specular components. The diffuse component of the reflected light field is subtracted out using the 9 parameter approximation for Lambertian surfaces. The object is then treated as a gazing sphere, with the illumination recovered from the specular component alone. A consistency condition ensures that the high frequency lighting recovered from the specular BRDF component is indeed consistent with the low frequency lighting used to subtract out the diffuse component of the reflected light field.

Our results are also in accordance with the perception literature, such as Land's retinex theory [46]. It is common in visual perception to associate lighting effects with low frequency variation, and texture with high frequency variation. Our results formalize this observation, showing that distant lighting effects can produce only low frequency variation, with respect to orientation, in the intensity of a homogeneous convex Lambertian surface. Therefore, it should be possible to estimate high frequency texture independently of the lighting. However, for accurate computational estimates as required for instance in computer graphics, there is still an ambiguity between low frequency texture and lighting-related effects.

Since the approximation of Lambertian surfaces is commonly used in graphics and vision, the above results are of interest for many other problems. For instance, we can simply use the first 9 terms of the lighting to compute the irradiance, i.e. the shading on a Lambertian surface. Further implementation details on computing and rendering with *irradiance environment maps* are found in chapter 4. The 9 parameter approximation also means that images of a diffuse object under all possible illumination conditions lie close to a 9D subspace. This is a step toward explaining many previous empirical observations of the low-dimensional effects of lighting made in the computer vision community, such as by Hallinan [30] and Epstein et al. [18]. Basri and Jacobs [2] have derived a Lambertian formula similar to ours, and have applied this result to lighting invariant recognition, and more recently to photometric stereo under general unknown lighting [3].

3.2.6 Phong BRDF

The normalized Phong transfer function is

$$\hat{\rho} = \frac{s+1}{2\pi} \left(\vec{R} \cdot \vec{L} \right)^s , \tag{3.34}$$

where $\vec{R}$ is the reflection vector, $\vec{L}$ is the direction to the light source, and s is the *shininess*, or Phong exponent. The normalization ensures that the Phong lobe has unit energy. Technically, we must also zero the BRDF when the light vector is not in the upper hemisphere. However, the Phong BRDF is not physically based anyway, so others have often ignored this boundary effect, and we will do the same. This allows us to reparameterize by the reflection vector $\vec{R}$, making the transformations outlined in section 2.3.4 or at the end of section 3.1.1. In particular, $\vec{R} \cdot \vec{L} \rightarrow \cos\theta_i'$. Since the BRDF transfer function depends only on $\vec{R} \cdot \vec{L} = \cos\theta_i'$, the Phong BRDF after reparameterization is mathematically analogous to the Lambertian BRDF just discussed (they are both *radially symmetric*). In particular, equation 3.28 holds. However, note that while the Phong BRDF is mathematically analogous to the Lambertian case, it is not physically similar since we have reparameterized by the reflection vector. The BRDF coefficients depend on s, and are given by

$$\begin{aligned} \hat{\rho}_l &= (s+1) \int_0^{\pi/2} [\cos\theta_i']^s \, Y_{l0}(\theta_i') \, \sin\theta_i' \, d\theta_i' \\ B_{lm} &= \Lambda_l \hat{\rho}_l L_{lm}. \end{aligned} \tag{3.35}$$

To solve this integral, we substitute $u = \cos\theta_i'$ in equation 3.35. We also note that $Y_{l0}(\theta_i') = \Lambda_l^{-1} P_l(\cos\theta_i')$, where P_l is the legendre polynomial of order l. Then, equation 3.35 becomes

$$\hat{\rho}_l = \Lambda_l^{-1}(s+1) \int_0^1 u^s P_l(u) \, du. \tag{3.36}$$

An analytic formula is given by MacRobert [52] in equations 19 and 20 of chapter 5,

$$\begin{aligned}
\text{ODD } l \quad \Lambda_l \hat{\rho}_l &= \frac{(s+1)(s-1)(s-3)\dots(s-l+2)}{(s+l+1)(s+l-1)\dots(s+2)} \\
\text{EVEN } l \quad \Lambda_l \hat{\rho}_l &= \frac{s(s-2)\dots(s-l+2)}{(s+l+1)(s+l-1)\dots(s+3)}.
\end{aligned} \tag{3.37}$$

This can be expressed using Euler's Gamma function, which for positive integers is simply the factorial function, $\Gamma(n) = (n-1)!$. Neglecting constant terms, we obtain for large s and $s > l - 1$,

$$\Lambda_l \hat{\rho}_l = \frac{\left[\Gamma\left(\frac{s}{2}\right)\right]^2}{\Gamma\left(\frac{s}{2} - \frac{l}{2}\right)\Gamma\left(\frac{s}{2} + \frac{l}{2}\right)}. \tag{3.38}$$

If $l \ll s$, we can expand the logarithm of this function in a Taylor series about $l = 0$. Using Stirling's formula, we obtain

$$\log(\Lambda_l \hat{\rho}_l) = -l^2 \left(\frac{1}{2s} - \frac{1}{2s^2}\right) + O\left(\frac{l^4}{s^2}\right). \tag{3.39}$$

For large s, $1/s \gg 1/s^2$, and we may derive the approximation

$$\Lambda_l \hat{\rho}_l \approx \exp\left[-\frac{l^2}{2s}\right]. \tag{3.40}$$

The coefficients fall off as a gaussian with width of order $\sqrt{s}$. The Phong BRDF behaves in the frequency domain like a gaussian filter, with the filter width controlled by the shininess. Therefore, inverse lighting calculations will be well-conditioned only up to order $\sqrt{s}$. As s approaches infinity, $\Lambda_l \hat{\rho}_l = 1$, and the frequency spectrum becomes constant, corresponding to a perfect mirror. Note that the frequency domain width of the filter varies inversely with the angular domain extent of the BRDF filter. A plot of the BRDF coefficients and the approximation in equation 3.40 is shown in figure 3.3.

We should also note that for $l > s$, $\hat{\rho}_l$ vanishes if l and s are both odd or both even. It can be shown that for $l \gg s$, the nonzero coefficients fall off very rapidly as $\hat{\rho}_l \sim l^{-(s+1)}$. This agrees with the result for the mathematically analogous Lambertian case, where $s = 1$ and $\hat{\rho}_l \sim l^{-2}$. Note that $s \gg \sqrt{s}$, so $\hat{\rho}_l$ is already nearly 0 when $l \approx s$.

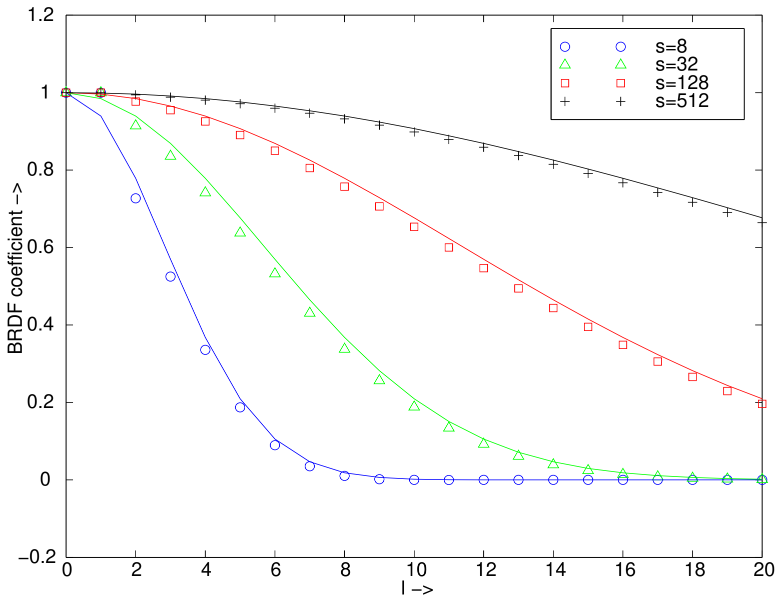

Figure 3.3: *Numerical plots of the Phong coefficients $\Lambda_l \hat{\rho}_l$, as defined by equation 3.37. The solid lines are the gaussian filter approximations in equation 3.40. As the Phong exponent s increases, corresponding to increasing the angular width of the BRDF filter, the frequency width of the BRDF filter decreases.*

Associativity of Convolution: With respect to factorization of light fields with surfaces approximated by Phong BRDFs, we obtain the same results for reflective surfaces as we did for mirror BRDFs. From the form of equation 3.35, it is clear that there is an unrecoverable scale factor for each order l. In physical terms, using the real BRDF and real lighting is equivalent to using a blurred version of the illumination and a mirror BRDF. In signal processing terminology, associativity of convolution allows us to sharpen the BRDF while blurring the illumination without affecting the reflected light field. To be more precise, we may rewrite equation 3.35 as

$$\begin{aligned} L'_{lm} &= \Lambda_l \hat{\rho}_l L_{lm} \\ \Lambda_l \hat{\rho}'_l &= 1 \\ B_{lm} &= \Lambda_l \hat{\rho}'_l L'_{lm} \\ &= L'_{lm}. \end{aligned} \tag{3.41}$$

These equations say that we may blur the illumination using the BRDF filter, while treating the BRDF as a mirror. This formulation also allows us to analyze the conditioning in estimating the parameters of a Phong BRDF model under arbitrary illumination. The form of equation 3.40 for the Phong BRDF coefficients indicates that for $l \ll \sqrt{s}$, the effects of the filtering in equation 3.41 are minimal. The BRDF filter passes through virtually all the low-frequency energy with practically no attenuation. Thus, under low-frequency lighting, the reflected light field is essentially independent of the Phong exponent s. This means that **under low-frequency lighting, estimation of the exponent s of a Phong BRDF model is ill-conditioned**. In physical terms, it is difficult to determine the shininess of an object under diffuse lighting. In order to do so robustly, we must have high-frequency lighting components like directional sources. This observation holds for many reflective BRDF models. In particular, we shall see that a similar result can be derived for microfacet BRDFs; estimation of the surface roughness is ill-conditioned under low-frequency lighting.

3.2.7 Microfacet BRDF

Consider a simplified Torrance-Sparrow [84] model,

$$\rho = \frac{1}{4\pi\sigma^2 \cos\theta_i' \cos\theta_o'} \exp\left[-\left(\frac{\theta_h'}{\sigma}\right)^2\right]. \tag{3.42}$$

The subscript h stands for the half-way vector, while σ corresponds to the surface roughness parameter. For simplicity, we have omitted Fresnel and geometric shadowing terms, as well as the Lambertian component usually included in such models.

It is convenient to reparameterize by the reflection vector, as we did for the Phong BRDF. However, it is important to note that microfacet BRDFs are not symmetric about the reflection vector. Unlike for the Phong BRDF, there is a preferred direction, determined by the exitant angle. However, it can be shown by Taylor-series expansions and verified numerically that it is often reasonable to treat the microfacet BRDF using the same machinery as for the Phong case, assuming no outgoing angular dependence. Even under this assumption, it is somewhat difficult to derive precise analytic formulae. However, we may make good approximations.

We analyze the microfacet BRDF by fixing the outgoing angle and reparameterizing by the reflection vector. That is, we set the outgoing angle to $(\theta'_o, 0)$, corresponding to an angle of $2\theta'_o$ with respect to the reflection vector. We now write the BRDF as

$$\hat{\rho} = \sum_{l=0}^{\infty} \sum_{q=-l}^{l} \hat{\rho}_{lq}(\theta'_o) Y_{lq}(\theta'_i, \phi'_i). \tag{3.43}$$

Note that we have reparameterized with respect to the reflection vector, so θ'_i refers to the angle made with the reflection vector. Our goal is to show that azimuthally dependent terms, i.e. those with $q \neq 0$ are small, at least for small angles θ'_o. Furthermore, we would like to find the forms of the terms with $q = 0$. We start off by considering the simplest case, i.e. $\theta'_o = 0$. This corresponds to normal exitance, with the reflection vector also being normal to the surface. We then show how Taylor series expansions can be used to generalize the results.

Normal Exitance

For normal exitance, there is no azimuthal dependence, and the half angle, $\theta'_h = \theta'_i/2$,

$$\hat{\rho}_l = 2\pi \int_0^{\pi/2} \frac{\exp\left[-\theta_i'^2/4\sigma^2\right]}{4\pi\sigma^2} Y_{l0}(\theta'_i) \sin\theta'_i \, d\theta'_i. \tag{3.44}$$

The expansion of $Y_{l0}(t)$ near $t = 0$ for small l is

$$Y_{l0}(t) = \Lambda_l^{-1} \left(1 - \frac{l(l+1)}{4} t^2 + O(t^4)\right). \tag{3.45}$$

The asymptotic form of $Y_{l0}(t)$ near $t = 0$ for large l is

$$Y_{l0}(t) \sim \Lambda_l^{-1} \left(\frac{1}{\sqrt{t}} \cos[(l + 1/2)t - \pi/4]\right). \tag{3.46}$$

To integrate equation 3.44, we substitute $u = \theta'_i/2\sigma$. Then, $\theta'_i = 2\sigma u$. Assuming $\sigma \ll 1$, as it is for most surfaces, the upper limit of the integral becomes infinite, and we have that

$\sin\theta_i' \, d\theta_i' = \theta_i' \, d\theta_i' = 4\sigma^2 u \, du$,

$$\hat{\rho}_l = \int_0^\infty 2e^{-u^2} Y_{l0}(2\sigma u) u \, du. \tag{3.47}$$

We therefore set $t = 2\sigma u$ in equations 3.45 and 3.46. When $\sigma l \ll 1$, we use equation 3.45 to obtain to $O([\sigma l]^4)$,

$$\Lambda_l \hat{\rho}_l = \left(\int_0^\infty 2ue^{-u^2} \, d\, u - (\sigma l)^2 \int_0^\infty 2u^3 e^{-u^2} \, du \right). \tag{3.48}$$

Substituting, $v = u^2$, both integrals evaluate to 1, so we obtain

$$\Lambda_l \hat{\rho}_l = 1 - (\sigma l)^2 + O\left([\sigma l]^4 \right). \tag{3.49}$$

We note that these are the first terms of the Taylor series expansion of $\exp[-(\sigma l)^2]$. When $\sigma l \gg 1$, we use equation 3.46 to obtain (Φ is a phase that encapsulates the lower-order terms)

$$\Lambda_l \hat{\rho}_l \sim \int_0^\infty e^{-u^2} \sqrt{u} \cos[(2\sigma l)u + \Phi] \, du. \tag{3.50}$$

The dominant term can be shown to be $\exp[-(2\sigma l)^2/4] = \exp[-(\sigma l)^2]$. Therefore, we can simply use $\exp[-(\sigma l)^2]$ as a valid approximation in both domains, giving rise to an approximation of the form

$$\Lambda_l \hat{\rho}_l \approx \exp\left[-\left(\sigma l\right)^2 \right]. \tag{3.51}$$

We have also verified this result numerically.

For normal exitance, the BRDF is symmetric about the reflection vector and gaussian, so in that case, equation 3.51 simply states that even in the spherical-harmonic basis, the frequency spectrum of a gaussian is also approximately gaussian, with the frequency width related to the reciprocal of the angular width. For non-normal exitance, we will see that this is still a good approximation. The corrections are small except when l is large (corresponding to directional light sources), and at large viewing angles. These statements are made more precise in the next subsection.

Non-Normal Exitance

For non-normal exitance, we first expand θ'_h in a Taylor series in terms of θ'_i. After some tedious manipulation, we can verify that to first order,

$$\theta_h'^2 = \left(\frac{\theta'_i}{2}\right)^2 \left(1 + \sin^2\phi'_i \tan^2\theta'_o\right). \tag{3.52}$$

When $\theta'_o = 0$ (normal exitance), this is the result we obtained earlier. When $\theta'_o \neq 0$, there is some asymmetry, and the half-angle depends on the azimuthal angle between the light and viewing vector, as defined by the formula above. In angular-space, the BRDF behaves as an "anisotropic" filter over the incident illumination. Our goal here is to bound the extent of this "anisotropy", or asymmetry.

We first consider the coefficients for the isotropic or azimuthally symmetric term Y_{l0}. For l small, we can expand as $1 - (\sigma l)^2 + O(\sigma l)^4$. Now, the constant term is just the area of the microfacet distribution, but when normalized by $\cos\theta'_o$ must integrate to 1. Therefore, corrections are at least of $O(\theta_o'^2(\sigma l)^2)$. In fact, it is possible to obtain a tighter bound by deeper analysis.

Therefore, these corrections are not significant for small outgoing angles, and within that domain, we may use equation 3.51 as a good approximation. It should be noted that in practical applications, measurements made at wide or near-grazing angles are given low confidence anyway.

To make this more concrete, consider what happens to equation 3.47. Considering the first term in the Taylor-series expansion, and including the $\cos\theta'_o$ term in the denominator, we get

$$\hat{\rho}_l = \frac{1}{\cos\theta'_o}\int_0^\infty 2ue^{-u^2} Y_{l0}(2\sigma u)\left(1 - \frac{u^2\tan^2\theta'_o}{2}\right) du. \tag{3.53}$$

The factor of 2 in the denominator of the last (correction) term is to take care of the integration of $\sin^2\phi'_i$. Next, we expand both the cosine and tangent functions about $\theta'_o = 0$, i.e. $\cos\theta'_o = 1 - \theta_o'^2/2 + O\left(\theta_o'^4\right)$ and $\tan^2\theta'_o = \theta_o'^2 + O\left(\theta_o'^4\right)$. Upon doing this, it can be verified that, as expected from physical arguments, there is no net correction to the integral and equation 3.53 evaluates to equation 3.49.

Now, consider the anisotropic terms Y_{lq} with $q \neq 0$. If we Taylor expand as in equation 3.53, we are again going to get something at least of $O\left(\theta_o'^2\right)$, and because the Y_{lq} constant term vanishes, there will be another factor of σ. In fact, if we Taylor-expand, we get an expansion of the form $\sin^2\phi_i' + \sin^4\phi_i' + \cdots$ and it is clear that the azimuthal integral against Y_{lq} vanishes unless there is a term of type $\sin^q\phi_i'$. Therefore, the corrections are actually $O\left(\theta_o'^q\sigma\right)$. What this means in practice is that the azimuthally dependent corrections are only important for large viewing angles and large orders l (l must be large if q is large, and for small q, $\theta_o'^q\sigma$ is small for $\sigma \ll 1$ even if θ_o' is large).

But, that situation will arise only for observations made at large angles of lights that have broad spectra, i.e. directional sources. Therefore, these should be treated separately, by separating the illumination into slow and fast varying components. Equation 3.51 in the frequency domain is a very good approximation for the slow-varying lighting component, while we may approximate the fast-varying lighting component using one or more directional sources in the angular domain.

Conditioning Properties

Since equation 3.51 has many similarities to the equations for the Phong BRDF, most of those results apply here too. In particular, under low-frequency lighting, there is an ambiguity with respect to estimation of the surface roughness σ. Also, inverse lighting is well-conditioned only up to order σ^{-1}. With respect to factorization, there are ambiguities between illumination and reflectance, similar to those for mirror and Phong BRDFs. Specifically, we may blur the illumination while sharpening the BRDF. However, it is important to note that while these ambiguities are exact for Phong and mirror BRDFs, they are only a good approximation for microfacet BRDFs since equation 3.51 does not hold at grazing angles or for high-frequency lighting distributions. In these cases, the ambiguity can be broken, and we have used this fact in an algorithm to simultaneously determine the lighting and the parameters of a microfacet model [73].

In summary, while it is difficult to derive precise analytic formulae, we can derive good approximations to the frequency-space behavior of a microfacet BRDF model. The results are rather similar to those for the Phong BRDF with the Phong exponent s replaced by the physically-based surface roughness parameter σ.

In this subsection, we have seen analytic formulae derived for a variety of common lighting and BRDF models, demonstrating the implications of the theoretical analysis. We end this chapter by summarizing its key contributions, discussing avenues for future work, and outlining the rest of the dissertation, which discusses some practical applications of the theoretical analysis in chapters 2 and 3.

3.3 Conclusions and Future Work

We have presented a theoretical analysis of the structure of the reflected light field from a convex homogeneous object under a distant illumination field. In chapter 2, we have shown that the reflected light field can be formally described as a convolution of the incident illumination and the BRDF, and derived an analytic frequency space convolution formula. This means that, under our assumptions, reflection can be viewed in signal processing terms as a filtering operation between the lighting and the BRDF to produce the output light field. Furthermore, inverse rendering to estimate the lighting or BRDF from the reflected light field can be understood as deconvolution. This result provides a novel viewpoint for many forward and inverse rendering problems.

In this chapter, we have derived analytic formulae for the frequency spectra of many common BRDF and lighting models, and have demonstrated the implications for inverse problems such as lighting recovery, BRDF recovery, and light field factorization. We have shown in frequency-space why a gazing sphere is well-suited for recovering the lighting—the frequency spectrum of the mirror BRDF (a delta function) is constant—and why a directional source is well-suited for recovering the BRDF—we are estimating the BRDF filter by considering its impulse response. With the aid of our theory, we have been able to quantitatively determine the well-posedness and conditioning of many inverse problems. The ill-conditioning observed by Marschner and Greenberg [54] in estimating the lighting from a Lambertian surface has been explained by showing that only the first 9 coefficients of the lighting can robustly be recovered, and we have shown that factorization of lighting effects and low-frequency texture is ambiguous. All these results indicate that the theory provides a useful analytical tool for studying the properties of inverse problems.

Of course, all the results presented in this chapter depend on our assumptions. Furthermore, the results for well-posedness of inverse problems depend on having all the reflected light field coefficients, or the entire reflected light field available. It is an interesting future direction to consider how these results change when we have only a limited fraction of the reflected light field available, as in most practical applications, or can move our viewpoint only in a narrow range. More generally, we believe a formal analysis of inverse problems under more general assumptions is of significant and growing interest in many areas of computer graphics, computer vision, and visual perception.

The remainder of this dissertation describes practical applications of the theoretical framework for forward and inverse rendering. Chapters 4 and 5 describe frequency domain algorithms for environment map prefiltering and rendering. Chapter 4 discusses the case of Lambertian BRDFs (irradiance maps), while chapter 5 extends these ideas to the general case with arbitrary isotropic BRDFs. Finally, chapter 6 discusses how the theoretical analysis can be extended and applied to solve practical inverse rendering problems under complex illumination.

Chapter 4

Irradiance Environment Maps

In the introduction to this dissertation, we noted that complex realistic lighting environments are rarely used in either forward or inverse rendering. We also stated our thesis that a deeper understanding of the computational nature of reflection and illumination helps to address these difficulties and restrictions in a number of areas in computer graphics and vision. Subsequently, in chapters 2 and 3, we have developed a new way of looking at reflection, formalizing the idea of reflection as a spherical convolution of the incident illumination and BRDF. The insights from these two chapters lead to the possibility of attacking a number of difficult forward and inverse rendering problems in the frequency domain.

Chapters 4, 5 and 6 of this dissertation are devoted to practical applications of the signal-processing ideas developed theoretically in the previous two chapters. Chapters 4 and 5 deal with efficient representations and algorithms for forward rendering using environment maps, which are representations of the (distant) incident illumination distribution at a point. This chapter considers the case of *Irradiance Environment Maps*, corresponding to the reflection from diffuse or Lambertian objects. We show that frequency-space analysis can be used to reduce the effects of arbitrarily complex (but distant) incident illumination to a simple analytic low-dimensional formula. In the next chapter, we will extend these ideas, using similar methods for general BRDFs, further showcasing the practical benefits of frequency-space concepts like sampling rate analysis and efficient frequency domain convolutions. Finally, chapter 6 presents practical algorithms for inverse rendering—estimation of illumination and material properties.

In this chapter, we consider the rendering of diffuse objects under distant illumination, as specified by an environment map. Using an analytic expression for the irradiance in terms of spherical harmonic coefficients of the lighting, derived in chapter 3.2.5, we show that one needs to compute and use only 9 coefficients, corresponding to the lowest-frequency modes of the illumination, in order to achieve average errors of only 1%. In other words, the irradiance is insensitive to high frequencies in the lighting, and is well approximated using only 9 parameters. In fact, we show that the irradiance can be procedurally represented simply as a quadratic polynomial in the cartesian components of the surface normal, and give explicit formulae. These observations lead to a simple and efficient procedural rendering algorithm amenable to hardware implementation, a prefiltering method up to three orders of magnitude faster than previous techniques, and new representations for lighting design and image-based rendering.

The rest of this chapter is organized as follows. After an introduction to the specific problem of interest here, in section 1, we briefly describe the relevant background and practical details from the previous theoretical analysis required here in section 2. Section 3 discusses practical implementation of our algorithms. Finally, section 4 concludes this paper and suggests directions for future work. This chapter corresponds to our SIGGRAPH paper on *An Efficient Representation for Irradiance Environment Maps* [71].

4.1 Introduction and Previous Work

Lighting in most real scenes is complex, coming from a variety of sources including area lights and large continuous lighting distributions like skylight. But current graphics hardware only supports point or directional light sources. One reason is the lack of simple procedural formulas for general lighting distributions. Instead, an integration over the upper hemisphere must be done for each pixel.

One approach to using general lighting distributions is the method of environment maps. Environment maps are representations of the incident illumination at a point. Blinn and Newell [5] used them to efficiently find reflections of distant objects. Miller and Hoffman [59], and Greene [22] *prefiltered* environment maps, precomputing separate reflection

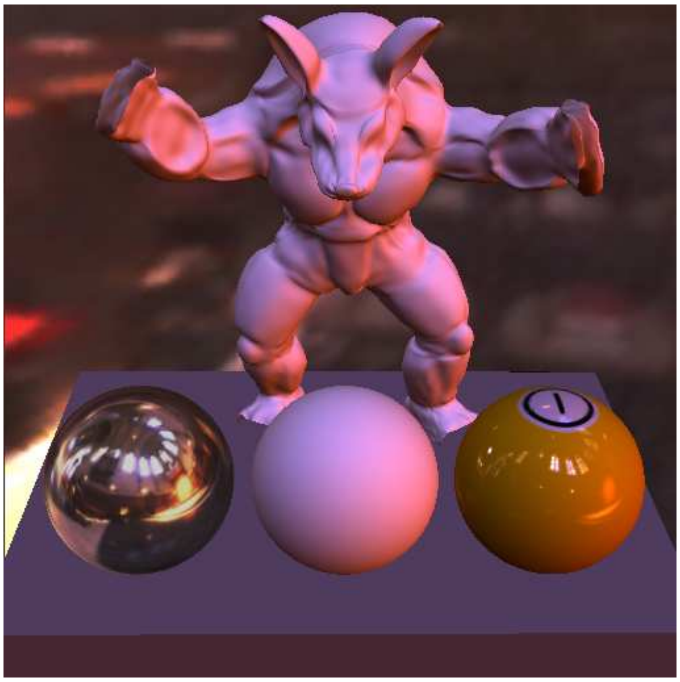

Figure 4.1: *The diffuse shading on all the objects is computed procedurally in real-time using our method. The middle sphere, armadillo, and table are white diffuse reflectors. The colors come from the environment—owing to a variety of colored sources, including blue stained-glass windows. Our method can also be combined with standard texture mapping—used to modulate the albedo of the pool-ball on the right—and reflection mapping—used for specular highlights on the pool-ball, and for the mirror sphere on the left. The environment is a light probe of the Grace Cathedral. Tone mapping is used to convey high dynamic range for the background and the mirror sphere; the remaining objects are shaded using a linear scale.*

maps for the diffuse and specular components of the BRDF. Cabral et al. [8] handled general BRDFs by using a 2D set of prerendered images. Prefiltering is generally an offline, computationally expensive process. After prefiltering, rendering can usually be performed at interactive rates with graphics hardware using texture-mapping.

Of course, environment maps, and the relevant techniques presented in this dissertation, are only an approximation and do not account for near-field illumination, cast shadows, or interreflection. Nevertheless, they have proven an effective tool for interactive rendering with realistic lighting effects.

This chapter focuses on the Lambertian component of the BRDF. We use the term *irradiance environment map* for a diffuse reflection map indexed by the surface normal, since each pixel simply stores the irradiance for a particular orientation of the surface. For applications like games, irradiance maps are often stored directly on the surface, instead of as a function of the normal vector, and are called *light maps*. Irradiance environment maps can also be extended to spatially varying illumination by computing an *irradiance volume*, as done by Greger et al. [23]. Many of the same ideas can be applied to speeding up global illumination algorithms. The slowly varying nature of irradiance has led to Ward and Heckbert [85] proposing interpolation using irradiance gradients, while the idea of storing irradiance as a function of surface orientation in *orientation lightmaps* has been proposed by Wilkie et al. [87].

Our approach relies on the rapid computation of an analytic approximation to the irradiance environment map. For rendering, we demonstrate a simple procedural algorithm that runs at interactive frame rates, and is amenable to hardware implementation. The procedural approach is preferable to texture-mapping in some applications. Since irradiance varies slowly with orientation, it need only be computed per-vertex and interpolated across triangles. Further, we require only a single texturing pass to render textured objects with irradiance environment maps, since the irradiance is computed procedurally. On the other hand, the standard approach requires a separate texture for the irradiance, and needs *multitexturing* support or multiple texturing passes. In other applications, where per-fragment texture-mapping is relatively inexpensive, our method can be used to very efficiently compute the irradiance environment map texture. Our novel representation also suggests new approaches to lighting design and image-based rendering.

4.2 Background

Empirically, it is well known that the reflected intensity from a diffuse surface varies slowly as a function of surface orientation. This qualitative observation has been used to justify representing irradiance environment maps at low resolutions [59], and in efficiently computing the shading hierarchically [39, 45]. Our goal is to use an analytic quantitative formula for the irradiance, derived in section 3.2.5, which formalizes these observations, and allows for principled approximations.

Let L denote the distant lighting distribution. As is common with environment map algorithms, we neglect the effects of cast shadows and near-field illumination. The irradiance E is then a function of the surface normal only and is given by an integral over the upper or visible hemisphere,

$$E(\alpha, \beta) = \int_{\theta_i'=0}^{\pi/2} \int_{\phi_i'=0}^{2\pi} L\left(R_{\alpha,\beta}(\theta_i', \phi_i')\right) \cos\theta_i' \, d\theta_i' d\phi_i'. \tag{4.1}$$

We must scale E by the surface albedo[1], which may be dependent on position $\vec{X}$ and be described by a texture $T(\vec{X})$, to find the reflected light field B, which corresponds directly to the image intensity,

$$B(\vec{X}; \alpha, \beta) = T(\vec{X}) E(\alpha, \beta). \tag{4.2}$$

Our main concern will be approximating the irradiance E. A texture map $T(\vec{X})$ may be used later to simply modulate the reflected light field computed. Note that the form of equation 4.1 is simply a special case of the reflection equation 2.11 for isotropic surfaces with no outgoing angular dependence. The limits of the θ_i' integral range from 0 to $\pi/2$ because we consider only the front hemisphere, where the cosine of the incident angle is positive. The transfer function corresponding to the Lambertian BRDF is the *clamped cosine* function $\hat{\rho}(\theta_i') = \max(\cos\theta_i', 0)$.

In section 3.2.5 (a more detailed version of which is published in [72]), we have been able to derive an analytic formula for the irradiance by determining the spherical harmonic filter coefficients for the Lambertian clamped-cosine function. Similar results have been obtained independently by Basri and Jacobs [2] in simultaneous work on face recognition.

For the purposes of implementation, it is often convenient to use real-valued functions where possible, rather than the complex forms of the spherical harmonics given in equation 2.27. It is easy to define real forms of the spherical harmonics, simply by considering the real and complex parts separately. For this purpose, we define the real form of the

[1] Technically, for Lambertian objects, the BRDF is given by $1/\pi$ times the albedo, so the textures should be multiplied by $1/\pi$.

spherical harmonics as follows (c.f. equation 2.26),

$$\begin{aligned} N_{lm} &= \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} \\ Y_{lm}(\theta,\phi) &= N_{lm}P_{lm}(\cos\theta)az_m(\phi), \end{aligned} \tag{4.3}$$

where the azimuthal basis functions are defined by

$$\begin{aligned} az_{+m}(\phi) &= \sqrt{2}\cos\phi \\ az_0(\phi) &= 1 \\ az_{-m}(\phi) &= \sqrt{2}\sin\phi. \end{aligned} \tag{4.4}$$

While this is essentially the standard definition of the real form of the spherical harmonics, the sign conventions used are not always consistent. For that reason, we will make explicit the numerical values used here to fix the precise conventions used by us.

Recall that the spherical harmonics may be written as polynomials of the cartesian components (x, y, z). Below, we give the numeric values of the the first 9 spherical harmonics (with $l \leq 2$), which are simply constant ($l = 0$), linear ($l = 1$), and quadratic ($l = 2$) polynomials (c.f. figure 2.3),

$$\begin{aligned} (x,y,z) &= (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta) \\ Y_{00}(\theta,\phi) &= 0.282095 \\ (Y_{11}; Y_{10}; Y_{1-1})\,(\theta,\phi) &= 0.488603\,(x; z; y) \\ (Y_{21}; Y_{2-1}; Y_{2-2})\,(\theta,\phi) &= 1.092548\,(xz; yz; xy) \\ Y_{20}(\theta,\phi) &= 0.315392\left(3z^2-1\right) \\ Y_{22}(\theta,\phi) &= 0.546274\left(x^2-y^2\right). \end{aligned} \tag{4.5}$$

Note that these basis functions are closely related to the spherical polynomials used by Arvo [1] in his irradiance tensor formulation.

$E(\alpha, \beta)$ and $L(\theta, \phi)$ can be represented by the coefficients—E_{lm} and L_{lm}—in their spherical harmonic expansion,

$$\begin{aligned} L(\theta,\phi) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l} L_{lm}Y_{lm}(\theta,\phi) \\ E(\alpha,\beta) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l} E_{lm}Y_{lm}(\alpha,\beta). \end{aligned} \tag{4.6}$$

We may also expand the Lambertian transfer function $\hat{\rho}(\theta_i') = \max(\cos\theta_i', 0)$, i.e. the clamped cosine, in terms of spherical harmonics. Since $\hat{\rho}$ has no azimuthal dependence, $m = 0$ and we use only the l index,

$$\hat{\rho}(\theta) = \max\left[\cos\theta, 0\right] = \sum_l \hat{\rho}_l Y_{l0}(\theta). \tag{4.7}$$

With these definitions, one can directly apply equation 2.62 or equation 3.10,

$$E_{lm} = \Lambda_l \hat{\rho}_l L_{lm}, \tag{4.8}$$

where $\Lambda_l = \sqrt{4\pi/(2l+1)}$. The only difference in equation 3.28 in section 3.2.5 is that we used there the reflected light field B, which is simply a scaled version of the irradiance E for Lambertian surfaces.

It will be convenient to define a new variable A_l by

$$A_l = \Lambda_l \hat{\rho}_l, \tag{4.9}$$

and to expand out the irradiance for rendering,

$$E(\alpha,\beta) = \sum_{l,m} A_l L_{lm} Y_{lm}(\alpha,\beta). \tag{4.10}$$

An analytic formula for $\hat{\rho}_l$ (and hence A_l) has been derived in section 3.2.5. It can be shown that A_l vanishes for odd values of $l > 1$, and even terms fall off very rapidly as $l^{-5/2}$. The analytic formulae are given by (c.f. equation 3.32)

$$\begin{array}{ll} l = 1 & A_l = \frac{2\pi}{3} \\ l > 1, \textbf{odd} & A_l = 0 \\ l \textbf{ even} & A_l = (-1)^{\frac{l}{2}-1} \dfrac{2\pi}{(l+2)(l-1)} \left[\dfrac{l!}{2^l \left(\frac{l}{2}!\right)^2} \right]. \end{array} \tag{4.11}$$

Numerically, the first few terms are

$$\begin{gathered} A_0 = 3.141593 \quad A_1 = 2.094395 \quad A_2 = 0.785398 \\ A_3 = 0 \quad A_4 = -0.130900 \quad A_5 = 0 \quad A_6 = 0.049087. \end{gathered} \tag{4.12}$$

Approximation: For rendering, we make use of the observation that A_l decays so fast that we need consider only low-frequency lighting coefficients, of order $l \leq 2$. Equivalently, **the irradiance is well approximated by only 9 parameters—**1 for $l = 0, m = 0$, 3 for $l = 1, -1 \leq m \leq 1$, and 5 for $l = 2, -2 \leq m \leq 2$. By working in frequency-space, we exploit the low-frequency character of the Lambertian BRDF filter, using a few coefficients instead of a full hemispherical integral. The simple form of the first 9 spherical harmonics, given in equation 4.5, makes implementation straightforward.

4.3 Algorithms and Results

In this section, we discuss three applications of this result. First, we show how to rapidly prefilter the lighting distribution, computing the coefficients L_{lm}. Next, we develop a simple real-time procedural shader for rendering that takes these coefficients as inputs. Finally, we discuss other applications of our representation.

4.3.1 Prefiltering

For a given environment map, we first find the 9 lighting coefficients, L_{lm} for $l \leq 2$, by integrating against the spherical harmonic basis functions. Each color channel is treated separately, so the coefficients can be thought of as RGB values,

$$L_{lm} = \int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} L(\theta, \phi) Y_{lm}(\theta, \phi) \sin\theta \, d\theta d\phi. \tag{4.13}$$

The expressions for the Y_{lm} are found in equation 4.5. The integrals are simply sums of the pixels in the environment map L, weighted by the functions Y_{lm}. The integrals can also be viewed as moments of the lighting, or as inner-products of the functions L and Y_{lm}.

Since we compute 9 numbers, the prefiltering step takes O(9S) time, where S is the size (total number of pixels) of the environment map. By comparison, the standard method of computing an irradiance environment map texture takes $O(\mid T \mid \cdot S)$ time, where $\mid T \mid$ is the number of texels in the irradiance environment map. Our method will therefore be approximately $\mid T \mid /9$ times faster[2]. Even if a conventional irradiance environment map is computed at a very low resolution of 64×64, corresponding to $\mid T \mid = 4096$, our method will be nearly 500 times faster.

We have implemented prefiltering as a preprocessing step for a given environment map. Values of L_{lm} for a few light probes are tabulated in figure 4.1. The computation time for a 300x300 environment map was less than a second. This indicates that our approach might be able to handle scenes with dynamic lighting in the future. By contrast, the standard method of performing a hemispherical integral for each pixel to compute the irradiance environment map took approximately two hours. In fact, if an explicit representation of the irradiance environment map texture is required, we believe the best way of computing it is to first compute the 9 coefficients L_{lm} using our method, and then use these to very rapidly generate the irradiance environment map using the rendering method described below.

It is important to know what errors result from our 9 parameter approximation. The maximum error for any pixel, as a fraction of the total intensity of the illumination, is 9% and corresponds to the maximum error in the order 2 approximation of the clamped cosine function. Furthermore, the average error over all surface orientations can be shown to be under 3% for any physical input lighting distribution [2]. For the environment maps used in our examples, corresponding to complex natural illumination, the results are somewhat

[2]It may be possible to use a hierarchical integration scheme, as demonstrated by Kautz et al. [39] for Phong BRDFs, to speed up both our method and the conventional approach. Hardware acceleration may also be possible.

	Grace Cathedral			Eucalyptus Grove			St. Peters Basilica		
L_{00}	.79	.44	.54	.38	.43	.45	.36	.26	.23
L_{1-1}	.39	.35	.60	.29	.36	.41	.18	.14	.13
L_{10}	-.34	-.18	-.27	.04	.03	.01	-.02	-.01	-.00
L_{11}	-.29	-.06	.01	-.10	-.10	-.09	.03	.02	.01
L_{2-2}	-.11	-.05	-.12	-.06	-.06	-.04	.02	.01	.00
L_{2-1}	-.26	-.22	-.47	.01	-.01	-.05	-.05	-.03	-.01
L_{20}	-.16	-.09	-.15	-.09	-.13	-.15	-.09	-.08	-.07
L_{21}	.56	.21	.14	-.06	-.05	-.04	.01	.00	.00
L_{22}	.21	-.05	-.30	.02	-.00	-.05	-.08	-.06	.00

Table 4.1: *Scaled RGB values of lighting coefficients for a few environments. These may be used directly for rendering, and for checking the correctness of implementations.*

better than the worst-case bounds—the average error is under 1%, and the maximum pixel error is under 5%. Finally, figure 4.2 provides a visual comparison of the quality of our results with standard prefiltering, showing that our method produces a perceptually accurate answer.

4.3.2 Rendering

For rendering, we can find the irradiance using equation 4.10. Since we are only considering $l \leq 2$, the irradiance is simply a quadratic polynomial of the coordinates of the (normalized) surface normal. Hence, with $\vec{N}^t = (x\ y\ z\ 1)$, we can write

$$E(\vec{N}) = \vec{N}^t M \vec{N}. \tag{4.14}$$

M is a symmetric 4x4 matrix. Each color has an independent matrix M. Equation 4.14 is particularly useful for rendering, since we require only a matrix-vector multiplication and a dot-product to compute E. The matrix M is obtained by expanding equation 4.10,

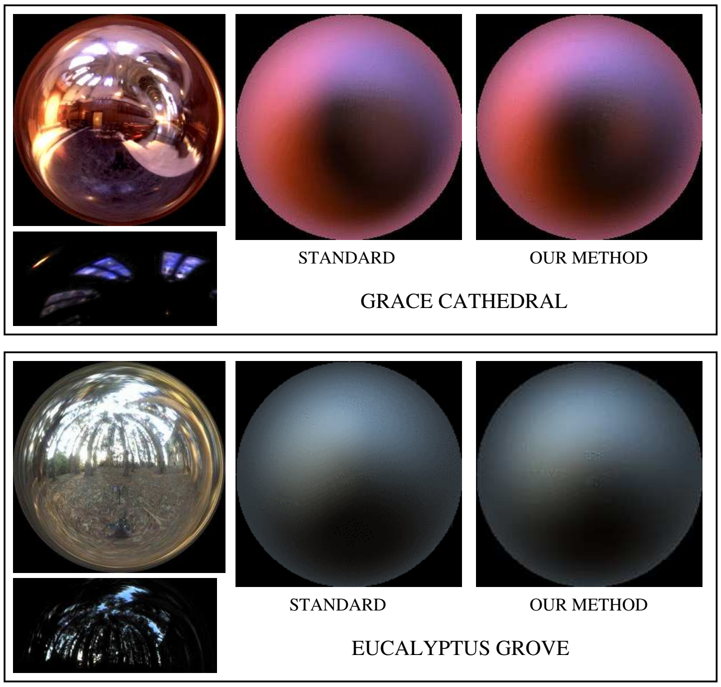

Figure 4.2: *A comparison of irradiance maps from our method to standard prefiltering. The irradiance map resolutions are 256x256. For each light probe, the left image is a tone-mapped version of the environment. Below that, we show the brightest parts of the environment on a linear scale. Both environments have bright bluish lights—from stained-glass windows, and the sky respectively—which are not apparent in the tone-mapped images. This accounts for the bluish portions of the irradiance maps. It can be seen that our method produces a result very close to the correct answer. Note that our rendering algorithm does not actually use irradiance maps; we computed them here solely for the purposes of the quality comparison. The coordinate mapping in the images is such that the center of the image is straight forward ($\theta = 0$, the north pole or +Z), the circumference of the image is straight backwards ($\theta = \pi$, the south pole or -Z), and θ varies uniformly in the radial direction from 0 to π. The azimuthal angle ϕ corresponds to the image polar angle.*

$$M = \begin{pmatrix} c_1 L_{22} & c_1 L_{2-2} & c_1 L_{21} & c_2 L_{11} \\ c_1 L_{2-2} & -c_1 L_{22} & c_1 L_{2-1} & c_2 L_{1-1} \\ c_1 L_{21} & c_1 L_{2-1} & c_3 L_{20} & c_2 L_{10} \\ c_2 L_{11} & c_2 L_{1-1} & c_2 L_{10} & c_4 L_{00} - c_5 L_{20} \end{pmatrix}$$

$$c_1 = 0.429043 \quad c_2 = 0.511664$$

$$c_3 = 0.743125 \quad c_4 = 0.886227 \quad c_5 = 0.247708. \tag{4.15}$$

The entries of M depend[3] on the 9 lighting coefficients L_{lm} and the expressions for the spherical harmonics. The constants come from the numerical values of A_l given in equation 4.12, and the spherical harmonic normalizations given in equation 4.5.

On systems not optimized for matrix and vector operations, it may be more efficient to explicitly write out equation 4.10 for the irradiance as a sum of terms,i.e. expand equation 4.15,

$$\begin{aligned} E(\vec{N}) &= c_1 L_{22} \left(x^2 - y^2\right) + c_3 L_{20} z^2 + c_4 L_{00} - c_5 L_{20} \\ &+ 2c_1 \left(L_{2-2}xy + L_{21}xz + L_{2-1}yz\right) \\ &+ 2c_2 \left(L_{11}x + L_{1-1}y + L_{10}z\right). \end{aligned} \tag{4.16}$$

We implemented equations 4.14 and 4.16 as procedural shaders in the Stanford real-time programmable shading system [68]. We used the ability of that system to perform computations per-vertex. Since E varies slowly, this is adequate and the shading is insensitive to how finely the surfaces are tessellated. The irradiance computations may be performed in software or compiled to vertex programming hardware, if available. The simple forms of equations 4.14 and 4.16 indicate that a per-fragment method could also be implemented in programmable hardware.

We were able to achieve real-time frame rates on PCs and SGIs. As shown in the SIGGRAPH 2001 conference proceedings videotape, we can interactively rotate objects and move our viewpoint, with the irradiance being procedurally recomputed at every frame. We can also rotate the lighting by applying the inverse rotation to the normal $\vec{N}$. Images rendered using our method look identical to those obtained by texture-mapping after precomputing irradiance environment maps.

4.3.3 Representation

Conceptually, the final image is composed of a sum of spherical harmonic basis functions, scaled by the lighting coefficients L_{lm}. These 3D irradiance basis functions depend on the surface normal and are defined over the entire object, making it possible to generate an

[3] A symmetric 4x4 matrix has 10 degrees of freedom. One additional degree is removed since $\vec{N}$ lies on the unit sphere.

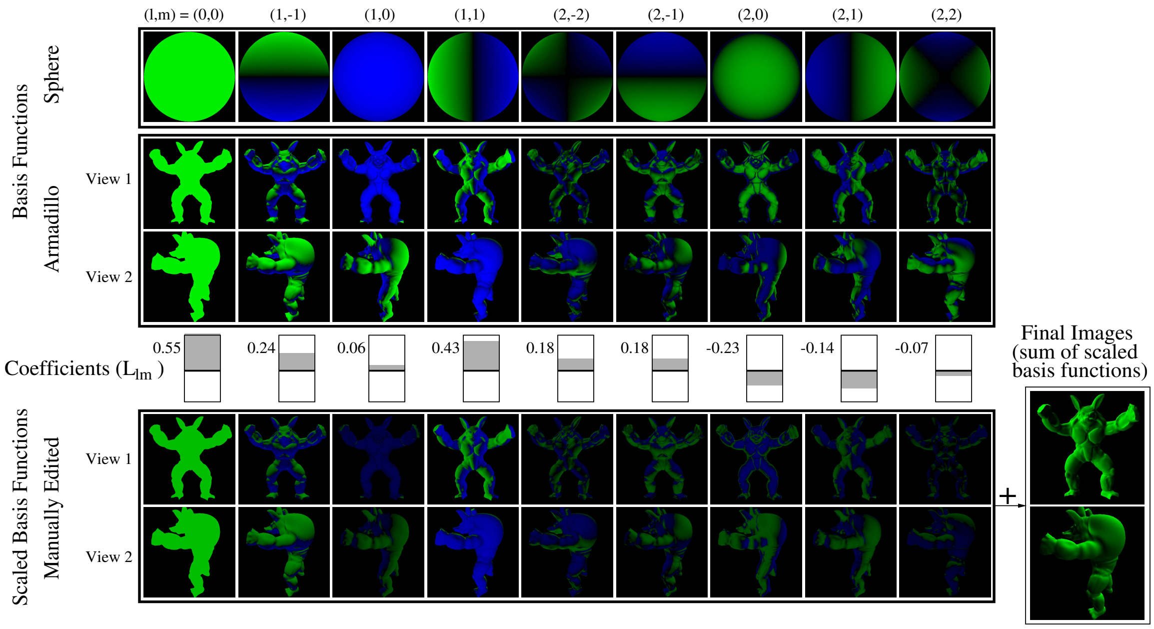

Figure 4.3: *Illustration of our representation, and applications to control appearance. The basis functions have both positive values, shown in green, and negative values, shown in blue. Topmost, we show the spherical harmonic basis functions on a sphere—note that these are actual images, not the coordinate mappings of figure 4.2—and the armadillo. The basis functions are defined over the entire object surface; we show only two views. The rightmost 5 functions are dimmer since they have the highest frequency* $(l = 2)$ *and contribute the least. Conceptually, the basis functions are then scaled by the lighting coefficients* L_{lm} *and added to produce renderings.* L_{lm} *are actually RGB values; for simplicity, we show the coefficients for only one color (green). The coefficients* L_{lm} *may be adjusted manually to manipulate appearance. This editing can be fairly intuitive—for instance, we make* L_{11} *large and positive to darken the right side (with respect to us) and left arm of the armadillo image, since the basis function* $(1, 1)$ *is negative in that region.*

image from any viewpoint. We may also manually adjust the 9 lighting coefficients L_{lm} to directly control appearance, as shown in figure 4.3. The lighting coefficients can often be assigned intuitive meanings. For instance, L_{1-1} is the moment about the vertical or y-axis, and measures the extent to which the upper hemisphere is brighter than the lower hemisphere. As can be seen from figure 4.1, L_{1-1} is usually large and positive, since most scenes are lit from above. By making this value negative, we could give the appearance of the object being lit from below.

Our representation may also be useful in the future for *image-based rendering with varying illumination*. Hallinan [30] and Epstein et al. [18] have observed empirically that, for a given view, images of a matte object under variable lighting lie in a low-dimensional subspace. Our theory explains this observation, and indicates that a 9D subspace suffices. Basri and Jacobs [2] have obtained similar theoretical results. To synthesize images of a diffuse object under arbitrary illumination, we therefore need only the 9 basis functions, which could be computed from a small number of photographs. Such an approach would significantly speed up both acquisition and rendering in a method such as Debevec et al. [15].

4.4 Conclusions and Future Work

We have described a novel analytic representation for environment maps used to render diffuse objects, and have given explicit formulae for implementation. Our approach allows us to use an arbitrary illumination distribution for the diffuse component of the BRDF, instead of the limitation of current graphics hardware to point or directional sources. We simply specify or compute the first 9 moments of the lighting. Even where more conventional texture-mapping methods are desired, our approach allows us to very efficiently compute irradiance environment map textures.

It should be noted that environment mapping in general, and the methods described in this and the next chapter in particular, are restricted to distant illumination without cast shadows or interreflection. An obvious question is how we can modify and utilize our results when the theoretical assumptions don't exactly hold, i.e. we want to account for some spatially-varying illumination, cast shadows, and interreflection. It is our belief that low-dimensional subspaces may still be appropriate, even if they are not specified by simple analytic formulae in terms of spherical harmonics, or can be described using only 9 parameters. Some preliminary work in this area has already been demonstrated by Sloan et al. [66].

Another interesting question is whether we can use similar frequency-space methods for the specular BRDF component, and more general non-Lambertian reflection functions. One solution to this problem will be presented in the next chapter. In the future, we would also like to further explore the applications to lighting design and image-based rendering discussed in this chapter.

Chapter 5

Frequency Space Environment Map Rendering

In the previous chapter on irradiance environment maps, we have seen how frequency-space analysis can be used to efficiently compute and represent convolutions of the incident illumination with the low-pass Lambertian reflection function, i.e. clamped cosine. An obvious question is whether a similar approach is possible for general BRDFs. In this chapter, we will present such a method. We will therefore have shown that the signal-processing ideas can be used to significantly increase the quality of interactive rendering. Since we are using frequency space methods to render with environment maps, we have entitled this chapter *Frequency Space Environment Map Rendering*. The next chapter of the dissertation will switch from forward to inverse rendering, extending and applying the theoretical framework to determining lighting and BRDFs from sequences of real photographs.

This chapter presents a new method for real-time rendering of objects with complex isotropic BRDFs under distant natural illumination, as specified by an environment map. Our approach is based on spherical frequency space analysis and includes three significant improvements over previous work. Firstly, we are able to theoretically analyze required sampling rates and resolutions, which have traditionally been determined in an ad-hoc manner. We also introduce a new compact representation, which we call a *spherical harmonic reflection map (SHRM)*, for efficient representation and rendering. Finally, we show how to

Figure 5.1: *These images, showing many different lighting conditions and BRDFs, were each rendered at approximately 30 frames per second using our Spherical Harmonic Reflection Map (SHRM) representation. From left to right, a simplified microfacet BRDF, krylon blue (using McCool et al.'s reconstruction from measurements at Cornell), orange and velvet (CURET database), and an anisotropic BRDF (based on the Kajiya-Kay model). The environment maps are the Grace Cathedral, St. Peter's Basilica, the Uffizi gallery, and a Eucalyptus grove, courtesy Paul Debevec. The armadillo model is from Venkat Krishnamurthy.*

rapidly prefilter the environment map to compute the *SHRM*—our frequency domain prefiltering algorithm is generally orders of magnitude faster than previous angular (spatial) domain approaches.

The rest of this chapter is organized as follows. We introduce the problem in section 1, followed by a discussion of previous work in section 2. Section 3 introduces the preliminaries and background required. Section 4 describes our new *SHRM* representation, section 5 discusses an analysis of sampling rates in order to set frequency space orders and resolutions, and section 6 describes our efficient frequency domain prefiltering algorithms. Section 7 presents our results using a number of different environment maps and BRDFs. Finally, we conclude this chapter and discuss future work in section 8.

5.1 Introduction

Our goals are real-time rendering with complex natural illumination and realistic, possibly measured, BRDFs. The closest previous work is that of Cabral et al. [8], who extended standard environment maps by interactively warping and combining a sparse 2D set of prerendered images. These precomputed images were obtained by *prefiltering* the environment map, i.e. integrating the product of the BRDF and lighting over the visible (upper) hemisphere for each image pixel, with each pixel corresponding to a particular surface normal direction. Subsequently, Kautz et al. [38, 40] proposed alternative implementations and improved prefiltering methods.

This paper introduces a new frequency space paradigm for prefiltering and rendering environment mapped images with general isotropic BRDFs. We show that frequency space analysis allows for setting sampling rates accurately, and enables compact representations. Further, just as image convolutions are often computed in the Fourier rather than the spatial domain, prefiltering is more efficient in frequency rather than angular space. Our main contributions are:

Theoretical analysis of sampling rates and resolutions: Most previous work has determined reflection map resolutions, or the number of reflection maps required, in an ad-hoc manner. By using a signal-processing framework, we are able to perform error analysis, that allows us to set sampling rates and resolutions accurately.

Efficient representation and rendering with Spherical Harmonic Reflection Maps: We introduce *spherical harmonic reflection maps (SHRMs)* as a compact representation. Instead of a single color, each pixel stores coefficients of a spherical harmonic expansion encoding view-dependence of the reflection map. An observation that emerges from the theoretical analysis is that for almost all BRDFs, a very low order spherical harmonic expansion suffices. Thus, SHRMs can be evaluated in real-time for rendering. Further, they are significantly more compact and accurate than previous methods [8, 38] that use an explicit 1D or 2D set of images.

Fast prefiltering: One of the drawbacks of current environment mapping techniques is the significant computational time required for prefiltering, which can run into hours, and preclude the use of these approaches in applications involving lighting and material design, or dynamic lighting. We introduce new prefiltering methods based on spherical harmonic transforms, and show both empirically, and analytically by computational complexity analysis, that our algorithms are orders of magnitude faster than previous work.

We present a complete theoretical analysis and practical algorithm pipeline, incorporating all three contributions. It is also possible to separately (incrementally) incorporate any one (or two) of the improvements into previous methods.

5.2 Related Work

Angular space environment map rendering has a long history in graphics, including early work by Blinn and Newell [5], Miller and Hoffman [59], and Greene [22]. Hakura et al. [28] propose location and geometry-dependent environment maps for local reflections. Our goals are different in that we want to capture the effects of complex BRDFs and use any object geometry, but it should be possible in future to combine the methods for local reflections and complex BRDFs.

As noted by Cabral et al. [8], environment mapping can be viewed as reflection-space image-based rendering, and is therefore related to a number of IBR methods like surface light fields [60, 64, 88]. A surface light field stores the outgoing radiance distribution for each point on a geometric model. Our representation is essentially equivalent to the surface light field of a sphere. We store the reflected radiance distribution for each normal direction, allowing our representation to be mapped on to any object geometry.

Our work also relates to recent research on hardware rendering with factored BRDFs [37, 57]. However, these methods require the BRDF to at least approximately satisfy a particular factored form. These previous methods also do not support complex illumination.

5.3 Preliminaries

In this section, we first discuss the reflection equation and introduce the basic framework for our method. We then describe our BRDF parameterization, and discuss previous 4D function representations. We use the same notation as in the rest of the dissertation, summarized in table 1.1. In addition, there is some new notation used in this chapter, summarized in table 5.1.

Assumptions: We make a number of simplifying assumptions common in real-time rendering in general, and environment mapping [8, 38, 40, 59] in particular. Specifically, we assume distant illumination and isotropic BRDFs, and restrict ourselves to direct lighting, ignoring interreflection and self-shadowing. We will also not explicitly considered textured objects, but it is easy to use texture-mapping to modulate the net reflectance, simply

$\hat{B}(\alpha,\beta,\theta_o,\phi_o)$	Reflected light field in terms of fixed global viewing direction
F	Maximum order l of coefficients $\hat{\rho}_{lpq}$,B_{lmpq}
P	Maximum order p in spherical harmonic expansion
S	Angular resolution $(S > F)$
T	Number of images in angular space $(T > P)$
W	Angular width of BRDF
ϵ	Error (unaccounted energy in approximation)
C_a	Angular space computational cost
C_f	Frequency domain computational cost

Table 5.1: *Notation used in chapter 5.*

by multiplying the texture and the reflected light field computed by our methods. It should be noted that it is not easy to modulate BRDF parameters, only the net reflectance.

Reflection Equation: Given our assumptions, we may drop spatial dependence, and write the standard reflection equation (c.f. equation 2.9),

$$B(\alpha,\beta,\gamma,\theta'_o) = \int_{\Omega'_i} L(\theta_i,\phi_i)\hat{\rho}(\theta'_i,\phi'_i,\theta'_o,\phi'_o)\,d\omega'_i. \tag{5.1}$$

As usual, L is the incident radiance, $\hat{\rho}$ is the BRDF transfer function, and B is the reflected light field. (α,β) correspond to the *global* coordinates of the surface normal, (θ_i,ϕ_i) are the *global* coordinates of the incident light direction, and (θ'_i,ϕ'_i) and (θ'_o,ϕ'_o) are *local* incident and outgoing directions. Applying the standard rotation to align local and global coordinates, we obtain (c.f. equation 2.11),

$$B(\alpha,\beta,\theta'_o,\phi'_o) = \int_{\Omega'_i} L\left(R_{\alpha,\beta}(\theta'_i,\phi'_i)\right)\hat{\rho}(\theta'_i,\phi'_i,\theta'_o,\phi'_o)\,d\omega'_i. \tag{5.2}$$

When the viewer is distant, it is often useful to rewrite the reflected light field in terms of the global viewing direction $\vec{V} = (\theta_o,\phi_o)$. Note that so far in this dissertation, we have considered local outgoing angles only, so this is a new notation. To distinguish it from the familiar notation in terms of local angles, we denote the reflected light field, parameterized by global outgoing direction as $\hat{B}$. We now have,

$$
\begin{aligned}
(\theta'_o, \phi'_o) &= R^{-1}_{\alpha,\beta}(\theta_o, \phi_o) \\
\hat{B}(\alpha, \beta; \theta_o, \phi_o) &= B(\alpha, \beta, \theta'_o, \phi'_o) \\
&= B\left(\alpha, \beta; R^{-1}_{\alpha,\beta}(\theta_o, \phi_o)\right).
\end{aligned} \tag{5.3}
$$

Our general approach (and that of previous work [8, 38]) is to represent the incident lighting L by an environment map. The environment map is *prefiltered* to compute some representation of B (or $\tilde{B}$), followed by interactive *rendering* with this representation.

The rest of this chapter covers a number of issues that must be addressed. First, we must find the appropriate (re)parameterization for B and $\hat{\rho}$. Next, we must determine how to represent B in a compact manner suitable for interactive rendering. For this, it is important to analyze the required sampling rates and output resolutions. Finally, we must determine how to efficiently compute our representation of B, i.e. rapidly prefilter the environment map. An overview of our entire algorithm pipeline is shown in figure 5.2.

5.3.1 Reparameterization by central BRDF direction

Our goal is to reparameterize the BRDF and reflected light field so that they become relatively simple and compact, and possibly lower-dimensional functions. Reparameterizing also allows us to eliminate the warping step required by Cabral et al. [8]. To do this, we reparameterize by the central BRDF direction, commonly the reflection vector, as per the discussion in section 2.3.4. We repeat below the important points and discuss some relevant details. Please refer to figure 2.4 in chapter 2 for an illustration of the ideas.

Consider first the special case of *radially symmetric or 1D BRDFs*, where the BRDF consists of a single symmetric lobe of fixed shape, whose orientation depends only on a well-defined central direction $\vec{C}$. In other words, the BRDF is given by a 1D function u as $\hat{\rho} = u(\vec{C} \cdot \vec{L})$. Examples are Lambertian $\hat{\rho} = \vec{N} \cdot \vec{L}$ and Phong $\hat{\rho} = (\vec{R} \cdot \vec{L})^s$ models. If we reparameterize by $\vec{C}$, the BRDF becomes a function of only 1 variable (θ_i with $\cos\theta_i = \vec{C} \cdot \vec{L}$) instead of 3. Further, the reflected light field can be represented simply by a 2D reflection map $B(\alpha, \beta)$ parameterized by $\vec{C} = (\alpha, \beta)$. Note that we will often use

$\vec{R}$, the reflection of the viewing direction about the surface normal, as a synonym for $\vec{C}$ since that is the most common case; however our analysis applies generally.

For general BRDFs, the radial symmetry property does not hold precisely, so they cannot be reduced exactly to 1D functions, nor can B be written simply as a 2D reflection map. Nevertheless, a reparameterization by the reflection vector still yields compact forms. As can be seen for the $\vec{N} \cdot \vec{H}$ model shown in the lower part of figure 2.4, most of the variation in the BRDF is still over only a single variable (θ_i) after reparameterization, while there is very little variation over θ_o (or ϕ) for fixed θ_i. Further, most of the variation in B remains over (α, β), with only low-frequency variation over the other two variables (θ_o, ϕ_o).

To reparameterize, we simply recenter the BRDF (and the reflection integral) about the reflection vector, rather than the surface normal, as shown in figure 2.4. The reflection vector now takes the place of the surface normal, i.e. $\vec{R} = (\alpha, \beta)$, and the dependence on the surface normal becomes indirect (just as the dependence on $\vec{R}$ is indirect in the standard parameterization). The angles θ'_i and θ'_o are now given with respect to $\vec{R}$ by $\cos\theta'_i = \vec{R} \cdot \vec{L}$ and $\cos\theta'_o = \vec{R} \cdot \vec{V}$, with $B(\alpha, \beta, \theta'_o, \phi'_o)$ a function of $\vec{R} = (\alpha, \beta)$ and (θ'_o, ϕ'_o).

Note that, although this paper does not address general 4D anisotropic BRDFs, reparameterization by the tangent vector $\vec{T}$ can be used in special cases to reduce anisotropic BRDFs to 1D or 2D functions, amenable to treatment by our algorithm pipeline. For instance, consider the Kajiya-Kay [35] model. This BRDF is an anisotropic extension of the Phong model, and depends on $\vec{T} \cdot \vec{L}$ and $\vec{T} \cdot \vec{V}$. We may reparameterize by the tangent vector, just as we did above with the reflection vector. Then, $\cos\theta'_i = \vec{T} \cdot \vec{L}$ and $\cos\theta'_o = \vec{T} \cdot \vec{V}$, with the BRDF being given by $\hat{\rho} = \cos^s(\theta'_i - \theta'_o)$. The BRDF is now a simple 2D function, which is only slightly more complex than the reparameterized 1D Phong BRDF.

An important requirement of our reparameterization is that it be suitable for *both* the BRDF $\hat{\rho}$ and the reflected light field B. Thus, Rusinkiewicz's [74] reparameterization of BRDFs, based on the half angle vector $\vec{H}$, cannot be used since it is unsuitable for the reflected light field. $\vec{H}$ depends on the incident light direction, while B depends only on the the viewing direction and surface normal (or reflection vector). Our approach is motivated by the "reflective transformation" used by Wood et al. [88] to reparameterize the outgoing radiance of surface light fields by $\vec{R}$. However, our final representations differ significantly. In their case, the surface light field is parameterized by object geometry and reflection vector $\vec{R}$, while in our case, the "orientation light field" is parameterized by the reflection vector $\vec{R}$ and viewing vector $\vec{V}$.

5.3.2 4D function representations

Our goal is to compactly represent $B(\alpha, \beta, \theta'_o, \phi'_o)$ in a manner suitable for interactive rendering, while exploiting its characteristics. As noted previously [8], the variation over (α, β) may be rapid (high-frequency), while that over (θ'_o, ϕ'_o) is usually slow (low-frequency), reducing to constant (no variation) for radially symmetric 1D BRDFs. There have been a number of representations for 4D functions proposed in the graphics community, primarily for image-based rendering. The main categories are listed below. Table 5.2 compares tabular, compressed, factored, coefficient-based, as well as our SHRM representation, in terms of simplicity, compactness of the representation, efficiency for rendering, ease of error analysis, and speed for computation (prefiltering).

Explicit tabular representation: We may simply tabulate $B(\alpha, \beta, \theta'_o, \phi'_o)$ on a 4D grid. Cabral et al. [8] use a sparse 2D set of standard 2D reflection maps. However, a very large amount of data will be required to accurately tabulate a 4D light field. Cabral et al. [8] use only 12 reflection maps, trading accuracy for compactness. Kautz and McCool [38] approximate the BRDF as a 2D function $\hat{\rho} = u(\theta'_i, \theta'_o)$ having no azimuthal dependence to create a 3D texture $B(\alpha, \beta, \theta'_o)$. This lower-dimensional representation is more compact, but loses some generality, and can still require a large amount of data. Interactive rendering with these methods usually simply involves a texture lookup and interpolation.

Compressed forms: Compression based on vector-quantization or MPEG-like methods can be used to reduce the size of a tabular representation, as done for surface light fields [60, 88]. or multidimensional animations [27]. Care must be taken that the compressed form can be rendered interactively. Note that both computation of and error analysis on the compressed form require us to first compute the dense uncompressed 4D light field, which can be a computation and data-intensive operation.

Method	Simple	Compact	Rendering	Analysis	Speed
Tabular	Yes	No	Yes	No	No
Compressed	No	Yes	Maybe	*	No
Factored	Yes	Yes	Yes	*	No
Coefficients	Yes	Yes	No	Yes	Yes
SHRM	Yes	Yes	Yes	Yes	Yes

Table 5.2: *Comparison of different 4D representations. The columns stand for simplicity/intuitiveness, compactness, efficiency for rendering, ease of error analysis, and speed of computation (prefiltering). We use $*$ for error analysis of compressed and factored representations because, while error analysis is straightforward, it requires knowledge of a densely sampled 4D light field, and cannot be applied directly.*

Factored representation: Factorization can be seen as a simple compression technique that yields compact results suitable for interactive rendering using texture mapping. Previous methods that can be seen as factorizations include eigen-textures [64], polynomial texture maps [53], and BRDF factorizations [37, 57]. The 4D light field is written as the sum of a few terms, each being the product of two 2D functions (textures),

$$B(\alpha, \beta, \theta_o', \phi_o') = \sum_a g_a(\alpha, \beta) h_a(\theta_o', \phi_o'). \tag{5.4}$$

The trick is to find the factors (2D textures) g_a and h_a so only a small number of terms are needed. Rendering can be done interactively by multiplying together texture maps g and h.

Basis function coefficients: Factored representations can be viewed as a basis-function expansion. To see this, we first introduce a representation purely in terms of basis function coefficients:

$$B(\alpha, \beta, \theta_o', \phi_o') = \sum_a \sum_b c_{ab} d_b(\alpha, \beta) h_a(\theta_o', \phi_o'). \tag{5.5}$$

The basis functions are d_b and h_a with coefficients c_{ab}. We need only store c_{ab}, and can evaluate the basis functions procedurally. This is a simple compact form. However, interactive rendering is difficult since there will usually be a large number of coefficients. We may reduce the number of terms by doing the summation over b to get a factored representation identical to equation 5.4, defining

$$g_a(\alpha, \beta) = \sum_b c_{ab} d_b(\alpha, \beta). \tag{5.6}$$

5.4 Spherical Harmonic Reflection Maps

In this section, we introduce the *spherical harmonic reflection map* or SHRM representation, which is a compact factored representation derived from a spherical harmonic basis function expansion. Figure 5.2 shows an overview of our entire pipeline. S and T stand

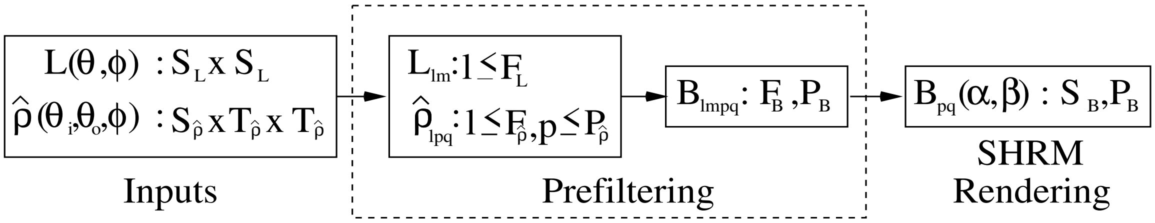

Figure 5.2: *An overview of our entire pipeline. S and T are angular resolutions, while F and P are orders of the spherical harmonic expansions.*

for angular resolutions while F and P stand for orders of the spherical harmonic expansions, which are determined using the theoretical analysis in section 5.5. The inputs to the algorithm are tabulated values of the lighting $L(\theta, \phi)$ and 3D isotropic BRDF $\hat{\rho}(\theta_i', \theta_o', \phi)$.

We then use our fast prefiltering algorithm, described in detail in section 5.6, to compute the SHRM. This is done by first computing the spherical harmonic lighting coefficients L_{lm} and BRDF coefficients $\hat{\rho}_{lpq}$. We then use the spherical frequency-space convolution formula derived in chapter 2, which may be viewed as the frequency domain analog of equation 5.2, to compute reflected light field coefficients B_{lmpq} (c.f. equation 2.55),

$$B_{lmpq} = \Lambda_l L_{lm} \hat{\rho}_{lpq}. \tag{5.7}$$

Now, if we want to expand the coefficients to get the entire reflected light field, we know that (c.f. equation 2.54)

$$\begin{aligned} B(\alpha,\beta,\theta'_o,\phi'_o) &= \sum_{l,m,p,q} \Lambda_l L_{lm}\hat{\rho}_{lpq}\Lambda_l^{-1} D^l_{mq}(\alpha,\beta) Y_{pq}(\theta'_o,\phi'_o) \\ &= \sum_{lmpq} L_{lm}\hat{\rho}_{lpq} \left(D^l_{mq}(\alpha,\beta) Y_{pq}(\theta'_o,\phi'_o) \right). \end{aligned} \tag{5.8}$$

We expand the right-hand side partially (over indices l and m), defining

$$B_{pq}(\alpha,\beta) = \sum_{l,m} L_{lm}\hat{\rho}_{lpq} D^l_{mq}(\alpha,\beta). \tag{5.9}$$

Finally, we may write down the equation for the SHRM as

$$B(\alpha,\beta,\theta_o,\phi_o) = \sum_{p=0}^{P_B} \sum_{q=-p}^{p} B_{pq}(\alpha,\beta) Y_{pq}(\theta'_o,\phi'_o). \tag{5.10}$$

In this equation, $B_{pq}(\alpha,\beta)$ is one coefficient in the SHRM, and $P_B \geq 0$ is the maximum order of the expansion, with the SHRM containing a total of $(P_B+1)^2$ terms. Figure 5.3 illustrates the idea behind SHRMs. Each pixel (α,β) in a reflection (cube)map has a particular distribution of outgoing radiance $B(\alpha,\beta,\theta'_o,\phi'_o)$. This distribution is encoded by the SHRM as a spherical harmonic expansion in (θ'_o,ϕ'_o), with coefficients $B_{pq}(\alpha,\beta)$. For the special case of radially symmetric 1D BRDFs, there is no dependence on (θ'_o,ϕ'_o) after reparameterization, so we need only the DC or constant term $B_{00}(\alpha,\beta)$, and the SHRM reduces to a simple 2D reflection map.

So far, we have considered *local* SHRMs, depending on local outgoing angles (θ'_o,ϕ'_o), which are different for each (α,β). It is often convenient to assume the viewer is distant and compute *global* SHRMs, dependent on a global viewing direction (θ_o,ϕ_o).

$$\hat{B}(\alpha,\beta,\theta_o,\phi_o) = \sum_{p=0}^{P_B} \sum_{q=-p}^{p} \hat{B}_{pq}(\alpha,\beta) Y_{pq}(\theta_o,\phi_o). \tag{5.11}$$

The advantage of equation 5.11 over equation 5.10 lies in ease of evaluation for rendering, since $Y_{pq}(\theta_o,\phi_o)$ can be evaluated once per frame, instead of per pixel. In fact, we will show in section 5.7.5 that this allows global SHRMs to be rendered using a single dynamic

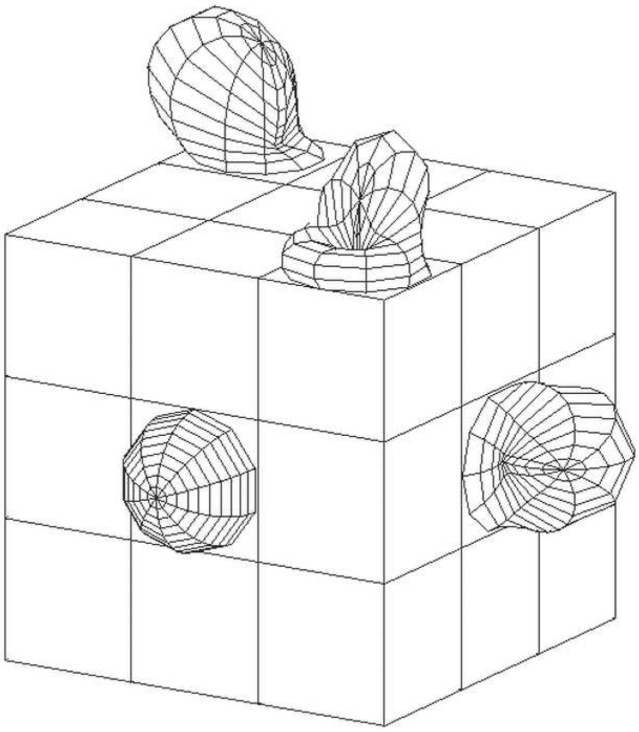

Figure 5.3: *The idea behind SHRMs. Each pixel (α, β) in the reflection cubemap has some distribution of reflected light. This distribution is encoded as a low-order spherical harmonic expansion in (θ'_o, ϕ'_o) for every pixel, with coefficients $B_{pq}(\alpha, \beta)$. To avoid clutter, this diagram uses only a 3x3 resolution in the reflection map and shows the outgoing distribution for only four of the pixels.*

reflection map, with standard reflection mapping hardware.

We still need to know how to determine global SHRM coefficients $\hat{B}_{pq}(\alpha, \beta)$. The spherical convolution formula in equations 5.7 or 2.55 apply only to local SHRMs. However, we may rotate coefficients to compute the global SHRM. We make use of equations 5.3 and 5.10, with the subscript q changed to s for later convenience,

$$\hat{B}\left(\alpha, \beta, \theta_o, \phi_o\right) = \sum_{p=0}^{P_B} \sum_{s=-p}^{p} B_{ps}(\alpha, \beta) Y_{ps}\left(R^{-1}_{\alpha,\beta}(\theta_o, \phi_o)\right). \tag{5.12}$$

From this expression, we can determine the modified (rotated) coefficients of the global SHRM separately for each orientation (α, β). For this, we need to know how to apply the inverse rotation in the equation above, analogous to the forward rotation formula used in equation 2.29. We may derive

$$\begin{aligned} Y_{ps}\left(R^{-1}_{\alpha,\beta}(\theta_o,\phi_o)\right) &= Y_{ps}\left(R_y(-\alpha)R_z(-\beta)\{\theta_o,\phi_o\}\right) \\ &= \sum_{q=-p}^{p} D^p_{sq}(-\alpha)Y_{pq}(R_z(-\beta)\{\theta_o,\phi_o\}) \\ &= \sum_{q=-p}^{p} D^p_{sq}(-\alpha)e^{-Iq\beta}Y_{pq}(\theta_o,\phi_o). \end{aligned} \tag{5.13}$$

Finally, putting equation 5.12 and 5.13, we obtain

$$\hat{B}_{pq}(\alpha,\beta) = \sum_{s=-p}^{p} \left(D^p_{sq}(-\alpha)e^{-Iq\beta}\right) B_{ps}(\alpha,\beta). \tag{5.14}$$

Advantages: SHRMs are a hybrid form, midway between a pure coefficient-based approach, and an explicit tabular representation. We believe this is a good point in the design space, and our representation has the following significant advantages:

- **Compact, Efficient and Accurate:** An important observation from the theoretical analysis is that for essentially all BRDFs, a very low value of P_B (usually ≤ 3) suffices for high accuracy. This is the formal basis for using a low order spherical harmonic expansion in the SHRM, and ensures that our representation is very compact and accurate compared to previous approaches, as well as being efficient to evaluate for real-time rendering.

- **Error analysis and number of coefficients/resolutions:** Unlike for other compression and factorization techniques, the error analysis in section 5.5 does not first require computation of a dense 4D reflected light field, and allows us to easily determine the correct order P_B of the spherical harmonic expansion and the resolutions of the reflection maps.

- **Rapid computation:** In section 5.6, we show how the SHRM can be computed using frequency domain prefiltering, orders of magnitude faster than previous approaches.

5.5 Analysis of sampling rates/resolutions

In this section, we present our framework for analyzing the required sampling rates, i.e. the number of coefficients needed in our spherical harmonic expansions. At the end of the section, we will justify the SHRM representation based on our analysis.

The sampling rates will depend on the frequency content of the lighting and BRDF. Figure 5.4 shows spheres rendered with progressively blurred illumination (along the y axis) and a progressively more diffuse BRDF (along the x axis). It can be seen that the highest frequencies in the reflected light field are determined approximately by the minimum of the highest frequencies in the lighting and BRDF. This is not surprising, since we may view the BRDF as a low pass filter acting on the lighting signal.

As summarized in figure 5.2, we assume the input lighting $L(\theta, \phi)$ is represented on an $S_L \times S_L$ grid, where S_L is the grid angular resolution, and that the 3D isotropic BRDF $\hat{\rho}(\theta'_i, \theta'_o, \phi)$ is represented on a grid of size $S_{\hat{\rho}} \times T_{\hat{\rho}} \times T_{\hat{\rho}}$ where $S_{\hat{\rho}}$ is the angular resolution with respect to θ'_i and $T_{\hat{\rho}}$ is the angular resolution with respect to (θ_o, ϕ). For simplicity, we will consider the lighting and BRDF to be represented in latitude-longitude form, i.e. simply as tabulated values on an equally-spaced grid $0 \leq \theta \leq \pi$ and $0 \leq \phi \leq 2\pi$. Converting to less distorted forms like cubemaps for practical use on graphics hardware is a straightforward table lookup.

Our prefiltering algorithm computes the lighting coefficients L_{lm} to order F_L (i.e. F_L is the maximum value of l) and BRDF coefficients $\hat{\rho}_{lpq}$ to orders $F_{\hat{\rho}}$ and $P_{\hat{\rho}}$ (i.e. $l \leq F_{\hat{\rho}}$ and $p \leq P_{\hat{\rho}}$). The light field coefficients B_{lmpq} are computed to orders F_B and P_B. Finally, we generate the SHRM $B_{pq}(\alpha, \beta)$, with the angular size in (α, β) being $S_B \times S_B$, and the spherical harmonic expansion up to order P_B. Radially symmetric 1D BRDFs can be seen as special cases of this general framework with $T_{\hat{\rho}} = 1$, and $P_{\hat{\rho}} = P_B = 0$.

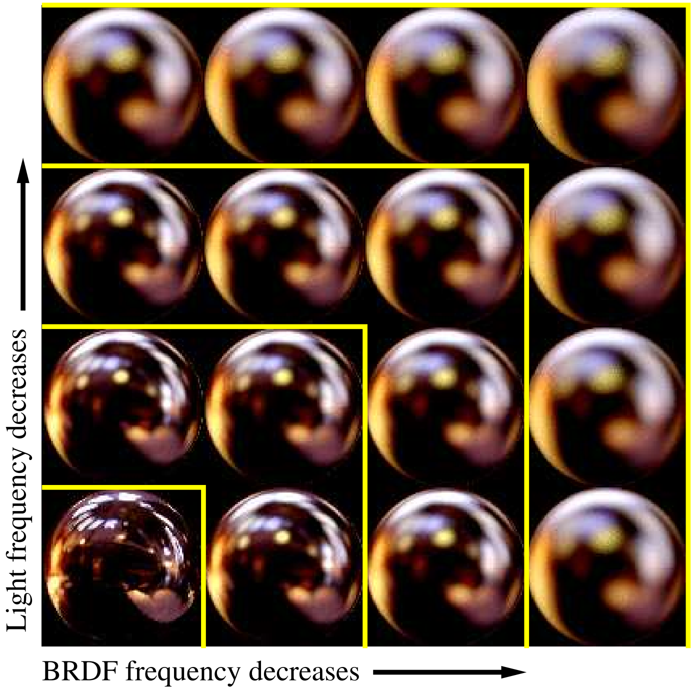

Figure 5.4: *Renderings with different lighting and BRDF conditions. The highest frequency in the images is approximately the minimum of the highest frequencies in the lighting and BRDF. Specifically, all spheres inside a yellow delimiter look similar.*

5.5.1 Order of spherical harmonic expansions

We now analyze required orders F and P in our spherical harmonic expansions. First, consider the lighting. The total energy of the lighting signal is given by

$$\int_{\theta=0}^{\pi}\int_{\phi=0}^{2\pi} L^2(\theta,\phi)\,\sin\theta\,d\theta d\phi = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} \mid L_{lm}\mid^2 . \tag{5.15}$$

Once we compute the total energy in the lighting, we can estimate the error ϵ in an order F_L spherical harmonic expansion by considering what fraction of the total lighting energy is captured up to order F_L. To obtain an accuracy $1-\epsilon$, we require that

$$\sum_{l=0}^{F_L}\sum_{m=-l}^{l} \mid L_{lm}\mid^2 \geq (1-\epsilon)\int_{\theta=0}^{\pi}\int_{\phi=0}^{2\pi} L^2(\theta,\phi)\,\sin\theta\,d\theta d\phi. \tag{5.16}$$

For any given input illumination $L(\theta,\phi)$ and frequency F_L, it is easy to determine what the error ϵ is, and check if it is below threshold. Conversely, if we fix the error threshold ϵ, we can compute the required sampling rate or order F_L as the minimum frequency for which the above equation holds. Note that the number of coefficients required measures the frequency width of the signal, and ϵ measures the missing (residual) information content.

A similar method may be used for analyzing the BRDF. We give below the separate results for the 1D radially symmetric and 3D isotropic cases,

$$\sum_{l=0}^{F_{\hat{\rho}}} \mid \hat{\rho}_l \mid^2 \geq (1-\epsilon) 2\pi \int_{\theta_i=0}^{\pi/2} \hat{\rho}^2(\theta_i) \sin\theta_i \, d\theta_i. \tag{5.17}$$

$$\sum_{l=0}^{F_{\hat{\rho}}} \sum_{p=0}^{P_{\hat{\rho}}} \sum_{q=-\min(l,p)}^{\min(l,p)} \mid \hat{\rho}_{lpq} \mid^2 \geq$$
$$(1-\epsilon) 2\pi \int_0^{\pi} \int_0^{\pi} \int_0^{2\pi} \hat{\rho}^2(\theta_i, \theta_o, \phi) \sin\theta_i \sin\theta_o \, d\theta_i d\theta_o d\phi. \tag{5.18}$$

The remaining issue is how to combine the information for lighting and BRDFs to determine appropriate orders for the reflected light field B. We list below two possible approaches.

- **Minimum of orders or errors:** Consider the case where $\epsilon = 0$ for either the lighting or BRDF, i.e. one or both is bandlimited. The reflected light field is then exactly reproduced by using an expansion to order $(F_B, P_B) = (\min(F_L, F_{\hat{\rho}}), P_{\hat{\rho}})$. This formalizes the intuition that we need to sample densely enough to catch the highest frequency present simultaneously in both the lighting signal and BRDF filter. This analysis does not apply rigorously when neither signal is bandlimited, but simply decays with increasing frequency. Nevertheless, in practice, taking the minimum of orders for a given error ϵ is still a good heuristic. Conversely, for a given order of expansion, we can estimate the error $\epsilon_B = \min(\epsilon_L, \epsilon_{\hat{\rho}})$.

  Since the lighting signal usually contains substantial high frequency content, while the BRDF acts as a low-pass filter, this method often reduces simply to capturing $1-\epsilon$ of the BRDF energy, i.e. choosing $F_B, P_B = F_{\hat{\rho}}, P_{\hat{\rho}}$ or setting $\epsilon_B = \epsilon_{\hat{\rho}}$.

- **Bound residual energy:** For completeness, we discuss a more rigorous numerical scheme, which can be proven to give conservative estimates. The scheme is based on bounding the residual unaccounted for energy in the reflected light field. One disadvantage of this method is that, unlike the previous method, we first need to actually calculate the coefficients B_{lmpq} of the reflected light field. Thus, this method is most useful as a final sanity check on the validity of the earlier heuristic.

To derive this bound, we consider an expansion to order F with errors for lighting and BRDF given by ϵ_L and $\epsilon_{\hat{\rho}}$. Denote the total BRDF and lighting energies by $\hat{\rho}_{tot}$ and L_{tot}. Since the reflected light field coefficients are simply a product of lighting and BRDF terms, the worst case for the residual energy occurs when it is all concentrated in mode $F+1$. This residual energy, denoted by B_{res}, and a conservative error estimate ϵ_B are

$$\begin{aligned} B_{res} &= \Lambda_{F+1}\left(\epsilon_L L_{tot}\right)\left(\epsilon_{\hat{\rho}}\hat{\rho}_{tot}\right) \\ \epsilon_B &\leq \frac{B_{res}}{\sum_{l=0}^{F}\sum_{m=-l}^{l}\left|B_{lm}\right|^2 + B_{res}}. \end{aligned} \tag{5.19}$$

This is for the radially symmetric case; in general, we simply use B_{lmpq} in place of B_{lm}. Note that ϵ_B tends to 0 as B_{res} tends to 0. But the latter quantity is a product of $\epsilon_{\hat{\rho}}$ and ϵ_L, and therefore always tends to 0 as F increases.

We use a simple example to illustrate these methods. For a particular illumination (the Grace Cathedral), and a Phong BRDF (exponent $s = 32$), we computed approximations to the lighting, BRDF, and reflected light field for increasing values of order $F = F_L = F_{\hat{\rho}} = F_B$. Since the BRDF is radially symmetric, $P = P_{\hat{\rho}} = P_B = 0$. We also computed the reflected light field accurately, by using a very high order $F = 30$, so we could determine the errors of lower-order approximations. Figure 5.5 plots the accuracy (top curve) of an order F approximation of B, as well as estimates of this accuracy obtained by taking the minimum of BRDF and light errors at order F, and by bounding the residual energy. We see that both accuracy estimates are conservative but fairly tight, especially for small errors or high accuracies (at higher frequencies). Further, taking the minimum of lighting and BRDF errors is almost always equivalent simply to using the BRDF error. Therefore, we choose the simplest approach of using the BRDF error, requiring $\epsilon_{\hat{\rho}}$ be lower than a tolerance that the user has selected.

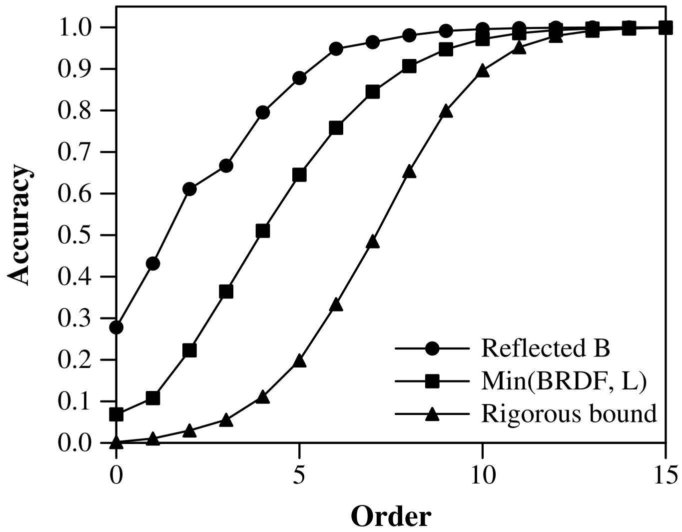

Figure 5.5: *Accuracy $(1-\epsilon)$ versus frequency F for an order F approximation of the reflected light field B , and estimates of that accuracy obtained by taking the minimum error for BRDF and lighting, and by using the conservative bound based on residual energy. We have not separately plotted using the BRDF error only, as this gives almost exactly the same curve as taking the minimum error for BRDF and lighting.*

5.5.2 Justification for SHRM representation

We seek to determine the best point in the spectrum of time/space or angular/frequency tradeoffs. For this, we must understand how to relate angular space resolutions S and T to frequency-space orders F and P. As a simple illustration, consider irradiance maps from Lambertian BRDFs. It has been shown [2, 72] that an order 2 spherical harmonic expansion suffices. However, a 3×3 irradiance map will clearly be inadequate. In practice, irradiance maps are usually represented at angular resolutions higher than 16×16. Experimenting with different resolutions, we have found that in general, one requires $S \sim 10F$ (and $T \sim 10P$).

Therefore, a significantly more compact size for B is obtained using spherical harmonic coefficients rather than an explicit 4D tabular representation. The other extreme in the spectrum of time-space tradeoffs—using a purely coefficient-based approach—is also usually[1] undesirable. Efficient rendering of 2D reflection maps having high frequency content, such as specular reflection maps from Phong BRDFs, is difficult directly from the spherical harmonic expansion, since $O(F^2)$ terms must be added per pixel, with F generally larger than 10. Rendering the 4D light field purely from coefficients is even harder, requiring $O(F^2P^2)$ terms.

Hence, we believe an intermediate representation, allowing for both compact representation, and fast rendering, is optimal. In order to determine the best representation for B, we must know the common values for orders F and P (and hence resolutions S and T). Our results in section 5.7 show that for practically all currently available analytic and measured BRDFs, values of $F \leq 30$ and $P \leq 5$ suffice for an accuracy greater than 90%. Therefore, it is best to encode the view dependence (θ_o', ϕ_o') as a compact (and easily evaluated) spherical harmonic expansion consisting of $(P+1)^2$ terms, while explicitly representing the high-frequency dependence on (α, β). This is the approach taken by SHRMs, where each pixel (α, β) stores coefficients $B_{pq}(\alpha, \beta)$ of an order P spherical harmonic expansion.

5.6 Prefiltering

We now describe our efficient frequency space prefiltering algorithms to create the SHRM and efficiently implement the pipeline in figure 5.2. Frequency space representations allow for very efficient computations and are also very compact, requiring evaluation of fewer terms. Hence our algorithms are much faster than previous angular-space methods. We will present an analysis of the computational complexity of our algorithms, and end this section by validating our conclusions on the Phong BRDF.

[1] For very diffuse BRDFs ($F_{\hat{\rho}}$ and $P_{\hat{\rho}}$ both very small), a purely coefficient-based approach may be acceptable. The most notable example is the Lambertian BRDF ($F_{\hat{\rho}} = 2$, $P_{\hat{\rho}} = 0$), where a 9 term spherical harmonic expansion suffices, as seen in the previous chapter.

5.6.1 Main steps and insights

Our prefiltering method has two main components. First, we must efficiently convert between angular and frequency space descriptions. Second, we must efficiently compute coefficients of the reflected light field from lighting and BRDF coefficients. Both components can be performed rapidly because of the insights below. The algorithm itself is just a direct three step efficient implementation of the pipeline of figure 5.2. Implementation details, and the time complexities of the various steps, are found in section 5.6.2.

Linear time complexity of convolution formula: The reflected light field coefficients B_{lmpq} can be computed in time *linear* in the number of output coefficients B_{lmpq} simply by applying the spherical convolution formula in equation 5.7.

Fast conversion to and from spherical harmonics: We still need to convert from an angular space representation of $L(\theta, \phi)$ and $\hat{\rho}(\theta_i, \theta_o, \phi)$ to the spherical harmonic coefficients, as well as generate the SHRM from B_{lmpq}. As an example, consider computation of lighting coefficients L_{lm}. For any l, m we have

$$L_{lm} = \int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} L(\theta, \phi) Y_{lm}^{*}(\theta, \phi) \sin\theta \, d\theta d\phi. \tag{5.20}$$

The cost of performing this integral is $O(S_L^2)$. Since we must do this for all coefficients, it would appear the total cost would be $O(F_L^2 S_L^2)$. In fact, we can amortize the cost to compute all the coefficients in $O(F_L S_L^2)$ time by writing the spherical harmonics as products of functions in θ and ϕ, and then separating the computations in θ and ϕ. The basic idea is to compute in succession:

$$\begin{aligned} \forall m, \theta: \quad L_m(\theta) &= \int_{\phi=0}^{2\pi} L(\theta, \phi) e^{-Im\phi} \, d\phi \\ \forall l, m: \quad L_{lm} &= \int_{\theta=0}^{\pi} L_m(\theta) f_{lm}(\theta) \sin\theta \, d\theta. \end{aligned} \tag{5.21}$$

Here, the spherical harmonic $Y_{lm}(\theta, \phi) = f_{lm}(\theta) e^{Im\phi}$. The first step involves a loop over $(2F_L+1)S_L$ elements, each step of which requires numerical integration by adding together S_L values. Thus, the cost is $O(F_L S_L^2)$. A similar argument shows that the second step takes time of $O(F_L^2 S_L)$. Since $F_L < S_L$, the first step dominates, and the total complexity is $O(F_L S_L^2)$. Note that a similar approach can also be used to efficiently compute BRDF coefficients, and for the inverse transform to compute the SHRM. Fast spherical harmonic transform methods [61], analogous to the Fast Fourier Transform, may reduce the cost further[2] to logarithmic in F_L, i.e $O(S_L^2 \log^2 S_L) \sim O(S_L^2 \log^2 F_L)$. However, these methods are complicated, and although asymptotically faster, have relatively large constant cost factors. Therefore, they are unlikely to be significantly faster for the low frequencies $F \sim 30$ relevant for us.

5.6.2 Prefiltering Algorithm

Below, we give full details of our prefiltering algorithm, which is a three-step process that implements the pipeline in figure 5.2.

Step 1. Compute lighting and BRDF coefficients: We go from the inputs $L(\theta, \phi)$ and $\hat{\rho}(\theta_i, \theta_o, \phi)$ to the spherical harmonic coefficients L_{lm} and $\hat{\rho}_{lpq}$. We first compute the spherical harmonic coefficients of the BRDF using the following three step algorithm.

$$\begin{aligned}
\hat{\rho}_q(\theta_i, \theta_o) &= 2\pi \int_0^{2\pi} \hat{\rho}(\theta_i, \theta_o, \phi) \cos q\phi \, d\phi \\
\hat{\rho}_{pq}(\theta_i) &= \int_0^{\pi} \hat{\rho}_q(\theta_i, \theta_o) f_{pq}(\theta_o) \sin\theta_o \, d\theta_o \\
\hat{\rho}_{lpq} &= \int_0^{\pi} \hat{\rho}_{pq}(\theta_i) f_{lq}(\theta_i) \sin\theta_i \, d\theta_i.
\end{aligned} \tag{5.22}$$

The computational costs of the three terms in the above sequence are given by, respectively, $O(P_{\hat{\rho}} S_{\hat{\rho}} T_{\hat{\rho}}^2)$, $O(P_{\hat{\rho}}^2 S_{\hat{\rho}} T_{\hat{\rho}})$, and $O(F_{\hat{\rho}} P_{\hat{\rho}}^2 S_{\hat{\rho}})$. Since $P_{\hat{\rho}} < T_{\hat{\rho}}$, the first term dominates the second, and the net cost is $O\left(P_{\hat{\rho}} S_{\hat{\rho}} (T_{\hat{\rho}}^2 + F_{\hat{\rho}} P_{\hat{\rho}})\right)$. For most non radially-symmetric BRDFs, $T_{\hat{\rho}}^2 > F_{\hat{\rho}} P_{\hat{\rho}}$ (for instance, use $T_{\hat{\rho}} \sim 10$, $F_{\hat{\rho}} \sim 20$ and $P_{\hat{\rho}} \sim 3$), so the first term dominates and the total cost is $O(P_{\hat{\rho}} S_{\hat{\rho}} T_{\hat{\rho}}^2)$.

If our error tolerance ϵ is satisfied, we see how far $P_{\hat{\rho}}$ can be reduced to still satisfy the error tolerance, and then also reduce $F_{\hat{\rho}}$ as much as possible. We can then set F_B and

[2] Since $S_L \sim 10F_L$, $\log S_L \sim \log F_L$. Also note that simply using an FFT in step 1 will not suffice, since step 2 does not have logarithmic complexity.

P_B according to the minimal values of $F_{\hat{\rho}}$ and $P_{\hat{\rho}}$. If the error is initially larger than ϵ, we repeat the algorithm with larger values for $F_{\hat{\rho}}$ and $P_{\hat{\rho}}$. Since computing BRDF coefficients is not the dominant algorithm cost, this recomputation does not significantly affect the total time, nor does using large initial values for $F_{\hat{\rho}}$ and $P_{\hat{\rho}}$.

Finally, we compute the lighting coefficients in time $O(F_B S_L^2)$. Note that we have already determined F_B, so we are not required to consider higher frequencies for the lighting, which is why we use F_B instead of F_L.

Step 2. Find reflected light field coefficients: We now find $B_{lmpq} = L_{lm}\hat{\rho}_{lpq}$ in time $O(F_B^2 P_B^2)$. Note that we have dropped the constant Λ_l compared to equation 5.7 to avoid dealing with the factors of Λ_l^{-1} and Λ_l.

Step 3. Compute SHRM: We now compute the local SHRM. We need to compute

$$B_{pq}(\alpha, \beta) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lmpq} \left(D_{mq}^{l}(\alpha) e^{Im\beta} \right), \tag{5.23}$$

which can be done efficiently as shown below,

$$\begin{aligned} B_{mpq}(\alpha) &= \sum_{l=|m|}^{F_B} B_{lmpq} D_{mq}^{l}(\alpha) \\ B_{pq}(\alpha, \beta) &= \sum_{m=-F_B}^{F_B} B_{mpq}(\alpha) e^{Im\beta}. \end{aligned} \tag{5.24}$$

From this, we can compute the *global* SHRM $\tilde{B}_{pq}(\alpha, \beta)$ using equation 5.14. The costs of the two terms in the above sequence are $O(F_B^2 P_B^2 S_B)$ and $O(F_B P_B^2 S_B^2)$. Since $S_B > F_B$, the net cost is $O(F_B P_B^2 S_B^2)$. The cost for this step is also the dominant cost for the entire algorithm.

Radially symmetric BRDFs: For the special case of BRDFs like Lambertian and Phong models, $T_{\hat{\rho}} = 1$ and $P_{\hat{\rho}} = P_B = 0$. Technically, the complexity formulae above should use $P + 1$ instead of P, to yield meaningful results for radially symmetric BRDFs. For these models, step 1 takes time of $O(F_{\hat{\rho}} S_{\hat{\rho}})$ to compute BRDF coefficients $\hat{\rho}_l$, and time $O(F_B S_L^2)$

to compute lighting coefficients L_{lm}. Step 2 takes $O(F_B^2)$ time. Finally, the SHRM in step 3 includes only the constant term and is therefore a simple reflection map $B(\alpha, \beta)$, computed in time $O(F_B S_B^2)$. The dominant cost here is to convert to and from spherical harmonic representations. Assuming we downsample the environment map if necessary so $S_L \sim S_B$, the total time is $O(F_B S_B^2)$ or $O(F_B)$ per output image pixel.

Conversion between SHRMs and explicit forms: It is possible to incorporate the prefiltering and rendering phases of our algorithm separately into existing systems. SHRMs may be created from explicit representations simply by fitting coefficients or integrating. If the implementer wants to use only our fast prefiltering method, but render using previous explicit representations, they can compute tabular representations from SHRMs. Cabral's twelve prerendered reflection maps may be computed very rapidly using equation 5.35, with $(\tilde{\theta_o}, \tilde{\phi_o})$ set to vertices of an icosahedron. Kautz and McCool's [38] 3D texture is computed by expanding

$$B(\alpha, \beta, \theta_o) = \sum_{p=0}^{P_B} B_{p0}(\alpha, \beta) Y_{p0}(\theta_o). \tag{5.25}$$

This takes time $O(P_B)$ per output texel. Using fast conversion methods, we can also explicitly generate $T_B \times T_B$ reflection maps (a full 4D light field) in time $O(P_B)$ per output pixel, for a total cost of $O(P_B T_B^2 S_B^2)$.

5.6.3 Computational complexity

The cost of previous angular domain algorithms is $O(W S_L^2)$ per pixel in the output, since they perform a hemispherical integral for each pixel. Here, W is the fraction of the illumination pixels that need be considered, corresponding to the angular width of the BRDF, with $W \to 0$ for a mirror, and $W = 1/2$ if one considers the entire visible hemisphere. In appendix C, we derive the cost for our frequency space algorithm[3], which is much lower, being $O(F_B)$ or $O(P_B)$ per pixel. Table 5.3 summarizes our main results.

[3]If we were to use fast spherical harmonic transform methods [61], the asymptotic complexity per output pixel would be $O(\log^2 F_B)$ or $O(\log^2 P_B)$ instead.

Type	Angular		Frequency	
	Cost	/pixel	Cost	/pixel
1D BRDF	$WS_L^2S_B^2$	WS_L^2	$F_BS_B^2$	F_B
SHRM			$F_BP_B^2S_B^2$	F_B
3D BRDF	$WS_L^2T_B^2S_B^2$	WS_L^2	$P_BT_B^2S_B^2$	P_B

Table 5.3: *Computational complexity of prefiltering. We show both total costs and costs per pixel. The angular costs correspond to hemispherical integration, with W being the BRDF angular width. The frequency space costs are for our method.*

Radially Symmetric 1D BRDFs: The output reflection map size is $S_B \times S_B$. Standard hemispherical integration is *quadratic* in the output size, since we must examine $O(WS_L^2)$ illumination pixels per output pixel and[4] $S_L \geq S_B$. By contrast, our frequency space prefiltering algorithm requires only $O(F_B)$ time per output pixel. Since $F_B \ll S_B < S_L$, this method is only slightly *super-linear*, being substantially *sub-quadratic* in the output size.

SHRM creation: The SHRM is a new representation, not created by traditional prefiltering algorithms. We can still compare our prefiltering cost with the theoretical minimal complexity. As for 1D BRDFs, the cost of our method is $O(F_B)$ per output pixel per coefficient.

3D BRDFs, explicit reflection maps: In order to compare to angular domain methods, we must produce the same output. We can use the SHRM to explicitly compute a set of $T_B \times T_B$ reflection maps (with $T_B \sim 10P_B$), similar to the explicit representations of Cabral et al. [8] or Kautz and McCool [38][5]. The cost of traditional prefiltering remains $O(WS_L^2)$ per output pixel. On the other hand, our method takes $O(P_B)$ time per pixel. Since $P_B \leq 3$ in most cases, it can be regarded a constant. Hence, our method is *quasi-linear* in the output size. This is a speedup of three to four orders of magnitude—the difference between near-interactive computation in a few seconds, and prefiltering times in hours reported by Cabral et al. [8] and other authors.

[4] One way to compute F_B is $\min(F_L, F_{\hat{\rho}})$ so $F_B \leq F_L$ and $S_B \leq S_L$.

[5] Since their representation is 3D, we should compute only T_B reflection maps.

5.6.4 Validation with Phong BRDF

In this subsection, we validate our theoretical computational complexity analysis on the simple radially symmetric Phong model. In this case, $P_{\hat{\rho}} = P_B = 0$ and the SHRM reduces to a standard 2D reflection map $B(\alpha, \beta)$. In the results section, we show timings, including for more general 3D isotropic BRDFs.

The normalized and reparameterized Phong BRDF is defined by

$$\hat{\rho} = \frac{s+1}{2\pi} \cos^s \theta_i, \tag{5.26}$$

where $\cos^s \theta_i = (\vec{R} \cdot \vec{L})^s$. BRDF coefficients $\hat{\rho}_l$ can be derived analytically [73], and an accurate approximation is

$$\hat{\rho}_l \approx \Lambda_l^{-1} \exp\left[-\frac{l^2}{2s}\right]. \tag{5.27}$$

In the frequency domain, the order $F = F_B = F_{\hat{\rho}}$ for an error ϵ is found by the following sequence of steps,

$$\begin{aligned} \sum_{l=0}^{F} \hat{\rho}_l^2 &= (1-\epsilon) \sum_{l=0}^{\infty} \hat{\rho}_l^2 \\ \int_0^F l e^{-l^2/s}\, dl &\approx (1-\epsilon) \int_0^{\infty} l e^{-l^2/s}\, dl \\ 1 - e^{-F^2/s} &\approx 1 - \epsilon \\ F &\approx \sqrt{-s \log \epsilon}. \end{aligned} \tag{5.28}$$

Ignoring the constant $\sqrt{-\log \epsilon}$, the prefiltering cost of our frequency space algorithm is therefore

$$C_f = O(F_B S_B^2) = O(S_B^2 \sqrt{s}). \tag{5.29}$$

In the angular domain, we may truncate the BRDF, so $(1-\epsilon)$ of the angular domain energy lies in $\theta_i \leq \theta_i^*$. We find the angular width W and θ_i^* by the following steps,

$$
\begin{aligned}
\int_0^{\theta_i^*} \cos^s \theta_i \sin \theta_i \, d\theta_i &= (1-\epsilon) \int_0^{\pi/2} \cos^s \theta_i \sin \theta_i \, d\theta_i \\
1 - \cos^{s+1} \theta_i^* &= 1 - \epsilon \\
\cos \theta_i^* &= \epsilon^{1/(s+1)} \\
W = \frac{2\pi}{4\pi} \int_0^{\theta_i^*} \sin \theta_i \, d\theta_i &= \frac{1}{2} \left(1 - \epsilon^{1/(s+1)}\right). \qquad (5.30)
\end{aligned}
$$

We now assume s is large ($1/s \to 0$) and perform a Taylor series expansion,

$$
W \sim \left(1 - \epsilon^{1/(s+1)}\right) = \left(1 - \left[1 + \frac{\log \epsilon}{s+1}\right]\right) \approx \frac{-\log \epsilon}{s}. \qquad (5.31)
$$

Ignoring the factor of $\log \epsilon$, the angular domain prefiltering cost is

$$
C_a = O(W S_L^2 S_B^2) = O\left(\frac{S_L^2 S_B^2}{s}\right). \qquad (5.32)
$$

We note that the frequency space cost C_f increases with increasing Phong exponent as $\sqrt{s}$, while the angular space cost C_a decreases with increasing Phong exponent as $1/s$. This is entirely expected, since sharp specular surfaces (large s) have a BRDF which is very local in the angular domain but requires a large number of coefficients to represent in the frequency domain. Conversely, rough surfaces (small s) are very easily handled in the frequency domain, but their BRDFs have a large angular width. Therefore, for $s < s^*$, our frequency domain methods are to be preferred and for $s > s^*$, conventional angular domain techniques are preferable. s^* can be found by equating equations 5.29 and 5.32,

$$
S_B^2 \sqrt{s^*} \sim S_L^2 S_B^2 / s^* \Rightarrow s^* \sim S_L^{4/3}. \qquad (5.33)
$$

What does this mean numerically? Assume a small size of $S_L = 100$. We then obtain $s^* \approx 464$. Therefore, in essentially all practical cases of interest, our frequency domain algorithm is superior to conventional angular domain methods, often by one to two orders of magnitude. Of course, the actual numerical value for s^* depends on the constant cost factors associated with the respective implementations. Our empirical tests, discussed in section 5.7.4, show that the practical value of s^* is actually even higher than predicted.

5.7 Results

We have tested our method using a number of different lighting conditions and BRDFs. This section reports our main results.

5.7.1 Number of coefficients for analytic BRDFs

The practical values of the orders in the spherical harmonic expansion of the BRDF $F_{\hat{\rho}}$ and $P_{\hat{\rho}}$ (and hence F_B and P_B) will depend on the form of the BRDF, with slowly varying BRDFs requiring fewer coefficients. As a basic test to determine reasonable empirical values, we considered three general analytic BRDFs.

Microfacet: Consider a simplified microfacet [84] model,

$$\hat{\rho} = \frac{1}{4\pi\sigma^2} e^{-(\theta_h/\sigma)^2}, \tag{5.34}$$

where $\theta_h = \cos^{-1}(\vec{N} \cdot \vec{H})$ is the angle between the normal and the half-angle vector. Approximations to reflection maps with different values of $P = P_{\hat{\rho}} = P_B$ are shown in figure 5.6. As expected, the accuracy improves as we use higher values of P. Specifically, $P = 2$ suffices to produce very accurate results, with the BRDF error $\epsilon < .03$. Recall from section 5.5.1 that we use the BRDF accuracy as a conservative estimate of the accuracy of the reflected light field. For this BRDF, $F = F_{\hat{\rho}} = F_B$ is given approximately by $F \sim \sigma^{-1}$, and ranges from 10 to 30 for common values of $\sigma \sim 0.1$. These values of F and P are typical for most BRDFs. In general, P is very small, while F is usually much larger.

Lafortune BRDF: We also tested the model of [44], with coefficients obtained from the skin measurements of Marschner et al. [55]. Although the behavior is more interesting, with much stronger specularities exhibited toward grazing angles, a value of $P = 4$ still suffices for an error $\epsilon < .03$.

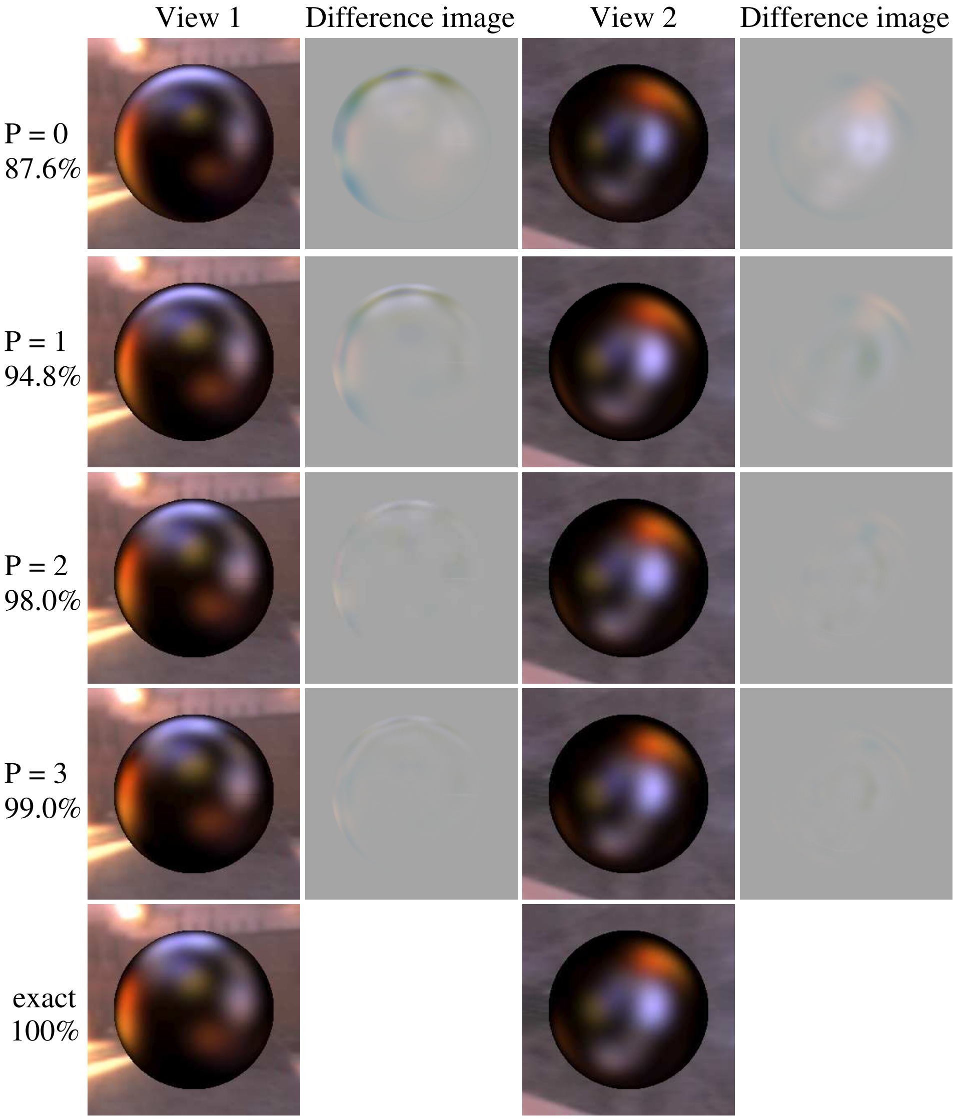

Figure 5.6: *Comparing images obtained with different values for P for a simplified microfacet BRDF model with surface roughness $\sigma = 0.2$. These images correspond to two particular views, i.e. values of (θ_o, ϕ_o). The percentages are fractions of the total energy $(1-\epsilon)$ of the BRDF captured for that P, which we use as a conservative estimate of the accuracy of the reflected light field. The exact images at the bottom were computed by a full hemispherical angular-space integral for each image pixel. For this and subsequent figures, the difference images are not amplified, and we used $F_B = 30$ and $S_B = 128$.*

Kajiya-Kay model: Finally, we tried the Kajiya-Kay [35] model, which is an anisotropic variant of the Phong BRDF, and depends on incident and outgoing angles with respect to the tangent vector. As discussed in section 5.3.1, we may reparameterize by the tangent vector, to derive $\hat{\rho} = \cos^s(\theta_i - \theta_o)$. While this paper does not consider general 4D anisotropic BRDFs, we can handle the Kajiya-Kay BRDF, since it is mathematically analogous to a (2D) isotropic BRDF after reparameterization. Unlike for ordinary Phong-like BRDFs, we cannot apply any further reflective reparameterization. Therefore, the value of P required is large (we found $P = 8$ for $s = 32$ and $\epsilon < .03$). However, there is no azimuthal dependence, so we require only $P + 1$ terms $B_{p0}(\alpha, \beta)$ in the SHRM instead of $(P + 1)^2$ (i.e. $q = 0$, with no dependence on ϕ_o). Hence, the SHRM is still a very efficient and compact representation.

5.7.2 Number of coefficients for measured BRDFs

To further evaluate the accuracy of our approximations, we used the CURET database [14]. This database consists of 61 BRDFs and BTFs, corresponding to a variety of materials. For each sample, there are 205 BRDF measurements, which may be interpolated by fitting order 8 Zernike polynomials to create a complete BRDF description [43].

Figure 5.7 is a bar chart showing, for each of the 61 samples, the accuracy of a BRDF approximation[6] with $F_{\hat{\rho}} = 30$ and values of $P_{\hat{\rho}}$ ranging from 0 to 5. In 56 cases, the accuracy for $F_{\hat{\rho}} = 30$ and $P_{\hat{\rho}} = 5$ was greater than 90% ($\epsilon < 0.1$), and was usually significantly higher (in most cases, $\epsilon < .05$ for $P_{\hat{\rho}} = 3$). The remaining 5 examples (9-frosted glass, 23-lettuce leaf, 33-slate a, 41-brick b, 57-peacock feather) were all significantly anisotropic.

Therefore, we conclude that for almost all BRDFs of interest, an order $P_B \leq 5$ suffices for the SHRM, with $F \leq 30$. In fact, for most BRDFs, a quadratic or cubic (second or third order with $P_B = 2$ or 3) spherical harmonic expansion in the SHRM suffices.

[6]We reparameterized all BRDFs by the reflection vector. Our results demonstrate that this reparameterization is suitable even if the BRDF is not primarily reflective, or consists of both diffuse and specular components. The specular components are compactly represented, while the diffuse components are low frequency anyway.

5.7.3 SHRM accuracy

We now compare images created with SHRMs to the correct image, and to previous approaches. First, figure 5.8 compares our method to Kautz and McCool's [38] 3D texture-mapping technique[7], where they approximate the BRDF—and hence, the reflected light field—as having no azimuthal dependence. For the relatively complex velvet BRDF (CURET database) in figure 5.8, their approximation introduces large errors, while the SHRM with $P_B = 5$ is accurate.

Figure 5.9 compares our approach to the correct image and Cabral's icosahedral interpolation. For sharply-varying BRDFs, such as those exhibiting strong near-grazing specularities, or complex anisotropic behavior, Cabral's approximation can lead to large errors, while our approach still gives accurate results.

In our approach, the theoretical analysis can be used to systematically trade off accuracy for compactness and efficiency. Specifically, if Kautz and McCool's [38] approximation of 2D BRDFs with no azimuthal dependence suffices ($q = 0$), we get a 3D SHRM with only $P_B + 1$ terms instead of $(P_B + 1)^2$. If Cabral et al.'s [8] icosahedral set of 12 reflection maps suffices, we can use a very small number of terms ($P_B = 1$ or 2) in the SHRM.

5.7.4 Speed of prefiltering

We first consider Phong BRDFs, experimentally validating the theoretical conclusions of section 5.6.4. For our frequency domain algorithm, we used $\epsilon = .01$, conservatively setting $F_B = 1 + \sqrt{6s}$. For the angular domain, we were more aggressive, setting $\epsilon = .05$. The resolution S_L and S_B of the inputs and final results were 128, i.e. we generated output Phong reflection maps at 128×128 resolution. According to the theory, this is an appropriate resolution for $s = 32$ and $s = 64$ (i.e. $F_B \approx 12$), and is therefore the most suitable single resolution for the entire range of Phong exponents. The numerical running times reported in table 5.4 obviously depend on our implementation and hardware. However, we believe the *ratio* in running times of angular and frequency domain methods is quite representative. Furthermore, the timing data can be fit almost precisely to the theoretical

[7] We use their single lobe model, with the BRDF being an arbitrary 2D function $\hat{\rho} = u(\theta_i, \theta_o)$. This is essentially equivalent to setting $q = 0$ in the local SHRM, using only the azimuthally independent terms.

predictions of equations 5.29 and 5.32. The results indicate that our prefiltering method is usually two orders of magnitude faster than angular-space methods, and that Phong BRDFs can be prefiltered at close to real-time rates using our approach.

Exponent s	Time (sec)		Ratio
	Angular	Frequency	(Ang/Freq)
8	67.28	0.081	830.6
16	38.03	0.114	333.6
32	21.80	0.159	137.1
64	11.94	0.227	52.6
128	7.17	0.328	21.9
256	3.55	0.461	7.7
512	2.28	0.686	3.3

Table 5.4: *Comparison of timings of angular and frequency-space prefiltering for different values of the Phong exponent s. The timings are on a 1.4GHz Pentium IV.*

As predicted by complexity analysis, the speedups are even more dramatic for the general case—illustrated using the microfacet model of equation 5.34. In table 5.5, we compare computation time for our approach and conventional methods. It can be seen that even the cost for creating the entire SHRM is much less than the cost of hemispherical integration for a single reflection map. When the cost to explicitly create multiple reflection maps is considered, our approach is three to four orders of magnitude faster.

σ	F_B, P_B	angular-space		frequency-space	
		Time (s)	Time/Image (s)	Time (s)	SHRM (s)
.1	24,3	923	9.23	2.70	1.55
.2	12,2	2744	27.44	1.55	0.72
.3	7,2	5731	57.31	1.49	0.67
.4	5,2	9034	90.34	1.47	0.65
.5	5,2	12580	125.80	1.45	0.64

Table 5.5: *Times for angular-space and our frequency-space prefiltering , with $T_B = 10$. The six columns are the value of the roughness σ, the order of expansion F_B, P_B for $\epsilon < .03$, the total angular-space computational time to create $T_B \times T_B = 100$ reflection maps, the angular-space time per reflection map, the total frequency-space time, and the frequency-space time for SHRM creation (but not explicit generation of reflection maps). Our approach is orders of magnitude faster, and even creation of the entire SHRM is usually faster than generating only a single image in angular space.*

5.7.5 Real-time rendering

There are several possibilities for real-time rendering. We could simply evaluate equation 5.10 in software for each pixel of the final image. If hardware multitexturing support is available, we may represent the spherical harmonics $Y_{pq}(\theta_o, \phi_o)$ and the local SHRM coefficients $B_{pq}(\alpha, \beta)$ by 2D texture maps. Since we are reparameterizing by the reflection vector, we will sometimes also refer to $B_{pq}(\alpha, \beta)$ as reflection maps. If $P_B = 2$, there would be 9 terms in the SHRM, corresponding to a total of 18 texture maps. We would then use graphics hardware to accumulate 9 terms, with each term being the product of two texture maps, i.e. $B_{pq}(\alpha, \beta)Y_{pq}(\theta_o, \phi_o)$. Since this algorithm is essentially that previously used for rendering factored BRDFs [37, 57], the same code can now be easily adapted for arbitrary isotropic BRDFs and complex illumination.

A simpler approach is possible when the viewer can be assumed distant, using the global SHRM in equation 5.11. The spherical harmonics $Y_{pq}(\tilde{\theta}_o, \tilde{\phi}_o)$ need be evaluated only once per frame, for given viewpoint $(\tilde{\theta}_o, \tilde{\phi}_o)$, instead of at each vertex or pixel. In fact, it is possible to render the scene using only a single reflection mapping pass. The key idea is to explicitly sum equation 5.11 to create a single dynamic 2D reflection map $\tilde{B}(\alpha, \beta)$, which is updated for every frame, i.e. each new viewpoint $(\tilde{\theta}_o, \tilde{\phi}_o)$,

$$\tilde{B}(\alpha, \beta) = \sum_{p=0}^{P_B} \sum_{q=-p}^{p} \tilde{B}_{pq}(\alpha, \beta) Y_{pq}(\tilde{\theta}_o, \tilde{\phi}_o). \tag{5.35}$$

Our implementation extends the Stanford real-time programmable shading system [68] to render with global SHRMs using equation 5.35. An advantage of our approach is that standard reflection maps can be upgraded to SHRMs with no change in the external shader programs. Internally, we simply update the reflection map for each frame. We compute equation 5.35 in software, which allows us to easily consider high-dynamic range, and avoids hardware precision and clamping issues. In the figures, the high-dynamic range backgrounds are tone-mapped, but the objects themselves are computed and shaded using a linear scale.

We used a 1.4 GHz Pentium IV running Linux, with an NVIDIA Geforce2 GTS graphics card, for our tests. The reflection (cube)map $\tilde{B}(\alpha, \beta)$ was computed at a resolution of

$64 \times 64 \times 6$, which is an appropriate resolution for most BRDFs, i.e. $F \sim 20$. Since there is real-time cube mapping hardware, the major cost is that for computing equation 5.35 in software per frame. We are able to achieve frame rates of approximately 30 frames per second, with real-time speeds even in scenes with multiple SHRMs. Figure 5.1 shows a number of examples.

5.8 Conclusions and Future Work

In this chapter, we have presented new frequency-space algorithms for real-time rendering of complex isotropic BRDFs under arbitrary distant illumination, and have validated our approach using many different BRDFs and lighting conditions. Our contributions include theoretical analysis that allows us to precisely determine the orders of our spherical harmonic expansions, the new compact and efficient SHRM representation for the reflected light field, and very fast prefiltering algorithms based on spherical harmonic transforms. We have integrated the three contributions into a complete frequency-space pipeline, as per figure 5.2. However, it is also easy to convert between SHRMs and previous explicit representations. Therefore, the contributions of this paper are relatively independent, and can also be incorporated separately.

There are several interesting similarities and differences between the SHRM and surface light field representations. In fact, the SHRM can be seen as a surface light field on a sphere. The main advantage of the SHRM is that it is independent of geometry, i.e. it can be mapped on to any object geometry using the surface normal.

On the other hand, surface light fields have a number of advantages over SHRMs. They capture the effects of spatially varying illumination, which includes non-uniform lighting, as well as geometry-specific interreflection and self-shadowing effects.

Another difference between SHRMS and surface light fields is that our method is primarily a synthetic approach, while surface light fields represent collected data. Of course, it is possible to convert a surface light field on a sphere to an SHRM, or to synthesize a surface light field given the geometry, reflectance and illumination.

Since surface light fields are so large, they must be compressed. It has been widely assumed that the optimal compression algorithm for light fields will be based on model-based prediction and coding. SHRMs expose the structure of the reflected light field in terms of the frequency properties of the illumination and BRDF. It is therefore interesting to compare spherical harmonic basis functions to PCA-based compression and factorization methods used for BRDFs [37] and surface light fields [64, 88].

The main advantage of SHRMs is that the theoretical analysis gives insight into the intrinsic complexity of the reflected light field, given the complexities of the illumination and reflection functions. This allows us to directly compute the SHRM with the right order and resolution. Furthermore, our approach is much more efficient than computing a full PCA on the entire reflected light field. This is because brute-force PCA computation requires a dense 4D reflected light field as input, and requires computing a singular-value decomposition, which is an expensive operation. Our approach, in contrast, only computes the terms that are needed and uses very fast prefiltering methods.

On the other hand, our method uses a predetermined basis—the spherical harmonics, while PCA-based methods have the advantage of finding an optimal basis to represent a data set, assuming no a priori knowledge. As future work, it would be interesting to see if PCA can be run directly on the SHRM in order to compute a compact final representation efficiently, thus getting the best features of both PCA, and our SHRM representation.

Several aspects of this paper are likely to have broader implications in the general context of rendering. SHRMs are related to Sillion et al.'s [79] spherical harmonic representation of outgoing light for radiosity computations, but differ in storing B_{pq} as a function of orientation (α, β), rather than at each vertex of a geometric model. Because we reparameterize by the reflection vector, our representation is much more compact. SHRMs are also similar in spirit to the representation of Malzbender et al. [53]. Like us, they have a 4D function with rapid variation over two dimensions (for them, texture), and slow variation over the other two (for them, illumination). Since spherical harmonics are polynomials of the cartesian components, their (quadratic) polynomial texture maps can be seen as a subset of our representation.

One drawback of synthetic IBR is the long time required for precomputation, which precludes dynamic lighting or interactive manipulation of material properties. Our new prefiltering method takes an important step in addressing this problem for environment maps. However, we do not currently take into account near-field lighting, interreflection or self-shadowing. In the future, we wish to address these limitations, adapting our prefiltering method to rapidly compute synthetic surface light fields.

For the special case of radially symmetric BRDFs, Kautz et al. [40] have proposed using hardware-assisted 2D image convolution. However, while BRDFs are shift-invariant filters on the spherical domain, they are not shift-invariant in the plane, since projection on to a 2D image introduces distortion [40], and may lead to inconsistencies—for instance, rotating the lighting may not correspond simply to rotating the prefiltered image. Our prefiltering algorithm can be viewed as spherical image processing on the incident illumination, convolving it with the BRDF filter. Our speedups are not surprising, given that planar image convolutions are often more efficiently computed in the Fourier domain. Other approaches to speed up prefiltering are hierarchical methods [40] and spherical wavelets [78]. However, *there is no wavelet or hierarchical convolution formula*, so frequency domain methods are more appropriate for environment mapping. and will give asymptotically better results according to our computational complexity analysis. Further, hierarchical methods can be seen as effectively determining angular resolutions appropriately for a given error tolerance, which we are able to formally analyze.

There has been considerable recent interest in real-time rendering with complex illumination and material models. In work simultaneous with ours, Latta and Kolb [47] have used homomorphic factorization to represent the reflected light field as aa product of textures. As pointed out in the text, the SHRM may be seen as a sum of a number of terms, each being the product of factors depending on the reflected and viewing directions. Having a small number of terms, determined by error analysis, allows greater accuracy. A further significant advantage of the SHRM is that it can be computed extremely efficiently. On the other hand, Latta and Kolb's approach [47] of factoring after computing a dense reflected light field representation has the advantage of allowing for a number of different factorizations, such as using the surface normal instead of the viewing direction, which may be more suitable for diffuse BRDFs. Also, simultaneous with our work, Sloan et al. [66] have restricted themselves to low-frequency lighting, but shown how complex light transport such as shadows and interreflection may be precomputed and then rendered in real-time.

The SHRM differs in that we consider natural illumination, which has significant high-frequency content, but assume no shadowing or interreflection.

In the future, we would like to combine and extend the functionality of these different approaches, so we may handle general near-field illumination, spatially varying realistic materials, and physically-correct light transport including shadows and interreflection. We would also like to adapt the methods of this paper, such as the sampling rate analysis and the SHRM representation, to other image-based rendering problems in graphics.

In summary, natural illumination and accurate BRDFs are of growing importance in interactive applications, and this chapter has presented a complete frequency-space pipeline to enable this.

In the these last two chapters, we have seen the application of the signal-processing framework to problems in interactive forward rendering. The next chapter applies the framework to inverse rendering—estimating illumination and material properties from images.

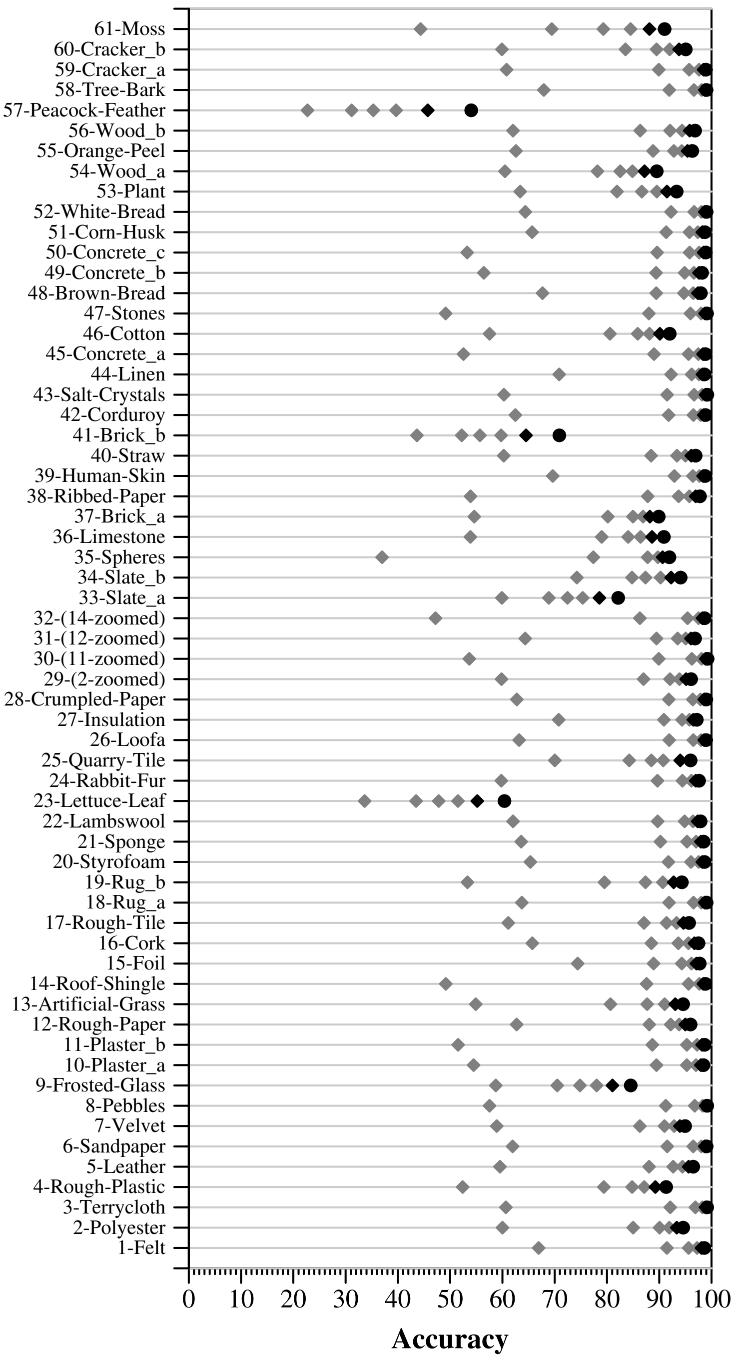

Figure 5.7: *Accuracy of a spherical harmonic BRDF approximation for all 61 BRDFs in the CURET database. We show 6 values of $P_{\hat{\rho}}$ ranging from 0 to 5 from left to right. The low orders for $P_{\hat{\rho}}$ are shown with light gray diamonds, while a black circle shows the highest order $P_{\hat{\rho}} = 5$. Note that the rightmost circle corresponds to an accuracy greater than* 90% *in 56 of the 61 rows.*

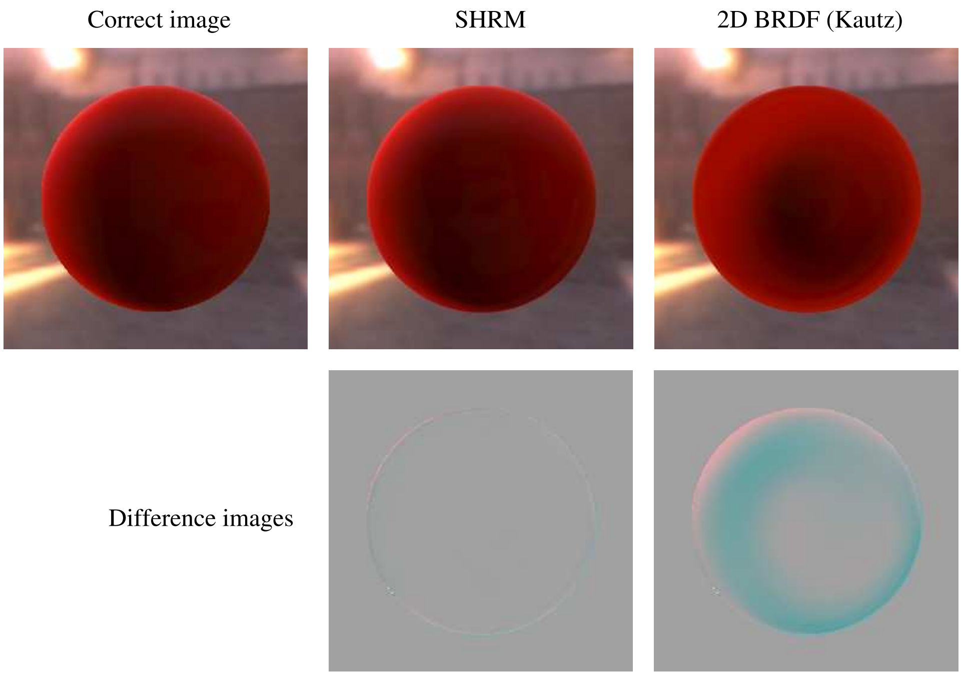

Figure 5.8: *Comparing the correct image on the left to those created using SHRMs (middle) and the 2D BRDF approximation of Kautz and McCool (right).*

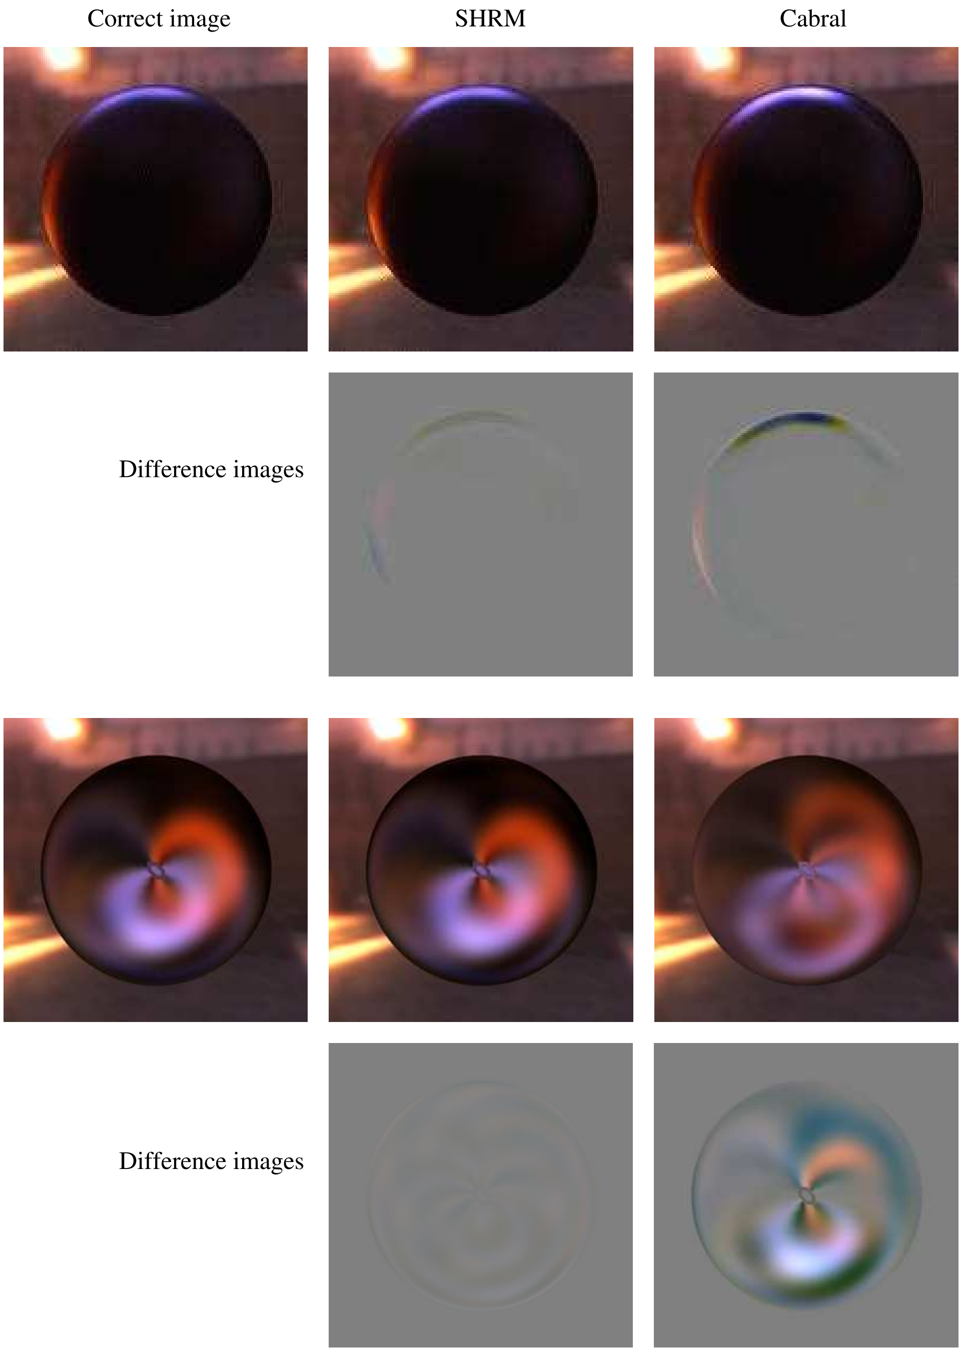

Figure 5.9: *Comparing the correct image to those created using SHRMs and icosahedral interpolation (Cabral's method). We see that the SHRM image is accurate, while Cabral's approximation is inadequate for sharp near-grazing reflections (top), and for complex BRDFs like the anisotropic Kajiya-Kay model (bottom).*

Chapter 6

Inverse Rendering Under Complex Illumination

The previous two chapters have applied our signal processing framework to efficient rendering with environment maps. In this chapter, we switch from forward to inverse rendering, addressing estimation of lighting and material properties from real photographs.

Accurate modeling of the visual world requires accurate models for object geometry and appearance. There has been a significant body of computer vision research over the last two decades on determining shape from observations. However, much less attention has been paid to determining illumination and reflectance properties, even though the perception of materials may be considered as important as the perception of shape in visual modeling. Similarly, until recently, computer graphics has focussed on geometric modeling and the development of accurate physically-based light transport algorithms. However, creation of realistic computer-generated images also requires accurate input models of illumination and reflective properties (BRDFs and textures) of surfaces. In the past, the illumination and reflective properties have usually been set in an ad-hoc manner. This is often now the limiting factor in the realism of synthetic imagery, and is one reason for the growth of *image-based rendering* techniques. In its simplest form, image-based rendering uses view interpolation to construct new images from acquired images without constructing a conventional scene model.

The quality of view interpolation may be significantly improved if it is coupled with *inverse rendering*. Inverse rendering *measures* rendering attributes—lighting, textures, and BRDF—from photographs. Whether traditional or image-based rendering algorithms are used, rendered images use measurements from real objects, and therefore appear very similar to real scenes. Measuring scene attributes also introduces structure into the raw imagery, making it easier to manipulate the scene in intuitive ways. For example, an artist can change independently the material properties or the lighting.

In the last few years, there has been a significant amount of research in inverse rendering. The methods of Debevec et al. [15], Marschner et al. [55], Sato et al. [77], and others have produced high quality measurements, leading to the creation of very realistic images. However, most previous work has been conducted in highly controlled lighting conditions, usually by careful active positioning of a single point source. Even methods that work in outdoor conditions, such as those of Yu and Malik [90], Sato and Ikeuchi [76] and Love [50], are designed specifically for natural illumination, and assume a simple parametric model for skylight.

The usefulness of inverse rendering would be greatly enhanced if it could be applied under general uncontrolled, and possibly unknown, lighting. For instance, this would allow for application in general unconstrained indoor or outdoor settings, or for estimation of BRDFs under unknown illumination. There are also a number of applications to human vision and perception. For instance, Dror et al. [16] have studied reflectance classification from a single image of a sphere under complex illumination to clarify how well the human visual system perceives materials, and to develop computational vision methods for the same task.

One reason there has previously been relatively little work in considering complex illumination is the lack of a common theoretical framework for determining under what conditions inverse problems can and cannot be solved, and for making principled approximations. Recently, we [72, 73] have developed a signal-processing framework for reflection on a curved surface, whereby the reflected light field can be viewed as a spherical convolution of the incident illumination and the BRDF. This framework can be used to determine the well-posedness of inverse problems, i.e. analyze which inverse problems can be solved, and to make appropriate approximations.

In this chapter, we first develop a taxonomy of inverse problems under complex illumination, indicating which problems have been addressed, which will be solved in this article, and which remain subjects for future investigations. We will then use the insights from our previous theoretical analysis to derive new representations and algorithms for inverse rendering under complex illumination. Our contributions include a new dual angular and frequency-space low parameter representation, and new methods for estimating BRDFs, illumination, and *factoring* the reflected light field to simultaneously determine the lighting and BRDF.

This chapter is an expanded version of the practical aspects of our earlier SIGGRAPH paper [73]. That earlier work derived the theoretical signal-processing framework and briefly described applications to inverse rendering. Here, we assume the theoretical framework and focus in more detail on the practical algorithms for inverse rendering under complex illumination.

The rest of this chapter is organized as follows. In section 6.1, we develop a taxonomy of inverse problems, classifying previous work, and indicating unsolved problems that are the subject of this paper. In section 6.2, we give some background on our assumptions and the signal-processing framework we apply here, introducing the main practical implications. In section 6.3, we use these ideas to derive a new low-parameter dual angular and frequency space representation applied in the next section. Section 6.4 presents our new practical algorithms, and illustrates the concepts using spheres of different materials. Section 6.5 presents our results using complex geometric objects, demonstrating improved and more general methods for inverse rendering under complex illumination. Finally, section 6.6 concludes the chapter and suggests directions for future work.

6.1 Taxonomy of Inverse problems and Previous Work

We introduce a taxonomy of inverse problems and algorithms based on a number of factors. To motivate the taxonomy, we first write a simplified version of the reflection equation, omitting visibility. This is the standard reflection equation (c.f. equation 2.8), except that we have added a single spatially varying texture $T(\vec{X})$ to modulate the reflectance,

$$B(\vec{X}, \theta'_o, \phi'_o) = \int_{\Omega'_i} T(\vec{X}) L(\theta_i, \phi_i) \rho(\theta'_i, \phi'_i, \theta'_o, \phi'_o) \cos\theta'_i \, d\omega'_i. \tag{6.1}$$

In practice, we would use separate textures for the diffuse and specular components of the BRDF, and more generally, the BRDF could be spatially varying, i.e. $\rho(\vec{X}, \theta'_i, \phi'_i, \theta'_o, \phi'_o)$.

The integrand is a product of terms—the texture $T(\vec{X})$, the BRDF $\rho(\theta'_i, \phi'_i, \theta'_o, \phi'_o)$, and the lighting $L(\vec{X}, \theta'_i, \phi'_i)$. Inverse rendering, assuming known geometry, involves inverting the integral in equation 6.1 to recover one or more of ρ, L, or T. If two or more quantities are unknown, inverse rendering involves *factoring* the reflected light field. There are a number of axes along which we can classify inverse problems and solutions, as described below.

Unknown quantities—Lighting/BRDF/Texture: We may classify inverse problems depending on how many of the three quantities—lighting, BRDF and texture—are unknown. Considering all possible combinations, this gives rise to a total of seven problem classes.

BRDF representation—low parameter/factored/full measured: Next, we may consider the assumptions made about the form of the illumination and BRDF. Since we are considering complex illumination, we will assume the lighting to be represented as a 2D distant illumination field or environment map, although we will also discuss previous work that makes the assumption of point light sources only. The more interesting axis for us will be the assumptions made for the BRDF. We may assume the BRDF to be a parametric low-parameter representation such as the Torrance-Sparrow [84] model. Alternatively, we may use or estimate a full measured BRDF. In between these two alternatives is the largely unexplored area of lower-dimensional factored representations [37, 57].

Acquired image dataset—2D/3D/4D: In this chapter, we will assume static scenes, with multiple images acquired simply by using a new viewing direction. The image datasets we use could be 2D—corresponding to a single image or a small number of images, 3D—corresponding to a 1D sequence of 2D images, or 4D—corresponding to a dense sampling of the entire reflected light field.

Algorithm and Representation—angular/frequency/dual: Finally, we may classify the solution methods and representations as working in the spatial or angular domain (as with most previous work), working purely in the frequency domain, or using a combination of angular and frequency domains. We develop new algorithms that use either the frequency domain or a combination of angular and frequency domains.

6.1.1 Previous Work on Inverse Rendering

We now proceed to classify previous work according to the taxonomy above. To fully describe the taxonomy above in terms of each of the four categories, we would need a four-dimensional representation. Instead, we will organize the work according to the unknown quantities. For each of the seven classes, we will discuss the remaining axes of the taxonomy. We will then consider unsolved problems, some of which will be addressed in this chapter.

1. Unknown Texture: We first consider the case when we seek to estimate only the texture, with the lighting and BRDF known. Previous methods have recovered the diffuse texture on a surface using a single point light source by dividing by the irradiance in order to estimate the albedo at each point. Details are given by Marschner [81] and Levoy et al. [49]. More complex methods that also make an estimate of the specular component of the BRDF will be covered in other categories. Since the texture is simply an albedo map, it is easy to calculate given the lighting, BRDF, and image data. Therefore, this problem could be solved simply with any BRDF representation. A single image (2D slice of reflected light field) suffices in principle, though better estimates may be obtained with more images. No particular benefit has so far been demonstrated of considering this problem in the frequency domain, except that irradiance calculations may be more efficient.

2. Unknown BRDF: We now consider the case when the lighting and texture are known, and we seek to estimate the BRDF. Essentially all work in this category has assumed homogeneous untextured surfaces, since it is difficult to independently determine the texture. The BRDF [62] is a fundamental intrinsic surface property. Active measurement methods, known as gonioreflectometry, involving a single point source and a single observation at a time, have been developed. Improvements are suggested by Ward [85] and Karner et al. [36]. More recently, image-based BRDF measurement methods have been proposed by Lu et al. [51] and Marschner et al. [55]. These methods work with a single point light source, and estimate a full BRDF representation. Therefore, they use a large number of input images. A 3D dataset or 1D sequence of images is required for an isotropic BRDF. A 4D dataset (or 2D sequence of 2D images) would be required for anisotropic BRDFs. While some such measurements have been made by Dana et al. [14], as part of the BRDF data in the CURET database, this data is still very sparse (only 205 measurements for each sample), and there is room for future work on dense BRDF measurements of anisotropic materials.

An alternative representation is by low-parameter models such as those of Ward [85] or Torrance and Sparrow [84]. The *parametric* BRDF will generally not be as accurate as a full *measured* BRDF. However, parametric models are often preferred in practice since they are compact, and are simpler to estimate. Often, a small number of images suffices (2D data), and even a single image may be used. There has been some previous work on determining parametric BRDFs under nontrivial lighting conditions. Love [50] estimates parametric BRDFs under natural illumination, assuming a low-parameter model for skylight and sunlight. Dror et al. [16] use a single image of a homogeneous sphere to classify the surface reflectance as one of a small number of predetermined BRDFs, making use of assumed statistical characteristics of natural lighting.

The *inverse BRDF* problem has not been solved for general illumination. Within this context, there are a number of open questions, including estimation of low parameter, factored and full measured representations. It is not obvious how much data (2D/3D/4D) one needs for each of these tasks or what the best ways of solving the problem (angular vs frequency domain) are. Some of these problems are addressed in this chapter.

3. Unknown Lighting: A common solution is to use a mirrored ball, as done by Miller and Hoffman [59]. Marschner and Greenberg [54] find the lighting from a Lambertian surface. D'Zmura [17] proposes, but does not demonstrate, estimating spherical harmonic coefficients. For Lambertian objects, we [72] have shown how to recover the first 9 spherical harmonics. Since we're assuming here that the lighting is distant and can be described by a 2D environment map, a small number of images suffices with any BRDF representation. However, a single image is usually inadequate because of sampling and conditioning problems. Previous work has not estimated the lighting from curved surfaces with general parametric or measured BRDFs. We will address this question here, and demonstrate the benefits of frequency domain and dual angular/frequency space algorithms.

4. Factorization—Unknown Lighting and BRDF: For the special case when the lighting consists of a single source of unknown direction, BRDF estimation methods have been proposed by Ikeuchi and Sato [31] and Tominaga and Tanaka [83]. Sato et al. [75] use shadows on a plane to estimate the illumination distribution and the surface reflectance properties.

However, this problem remains unsolved for complex lighting distributions and curved surfaces. There are a number of issues to be addressed, including both parametric and measured BRDF models. As for BRDF estimation, it is not obvious how much data (2D/3D/4D) is required for each of these cases, nor what the best algorithms (angular/freqeuency) are.

5. Factorization—Unknown Texture and BRDF: This corresponds to recovering *textured*, or spatially-varying BRDFs. For estimation of textured parametric BRDFs, a small number of input images suffices, though using more images gives greater accuracy, and allows for observation of specular highlights over most of the surface, in at least one of the input images. Kay and Caelli [41] use a few images, taken under point sources at different locations, to estimate a simplified Torrance-Sparrow model for each pixel. Sato et al. [77] rotate an object on a turntable, using a single point source, to recover BRDF parameters and texture. Yu et al. [89] recover a texture only for the diffuse BRDF component, but account for interreflections. Sato and Ikeuchi [76] and Yu and Malik [90] recover BRDFs and diffuse textures under natural illumination, assuming a simple parametric model for skylight, and using a sequence of images acquired under different illumination conditions.

Most of the above methods recover only diffuse textures; constant values, or relatively low-resolution textures, are used for the specular parameters. If more detailed models are sought, a small number of images is no longer sufficient and 4D or larger image datasets are required. Using a large number of images obtained by moving a point source around a sphere surrounding the subject, Debevec et al. [15] acquire the reflectance field of a human face, and recover parameters of a microfacet BRDF model for each surface location. Dana et al. [14] generalize BRDFs to a 6D bi-directional texture function (BTF).

6. Factorization—Unknown Lighting and Texture: We can also try to simultaneously determine the lighting and texture, given a known (in the simplest case, Lambertian) BRDF. This corresponds to texture estimation under unknown lighting. We have shown [72] that a distant illumination field can cause only low frequency —with respect to curvature— variation in the radiosity of a convex Lambertian surface. This implies that, for a diffuse object, high-frequency texture can be recovered independently of lighting. These observations are in agreement with the perception literature, such as Land's retinex theory [46], wherein high-frequency variation is usually attributed to texture, and low-frequency variation associated with illumination. However, note that there is a fundamental ambiguity between low-frequency texture and lighting effects. Therefore, lighting and texture cannot be factored without using active methods or making further assumptions regarding their expected statistical characteristics.

For non-Lambertian BRDFs, it would be possible in principle to separate the diffuse and specular components of the reflected light, based on the change of specular intensity for different viewing directions. This could then be used to determine the irradiance and hence, the diffuse texture. We have not in practice found this to be a viable solution technique because the effects are subtle and the assumed reflection models are not exact. Some recent results along these lines are reported by Nishino et al. [64], but they also appear to have difficulty obtaining accurate results.

7. Factorization—Unknown Lighting, Texture, BRDF: Ultimately, we wish to recover textured BRDFs under unknown lighting. We cannot solve this problem without further assumptions, because we must first resolve the lighting-texture ambiguity.

Our approach differs from much of the previous work in that it is derived from a mathematical theory of inverse rendering. As such, it has similarities to inverse methods used in areas of radiative transfer and transport theory such as hydrologic optics [67] and neutron scattering. See McCormick [58] for a review. Our results are based on recent theoretical work, where we have formalized the notion of reflection on a curved surface as a spherical convolution for flatland light fields [70], Lambertian surfaces [72], and finally for general isotropic BRDFs [73] . For the Lambertian case, similar results have been derived independently by Basri and Jacobs [2].

6.1.2 Open Problems

Based on the taxonomy introduced in the previous subsection, we may identify a number of open problems, some of which will be addressed in this chapter, and some of which identify directions for future research. In this subsection, we identify some important classes of open problems, discuss our contributions, and the questions that remain unanswered. In the next subsection, we will give an overview of our new algorithms and important future directions of work.

(Textured) BRDF estimation under complex illumination: The *inverse-BRDF* problem remains largely unexplored for general complex illumination, as does estimation of textured BRDFs, although considerable progress has been made for specific models of skylight illumination [50, 76, 90]. We address this question for parametric BRDF models, using a small number of views. An important future direction is estimating factored or measured BRDF representations.

Factorization of Lighting and BRDF: Simultaneous determination of BRDFs and lighting under complex uncontrolled illumination for complex geometric surfaces has not been fully addressed. One of the main practical contributions of this chapter is one solution to this problem for curved surfaces, allowing us to estimate parametric BRDFs under general unknown illumination, while also determining the lighting, from a small number of input photographs. Estimation of higher-dimensional or measured BRDFs remains an open problem.

Factorization of lighting/BRDF/texture: We have already discussed how factoring lighting and texture (and consequently determining lighting, BRDFs and texture simultaneously) is an ill-posed problem. However, it is possible to make statistical assumptions or assume one or more quantities may be represented by a low-parameter function. A full exploration of these ideas is a subject of future work.

Single-view estimation of lighting and BRDF: While a number of inverse rendering problems can be solved using a 2D slice of the reflected light field, there has been relatively less work on single view modeling of reflectance and lighting properties. Boivin and Gagalowicz [6] take a first step, including interreflections, with known lighting from a small number of point sources. There has also recently been considerable interest in single-view geometric modeling [65, 92], and single-view reflectance modeling is a natural extension.

Frequency domain and hybrid angular and frequency space algorithms: Most previous work has used spatial or angular domain method. By using methods based on our signal-processing framework [73], we can develop new frequency-space and hybrid angular and frequency-domain representations and methods. We demonstrate the improved quality of these new approaches.

Our practical contributions include five algorithms for BRDF and lighting estimation under complex illumination. We present two types of methods—algorithms that recover coefficients of a purely frequency-space description of the lighting or BRDF by representing these quantities as a sum of spherical harmonic terms, and algorithms that estimate parameters corresponding to a new dual angular and frequency-space representation introduced later in this chapter.

It should be noted that a number of the open problems discussed above remain unanswered in this chapter, and are directions for future investigation. Specifically, estimation of factored and full measured BRDF representations under complex illumination remains an open problem, under both known and unknown lighting. Single view estimation and modeling of reflectance and illumination also remains an open problem, as does the use of statistical and other assumptions for a full factorization into lighting, BRDF and texture.

6.2 Preliminaries

The input to our algorithms consists of object geometry (acquired using a laser range scanner and a volumetric merging algorithm [13]) and photographs from a number of different viewing directions, with known extrinsic and intrinsic camera parameters. We assume static scenes, i.e. that the object remains stationary and the lighting remains the same between views. Our method is a passive-vision approach; we do not actively disturb the environment. In this chapter, we will also assume the illumination comes from distant sources, as we have been doing throughout this dissertation, and is a function only of the global incident direction, which can be represented with an environment map. For simplicity, we will restrict ourselves to isotropic BRDFs and neglect the effects of interreflection. Our theoretical analysis also discounts self-shadowing for concave surfaces, although our practical algorithms will account for it where necessary. Our assumptions (known geometry, distant illumination, isotropic BRDFs and no interreflections) are commonly made in computer vision and interactive computer graphics.

In the theoretical part of this dissertation, in chapters 2 and 3, we have developed a signal-processing framework for reflection based on the assumptions outlined above, ignoring concavities and self-shadowing. Thus, the reflected light field can be expressed as a spherical convolution of the incident illumination and the BRDF, and expressed as a product of spherical harmonic coefficients of the lighting and BRDF. This allows us to view inverse rendering as *deconvolution*, or as a *factorization* of the reflected light field into the lighting and BRDF. Our analysis also allows us to formally determine which inverse problems are ill-conditioned or ill-posed versus well-conditioned and well-posed. In particular, we may view the incident lighting as a signal and the BRDF as a filter, so inverse problems are ill-posed when certain modes in the signal or filter vanish. For instance, we may formally say that determining the surface roughness on a cloudy day is ill-posed since the incident illumination does not include high frequencies, and the high frequencies of the BRDF cannot therefore be estimated.

Our theory leads to several new insights by reformulating reflection in the frequency domain. However, the frequency-space ideas must be put into practice carefully. This is analogous to practical implementation of the Fourier-space theory of aliasing. The ideal Fourier-space bandpass filter in the spatial domain, the sinc function, is usually modified for practical purposes because it has infinite extent and leads to ringing. Similarly, representing BRDFs purely as a linear combination of spherical harmonics leads to ringing. Moreover, it is difficult to compute Fourier spectra from sparse irregularly sampled data. Similarly, it is difficult to compute the reflected light field coefficients from a few photographs; we would require a very large number of input images, densely sampling the entire sphere of possible directions.

Here, we consider the implications of the theoretical analysis for practical inverse rendering algorithms. We first briefly discuss a number of practical implications of the theory. We then use these ideas to derive a simple practical model of the reflected light field for the microfacet BRDF. A similar form can be derived for other common BRDFs like the Phong reflection model. This representation will be used extensively in the practical algorithms of section 6.4.

6.2.1 Practical implications of theory

We now discuss a number of ideas and quantitative results obtained from the theory that influence our practical representations.

Dual Angular and Frequency-Space Representations: Quantities local in angular space have broad frequency spectra and vice-versa. By developing a frequency-space view of reflection, we ensure that we can use either the angular-space or frequency-space representation, or even a combination of the two. The diffuse BRDF component is slowly varying in angular-space, but is local in frequency-space, while the specular BRDF component is local in the angular domain. For representing the lighting, the frequency-space view is appropriate for the diffuse BRDF component, while the angular-space view is appropriate for the specular component.

Irradiance formula: For the Lambertian BRDF component, we have derived [72] a simple analytic formula, and have shown that the irradiance at all surface orientations can be approximated to within 1% using only 9 parameters, i.e. coefficients of spherical harmonics up to order 2. Thus, it makes sense to apply this simple formula where possible, representing the diffuse component of the reflected light field in the frequency domain.

Associativity of convolution: Because the coefficients of the reflected light field in the frequency domain are simply a product of the spherical harmonic coefficients of the incident illumination and the BRDF, we may apply the associativity of convolution. Thus, we can blur the illumination and sharpen the BRDF without changing the final results. In the extreme case, for specular models like the Phong BRDF, we may treat the BRDF as a mirror, while blurring the illumination, convolving it with the BRDF filter. Within the context of environment map rendering [22, 59], this is known as prefiltering. Besides increased efficiency, this approach also allows for very efficient approximate computation of shadows. One need simply check the reflected ray, as if the surface were a mirror, which is a simple operation in a raytracer.

Separation of slow and fast-varying lighting: In general, because the lighting and BRDF are not one-dimensional quantities, applying the associativity property above destroys the symmetries and reciprocity of the BRDF, so we cannot simply blur the illumination and treat the BRDF as a perfect mirror. However, for radially symmetric specular BRDFs, like the Phong model, where the BRDF depends only on the angle between the incident illumination and the reflection of the viewing direction about the surface normal, this is a valid operation. Therefore, we separate the illumination into slow and fast-varying components, corresponding to area sources and point sources. It can be shown that for low-frequency lighting, models like the microfacet BRDF (Torrance-Sparrow [84] model) behave much like a Phong model (the dominant term is Phong-like reflection), so that we may blur the illumination and treat the BRDF as a mirror. Furthermore, the largest errors in this approximation occur for grazing angles, where measurements are accorded low confidence in practical applications anyway. The fast-varying lighting components may be treated as point sources, which makes it easy to find angular-space formulae for the reflected light field.

It should be noted that the theoretical analysis is conducted without taking concavities into account. We will derive our representation in the next section under the convex-surface assumption. However, we will also show there how the representation can be simply extended to account for textured objects and cast shadows.

6.3 Dual angular and frequency-space representation

In a sense, the practical implications discussed above simply formalize a reflection model commonly used when rendering with environment maps. In that context, the BRDF is assumed to be a combination of Lambertian diffuse and Phong specular reflection. The reflected light is then the combination of a diffuse irradiance map due to the Lambertian BRDF component and a specular reflection map due to the specular Phong lobe. Our theoretical analysis allows for two practical improvements to be made. Firstly, the irradiance map can be represented using only 9 parameters in the frequency domain, which makes computations more efficient and compact. Secondly, we may use a single angular-space reflection map as a good approximation for the specular reflections, even for more complex physically-based BRDFs like the microfacet model [84], provided we first separate the lighting into slow and fast-varying components.

In this chapter, we will use a simplified Torrance-Sparrow [84] model, defined as follows. Please note that to simplify the notation, we use $\vec{\omega}_i'$ and $\vec{\omega}_o'$ to denote the (unit vector) incident (θ_i', ϕ_i') and outgoing (θ_o', ϕ_o') directions,

$$\begin{aligned}
\rho(\vec{\omega}_i', \vec{\omega}_o') &= K_d + K_s \frac{FS}{4\cos\theta_i' \, \cos\theta_o'} \\
\vec{\omega}_h' &= \frac{\vec{\omega}_i' + \vec{\omega}_o'}{\| \vec{\omega}_i' + \vec{\omega}_o' \|} \\
F &= \frac{F(\mu, \theta_o)}{F(\mu, 0)} \\
S &= \frac{1}{\pi\sigma^2} \exp\left[-\left(\theta_h'/\sigma\right)^2 \right].
\end{aligned} \tag{6.2}$$

Here, ρ is the BRDF, and σ is the surface roughness parameter. The subscript h stands for the half-way vector. $F(\mu, \theta_o')$ is the Fresnel term for refractive index μ; we normalize it to be 1 at normal exitance. Actually, F depends on the angle with respect to the half-way vector; in practice, this angle is usually very close to θ'_o. For simplicity in the analysis, we have omitted the geometric attenuation factor G. In practice, this omission is not very significant except for observations made at grazing angles, which are usually assigned low confidence anyway in practical applications.

6.3.1 Model for reflected light field

Our model for the reflected light from the microfacet BRDF now includes three terms,

$$B = B_d + B_{s,slow} + B_{s,fast}. \tag{6.3}$$

Here, B is the net reflected light field. The component because of the diffuse part in the BRDF is denoted B_d. $B_{s,slow}$ represents specularities from the slowly-varying lighting, and $B_{s,fast}$ specular highlights from the fast varying lighting component.

We may represent and compute B_d in the frequency domain by using the irradiance formula (which corresponds directly to the reflection from a Lambertian surface). We use the 9 parameter representation, explicitly noting the frequency $l \leq 2$,

$$\begin{aligned} B_d &= K_d E(\alpha, \beta) \\ E(\alpha, \beta) &= \sum_{l=0}^{2} \left(\hat{\rho}_l \sum_{m=-l}^{+l} L_{lm} Y_{lm}(\alpha, \beta) \right). \end{aligned} \tag{6.4}$$

Here, E is the irradiance, and K_d is the albedo or coefficient for diffuse reflection. The surface is parameterized by its orientation or surface normal in spherical coordinates (α, β). The spherical harmonics are denoted by Y_{lm}, and the spherical harmonic coefficients of the lighting by L_{lm}. The numerical values of $\hat{\rho}_l$ are given by

$$\hat{\rho}_0 = \pi \qquad \hat{\rho}_1 = 2\pi/3 \qquad \hat{\rho}_2 = \pi/4. \tag{6.5}$$

For $B_{s,slow}$, we filter the lighting, and treat the BRDF as a mirror. With $\vec{R}$ denoting the reflected direction, and L_{slow} the filtered version of the lighting, we obtain

$$B_{s,slow} = K_s F(\mu, \theta'_o) L_{slow}(\vec{R}). \tag{6.6}$$

The filtered version of the illumination L_{slow} is obtained by multiplying the illumination coefficients by those of a filter corresponding to the term S in the microfacet BRDF of equation 6.2, i.e.

$$L_{lm}^{slow} = \exp\left[-\left(\sigma l\right)^2\right] L_{lm}. \tag{6.7}$$

In the angular domain, this corresponds to convolving with a filter of angular width approximately σ^{-1}, or using a normalized Phong lobe with Phong exponent $\frac{1}{2}\sigma^{-2}$.

For the *fast varying* portion of the lighting—corresponding to sources of angular width $\ll \sigma$—we treat the total energy of the source, given by an integral over the (small) solid angle subtended, as located at its center, so the lighting is a delta function. $B_{s,fast}$ is given by the standard equation for the specular highlight from a directional source. The extra factor of $4\cos\theta'_o$ in the denominator as compared to equation 6.6 comes from the relation between differential microfacet and global solid angles,

$$\begin{aligned} B_{s,fast} &= \frac{K_s F(\mu, \theta'_o)}{4\cos\theta'_o} \sum_j T_j \\ T_j &= \exp\left[-\left(\theta'_h/\sigma\right)^2\right] \left(\frac{L_{j,fast}}{\pi\sigma^2}\right). \end{aligned} \tag{6.8}$$

The subscript j denotes a particular directional source; there could be several. Note that $L_{j,fast}$ is now the *total energy* of the source.

For BRDF estimation, it is convenient to expand out these equations, making dependence on the BRDF parameters explicit,

$$B = K_d \sum_{l=0}^{2} \left(\hat{\rho}_l \sum_{m=-l}^{+l} L_{lm} Y_{lm}(\alpha, \beta) \right) + K_s F(\mu, \theta'_o) \left[L_{slow}(\vec{R}) + \frac{1}{4\cos\theta'_o} \sum_j T_j(\sigma) \right]. \tag{6.9}$$

6.3.2 Textures and shadowing

We now show how to extend our representation to account for object textures and self-shadowing on complex concave geometry. The representation can be extended to textured surfaces simply by letting the BRDF parameters (such as K_d and K_s) be functions of surface location. It would appear that concave regions, where one part of the surface may shadow another, are a more serious problem since our theory is developed for convex objects and assumes no self-shadowing. However, in the remainder of this section, we will see that the extensions necessary mainly just involve checking for shadowing of the reflected ray and directional sources, which are routine operations in a raytracer.

We consider each of the three terms in our model of the reflected light field. In the presence of shadows, the 9 parameter model can no longer be used to directly compute B_d. Instead, the irradiance may be computed in the more conventional angular-space way by integrating the scene lighting while considering visibility. Alternatively, we can continue to use a spherical harmonic approximation, making use of the linearity of light transport. Note that the irradiance can still be written as a linear combination of lighting coefficients. Thus, we may replace equation 6.4 by

$$\begin{aligned} B_d &= K_d E(\vec{x}) \\ E(\vec{x}) &= \sum_{l=0}^{l_{max}} \sum_{m=-l}^{+l} L_{lm} \tilde{Y}_{lm}(\vec{x}). \end{aligned} \tag{6.10}$$

Here, we have increased the maximum frequency from 2 to l_{max}, where l_{max} can be larger than 2. Further, we have replaced the spherical harmonics with $\tilde{Y}_{lm}$. $\tilde{Y}_{lm}$ is the effect of the illumination spherical harmonic Y_{lm}. Since this effect now depends on the specific shadowing patterns, we have replaced the surface normal (α, β) with the position $\vec{x}$. For convex objects, as per equation 6.4, $\tilde{Y}_{lm}(\vec{x}) = \hat{\rho}_l Y_{lm}(\alpha, \beta)$.

For the specular components of the reflected light field, we simply check if the reflected ray (for the "slow" component) or the point sources (for the "fast" component) are shadowed. The main benefit is for slow specularities, where instead of a complex integration including visibility, the effects of shadowing are approximated simply by checking the reflected ray. It should be emphasized that in all cases, the corrections for visibility depend only on object geometry and viewing configuration (to determine the reflected direction), and can be precomputed for each point on the object using a ray tracer. Thus, we may replace equation 6.6 by

$$B_{s,slow} = K_s F(\mu, \theta'_o) L_{slow}(\vec{R}) V(\vec{R}). \tag{6.11}$$

where V is a binary value specifying if the reflected ray is unshadowed. Similarly, a visibility term needs to multiply T_j in equation 6.8. Putting it all together, and including the effects of textures, by making the diffuse and specular reflectances function of position $\vec{x}$, equation 6.9 becomes

$$B = K_d(\vec{x}) \sum_{l=0}^{l_{max}} \sum_{m=-l}^{+l} L_{lm} \tilde{Y}_{lm}(\vec{x}) + K_s(\vec{x}) F(\mu, \theta'_o) \left[V(\vec{R}) L_{slow}(\vec{R}) + \frac{1}{4 \cos \theta'_o} \sum_j V_j T_j(\sigma) \right]. \tag{6.12}$$

6.4 Algorithms

This section presents our practical algorithms for a broad range of inverse rendering problems under complex illumination, with simple illustrations using spheres of different materials. Our results using more complex geometric objects are presented in the next section. We describe two types of methods—algorithms that recover coefficients of a purely frequency-space description of the lighting or BRDF by representing these quantities as a sum of spherical harmonic terms, and algorithms that estimate parameters corresponding to our dual angular and frequency-space model of section 6.3. Section 6.4.2 on BRDF estimation discusses direct recovery of spherical harmonic BRDF coefficients, as well as estimation of parametric microfacet BRDFs using equations 6.9 and 6.12. Similarly, section 6.4.3 demonstrates direct recovery of spherical harmonic lighting coefficients, as well as estimation of a dual angular and frequency-space lighting description as per the model of section 6.2. Finally, section 6.4.4 shows how to combine BRDF and lighting estimation techniques to simultaneously recover the lighting and BRDF parameters, when both are unknown. In this case, we do not show direct recovery of spherical harmonic coefficients, as we have thus far found this to be impractical.

6.4.1 Data Acquisition

To experimentally test our methods, we first used homogeneous spheres[1] of different materials. Spheres are naturally parameterized with spherical coordinates, and therefore correspond directly to our theory. Later, we also used complex objects—the next section demonstrates results using a white cat sculpture, and a textured wooden doll—to show the generality of our methods.

We used a mechanical gantry to position an inward-looking Toshiba IK-TU40A CCD(x3) camera on an arc of radius 60cm. Calibration of intrinsics was done by the method of Zhang [93]. Since the camera position was computer-controlled, extrinsics were known. The mapping between pixel and radiance values was also calibrated. We acquired 60 images of the target sphere, taken at 3 degree intervals on a great-circle (or equatorial) arc. A schematic and photograph of our setup are in figure 6.1. To map from image pixels to

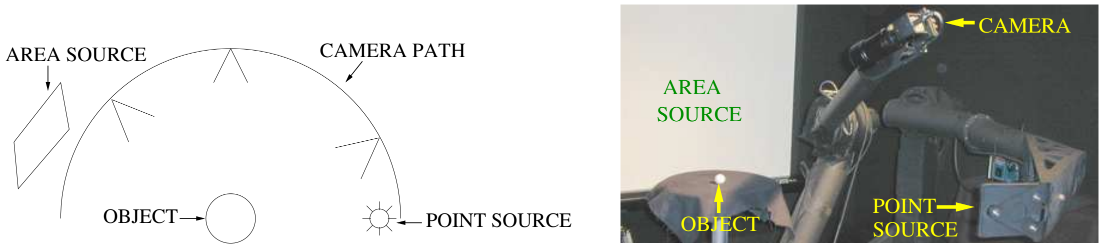

Figure 6.1: **Left:** *Schematic of experimental setup* **Right:** *Photograph*

angular coordinates $(\alpha, \beta, \phi'_o, \phi'_o)$, we used image silhouettes to find the geometric location of the center of the sphere and its radius.

Our gantry also positioned a 150W white fiberoptic point source along an arc. Since this arc radius (90 cm) was much larger than the sphere radii (between 1.25 and 2cm), we treated the point source as a directional light. A large area source, with 99% of its energy in low-frequency modes of order $l \leq 6$, was obtained by projecting white light on a projection screen. The lighting distribution was determined using a mirror sphere. This information was used directly for experiments assuming known illumination, and as a reference solution for experiments assuming unknown illumination.

We also used the same experimental setup, but with only the point source, to measure

[1] Ordered from the McMaster-Carr catalog http://www.mcmaster.com

the BRDF of a white teflon sphere using the image-based method of Marschner et al. [55]. This independent measurement was used to verify the accuracy of our BRDF estimation algorithms under complex illumination.

6.4.2 Inverse BRDF with known lighting

Estimation of Spherical Harmonic BRDF coefficients: Spherical harmonics and Zernike polynomials have been fit [43] to measured BRDF data, but previous work has not tried to estimate coefficients directly. Since the BRDF is linear in the coefficients $\hat{\rho}_{lpq}$, we simply solve a linear system to determine $\hat{\rho}_{lpq}$, to minimize the RMS error with respect to image observations[2]. It should be noted that in so doing, we are effectively interpolating (and extrapolating) the reflected light field to the entire 4D space, from a limited number of images.

Figure 6.2 compares the parametric BRDFs estimated under complex lighting to BRDFs measured using a single point source with the method of Marschner et al. [55]. As expected [43], the recovered BRDFs exhibit ringing. One way to reduce ringing is to attenuate high-frequency coefficients. According to our theory, this is equivalent to using low frequency lighting. Therefore, as seen in figure 6.2, images rendered with low-frequency lighting do not exhibit ringing and closely match real photographs, since only the low-frequency components of the BRDF are important. However, images rendered using directional sources show significant ringing.

For practical applications, it is usually more convenient to recover low-parameter BRDF models since these are compact, can be estimated from relatively fewer observations, and do not exhibit ringing. In the rest of this section, we will derive improved inverse rendering algorithms, assuming a parametric microfacet BRDF model.

Estimation of parametric BRDF model: We estimate BRDF parameters under general known lighting distributions using equation 6.9. The inputs are images that sample the reflected light field B. We perform the estimation using nested procedures. In the outer procedure, a simplex algorithm adjusts the nonlinear parameters μ and σ to minimize RMS

[2]Since the number of image pixels in a number of views can be very large, we randomly subsample the data for computational simplicity. We have used 12000 randomly selected image pixels.

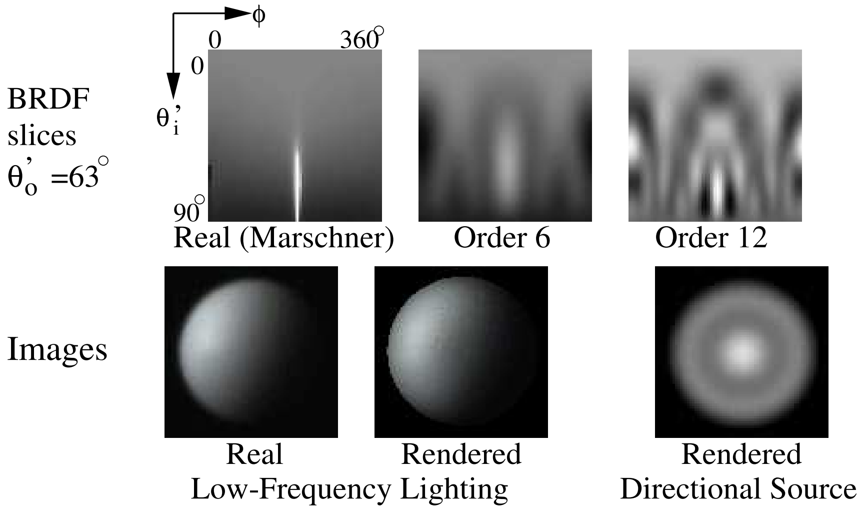

Figure 6.2: *Direct recovery of BRDF coefficients.* **Top:** *Slices of the BRDF transfer function of a teflon sphere for fixed exitant angle of* $63°$. ϕ'_o *varies linearly from* $0°$ *to* $90°$ *from top to bottom, and* $\mid \phi'_o - \phi_i \mid$ *linearly from* $0°$ *to* $360°$ *from left to right. The central bright feature is the specular highlight. Left is the BRDF slice independently measured using the approach of Marschner et al. [55], middle is the recovered value using a maximum order 6, and right is the recovered version for order 12. Ringing is apparent in both recovered BRDFs. The right version is sharper, but exhibits more pronounced ringing.* **Bottom:** *Left is an actual photograph; the lighting is low-frequency from a large area source. Middle is a rendering using the recovered BRDF for order 6 and the same lighting. Since the lighting is low-frequency, only low-frequency components of the BRDF are important, and the rendering appears very similar to the photograph even though the recovered BRDF does not include frequencies higher than order 6. Right is a rendering with a directional source at the viewpoint, and exhibits ringing.*

error with respect to image pixels. In the inner procedure, a linear problem is solved for K_d and K_s. For numerical work, we use the simplex method `e04ccc` and linear solvers `f01qcc and f01qdc` in the NAG [25] C libraries. The main difference from previous work is that equation 6.9 provides a principled way of accounting for all components of the lighting and BRDF, allowing for the use of general illumination conditions.

We tested our algorithm on the spheres. Since the lighting includes high and low-frequency components (a directional source and an area source), the theory predicts that parameter estimation is well-conditioned. To validate our algorithm, we compared parameters recovered under complex lighting for one of the samples, a white teflon sphere, to those obtained by fitting to the full BRDF separately measured by us using the method of

Marschner et al. [55]. Unlike most previous work on BRDF estimation, we consider the Fresnel term. It should be noted that accurate estimates for the refractive index μ require correct noise-free measurements at grazing angles. Since these measurements tend to be the most error-prone, there will be small errors in the estimated values of μ for some materials. Nevertheless, we find the Fresnel term important for reproducing accurate specular highlights at oblique angles. It should also be noted that while the results are quite accurate, there is still potential for future work on appropriate error metrics, especially for estimation of the roughness σ; a linear RMS error may not always be optimal.

Parameter	**Our Method**	**Fit to Data**
Reflectance	0.86	0.87
$K_d/(K_d + K_s)$	0.89	0.91
$K_s/(K_d + K_s)$	0.11	0.09
μ	1.78	1.85
σ	0.12	0.13
RMS	9.3%	8.5%

Figure 6.3: *Comparison of BRDF parameters recovered by our algorithm under complex lighting to those fit to measurements made by the method of Marschner et al. [55].*

The results in figure 6.3 show that the estimates of BRDF parameters from our method are quite accurate, and there is only a small increase in the error-of-fit when using parameters recovered by our algorithm to fit the measured BRDF. We also determined percentage RMS errors between images rendered using recovered BRDFs and real photographs to be between 5 and 10%. A visual comparison is shown in the first and third rows of figure 6.8. All these results indicate that, as expected theoretically, we can accurately estimate BRDFs even under complex lighting.

Textured objects with complex geometry: Handling concavities in complex geometric objects is not significantly more difficult, since we simply need to take visibility into account, and use equation 6.12 instead of equation 6.9. Equation 6.12 can also be used directly to estimate textured BRDFs. However, there are a number of subtle differences from direct BRDF estimation, which are noted below.

In considering textured surfaces, we essentially wish to consider each point on the surface separately, estimating a BRDF for each point independently from observations of that point alone. However, we now have only a few observations for each point (the number of images used). If there were no image noise, and our simplified four parameter microfacet model were a perfectly accurate description of the surface, this would still be sufficient. However, in practice, we are not able to reliably estimate the nonlinear parameters from such sparse data. This is true even for point source illumination, and has been observed by many authors. In our case, since we have complex illumination, the problem is even harder. Therefore, like much previous work, we assume the nonlinear parameters σ and μ are constant across the surface. A weaker assumption would be to allow them to vary slowly, or break the surface into regions of constant μ and σ.

Therefore, we will solve for the global nonlinear parameters σ and μ, as well as the diffuse and specular textures, $K_d(\vec{x})$ and $K_s(\vec{x})$. The corresponding radiance values for each image observation can be written as

$$B = K_d(\vec{x})D + K_s(\vec{x})S(\mu, \sigma), \tag{6.13}$$

where D and S stand for the diffuse and specular components computed from equation 6.12. These depend only on the lighting and viewing configuration, and S also depends on the nonlinear parameters μ and σ. It should be noted that much previous work has assumed constant values for the specular coefficient. The reason is that specularities are not usually observed over the whole object surface. By using complex illumination, we alleviate this problem somewhat, since large regions of the object can exhibit specularity in a single image. Nevertheless, there might be dimly lit regions or places where no specularities are observed in a sequence of views, and we will not be able to estimate coefficients in these regions. Therefore, we introduce confidence measures to enscapsulate the importance of each observation,

$$\begin{aligned} W_d &= \frac{D\cos\theta'_o}{\epsilon + S} \\ W_s &= S\cos\theta'_o. \end{aligned} \tag{6.14}$$

Here, W_d and W_s are the confidence parameters for diffuse and specular reflection respectively. The multiplication by $\cos\theta'_o$ is to give less weight to observations made at grazing exitant angles. ϵ is a small constant to avoid divisions by 0. In the diffuse weight W_d, we give greater importance to well illuminated pixels (high values of D) without too much specularity. In the specular weight W_s, we give importance to pixels observing strong specular reflections S.

Parameter estimation now proceeds much as BRDF estimation for untextured surfaces. Initially, we solve for values of the nonlinear parameters μ and σ using a simplex algorithm (outer loop). To account for regions where specularity is not strongly observed, in this phase, we include K_s as a global parameter to be solved for. In the inner loop of the procedure, we solve (at each point separately) for $K_d(\vec{x})$ to minimize the RMS error over all views. The output from this stage are parameters K_s, μ and σ, as well as an initial estimate of the diffuse texture $K_d(\vec{x})$. We use these global values of μ and σ. The global estimated value of K_s will be used in regions where a better estimate is not possible, but will in general be refined. In this first pass of the algorithm, we weight each observation using the confidence weight W_d.

We then use an iterative scheme to refine the estimates of K_s and K_d. While we could simply solve a linear system, corresponding to equation 6.13, for each vertex on the object, we have obtained better results using an iterative scheme, alternatively solving for $K_d(\vec{x})$ and $K_s(\vec{x})$ while keeping the other fixed. Since we use the dielectric model, K_s has no color, and we recover 4 linear texture parameters for each pixel (a diffuse RGB color and a specular coefficient). It should be noted that different confidence weights (W_d or W_s) are used in the iteration, depending on whether we are estimating the diffuse or specular component of the texture. We start by using a constant value of K_s, and the corresponding value of $K_d(\vec{x})$ recovered in the first phase, where we solved for μ and σ. We then hold $K_d(\vec{x})$ fixed and solve for $K_s(\vec{x})$. Thereafter, we hold $K_s(\vec{x})$ fixed and solve for $K_d(\vec{x})$, and repeat this process till convergence to the desired tolerance, which usually takes a few iterations.

There can of course be cases where $\sum W_s$ or $\sum W_d$ (the summation is over all views of that point) are too low (numerically zero) to accurately estimate specular or diffuse textures respectively. This corresponds to not observing specularities (when $\sum W_s$ is close to 0), or having the point being so dimly lit that the texture isn't discernible (when $\sum W_d$ is close to 0). In the former case, we simply use the mean value of the specular texture, while in the latter case, we mark the diffuse texture estimate as unreliable. It should be noted that using complex illumination greatly reduces the number of points where this is an issue, since much more of the object receives illumination and exhibits specularities than with a point light source.

6.4.3 Inverse Lighting with Known BRDF

Previous methods for estimating the lighting have been developed only for the special cases of mirror BRDFs (a gazing sphere), Lambertian BRDFs (Marschner and Greenberg [54]), and when shadows are present (Sato et al. [75]). Previous methods [54, 75] have also required regularization using penalty terms with user-specified weights, and have been limited by the computational complexity of their formulations to a coarse discretization of the sphere. We present two new algorithms for curved surfaces with general BRDFs. The first method directly recovers spherical harmonic lighting coefficients L_{lm}. The second algorithm estimates parameters of the dual angular and frequency-space lighting model of section 6.2. This method requires no explicit regularization, and yields high-resolution results that are sharper than those from the first algorithm, but is more difficult to extend to concave surfaces.

The theory tells us that inverse lighting is ill-conditioned for high-frequencies. Therefore, we will recover only low-frequency continuous lighting distributions, and will not explicitly account for directional sources, i.e. we assume that $B_{s,fast} = 0$. The reflected light field is essentially independent of the surface roughness σ under these conditions, so our algorithms do not explicitly use σ. The theory predicts that the recovered illumination will be a filtered version of the real lighting. Directional sources will appear as filtered into continuous distributions of angular width approximately σ.

Estimation of Spherical Harmonic Lighting coefficients: We may represent the lighting entirely in frequency-space by coefficients L_{lm} with $l \leq l^*$, and solve a linear least-squares system for L_{lm}. The first term in parentheses below corresponds to B_d, and the second to $B_{s,slow}$. The cutoff l^* is used for regularization, and should be of order $l^* \sim \sigma^{-1}$. Since most materials have $\sigma \sim .1$, we use $l^* = 12$,

$$B = \sum_{l=0}^{l^*} \sum_{m=-l}^{l} L_{lm} \left(K_d \hat{\rho}_l Y_{lm}(\alpha, \beta) + K_s F Y_{lm}(\theta_R, \phi_R) \right). \tag{6.15}$$

To extend this to concave surfaces, we simply need to add terms corresponding to visibility and shadowing, following equation 6.12, but the problem remains a linear system,

$$B = \sum_{l=0}^{l^*} \sum_{m=-l}^{l} L_{lm} \left(K_d \tilde{Y}_{lm}(\vec{x}) + K_s F V(\theta_R, \phi_R) Y_{lm}(\theta_R, \phi_R) \right). \tag{6.16}$$

Estimation of Parametric Dual Lighting Model: Another approach is to estimate the dual angular and frequency-space lighting model of section 6.2. Our algorithm is based on subtracting out the diffuse component B_d of the reflected light field. After this, we treat the object as a mirror sphere, recovering a high-resolution angular-space version of the illumination from the specular component alone. To determine B_d, we need only the 9 lowest frequency-space coefficients L_{lm} with $l \leq 2$. Our algorithm uses the following methods to convert between angular and frequency-space:

1. **9 parameters to High-Resolution Lighting:** The inputs to phase 1 are the coefficients L^1_{lm}. These suffice to find B^1_d by equation 6.4. Since we assumed that $B_{s,fast} = 0$,

$$\begin{aligned} B_{s,slow} &= K_s F(\mu, \theta'_o) L_{slow}(\vec{R}) \quad = B - B^1_d(L^1_{lm}) \\ L_{slow}(\vec{R}) &= \frac{B - B^1_d(L^1_{lm})}{K_s F(\mu, \theta'_o)}. \end{aligned} \tag{6.17}$$

   We assume the BRDF parameters are known, and B is the input to the algorithm, so the right-hand side can be evaluated.

   In practice, we will have several observations corresponding to the reflected direction, and these will be weighted by the appropriate confidence and combined. For

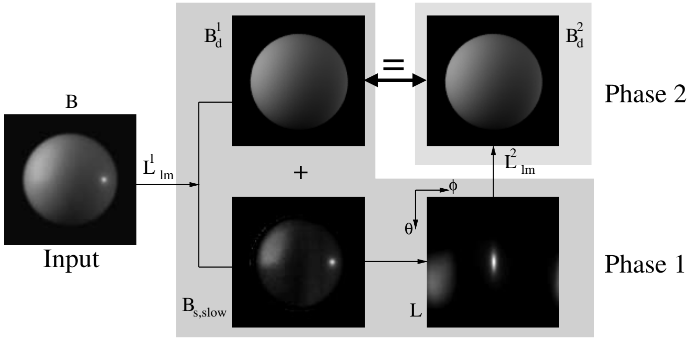

Figure 6.4: *Estimation of dual lighting representation. In phase 1, we use frequency-space parameters L^1_{lm} to compute diffuse component B^1_d. This is subtracted from the input image, leaving the specular component, from which the angular-space lighting is found. In phase 2, we compute coefficients L^2_{lm}, which can be used to determine B^2_d. The consistency condition is that $B^1_d = B^2_d$ or $L^1_{lm} = L^2_{lm}$. In this and all subsequent figures, the lighting is visualized by unwrapping the sphere so θ ranges in equal increments from 0 to π from top to bottom, and ϕ ranges in equal increments from 0 to 2π from left to right (so the image wraps around in the horizontal direction).*

simplicity, the rest of the mathematical discussion will assume without loss of generality, that there is a single image observation for each reflected direction.

2. **High-Resolution Lighting to 9 parameters:** Using the angular space values L found from the first phase, we can easily find the 9 frequency-space parameters of the lighting L^2_{lm}.

Now, assume we run phase 1 (with inputs L^1_{lm}) and phase 2 (with outputs L^2_{lm}) sequentially. The *consistency condition* is that $L^1_{lm} = L^2_{lm}$—converting from frequency to angular to frequency space must not change the result. Equivalently, the computed diffuse components must match, i.e. $B^1_d(L^1_{lm}) = B^2_d(L^2_{lm})$. This is illustrated in figure 6.4. Since everything is linear in terms of the lighting coefficients, the consistency condition reduces to a system of 9 simultaneous equations. After solving for L_{lm}, we run phase 1 to determine the high-resolution lighting in angular space.

More formally, phase 1 can be written as a linear system in terms of constants U and W_{lm}, with (α, β) the coordinates of the surface normal,

$$
\begin{aligned}
L_{slow}(\theta_R, \phi_R) &= U(\theta_R, \phi_R) - \sum_{l=0}^{2} \sum_{m=-l}^{l} W_{lm}(\theta_R, \phi_R) L^1_{lm} \\
U(\theta_R, \phi_R) &= \frac{B}{K_s F(\mu, \theta'_o)} \\
W_{lm}(\theta_R, \phi_R) &= \frac{K_d \hat{\rho}_l Y_{lm}(\alpha, \beta)}{K_s F(\mu, \theta'_o)}.
\end{aligned} \tag{6.18}
$$

Phase 2 to compute the lighting coefficients can also be written as a linear expression in terms of all the (discretized) reflected directions,

$$
L^2_{lm} = \frac{2\pi^2}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sin\theta_i L_{slow}(\theta_i, \phi_j) Y^*_{lm}(\theta_i, \phi_j). \tag{6.19}
$$

Here, N is the angular resolution, with the summation being a discrete version of the integral to find lighting coefficients.

But, the summation on the right hand side can be written in terms of lighting coefficients L^1_{lm}, simply by plugging in the formula for L_{slow}. We now obtain

$$
L^2_{lm} = \frac{2\pi^2}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sin\theta_i \left(U(\theta_i, \phi_j) - \sum_{l'=0}^{2} \sum_{m'=-l'}^{l'} W_{l'm'}(\theta_i, \phi_j) L^1_{l'm'} \right) Y^*_{lm}(\theta_i, \phi_j). \tag{6.20}
$$

Mathematically, the consistency condition allows us to drop the superscripts, reducing the above to a linear system for L_{lm}. This will involve a simple 9×9 linear system expressed in terms of a matrix $Q_{l'm',lm}$,

$$
\begin{aligned}
\sum_{l'=0}^{2} \sum_{m'=-l'}^{l'} Q_{lm,l'm'} L_{l'm'} &= P_{lm} \\
P_{lm} &= \frac{2\pi^2}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sin\theta_i \, U(\theta_i, \phi_j) Y^*_{lm}(\theta_i, \phi_j) \\
Q_{lm,l',m'} &= \delta_{lm,l'm'} + \frac{2\pi^2}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sin\theta_i \, W_{l'm'}(\theta_i, \phi_j) Y^*_{lm}(\theta_i, \phi_j).
\end{aligned} \tag{6.21}
$$

The summations are just discrete versions of integrals that determine the appropriate spherical harmonic coefficients. The above equation has a very intuitive explanation. It may be derived direction from equation 6.9, considering the linear system that results for the first 9 lighting terms. The key idea is that we have reparameterized by the reflection vector, so we may simply take the first 9 coefficients of the reflected light field. The formula for the irradiance becomes more complicated (because of the reparameterization) but can still be expressed in terms of the first 9 lighting coefficients. Mathematically, we can rewrite equation 6.9 for our purposes as

$$B(\theta_R, \phi_R) = K_d \sum_{l'=0}^{2} \sum_{m'=-l'}^{l'} \hat{\rho}_{l'} L_{l'm'} Y_{l'm'}(\alpha, \beta) + K_s F(\mu, \theta'_o) L_{slow}(\theta_R, \phi_R). \tag{6.22}$$

Here, we have simply parameterized the reflected light field by the reflected direction (θ_R, ϕ_R). Remember that for simplicity, we're assuming a single image, i.e. one value of (α, β) corresponding to each (θ_R, ϕ_R), with (α, β) a function of (θ_R, ϕ_R). With multiple images, we would have to weight contributions appropriately.

Now, it's a simple enough matter to compute coefficients obtaining

$$B_{lm} = K_d \sum_{l', m'} \hat{\rho}_{l'} L_{l'm'} \left\langle Y^*_{lm}(\theta_R, \phi_R), Y_{l'm'}(\alpha, \beta) \right\rangle + K_s F(\mu, \theta'_o) L_{lm}. \tag{6.23}$$

Here, we have used the notation $< \cdot, \cdot >$ for the integral or inner product over the spherical domain of integration. This is what is computed discretely in equation 6.22. It can now be seen that equation 6.23 has the same form as equation 6.22. Note that in equation 6.23, we have multiplied out the denominators, and we use B_{lm} here instead of P_{lm}.

This inverse lighting method is difficult to extend to concave surfaces, since the 9 parameter diffuse model is no longer entirely valid. It is a subject of future work to see if it can be applied simply by increasing the number of parameters and the size of the matrix of simultaneous equations to be solved.

Positive regularization: So far, we have not explicitly tried to ensure positivity of the illumination. In practical applications, the methods above when applied naively will result in negative values, especially where the illumination is dark, and there is uncertainty about the precise value. Regularizing so the results are positive can also substantially improve the quality of the results by reducing high-frequency noise centered close to the zero point in dark regions.

We apply positive regularization to the unregularized solution from either of the previous two methods. For the first method (direct solution of linear system to determine lighting coefficients), we simply add another term to the RMS error which penalizes negative regions. While this method is a soft constraint and can still leave some negative regions, we have observed that it works well in practice. We use a conjugate gradient method to minimize

$$\nu' = \nu + \lambda \left(\frac{\nu_0}{\kappa_0} \right) \kappa. \tag{6.24}$$

Here, ν is the RMS error corresponding to equation 6.15 or 6.16. κ is a new penalty term added to penalize negative values of the lighting. ν_0 and κ_0 are initial values (for the unregularized solution), and λ weights the importance of the penalty term. λ is a dimensionless quantity, and we have found experimentally that $\lambda = 1$ works well. The penalty term κ is simply the sum of squares of all lighting pixels having negative values. Thus, negative values are penalized, but no penalty is imposed for positive values.

For our second method (using a dual angular and frequency-space method to estimate the lighting), regularization may be enforced (in step 1) simply by clamping L_{slow} to 0 if the right hand side in the first line of equation 6.18 is negative. This must be taken into account in the 9×9 simultaneous equations, and we solve the *positivity enforced* equations with a conjugate gradient method, using as a starting guess the solution without enforced positivity.

Comparison: Figure 6.5 compares the methods to each other, and to a reference solution from a gazing sphere. Both algorithms give reasonably accurate results. As predicted by the theory, high-frequency components are filtered by the roughness σ. In the first method, involving direct recovery of L_{lm}, there will still be some residual energy for $l > l^*$. Since we regularize by not considering higher frequencies—we could increase l^*, but this makes the result noisier—the recovered lighting is somewhat blurred compared to our dual angular and frequency-space algorithm (second method). As expected, positive regularization in algorithm 2 results in a smoother solution.

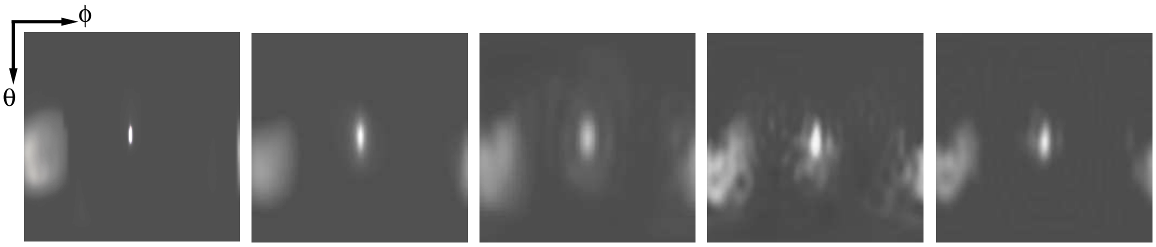

Figure 6.5: *Comparison of inverse lighting methods. From left to right, real lighting (from a gazing sphere), recovered illumination by direct estimation of spherical harmonic coefficients with* $l^* = 12$ *and* $l^* = 20$*, and estimation of dual angular and frequency-space lighting model. To make the artifacts more apparent, we have set* 0 *to gray. The results from the dual algorithm are sharper, but still somewhat blurred because of filtering by* σ*. A small amount of ringing occurs for direct coefficient recovery, and can be seen for* $l^* = 12$*. Using* $l^* = 20$ *makes the solution very noisy. Positive regularization (rightmost) gives a smoother solution.*

6.4.4 Factorization—Unknown Lighting and BRDF

We can combine the inverse-BRDF and inverse-lighting methods to *factor* the reflected light field, simultaneously recovering the lighting and BRDF when both are unknown. Therefore, we are able to accurately recover BRDFs of curved surfaces under unknown complex illumination, something which has not previously been demonstrated. There is an unrecoverable global scale factor, so we set $K_d + K_s = 1$; we cannot find absolute reflectance. Also, the theory predicts that for low-frequency lighting, estimation of the surface roughness σ is ill-conditioned—blurring the lighting while sharpening the BRDF does not significantly change the reflected light field. However, for high-frequency lighting, this ambiguity can be removed. We will use a single manually specified directional source in the recovered lighting distribution to estimate σ.

Algorithm: The algorithm consists of nested procedures. In the outer loop, we effectively solve an inverse-BRDF problem—a nonlinear simplex algorithm adjusts the BRDF parameters to minimize error with respect to image pixels. Since $K_d + K_s = 1$, and σ will not be solved for till after the lighting and other BRDF parameters have been recovered, there are only 2 free parameters, K_s and μ. In the inner procedure, a linear problem is solved to estimate the lighting for a given set of BRDF parameters, using the methods of the previous subsection. Pseudocode is given below.

global B_{input}	// Input images
global K_d,K_s,μ,σ	// BRDF parameters
global L	// Lighting
procedure Factor	
Minimize(K_s,μ,ObjFun)	// Simplex Method
σ = FindRoughness(L)	// Figure 6.6, Equation 6.25
function ObjFun(K_s,μ)	
$K_d = 1 - K_s$	// $K_d + K_s = 1$
L = Lighting(K_d,K_s,μ)	// Inverse Lighting
B_{pred} = Predict(L,K_d,K_s,μ)	// Predicted Light Field
return RMS(B_{input},B_{pred})	// RMS Error

Finding σ using a directional source: If a directional source is present—and manually specified by us in the recovered lighting—we can estimate σ by equating specular components predicted by equations 6.6 and 6.8 for the center, i.e. brightest point, of the light source at normal exitance. An illustration is in figure 6.6,

$$L_{cen} \approx \frac{L_{total}}{4\pi\sigma^2}. \tag{6.25}$$

Color: We have so far ignored issues of color, assuming the three color channels are considered separately. However, in the case of BRDF recovery under unknown lighting, there is a separate scale factor associated with each color channel. In order to obtain accurate colors for the BRDF and lighting components, we need some way to relate these 3 scale factors. For dielectrics, the specular component K_s is not spectrally sensitive, i.e. it is the same for red,green, and blue channels. The recovered BRDFs are scaled in order to make this hold. The issue is trickier for metals. There is a fundamental ambiguity between the color of the BRDF and the color of the lighting. We resolve this by considering the average

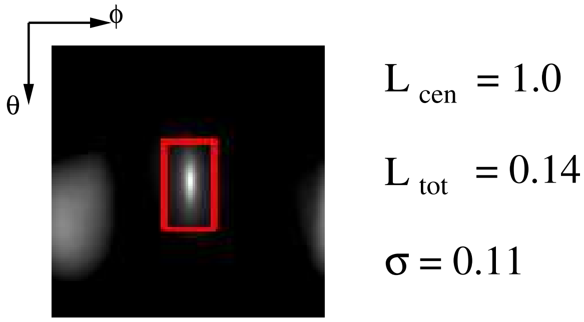

Figure 6.6: *Determining surface roughness parameter σ. We manually specify (red box) the region corresponding to the directional source in a visualization of the lighting. The algorithm then determines L_{cen}, the intensity at the center (brightest point), L_{tot}, the total energy integrated over the region specified by the red box, and computes σ using equation 6.25. The method does not depend on the size of the red box—provided it encloses the entire (filtered) source—nor the precise shape into which the source is filtered in the recovered lighting.*

color of the metallic surface as corresponding to white light. The use of more sophisticated color-space separation methods such as that of Klinker et al [42] might bring further benefits.

Results: We used the method of this subsection—with the dual angular and frequency-space algorithm for inverse lighting—to factor the light field for the spheres, simultaneously estimating the BRDF and lighting. The same setup and lighting were used for all the spheres so we could compare the recovered illumination distributions.

We see from figure 6.7 that the BRDF estimates under unknown lighting are accurate. Absolute errors are small, compared to parameters recovered under known lighting. The only significant anomalies are the slightly low values for the refractive index μ—caused because we cannot know the high-frequency lighting components, which are necessary for more accurately estimating the Fresnel term. We are also able to estimate a filtered version of the lighting. As shown in figure 6.8, the recovered lighting distributions from all the samples are largely consistent. As predicted by the theory, the directional source is spread out to different extents depending on how rough the surface is, i.e. the value of σ. Finally, figure 6.8 shows that rendered images using the estimated lighting and BRDF are almost indistinguishable from real photographs.

Material	K_d		K_s		μ		σ	
	Known	Unknown	Known	Unknown	Known	Unknown	Known	Unkn.
Teflon	0.89	0.87	0.11	0.13	1.78	1.48	0.12	0.14
Delrin	0.87	0.88	0.13	0.12	1.44	1.35	0.10	0.11
Neoprene Rubber	0.92	0.93	0.08	0.07	1.49	1.34	0.10	0.10
Sandblasted Steel	0.20	0.14	0.80	0.86			0.20	0.19
Bronze	(.15,.08,.05)	(.09,.07,.07)	(.85,.68,.59)	(.91,.69,.55)			0.12	0.10
Painted	(.62,.71,.62)	(.67,.75,.64)	0.29	0.25	1.38	1.15	0.15	0.15

Figure 6.7: *BRDFs of various spheres, recovered under known (section 6.4.2) and unknown (section 6.4.4) lighting. The reported values are normalized so $K_d + K_s = 1$. RGB values are reported for colored objects. We see that K_s is much higher for the more specular metallic spheres, and that σ is especially high for the rough sandblasted sphere. The Fresnel effect is very close to 1 for metals, so we do not consider the Fresnel term for these spheres.*

6.5 Results on Complex Geometric Objects

In the previous section, we presented our new algorithms for inverse rendering with complex illumination, illustrating their performance using spheres of different materials. To demonstrate the practical applicability of these methods, in this section, we report on two experiments using complex geometric objects that include concavities and self-shadowing. The previous section has already discussed, where appropriate, how the algorithms for BRDF and lighting estimation can be extended to concave and textured surfaces.

Our first experiment uses a white cat sculpture of approximately uniform material properties to demonstrate BRDF estimation under known and unknown complex illumination. This is the first demonstration of accurate BRDF estimation under complex unknown illumination for geometrically complex objects. Geometry was acquired using a Cyberware range scanner and aligned to the images by manually specifying correspondences. The lighting was slightly more complex than that for the *spheres* experiment; we used a second directional source in addition to the area source.

To show that we can recover BRDFs using a small number of images, we used only 3 input photographs. We recovered BRDFs under both known lighting, using the method of section 6.4.2, and unknown lighting—using the factorization method of section 6.4.4, with the inverse lighting component being direct recovery of spherical harmonic coefficients using $l^* = 12$. Comparisons of photographs and renderings are in figures 6.9 and 6.10. BRDF and lighting parameters are tabulated in figure 6.11. This experiment indicates that our methods for BRDF recovery under known and unknown lighting are consistent, and

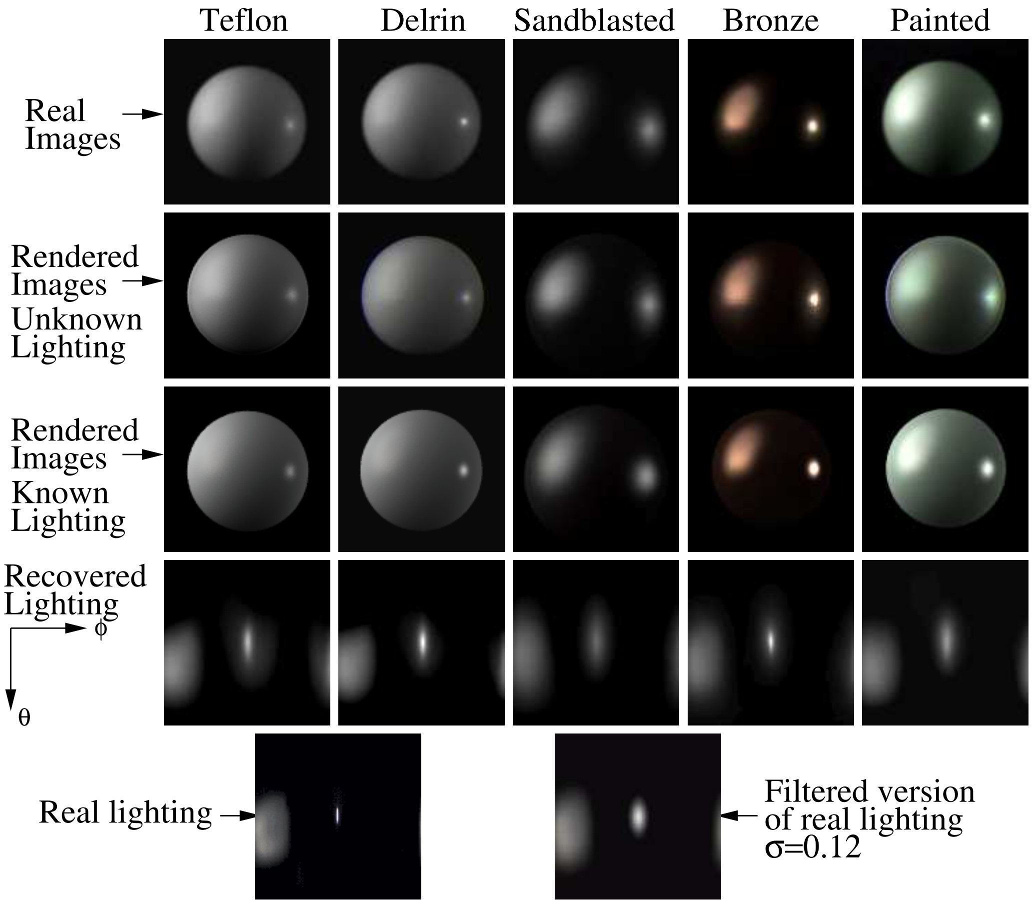

Figure 6.8: *Spheres rendered using BRDFs estimated under known (section 6.4.2) and unknown (section 6.4.4) lighting. The algorithm in section 6.4.4 also recovers the lighting. Since there is an unknown global scale, we scale the recovered lighting distributions in order to compare them. The recovered illumination is largely consistent between all samples, and is similar to a filtered version of the real lighting. As predicted by the theory, the different roughnesses σ cause the directional source to be spread out to different extents. The filtered source is slightly elongated or asymmetric because the microfacet BRDF is not completely symmetric about the reflection vector.*

are accurate even for complex lighting and geometry. The rendered images are very close to the original photographs, even under viewing and lighting conditions not used for BRDF recovery. The most prominent artifacts are because of imprecise geometric alignment and insufficient geometric resolution. For instance, since our geometric model does not include the eyelids of the cat, that feature is missing from the rendered images.

Figure 6.9: *Comparison of photograph and rendered image for cat sculpture.* **Left:** *Photograph* **Right:** *Rendered image. The BRDF used for the rendered image was estimated under complex unknown illumination from 3 photographs of a cat sculpture with known geometry. Our algorithm also recovered the lighting distribution, which consisted of two directional sources and an area source. The images above show a new view not used in BRDF recovery; the lighting is also new, being composed of a single directional source (with known direction) not used in BRDF estimation. These images show that the recovered BRDF accurately predicts appearance even under novel viewing and lighting conditions.*

Textured BRDFs: Since the theory shows that factorization of lighting and texture is ambiguous, we consider only recovery of textured BRDFs under known lighting, using the method described at the end of section 6.4.2. As an experimental test, we used a wooden doll. We compared the real input photographs with images rendered using the recovered textured BRDF. We also took a photograph of the same object under a single directional source—directional sources have all the high-frequency components, and therefore bring out all the features of the BRDF filter—and compared this to a rendering using the textured BRDF recovered under complex illumination. The results in figure 6.12 show that our renderings closely resemble real photographs. The main artifact is blurring of texture because of geometry-image misregistration. This is unrelated to the use of complex lighting, and is

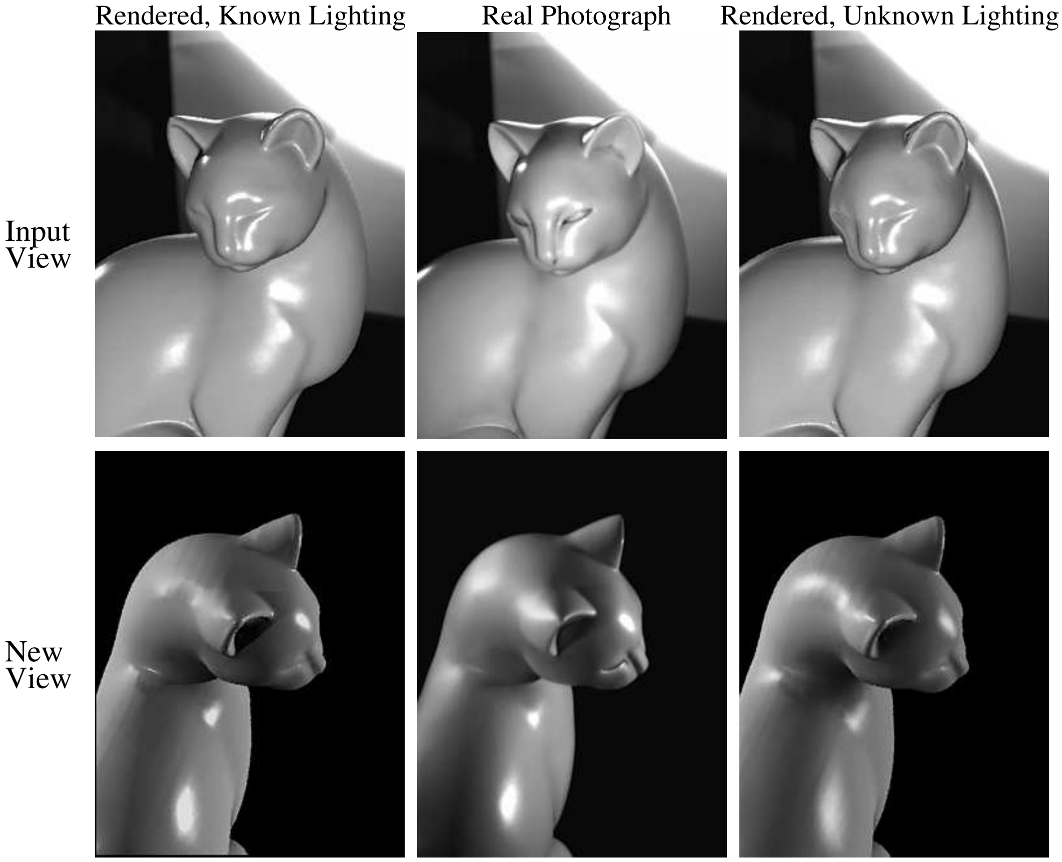

Figure 6.10: *Comparison of photographs (middle column) to images rendered using BRDFs recovered under known lighting (left column), and using BRDFs (and lighting) estimated under unknown lighting (right column). The top row is one of the 3 input views. The bottom row is a new view, not used for BRDF estimation. Note that in the top row, we have composited the left and right renderings over the same background as the middle photograph in order to make a meaningful comparison.*

also noted by Sato et al. [77].

6.6 Conclusions and Future Work

This chapter has considered the problem of complex illumination in inverse rendering. The use of realistic lighting and materials is likely to be of increasing importance in graphics and vision, and inverse rendering to acquire real-world material properties is likely to be a significant future direction. The use of natural lighting is becoming increasingly common

Parameter	Known Lighting	Unknown Lighting
BRDF Parameters		
K_d	0.88	0.90
K_s	0.12	0.10
μ	1.68	1.47
σ	0.12	0.14
Lighting Coefficients (l,m)		
(l,m) = (0,0)	0.68	0.68
(1,-1)	-0.06	-0.02
(1,0)	-0.17	-0.15
(1,1)	-0.02	-0.06
(2,-2)	0.10	0.04
(2,-1)	0.03	0.09
(2,0)	-0.55	-0.51
(2,1)	0.30	0.28
(2,2)	0.84	0.88

Figure 6.11: *BRDF and lighting parameters for the cat sculpture. We see good agreement between BRDF parameter values recovered with known and unknown lighting, showing our methods are consistent. Note that we normalize so $K_d + K_s = 1$. We may also check the accuracy of the recovered lighting. Since there is an unknown global scale for the recovered values, we report normalized lighting coefficient values for the first 9 spherical harmonic coefficients (in real form), which are the most important, because they significantly affect the diffuse component of the BRDF.*

in computer graphics and vision, and this chapter allows for inverse rendering techniques to be applied in arbitrary uncontrolled conditions rather than a laboratory setting. Furthermore, in certain cases, complex illumination may help in solving inverse problems, such as by allowing a much larger fraction of an image to exhibit specularity. Dror et al. [16] have also shown that people perceive reflectance properties much more easily under natural illumination.

This chapter first presents a taxonomy of inverse problems, identifying the many unexplored directions of work in solving inverse problems under complex illumination. We then address some of these areas. The insights gained from our signal-processing framework for reflection lead to a new practical representation. We can numerically represent quantities in angular or frequency space, depending on where they are more local.

This leads to new algorithms which are often expressed in a combination of angular and

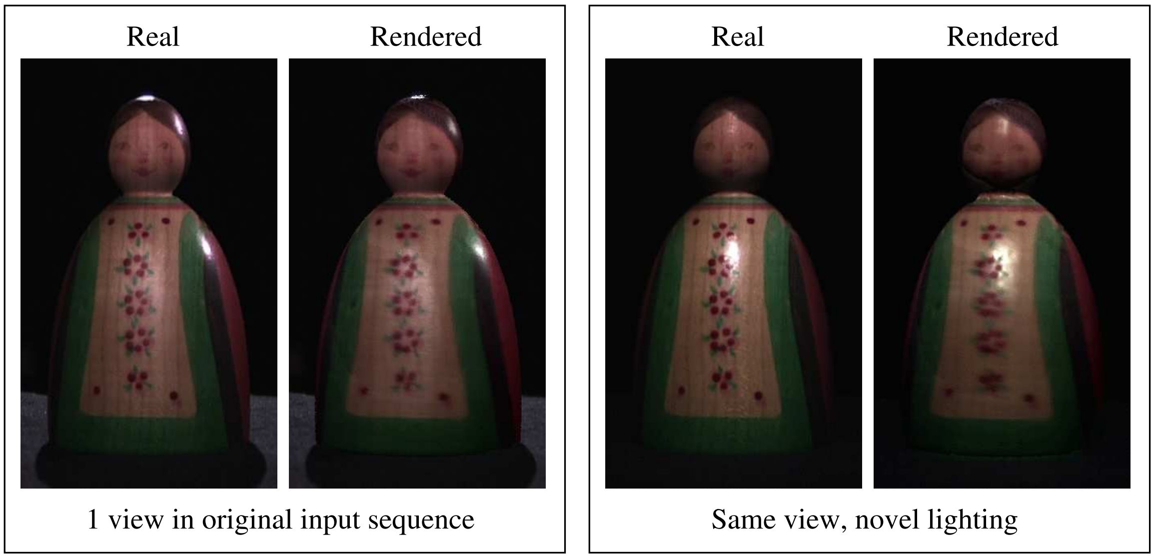

Figure 6.12: *Recovering textured BRDFs under complex lighting. The rendered images closely resemble the real photographs, even under novel lighting.*

frequency-space. We can determine which BRDF and lighting parameters are important, and can handle the various components appropriately. For BRDF estimation, the parametric recovery algorithms of Yu and Malik [90], Sato and Ikeuchi [76], and Love [50]—which are designed specifically for natural lighting—can be seen as special cases of this general approach; they treat sunlight (high-frequency) and skylight (low-frequency) separately. We provide a general framework for arbitrary illumination, and also determine conditions under which parameter recovery is robust. For instance, our theory predicts that estimation of σ is ill-conditioned on a cloudy day, with only low-frequency lighting. Our framework can also be applied to developing new frequency-space algorithms to estimate the lighting from objects with general BRDFs. The use of frequency-space naturally handles continuous lighting distributions. Our dual angular and frequency-space algorithm effectively reduces the problem for general BRDFs to that for a gazing sphere, requires no explicit regularization, and allows much higher angular resolutions to be obtained than with previous purely angular-space methods [54, 75]. Finally, we demonstrate a method for *factoring* the light field to simultaneously estimate the lighting and BRDF. This allows us to estimate BRDFs of geometrically complex objects under unknown general lighting, which has not previously been demonstrated.

Our results presented here do contain a number of limitations. The examples presented here, while intended to show the proof of concept, have relatively simple lighting configurations. In the future, we would like to apply the methods of this chapter under more complex outdoor and indoor illumination conditions. The BRDF model used by us is also a relatively simple microfacet model. As noted, we have observed a number of difficulties using a more general model based on spherical harmonic coefficients. In terms of algorithmic issues, we do not currently handle interreflections. We believe that in the future, they could be addressed by an iterative method in a manner similar to that of Yu et al. [89]. We have also presented results using only a single object and relatively distant illumination. In the future, we wish to extend our results to entire scenes and possibly near-field illumination, as in the work of Boivin and Gagalowicz [6] or Yu et al. [89], thereby extending those methods to complex illumination conditions. Finally, it should be pointed out that the results obtained by us are not completely faithful to the original objects, which indicates that further work is required on relaxing many of the assumptions made by us.

Note that the techniques presented in this chapter are only the first step in solving inverse problems under complex illumination. A number of the open problems identified in our taxonomy remain subjects for future work. There has been relatively little work on BRDFs in between low-parameter models and full measured representations. It is not clear what the best way to measure factored or separable BRDF representations is, or how to estimate these or higher-dimensional representations under complex illumination. We have not considered statistical properties of the illumination, that may allow us to simultaneously determine the lighting, BRDF and texture. All of our examples use a small number of images, and an interesting future direction is whether a single image suffices. We also believe there are more insights to be obtained from frequency space analysis and new frequency domain algorithms to be explored for inverse problems. Another interesting future direction is to preserve and extrapolate the original images to new conditions, using inverse rendering only as a guide in this extrapolation, but without explicitly computing lighting and reflectance parameters. This would have the potential to unify inverse and image-based (interpolatory) rendering methods.

Finally, the use of measured reflectance and illumination properties is of growing importance in graphics and vision, and we have taken a step towards finding general inverse methods to determine these from image observations.

In chapters 4 and 5 of this dissertation, we have applied the signal-processing framework developed in chapters 2 and 3 to forward rendering problems. In this chapter, we have demonstrated the applicability to inverse rendering. The next chapter concludes this dissertation, presenting a unified discussion of forward and inverse rendering, and suggests directions for future work.

Chapter 7

Conclusions and Future Work

In this dissertation, we have presented a new way of analyzing a basic building block in computer graphics rendering algorithms—the computational interaction between illumination and the reflective properties of a surface. We formalize the notion of reflection on a curved surface as a spherical convolution of the incident illumination and the bi-directional reflectance distribution function of the surface. This allows us to develop a signal-processing framework for reflection, leading to new frequency domain and combined angular and frequency-space methods for forward and inverse problems in computer graphics, taking into account the effects of complex natural illumination, and physically realistic BRDFs. We use the theoretical framework to develop new practical algorithms for interactive forward rendering using environment maps and inverse rendering to estimate lighting and material properties under complex illumination conditions.

The main thesis validated by this dissertation is that a deeper understanding of the computational nature of reflection and illumination is important and leads to new practical algorithms. Secondly, since the properties of the reflection operator impact a number of different areas, the same fundamental insights lead to novel algorithms for a number of different problems in forward and inverse rendering in graphics, as well as related areas in computer vision. Since this dissertation has both a theoretical and practical component, we discuss below our main conclusions, the bigger picture, and future directions for each of the components.

7.1 Computational Fundamentals of Reflection

Although the physical phenomena and mathematical framework for reflection in the spatial or angular domain are well known, there has not so far been a deep understanding of the computational issues. Therefore, it is important to lay the mathematical foundations for this area, which I refer to as the computational fundamentals of reflection. This dissertation has presented one such analysis in terms of signal-processing. Although the qualitative notion of reflection as convolution has a fairly long history of nearly 20 years in the graphics community and a number of people have tried to formalize this idea, this is the first time a precise mathematical description has been derived, with an associated convolution theorem. The analysis in chapters 2 and 3 has led to a number of basic questions being answered including the fundamental limits and conditioning of inverse problems, i.e. from observations of an object, what can we know about the lighting/reflectance? Our analysis also leads to a number of interesting observations, some of which may be lines for future work.

Relationship between Forward and Inverse Problems: Firstly, we have demonstrated a significant duality between forward and inverse problems. An ill-posed or ill-conditioned inverse problem is one where the input signal cannot be estimated from the output, because it is attenuated or truncated by the reflection operator (or BRDF filter). An example is illumination estimation from a diffuse or Lambertian surface. However, this also means that in the forward direction, most of the information in the input signal is not required to synthesize the output image. In chapter 4, we saw that this made it very efficient to represent and compute irradiance maps corresponding to reflections from diffuse or Lambertian surfaces. It is likely that similar considerations of duality can be profitably employed to analyze and create efficient and robust algorithms for forward and inverse problems in the future.

Dual Angular and Frequency-Space Representations: The Heisenberg uncertainty principle can effectively be interpreted as saying that quantities cannot simultaneously be localized in the spatial (angular) and frequency domains. Thus, quantities local in angular space have broad frequency spectra and vice-versa. However, we can turn this observation to our advantage by making use of both angular and frequency domains simultaneously.

By developing a frequency-space view of reflection, we ensure that we can use either the angular-space or frequency-space representation, or even a combination of the two. The diffuse BRDF component is slowly varying in angular-space, so we represent it in the frequency domain where it is a very local low-pass filter. Similarly, it is often useful to represent the specular BRDF component in the angular domain, where it is local, as opposed to the frequency domain where a large number of terms are required. This is the argument behind our representation in section 6.3. This principle of duality is likely to have many further implications, with algorithms being designed separately for high and low-frequency components, and investigation of tradeoffs and crossover points between the two.

If we seek to extend and apply our signal-processing framework in cases where there is no specific convolution formula, then it is likely that significant benefit can be obtained using wavelet or multiresolution representations, one of whose properties is to create basis functions local in both spatial (angular) and frequency domains. Such representations might provide a unified way to account for both high frequency (local in the angular domain) and low-frequency (local in the frequency domain) components.

Sampling Theory: An obvious extension of the signal-processing framework is to develop a sampling theory for various problems in graphics, determining how many observations or samples are needed for reconstruction. For instance, in the Lambertian case, our theory indicates that only 9 samples are required to reconstruct the irradiance map, if resampling is done correctly. A general framework for determining reflection map resolutions and sampling rates is described in chapter 5. More generally, we seek in future to address some important problems in graphics and vision. Within the context of image-based rendering, a long-standing question is how many images are needed to produce the effects we desire. Too many, and data acquisition and manipulation become unwieldy. Too few, and there are artifacts and we fail to get the effects we want. Some work on this topic under the assumption of Lambertian surfaces and with imprecise geometry has been presented by Chai et al. [33]. It is interesting to consider a framework based on our analysis, assuming good or perfect geometry, but taking into account the effects of complex illumination and reflection models. A number of questions need to be answered such as the number of images required for good reconstruction and the right interpolation approach taking into account the frequency variation of the various quantities.

Low-Dimensional Subspaces: One of the results of our theory is formalizing the notion of low-dimensional subspaces for images of a Lambertian object (9D) and similar results for other reflection functions. Similar ideas have been used by Nimeroff et al. [63] and Teo et al. [82] for lighting design. One interesting question is whether the theoretical analysis in this dissertation can be seen as a special case of a general framework for determining good low-dimensional subspace approximations of the reflection operator, and whether our theoretical analysis can be adapted to cases where information on probable lighting conditions, reflection functions and object geometry are known. One approach is to find the principal components or basis function images of the space of all possible images of an object. This is a common approach used for compression in graphics and vision, but so far without much understanding of the number of principal components required. A preliminary step has been taken by us [69], where we have shown how to analytically construct the PCA for the space of images of a Lambertian object under all possible distant illumination conditions. We have shown that under idealized assumptions, the principal components reduce to the spherical harmonic basis functions, but must be modified for a single viewpoint where only the front-facing normals are visible. We have also shown how to adapt the 9D Lambertian result to derive lower-dimensional subspaces when we have only a single viewpoint. A promising future direction is to develop a general theory allowing for derivation of low-dimensional results under arbitrary assumptions about lighting, reflectance, geometry, and the effects to be considered in the reflection operator.

Integral Equations: It may be possible to view our convolution result as an analytic formula for a special case of the integral equation corresponding to the reflection operator. Similarly, the convolution result and the spherical harmonic basis functions may be the optimal low-dimensional approximations under idealized assumptions. This leads to the possibility of analyzing the integral equation corresponding to the reflection and rendering operators to determine the ill-posedness and conditioning of inverse problems, and the rank of the kernel directly to determine low-dimensional subspaces.

Differential Framework: Of course, signal-processing is only one way of conducting a computational analysis of the reflection operator. Another approach is to consider a differential framework for reflection. This answers questions about how to interpolate and extrapolate neighbouring views, or how much information is available from nearby observations, such as the slightly separated two eyes. A significant amount of work has been carried out in the image domain in the computer vision community under the aegis of optical flow methods. However, we believe significant insight and practical algorithms may emerge from considering an analysis on the object surface, taking into account complex illumination and reflection functions.

Perception: This dissertation has focussed on the physical computation of light transport. However, humans do not perceive on a linear scale; rather, the perceptual response is logarithmically related to the physical intensity. This has been exploited within the context of perceptual theories like the retinex framework [46] that often perform a logarithmic transformation before applying particular algorithms. In this case, separation of high-frequency texture and low-frequency illumination is made easier after a logarithmic transform. Similarly, in computer graphics, a number of homomorphic factorization techniques [47, 57] have been applied to represent reflected light fields and BRDFs by considering the error and doing computations in a logarithmic space. All of these results lead to a number of interesting questions about how approximations should be done, whether nonlinearities are introduced and whether a perceptual-space analysis and convolution relation can be derived. Considering perceptual aspects remains an important direction of future work.

7.2 High Quality Interactive Rendering

The first practical application discussed in the dissertation is of interactive rendering with natural illumination and physically based BRDFs. High quality interactive rendering is important in many applications including visualization, simulation and training, and video games. However, current graphics hardware rendering is usually limited to very simple lighting models (usually point and directional sources only) and reflection functions (Phong and Lambertian BRDFs). This imposes a gap between photorealism and interactivity that we have seeked to bridge. Within this context, we expect a number of future improvements.

Interactive rendering with complex illumination, shading, visibility: In this dissertation, we have addressed interactive rendering with complex illumination and reflectance functions. Several further steps may be taken to improve the quality of interactive computer graphics. Future work will encompass effects like spatially varying illumination, interreflection, and cast shadows. In the long term, we would like to see all the effects currently expensive to compute with global illumination simulations, incorporated into interactive renderings.

Multidimensional rendering: Another area of future work concerns multidimensional rendering. Most current computer graphics algorithms are optimized for creating a single image. However, many current computer graphics techniques deal with multi-view, multiple-light or other dimensions of data. While inverse rendering techniques to infer parameters for hardware rendering of multidimensional image spaces have been demonstrated by Hakura [26], there has been very little work on efficient multidimensional rendering. Halle [29] has shown how to efficient scan-convert a scene for multiple views, but much future work needs to be done in the context of efficiently rendering surface light fields or other multidimensional animations. We believe the fast prefiltering algorithms in sections 4 and 5 provide an important first step in this endeavor.

Application of computational fundamentals: Many of the theoretical analyses discussed in the previous subsection can be applied to a number of problems in rendering. For instance, differential analysis can be applied to accurate interpolation, and incremental changes to programmable shading calculations. Determining appropriate sampling rates can lead to new methods for antialiasing or filtering bump maps. Low-dimensional subspace methods and angular-frequency tradeoffs could be profitably applied in a number of domains in rendering.

7.3 Inverse Rendering

Realism in computer-generated images requires accurate input models for lighting, textures and reflectances (BRDFs). One of the best ways of obtaining high-quality data is through measurements of scene attributes from real photographs by **inverse rendering**. Measuring scene attributes also introduces structure into the raw imagery, allowing an artist to independently manipulate the material properties or lighting. This dissertation has presented a coherent mathematical framework for inverse rendering under general illumination conditions. Besides the formal study of the well-posedness and conditioning of inverse problems, we have derived new practical representations and novel frequency domain and dual spatial and frequency domain algorithms for BRDF and lighting estimation. Chapter 6 demonstrates examples of all these algorithms, and shows that the rendered images appear nearly identical to real photographs.

There remain some limitations in our work that we seek to address in the future, such as handling interreflections and extending the results to more complex illumination conditions and entire scenes. Besides this, there remain a number of problems to be solved.

Structured BRDF and higher-order representations: There are a number of methods for estimating BRDFs. These lead to incomplete information; from this, a full BRDF must be reconstructed for practical use in applications. One future direction is to explore how various representations—parametric, generative, and tabulated can be fit as data of increasing precision arrives. We also wish to study ideas with respect to adaptive reconstruction. Also, based on differential analysis, a number of higher-order differential properties are worth estimating. Finally, structured representations of the high-dimensional datasets in graphics today are very important for representation, understanding, manipulation and compression.

Factored BRDF representations: There has been relatively little theoretical or practical work on BRDFs in between low-parameter models and full measured representations. This remains a subject of future work. It is interesting to consider curve-based or factored BRDF models, and how they can be estimated from a small number of photographs, under both simple and complex illumination conditions. More generally, the idea of *factoring* the reflected light field, or of constructing structured representations of the high-dimensional datasets currently prevalent in graphics (such as BRDFs, BTFs and light fields) is one of growing interest.

Combined estimation of illumination, materials and geometry: Still another interesting question is whether we can combine inverse rendering to find illumination and material properties with geometry estimation, or at least use reflection information to determine improved estimates of geometric properties. In any case, our work is clearly applicable in the context of computer vision algorithms that determine geometry; we seek to extend these methods to incorporate complex illumination and reflectance properties.

General frameworks for extrapolatory rendering: Taking a step back, our goals really extend beyond estimating material properties. In computer graphics, we often wish to actually use the difference (or "error") between the photographs and the predicted model. It is often this difference which gives the photograph its real-world, as opposed to computer-generated, appearance. A simple example is the use of texture maps, wherein the real-world texture can be used to modulate or multiply shading computations from a light transport model that assumes homogeneous objects.

This is very different from most use of empirical data elsewhere in science—where we expect the theory to fit the data, treating errors as experimental noise to be ignored. Thus, one goal of future research is to derive a general framework for combining empirical observations in the form of real photographs with a model-based approach, perhaps relying on inverse rendering, to get very realistic images that can be easily manipulated using standard graphics operations. For instance, we could manipulate the low-frequency information predicted by our rendering model, such as illumination effects, while preserving the high-frequency information in the photograph, such as object texture.

A general framework along these lines would allow us to unify inverse and image-based (interpolatory) rendering methods. Similar ideas are likely to be applicable to animation, where we want to effectively edit or manipulate motion-capture data, and modeling, where we want to manipulate range data.

7.4 Summary

In summary, this dissertation has presented a new approach toward forward and inverse problems in graphics and vision. We have analyzed the computational properties of the reflection operator, in this case in terms of signal-processing. The ideas from this analysis lead to new robust and efficient algorithms for domains that have hitherto been considered largely separate—interactive high quality rendering, and measurement of lighting and material properties from photographs by inverse rendering. Basri and Jacobs [2, 3] have shown the further application of these ideas to the problems of photometric stereo and lighting-invariant recognition in computer vision. This suggests a new approach to problems in graphics and vision. We first make progress in understanding the *computational fundamentals of reflection*. This in turn leads to broad practical impact for forward and inverse problems in rendering and computer vision—domains that have hitherto been considered largely distinct. This dissertation has presented one such investigation in terms of signal-processing, and we believe there is much progress to be made.

Appendix A

Properties of the Representation Matrices

In this appendix, we derive the two properties of representation matrices listed in equation 2.35. The first property follows from the addition theorem for spherical harmonics (see for instance, Jackson [34] equation 3.62),

$$Y_{l0}(u,v) = \Lambda_l \sum_{m=-l}^{l} Y^*_{lm}(\theta,\phi) Y_{lm}(\theta',\phi'). \tag{A.1}$$

Here, v is a dummy-variable since Y_{l0} has no azimuthal dependence, and u is the angle between (θ,ϕ) and (θ',ϕ'), i.e.

$$\cos u = \cos\theta\cos\theta' + \sin\theta\sin\theta'\cos(\phi-\phi'). \tag{A.2}$$

Now, let $(u,v) = R_\alpha(\theta',\phi')$. Here, $R_\alpha = R_y(\alpha)$. We omit the z rotation since that does not affect Y_{l0} which has no azimuthal dependence. The vector corresponding to coordinates (u,v) is then given by

$$\begin{pmatrix} \sin u\cos v \\ \sin u\sin v \\ \cos u \end{pmatrix} = \begin{pmatrix} \cos\alpha & 0 & \sin\alpha \\ 0 & 1 & 0 \\ -\sin\alpha & 0 & \cos\alpha \end{pmatrix} \begin{pmatrix} \sin\theta'\cos\phi' \\ \sin\theta'\sin\phi' \\ \cos\theta' \end{pmatrix} = \begin{pmatrix} \cos\alpha\sin\theta'\cos\phi' + \sin\alpha\cos\theta' \\ \sin\theta'\sin\phi' \\ \cos\alpha\cos\theta' + \sin\alpha\sin\theta'\left(-\cos\phi'\right) \end{pmatrix}. \tag{A.3}$$

Since $(-\cos\phi') = \cos(\pi - \phi')$, we know from equation A.2 that u corresponds to the angle between (α, π) and (θ', ϕ'). In other words, we may set $(\theta, \phi) = (\alpha, \pi)$. To summarize,

$$Y_{l0}\left(R_\alpha(\theta', \phi')\right) = \Lambda_l \sum_{m=-l}^{l} Y^*_{lm}(\alpha, \pi) Y_{lm}(\theta', \phi'). \tag{A.4}$$

To proceed further, we write the rotation formula for spherical harmonics, omitting the z rotation by β, since that has no significance for azimuthally symmetric harmonics.

$$Y_{l0}\left(R_\alpha(\theta', \phi')\right) = \sum_{m=-l}^{l} d^l_{0m}(\alpha) Y_{lm}(\theta', \phi') \tag{A.5}$$

A comparision of equations A.4 and A.5 yields the first property of representation matrices in equation 2.35, i.e.

$$d^l_{0m}(\alpha) = \Lambda_l Y^*_{lm}(\alpha, \pi). \tag{A.6}$$

To obtain the second property in equation 2.35, we use the form of the spherical harmonic expansion when the elevation angle is 0, i.e. we are at the north pole. Specifically, we note that $Y_{lm'}(0', \phi') = \Lambda_l^{-1} \delta_{m'0}$. With this in mind, the derivation is as follows,

$$\begin{aligned} Y_{lm}(\alpha, \beta) &= Y_{lm}\left(R_{\alpha,\beta,\gamma}(0', \phi')\right) \\ &= \sum_{m'=-l}^{l} D^l_{mm'}(\alpha, \beta, \gamma) Y_{lm'}(0', \phi') \\ &= \Lambda_l^{-1} D^l_{m0}(\alpha, \beta, \gamma). \end{aligned} \tag{A.7}$$

This brings us to the second property stated in equation 2.35,

$$D^l_{m0}(\alpha, \beta, \gamma) = \Lambda_l Y_{lm}(\alpha, \beta). \tag{A.8}$$

Bibliography

[1] J. Arvo. Applications of irradiance tensors to the simulation of non-lambertian phenomena. In *SIGGRAPH 95*, pages 335–342, 1995.

[2] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. In *International Conference on Computer Vision*, pages 383–390, 2001.

[3] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. In *CVPR 01*, pages II–374–II–381, 2001.

[4] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible illumination conditions? *IJCV*, 28(3):245–260, 1998.

[5] J. Blinn and M. Newell. Texture and reflection in computer generated images. *Communications of the ACM*, 19:542–546, 1976.

[6] S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *SIGGRAPH 01*, pages 107–116, 2001.

[7] B. Cabral, N. Max, and R. Springmeyer. Bidirectional reflection functions from surface bump maps. In *SIGGRAPH 87*, pages 273–281, 1987.

[8] B. Cabral, M. Olano, and P. Nemec. Reflection space image based rendering. In *SIGGRAPH 99*, pages 165–170, 1999.

[9] M. Chen and J. Arvo. Simulating non-lambertian phenomena involving linearly-varying luminaires. In *Eurographics Workshop on Rendering*, pages 25–38, 2001.

[10] G. Chirikjian and A. Kyatkin. *Engineering applications of noncommutative harmonic analysis: with emphasis on rotation and motion groups*. CRC press, 2000.

[11] J. Cochran. *The analysis of linear integral equations*. McGraw-Hill, 1972.

[12] M. F. Cohen and J. R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press, 1993.

[13] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH 96*, pages 303–312, 1996.

[14] K. Dana, B. Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.

[15] P. Debevec, T. Hawkins, C. Tchou, H.P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *SIGGRAPH 00*, pages 145–156, 2000.

[16] R. Dror, E. Adelson, and A. Willsky. Estimating surface reflectance properties from images under unknown illumination. In *SPIE Photonics West: Human Vision and Electronic Imaging VI*, pages 231–242, 2001.

[17] M. D'Zmura. *Computational Models of Visual Processing*, chapter Shading Ambiguity: Reflectance and Illumination, pages 187–207. MIT Press, 1991.

[18] R. Epstein, P.W. Hallinan, and A. Yuille. 5 plus or minus 2 eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *IEEE Workshop on Physics-Based Modeling in Computer Vision*, pages 108–116, 1995.

[19] W. Fulton and J. Harris. *Representation Theory: A first course*. Springer-Verlag, 1991.

[20] A. Gershun. The light field. *Journal of Mathematics and Physics*, XVIII:51–151, 1939. Translated by P. Moon and G. Timoshenko.

[21] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *SIGGRAPH 96*, pages 43–54, 1996.

[22] N. Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics & Applications*, 6(11):21–29, 1986.

[23] G. Greger, P. Shirley, P. Hubbard, and D. Greenberg. The irradiance volume. *IEEE Computer Graphics & Applications*, 18(2):32–43, 1998.

[24] H. Groemer. *Geometric applications of Fourier series and spherical harmonics*. Cambridge University Press, 1996.

[25] Numerical Algorithms Group. *NAG C Library Manual, Mark 5*. 1999.

[26] Z. Hakura. *Inverse Rendering Methods for Hardware-Accelerated Display of Parameterized Image Spaces*. PhD thesis, Stanford University, Oct 2001.

[27] Z. Hakura, J. Snyder, and J. Lengyel. Parameterized animation compression. In *EuroGraphics Rendering Workshop 00*, pages 101–112, 2000.

[28] Z. Hakura, J. Snyder, and J. Lengyel. Parameterized environment maps. In *ACM symposium on interactive 3D graphics*, pages 203–208, 2001.

[29] M. Halle. Multiple viewpoint rendering. In *SIGGRAPH 98*, pages 243–254, 1998.

[30] P.W. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *CVPR 94*, pages 995–999, 1994.

[31] K. Ikeuchi and K. Sato. Determining reflectance properties of an object using range and brightness images. *PAMI*, 13(11):1139–1153, 1991.

[32] T. Inui, Y. Tanabe, and Y. Onodera. *Group theory and its applications in physics*. Springer Verlag, 1990.

[33] X. Tong J. Chai and H. Shum. Plenoptic sampling. In *SIGGRAPH 00*, pages 307–318, 2000.

[34] J.D. Jackson. *Classical Electrodynamics*. John Wiley, 1975.

[35] J. Kajiya and T. Kay. Rendering fur with three dimensional textures. In *SIGGRAPH 89*, pages 271–280, 1989.

[36] K. F. Karner, H. Mayer, and M. Gervautz. An image based measurement system for anisotropic reflection. *Computer Graphics Forum*, 15(3):119–128, 1996.

[37] J. Kautz and M. McCool. Interactive rendering with arbitrary BRDFs using separable approximations. In *EGRW 99*, pages 247–260, 1999.

[38] J. Kautz and M. McCool. Approximation of glossy reflection with prefiltered environment maps. In *Graphics Interface*, pages 119–126, 2000.

[39] J. Kautz, P. Vázquez, W. Heidrich, and H.P. Seidel. A unified approach to prefiltered environment maps. In *EuroGraphics Rendering Workshop 00*, pages 185–196, 2000.

[40] J. Kautz, P. Vázquez, W. Heidrich, and H.P. Seidel. A unified approach to prefiltered environment maps. In *EGRW 00*, pages 185–196, 2000.

[41] G. Kay and T. Caelli. Inverting an illumination model from range and intensity maps. *CVGIP-Image Understanding*, 59(2):183–201, 1994.

[42] G.J. Klinker, S.A. Shafer, and T. Kanade. The measurement of highlights in color images. *IJCV*, 2(1):7–32, 1988.

[43] J. Koenderink and A. van Doorn. Phenomenological description of bidirectional surface reflection. *JOSA A*, 15(11):2903–2912, 1998.

[44] E. Lafortune, S. Foo, K. Torrance, and D. Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH 97*, pages 117–126, 1997.

[45] P. Lalonde and A. Fournier. Filtered local shading in the wavelet domain. In *EGRW 97*, pages 163–174, 1997.

[46] E. Land and J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.

[47] L. Latta and A. Kolb. Homomorphic factorization of brdf-based lighting computation. In *SIGGRAPH 02*, pages 509–516, 2002.

[48] M. Levoy and P. Hanrahan. Light field rendering. In *SIGGRAPH 96*, pages 31–42, 1996.

[49] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital michelangelo project: 3D scanning of large statues. In *SIGGRAPH 00*, pages 131–144, 2000.

[50] R. Love. *Surface Reflection Model Estimation from Naturally Illuminated Image Sequences*. PhD thesis, Leeds, 1997.

[51] R. Lu, J.J. Koenderink, and A.M.L. Kappers. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics*, 37(25):5974–5984, 1998.

[52] T. MacRobert. *Spherical harmonics; an elementary treatise on harmonic functions, with applications*. Dover Publications, 1948.

[53] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *SIGGRAPH 01*, pages 519–528, 2001.

[54] S.R. Marschner and D.P. Greenberg. Inverse lighting for photography. In *Fifth Color Imaging Conference*, pages 262–265, 1997.

[55] S.R. Marschner, S.H. Westin, E.P.F. Lafortune, and K.E. Torrance. Image-Based BRDF measurement. *Applied Optics*, 39(16):2592–2600, 2000.

[56] R. McCluney. *Introduction to Radiometry and Photometry*. Artech House, 1994.

[57] M. McCool, J. Ang, and A. Ahmad. Homomorphic factorization of BRDFs for high-performance rendering. In *SIGGRAPH 01*, pages 171–178, 2001.

[58] N. McCormick. Inverse radiative transfer problems: a review. *Nuclear Science and Engineering*, 112:185–198, 1992.

[59] G. Miller and C. Hoffman. Illumination and reflection maps: Simulated objects in simulated and real environments. *SIGGRAPH 84 Advanced Computer Graphics Animation seminar notes*, 1984.

[60] G. Miller, S. Rubin, and D. Poncelon. Lazy decompression of surface light fields for precomputed global illumination. In *EGRW 98*, pages 281–292, 1998.

[61] M. Mohlenkamp. A fast transform for spherical harmonics. *The Journal of Fourier Analysis and Applications*, 5(2/3):159–184, 1999.

[62] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. *Geometric Considerations and Nomenclature for Reflectance*. National Bureau of Standards (US), 1977.

[63] J. Nimeroff, E. Simoncelli, and J. Dorsey. Efficient re-rendering of naturally illuminated environments. In *Fifth Eurographics Workshop on Rendering*, pages 359–373, June 1994.

[64] K. Nishino, Y. Sato, and K. Ikeuchi. Eigen-texture method: Appearance compression based on 3D model. In *CVPR 99*, pages 618–624, 1999.

[65] B. Oh, M. Chen, J. Dorsey, and F. Durand. Image-based modeling and photo editing. In *SIGGRAPH 01*, pages 433–442, 2001.

[66] J. Kautz P. Sloan and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *SIGGRAPH 02*, pages 527–536, 2002.

[67] R.W. Preisendorfer. *Hydrologic Optics*. US Dept Commerce, 1976.

[68] K. Proudfoot, W. Mark, S. Tzvetkov, and P. Hanrahan. A real-time procedural shading system for programmable graphics hardware. In *SIGGRAPH 01*, pages 159–170, 2001.

[69] R. Ramamoorthi. Analytic PCA construction for theoretical analysis of lighting variability, including attached shadows, in a single image of a convex lambertian object. In *CVPR workshop on identifying objects across variations in lighting: psychophysics and computation*, pages 48–55, 2001.

[70] R. Ramamoorthi and P. Hanrahan. Analysis of planar light fields from homogeneous convex curved surfaces under distant illumination. In *SPIE Photonics West: Human Vision and Electronic Imaging VI*, pages 185–198, 2001.

[71] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *SIGGRAPH 01*, pages 497–500, 2001.

[72] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *JOSA A*, 18(10):2448–2459, 2001.

[73] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, pages 117–128, 2001.

[74] S. Rusinkiewicz. A new change of variables for efficient BRDF representation. In *EGRW 98*, pages 11–22, 1998.

[75] I. Sato, Y. Sato, and K. Ikeuchi. Illumination distribution from brightness in shadows: adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. In *ICCV 99*, pages 875 – 882, 1999.

[76] Y. Sato and K. Ikeuchi. Reflectance analysis under solar illumination. Technical Report CMU-CS-94-221, CMU, 1994.

[77] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *SIGGRAPH 97*, pages 379–388, 1997.

[78] P. Schröder and W. Sweldens. Spherical wavelets: Texture processing. In *EGRW 95*, pages 252–263, 1995.

[79] F. Sillion, J. Arvo, S. Westin, and D. Greenberg. A global illumination solution for general reflectance distributions. In *SIGGRAPH 91*, pages 187–196, 1991.

[80] C. Soler and F. Sillion. Fast calculation of soft shadow textures using convolution. In *SIGGRAPH 98*, pages 321–332, 1998.

[81] S.R.Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell, 1998.

[82] P. Teo, E. Simoncelli, and D. Heeger. Efficient linear rerendering for interactive lighting design. Technical Report STAN-CS-TN-97-60, Stanford, 1997.

[83] S. Tominaga and N. Tanaka. Estimating reflection parameters from a single color image. *IEEE Computer Graphics & Applications*, 20(5):58–66, 2000.

[84] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967.

[85] G. Ward and P. Heckbert. Irradiance gradients. In *EGRW 92*, pages 85–98, 1992.

[86] S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92*, pages 255–264, 1992.

[87] A. Wilkie, R. Tobler, and W. Purgathofer. Orientation lightmaps for photon radiosity in complex environments. In *CGI 00*, pages 279–286, 2000.

[88] D. Wood, D. Azuma, K. Aldinger, B. Curless, T. Duchamp, D. Salesin, and W. Stuetzle. Surface light fields for 3D photography. In *SIGGRAPH 00*, pages 287–296, 2000.

[89] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *SIGGRAPH 99*, pages 215–224, 1999.

[90] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *SIGGRAPH 98*, pages 207–218, 1998.

[91] A. Yuille, D. Snow, R. Epstein, and P. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using SVD and integrability. *IJCV*, 35(3):203–222, 1999.

[92] L. Zhang, G. Dugas-Phocion, J. Samson, and S. Seitz. Single view modeling of free-form scenes. In *CVPR 01*, pages I–990–I–997, 2001.

[93] Z. Zhang. A flexible new technique for camera calibration. Technical Report MSR-TR-98-71, Microsoft Research, 1998.