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

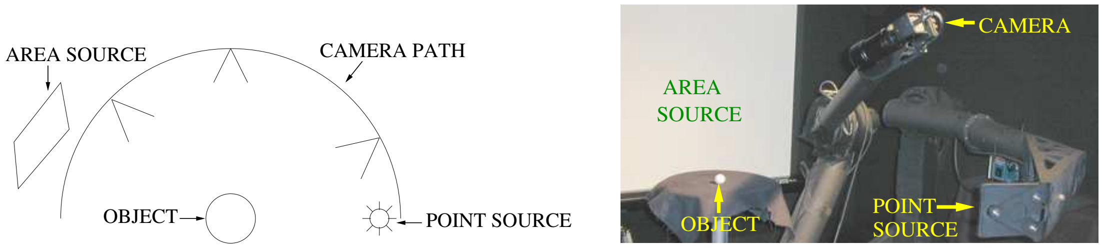

Figure 6.1: **Left:** *Schematic of experimental setup* **Right:** *Photograph*

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.

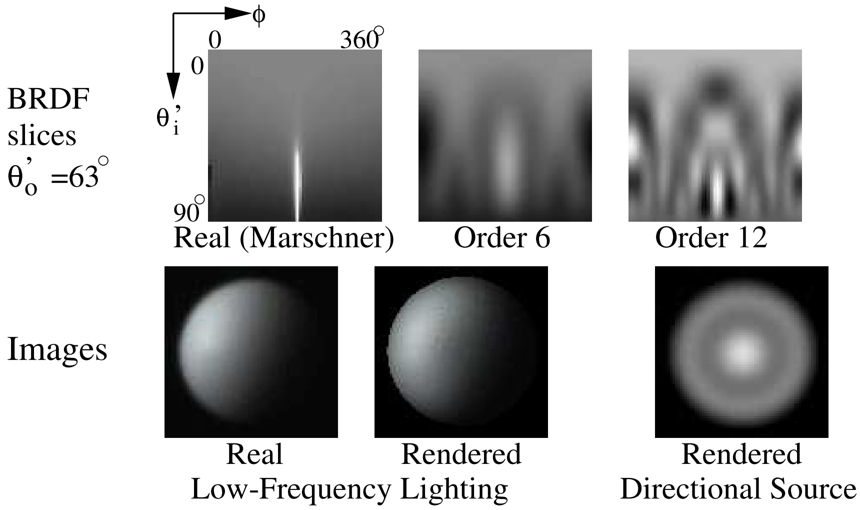

Figure 6.2: *Direct recovery of BRDF coefficients.* **Top:** *Slices of the BRDF transfer function of a teflon sphere for fixed exitant angle of* $63°$. ϕ'_o *varies linearly from* $0°$ *to* $90°$ *from top to bottom, and* $\mid \phi'_o - \phi_i \mid$ *linearly from* $0°$ *to* $360°$ *from left to right. 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 $B_{s,fast} = 0$,

$$\begin{aligned} B_{s,slow} &= K_s F(\mu, \theta'_o) L_{slow}(\vec{R}) \;\; = B - B^1_d(L^1_{lm}) \\ L_{slow}(\vec{R}) &= \frac{B - B^1_d(L^1_{lm})}{K_s F(\mu, \theta'_o)}. \end{aligned} \tag{6.17}$$

   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

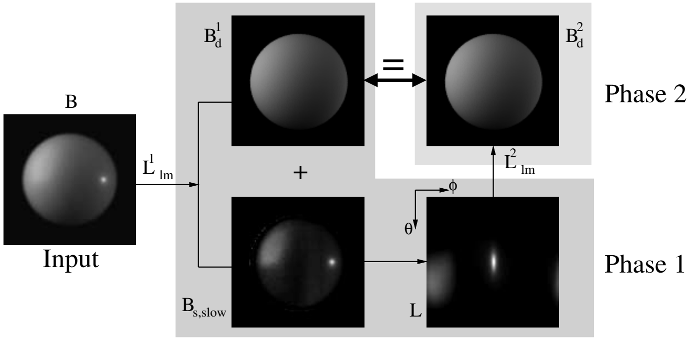

Figure 6.4: *Estimation of dual lighting representation. In phase 1, we use frequency-space parameters L^1_{lm} to compute diffuse component B^1_d. This is subtracted from the input image, leaving the specular component, from which the angular-space lighting is found. In phase 2, we compute coefficients 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,

$$
\begin{aligned}
L_{slow}(\theta_R, \phi_R) &= U(\theta_R, \phi_R) - \sum_{l=0}^{2} \sum_{m=-l}^{l} W_{lm}(\theta_R, \phi_R) L_{lm}^1 \\
U(\theta_R, \phi_R) &= \frac{B}{K_s F(\mu, \theta_o')} \\
W_{lm}(\theta_R, \phi_R) &= \frac{K_d \hat{\rho}_l Y_{lm}(\alpha, \beta)}{K_s F(\mu, \theta_o')}.
\end{aligned}
\tag{6.18}
$$

Phase 2 to compute the lighting coefficients can also be written as a linear expression in terms of all the (discretized) reflected directions,

$$
L_{lm}^2 = \frac{2\pi^2}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sin\theta_i L_{slow}(\theta_i, \phi_j) Y_{lm}^*(\theta_i, \phi_j). \tag{6.19}
$$

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_{lm}^1, simply by plugging in the formula for L_{slow}. We now obtain

$$
L_{lm}^2 = \frac{2\pi^2}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sin\theta_i \left(U(\theta_i, \phi_j) - \sum_{l'=0}^{2} \sum_{m'=-l'}^{l'} W_{l'm'}(\theta_i, \phi_j) L_{l'm'}^1 \right) Y_{lm}^*(\theta_i, \phi_j). \tag{6.20}
$$

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.

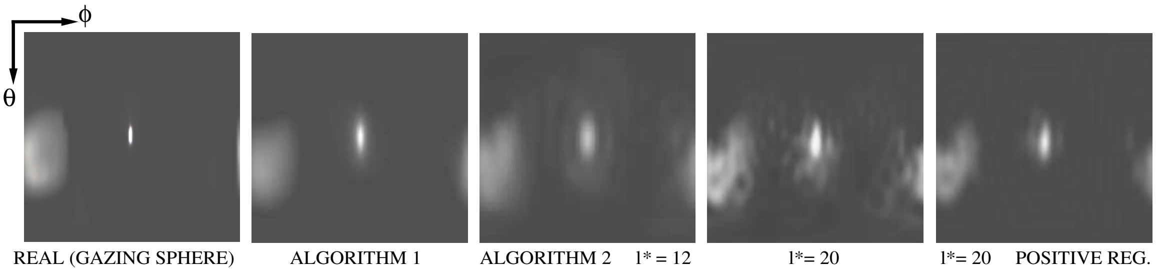

Figure 6.5: *Comparison of inverse lighting methods. From left to right, real lighting (from a gazing sphere), recovered illumination by direct estimation of spherical harmonic coefficients with* $l^* = 12$ *and* $l^* = 20$*, and estimation of dual angular and frequency-space lighting model. To make 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,

$$L_{cen} \approx \frac{L_{total}}{4\pi\sigma^2}. \qquad (6.25)$$

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

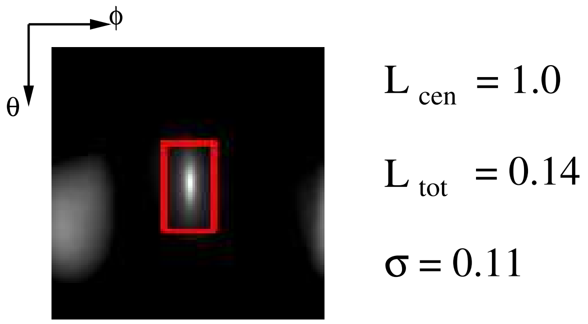

Figure 6.6: *Determining surface roughness parameter σ. We manually specify (red box) the region corresponding to the directional source in a visualization of the lighting. The algorithm then determines L_{cen}, the intensity at the center (brightest point), L_{tot}, the total energy integrated over 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

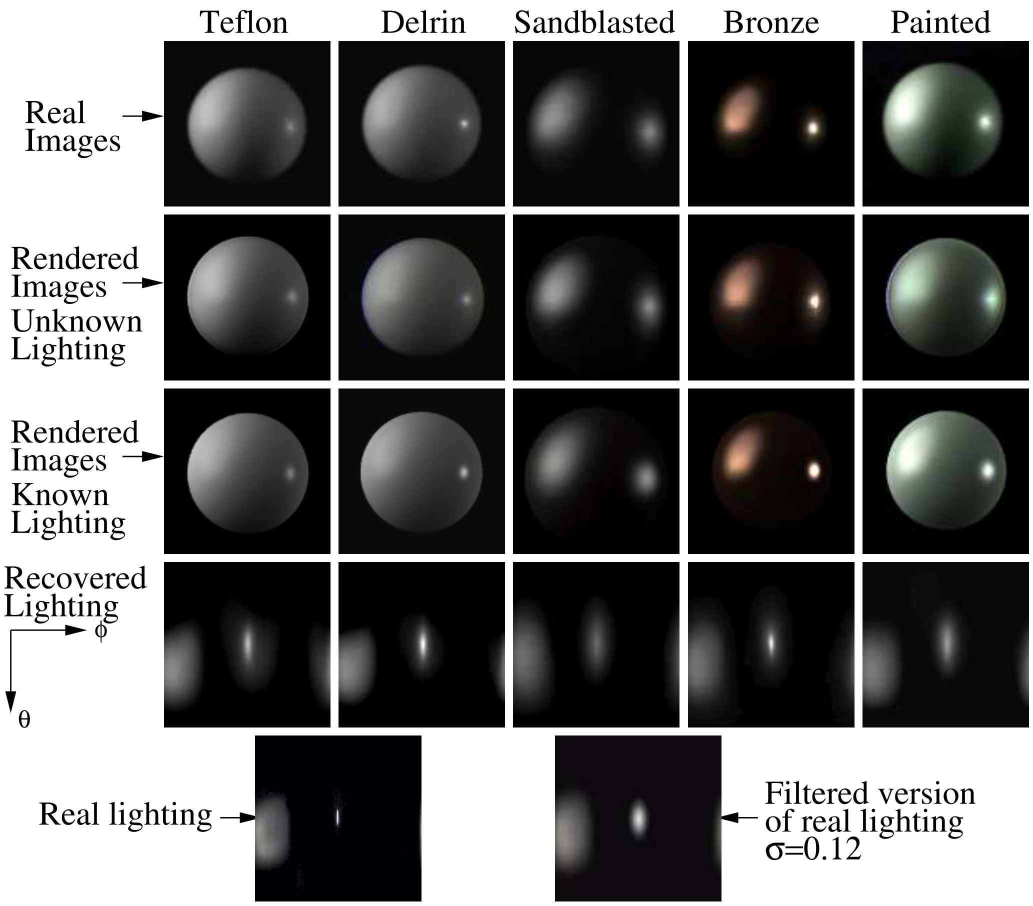

Figure 6.8: *Spheres rendered using BRDFs estimated under known (section 6.4.2) and unknown (section 6.4.4) lighting. The algorithm in section 6.4.4 also recovers the lighting. Since there is an unknown global scale, we scale the recovered lighting distributions in order to compare them. The recovered 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.

Real Photograph Rendered Image

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

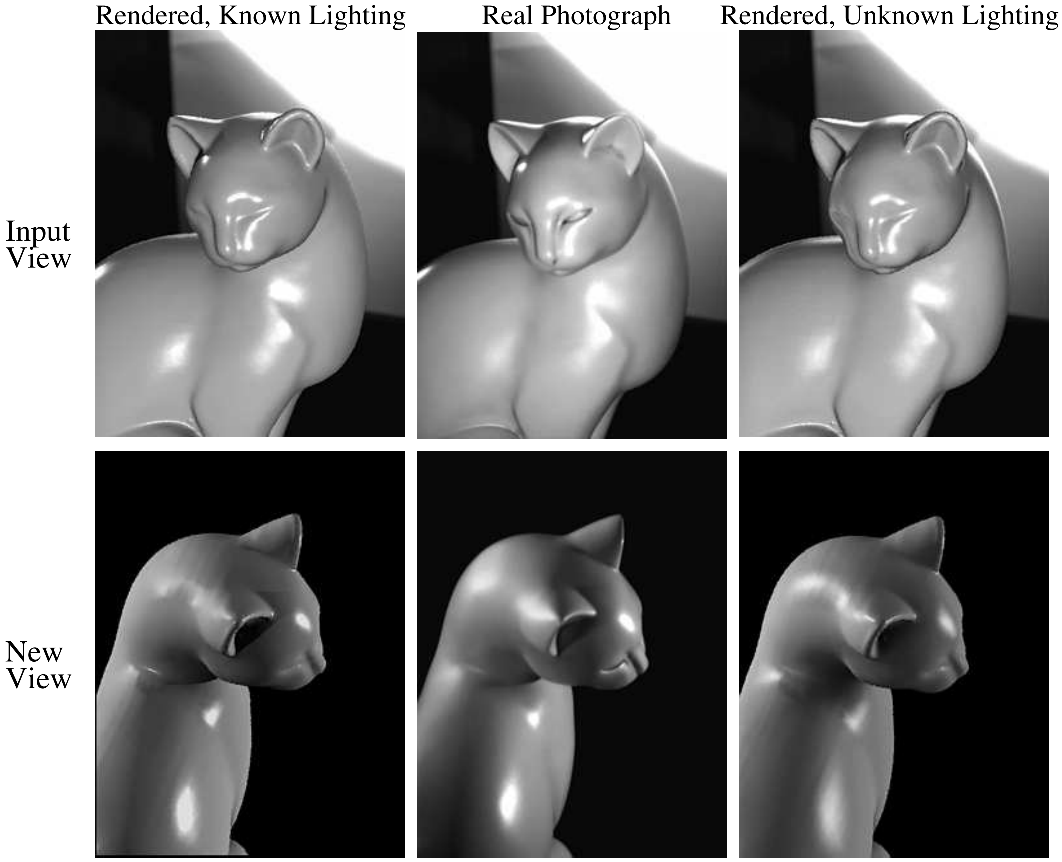

Figure 6.10: *Comparison of photographs (middle column) to images rendered using BRDFs recovered under known lighting (left column), and using BRDFs (and lighting) estimated under unknown lighting (right column). The top row is one of the 3 input views. The bottom row is a new view, not used for BRDF 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

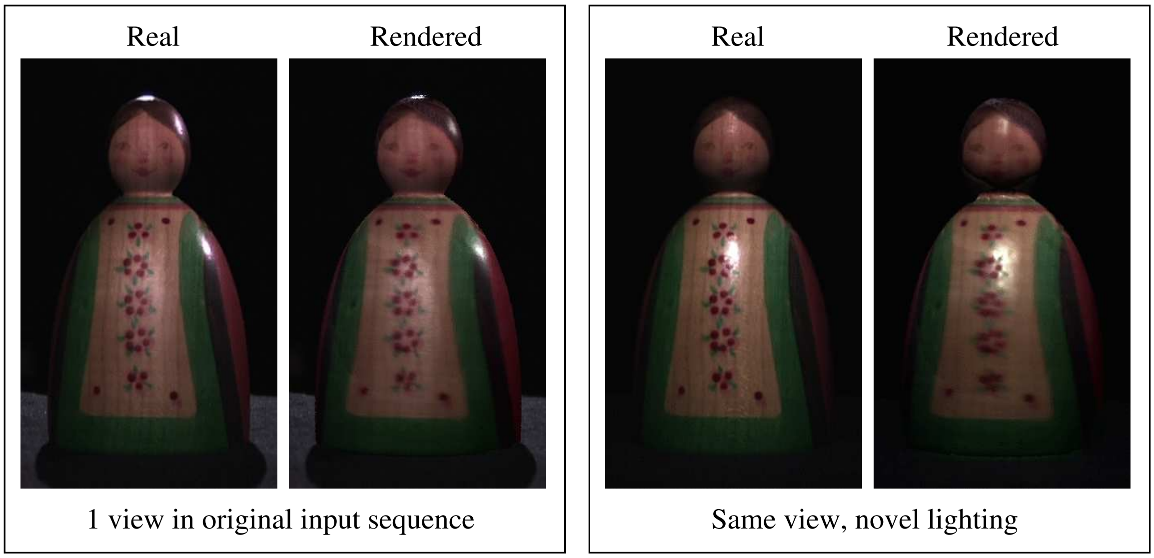

Figure 6.12: *Recovering textured BRDFs under complex lighting. The rendered images closely resemble the real photographs, even under novel lighting.*

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.