

A Signal-Processing Framework for Inverse Rendering

Ravi Ramamoorthi

Pat Hanrahan

Stanford University *

Abstract

Realism in computer-generated images requires accurate input models for lighting, textures and BRDFs. One of the best ways of obtaining high-quality data is through measurements of scene attributes from real photographs by *inverse rendering*. However, inverse rendering methods have been largely limited to settings with highly controlled lighting. One of the reasons for this is the lack of a coherent mathematical framework for inverse rendering under general illumination conditions. Our main contribution is the introduction of a signal-processing framework which describes the reflected light field as a convolution of the lighting and BRDF, and expresses it mathematically as a product of spherical harmonic coefficients of the BRDF and the lighting. Inverse rendering can then be viewed as deconvolution. We apply this theory to a variety of problems in inverse rendering, explaining a number of previous empirical results. We will show why certain problems are ill-posed or numerically ill-conditioned, and why other problems are more amenable to solution. The theory developed here also leads to new practical representations and algorithms. For instance, we present a method to *factor* the lighting and BRDF from a small number of views, i.e. to estimate both simultaneously when neither is known.

CR Categories: I.3.7 [Computer Graphics]: Realism; I.4.8 [Computer Vision]: Scene Analysis—Photometry

Keywords: Signal Processing, Spherical Harmonics, Inverse Rendering, Radiance, Light Field, Irradiance, Illumination, BRDF

1 Introduction

To create a realistic computer-generated image, we need both an accurate, physically-based rendering algorithm and a detailed model of the scene including light sources and objects specified by their geometry and material properties—texture and reflectance (BRDF). There has been substantial progress in the development of rendering algorithms, and nowadays, realism is often limited by the quality of input models. As a result, *image-based rendering* is becoming widespread. 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. For example, an artist can change independently the material properties or the lighting.

Inverse rendering methods such as those of Debevec et al. [6], Marschner et al. [21], and Sato et al. [32], have produced high




Figure 1: **Left:** Real 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.

quality measurements. 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 [39], Sato and Ikeuchi [31] and Love [17], are designed specifically for natural illumination, and assume a simple parametric model for skylight. Previous methods have also usually tried to recover only one of the unknowns—texture, BRDF or lighting. The usefulness of inverse rendering would be greatly enhanced if it could be applied under general uncontrolled lighting, and if we could simultaneously estimate more than one unknown. For instance, if we could recover both the lighting and BRDF, we could determine BRDFs under unknown illumination. One reason there has been relatively little work in these areas 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.

Our main contribution is a theoretical framework for analyzing the reflected light field from a curved convex homogeneous surface under distant illumination. We believe this framework provides a solid mathematical foundation for many areas of graphics. With respect to inverse rendering, we obtain the following results:

Reflection as Convolution: It has been observed qualitatively by Miller and Hoffman [23], Cabral et al. [3], Bastos et al. [2] and others that the reflection operator behaves like a convolution in the angular domain. We formalize these notions mathematically. 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 illumination using the BRDF. We believe this is a useful way of analyzing many computer graphics problems. In particular, inverse rendering can be viewed as *deconvolution*.

Well-posedness and Conditioning of Inverse Problems: Inverse problems can be ill-posed—there may be no solutions or several solutions. They are also often numerically ill-conditioned, i.e. extremely sensitive to noisy input data. From our theory, we are able to analyze the well-posedness and conditioning of a number of inverse problems, explaining many previous empirical observations. This analysis can serve as a guideline for future research.

New Practical Representations and Algorithms: Insights from the theory lead to the derivation of a simple practical representation, which can be used to estimate BRDFs under complex lighting. The

* (ravir.hanrahan)@graphics.stanford.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

theory also leads to novel frequency space and hybrid angular and frequency space methods for inverse problems, including two new algorithms for estimating the lighting, and an algorithm for simultaneously determining the lighting and BRDF. Therefore, we can recover BRDFs under general, unknown lighting conditions.

2 Previous Work

To describe previous work, we will introduce a taxonomy based on how many of the three quantities—lighting, BRDF and texture—are unknown. To motivate the taxonomy, we first write a simplified version of the reflection equation, omitting visibility.

$$B(\mathbf{x}, \vec{\omega}_o) = \int_{\Omega_i} T(\mathbf{x}) \rho(\vec{\omega}_i, \vec{\omega}_o) L(\mathbf{x}, \vec{\omega}_i) (\vec{\omega}_i \cdot \vec{n}) d\omega_i \quad (1)$$

Here, B is the reflected light field, expressed as a function of the surface position \mathbf{x} and outgoing direction $\vec{\omega}_o$. The normal vector is \vec{n} . For simplicity, we assume that a single texture T modulates the BRDF. In practice, we would use separate textures for the diffuse and specular components of the BRDF.

The integrand is a product of terms—the texture $T(\mathbf{x})$, the BRDF $\rho(\vec{\omega}_i, \vec{\omega}_o)$, and the lighting $L(\mathbf{x}, \vec{\omega}_i)$. Inverse rendering, assuming known geometry, involves inverting the integral 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.

One Unknown

1. Unknown Texture: Previous methods have recovered the diffuse texture of 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 [34] and Levoy et al. [16].

2. Unknown BRDF: The BRDF [24] is a fundamental intrinsic surface property. Active measurement methods, known as gonio-reflectometry, involving a single point source and a single observation at a time, have been developed. Improvements are suggested by Ward [37] and Karner et al. [12]. More recently, image-based BRDF measurement methods have been proposed by Lu et al. [18] and Marschner et al. [21]. 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 [37] or Torrance and Sparrow [36]. 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. Love [17] estimates parametric BRDFs under natural illumination, assuming a low-parameter model for skylight and sunlight. Dror et al. [7] classify the surface reflectance as one of a small number of predetermined BRDFs, making use of assumed statistical characteristics of natural lighting. However, the *inverse BRDF* problem has not been solved for general illumination.

3. Unknown Lighting: A common solution is to use a mirrored ball, as done by Miller and Hoffman [23]. Marschner and Greenberg [20] find the lighting from a Lambertian surface. D’Zmura [8] proposes, but does not demonstrate, estimating spherical harmonic coefficients. For Lambertian objects, we [29] have shown how to recover the first 9 spherical harmonics. Previous work has not estimated the lighting from curved surfaces with general BRDFs.

Two Unknowns

4. Factorization—Unknown Lighting and BRDF: BRDF estimation methods have been proposed by Ikeuchi and Sato [10] and Tominaga and Tanaka [35] for the special case when the lighting consists of a single source of unknown direction. However, these methods cannot simultaneously recover a complex lighting distribution and the object BRDF. One of the main practical contributions of this paper is a solution to this problem for curved surfaces, allowing us to estimate BRDFs under general unknown illumination, while also determining the lighting. The closest previous work is that of Sato et al. [30] who use shadows to estimate the illumination distribution and the surface reflectance properties. We

extend this work by not requiring shadow information, and presenting improved methods for estimating the illumination.

5. Factorization—Unknown Texture and BRDF: This corresponds to recovering *textured*, or spatially-varying BRDFs. Sato et al. [32] rotate an object on a turntable, using a single point source, to recover BRDF parameters and texture. Yu et al. [38] recover a texture only for the diffuse BRDF component, but account for interreflections. Using a large number of images obtained by moving a point source around a sphere surrounding the subject, Debevec et al. [6] acquire the reflectance field of a human face, and recover parameters of a microfacet BRDF model for each surface location. Sato and Ikeuchi [31] and Yu and Malik [39] 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 these methods recover only diffuse textures; constant values, or relatively low-resolution textures, are used for the specular parameters. A notable exception is the work of Dana et al. [5] who generalize BRDFs to a 6D bi-directional texture function (BTF).

6. Factorization—Unknown Lighting and Texture: We have shown [29] that a distant illumination field can cause only low frequency 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 [15], 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 characteristics.

General Case: Three Unknowns

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 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 [26] and neutron scattering. See McCormick [22] for a review.

In previous theoretical work, D’Zmura [8] has analyzed reflection as a linear operator in terms of spherical harmonics, and discussed some resulting perceptual ambiguities between reflectance and illumination. In computer graphics, Cabral et al. [3] first demonstrated the use of spherical harmonics to represent BRDFs. We extend these methods by explicitly deriving the frequency-space reflection equation (i.e. convolution formula), and by providing quantitative results for various special cases. We have earlier reported on theoretical results for planar or *flatland* light fields [27], and for determining the lighting from a Lambertian surface [29]. For the Lambertian case, similar results have been derived independently by Basri and Jacobs [1] in simultaneous work on face recognition. This paper extends these previous results to the general 3D case with arbitrary isotropic BRDFs, and applies the theory to developing new practical inverse-rendering algorithms.

3 Assumptions

The input to our algorithms consists of object geometry and photographs from a number of different locations, 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. Our assumptions are:

B	Reflected radiance
$B_{lm pq}$	Coefficients of basis-function expansion of B
L	Incoming radiance
L_{lm}	Coefficients of spherical-harmonic expansion of L
ρ	Surface BRDF
$\hat{\rho}$	BRDF multiplied by cosine of incident angle
$\hat{\rho}_{tpq}$	Coefficients of spherical-harmonic expansion of $\hat{\rho}$
θ'_i, θ_i	Incident elevation angle in <i>local, global</i> coordinates
ϕ'_i, ϕ_i	Incident azimuthal angle in <i>local, global</i> coordinates
θ'_o, θ_o	Outgoing elevation angle in <i>local, global</i> coordinates
ϕ'_o, ϕ_o	Outgoing azimuthal angle in <i>local, global</i> coordinates
Ω'_i, Ω_i	Hemisphere of integration in <i>local, global</i> coordinates
\mathbf{x}	Surface position
α	Surface normal parameterization—elevation angle
β	Surface normal parameterization—azimuthal angle
$R_{\alpha, \beta}$	Rotation operator for surface normal (α, β)
$D^l_{mm'}$	Matrix related to Rotation Group $SO(3)$
Y_{lm}	Spherical Harmonic basis function
Y_{lm}^*	Complex Conjugate of Spherical Harmonic
Λ_l	Normalization constant, $\sqrt{4\pi/(2l+1)}$
I	$\sqrt{-1}$

Figure 2: Notation

Known Geometry: We use a laser range scanner and a volumetric merging algorithm [4] to obtain object geometry. By assuming known geometry, we can focus on lighting and material properties.

Curved Objects: Our theoretical analysis requires curved surfaces, and assumes knowledge of the entire 4D reflected light field, corresponding to the hemisphere of outgoing directions for all surface orientations. However, our practical algorithms will require only a small number of photographs.

Distant Illumination: The illumination field will be assumed to be homogeneous, i.e. generated by distant sources, allowing us to use the same lighting function regardless of surface location. We treat the lighting as a general function of the incident angle.

Isotropic BRDFs: We will consider only surfaces having isotropic BRDFs. The BRDF will therefore be a function of only 3 variables, instead of 4, i.e. 3D instead of 4D.

No Interreflection: For concave surfaces, interreflection will be ignored. Also, shadowing is not considered in our theoretical analysis, which is limited to convex surfaces. However, we will account for shadowing in our practical algorithms, where necessary.

4 Theory of Reflection as Convolution

This section presents a signal-processing framework wherein reflection can be viewed as convolution, and inverse rendering as deconvolution. First, we introduce some preliminaries, defining the notation and deriving a version of the reflection equation. We then expand the lighting, BRDF and reflected light field in spherical harmonics to derive a simple equation in terms of spherical harmonic coefficients. The next section explores implications of this result.

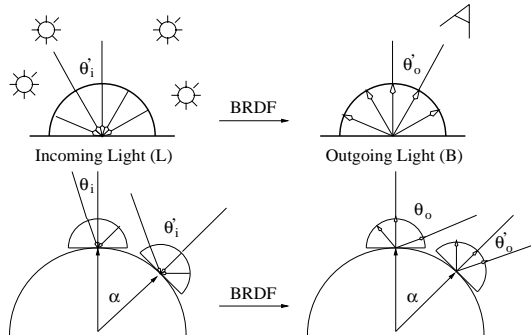


Figure 3: Schematic of reflection. On top, we show the situation with respect to the local surface. The BRDF maps the incoming light distribution L to an outgoing light distribution B . The bottom figure shows how the rotation α affects the situation. Different orientations of the surface correspond to rotations of the upper hemisphere and BRDF, with global directions (θ_i, ϕ_i) corresponding to local directions (θ'_i, ϕ'_i) .

4.1 Preliminaries

For the purposes of theoretical analysis, we assume curved convex isotropic surfaces. We also assume homogeneous objects, i.e. untextured surfaces, with the same BRDF everywhere. We parameterize the surface by the spherical coordinates of the normal vector (α, β) , using the standard convention that $(\alpha, \beta) = (0, 0)$ corresponds to the *north pole* or +Z axis. Notation used in this section is listed in figure 2, and a diagram is in figure 3. 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. These two coordinate systems are related simply by a rotation, to be defined shortly.

Reflection Equation: We modify equation 1 based on our assumptions, dropping the texturing term, and using the surface normal (α, β) instead of the position \mathbf{x} to parameterize B . Since $L(\mathbf{x}, \vec{\omega}_i)$ is assumed to be independent of \mathbf{x} , we write it as $L(\theta_i, \phi_i)$. Finally, $(\vec{\omega}_i \cdot \vec{n})$ can be written simply as $\cos \theta'_i$, the cosine of the incident angle in local coordinates.

$$B(\alpha, \beta, \theta'_o, \phi'_o) = \int_{\Omega'_i} L(\theta_i, \phi_i) \rho(\theta'_i, \phi'_i, \theta'_o, \phi'_o) \cos \theta'_i d\omega'_i \quad (2)$$

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. The reflected radiance B can be expressed conveniently in either local or global coordinates; we have used local coordinates to match the BRDF. Similarly, integration can be conveniently done over either local or global coordinates, but the upper hemisphere is easier to express in local coordinates.

We now define a transfer¹ function $\hat{\rho} = \rho \cos \theta'_i$ in order to absorb the cosine term. With this modification, equation 2 becomes

$$B(\alpha, \beta, \theta'_o, \phi'_o) = \int_{\Omega'_i} L(\theta_i, \phi_i) \hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) d\omega'_i \quad (3)$$

Rotations—Converting Local and Global coordinates:

Local and global coordinates are related by a rotation corresponding to the surface normal (α, β) . The *north pole* in local coordinates, $(0', 0')$ is the surface normal. The corresponding global coordinates are clearly (α, β) . We define $R_{\alpha, \beta}$ as a rotation operator² on column vectors that rotates (θ'_i, ϕ'_i) into global coordinates, and is given by $R_{\alpha, \beta} = R_z(\beta)R_y(\alpha)$ where R_z is a rotation about the Z axis and R_y a rotation about the Y axis.

$$\begin{aligned} (\theta_i, \phi_i) &= R_z(\beta)R_y(\alpha)(\theta'_i, \phi'_i) = R_{\alpha, \beta}(\theta'_i, \phi'_i) \\ (\theta'_i, \phi'_i) &= R_y(-\alpha)R_z(-\beta)(\theta_i, \phi_i) = R_{\alpha, \beta}^{-1}(\theta_i, \phi_i) \end{aligned}$$

We can now write the dependence on incident angle in equation 3 entirely in global coordinates, or entirely in local coordinates.

$$\begin{aligned} B(\alpha, \beta, \theta'_o, \phi'_o) &= \int_{\Omega_i} L(\theta_i, \phi_i) \hat{\rho}(R_{\alpha, \beta}^{-1}(\theta_i, \phi_i), \theta'_o, \phi'_o) d\omega_i \quad (4) \\ &= \int_{\Omega'_i} L(R_{\alpha, \beta}(\theta'_i, \phi'_i)) \hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) d\omega'_i \quad (5) \end{aligned}$$

¹If we want the transfer function to be reciprocal, i.e. symmetric with respect to incident and outgoing angles, we may multiply both the transfer function and the reflected light field by $\cos \theta'_o$. See equation 13.

²For anisotropic surfaces, we need an initial rotation about Z to set the local tangent frame. We would then have rotations about Z, Y and Z—the familiar Euler-Angle parameterization. Since we are dealing with isotropic surfaces, we have ignored this initial Z rotation, which has no physical significance. It is not difficult to derive the theory for the more general anisotropic case.

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

$$(f \otimes g)(a) = \int_t f(t)g(T_a(t)) dt$$

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 4 and 5 represent rotational convolutions. Equation 4 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 5 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.

4.2 Spherical Harmonic Representation

For the translational case, the well-known frequency-space convolution formula is given in terms of Fourier transforms. For a general operator, an analogous formula can be obtained in terms of group representations and the associated basis functions. For translations, these basis functions are sines and cosines—the familiar Fourier basis. For rotations, the corresponding basis functions are spherical harmonics, and we now proceed to derive the frequency-space rotational convolution formula in terms of spherical harmonics.

Inui et al. [11] is a good reference for background on spherical harmonics and their relationship to rotations. Our use of spherical harmonics to represent the lighting is similar in some respects to previous methods [25] that use steerable linear basis functions. Spherical harmonics, as well as the closely related Zernike Polynomials, have been used before to represent BRDFs [3, 14, 33].

Spherical harmonics are the analog on the sphere to the Fourier basis on the line or circle. The spherical harmonic Y_{lm} is given by

$$N_{lm} = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}}$$

$$Y_{lm}(\theta, \phi) = N_{lm} P_l^m(\cos \theta) e^{Im\phi}$$

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_l^m . The indices obey $l \geq 0$ and $-l \leq m \leq l$.

The rotation formula for spherical harmonics is

$$Y_{lm}(R_{\alpha, \beta}(\theta'_i, \phi'_i)) = \sum_{m'=-l}^l D_{mm'}^l(\alpha) e^{Im\beta} Y_{lm'}(\theta'_i, \phi'_i) \quad (6)$$

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 is a matrix that tells us how a spherical harmonic transforms under rotation about the y -axis, i.e. how to rewrite a rotated spherical harmonic as a linear combination of all the spherical harmonics of the same order.

We begin by expanding the lighting in global coordinates.

$$L(\theta_i, \phi_i) = \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm} Y_{lm}(\theta_i, \phi_i) \quad (7)$$

Here, the coefficients L_{lm} can be computed in the standard way by integrating against the complex conjugate Y_{lm}^*

$$L_{lm} = \int_{\theta_i=0}^{\pi} \int_{\phi_i=0}^{2\pi} L(\theta_i, \phi_i) Y_{lm}^*(\theta_i, \phi_i) \sin \theta_i d\theta_i d\phi_i$$

We now represent the transfer function $\hat{\rho} = \rho \cos(\theta'_i)$ in terms of spherical harmonics. Note that $\hat{\rho}$ is nonzero only over the upper hemisphere, i.e. when $\cos \theta'_i > 0$ and $\cos \theta'_o > 0$.

$$\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \sum_{l,m,p,q} \hat{\rho}_{lm,pq} Y_{lm}^*(\theta'_i, \phi'_i) Y_{pq}(\theta'_o, \phi'_o)$$

We are interested in isotropic BRDFs, which depend only on $|\phi'_o - \phi'_i|$. This implies that the BRDF is invariant with respect to adding a constant angle $\Delta\phi'$ to both incident and outgoing azimuthal angles. It can be shown from the form of the spherical harmonics that this condition forces all terms to vanish unless $m = q$. The use of the complex conjugate for Y_{lm}^* in the expansion above is to make $m = q$ instead of $m = -q$. We now write

$$\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \sum_{l,p,q} \hat{\rho}_{lq,pq} Y_{lq}^*(\theta'_i, \phi'_i) Y_{pq}(\theta'_o, \phi'_o) \quad (8)$$

Furthermore, invariance of the BRDF with respect to negating both incident and outgoing azimuthal angles requires that $\hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q),p(-q)}$. Finally, we use only three indices for the BRDF.

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

To represent the reflected light field, we define a new set of orthonormal basis functions. The normalization and form of these functions are derived in the appendix. In particular, the matrix D comes from the rotation formula for spherical harmonics, equation 6. It will be convenient to first define a normalization constant.

$$\Lambda_l = \sqrt{\frac{4\pi}{2l+1}} \quad \Lambda_l^{-1} = \sqrt{\frac{2l+1}{4\pi}}$$

The new basis functions can then be written

$$C_{lmq}(\alpha, \beta, \theta'_o, \phi'_o) = \Lambda_l^{-1} D_{mq}^l(\alpha) e^{Im\beta} Y_{pq}(\theta'_o, \phi'_o) \quad (9)$$

The expansion of the reflected light field is now

$$B(\alpha, \beta, \theta'_o, \phi'_o) = \sum_{l,m,p,q} B_{lmq} C_{lmq}(\alpha, \beta, \theta'_o, \phi'_o)$$

The translational convolution theorem expresses convolution in frequency-space as a product of Fourier coefficients. For the rotational case, an analogous result is derived in the appendix, using spherical harmonics instead of complex exponentials. The *frequency-space* reflection equation (or rotational convolution formula) is a similar product of basis-function coefficients.

$$B_{lmq} = \Lambda_l L_{lm} \hat{\rho}_{lpq} \quad (10)$$

5 Implications

This section explores the implications of our results for problems in inverse rendering, and works out some special cases in detail. Our theory indicates which inverse problems are tractable, as opposed to being ill-posed or ill-conditioned. Finally, we will use the insights gained to develop a new practical representation.

5.1 General Observations

Inverse BRDF: Equation 10 can be manipulated to yield

$$\hat{\rho}_{lpq} = \Lambda_l^{-1} \frac{B_{lmq}}{L_{lm}} \quad (11)$$

We may use any index m in inverse BRDF computation. Therefore, BRDF recovery is well-posed unless the denominator vanishes for all m , i.e. all terms for some order l in the spherical harmonic expansion of the lighting vanish. 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 lighting contains high frequencies like directional sources, and is ill-conditioned for soft lighting.

Inverse Lighting: Equation 10 can also be manipulated to yield

$$L_{lm} = \Lambda_l^{-1} \frac{B_{lm pq}}{\hat{\rho}_{lpq}} \quad (12)$$

Similarly as for BRDF recovery, any p, q can be used for inverse lighting. The problem is well-posed unless the denominator $\hat{\rho}_{lpq}$ vanishes for all p, q for some l . 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 components like sharp specularities, and is ill-conditioned for diffuse surfaces.

Light Field Factorization—Lighting and BRDF: We now consider the problem of *factorizing* the light field, i.e simultaneously recovering the lighting and BRDF when both are unknown. The reflected light field is defined on a four-dimensional domain while the lighting is a function of two dimensions and the isotropic BRDF is defined on a three-dimensional domain. This seems to indicate that we have more knowns (in terms of coefficients of the reflected light field) than unknowns (lighting and BRDF coefficients).

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 to bootstrap this process. It would appear from equation 10, however, that there is an unrecoverable scale factor for each order l , corresponding to the known coefficient we require. But, we can also use reciprocity of the BRDF. To make the transfer function symmetric, we multiply it, as well as the reflected light field B , by $\cos \theta'_o$.

$$\begin{aligned} \tilde{\rho} &= \hat{\rho} \cos \theta'_o = \rho \cos \theta'_i \cos \theta'_o \\ \tilde{B} &= B \cos \theta'_o \\ \tilde{B}_{lm pq} &= \Lambda_l L_{lm} \tilde{\rho}_{lpq} \end{aligned} \quad (13)$$

The new transfer function $\tilde{\rho}$ is symmetric with respect to incident and outgoing directions, and corresponding indices: $\tilde{\rho}_{pq} = \tilde{\rho}_{p'q'}$.

There is a global scale factor we cannot recover, since \tilde{B} is not affected if we multiply the lighting and divide the BRDF by the same amount. Therefore, we scale the lighting so the DC term $L_{00} = \Lambda_0^{-1} = \sqrt{1/(4\pi)}$. Now, using equations 11, 12, and 13,

$$\begin{aligned} L_{00} &= \Lambda_0^{-1} \\ \tilde{\rho}_{0p0} &= \tilde{B}_{00p0} \\ L_{lm} &= \Lambda_l^{-1} \left(\frac{\tilde{B}_{lm pq}}{\tilde{\rho}_{lpq}} = \frac{\tilde{B}_{lm 00}}{\tilde{\rho}_{l00}} = \frac{\tilde{B}_{lm 00}}{\tilde{\rho}_{0l0}} \right) \\ &= \Lambda_l^{-1} \frac{\tilde{B}_{lm 00}}{\tilde{B}_{00l0}} \\ \tilde{\rho}_{lpq} &= \Lambda_l^{-1} \frac{\tilde{B}_{lm pq}}{L_{lm}} \\ &= \frac{\tilde{B}_{lm pq} \tilde{B}_{00l0}}{\tilde{B}_{lm 00}} \end{aligned}$$

In the last line, we can use any value of m . This gives an explicit formula for the lighting and BRDF in terms of coefficients of the output light field. 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.

5.2 Special Cases

Mirror BRDF: The mirror BRDF corresponds to a gazing sphere. Just as the inverse lighting problem is easily solved in angular space in this case, we will show that it is well-posed and easily solved in frequency space. The BRDF involves a delta function,

$$\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \delta(\cos \theta'_i - \cos \theta'_o) \delta(\phi'_i - \phi'_o \pm \pi)$$

Note that the BRDF is nonzero only when $\theta'_i \leq \pi/2$ and $\theta'_o \leq \pi/2$. The coefficients for the BRDF, reflected light field, and lighting are

$$\begin{aligned} \hat{\rho}_{lpq} &= (-1)^q \delta_{lp} \\ B_{lm pq} &= \Lambda_l (-1)^q \delta_{lp} L_{lm} \\ \forall q: L_{lm} &= \Lambda_l^{-1} (-1)^q B_{lm lq} \end{aligned} \quad (14)$$




Figure 4: Left: Successive approximations to the clamped cosine function by adding more spherical harmonic terms. For $l = 2$, we get a very good approximation. Right: The solid line is a plot of spherical harmonic coefficients $A_l = \Lambda_l \hat{\rho}_l$. For $l > 1$, odd terms vanish, and even terms decay rapidly.

The factor of $(-1)^q$ is because the azimuthal angle changes by π upon reflection. We see that the lighting coefficients correspond in a very direct way to the coefficients of the reflected light field. In signal processing terminology, the inverse lighting problem is well conditioned because the frequency spectrum of a delta function remains constant with increasing order l , and does not decay.

Single Directional Source: For convenience, we position the coordinate axes so that the source is located at $+Z$, i.e. at $(0, 0)$. Because the directional source is described by a delta function, the spherical harmonic coefficients are given simply by $L_{lm} = Y_{lm}^*(0)$, which vanishes for $m \neq 0$. Thus,

$$\begin{aligned} L_{lm} &= \delta_{m0} Y_{l0}^*(0) = \delta_{m0} \Lambda_l^{-1} \\ B_{lm pq} &= \delta_{m0} \hat{\rho}_{lpq} \\ \hat{\rho}_{lpq} &= B_{l0 pq} \end{aligned}$$

In angular space, a single observation corresponds to a single BRDF measurement. This property is used in image-based BRDF measurement [18, 21]. We see that in frequency space, there is a similar straightforward relation between BRDF coefficients and reflected light field coefficients. BRDF recovery is well-conditioned since we are estimating the BRDF filter from its impulse response.

Lambertian BRDF: For a Lambertian object, the transfer function is a scaled *clamped cosine* function, since it is proportional to the cosine of the incident angle over the upper hemisphere when $\cos \theta'_i \geq 0$ and is equal to 0 over the lower hemisphere. Plots of spherical-harmonic fits to the clamped cosine function and the magnitude of the coefficients are shown in figure 4. Because there is no dependence on outgoing angle, we can drop the indices p and q . Further, the reflected light field is now effectively the surface radiosity function, and can be expanded³ in spherical harmonics.

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

We [29] have shown that with the definitions,

$$\begin{aligned} \hat{\rho}(\theta'_i) &= \max[\cos \theta'_i, 0] = \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta'_i) \\ \hat{\rho}_l &= 2\pi \int_0^{\pi/2} \cos \theta'_i Y_{l0}(\theta'_i) \sin \theta'_i d\theta'_i \end{aligned}$$

one can derive

$$\begin{aligned} B_{lm} &= \Lambda_l \hat{\rho}_l L_{lm} \\ L_{lm} &= \Lambda_l^{-1} \frac{B_{lm}}{\hat{\rho}_l} \end{aligned} \quad (15)$$

We define $\hat{A}_l = \Lambda_l \hat{\rho}_l$. An analytic formula for A_l may be derived [29]. It can be shown that \hat{A}_l vanishes for odd values of $l > 1$,

³The basis functions $C_{lm pq}$ in equation 9 become $\Lambda_l^{-1} D_{m0}^l(\alpha) e^{l m \beta}$ if we ignore output dependence, and set $q = 0$ (the BRDF is azimuthally symmetric). It can be shown that this is simply $Y_{lm}(\alpha, \beta)$. Equation 15 now follows naturally from equation 10 upon dropping indices p and q . Our previous derivation [29] was specialized to the Lambertian case, and ignored the output dependence from the onset.

and even terms fall off very rapidly as $l^{-5/2}$. More than 99% of the energy of the BRDF filter is captured by $l \leq 2$. Numerically,

$$\hat{A}_0 = 3.14 \quad \hat{A}_1 = 2.09 \quad \hat{A}_2 = 0.79 \quad \hat{A}_3 = 0 \quad \hat{A}_4 = -0.13 \quad (16)$$

Thus, the Lambertian BRDF acts like a low-pass filter, truncating or severely attenuating frequencies with $l > 2$. Therefore, from observations of a Lambertian surface, estimation of the illumination is formally ill-posed, and is well-conditioned only for the lighting coefficients with $l \leq 2$, corresponding to 9 parameters—1 for order 0 ($(l, m) = (0, 0)$), 3 for order 1 ($l = 1, -1 \leq m \leq 1$), and 5 for order 2 ($l = 2, -2 \leq m \leq 2$). This explains the ill-conditioning observed by Marschner and Greenberg [20] in trying to solve the *inverse lighting* problem from a Lambertian surface. Furthermore, for practical applications, including forward rendering [28], **the reflected light field from a Lambertian surface can be characterized using only its first 9 spherical harmonic coefficients**; lighting effects cannot produce high-frequency variation in intensity with respect to surface curvature.

Phong BRDF: The normalized Phong transfer function is

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

where \vec{R} is the reflection of the outgoing (viewing) direction about the surface normal, \vec{L} is the direction to the light source, and s is the *shininess*, or Phong exponent. The normalization ensures the Phong lobe has unit energy. Technically, we must also zero the BRDF when the light is not in the upper hemisphere. However, the Phong BRDF is not physically based, so others have often ignored this boundary effect, and we will do the same.

We now reparameterize by the reflection vector \vec{R} , transforming the integral over the upper hemisphere centered on the surface normal to an integral centered on \vec{R} . The reflection vector takes the place of the normal in the analysis, with (α, β) referring to \vec{R} , and $\vec{R} \cdot \vec{L} = \cos \theta'_i$. The Phong BRDF after reparameterization is mathematically analogous to the Lambertian BRDF just discussed. In fact, the properties of convolution can be used to show that **for the Phong BRDF, blurring the lighting and using a mirror BRDF is equivalent to using the real lighting and real BRDF**. This formalizes the transformation often made in rendering with environment maps [23]. Specifically, equation 15 can be written as

$$\begin{aligned} L'_{lm} &= \Lambda_l \hat{\rho}_l L_{lm} \\ \Lambda_l \hat{\rho}'_l &= 1 \\ B_{lm} &= \Lambda_l \hat{\rho}'_l L'_{lm} = L'_{lm} \end{aligned} \quad (17)$$

Here, L'_{lm} is the blurred illumination and $\hat{\rho}'_l$ is the mirror BRDF⁴.

The BRDF coefficients depend on s , and are given by

$$\hat{\rho}_l = (s+1) \int_0^{\pi/2} [\cos \theta'_i]^s Y_{10}(\theta'_i) \sin \theta'_i d\theta'_i \quad (18)$$

This integral may be solved analytically. Formulae are in the appendix, and numerical plots are in figure 5.




Figure 5: Numerical plots of the Phong coefficients $\Lambda_l \hat{\rho}_l$, as defined by equation 18. The solid lines are the approximations in equation 19.

⁴The formula $\Lambda_l \hat{\rho}'_l = 1$ is not identical to equation 14 since we have now reparameterized by the reflection vector. This accounts for the slightly different normalization.

For large s and $l \ll s$, a good approximation is

$$\Lambda_l \hat{\rho}_l \approx \exp \left[-\frac{l^2}{2s} \right] \quad (19)$$

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.

Microfacet BRDF: We now consider a simplified 4-parameter Torrance-Sparrow [36] model, with parameters K_d , K_s , μ and σ . This microfacet model is widely used in computer graphics.

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

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.

We focus on the specular component, reparameterizing by the reflection vector, as for the Phong BRDF. It will also simplify matters to fix the exitant direction, and focus on the frequency-space representation of the incident-angle dependence. Precise analytic formulae are difficult to derive, but we can make a good approximation, as shown in the appendix. For normal exitance,

$$\Lambda_l \hat{\rho}_l \approx \exp \left[-(\sigma l)^2 \right] \quad (20)$$

For normal exitance, the specular part of the BRDF is a gaussian, so equation 20 simply states that even in the spherical-harmonic basis, the frequency spectrum of a gaussian is also gaussian, with the frequency width related to the reciprocal of the angular width.

For non-normal exitance, 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, the BRDF filter is essentially symmetric about the reflected direction for small viewing angles, as well as for low frequencies l . Hence, it can be shown by Taylor-series expansions and verified numerically, that the corrections to equation 20 are small under these conditions. Finally, we approximate the effects of the Fresnel factor at non-normal exitance by multiplying our expressions by $F(\mu, \theta'_o)$.

With respect to the conditioning of inverse problems, equation 20 indicates that **inverse lighting from a microfacet BRDF is well-conditioned only for frequencies up to order $l \sim \sigma^{-1}$** . Equation 20 also indicates that BRDF estimation is ill-conditioned under low-frequency lighting. For low-frequency lighting, we may apply the properties of convolution as we did for Phong BRDFs, filtering the lighting using equations 17 and 20, while using a mirror BRDF. Note that for frequencies $l \ll \sigma^{-1}$, the effects of this filtering are insignificant. The BRDF passes through virtually all the low-frequency energy. Therefore, **if the lighting contains only low frequencies, the reflected light field from a microfacet BRDF is essentially independent of the BRDF filter width σ^{-1} ; this makes estimation of the surface roughness σ ill-conditioned**.

5.3 Practical Representation

Thus far, we have presented the theoretical foundation for, and some implications of, a frequency-space view of reflection. A sig-

nal processing approach has been used before in some other areas of computer graphics, notably the theory of aliasing. Just as a frequency-space analysis of aliasing gives many insights difficult to obtain by other means, the last two sections lead to new ways of analyzing inverse rendering problems. However, the Fourier-space theory of aliasing is not generally used directly for antialiasing. 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 B_{lmpq} from a few photographs; we would require a very large number of input images, densely sampling the entire sphere of possible directions.

For these reasons, the frequency-space ideas must be put into practice carefully. Here, we first discuss two useful practical techniques—dual angular and frequency-space representations, and the separation of the lighting into slow and fast-varying components. Finally, we use these ideas, and the insights gained from the previous subsection, to derive a simple practical model of the reflected light field for the microfacet BRDF. This representation will be used extensively in the practical algorithms of section 6.

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.

Separation of slow and fast-varying lighting: For the angular-space description of the lighting, used in computing the reflected light field from the specular BRDF component, we separate the lighting into a *slow* varying component corresponding to low frequencies or area sources—for which we filter the lighting and use a mirror BRDF—and a *fast* varying component corresponding to high frequencies or directional sources. For the frequency-space lighting description, used for the diffuse BRDF component, this distinction need not be made since the formulae for the Lambertian BRDF are the same for both slow and fast varying components.

Model for Reflected Light Field: Our model for the reflected light field from the microfacet BRDF includes three components.

$$B = B_d + B_{s,slow} + B_{s,fast}$$

B_d is from the diffuse component of the BRDF. $B_{s,slow}$ represents specularities from the slowly-varying lighting, and $B_{s,fast}$ specular highlights from the fast varying lighting component.

To write B_d , corresponding to the Lambertian BRDF component, we use the 9 parameter frequency-space representation of the lighting. Explicitly noting $l \leq 2$, and with E being the irradiance,

$$B_d = K_d E(\alpha, \beta)$$

$$E(\alpha, \beta) = \sum_{l=0}^2 \left(\Lambda_l \hat{\rho}_l \sum_{m=-l}^{+l} L_{lm} Y_{lm}(\alpha, \beta) \right) \quad (21)$$

The numerical values of $\Lambda_l \hat{\rho}_l$ are given in equation 16.

For $B_{s,slow}$, we filter the lighting, using equations 17 and 20, 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

$$B_{s,slow} = K_s F(\mu, \theta'_o) L_{slow}(\vec{R}) \quad (22)$$

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 22 comes from the relation between differential microfacet and global solid angles.

$$B_{s,fast} = \frac{K_s F(\mu, \theta'_o)}{4 \cos \theta'_o} \sum_j T_j$$

$$T_j = \exp \left[-(\theta'_h / \sigma)^2 \right] \left(\frac{L_{j,fast}}{\pi \sigma^2} \right) \quad (23)$$

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.

$$B = K_d E + K_s F(\mu, \theta'_o) \left[L_{slow}(\vec{R}) + \frac{1}{4 \cos \theta'_o} \sum_j T_j(\sigma) \right] \quad (24)$$

6 Algorithms and Results

This section shows how the theory, and in particular the model just derived in section 5.3, can be applied to a broad range of practical inverse rendering problems. 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 our model of section 5.3. Section 6.1 on BRDF estimation demonstrates direct recovery of spherical harmonic BRDF coefficients, as well as estimation of parametric microfacet BRDFs using equation 24. Similarly, section 6.2 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 5.3. Section 6.3 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. Finally, section 6.4 demonstrates our algorithms on geometrically complex objects, showing how it is straightforward to extend our model to handle textures and shadowing.

To test our methods, we first used homogeneous spheres⁵ of different materials. Spheres are naturally parameterized with spherical coordinates, and therefore correspond directly to our theory. Later, we also used complex objects—a white cat sculpture, and a textured wooden doll—to show the generality of our algorithms.

Data Acquisition: 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 [40]. 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. To map from image pixels to angular coordinates $(\alpha, \beta, \theta'_o, \phi'_o)$, 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 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 gazing 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 the BRDF of a white teflon sphere using the image-based method of Marschner et al. [21]. This independent measurement was used to verify the accuracy of our BRDF estimation algorithms under complex illumination.

⁵Ordered from the McMaster-Carr catalog <http://www.mcmaster.com>

6.1 Inverse BRDF with Known Lighting

Estimation of Spherical Harmonic BRDF coefficients:

Spherical harmonics and Zernike polynomials have been fit [14] 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 $\hat{\rho}_{lpq}$.

Figure 6 compares the parametric BRDFs estimated under complex lighting to BRDFs measured using a single point source with the method of Marschner et al. [21]. As expected [14], 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, 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.

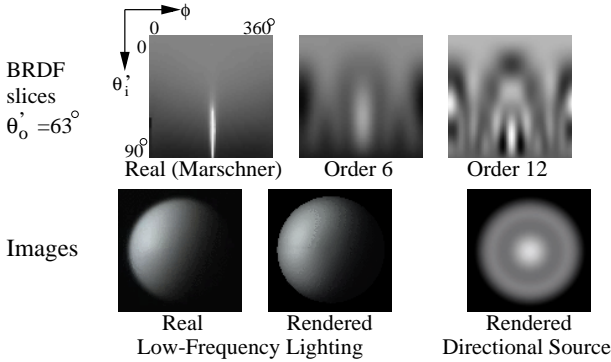


Figure 6: **Top:** Slices of the BRDF transfer function of a teflon sphere for fixed exitant angle of 63° . θ'_i varies linearly from 0° to 90° from top to bottom, and $|\phi'_o - \phi'_i|$ linearly from 0° to 360° from left to right. The central bright feature is the specular highlight. Left is the BRDF slice independently measured using the approach of Marschner et al. [21], 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.

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

Estimation of Parametric BRDF Model: We estimate BRDF parameters under general known lighting distributions using equation 24. 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 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 [9] C libraries. The main difference from previous work is that equation 24 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. [21]. 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.

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 7: Comparison of BRDF parameters recovered by our algorithm under complex lighting to those fit to measurements made by the method of Marschner et al. [21].

The results in figure 7 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 12. All these results indicate that, as expected theoretically, we can accurately estimate BRDFs even under complex lighting.

6.2 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 [20]), and when shadows are present (Sato et al. [30]). Previous methods [20, 30] 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 5.3. 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 continuous distributions of angular width approximately σ .

Estimation of Spherical Harmonic Lighting coefficients:

We represent the lighting 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 $l^* \sim \sigma^{-1}$.

$$B = \sum_{l=0}^{l^*} \sum_{m=-l}^l L_{lm} (K_d \Lambda_l \hat{\rho}_l Y_{lm}(\alpha, \beta) + K_s F Y_{lm}(\theta_R, \phi_R)) \quad (25)$$

Estimation of Parametric Dual Lighting Model: Another approach is to estimate the dual angular and frequency-space lighting model of section 5.3. 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:




Figure 8: Estimation of dual lighting representation. In phase 1, we use frequency-space parameters L_{lm}^1 to compute diffuse component B_d^1 . 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_{lm}^2 , which can be used to determine B_d^2 . The consistency condition is that $B_d^1 = B_d^2$ or $L_{lm}^1 = L_{lm}^2$. 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).

1. **9 parameters to High-Resolution Lighting:** The inputs to phase 1 are the coefficients L_{lm}^1 . These suffice to find B_d^1 by equation 21. Since we assumed that $B_{s,fast} = 0$,

$$B_{s,slow} = K_s F(\mu, \theta'_o) L_{slow}(\vec{R}) = B - B_d^1(L_{lm}^1)$$

$$L_{slow}(\vec{R}) = \frac{B - B_d^1(L_{lm}^1)}{K_s F(\mu, \theta'_o)}$$

We assume the BRDF parameters are known, and B is the input to the algorithm, so the right-hand side can be evaluated.

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_{lm}^2 .

Now, assume we run phase 1 (with inputs L_{lm}^1) and phase 2 (with outputs L_{lm}^2) sequentially. The consistency condition is that $L_{lm}^1 = L_{lm}^2$ —converting from frequency to angular to frequency space must not change the result. Equivalently, the computed diffuse components must match, i.e. $B_d^1(L_{lm}^1) = B_d^2(L_{lm}^2)$. This is illustrated in figure 8. 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.

Figure 9 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).




Figure 9: 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.

6.3 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 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  $K_d, K_s, \mu, \sigma$  // BRDF parameters
global  $L$  // Lighting
procedure Factor
  Minimize( $K_s, \mu, \text{ObjFun}$ ) // Simplex Method
   $\sigma = \text{FindRoughness}(L)$  // Figure 10, Equation 26
function ObjFun( $K_s, \mu$ )
   $K_d = 1 - K_s$  //  $K_d + K_s = 1$ 
   $L = \text{Lighting}(K_d, K_s, \mu)$  // Inverse Lighting
   $B_{pred} = \text{Predict}(L, K_d, K_s, \mu)$  // 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 22 and 23 for the center, i.e. brightest point, of the light source at normal exitance. An illustration is in figure 10.

$$L_{cen} \approx \frac{L_{total}}{4\pi\sigma^2} \quad (26)$$




Figure 10: 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 26. The method does not depend on the size of the red box—provided it encloses the entire (filtered) source—the precise shape into which the source is filtered in the recovered lighting.

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.

We see from figure 11 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 12, 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 12 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	Unknown
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 11: BRDFs of various spheres, recovered under known (section 6.1) and unknown (section 6.3) 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.

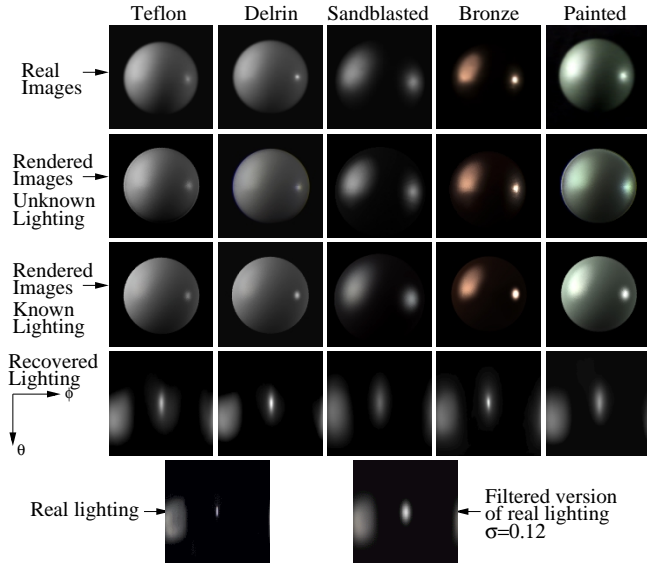


Figure 12: Spheres rendered using BRDFs estimated under known (section 6.1) and unknown (section 6.3) lighting. The algorithm in section 6.3 also recovers the lighting. Since there is an unknown global scale, we scale the recovered lighting distributions in order to compare them. The recovered illumination 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.

6.4 Complex Objects—Texture and Shadowing

We now demonstrate our algorithms on objects with complex geometry, and discuss extensions to handle concave surfaces and textured objects. Although the theory is developed for homogeneous surfaces, our algorithms can be extended to textured objects simply by letting the BRDF parameters be functions of surface position. 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, using our new practical model of section 5.3, 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.

Shadowing—Concave Surfaces: In our practical model, the reflected light field consists of 3 parts— B_d , $B_{s,slow}$, and $B_{s,fast}$. $B_{s,slow}$ depends on $L_{slow}(\vec{R})$, the slowly-varying component of the lighting evaluated at the reflection vector. Our model allows us to approximate the effects of shadowing simply by checking if the reflected ray is shadowed. The other components are handled in the standard manner. To consider shadowing when computing $B_{s,fast}$, corresponding to specularities from directional sources, we check if these sources are shadowed. B_d depends on the irradiance E , which should now be computed in the more conventional angular-space way by integrating the scene lighting while considering visibility, instead of using the 9-parameter lighting approximation of equation 21. It should be emphasized that in all cases, the corrections for visibility depend only on object geometry, and can be precomputed for each point on the object using a ray tracer.

For parametric BRDF estimation, we modify each component of equation 24 to consider visibility, as discussed above. Our first inverse lighting method, that directly recovers the coefficients L_{lm} , is modified similarly. In equation 25, we check if the reflected ray is shadowed, and consider shadowing when computing the irradiance due to each Y_{lm} . Note that B is still a linear combination of the lighting coefficients, so we will still solve a linear system for L_{lm} . However, it is difficult to extend our dual angular and frequency-space method for inverse lighting to handle concave surfaces because B_d no longer depends only on the 9 lighting coefficients L_{lm} with $l \leq 2$. For light field factorization, we simply extend both the inverse-BRDF and inverse-lighting methods as discussed.

A white cat sculpture was used to test our algorithms on complex geometric objects that include concavities. 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 6.1, and unknown lighting—using the factorization method of section 6.3, with the inverse lighting component being direct recovery of spherical harmonic coefficients using $l^* = 12$. Comparisons of photographs and renderings are in figures 1 and 13. BRDF and lighting parameters are tabulated in figure 14. This experiment indicates that our methods for BRDF recovery under known and unknown lighting are consistent, and 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.

Textured BRDFs: Since the theory shows that factorization of lighting and texture is ambiguous, we consider only recovery of textured BRDFs under known lighting. It is fairly straightforward to allow $K_d(\mathbf{x})$ and $K_s(\mathbf{x})$ to be described by textures that depend on surface position \mathbf{x} . In the inner procedure of the parametric BRDF estimation algorithm of section 6.1, we simply solve a separate linear problem for each point \mathbf{x} to estimate $K_d(\mathbf{x})$ and $K_s(\mathbf{x})$.

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 and compared this to a rendering using the textured BRDF recovered under complex illumination. The results in figure 15 show that our renderings closely resemble real photographs. The main artifact is blurring of texture because of geometry-image misregistration.

7 Conclusions and Future Work

This paper has developed a signal-processing framework for inverse rendering. The qualitative observation that the reflected light field is a convolution of the lighting and BRDF has been formalized mathematically. 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




Figure 13: Comparison of real 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.

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 14: 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.

why a directional source is well-suited for recovering the BRDF—we are estimating the BRDF filter by considering its impulse response. The conditioning properties and well-posedness of BRDF and lighting estimation under various conditions have been derived, as well as an explicit formula for *factoring* the reflected light field into the lighting and BRDF. The ill-conditioning observed by Marschner and Greenberg [20] in estimating the lighting from a Lambertian surface has been explained, 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.

The insights gained from the theory also 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 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 [39], Sato and Ikeuchi [31], and Love [17]—which are designed specifically for natural lighting—can be seen as special cases of this general




Figure 15: Recovering textured BRDFs under complex lighting. The rendered images closely resemble the real photographs, even under novel lighting.

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 resolutions to be obtained than with previous purely angular-space methods [20, 30]. 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.

We have only scratched the surface of possible applications. In the future, it is likely that many more algorithms can be derived using the basic approaches outlined here. Possible algorithmic improvements include extending the consistency condition for inverse lighting so we can use color-space methods [13] to help separate diffuse and specular components for colored objects. Finally, while we have discussed only inverse rendering applications, we believe the convolution-based approach is of theoretical and practical importance in many other areas of computer graphics. We have already shown [28] how to use the 9 term irradiance approximation for efficient forward rendering of diffuse objects with environment maps, and we believe there are many further applications.

Acknowledgements: We are grateful to Marc Levoy for many helpful initial discussions regarding both the interpretation of reflection as convolution and the practical issues in inverse rendering, as well as for reading a draft of the paper. We thank Szymon Rusinkiewicz for many suggestions regarding this research, for reading early versions of the paper, and for being ever willing to assist us in debugging problems with the gantry. Steve Marschner also deserves thanks for detailed comments on many early drafts. Jean Gleason at Bal-tec assisted in supplying us with our mirror sphere. Finally, we thank John Parisseenti at Polytec Products for much useful advice on obtaining uniform spheres, and for getting one of our steel spheres sandblasted. The cat is a range scan of the sculpture *Serenity* by Sue Dawes. The work described in this paper was supported in part by a Hodgson-Reed Stanford graduate fellowship and NSF ITR grant #0085864 “Interacting with the Visual World.”

References

- [1] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. In *International Conference on Computer Vision*, 2001.
- [2] R. Bastos, K. Hoff, W. Wynn, and A. Lastra. Increased photorealism for interactive architectural walkthroughs. In *ISD 99*, pages 183–190, 1999.
- [3] B. Cabral, N. Max, and R. Springmeyer. Bidirectional reflection functions from surface bump maps. In *SIGGRAPH 87*, pages 273–281, 1987.
- [4] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH 96*, pages 303–312, 1996.
- [5] 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.
- [6] 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.
- [7] 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*, 2001.
- [8] M. D’Zmura. *Computational Models of Visual Processing*, chapter Shading Ambiguity: Reflectance and Illumination, pages 187–207. MIT Press, 1991.

- [9] Numerical Algorithms Group. *NAG C Library Manual, Mark 5*. 1999.
- [10] K. Ikeuchi and K. Sato. Determining reflectance properties of an object using range and brightness images. *PAMI*, 13(11):1139–1153, 1991.
- [11] T. Inui, Y. Tanabe, and Y. Onodera. *Group theory and its applications in physics*. Springer Verlag, 1990.
- [12] K. F. Karner, H. Mayer, and M. Gervautz. An image based measurement system for anisotropic reflection. *Computer Graphics Forum*, 15(3):119–128, 1996.
- [13] G.J. Klinker, S.A. Shafer, and T. Kanade. The measurement of highlights in color images. *IJCV*, 2(1):7–32, 1988.
- [14] J. J. Koenderink and A. J. van Doorn. Phenomenological description of bidirectional surface reflection. *JOSA A*, 15(11):2903–2912, 1998.
- [15] E. Land and J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.
- [16] 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.
- [17] R.C. Love. *Surface Reflection Model Estimation from Naturally Illuminated Image Sequences*. PhD thesis, Leeds, 1997.
- [18] R. Lu, J. Koenderink, and A. Kappers. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics*, 37(25):5974–5984, 1998.
- [19] T.M. MacRobert. *Spherical harmonics; an elementary treatise on harmonic functions, with applications*. Dover Publications, 1948.
- [20] S.R. Marschner and D.P. Greenberg. Inverse lighting for photography. In *Fifth Color Imaging Conference*, pages 262–265, 1997.
- [21] 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.
- [22] N. McCormick. Inverse radiative transfer problems: a review. *Nuclear Science and Engineering*, 112:185–198, 1992.
- [23] 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.
- [24] 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.
- [25] J. Nimeroff, E. Simoncelli, and J. Dorsey. Efficient re-rendering of naturally illuminated environments. In *EGWR 94*, pages 359–373.
- [26] R.W. Preisendorfer. *Hydrologic Optics*. US Dept Commerce, 1976.
- [27] 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*, 2001.
- [28] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *SIGGRAPH 01*, 2001.
- [29] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. To appear, *Journal of the Optical Society of America A*, 2001.
- [30] 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.
- [31] Y. Sato and K. Ikeuchi. Reflectance analysis under solar illumination. Technical Report CMU-CS-94-221, Carnegie Mellon University, 1994.
- [32] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *SIGGRAPH 97*, pages 379–388, 1997.
- [33] F. X. Sillion, J. Arvo, S. H. Westin, and D. Greenberg. A global illumination solution for general reflectance distributions. In *SIGGRAPH 91*, pages 187–196.
- [34] S.R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell, 1998.
- [35] S. Tominaga and N. Tanaka. Estimating reflection parameters from a single color image. *IEEE Computer Graphics & Applications*, 20(5):58–66, 2000.
- [36] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967.
- [37] G. J. Ward. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92*, pages 265–272, 1992.
- [38] 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.
- [39] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *SIGGRAPH 98*, pages 207–218, 1998.
- [40] Z. Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1330–1334, 2000.

Appendix: Mathematical Derivations

Frequency-space Reflection Equation: We rewrite the spherical harmonic lighting expansion in equation 7, applying the rotation formula in equation 6.

$$\begin{aligned} L(\theta_i, \phi_i) &= L(R_{\alpha, \beta}(\theta'_i, \phi'_i)) = \sum_{l, m} L_{lm} Y_{lm}(R_{\alpha, \beta}(\theta'_i, \phi'_i)) \\ &= \sum_{l, m, m'} L_{lm} D_{m m'}^l(\alpha) e^{I m \beta} Y_{l m'}(\theta'_i, \phi'_i) \end{aligned}$$

Next, we reproduce the BRDF expansion in equation 8, using l' instead of l , in order not to conflict with the index l , already in use for the lighting.

$$\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \sum_{l', p, q} \hat{\rho}_{l' p q} Y_{l' p}^*(\theta'_i, \phi'_i) Y_{p q}(\theta'_o, \phi'_o)$$

We can now write out the reflection integral, equation 5. The integrand is simply the product of the above two results for the lighting and BRDF. After taking terms not depending on (θ'_i, ϕ'_i) outside the integral, we obtain

$$\begin{aligned} B(\alpha, \beta, \theta'_o, \phi'_o) &= \\ &\sum_{l, l', m, m', p, q} [L_{lm} \hat{\rho}_{l' p q} D_{m m'}^l(\alpha) e^{I m \beta} Y_{p q}(\theta'_o, \phi'_o) \cdot \\ &\int_{\Omega'_i} Y_{l m'}(\theta'_i, \phi'_i) Y_{l' q}^*(\theta'_i, \phi'_i) d\omega'_i] \end{aligned}$$

By orthonormality of the spherical harmonics, the value of the integral is $\delta_{l l'} \delta_{m' q}$. We can therefore set $m' = q$ and $l' = l$ and eliminate l' and m' .

$$B(\alpha, \beta, \theta'_o, \phi'_o) = \sum_{l, m, p, q} (L_{lm} \hat{\rho}_{l p q}) (D_{m q}^l(\alpha) e^{I m \beta} Y_{p q}(\theta'_o, \phi'_o))$$

From this, we derive the form of $C_{l m p q}$ in equation 9, and the frequency-space reflection formula, equation 10. The prefactor Λ_l^{-1} in equation 9, and the resulting factor Λ_l in equation 10, is for normalization. The orthonormality of the basis functions $C_{l m p q}$, as well as the value of the normalization constant Λ_l , can be derived from the orthogonality of the group representation matrices D^l (equation 7.73 of [11]).

Phong BRDF: We substitute $u = \cos \theta'_i$ in equation 18, and note that $Y_{10}(\theta'_i) = \Lambda_l^{-1} P_l(\cos \theta'_i)$, where P_l is the legendre polynomial of order l .

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

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

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

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 $s > l-1$,

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

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

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

For large s , $1/s \gg 1/s^2$, and we derive equation 19.

Microfacet BRDF: Consider normal exitance, and ignore Fresnel variation with incident angle. There is no azimuthal dependence, and $\theta'_h = \theta'_i/2$.

$$\hat{\rho}_l = 2\pi \int_0^{\pi/2} \frac{\exp[-\theta'^2/4\sigma^2]}{4\pi\sigma^2} Y_{10}(\theta'_i) \sin \theta'_i d\theta'_i \quad (27)$$

The expansion of $Y_{10}(t)$ near $t=0$ for small l is

$$Y_{10}(t) = \Lambda_l^{-1} \left(1 - \frac{l(l+1)}{4} t^2 + O(t^4) \right) \quad (28)$$

The asymptotic form of $Y_{10}(t)$ near $t=0$ for large l is

$$Y_{10}(t) \sim \Lambda_l^{-1} \left(\frac{1}{\sqrt{t}} \cos[(l+1/2)t - \pi/4] \right) \quad (29)$$

To integrate equation 27, we substitute $\theta'_i = 2\sigma u$. Assuming $\sigma \ll 1$, the upper limit of the integral becomes infinite, and $\sin \theta'_i d\theta'_i = \theta'_i d\theta'_i = 4\sigma^2 u du$.

$$\hat{\rho}_l = \int_0^\infty 2e^{-u^2} Y_{10}(2\sigma u) u du$$

We therefore set $t = 2\sigma u$ in equations 28 and 29. When $\sigma l \ll 1$, we use equation 28.

$$\begin{aligned} \Lambda_l \hat{\rho}_l &= \left(\int_0^\infty 2ue^{-u^2} du - (\sigma l)^2 \int_0^\infty 2u^3 e^{-u^2} du \right) \\ &= 1 - (\sigma l)^2 + O([\sigma l]^4) \end{aligned}$$

Note that these are the first terms in the Taylor series of $\exp[-(\sigma l)^2]$. When $\sigma l \gg 1$, we use equation 29 to obtain (Φ is a phase that encapsulates the lower-order terms):

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

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 equation 20. We have also verified this result numerically.