

A Signal-Processing Framework for Reflection—Part 2: Analytic Formulae for Common Lighting and BRDF Models

Ravi Ramamoorthi, Columbia University*

Pat Hanrahan, Stanford University

Abstract

The first part of this paper presented a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. The notion of reflection as convolution was formalized, with a quantitative analysis of the reflected light field as obtained by a spherical convolution of the incident illumination and the BRDF. In this part of the paper, we derive analytic formulae for spherical harmonic coefficients of many common lighting and BRDF models. In this way, we explore the implications of the theoretical analysis for many special cases of interest, determining the well-posedness and conditioning of many forward and inverse rendering problems, and explaining previous empirical results.

Keywords: Reflection, Illumination, BRDF, Signal Processing, Spherical Harmonics, Environment Maps, Inverse Rendering,

1 Introduction

An analysis of the computational properties of the reflection operator is of interest in both computer graphics and vision, for analyzing forward and inverse problems in rendering. Building on previous qualitative observations by Miller and Hoffman [13], Cabral et al. [3, 4], D’Zmura [5] and others, the first part of this paper formalized the notion of reflection as a spherical convolution of the incident illumination and the BRDF.

Specifically, we were able to develop a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. Under these

* (ravir@cs.columbia.edu hanrahan@cs.stanford.edu) Department of Computer Science, Columbia University, 1214 Amsterdam Ave., New York, NY 10027. 917-375-8378, 212-666-0140 (fax)

assumptions, we were able to derive an analytic formula for the reflected light field in terms of the product of spherical harmonic coefficients of the BRDF and the lighting. Our formulation allows us to view forward rendering as *convolution* and inverse rendering as *deconvolution*.

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

The rest of this paper is organized as follows. In section 2, we briefly summarize the main results from the first part of the paper. Section 3 is the main part of this paper, and works out analytic formulae for spherical harmonic coefficients of many common lighting and BRDF models, demonstrating the implications of the theoretical analysis. This section is a more detailed version of the derivations in our previous SIGGRAPH paper [17]. New material includes verification of the spherical harmonic formulae from first principles, as well as a discussion of light field factorization for the special cases of interest. Finally, section 4 concludes the paper and discusses future work.

2 Background

In this section, we briefly summarize the main theoretical results derived in the first part of the paper, introducing the notation and terminology required in the next section. In this paper, we will only discuss results in 3D, since that is of greater practical importance. For simplicity, we will also restrict ourselves to isotropic BRDFs. A more complete derivation of the convolution formula, along with a number of alternative forms, is found in the first part of the paper. Notation used in this paper is in figure 1.

2.1 Reflection Equation and Convolution Formula

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

$$B(\alpha, \beta, \theta'_o, \phi'_o) = \int_{\Omega'_i} L(R_{\alpha, \beta}(\theta'_i, \phi'_i)) \hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) d\omega'_i, \quad (1)$$

where (α, β) are the spherical coordinates of the surface normal, (θ'_o, ϕ'_o) are the angular coordinates of the viewing vector in local coordinates (i.e. with respect to the local surface normal), and (θ'_i, ϕ'_i) are the angular coordinates of the local incident direction. B is the reflected light field, L is the incident illumination, and $\hat{\rho}$ denotes the BRDF multiplied by the cosine of the incident angle. R

B	Reflected radiance
B_{lmpq}	Coefficients of isotropic basis-function expansion of B
L	Incoming radiance
L_{lm}	Coefficients of spherical-harmonic expansion of L i
$\hat{\rho}$	BRDF multiplied by cosine of incident angle
$\hat{\rho}_{lpq}$	Coefficients of isotropic 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
α, β	Surface normal parameterization—elevation, azimuthal angles
$R_{\alpha, \beta}$	Rotation operator for surface normal (α, β)
Y_{lm}	Spherical Harmonic
Y_{lm}^*	Complex Conjugate of Spherical Harmonic
$D_{mm'}^l$	Representation matrix of dimension $2l + 1$ for rotation group $SO(3)$
Λ_l	Normalization constant, $\sqrt{4\pi/(2l + 1)}$
I	$\sqrt{-1}$

Figure 1: *Notation*

is a rotation operator corresponding to the surface normal that rotates the local incident direction into global coordinates in which the lighting is measured, i.e. $R_{\alpha,\beta} = R_z(\beta)R_y(\alpha)$.

It is possible to derive a frequency-space convolution formula corresponding to equation 1. For this purpose, we must expand quantities in terms of spherical harmonics. Spherical harmonics are the analogue on the sphere to the Fourier basis on the line or circle. The spherical harmonic Y_{lm} is given by

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

where N_{lm} is a normalization factor. In the above equation, the azimuthal dependence is expanded in terms of Fourier basis functions. The θ dependence is expanded in terms of the associated Legendre functions P_{lm} . The indices obey $l \geq 0$ and $-l \leq m \leq l$. Thus, there are $2l + 1$ basis functions for given order l .

The rotation formula for spherical harmonics is

$$Y_{lm}(R_{\alpha,\beta}(\theta, \phi)) = \sum_{m'=-l}^l D_{mm'}^l(\alpha, \beta) Y_{lm'}(\theta, \phi). \quad (3)$$

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, i.e. how to rewrite a rotated spherical harmonic as a linear combination of all the spherical harmonics of the same order. In terms of group theory, the matrix D^l is analogous to the $(2l + 1)$ -dimensional representation of the rotation group $SO(3)$. Two important properties obeyed by the representation matrices are

$$\begin{aligned} D_{0m'}^l(\alpha, \beta) &= \Lambda_l Y_{lm'}^*(\alpha, \pi) \\ D_{m0}^l(\alpha, \beta) &= \Lambda_l Y_{lm}(\alpha, \beta), \end{aligned} \quad (4)$$

where we have used the constant $\Lambda_l = \sqrt{4\pi/(2l + 1)}$.

The spherical harmonics form an orthonormal basis in terms of which functions on the sphere can be expanded. Specifically, the illumination can be written as

$$\begin{aligned} L(\theta_i, \phi_i) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm} Y_{lm}(\theta_i, \phi_i) \\ L(\theta_i, \phi_i) = L(R_{\alpha,\beta}(\theta'_i, \phi'_i)) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{+l} \sum_{m'=-l}^l L_{lm} D_{mm'}^l(\alpha, \beta) Y_{lm'}(\theta'_i, \phi'_i). \end{aligned} \quad (5)$$

Then, we write the expansion of the isotropic BRDF,

$$\hat{\rho}(\theta'_i, \theta'_o, | \phi'_o - \phi'_i |) = \sum_{l=0}^{\infty} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} \hat{\rho}_{lpq} Y_{lq}^*(\theta'_i, \phi'_i) Y_{pq}(\theta'_o, \phi'_o). \quad (6)$$

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

$$\begin{aligned} C_{lmpq}(\alpha, \beta, \theta'_o, \phi'_o) &= \Lambda_l^{-1} D_{mq}^l(\alpha, \beta) Y_{pq}(\theta'_o, \phi'_o) \\ B(\alpha, \beta, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} B_{lmpq} C_{lmpq}(\alpha, \beta, \theta'_o, \phi'_o) \\ B_{lmpq} &= \int_{\phi'_o=0}^{2\pi} \int_{\theta'_o=0}^{\pi} \int_{\beta=0}^{2\pi} \int_{\alpha=0}^{\pi} B(\alpha, \beta, \theta'_o, \phi'_o) C_{lmpq}^*(\alpha, \beta, \theta'_o, \phi'_o) \sin \alpha \sin \theta'_o d\alpha d\beta d\theta'_o d\phi'_o. \end{aligned} \quad (7)$$

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

$$B_{lmpq} = \Lambda_l L_{lm} \hat{\rho}_{lpq} \quad (8)$$

We may also derive an alternative form, holding the outgoing elevation angle θ'_o fixed. We first define the modified expansions,

$$\begin{aligned} \hat{\rho}(\theta'_i, \theta'_o, | \phi'_o - \phi'_i |) &= \sum_{l=0}^{\infty} \sum_{q=-l}^l \hat{\rho}_{lq}(\theta'_o) \left(\frac{1}{\sqrt{2\pi}} Y_{lq}^*(\theta'_i, \phi'_i) \exp(Iq\phi'_o) \right) \\ B(\alpha, \beta, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{q=-l}^l B_{lmq}(\theta'_o) \left(\frac{1}{\sqrt{2\pi}} \Lambda_l^{-1} D_{mq}^l(\alpha, \beta) \exp(Iq\phi'_o) \right). \end{aligned} \quad (9)$$

Then, we may write down an alternative version of the isotropic convolution formula.

$$B_{lmq}(\theta'_o) = \Lambda_l L_{lm} \hat{\rho}_{lq}(\theta'_o) \quad (10)$$

If we seek to preserve the *reciprocity* of the BRDF, i.e. symmetry with respect to incident and outgoing angles, we may multiply both the transfer function and the reflected light field by $\cos \theta'_o$, defining

$$\begin{aligned} \tilde{\rho} &= \hat{\rho} \cos \theta'_o = \rho \cos \theta'_i \cos \theta'_o \\ \tilde{B} &= B \cos \theta'_o. \end{aligned} \quad (11)$$

With these derivations, equation 8 becomes

$$\tilde{B}_{lmpq} = \Lambda_l L_{lm} \tilde{\rho}_{lpq} \quad (12)$$

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

Another interesting special case occurs for specular BRDFs, or those that are nearly symmetric about a central direction such as the reflection vector. In this case, we may simplify the BRDF by reparameterizing by the reflection vector.

Many models, such as Lambertian and Phong BRDFs are *radially symmetric or 1D BRDFs*, where the BRDF consists of a single symmetric lobe of fixed shape, whose orientation depends only on a well-defined central direction \vec{C} . If we reparameterize by \vec{C} , the BRDF becomes a function of only 1 variable (θ'_i with $\cos \theta'_i = \vec{C} \cdot \vec{L}$) instead of 3. Further, the reflected light field can be represented simply by a 2D reflection map $B(\alpha, \beta)$ parameterized by $\vec{C} = (\alpha, \beta)$. In other words, after reparameterization, there is no explicit exitant (outgoing) angular dependence for either the BRDF or reflected light field. In this case, we may write the BRDF and equations for the reflected light field as

$$\begin{aligned}
\hat{\rho}(\theta'_i) &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta'_i) \\
\hat{\rho}_l &= 2\pi \int_0^{\pi/2} \hat{\rho}(\theta'_i) Y_{l0}(\theta'_i) \sin \theta'_i d\theta'_i \\
B(\alpha, \beta) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{q=-l}^l \hat{\rho}_l L_{lm} D_{mq}^l(\alpha, \beta) \int_0^{2\pi} \int_0^{\pi} Y_{lq}(\theta'_i, \phi'_i) Y_{l0}(\theta'_i) \sin \theta'_i d\theta'_i d\phi'_i \\
&= \sum_{l=0}^{\infty} \sum_{m=-l}^l \hat{\rho}_l L_{lm} D_{m0}^l(\alpha, \beta). \tag{13}
\end{aligned}$$

Now, we use the second property of the matrices D from equation 4, i.e. $D_{m0}^l(\alpha, \beta) = \Lambda_l Y_{lm}(\alpha, \beta)$. Therefore, the reflected light field can be expanded simply in terms of spherical harmonics,

$$B(\alpha, \beta) = \sum_{l=0}^{\infty} \sum_{m=-l}^l B_{lm} Y_{lm}(\alpha, \beta). \tag{14}$$

The required convolution formula now becomes

$$B_{lm} = \Lambda_l \hat{\rho}_l L_{lm}. \tag{15}$$

2.2 Analysis of Inverse Problems

The convolution formula in equation 8 can be used to analyze the well-posedness and numerical conditioning of inverse problems. For the inverse-BRDF problem, we manipulate equation 8 to yield

$$\hat{\rho}_{lpq} = \Lambda_l^{-1} \frac{B_{lmpq}}{L_{lm}}. \tag{16}$$

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

A similar analysis can be done for estimation of the lighting,

$$L_{lm} = \Lambda_l^{-1} \frac{B_{lm pq}}{\hat{\rho}_{lpq}}. \quad (17)$$

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

Finally, we can put these results together to derive an analytic formula for *factoring* the reflected light field, i.e. determining both the lighting and BRDF in terms of coefficients of the reflected light field. Since we cannot recover the global scale, we will arbitrarily scale the DC term of the lighting so $L_{00} = \Lambda_0^{-1} = \sqrt{1/(4\pi)}$. Using equations 12, 16, and 17, we obtain

$$\begin{aligned} L_{00} &= \Lambda_0^{-1} && : \text{Global Scale} \\ \tilde{\rho}_{0p0} &= \tilde{B}_{00p0} && : \text{Equation 16 } (l = q = 0) \\ L_{lm} &= \Lambda_l^{-1} \frac{\tilde{B}_{lm pq}}{\tilde{\rho}_{lpq}} && : \text{Equation 17} \\ &= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{l00}} && : \text{Set } p = q = 0 \\ &= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{0l0}} && : \text{Reciprocity, } \tilde{\rho}_{0l0} = \tilde{\rho}_{l00} \\ &= \Lambda_l^{-1} \frac{\tilde{B}_{lm00}}{B_{00l0}} && : \text{Plug in from 2}^{\text{nd}} \text{ line} \\ \tilde{\rho}_{lpq} &= \Lambda_l^{-1} \frac{\tilde{B}_{lm pq}}{L_{lm}} && : \text{Equation 16} \\ &= \frac{\tilde{B}_{lm pq} \tilde{B}_{00l0}}{\tilde{B}_{lm00}} && : \text{Substitute from above for } L_{lm}. \end{aligned} \quad (18)$$

Therefore, up to global scale, **the reflected light field can be factored into the lighting and the BRDF**, provided the appropriate coefficients of the reflected light field do not vanish, i.e. the denominators above are nonzero.

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

3 Analytic Formulae for Common Lighting, BRDF Models

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

3.1 Directional Source

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

$$\begin{aligned}
 L(\theta_i, \phi_i) &= \delta(\cos \theta_i - \cos \theta_s) \delta(\phi_i - \phi_s) \\
 L_{lm} &= \int_0^{2\pi} \int_0^\pi \delta(\cos \theta_i - \cos \theta_s) \delta(\phi_i - \phi_s) Y_{lm}^*(\theta_i, \phi_i) \sin \theta_i d\theta_i d\phi_i \\
 &= Y_{lm}^*(\theta_s, \phi_s).
 \end{aligned} \tag{19}$$

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

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

$$\begin{aligned}
 L_{lm} &= Y_{lm}^*(0) \\
 &= \Lambda_l^{-1} \delta_{m0} \\
 B_{lm pq} &= \delta_{m0} \hat{\rho}_{l pq} \\
 \hat{\rho}_{l pq} &= B_{l0 pq}.
 \end{aligned} \tag{20}$$

In angular space, a single observation of the reflected light field corresponds to a single BRDF measurement. This property is used in image-based BRDF measurement [9, 12]. We see that in frequency space, there is a similar straightforward relation between BRDF coefficients and reflected light field coefficients. BRDF recovery is well-posed and well-conditioned since we are estimating the BRDF filter from its impulse response.

It is instructive to verify equation 20 directly from first principles. We first note that in angular space, the reflected light field for a directional source at the north pole can be written as

$$B(\alpha, \beta, \theta'_o, \phi'_o) = \hat{\rho}(\alpha, \pi, \theta'_o, \phi'_o). \tag{21}$$

Note that the incident angle for surface normal (α, β) is given by (α, π) . Clearly the elevation angle of incidence must be α , and because of our standard right-handed sign convention, the azimuthal angle is π . We will now show that equation 20 is simply a frequency-space version of equation 21 by expanding out B and $\hat{\rho}$, using the expressions in equation 20. We will need to use the first property of the representation matrices from equation 4. The first line below simply derives the form of $C_{l0 pq}$, making use of equation 4. In the next line, we expand the left hand side of equation 21, $B(\alpha, \beta, \theta'_o, \phi'_o)$ in terms of $C_{l0 pq}$. Note that since the directional source is at the north pole, there is no azimuthal dependence and we can assume that $m = 0$. Finally, we expand the right hand side of equation 21, and equate coefficients,

$$\begin{aligned}
C_{l0pq}(\alpha, \beta, \theta'_o, \phi'_o) &= \Lambda_l^{-1} D_{0q}^l(\alpha, \beta) Y_{pq}(\theta'_o, \phi'_o) \\
&= Y_{lq}^*(\alpha, \pi) Y_{pq}(\theta'_o, \phi'_o) \\
B(\alpha, \beta, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} B_{l0pq} Y_{lq}^*(\alpha, \pi) Y_{pq}(\theta'_o, \phi'_o) \\
\hat{\rho}(\alpha, \pi, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} \hat{\rho}_{l0pq} Y_{lq}^*(\alpha, \pi) Y_{pq}(\theta'_o, \phi'_o) \\
B_{l0pq} &= \hat{\rho}_{l0pq}.
\end{aligned} \tag{22}$$

The last line comes from equating coefficients, and this confirms the correctness of equation 20, thereby verifying the convolution formula for the special case of a single directional source.

Finally, we may verify the factorization relations of equation 18 for the case when both the BRDF and lighting are unknown *a priori*, and the lighting is actually a single directional source.

$$\begin{aligned}
L_{lm} &= \Lambda_l^{-1} \frac{\tilde{B}_{lm00}}{\tilde{B}_{00l0}} \\
&= \Lambda_l^{-1} \delta_{m0} \left(\frac{\tilde{\rho}_{l00}}{\tilde{\rho}_{0l0}} \right) \\
\tilde{\rho}_{l0pq} &= \frac{B_{l0pq} B_{00l0}}{B_{l000}} \\
&= \tilde{\rho}_{l0pq} \left(\frac{\tilde{\rho}_{0l0}}{\tilde{\rho}_{l00}} \right).
\end{aligned} \tag{23}$$

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

In summary, we have derived an analytic frequency-space formula, that has been verified from first principles. A directional source corresponds to a delta function, whose frequency spectrum does not decay. Therefore, BRDF estimation is well-posed and well-conditioned for all frequencies. In effect, we are estimating the BRDF filter from its impulse response. This is a frequency-space explanation for the use of directional sources in BRDF measurement, especially in the newer image-based methods for curved surfaces [9, 12].

3.2 Axially Symmetric Distribution

We now consider a lighting distribution that is symmetric about some axis. For convenience, we position the coordinate system so that the Z axis is the axis of symmetry. Such a distribution closely approximates the illumination due to skylight on a cloudy day. We should also note that

a single directional source, as just discussed, is a trivial example of an axially symmetric lighting distribution. The general property of these configurations is that the lighting coefficients L_{lm} with $m \neq 0$ vanish, since they have azimuthal dependence. The frequency-space reflection formulas now become

$$\begin{aligned} B_{lmpq} &= \delta_{m0} \Lambda_l L_{l0} \hat{\rho}_{lpq} \\ \hat{\rho}_{lpq} &= \Lambda_l^{-1} \frac{B_{l0pq}}{L_{l0}}. \end{aligned} \quad (24)$$

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

3.3 Uniform Lighting

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

$$\begin{aligned} L_{lm} &= \delta_{l0} \delta_{m0} \Lambda_0^{-1} \\ B_{lmpq} &= \delta_{l0} \delta_{m0} \delta_{q0} \hat{\rho}_{0p0} \\ \hat{\rho}_{0p0} &= B_{00p0}. \end{aligned} \quad (25)$$

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

All this may be somewhat clearer if we use the convolution formula from equation 10, where the dependence on the outgoing elevation angle is not expanded into basis functions,

$$B_{lmq}(\theta'_o) = \delta_{l0} \delta_{m0} \delta_{q0} \hat{\rho}_{00}(\theta'_o). \quad (26)$$

This formulation makes it clear that only the lowest frequency, i.e. DC term of the transfer function contributes to the reflected light field, with B being independent of surface orientation. From

observations made under uniform lighting, we can estimate only the DC term of the BRDF transfer function; higher frequencies in the BRDF cannot be estimated. Hence, a mirror surface cannot be distinguished from a Lambertian object under uniform lighting.

We may also verify equation 26 by writing it in terms of angular space coordinates. First, we note the form of the basis functions, noting that $Y_{00} = \Lambda_0^{-1}$ and that $D_{00}^0 = 1$.

$$\begin{aligned}
C_{000} &= \sqrt{\frac{1}{2\pi}} \Lambda_0^{-1} \\
\hat{\rho}_{00}(\theta'_o) &= \sqrt{\frac{1}{2\pi}} \int_0^{2\pi} \int_0^{2\pi} \int_0^{\pi/2} \hat{\rho}(\theta'_i, \theta'_o, |\phi'_o - \phi'_i|) \sin \theta'_i d\theta'_i d\phi'_i d\phi'_o \\
&= \sqrt{2\pi} \int_0^{2\pi} \int_0^{\pi/2} \hat{\rho}(\theta'_i, \theta'_o, |\phi|) \sin \theta'_i d\theta'_i d\phi
\end{aligned} \tag{27}$$

In the last line, we have set $\phi = \phi'_o - \phi'_i$ and integrated over ϕ'_o , obtaining a factor of 2π . Substituting in the expansion of equation 26, we obtain

$$\begin{aligned}
B &= B_{000}(\theta'_o) C_{000} \\
&= \hat{\rho}_{00}(\theta'_o) \sqrt{\frac{1}{2\pi}} \Lambda_0^{-1} \\
&= \Lambda_0^{-1} \int_0^{2\pi} \int_0^{\pi/2} \hat{\rho}(\theta'_i, \theta'_o, |\phi|) \sin \theta'_i d\theta'_i d\phi,
\end{aligned} \tag{28}$$

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

We next consider factorization of the light field to estimate both the lighting and the BRDF. An examination of the formulas in equation 18 shows that we will indeed be able to estimate all the lighting coefficients, provided the BRDF terms $\tilde{\rho}_{0p0}$ do not vanish. We will thus be able to determine that the lighting is uniform, i.e. that all the lighting coefficients L_{lm} vanish unless $l = m = 0$. However, once we do this, the factorization will be ill-posed since BRDF estimation is ill-posed. To summarize, we have derived and verified analytic formulae for the case of uniform lighting. BRDF estimation is extremely ill-posed, and only the lowest-frequency terms of the BRDF can be found. Under uniform lighting, a mirror surface and a Lambertian surface will look identical, and cannot be distinguished. With respect to factoring the light field when the lighting is uniform, we will correctly be able to determine that the illumination is indeed uniform, but the factorization will remain ill-posed since further information regarding the BRDF cannot be obtained. In signal processing terms, the input signal is constant and therefore has no amplitude along nonzero frequencies of the BRDF filter. Therefore, these nonzero frequencies of the BRDF filter cannot be estimated.

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

3.4 Mirror BRDF

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

$$\begin{aligned}
\hat{\rho} &= \delta(\cos \theta'_o - \cos \theta'_i) \delta(\phi'_o - \phi'_i \pm \pi) \\
\hat{\rho}_{lpq} &= \int_0^{2\pi} \int_0^{2\pi} \int_0^{\pi/2} \int_0^{\pi/2} \delta(\cos \theta'_o - \cos \theta'_i) \delta(\phi'_o - \phi'_i \pm \pi) Y_{lq}(\theta'_i, \phi'_i) Y_{pq}^*(\theta'_o, \phi'_o) \sin \theta'_i \sin \theta'_o d\theta'_i d\theta'_o d\phi'_i d\phi'_o \\
&= \int_0^{2\pi} \int_0^{\pi/2} Y_{pq}^*(\theta'_i, \phi'_i \pm \pi) Y_{lq}(\theta'_i, \phi'_i) \sin \theta'_i d\theta'_i d\phi'_i \\
&= (-1)^q \delta_{lp}.
\end{aligned} \tag{29}$$

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

The reflected light field and BRDF are now related by

$$\begin{aligned}
B_{lmpq} &= \Lambda_l (-1)^q \delta_{lp} L_{lm} \\
L_{lm} &= \Lambda_l^{-1} (-1)^q B_{lmlq}.
\end{aligned} \tag{30}$$

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

It is easy to verify equation 30 from first principles. We may write down the appropriate equations in angular space and then expand them in terms of spherical harmonics.

$$\begin{aligned}
B(\alpha, \beta, \theta'_o, \phi'_o) &= L(R_{\alpha, \beta}(\theta'_o, \phi'_o \pm \pi)) \\
&= \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm} Y_{lm}(R_{\alpha, \beta}(\theta'_o, \phi'_o \pm \pi)) \\
&= \Lambda_l (-1)^q \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm} \left(\Lambda_l^{-1} D_{mq}^l(\alpha, \beta) Y_{lq}(\theta'_o, \phi'_o) \right),
\end{aligned} \tag{31}$$

from which equation 30 follows immediately.

Factorization: The result for factoring a light field with a mirror BRDF is interesting. We first note that, unlike in our previous examples, there is no need to perform the symmetrizing

transformations in equation 11 since the transfer function is already symmetric with respect to indices l and p ($\delta_{lp} = \delta_{pl}$). We next note that equation 18 seems to indicate that the factorization is very ill-posed. Indeed, both the denominators, B_{00l0} and B_{lm00} , vanish unless $l = 0$. In fact, it is possible to explain the ambiguity by constructing a suitable BRDF,

$$\begin{aligned}\hat{\rho}_l &= f(l)(-1)^q \delta_{lp} \\ \hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{q=-l}^l f(l) Y_{lq}^*(\theta'_i, \phi'_i) Y_{lq}(\theta'_o, \phi'_o),\end{aligned}\tag{32}$$

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

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

Note that if the lighting coefficients do not vanish, we will indeed be able to learn that the BRDF has the form in equation 32. However, we will not be able to make further progress without additional assumptions about the form of the BRDF, i.e. of the function $f(l)$. In certain applications, we may want to turn this ambiguity to our advantage by selecting $f(l)$ appropriately to give a simple form for the BRDF or the lighting, without affecting the reflected light field. The ambiguity, and its use in simplifying the form for the reflected light field, are common to many reflective BRDFs, and we will encounter this issue again for the Phong and microfacet models.

Reparameterization by Reflection Vector: For reflective BRDFs, it is often convenient to reparameterize by the reflection vector, as discussed at the end of section 2.1. The transfer function can then be written simply as a function of the incident angle (with respect to the reflection vector), and is still a delta function. Since there is no dependence on the outgoing angle after reparameterization, we obtain

$$\begin{aligned}\hat{\rho}(\theta'_i) &= \frac{\delta(\cos(0) - \cos \theta'_i)}{2\pi} \\ \hat{\rho}_l &= Y_{l0}(0) = \Lambda_l^{-1} \\ B_{lm} &= L_{lm}.\end{aligned}\tag{33}$$

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

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

3.5 Lambertian BRDF

For a Lambertian surface, the BRDF is a constant, corresponding to the surface albedo—which for purposes of energy conservation must not be greater than $1/\pi$. For simplicity, we will omit this constant in what follows. The transfer function is a *clamped cosine* since it is equal to the cosine of the incident angle over the upper hemisphere when $\cos \theta'_i > 0$ and is equal to 0 over the lower hemisphere when $\cos \theta'_i < 0$. A graph of this function, along with spherical harmonic approximations to it up to order 4 is in figure 2.




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

Since the reflected light field is proportional to the incident irradiance and is equal for all outgoing directions, we will drop the outgoing angular dependence and use the form of the convolution formula given in equation 15,

$$\begin{aligned}\hat{\rho} = \max(\cos \theta'_i, 0) &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta'_i) \\ B_{lm} &= \Lambda_l \hat{\rho}_l L_{lm}.\end{aligned}\tag{34}$$

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

We will need to use many formulas for representing integrals of spherical harmonics, for which a reference [10] will be useful. The spherical harmonic coefficients are given by

$$\hat{\rho}_l = 2\pi \int_0^{\frac{\pi}{2}} Y_{l0}(\theta'_i) \cos \theta'_i \sin \theta'_i d\theta'_i.\tag{35}$$

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

$$\hat{\rho}_l = 2\pi \Lambda_l^{-1} \int_0^1 P_l(u) P_1(u) du.\tag{36}$$

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

$$\int_{-1}^1 P_a(u) P_b(u) du = \frac{2}{2a+1} \delta_{a,b}.\tag{37}$$

From this, we can establish some results about equation 36. When l is equal to 1, the integral evaluates to half the norm above, i.e. $1/3$. When l is odd but greater than 1, the integral in equation 36 vanishes. This is because, for $a = l$ and $b = 1$, we can break the left-hand side of equation 37 using the oddness of a and b into two equal integrals from $[-1, 0]$ and $[0, 1]$. Therefore, both of these integrals must vanish, and the latter integral is the right-hand integral in equation 36. When l is even, the required formula is given by manipulating equation 20 in chapter 5 of MacRobert[10]. Putting it all together, we obtain,

$$\begin{aligned}l = 1 \quad \hat{\rho}_l &= \sqrt{\frac{\pi}{3}} \\ l > 1, \text{ odd} \quad \hat{\rho}_l &= 0 \\ l \text{ even} \quad \hat{\rho}_l &= 2\pi \sqrt{\frac{2l+1}{4\pi}} \frac{(-1)^{\frac{l}{2}-1}}{(l+2)(l-1)} \left[\frac{l!}{2^l (\frac{l}{2}!)^2} \right].\end{aligned}\tag{38}$$

This formula is plotted in figure 3.




Figure 3: The solid line is a plot of $\hat{\rho}_l$ versus l , as per equation 38. It can be seen that odd terms with $l > 1$ have $\hat{\rho}_l = 0$. Also, as l increases, the BRDF coefficients decay rapidly.

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

$$\begin{aligned}
 \Lambda_0 \hat{\rho}_0 &= 3.142 \\
 \Lambda_1 \hat{\rho}_1 &= 2.094 \\
 \Lambda_2 \hat{\rho}_2 &= 0.785 \\
 \Lambda_3 \hat{\rho}_3 &= 0 \\
 \Lambda_4 \hat{\rho}_4 &= -0.131 \\
 \Lambda_5 \hat{\rho}_5 &= 0 \\
 \Lambda_6 \hat{\rho}_6 &= 0.049
 \end{aligned} \tag{39}$$

We see that already for $l = 4$, the coefficient is only about 4% of what it is for $l = 0$. In fact, it can be shown that over 99% of the energy of the BRDF filter is captured by $l \leq 2$. By considering the fact that the lighting must remain positive everywhere [1], similar worst-case bounds can be shown for the approximation of the reflected light field by $l \leq 2$. **Therefore, the irradiance, or equivalently, the reflected light field from a Lambertian surface can be well approximated using only the first 9 terms of its spherical harmonic expansion**—1 term with order 0, 3 terms with order 1, and 5 terms with order 2. Note that the single order 0 mode

Y_{00} is a constant, the 3 order 1 modes are linear functions of the Cartesian coordinates—in real form, they are simply x , y , and z —while the 5 order 2 modes are quadratic functions of the Cartesian coordinates. Therefore, the irradiance can be well approximated as a quadratic polynomial of the Cartesian coordinates of the surface normal vector.

We first consider illumination estimation, or the inverse lighting problem. The fact that the odd frequencies greater than 1 of the BRDF vanish means the inverse lighting problem is formally ill-posed for a Lambertian surface. The filter zeros the odd frequencies of the input signal, so these terms cannot be estimated from images of a convex Lambertian object. This observation corrects a commonly held notion (see Preissendorfer [14], volume 2, pages 143–151) that radiance and irradiance are equivalent in the sense that irradiance can be formally inverted to recover the radiance. Different radiance distributions can give rise to the same irradiance distribution. For further details, see [16]. Moreover, in practical applications, we can robustly estimate only the first 9 coefficients of the incident illumination, those with $l \leq 2$. Thus, inverse lighting from a Lambertian surface is not just formally ill-posed for odd frequencies, but very ill-conditioned for even frequencies. This result explains the ill-conditioning observed by Marschner and Greenberg [11] in estimating the lighting from a surface assumed to be Lambertian.

The 9 parameter approximation also gives rise to a simple algorithm [17] for estimating the illumination at high angular resolution from surfaces having both diffuse and specular components. The diffuse component of the reflected light field is subtracted out using the 9 parameter approximation for Lambertian surfaces. The object is then treated as a gazing sphere, with the illumination recovered from the specular component alone. A consistency condition ensures that the high frequency lighting recovered from the specular BRDF component is indeed consistent with the low frequency lighting used to subtract out the diffuse component of the reflected light field.

Our results are also in accordance with the perception literature, such as Land’s retinex theory [8]. It is common in visual perception to associate lighting effects with low frequency variation, and texture with high frequency variation. Our results formalize this observation, showing that distant lighting effects can produce only low frequency variation, with respect to orientation, in the intensity of a homogeneous convex Lambertian surface. Therefore, it should be possible to estimate high frequency texture independently of the lighting. However, for accurate computational estimates as required for instance in computer graphics, there is still an ambiguity between low frequency texture and lighting-related effects.

Since the approximation of Lambertian surfaces is commonly used in graphics and vision, the above results are of interest for many other problems. For instance, we can simply use the first 9 terms of the lighting to compute the irradiance, i.e. the shading on a Lambertian surface. Further implementation details on computing and rendering with *irradiance environment maps* are found in [15]. The 9 parameter approximation also means that images of a diffuse object under all possible illumination conditions lie close to a 9D subspace. This is a step toward explaining many

previous empirical observations of the low-dimensional effects of lighting made in the computer vision community, such as by Hallinan [7] and Epstein et al. [6]. Basri and Jacobs [1] have derived a Lambertian formula similar to ours, and have applied this result to lighting invariant recognition, and more recently to photometric stereo under general unknown lighting [2].

3.6 Phong BRDF

The normalized Phong transfer function is

$$\hat{\rho} = \frac{s+1}{2\pi} (\vec{R} \cdot \vec{L})^s, \quad (40)$$

where \vec{R} is the reflection vector, \vec{L} is the direction to the light source, and s is the *shininess*, or Phong exponent. The normalization ensures that the Phong lobe has unit energy. Technically, we must also zero the BRDF when the light vector is not in the upper hemisphere. However, the Phong BRDF is not physically based anyway, so others have often ignored this boundary effect, and we will do the same. This allows us to reparameterize by the reflection vector \vec{R} , making the transformations outlined at the end of section 2.1. In particular, $\vec{R} \cdot \vec{L} \rightarrow \cos \theta'_i$. Since the BRDF transfer function depends only on $\vec{R} \cdot \vec{L} = \cos \theta'_i$, the Phong BRDF after reparameterization is mathematically analogous to the Lambertian BRDF just discussed (they are both *radially symmetric*). In particular, equation 34 holds. However, note that while the Phong BRDF is mathematically analogous to the Lambertian case, it is not physically similar since we have reparameterized by the reflection vector. The BRDF coefficients depend on s , and are given by

$$\begin{aligned} \hat{\rho}_l &= (s+1) \int_0^{\pi/2} [\cos \theta'_i]^s Y_{l0}(\theta'_i) \sin \theta'_i d\theta'_i \\ B_{lm} &= \Lambda_l \hat{\rho}_l L_{lm}. \end{aligned} \quad (41)$$

To solve this integral, we substitute $u = \cos \theta'_i$ in equation 41. We also note that $Y_{l0}(\theta'_i) = \Lambda_l^{-1} P_l(\cos \theta'_i)$, where P_l is the legendre polynomial of order l . Then, equation 41 becomes

$$\hat{\rho}_l = \Lambda_l^{-1} (s+1) \int_0^1 u^s P_l(u) du \quad (42)$$

An analytic formula is given by MacRobert [10] in equations 19 and 20 of chapter 5.

$$\begin{aligned} \text{ODD } l \quad \Lambda_l \hat{\rho}_l &= \frac{(s+1)(s-1)(s-3)\dots(s-l+2)}{(s+l+1)(s+l-1)\dots(s+2)} \\ \text{EVEN } l \quad \Lambda_l \hat{\rho}_l &= \frac{s(s-2)\dots(s-l+2)}{(s+l+1)(s+l-1)\dots(s+3)} \end{aligned} \quad (43)$$

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

$$\Lambda_l \hat{\rho}_l = \frac{\left[\Gamma\left(\frac{s}{2}\right)\right]^2}{\Gamma\left(\frac{s}{2} - \frac{l}{2}\right) \Gamma\left(\frac{s}{2} + \frac{l}{2}\right)}. \quad (44)$$




Figure 4: Numerical plots of the Phong coefficients $\Lambda_l \hat{\rho}_l$, as defined by equation 43. The solid lines are the gaussian filter approximations in equation 46. As the Phong exponent s increases, corresponding to increasing the angular width of the BRDF filter, the frequency width of the BRDF filter decreases.

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

$$\log(\Lambda_l \hat{\rho}_l) = -l^2 \left(\frac{1}{2s} - \frac{1}{2s^2} \right) + O\left(\frac{l^4}{s^2}\right). \quad (45)$$

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

$$\Lambda_l \hat{\rho}_l \approx \exp\left[-\frac{l^2}{2s}\right]. \quad (46)$$

The coefficients fall off as a gaussian with width of order \sqrt{s} . The Phong BRDF behaves in the frequency domain like a gaussian filter, with the filter width controlled by the shininess. Therefore, inverse lighting calculations will be well-conditioned only up to order \sqrt{s} . As s approaches infinity, $\Lambda_l \hat{\rho}_l = 1$, and the frequency spectrum becomes constant, corresponding to a perfect mirror. Note that the frequency domain width of the filter varies inversely with the angular domain extent of the BRDF filter. A plot of the BRDF coefficients and the approximation in equation 46 is shown in figure 4.

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

Associativity of Convolution: With respect to factorization of light fields with surfaces approximated by Phong BRDFs, we obtain the same results for reflective surfaces as we did for mirror BRDFs. From the form of equation 41, it is clear that there is an unrecoverable scale factor for each order l . In physical terms, using the real BRDF and real lighting is equivalent to using a blurred version of the illumination and a mirror BRDF. In signal processing terminology, associativity of convolution allows us to sharpen the BRDF while blurring the illumination without affecting the reflected light field. To be more precise, we may rewrite equation 41 as

$$\begin{aligned}
 L'_{lm} &= \Lambda_l \hat{\rho}_l L_{lm} \\
 \Lambda_l \hat{\rho}'_l &= 1 \\
 B_{lm} &= \Lambda_l \hat{\rho}'_l L'_{lm} \\
 &= L'_{lm}.
 \end{aligned} \tag{47}$$

These equations say that we may blur the illumination using the BRDF filter, while treating the BRDF as a mirror. This formulation also allows us to analyze the conditioning in estimating the parameters of a Phong BRDF model under arbitrary illumination. The form of equation 46 for the Phong BRDF coefficients indicates that for $l \ll \sqrt{s}$, the effects of the filtering in equation 47 are minimal. The BRDF filter passes through virtually all the low-frequency energy with practically no attenuation. Thus, under low-frequency lighting, the reflected light field is essentially independent of the Phong exponent s . This means that **under low-frequency lighting, estimation of the exponent s of a Phong BRDF model is ill-conditioned**. In physical terms, it is difficult to determine the shininess of an object under diffuse lighting. In order to do so robustly, we must have high-frequency lighting components like directional sources. This observation holds for many reflective BRDF models. In particular, we shall see that a similar result can be derived for microfacet BRDFs; estimation of the surface roughness is ill-conditioned under low-frequency lighting.

3.7 Microfacet BRDF

Consider a simplified Torrance-Sparrow [18] model,

$$\rho = \frac{1}{4\pi\sigma^2 \cos \theta'_i \cos \theta'_o} \exp \left[- \left(\frac{\theta'_h}{\sigma} \right)^2 \right]. \tag{48}$$

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

It is convenient to reparameterize by the reflection vector, as we did for the Phong BRDF. However, it is important to note that microfacet BRDFs are not symmetric about the reflection vector. Unlike for the Phong BRDF, there is a preferred direction, determined by the exitant

angle. However, it can be shown by Taylor-series expansions and verified numerically that it is often reasonable to treat the microfacet BRDF using the same machinery as for the Phong case, assuming no outgoing angular dependence. Even under this assumption, it is somewhat difficult to derive precise analytic formulae. However, we may make good approximations.

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

$$\hat{\rho} = \sum_{l=0}^{\infty} \sum_{q=-l}^l \hat{\rho}_{lq}(\theta'_o) Y_{lq}(\theta'_i, \phi'_i). \quad (49)$$

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

3.7.1 Normal Exitance

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

$$\hat{\rho}_l = 2\pi \int_0^{\pi/2} \frac{\exp[-\theta_i'^2/4\sigma^2]}{4\pi\sigma^2} Y_{l0}(\theta'_i) \sin \theta'_i d\theta'_i. \quad (50)$$

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

$$Y_{l0}(t) = \Lambda_l^{-1} \left(1 - \frac{l(l+1)}{4} t^2 + O(t^4) \right). \quad (51)$$

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

$$Y_{l0}(t) \sim \Lambda_l^{-1} \left(\frac{1}{\sqrt{t}} \cos[(l+1/2)t - \pi/4] \right). \quad (52)$$

To integrate equation 50, we substitute $u = \theta'_i/2\sigma$. Then, $\theta'_i = 2\sigma u$. Assuming $\sigma \ll 1$, as it is for most surfaces, the upper limit of the integral becomes infinite, and $\sin \theta'_i d\theta'_i = \theta'_i d\theta'_i = 4\sigma^2 u du$.

$$\hat{\rho}_l = \int_0^{\infty} 2e^{-u^2} Y_{l0}(2\sigma u) u du \quad (53)$$

We therefore set $t = 2\sigma u$ in equations 51 and 52. When $\sigma l \ll 1$, we use equation 51 to obtain to $O([\sigma l]^4)$,

$$\Lambda_l \hat{\rho}_l = \left(\int_0^{\infty} 2ue^{-u^2} du - (\sigma l)^2 \int_0^{\infty} 2u^3 e^{-u^2} du \right). \quad (54)$$

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

$$\Lambda_l \hat{\rho}_l = 1 - (\sigma l)^2 + O([\sigma l]^4). \quad (55)$$

We note that these are the first terms of the Taylor series expansion of $\exp[-(\sigma l)^2]$. When $\sigma l \gg 1$, we use equation 52 to obtain (Φ is a phase that encapsulates the lower-order terms)

$$\Lambda_l \hat{\rho}_l \sim \int_0^\infty e^{-u^2} \sqrt{u} \cos[(2\sigma l)u + \Phi] du. \quad (56)$$

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

$$\Lambda_l \hat{\rho}_l \approx \exp[-(\sigma l)^2]. \quad (57)$$

We have also verified this result numerically.

For normal exitance, the BRDF is symmetric about the reflection vector and gaussian, so in that case, equation 57 simply states that even in the spherical-harmonic basis, the frequency spectrum of a gaussian is also approximately gaussian, with the frequency width related to the reciprocal of the angular width. For non-normal exitance, we will see that this is still a good approximation. The corrections are small except when l is large (corresponding to directional light sources), and at large viewing angles. These statements are made more precise in the next subsection.

3.7.2 Non-Normal Exitance

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

$$\theta'_h{}^2 = \left(\frac{\theta'_i}{2}\right)^2 (1 + \sin^2 \phi'_i \tan^2 \theta'_o) \quad (58)$$

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

We first consider the coefficients for the isotropic or azimuthally symmetric term Y_{10} . For l small, we can expand as $1 - (\sigma l)^2 + O(\sigma l)^4$. Now, the constant term is just the area of the microfacet distribution, but when normalized by $\cos \theta'_o$ must integrate to 1. Therefore, corrections are at least of $O(\theta'_o{}^2(\sigma l)^2)$. In fact, it is possible to obtain a tighter bound by deeper analysis.

Therefore, these corrections are not significant for small outgoing angles, and within that domain, we may use equation 57 as a good approximation. It should be noted that in practical applications, measurements made at wide or near-grazing angles are given low confidence anyway.

To make this more concrete, consider what happens to equation 53. Considering the first term in the Taylor-series expansion, and including the $\cos \theta'_o$ term in the denominator, we get

$$\hat{\rho}_l = \frac{1}{\cos \theta'_o} \int_0^\infty 2ue^{-u^2} Y_{l0}(2\sigma u) \left(1 - \frac{u^2 \tan^2 \theta'_o}{2}\right) du. \quad (59)$$

The factor of 2 in the denominator of the last (correction) term is to take care of the integration of $\sin^2 \phi'_i$. Next, we expand both the cosine and tangent functions about $\theta'_o = 0$, i.e. $\cos \theta'_o = 1 - \theta'^2_o/2 + O(\theta'^4_o)$ and $\tan^2 \theta'_o = \theta'^2_o + O(\theta'^4_o)$. Upon doing this, it can be verified that, as expected from physical arguments, there is no net correction to the integral and equation 59 evaluates to equation 55.

Now, consider the anisotropic terms Y_{lq} with $q \neq 0$. If we Taylor expand as in equation 59, we are again going to get something at least of $O(\theta'^2_o)$, and because the Y_{lq} constant term vanishes, there will be another factor of σ . In fact, if we Taylor-expand, we get an expansion of the form $\sin^2 \phi'_i + \sin^4 \phi'_i + \dots$ and it is clear that the azimuthal integral against Y_{lq} vanishes unless there is a term of type $\sin^q \phi'_i$. Therefore, the corrections are actually $O(\theta'^q_o \sigma)$. What this means in practice is that the azimuthally dependent corrections are only important for large viewing angles and large orders l (l must be large if q is large, and for small q , $\theta'^q_o \sigma$ is small for $\sigma \ll 1$ even if θ'_o is large).

But, that situation will arise only for observations made at large angles of lights that have broad spectra, i.e. directional sources. Therefore, these should be treated separately, by separating the illumination into slow and fast varying components. Equation 57 in the frequency domain is a very good approximation for the slow-varying lighting component, while we may approximate the fast-varying lighting component using one or more directional sources in the angular domain.

3.7.3 Conditioning Properties

Since equation 57 has many similarities to the equations for the Phong BRDF, most of those results apply here too. In particular, under low-frequency lighting, there is an ambiguity with respect to estimation of the surface roughness σ . Also, inverse lighting is well-conditioned only up to order σ^{-1} . With respect to factorization, there are ambiguities between illumination and reflectance, similar to those for mirror and Phong BRDFs. Specifically, we may blur the illumination while sharpening the BRDF. However, it is important to note that while these ambiguities are exact for Phong and mirror BRDFs, they are only a good approximation for microfacet BRDFs since equation 57 does not hold at grazing angles or for high-frequency lighting distributions. In these cases, the ambiguity can be broken, and we have used this fact in an algorithm to simultaneously determine the lighting and the parameters of a microfacet model [17].

In summary, while it is difficult to derive precise analytic formulae, we can derive good approximations to the frequency-space behavior of a microfacet BRDF model. The results are rather

similar to those for the Phong BRDF with the Phong exponent s replaced by the physically-based surface roughness parameter σ .

In this subsection, we have seen analytic formulae derived for a variety of common lighting and BRDF models, demonstrating the implications of the theoretical analysis. We end this paper by summarizing our key contributions, and discussing avenues for future work.

4 Conclusions and Future Work

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

In this part of the paper, we have derived analytic formulae for the frequency spectra of many common BRDF and lighting models, and have demonstrated the implications for inverse problems such as lighting recovery, BRDF recovery, light field factorization. We have shown in frequency-space why a gazing sphere is well-suited for recovering the lighting—the frequency spectrum of the mirror BRDF (a delta function) is constant—and why a directional source is well-suited for recovering the BRDF—we are estimating the BRDF filter by considering its impulse response. With the aid of our theory, we have been able to quantitatively determine the well-posedness and conditioning of many inverse problems. The ill-conditioning observed by Marschner and Greenberg [11] in estimating the lighting from a Lambertian surface has been explained by showing that only the first 9 coefficients of the lighting can robustly be recovered, and we have shown that factorization of lighting effects and low-frequency texture is ambiguous. All these results indicate that the theory provides a useful analytical tool for studying the properties of inverse problems.

We believe this formal analysis of inverse problems will be increasingly important in deriving robust and efficient algorithms in the future. While we have made a first step in this direction, other issues such as how our results change when we have only a limited fraction of the reflected light field available, or can move our viewpoint only in a narrow range, need to be studied. Finally, we believe the formal analysis of inverse problems is of significant and growing interest in many areas of computer graphics, computer vision, and visual perception.

References

- [1] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. In *ICCV 01*, pages 383–390, 2001.
- [2] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. In *CVPR 01*, pages II–374–II–381, 2001.
- [3] B. Cabral, N. Max, and R. Springmeyer. Bidirectional reflection functions from surface bump maps. In *SIGGRAPH 87*, pages 273–281, 1987.
- [4] B. Cabral, M. Olano, and P. Nemeč. Reflection space image based rendering. In *SIGGRAPH 99*, pages 165–170, 1999.
- [5] M. D’Zmura. *Computational Models of Visual Processing*, chapter Shading Ambiguity: Reflectance and Illumination, pages 187–207. MIT Press, 1991.
- [6] R. Epstein, P.W. Hallinan, and A. Yuille. 5 plus or minus 2 eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *IEEE Workshop on Physics-Based Modeling in Computer Vision*, pages 108–116, 1995.
- [7] P. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *CVPR 94*, pages 995–999, 1994.
- [8] E. Land and J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.
- [9] R. Lu, J. Koenderink, and A. Kappers. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics*, 37(25):5974–5984, 1998.
- [10] T.M. MacRobert. *Spherical harmonics; an elementary treatise on harmonic functions, with applications*. Dover Publications, 1948.
- [11] S.R. Marschner and D.P. Greenberg. Inverse lighting for photography. In *Fifth Color Imaging Conference*, pages 262–265, 1997.
- [12] 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.
- [13] 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.

- [14] R.W. Preisendorfer. *Hydrologic Optics*. US Dept Commerce, 1976.
- [15] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *SIGGRAPH 01*, pages 497–500, 2001.
- [16] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *JOSA A*, 18(10):2448–2459, 2001.
- [17] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, pages 117–128, 2001.
- [18] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967.