A Signal-Processing Framework for Reflection

RAVI RAMAMOORTHI
Columbia University
and
PAT HANRAHAN
Stanford University

We present a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. This analysis is of theoretical interest in both graphics and vision and is also of practical importance in many computer graphics problems—for instance, in determining lighting distributions and bidirectional reflectance distribution functions (BRDFs), in rendering with environment maps, and in image-based rendering. It is well known that under our assumptions, the reflection operator behaves qualitatively like a convolution. In this paper, we formalize these notions, showing that the reflected light field can be thought of in a precise quantitative way as obtained by convolving the lighting and BRDF, i.e. by filtering the incident illumination using the BRDF. Mathematically, we are able to express the frequency-space coefficients of the reflected light field as a product of the spherical harmonic coefficients of the illumination and the BRDF. These results are of practical importance in determining the well-posedness and conditioning of problems in *inverse rendering*—estimation of BRDF and lighting parameters from real photographs. Furthermore, we are able to derive analytic formulae for the spherical harmonic coefficients of many common BRDF and lighting models. From this formal analysis, we are able to determine precise conditions under which estimation of BRDFs and lighting distributions are well posed and well-conditioned. Our mathematical analysis also has implications for *forward rendering*—especially the efficient rendering of objects under complex lighting conditions specified by environment maps. The results, especially the analytic formulae derived for Lambertian surfaces, are also relevant in *computer vision* in the areas of recognition, photometric stereo and structure from motion.

Categories and Subject Descriptors: I.3.0 [**Computer Graphics**]: General; I.4.0 [**Image Processing and Computer Vision**]: General

General Terms: Theory

Additional Key Words and Phrases: BRDF, environment maps, Fourier analysis, illumination, inverse rendering, reflection, signal processing, spherical harmonics

1. INTRODUCTION

The study of reflection is of fundamental importance in both computer graphics and vision. In computer graphics, the interaction between the incident illumination and the BRDF (bidirectional reflectance distribution function) is a basic building block in most rendering algorithms. In computer vision, we often want to undo the effects of the reflection operator: to invert the interaction between the BRDF and

This work was supported in part by grants from the National Science Foundation (#0085864, #0305322 and 0430528) and Intel Corporation.
Authors' addresses: R. Ramamoorthi, Department of Computer Science, Columbia University, 450 CS Building, 500 W. 120th Street, New York, NY 10027; email: ravir@cs.columbia.edu; P. Hanrahan, Stanford University, Gates Computer Science Building, Room 370 3B, Stanford, CA 94305-4070; email: hanrahan@cs.stanford.edu.

lighting. In other words, we often want to perform *inverse rendering*—the estimation of material and lighting properties from real photographs. It should be noted that inverse rendering is also of increasing importance in graphics, where we wish to obtain accurate input models for (forward) rendering. The goal of this article is to fill some of the gaps in the formal analysis of reflection, thereby putting algorithms for both forward and inverse rendering on a sounder mathematical foundation.

In computer vision, previous theoretical work has mainly focussed on the problem of estimating shape from images, with relatively little work on estimating material properties or lighting. In computer graphics, the theory for global illumination calculations has been fairly well developed. The foundation for this analysis is the rendering equation [Kajiya 1986], which has subsequently been analyzed extensively [Arvo 1995]. However, there has been relatively little theoretical work on the simpler reflection equation, which deals with the direct illumination incident on a surface. We believe that this lack of a formal mathematical understanding of the properties of the reflection equation is one of the reasons why complex, realistic lighting environments and reflection functions are rarely used either in forward or inverse rendering.

It should be noted that there is a significant amount of qualitative knowledge regarding the properties of the reflection operator. For instance, for environment map prefiltering and rendering, we usually represent the diffuse reflection map at low resolutions [Miller and Hoffman 1984] since the reflection from a Lambertian surface *blurs* the illumination. Similarly, we realize that it is essentially impossible to accurately estimate the lighting from an image of a Lambertian surface; instead, we use mirror surfaces like gazing spheres. The goal of this paper is to formalize these observations and present a mathematical theory of reflection for general complex lighting environments and arbitrary BRDFs.

Specifically, we describe a signal-processing framework for analyzing the reflected light field from a homogeneous convex curved surface under distant illumination. Under these assumptions, we are able to derive an analytic formula for the reflected light field in terms of the spherical harmonic coefficients of the BRDF and the lighting. Our formulation leads to the following theoretical results:

Signal-Processing Framework for Reflection as Convolution: It has been observed qualitatively by Miller and Hoffman [1984], Cabral et al. [1987, 1999], Nayar et al. [1990], D'Zmura [1991] and others that the reflection operator behaves like a convolution. We are able to 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, or filtering the incident illumination using the BRDF. Mathematically, we are able to express the frequency-space coefficients of the reflected light field as a product of the spherical harmonic coefficients of the illumination and the BRDF. We believe this is a useful way of analyzing many forward and inverse problems. In particular, forward rendering can be viewed as *convolution* and inverse rendering as *deconvolution*.

Well-posedness and Conditioning of Forward and Inverse Problems: Inverse problems can be ill-posed—there may be several solutions. They are also often numerically ill-conditioned, which may make devising practical algorithms infeasible. From our 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 in inverse rendering. We expect fruitful areas of research to be those problems that are well-conditioned. Additional assumptions will likely be required to address ill-conditioned or ill-posed inverse problems. This analysis is also of interest for forward rendering. An ill-conditioned inverse problem corresponds to a forward problem where the final results are not sensitive to certain components of the initial conditions. This often allows us to approximate the initial conditions—such as the incident illumination—in a principled way, giving rise to more efficient forward rendering algorithms.

Recently, there has been considerable theoretical and practical work based on some of these results, both by ourselves and by other authors. For instance, we have demonstrated the practical applications to both forward [Ramamoorthi and Hanrahan 2001b, 2002] and inverse rendering [Ramamoorthi and Hanrahan 2001d], as well as theoretical results for analytically constructing the principal components of the space of images of a Lambertian object for lighting-insensitive recognition [Ramamoorthi 2002a]. Sloan et al. [2002, 2003] have demonstrated real-time rendering with low-frequency lighting and complex light transport. In computer vision, Basri and Jacobs [2003] have applied these ideas to lighting-insensitive recognition and photometric stereo [Basri and Jacobs 2001b]. Subsequently, a number of other computer vision researchers have also extended our results, developing new practical algorithms for recognition, shape reconstruction and structure from motion [Nillius and Eklundh 2003a, 2003b, 2004; Sato et al. 2003; Lee et al. 2001; Ho et al. 2003; Zhang and Samaras 2003; Simakov et al. 2003; Zhang et al. 2003].

Therefore, we feel that it is worthwhile to present a complete, detailed and unified account of the theoretical framework for archival purposes, fully describing the mathematical foundation underlying these previous articles. This paper builds on and extends previous theoretical work by us [Ramamoorthi and Hanrahan 2001a, 2001c, 2001d]. As discussed at the end of the next section, we have included a significant amount of new material and detail that has not previously appeared.

2. PREVIOUS WORK

In this section, we briefly discuss previous work. Since the reflection operator is of fundamental interest in a number of fields, the relevant previous work is fairly diverse.

Forward Rendering by Environment Mapping: The theoretical analysis in this article employs essentially the same assumptions typically made in rendering with environment maps: distant illumination—allowing the lighting to be represented by a single environment map—incident on curved surfaces. Blinn and Newell [1976] first used environment maps to efficiently find reflections of distant objects. The technique was greatly generalized by Miller and Hoffman [1984] and Greene [1986] who precomputed diffuse and specular reflection maps, allowing for images with complex realistic lighting and BRDFs to be synthesized. Cabral et al. [1987] later extended this general method to computing reflections from bump-mapped surfaces, and to computing environment-mapped images with more general BRDFs [Cabral et al. 1999]. Similar results for more general BRDFs were also obtained by Kautz et al. [2000] and Kautz and McCool [2000], building on previous work by Heidrich and Seidel [1999]. It should be noted that both Miller and Hoffman [1984], and Cabral et al. [1987, 1999] qualitatively described the reflection maps as obtained by convolving the lighting with the BRDF. Kautz et al. [2000] actually used convolution to implement reflection, but on somewhat distorted planar projections of the environment map, and without full theoretical justification. In this article, we will formalize these ideas, making the notion of convolution precise, and derive analytic formulae.

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

Perhaps the simplest inverse rendering method is the use of a mirror sphere to find the lighting, first introduced by Miller and Hoffman [1984]. A more sophisticated *inverse lighting* approach is that of Marschner and Greenberg [1997], who try to find the lighting under the assumption of a Lambertian BRDF. D'Zmura [1991] proposes estimating spherical harmonic coefficients of the lighting.

Most work in inverse rendering has focused on BRDF [Nicodemus et al. 1977] estimation. Recently, image-based BRDF measurement methods have been proposed in 2D by Lu et al. [1998] and in 3D by Marschner et al. [2000]. 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 [1992] or Torrance and Sparrow [1967]. Parametric models are often preferred in practice since they are compact, and are simpler to estimate. A number of methods [Dana et al. 1999; Debevec et al. 2000; Sato et al. 1997; Yu et al. 1999] have been proposed to estimate parametric BRDF models, often along with a modulating texture.

However, it should be noted that all of the methods described above use a single point source. One of the main goals of the theoretical analysis in this paper is to enable the use of inverse rendering with complex lighting. Recently, there has been some work in this area [Dror et al. 2001; Love 1997; Nishino et al. 2001; Sato et al. 1999; Sato and Ikeuchi 1994; Yu and Malik 1998], although many of those methods are specific to a particular illumination model. Using the theoretical analysis described in this article, we [Ramamoorthi and Hanrahan 2001d] have presented a general method for complex illumination, that handles the various components of the lighting and BRDF in a principled manner to allow for BRDF estimation under general lighting conditions. Furthermore, we will show that it is possible in theory to separately estimate the lighting and BRDF, up to a global scale.

Frequency-Space Representations: We show reflection to be a convolution and analyze it in frequency space. Since we will be primarily concerned with analyzing quantities like the BRDF and distant lighting, which can be parameterized as a function on the unit sphere, the appropriate frequency-space representations are spherical harmonics [Inui et al. 1990; Jackson 1975; MacRobert 1948]. The use of spherical harmonics to represent the illumination and BRDF was pioneered by Cabral et al. [1987]. D'Zmura [1991] analyzed reflection as a linear operator in terms of spherical harmonics, and discussed some resulting ambiguities between reflectance and illumination. We extend his work by explicitly deriving the frequency-space reflection equation (i.e. convolution formula), and by providing quantitative results for various special cases. Our use of spherical harmonics to represent the lighting is similar in some respects to previous methods such as that of Nimeroff et al. [1994] that use steerable linear basis functions. Spherical harmonics have also been used before in computer graphics for representing BRDFs by a number of other authors [Sillion et al. 1991; Westin et al. 1992].

Formal Analysis of Reflection: This article presents a formal study of the reflection operator. As such, our approach is similar in spirit to mathematical methods used to study inverse problems in other areas of radiative transfer and transport theory such as hydrologic optics [Preisendorfer 1976] and neutron scattering. See McCormick [1992] for a review.

Within computer graphics and vision, the closest previous theoretical work lies in the object recognition community, where there has been a significant amount of interest in characterizing the appearance of a surface under all possible illumination conditions, usually under the assumption of Lambertian reflection. For instance, Belhumeur and Kriegman [1998] have theoretically described this set of images in terms of an illumination cone, while empirical results have been obtained by Hallinan [1994] and Epstein et al. [1995]. These results suggest that the space spanned by images of a Lambertian object under all (distant) illumination conditions lies very close to a low-dimensional subspace. We will see that our theoretical analysis will help in explaining these observations, and in extending the predictions to arbitrary reflectance models. In independent work on face recognition, simultaneous with our own, Basri and Jacobs [2001a] have described Lambertian reflection as a convolution and obtained similar analytic results for that particular case.

This article builds on previous theoretical work by us on planar or flatland light fields [Ramamoorthi and Hanrahan 2001a], on the reflected light field from a Lambertian surface [Ramamoorthi and

Hanrahan 2001c], and on the theory for the general 3D case with isotropic BRDFs [Ramamoorthi and Hanrahan 2001d]. The goal of this article is to present a unified, complete and more detailed account of the theory in the general case. In this article, we describe a unified view of the 2D and 3D cases, including general anisotropic BRDFs, a number of alternative forms of the convolution formula, and the relationship to the theory of Fredholm integral equations of the first kind, which have not been discussed in earlier work. Some further details can be found in chapters 2 and 3 of the PhD thesis of the first author [Ramamoorthi 2002b].

3. REFLECTION EQUATION

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

3.1 Assumptions

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

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

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

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

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

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

Discussion: Our assumptions are very similar to those made in most interactive graphics applications, including environment map rendering algorithms such as those of Miller and Hoffman [1984] and Cabral et al. [1999]. Our assumptions also accord closely with those usually made in computer vision and inverse rendering. The only significant additional assumption is that of homogeneous surfaces. However, this is not particularly restrictive since spatially varying BRDFs are often approximated in practical graphics or vision applications by using a spatially varying texture that simply modulates one or more components of the BRDF. This can be incorporated into the ensuing theoretical analysis by merely multiplying the reflected light field by a texture dependent on surface position. We believe that our assumptions are a good approximation to many real-world situations, while being simple enough to treat analytically. Furthermore, it is likely that the insights obtained from the analysis in

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

Fig. 1. Notation.

this article will be applicable even in cases where the assumptions are not exactly satisfied. We have already demonstrated [Ramamoorthi and Hanrahan 2001d] that in practical applications, it is possible to extend methods derived from these assumptions to be applicable in an even more general context.

We now proceed to derive the reflection equation for the 2D and 3D case under the assumptions outlined above. Notation used in the article is listed in Figure 1. We will use two types of coordinates. Unprimed global coordinates denote angles with respect to a global reference frame. On the other hand, primed local coordinates denote angles with respect to the local reference frame, defined by the local surface normal and a tangent vector. These two coordinate systems are related simply by a rotation.

3.2 Flatland 2D Case

In this subsection, we consider the *flatland* or 2D case. This is equivalent to assuming that all surface normals, measurements, and illumination are restricted to a single plane (which can be thought of as normal to the axis of a cylindrical surface). Considering the 2D case allows us to explain the key concepts clearly, and show how they generalize to 3D. A diagram illustrating the key concepts for the planar 2D case is in Figure 2.

In local coordinates, we can write the reflection equation as

$$B(\mathbf{x}, \theta_o') = \int_{-\pi/2}^{\pi/2} L(\mathbf{x}, \theta_i')\rho(\theta_i', \theta_o')\cos\theta_i'\, d\theta_i'. \tag{1}$$

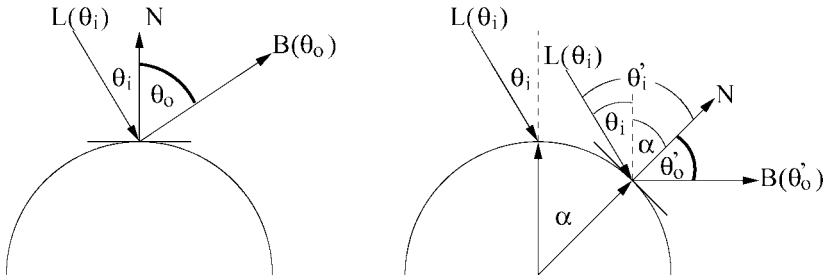

Fig. 2. Schematic of reflection in 2D. On the left, we show the situation with respect to one point on the surface (the north pole or 0° location, where global and local coordinates are the same). The right figure shows the effect of the surface orientation α. Different orientations of the surface correspond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

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

We now make a number of substitutions in Equation 1, based on our assumptions. First, consider the assumption of a convex surface. This ensures there is no shadowing or interreflection; this fact has implicitly been assumed in Equation 1. The reflected radiance therefore depends only on the distant illumination field L and the surface BRDF ρ. Next, consider the assumption of distant illumination. This implies that the reflected light field depends directly only on the surface orientation, as described by the surface normal $\mathbf{N}$, and does not directly depend on the position $\mathbf{x}$. We may therefore reparameterize the surface by its angular coordinates α, with $\mathbf{N} = [\sin\alpha, \cos\alpha]$; $B(\mathbf{x}, \theta_o') \rightarrow B(\alpha, \theta_o')$ and $L(\mathbf{x}, \theta_i') \rightarrow L(\alpha, \theta_i')$. The assumption of distant sources also allows us to represent the incident illumination by a single environment map for all surface positions using a single function L regardless of surface position. In other words, the lighting is a function only of the global incident angle, $L(\alpha, \theta_i') \rightarrow L(\theta_i)$. Finally, we define a transfer function $\hat{\rho} = \rho \cos\theta_i'$ to absorb the cosine term in the integrand. With these modifications, Equation 1 becomes

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

It is important to note that in Equation 2, we have mixed local (primed) and global (unprimed) coordinates. The lighting is a global function, and is naturally expressed in a global coordinate frame as a function of global angles. On the other hand, the BRDF is naturally expressed as a function of the local incident and reflected angles. When expressed in the local coordinate frame, the BRDF is the same everywhere for a homogeneous surface. Similarly, when expressed in the global coordinate frame, the lighting is the same everywhere, under the assumption of distant illumination. Integration can be conveniently done over either local or global coordinates, but the upper hemisphere is easier to keep track of in local coordinates.

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

we apply the inverse rotation: $R_{-\alpha}$. To summarize,

$$\begin{aligned} \theta_i &= R_\alpha(\theta_i') &= \alpha + \theta_i' \\ \theta_i' &= R_\alpha^{-1}(\theta_i) &= -\alpha + \theta_i. \end{aligned} \tag{3}$$

It should be noted that the signs of the various quantities are taken into account in Equation 3. Specifically, from the right of Figure 2, it is clear that $\mid \theta_i' \mid = \mid \theta_i \mid + \mid \alpha \mid$. In our sign convention, α is positive in Figure 2, while θ_i' and θ_i are negative. Substituting $\mid \theta_i' \mid = -\theta_i'$ and $\mid \theta_i \mid = -\theta_i$, we verify Equation 3.

With the help of Equation 3, we can express the incident angle dependence of Equation 2 in either local coordinates entirely, or global coordinates entirely. It should be noted that we always leave the outgoing angular dependence of the reflected light field in local coordinates in order to match the BRDF transfer function.

$$B(\alpha, \theta_o') = \int_{-\pi/2}^{\pi/2} L(R_\alpha(\theta_i'))\hat{\rho}(\theta_i', \theta_o')\, d\theta_i' \tag{4}$$

$$= \int_{-\pi/2+\alpha}^{\pi/2+\alpha} L(\theta_i)\, \hat{\rho}\left(R_\alpha^{-1}(\theta_i), \theta_o'\right) d\theta_i. \tag{5}$$

By plugging in the appropriate relations for the rotation operator from Equation 3, we can obtain

$$B(\alpha, \theta_o') = \int_{-\pi/2}^{\pi/2} L(\alpha + \theta_i')\hat{\rho}(\theta_i', \theta_o')\, d\theta_i' \tag{6}$$

$$= \int_{-\pi/2+\alpha}^{\pi/2+\alpha} L(\theta_i)\, \hat{\rho}\,(-\alpha + \theta_i, \theta_o')\, d\theta_i. \tag{7}$$

Interpretation as Convolution: Equations 6 and 7 (and the equivalent forms in Equations 4 and 5) are *convolutions*. The reflected light field can therefore be described formally as a convolution of the incident illumination and the BRDF transfer function. Equation 5 in global coordinates states that the reflected light field at a given surface orientation corresponds to *rotating* the BRDF to that orientation, and then integrating over the upper half-circle. In signal processing terms, the BRDF can be thought of as the filter, while the lighting is the input signal. The reflected light field is obtained by filtering the input signal (i.e. lighting) using the filter derived from the BRDF. Symmetrically, Equation 4 in local coordinates states that the reflected light field at a given surface orientation may be computed by *rotating* the lighting into the local coordinate system of the BRDF, and then doing the integration over the upper half-circle.

It is important to note that we are fundamentally dealing with rotations, as is brought out by Equations 4 and 5. For the 2D case, rotations are equivalent to translations, and Equations 6 and 7 are the familiar Equations for translational convolution. The main difficulty in formally generalizing the convolution interpretation to 3D is that the structure of rotations is more complex. In fact, we will need to consider a generalization of the notion of convolution in order to encompass rotational convolutions.

3.3 Generalization to 3D

The flatland development can be extended to 3D. In 3D, we can write the reflection equation, analogous to Equation 1, as

$$B(\mathbf{x}, \theta_o', \phi_o') = \int_{\Omega_i'} L(\mathbf{x}, \theta_i', \phi_i')\rho(\theta_i', \phi_i', \theta_o', \phi_o') \cos\theta_i'\, d\omega_i'. \tag{8}$$

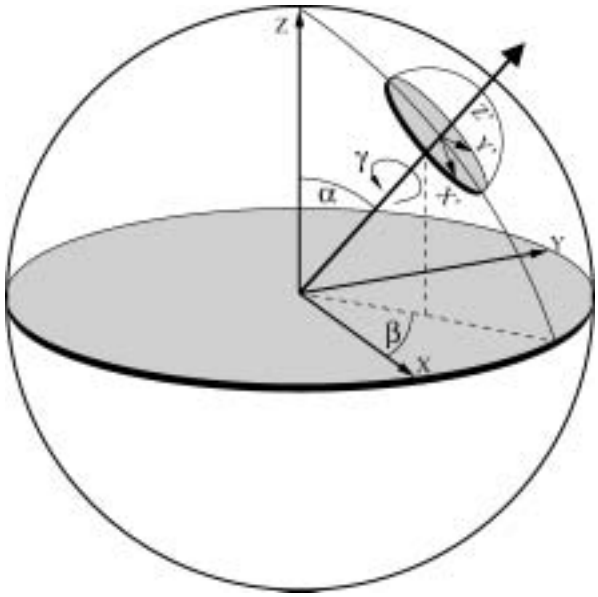

Fig. 3. Diagram showing how the rotation corresponding to (α, β, γ) transforms between local (primed) and global (unprimed) coordinates. The net rotation is composed of three independent rotations about Z, Y, and Z, with the angles α, β, and γ corresponding directly to the Euler angles.

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

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

$$B(\alpha, \beta, \gamma, \theta_o', \phi_o') = \int_{\Omega_i'} L(\theta_i, \phi_i)\hat{\rho}(\theta_i', \phi_i', \theta_o', \phi_o')\, d\omega_i'. \tag{9}$$

Rotations—Converting between Local and Global Coordinates: To do the integral above, we need to apply a rotation to convert between local and global coordinates, just as in 2D. The rotation operator is substantially more complicated in 3D, but the operations are conceptually very similar to those in flatland. The *north pole* $(0', 0')$ or $+$Z axis in local coordinates is the surface normal, and the corresponding global coordinates are (α, β). It can be verified that a rotation of the form $R_z(\beta)R_y(\alpha)$ correctly performs this transformation, where the subscript z denotes rotation about the Z axis and the subscript y denotes rotation about the Y axis. For full generality, the rotation between local and global coordinates should also specify the transformation of the local tangent frame, so the general rotation

operator is given by $R_{\alpha,\beta,\gamma} = R_z(\beta)R_y(\alpha)R_z(\gamma)$. This is essentially the Euler-angle representation of rotations in 3D. We may now summarize these results, obtaining the 3D equivalent of Equation 3,

$$\begin{aligned}(\theta_i, \phi_i) &= R_{\alpha,\beta,\gamma}(\theta_i', \phi_i') = R_z(\beta)R_y(\alpha)R_z(\gamma)\{\theta_i', \phi_i'\} \\ (\theta_i', \phi_i') &= R^{-1}_{\alpha,\beta,\gamma}(\theta_i, \phi_i) = R_z(-\gamma)R_y(-\alpha)R_z(-\beta)\{\theta_i, \phi_i\}.\end{aligned} \tag{10}$$

It is now straightforward to substitute these results into Equation 9, transforming the integral either entirely into local coordinates or entirely into global coordinates, and obtaining the 3D analogue of Equations 4 and 5,

$$\boxed{\begin{aligned}B(\alpha, \beta, \gamma, \theta_o', \phi_o') &= \int_{\Omega_i'} L(R_{\alpha,\beta,\gamma}(\theta_i', \phi_i'))\hat{\rho}(\theta_i', \phi_i', \theta_o', \phi_o')\, d\omega_i' \\ &= \int_{\Omega_i} L(\theta_i, \phi_i)\hat{\rho}\left(R^{-1}_{\alpha,\beta,\gamma}(\theta_i, \phi_i), \theta_o', \phi_o'\right) d\omega_i.\end{aligned}} \tag{11}$$

As we have written them, these equations depend on spherical coordinates. It might clarify matters somewhat to also present an alternate form in terms of rotations and unit vectors in a coordinate-independent way. We simply use R for the rotation, which could be written as a 3×3 rotation matrix, while ω_i and ω_o stand for unit vectors (3×1 column vectors) corresponding to the incident and outgoing directions (with primes added for local coordinates). Equation 11 may then be written as

$$B(R, \omega_o') = \int_{\Omega_i'} L(R\omega_i')\hat{\rho}(\omega_i', \omega_o')\, d\omega_i' \tag{12}$$

$$= \int_{\Omega_i} L(\omega_i)\hat{\rho}(R^{-1}\omega_i, \omega_o')\, d\omega_i, \tag{13}$$

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

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

$$(f \otimes g)(a) = \int_t f(T_a(t))g(t)\, dt. \tag{14}$$

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, and is a slightly simplified version of Equation 11.

4. FREQUENCY-SPACE ANALYSIS

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

4.1 Fourier Analysis in 2D

We now carry out a Fourier analysis of the 2D reflection equation. We will define the Fourier basis functions F_k (with complex conjugates $F_k^* = F_{-k}$) by

$$F_k(\theta) = \frac{1}{\sqrt{2\pi}} e^{Ik\theta}. \tag{15}$$

It is easy to see that the basis functions are orthonormal over the domain $[-\pi, \pi]$.

Decomposition into Fourier Series: We now consider the reflection equation, in the form of Equation 6. We will expand all quantities in terms of Fourier series.

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

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

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

$$\begin{aligned} L(\theta_i) = L(\alpha + \theta_i') &= \sum_{l=-\infty}^{\infty} L_l F_l(\alpha + \theta_i') \\ &= \sqrt{2\pi} \sum_{l=-\infty}^{\infty} L_l F_l(\alpha) F_l(\theta_i'). \end{aligned} \tag{16}$$

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

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

$$\begin{aligned} B(\alpha, \theta_o') &= \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} B_{lp} F_l(\alpha) F_p(\theta_o') \\ \hat{\rho}(\theta_i', \theta_o') &= \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} \hat{\rho}_{lp} F_l^*(\theta_i') F_p(\theta_o'). \end{aligned} \tag{17}$$

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

Fourier-Space Reflection Equation: We are now ready to write Equation 6 in terms of Fourier coefficients. For the purposes of summation, we want to avoid confusion of the indices for L and $\hat{\rho}$. For this purpose, we will use the indices L_l and $\hat{\rho}_{l'p}$. We now simply multiply out the expansions for L and $\hat{\rho}$. After taking the summations, and terms not depending on θ_i' outside the integral, Equation 6 now becomes

$$B(\alpha, \theta_o') = \sqrt{2\pi} \sum_{l=-\infty}^{\infty} \sum_{l'=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} L_l \hat{\rho}_{l'p} F_l(\alpha) F_p(\theta_o') \int_{-\pi}^{\pi} F_{l'}^*(\theta_i') F_l(\theta_i')\, d\theta_i'. \tag{18}$$

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

given as

$$\int_{-\pi}^{\pi} F_l'^*(\theta_i') F_l(\theta_i')\, d\theta_i' = \delta_{ll'}. \tag{19}$$

In other words, we can set $l' = l$ since terms not satisfying this condition vanish. Making this substitution in Equation 18, we obtain

$$B(\alpha, \theta_o') = \sqrt{2\pi} \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} L_l \hat{\rho}_{lp} F_l(\alpha) F_p(\theta_o'). \tag{20}$$

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

$$B_{lp} = \sqrt{2\pi} L_l \hat{\rho}_{lp}. \tag{21}$$

This result reiterates that the reflection equation can be viewed as a convolution of the incident illumination and BRDF, and becomes a simple product in Fourier space, with an analytic formula being given by Equation 21.

An alternative form of Equation 21 that may be more instructive results from holding the local outgoing angle fixed, instead of expanding it also in terms of Fourier coefficients, replacing the index p by the outgoing angle θ_o',

$$B_l(\theta_o') = \sqrt{2\pi} L_l \hat{\rho}_l(\theta_o'). \tag{22}$$

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

4.2 Spherical Harmonic Analysis in 3D

To extend our frequency-space analysis to 3D, we must consider the structure of rotations and vectors in 3D. In particular, the unit vectors corresponding to incident and reflected directions lie on a sphere of unit magnitude. The appropriate signal-processing tools for the sphere are spherical-harmonics, which are the equivalent for that domain to the Fourier series in 2D (on a circle). These basis functions arise in connection with many physical systems such as those found in quantum mechanics and electrodynamics. A summary of the properties of spherical harmonics can therefore be found in many standard physics textbooks [Inui et al. 1990; Jackson 1975; MacRobert 1948].

Key Properties of Spherical Harmonics: The spherical harmonic Y_{lm} is given by

$$\begin{aligned} N_{lm} &= \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \\ Y_{lm}(\theta, \phi) &= N_{lm} P_{lm}(\cos\theta) e^{Im\phi}, \end{aligned} \tag{23}$$

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. They may be written either as trigonometric functions of the spherical coordinates θ and ϕ or as polynomials of the cartesian components x, y and z, with $x^2 + y^2 + z^2 = 1$. In general, a spherical harmonic Y_{lm} is a polynomial of maximum degree l. Another useful relation is that $Y_{l-m} = (-1)^m Y_{lm}^*$.

$$\begin{pmatrix} Y_{00} \\ Y_{1-1} \\ Y_{10} \\ Y_{11} \\ Y_{2-2} \\ Y_{2-1} \\ Y_{20} \\ Y_{21} \\ Y_{22} \\ \cdots \end{pmatrix} = \begin{pmatrix} D^0_{00} & & & & & & & & & \\ & D^1_{-1-1} & D^1_{-10} & D^1_{-11} & & & & & & \\ & D^1_{0-1} & D^1_{00} & D^1_{01} & & & & & & \\ & D^1_{1-1} & D^1_{10} & D^1_{11} & & & & & & \\ & & & & D^2_{-2-2} & D^2_{-2-1} & D^2_{-20} & D^2_{-21} & D^2_{-22} & \\ & & & & D^2_{-1-2} & D^2_{-1-1} & D^2_{-10} & D^2_{-11} & D^2_{-12} & \\ & & & & D^2_{0-2} & D^2_{0-1} & D^2_{00} & D^2_{01} & D^2_{02} & \\ & & & & D^2_{1-2} & D^2_{1-1} & D^2_{10} & D^2_{11} & D^2_{12} & \\ & & & & D^2_{2-2} & D^2_{2-1} & D^2_{20} & D^2_{21} & D^2_{22} & \\ & & & & & & & & & \cdots \end{pmatrix} \begin{pmatrix} Y_{00} \\ Y_{1-1} \\ Y_{10} \\ Y_{11} \\ Y_{2-2} \\ Y_{2-1} \\ Y_{20} \\ Y_{21} \\ Y_{22} \\ \cdots \end{pmatrix}$$

Fig. 4. Depiction of Equation 25 as a matrix equation. Note the block-diagonal nature of the matrix, with only spherical harmonics of the same order (i.e. $l = 0, 1, 2$) mixed. The matrix elements D are functions of the angles of rotation.

The first 3 orders (we give only terms with $m \geq 0$) may be written as

$$\begin{aligned} Y_{00} &= \sqrt{\frac{1}{4\pi}} \\ Y_{10} &= \sqrt{\frac{3}{4\pi}}\cos\theta & &= \sqrt{\frac{3}{4\pi}}z \\ Y_{11} &= -\sqrt{\frac{3}{8\pi}}\sin\theta e^{I\phi} & &= -\sqrt{\frac{3}{8\pi}}(x+Iy) \\ Y_{20} &= \frac{1}{2}\sqrt{\frac{5}{4\pi}}(3\cos^2\theta - 1) & &= \frac{1}{2}\sqrt{\frac{5}{4\pi}}(3z^2-1) \\ Y_{21} &= -\sqrt{\frac{15}{8\pi}}\sin\theta\cos\theta e^{I\phi} & &= -\sqrt{\frac{15}{8\pi}}z(x+Iy) \\ Y_{22} &= \frac{1}{2}\sqrt{\frac{15}{8\pi}}\sin^2\theta e^{2I\phi} & &= \frac{1}{2}\sqrt{\frac{15}{8\pi}}\left(x+Iy\right)^2. \end{aligned} \tag{24}$$

The rotation formula for spherical harmonics is

$$Y_{lm}(R_{\alpha,\beta,\gamma}(\theta,\phi)) = \sum_{m'=-l}^{l} D^l_{m\,m'}(\alpha,\beta,\gamma)Y_{lm'}(\theta,\phi). \tag{25}$$

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: how to rewrite a rotated spherical harmonic as a linear combination of all the spherical harmonics of the same order. In terms of group theory, the matrix D^l is the $(2l+1)$-dimensional representation of the rotation group $SO(3)$. A pictorial depiction of Equation 25 as a matrix multiplication is found in Figure 4.

The matrices D^l therefore satisfy the formula,

$$D^l_{mm'}(\alpha, \beta, \gamma) = \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} Y_{lm}(R_{\alpha,\beta,\gamma}(\theta, \phi)) Y^*_{lm'}(\theta, \phi) \sin\theta \, d\theta \, d\phi. \tag{26}$$

An analytic form for the matrices D^l can be found in standard references, such as Inui et al. [1990]. In particular, since $R_{\alpha,\beta,\gamma} = R_z(\beta) R_y(\alpha) R_z(\gamma)$, the dependence of D^l on β and γ is simple, since rotation of the spherical harmonics about the z-axis is straightforward,

$$D^l_{mm'}(\alpha, \beta, \gamma) = d^l_{mm'}(\alpha) e^{Im\beta} e^{Im'\gamma}, \tag{27}$$

where d^l is a matrix that defines how a spherical harmonic transforms under rotation about the y-axis. For the purposes of the exposition, we will not generally need to be concerned with the precise formula for the matrix d^l, and numerical calculations can compute it using a simplified version of Equation 26 without the z rotations (i.e. $\beta = \gamma = 0$),

$$d^l_{mm'}(\alpha) = \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} Y_{lm}(R_y(\alpha)(\theta, \phi)) Y^*_{lm'}(\theta, \phi) \sin\theta \, d\theta \, d\phi. \tag{28}$$

For completeness, we give below the relatively complicated analytic formula, as derived in Equation 7.48 of Inui et al. [1990],

$$\begin{aligned} \xi &= \sin^2 \frac{\alpha}{2} \\ N(l, m, m') &= (-1)^{m-m'} \sqrt{\frac{(l+m)!}{(l-m)!(l+m')!(l-m')!}} \\ d^l_{mm'}(\alpha) &= N(l, m, m') \times \xi^{-(m-m')/2} (1-\xi)^{-(m+m')/2} \left(\frac{d}{d\xi}\right)^{l-m} \xi^{l-m'} (1-\xi)^{l+m'}, \end{aligned} \tag{29}$$

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

$$d^0(\alpha) = 1$$

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

$$d^2(\alpha) = \begin{pmatrix} \cos^4\frac{\alpha}{2} & 2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & 2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & \sin^4\frac{\alpha}{2} \\ -2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \cos^2\frac{\alpha}{2}(-1+2\cos\alpha) & \sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & (1+2\cos\alpha)\sin^2\frac{\alpha}{2} & 2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} \\ \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & -\sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \frac{1}{2}(3\cos^2\alpha - 1) & \sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha \\ -2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & (1+2\cos\alpha)\sin^2\frac{\alpha}{2} & -\sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \cos^2\frac{\alpha}{2}(-1+2\cos\alpha) & 2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} \\ \sin^4\frac{\alpha}{2} & -2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & -2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \cos^4\frac{\alpha}{2} \end{pmatrix}. \tag{30}$$

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

$$\begin{aligned} D^l_{0m'}(\alpha, \beta, 0) &= d^l_{0m'}(\alpha) = \sqrt{\frac{4\pi}{2l+1}} Y^*_{lm'}(\alpha, \pi) \\ D^l_{m0}(\alpha, \beta, \gamma) &= d^l_{m0}(\alpha) e^{Im\beta} = \sqrt{\frac{4\pi}{2l+1}} Y_{lm}(\alpha, \beta). \end{aligned} \tag{31}$$

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

$$L(\theta_i, \phi_i) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} Y_{lm}(\theta_i, \phi_i). \tag{32}$$

To obtain the lighting in local coordinates, we must rotate the above expression, just as we did in 2D. Using Equation 25, we get,

$$L(\theta_i, \phi_i) = L(R_{\alpha,\beta,\gamma}(\theta'_i, \phi'_i)) = \sum_{l=0}^{\infty} \sum_{m=-l}^{+l} \sum_{m'=-l}^{l} L_{lm} D^l_{mm'}(\alpha, \beta, \gamma) Y_{lm'}(\theta'_i, \phi'_i). \tag{33}$$

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

$$\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \sum_{l=0}^{\infty} \sum_{n=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-p}^{p} \hat{\rho}_{ln,pq} Y^*_{ln}(\theta'_i, \phi'_i) Y_{pq}(\theta'_o, \phi'_o) \tag{34}$$

Spherical Harmonic Reflection Equation: We can now write down the reflection equation, as given by Equation 11, in terms of the expansions just defined. As in 2D, we multiply the expansions for the lighting and BRDF. To avoid confusion between the indices in this intermediate step, we will use L_{lm} and $\hat{\rho}_{l'n,pq}$ to obtain

$$\begin{aligned} B(\alpha, \beta, \gamma, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{m'=-l}^{l} \sum_{l'=0}^{\infty} \sum_{n=-l'}^{l'} \sum_{p=0}^{\infty} \sum_{q=-p}^{p} L_{lm} \hat{\rho}_{l'n,pq} D^l_{mm'}(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o) T_{lm'l'n} \\ T_{lm'l'n} &= \int_{\phi'_i=0}^{2\pi} \int_{\theta'_i=0}^{\pi} Y_{lm'}(\theta'_i, \phi'_i) Y^*_{l'n}(\theta'_i, \phi'_i) \sin\theta'_i \, d\theta'_i d\phi'_i \\ &= \delta_{ll'} \delta_{m'n}. \end{aligned} \tag{35}$$

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

$$B(\alpha, \beta, \gamma, \theta'_o, \phi'_o) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-p}^{p} L_{lm} \hat{\rho}_{ln,pq} \big(D^l_{mn}(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o) \big). \tag{36}$$

This result suggests that we should expand the reflected light field B in terms of the new basis functions given by $C_{lmnpq} = D^l_{mn}(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o)$. The appearance of the matrix D^l in these basis functions is quite intuitive, coming directly from the rotation formula for spherical harmonics. These basis functions are *mixed* in the sense that they are a product of the matrices D^l and the spherical harmonics Y_{pq}. This

can be understood from realizing that the reflected direction is a unit vector described by two parameters (θ_o', ϕ_o'), while the surface parameterization is really a rotation, described by three parameters (α, β, γ). Finally, we need to consider the normalization of these new basis functions. The spherical harmonics are already orthonormal. The orthogonality relation for the matrices D^l is given in any standard text on group theory (for instance, Equation 7.73 of Inui et al. [1990]). Specifically,

$$\int_{\gamma=0}^{2\pi}\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi}\left(D^l_{mn}(\alpha,\beta,\gamma)\right)^*\left(D^{l'}_{m'n'}(\alpha,\beta,\gamma)\right)\sin\alpha\,d\alpha\,d\beta\,d\gamma = \frac{8\pi^2}{2l+1}\delta^{ll'}\delta_{mm'}\delta_{nn'}. \tag{37}$$

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

$$\begin{aligned}
C_{lmnpq} &= \sqrt{\frac{2l+1}{8\pi^2}}D^l_{mn}(\alpha,\beta,\gamma)Y_{pq}(\theta_o',\phi_o') \\
B &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{n=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-p}^{p}B_{lmnpq}C_{lmnpq}(\alpha,\beta,\gamma,\theta_o',\phi_o') \\
B_{lmnpq} &= \int_{\phi_o'=0}^{2\pi}\int_{\theta_o'=0}^{\pi}\int_{\gamma=0}^{2\pi}\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi}B(\alpha,\beta,\gamma,\theta_o',\phi_o')C^*_{lmnpq}(\alpha,\beta,\gamma,\theta_o',\phi_o')\sin\alpha\sin\theta_o'\,d\alpha\,d\beta\,d\gamma\,d\theta_o'\,d\phi_o'.
\end{aligned} \tag{38}$$

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

Finally, we can write down the frequency space reflection equation by comparing Equations 36 and 38 and equating coefficients. This result is comparable to its 2D counterpart, given in Equation 21, and as in 2D, is a convolution. In frequency-space, the reflected light field is obtained simply by multiplying together coefficients of the lighting and BRDF—by *convolving* the incident illumination with the BRDF,

$$\boxed{B_{lmnpq} = \sqrt{\frac{8\pi^2}{2l+1}}L_{lm}\hat{\rho}_{ln,pq}.} \tag{39}$$

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

$$B_{lmn}(\theta_o',\phi_o') = \sqrt{\frac{8\pi^2}{2l+1}}L_{lm}\hat{\rho}_{ln}(\theta_o',\phi_o'). \tag{40}$$

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

4.3 Alternative Forms

For the analysis of certain problems, it will be more convenient to rewrite Equation 39 in a number of different ways. We have already seen one example, of considering the outgoing angle fixed, as shown in Equation 40. In this subsection, we consider a few more alternative forms.

4.3.1 *Isotropic BRDFs.* Isotropic BRDFs are those where rotating the local tangent frame makes no difference: they are functions of only 3 variables, $\hat{\rho}(\theta_i', \phi_i', \theta_o', \phi_o') = \hat{\rho}(\theta_i', \theta_o', \mid \phi_o' - \phi_i' \mid)$. With respect to the reflected light field, the parameter γ, which controls the orientation of the local tangent frame, has no physical significance.

To consider the simplifications that result from isotropy, we first analyze the BRDF coefficients $\hat{\rho}_{ln,pq}$. In the BRDF expansion of Equation 34, only terms that satisfy isotropy, that are invariant with respect to adding an angle $\triangle\phi$ to both incident and outgoing azimuthal angles, are nonzero. From the form of the spherical harmonics, this requires that $n = q$. Furthermore, since we are considering BRDFs that depend only on $|\phi_o' - \phi_i'|$, we should be able to negate both incident and outgoing azimuthal angles without changing the result. This leads to the condition that $\hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q),p(-q)}$. Finally, we define a 3-index BRDF coefficient by

$$\hat{\rho}_{lpq} = \hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q),p(-q)}. \tag{41}$$

Note that isotropy reduces the dimensionality of the BRDF from 4D to 3D. This is reflected in the fact that we now have only three independent indices. Furthermore, half the degrees of freedom are constrained since we can negate the azimuthal angle.

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

$$R_{\alpha,\beta} = R_{\alpha,\beta,0} = R_z(\beta) R_y(\alpha). \tag{42}$$

Next, the representation matrices can be rewritten as

$$D^l_{mn}(\alpha, \beta) = D^l_{mn}(\alpha, \beta, 0) = d^l_{mn}(\alpha) e^{Im\beta}. \tag{43}$$

It should be noted that removing the dependence on γ weakens the orthogonality condition (Equation 37) on the representation matrices, since we no longer integrate over γ. The new orthonormality relation for these matrices is given by

$$\int_{\beta=0}^{2\pi} \int_{\alpha=0}^{\pi} \left(D^l_{mn}(\alpha, \beta)\right)^* \left(D^{l'}_{m'n}(\alpha, \beta)\right) \sin\alpha \, d\alpha \, d\beta = \frac{4\pi}{2l+1} \delta^{ll'} \delta_{mm'}. \tag{44}$$

In particular, the orthogonality relation for the index n no longer holds, which is why we have used the index n for both D and D' instead of using n and n', as in Equation 37. The absence of an integral over γ leads to a slight weakening of the orthogonality relation, as well as a somewhat different normalization than in Equation 37. For the future discussion, it will be convenient to define normalization constants by

$$\Lambda_l = \sqrt{\frac{4\pi}{2l+1}}. \tag{45}$$

We now have the tools necessary to rewrite Equation 39 for isotropic BRDFs. Since we will be using the equations for isotropic BRDFs extensively in the rest of this article, it will be worthwhile to briefly review the representations of the various quantities, specialized to the isotropic case.

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

$$\begin{aligned} L(\theta_i, \phi_i) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} Y_{lm}(\theta_i, \phi_i) \\ L(\theta_i, \phi_i) = L(R_{\alpha,\beta}(\theta_i', \phi_i')) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{+l} \sum_{m'=-l}^{l} L_{lm} D^l_{mm'}(\alpha, \beta) Y_{lm'}(\theta_i', \phi_i'). \end{aligned} \tag{46}$$

Then, we write the expansion of the isotropic BRDF,

$$\hat{\rho}(\theta_i', \theta_o', \mid \phi_o' - \phi_i' \mid) = \sum_{l=0}^{\infty} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} \hat{\rho}_{lpq} Y^*_{lq}(\theta_i', \phi_i') Y_{pq}(\theta_o', \phi_o'). \tag{47}$$

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

$$\begin{aligned} C_{lmpq}(\alpha, \beta, \theta_o', \phi_o') &= \Lambda_l^{-1} D^l_{mq}(\alpha, \beta) Y_{pq}(\theta_o', \phi_o') \\ B(\alpha, \beta, \theta_o', \phi_o') &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} B_{lmpq} C_{lmpq}(\alpha, \beta, \theta_o', \phi_o') \\ B_{lmpq} &= \int_{\phi_o'=0}^{2\pi} \int_{\theta_o'=0}^{\pi} \int_{\beta=0}^{2\pi} \int_{\alpha=0}^{\pi} B(\alpha, \beta, \theta_o', \phi_o') C^*_{lmpq}(\alpha, \beta, \theta_o', \phi_o') \sin\alpha \sin\theta_o' \, d\alpha \, d\beta \, d\theta_o' \, d\phi_o'. \end{aligned} \tag{48}$$

It should be noted that the basis functions C_{lmpq} are orthonormal in spite of the weakened orthogonality of the functions D^l_{mq}, as expressed in Equation 44. Note that the index q in the definition of C_{lmpq} is the same (coupled) for both factors D^l_{mq} and Y_{pq}. This is a consequence of isotropy, and is not true in the anisotropic case. Therefore, although the representation matrices D^l no longer satisfy orthogonality over the index q (corresponding to the index n in Equation 44), orthogonality over the index q follows from the orthonormality of the spherical harmonics Y_{pq}.

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

$$\boxed{B_{lmpq} = \Lambda_l L_{lm} \hat{\rho}_{lpq}} \tag{49}$$

Apart from a slightly different normalization, and the removal of γ and the corresponding index n, this is essentially the same as Equation 39. We will be using this equation for isotropic BRDFs extensively in Section 5 of this article, where we quantitatively analyze the reflection equation for many special cases of interest.

We may also try to derive an alternative form, analogous to Equation 40, by holding the outgoing elevation angle θ_o' fixed. Since the isotropic BRDF depends only on $|\phi_o' - \phi_i'|$, and not directly on ϕ_o', we do not hold ϕ_o' fixed, as we did in Equation 40. We first define the modified expansions,

$$\begin{aligned} \hat{\rho}(\theta_i', \theta_o', \mid \phi_o' - \phi_i' \mid) &= \sum_{l=0}^{\infty} \sum_{q=-l}^{l} \hat{\rho}_{lq}(\theta_o') \left(\frac{1}{\sqrt{2\pi}} Y^*_{lq}(\theta_i', \phi_i') \exp(Iq\phi_o') \right) \\ B(\alpha, \beta, \theta_o', \phi_o') &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{q=-l}^{l} B_{lmq}(\theta_o') \left(\frac{1}{\sqrt{2\pi}} \Lambda_l^{-1} D^l_{mq}(\alpha, \beta) \exp(Iq\phi_o') \right). \end{aligned} \tag{50}$$

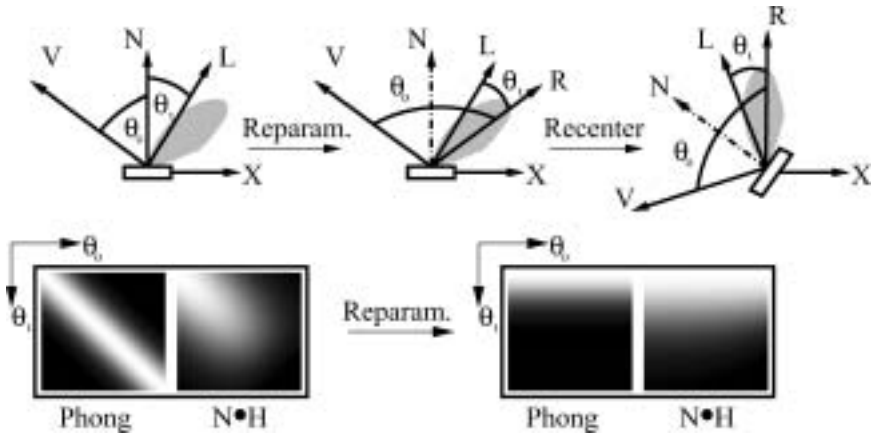

Fig. 5. Reparameterization involves recentering about the reflection vector. BRDFs become more compact, and in special cases (Phong) become 1D functions.

Then, we may write down the isotropic convolution formula corresponding to Equation 40.

$$B_{lmq}(\theta'_o) = \Lambda_l L_{lm} \hat{\rho}_{lq}(\theta'_o) \tag{51}$$

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

$$\begin{aligned} \tilde{\rho} &= \hat{\rho} \cos\theta'_o = \rho \cos\theta'_i \cos\theta'_o \\ \tilde{B} &= B \cos\theta'_o \end{aligned} \tag{52}$$

With these definitions, all of the derivations presented so far still hold. In particular, the convolution formulas in Equations 39 and 49 hold with the replacements $B \to \tilde{B}$, $\hat{\rho} \to \tilde{\rho}$. For example, Equation 49 for isotropic BRDFs becomes

$$\tilde{B}_{lmpq} = \Lambda_l L_{lm} \tilde{\rho}_{lpq} \tag{53}$$

The symmetry of the transfer function ensures that its coefficients are unchanged if the indices corresponding to incident and outgoing angles are interchanged, i.e. $\tilde{\rho}_{lpq} = \tilde{\rho}_{plq}$. In the more general anisotropic case, $\tilde{\rho}_{ln,pq} = \tilde{\rho}_{pq,ln}$. We will use the frequency-space reflection formula, as given by Equation 53, whenever explicitly maintaining the reciprocity of the BRDF is important.

4.3.3 *Reparameterization by Central BRDF Direction.* Consider first the special case of *radially symmetric or 1D BRDFs*, where the BRDF consists of a single symmetric lobe of fixed shape, whose orientation depends only on a well-defined central direction $\vec{C}$. In other words, the BRDF is given by a 1D function u as $\hat{\rho} = u(\vec{C} \cdot \vec{L})$. Examples are Lambertian $\hat{\rho} = \vec{N} \cdot \vec{L}$ and Phong $\hat{\rho} = (\vec{R} \cdot \vec{L})^s$ models. If we reparameterize the BRDF and reflected light field by $\vec{C}$, the BRDF becomes a function of only 1 variable (θ'_i with $\cos\theta'_i = \vec{C} \cdot \vec{L}$) instead of 3. Refer to Figure 5 for an illustration. Further, the reflected light field can be represented simply by a 2D reflection map $B(\alpha, \beta)$ parameterized by $\vec{C} = (\alpha, \beta)$. In other words, after reparameterization, there is no explicit exitant (outgoing) angular dependence for either the BRDF or reflected light field.

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

$$\begin{aligned}
\hat{\rho}(\theta_i') &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta_i') \\
\hat{\rho}_l &= 2\pi \int_0^{\pi/2} \hat{\rho}(\theta_i') Y_{l0}(\theta_i') \sin\theta_i' \, d\theta_i' \\
B(\alpha, \beta) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{q=-l}^{l} \hat{\rho}_l L_{lm} D^l_{mq}(\alpha, \beta) \int_0^{2\pi} \int_0^{\pi} Y_{lq}(\theta_i', \phi_i') Y_{l0}(\theta_i') \sin\theta_i' \, d\theta_i' d\phi_i' \\
&= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \hat{\rho}_l L_{lm} D^l_{m0}(\alpha, \beta).
\end{aligned} \tag{54}$$

In computing $\hat{\rho}_l$, we have integrated out the azimuthal dependence, accounting for the factor of 2π. In the last line, we have used orthonormality of the spherical harmonics. Now, we use the second property of the matrices D from Equation 31, i.e. $D^l_{m0}(\alpha, \beta) = \Lambda_l Y_{lm}(\alpha, \beta)$. Therefore, the reflected light field can be expanded simply in terms of spherical harmonics,

$$B(\alpha, \beta) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm} Y_{lm}(\alpha, \beta). \tag{55}$$

The required convolution formula now becomes

$$\boxed{B_{lm} = \Lambda_l \hat{\rho}_l L_{lm}}. \tag{56}$$

In the context of Lambertian BRDFs (for which no reparameterization is required), it has been noted by Basri and Jacobs [2001a, 2003] that Equation 56 is mathematically an instance of the Funk-Hecke theorem (as stated, for instance in Groemer [1996], page 98). However, that theorem does not generalize to the other relations previously encountered. With respect to Equation 49, we have essentially just dropped the indices p and q corresponding to the outgoing angular dependence. It is important to remember that the reflected light field is now expanded in terms of spherical harmonics. B is simply a filtered version of L, with each frequency l being attenuated by a different amount, corresponding to the BRDF transfer function $\hat{\rho}_l$.

For general BRDFs, the radial symmetry property does not hold precisely, so they cannot be reduced exactly to 1D functions, nor can B be written simply as a 2D reflection map. Nevertheless, a reparameterization of the specular BRDF components by the reflection vector (or other central BRDF direction) still yields compact forms. To reparameterize, we simply recenter the BRDF (and the reflection integral) about the reflection vector $\vec{R}$, rather than the surface normal, as shown in Figure 5. The reflection vector now takes the place of the surface normal, i.e. $\vec{R} = (\alpha, \beta)$, and the dependence on the surface normal becomes indirect (just as the dependence on $\vec{R}$ is indirect in the standard parameterization). The angles θ_i' and θ_o' are now given with respect to $\vec{R}$ by $\cos\theta_i' = \vec{R} \cdot \vec{L}$ and $\cos\theta_o' = \vec{R} \cdot \vec{V}$, with $B(\alpha, \beta, \theta_o', \phi_o')$ a function of $\vec{R} = (\alpha, \beta)$ and $\omega_o = (\theta_o', \phi_o')$. Once we have done this, we can directly apply the general convolution formulas, such as Equation 49.

5. IMPLICATIONS

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

We will deal here exclusively with the 3D case, since that is of greater practical importance. A preliminary analysis for the 2D case can be found in an earlier paper [Ramamoorthi and Hanrahan 2001a]. The quantitative results in 2D and 3D are closely related, although the fact that the 3D treatment is in terms of spherical harmonics, as opposed to the 2D treatment in terms of Fourier series, results in some important differences. For simplicity, we will also restrict the ensuing discussion to the case of isotropic BRDFs. The extension to anisotropic surfaces can be done using the equations derived earlier for the general anisotropic case.

5.1 General Observations

In this subsection, we briefly discuss the key ideas resulting from the theory, as applied to general lighting conditions and BRDF models.

5.1.1 *Forward Rendering with Environment Maps.* We first consider the problem of rendering with *environment maps*: general lighting distributions. For the purposes of rendering, it is convenient to explicitly write the formula for the reflected light field as

$$B(\alpha, \beta, \theta_o', \phi_o') = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} L_{lm} \hat{\rho}_{lpq} \left(D^l_{mq}(\alpha, \beta) Y_{pq}(\theta_o', \phi_o') \right). \tag{57}$$

If either the lighting or the BRDF is low frequency, the total number of terms in the summation will be relatively small, and it may be possible to use Equation 57 directly for shading a pixel. We have already demonstrated the practicality of this approach for Lambertian BRDFs [Ramamoorthi and Hanrahan 2001b], where we can set $p = q = 0$, and use $l \leq 2$—only 9 spherical harmonic terms.

In the general case, frequency space analysis allows for setting sampling rates accurately, and enables compact frequency domain representations. Further, just as image convolutions are often computed in the Fourier rather than the spatial domain, computing the reflected light field is more efficient in frequency space, using Equation 57, rather than in angular space. We have already demonstrated the practical implementation of these ideas [Ramamoorthi and Hanrahan 2002].

5.1.2 *Well-Posedness and Conditioning of Inverse Lighting and BRDF.* In this subsubsection, we briefly discuss how to apply ideas from the theoretical analysis to determine which inverse problems are well-posed, or solvable, versus ill-posed, that is unsolvable, and also determine the numerical conditioning properties. We note that the inverse problems we consider here differ from much work in computer vision in the sense that we assume known shape, estimating lighting and material properties, as opposed to solving for shape, given simple models of the lighting and materials (such as Lambertian reflectance). Later, we will also relate these results to the general theory of linear integral equations. An important duality should be noted here. Forward problems for which an efficient frequency domain solution is possible, such as those involving diffuse surfaces and/or soft lighting, have corresponding inverse problems that are ill-conditioned. Turned around, ill-conditioned inverse problems allow us to get a very good solution to the forward problem by using very coarse low-frequency approximations of the initial conditions. For instance, Lambertian surfaces act as low-pass filters, the precise form of which we will explore later in this section, blurring the illumination. Therefore, high-frequency components of the lighting are not essential to rendering images of diffuse objects, and we can make very coarse low-frequency approximations to the lighting without significantly affecting the final image.

Inverse BRDF: We first address the question of BRDF estimation. Our goal is to consider BRDF estimation under general illumination conditions, and understand when the BRDF can be recovered—BRDF estimation is well posed—and when the BRDF cannot be recovered—estimation is ill-posed. We would

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

For isotropic surfaces, a simple manipulation of Equation 49 yields

$$\hat{\rho}_{lpq} = \Lambda_l^{-1} \frac{B_{lmpq}}{L_{lm}}. \tag{58}$$

In general, BRDF estimation will be well-posed, or unambiguous, as long as the denominator on the right-hand side does not vanish. Of course, to be physically accurate, the numerator will also become 0 if the denominator vanishes, so the right-hand side will become indeterminate. From Equation 58, we see that BRDF estimation is well posed as long as for all l, there exists at least one value of m such that $L_{lm} \neq 0$. In other words, all orders in the spherical harmonic expansion of the lighting should have at least one coefficient with nonzero amplitude. If any order l completely vanishes, the corresponding BRDF coefficients cannot be estimated.

In signal processing terms, if the input signal (lighting) has no amplitude along certain modes of the filter (BRDF), those modes cannot be estimated. BRDF recovery is well conditioned when the spherical harmonic expansion of the lighting does not decay rapidly with increasing frequency—when the lighting contains high frequencies like directional sources or sharp edges—and is ill-conditioned for soft lighting. Equation 58 gives a precise mathematical characterization of the conditions for BRDF estimation to be well-posed and well-conditioned. These results are similar to those obtained by D'Zmura [1991] who states that there is an ambiguity regarding the BRDF in case of *inadequate illumination*. In our framework, inadequate illumination corresponds to certain frequencies l of the lighting completely vanishing.

Inverse Lighting: A similar analysis can be done for estimation of the lighting. Manipulation of Equation 49 yields

$$L_{lm} = \Lambda_l^{-1} \frac{B_{lmpq}}{\hat{\rho}_{lpq}}. \tag{59}$$

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

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

integral equations of the first kind may be written generally as

$$b(s) = \int_t K(s, t) f(t)\, dt, \tag{60}$$

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

The theory of linear integral equations, as for instance in Cochran [1972], analyzes Equation 60 based on the structure of the kernel. In particular, assume we may find a basis function expansion of the form

$$\begin{aligned} b(s) &= \sum_{i=1}^{n} b_i u_i(s) \\ K(s, t) &= \sum_{i=1}^{n} K_i u_i(s) v_i^*(t) \\ f(t) &= \sum_{i=1}^{\infty} f_i v_i(t), \end{aligned} \tag{61}$$

where each of the sets of functions u_i and v_i (with v_i^* being the complex conjugate) is linearly independent. Here, n is the number of terms in, or rank of the kernel, K. If n is finite, the kernel is referred to as degenerate. It should be noted that if the function sets u and v were orthonormal, then we would have a result of the form $b_i = K_i f_i$. In effect, we have constructed an expansion of the form of Equation 61 using orthonormal basis functions involving group representations and spherical harmonics, thereby deriving the convolution result.

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

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

We first note that there is a global scale factor that we cannot recover. Multiplying the lighting everywhere by some constant amount and dividing the BRDF uniformly by the same amount leaves the reflected light field, which is a product of the two, unchanged. Of course, physical considerations bound the scale factor, since the BRDF must remain energy preserving. Nevertheless, within this

general constraint, it is not possible to estimate the absolute magnitudes of the lighting and BRDF. However, we will demonstrate that apart from this ambiguity, the light field can indeed be factored, allowing us to simultaneously determine both the lighting and the BRDF.

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

It would appear from Equation 49 however, that there is an unrecoverable scale factor for each order l, corresponding to the known coefficient we require. In other words, we may multiply the lighting for each order l by some amount (which may be different for different frequencies l) while dividing the BRDF by the same amount. However, there is an important additional physical constraint. The BRDF must be reciprocal: symmetric with respect to incident and outgoing angles. The corresponding condition in the frequency domain is that the BRDF coefficients must be symmetric with respect to interchange of the indices corresponding to the incident and outgoing directions. To take advantage of this symmetry, we will use the reciprocal form of the frequency-space equations, as defined in Equation 53.

We now derive an analytic formula for the lighting and BRDF in terms of coefficients of the reflected light field. Since we cannot recover the global scale, we will arbitrarily scale the DC term of the lighting so $L_{00} = \Lambda_0^{-1} = \sqrt{1/(4\pi)}$. Note that this scaling is valid unless the DC term is 0, corresponding to no light—an uninteresting case. Using Equations 53, 58, and 59, we obtain

$$
\begin{aligned}
L_{00} &= \Lambda_0^{-1} && \text{: Global Scale} \\
\tilde{\rho}_{0p0} &= \tilde{B}_{00p0} && \text{: Equation 58 } (l = q = 0) \\
L_{lm} &= \Lambda_l^{-1} \frac{\tilde{B}_{lmpq}}{\tilde{\rho}_{lpq}} && \text{: Equation 59} \\
&= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{l00}} && \text{: Set } p = q = 0 \\
&= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{0l0}} && \text{: Reciprocity, } \tilde{\rho}_{0l0} = \tilde{\rho}_{l00} \\
&= \Lambda_l^{-1} \frac{\tilde{B}_{lm00}}{\tilde{B}_{00l0}} && \text{: Plug in from 2nd line} \\
\tilde{\rho}_{lpq} &= \Lambda_l^{-1} \frac{\tilde{B}_{lmpq}}{L_{lm}} && \text{: Equation 58} \\
&= \frac{\tilde{B}_{lmpq} \tilde{B}_{00l0}}{\tilde{B}_{lm00}} && \text{: Substitute from above for } L_{lm}.
\end{aligned}
\tag{62}
$$

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

Therefore, up to global scale, **the reflected light field can be factored into the lighting and the BRDF**, provided the appropriate coefficients of the reflected light field do not vanish, that is the denominators above are nonzero. If the denominators do vanish, the inverse-lighting or inverse-BRDF problems become ill-posed and consequently, the factorization becomes ill-posed. Note that the above relations are one possible factorization formula. We may still be able to factor the light field even if some of the $\tilde{\rho}_{l00}$ terms vanish in Equation 63, by using different values of $\tilde{\rho}_{lpq}$ with $p \neq 0$.

5.2 Analytic Formulae for Common Examples

We now illustrate these general ideas with several important special-cases. 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 two lighting conditions: a single directional source, and uniform lighting. Then, we consider four BRDF models: a mirror surface, a Lambertian surface, the Phong BRDF, and a microfacet model. A more detailed exposition of some parts can be found in Chapter 3 of the PhD thesis of the first author [Ramamoorthi 2002b].

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

$$\begin{aligned} L(\theta_i, \phi_i) &= \delta(\cos\theta_i - \cos\theta_s)\,\delta(\phi_i - \phi_s) \\ L_{lm} &= \int_0^{2\pi}\int_0^{\pi} \delta(\cos\theta_i - \cos\theta_s)\,\delta(\phi_i - \phi_s)\,Y^*_{lm}(\theta_i, \phi_i)\,\sin\theta_i\,d\theta_i\,d\phi_i \\ &= Y^*_{lm}(\theta_s, \phi_s). \end{aligned} \tag{63}$$

Note that in the above equation, 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; $\theta_s = 0$. It is now straightforward to write

$$\begin{aligned} L_{lm} &= Y^*_{lm}(0) \\ &= \Lambda_l^{-1}\delta_{m0} \\ B_{lmpq} &= \delta_{m0}\hat{\rho}_{lpq} \\ \hat{\rho}_{lpq} &= B_{l0pq}. \end{aligned} \tag{64}$$

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 [Lu et al. 1998; Marschner et al. 2000]. 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.

5.2.2 *Uniform Lighting.* This can be considered the canonical opposite to the case of a directional source. For uniform lighting, only the DC term of the lighting does not vanish, that is we set $L_{00} = \Lambda_0^{-1}$, with other coefficients being zero.

$$\begin{aligned} L_{lm} &= \delta_{l0}\delta_{m0}\Lambda_0^{-1} \\ B_{lmpq} &= \delta_{l0}\delta_{m0}\delta_{q0}\hat{\rho}_{0p0} \\ \hat{\rho}_{0p0} &= B_{00p0}. \end{aligned} \tag{65}$$

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. For instance, a mirror surface cannot be distinguished from a Lambertian object under uniform lighting.

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

5.2.3 *Mirror BRDF.* A mirror BRDF is analogous to the case of a directional source. The reflected light has the same intensity as the incident light. 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. We use the BRDF formula in Equation 2.29 of Cohen and Wallace [1993], multiplying by the cosine of the local incident angle, to get the appropriate formula for the transfer function. In this special case, the transfer function is also reciprocal since the cosines of the incident and outgoing angles are the same. The spherical harmonic coefficients can be written as

$$\begin{aligned}
\hat{\rho} &= \delta(\cos\theta_o' - \cos\theta_i')\delta(\phi_o' - \phi_i' \pm \pi) \\
\hat{\rho}_{lpq} &= \int_0^{2\pi}\int_0^{2\pi}\int_0^{\pi/2}\int_0^{\pi/2} \\
&\quad \delta(\cos\theta_o' - \cos\theta_i')\delta(\phi_o' - \phi_i' \pm \pi) Y_{lq}(\theta_i', \phi_i') Y_{pq}^*(\theta_o', \phi_o') \sin\theta_i' \sin\theta_o' \, d\theta_i' d\theta_o' d\phi_i' d\phi_o' \\
&= \int_0^{2\pi}\int_0^{\pi/2} Y_{pq}^*(\theta_i', \phi_i' \pm \pi) Y_{lq}(\theta_i', \phi_i') \, \sin\theta_i' \, d\theta_i' d\phi_i' \\
&= (-1)^q \delta_{lp}.
\end{aligned} \tag{66}$$

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 elevation 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

$$\begin{aligned}
B_{lmpq} &= \Lambda_l (-1)^q \delta_{lp} L_{lm} \\
L_{lm} &= \Lambda_l^{-1} (-1)^q B_{lmlq}.
\end{aligned} \tag{67}$$

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.

Reparameterization by Reflection Vector: For reflective BRDFs, it is often convenient to reparameterize by the reflection vector, as discussed in Section 4.3.3. 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.

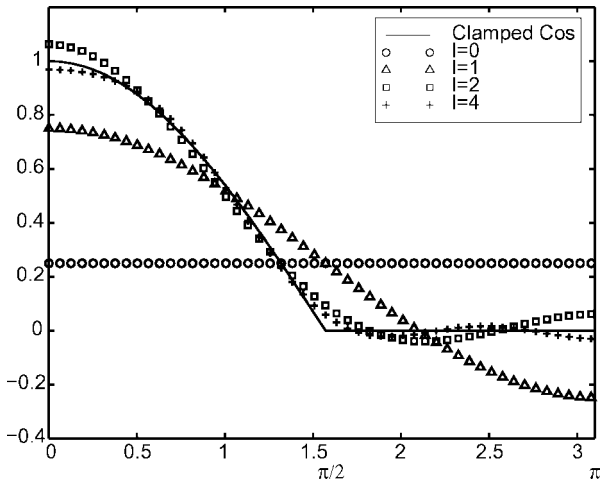

Fig. 6. The *clamped cosine* filter corresponding to the Lambertian BRDF and successive approximations obtained by adding more spherical harmonic terms. For $l = 2$, we already get a very good approximation.

Since there is no dependence on the outgoing angle after reparameterization, we obtain

$$\begin{aligned} \hat{\rho}(\theta_i') &= \frac{\delta(\cos(0) - \cos\theta_i')}{2\pi} \\ \hat{\rho}_l &= Y_{l0}(0) = \Lambda_l^{-1} \\ B_{lm} &= L_{lm}. \end{aligned} \tag{68}$$

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 directly corresponds 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.

5.2.4 *Lambertian BRDF.* For a Lambertian surface, the BRDF is a constant, corresponding to the scaled surface albedo—for purposes of energy conservation, the albedo must be less than 1, and the BRDF must consequently 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 6.

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 56,

$$\begin{aligned} \hat{\rho} = \max(\cos\theta_i', 0) &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta_i') \\ B_{lm} &= \Lambda_l \hat{\rho}_l L_{lm}. \end{aligned} \tag{69}$$

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 [MacRobert 1948] will be useful. The spherical harmonic coefficients are given by

$$\hat{\rho}_l = 2\pi \int_0^{\frac{\pi}{2}} Y_{l0}(\theta_i') \cos\theta_i' \sin\theta_i' \, d\theta_i'. \tag{70}$$

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

$$\hat{\rho}_l = 2\pi \Lambda_l^{-1} \int_0^1 P_l(u) P_1(u) \, du. \tag{71}$$

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

$$\int_{-1}^1 P_a(u) P_b(u) = \frac{2}{2a+1} \delta_{a,b}. \tag{72}$$

From this, we can establish some results about Equation 71. When l is equal to 1, the integral evaluates to half the above norm: 1/3. When l is odd but greater than 1, the integral in Equation 71 vanishes. This is because, for $a = l$ and $b = 1$, we can break the left-hand side of Equation 72 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 71. When l is even, the required formula is given by manipulating Equation 20 in Chapter 5 of MacRobert [1948]. Putting it all together, we obtain, as plotted in Figure 7,

$$\begin{aligned} l = 1 \quad \hat{\rho}_l &= \sqrt{\frac{\pi}{3}} \\ l > 1, odd \quad \hat{\rho}_l &= 0 \\ l \ \ even \quad \hat{\rho}_l &= 2\pi \sqrt{\frac{2l+1}{4\pi}} \frac{(-1)^{\frac{l}{2}-1}}{(l+2)(l-1)} \left[\frac{l!}{2^l \left(\frac{l}{2}!\right)^2} \right]. \end{aligned} \tag{73}$$

There are two important points to note here. First, the transfer function is identically zero for odd frequencies greater than 1. Second, 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,

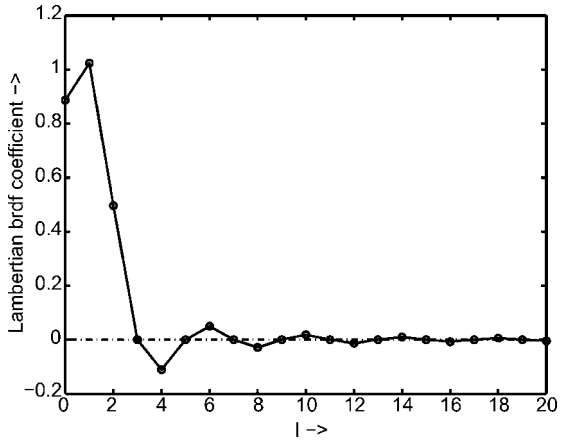

Fig. 7. The solid line is a plot of $\hat{\rho}_l$ versus l, as per Equation 74. It can be seen that odd terms with $l > 1$ have $\hat{\rho}_l = 0$. Also, as l increases, the BRDF coefficients decay rapidly.

we may write the first few terms for the BRDF filter as

$$\begin{aligned} \Lambda_0\hat{\rho}_0 &= 3.142 \\ \Lambda_1\hat{\rho}_1 &= 2.094 \\ \Lambda_2\hat{\rho}_2 &= 0.785 \\ \Lambda_3\hat{\rho}_3 &= 0 \\ \Lambda_4\hat{\rho}_4 &= -0.131 \\ \Lambda_5\hat{\rho}_5 &= 0 \\ \Lambda_6\hat{\rho}_6 &= 0.049. \end{aligned} \tag{74}$$

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 [Basri and Jacobs 2001a], 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 should note that, unlike the actual clamped cosine function, the order 2 spherical harmonic approximation is small but not exactly 0 in the lower or unseen hemisphere, and is also negative in some locations. It can therefore reflect a small amount of light from below the horizon. While we have not observed any significant practical problems as a result, some visual differences may be perceptible in some circumstances (for instance, the actual clamped cosine function has a C^1 discontinuity at the horizon, while the spherical harmonic approximation does not).

We first consider illumination estimation, or the inverse lighting problem. The fact that the odd frequencies greater than 1 of the BRDF vanish means the inverse lighting problem is formally ill-posed for a Lambertian surface. The filter zeros the odd frequencies of the input signal, so these terms cannot

be estimated from images of a convex Lambertian object. This observation corrects a commonly held notion (see Preisendorfer [1976], volume 2, pages 143–151) that radiance and irradiance are equivalent in the sense that irradiance can be formally inverted to recover the radiance. In fact, different radiance distributions can give rise to the same irradiance distribution. For further details, see Ramamoorthi and Hanrahan [2001c]. 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 [1997].

Our results are also in accordance with the perception literature, such as Land's retinex theory [Land and McCann 1971]. 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, 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—the shading on a Lambertian surface. Further implementation details on computing and rendering with *irradiance environment maps* are found in Ramamoorthi and Hanrahan [2001b]. 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 [1994] and Epstein et al. [1995]. Basri and Jacobs [2001a, 2003] have derived a Lambertian formula similar to ours, and have applied this result to lighting invariant recognition, and to photometric stereo under general unknown lighting [Basri and Jacobs 2001b]. A number of other researchers have also used these results to develop new practical algorithms for recognition and reconstruction in computer vision [Lee et al. 2001; Ho et al. 2003; Zhang and Samaras 2003; Simakov et al. 2003; Zhang et al. 2003].

5.2.5 *Phong BRDF.* The normalized Phong transfer function we use is

$$\hat{\rho} = \frac{s+1}{2\pi}(\vec{R} \cdot \vec{L})^s, \tag{75}$$

where $\vec{R}$ is the reflection vector, $\vec{L}$ is the direction to the light source, and s is the *shininess*, or Phong exponent. Lewis [1993] discusses various ways of making shaders, like the Phong BRDF, physically plausible. We adopt his normalization term, ensuring 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. Similarly, we also treat Equation 75 directly as the transfer function, rather than the BRDF, for simplicity. A more detailed analysis of Phong reflection using Clebsch-Gordan coefficients [Inui et al. 1990], taking into account horizon effects and multiplying by the cosine of the incident angle, is provided by Thornber and Jacobs [2001].

The form of Equation 75 allows us to reparameterize by the reflection vector $\vec{R}$, making the transformations outlined in Section 4.3.3. 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 69 holds. However, note that while the Phong BRDF is mathematically analogous to the Lambertian case,

it is not physically similar since we have reparameterized by the reflection vector. The BRDF coefficients depend on s, and are given by

$$\begin{aligned} \hat{\rho}_l &= (s+1)\int_0^{\pi/2} [\cos\theta_i']^s Y_{l0}(\theta_i') \sin\theta_i' \, d\theta_i' \\ B_{lm} &= \Lambda_l \hat{\rho}_l L_{lm}. \end{aligned} \tag{76}$$

To solve this integral, we substitute $u = \cos\theta_i'$ in Equation 76. 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 76 becomes

$$\hat{\rho}_l = \Lambda_l^{-1}(s+1)\int_0^1 u^s P_l(u)\, du. \tag{77}$$

An analytic formula is given by MacRobert [1948] in equations 19 and 20 of Chapter 5.

$$\begin{aligned} \text{ODD } l \quad \Lambda_l \hat{\rho}_l &= \frac{(s+1)(s-1)(s-3)\cdots(s-l+2)}{(s+l+1)(s+l-1)\cdots(s+2)} \\ \text{EVEN } l \quad \Lambda_l \hat{\rho}_l &= \frac{s(s-2)\cdots(s-l+2)}{(s+l+1)(s+l-1)\cdots(s+3)}. \end{aligned} \tag{78}$$

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 $s > l-1$,

$$\Lambda_l \hat{\rho}_l \sim \frac{\left[\Gamma\left(\frac{s}{2}\right)\right]^2}{\Gamma\left(\frac{s}{2}-\frac{l}{2}\right)\Gamma\left(\frac{s}{2}+\frac{l}{2}\right)}. \tag{79}$$

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

$$\log(\Lambda_l \hat{\rho}_l) = -l^2\left(\frac{1}{2s} - \frac{1}{2s^2}\right) + O\left(\frac{l^4}{s^2}\right). \tag{80}$$

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

$$\Lambda_l \hat{\rho}_l \approx \exp\left[-\frac{l^2}{2s}\right]. \tag{81}$$

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 81 is shown in Figure 8.

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: We now consider factorization of light fields for surfaces approximated by Phong BRDFs. From the form of Equation 76, 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

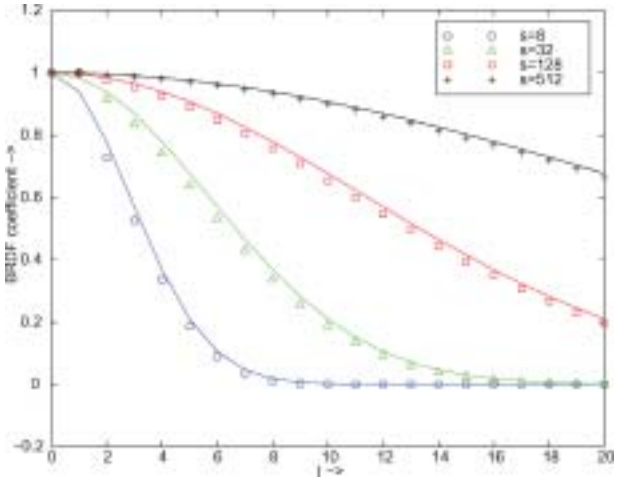

Fig. 8. Numerical plots of the Phong coefficients $\Lambda_l \hat{\rho}_l$, as defined by Equation 78. The solid lines are the Gaussian filter approximations in Equation 81. As the Phong exponent s increases, corresponding to increasing the angular width of the BRDF filter, the frequency width of the BRDF filter decreases.

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 76 as

$$\begin{aligned} L'_{lm} &= \Lambda_l \hat{\rho}_l L_{lm} \\ \Lambda_l \hat{\rho}'_l &= 1 \\ B_{lm} &= \Lambda_l \hat{\rho}'_l L'_{lm} \\ &= L'_{lm}. \end{aligned} \tag{82}$$

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 81 for the Phong BRDF coefficients indicates that for $l \ll \sqrt{s}$, the effects of the filtering in Equation 83 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.

5.2.6 *Microfacet BRDF.* Consider a simplified Torrance-Sparrow [1967] model,

$$\rho = \frac{1}{4\pi\sigma^2 \cos\theta'_i \cos\theta'_o} \exp\left[-\left(\frac{\theta'_h}{\sigma}\right)^2\right]. \tag{83}$$

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.

We analyze the microfacet BRDF by fixing the outgoing angle, and writing the BRDF as

$$\hat{\rho}(\theta_i', \theta_o', \phi_i') = \sum_{l=0}^{\infty} \sum_{q=-l}^{l} \hat{\rho}_{lq}(\theta_o') Y_{lq}(\theta_i', \phi_i'). \tag{84}$$

The advantage of this representation is that other effects like Fresnel can be incorporated essentially as simply a scaling function, dependent on θ_o', applied to $\hat{\rho}_{lq}$, instead of affecting the coefficients in more complicated ways.

It is convenient to also reparameterize by the reflection vector, as we did for the Phong BRDF (i.e. consider (θ_i', ϕ_i') with respect to the reflection of the outgoing direction about the surface normal). However, it is important to note that microfacet BRDFs are not symmetric about the reflection vector. As distinct from 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. This is in accord with observations made by others (such as Ikeuchi and Sato [1991]) that the Torrance-Sparrow BRDF is radially symmetric for small outgoing angles. Furthermore, in practical inverse rendering applications, measurements made at large or grazing angles are given low confidence anyway. Even under this assumption, it is somewhat difficult to derive precise analytic formulae. However, we may make good approximations.

In this article, we will only describe our derivation of the coefficients $\hat{\rho}_{lq}$ for normal exitance $\theta_o' = 0$. In this case, because of azimuthal symmetry, only terms with $q = 0$ will be nonzero. Since the BRDF is approximately radially symmetric for small outgoing angles, these results will also hold for these cases: we can apply the same mathematical machinery as for Lambertian and Phong. More details on approximations for non-normal exitance are found in Ramamoorthi [2002b], where it is shown more formally that the corrections to the results below are small except when l is large (corresponding to high-frequency lighting distributions or directional light sources), and at large viewing angles.

For normal exitance, there is no azimuthal dependence, and the half angle, $\theta_h' = \theta_i'/2$,

$$\hat{\rho}_l = 2\pi \int_0^{\pi/2} \frac{\exp\left[-\theta_i'^2/4\sigma^2\right]}{4\pi\sigma^2} Y_{l0}(\theta_i') \sin\theta_i' \, d\theta_i'. \tag{85}$$

The expansion of $Y_{l0}(t)$ near $t = 0$ for small l is (this can be verified using a symbolic mathematics program like Mathematica) is

$$Y_{l0}(t) = \Lambda_l^{-1} \left(1 - \frac{l(l+1)}{4} t^2 + O([lt]^4)\right). \tag{86}$$

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

$$Y_{l0}(t) \sim \Lambda_l^{-1} \left(\frac{1}{\sqrt{t}} \cos[(l + 1/2)t - \pi/4]\right). \tag{87}$$

To integrate Equation 85, 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 $\sin\theta_i' \, d\theta_i' = \theta_i' \, d\theta_i' = 4\sigma^2 u \, du$.

$$\hat{\rho}_l \sim \int_0^{\infty} 2e^{-u^2} Y_{l0}(2\sigma u) u \, du \tag{88}$$

We therefore set $t = 2\sigma u$ in Equations 86 and 87. When $\sigma l \ll 1$, we use Equation 86 to obtain to $O([\sigma l]^4)$,

$$\Lambda_l \hat{\rho}_l \sim \left(\int_0^{\infty} 2ue^{-u^2} \, du - (\sigma l)^2 \int_0^{\infty} 2u^3 e^{-u^2} \, du\right). \tag{89}$$

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

$$\Lambda_l \hat{\rho}_l = 1 - (\sigma l)^2 + O([\sigma l]^4). \tag{90}$$

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 87 to obtain (Φ is a phase that encapsulates the lower-order terms)

$$\Lambda_l \hat{\rho}_l \sim \int_0^\infty e^{-u^2} \sqrt{u} \cos[(2\sigma l)u + \Phi]\, du. \tag{91}$$

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

$$\Lambda_l \hat{\rho}_l \approx \exp[-(\sigma l)^2]. \tag{92}$$

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 92 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.

Conditioning Properties: Since Equation 92 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 the Phong 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 92 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 [Ramamoorthi and Hanrahan 2001d].

6. 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. We have shown that the reflected light field can be formally described as a convolution of the incident illumination and the BRDF, and derived an analytic frequency space convolution formula. This means that reflection can be viewed in signal processing terms as a filtering operation between the lighting and the BRDF to produce the output light field. Furthermore, inverse rendering to estimate the lighting or BRDF from the reflected light field can be understood as deconvolution. This result provides a novel viewpoint for many forward and inverse rendering problems, and allows us to understand the duality between forward and inverse problems, wherein an ill-conditioned inverse problem may lead to an efficient solution to a forward problem. We have discussed the implications for inverse problems such as lighting recovery, BRDF recovery, light field factorization, and forward rendering problems such as environment map prefiltering and rendering. We have also made these ideas concrete for many special cases, deriving analytic formulae for the frequency spectra of many common BRDF and lighting models.

In terms of the theoretical results for the Lambertian case, we have derived a low-dimensional representation composed of spherical harmonics up to order 2. This explains previous results in computer vision on low-dimensional subspaces for lighting variation [Hallinan 1994; Epstein et al. 1995]. Recent theoretical work has also attempted to directly link the common principal component methods used to derive low-dimensional subspaces with our spherical harmonic convolution approach [Ramamoorthi

2002a; Nillius and Eklundh 2003a]. Other theoretical and practical work has tried to remove some of the restrictions in our framework, such as the limitation to distant illumination [Annen et al. 2004; Frolova et al. 2004]. There has also been an effort to develop appropriate basis function methods for more general problems, where the convolution results with spherical harmonics may not hold, using hemispherical basis functions [Koenderink and van Doorn 1998; Gautron et al. 2004] or Haar wavelets [Ng et al. 2003, 2004].

More broadly, many of the ideas developed here have links to other areas of graphics and vision, and some work has already begun in exploring and extending these links. For instance, our factorization of the reflected light field to separate the illumination and BRDF bears some conceptual similarities to the use of factorization in many areas of computer vision, such as structure from motion [Tomasi and Kanade 1992]. In this article, we have explored the notion of reflection as convolution for convex opaque surfaces, but the idea of convolution has more general implications for rendering. Convolution results can also be derived for cast shadows under specific assumptions [Soler and Sillion 1998; Ramamoorthi et al. 2004]. These ideas are also very powerful for rendering volumetric materials, where direct numerical simulation is usually very slow. However, practical methods that conceptually use convolution of the incident irradiance with a point-spread function have been employed successfully for subsurface scattering [Jensen et al. 2001; Jensen and Buhler 2002], and we have recently developed such methods for rendering multiple scattering effects in volumetric participating media [Narasimhan et al. 2004; Premoze et al. 2004].

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

APPENDIX

Properties of the Representation Matrices

In this appendix, we derive the two properties of representation matrices listed in Equation 31. The first property follows from the addition theorem for spherical harmonics (see for instance, Jackson [1975] Equation 3.62),

$$Y_{l0}(u, v) = \Lambda_l \sum_{m=-l}^{l} Y^*_{lm}(\theta, \phi) Y_{lm}(\theta', \phi'). \tag{93}$$

Here, v is a dummy-variable since Y_{l0} has no azimuthal dependence, and u is the angle between (θ, ϕ) and (θ', ϕ'), i.e.

$$\cos u = \cos\theta \cos\theta' + \sin\theta \sin\theta' \cos(\phi - \phi'). \tag{94}$$

Now, let $(u, v) = R_\alpha(\theta', \phi')$. Here, $R_\alpha = R_y(\alpha)$. We omit the z rotation since that does not affect Y_{l0} which has no azimuthal dependence. The vector corresponding to coordinates (u, v) is then given by

$$\begin{pmatrix} \sin u \cos v \\ \sin u \sin v \\ \cos u \end{pmatrix} = \begin{pmatrix} \cos\alpha & 0 & \sin\alpha \\ 0 & 1 & 0 \\ -\sin\alpha & 0 & \cos\alpha \end{pmatrix} \begin{pmatrix} \sin\theta' \cos\phi' \\ \sin\theta' \sin\phi' \\ \cos\theta' \end{pmatrix} = \begin{pmatrix} \cos\alpha \sin\theta' \cos\phi' + \sin\alpha \cos\theta' \\ \sin\theta' \sin\phi' \\ \cos\alpha \cos\theta' + \sin\alpha \sin\theta' \left(-\cos\phi'\right) \end{pmatrix}. \quad (95)$$

Since $(-\cos\phi') = \cos(\pi - \phi')$, we know from Equation 94 that u corresponds to the angle between (α, π) and (θ', ϕ'). In other words, we may set $(\theta, \phi) = (\alpha, \pi)$. To summarize,

$$Y_{l0}(R_\alpha(\theta', \phi')) = \Lambda_l \sum_{m=-l}^{l} Y^*_{lm}(\alpha, \pi) Y_{lm}(\theta', \phi'). \quad (96)$$

To proceed further, we write the rotation formula for spherical harmonics, omitting the z rotation by β and γ, since they have no significance for azimuthally symmetric harmonics.

$$Y_{l0}(R_\alpha(\theta', \phi')) = \sum_{m=-l}^{l} d^l_{0m}(\alpha) Y_{lm}(\theta', \phi') \quad (97)$$

A comparision of Equations 96 and 97 yields the first property of representation matrices in Equation 31

$$d^l_{0m}(\alpha) = \Lambda_l Y^*_{lm}(\alpha, \pi). \quad (98)$$

To obtain the second property in Equation 31, we use the form of the spherical harmonic expansion when the elevation angle is 0, i.e. we are at the north pole. Specifically, we note [Jackson 1975] that $Y_{lm'}(0', \phi') = \Lambda_l^{-1} \delta_{m'0}$. With this in mind, the derivation is as follows,

$$\begin{aligned} Y_{lm}(\alpha, \beta) &= Y_{lm}(R_{\alpha,\beta,\gamma}(0', \phi')) \\ &= \sum_{m'=-l}^{l} D^l_{m\,m'}(\alpha, \beta, \gamma) Y_{lm'}(0', \phi') \\ &= \Lambda_l^{-1} D^l_{m0}(\alpha, \beta, \gamma). \end{aligned} \quad (99)$$

This brings us to the second property stated in Equation 31,

$$D^l_{m0}(\alpha, \beta, \gamma) = \Lambda_l Y_{lm}(\alpha, \beta). \quad (100)$$

ACKNOWLEDGMENTS

We thank the reviewers for their painstaking efforts in reading the submitted version of this paper and suggesting several improvements that have significantly improved the readability of the paper.

REFERENCES

Annen, T., Kautz, J., Durand, F., and Seidel, H. 2004. Spherical harmonic gradients for mid-range illumination. In *EuroGraphics Symposium on Rendering*.

Arvo, J. 1995. Analytic methods for simulated light transport. Ph.D. thesis, Yale University.

Basri, R. and Jacobs, D. 2001a. Lambertian reflectance and linear subspaces. In *ICCV 01*. 383–390.

Basri, R. and Jacobs, D. 2001b. Photometric stereo with general, unknown lighting. In *CVPR 01*. II–374–II–381.

Basri, R. and Jacobs, D. 2003. Lambertian reflectance and linear subspaces. *PAMI 25*, 2, 218–233.

Belhumeur, P. and Kriegman, D. 1998. What is the set of images of an object under all possible illumination conditions? *IJCV 28*, 3, 245–260.

BLINN, J. F. AND NEWELL, M. E. 1976. Texture and reflection in computer generated images. *Comm. ACM 19*, 542–546.

CABRAL, B., MAX, N., AND SPRINGMEYER, R. 1987. Bidirectional reflection functions from surface bump maps. In *SIGGRAPH 87*. 273–281.

CABRAL, B., OLANO, M., AND NEMEC, P. 1999. Reflection space image based rendering. In *SIGGRAPH 99*. 165–170.

COCHRAN, J. 1972. *The Analysis of Linear Integral Equations*. McGraw-Hill.

COHEN, M. F. AND WALLACE, J. R. 1993. *Radiosity and Realistic Image Synthesis*. Academic Press.

DANA, K., GINNEKEN, B., NAYAR, S., AND KOENDERINK, J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph. 18*, 1 (January), 1–34.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *SIGGRAPH 00*. 145–156.

DROR, R., ADELSON, E., AND WILLSKY, A. 2001. Estimating surface reflectance properties from images under unknown illumination. In *SPIE Photonics West: Human Vision and Electronic Imaging VI*. 231–242.

D'ZMURA, M. 1991. *Computational Models of Visual Processing*. MIT Press, Chapter Shading Ambiguity: Reflectance and Illumination, 187–207.

EPSTEIN, R., HALLINAN, P., AND YUILLE, A. 1995. 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*. 108–116.

FROLOVA, D., SIMAKOV, D., AND BASRI, R. 2004. Accuracy of spherical harmonic approximations for images of lambertian objects under far and near lighting. In *ECCV*. I–574–I–587.

GAUTRON, P., KRIVANEK, J., PATTANAIK, S., AND BOUATOUCH, K. 2004. A novel hemispherical basis for accurate and efficient rendering. In *EuroGraphics Symposium on Rendering*.

GREENE, N. 1986. Environment mapping and other applications of world projections. *IEEE Comput. Graph. Appl. 6*, 11, 21–29.

GROEMER, H. 1996. *Geometric Applications of Fourier Series and Spherical Harmonics*. Cambridge University Press.

HALLINAN, P. 1994. A low-dimensional representation of human faces for arbitrary lighting conditions. In *CVPR 94*. 995–999.

HEIDRICH, W. AND SEIDEL, H. P. 1999. Realistic, hardware-accelerated shading and lighting. In *SIGGRAPH 99*. 171–178.

HO, J., YANG, M., LIM, J., LEE, K., AND KRIEGMAN, D. 2003. Clustering appearances of objects under varying illumination conditions. In *CVPR*. Vol. 1. 11–18.

IKEUCHI, K. AND SATO, K. 1991. Determining reflectance properties of an object using range and brightness images. *PAMI 13*, 11, 1139–1153.

INUI, T., TANABE, Y., AND ONODERA, Y. 1990. *Group Theory and its Applications in Physics*. Springer Verlag.

JACKSON, J. 1975. *Classical Electrodynamics*. John Wiley.

JENSEN, H. AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graph. (SIGGRAPH 2002) 21*, 3, 576–581.

JENSEN, H., MARSCHNER, S., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *SIGGRAPH 2001*. 511–518.

KAJIYA, J. 1986. The rendering equation. In *SIGGRAPH 86*. 143–150.

KAUTZ, J. AND MCCOOL, M. 2000. Approximation of glossy reflection with prefiltered environment maps. In *Graphics Interface*. 119–126.

KAUTZ, J., VÁZQUEZ, P., HEIDRICH, W., AND SEIDEL, H. 2000. A unified approach to prefiltered environment maps. In *11th Eurographics Workshop on Rendering*. 185–196.

KOENDERINK, J. J. AND VAN DOORN, A. J. 1998. Phenomenological description of bidirectional surface reflection. *JOSA A 15*, 11, 2903–2912.

LAND, E. AND MCCANN, J. 1971. Lightness and retinex theory. *J. Optic. Soc. Amer. 61*, 1, 1–11.

LEE, K., HO, J., AND KRIEGMAN, D. 2001. Nine points of light: Acquiring subspaces for face recognition under variable lighting. In *CVPR*. 519–526.

LEWIS, R. 1993. Making shaders more physically plausible. In *Eurographics Workshop on Rendering*. 47–62.

LOVE, R. 1997. Surface reflection model estimation from naturally illuminated image sequences. Ph.D. thesis, Leeds.

LU, R., KOENDERINK, J., AND KAPPERS, A. 1998. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics 37*, 25, 5974–5984.

MACROBERT, T. 1948. *Spherical Harmonics; an Elementary Treatise on Harmonic Functions, with Applications*. Dover Publications.

MARSCHNER, S. AND GREENBERG, D. 1997. Inverse lighting for photography. In *Fifth Color Imaging Conference*. 262–265.

MARSCHNER, S., WESTIN, S., LAFORTUNE, E., AND TORRANCE, K. 2000. Image-Based BRDF measurement. *Applied Optics 39*, 16, 2592–2600.

McCormick, N. 1992. Inverse radiative transfer problems: a review. *Nucl. Sci. Eng. 112*, 185–198.

Miller, G. and Hoffman, C. 1984. Illumination and reflection maps: Simulated objects in simulated and real environments. *SIGGRAPH 84 Advanced Computer Graphics Animation seminar notes*.

Narasimhan, S., Ramamoorthi, R., and Nayar, S. 2004. Analytic rendering of multiple scattering in participating media. *Submitted to ACM Trans. Graph.*.

Nayar, S., Ikeuchi, K., and Kanade, T. 1990. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Trans. Robotics Auto. 6*, 4 (Aug), 418–430.

Ng, R., Ramamoorthi, R., and Hanrahan, P. 2003. All-frequency shadows using nonlinear wavelet lighting approximation. *ACM Trans. Graph. (SIGGRAPH 2003) 22*, 3.

Ng, R., Ramamoorthi, R., and Hanrahan, P. 2004. Triple product wavelet integrals for all-frequency relighting. *ACM Trans. Graph. (SIGGRAPH 2004) 23*, 3.

Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., and Limperis, T. 1977. *Geometric Considerations and Nomenclature for Reflectance*. National Bureau of Standards (US).

Nillius, P. and Eklundh, J. 2003a. Low-dimensional representations of shaded surfaces under varying illumination. In *CVPR 03*. II:185–II:192.

Nillius, P. and Eklundh, J. 2003b. Phenomenological eigenfunctions for irradiance. In *ICCV 03*. I:568–I:575.

Nillius, P. and Eklundh, J. 2004. Classifying materials from their reflectance properties. In *ECCV 04*. IV:366–IV:376.

Nimeroff, J., Simoncelli, E., and Dorsey, J. 1994. Efficient re-rendering of naturally illuminated environments. In *EGWR 94*. 359–373.

Nishino, K., Zhang, Z., and Ikeuchi, K. 2001. Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis. In *ICCV 01*. 599–606.

Preisendorfer, R. 1976. *Hydrologic Optics*. US Dept Commerce.

Premoze, S., Ashikhmin, M., Ramamoorthi, R., and Nayar, S. 2004. Practical rendering of multiple scattering effects in participating media. In *EuroGraphics Symposium on Rendering*.

Ramamoorthi, R. 2002a. Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI) 24*, 10 (Oct), 1322–1333.

Ramamoorthi, R. 2002b. A signal-processing framework for forward and inverse rendering. Ph.D. thesis, Stanford University.

Ramamoorthi, R. and Hanrahan, P. 2001a. Analysis of planar light fields from homogeneous convex curved surfaces under distant illumination. In *SPIE Photonics West: Human Vision and Electronic Imaging VI*. 185–198. http://graphics.stanford.edu/papers/planarlf/.

Ramamoorthi, R. and Hanrahan, P. 2001b. An efficient representation for irradiance environment maps. In *SIGGRAPH 01*. 497–500.

Ramamoorthi, R. and Hanrahan, P. 2001c. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *JOSA A 18*, 10, 2448–2459.

Ramamoorthi, R. and Hanrahan, P. 2001d. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*. 117–128.

Ramamoorthi, R. and Hanrahan, P. 2002. Frequency space environment map rendering. In *SIGGRAPH 02*. 517–526.

Ramamoorthi, R., Koudelka, M., and Belhumeur, P. 2004. A Fourier theory for cast shadows. In *ECCV*. I–146–I–162.

Sato, I., Okabe, T., Sato, Y., and Ikeuchi, K. 2003. Appearance sampling for obtaining a set of basis images for variable illumination. In *ICCV*. II–800–II–807.

Sato, I., Sato, Y., and Ikeuchi, K. 1999. Illumination distribution from brightness in shadows: adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. In *ICCV 99*. 875 – 882.

Sato, Y. and Ikeuchi, K. 1994. Reflectance analysis under solar illumination. Tech. Rep. CMU-CS-94-221, CMU.

Sato, Y., Wheeler, M. D., and Ikeuchi, K. 1997. Object shape and reflectance modeling from observation. In *SIGGRAPH 97*. 379–388.

Sillion, F. X., Arvo, J., Westin, S. H., and Greenberg, D. 1991. A global illumination solution for general reflectance distributions. In *SIGGRAPH 91*. 187–196.

Simakov, D., Frolova, D., and Basri, R. 2003. Dense shape reconstruction of a moving object under arbitrary, unknown lighting. In *ICCV 03*. 1202–1209.

Sloan, P., Kautz, J. and Snyder, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *SIGGRAPH 02*. 527–536.

Sloan, P., Hall, J., Hart, J., and Snyder, J. 2003. Clustered principal components for precomputed radiance transfer. *ACM Tran. Graph. (SIGGRAPH 03 proceedings) 22*, 3.

SOLER, C. AND SILLION, F. 1998. Fast calculation of soft shadow textures using convolution. In *SIGGRAPH 98*. 321–332.

THORNBER, K. AND JACOBS, D. 2001. Broadened, specular reflection and linear subspaces. Tech. Rep. TR#2001-033, NEC.

TOMASI, C. AND KANADE, T. 1992. Shape and motion from image streams under orthography: a factorization method. *Int. J. Comput. Vision 9*, 2, 137–154.

TORRANCE, K. E. AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *JOSA 57*, 9, 1105–1114.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92*. 265–272.

WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92*.

YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *SIGGRAPH 99*. 215–224.

YU, Y. AND MALIK, J. 1998. Recovering photometric properties of architectural scenes from photographs. In *SIGGRAPH 98*. 207–218.

ZHANG, L., CURLESS, B., HERTZMANN, A., AND SEITZ, S. 2003. Shape and motion under varying illumination: Unifying multiview stereo, photometric stereo, and structure from motion. In *International Conference on Computer Vision*. 618–625.

ZHANG, L. AND SAMARAS, D. 2003. Face recognition under variable lighting using harmonic image exemplars. In *CVPR*. I:19–I:25.

Received August 2002; revised February 2004; accepted June 2004