A Signal-Processing Framework for Reflection—Part 1: Reflection as Convolution

Ravi Ramamoorthi, Columbia University*
Pat Hanrahan, Stanford University

Abstract

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 first part of the 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. Our mathematical analysis also has implications for *forward rendering*—especially the efficient rendering of objects under complex lighting conditions specified by environment maps.

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

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

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 is a basic building block in most rendering algorithms. In computer vision, we often want to undo the effects of the reflection operator, i.e. to invert the interaction between the BRDF and 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 paper 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 [20]. 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 [28] 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, i.e. 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 [28], Cabral et al. [5, 6], D'Zmura [12] 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, i.e. by filtering the incident illumination using the BRDF. Mathematically, we are able to express the frequency-space coefficients of the reflected light field as a product of the

spherical harmonic coefficients of the illumination and the BRDF. We believe this is a useful way of analyzing many forward and inverse problems. In particular, forward rendering can be viewed as *convolution* and inverse rendering as *deconvolution*.

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 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 [35, 38] and inverse rendering [37]. Basri and Jacobs have applied these ideas in computer vision to lighting-invariant recognition [1] and photometric stereo [2], while Sloan et al. [32] have demonstrated real-time rendering with low-frequency lighting and complex light transport. Therefore, we feel that it would be worthwhile to present a complete, detailed and unified account of the theoretical framework for archival purposes. This paper builds on and extends previous theoretical work by us [34, 36, 37]. As discussed at the end of the next section, we have also included a significant amount of new material and detail that has not previously appeared.

The goal of this paper is to present a unified, detailed and complete description of the mathematical foundation underlying these previous articles. We will briefly point out the practical implications of the results derived in this paper, but will refer the reader to previously published work for implementation details. The rest of this paper is organized as follows. In section 2, we discuss previous work. Section 3 introduces the reflection equation in 2D and 3D, showing how it can be viewed as a convolution. Section 4 carries out a formal frequency-space analysis of the reflection equation. Section 5 applies these results to many problems in inverse rendering, showing which problems are well posed and well conditioned versus ill-posed or ill-conditioned. Finally, section 6 concludes the paper and discusses future work.

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

Forward Rendering by Environment Mapping: The theoretical analysis in this paper employs essentially the same assumptions typically made in rendering with environment maps, i.e. distant illumination—allowing the lighting to be represented by a single environment map—incident on curved surfaces. Blinn and Newell [4] first used environment maps to efficiently find reflections of distant objects. The technique was generalized by Miller and Hoffman [28] and Greene [15] who precomputed diffuse and specular reflection maps, allowing for images with complex realistic lighting and BRDFs to be synthesized. Cabral et al. [5] later extended this general method to computing reflections from bump-mapped surfaces, and to computing environment-mapped images with arbitrary BRDFs [6]. It should be noted that both Miller and Hoffman [28], and Cabral et al. [5, 6] qualitatively described the reflection maps as obtained by convolving the lighting with the BRDF. In this paper, we will formalize these ideas, making the notion of convolution precise, and derive analytic formulae. We have already shown [35, 38] how this theoretical analysis leads practically to a significantly more efficient method for prefiltering and rendering with environment maps.

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 [28]. A more sophisticated *inverse lighting* approach is that of Marschner and Greenberg [25], who try to find the lighting under the assumption of a Lambertian BRDF. D'Zmura [12] proposes estimating spherical harmonic coefficients of the lighting.

Most work in inverse rendering has focused on BRDF [29] estimation. Recently, image-based BRDF measurement methods have been proposed in 2D by Lu et al. [23] and in 3D by Marschner et al. [26]. 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 [44] or Torrance and

Sparrow [43]. Parametric models are often preferred in practice since they are compact, and are simpler to estimate. A number of methods [9, 10, 41, 46] 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 [11, 22, 31, 39, 40, 47], although many of those methods are specific to a particular illumination model. Using the theoretical analysis described in this paper, we [37] have presented a general method for complex illumination, that handles the various components of the lighting and BRDF in a principled manner to allow for BRDF estimation under general lighting conditions. Furthermore, we will show that it is possible in theory to separately estimate the lighting and BRDF, up to a global scale factor. We have been able to use these ideas to develop a practical method [37] of *factoring* the light field to simultaneously determine the lighting and BRDF for geometrically complex objects.

Frequency-Space Representations: Since we are going to treat reflection as a convolution and analyze it in frequency-space, we will briefly discuss previous work on frequency-space representations. Since we will be primarily concerned with analyzing quantities like the BRDF and distant lighting which can be parameterized as a function on the unit sphere, the appropriate frequency-space representations are spherical harmonics [18, 19, 24]. The use of spherical harmonics to represent the illumination and BRDF was pioneered by Cabral et al. [5]. D'Zmura [12] analyzed reflection as a linear operator in terms of spherical harmonics, and discussed some resulting ambiguities between reflectance and illumination. We extend his work by explicitly deriving the frequency-space reflection equation (i.e. convolution formula) in this first part of the paper, and by providing quantitative results for various special cases in the second part of the paper. Our use of spherical harmonics to represent the lighting is similar in some respects to previous methods such as that of Nimeroff et al. [30] that use steerable linear basis functions. Spherical harmonics, as well as the closely related Zernike polynomials, have also been used before in computer graphics for representing BRDFs by a number of other authors [21, 42, 45].

Formal Analysis of Reflection: This paper conducts a formal study of the reflection operator by showing mathematically that it can be described as a convolution, deriving an analytic formula for the resulting convolution equation, and using this result to study the well-posedness and conditioning of several inverse problems. As such, our approach is similar in spirit to mathematical methods used to study inverse problems in other areas of radiative transfer and transport theory such as hydrologic optics [33] and neutron scattering. See McCormick [27] 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 [3] have theoretically described this set of images in terms of an illumination cone, while empirical results have been obtained by Hallinan [17] and Epstein et al. [13]. 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 [1] have described Lambertian reflection as a convolution and obtained similar analytic results for that particular case.

This paper builds on previous theoretical work by us on planar or flatland light fields [34], on the reflected light field from a Lambertian surface [36], and on the theory for the general 3D case with isotropic BRDFs [37]. The goal of this paper is to present a unified, complete and 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, a group-theoretic interpretation in terms of generalized convolutions, the relationship to the theory of Fredholm integral equations of the first kind, and explicit verification of many of the frequency-space formulae from first principles, which have not been discussed in earlier work. Besides the new material added, we have also significantly expanded our previous accounts with details and illustrations.

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 paper 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 goal is to analyze this variation in terms of the incident illumination and the surface BRDF. Our theory will be based on the fact that different orientations of a curved surface correspond to different orientations of the upper hemisphere and

BRDF. Equivalently, each orientation of the surface corresponds to a different integral over the lighting, and the reflected light field will therefore be a function of surface orientation.

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

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

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

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

Discussion: We note that for the most part, our assumptions are very similar to those made in most interactive graphics applications, including environment map rendering algorithms such as those of Miller and Hoffman [28] and Cabral et al. [6]. Our assumptions also accord closely with those usually made in computer vision and inverse rendering. The only significant additional assumption is that of homogeneous surfaces. However, this is not particularly restrictive since spatially varying BRDFs are often approximated in practical graphics or vision applications by using a spatially varying texture that simply modulates one or more components of the BRDF. This can be incorporated into the ensuing theoretical analysis by merely multiplying the reflected light field by a texture dependent on surface position. We believe that our assumptions are a good approximation to many real-world situations, while being simple enough to treat analytically. Furthermore, it is likely that the insights obtained from the analysis in this paper will be applicable even in cases where the assumptions are not exactly satisfied. We have already demonstrated [37] 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 paper 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

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	$\sqrt{-1}$

Figure 1: *Notation*

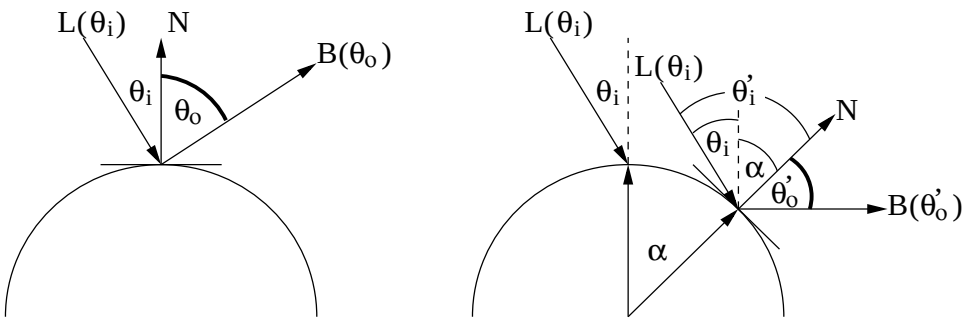

Figure 2: *Schematic of reflection in 2D. On the left, we show the situation with respect to one point on the surface (the north pole or* 0° *location, where global and local coordinates are the same). The right figure shows the effect of the surface orientation* α*. Different orientations of the surface 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*

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

In local coordinates, we can write the reflection equation as

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

Here, B is the reflected radiance, L is the incident radiance, i.e 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 $\mathbf{N} = [\sin\alpha, \cos\alpha]$, i.e. $B(\mathbf{x}, \theta_o') \rightarrow B(\alpha, \theta_o')$ 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, i.e. use a single function L regardless of surface position. In other words, the lighting is a function only of the global incident angle, $L(\alpha, \theta_i') \rightarrow L(\theta_i)$. Finally, we define a transfer function $\hat{\rho} = \rho \cos\theta_i'$ to absorb the cosine term in the integrand. With these modifications, equation 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, i.e $0'$ is the surface normal. The corresponding global coordinates are clearly α. We define R_α as an operator that rotates θ_i' into global coordinates, and is given in 2D simply by $R_\alpha(\theta_i') = \alpha + \theta_i'$. To convert from global to local coordinates, we apply the inverse rotation, i.e. $R_{-\alpha}$. To summarize,

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

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.

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

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

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

$$B(\alpha, \theta'_o) = \int_{-\pi/2}^{\pi/2} L\left(\alpha + \theta'_i\right) \hat{\rho}\left(\theta'_i, \theta'_o\right)\, d\theta'_i \tag{6}$$

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

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

$$B(\mathbf{x}, \theta'_o, \phi'_o) = \int_{\Omega'_i} L(\mathbf{x}, \theta'_i, \phi'_i) \rho(\theta'_i, \phi'_i, \theta'_o, \phi'_o) \cos\theta'_i \, d\omega'_i. \tag{8}$$

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

We can make the same substitutions that we did in 2D. We reparameterize the surface position $\mathbf{x}$ 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, $B(\mathbf{x}, \theta'_o, \phi'_o) \rightarrow$

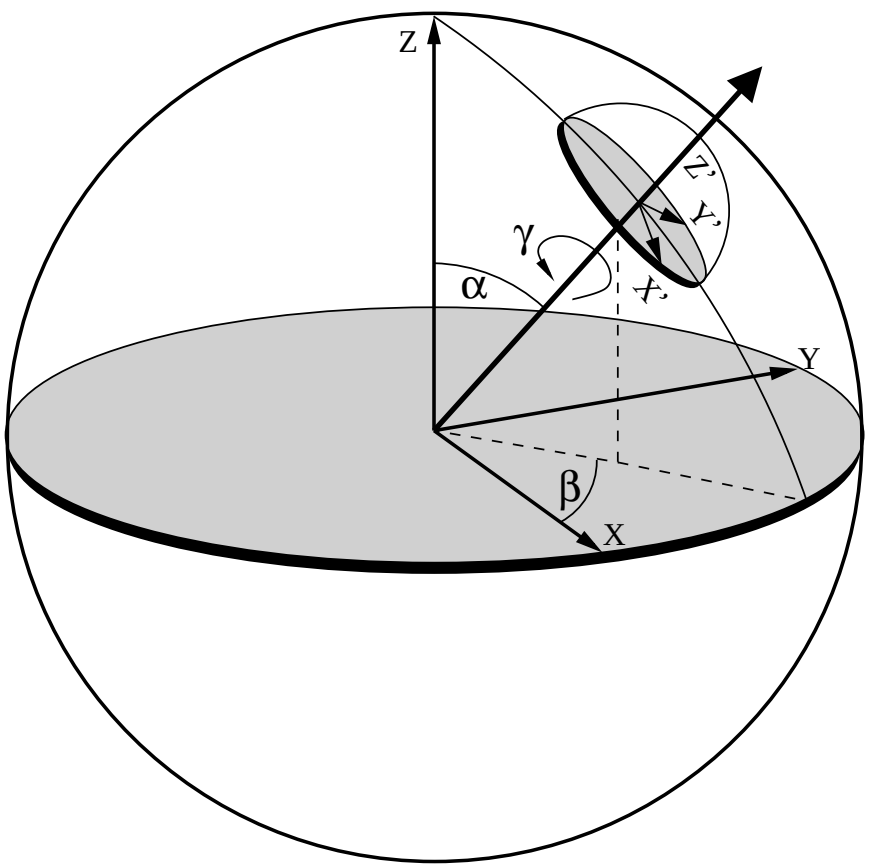

Figure 3: *Diagram showing how the rotation corresponding to* (α, β, γ) *transforms between local (primed) and global (unprimed) coordinates. The net rotation is composed of three independent rotations about Z, Y, and Z, with the angles* α*,* β*, and* γ *corresponding directly to the Euler angles.*

$B(\alpha, \beta, \gamma, \theta'_o, \phi'_o)$ 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,

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

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,

$$B(\alpha, \beta, \gamma, \theta_o', \phi_o') = \int_{\Omega_i'} L\left(R_{\alpha,\beta,\gamma}(\theta_i', \phi_i')\right) \hat{\rho}(\theta_i', \phi_i', \theta_o', \phi_o')\, d\omega_i' \tag{11}$$

$$= \int_{\Omega_i} L(\theta_i, \phi_i)\hat{\rho}\left(R^{-1}_{\alpha,\beta,\gamma}(\theta_i, \phi_i), \theta_o', \phi_o'\right)\, d\omega_i. \tag{12}$$

As we have written them, these equations depend on spherical coordinates. It might clarify matters somewhat to also present an alternate form in terms of rotations and unit vectors in a coordinate-independent way. We simply use R for the rotation, which could be written as a 3×3 rotation matrix, while ω_i and ω_o stand for unit vectors corresponding to the incident and outgoing directions (with primes added for local coordinates). Equations 11 and 12 may then be written as

$$B(R, \omega_o') = \int_{\Omega_i'} L\left(R\omega_i'\right) \hat{\rho}(\omega_i', \omega_o')\, d\omega_i' \tag{13}$$

$$= \int_{\Omega_i} L(\omega_i)\hat{\rho}\left(R^{-1}\omega_i, \omega_o'\right)\, d\omega_i, \tag{14}$$

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

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

$$(f \otimes g)(a) = \int_t f\left(T_a(t)\right) g(t)\, dt. \tag{15}$$

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

Therefore, equations 11 and 12 (or 13 and 14) represent rotational convolutions. Equation 12 in global coordinates states that the reflected light field at a given surface orientation corresponds to *rotating* the BRDF to that orientation, and then integrating over the upper hemisphere. The BRDF can be thought of as the filter, while the lighting is the input signal. Symmetrically, equation 11 in local coordinates states that the reflected light field at a given surface orientation may be computed by *rotating* the lighting into the local coordinate system of the BRDF, and then doing the hemispherical integration. These observations are similar to those we made earlier for the 2D case.

Group-theoretic Interpretation as Generalized Convolution: In fact, it is possible to formally generalize the notion of convolution to groups. Within this context, the standard Fourier convolution formula can be seen as a special case for $SO(2)$, the group of rotations in 2D. More information may be found in books on group representation theory, such as Fulton and Harris [14] (especially note exercise 3.32). One reference that focuses specifically on the rotation group is Chirikjian and Kyatkin [7]. In the general case, we may modify equation 15 slightly to write for compact groups,

$$(f \otimes g)(s) = \int_t f(s \circ t) g(t)\, dt, \tag{16}$$

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

It is also possible to generalize the Fourier convolution formula in terms of representation matrices of the group in question. In our case, the relations do not exactly satisfy equation 16, since we have both rotations (in the rotation group $SO(3)$) and unit vectors. Therefore, for frequency space analysis in the 3D case, we will need both the representation matrices of $SO(3)$, and the associated basis functions for unit vectors on a sphere, which are the spherical harmonics.

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 series of a function f by

$$\begin{aligned} F_k(\theta) &= \frac{1}{\sqrt{2\pi}} e^{Ik\theta} \\ f(\theta) &= \sum_{k=-\infty}^{\infty} f_k F_k(\theta) \\ f_k &= \int_{-\pi}^{\pi} f(\theta) F_k^*(\theta) d\theta. \end{aligned} \tag{17}$$

In the last line, the $*$ in the superscript stands for the complex conjugate. For the Fourier basis functions, $F_k^* = F_{-k} = (1/\sqrt{2\pi}) \exp(-Ik\theta)$. It should be noted that the relations in equation 17 are similar for any orthonormal basis functions F, and we will later be able to use much of the same machinery to define spherical harmonic expansions in 3D.

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

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

$$L(\theta_i) = \sum_{l=-\infty}^{\infty} L_l F_l(\theta_i). \tag{18}$$

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

$$\begin{aligned} L(\theta_i) = L(\alpha + \theta_i') &= \sum_{l=-\infty}^{\infty} L_l F_l(\alpha + \theta_i') \\ &= \sqrt{2\pi} \sum_{l=-\infty}^{\infty} L_l F_l(\alpha) F_l(\theta_i'). \end{aligned} \tag{19}$$

The last line follows from the form of the complex exponentials, i.e. $F_l(\alpha+\theta_i') = (1/\sqrt{2\pi}) \exp(Il(\alpha + \theta_i'))$. 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,

$$\begin{aligned} B(\alpha, \theta_o') &= \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} B_{lp} F_l(\alpha) F_p(\theta_o') \\ \hat{\rho}(\theta_i', \theta_o') &= \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} \hat{\rho}_{lp} F_l^*(\theta_i') F_p(\theta_o'). \end{aligned} \tag{20}$$

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

Fourier-Space Reflection Equation: We are now ready to write equation 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

$$B(\alpha, \theta_o') = \sqrt{2\pi} \sum_{l=-\infty}^{\infty} \sum_{l'=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} L_l \hat{\rho}_{l'p} F_l(\alpha) F_p(\theta_o') \int_{-\pi}^{\pi} F_{l'}^*(\theta_i') F_l(\theta_i')\, d\theta_i'. \tag{21}$$

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

$$\int_{-\pi}^{\pi} F_l'^*(\theta_i') F_l(\theta_i')\, d\theta_i' = \delta_{ll'}. \tag{22}$$

In other words, we can set $l' = l$ since terms not satisfying this condition vanish. Making this substitution in equation 21, we obtain

$$B(\alpha, \theta'_o) = \sqrt{2\pi} \sum_{l=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} L_l \hat{\rho}_{lp} F_l(\alpha) F_p(\theta'_o). \tag{23}$$

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

$$B_{lp} = \sqrt{2\pi} L_l \hat{\rho}_{lp}. \tag{24}$$

This result reiterates once more that the reflection equation can be viewed as a convolution of the incident illumination and BRDF, and becomes a simple product in Fourier space, with an analytic formula being given by equation 24.

An alternative form of equation 24 that may be more instructive results from holding the local outgoing angle fixed, instead of expanding it also in terms of Fourier coefficients, i.e. replacing the index p by the outgoing angle θ'_o,

$$B_l(\theta'_o) = \sqrt{2\pi} L_l \hat{\rho}_l(\theta'_o). \tag{25}$$

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

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

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 [18, 19, 24].

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

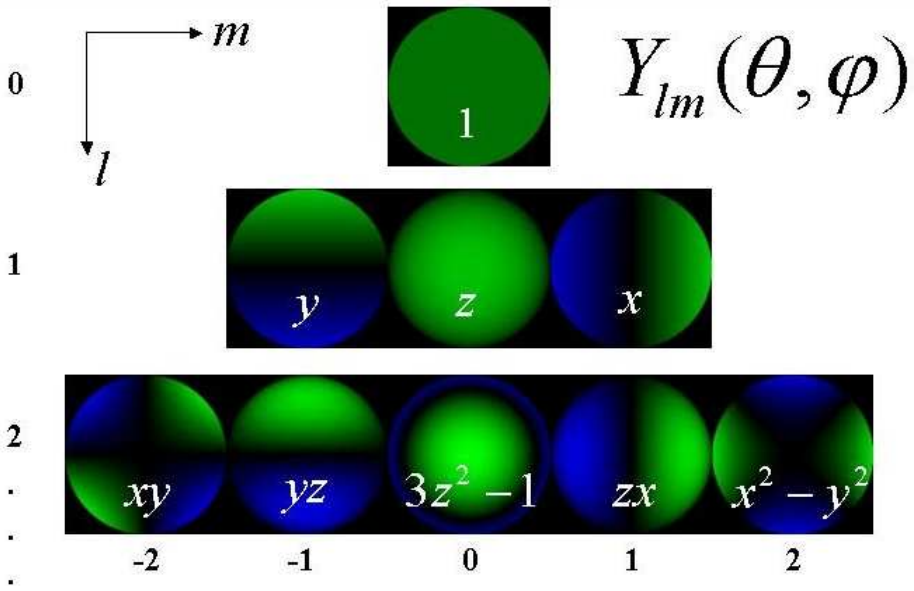

Figure 4: *The first 3 orders of real spherical harmonics ($l = 0, 1, 2$) corresponding to a total of 9 basis functions. The spherical harmonics Y_{lm} may be written either as trigonometric functions of the spherical coordinates θ and ϕ or as polynomials of the cartesian components x, y and z, with $x^2+y^2+z^2 = 1$. In general, a spherical harmonic Y_{lm} is a polynomial of maximum degree l. In these images, we show only the front the sphere, with green denoting positive values and blue denoting negative values. Also note that these images show the real form of the spherical harmonics. The complex forms are given in equation 27.*

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

$$\begin{aligned} N_{lm} &= \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} \\ Y_{lm}(\theta,\phi) &= N_{lm} P_{lm}(\cos\theta) e^{Im\phi}, \end{aligned} \tag{26}$$

where N_{lm} is a normalization factor. In the above equation, the azimuthal dependence is expanded in terms of Fourier basis functions. The θ dependence is expanded in terms of the associated Legendre functions P_{lm}. The indices obey $l \geq 0$ and $-l \leq m \leq l$. Thus, there are $2l + 1$ basis functions for given order l. Figure 4 shows the first 3 orders of spherical harmonics, i.e. the first 9 basis functions corresponding to $l = 0, 1, 2$. They may be written either as trigonometric functions of the spherical coordinates θ and ϕ or as polynomials of the cartesian components x, y and z, with $x^2 + y^2 + z^2 = 1$. In general, a spherical harmonic Y_{lm} is a polynomial of maximum degree l. Another useful relation is that $Y_{l-m} = (-1)^m Y^*_{lm}$. The first 3 orders (we give only terms with

$m \geq 0$) may be written as

$$\begin{aligned}
Y_{00} &= \sqrt{\frac{1}{4\pi}} \\
Y_{10} &= \sqrt{\frac{3}{4\pi}}\cos\theta &&= \sqrt{\frac{3}{4\pi}}z \\
Y_{11} &= -\sqrt{\frac{3}{8\pi}}\sin\theta e^{I\phi} &&= -\sqrt{\frac{3}{8\pi}}\left(x+Iy\right) \\
Y_{20} &= \frac{1}{2}\sqrt{\frac{5}{4\pi}}\left(3\cos^2\theta-1\right) &&= \frac{1}{2}\sqrt{\frac{5}{4\pi}}\left(3z^2-1\right) \\
Y_{21} &= -\sqrt{\frac{15}{8\pi}}\sin\theta\cos\theta e^{I\phi} &&= -\sqrt{\frac{15}{8\pi}}z\left(x+Iy\right) \\
Y_{22} &= \frac{1}{2}\sqrt{\frac{15}{8\pi}}\sin^2\theta e^{2I\phi} &&= \frac{1}{2}\sqrt{\frac{15}{8\pi}}\left(x+Iy\right)^2.
\end{aligned} \tag{27}$$

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

$$\begin{aligned}
f(\theta,\phi) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l} f_{lm}Y_{lm}(\theta,\phi) \\
f_{lm} &= \int_{\phi=0}^{2\pi}\int_{\theta=0}^{\pi} f(\theta,\phi)Y^*_{lm}(\theta,\phi)\sin\theta\, d\theta d\phi.
\end{aligned} \tag{28}$$

Note the close parallel with equation 17.

The rotation formula for spherical harmonics is

$$Y_{lm}\left(R_{\alpha,\beta,\gamma}(\theta,\phi)\right) = \sum_{m'=-l}^{l} D^l_{mm'}(\alpha,\beta,\gamma)Y_{lm'}(\theta,\phi). \tag{29}$$

The important thing to note here is that the m indices are *mixed*—a spherical harmonic after rotation must be expressed as a combination of other spherical harmonics with different m indices. However, the l indices are not mixed; rotations of spherical harmonics with order l are composed entirely of other spherical harmonics with order l. For given order l, D^l is a matrix that tells us how a spherical harmonic transforms under rotation, i.e. how to rewrite a rotated spherical harmonic as a linear combination of all the spherical harmonics of the same order. In terms of group theory, the matrix D^l is the $(2l+1)$-dimensional representation of the rotation group $SO(3)$. A pictorial depiction of equation 29 as a matrix multiplication is found in figure 5.

The matrices D^l therefore satisfy the formula,

$$D^l_{mm'}(\alpha,\beta,\gamma) = \int_{\phi=0}^{2\pi}\int_{\theta=0}^{\pi} Y_{lm}(R_{\alpha,\beta,\gamma}(\theta,\phi))Y^*_{lm'}(\theta,\phi)\sin\theta\, d\theta\, d\phi. \tag{30}$$

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

Figure 5: *Depiction of equation 29 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.*

An analytic form for the matrices D^l can be found in standard references, such as Inui et al. [18]. 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, i.e.

$$D^l_{mm'}(\alpha, \beta, \gamma) = d^l_{mm'}(\alpha) e^{Im\beta} e^{Im'\gamma}, \tag{31}$$

where d^l is a matrix that defines how a spherical harmonic transforms under rotation about the $y-$axis. For the purposes of the exposition, we will not generally need to be concerned with the precise formula for the matrix d^l, and numerical calculations can compute it using a simplified version of equation 30 without the z rotations (i.e. $\beta = \gamma = 0$),

$$d^l_{mm'}(\alpha) = \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} Y_{lm}(R_y(\alpha)(\theta, \phi)) Y^*_{lm'}(\theta, \phi) \sin\theta \, d\theta \, d\phi. \tag{32}$$

For completeness, we give below the relatively complicated analytic formula, as derived in equation 7.48 of Inui et al. [18],

$$\begin{aligned} \xi &= \sin^2 \frac{\alpha}{2} \\ N(l, m, m') &= (-1)^{m-m'} \sqrt{\frac{(l+m)!}{(l-m)!(l+m')!(l-m')!}} \\ d^l_{mm'}(\alpha) &= N(l, m, m') \times \xi^{-(m-m')/2} (1-\xi)^{-(m+m')/2} \left(\frac{d}{d\xi}\right)^{l-m} \xi^{l-m'} (1-\xi)^{l+m'}, \end{aligned} \tag{33}$$

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

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

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

$$d^2(\alpha) = \begin{pmatrix} \cos^4\frac{\alpha}{2} & 2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & 2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & \sin^4\frac{\alpha}{2} \\ -2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \cos^2\frac{\alpha}{2}(-1+2\cos\alpha) & \sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & (1+2\cos\alpha)\sin^2\frac{\alpha}{2} & 2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} \\ \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & -\sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \frac{1}{2}(3\cos^2\alpha-1) & \sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha \\ -2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & (1+2\cos\alpha)\sin^2\frac{\alpha}{2} & -\sqrt{\frac{3}{2}}\cos\alpha\sin\alpha & \cos^2\frac{\alpha}{2}(-1+2\cos\alpha) & 2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} \\ \sin^4\frac{\alpha}{2} & -2\cos\frac{\alpha}{2}\sin^3\frac{\alpha}{2} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\alpha & -2\cos^3\frac{\alpha}{2}\sin\frac{\alpha}{2} & \cos^4\frac{\alpha}{2} \end{pmatrix}. \tag{34}$$

To derive some of the quantitative results in section 5, we will require two important properties of the representation matrices D^l, which are derived in the appendix,

$$\begin{aligned} D^l_{0m'}(\alpha,\beta,0) = &\quad d^l_{0m'}(\alpha) &&= \sqrt{\frac{4\pi}{2l+1}}Y^*_{lm'}(\alpha,\pi) \\ D^l_{m0}(\alpha,\beta,\gamma) = &\quad d^l_{m0}(\alpha)e^{Im\beta} &&= \sqrt{\frac{4\pi}{2l+1}}Y_{lm}(\alpha,\beta). \end{aligned} \tag{35}$$

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

$$L(\theta_i,\phi_i) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} L_{lm}Y_{lm}(\theta_i,\phi_i). \tag{36}$$

To obtain the lighting in local coordinates, we must rotate the above expression, just as we did in 2D. Using equation 29, we get,

$$L(\theta_i,\phi_i) = L\left(R_{\alpha,\beta,\gamma}(\theta'_i,\phi'_i)\right) = \sum_{l=0}^{\infty}\sum_{m=-l}^{+l}\sum_{m'=-l}^{l} L_{lm}D^l_{mm'}(\alpha,\beta,\gamma)Y_{lm'}(\theta'_i,\phi'_i). \tag{37}$$

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

$$\hat{\rho}(\theta'_i,\phi'_i,\theta'_o,\phi'_o) = \sum_{l=0}^{\infty}\sum_{n=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-p}^{p}\hat{\rho}_{ln,pq}Y^*_{ln}(\theta'_i,\phi'_i)Y_{pq}(\theta'_o,\phi'_o) \tag{38}$$

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

$$\begin{aligned} B(\alpha,\beta,\gamma,\theta'_o,\phi'_o) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{m'=-l}^{l}\sum_{l'=0}^{\infty}\sum_{n=-l'}^{l'}\sum_{p=0}^{\infty}\sum_{q=-p}^{p} L_{lm}\hat{\rho}_{l'n,pq}D^l_{mm'}(\alpha,\beta,\gamma)Y_{pq}(\theta'_o,\phi'_o)T_{lm'l'n} \\ T_{lm'l'n} &= \int_{\phi'_i=0}^{2\pi}\int_{\theta'_i=0}^{\pi} Y_{lm'}(\theta'_i,\phi'_i)Y^*_{l'n}(\theta'_i,\phi'_i)\,\sin\theta'_i\,d\theta'_i d\phi'_i \\ &= \delta_{ll'}\delta_{m'n}. \end{aligned} \tag{39}$$

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

$$B(\alpha,\beta,\gamma,\theta'_o,\phi'_o) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{n=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-p}^{p} L_{lm}\hat{\rho}_{ln,pq}\left(D^l_{mn}(\alpha,\beta,\gamma)Y_{pq}(\theta'_o,\phi'_o)\right). \tag{40}$$

This result suggests that we should expand the reflected light field B in terms of the new basis functions given by $C_{lmnpq} = D^l_{mn}(\alpha,\beta,\gamma)Y_{pq}(\theta'_o,\phi'_o)$. The appearance of the matrix D^l in these basis functions is quite intuitive, coming directly from the rotation formula for spherical harmonics. These basis functions are *mixed* in the sense that they are a product of the matrices D^l and the spherical harmonics Y_{pq}. This can be understood from realizing that the reflected direction is a unit vector described by two parameters (θ'_o,ϕ'_o), while the surface parameterization is really a rotation, described by three parameters (α,β,γ). Finally, we need to consider the normalization of these new basis functions. The spherical harmonics are already orthonormal. The orthogonality relation for the matrices D^l is given in any standard text on group theory (for instance, equation 7.73 of Inui et al. [18]). Specifically,

$$\int_{\gamma=0}^{2\pi}\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi}\left(D^l_{mn}(\alpha,\beta,\gamma)\right)^*\left(D^{l'}_{m'n'}(\alpha,\beta,\gamma)\right)\sin\alpha\,d\alpha\,d\beta\,d\gamma = \frac{8\pi^2}{2l+1}\delta^{ll'}\delta_{mm'}\delta_{nn'}. \tag{41}$$

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

$$\begin{aligned} C_{lmnpq} &= \sqrt{\frac{2l+1}{8\pi^2}}D^l_{mn}(\alpha,\beta,\gamma)Y_{pq}(\theta'_o,\phi'_o) \\ B &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{n=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-p}^{p} B_{lmnpq}C_{lmnpq}(\alpha,\beta,\gamma,\theta'_o,\phi'_o) \\ B_{lmnpq} &= \int_{\phi'_o=0}^{2\pi}\int_{\theta'_o=0}^{\pi}\int_{\gamma=0}^{2\pi}\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi} B(\alpha,\beta,\gamma,\theta'_o,\phi'_o)C^*_{lmnpq}(\alpha,\beta,\gamma,\theta'_o,\phi'_o)\sin\alpha\sin\theta'_o\,d\alpha\,d\beta\,d\gamma\,d\theta'_o\,d\phi'_o. \end{aligned} \tag{42}$$

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

Finally, we can write down the frequency space reflection equation by comparing equations 40 and 42 and equating coefficients. This result is comparable to its 2D counterpart, given in equation 24, and as in 2D, is a convolution. In frequency-space, the reflected light field is obtained simply by multiplying together coefficients of the lighting and BRDF, i.e. by *convolving* the incident illumination with the BRDF,

$$B_{lmnpq} = \sqrt{\frac{8\pi^2}{2l+1}} L_{lm} \hat{\rho}_{ln,pq}. \tag{43}$$

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

$$B_{lmn}(\theta'_o, \phi'_o) = \sqrt{\frac{8\pi^2}{2l+1}} L_{lm} \hat{\rho}_{ln}(\theta'_o, \phi'_o). \tag{44}$$

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

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

4.3 Group-theoretic Unified Analysis

While not required for understanding the rest of this paper, it is insightful to attempt to analyze the 2D and 3D derivations as special cases of a more general convolution formula in terms of the

representation theory of compact groups. Our analysis in this subsection will be based on that in Fulton and Harris [14] and Chirikjian and Kyatkin [7].

Convolution can be defined on general compact groups using equation 16. To analyze this in the frequency domain, we need a generalization of the Fourier transform. It is possible to define

$$f_l = \int_G f(g) D^l(g)\, dg. \tag{45}$$

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

To obtain some intuition, consider the flatland case where the group corresponds to rotations in 2D, i.e. $G = SO(2)$. The elements g are then simply the angles ϕ, and the representation matrices are all 1-dimensional and correspond to the standard Fourier series, i.e. $D^l = e^{Il\phi}$. Thus, equation 45 corresponds directly to the standard Fourier series in 2D. Now, consider the case where the group is that of 3D rotations, i.e. $G = SO(3)$. In this case, D^l corresponds to the $2l + 1$-dimensional representation, and is a $(2l + 1) \times (2l + 1)$ representation matrix. The generalized Fourier transform is therefore *matrix-valued*. For a general compact group, we can generalize the notion of the Fourier transform to a matrix-valued function labeled by indices corresponding to the group representation. This reduces to the standard Fourier series for the 2D or flatland case, since the group representations are all one-dimensional and correspond directly to complex exponentials.

Once we have the generalization of the Fourier transform, one can derive [14] a convolution formula corresponding to equation 16,

$$(f \otimes g)_l = f_l \times g_l. \tag{46}$$

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

However, it is important to note that the 3D case discussed in the previous subsection does not correspond exactly either to equation 16 or 46. Specifically, all operations are not carried out in the rotation group $SO(3)$. Instead, we have rotations operating on unit vectors. Thus, it is not possible to apply equation 46 directly, and a separate convolution formula must be derived, as we have done in the previous subsection.

Note that we use the associated basis functions, i.e. spherical harmonics, and not the group representations directly, as basis functions for the lighting and BRDF, since these quantities are functions of directions or unit vectors, and not rotations. For the reflected light field, which is a

function of the rotation applied as well as the outgoing direction (a unit vector), we use mixed basis functions that are a product of group representations of $SO(3)$ and the spherical harmonics. The convolution formula we derive in equation 43 is actually simpler than equation 46, since it does not require matrix multiplication.

In the remainder of this section, we will derive a number of alternative forms for equation 43 that may be more suitable for particular cases. Then, in the next section, we will discuss the implications of equations 43 and 44 for forward and inverse problems in computer graphics.

4.4 Alternative Forms

For the analysis of certain problems, it will be more convenient to rewrite equation 43 in a number of different ways. We have already seen one example, of considering the outgoing angle fixed, as shown in equation 44. In this subsection, we consider a few more alternative forms.

4.4.1 Isotropic BRDFs

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

To consider the simplifications that result from isotropy, we first analyze the BRDF coefficients $\hat{\rho}_{ln,pq}$. In the BRDF expansion of equation 38, only terms that satisfy isotropy, i.e. are invariant with respect to adding an angle $\triangle\phi$ to both incident and outgoing azimuthal angles, are nonzero. From the form of the spherical harmonics, this requires that $n = q$. Furthermore, since we are considering BRDFs that depend only on $\mid \phi'_o - \phi'_i \mid$, we should be able to negate both incident and outgoing azimuthal angles without changing the result. This leads to the condition that $\hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q)p(-q)}$. Finally, we define a 3-index BRDF coefficient by

$$\hat{\rho}_{lpq} = \hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q)p(-q)}. \tag{47}$$

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

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

$$R_{\alpha,\beta} = R_{\alpha,\beta,0} = R_z(\beta)R_y(\alpha). \tag{48}$$

Next, the representation matrices can be rewritten as

$$D^l_{mn}(\alpha,\beta) = D^l_{mn}(\alpha,\beta,0) = d^l_{mn}(\alpha)e^{Im\beta}. \tag{49}$$

It should be noted that removing the dependence on γ weakens the orthogonality condition (equation 41) on the representation matrices, since we no longer integrate over γ. The new orthonormality relation for these matrices is given by

$$\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi}\left(D^l_{mn}(\alpha,\beta)\right)^*\left(D^{l'}_{m'n}(\alpha,\beta)\right)\sin\alpha\,d\alpha\,d\beta = \frac{4\pi}{2l+1}\delta^{ll'}\delta_{mm'}. \tag{50}$$

In particular, the orthogonality relation for the index n no longer holds, which is why we have used the index n for both D and D' instead of using n and n', as in equation 41. The absence of an integral over γ leads to a slight weakening of the orthogonality relation, as well as a somewhat different normalization than in equation 41. For the future discussion, it will be convenient to define normalization constants by

$$\Lambda_l = \sqrt{\frac{4\pi}{2l+1}}. \tag{51}$$

We now have the tools necessary to rewrite equation 43 for isotropic BRDFs. Since we will be using the equations for isotropic BRDFs extensively in the rest of this paper, it will be worthwhile to briefly review the representations of the various quantities, specialized to the isotropic case.

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

$$\begin{aligned} L(\theta_i,\phi_i) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l} L_{lm}Y_{lm}(\theta_i,\phi_i) \\ L(\theta_i,\phi_i) = L\left(R_{\alpha,\beta}(\theta'_i,\phi'_i)\right) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{+l}\sum_{m'=-l}^{l} L_{lm}D^l_{mm'}(\alpha,\beta)Y_{lm'}(\theta'_i,\phi'_i). \end{aligned} \tag{52}$$

Then, we write the expansion of the isotropic BRDF,

$$\hat{\rho}(\theta'_i,\theta'_o,\mid\phi'_o-\phi'_i\mid) = \sum_{l=0}^{\infty}\sum_{p=0}^{\infty}\sum_{q=-\min(l,p)}^{\min(l,p)}\hat{\rho}_{lpq}Y^*_{lq}(\theta'_i,\phi'_i)Y_{pq}(\theta'_o,\phi'_o). \tag{53}$$

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

$$\begin{aligned} C_{lmpq}(\alpha,\beta,\theta'_o,\phi'_o) &= \Lambda_l^{-1}D^l_{mq}(\alpha,\beta)Y_{pq}(\theta'_o,\phi'_o) \\ B(\alpha,\beta,\theta'_o,\phi'_o) &= \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-\min(l,p)}^{\min(l,p)} B_{lmpq}C_{lmpq}(\alpha,\beta,\theta'_o,\phi'_o) \\ B_{lmpq} &= \int_{\phi'_o=0}^{2\pi}\int_{\theta'_o=0}^{\pi}\int_{\beta=0}^{2\pi}\int_{\alpha=0}^{\pi} B(\alpha,\beta,\theta'_o,\phi'_o)C^*_{lmpq}(\alpha,\beta,\theta'_o,\phi'_o)\sin\alpha\sin\theta'_o\,d\alpha\,d\beta\,d\theta'_o\,d\phi'_o. \end{aligned} \tag{54}$$

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 50. Note that the index q in the definition of C_{lmpq} is the same (coupled) for both factors D^l_{mq} and Y_{pq}. This is a consequence of isotropy, and is not true in the anisotropic case. Therefore, although the representation matrices D^l no longer satisfy orthogonality over the index q (corresponding to the index n in equation 50), orthogonality over the index q follows from the orthonormality of the spherical harmonics Y_{pq}.

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

$$B_{lmpq} = \Lambda_l L_{lm} \hat{\rho}_{lpq} \tag{55}$$

Apart from a slightly different normalization, and the removal of γ and the corresponding index n, this is essentially the same as equation 43. We will be using this equation for isotropic BRDFs extensively in the second part of this paper, 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 44, by holding the outgoing elevation angle θ'_o fixed. Since the isotropic BRDF depends only on $\mid \phi'_o - \phi'_i \mid$, and not directly on ϕ'_o, we do not hold ϕ'_o fixed, as we did in equation 44. We first define the modified expansions,

$$\begin{aligned}
\hat{\rho}(\theta'_i, \theta'_o, \mid \phi'_o - \phi'_i \mid) &= \sum_{l=0}^{\infty} \sum_{q=-l}^{l} \hat{\rho}_{lq}(\theta'_o) \left(\frac{1}{\sqrt{2\pi}} Y^*_{lq}(\theta'_i, \phi'_i) \exp(Iq\phi'_o) \right) \\
B(\alpha, \beta, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{q=-l}^{l} B_{lmq}(\theta'_o) \left(\frac{1}{\sqrt{2\pi}} \Lambda_l^{-1} D^l_{mq}(\alpha, \beta) \exp(Iq\phi'_o) \right).
\end{aligned} \tag{56}$$

Then, we may write down the isotropic convolution formula corresponding to equation 44.

$$B_{lmq}(\theta'_o) = \Lambda_l L_{lm} \hat{\rho}_{lq}(\theta'_o) \tag{57}$$

4.4.2 Reciprocity Preserving

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

$$\begin{aligned}
\tilde{\rho} &= \hat{\rho} \cos\theta'_o = \rho \cos\theta'_i \cos\theta'_o \\
\tilde{B} &= B \cos\theta'_o
\end{aligned} \tag{58}$$

With these definitions, all of the derivations presented so far still hold. In particular, the convolution formulas in equation 43 and 55 hold with the replacements $B \rightarrow \tilde{B}, \hat{\rho} \rightarrow \tilde{\rho}$. For example, equation 55 for isotropic BRDFs becomes

$$\tilde{B}_{lmpq} = \Lambda_l L_{lm} \tilde{\rho}_{lpq} \tag{59}$$

The symmetry of the transfer function ensures that its coefficients are unchanged if the indices corresponding to incident and outgoing angles are interchanged, i.e. $\tilde{\rho}_{lpq} = \tilde{\rho}_{plq}$. In the more general anisotropic case, $\tilde{\rho}_{ln,pq} = \tilde{\rho}_{pq,ln}$. We will use the frequency-space reflection formula, as given by equation 59, whenever explicitly maintaining the reciprocity of the BRDF is important.

4.4.3 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 6 for an illustration. Further, the reflected light field can be represented simply by a 2D reflection map $B(\alpha, \beta)$ parameterized by $\vec{C} = (\alpha, \beta)$. In other words, after reparameterization, there is no explicit exitant (outgoing) angular dependence for either the BRDF or reflected light field.

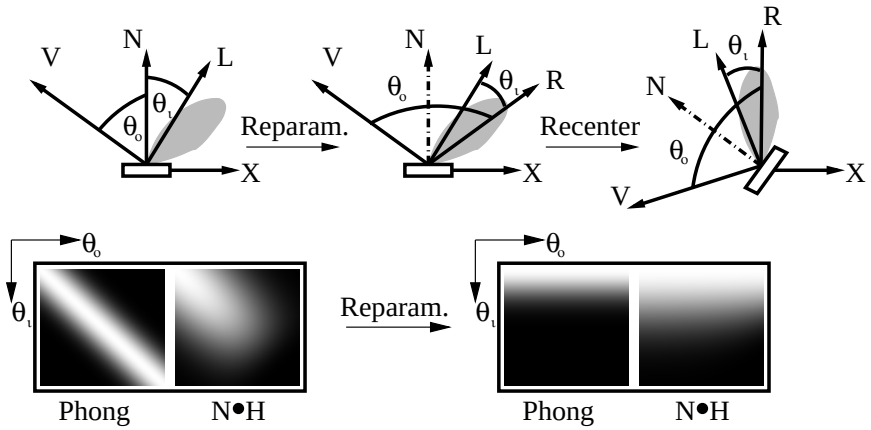

Figure 6: *Reparameterization involves recentering about the reflection vector. BRDFs become more compact, and in special cases (Phong) become 1D functions.*

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

$$\begin{aligned}
\hat{\rho}(\theta_i') &= \sum_{l=0}^{\infty} \hat{\rho}_l Y_{l0}(\theta_i') \\
\hat{\rho}_l &= 2\pi \int_0^{\pi/2} \hat{\rho}(\theta_i') Y_{l0}(\theta_i') \sin\theta_i' \, d\theta_i' \\
B(\alpha,\beta) &= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{q=-l}^{l} \hat{\rho}_l L_{lm} D^l_{mq}(\alpha,\beta) \int_0^{2\pi} \int_0^{\pi} Y_{lq}(\theta_i', \phi_i') Y_{l0}(\theta_i') \sin\theta_i' \, d\theta_i' d\phi_i' \\
&= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \hat{\rho}_l L_{lm} D^l_{m0}(\alpha,\beta). \qquad (60)
\end{aligned}$$

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 35, 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,

$$B(\alpha, \beta) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm} Y_{lm}(\alpha, \beta). \tag{61}$$

The required convolution formula now becomes

$$B_{lm} = \Lambda_l \hat{\rho}_l L_{lm}. \tag{62}$$

In the context of Lambertian BRDFs (for which no reparameterization is required), it has been noted by Basri and Jacobs [1] that equation 62 is mathematically an instance of the Funk-Hecke theorem (as stated, for instance in Groemer [16], page 98). However, that theorem does not generalize to the other relations previously encountered. With respect to equation 55, we have essentially just dropped the indices p and q corresponding to the outgoing angular dependence. It is important to remember that the reflected light field is now expanded in terms of spherical harmonics. B is simply a filtered version of L, with each frequency l being attenuated by a different amount, corresponding to the BRDF transfer function $\hat{\rho}_l$.

For general BRDFs, the radial symmetry property does not hold precisely, so they cannot be reduced exactly to 1D functions, nor can B be written simply as a 2D reflection map. Nevertheless, a reparameterization of the specular BRDF components by the reflection vector (or other central BRDF direction) still yields compact forms. To reparameterize, we simply recenter the BRDF (and the reflection integral) about the reflection vector $\vec{R}$, rather than the surface normal, as shown in figure 6. 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 55.

This section has presented a frequency-space analysis of the reflection equation. We have shown the simple quantitative form that results from this analysis, as embodied by equations 43 and 55. The mathematical analysis leading to these results is the main contribution of this paper, showing quantitatively that reflection can be viewed as a convolution. The next section gives an overview of the implications of these results for forward and inverse problems in rendering. The second part of the paper will work out a number of special cases of interest.

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. In this section, we make some general observations. In the second part of the paper, we will quantitatively analyze a number of special cases of interest.

We will deal here exclusively with the 3D case, since that is of greater practical importance. A preliminary analysis for the 2D case can be found in an earlier paper [34]. 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 Forward Rendering with Environment Maps

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

$$B(\alpha, \beta, \theta'_o, \phi'_o) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} L_{lm} \hat{\rho}_{lpq} \left(D^l_{mq}(\alpha, \beta) Y_{pq}(\theta'_o, \phi'_o) \right). \tag{63}$$

If either the lighting or the BRDF is low frequency, the total number of terms in the summation will be relatively small, and it may be possible to use equation 63 directly for shading a pixel. We have already demonstrated the practicality of this approach for Lambertian BRDFs [35], where we can set $p = q = 0$, and use $l \leq 2$, i.e. only 9 spherical harmonic terms.

In the general case, frequency space analysis allows for setting sampling rates accurately, and enables compact frequency domain representations. Further, just as image convolutions are often computed in the Fourier rather than the spatial domain, computing the reflected light field is more efficient in frequency space, using equation 63, rather than in angular space. We have already demonstrated the practical implementation of these ideas [38].

5.2 Well-posedness and conditioning of Inverse Lighting and BRDF

In this subsection, we briefly discuss how to apply ideas from the theoretical analysis to determine which inverse problems are well-posed, i.e. solvable, versus ill-posed, i.e. unsolvable, and also determine the numerical conditioning properties. At the end of this subsection, we will also relate these results to the general theory of linear integral equations. An important duality should be

noted here. Forward problems for which an efficient frequency domain solution is possible, such as those involving diffuse surfaces and/or soft lighting, have corresponding inverse problems that are ill-conditioned. Turned around, ill-conditioned inverse problems allow us to get a very good solution to the forward problem by using very coarse low-frequency approximations of the initial conditions. For instance, Lambertian surfaces act as low-pass filters, the precise form of which we will explore in the second part of this paper, blurring the illumination. Therefore, high-frequency components of the lighting are not essential to rendering images of diffuse objects, and we can make very coarse low-frequency approximations to the lighting without significantly affecting the final image. This leads to more efficient algorithms for computer graphics, and illustrates one of the benefits in considering a signal-processing view of reflection.

5.2.1 Inverse-BRDF

We first address the question of BRDF estimation. Our goal is to consider this problem under general illumination conditions, and understand when the BRDF can be recovered, i.e. BRDF estimation is well posed, and when the BRDF cannot be recovered, i.e. estimation is ill-posed. We would also like to know when BRDF recovery will be well-conditioned, i.e. numerically robust. An understanding of these issues is critical in designing BRDF estimation algorithms that work under arbitrary lighting. Otherwise, we may devise algorithms that attempt to estimate BRDF components that cannot be calculated, or whose estimation is ill-conditioned.

For isotropic surfaces, a simple manipulation of equation 55 yields

$$\hat{\rho}_{lpq} = \Lambda_l^{-1} \frac{B_{lmpq}}{L_{lm}}. \tag{64}$$

In general, BRDF estimation will be well-posed, i.e. unambiguous as long as the denominator on the right-hand side does not vanish. Of course, to be physically accurate, the numerator will also become 0 if the denominator vanishes, so the right-hand side will become indeterminate. From equation 64, we see that BRDF estimation is well posed as long as for all l, there exists at least one value of m so that $L_{lm} \neq 0$. In other words, all orders in the spherical harmonic expansion of the lighting should have at least one coefficient with nonzero amplitude. If any order l completely vanishes, the corresponding BRDF coefficients cannot be estimated.

In signal processing terms, if the input signal (lighting) has no amplitude along certain modes of the filter (BRDF), those modes cannot be estimated. BRDF recovery is well conditioned when the spherical harmonic expansion of the lighting does not decay rapidly with increasing frequency, i.e. when the lighting contains high frequencies like directional sources or sharp edges, and is ill-conditioned for soft lighting. Equation 64 gives a precise mathematical characterization of the conditions for BRDF estimation to be well-posed and well-conditioned. These results are similar to those obtained by D'Zmura [12] 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.

5.2.2 Inverse Lighting

A similar analysis can be done for estimation of the lighting. Manipulation of equation 55 yields

$$L_{lm} = \Lambda_l^{-1} \frac{B_{lmpq}}{\hat{\rho}_{lpq}}. \tag{65}$$

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

5.2.3 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

$$b(s) = \int_t K(s,t) f(t)\, dt, \tag{66}$$

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

The theory of linear integral equations, as for instance in Cochran [8], analyzes equation 66 based on the structure of the kernel. In particular, assume we may find a basis function expansion

of the form

$$\begin{aligned} b(s) &= \sum_{i=1}^{n} b_i u_i(s) \\ K(s,t) &= \sum_{i=1}^{n} K_i u_i(s) v_i^*(t) \\ f(t) &= \sum_{i=1}^{\infty} f_i v_i(t), \end{aligned} \tag{67}$$

where each of the sets of functions u_i and v_i (with v_i^* being the complex conjugate) is linearly independent. Here, n is the number of terms in, or rank of the kernel, K. If n is finite, the kernel is referred to as degenerate. It should be noted that if the function sets u and v were orthonormal, then we would have a result of the form $b_i = K_i f_i$. In effect, we have constructed an expansion of the form of equation 67 using orthonormal basis functions involving group representations and spherical harmonics, thereby deriving the convolution result.

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

5.3 Light Field Factorization

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

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

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

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

It would appear from equation 55 however, that there is an unrecoverable scale factor for each order l, corresponding to the known coefficient we require. In other words, we may multiply the lighting for each order l by some amount (which may be different for different frequencies l) while dividing the BRDF by the same amount. However, there is an important additional physical constraint. The BRDF must be reciprocal, i.e. symmetric with respect to incident and outgoing angles. The corresponding condition in the frequency domain is that the BRDF coefficients must be symmetric with respect to interchange of the indices corresponding to the incident and outgoing directions. To take advantage of this symmetry, we will use the reciprocal form of the frequency-space equations, as defined in equation 59.

We now derive an analytic formula for the lighting and BRDF in terms of coefficients of the reflected light field. Since we cannot recover the global scale, we will arbitrarily scale the DC term of the lighting so $L_{00} = \Lambda_0^{-1} = \sqrt{1/(4\pi)}$. Note that this scaling is valid unless the DC term is 0, corresponding to no light—an uninteresting case. Using equations 59, 64, and 65, we obtain

$$
\begin{aligned}
L_{00} &= \Lambda_0^{-1} && \text{: Global Scale} \\
\tilde{\rho}_{0p0} &= \tilde{B}_{00p0} && \text{: Equation 64 } (l = q = 0) \\
L_{lm} &= \Lambda_l^{-1} \frac{\tilde{B}_{lmpq}}{\tilde{\rho}_{lpq}} && \text{: Equation 65} \\
&= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{l00}} && \text{: Set } p = q = 0 \\
&= \frac{\tilde{B}_{lm00}}{\tilde{\rho}_{0l0}} && \text{: Reciprocity, } \tilde{\rho}_{0l0} = \tilde{\rho}_{l00} \\
&= \Lambda_l^{-1} \frac{\tilde{B}_{lm00}}{\tilde{B}_{00l0}} && \text{: Plug in from } 2^{\text{nd}} \text{ line} \\
\tilde{\rho}_{lpq} &= \Lambda_l^{-1} \frac{\tilde{B}_{lmpq}}{L_{lm}} && \text{: Equation 64} \\
&= \frac{\tilde{B}_{lmpq} \tilde{B}_{00l0}}{\tilde{B}_{lm00}} && \text{: Substitute from above for } L_{lm}. \qquad (68)
\end{aligned}
$$

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

Therefore, up to global scale, **the reflected light field can be factored into the lighting and the BRDF**, provided the appropriate coefficients of the reflected light field do not vanish, i.e. the denominators above are nonzero. If the denominators do vanish, the inverse-lighting or inverse-BRDF problems become ill-posed and consequently, the factorization becomes ill-posed. Note that the above relations are one possible factorization formula. We may still be able to factor the light field even if some of the $\tilde{\rho}_{l00}$ terms vanish in equation 68, by using different values of $\tilde{\rho}_{lpq}$ with $p \neq 0$.

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

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 also discussed the implications for inverse problems such as lighting recovery, BRDF recovery, light field factorization, and forward rendering problems such as environment map prefiltering and rendering. The second part of this paper will make these ideas concrete for many special cases, deriving analytic formulae for the frequency spectra of many common BRDF and lighting models.

The present paper 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 35. The first property follows from the addition theorem for spherical harmonics (see for instance, Jackson [19] equation 3.62),

$$Y_{l0}(u,v) = \Lambda_l \sum_{m=-l}^{l} Y^*_{lm}(\theta,\phi) Y_{lm}(\theta',\phi'). \tag{69}$$

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.

$$\cos u = \cos\theta\cos\theta' + \sin\theta\sin\theta'\cos(\phi - \phi'). \tag{70}$$

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

$$\begin{pmatrix} \sin u\cos v \\ \sin u\sin v \\ \cos u \end{pmatrix} = \begin{pmatrix} \cos\alpha & 0 & \sin\alpha \\ 0 & 1 & 0 \\ -\sin\alpha & 0 & \cos\alpha \end{pmatrix} \begin{pmatrix} \sin\theta'\cos\phi' \\ \sin\theta'\sin\phi' \\ \cos\theta' \end{pmatrix} = \begin{pmatrix} \cos\alpha\sin\theta'\cos\phi' + \sin\alpha\cos\theta' \\ \sin\theta'\sin\phi' \\ \cos\alpha\cos\theta' + \sin\alpha\sin\theta'\,(-\cos\phi') \end{pmatrix}. \tag{71}$$

Since $(-\cos\phi') = \cos(\pi - \phi')$, we know from equation 70 that u corresponds to the angle between (α, π) and (θ', ϕ'). In other words, we may set $(\theta, \phi) = (\alpha, \pi)$. To summarize,

$$Y_{l0}\left(R_\alpha(\theta', \phi')\right) = \Lambda_l \sum_{m=-l}^{l} Y^*_{lm}(\alpha, \pi) Y_{lm}(\theta', \phi'). \tag{72}$$

To proceed further, we write the rotation formula for spherical harmonics, omitting the z rotation by β, since that has no significance for azimuthally symmetric harmonics.

$$Y_{l0}\left(R_\alpha(\theta', \phi')\right) = \sum_{m=-l}^{l} d^l_{0m}(\alpha) Y_{lm}(\theta', \phi') \tag{73}$$

A comparision of equations 72 and 73 yields the first property of representation matrices in equation 35, i.e.

$$d^l_{0m}(\alpha) = \Lambda_l Y^*_{lm}(\alpha, \pi). \tag{74}$$

To obtain the second property in equation 35, we use the form of the spherical harmonic expansion when the elevation angle is 0, i.e. we are at the north pole. Specifically, we note that $Y_{lm'}(0', \phi') = \Lambda_l^{-1}\delta_{m'0}$. With this in mind, the derivation is as follows,

$$\begin{aligned} Y_{lm}(\alpha, \beta) &= Y_{lm}\left(R_{\alpha,\beta,\gamma}(0', \phi')\right) \\ &= \sum_{m'=-l}^{l} D^l_{mm'}(\alpha, \beta, \gamma) Y_{lm'}(0', \phi') \\ &= \Lambda_l^{-1} D^l_{m0}(\alpha, \beta, \gamma). \end{aligned} \tag{75}$$

This brings us to the second property stated in equation 35,

$$D^l_{m0}(\alpha, \beta, \gamma) = \Lambda_l Y_{lm}(\alpha, \beta). \tag{76}$$

References

[1] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. In *ICCV 01*, pages 383–390, 2001.

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

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

[4] J. F. Blinn and M. E. Newell. Texture and reflection in computer generated images. *Communications of the ACM*, 19:542–546, 1976.

[5] B. Cabral, N. Max, and R. Springmeyer. Bidirectional reflection functions from surface bump maps. In *SIGGRAPH 87*, pages 273–281, 1987.

[6] B. Cabral, M. Olano, and P. Nemec. Reflection space image based rendering. In *SIGGRAPH 99*, pages 165–170, 1999.

[7] G. Chirikjian and A. Kyatkin. *Engineering applications of noncommutative harmonic analysis: with emphasis on rotation and motion groups*. CRC press, 2000.

[8] J. Cochran. *The analysis of linear integral equations*. McGraw-Hill, 1972.

[9] K. Dana, B. Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.

[10] P. Debevec, T. Hawkins, C. Tchou, H.P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *SIGGRAPH 00*, pages 145–156, 2000.

[11] R. Dror, E. Adelson, and A. Willsky. Estimating surface reflectance properties from images under unknown illumination. In *SPIE Photonics West: Human Vision and Electronic Imaging VI*, pages 231–242, 2001.

[12] M. D'Zmura. *Computational Models of Visual Processing*, chapter Shading Ambiguity: Reflectance and Illumination, pages 187–207. MIT Press, 1991.

[13] R. Epstein, P.W. Hallinan, and A. Yuille. 5 plus or minus 2 eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *IEEE Workshop on Physics-Based Modeling in Computer Vision*, pages 108–116, 1995.

[14] W. Fulton and J. Harris. *Representation Theory: A first course*. Springer-Verlag, 1991.

[15] N. Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics & Applications*, 6(11):21–29, 1986.

[16] H. Groemer. *Geometric applications of Fourier series and spherical harmonics*. Cambridge University Press, 1996.

[17] P. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *CVPR 94*, pages 995–999, 1994.

[18] T. Inui, Y. Tanabe, and Y. Onodera. *Group theory and its applications in physics*. Springer Verlag, 1990.

[19] J.D. Jackson. *Classical Electrodynamics*. John Wiley, 1975.

[20] J. Kajiya. The rendering equation. In *SIGGRAPH 86*, pages 143–150, 1986.

[21] J. J. Koenderink and A. J. van Doorn. Phenomenological description of bidirectional surface reflection. *JOSA A*, 15(11):2903–2912, 1998.

[22] R. Love. *Surface Reflection Model Estimation from Naturally Illuminated Image Sequences*. PhD thesis, Leeds, 1997.

[23] R. Lu, J. Koenderink, and A. Kappers. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics*, 37(25):5974–5984, 1998.

[24] T.M. MacRobert. *Spherical harmonics; an elementary treatise on harmonic functions, with applications*. Dover Publications, 1948.

[25] S.R. Marschner and D.P. Greenberg. Inverse lighting for photography. In *Fifth Color Imaging Conference*, pages 262–265, 1997.

[26] S.R. Marschner, S.H. Westin, E.P.F. Lafortune, and K.E. Torrance. Image-Based BRDF measurement. *Applied Optics*, 39(16):2592–2600, 2000.

[27] N. McCormick. Inverse radiative transfer problems: a review. *Nuclear Science and Engineering*, 112:185–198, 1992.

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

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

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

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

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

[33] R.W. Preisendorfer. *Hydrologic Optics*. US Dept Commerce, 1976.

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

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

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

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

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

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

[40] Y. Sato and K. Ikeuchi. Reflectance analysis under solar illumination. Technical Report CMU-CS-94-221, CMU, 1994.

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

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

[43] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967.

[44] G. J. Ward. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92*, pages 265–272, 1992.

[45] Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92*, 1992.

[46] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *SIGGRAPH 99*, pages 215–224, 1999.

[47] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *SIGGRAPH 98*, pages 207–218, 1998.