

Analysis of Planar Light Fields From Homogeneous Convex Curved Surfaces Under Distant Illumination

Ravi Ramamoorthi and Pat Hanrahan
Stanford University

ABSTRACT

We consider the *flatland* or 2D properties of the light field generated when a homogeneous convex curved surface reflects a distant illumination field. Besides being of considerable theoretical interest, this problem has applications in computer vision and graphics—for instance, in determining lighting and bidirectional reflectance distribution functions (BRDFs), in rendering environment maps, and in image-based rendering. We demonstrate that the integral for the reflected light transforms to a simple product of coefficients in Fourier space. Thus, the operation of rendering can be viewed in simple signal processing terms as a filtering operation that convolves the incident illumination with the BRDF. This analysis leads to a number of interesting observations for computer graphics, computer vision, and visual perception.

Keywords: Light Field, BRDF, Fourier analysis, Convolution, Radiance, Irradiance, Inverse rendering, Environment Maps

1. INTRODUCTION

This paper considers the theoretical properties of the light field generated when a homogeneous convex curved surface reflects a distant illumination field. Besides being of considerable theoretical interest, this problem has applications in computer vision and graphics—for instance, in determining lighting and bidirectional reflectance distribution functions (BRDFs) from photographs, in rendering environment maps, and in image-based rendering. The utility of curved surfaces and a distant light source in an image-based BRDF measurement system has been demonstrated recently for the planar case by Lu et al.¹⁰ in measuring the plane-of-incidence BRDF of velvet, and by Marschner et al.¹² for image-based measurement of general isotropic BRDFs.

We seek to put future work in inverse rendering on a sound mathematical foundation by analyzing the structure of the light field, thereby giving insight into the well-posedness and conditioning properties of many inverse problems. The work is also likely to have application in traditional computer graphics, since an ill-posed inverse problem is likely to have a corresponding forward problem where accurate initial conditions are relatively unimportant, and may therefore be simplified to allow for more efficient algorithms.

In this paper, we restrict ourselves to the *flatland* case, assuming all illumination and measurements occur in the same plane. We have chosen to report on the flatland case because the mathematics is much simpler, while capturing most of the important insights. A subsequent paper will generalize these results to 3D.

We analyze the reflected light field generated by a convex curved surface. For our analysis, we construct the Fourier transform of both the input lighting and the BRDF. We demonstrate that the integral for the reflected light transforms to a simple product of Fourier coefficients. Thus, the operation of rendering can be viewed in simple signal processing terms as *convolving* the lighting with the BRDF. This Fourier analysis of the structure of the light field leads to a number of interesting observations. In particular, inverse problems such as recovering the lighting or BRDF can be seen as problems in *deconvolution*, and the conditions for this deconvolution to be well-posed and well-conditioned can be derived.

For instance, in general, the BRDF can be completely determined in theory, given the input lighting and the complete output radiance function. This can be viewed as a problem of filter estimation, since we are estimating a filter (BRDF) given the output (radiance) and the input (lighting). The problem is unsolvable when terms in the Fourier expansion of the lighting are zero, corresponding to the signal having no amplitude for certain modes of the BRDF filter. Furthermore, this explains the use of a point light source for image-based BRDF measurement—the Fourier transform of a delta function is constant and nonzero everywhere. We can also think of this as determining the

BRDF filter by considering its impulse-response. This theoretical result may have practical applications in efficiently recovering BRDFs under uncontrolled illumination conditions.

A similar result holds for the lighting, viz. we can recover the lighting given the BRDF and the output radiance. Similarly as above, this becomes impossible when terms in the Fourier transform of the BRDF are zero. The theory allows us to easily derive a closed-form formula for the action of distant illumination on a Lambertian object, an important special case. We will show that inverse lighting from a Lambertian object is in general ill-conditioned, allowing only the lowest-order modes of the lighting to be recovered. On the other hand, this result may allow us to efficiently prefilter and render environment maps in computer graphics, since only a very low frequency representation of the environment need be used.

Our work may also have applications in visual perception. We demonstrate that lighting effects from a distant illumination field—without considering shadowing and interreflection—can usually only induce low-frequency variation in the intensity of a Lambertian surface. Therefore, all high-frequency variation with respect to surface orientation is because of texture. This explains why we can perceive—high-frequency or rapidly varying—texture on surfaces independently of lighting effects, but find it difficult to distinguish lighting effects from a slowly varying texture. We also demonstrate that if reciprocity is taken into account, we can, in general, factor the light field into a product of BRDF and lighting terms, without knowing either. Besides obvious applications in computer vision and graphics, this accords with our perceptual ability to ascertain the BRDF (or degree of shininess) of a surface, independently of the lighting conditions.

The rest of this paper is organized as follows. In section 2, we briefly discuss some previous work. In section 3, we introduce the mathematical and physical preliminaries. In section 4, we obtain the Fourier space equation for the light field. Section 5 discusses the implications of these results with several examples. Section 6 briefly discusses applications to visual perception and extensions to 3D. We conclude the paper in section 7.

2. PREVIOUS WORK

The light field³ or plenoptic function¹ is a fundamental quantity in light transport and therefore has wide applicability for both forward and inverse problems in a number of fields. Inverse problems in radiative transfer and transport theory have been studied in many areas such as hydrologic optics¹⁶ and neutron scattering. See McCormick¹³ for a review. Until recently, computer vision has focused on shape recovery assuming a simple model for the lighting and BRDF. However, recently there has been work on so-called *inverse rendering* problems, notably by Marschner^{11,12} where lighting and optical properties are estimated and used to render novel images. Light fields have also been used directly for rendering images from photographs in computer graphics, without considering the underlying geometry,^{4,9} or by parameterizing the light field on the object surface.²¹

There has been some previous work on the specific research problems we discuss in this paper as possible applications of our methods—inverse lighting, BRDF estimation, light field factorization, and environment map prefiltering and rendering. Marschner and Greenberg¹¹ have considered the inverse lighting problem, assuming Lambertian surfaces. BRDF¹⁵ measurement has a fairly long history,^{6,20} but image-based methods utilizing surface curvature to rapidly recover the BRDF have only recently been proposed.^{10,12} While an explicit general factorization into lighting and BRDF has not been undertaken, various constrained partial factorizations or separations have been performed. Klinker et al.⁸ and Sato et al.¹⁸ have used color-space methods to separate diffuse and specular components. Sato et al.¹⁷ use shadows to recover the lighting in a scene, and also estimate the reflectance parameters of a planar surface. Efficient methods to prefilter and render environment maps are discussed by Miller and Hoffman,¹⁴ Greene,⁵ and more recently by Cabral et al.² and Kautz et al.⁷ Our paper attempts to offer insights into all of these areas, helping to intuitively explain successes and failures of earlier methods, and to suggest the way towards more efficient and general algorithms.

3. PRELIMINARIES

For the purposes of this paper, we will deal exclusively with the 2-dimensional or flatland case, wherein all surfaces (effectively, curves) and light sources are constrained to lie in a single plane. We will assume the illuminating light sources are distant, so the illumination does not depend upon surface position, but only on surface orientation. Also, surfaces will be assumed to be homogeneous, i.e. there is a single bi-directional reflectance distribution function (BRDF) for the surface. Since we are considering flatland, the BRDF will be two dimensional with arguments being

B	Outgoing radiance
$B_{p,q}$	Fourier coefficients of outgoing radiance
\tilde{B}	Outgoing radiance multiplied by cosine of exitant angle
$\tilde{B}_{p,q}$	Fourier coefficients of \tilde{B}
L	Incoming radiance
L_p	Fourier coefficients of incoming radiance
ρ	The BRDF
$\hat{\rho}$	BRDF multiplied by cosine of incident angle
$\hat{\rho}_{p,q}$	Fourier coefficients of $\hat{\rho}$
$\tilde{\rho}$	BRDF multiplied by cosines of incident and outgoing angles
$\tilde{\rho}_{p,q}$	Fourier coefficients of $\tilde{\rho}$
θ_i	Incident angle in <i>global</i> coordinates
θ'_i	Incident angle in <i>local</i> coordinates
θ_o	Outgoing angle in <i>global</i> coordinates
θ'_o	Outgoing angle in <i>local</i> coordinates
\mathbf{x}	Surface position
α	Surface normal parameterization by angle
I	$\sqrt{-1}$

Figure 1. Notation used in the paper




Figure 2. Diagram showing the geometry, with local and global coordinate angles marked. Note that angles are signed; in particular, θ'_i is negative. The fact that α is nearly a right angle has no significance and is only for convenience in illustration.

a single incident and reflected angle. We will further assume that the surfaces under consideration are convex, so that interreflection and shadowing can be ignored. This allows us to parameterize the surface by the surface orientation only, as specified by the surface normal (α). Notation used in the paper is listed in figure 1. A diagram of the geometry of the situation is shown in figure 2.

Local and Global coordinates, Rotations 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 angle with respect to the local reference frame (i.e. surface normal). The two are related by

$$\theta' = \theta - \alpha \quad (1)$$

where θ' is the (signed) local angle, related to the (signed) global angle θ by subtracting out the orientation of the surface (α). This subtraction is merely the *rotation* operator in the plane; we go from local to global coordinates by a rotation corresponding to the surface normal.

The Reflection Equation In local coordinates, we can relate the outgoing radiance to the incoming radiance by

$$B(\mathbf{x}, \theta'_o) = \int_{-\pi/2}^{\pi/2} L(\mathbf{x}, \theta'_i) \rho(\theta'_i, \theta'_o) \cos(\theta'_i) d\theta'_i \quad (2)$$

The limits of integration correspond to the *visible half-circle*—the 2D analogue of the upper hemisphere in 3D. B is the output light field caused by the input light field L and BRDF ρ .

We will often want to consider the cosine term in the integral together with the BRDF, and we can define a new transfer function for this purpose.

$$\hat{\rho}(\theta'_i, \theta'_o) = \rho(\theta'_i, \theta'_o) \cos(\theta'_i) \quad (3)$$

For further analysis, we will want to treat the incoming illumination L in global coordinates, but we would like to keep the BRDF in local coordinates, and it will be simpler to also keep the outgoing radiance B in local coordinates. By combining equations 1 and 2, noting that $L(\mathbf{x}, \theta'_i) = L(\theta_i) = L(\alpha + \theta'_i)$, and by incorporating the cosine into the transfer function using equation 3, we obtain

$$B(\alpha, \theta'_o) = \int_{-\pi/2}^{\pi/2} L(\alpha + \theta'_i) \hat{\rho}(\theta'_i, \theta'_o) d\theta'_i \quad (4)$$

Physical BRDFs must be reciprocal i.e. symmetric with respect to incoming and outgoing directions. In certain cases, it will be desirable to preserve the reciprocity of the BRDF with respect to the incident and outgoing angle in the newly defined transfer function. We may multiply both the outgoing radiance and the transfer function by $\cos(\theta'_o)$, to obtain

$$\begin{aligned} \tilde{B}(\alpha, \theta'_o) &= B(\alpha, \theta'_o) \cos(\theta'_o) \\ \tilde{\rho}(\theta'_i, \theta'_o) &= \rho(\theta'_i, \theta'_o) \cos(\theta'_i) \cos(\theta'_o) \\ \tilde{B}(\alpha, \theta'_o) &= \int_{-\pi/2}^{\pi/2} L(\alpha + \theta'_i) \tilde{\rho}(\theta'_i, \theta'_o) d\theta'_i \end{aligned} \quad (5)$$

where the newly defined function $\tilde{\rho}$ preserves reciprocity i.e. $\tilde{\rho}(\theta'_i, \theta'_o) = \tilde{\rho}(\theta'_o, \theta'_i)$.

Equations 4 and 5 are simple equations for direct illumination. The outgoing radiance is an integral of the incident illumination with the BRDF (which incorporates the cosine term). The only interesting feature is the rotation of the global incoming illumination field corresponding to the surface normal α , in order to align it with local coordinates. For a given θ'_o , each observation with a different value of α corresponds to a different orientation of the surface normal, but can also be viewed as corresponding to a different rotation of the incident light field. Although we are dealing with rotations, and not translations, the equation above shows that the reflected light field is a *convolution* of the incident illumination and the transfer function. This will be crucial for the Fourier series development in the next section.

4. FOURIER ANALYSIS

The goal of this section is to carry out a Fourier analysis of equation 4. For this purpose, we will form the Fourier series of a function.

$$\begin{aligned} f(\theta) &= \sum_k f_k e^{Ik\theta} \\ f_k &= \frac{1}{2\pi} \int_{-\pi}^{\pi} f(\theta) e^{-Ik\theta} d\theta \end{aligned} \quad (6)$$

Decomposition into Fourier Series: For the term L from equation 4, the Fourier series coefficients become

$$\begin{aligned} L(\theta_i) &= \sum_m L_m e^{Im\theta_i} \\ L(\alpha + \theta'_i) &= \sum_m L_m e^{Im(\alpha + \theta'_i)} \\ &= \sum_m L_m e^{Im\alpha} e^{Im\theta'_i} \end{aligned} \quad (7)$$

This result indicates that the effect of rotating the lighting to align it with the local coordinate system is simply to multiply the Fourier frequency coefficients by $e^{Im\alpha}$.

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

$$\begin{aligned} B(\alpha, \theta'_o) &= \sum_p B_p(\theta'_o) e^{Ip\alpha} \\ \hat{\rho}(\theta'_i, \theta'_o) &= \sum_n \hat{\rho}_n(\theta'_o) e^{In\theta'_i} \end{aligned} \quad (8)$$

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 of $[-\frac{\pi}{2}, \frac{\pi}{2}]$.

We are now ready to write equation 4 in terms of frequency coefficients.

$$\sum_p B_p(\theta'_o) e^{Ip\alpha} = \sum_m \sum_n L_m \hat{\rho}_n(\theta'_o) e^{Im\alpha} \int_{-\pi}^{\pi} e^{Im\theta'_i} e^{In\theta'_i} d\theta'_i \quad (9)$$

This can be simplified using orthogonality of the exponentials with respect to θ'_i in the integral. Setting $n = -m$ (other terms vanish by orthogonality), we obtain

$$\sum_p B_p(\theta'_o) e^{Ip\alpha} = 2\pi \sum_m L_m \hat{\rho}_{-m}(\theta'_o) e^{Im\alpha} \quad (10)$$

Now, it is a simple matter to equate coefficients for α . This sets $m = p$, and we obtain

$$B_p(\theta'_o) = 2\pi L_p \hat{\rho}_{-p}(\theta'_o) \quad (11)$$

It is also possible to perform another Fourier series expansion of the θ'_o dependence. Again, since the Fourier series is defined over $[-\pi, \pi]$, the function values are treated as 0 outside the range for θ'_o of $[-\frac{\pi}{2}, \frac{\pi}{2}]$. Equation 8 becomes

$$\begin{aligned} B(\alpha, \theta'_o) &= \sum_p \sum_q B_{p,q} e^{Ip\alpha} e^{Iq\theta'_o} \\ \hat{\rho}(\theta'_i, \theta'_o) &= \sum_n \sum_s \hat{\rho}_{n,s} e^{In\theta'_i} e^{Is\theta'_o} \end{aligned} \quad (12)$$

Alternatively,

$$\begin{aligned} B_p(\theta'_o) &= \sum_q B_{p,q} e^{Iq\theta'_o} \\ \hat{\rho}_n(\theta'_o) &= \sum_s \hat{\rho}_{n,s} e^{Is\theta'_o} \end{aligned} \quad (13)$$

It is now possible to rewrite equation 11 by using linear independence of exponentials to set $n = -p$ and $s = q$ in order to equate similar terms with respect to θ'_o . We thus obtain our fundamental equation:

$$B_{p,q} = 2\pi L_p \hat{\rho}_{-p,q} \quad (14)$$

Equation 14 is remarkable in that it states that the standard direct illumination integral in equation 4 can be viewed instead as a simple product in Fourier space. This is not really surprising, considering that equation 4 shows that the reflected light field is a *convolution* of the incident illumination and the transfer function, with different observations B for given θ'_o corresponding to different rotations of the incident light field. Since equation 14 is in terms of Fourier coefficients, it can be viewed in signal processing terms as a filtering operation; the output light field can be obtained by filtering the input lighting using the BRDF. As is made explicit in equation 11, the filtering operation in equation 14 takes place only over the index p corresponding to the output radiance parameter α , and the BRDF parameter θ'_i . The two are joined together by the lighting parameter $\theta_i = \alpha + \theta'_i$. Note the symmetry between the BRDF and the lighting. In signal processing terms, we can view either as the input, with the other being viewed as the filter. We believe these novel viewpoints can lead to many new insights and algorithms for problems in computer vision and graphics. The remaining sections of this paper are devoted to exploring the implications of this result.

As presented so far, the transfer function does not preserve symmetry with respect to the incident and outgoing angles. If we wish to preserve this property of the BRDF, we may make the transformations in equation 5 and rewrite the above equation:

$$\tilde{B}_{p,q} = 2\pi L_p \tilde{\rho}_{-p,q} \quad (15)$$

The condition $\tilde{\rho}(\theta'_i, \theta'_o) = \tilde{\rho}(\theta'_o, \theta'_i)$ expressing the symmetry between incident and outgoing angles can be expressed as a corresponding symmetry of the Fourier coefficients $\tilde{\rho}_{p,q} = \tilde{\rho}_{q,p}$. This follows from inspection of equation 12, and is also intuitively understood as symmetry between indices for expansion in terms of incoming and outgoing angles.

5. APPLICATIONS

This section explores the implications of equation 14 with reference to several applications in computer vision, graphics and visual perception. We start by considering inverse problems, i.e recovering the lighting and/or the BRDF knowing the output light field. We have already seen that the reflected light field can be viewed as being generated by convolving the incident light field with the BRDF transfer function. Therefore, the inverse problems can be viewed as problems involving *deconvolution*. Besides giving explicit formulas or algorithms, our goals will be to understand which problems are well-posed versus ill-posed and which problems are well-conditioned versus ill-conditioned. This understanding can likely be carried forward as a guide to practical algorithms for real 3D surfaces.

5.1. Inverse Lighting

Our goal here is to recover the lighting, given the BRDF and the output light field. From equation 14, it is easy to read off the lighting coefficients as

$$L_p = \frac{1}{2\pi} \frac{B_{p,q}}{\hat{\rho}_{-p,q}} \quad (16)$$

In theory, this problem is well-posed unless the denominator vanishes* for each q i.e. $\forall q : \hat{\rho}_{-p,q} = 0$. In practice, however, the problem is ill-conditioned if $\hat{\rho}_{-p,q}$ is small. As we will see, the Fourier spectrum of most BRDFs will have declining higher-order terms. This indicates that numerically, the higher frequencies of the lighting will be difficult to recover.

In signal processing terms, this is a deconvolution problem of estimating the input, knowing the filter (the BRDF) and the output. When the BRDF filter truncates certain frequencies in the input lighting signal to 0 (for instance, if it were a low-pass filter), we cannot determine the corresponding frequencies of the lighting from the output signal.

*Note that by physical considerations, if the denominator of equation 16 vanishes, the numerator must also vanish, leaving the right-hand side indeterminate, but not infinite. Since we can use any value of q to find L_p , the problem is ill-posed only if the denominator is 0 for all q .

Mirror BRDF: An important special case to consider is that of a perfectly reflective object or mirror. In this case, the BRDF involves a delta function[†] of the sum of the (signed) incident and reflected angles. The Fourier series is obtained simply by integrating using the definition in equation 6, and reflects the fact that the Fourier transform of a delta function is a constant.

$$\begin{aligned}\hat{\rho}(\theta'_i, \theta'_o) &= \delta(\theta'_i + \theta'_o) \\ \hat{\rho}_{p,q} &= \frac{\delta_{p,q}}{2\pi}\end{aligned}\tag{17}$$

Plugging this into equation 16, we obtain

$$\begin{aligned}L_p &= \frac{1}{2\pi} \frac{B_{p,q}}{\frac{\delta_{-p,q}}{2\pi}} \\ L_p &= B_{p,-p}\end{aligned}\tag{18}$$

This is just the Fourier series form of the identity for a mirror BRDF:

$$L(\alpha - \theta'_o) = B(\alpha, \theta'_o)\tag{19}$$

Equation 18 indicates that we can obtain the lighting terms simply by reading off coefficients of the output light field. Therefore, this is a well posed and well conditioned problem. These results explain in Fourier space why a mirrored sphere is so often used in graphics to obtain an estimate of the lighting in a scene. In fact, as expected, the mirrored sphere is the ideal BRDF for lighting estimation since its frequency spectrum remains constant.

Lambertian BRDF: Another important case is when the BRDF is Lambertian. Then, the transfer function for a Lambertian object with reflectance 1 is given by $\hat{\rho}(\theta'_i, \theta'_o) = \frac{1}{2} \cos(\theta'_i)$. The factor of 1/2 comes from energy conservation. Using the alternative form of equation 11 instead of equation 14 and dropping the dependence on θ'_o since the surface is Lambertian, we obtain the analogue to equations 11 and 16 for Lambertian surfaces.

$$\begin{aligned}B_p &= 2\pi L_p \hat{\rho}_{-p} \\ L_p &= \frac{1}{2\pi} \frac{B_p}{\hat{\rho}_{-p}}\end{aligned}\tag{20}$$

It remains to determine the form of the Fourier series for the Lambertian transfer function $\hat{\rho} = \frac{1}{2} \cos(\theta'_i)$. Note that since the BRDF is nonzero only over the visible halfcircle, the range of integration is $[-\pi/2, \pi/2]$, instead of the interval $[-\pi, \pi]$ over which the Fourier series is defined.

$$\hat{\rho}_{-p} = \frac{1}{2\pi} \int_{-\pi/2}^{\pi/2} \frac{1}{2} \cos(\theta'_i) e^{Ip\theta'_i} d\theta'_i\tag{21}$$

This evaluates to 1/8 when $p = \pm 1$ and is 0 for all other odd values of p . When p is even, the value is given by

$$\hat{\rho}_{-2p} = \frac{(-1)^{p+1}}{2\pi(4p^2 - 1)}\tag{22}$$

Note that the magnitude of the Fourier coefficient falls off as $1/p^2$.

There are several noteworthy points here. Firstly, for p odd and not equal to ± 1 , B_p and $\hat{\rho}_{-p}$ are both 0, and the value of L_p cannot be determined. Thus, there is a fundamental theoretical obstacle to the inverse lighting problem from a Lambertian surface.

As an example, consider perturbing some existing incoming lighting distribution by a term proportional to $\cos(3\theta_i)$. This term must be a perturbation, and not the entire lighting, since it includes negative regions. Our goal

[†] Actually, the BRDF is 0 when $|\theta'_i| > \pi/2$. For a mirror only, the definition below does not change the output light field where defined—when $|\theta'_o| \leq \pi/2$. For other BRDFs, we must explicitly zero them when $|\theta'_i| > \pi/2$.

is to show that this perturbation of the initial lighting does not change the output light field. We can directly use equation 4 to compute the perturbation in the light field:

$$\begin{aligned}
 \Delta B(\alpha, \theta'_o) &= \int_{-\pi/2}^{\pi/2} (\cos[3(\alpha + \theta'_i)] \cos(\theta'_i)) d\theta'_i \\
 &= \int_{-\pi/2}^{\pi/2} (\cos(3\alpha) \cos(3\theta'_i) \cos(\theta'_i) - \sin(3\alpha) \sin(3\theta'_i) \cos(\theta'_i)) d\theta'_i \\
 &= \frac{1}{2} \int_{-\pi/2}^{\pi/2} (\cos(3\alpha)[\cos(4\theta'_i) + \cos(2\theta'_i)] - \sin(3\alpha)[\sin(4\theta'_i) - \sin(2\theta'_i)]) d\theta'_i \\
 &= 0
 \end{aligned} \tag{23}$$

A second interesting point about the frequency spectrum of the Lambertian BRDF in equation 22 is the rapid falloff with increasing $|p|$. This means that the inverse-lighting problem is ill-conditioned even for coefficients for which it is well-posed. In practice, we cannot hope to recover more than the lowest order terms of the lighting. This observation helps explain the results of Marschner and Greenberg.¹¹ In that paper, an attempt was made to solve the inverse lighting problem, treating the surfaces as Lambertian. The authors noted that the problem appeared ill-conditioned, and not amenable to accurate solution. Therefore, they had to rely heavily on a regularizing term that preserved the smoothness of the solution. The preceding arguments show why the problem is ill-conditioned, and suggest a different regularization scheme. We can assume the high-frequency lighting coefficients to be inaccurate, so we do not attempt to recover them, and merely set them to 0. This indicates that a Fourier basis is ideal for recovering the lighting.




Figure 3. Plot of magnitude of Fourier coefficient versus frequency p for a phong BRDF model when $\theta'_o = 0$. The solid line is for a Lambertian object with $\beta = 0$. The coefficients rapidly diminish and there are zeros for odd coefficients greater than 1. The dashdot line below it is for $\beta = 8$ which has a more gradual decay of coefficients, but we would still be able to recover only the first few lighting terms. The lowest dashed line is for $\beta = 128$ which shows a much more gradual decay of coefficients, making inverse lighting better conditioned. In general, we can only recover lighting coefficients up to order approximately $\sqrt{\beta}$.

Phong BRDF: The Phong BRDF has a specular lobe of the form $(R \cdot L)^\beta$ where β is the Phong exponent, L is the light vector and R is the reflection of the viewing vector about the surface normal. Two extremes of the Phong

BRDF are the Lambertian surface and mirror, both of which we have just dealt with. For the Lambertian surface, $\beta = 0$, while $\beta = \infty$ for a mirror surface. The appendix lists Fourier series expansions of general Phong BRDF functions, and figure 3 shows plots of the magnitudes of the coefficients.

In flatland, the Phong BRDF can be written in terms of the signed incident and outgoing angles as $\cos^\beta(\theta'_i + \theta'_o)$. Thus, the transfer function $\hat{\rho}$ would be of the form $\cos(\theta'_i) \cos^\beta(\theta'_i + \theta'_o)$. In the appendix, we show that for $|p| \geq \beta + 3$, the Fourier coefficient is either 0 or falls off sharply with increasing p . Furthermore, for large β , the Fourier coefficients decay like a gaussian of width approximately $\sqrt{2(\beta + 3)}$. Therefore, for large β , only lighting coefficients up to order $\sqrt{\beta}$, corresponding to the square-root of the Phong exponent, can be reliably recovered. In terms of signal-processing, the BRDF is viewed as a low-pass filter that truncates to zero, or severely diminishes, all high-frequency terms in the lighting. Therefore, these terms cannot be recovered by observing the output light field.

For a Lambertian surface, $\beta = 0$, and we expect that lighting coefficients with $|p| \geq 3$ cannot be easily recovered. As seen from the plot in figure 3, the Fourier coefficient for the Lambertian BRDF is 0 when $|p| = 3$ and falls off sharply for higher frequencies. For a mirror surface, $\beta = \infty$, indicating that we should be able to recover all the lighting coefficients, which is indeed the case.

Microfacet BRDFs: For microfacet BRDF models such as the Torrance-Sparrow model,¹⁹ the specular term is of the form $F \cdot G \cdot D$ where F is the Fresnel term, G is a geometric term, and D depends on the distribution of the microfacets. We will focus here on D , the distribution term. In general, D will be a function of $(N \cdot H)$ where N is the surface normal and H is the bisector of light and viewing directions—the half angle, given in flatland by $(\theta'_i + \theta'_o)/2$. A typical form for D would be $\exp[-\{(\theta'_i + \theta'_o)/2\}^2/2\alpha^2]$ where α is a measure of the surface roughness and corresponds to the width of the distribution. Because the distribution has a Gaussian form, the Fourier series will also have a Gaussian-like form with the width of the Gaussian being of the order $1/\alpha$. Therefore, the higher α is, or the rougher and more Lambertian the surface becomes, the faster the Fourier coefficients will diminish, making inverse lighting ill-conditioned. On the other hand, if α is very small, corresponding to a smooth mirror-like surface, the Fourier series coefficients will decrease slowly, and the inverse lighting problem will be well-conditioned.

5.2. Inverse BRDF computation

This subsection addresses the problem of recovering the BRDF, knowing the output light field and the incoming illumination field. The BRDF is a fundamental quantity, so the solution of this problem is of considerable practical interest. Recently, Lu et al.¹⁰ and Marschner et al.¹² have addressed this problem for a controlled point light source. We wish to work out the general case for uncontrolled illumination conditions. Because of the near-symmetry of the BRDF and lighting in our formulation, this problem can essentially be viewed as a dual to the inverse lighting problem, and gives essentially the same results.

The analogue to equation 16 is

$$\hat{\rho}_{p,q} = \frac{1}{2\pi} \frac{B_{-p,q}}{L_{-p}} \quad (24)$$

This equation says that BRDF recovery is well posed when no terms in the Fourier expansion of the lighting are 0. It is well conditioned when the Fourier expansion of the lighting does not decay rapidly with increasing frequency. In signal processing terms, the lighting can be viewed as the input signal, and the BRDF filter cannot reliably be estimated if certain frequencies in the signal are near zero. We will now consider several special cases for the input lighting and see how well-posed and well-conditioned the BRDF recovery problem is.

Single Directional Source: Consider a single directional light source at $\theta_i = 0$, so the lighting is of the form:

$$\begin{aligned} L(\theta_i) &= \delta(\theta_i) \\ L_{-p} &= \frac{1}{2\pi} \end{aligned} \quad (25)$$

Plugging into equation 24, we see that

$$\hat{\rho}_{p,q} = B_{-p,q} \quad (26)$$

which is simply the Fourier series expansion of

$$B(\alpha, \theta'_o) = \hat{\rho}(-\alpha, \theta'_o) \quad (27)$$

Therefore, for a directional source, the BRDF recovery problem is well-conditioned and well-posed in analogy with the inverse lighting problem for a mirror BRDF. Again, we are effectively estimating the BRDF filter by considering its impulse response. Therefore, BRDF measurement methods^{10,12} will give accurate results.

Uniform Lighting: We next consider lighting that is uniform everywhere i.e. $L(\theta_i) = 1/2\pi$. It can easily be verified that $L_0 = 1$ and that $L_p = 0$ for $p \neq 0$. Therefore, only the first (constant) term in the BRDF expansion over the incident angle can be recovered i.e. the only recoverable coefficients are

$$\begin{aligned} \hat{\rho}_{0,q} &= B_{0,q} \\ \hat{\rho}_0(\theta'_o) &= B_0(\theta'_o) \end{aligned} \quad (28)$$

Nothing can be said about higher-order terms[‡]. The lighting is a constant input signal, so only the constant term of the BRDF filter can be estimated. It can be easily verified that under constant lighting, a mirror surface and a Lambertian surface will both have the same output radiance regardless of the location and direction viewed, and it will be impossible to distinguish the two.

In between the two extremes of a single directional source and uniform lighting, the accuracy of BRDF recovery will depend on the frequency spectrum of the lighting. If the coefficients decay quickly, BRDF recovery will be ill-conditioned. Therefore, we want a nearly flat frequency spectrum, which is best achieved using a single directional source. However, our results indicate that it should still be possible to recover the BRDF even in the presence of an extraneous area source of relatively low magnitude, as might be expected, for instance, under skylight and sunlight.

5.3. Light Field Factorization

We now consider the problem of *factorizing* the light field i.e recovering both the lighting and BRDF when both are unknown. The output light field is two-dimensional while the lighting is one-dimensional and the BRDF is two-dimensional, but has half its parameters determined by symmetry between incident and outgoing angles. This seems to indicate that factorization is tractable, since we have more outputs than unknowns.

This problem is of theoretical interest and also has applications in a number of areas. It would be useful to have a passive method to recover the BRDF under uncontrolled conditions when we cannot measure the lighting. Similarly, it would be useful to be able to recover the lighting from an object of unknown BRDF. Reducing the dimensionality is very useful in compression when we seek to compress light fields that are usually very large. Further, understanding the structure of the light field can lead to novel intuitive ways to edit photographs.

For the purposes of this subsection, the reciprocity of the BRDF will prove to be crucial. Therefore, we use the reciprocal form of the transfer function given in equation 5, defining

$$\begin{aligned} \tilde{\rho}(\theta'_i, \theta'_o) &= \rho(\theta'_i, \theta'_o) \cos(\theta'_i) \cos(\theta'_o) \\ \tilde{B}(\alpha, \theta'_o) &= B(\alpha, \theta'_o) \cos(\theta'_o) \end{aligned}$$

The condition of reciprocity or symmetry in the BRDF implies that $\tilde{\rho}_{p,q} = \tilde{\rho}_{q,p}$. We can then use equation 15, reproduced below.

$$\tilde{B}_{p,q} = 2\pi L_p \tilde{\rho}_{-p,q}$$

There is a global scale factor that we cannot recover. That is, if we multiply the lighting everywhere by this factor, and simultaneously divide the BRDF everywhere by the same factor, we obtain an identical reflected light field. We are therefore allowed to arbitrarily set any one nonzero term in order to fix the scale factor. We will choose $L_0 = 1$. L_0 is simply the coefficient of the DC term in the Fourier expansion of the lighting, i.e. a measure of the

[‡]If reciprocity of the BRDF is considered, we are able to determine the corresponding terms with θ'_i and θ'_o exchanged, giving us slightly more information. The transfer function above is not reciprocal but can be made so using equation 5.

total energy of the incident illumination. Note that since the lighting is everywhere non-negative, L_0 cannot be 0 unless the lighting is 0 everywhere, an uninteresting case. It is now easy to use the condition of reciprocity $\tilde{\rho}_{p,q} = \tilde{\rho}_{q,p}$ to explicitly write

$$\begin{aligned}
L_0 &= 1 \\
\tilde{\rho}_{0,q} &= \frac{\tilde{B}_{0,q}}{2\pi} \\
L_p &= \frac{\tilde{B}_{p,q}}{2\pi\tilde{\rho}_{-p,q}} = \frac{\tilde{B}_{p,0}}{2\pi\tilde{\rho}_{-p,0}} = \frac{\tilde{B}_{p,0}}{2\pi\tilde{\rho}_{0,-p}} = \frac{\tilde{B}_{p,0}}{\tilde{B}_{0,-p}} \\
\tilde{\rho}_{p,q} &= \frac{\tilde{B}_{-p,q}}{2\pi L_{-p}} = \frac{\tilde{B}_{-p,q}}{2\pi\frac{\tilde{B}_{-p,0}}{\tilde{B}_{0,p}}} = \frac{\tilde{B}_{0,p}\tilde{B}_{-p,q}}{2\pi\tilde{B}_{-p,0}}
\end{aligned} \tag{29}$$

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.

The fact that this factorization is possible is likely to have important implications; it shows that lighting and BRDF can be factored, at least in principle, if the entire reflected light field is known. Of course, the results will be more and more ill-conditioned, the closer the output 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. Note that assuming reciprocity of the BRDF is critical. Without it, we would not be able to relate $\tilde{\rho}_{-p,0}$ and $\tilde{\rho}_{0,-p}$ above, and we would need a separate scale factor for each p (the index over which the lighting and BRDF interact, or filtering occurs) i.e. we would effectively need to know the lighting, or an equivalent set of information regarding the BRDF.

It should be noted that even if some of the $\tilde{\rho}_{0,q}$ terms vanish in equation 29, by using a different q in the recovery formula for L_p , we may still be able to factor the light field. Thus, L_p remains indeterminate only if for all q , $\tilde{\rho}_{-p,q}$ and $\tilde{\rho}_{q,-p}$ are both 0 or cannot be determined. Similarly, $\tilde{\rho}_{p,q}$ is undetermined only if both L_p and L_q cannot be determined or are 0.

5.4. Environment Map Prefiltering and Rendering

Since computing an image by calculating the outgoing radiance for each point on the surface in real time is often too computationally intensive, environment maps are often prefiltered.⁵ For instance, for a Lambertian convex surface, we prefilter to obtain an intensity for each of a set of discretized directions of the surface normal. Once prefiltering has been done, the environment can be viewed in real time.

Prefiltering remains an expensive operation, which means dynamic lighting or lighting design applications are difficult to support. However, equation 14 or equation 11 can be used to efficiently prefilter and render environment maps in Fourier space. For smooth BRDFs, only a few terms of the Fourier expansion need be considered. This method will not work well for highly specular or otherwise peaked BRDF distributions. As in much of the previous work on environment map prefiltering, we can separate these into diffuse and specular components, treating both separately. The smooth or near-diffuse component is well handled by a few terms of the Fourier expansion and is treated much more efficiently in Fourier space than conventionally, where an integral over the entire incoming hemisphere must be performed. On the other hand, the Fourier coefficients for specular terms will not decay rapidly and many terms will be needed to approximate their effect correctly. However, they are well handled in the conventional way, since only a very small region of the incoming illumination must be considered. Thus, we can choose angular or Fourier space, depending on whether we want to consider a near-diffuse or near-specular BRDF component. This follows from the well-known observation that the Fourier transform of a localized function is highly non-local and near-constant while the Fourier transform of a near-constant non-local function is localized.

Consider the special case of a Lambertian surface, for which equation 11 becomes—using equation 22 for even p , and dropping the dependence on exitant angle:

$$B_{2p} = L_{2p} \left(\frac{(-1)^{-p+1}}{4p^2 - 1} \right) \tag{30}$$

This falls off rapidly as $1/p^2$, so a handful of terms will suffice to give accurate results. The most computationally intensive part is converting to and from the Fourier representation, but the Fourier basis functions can be precomputed and stored as images, so graphics hardware can be used to perform the multiplications and additions required.

6. DISCUSSION

This section discusses some applications to visual perception, and some issues involved in extending this work to concave surfaces and full 3D light fields.

Implications for Visual Perception: From equation 30, we know that the Lambertian BRDF acts as a low pass filter. This means that regardless of the lighting distribution, the output light field will not include high-frequency terms in the absence of shadows. Therefore, neglecting shadowing effects, when we perceive high-frequency variation in a curved surface, this can only be due to specularities or gloss of the surface, or a surface texture. This may help explain why we can determine the shininess of a (textureless) surface independently of the lighting. From a single image, texture and shininess are fundamentally ambiguous and cannot be distinguished; the visual system probably employs heuristics—for instance, specularities are often much brighter than any surrounding texture—to disambiguate the two. From multiple images from different viewpoints, however, the strong view-dependence of specularities allows it to be distinguished from texture variations. This suggests, that at least in theory, one should be able to separate texture, lighting and the BRDF from the output light field. In practice, the separation of lighting and texture-related effects appears to be ill-conditioned—humans are not easily able to distinguish slow subtle texture variations from lighting-related effects. Of course, the visual system also makes various heuristic assumptions to aid in factorization, but we believe that our results will help clarify where assumptions are and are not needed. For well conditioned inverse problems, extraneous assumptions are not needed, while for ill-conditioned or fundamentally ambiguous problems, it is likely that strong heuristics must be employed.

Concave Surfaces: Our treatment has considered only convex surfaces. For a concave surface, we must also consider self-shadowing and interreflection. Therefore, the surface orientation is no longer sufficient to determine all characteristics of the incident and outgoing radiance. We can no longer parameterize using the surface normal only, and must explicitly consider location on the surface \mathbf{x} as well as orientation. The shadowing function is a binary term that must be included in the integral in equation 3. Since the shadowing function is highly discontinuous, it makes the kernel of the integral much less smooth. This makes the forward problem (for instance, environment map rendering) significantly harder. However, it makes inverse problems better conditioned (albeit potentially harder computationally and much less tractable to mathematical analysis). Sato et al.¹⁷ have exploited this in order to solve an inverse lighting problem using shadowing information on a plane—a surface without any curvature at all. Interreflection is much harder to consider and is usually ignored. A notable exception is the work of Yu et al.²²

Extension to 3D: The theory presented here is limited to planar configurations. By restricting ourselves to 2D, we have been able to present the key ideas while keeping the mathematics more familiar and straightforward. In real-world applications, we must deal with a full three-dimensional world. We have worked out the generalization to 3D, and will report on this in a future paper. The 3D case is conceptually a fairly straightforward extension of the 2D case, with most of the key insights transferring with only minor modifications, but is mathematically somewhat more involved. In particular, the rotation operator is much more complicated in 3D. The 3D treatment can be unified conceptually with the treatment presented in this paper by considering the structures of the respective rotation groups— $SO(2)$ in 2D and $SO(3)$ in 3D—and their associated representations. The analogue of the 2D Fourier basis in 3D is the spherical harmonic series, and the simple rotation of the lighting in 2D is replaced with convolution using representations of $SO(3)$. We emphasize that while somewhat more complex, the 3D theory leads to essentially the same conclusions and observations as those reported in this paper for 2D.

7. CONCLUSIONS

We have presented a theoretical analysis in flatland of the structure of the reflected light field from a convex homogeneous object under a distant illumination field. We have shown that the direct illumination integral can be viewed in signal processing terms as a filtering operation between the lighting and the BRDF to produce the output light field.

This result provides a novel viewpoint for many forward and inverse rendering problems. We have demonstrated 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.

With the aid of our theory, we have been able to determine the well-posedness and conditioning of many inverse problems. This serves as a guide to future endeavors in inverse rendering for computer vision and graphics. In a few cases, we have been able to show in signal processing terms why approaches taken by other researchers fail to perform as expected. In other cases, we have shown the optimality of commonly-used methods such as using a mirrored sphere and point light sources. We hope that in the future, the theory can guide us toward efficient algorithms and help explain the success or failure of various methods.

Acknowledgements

We thank Marc Levoy for many helpful discussions, especially regarding the interpretation of reflection as a convolution operation. Steve Marschner and Szymon Rusinkiewicz read early drafts, and provided many insightful comments. This work was supported in part by a Hodgson-Reed Stanford graduate fellowship.

REFERENCES

1. E.H. Adelson and J.R. Bergen. *Computational Models of Visual Processing*, chapter The Plenoptic Function and the Elements of Early Vision. MIT Press, 1991.
2. B. Cabral, M. Olano, and P. Nemecek. Reflection space image based rendering. In *SIGGRAPH 99*, pages 165–170, 1999.
3. A. Gershun. The light field. *Journal of Mathematics and Physics*, XVIII:51–151, 1939. Translated by P. Moon and G. Timoshenko.
4. S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *SIGGRAPH 96*, pages 43–54, 1996.
5. N. Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics & Applications*, 6(11):21–29, 1986.
6. K. F. Karner, H. Mayer, and M. Gervautz. An image based measurement system for anisotropic reflection. *Computer Graphics Forum*, 15(3):119–128, 1996.
7. J. Kautz, P. Vázquez, W. Heidrich, and H.P. Seidel. A unified approach to prefiltered environment maps. In *11th Eurographics Workshop on Rendering*, pages 185–196, 2000.
8. G.J. Klinker, S.A. Shafer, and T. Kanade. The measurement of highlights in color images. *IJCV*, 2(1):7–32, 1988.
9. M. Levoy and P. Hanrahan. Light field rendering. In *SIGGRAPH 96*, pages 31–42, 1996.
10. R. Lu, J.J. Koenderink, and A.M.L. Kappers. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics*, 37(25):5974–5984, 1998.
11. S.R. Marschner and D.P. Greenberg. Inverse lighting for photography. In *Fifth Color Imaging Conference*, pages 262–265, 1997.
12. S.R. Marschner, S.H. Westin, E.P.F. Lafortune, and K.E. Torrance. Image-Based BRDF measurement. *Applied Optics*, 39(16):2592–2600, 2000.
13. N. McCormick. Inverse radiative transfer problems: a review. *Nuclear Science and Engineering*, 112:185–198, 1992.
14. 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.
15. 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.
16. R.W. Preisendorfer. *Hydrologic Optics*. US Dept Commerce, 1976.
17. 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.
18. Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *SIGGRAPH 97*, pages 379–388, 1997.

19. K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1114, 1967.
20. G. J. Ward. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92*, pages 265–272, 1992.
21. D. N. Wood, D. I. Azuma, K. Aldinger, B. Curless, T. Duchamp, D. H. Salesin, and W. Stuetzle. Surface light fields for 3D photography. In *SIGGRAPH 2000*, pages 287–296, 2000.
22. 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.

Appendix: Fourier Series for Phong BRDFs

This appendix considers the Fourier series representations for Phong BRDFs, giving analytic formulae in support of the discussion regarding Phong BRDFs in subsection 5.1. For simplicity, and ease of plotting, we restrict ourselves[§] to the case of $\theta'_o = 0$ i.e. normal viewing angle.

For exponent β , ignoring normalization, the Phong BRDF is given by $\cos^\beta(\theta'_i + \theta'_o)$. We set $\theta'_o = 0$ for normal viewing angle, and find the Fourier coefficients

$$\hat{\rho}_{-p}(0) = \frac{1}{2\pi} \int_{-\pi/2}^{\pi/2} \cos(\theta'_i)^{\beta+1} e^{Ip\theta'_i} d\theta'_i \quad (31)$$

which evaluates to

$$\hat{\rho}_{-p}(0) = \frac{\Gamma(2 + \beta)}{(4 \cdot 2^\beta) \Gamma(\frac{3-p+\beta}{2}) \Gamma(\frac{3+p+\beta}{2})} \quad (32)$$

where Γ is Euler’s Gamma function, which is the factorial function for positive integers: $\Gamma(n) = (n - 1)!$, and is infinite for non-positive integers. It should be noted that the above formula is symmetric about $p = 0$.

We observe that when $p = \pm(\beta + 3 + 2k)$ for a non-negative integer k , the corresponding Fourier coefficient $\hat{\rho}_{\mp p}(0)$ vanishes. This is because the argument of one of the gamma functions in the denominator becomes a negative integer, causing the denominator to become infinite. This is why odd frequencies with $|p| > 1$ vanish for the case of a Lambertian surface, where $\beta = 0$.

We may also observe that for $|p| \gg \beta + 3$, the asymptotic behavior of the Fourier coefficients will depend on the asymptotic behavior of the gamma functions in the denominator. The asymptotic behavior of the Fourier coefficients goes as $1/[\Gamma(a + x)\Gamma(a - x)]$ where $a = (\beta + 3)/2$ and $x = p/2$. The asymptotic behavior of this quantity can be shown to depend on x/x^{2a} which implies that for large $|p|$,

$$|p| > 3 + \beta \Rightarrow \hat{\rho}_{\pm p}(0) \sim |p|^{-(2+\beta)} \quad (33)$$

which explains the $1/p^2$ falloff for Lambertian surfaces.

For $|p| \ll \beta$, $1/[\Gamma(a + x)\Gamma(a - x)]$ for $a = (\beta + 3)/2$ and $x = p/2$ goes as $\exp[-x^2(1/a - 1/2a^2)]$. Thus, the coefficients fall off like a gaussian with width $\sqrt{2(\beta + 3)}$.

As far as practical issues are concerned, this means inverse lighting is ill-conditioned for $|p| \sim \sqrt{\beta}$, and ill-posed for $|p| \geq \beta + 3$. As far as possible, specular surfaces with larger values of β should be used. Plots of the coefficients for some of these BRDFs are shown in figure 3.

[§]If we ignore boundary effects, this can be shown to be equivalent to reparameterizing by the reflection vector instead of the surface normal, and using a transfer function with exponent $\beta + 1$.