

# Modeling Illumination Variation with Spherical Harmonics

Ravi Ramamoorthi  
Columbia University  
Email: ravir@cs.columbia.edu

## 0.1 Introduction

Illumination can have a significant impact on the appearance of surfaces, as the patterns of shading, specularities and shadows change. For instance, some images of a face under different lighting conditions are shown in figure 1. Differences in lighting can often play a much greater role in image variability of human faces than differences between individual people. Lighting designers in movies can often set the mood of a scene with carefully chosen lighting. To achieve a sinister effect, for instance, one can use illumination from below the subject—a sharp contrast to most natural indoor or outdoor scenes where the dominant light sources are above the person.

Characterizing the variability in appearance with lighting is a fundamental problem in many areas of computer vision, face modeling, and computer graphics. One of the great challenges of computer vision is to produce systems that can work in uncontrolled environments. To be robust, recognition systems must be able to work outdoors in a lighting-insensitive manner. In computer graphics, the challenge is to be able to efficiently create the visual appearance of a scene under realistic, possibly changing, illumination.

At first glance, modeling the variation with lighting may seem intractable. For instance, a video projector can illuminate an object like a face with essentially any pattern. In this chapter, we will stay away from such extreme examples, making a set of assumptions that are approximately true in many common situations. One assumption we make is that illumination is distant. By this, we mean that the direction to, and intensity of, the light sources is approximately the same throughout the region of interest. This explicitly rules out cases like slide projectors. This is a reasonably good approximation in outdoor scenes, where the sky can be assumed to be far away. It is also fairly accurate in many indoor environments, where the light sources can be considered much further away relative to the size of the object.

Even under the assumption of distant lighting, the variations may seem intractable. The illumination can come from any incident direction, and can be composed of multiple illuminants including localized light sources like sunlight and broad area distributions like skylight. In the general case, we would need to model the intensity from each of infinitely many incident lighting directions. Thus, the space we are dealing with appears to be infinite-dimensional. By contrast, a number of other causes of appearance variation are low-dimensional. For instance, appearance varies with viewing direction as well. However, unlike lighting, there can only be a single view direction. In general, variation because of pose and translation can be described using only six degrees of freedom.

Given the daunting nature of this problem, it is not surprising that most previous analytic models have been restricted to the case of a single directional (distant) light source, usually without considering shadows. In computer graphics, this model is sufficient (but not necessarily efficient) for numerical Monte Carlo simulation of appearance, since one can simply treat multiple light sources separately and add up the results. In computer vision, such models can be reasonable approximations under some circumstances, such as controlled laboratory conditions. However, they do not suffice for modeling illumination variation in uncontrolled conditions like the outdoors.

Fortunately, there is an obvious empirical method of analyzing illumination variation. One can simply record a number of images of an object under light sources from all different directions. In practice, this usually corresponds to moving a light source in a sphere around the object or person, while keeping camera and pose fixed. Because of the linearity of light transport—the image under multiple lights is the sum of that under the individual sources—an image under arbitrary distant illumination can be written as a linear combination of these source images. This observation in itself provides an attractive direct approach for relighting in computer graphics. Furthermore, instead of using the source images, we can try to find linear combinations or basis images that best explain the




Figure 1: Images of a face, lit from a number of different directions. Note the vast variation in appearance due to illumination. Images courtesy of Debevec et al. [13].

variability due to illumination. This is the standard technique of finding low-dimensional subspaces or principal components. The surprising result in these experiments is that most of the variability due to lighting is modeled using a very low-dimensional subspace—usually only five basis functions.

Given the discussion about the infinite-dimensionality of illumination, this observation seems strikingly counter-intuitive. However, some insight may be obtained by looking at an (untextured) matte or diffuse surface. One will notice that even if the lighting is complicated, the surface has a smooth appearance. In effect, it is blurring or low-pass filtering the illumination. In this chapter, we will make these ideas formal, explaining previous empirical results. A diffuse or matte surface (technically Lambertian) can be viewed as a low-pass filter acting on the incident illumination signal, with the output image obtained by convolving the lighting with the surface reflectance. This leads to a frequency-domain view of reflection, that has not previously been explored, and that leads to many interesting insights. In particular, we can derive a product convolution formula using spherical harmonic basis functions. The theoretical results have practical implications for computer graphics rendering, illuminant estimation, and recognition and reconstruction in computer vision.

## 0.2 Background and Previous Work

Lighting and appearance have been studied in many forms almost since the beginning of research in computer vision and graphics, as well as in a number of other areas. Horn’s classic text [27] provides background on work in vision. In this section, we only survey the most relevant previous work, that relates to the theoretical developments in this chapter. This will also be a vehicle to introduce some of the basic concepts on which we will build.

### 0.2.1 Environment Maps in Computer Graphics

A common approximation in computer graphics, especially for interactive hardware rendering, is to assume distant illumination. The lighting can then be represented as a function of direction, known in the literature as an environment map. Practically, environment maps can be acquired by photographing a chrome-steel or mirror sphere, that simply reflects the incident lighting. Environment mapping corresponds exactly to the distant illumination assumption we make in this chapter. Another common assumption, which we also make, is to neglect *cast shadows* from one part of the object on another. These should be distinguished from *attached shadows* when a point is in shadow because the light source is behind the surface (below the horizon). We will explicitly take attached shadows into account.

In terms of previous work, Blinn and Newell [7] first used environment maps to efficiently find reflections of distant objects. This technique, known as reflection mapping, is still widely in use today for interactive computer graphics rendering. The method was greatly generalized by Miller and Hoffman [46] (and later Greene [21]). They introduced the idea that one could acquire a real environment by photographing a mirror sphere. They also precomputed diffuse and specular reflection maps, for the corresponding components of surface reflection. Cabral et al. [9] later extended this general method to computing reflections from bump-mapped surfaces, and to computing environment-mapped images with more general reflective properties [10] (technically, the bi-directional reflectance

distribution function or BRDF [52]). Similar results for more general materials were also obtained by Kautz et al. [32, 33], building on previous work by Heidrich and Seidel [25]. It should be noted that both Miller and Hoffman [46], and Cabral et al. [9, 10] qualitatively described the reflection maps as obtained by convolving the lighting with the reflective properties of the surface. Kautz et al. [33] actually used convolution to implement reflection, but on somewhat distorted planar projections of the environment map, and without full theoretical justification. In this chapter, we will formalize these ideas, making the notion of convolution precise, and derive analytic formulas.

### 0.2.2 Lighting-Insensitive Recognition in Computer Vision

Another important area of research is in computer vision, where there has been much work on modeling the variation of appearance with lighting for robust lighting-insensitive recognition algorithms. Some of this prior work is discussed in detail in the excellent chapter *on the effect of illumination and face recognition* by Ho and Kriegman in this volume. Illumination modeling is also important in many other vision problems, such as photometric stereo and structure from motion.

The work in this area has taken two directions. One approach is to apply an image processing operator that is relatively insensitive to the lighting. Work in this domain includes image gradients [8], the direction of the gradient [11], and Gabor jets [37]. While these methods reduce the effects of lighting variation, they operate without knowledge of the scene geometry, and so are inherently limited in the ability to model illumination variation. A number of empirical studies have suggested the difficulties of pure image-processing methods, in terms of achieving lighting-insensitivity [47]. Our approach is more related to a second line of work that seeks to explicitly model illumination variation, analyzing the effects of the lighting on a 3D model of the object (here, a face).

**Linear 3D Lighting Model Without Shadows:** Within this category, a first important result (Shashua [69], Murase and Nayar [48], and others) is that for Lambertian objects, in the absence of *both* attached and cast shadows, the images under all possible illumination conditions lie in a three-dimensional subspace. To obtain this result, let us first consider a model for the illumination from a single directional source on a Lambertian object.

$$B = \rho L(\boldsymbol{\omega} \cdot \mathbf{n}) = \mathbf{L} \cdot \mathbf{N}, \quad (1)$$

where  $L$  in the first relation corresponds to the intensity of the illumination in direction  $\boldsymbol{\omega}$ ,  $\mathbf{n}$  is the surface normal at the point,  $\rho$  is the surface albedo and  $B$  is the outgoing radiosity (radiant exitance) or reflected light. For a nomenclature of radiometric quantities, see a textbook, like McCluney [45] or chapter 2 of Cohen and Wallace [12].<sup>1</sup> It is sometimes useful to incorporate the albedo into the surface normal, defining vectors  $\mathbf{N} = \rho \mathbf{n}$  and correspondingly  $\mathbf{L} = L\boldsymbol{\omega}$ , so we can simply write (for Lambertian surfaces in the absence of all shadows, attached or cast)  $B = \mathbf{L} \cdot \mathbf{N}$ .

Now, consider illumination from a variety of light sources  $\mathbf{L}_1, \mathbf{L}_2$  and so on. It is straightforward to use the linearity of light transport to write the net reflected light as

$$B = \sum_i \mathbf{L}_i \cdot \mathbf{N} = \left( \sum_i \mathbf{L}_i \right) \cdot \mathbf{N} = \hat{\mathbf{L}} \cdot \mathbf{N}, \quad (2)$$

where  $\hat{\mathbf{L}} = \sum_i \mathbf{L}_i$ . But this has essentially the same form as equation 1. Thus, in the absence of all shadows, there is a very simple linear model for lighting in Lambertian surfaces, where we can replace a complex lighting distribution by the weighted sum of the individual light sources. We can then treat the object as if lit by a *single* effective light source  $\hat{\mathbf{L}}$ .

Finally, it is easy to see that images under all possible lighting conditions lie in a 3D subspace, being linear combinations of the cartesian components  $N_x, N_y, N_z$ , with  $B = \hat{L}_x N_x + \hat{L}_y N_y + \hat{L}_z N_z$ . Note that this entire analysis applies to all points on the object. The basis images of the 3D lighting subspace are simply the Cartesian components of the surface normals over the objects, scaled by the albedos at those surface locations<sup>2</sup>.

<sup>1</sup>In general, we will use  $B$  for the reflected radiance, and  $L$  for the incident radiance, while  $\rho$  will denote the BRDF. For Lambertian surfaces, it is conventional to use  $\rho$  to denote the albedo (technically, the surface reflectance that lies between 0 and 1), which is  $\pi$  times the BRDF. Interpreting  $B$  as the radiosity accounts for this factor of  $\pi$ .

<sup>2</sup>Color is largely ignored in this chapter; we assume each color band, such as red, green and blue, is treated separately.

The 3D linear subspace has been used in a number of works on recognition, as well as other areas. For instance, Hayakawa [24] used factorization based on the 3D subspace to build models using photometric stereo. Koenderink and van Doorn [35] added an extra ambient term to make this a 4D subspace. The extra term corresponds to perfectly diffuse lighting over the whole sphere of incoming directions. This corresponds to adding the albedos over the surface themselves to the previous 3D subspace, i.e. adding  $\rho = \| \mathbf{N} \| = \sqrt{N_x^2 + N_y^2 + N_z^2}$ .

**Empirical Low-Dimensional Models:** These theoretical results have inspired a number of authors to develop empirical models for lighting variability. As described in the introduction, one takes a number of images with different lighting directions, and then uses standard dimensionality reduction methods like Principal Component Analysis. PCA-based methods were pioneered for faces by Kirby and Sirovich [34, 72], and for face recognition by Turk and Pentland [78], but these authors did not account for variations due to illumination.

The effects of lighting alone were studied in a series of experiments by Hallinan [23], Epstein et al. [15], and Yuille et al. [80]. They found that for human faces, and other diffuse objects like basketballs, a 5D subspace sufficed to approximate lighting variability very well. That is, with a linear combination of a mean and 5 basis images, we can accurately predict appearance under arbitrary illumination. Furthermore, the form of these basis functions, and even the amount of variance accounted for, were largely consistent across human faces. For instance, Epstein et al. [15] report that for images of a human face, three basis images capture 90% of image variance, while five basis images account for 94%.

At the time however, these results had no complete theoretical explanation. Furthermore, they indicate that the 3D subspace given above is inadequate. This is not difficult to understand. If we consider the appearance of a face in outdoor lighting from the entire sky, there will often be attached shadows or regions in the environment that are not visible to a given surface point (these correspond to lighting below the horizon for that point, where  $\boldsymbol{\omega} \cdot \mathbf{n} < 0$ ).

**Theoretical Models:** The above discussion indicates the value of developing an analytic model to account for lighting variability. Theoretical results can give new insights, and also lead to simpler and more efficient and robust algorithms.

Belhumeur and Kriegman [6] have taken a first step in this direction, developing the illumination cone representation. Under very mild assumptions, the images of an object under arbitrary lighting form a convex cone in image space. In a sense, this follows directly from the linearity of light transport. Any image can be scaled by a positive value simply by scaling the illumination. The convexity of the cone is because one can add two images, simply by adding the corresponding lighting. Formally, even for a Lambertian object with only attached shadows, the dimension of the illumination cone can be infinite (this grows as  $O(n^2)$ , where  $n$  is the number of distinct surface normals visible in the image). Georgiades et al. [19, 20] have developed recognition algorithms based on the illumination cone. One approach is to sample the cone using extremal rays, corresponding to rendering or imaging the face using directional light sources. It should be noted that exact recognition using the illumination cone involves a slow complex optimization (constrained optimization must be used to enforce a non-negative linear combination of the basis images), and methods using low-dimensional linear subspaces and unconstrained optimization (which essentially reduces to a simple linear system) are more efficient and usually required for practical applications.

Another approach is to try to analytically construct the principal component decomposition, analogous to what was done experimentally. Numerical PCA techniques could be biased by the specific lighting conditions used, so an explicit analytic method is helpful. It is only recently that there has been progress in analytic methods for extending PCA from a discrete set of images to a continuous sampling [40, 83]. These approaches demonstrate better generalization properties than pure empirical techniques. In fact, Zhao and Yang [83] have analytically constructed the co-variance matrix for PCA of lighting variability, but under the assumption of no shadows.

**Summary:** To summarize, prior to the work reported in this chapter, the illumination variability could be described theoretically by the illumination cone. It was known from numerical and real experiments that the illumination cone lay close to a linear low-dimensional space for Lambertian objects with attached shadows. However, a theoretical explanation of these results was not available. In this chapter, we develop a simple linear lighting model using spherical harmonics. These theoretical results were first introduced for Lambertian surfaces simultaneously by Basri and Jacobs [2, 4] and Ramamoorthi and Hanrahan [64]. Much of the work of Basri and Jacobs is also summarized in an excellent book chapter on illumination modeling for face recognition [5].

### 0.2.3 Frequency-Space Representations: Spherical Harmonics

We show reflection to be a convolution and analyze it in frequency space. We will primarily be concerned with analyzing quantities like the BRDF and distant lighting, which can be parameterized as functions on the unit sphere. Hence, the appropriate frequency-space representations are spherical harmonics [28, 29, 42]. Spherical harmonics can be thought of as signal-processing tools on the unit sphere, analogous to the Fourier series or sines and cosines on the line or circle. They can be written either as trigonometric functions of the spherical coordinates, or as simple polynomials of the Cartesian components. They form an orthonormal basis on the unit sphere, in terms of which quantities like the lighting or BRDF can be expanded and analyzed.

The use of spherical harmonics to represent the illumination and BRDF was pioneered in computer graphics by Cabral et al. [9]. In perception, D’Zmura [14] analyzed reflection as a linear operator in terms of spherical harmonics, and discussed some resulting ambiguities between reflectance and illumination. Our use of spherical harmonics to represent the lighting is also similar in some respects to previous methods such as that of Nimeroff et al. [56] that use steerable linear basis functions. Spherical harmonics have also been used before in computer graphics for representing BRDFs by a number of other authors [70, 79].

The results described in this chapter are based on a number of papers by us. This includes theoretical work in the planar 2D case or flatland [62], on the analysis of the appearance of a Lambertian surface using spherical harmonics [64], the theory for the general 3D case with isotropic BRDFs [65], and a comprehensive account including a unified view of 2D and 3D cases including anisotropic materials [67]. More details can also be found in the PhD thesis of the author [61]. Recently, we have also linked the convolution approach using spherical harmonics with principal component analysis, quantitatively explaining previous empirical results on lighting variability [60].

## 0.3 Analyzing Lambertian Reflection using Spherical Harmonics

In this section, we analyze the important case of Lambertian surfaces under distant illumination [64], using spherical harmonics. We will derive a simple linear relationship, where the coefficients of the reflected light are simply filtered versions of the incident illumination. Furthermore, a low frequency approximation with only 9 spherical harmonic terms is shown to be accurate.

**Assumptions:** Mathematically, we are simply considering the relationship between the irradiance (a measure of the intensity of incident light, reflected equally in all directions by diffuse objects), expressed as a function of surface orientation, and the incoming radiance or incident illumination, expressed as a function of incident angle. The corresponding physical system is a curved convex homogeneous Lambertian surface reflecting a distant illumination field. For the physical system, we will assume that the *surfaces under consideration are convex*, so they may be parameterized by the surface orientation, as described by the surface normal, and so that *interreflection and cast shadowing (but not attached shadows) can be ignored*. Also, surfaces will be assumed to be *Lambertian and homogeneous*, so the reflectivity can be characterized by a constant albedo. We will further assume here that the *illuminating light sources are distant*, so the illumination or incoming radiance can be represented as a function of direction only.

Notation used in the chapter (some of which pertains to later sections) is listed in table 2. A diagram of the local geometry of the situation is shown in figure 3. We will use two types of coordinates. Unprimed global coordinates denote angles with respect to a global reference frame. On the other hand, primed local coordinates denote angles with respect to the local reference frame,

$B$	Reflected radiance
$B_{lmn,pq}, B_{lm pq}$	Coefficients of basis-function expansion of $B$
$L$	Incoming radiance
$L_{lm}$	Coefficients of spherical-harmonic expansion of $L$
$E$	Incident irradiance (for Lambertian surfaces)
$E_{lm}$	Coefficients of spherical-harmonic expansion of $E$
$\rho$	Surface BRDF
$\hat{\rho}$	BRDF multiplied by cosine of incident angle
$\hat{\rho}_{ln,pq}, \hat{\rho}_{lpq}$	Coefficients of spherical-harmonic expansion of $\hat{\rho}$
$\theta'_i, \theta_i$	Incident elevation angle in <i>local</i> , <i>global</i> coordinates
$\phi'_i, \phi_i$	Incident azimuthal angle in <i>local</i> , <i>global</i> coordinates
$\theta'_o, \theta_o$	Outgoing elevation angle in <i>local</i> , <i>global</i> coordinates
$\phi'_o, \phi_o$	Outgoing azimuthal angle in <i>local</i> , <i>global</i> coordinates
$A(\theta'_i)$	Half-cosine Lambertian transfer function $A(\theta'_i) = \max(\cos(\theta'_i), 0)$
$A_l$	Spherical harmonic coefficients of Lambertian transfer function
$\mathbf{x}$	Surface position
$\alpha$	Surface normal parameterization—elevation angle
$\beta$	Surface normal parameterization—azimuthal angle
$\gamma$	Orientation of tangent frame for anisotropic surfaces
$R_{\alpha,\beta,\gamma}$	Rotation operator for tangent frame orientation $(\alpha, \beta, \gamma)$
$Y_{lm}$	Spherical Harmonic
$Y_{lm}^*$	Complex Conjugate of Spherical Harmonic
$D_{mm'}^l$	Representation matrix of dimension $2l + 1$ for rotation group $SO(3)$
$\Lambda_l$	Normalization constant, $\sqrt{4\pi/(2l + 1)}$
$I$	$\sqrt{-1}$

Figure 2: Notation

defined by the local surface normal and an arbitrarily chosen tangent vector. These two coordinate systems are related simply by a rotation, and this relationship will be detailed shortly.

### 0.3.1 Reflection Equation

In local coordinates, we can relate the irradiance to the incoming radiance by

$$E(\mathbf{x}) = \int_{\Omega'_i} L(\mathbf{x}, \theta'_i, \phi'_i) \cos \theta'_i d\Omega'_i, \quad (3)$$

where  $E$  is the irradiance, as a function of position  $\mathbf{x}$  on the object surface, and  $L$  is the radiance of the incident light field. As noted earlier, primes denote quantities in local coordinates. The integral is over the upper hemisphere with respect to the local surface normal. For the purposes of this derivation, we will be interested in the relationship of the irradiance to the radiance. In practice, the reflected light (here, radiant exitance or radiosity) can be related to the irradiance using  $B(\mathbf{x}) = \rho E(\mathbf{x})$ , where  $\rho$  is the surface reflectance or albedo that lies between 0 and 1. This last relation also holds when we interpret  $\rho$  as the BRDF (obtained by scaling the reflectance by  $1/\pi$ ) and  $B$  as the reflected radiance (obtained by scaling the radiant exitance by  $1/\pi$ ).

We now manipulate equation 3 by performing a number of substitutions. First, the assumption of distant illumination means the illumination field is homogeneous over the surface, i.e. independent of surface position  $\mathbf{x}$ , and depends only on the *global* incident angle  $(\theta_i, \phi_i)$ . This allows us to replace  $L(\mathbf{x}, \theta'_i, \phi'_i) \rightarrow L(\theta_i, \phi_i)$ . Second, consider the assumption of a curved convex surface. This ensures there is no shadowing or interreflection, so that the irradiance is only because of the distant illumination field  $L$ . This fact is implicitly assumed in equation 3. Furthermore, since the illumination is distant, we may reparameterize the surface simply by the surface normal  $\mathbf{n}$ . Equation 3 now becomes

$$E(\mathbf{n}) = \int_{\Omega'_i} L(\theta_i, \phi_i) \cos \theta'_i d\Omega'_i. \quad (4)$$




Figure 3: Diagram showing the local geometry. Quantities are primed because they are all in local coordinates.

To proceed further, we will parameterize  $\mathbf{n}$  by its spherical angular-coordinates  $(\alpha, \beta, \gamma)$ . Here,  $(\alpha, \beta)$  define the angular coordinates of the local normal vector, i.e.

$$\mathbf{n} = [\sin \alpha \cos \beta, \sin \alpha \sin \beta, \cos \alpha]. \quad (5)$$

$\gamma$  defines the local tangent frame, i.e. rotation of the coordinate axes about the normal vector. For isotropic surfaces—those where there is no preferred tangential direction, i.e. where rotation of the tangent frame about the surface normal has no physical effect—the parameter  $\gamma$  has no physical significance, and we have therefore not explicitly considered it in equation 5. We will include  $\gamma$  for completeness in the ensuing discussion on rotations, but will eventually eliminate it from our equations after showing mathematically that it does in fact have no effect on the final results. Finally, for convenience, we will define a transfer function  $A(\theta'_i) = \cos \theta'_i$ . With these modifications, equation 4 becomes

$$E(\alpha, \beta, \gamma) = \int_{\Omega'_i} L(\theta_i, \phi_i) A(\theta'_i) d\Omega'_i. \quad (6)$$

Note that local and global coordinates are mixed. The lighting is expressed in global coordinates, since it is constant over the object surface when viewed with respect to a global reference frame, while the transfer function  $A = \cos \theta'_i$  is expressed naturally in local coordinates. 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 6, we must relate local and global coordinates. The *north pole*  $(0', 0')$  or  $+Z$  axis in local coordinates is the surface normal, and the corresponding global coordinates are  $(\alpha, \beta)$ . 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. Refer to figure 4 for an illustration. The relevant transformations are given below,

$$\begin{aligned} (\theta_i, \phi_i) &= R_{\alpha, \beta, \gamma}(\theta'_i, \phi'_i) = R_z(\beta)R_y(\alpha)R_z(\gamma) \{\theta'_i, \phi'_i\} \\ (\theta'_i, \phi'_i) &= R_{\alpha, \beta, \gamma}^{-1}(\theta_i, \phi_i) = R_z(-\gamma)R_y(-\alpha)R_z(-\beta) \{\theta_i, \phi_i\}. \end{aligned} \quad (7)$$




Figure 4: Diagram showing how the rotation corresponding to  $(\alpha, \beta, \gamma)$  transforms between local (primed) and global (unprimed) coordinates.

Note that the angular parameters are rotated as if they were a unit vector pointing in the appropriate direction. It should also be noted that this rotation of parameters is equivalent to an inverse rotation of the function, with  $R^{-1}$  being given by  $R_z(-\gamma)R_y(-\alpha)R_z(-\beta)$ .

Finally, we can substitute equation 7 into equation 6 to derive

$$E(\alpha, \beta, \gamma) = \int_{\Omega'_i} L(R_{\alpha, \beta, \gamma}(\theta'_i, \phi'_i)) A(\theta'_i) d\Omega'_i. \quad (8)$$

As we have written it, this equation depends 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  $3 \times 3$  rotation matrix, while  $\omega_i$  is a unit vector ( $3 \times 1$  column vector) corresponding to the incident direction (with primes added for local coordinates). Equation 8 may then be written as

$$E(R) = \int_{\Omega'_i} L(R\omega'_i) A(\omega'_i) d\omega'_i, \quad (9)$$

where  $R\omega'_i$  is simply a matrix-vector multiplication.

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

$$(f \otimes g)(a) = \int_t f(T_a(t)) g(t) dt. \quad (10)$$

When  $T_a$  is a translation by  $a$ , we obtain the standard expression for spatial convolution. When  $T_a$  is a rotation by the angle  $a$ , the above formula defines convolution in the angular domain, and is a slightly simplified version of equations 8 and 9.

The irradiance can therefore be viewed as obtained by taking the incident illumination signal  $L$  and filtering it using the transfer function  $A = \cos \theta'_i$ . Different observations of the irradiance  $E$ , at points on the object surface with different orientations, correspond to different rotations of the






Figure 5: The first 3 orders of real spherical harmonics ( $l = 0, 1, 2$ ) corresponding to a total of 9 basis functions. 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 12.

transfer function—since the local upper hemisphere is rotated—which can also be thought of as different rotations of the incident light field. We will see that this integral becomes a simple product when transformed to spherical harmonics, further stressing the analogy with convolution.

### 0.3.2 Spherical Harmonic Representation

We now proceed to construct a closed-form description of the irradiance. Since we are dealing with convolutions, it is natural to analyze them in the frequency domain. Since unit vectors corresponding to the surface normal or the incident direction lie on a sphere of unit magnitude, the appropriate signal-processing tools 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 [28, 29, 42].

**Key properties of spherical harmonics:** The spherical harmonic  $Y_{lm}$  is given by (see figure 5 for an illustration)

$$\begin{aligned}
 N_{lm} &= \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \\
 Y_{lm}(\theta, \phi) &= N_{lm} P_{lm}(\cos \theta) e^{Im\phi},
 \end{aligned}
 \tag{11}$$

where  $N_{lm}$  is a normalization factor. In the above equation, the azimuthal dependence is expanded in terms of Fourier basis functions. The  $\theta$  dependence is expanded in terms of the associated Legendre functions  $P_{lm}$ . The indices obey  $l \geq 0$  and  $-l \leq m \leq l$ . Thus, there are  $2l + 1$  basis functions for given order  $l$ . They may be written either as trigonometric functions of the spherical coordinates  $\theta$  and  $\phi$  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  $Y_{l-m} = (-1)^m Y_{lm}^*$ , where  $Y_{lm}^*$  denotes the complex conjugate. The first 3 orders (we give only terms with  $m \geq 0$ ) may be written as

$$\begin{aligned}
Y_{00} &= \sqrt{\frac{1}{4\pi}} \\
Y_{10} &= \sqrt{\frac{3}{4\pi}} \cos \theta &= \sqrt{\frac{3}{4\pi}} z \\
Y_{11} &= -\sqrt{\frac{3}{8\pi}} \sin \theta e^{I\phi} &= -\sqrt{\frac{3}{8\pi}} (x + Iy) \\
Y_{20} &= \frac{1}{2} \sqrt{\frac{5}{4\pi}} (3 \cos^2 \theta - 1) &= \frac{1}{2} \sqrt{\frac{5}{4\pi}} (3z^2 - 1) \\
Y_{21} &= -\sqrt{\frac{15}{8\pi}} \sin \theta \cos \theta e^{I\phi} &= -\sqrt{\frac{15}{8\pi}} z (x + Iy) \\
Y_{22} &= \frac{1}{2} \sqrt{\frac{15}{8\pi}} \sin^2 \theta e^{2I\phi} &= \frac{1}{2} \sqrt{\frac{15}{8\pi}} (x + Iy)^2.
\end{aligned} \tag{12}$$

The spherical harmonics form an orthonormal and complete basis for functions on the unit sphere,

$$\int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} Y_{lm}^*(\theta, \phi) Y_{l'm'}(\theta, \phi) \sin \theta d\theta d\phi = \delta_{ll'} \delta_{mm'}. \tag{13}$$

To find the spherical harmonic coefficients of an arbitrary function  $f$ , we simply integrate against the complex conjugate, as with any orthonormal basis set. That is,

$$\begin{aligned}
f(\theta, \phi) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l f_{lm} Y_{lm}(\theta, \phi) \\
f_{lm} &= \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} f(\theta, \phi) Y_{lm}^*(\theta, \phi) \sin \theta d\theta d\phi.
\end{aligned} \tag{14}$$

Let us now build up the rotation operator on the spherical harmonics. Rotation about the  $z$  axis is simple, with

$$Y_{lm}(R_z(\beta)\{\theta, \phi\}) = Y_{lm}(\theta, \phi + \beta) = \exp(I m \beta) Y_{lm}(\theta, \phi). \tag{15}$$

Rotation about other axes is more complicated, and the general rotation formula can be written as

$$Y_{lm}(R_{\alpha, \beta, \gamma}(\theta, \phi)) = \sum_{m'=-l}^l D_{mm'}^l(\alpha, \beta, \gamma) Y_{lm'}(\theta, \phi). \tag{16}$$

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 16 as a matrix multiplication is found in figure 6. An analytic form for the matrices  $D^l$  can be found in standard references, such as Inui et al. [28]. In particular, since  $R_{\alpha, \beta, \gamma} = R_z(\beta) R_y(\alpha) R_z(\gamma)$ , the dependence of  $D^l$  on  $\beta$  and  $\gamma$  is simple, since rotation of the spherical harmonics about the  $z$ -axis is straightforward, i.e.

$$D_{mm'}^l(\alpha, \beta, \gamma) = d_{mm'}^l(\alpha) e^{Im\beta} e^{Im'\gamma}, \tag{17}$$

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  $d^l$ . The analytic formula is rather complicated, and is derived in






Figure 7: The solid line is a plot of  $A_l$  versus  $l$ . It can be seen that odd terms with  $l > 1$  have  $A_l = 0$ . Also, as  $l$  increases, the BRDF coefficient or transfer function rapidly decays.

**Spherical Harmonic Reflection Equation:** We can now write down the reflection equation, as given by equation 8, in terms of the expansions just defined. To do so, we multiply the expansions for the lighting and BRDF, and integrate. By orthonormality of the spherical harmonics, we require  $m' = 0$ . Hence, we obtain

$$\begin{aligned}
 E(\alpha, \beta, \gamma) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l A_l L_{lm} D_{m0}^l(\alpha, \beta, \gamma) \\
 E(\alpha, \beta) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l \sqrt{\frac{4\pi}{2l+1}} A_l L_{lm} Y_{lm}(\alpha, \beta).
 \end{aligned} \tag{23}$$

Note that, as expected, the tangential rotation  $\gamma$ , that has no physical meaning here, has vanished from the equations.

Finally, we can equate spherical harmonic coefficients of the irradiance, and use a symbol for the normalization,  $\Lambda_l = \sqrt{4\pi/(2l+1)}$ , to obtain the key equation for Lambertian reflection,

$$\boxed{E_{lm} = \Lambda_l A_l L_{lm}.} \tag{24}$$

This equation states that the standard direct illumination integral in equation 8 can be viewed as a simple product in terms of spherical harmonic coefficients. This is not really surprising<sup>3</sup>, considering that equation 8 can be interpreted as showing that the irradiance is a *convolution* of the incident illumination and the transfer function. We have simply derived the product formula for convolution in the frequency domain, analogous to the standard formula in terms of Fourier coefficients.

**Representing the Transfer Function:** The one remaining component is to explicitly find the spherical harmonic coefficients of the half-cosine or clamped-cosine transfer function  $A_l$ . The coefficients are given by

$$A_l = 2\pi \int_0^{\pi/2} Y_{l0}(\theta'_i) \cos \theta'_i \sin \theta'_i d\theta'_i, \tag{25}$$

where the factor of  $2\pi$  comes from integrating 1 over the azimuthal dependence. It is important to note that the limits of the integral range from 0 to  $\pi/2$  and not  $\pi$  because we are considering only

<sup>3</sup>Basri and Jacobs [2, 4] have noticed that this result follows directly from equation 8 by the Funk-Hecke theorem (as stated, for instance in Groemer [22], page 98). However, that theorem does not generalize to more complex materials (where the BRDF lobe is not radially symmetric, as the half-cosine function is). The derivation above enables us to easily generalize the results to arbitrary materials, as discussed later in the chapter.




Figure 8: Successive approximations to the *clamped cosine* function by adding more spherical harmonic terms. For  $l = 2$ , we already get a very good approximation.

the upper hemisphere ( $A(\theta'_i) = 0$  in the lower hemisphere). The expression above may be simplified by writing in terms of Legendre polynomials  $P(\cos \theta'_i)$ . Putting  $u = \cos \theta'_i$  in the above integral and noting that  $P_1(u) = u$  and that  $Y_{l0}(\theta'_i) = \sqrt{(2l+1)/(4\pi)}P_l(\cos \theta'_i)$ , we obtain

$$A_l = 2\pi\Lambda_l^{-1} \int_0^1 P_l(u)P_1(u) du. \quad (26)$$

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

$$\int_{-1}^1 P_a(u)P_b(u) = \frac{2}{2a+1} \delta_{a,b} \quad (27)$$

From this, we can establish some results about equation 26. When  $l$  is equal to 1, the integral evaluates to half the norm above, i.e.  $1/3$ . When  $l$  is odd but greater than 1, the integral in equation 26 vanishes. This is because for  $a = l$  and  $b = 1$ , we can break the left-hand side of equation 27 using the oddness of  $a$  and  $b$  into two equal integrals from  $[-1, 0]$  and  $[0, 1]$ . Therefore, both of these integrals must vanish, and the latter integral is the right-hand integral in equation 26. When  $l$  is even, the required formula is given by manipulating equation 20 in chapter 5 of MacRobert[42]. Putting it all together, we have,

$$\boxed{\begin{array}{ll} l = 1 & A_l = \sqrt{\frac{\pi}{3}} \\ l > 1, \text{ odd} & A_l = 0 \\ l \text{ even} & A_l = 2\pi\sqrt{\frac{2l+1}{4\pi}} \frac{(-1)^{\frac{l}{2}-1}}{(l+2)(l-1)} \left[ \frac{l!}{2^l(\frac{l}{2}!)^2} \right]. \end{array}} \quad (28)$$

We can determine the asymptotic behavior of  $A_l$  for large even  $l$  by using Stirling's formula. The bracketed term goes as  $l^{-1/2}$ , which cancels the term in the square root. Therefore, the asymptotic behavior for even terms is  $A_l \sim l^{-2}$ . A plot of  $A_l$  for the first few terms is shown in figure 7, and approximation of the *clamped cosine* by spherical harmonic terms as  $l$  increases is shown in figure 8.

Since  $A_l$  vanishes for odd values of  $l > 1$ , as seen in equation 28, the irradiance (per equation 24) has zero projection onto odd-order modes, i.e.  $E_{lm} = 0$  when  $l > 1$  and odd. In terms of the filtering

analogy in signal-processing, since the filter corresponding to  $A(\theta'_i) = \cos \theta'_i$  destroys high-frequency odd terms in the spherical harmonic expansion of the lighting, the corresponding terms are not found in the irradiance. Further, for large even  $l$ , the asymptotic behavior of  $E_{lm} \sim l^{-5/2}$  since  $A_l \sim l^{-2}$ . The transfer function  $A$  acts as a low pass filter, causing the rapid decay of high-frequency terms in the lighting.

It is instructive to explicitly write out numerically the first few terms for the irradiance per equations 24 and 28,

$$\begin{aligned}
 E_{00} &= \pi L_{00} && \approx 3.142L_{00} \\
 E_{1m} &= \frac{2\pi}{3} L_{1m} && \approx 2.094L_{1m} \\
 E_{2m} &= \frac{\pi}{4} L_{2m} && \approx 0.785L_{2m} \\
 E_{3m} &= 0 \\
 E_{4m} &= -\frac{\pi}{24} L_{4m} && \approx -0.131L_{4m} \\
 E_{5m} &= 0 \\
 E_{6m} &= \frac{\pi}{64} L_{6m} && \approx 0.049L_{6m}.
 \end{aligned} \tag{29}$$

We see that already for  $E_{4m}$ , the coefficient is only about 1% of what it is for  $E_{00}$ . Therefore, in real applications, where surfaces are only approximately Lambertian, and there are errors in measurement, we can obtain good results using an order 2 approximation of the illumination and irradiance. Since there are  $2l + 1$  indices (values of  $m$ , which ranges from  $-l$  to  $+l$ ) for order  $l$ , this corresponds to 9 coefficients for  $l \leq 2$ —1 term with order 0, 3 terms with order 1, and 5 terms with order 2. Note that the single order 0 mode  $Y_{00}$  is a constant, the 3 order 1 modes are linear functions of the Cartesian coordinates—in real form, they are simply  $x$ ,  $y$ , and  $z$ —while the 5 order 2 modes are quadratic functions of the Cartesian coordinates.

**Therefore, the irradiance—or equivalently, the reflected light field from a convex Lambertian object—can be well approximated using spherical harmonics up to order 2 (a 9 term representation), i.e. as a quadratic polynomial of the Cartesian coordinates of the surface normal vector.** This is one of the key results of this chapter. By a careful formal analysis, Basri and Jacobs [2, 4] have shown that over 99.2% of the energy of the filter is captured in orders 0,1 and 2, and 87.5% is captured even in an order 1 four term approximation. Furthermore, by enforcing non-negativity of the illumination, one can show that for *any* lighting conditions, the average accuracy of an order 2 approximation is at least 98%.

### 0.3.3 Explanation of Experimental Results

Having obtained the theoretical result connecting irradiance and radiance, we can now address the empirical results on low dimensionality we discussed in the background section. First, consider images under all possible illuminations. Since any image can be well described by a linear combination of 9 terms—the images corresponding to the 9 lowest frequency spherical harmonic lights—we expect the illumination variation to lie close to a 9-dimensional subspace. In fact, a 4D subspace consisting of the linear and constant terms itself accounts for almost 90% of the energy. This 4D subspace is exactly that of Koenderink and van Doorn [35]. Similarly, the 3D subspace of Shashua [69] can be thought of as corresponding to the lighting from order 1 spherical harmonics only. However, this neglects the ambient or order 0 term, which is quite important from our analytic results.

There remain at this point two interesting theoretical questions. First, how does the spherical harmonic convolution result connect to principal component analysis? Second, the 9D subspace predicted by the above analysis does not completely agree with the 5D subspace observed in a number of experiments [15, 23]. Recently, we have shown [60] that under appropriate assumptions, the principal components or eigenvectors are equal to the spherical harmonic basis functions and the eigenvalues, corresponding to how important a principal component is in explaining image variability, are equal to the spherical harmonic coefficients  $\Lambda_l A_l$ . Furthermore, if we see only a single image, as in previous experimental tests, only the front-facing normals, and not the entire sphere of surface normals is visible. This allows approximation by a lower-dimensional subspace. We have shown how to analyze this effect with spherical harmonics, deriving analytic forms for the principal components of canonical shapes. The results for a human face are shown in figure 9, and are in qualitative and quantitative agreement with previous empirical work, providing the first full theoretical explanation.




Figure 9: The first 5 principal components of a face, computed by our method [60] for analytically constructing the PCA for lighting variability using spherical harmonics. The form of these eigenmodes is strikingly similar to those derived empirically by previous researchers (see figure 1 in Hallinan [23]). The corresponding eigenvalues are in the ratio of .42, .33, .16, .035, .021 and are in agreement with empirical observation. The first 3 and 5 eigenmodes respectively account for 91% and 97% of the variance, compared to empirical values of 90% and 94% respectively [15]. The slight discrepancy is likely due to specularities, cast shadows, and noisy or biased measurements. The principal components contain both positive values (bright regions) and negative values (dark regions), with zero set to the neutral grey of the background. The face is an actual range scan courtesy of Cyberware.

## 0.4 Applications of Lambertian 9 Term Spherical Harmonic Model

In the previous section, we have derived an analytic formula for the irradiance, or reflected light from a Lambertian surface, as a convolution of the illumination and a low-pass filter corresponding to the clamped cosine function. Though the mathematical derivation is rather involved, the final result is very simple, as given in equations 24 and 28. This simple 9 term Lambertian model allows one to reduce the effects of very complex lighting distributions, involving multiple point and extended sources, to a simple quadratic formula. It therefore enables computer vision applications, that had so far to rely on the simple point or directional light source model, to handle arbitrary distant lighting with attached shadows. In computer graphics, lighting calculations that had previously required expensive finite element or Monte Carlo integration, can now be done with a simple analytic formula. Because of the simplicity and effectiveness of the final model, it has been widely adopted in both graphics and vision for a number of practical applications.

### 0.4.1 Computer Vision: Recognition

In computer vision, a critical problem is being able to recognize objects in a way that is lighting-insensitive. This is especially important to build robust face detection and recognition systems. In fact, face recognition was the first application of the Lambertian 9 term model (Basri and Jacobs [2]) to vision. To apply the results of the previous section, we assume the availability of 3D geometric models and albedos of the faces of interest. With the ready availability of range scanners that give detailed 3D models, this is an increasingly reasonable assumption.

Once the position of the object and pose have been estimated, the lighting-insensitive recognition problem can be stated as follows. Given a 3D geometric model and albedos for an object, determine if the query image observed can correspond to the object under some lighting condition. To solve this problem, we note that images of a Lambertian object lie close to a 9D subspace. We simply need to take the query image and project it onto this subspace, to see if it lies close enough<sup>4</sup>. Furthermore, our spherical harmonic formula gives an easy way of analytically constructing the subspace, simply by analytically rendering images of the object under the 9 spherical harmonic basis lighting conditions, using equation 24. This approach is efficient and robust, compared to numerical PCA methods. The results are excellent (recognition rates > 95% on small datasets), indicating that human faces are reasonably well approximated as Lambertian convex objects for the purpose of modeling illumination variation in recognition. Hence, the Lambertian 9 term model has significant practical relevance for robust face modeling and recognition systems.

<sup>4</sup>In practice, the object that lies closest to the subspace can be chosen—we recognize the best matching face.



Figure 10: Some of the objects which we have been able to recognize [57], using a combination of the Lambertian 9 term model, and compact light source models to account for specularities.

Recent work has addressed a number of subproblems in recognition. Lee et al. [39] show how to construct the 9D Lambertian subspace using nine images taken with point light sources, while Zhang and Samaras [82] have shown how to learn spherical harmonic subspaces for recognition. Ho et al. [26] have used the harmonic subspace for image clustering. We (Osadchy et al. [57]) have also built on the spherical harmonic representation, combining it with compact models for light sources to tackle the problem of recognizing specular and transparent objects (figure 10).

#### 0.4.2 Modeling

An understanding of lighting is also critical for many modeling problems. For instance, in computer vision, we can obtain the shape of an object to build 3D models using well known and well studied techniques like photometric stereo, stereo, or structure from motion. However, all of these methods usually make very simple assumptions about the lighting conditions, such as point light sources without shadows. Our Lambertian results enable solving these problems under realistic complex lighting, which has implications for shape modeling in uncontrolled environments.

In photometric stereo, we have a number of images under different lighting conditions (previously, these were always point sources). We seek to determine the shape (surface normals). Basri and Jacobs [3] have shown how to solve this problem for the first time under unknown complex (but distant) illumination, under the assumption of Lambertian reflectance and convex objects.

The mathematics of their solution is conceptually simple. Consider taking  $f$  images of  $p$  pixels each. One can create a large matrix  $M$  of size  $f \times p$  holding all the image information. The low-dimensional Lambertian results indicate that this matrix can be approximated as  $M \approx LS$ , where  $L$  is a  $f \times 9$  lighting matrix, that gives the lowest-frequency lighting coefficients for the  $f$  images, and  $S$  is a  $9 \times p$  matrix, that corresponds to the spherical harmonic basis images of the object under the 9 lighting conditions. Of course, once the basis images are recovered, the surface normals and albedos can be trivially estimated from the order 0 and order 1 terms, which are simply the albedos and scaled surface normals themselves. The computation of  $M \approx LS$  can be done using singular value decomposition, but is not generally unique. Indeed, there is a  $9 \times 9$  ambiguity matrix ( $4 \times 4$  if using a linear 4D subspace approximation). Basri and Jacobs [3] analyze these ambiguities in some detail. It should be noted that their general approach is quite similar to previous work on photometric stereo with point sources, such as the factorization method proposed by Hayakawa [24]. The Lambertian low-dimensional lighting model enables much of the same machinery to be used, now taking complex illumination and attached shadows into account.

In photometric stereo, we assume the scene is static, and only the lighting changes. This is not always practical, for instance when building models of human faces. A far less restrictive acquisition scenario is to directly model objects from video, where the subject is moving normally. We would like to take your home videos, and construct a 3D geometric model of your face. This problem also relates to structure from motion, a classic computer vision method. Simakov et al. [71] have recently addressed this problem. They assume the motion is known, such as from tracking features like the eyes or nose. They then seek to recover the 3D shape and albedos of the rigid moving object, as well as a spherical harmonic description of the illumination in the scene. They use the Lambertian spherical harmonic model to create a consistency measure within a stereo algorithm. Zhang et al. [81] do not assume known motion, but assume infinitesimal movement between frames.






Figure 11: The diffuse shading on all the objects is computed procedurally in real-time using the quadratic formula for Lambertian reflection [63]. The middle sphere, armadillo and table are diffuse reflectors. Our method can also be combined with standard texture mapping, used to modulate the albedo of the pool ball on the right, and reflection mapping, used for specular highlights on the pool ball and mirror sphere on the left. The environment is a light probe of the Grace Cathedral in San Francisco, courtesy Paul Debevec.

### 0.4.3 Inverse Problems: Estimating Illumination

An important theoretical and practical problem is to estimate the illumination from observations. Once we have done this, we can add synthetic objects to a real scene with realistic lighting for computer graphics, or use knowledge of the lighting to solve computer vision problems. The study of inverse problems can be seen in terms of our results as a deconvolution.

For estimating the illumination from a convex Lambertian object, our theoretical results give a simple practical solution. We just need to use equation 24, solving for lighting coefficients using  $L_{lm} = \Lambda_l^{-1} E_{lm} / A_l$ . The form of the Lambertian filter coefficients indicates that inverse lighting is ill-posed (numerator and denominator vanish) for odd illumination frequencies greater than one. These frequencies of the signal are completely annihilated by the Lambertian filter, and do not appear in the output irradiance at all. Furthermore, inverse lighting is an ill-conditioned inverse problem, because the Lambertian filter is low-pass, and the irradiance cannot easily be deconvolved. In fact, we will be able to reliably estimate only the *first 9 low-order spherical harmonic terms* of the incident lighting from images of a Lambertian surface.

These results settle a theoretical question on whether radiance and irradiance are equivalent [64], by showing formally that irradiance is not invertible, since the Lambertian kernel has zero entries. This corrects a long-standing (but incorrect) conjecture (Preisendorfer [58], vol. 2, pp 143–151) in the radiative transfer community. It also explains the ill-conditioning in inverse lighting calculations observed by previous authors like Marschner and Greenberg [44]. Finally, the results are in accord with important perceptual observations like the retinex theory [38]. Since lighting cannot produce high-frequency patterns on Lambertian surfaces, such patterns must be attributable to texture, allowing perceptual separation of illumination effects and those due to texture or albedo.

### 0.4.4 Computer Graphics: Irradiance Environment Maps for Rendering

The OpenGL standard for graphics hardware currently has native support only for point or directional light sources. The main reason is the lack of simple procedural formulas for general lighting distributions. Instead, the lighting from all sources must be summed or integrated. Our Lambertian 9 term model makes it straightforward to compute, represent and render with the irradiance environment map  $E(\alpha, \beta)$ . This was our first practical application of the Lambertian 9 term model in computer graphics [63].

First, we *prefilter* the environment map, computing the lighting coefficients  $L_{lm}$ . From this, we can

compute the 9 irradiance coefficients  $E_{lm}$  per equation 24. This computation is very efficient, being linear in the size of the environment map. We can now adopt one of two approaches. The first is to explicitly compute the irradiance map  $E(\alpha, \beta)$  by evaluating it from the irradiance coefficients. Then, for any surface normal  $(\alpha, \beta)$ , we can simply look up the irradiance value and multiply by the albedo to shade the Lambertian surface. This was the approach taken in previous work on environment mapping [21, 46]. Those methods needed to explicitly compute  $E(\alpha, \beta)$  with a hemispherical integral of all the incident illumination for each surface normal. This was a slow procedure (it is  $O(n^2)$  where  $n$  is the size of the environment map in pixels). By contrast, our spherical harmonic prefiltering method is linear time  $O(n)$ . Optimized implementations can run in near real-time, enabling our method to potentially be used for dynamically changing illumination. Software for preconvolving environment maps is available on our website and widely used by industry.

However, we can go much further than simply computing an explicit representation of  $E(\alpha, \beta)$  quickly. In fact, we can use a simple procedural formula for shading. This allows us to avoid using textures for irradiance calculations. Further, the computations can be done at each vertex, instead of at each pixel, since irradiance varies slowly. They are very simple to implement in either vertex shaders on the graphics card, or in software. All we need to do is explicitly compute

$$E(\alpha, \beta) \approx \sum_{l=0}^2 \sum_{m=-l}^l E_{lm} Y_{lm}(\alpha, \beta). \quad (30)$$

This is straightforward, because the spherical harmonics up to order 2 are simply constant, linear and quadratic polynomials. In fact, we can write an alternative simple quadratic form,

$$E(\mathbf{n}) = \mathbf{n}^t M \mathbf{n}, \quad (31)$$

where  $\mathbf{n}$  is the  $4 \times 1$  surface normal in homogeneous cartesian coordinates and  $M$  is a  $4 \times 4$  matrix. Since this involves only a matrix-vector multiplication and dot-product, and hardware is optimized for these operations with  $4 \times 4$  matrices and vectors, the calculation is very efficient. An explicit form for the matrix  $M$  is given in [63]. The image in figure 11 was rendered using this method, which is now widely adopted for real-time rendering in video games.

Note the **duality between forward and inverse problems**. The inverse problem, inverse lighting, is ill-conditioned, with high frequencies of the lighting not easy to estimate. Conversely, the forward problem, irradiance environment maps, can be computed very efficiently, by approximating only the lowest frequencies of the illumination and irradiance. This duality between forward and inverse problems has rarely been used before, and we believe it can lead to many new applications.

## 0.5 Specularities: Convolution Formula for General Materials

In this section, we discuss the extension to general materials, including specularities, and briefly explore the implications and practical applications of these general convolution results. We first briefly derive the fully general case with anisotropic materials [61, 67]. An illustrative diagram is found in figure 12. For practical applications, it is often convenient to assume isotropic materials. The convolution result for isotropic BRDFs, along with implications are developed in [65, 67], and we will quickly derive it here as a special case of the anisotropic theory.

### 0.5.1 General Convolution Formula

The reflection equation, analogous to equation 8 in the Lambertian case, now becomes

$$B(\alpha, \beta, \gamma, \theta'_o, \phi'_o) = \int_{\Omega'_i} L(R_{\alpha, \beta, \gamma}(\theta'_i, \phi'_i)) \hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) d\omega'_i, \quad (32)$$

where the reflected light field  $B$  now also depends on the (local) outgoing angles  $\theta'_o, \phi'_o$ .  $\hat{\rho}$  is a transfer function (the BRDF multiplied by the cosine of the incident angle  $\hat{\rho} = \rho \max(\cos \theta'_i, 0)$ ), which depends on local incident and outgoing angles.

The lighting can be expanded as per equations 19 and 20 in the Lambertian derivation. The expansion of the transfer function is now more general than equation 21, being given in spherical




Figure 12: Schematic of reflection in 2D. On the left, we show the situation with respect to one point on the surface (the north pole or  $0^\circ$  location, where global and local coordinates are the same). The right figure shows the effect of the surface orientation  $\alpha$ . Different orientations of the surface correspond to rotations of the upper hemisphere and BRDF, with the global incident direction  $\theta_i$  corresponding to a rotation by  $\alpha$  of the local incident direction  $\theta'_i$ . Note that we also keep the local outgoing angle (between  $N$  and  $B$ ) fixed between the two figures.

harmonics by

$$\hat{\rho}(\theta'_i, \phi'_i, \theta'_o, \phi'_o) = \sum_{l=0}^{\infty} \sum_{n=-l}^l \sum_{p=0}^{\infty} \sum_{q=-p}^p \hat{\rho}_{ln,pq} Y_{ln}^*(\theta'_i, \phi'_i) Y_{pq}(\theta'_o, \phi'_o), \quad (33)$$

where the complex conjugate for the first factor is to simplify the final results.

Finally, we need to expand the reflected light field in basis functions. For the Lambertian case, we simply used spherical harmonics over  $(\alpha, \beta)$ . In the general case, we must consider the full tangent frame  $(\alpha, \beta, \gamma)$ . Thus, the expansion is over mixed basis functions,  $D_{mn}^l(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o)$ . With the correct normalization (equation 7.73 of [28]),

$$B(\alpha, \beta, \gamma, \theta'_o, \phi'_o) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{n=-l}^l \sum_{p=0}^{\infty} \sum_{q=-p}^p B_{lmn,pq} C_{lmn,pq}(\alpha, \beta, \gamma, \theta'_o, \phi'_o)$$

$$C_{lmn,pq}(\alpha, \beta, \gamma, \theta'_o, \phi'_o) = \sqrt{\frac{2l+1}{8\pi^2}} D_{mn}^l(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o). \quad (34)$$

We can now write down the reflection equation, as given by equation 32, in terms of the expansions just defined. As in the Lambertian case, we multiply the expansions for the lighting (equation 20) and BRDF (equation 33) and integrate. To avoid confusion between the indices in this intermediate step, we will use  $L_{lm}$  and  $\hat{\rho}_{l'n,pq}$  to obtain

$$B(\alpha, \beta, \gamma, \theta'_o, \phi'_o) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{m'=-l}^l \sum_{l'=0}^{\infty} \sum_{n=-l'}^l \sum_{p=0}^{\infty} \sum_{q=-p}^p L_{lm} \hat{\rho}_{l'n,pq} D_{mm'}^l(\alpha, \beta, \gamma) Y_{pq}(\theta'_o, \phi'_o) T_{lm'l'n}$$

$$T_{lm'l'n} = \int_{\phi'_i=0}^{2\pi} \int_{\theta'_i=0}^{\pi} Y_{lm'}(\theta'_i, \phi'_i) Y_{l'n}^*(\theta'_i, \phi'_i) \sin \theta'_i d\theta'_i d\phi'_i$$

$$= \delta_{ll'} \delta_{m'n}. \quad (35)$$

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

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

Finally, equating coefficients, we obtain the spherical harmonic reflection equation or convolution formula, analogous to equation 24. The reflected light field can be viewed as taking the incident

illumination signal, and filtering it with the material properties or BRDF of the surface,

$$B_{lmn,pq} = \sqrt{\frac{8\pi^2}{2l+1}} L_{lm} \hat{\rho}_{ln,pq}. \quad (37)$$

**Isotropic BRDFs:** An important special case are isotropic BRDFs, where the orientation of the local tangent frame  $\gamma$  does not matter. To consider the simplifications that result from isotropy, we first analyze the BRDF coefficients  $\hat{\rho}_{ln,pq}$ . In the BRDF expansion of equation 33, only terms that satisfy isotropy, i.e. are invariant with respect to adding an angle  $\Delta\phi$  to both incident and outgoing azimuthal angles, are nonzero. From the form of the spherical harmonics, this requires that  $n = q$  and we write the now 3D BRDF coefficients as  $\hat{\rho}_{lpq} = \hat{\rho}_{lq,pq}$ .

Next, we remove the dependence of the reflected light field on  $\gamma$  by arbitrarily setting  $\gamma = 0$ . The reflected light field can now be expanded in coefficients [65],

$$\begin{aligned} B(\alpha, \beta, \theta'_o, \phi'_o) &= \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{p=0}^{\infty} \sum_{q=-\min(l,p)}^{\min(l,p)} B_{lmpq} C_{lmpq}(\alpha, \beta, \theta'_o, \phi'_o) \\ C_{lmpq}(\alpha, \beta, \theta'_o, \phi'_o) &= \sqrt{\frac{2l+1}{4\pi}} D_{mq}^l(\alpha, \beta, 0) Y_{pq}(\theta'_o, \phi'_o). \end{aligned} \quad (38)$$

Finally, the convolution formula for isotropic BRDFs becomes

$$B_{lmpq} = \Lambda_l L_{lm} \hat{\rho}_{lpq}. \quad (39)$$

Finally, it is instructive to derive the Lambertian results as a special case of the above formulation. In this case,  $p = q = 0$  (no exitant angular dependence), and equation 24 is effectively obtained simply by dropping the indices  $p$  and  $q$  in the above equation.

### 0.5.2 Implications

We now briefly discuss the theoretical implications of these results. A more detailed discussion can be found in [65, 67]. First, we consider inverse problems like illuminant and BRDF estimation. From the formula above, we can solve them simply by dividing the reflected light field coefficients by the known lighting or BRDF coefficients, i.e.  $L_{lm} = \Lambda_l^{-1} B_{lmpq} / \hat{\rho}_{lpq}$  and  $\hat{\rho}_{lpq} = \Lambda_l^{-1} B_{lmpq} / L_{lm}$ . These problems will be well conditioned when the denominators do not vanish and contain high frequency elements. In other words, for inverse lighting to be well conditioned, the BRDF should be an all-pass filter, ideally a mirror surface. Mirrors are delta functions in the spatial or angular domain, and include all frequencies in the spherical harmonic domain. Estimating the illumination from a chrome-steel or mirror sphere is in fact a common approach. On the other hand, if the filter is low-pass, like a Lambertian surface, inverse lighting will be ill-conditioned. Similarly, for BRDF estimation to be well conditioned, the lighting should include high frequencies like sharp edges. The ideal is a directional or point light source. In fact, Lu et al. [41] and Marschner et al. [43] have developed image-based BRDF measurement methods from images of a homogeneous sphere lit by a point source. More recently, Gardner et al. [17] have used a linear light source. On the other hand, BRDF estimation is ill-conditioned under soft diffuse lighting. On a cloudy day, we will not be able to estimate the widths of specular highlights accurately, since they will be blurred out.

Interesting special cases are Lambertian, and Phong and Torrance-Sparrow [77] BRDFs. We have already discussed the Lambertian case in some detail. The frequency domain filters corresponding to Phong or microfacet models are approximately Gaussian [65]. In fact, the width of the Gaussian in the frequency domain is inversely related to the width of the filter in the spatial or angular domain. The extreme cases are that of a mirror surface, that is completely localized (delta function) in the angular domain and passes through all frequencies (infinite frequency width), and of a very rough or near-Lambertian surface that has full width in the angular domain, but is a very compact low-pass filter in the frequency domain. This duality between angular and frequency-space representations is similar to that for Fourier basis functions. In some ways, it is also analogous to the Heisenberg uncertainty principle—we cannot simultaneously localize a function in both angular and frequency

domains. For representation and rendering however, we can turn this to our advantage, using *either* the frequency or angular domain, wherever the different components of lighting and material properties can be more efficiently and compactly represented.

Another interesting theoretical question concerns *factorization* of the 4D reflected light field, simultaneously estimating the illumination (2D) and isotropic BRDF (3D). Since the lighting and BRDF are lower-dimensional entities, it seems that we have redundant information in the reflected light field. We show [65] that an analytic formula for both illumination and BRDF can be derived from reflected light field coefficients. Thus, up to a global scale and possible ill-conditioning, the reflected light field can in theory be factored into illumination and BRDF components.

### 0.5.3 Applications

**Inverse Rendering:** Inverse rendering refers to inverting the normal rendering process to estimate illumination and materials from photographs. Many previous inverse rendering algorithms have been limited to controlled laboratory conditions using point light sources. The insights from our convolution analysis have enabled the development of a general theoretical framework for handling complex illumination. We [61, 65] have developed a number of frequency domain and dual angular and frequency-domain algorithms for inverse rendering under complex illumination, including simultaneous estimation of the lighting and BRDFs from a small number of images.

**Forward Rendering (Environment Maps with Arbitrary BRDFs):** In the forward problem, all of the information on lighting and BRDFs is available, so we can directly apply the convolution formula [66]. We have shown how this can be applied to rendering scenes with arbitrary distant illumination (environment maps) and BRDFs. As opposed to the Lambertian case, we no longer have a simple 9 term formula. However, just as in the Lambertian case, convolution makes prefiltering the environment extremely efficient, three to four orders of magnitude faster than previous work. Furthermore, it allows us to use a new compact and efficient representation for rendering, the Spherical Harmonic Reflection Map or SHRM. The frequency domain analysis helps to precisely set the number of terms and sampling rates and frequencies used for a given accuracy.

**Material Representation and Recognition:** Nillius and Eklundh [54, 55] derive a low-dimensional representation using PCA for images with complex lighting and materials. This can be seen as a generalization of the PCA results for diffuse surfaces (to which they have also made an important contribution [53]). They can use these results to recognize or classify the materials in an image, based on their reflectance properties.

## 0.6 Relaxing and Broadening the Assumptions: Recent Work

So far, we have seen how to model illumination, and the resulting variation in appearance of objects using spherical harmonics. There remain a number of assumptions as far as the theoretical convolution analysis is concerned. Specifically, we assume distant lighting and convex objects without cast shadows or interreflections. Furthermore, we are dealing only with opaque objects in free space, and are not handling translucency or participating media. It has not so far been possible to extend the theoretical analysis to relax these assumptions. However, a large body of recent theoretical and practical work has indicated that the key concepts from the earlier analysis—such as convolution, filtering and spherical harmonics—can be used to analyze and derive new algorithms for all of the effects noted above. The convolution formulas presented earlier in the chapter will of course no longer hold, but the insights can be seen to be much more broadly applicable.

**Cast Shadows:** For certain special cases, it is possible to derive a convolution result for cast shadows (though not in terms of spherical harmonics). For parallel planes, Soler and Sillion [75] have derived a convolution result, using it as a basis for a fast algorithm for soft shadows. We [68] have developed a similar result for V-grooves to explain lighting variability in natural 3D textures.

An alternative approach is to incorporate cast shadows in the integrand of the reflection equation. The integral then includes three terms, instead of two—the illumination, the BRDF (or Lambertian half-cosine function), and a binary term for cast shadows. Since cast shadows vary spatially, reflection is no longer a convolution in the general case. However, it is still possible to use spherical harmonics

to represent the terms of the integrand. Three term or triple product integrals can be analyzed using the Clebsch-Gordan series for spherical harmonics [28], that is common for analyzing angular momentum in quantum mechanics. A start at working out the theory is made by Thornber and Jacobs [76]. Most recently, we [51] have investigated generalizations of the Clebsch-Gordan series for other bases, like Haar wavelets, in the context of real-time rendering with cast shadows.

**Pre-Computed Light Transport:** Sloan et al. [74] recently introduced a new approach to real-time rendering called precomputed radiance transfer. Their method uses low-frequency spherical harmonic representations (usually order 4 or 25 terms) of the illumination and BRDF, and allows for a number of complex light-transport effects including soft shadows, interreflections, global illumination effects like caustics, and dynamically changing lighting and viewpoint. Because of the high level of realism these approaches provide compared to standard hardware rendering, spherical harmonic lighting is now widely used in interactive applications like video games, and is incorporated into Microsoft’s DirectX API.

We introduce the theory only for Lambertian surfaces—relatively straightforward extensions are possible for complex materials, including view-dependence. The basic idea is to define a transport operator  $T$  that includes half-cosine, visibility and interreflection effects. We can then write, in analogy with equations 3 and 6,

$$E(\mathbf{x}) = \int_{\Omega_i} L(\boldsymbol{\omega}_i)T(\mathbf{x}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \quad (40)$$

where  $T$  denotes the irradiance at  $\mathbf{x}$  due to illumination from direction  $\boldsymbol{\omega}_i$ , and incorporates all light-transport effects including cast shadows and interreflections. An offline precomputation step (that may involve ray-tracing for instance) is used to determine  $T$  for a static scene. We can expand all quantities in spherical harmonics, to write

$$E(\mathbf{x}) \approx \sum_{l=0}^{l^*} \sum_{m=-l}^l T_{lm}(\mathbf{x})L_{lm}, \quad (41)$$

where  $l^*$  denotes the maximum order used for calculations (2 would be sufficient in the absence of cast shadows). We repeat this calculation for each vertex  $\mathbf{x}$  for rendering. The expression above is simply a dot product which can be implemented efficiently in current programmable graphics hardware, allowing real-time manipulation of illumination. Somewhat more complicated, but conceptually similar expressions, are available for glossy materials. For real-time rendering with varying lighting and viewpoint, Sloan et al. [73] precompute a view-dependent operator (a  $25 \times 25$  matrix of spherical harmonic coefficients at each vertex) and compress the information over the surface using clustered principal component analysis.

**New Basis Functions:** Spherical harmonics are the ideal signal-processing representation when we have a convolution formula or want to analyze functions in the frequency domain. However, they are not the best representation for “all-frequency” effects—an infinite number of spherical harmonics will be needed to accurately represent a point source or delta function. Furthermore, some quantities like the BRDF are better defined over the hemisphere rather than the full sphere. In cases where cast shadows and interreflections are involved, we no longer have a convolution formula, so spherical harmonics are sometimes not the preferred representation. Recently, we [50, 51] have introduced all-frequency relighting methods using Haar wavelets. These methods can be orders of magnitude more efficient than previous precomputed light transport techniques using spherical harmonics. Hemispherical basis functions for BRDF representation have been proposed by Gautron et al. [18] and Koenderink and van Doorn [36].

**Near-Field Illumination:** Recent theoretical work [16] has tried to remove the limitation of distant illumination, by showing a convolution result (with a different kernel) to also hold for near-field lighting under specific assumptions for Lambertian surfaces. A practical solution using spherical harmonic gradients for rendering is proposed by Annen et al. [1].

**Translucency and Participating Media:** Like most of computer vision and graphics, we have so far assumed clear-day conditions, without translucent materials like marble and skin, or participating media like milk, smoke, clouds or fog. Both spherical harmonics and convolution are popular tools for subsurface scattering, volumetric rendering and radiative transport problems. Practical methods that conceptually use convolution of the incident irradiance with a point-spread function have been employed successfully for subsurface scattering [30, 31], and we have recently developed such methods for rendering multiple scattering effects in volumetric participating media [49, 59].

## 0.7 Conclusion

Much of the richness of the visual appearance of our world arises from the effects of illumination, and the varying appearance of surfaces with lighting. In this chapter, we have seen that there are many situations in the real world, such as for diffuse Lambertian surfaces, where the appearance of an object can be relatively simple to describe, even if the lighting is complicated. We have introduced new tools to model illumination and appearance using spherical harmonics. With these tools, we can derive a signal-processing approach to reflection, where the reflected light can be thought of as obtained by filtering or convolving the incident illumination signal by the material properties of the surface. This insight leads to a number of new results, and explanations of previous experimental work, such as the 9 parameter low-dimensional subspace for Lambertian objects. There are implications for face modeling, recognition and rendering, as well as a number of other problems in computer vision, graphics and perception. Already, a number of new algorithms have been developed in computer vision for recognition and modeling, and in computer graphics for both forward and inverse problems in rendering, with the promise of many further developments in the near future.

**Acknowledgements:** Special thanks to all my co-authors and collaborators, with particular gratitude to Pat Hanrahan, David Jacobs, Ronen Basri and Ren Ng. This work was supported in part by grants #0305322, #0430258 and #0446916 from the National Science Foundation and from Intel Corporation.





# Bibliography

- [1] T. Annen, J. Kautz, F. Durand, and H. Seidel. Spherical harmonic gradients for mid-range illumination. In *EuroGraphics Symposium on Rendering*, 2004.
- [2] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. In *International Conference on Computer Vision*, pages 383–390, 2001.
- [3] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. In *CVPR 01*, pages II-374–II-381, 2001.
- [4] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, 2003.
- [5] R. Basri and D. Jacobs. *Illumination Modeling for Face Recognition*, chapter 5, Face Recognition Handbook. Springer Verlag, 2004.
- [6] 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.
- [7] J. Blinn and M. Newell. Texture and reflection in computer generated images. *Communications of the ACM*, 19:542–546, 1976.
- [8] R. Brunelli and T. Poggio. Face recognition: Features versus templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(10):1042–1062, 1993.
- [9] B. Cabral, N. Max, and R. Springmeyer. Bidirectional reflection functions from surface bump maps. In *SIGGRAPH 87*, pages 273–281, 1987.
- [10] B. Cabral, M. Olano, and P. Nemeč. Reflection space image based rendering. In *SIGGRAPH 99*, pages 165–170, 1999.
- [11] H. Chen, P. Belhumeur, and D. Jacobs. In search of illumination invariants. In *CVPR 00*, pages 254–261, 2000.
- [12] M. Cohen and J. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press, 1993.
- [13] 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.
- [14] M. D’Zmura. *Computational Models of Visual Processing*, chapter Shading Ambiguity: Reflectance and Illumination, pages 187–207. MIT Press, 1991.
- [15] 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.
- [16] D. Frolova, D. Simakov, and R. Basri. Accuracy of spherical harmonic approximations for images of lambertian objects under far and near lighting. In *ECCV*, pages I-574–I-587, 2004.
- [17] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec. Linear light source reflectometry. *ACM TOG (SIGGRAPH 2003)*, 22(3):749–758, 2003.
- [18] P. Gautron, J. Krivanek, S. Pattanaik, and K. Bouatouch. A novel hemispherical basis for accurate and efficient rendering. In *EuroGraphics Symposium on Rendering*, 2004.
- [19] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Generative models for recognition under variable pose and illumination. In *Fourth International Conference on Automatic Face and Gesture Recognition*, pages 277–284, 2000.
- [20] A. Georghiades, D. Kriegman, and P. Belhumeur. Illumination cones for recognition under variable lighting: Faces. In *CVPR 98*, pages 52–59, 1998.
- [21] N. Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics & Applications*, 6(11):21–29, 1986.

- [22] H. Groemer. *Geometric applications of Fourier series and spherical harmonics*. Cambridge University Press, 1996.
- [23] P.W. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *CVPR 94*, pages 995–999, 1994.
- [24] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *Journal of the Optical Society of America A*, 11(11):3079–3089, Nov 1994.
- [25] W. Heidrich and H. P. Seidel. Realistic, hardware-accelerated shading and lighting. In *SIGGRAPH 99*, pages 171–178, 1999.
- [26] J. Ho, M. Yang, J. Lim, K. Lee, and D. Kriegman. Clustering appearances of objects under varying illumination conditions. In *CVPR*, volume 1, pages 11–18, 2003.
- [27] B. Horn. *Robot Vision*. MIT Press, 1986.
- [28] T. Inui, Y. Tanabe, and Y. Onodera. *Group theory and its applications in physics*. Springer Verlag, 1990.
- [29] J. Jackson. *Classical Electrodynamics*. John Wiley, 1975.
- [30] H. Jensen and J. Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics (SIGGRAPH 2002)*, 21(3):576–581, 2002.
- [31] H. Jensen, S. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *SIGGRAPH 2001*, pages 511–518, 2001.
- [32] J. Kautz and M. McCool. Approximation of glossy reflection with prefiltered environment maps. In *Graphics Interface*, pages 119–126, 2000.
- [33] J. Kautz, P. Vázquez, W. Heidrich, and H.P. Seidel. A unified approach to prefiltered environment maps. In *Eurographics Rendering Workshop 00*, pages 185–196, 2000.
- [34] M. Kirby and L. Sirovich. Application of the Karhunen-Loeve procedure for the characterization of human faces. *PAMI*, 12(1):103–108, Jan 1990.
- [35] J. Koenderink and A. van Doorn. The generic bilinear calibration-estimation problem. *IJCV*, 23(3):217–234, 1997.
- [36] J. Koenderink and A. van Doorn. Phenomenological description of bidirectional surface reflection. *JOSA A*, 15(11):2903–2912, 1998.
- [37] M. Lades, J. Vorbruggen, J. Buhmann, J. Lange, C. von der Malsburg, R. Wurtz, and W. Konen. Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers*, 42(3):300–311, 1992.
- [38] E. Land and J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.
- [39] K. Lee, J. Ho, and D. Kriegman. Nine points of light: Acquiring subspaces for face recognition under variable lighting. In *CVPR*, pages 519–526, 2001.
- [40] A. Levin and A. Shashua. Principal component analysis over continuous subspaces and intersection of half-spaces. In *To appear, ECCV 02*, 2002.
- [41] 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.
- [42] T. MacRobert. *Spherical harmonics; an elementary treatise on harmonic functions, with applications*. Dover Publications, 1948.
- [43] S. Marschner, S. Westin, E. Lafortune, and K. Torrance. Image-Based BRDF measurement. *Applied Optics*, 39(16):2592–2600, 2000.
- [44] S.R. Marschner and D.P. Greenberg. Inverse lighting for photography. In *Fifth Color Imaging Conference*, pages 262–265, 1997.
- [45] R. McCluney. *Introduction to Radiometry and Photometry*. Artech House, 1994.
- [46] 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.
- [47] Y. Moses, Y. Adini, and S. Ullman. Face recognition: the problem of compensating for changes in illumination direction. In *ECCV 94*, pages 286–296, 1994.

- [48] H. Murase and S. Nayar. Visual learning and recognition of 3-D objects from appearance. *IJCV*, 14(1):5–24, 1995.
- [49] S. Narasimhan, R. Ramamoorthi, and S. Nayar. Analytic rendering of multiple scattering in participating media. *Submitted to ACM Transactions on Graphics*, 2004.
- [50] R. Ng, R. Ramamoorthi, and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics (SIGGRAPH 2003)*, 22(3), 2003.
- [51] R. Ng, R. Ramamoorthi, and P. Hanrahan. Triple product wavelet integrals for all-frequency relighting. *ACM Transactions on Graphics (SIGGRAPH 2004)*, 23(3):475–485, 2004.
- [52] 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.
- [53] P. Nillius and J. Eklundh. Low-dimensional representations of shaded surfaces under varying illumination. In *CVPR 03*, pages II:185–II:192, 2003.
- [54] P. Nillius and J. Eklundh. Phenomenological eigenfunctions for irradiance. In *ICCV 03*, pages I:568–I:575, 2003.
- [55] P. Nillius and J. Eklundh. Classifying materials from their reflectance properties. In *ECCV 04*, pages IV–366–IV–376, 2004.
- [56] J. Nimeroff, E. Simoncelli, and J. Dorsey. Efficient re-rendering of naturally illuminated environments. In *Eurographics Workshop on Rendering 94*, pages 359–373, June 1994.
- [57] M. Osadchy, D. Jacobs, and R. Ramamoorthi. Using specularities for recognition. In *ICCV*, pages 1512–1519, 2003.
- [58] R.W. Preisendorfer. *Hydrologic Optics*. US Dept Commerce, 1976.
- [59] S. Premoze, M. Ashikhmin, R. Ramamoorthi, and S. Nayar. Practical rendering of multiple scattering effects in participating media. In *EuroGraphics Symposium on Rendering*, 2004.
- [60] R. Ramamoorthi. Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 24(10):1322–1333, Oct 2002.
- [61] R. Ramamoorthi. *A Signal-Processing Framework for Forward and Inverse Rendering*. PhD thesis, Stanford University, 2002.
- [62] 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.
- [63] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *SIGGRAPH 01*, pages 497–500, 2001.
- [64] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *Journal of the Optical Society of America A*, 18(10):2448–2459, 2001.
- [65] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, pages 117–128, 2001.
- [66] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. *ACM Transactions on Graphics (SIGGRAPH 02 proceedings)*, 21(3):517–526, 2002.
- [67] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for reflection. *ACM Transactions on Graphics*, 23(4):1004–1042, 2004.
- [68] R. Ramamoorthi, M. Koudelka, and P. Belhumeur. A fourier theory for cast shadows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(2):288–295, 2005.
- [69] A. Shashua. On photometric issues in 3D visual recognition from a single 2D image. *IJCV*, 21:99–122, 1997.
- [70] F. Sillion, J. Arvo, S. Westin, and D. Greenberg. A global illumination solution for general reflectance distributions. In *SIGGRAPH 91*, pages 187–196, 1991.
- [71] D. Simakov, D. Frolova, and R. Basri. Dense shape reconstruction of a moving object under arbitrary, unknown lighting. In *ICCV 03*, pages 1202–1209, 2003.
- [72] L. Sirovich and M. Kirby. Low-dimensional procedure for the characterization of human faces. *JOSA A*, 4(3):519–524, Mar 1987.

- [73] P. Sloan, J. Hall, J. Hart, and J. Snyder. Clustered principal components for precomputed radiance transfer. *ACM Transactions on Graphics (SIGGRAPH 03 proceedings)*, 22(3), 2002.
- [74] P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics (SIGGRAPH 02 proceedings)*, 21(3):527–536, 2002.
- [75] C. Soler and F. Sillion. Fast calculation of soft shadow textures using convolution. In *SIGGRAPH 98*, pages 321–332, 1998.
- [76] K. Thornber and D. Jacobs. Broadened, specular reflection and linear subspaces. Technical Report TR#2001-033, NEC, 2001.
- [77] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967.
- [78] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–96, 1991.
- [79] S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92*, pages 255–264, 1992.
- [80] A. Yuille, D. Snow, R. Epstein, and P. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using SVD and integrability. *IJCV*, 35(3):203–222, 1999.
- [81] L. Zhang, B. Curless, A. Hertzmann, and S. Seitz. Shape and motion under varying illumination: Unifying multiview stereo, photometric stereo, and structure from motion. In *International Conference on Computer Vision*, pages 618–625, 2003.
- [82] L. Zhang and D. Samaras. Face recognition under variable lighting using harmonic image exemplars. In *CVPR*, pages I:19–I:25, 2003.
- [83] L. Zhao and Y. Yang. Theoretical analysis of illumination in PCA-based vision systems. *Pattern Recognition*, 32:547–564, 1999.