

A Fourier Theory for Cast Shadows

Ravi Ramamoorthi, Melissa Koudelka, and
Peter Belhumeur

Abstract—Cast shadows can be significant in many computer vision applications, such as lighting-insensitive recognition and surface reconstruction. Nevertheless, most algorithms neglect them, primarily because they involve nonlocal interactions in nonconvex regions, making formal analysis difficult. However, many real instances map closely to canonical configurations like a wall, a V-groove type structure, or a pitted surface. In particular, we experiment with 3D textures like moss, gravel, and a kitchen sponge, whose surfaces include canonical configurations like V-grooves. This paper takes a first step toward a formal analysis of cast shadows, showing theoretically that many configurations can be mathematically analyzed using convolutions and Fourier basis functions. Our analysis exposes the mathematical convolution structure of cast shadows and shows strong connections to recent signal-processing frameworks for reflection and illumination.

Index Terms—Cast shadows, convolution, Fourier analysis, eigenmodes, V-grooves.

1 INTRODUCTION

CAST shadows are an important feature of appearance. For instance, buildings may cause the sun to cast shadows on the ground, the nose can cast a shadow onto the face, and local concavities in rough surfaces or textures can lead to interesting shadowing effects. However, most current vision algorithms do not explicitly consider cast shadows. The primary reason is the difficulty in formally analyzing them since cast shadows involve nonlocal interactions in concave regions. In general, shadowing can be very complicated, such as sunlight passing through the leaves of a tree, and mathematical analysis seems hopeless. However, we believe many common shadowing situations have simpler structures, some of which are illustrated in Fig. 1.

Our theory is motivated by some surprising practical results. In particular, we focus on the appearance of natural 3D textures like moss, gravel, and kitchen sponge, shown in Fig. 2 and Fig. 6. These objects have fine-scale structures similar to the canonical configurations shown in Fig. 1. Hence, they exhibit interesting illumination and view-dependence, which is often described using a bidirectional texture function (BTF) [3]. In this paper, we analyze lighting variability, assuming a fixed view. Since these surfaces are nearly flat and diffuse, one might expect illumination variation to correspond to simple Lambertian cosine-dependence. However, cast shadows play a major role, leading to effects that are quantitatively described and mathematically explained here.

We show that, in many canonical cases, cast shadows have a simple convolution structure amenable to Fourier analysis. This indicates a strong link between the mathematical properties of visibility and those of reflection and illumination (but ignoring cast shadows) for which Basri and Jacobs [1] and Ramamoorthi and Hanrahan [15], [16] have recently derived signal-processing frameworks. In particular, they [1], [15] show that the irradiance

is a *convolution* of the lighting and the *clamped cosine* Lambertian reflection function. We derive an analogous result for cast shadows as convolution of the lighting with a *Heaviside step* function. Our results also generalize Soler and Sillion's [20] convolution result for shadows when source, blocker, and receiver are all in parallel planes—for instance, V-grooves (Fig. 1b, as well as Fig. 1a and Fig. 1d) do not contain any parallel planes.

Our paper builds on a rich history of previous work on reflection models, such as Oren and Nayar [12], Torrance and Sparrow [23], and Koenderink et al. [7], as well as several recent articles on the properties of 3D textures [2], [21]. Our analytic formulae are derived considering the standard V-grooves used in many of these previous reflection models [12], [23]. Note that many of these models include a complete analysis of visibility in V-grooves or similar structures for any single light source direction. We differ in considering cast shadows because of complex illumination, deriving a convolution framework, and analyzing the *eigenstructure* of visibility. Our work also differs from that of Thornber and Jacobs [22] in that we derive a particular convolution formula for shadows directly, rather than considering a general integral involving the illumination, reflective properties of the surface, and visibility.

2 THE STRUCTURE OF CAST SHADOWS

First, we briefly make some theoretical observations. Consider Fig. 1a and Fig. 1b. There is a *single extreme point* B . As we move from O to A' to A (with the extremal rays being OB , $A'B$, and AB), the visible region of the illumination *monotonically increases*. This local shadowing situation, with a single extreme point B , and monotonic variation of the visible region of the illumination as one moves along the surface is one of the main ideas in our derivation. Furthermore, multiple extreme points or blockers can often be handled independently. For instance, in Fig. 1c and Fig. 1d, we have two extreme points, B and C . The net shadowing effect is essentially the superposition of the effects of extreme rays through B and C .

Second, we describe some new experimental results on the variability of appearance in 3D textures with illumination, a major component of which are cast shadowing interactions similar to the canonical examples in Fig. 1. In Fig. 2, we show an initial experiment. We illuminated a sample of gravel along an arc (with the angle ranging from -90 degrees to $+64$ degrees, limited by the specifics of the acquisition). The varying appearance with illumination clearly suggests that cast shadows are an important visual feature. The figure also shows a conceptual diagrammatic representation of the profile of a cross-section of the surface with many points shadowed in a manner similar to Fig. 1a, Fig. 1b, and Fig. 1d.

3 2D ANALYSIS OF CAST SHADOWS

For mathematical analysis, we begin in flatland, i.e., a 2D slice through the viewpoint. We will consider a V-groove model, shown in Fig. 3, that corresponds to Fig. 1b. However, the derivation will be similar for any other shadowing situation, such as those in Fig. 1, where the visibility is locally *monotonically changing*. The V-groove model in Fig. 3 can model Fig. 1a and Fig. 1b ($\beta_1 = 0$, $\beta_2 = \pi/2$, and $\beta_1 = \beta_2$), and each of the extreme points of Fig. 1c and Fig. 1d.

3.1 Convolution Formula for Shadows in a V-groove

Our goal is to find the irradiance¹ $E(x, \beta)$ as a function of groove angle $\beta = [-\beta_1, +\beta_2]$ and the distance along the groove x . Without loss of generality, we consider the right side of the groove only. For

1. Since we focus on cast shadows, we will assume Lambertian surfaces and will neglect the incident cosine term. This cosine term may be folded into the illumination function if desired, as the surface normal over a particular face (side) of the V-groove is constant.

- R. Ramamoorthi and P. Belhumeur are with the Computer Science Department, Columbia University, 450 Computer Science Building, 500 W. 120 St., New York, NY 10027. E-mail: ravir@cs.columbia.edu.
- M. Koudelka is with Yale University, 140 Mill St. #423, East Haven, CT 06512. E-mail: melissa.koudelka@yale.edu.

Manuscript received 31 Jan. 2004; revised 25 Aug. 2004; accepted 1 Sept. 2004; published online 13 Dec. 2004.

Recommended for acceptance by S. Seitz.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0059-0104.




Fig. 1. Four common shadowing situations. We show that these all have similar structures, amenable to treatment using convolution and Fourier analysis. The red lines indicate extremal rays, corresponding to shadow boundaries for distant light sources.




Fig. 2. (a) Gravel texture, which exhibits strong shadowing. (b) Images with different light directions clearly show cast shadow appearance effects, especially at large angles. The light directions correspond to the red marks in (d). (c) Conceptual representation of a profile of a cross section through surface (drawn in black in (a)). (d) Schematic of experimental setup.

a particular groove (fixed β), pixels in a single image correspond directly to different values of x and the irradiance $E(x)$ is directly proportional to pixel brightness.

$$E(x, \beta) = \int_{-\pi/2}^{\pi/2} L(\omega) V(x, \omega, \beta) d\omega, \quad (1)$$

where $L(\omega)$ is the incident illumination intensity, which is a function of the incident direction ω . We make no restrictions on the lighting, except that it is assumed distant, so the angle ω does not depend on location x . This is a standard assumption in environment map rendering in graphics and has been used in previous derivations of analytic convolution formulae [1], [16]. V is the binary visibility in direction ω at location $A(x)$.

Monotonic Variation of Visibility: As per the geometry in Fig. 3, the visibility is 1 in the range from $-\beta_1$ to $\beta_2 + \alpha(x)$ and 0 or (cast) shadowed otherwise. It is important to note that $\alpha(x)$ is a *monotonically increasing* function of x , i.e., the portion of the illumination visible increases as one moves along the right side of




Fig. 3. Diagram of a V-groove with the groove angle ranging from $-\beta_1$ to $+\beta_2$.

the groove from O to A' to A (with corresponding extremal rays OB , $A'B$, and AB).

Reparameterization by α : We now simply use α to parameterize the V-groove. This is just a change of variables and is valid as long as α *monotonically varies* with x . Locally, α is always proportional to x since we may do a local Taylor series expansion, keeping only the first or linear term.

Representation of Visibility: We may now write down the function $V(x, \omega, \beta)$, newly reparameterized as $V(\alpha, \omega, \beta)$. Noting that V is 1 only in the range from $[-\beta_1, \beta_2 + \alpha]$,

$$V(\alpha, \omega, \beta) = H(-\beta_1 - \omega) - H((\beta_2 + \alpha) - \omega), \quad (2)$$

$$H(u) = 1 \text{ if } u < 0, \quad 0 \text{ if } u > 0,$$

where $H(u)$ is the Heaviside step function. The first term on the right hand side zeros the visibility when $\omega < -\beta_1$ and the second term when $\omega > \beta_2 + \alpha$. Fig. 4 illustrates this diagrammatically. In the limit of a perfectly flat Lambertian surface, $\beta_1 = \beta_2 = \pi/2$ and $\alpha = 0$. In that case, the first term on the right of (2) is always 1, the second term is 0, and $V = 1$ (no cast shadowing).

Convolution Formula: Plugging (2) back into (1), we obtain

$$E(\alpha, \beta) = \int_{-\pi/2}^{\pi/2} L(\omega) H(-\beta_1 - \omega) d\omega - \int_{-\pi/2}^{\pi/2} L(\omega) H((\beta_2 + \alpha) - \omega) d\omega. \quad (3)$$

E is the sum of two terms, the first of which depends only on groove angle β_1 and the second of which also depends on groove location or image position α . In the limit of a flat diffuse surface, the second term vanishes, while the first corresponds to convolution with unity and is simply the (unshadowed) irradiance or




Fig. 4. Illustration of the visibility function as per (2). The black portions of the graphs where $\omega < +\beta_2$ are independent of α or groove location, while the red portions with $\alpha > +\beta_2$ vary linearly with α , leading to the convolution structure.




Fig. 5. (a) Comparison of Fourier coefficients for the Heaviside step function for cast shadows and (b) the clamped cosine Lambertian function for irradiance. (c) is a loglog plot of the absolute values of the nonzero eigenvalues. The graphs are straight lines with slope -1 for cast shadows, compared to the quadratic decay (slope -2) for irradiance.

integral of the illumination. We now separate the two terms to simplify this result as ($\otimes$ is the convolution operator)

$$E(\alpha, \beta) = \tilde{E}(-\beta_1) - \tilde{E}(\beta_2 + \alpha)$$

$$\tilde{E}(u) = \int_{-\pi/2}^{\pi/2} L(\omega) H(u - \omega) d\omega = L \otimes H. \quad (4)$$

Fourier Analysis: Equation (4) makes clear that the net visibility or irradiance is a simple *convolution* of the incident illumination with the Heaviside step function that accounts for cast shadow effects. This is our main analytic result, deriving a new convolution formula that sheds theoretical insight on the structure of cast shadows. It is therefore natural to also derive a product formula in the Fourier or frequency domain,

$$\tilde{E}_k = \sqrt{\pi} L_k H_k, \quad (5)$$

where L_k are the Fourier illumination coefficients and H_k are Fourier coefficients of the Heaviside step function, plotted in Fig. 5. The analytic formula can be derived as follows:

$$H_k = \int_{-\pi/2}^0 \frac{1}{\sqrt{\pi}} e^{-2ikv} dv, \quad (6)$$

where the integral limits are set because we are considering the heaviside function and the integrand is the orthonormal Fourier basis on $[-\pi, +\pi]$.

The even coefficients H_{2k} vanish, while the odd coefficients decay as $1/k$:

$$k = 0 : H_0 = \frac{\sqrt{\pi}}{2}$$

$$\text{odd } k : H_k = \frac{i}{\sqrt{\pi} k}. \quad (7)$$

3.2 Eigenvalue Spectrum and Illumination Eigenmodes for Cast Shadows

Our convolution formula is conceptually quite similar to the convolution formula and signal-processing analysis done for convex curved Lambertian surfaces or irradiance by Basri and Jacobs [1] and Ramamoorthi and Hanrahan [14], [15]. In this subsection, we analyze

our results further in terms of the illumination eigenmodes that indicate the lighting distributions that have the most effect and the corresponding eigenvalues, or singular values, that determine the relative importance of the modes.

Illumination eigenmodes are usually found empirically by considering the SVD of a large number of images under different (directional source) illuminations, as in lighting-insensitive face and object recognition [4], [6]. We construct the PCA analytically, similar to previous work for Lambertian objects [13]. Specifically, we must relate the convolution formula above that applies to a single image with complex illumination to the eigenfunctions derived from a number of images taken assuming directional source lighting. We analyze $V(\alpha, \omega, \beta)$ for a particular groove (fixed β). Then, $V(\alpha, \omega)$ is a matrix with rows corresponding to groove locations (image pixels) α and columns corresponding to illumination directions ω . A singular-value decomposition (SVD) will give the eigenvalues (singular values) and illumination eigenmodes.

Equation (2) provides a formula for V . We are interested in the region $\omega > \beta_2$, where V depends on α . It is simplest to let $\omega' = \omega - \beta_2$ so both α and ω' lie in the range of $[0, \pi/2 - \beta_2]$ and we are considering the simple matrix $V(\alpha, \omega') = H(\omega' - \alpha)$.

Let us postulate (left and right) singular vectors or functions of the form $f(\alpha) = \sin(k\alpha)$ and $g(\omega') = \cos(k\omega')$, i.e., we can write the (nonsymmetric) matrix V as $F X G^t$, where X is a diagonal matrix of singular values and F and G are orthonormal matrices with $f(\alpha)$ and $g(\omega')$ corresponding to their columns. Standard linear algebra requires that $Vg = \sigma f$ and $V^t f = \sigma g$, where f and g are corresponding left and right singular vectors and σ is the corresponding singular value. In the continuous case, the matrix multiplication is replaced by an integration,

$$\int_0^{\pi/2 - \beta_2} H(\omega' - \alpha) \cos(k\omega') d\omega' = \sigma \sin(k\alpha),$$

$$\int_0^{\pi/2 - \beta_2} H(\omega' - \alpha) \sin(k\alpha) d\alpha = \sigma \cos(k\omega'). \quad (8)$$




Fig. 6. (a) Moss 3D texture with significant shadowing. Experimental setup is as in Fig. 2. (b) Six images of the moss with different lighting directions, as well as a control experiment of paper (a flat near-Lambertian surface). Note the variation of appearance of moss with illumination direction due to cast shadows, especially for large angles. (c) Decay of singular values for illumination eigenmodes for 3D textures is a straight line with slope approximately -1.5 on a logarithmic scale. In contrast, for a flat near-Lambertian surface, all of the energy is in the first eigenmode with a very rapid falloff. (d) The first four illumination eigenfunctions for moss, which are simply sines and cosines.

Now, the first equation above becomes (realizing that $H = 1$ when $\omega' < \alpha$),

$$\int_0^\alpha \cos(k\omega') d\omega' = \sigma \sin(k\alpha), \quad (9)$$

which is trivially satisfied for $\sigma = 1/k$. The second part of (8) becomes

$$\begin{aligned} \int_{\omega'}^{\pi/2 - \beta_2} \sin(k\alpha) d\alpha &= \frac{1}{k} \cos(k\omega') \\ \Rightarrow \frac{\cos(k\omega') - \cos(k(\pi/2 - \beta_2))}{k} &= \frac{\cos(k\omega')}{k} \\ \Rightarrow k &= \frac{(n + \frac{1}{2})\pi}{\pi/2 - \beta_2}, \end{aligned} \quad (10)$$

where n is a positive integer. We have also verified these results by numerical simulation. This derivation gives us the desired results for the eigenvalues σ and illumination eigenmodes g .

Eigenvalue Spectrum: The eigenvalues decay as $1/k$, corresponding to the Heaviside coefficients, as shown in Fig. 5. Because of the relatively slow $1/k$ decay, our frequencies need to be quite high (many terms) for a good approximation of cast shadows. On the other hand,² for irradiance on a convex curved surface, we convolve with the clamped cosine function $\max(\cos\theta, 0)$, whose Fourier coefficients falloff quadratically as $1/k^2$ with very few terms needed [1], [15].

In actual experiments on 3D textures, the eigenvalues decay somewhat faster. First, as explained in Section 4.1, the eigenvalues for cast shadows decay as $1/k^{3/2}$ (loglog slope -1.5) in 3D. Second, in the Lambertian case, since we are dealing with flat (as opposed to spherical) surfaces, the eigenvalues for irradiance drop off much faster than $1/k^2$. In fact, for an ideal flat diffuse surface, all of the

energy is in the first eigenmode that corresponds simply to Lambertian cosine-dependence.

Illumination Eigenmodes: The illumination eigenmodes are simply Fourier basis functions—sines and cosines. This is the case for irradiance on a curved surface in 2D as well [14], reinforcing the mathematically similar convolution structure.

3.3 Experimental Validation

In this section, we present an initial quantitative experimental result. The next sections generalize these results to 3D and present more thorough experimental validations. Note that the real images we observe can include global illumination and specular reflection effects not captured by the theory. Furthermore, some pixels in the images may not correspond to a V-groove-like structure at all. Nevertheless, we will see that the real data behaves much like the predictions of the model, motivating and validating our derivation.




Fig. 7. (a) Singular values for illumination basis functions due to cast shadows in a simulated 3D texture (randomly oriented V-grooves) plotted on a linear scale. A number of singular values cluster together. (b) Decay of singular values (value versus frequency or square root of singular value number) on a logarithmic scale (with natural logarithms included as axis labels). We get a straight line with slope approximately -1.5.

2. The Heaviside function has a position, or C^0 , discontinuity at the step, while the clamped cosine has a derivative, or C^1 , discontinuity at $\cos\theta = 0$. It is known in Fourier analysis [10] that a C^n discontinuity will generally result in a spectrum that falls off as $1/k^{n+1}$.




Fig. 8. Three-dimensional hemispherical basis functions obtained from numerical simulations of V-grooves. Green denotes positive values and red denotes negative values. θ and ϕ are a standard spherical parameterization.




Fig. 9. The functions in Fig. 8 are simple products of 1D basis functions along elevation θ and azimuthal ϕ directions, as per (13). Note that the $V_{\pm n}$ are sines and cosines, while the U_m are approximately Legendre polynomials (P_3 for $m = 1$, P_5 for $m = 2$). Fig. 12 shows corresponding experimental results on an actual 3D texture.

We used the experimental setup of Fig. 2, determining the eigenvalue spectrum and illumination eigenmodes for both a sample of moss and a flat piece of paper. The paper serves as a control experiment on a nearly Lambertian surface. Our results are shown in Fig. 6.

Eigenvalue Spectrum: As seen in Fig. 6c, the eigenvalues (singular values) for moss, when plotted on a log-log scale, lie on a straight line with slope approximately -1.5, as expected. This contrasts with the expected result for a flat Lambertian surface, where we should, in theory, see a single eigenmode (simply the cosine term). Indeed, in our control experiment with a piece of paper, also shown in Fig. 6c, 99.9 percent of the energy for the paper is in the first eigenmode.

Illumination Eigenmodes: As predicted, the illumination eigenmodes are simply Fourier basis functions—sines and cosines. This indicates that a common set of illumination eigenfunctions may describe lighting-dependence in many 3D textures.

4 3D NUMERICAL ANALYSIS OF CAST SHADOWS

In 3D, V-grooves can be rotated to any orientation about the vertical; hence, the *direction* of the Fourier basis functions can also be rotated. For a *given* V-groove direction, the 2D derivation essentially still holds since it depends on the monotonic increase in visibility as one moves along the groove, which still holds in 3D. The interesting question is what is the set of illumination basis functions that encompasses all V-groove (and correspondingly Fourier) orientations in 3D? In this section, we report on the results of numerical simulations, shown in Fig. 7, Fig. 8, and Fig. 9. We then verify these results with experiments on real 3D textures including moss, gravel, and a kitchen sponge.

4.1 Numerical Eigenvalue Spectrum and Illumination Eigenmodes

For numerical simulation, we consider V-grooves oriented at (rotated by) arbitrary angles about the vertical, ranging from 0 to 2π . An interesting future direction would be to consider alternative geometries, like general gaussian surfaces or height fields. For each orientation, we consider a number of V-groove angles with β ranging from 0 to $\pi/2$. In essence, we have an ensemble of a large number of V-grooves (1,000 in our simulations). Each point on each V-groove has a binary visibility value for each point on the illumination hemisphere. We assemble this information into a large visibility matrix, where rows correspond to V-groove points (image pixels) and columns to illumination directions. Then, as in experiments with real textures, we do an SVD³ to find the illumination eigenmodes.

Numerical Eigenvalue Spectrum: We first consider the eigenvalues or singular values, plotted in the left of Fig. 7, on a linear scale. At first glance, this plot is rather surprising. Even though the singular values decrease with increasing frequency, a number of them cluster together. Actually, these results are very similar to those for irradiance and spherical harmonics [1], [13], [15], where $2k + 1$ basis functions of order k are similar. We show k ranging from 1 to 15 in the right of Fig. 7. As expected, the curve is almost exactly a straight line on a log-log plot, with a slope of approximately -1.5. The higher slope (-1.5 compared to -1 in 2D) is a natural consequence of the properties of Fourier series of a function with a curve discontinuity [10]. The total energy (sum of squared singular values)

3. Owing to the large size of the matrices both here and in our experiments with real data, SVD is performed in a two step procedure in practice. First, we find the basis functions and eigenvalues for each V-groove. A second SVD is then performed on these weighted basis functions.

at each order k goes as $1/k^2$ in both 2D and 3D cases. However, in 3D, each frequency band contains $2k + 1$ functions, so the energy in each individual basis function decays as $1/k^3$, with the singular values therefore falling off as $1/k^{3/2}$.

Numerical Illumination Eigenmodes: The first nine eigenmodes are plotted in Fig. 8, where we label the eigenmodes using (m, n) with the net frequency given by $k = m + |n|$, with $k \geq 0$, $-k \leq n \leq k$, and $m = k - |n|$. This labeling anticipates the ensuing discussion and is also quite similar to that used for spherical harmonics.

To gain further insights, we attempt to *factor* these basis functions into a separable form. Most 2D basis functions are factorizable. For instance, consider the 2D Fourier transform,

$$W_{mn}(x, y) = U_m(x)V_n(y), \quad (11)$$

where W is the (complex) 2D basis function $\exp(imx)\exp(iny)$ and U_m and V_n are 1D Fourier functions ($\exp(imx)$ and $\exp(iny)$, respectively). Spherical harmonics and Zernike polynomials are also factorizable, but doing so is somewhat more complicated:

$$W_{mn}(\theta, \phi) = U_n^m(\theta)V_n(\phi), \quad (12)$$

where V_n is still a Fourier basis function $\exp(in\phi)$ (this is because of azimuthal symmetry in the problem and will be true in our case too) and U_n^m are associated Legendre polynomials for spherical harmonics, or Zernike polynomials. Note that U_n^m now has two indices, unlike the simpler Fourier case, and also depends on azimuthal index n .

We now factor our eigenmodes. The first few eigenfunctions are almost completely factorizable and representable in a form similar to (11), i.e., like a 2D Fourier transform, and simpler⁴ than spherical harmonics or Zernike polynomials,

$$W_{mn}(\theta, \phi) = U_m(\theta)V_n(\phi). \quad (13)$$

Fig. 9 shows factorization into 1D functions $U_m(\theta)$ and $V_n(\phi)$. It is observed that the U_m correspond closely to odd Legendre polynomials P_{2m+1} . This is not surprising since Legendre polynomials are spherical frequency-space basis functions. We observe only odd terms, $2m + 1$, since they correctly vanish at $\theta = \pi/2$ when a point is always shadowed. V_n are simply Fourier azimuthal functions, or sines and cosines. The net frequency $k = m + |n|$ with there being $2k + 1$ basis functions at order k .

4.2 Results of Experiments with Real 3D Textures

In this subsection, we report on empirical results in 3D, showing that the experimental observations are consistent with, and therefore validate, the theoretical and numerical analysis. We considered three different 3D textures—the moss and gravel, shown in Fig. 6, and a kitchen sponge. We report in this section primarily on results for the sponge; results for the other samples are similar.

For each texture, we took a number of images with a fixed overhead camera view and varying illumination direction. The set up in Fig. 2 shows a 2D slice of illumination directions. For the experiments in this section, the lighting ranged over the full 3D hemisphere. That is, θ ranged from $[14^\circ, 88^\circ]$ in 2 degree increments (38 different elevation angles) and ϕ from $[-180^\circ, 178^\circ]$ also in 2 degree increments (180 different azimuthal angles). Hence, we captured 6,840 images (38×180). This is a two order of magnitude denser sampling than the 205 images acquired by Dana et al. [3] to represent both light and view variation.

4. Mathematically, functions in the form of (13) can have a discontinuity at the pole $\theta = 0$. However, in our numerical simulations and experimental tests, we have found that this form closely approximates observed results and does not appear to create practical difficulties.




Fig. 10. (a) Plot of singular values for the sponge on a linear scale. (b) Singular values versus frequencies on a logarithmic scale with natural log axis labels. These experimental results should be compared to the predicted results from the numerical simulation in Fig. 7.

Experimental eigenvalue spectrum: Fig. 10 plots the experimentally observed falloff of eigenvalues. We see on the left that eigenmodes 2-4 (the first three after the mean term) cluster together as predicted by our numerical simulations. One can see a rather subtle effect of clustering in second order eigenmodes as well, but beyond that, the degeneracy is broken. This is not surprising for real data and is consistent with similar results for PCA analysis in Lambertian shading [13]. For low orders (corresponding to eigenmodes 2-16, or orders 1-3), the slope on a loglog plot is approximately -1.6, as shown in the right of Fig. 10, in agreement with the expected result of -1.5.

Experimental illumination eigenmodes: We next analyze the forms of the eigenmodes; the order 1 modes for moss, gravel, and sponge are shown in Fig. 11. The first order eigenmodes observed are linear combinations of the actual separable functions—this is expected and just corresponds to a rotation. We next found the separable functions U_m and V_n along θ and ϕ by using an SVD of the 2D eigenmodes. As expected, the U and V basis functions found separately from the three order 1 eigenmodes were largely similar and matched those obtained from numerical simulation. Our plots in Fig. 12 show both the *average* basis functions (in black) and the individual functions from the three eigenmodes (in red, blue, and green) for the sponge data set. We see that these have the expected forms and the eigenmodes are well described as a linear combination of separable basis functions.

4.3 Analytic Hemispherical Bases and Cast Shadows

Our results have indicated that a common set of hemispherical illumination basis functions may be appropriate for many natural 3D textures. Basis functions over the hemisphere are also important in a number of other computer graphics and vision applications with much recent work [5], [8], [9].

The illumination eigenmodes we observe suggest that the following set of orthonormal basis functions (per (13)) may be useful for analyzing cast shadows in 3D textures:

$$W_{mn}(\theta, \phi) = \sqrt{\frac{4m+3}{\pi}} P_{2m+1}(\cos \theta) a_n(\phi), \quad (14)$$

where $a_n(\phi)$ stands for $\cos n\phi$ or $\sin n\phi$, depending on whether n is plus or minus (and is $\sqrt{1/2}$ for $n = 0$), while P_{2m+1} are odd Legendre Polynomials.

These basis functions have some advantages as well as a few limitations. Since they involve Legendre polynomials over θ and Fourier basis functions over ϕ , they are closely related to spherical harmonics but specialized to the hemisphere. Their form, as per (13) and (14), is a simple product of 1D functions in θ and ϕ and is




Fig. 11. Order 1 eigenmodes experimentally observed for moss, gravel, and sponge. Note the similarity between the three textures and to the basis functions in Fig. 8. The numbers below represent each eigenmode as a linear combination of separable basis functions.




Fig. 12. Factored basis functions $U_m(\theta)$ and $V_n(\phi)$ for sponge. The top row shows the mean eigenmode and the functions $U_0(\theta)$ and $U_1(\theta)$. Below that are the nearly constant $V_0(\phi)$ and the sinusoidal $V_1(\phi)$, $V_{-1}(\phi)$. The colors red, blue, and green, respectively, are used to refer to the three order 1 eigenmodes that are factored to obtain U_m and V_n . We use black to denote the mean value across the three eigenmodes. It is seen that all the eigenmodes have very similar curves, which also match the results in Fig. 9.

simpler than (12) for Zernike polynomials, spherical harmonics, or hemispherical harmonics [5], [9]. Furthermore, for diffuse textures, due to visibility and shading effects, the intensity goes to 0 for lighting at grazing angles. These boundary conditions are automatically satisfied since odd Legendre polynomials vanish at $\theta = \pi/2$ or $\cos \theta = 0$. In some ways, our use of only odd terms is similar to some spherical harmonic constructions over the hemisphere [24]. Some simple experiments illustrating the ability of the basis functions above to represent real 3D texture data are reported in [17].

Note that the basis above does have limitations in terms of representing general hemispherical functions. The use of only odd Legendre polynomials prevents us from easily representing functions that are nonzero at $\theta = \pi/2$ (this can be partially corrected by adding a constant or zeroth order basis function). Also, slow convergence and inaccuracies may occur when representing some functions (even functions that go to 0 at $\theta = \pi/2$ can include even Legendre polynomial terms). Simply adding extra basis functions for the even Legendre polynomials is not easy analytically since the basis is no longer orthonormal (however, a numerical Gram-Schmidt orthogonalization procedure can enable this in practice).

Finally, we would like to note that there are now three possible ways of constructing hemispherical basis functions. The functions above are adapted from a domain of cylindrical topology. The Zernike polynomials [8] are adapted from the disk, while hemispherical harmonics [5], [9] are adapted directly from spherical harmonics on the sphere. All three techniques involve some undesirable stretching of the base domain to adapt to the hemisphere and have both advantages and disadvantages. An interesting future direction is to analyze the benefits of the various constructions and understand the best hemispherical basis functions to use for different practical problems.

5 IMPLICATIONS

While this paper is primarily theoretical, there are many potential future practical insights and implications from our work in computer vision and graphics.

Inverse Lighting: It has been shown [15] that illumination estimation from a Lambertian surface is ill-posed since only the first two orders of illumination can be reliably estimated. On the other hand, recent work by Sato et al. [18] has shown that illumination can often be estimated from cast shadows. Our results explain why it is feasible to estimate much higher frequencies of the illumination (up to approximately order 10-20) from the effects of cast shadows—the filter or kernel because of cast shadows decays only as $1/k$ instead of $1/k^2$.

Lighting-Insensitive Recognition: There has been much theoretical and practical work [1], [4], [6], [13] on low-dimensional subspaces for lighting-insensitive recognition of nearly Lambertian objects. However, these results do not consider cast shadows. Our results suggest that it is possible to develop, in a similar fashion, subspaces that include the effects of cast shadows simply by considering more terms.

BTF modeling and rendering: BTFs [3] are 6D functions that are very difficult to acquire a dense sampling of because of the sheer size of the data. Furthermore, BTFs must encode the effects of cast shadows, making simple quadratic bases such as order 2 spherical harmonics [1], [15] or polynomial texture maps [11] inadequate. Our results suggest that a common Fourier basis can be used for compactly modeling illumination variation in many BTFs.

Real-Time Rendering: Our results apply to real time rendering applications based on precomputed images or radiance functions on geometry, such as recent work by Sloan et al. [19]. While spherical harmonics are appropriate for irradiance and other general reflection functions [16], there has previously been no

theoretical foundation for using these basis functions to consider the effects of cast shadows nor for determining the number of terms to use. In this paper, we have formalized the idea of convolution for cast shadows, suggesting the use of a frequency space basis and seen that we need many more basis functions than the order 2 harmonics needed for irradiance.

6 CONCLUSIONS

This paper formally analyzes cast shadows, showing that a simple Fourier signal-processing framework can be derived in many common cases. Our results indicate a theoretical link between cast shadows and convolution formulae for irradiance and more general non-Lambertian materials [1], [15], [16]. This paper is also a first step in quantitatively understanding the effects of lighting in 3D textures, where cast shadows play a major role.

ACKNOWLEDGMENTS

The authors would like to thank the reviewers, especially reviewer 2, for their detailed reading of the paper. This work was supported in part by grants from the National Science Foundation (#0305322, #0430258, #0085864) and Intel Corporation.

REFERENCES

- [1] R. Basri and D. Jacobs, "Reflectance and Linear Subspaces," *Proc. IEEE Int'l Conf. Computer Vision (ICCV '01)*, pp. 383-390, 2001.
- [2] K. Dana and S. Nayar, "Histogram Model for 3D Textures," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition (CVPR '98)*, pp. 618-624, 1998.
- [3] K. Dana, B. van Ginneken, S. Nayar, and J. Koenderink, "Reflectance and Texture of Real-World Surfaces," *ACM Trans. Graphics*, vol. 18, no. 1, pp. 1-34, Jan. 1999.
- [4] R. Epstein, P. Hallinan, and A. Yuille, "5 Plus or Minus 2 Eigenimages Suffice: An Empirical Investigation of Low-Dimensional Lighting Models," *Proc. IEEE Workshop Physics-Based Modeling in Computer Vision*, pp. 108-116, 1995.
- [5] P. Gautron, J. Krivanek, S. Pattanaik, and K. Bouatouch, "A Novel Hemispherical Basis for Accurate and Efficient Rendering," *Proc. EuroGraphics Symp. Rendering (EGSR '04)*, 2004.
- [6] P. Hallinan, "A Low-Dimensional Representation of Human Faces for Arbitrary Lighting Conditions," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition (CVPR '94)*, pp. 995-999, 1994.
- [7] J. Koenderink, A. Doorn, K. Dana, and S. Nayar, "Bidirectional Reflection Distribution Function of Thoroughly Pitted Surfaces," *Int'l J. Computer Vision*, vol. 31, nos. 2/3, pp. 129-144, 1999.
- [8] J. Koenderink and A. van Doorn, "Phenomenological Description of Bidirectional Surface Reflection," *J. Optical Soc. Am. A*, vol. 15, no. 11, pp. 2903-2912, 1998.
- [9] O. Makhutkin, "Analysis of Radiative Transfer between Surfaces by Hemispherical Harmonics," *J. Quantitative Spectroscopy*, vol. 56, no. 6, pp. 869-879, 1996.
- [10] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [11] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial Texture Maps," *Proc. SIGGRAPH '01*, pp. 519-528, 2001.
- [12] M. Oren and S. Nayar, "Generalization of Lambert's Reflectance Model," *Proc. SIGGRAPH '94*, pp. 239-246, 1994.
- [13] R. Ramamoorthi, "Analytic PCA Construction for Theoretical Analysis of Lighting Variability in Images of a Lambertian Object," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 10, pp. 1322-1333, Oct. 2002.
- [14] R. Ramamoorthi and P. Hanrahan, "Analysis of Planar Light Fields from Homogeneous Convex Curved Surfaces under Distant Illumination," *SPIE Photonics West: Human Vision and Electronic Imaging VI*, pp. 185-198, 2001.
- [15] R. Ramamoorthi and P. Hanrahan, "On the Relationship between Radiance and Irradiance: Determining the Illumination from Images of a Convex Lambertian Object," *J. Optical Soc. Am. A*, vol. 18, no. 10, pp. 2448-2459, 2001.
- [16] R. Ramamoorthi and P. Hanrahan, "A Signal-Processing Framework for Inverse Rendering," *Proc. SIGGRAPH '01*, pp. 117-128, 2001.
- [17] R. Ramamoorthi, M. Koudelka, and P. Belhumeur, "A Fourier Theory for Cast Shadows," *Proc. European Conf. Computer Vision (ECCV '04)*, pp. 1146-1162, 2004.
- [18] I. Sato, Y. Sato, and K. Ikeuchi, "Illumination Distribution from Brightness in Shadows: Adaptive Estimation of Illumination Distribution with Unknown Reflectance Properties in Shadow Regions," *Proc. IEEE Int'l Conf. Computer Vision (ICCV '99)*, pp. 875-882, 1999.
- [19] P. Sloan, J. Kautz, and J. Snyder, "Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments," *ACM Trans. Graphics (SIGGRAPH '02)*, vol. 21, no. 3, pp. 527-536, 2002.
- [20] C. Soler and F. Sillion, "Fast Calculation of Soft Shadow Textures Using Convolution," *Proc. SIGGRAPH '98*, pp. 321-332, 1998.
- [21] P. Suen and G. Healey, "Analyzing the Bidirectional Texture Function," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition (CVPR '98)*, pp. 753-758, 1998.
- [22] K. Thornber and D. Jacobs, "Cast Shadows and Linear Subspaces," Technical Report NEC TR #2001-100, NEC, 2001.
- [23] K. Torrance and E. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces," *J. Optical Soc. Am.*, vol. 57, no. 9, pp. 1105-1114, 1967.
- [24] S. Westin, J. Arvo, and K. Torrance, "Predicting Reflectance Functions from Complex Surfaces," *Proc. SIGGRAPH '92*, pp. 255-264, 1992.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.