

A Fourier Theory for Cast Shadows

Ravi Ramamoorthi¹, Melissa Koudelka², and Peter Belhumeur¹

¹ Columbia University, {ravir, belhumeur}@cs.columbia.edu

² Yale University, melissa.koudelka@yale.edu

Abstract. Cast shadows can be significant in many computer vision applications such as lighting-insensitive recognition and surface reconstruction. However, most algorithms neglect them, primarily because they involve non-local interactions in non-convex regions, making formal analysis difficult. While general cast shadowing situations can be arbitrarily complex, many real instances map closely to canonical configurations like a wall, a V-groove type structure, or a pitted surface. In particular, we experiment on 3D textures like moss, gravel and a kitchen sponge, whose surfaces include canonical cast shadowing situations like V-grooves. This paper shows theoretically that many shadowing 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 recently developed signal-processing frameworks for reflection and illumination. An analytic convolution formula is derived for a 2D V-groove, which is shown to correspond closely to many common shadowing situations, especially in 3D textures. Numerical simulation is used to extend these results to general 3D textures. These results also provide evidence that a common set of illumination basis functions may be appropriate for representing lighting variability due to cast shadows in many 3D textures. We derive a new analytic basis suited for 3D textures to represent illumination on the hemisphere, with some advantages over commonly used Zernike polynomials and spherical harmonics. New experiments on analyzing the variability in appearance of real 3D textures with illumination motivate and validate our theoretical analysis. Empirical results show that illumination eigenfunctions often correspond closely to Fourier bases, while the eigenvalues drop off significantly slower than those for irradiance on a Lambertian curved surface. These new empirical results are explained in this paper, based on our theory.

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 non-local 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




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.

in Figure 1. From left to right, shadowing by a wall, a V-groove like structure, a plane such as a desk above, and a pitted or curved surface. Though the figure is in 2D, similar patterns often apply in 3D along the radial direction, with little change in the extent of shadowing along transverse or azimuthal directions.

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 Figures 2 and 6. These objects have fine-scale structures similar to the canonical configurations shown in Figure 1. Hence, they exhibit interesting illumination and view-dependence, which is often described using a bi-directional texture function (BTF) [3]. In this paper, we analyze lighting variability, assuming 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 [19] convolution result for shadows when source, blocker and receiver are all in parallel planes—for instance, V-grooves (b in Figure 1, as well as a and d) do not contain any parallel planes. Our specific technical contributions include the following:

- We derive an analytic convolution formula for a 2D V-groove, and show that it applies to many canonical shadowing situations, such as those in Figure 1.
- We analyze the illumination eigenmodes, showing how they correspond closely to Fourier basis functions. We also analyze the eigenvalue spectrum, discussing similarities and differences with convolution results for Lambertian curved surfaces and irradiance, and showing why the falloff is slower in the case of cast shadows.
- We explain important lighting effects in 3D textures, documented quantitatively here for the first time. Experimental results confirm the theoretical analysis.
- We introduce new illumination basis functions over the hemisphere for lighting variability due to cast shadows in 3D textures, potentially applicable to compression,

interpolation and prediction. These bases are based on analytic results and numerical simulation, and validated by empirical results. They have some advantages over the commonly used spherical harmonics and Zernike polynomials.

Our paper builds on a rich history of previous work on reflection models, such as Oren-Nayar [12], Torrance-Sparrow [21], Wolff et al. [23] and Koenderink et al. [6], as well as several recent articles on the properties of 3D textures [2,20]. Our analytic formulae are derived considering the standard V-grooves used in many of these previous reflection models [12,21]. 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 relates to recent approaches to real-time rendering, such as the precomputed transfer method of Sloan et al. [18], that represents appearance effects including cast shadows, due to low-frequency illumination, represented in spherical harmonics. However, there is no analytic convolution formula or insight in their work as to the optimal basis functions or the number of terms needed for good approximation. We seek to put future real-time rendering methods on a strong theoretical footing by formalizing the idea of convolution for cast shadows, analyzing the form of the eigenvalue spectrum, showing that the decay is much slower than for Lambertian irradiance, and that we therefore need many more basis functions to capture sharp shadows than the low order spherical harmonics and polynomials used by Sloan et al. [18] and Malzbender et al. [11].

2 The Structure of Cast Shadows

In this section, we briefly discuss the structure of cast shadows, followed in the next section by a derivation of an analytic convolution formula for a 2D V-groove, Fourier and principal component analysis, and initial experimental observations and validation.

First, we briefly make some theoretical observations. Consider Figures 1 a and b. 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 Figures 1 c and d, we have two extreme points B and C . The net shadowing effect is essentially just 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 Figure 1. In Figure 2, we show an initial experiment. We illuminated a sample of gravel along an arc (angle ranged from -90° to $+64^\circ$, limited by specifics of the acquisition). The varying appearance with illumination clearly suggests 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 Figure 1 (a), (b) and (d).

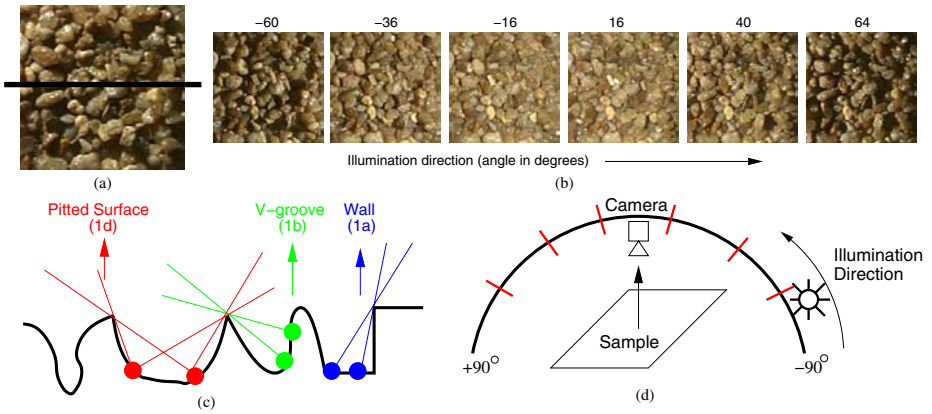


Fig. 2. (a): Gravel texture, which exhibits strong shadowing. (b): Images with different light directions clearly show cast shadow appearance effects, especially at large angles. The light directions correspond to the red marks in (d). (c): Conceptual representation of a profile of a cross section through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

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 Figure 3, corresponding to Figure 1 b. However, the derivation will be similar for any other shadowing situation, such as those in Figure 1, where the visibility is locally *monotonically changing*. Note that the V-groove model in Figure 3 can model the examples in Figures 1 a and b ($\beta_1 = 0, \beta_2 = \pi/2$ and $\beta_1 = \beta_2$), and each of the extreme points of Figures 1 c and d.




Fig. 3. Diagram of V-groove with groove angle ranging from $-\beta_1$ to $+\beta_2$. While the figure shows $\beta_1 = \beta_2$, as is common for previous V-groove models, there is no requirement of symmetry. We will be interested in visibility for points $A(x)$ where x is the distance along the groove (the labels are as in Figure 1; the line $A'B$ is omitted for clarity). Note that the visible region of $A(x)$, determined by $\alpha(x)$, increases monotonically with x along the groove.

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. The left side can be treated similarly. For 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 Figure 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 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 $[-\beta_1, \beta_2 + \alpha]$,

$$V(\alpha, \omega, \beta) = H(-\beta_1 - \omega) - H((\beta_2 + \alpha) - \omega), \\ H(u) = 1 \text{ if } u < 0, \quad 0 \text{ if } u > 0, \quad (2)$$

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 $\omega > \beta_2 + \alpha$. Figure 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 Equation 2 is always 1, the second term is 0, and $V = 1$ (no cast shadowing), as expected.

For a particular groove (fixed β), V is given by the following intervals.

$$\begin{array}{lll} -\pi/2 < \omega < -\beta_1 & V = 0 & \text{independent of } \alpha \\ -\beta_1 < \omega < +\beta_2 & V = 1 & \text{independent of } \alpha \\ +\beta_2 < \omega < \beta_2 + \alpha & V = 1 & \text{interval depends on } \alpha \\ \beta_2 + \alpha < \omega < \pi/2 & V = 0 & \text{interval depends on } \alpha. \end{array} \quad (3)$$

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.




Fig. 4. Illustration of the visibility function as per Equation 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.

Convolution Formula: Plugging Equation 2 back into Equation 1, we obtain

$$E(\alpha, \beta) = \int_{-\pi/2}^{\pi/2} L(\omega)H(-\beta_1 - \omega) d\omega - \int_{-\pi/2}^{\pi/2} L(\omega)H((\beta_2 + \alpha) - \omega) d\omega. \quad (4)$$

E is the sum of two terms, the first of which depends only on groove angle β_1 , and the second that 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 integral of the illumination. We now separate the two terms to simplify this result as (\otimes is the convolution operator)

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

$$\tilde{E}(u) = \int_{-\pi/2}^{\pi/2} L(\omega)H(u - \omega) d\omega = L \otimes H. \quad (5)$$

Fourier Analysis: Equation 5 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,

$$\tilde{E}_k = \sqrt{\pi}L_kH_k, \quad (6)$$

where L_k are the Fourier illumination coefficients, and H_k are Fourier coefficients of the Heaviside step function, plotted in Figure 5. The even coefficients H_{2k} vanish, while the odd **coefficients decay** as $1/k$. The analytic formula is

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

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

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




Fig. 5. Comparison of Fourier coefficients for the Heaviside step function for cast shadows (left) and the clamped cosine Lambertian function for irradiance (middle). For the step function, even terms vanish, while odd terms decay as $1/k$. For the clamped cosine, odd terms greater than 1 vanish, while even terms decay much faster as $1/k^2$. On the right 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.

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. We also compare to similar analyses for irradiance on a curved surface [1,13,14,15].

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 [4,5]. It seems intuitive in our case that the eigenfunctions will be sines and cosines. To formalize this analytically, 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. Our approach is conceptually similar to Ramamoorthi’s work on analytic PCA construction [13] for images of a convex curved Lambertian object.

Specifically, 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. It can be formally shown (details omitted here) that the following results hold, as expected.

Eigenvalue Spectrum: The eigenvalues decay as $1/k$, corresponding to the Heaviside coefficients, as shown in Figure 5. Because of the relatively slow $1/k$ decay, we need quite high frequencies (many terms) for good approximation of cast shadows. On the other hand², 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 for accurate representation [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

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 $1/k^{n+1}$.

slope -1.5) in 3D, as opposed to $1/k$ in 2D. 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.

Implications: There are many potential implications of these results, to explain empirical observations and devise practical algorithms. For instance, it has been shown [15] that illumination estimation from a convex Lambertian surface is ill-posed since only the first two orders can be estimated. But Sato et al. [17] have 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 from the effects of cast shadows. In lighting-insensitive recognition, there has been much work on low-dimensional subspaces for Lambertian objects [1,4,5,13]. Similar techniques might be applied, simply using more basis functions, and including cast shadow effects, since cast shadows and irradiance have the same mathematical structure. Our results have direct implications in BTF modeling and rendering for representing illumination variability, and providing appropriate basis functions for compression and synthesis.

3.3 Experimental Validation

In this subsection, we present an initial quantitative experimental result motivating and validating our derivation. The next sections generalize these results to 3D, and present more thorough experimental validations. We used the experimental setup of Figure 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 Figure 6.

Eigenvalue Spectrum: As seen in Figure 6 (c), 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 Figure 6 (c), 99.9% of the energy for the paper is in the first eigenmode, with a very fast decay after that.

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

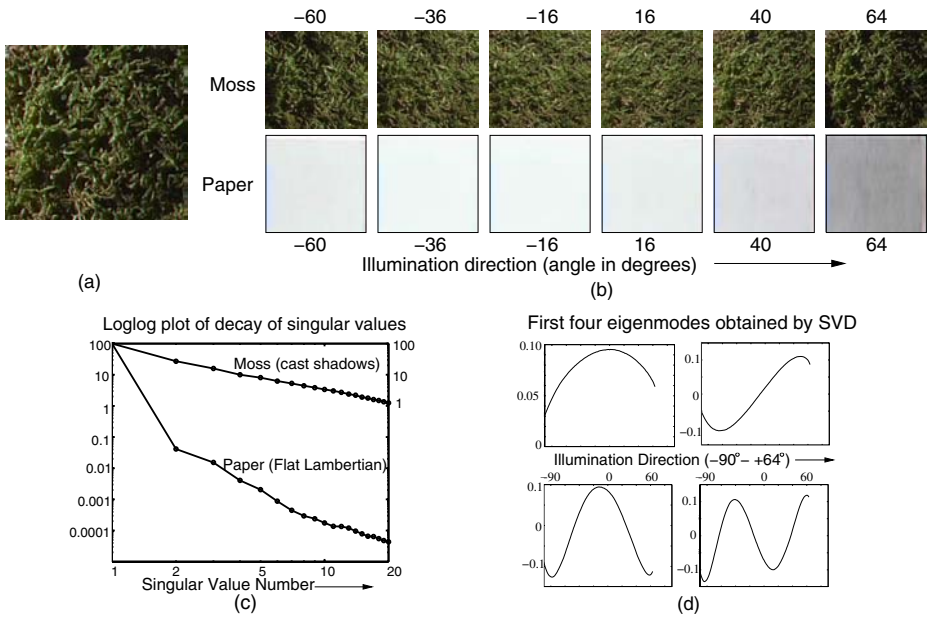


Fig. 6. (a): Moss 3D texture with significant shadowing. Experimental setup is as in Figure 2. (b): 6 images of the moss with different lighting directions, as well as a control experiment of paper (a flat near-Lambertian surface). Note the variation of appearance of moss with illumination direction due to cast shadows, especially for large angles. In contrast, while the overall intensity changes for the paper, there is almost no variation on the surface. (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.

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?

One might expect the basis functions to be close to spherical harmonics [9], the natural extension of the Fourier basis to the sphere. However, we are considering only the visible upper hemisphere, and we will see that our basis functions take a somewhat simpler form than spherical harmonics or Zernike polynomials [7], corresponding closely to 2D Fourier transforms. In this section, we report on the results of numerical simulations, shown in Figures 7, 8 and 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 . For each orientation, we consider a number

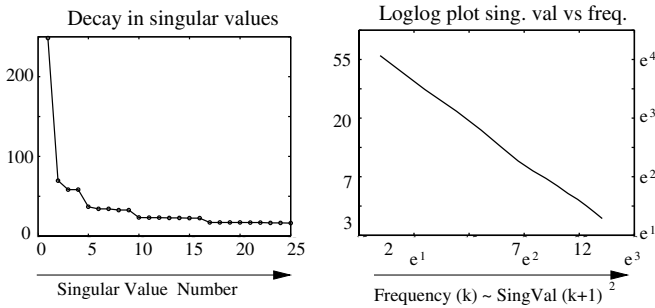


Fig. 7. *Left:* Singular values for illumination basis functions due to cast shadows in a simulated 3D texture (randomly oriented V-grooves), plotted on a linear scale. A number of singular values cluster together. *Right:* Decay of singular values [value vs frequency or square root of singular value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

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 (1000 in our simulations). Each point on each V-groove has a binary visibility value for each point on the illumination hemisphere. We can assemble all this information into a large visibility matrix, where the rows correspond to V-groove points (image pixels), and the columns to illumination directions. Then, as in experiments with real textures like Figure 6, we do an SVD³ to find the illumination eigenmodes and eigenvalues.

Numerical Eigenvalue Spectrum: We first consider the eigenvalues or singular values, plotted in the left of Figure 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. Similarly, our eigenmodes are Fourier-like, with $2k + 1$ eigenmodes at order k (with a total of $(k + 1)^2$ eigenmodes up to order k). Therefore, to determine the decay of singular values, it is more appropriate to consider them as a function of order k . We show k ranging from 1 to 15 in the right of Figure 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], as is the case in 3D visibility. The total energy (sum of *squared* singular values) 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 Figure 8, where we label the eigenmodes using (m, n) with the net frequency given by

³ Owing to the large size of the matrices both here and in our experiments with real data, SVD is performed in a 2 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.

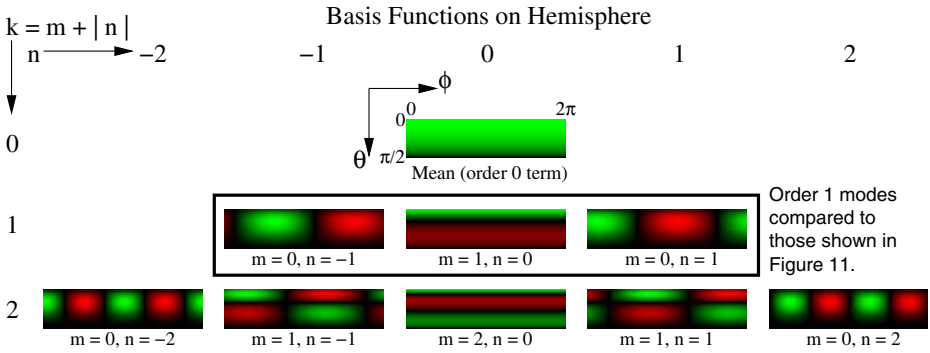


Fig. 8. 3D hemispherical basis functions obtained from numerical simulations of V-grooves. Green denotes positive values and red denotes negative values. θ and ϕ are a standard spherical parameterization, with the cartesian $(x, y, z) = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$.

$k = m + |n|$, 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 commonly used 2D or (hemi)spherical basis functions are factorizable. For instance, consider the 2D Fourier transform. In this case,

$$W_{mn}(x, y) = U_m(x)V_n(y), \tag{8}$$

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.

$$W_{mn}(\theta, \phi) = U_n^m(\theta)V_n(\phi), \tag{9}$$

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 Equation 8, i.e., like a 2D Fourier transform, and simpler ⁴ than spherical harmonics or Zernike polynomials,

$$W_{mn}(\theta, \phi) = U_m(\theta)V_n(\phi). \tag{10}$$

Figure 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

⁴ Mathematically, functions of the form of Equation 10 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. 9. The functions in Figure 8 are simple products of 1D basis functions along elevation θ and azimuthal ϕ directions, as per Equation 10. 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 $m = 2$). Figure 12 shows corresponding experimental results on an actual 3D texture.

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 Figure 6, and a kitchen sponge, shown in Figure 14. 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 setup in Figure 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). The acquisition setup restricted imaging near the pole. Hence we captured 6840 images (38×180) for each texture. 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, and provides a good testbed for comparison with simulations.

For numerical work, we then assembled all of this information in a large matrix, the rows of which were image pixels, and the columns of which were light source directions. Just as in our numerical simulations, we then used SVD to find the illumination eigenmodes and eigenvalues. We validate the numerical simulations by comparing the experimental results for real data to the expected (i.e., numerical) results just described.




Fig. 10. *Left:* Plot of singular values for the sponge on a linear scale. *Right:* Singular values vs frequencies on a logarithmic scale, with natural log axis labels. These experimental results should be compared to the predicted results from numerical simulation in Figure 7.

Experimental eigenvalue spectrum: Figure 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 consistent with similar results for PCA analysis in Lambertian shading [13].

Computing the slope for singular value dropoff is difficult because of insufficiency of accurate data (the first 20 or so eigenmodes correspond only to the first 5 orders, and noise is substantial for higher order eigenmodes). 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 Figure 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 Figure 11. The first order eigenmodes observed are linear combinations of the actual separable functions—this is expected, and just corresponds to a rotation.

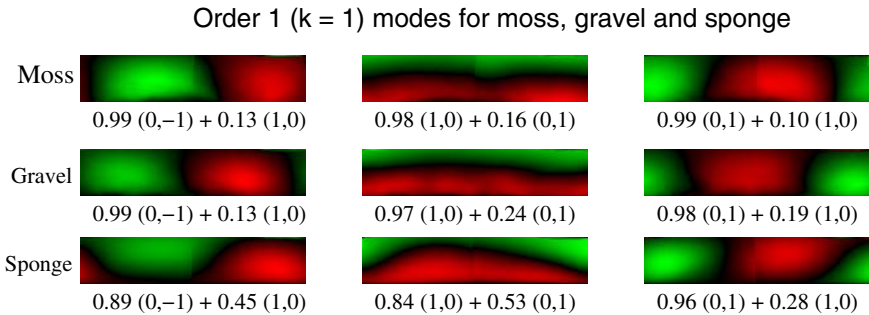


Fig. 11. Order 1 eigenmodes experimentally observed for moss, gravel and sponge. Note the similarity between the 3 textures, and to the basis functions in Figure 8. The numbers below represent each eigenmode as a linear combination of separable basis functions.

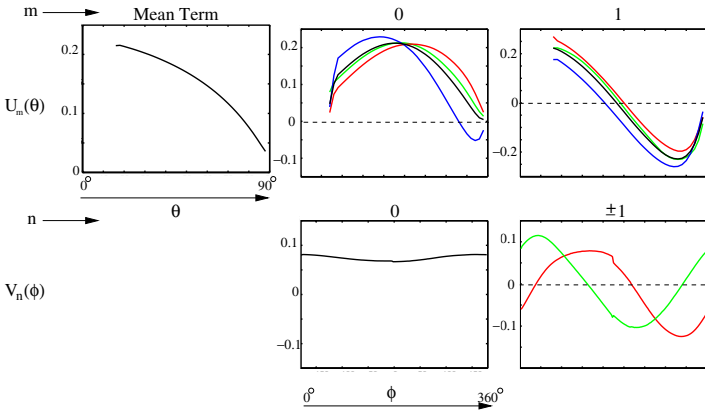


Fig. 12. Factored basis functions $U_m(\theta)$ and $V_n(\phi)$ for sponge. The top row shows the mean eigenmode, and the functions $U_0(\theta)$ and $U_1(\theta)$. Below that are the nearly constant $V_0(\phi)$ and the sinusoidal $V_1(\phi)$, $V_{-1}(\phi)$. The colors red, blue and green respectively 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 Figure 9.

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 Figure 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 dataset. We see that these have the expected forms, and the eigenmodes are well described as a linear combination of separable basis functions.

5 Representation of Cast Shadow Effects in 3D Textures

The previous sections have shown how to formally analyze cast shadow effects, numerically simulated illumination basis functions, and experimentally validated the results. In this section, we make a first attempt at using this knowledge to efficiently represent lighting variability due to cast shadows in 3D textures.

In particular, our results indicate that a common set of illumination basis functions may be appropriate for many natural 3D textures. We will use an analytic basis motivated by the form of the illumination eigenmodes observed in the previous section. We use Equation 10, with the normalized basis functions written as

$$W_{mn}(\theta, \phi) = \sqrt{\frac{4m+3}{\pi}} P_{2m+1}(\cos \theta) a z_n(\phi), \tag{11}$$

where $a z_n(\phi)$ 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.

This basis has some advantages over other possibilities such as spherical harmonics or Zernike polynomials for representing illumination over the hemisphere in 3D textures.

- The basis is specialized to the hemisphere, unlike spherical harmonics.
- Its form, as per Equations 10 and 11 is a simple product of 1D functions in θ and ϕ , simpler than Equation 9 for Zernike polynomials and spherical harmonics.
- For diffuse textures, due to visibility and shading effects, the intensity goes to 0 at grazing angles. These boundary conditions are automatically satisfied, since odd Legendre polynomials vanish at $\theta = \pi/2$ or $\cos \theta = 0$.
- Our basis seems consistent with numerical simulations and real experiments.

Figure 13 compares the resulting error with our basis to that for spherical harmonics, Zernike polynomials and the numerically computed optimal SVD basis for the sponge example. Note that the SVD basis performs best because it is tailored to the particular dataset and is by definition optimal. However, it requires prior knowledge of the data on a specific 3D texture, while we seek an analytic basis suitable for all 3D textures. These results demonstrate that our basis is competitive with other possibilities and can provide a good compact representation of measured illumination data in textures.

We demonstrate two simple applications of our analytic basis in Figure 14. In both cases, we use our basis to fit a function over the hemisphere. For 3D textures, this function is the illumination-dependence, fit separately at each pixel. The first application is to **compression**, wherein the original 6840 images are represented using 100 basis




Fig. 13. Comparison of errors from different bases on sponge example (the left shows the first 6 terms, while the right shows larger numbers of terms). The SVD basis is tailored to this particular dataset and hence performs best; however it requires full prior knowledge.



Fig. 14. On the left is one of the actual sponge images. In the middle is a reconstruction using 100 analytic basis functions, achieving a compression of 70:1. In the right, we reconstruct from a sparse set of 390 images, using our basis for interpolation and prediction. Note the subtle features of appearance, like accurate reconstruction of shadows, that are preserved.

function coefficients at each pixel. A compression of 70:1 is thus achieved, with only marginal loss in sharpness. This compression method was presented in [8], where a numerically computed SVD basis for textures sampled in both lighting and viewpoint was used. Using a standard analytic basis is simpler, and the same basis can now be used for all 3D textures. Further, note that once continuous basis functions have been fit, they can be evaluated for intermediate light source directions not in the original dataset. Our second application is to **interpolation** from a sparse sampling of 390 images. As shown in Figure 14, we are able to accurately reconstruct images not in the sparse dataset, potentially allowing for much faster acquisition times (an efficiency gain of 20 : 1 in this case), without sacrificing the resolution or quality of the final dataset.

It is important to discuss some limitations of our experiments and the basis proposed in equation 11. First, this is an initial experiment, and a full quantitative conclusion would require more validation on a variety of materials. Second, the basis functions in equation 11 are a good approximation to the eigenmodes derived from our numerical simulations, but the optimal basis will likely be somewhat different for specific shadowing configurations or 3D textures. Also, our basis functions are specialized to hemispherical illumination for macroscopically flat textures; the spherical harmonics or Zernike polynomials may be preferred in other applications. Another point concerns the use of our basis for representing general hemispherical functions. In particular, our basis functions go to 0 as $\theta = \pi/2$, which is appropriate for 3D textures, and similar to some spherical harmonic constructions over the hemisphere [22]. However, this makes it unsuitable for other applications, where we want a general hemispherical basis.

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. In that context, we have derived new illumination basis functions over the hemisphere, which are simply a separable basis written as a product of odd Legendre polynomials and Fourier azimuthal functions.

Acknowledgements. We thank the reviewers for pointing out several important references we had missed. This work was supported in part by grants from the National Science Foundation (ITR #0085864 Interacting with the Visual world, and CCF # 0305322 Real-Time Visualization and Rendering of Complex Scenes) and Intel Corporation (Real-Time Interaction and Rendering with Complex Illumination and Materials).

References

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

- [2] K. Dana and S. Nayar. Histogram model for 3d textures. In *CVPR 98*, pages 618–624, 1998.
- [3] K. Dana, B. van Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.
- [4] R. Epstein, P. Hallinan, and A. Yuille. 5 plus or minus 2 eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *IEEE 95 Workshop Physics-Based Modeling in Computer Vision*, pages 108–116.
- [5] P. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *CVPR 94*, pages 995–999, 1994.
- [6] J. Koenderink, A. Doorn, K. Dana, and S. Nayar. Bidirectional reflection distribution function of thoroughly pitted surfaces. *IJCV*, 31(2/3):129–144, 1999.
- [7] J. Koenderink and A. van Doorn. Phenomenological description of bidirectional surface reflection. *JOSA A*, 15(11):2903–2912, 1998.
- [8] M. Koudelka, S. Magda, P. Belhumeur, and D. Kriegman. Acquisition, compression, and synthesis of bidirectional texture functions. In *ICCV 03 Workshop on Texture Analysis and Synthesis*, 2003.
- [9] T.M. MacRobert. *Spherical harmonics; an elementary treatise on harmonic functions, with applications*. Dover Publications, 1948.
- [10] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [11] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *SIGGRAPH 01*, pages 519–528, 2001.
- [12] M. Oren and S. Nayar. Generalization of lambert’s reflectance model. In *SIGGRAPH 94*, pages 239–246, 1994.
- [13] R. Ramamoorthi. Analytic PCA construction for theoretical analysis of lighting variability in images of a lambertian object. *PAMI*, 24(10):1322–1333, 2002.
- [14] 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.
- [15] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *JOSA A*, 18(10):2448–2459, 2001.
- [16] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, pages 117–128, 2001.
- [17] I. Sato, Y. Sato, and K. Ikeuchi. Illumination distribution from brightness in shadows: adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. In *ICCV 99*, pages 875–882, 1999.
- [18] 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 2002)*, 21(3):527–536, 2002.
- [19] C. Soler and F. Sillion. Fast calculation of soft shadow textures using convolution. In *SIGGRAPH 98*, pages 321–332.
- [20] P. Suen and G. Healey. Analyzing the bidirectional texture function. In *CVPR 98*, pages 753–758, 1998.
- [21] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967.
- [22] S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92*, pages 255–264, 1992.
- [23] L. Wolff, S. Nayar, and M. Oren. Improved diffuse reflection models for computer vision. *IJCV*, 30:55–71, 1998.