

Reflectance Sharing: Predicting Appearance from a Sparse Set of Images of a Known Shape

Todd Zickler, *Member, IEEE*, Ravi Ramamoorthi, Sebastian Enrique,
and Peter N. Belhumeur

Abstract

Three-dimensional appearance models consisting of spatially-varying reflectance functions defined on a known shape can be used in analysis-by-synthesis approaches to a number of visual tasks. The construction of these models requires the measurement of reflectance, and the problem of recovering spatially-varying reflectance from images of known shape has drawn considerable interest. To date, existing methods rely on either: i) low-dimensional (e.g., parametric) reflectance models, or ii) large datasets involving thousands of images (or more) per object. Appearance models based on the former have limited accuracy and generality, since they require the selection of a specific reflectance model *a priori*; and while approaches based on the latter may be suitable for certain applications, they are generally too costly and cumbersome to be used for image analysis. We present an alternative approach that seeks to combine the benefits of existing methods by enabling the estimation of a *non-parametric* spatially-varying reflectance function from a small number of images. We frame the problem as scattered-data interpolation in a mixed spatial and angular domain, and we present a theory demonstrating that the angular accuracy of a recovered reflectance function can be increased in exchange for a decrease in its spatial resolution. We also present a practical solution to this interpolation problem using a new representation of reflectance based on radial basis functions. This representation is evaluated experimentally by testing its ability to predict appearance under novel view and lighting conditions. Our results suggest that, since reflectance typically varies slowly from point to point over much of an object's surface, we can often obtain a non-parametric reflectance function from a sparse set of images. In fact, in some cases, we can obtain reasonable results in the limiting case of only a single input image.

Index Terms

Reflectance, BRDF, image synthesis, image-based rendering, radial basis functions.

T. Zickler is with the Division of Engineering and Applied Sciences at Harvard University. E-mail: zickler@eecs.harvard.edu. R. Ramamoorthi, S. Enrique, and P. N. Belhumeur are with the Computer Science Department at Columbia University. Email: {ravir,senrique,belhumeur}@cs.columbia.edu

I. INTRODUCTION

Three-dimensional appearance models consisting of explicit shape models and spatially-varying reflectance functions defined on these shapes are effective tools for analysis-by-synthesis approaches to recognition, tracking, and other visual tasks. This approach to image analysis provides the advantage of separating intrinsic object properties (i.e., shape and material properties) from extrinsic scene properties, such as pose, illumination, and camera parameters.

Constructing an appearance model of a given object requires two stages: recovery of both shape and reflectance. Yet, while great strides have been made at recovering object shape (laser-scanners, structured-light systems, multi-view stereo, photometric stereo, Helmholtz stereo, etc.), less progress has been made at recovering reflectance properties. Recovering and efficiently representing reflectance is difficult because of its inherent high-dimensionality. At each surface point, reflectance is described by a four-dimensional function of the view and lighting directions termed the bi-directional reflectance distribution function (BRDF) [1]. The BRDF generally changes spatially over an object’s surface, and recovering this spatially-varying BRDF (or 6D SBRDF) without further assumptions typically requires a set of images large enough to densely sample high-frequency radiometric events, such as sharp specular highlights, at each point on the surface. This set consists of an exhaustive sampling of images of the object from all viewpoints and lighting directions, which can be tens of thousands of images or more.

For vision applications, recovering and representing spatial reflectance is often made tractable by approximating it with an analytic BRDF model. This reduces spatial reflectance from a 4D function at each point to either a handful of parameters (see [2]–[4] for recent examples), or in the Lambertian case, to a single parameter at each point. The parametric approach is also common for image-based rendering applications in computer graphics [5]–[7]. From an acquisition standpoint, since they only require a few parameters at each point, these methods are able to provide reflectance estimates from a small number of input images. They require the selection of a specific parametric BRDF model *a priori*, however, which limits their accuracy and generality. For example, the Cook-Torrance model [8] represents some plastics and metals quite well but is unable to capture retro-reflection effects observed in other materials [9], [10].

Another approach—one that avoids the restrictions of parametric reflectance models—is purely data-driven and relies on densely sampling the BRDF at each point. This approach is more com-

mon in graphics applications, in which hundreds [11], thousands [12] or tens of thousands [13] of images are used to capture the reflectance information for a given object.¹ While these data-driven appearance models are useful for image-based rendering applications, they are too cumbersome to be used for image analysis, and the high cost of acquisition means that they cannot be used, for example, for large-scale enrollment in face/object recognition databases.

In this paper, we present an alternative approach that moves toward the acquisition of non-parametric spatial reflectance from a small number of images. This approach seeks to combine the benefits of both parametric methods (i.e., sparse images) and purely data-driven, non-parametric methods (i.e., arbitrary reflectance functions.) The central idea is that SBRDF estimation can be framed as a scattered-data interpolation problem, with images providing dense 2D slices of data embedded in the mixed spatial and angular domain. This approach is very different from previous non-parametric techniques (e.g., [11]–[13]) that interpolate reflectance only in the angular dimensions, estimating a unique reflectance function at each point. Instead, we simultaneously interpolate in both the spatial and angular dimensions. This enables a controlled exchange between spatial and angular information, effectively giving up some of the spatial resolution in order to fill the holes between sparsely observed view and illumination conditions. Since reflectance (especially the specular component) typically varies slowly from point to point over much of an object’s surface, this means that we can often obtain good results from a drastically reduced set of images.

We first present a theoretical analysis showing how the angular accuracy of the SBRDF can be increased in exchange for a reduction in its spatial resolution. We then provide a practical solution to the SBRDF estimation problem by introducing: i) a new parameterization of the BRDF domain, and ii) a non-parametric representation of reflectance based on radial basis functions (RBFs). In addition to providing an efficient solution to our SBRDF interpolation problem, this representation has two benefits. First, the resulting appearance model degrades gracefully (in the sense of providing visually-pleasing results) as the number of input images decreases. It can even be applied in the limiting case of a single input image. Second, it is relatively dimension-independent and can therefore be adapted to represent homogeneous BRDF data, making it

¹These methods actually capture only a subset of the reflectance information, since even these large numbers of images are insufficient to densely sample the complete 4D BRDF at each point.

useful for image-based BRDF acquisition systems [10], [14], [15].

II. MAIN IDEAS AND RELATED WORK

The proposed method for SBRDF estimation builds on three principal observations.

Smooth Spatial Variation. Most existing methods recover a unique BRDF at each point and thereby provide an SBRDF with very high spatial resolution. Many parametric methods have demonstrated, however, that the number of input images can be reduced if one is willing to accept a decrease in spatial resolution. This has been exploited, for example, by Yu et al. [7], Georgiades [2], and Hara et al. [4], who assume that specular BRDF parameters are constant across a surface.² (The morphable model introduced by Blanz and Vetter [3] relies on this assumption as well.) Similarly, Sato et al. [6] estimate the specular parameters at only a small set of points, later interpolating these parameters across the surface. Lensch et al. [16] present a novel technique in which reflectance samples at clusters of surface points are used to estimate a basis of (1-lobe) Lafortune models. The reflectance at each point is then uniquely expressed as a linear combination of these basis BRDFs.

Similar to these approaches, our method trades spatial resolution for an increase in angular resolution. The difference, however, is that we implement this exchange using a *non-parametric* representation. We begin by assuming that the SBRDF varies smoothly in the spatial dimensions, but we also demonstrate how this can be relaxed to handle rapid spatial variation in terms of a multiplicative texture. (In cases where the shape of the BRDF itself changes rapidly, we currently assume that discontinuities are given as input.)

Curved Surfaces. Techniques for image-based BRDF measurement [10], [14], [15] exploit the fact that a single image of a curved, homogeneous surface represents a very dense sampling of a 2D slice of the 4D BRDF. In this paper, we extend this idea to the spatially-varying case, where an image provides a 2D slice in the higher-dimensional SBRDF domain. Our results demonstrate that, like the homogeneous case, surface curvature (along with smooth spatial variation) can be exploited to increase the angular resolution of the SBRDF.³

²When valid, this assumption also enables the recovery of additional information, such as shape and/or illumination [2], [4].

³For near-planar surfaces where curvature is not available, more angular reflectance information can be obtained using near-field illumination and perspective views [17].

Angular Compressibility. While it is a multi-dimensional function, a typical BRDF varies slowly over much of its angular domain. This property has been exploited for 3D shape reconstruction [18], efficient BRDF acquisition [15], and efficient image synthesis with complex BRDFs [19]–[22]. Here, we exploit compressibility by assuming that the BRDF typically varies rapidly only in certain dimensions, such as the half-angle.

The three ideas discussed in this section have been developed in very different contexts. This paper combines and expands these ideas in order to solve a novel problem: estimating non-parametric SBRDFs from sparse images. Section III discusses the fusion of these ideas in terms of Fourier theory, which provides intuition for the trade-off between spatial and angular resolution. In a more practical sense, a computational method for recovering SBRDFs is realized using the BRDF parameterization of Sect. IV and an interpolation approach that unifies the treatment of spatial and angular dimensions (Sects. V–VII). An abridged version of the practical aspects of this paper (Sects. IV–VII) appears in an earlier conference paper [23], where it is presented in the context of computer graphics.

A. Assumptions

We exploit scene geometry to reduce the number of input images required to accurately represent appearance. Thus, unlike pure light field techniques [24]–[26], the method requires a set of images of an object with known geometry, known viewpoint, and point-source or directional illumination. Many suitable acquisition systems have been presented (e.g., [6], [12], [13], [27]).

In addition, global effects such as sub-surface scattering and interreflection are not explicitly considered in our formulation. For directional illumination and orthographic views, however, some of these effects will be absorbed into our representation and can be reproduced when rendered under the same conditions. (See Sect. VIII.) In this case, our use of the term *SBRDF* is synonymous with the *non-local reflectance field* defined by Debevec et al. [12].

Finally, in this paper we restrict our attention to isotropic BRDFs. While the ideas of exploiting spatial coherence and using RBFs to interpolate scattered reflectance samples can be applied to the anisotropic case, this would require a parameterization which is different from that presented in Sect. IV and is left for future work.




Fig. 1. (a) Input/output parameterization of the BRDF domain defined in a local coordinate system defined by the surface normal and tangent vector. (b) Halfway/difference parameterization of Rusinkiewicz. In the isotropic case, the BRDF domain is parameterized by $(\theta_h, \phi_d, \theta_d)$ [28].

B. Notation and Terminology

The SBRDF is a function of six dimensions and is written $f(\vec{x}, \vec{\theta})$, where $\vec{x} = (x, y) \in \mathbb{R}^2$ is the pair of spatial coordinates that parameterize the surface geometry (a surface point is written $\vec{s}(x, y)$), and $\vec{\theta} \in \Omega \times \Omega$ are the angular coordinates that parameterize the double-hemisphere of view/illumination directions in a local coordinate frame defined on the tangent plane at a surface point (i.e., the BRDF domain). A common parameterization of the BRDF domain is $\vec{\theta} = (\theta_i, \phi_i, \theta_o, \phi_o)$, which represent the spherical coordinates of the light and view directions in the local frame (see Fig. 1(a).) When the BRDF is isotropic, the angular variation reduces to a function of three dimensions, commonly parameterized by $(\theta_i, \theta_o, \phi_o - \phi_i)$. As mentioned above, we restrict ourselves to this isotropic case in this paper and consider the SBRDF to be a function defined on a 5D domain. In the special case when the SBRDF is a constant function of the spatial dimensions (i.e., $f(\vec{x}, \vec{\theta}) = f(\vec{\theta})$) we say that the surface is *homogeneous* and is described by a 4D (or isotropic 3D) function. In contrast, a surface with significant spatial variation is said to have significant *texture*. A *Lambertian* surface is one whose SBRDF is a constant function of the angular dimensions (i.e., $f(\vec{x}, \vec{\theta}) = f(\vec{x})$).

The BRDF domain can be parameterized in many ways, and one good choice is Rusinkiewicz's halfway/difference parameterization [28], shown in Fig. 1(b). Using this parameterization in the isotropic case, the BRDF domain is $\vec{\theta} = (\theta_h, \phi_d, \theta_d) \subset [0, \frac{\pi}{2}) \times [0, \pi) \times [0, \frac{\pi}{2})$. (ϕ_d is restricted to $[0, \pi)$ since $\phi_d \mapsto \phi_d + \pi$ by reciprocity.) For glossy surfaces, specular peaks occur at small half-angles (i.e., $\theta_h \approx 0$), but variation with respect to the difference angle (θ_d) is governed primarily by Fresnel reflection and tends to be limited for small and moderate values of θ_d [29]. Since these common BRDF features are roughly aligned with the coordinate axes, the halfway/difference parameterization enables tabulated BRDFs to be significantly compressed [28].

III. THEORETICAL MOTIVATION: EXCHANGING SPATIAL AND ANGULAR RESOLUTION

A small set of images provides only a sparse sampling of the SBRDF. This section shows that there is a trade-off between spatial and angular resolution, and that the high spatial resolution of images can be used to increase the angular resolution of a recovered SBRDF. This trade-off is described in the frequency domain, using analysis that is conceptually similar to studies of light field sampling and light transport [30], [31]. While it is difficult to apply these signal-processing ideas to our problem directly in the frequency domain (see Sect. III-D), the intuition they provide motivates the practical algorithms described in subsequent sections.

For simplicity, in this section we restrict ourselves to the two-dimensional case (referred to as *flatland*) in which surfaces reduce to plane curves, and images are one-dimensional. Using the 2D analog of the isotropic halfway/difference parameterization, the SBRDF is written $f(x, \theta_h, \theta_d)$, with $\theta_h, \theta_d \in [-\frac{\pi}{2}, \frac{\pi}{2})$ referred to as the half-angle and difference-angle, respectively. Motivated by the observations of Sect. II-B, we further simplify analysis by assuming that the SBRDF is a constant function of θ_d and varies only with the half-angle.⁴ In this case, the BRDF is a function of one angular dimension and one spatial dimension and is written $f(x, \theta_h)$, where

$$\theta_h = \frac{\theta_i}{2} + \frac{\theta_o}{2},$$

with θ_i and θ_o denoting the local lighting and view directions (i.e., relative to the surface normal at each point). We emphasize that this assumption is made for purely illustrative purposes and is not required for the practical methods developed later.

A. Planar surfaces

We begin by considering a planar surface (a line in flatland), with distant light sources and an orthographic viewer. As shown in Fig. 2(a), the surface is observed from a fixed view with directional light sources separated by $\Delta\theta$. Each image is acquired with illumination from a single source, and since the geometry of the surface and the view and illumination directions are known, each pixel in these images provides a direct sample of the unknown SBRDF.

Figure 3(a) shows the resulting sampling pattern in the two-dimensional SBRDF domain. Since the SBRDF is parameterized by the half-angle, the separation between samples in the

⁴In practice, this is often a good approximation for specular surfaces when $\theta_o, \theta_i < 60^\circ$ [29].




Fig. 2. Acquisition geometry for a planar (a) and curved (b) surfaces. For curved surfaces, emitted radiance is a function of α , which parameterizes the spatial location, and half-angle θ_h , which depends on the local lighting and view directions θ_i and θ_o .

angular dimension (along θ_h) is $\Delta\theta/2$. In addition, since each image provides a very dense (almost continuous) sampling of the spatial dimension, the separation between samples in that dimension is generally very small, so that $\Delta\theta \gg \Delta x$.

To investigate this sampling pattern in the frequency domain, we define the 2D Fourier transform in the conventional way⁵ (and with $I = \sqrt{-1}$ as usual),

$$F(\Omega_u, \Omega_v) = \int \int f(u, v) e^{-2\pi I \Omega_u u} e^{-2\pi I \Omega_v v} du dv,$$

where we are interested in the spatially-varying BRDF $f(x, \theta_h)$ with Fourier transform $F(\Omega_x, \Omega_\theta)$.

Sampling in the spatial domain is represented by multiplication by the comb function

$$f_{sampled}(x, \theta_h) = f(x, \theta_h) \sum_{n_1, n_2} \delta(x - n_1 \Delta x) \delta(\theta_h - n_2 \Delta\theta/2),$$

where n_1 and n_2 are integers. The Fourier transform of this sampled signal is obtained by convolution with a comb function, which yields

$$F_{sampled}(\Omega_x, \Omega_\theta) = \sum_{n_1, n_2} F\left(\Omega_x - \frac{2\pi n_1}{\Delta x}, \Omega_\theta - \frac{4\pi n_2}{\Delta\theta}\right).$$

In the frequency domain, the sampling process produces shifted replicas of the SBRDF spectrum. These replicas are at intervals of $2\pi/\Delta x$ spatially and $4\pi/\Delta\theta$ in the angular dimension. This is shown in Fig. 3(b), assuming a sufficiently band-limited SBRDF.

In order to synthesize novel images of the surface (i.e., to predict its appearance), the SBRDF must be reconstructed from its samples. Accurate reconstruction that avoids aliasing during

⁵In general, one would need to use a Fourier series for periodic angular values and a continuous infinite domain for the spatial component. However, the core analysis is not affected by these issues, and in what follows, we omit the limits of integration, allowing whichever interpretation is more convenient for the reader.

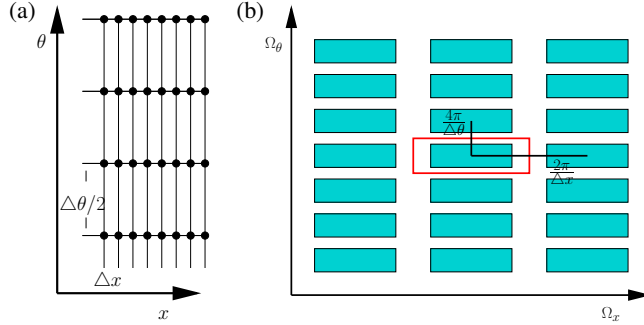


Fig. 3. (a) Spatial/angular sampling pattern for SBRDF acquisition using a planar surface. Sample spacing is high along the spatial dimension (Δx) but lower along the angular dimension ($\Delta\theta/2$). (b) Fourier spectrum of the sampled SBRDF. Green rectangles correspond to the original spectrum, which is replicated along spatial and angular frequencies. The red box in the middle corresponds to the ideal reconstruction filter.

image synthesis requires a sampling rate satisfying the Nyquist conditions. Letting Ω_x^{\max} and Ω_θ^{\max} denote the spatial and angular bandwidths⁶ of the SBRDF, we require

$$\Omega_x^{\max} < \frac{\pi}{\Delta x} \quad \Omega_\theta^{\max} < \frac{2\pi}{\Delta\theta}. \quad (1)$$

These sampling conditions are quite intuitive. The maximum spatial frequency of the surface cannot exceed the pixel resolution of the camera, and the maximum angular frequency (which is a measure of the surface ‘shininess’) cannot exceed the sampling resolution of the light sources. Furthermore, the optimal reconstruction filter in the Fourier domain is the elongated box shown in red in Fig. 3(b), which simply band-limits differently in the spatial and angular dimensions according to the relations above. Intuitively, when the SBRDF is reconstructed, each pixel is treated separately, and the ‘sharpness’ of the recovered specular highlights is bounded by the number of light source locations as noted in previous non-parametric acquisition methods [11]–[13]. This result is the foundation for the analysis of curved surfaces, discussed next.

B. Curved surfaces

Acquisition for curved surfaces (plane curves in flatland) is shown in Fig. 2(b). Assuming constant curvature, we can parameterize the spatial dimension x by angle α , since there is a linear relationship between the two. The goal is to reconstruct the spatially-varying BRDF $f(\alpha, \theta_h)$ from the samples obtained from a set of images similar to those in the previous section.

⁶We assume rectangular bandlimits for simplicity. There does not appear to be a physical basis for assuming, for example, an elliptical spectrum, nor would it significantly affect the results.




Fig. 4. (a) Sampling pattern in the SBRDF domain for acquisition using a curved surface. Unlike the planar case of Fig. 3(a), the sampling pattern is oriented at 45° to the axes, as per (III-B). Spacing in the two directions is $\sqrt{2}\Delta\alpha$ and $\Delta\theta/(2\sqrt{2})$ with $\Delta\theta \gg \Delta\alpha$. (b) The comb function (in the frequency domain) is described by (2). Sampling the SBRDF according to (a) corresponds to convolution in the frequency domain by this comb function, which has the same rotated form.

To derive an expression for θ_h in terms of the global positions of the camera and light source (θ_c and θ_l), we use the convention that α is positive in the direction shown in Fig. 2(b). Then,

$$\theta_h = \frac{\theta_l - \alpha + \theta_c - \alpha}{2} = \frac{\theta_l + \theta_c}{2} - \alpha = h - \alpha,$$

where $h = (\theta_l + \theta_c)/2$ can be thought of as the ‘global half-angle’. Thus, we can re-parameterize⁷ the spatially-varying BRDF:

$$f(\alpha, \theta_h) = f(\alpha, h - \alpha).$$

Figure 4(a) depicts the resulting sampling pattern in the SBRDF domain. Assuming uniformly-spaced samples on the surface (but see Sect. III-D), the spacing between samples is denoted $\Delta\alpha$ and is analogous to Δx in the planar case. The spacing between successive global half-angles (Δh) is simply $\Delta\theta/2$, just as in the planar case, and as before, $\Delta h \gg \Delta\alpha$. The sampling pattern is now rotated, however, as shown in Fig. 4(a). In this figure, each image corresponds to a line with negative slope and a y -intercept at h . Along each of these image lines, the net sample spacing is $\sqrt{2}\Delta\alpha$. The perpendicular distance between lines is $\Delta h/\sqrt{2} = \Delta\theta/(2\sqrt{2})$. As we now show, this shearing pattern leads to interesting consequences in the frequency domain.

In the SBRDF domain, the continuous signal is multiplied by a comb function with the pattern shown in Fig. 4(a). The effect in the frequency domain is described by convolution with a similar comb function, but with different periods along the two axes. To determine these periods, we note that the sampling pattern in Fig. 4(a) is obtained by rotating the planar sampling

⁷It is important to note that while the half-angle parameterization is helpful in representing real BRDFs, it is not critical to enable reflectance sharing. For example, if the SBRDF were a function of $f(\alpha, \theta_i)$ instead of $f(\alpha, \theta_h)$, we would simply write $\theta_i = \theta_l - \alpha$, with h replaced by θ_l in (III-B), leading to essentially the same form.

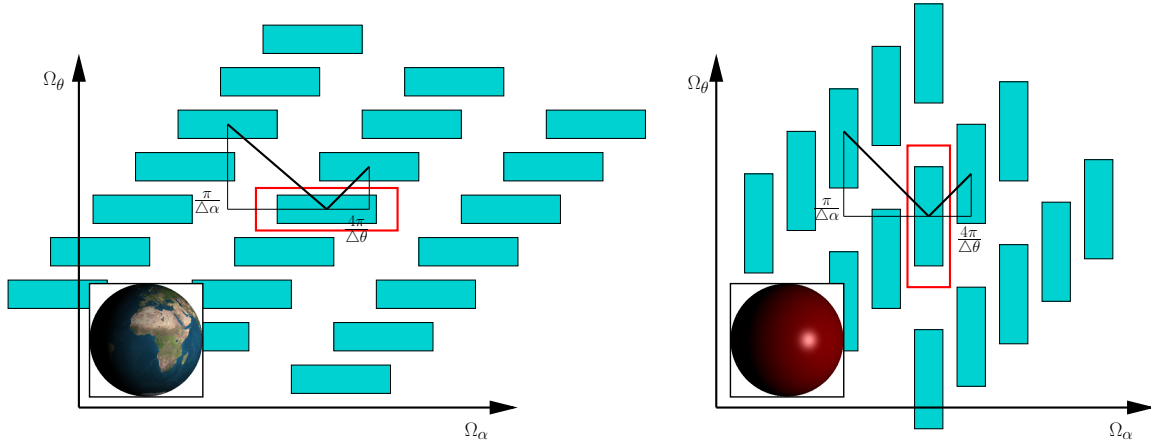


Fig. 5. Fourier transforms of sampled SBRDFs obtained from observations of curved surfaces. Green rectangles correspond to the original frequency spectra for predominantly diffuse (left) and specular (right) reflectance, and central red rectangles are the ideal reconstruction filters. Insets contain examples of diffuse and specular surfaces under directional illumination.

pattern (Fig. 3(a)) by $+45^\circ$. The linear transformation theorem [32] then tells us that the Fourier transform of the planar comb function will also be rotated by $+45^\circ$, and that the sampling rates will be scaled appropriately. More formally, the comb function in the frequency domain is (with m_1 and m_2 integers),

$$COMB(u, v) = \frac{4\sqrt{2}\pi m_1}{\Delta\theta} \left(\frac{u}{\sqrt{2}} + \frac{v}{\sqrt{2}} \right) + \frac{\sqrt{2}\pi m_2}{\Delta\alpha} \left(-\frac{u}{\sqrt{2}} + \frac{v}{\sqrt{2}} \right), \quad (2)$$

and the Fourier transform of the sampled curved-surface SBRDF is

$$F_{sampled}(\Omega_\alpha, \Omega_\theta) = \sum_{m_1, m_2} F\left(\Omega_\alpha - \frac{4\pi m_1}{\Delta\theta} + \frac{\pi m_2}{\Delta\alpha}, \Omega_\theta - \frac{4\pi m_1}{\Delta\theta} - \frac{\pi m_2}{\Delta\alpha}\right). \quad (3)$$

To provide some intuition, consider the set of replicas with $m_2 = 0$. In this case, the replicas indexed by m_1 lie on a 45° line with a sample spacing of $4\sqrt{2}\pi/\Delta\theta$ as calculated earlier. Replicas with $m_2 \neq 0$ lie on parallel lines, and since the spatial sampling rate is typically large relative to the angular sampling rate, these lines are well separated. This pattern of replicas—analogue to Fig. 3(b)—is shown in the left of Fig. 5. Note from this figure that the conditions for alias-free synthesis remain similar to the planar case, as does the optimal reconstruction filter. In particular, given sampling rates $\Delta\alpha$ and $\Delta\theta$ the conditions for alias-free image synthesis are

$$\Omega_\alpha^{\max} < \frac{\pi}{2\Delta\alpha} \quad \Omega_\theta^{\max} < \frac{2\pi}{\Delta\theta}. \quad (4)$$

C. Reflectance Sharing

As in the planar case, the bounds in (4) suggest that we can accurately recover an SBRDF from a small number of images (i.e., a large $\Delta\theta$) provided that there is only low-frequency variation

in the angular dimension. A synthetic example of an object that satisfies this criteria is shown in the inset on the left of Fig. 5, where even though there is high-frequency spatial variation (significant texture), the reflectance at each point is quite diffuse (i.e., it is nearly Lambertian).

What is very different from the planar case is that in addition to diffuse surfaces like that shown in the left of Fig. 5, we can also recover reflectance like that shown in the inset on the right of Fig. 5. This surface is characterized by low-frequency spatial variation (little texture) and high-frequency angular variation (sharp highlights). The frequency-domain effects of sampling this type of SBRDF are depicted in the graph on the right of Fig. 5. The spectrum of the SBRDF is a rectangle that is elongated in the angular frequency dimension, and the optimal reconstruction filter (shown in red) maintains reflects this shape. The limits on the SBRDF spectrum for alias-free synthesis are:

$$\Omega_{\alpha}^{\max} < \frac{2\pi}{\Delta\theta} \quad \Omega_{\theta}^{\max} < \frac{\pi}{2\Delta\alpha}. \quad (5)$$

Here, the maximum spatial bandwidth (i.e., in the α -dimension) is constrained by the separation of light source directions $\Delta\theta$, while the angular bandwidth is constrained by the spatial sampling $\Delta\alpha$. This is in direct opposition to the limits of (4), and it means, for example, that we can recover very sharp specular highlights from a *single* image (with dense spatial sampling) like that shown on the right of Fig. 5. The optimal reconstruction filter is a low-pass filter with a small cutoff frequency in the spatial dimension, which in the SBRDF domain, corresponds to an averaging over a large spatial extent. We refer to this averaging process as *reflectance sharing*, since it enables the use of the high spatial resolution available in images to increase the angular resolution of the recovered reflectance function.

The limits in (5) also indicate that the recovered spatial frequencies are constrained by the separation of light sources. Given only a single image, for example, the spatial bandwidth is limited to zero (since $\Delta\theta \rightarrow \infty$), and we can only recover accurate reflectance for *homogeneous* surfaces without texture. Information about spatial variation or inhomogeneity in the SBRDF can only be obtained with additional images that provide observations of specular highlights in multiple regions of the surface.

This spatial/angular trade-off can be further understood by noting that the product of the bounds on bandwidth in the spatial and angular dimensions is always given by

$$\Omega_{\alpha}^{\max} \Omega_{\theta}^{\max} < \frac{\pi^2}{\Delta\alpha \Delta\theta}. \quad (6)$$

This expression tells us that we can recover high SBRDF frequencies in one of the spatial or angular dimensions from a sparse set of images, but not both simultaneously. Conventional SBRDF acquisition methods [11]–[13] present only one possible approach. They treat each spatial sample separately and provide very high spatial resolution with limited angular resolution. By contrast, we seek to share reflectance spatially, estimating an SBRDF with high angular frequency at the cost of a decrease in spatial resolution.

It is also interesting to note that if the surface curvature k is explicitly represented through the relation $\alpha = kx$, the bandwidth limits are given by $\Omega_\alpha^{\max} < 2\pi k/\Delta\theta$ and $\Omega_\theta^{\max} < \pi/k\Delta\alpha$ (and (6) still holds). Thus, as long as there is some curvature ($k > 0$), we can recover very high angular frequencies, provided the spatial frequencies are strongly bandlimited.

Finally, we emphasize that glossy reflectance like that in on the right of Fig. 5 generally cannot be recovered from sparse images in the planar case.⁸ Indeed, given the sampling pattern of Fig. 3(a), the repeated spectra for a glossy surface would overlap. It is the surface curvature that causes the rotated patterns in Fig. 3, which in turn allows the recovery of high-frequency angular reflectance from sparse images.

D. Images and Irregular Sampling

The analysis of this section relies on the assumption that samples in the spatial dimension are uniformly spaced on the surface. For image-based data, however, samples are uniformly spaced *on the image-plane*, and for curved objects this leads to a very irregular sampling of the object’s surface. Thus, while Fourier analysis provides insight and motivation for the reflectance sharing approach, it is difficult to implement a practical frequency-based method using these ideas. Instead, the remainder of this paper develops a scattered data interpolation method using radial basis functions, with weights to trade off the extent of sharing in the spatial and angular dimensions. In the context of the preceding analysis, this is analogous to building a reconstruction filter by selecting appropriate cut-off frequencies in the spatial and angular dimensions.

IV. SBRDF PARAMETERIZATION

At the core of our approach is the interpolation of scattered data in multiple dimensions, the success of which depends on how the SBRDF is parameterized. The previous section used

⁸Again, using perspective cameras and near-field illumination, angular information *can* be obtained in the planar case [17].




Fig. 6. The mapping in (7) creates a BRDF parameterization suitable for interpolation. The BRDF $f(u, v, w)$ is guaranteed to satisfy reciprocity; the parameterization is defined for all values of θ_h ; BRDF samples from a single image lie on a plane of constant w ; and specular events are clustered near the w -axis, enabling significant compression.

the special case of plane curves to describe how the assumption of angular compressibility and the halfway/difference parameterization are important for correctly framing the problem. Inspired by this theory, this section considers the three-dimensional case, and introduces a new parameterization for the angular dimensions of an SBRDF. Based on this parameterization, our interpolation technique is discussed in Sects. VI and VII.

The halfway/difference BRDF parameterization of Rusinkiewicz [28] was described in Sect. II-B. The existence of singularities at $\theta_h = 0$ and $\theta_d = 0$ and the required periodicity ($\phi_d \mapsto \phi_d + \pi$) make this parameterization unsuitable for most interpolation techniques, and in order to address this, we define a mapping $(\theta_h, \phi_d, \theta_d) \mapsto (u, v, w)$, as

$$(u, v, w) = \left(\sin \theta_h \cos 2\phi_d, \sin \theta_h \sin 2\phi_d, \frac{2\theta_d}{\pi} \right), \quad (7)$$

which is shown in Fig. 6. This mapping defines a new parameterization for the angular dimensions of an SBRDF (i.e., the BRDF domain.) It eliminates the singularity at $\theta_h = 0$ and ensures that the BRDF $f(u, v, w)$ satisfies reciprocity. In addition, the mapping is such that the remaining singularity occurs at $\theta_d = 0$, where the light and view directions are equivalent. This is desirable because this configuration is difficult to create in practice, which makes it unlikely to occur during acquisition. During synthesis, however, it must be handled with care.

A. Considerations for Image-based Acquisition

As mentioned previously, the halfway/difference parameterization increases compression rates since common features such as specular and retro-reflective peaks are aligned with the coordinate axes [28]. The modified parameterization of (7) maintains this property, since specular events cluster along the w -axis, and retro-reflective peaks occur in the plane $w = 0$.

These parameterizations are useful for image-based data for an additional reason: they separate the sparsely- and densely-sampled dimensions of the BRDF. (Marschner's parameterization [33]

also shares this property.) To see this, note that for orthographic projection and distant lighting—or more generally, when scene relief is relatively small—a single image of a curved surface provides BRDF samples lying in a plane of constant θ_d , since this angle is independent of the surface normal. Indeed, for the flatland case of Sect. III-B, the difference angle is

$$\theta_d = \frac{\theta_c - \alpha - \theta_l + \alpha}{2} = \frac{\theta_c - \theta_l}{2}, \quad (8)$$

which depends on the global camera and light positions but not the spatial location α .

While each (orthographic) image provides only one sample of the θ_d dimension, it represents a nearly continuous sampling of θ_h and ϕ_d . As a result, a *set* of images provides dense sampling of (θ_h, ϕ_d) but only as many samples of θ_d as there are images.⁹ Conveniently, this irregular sampling obtained from image-based data corresponds well with the behavior of general BRDFs, which vary slowly in the sparsely sampled θ_d -dimension, especially when θ_d is small [29]. At the same time, by imaging curved surfaces, we ensure that the sampling rate of the half-angle θ_h is high enough to accurately recover the high-frequency variation (e.g., due to specular highlights) that is generally observed in that dimension.

V. SCATTERED DATA INTERPOLATION

Recall that our goal is to estimate a continuous SBRDF $f(\vec{x}, \vec{\theta})$ from a set of samples $f_i \in \mathbb{R}^5$ drawn from images of a surface with known geometry. Our task is complicated by the fact that, as discussed in Sect. III-D, the input samples are very non-uniformly distributed.

There are many methods for interpolating scattered data in this relatively high-dimensional space, but for our problem, interpolation using radial basis functions provides the most attractive choice. Given a set of samples, an RBF interpolant is computed by solving a linear system of equations, and the existence and uniqueness is guaranteed with few restrictions on the sample points. Thus, unlike homogeneous BRDF representations such as spherical harmonics, Zernike polynomials, wavelets, and the basis of Matusik *et al.* [15], an RBF representation does not require a local preprocessing step to resample the input data at regular intervals.

Additionally, for a fixed number of samples, the required computation and the resulting size of an RBF representation grow relatively slowly as the dimension increases. This is in contrast

⁹The orthographic/directional case is considered for illustrative purposes; it is not required by the method.

to methods such as piecewise polynomial splines (e.g., [34]) and local methods like polynomial regression (e.g., [14]) and the *push/pull* algorithm of Gortler et al. [24]. These alternative methods require either a triangulation of the domain or a tabulation of function values, both of which become computationally prohibitive in high dimensions. (Jaroszkiewicz and McCool [20] handle this by approximating the high-dimensional SBRDF by a product of 2D functions, each of which is triangulated independently.)

A. Radial Basis Functions

To briefly review RBF interpolation (see, e.g., [35], [36]), consider a general function $g(\vec{x})$, $\vec{x} \in \mathbb{R}^d$ from which we have N samples $\{g_i\}$ at sample points $\{\vec{x}_i\}$. This function is approximated by a sum of a low-order polynomial and a set of scaled, radially symmetric basis functions centered at the sample points:

$$g(\vec{x}) \approx \tilde{g}(\vec{x}) = p(\vec{x}) + \sum_{i=1}^N \lambda_i \psi(\|\vec{x} - \vec{x}_i\|), \quad (9)$$

where $p(\vec{x})$ is a polynomial of order n or less, $\psi : \mathbb{R}^+ \rightarrow \mathbb{R}$ is a continuous function, and $\|\cdot\|$ is the Euclidean norm. The sample points \vec{x}_i are referred to as *centers*, and the RBF interpolant \tilde{g} satisfies the interpolation conditions $\tilde{g}(\vec{x}_i) = g(\vec{x}_i)$.

Given a choice of n , an RBF ψ , and a basis for the polynomials of order n or less, the coefficients of the interpolant are given by the solution of the linear system

$$\begin{bmatrix} \Psi & P \\ P^\top & 0 \end{bmatrix} \begin{bmatrix} \vec{\lambda} \\ \vec{c} \end{bmatrix} = \begin{bmatrix} \vec{g} \\ 0 \end{bmatrix}, \quad (10)$$

where $\Psi_{ij} = \psi(\|\vec{x}_i - \vec{x}_j\|)$, $\vec{\lambda}_i = \lambda_i$, $\vec{g}_i = g_i$, $P_{ij} = p_j(\vec{x}_i)$ where $\{p_j\}$ are the polynomial basis functions, and $\vec{c}_i = c_i$ are the coefficients in this basis of the polynomial term in \tilde{g} . This system is invertible (and the RBF interpolant is uniquely determined) in arbitrary dimensions for many choices of ψ , with only mild conditions on n and the locations of the data points [37], [38].

In practical cases the samples $\{g_i\}$ are affected by noise, and it is desirable to allow the interpolant to deviate from the data points, balancing the smoothness of the interpolant with its fidelity to the data. This is accomplished by replacing Ψ in (10) with $\Psi - \rho N\mathbf{I}$, where \mathbf{I} is the identity matrix and ρ is a stiffness parameter. Further details are provided by Wahba [39].

In many cases, we can benefit from using *radially asymmetric* basis functions which are stretched in certain directions, and here we use them to: 1) control the trade-off between spatial

and angular resolution, and 2) manage the irregularity in our sampling pattern. (Recall that the (u, v) dimensions are sampled almost continuously while we have only as many samples of w as we have images.) Following Dinh et al. [40], an asymmetric radial function is created by scaling the Euclidean distance in (9) so that the basis functions become

$$\psi(\|M(\vec{x} - \vec{x}_i)\|), \quad (11)$$

where $M \in \mathbb{R}^{d \times d}$. In our case, we choose M to be a diagonal matrix whose entries reflect the expected relative rates of change in the spatial and angular dimensions, and the relative sampling rates of the three angular dimensions. (See Sects. VI and VII.)

When the number of samples is large ($N > 10,000$), solving (10) requires care and can be difficult—or impossible—using direct methods. This limitation has been addressed quite recently, and iterative fitting methods [41] and fast multipole methods (FMMs) for efficient evaluation [42] exist for many choices of ψ in many dimensions. In some cases, solutions for systems with over half a million centers have been reported [43]. The next sections include investigations of the number of RBF centers required to accurately represent image-based reflectance data, and we find this number to be sufficiently small to allow the use of direct methods.

VI. HOMOGENEOUS SURFACES

This section applies RBF interpolation to homogeneous surfaces, where we seek to estimate a global BRDF that is not spatially-varying. The derived representation may be useful for interpolating image-based BRDF data [10], [14], [15]. The theoretical analysis in section III shows that the variation with respect to the half-angle can be estimated from a single image of a homogeneous, curved surface. While more images are needed to acquire the full BRDF, this section shows that a sparse set often suffices.

As discussed in Sect. IV, for homogeneous BRDF data, reflectance is a function of three dimensions, (u, v, w) . In \mathbb{R}^3 , a good choice for ψ is the linear (or biharmonic) RBF, $\psi(r) = r$, with $n = 1$, since in this case, the interpolant from (10) exists for any non-coplanar data, is unique, minimizes a generalization of the thin-plate energy, and is therefore the smoothest in some sense [37], [43]. The BRDF is expressed as

$$\tilde{f}(\vec{\theta}) = c_1 + c_2u + c_3v + c_4w + \sum_{i=1}^N \lambda_i \|\vec{\theta} - \vec{\theta}_i\|, \quad (12)$$

where $\vec{\theta}_i = (u_i, v_i, w_i)$ represents a BRDF sample point obtained from the input images, and $\vec{\lambda}$ and \vec{c} are found by solving (10). As discussed in the previous section, we can benefit from the use of radially asymmetric basis functions, and for homogeneous BRDF data, this is accomplished using $M = \text{diag}(1, 1, m_w)$. For $m_w < 1$, the basis functions are elongated in the w dimension, which is appropriate since our sampling rate is much lower in that dimension. The appropriate value of this parameter depends on the angular density of the input images, and empirically we find that typical values for m_w are between 0.1 and 0.5.

As a practical consideration, since each pixel represents a sample point $\vec{\theta}_i$, even with modest image resolution, using all available samples as RBF centers is computationally prohibitive. Much of this data is redundant, however, and an accurate BRDF representation can be achieved using only a small fraction of these centers. A sufficient subset of centers could be chosen using knowledge of typical reflectance phenomena. (To represent sharp specular peaks, for example, RBF centers are generally required near $\theta_h = 0$.) Alternatively, Carr et al. [43] present an effective greedy algorithm for choosing this subset without assuming prior knowledge, and a slightly modified version of the same algorithm is applied here. The procedure begins by randomly selecting a small subset of the sample points $\vec{\theta}_i$ and fitting an RBF interpolant to these. Next, this interpolant is evaluated at all sample points and used to compute the radiance residuals, $\varepsilon_i = (f_i - \tilde{f}(\vec{\theta}_i)) \cos \theta_i$, where θ_i is the angle between the surface normal at the sample point and the illumination direction. Finally, points where ε_i is large are appended as additional RBF centers, and the process is repeated until the desired fitting accuracy is achieved.

It should be noted that an algorithmic choice of center locations could increase the efficiency of the resulting representation, since center locations would not necessarily need to be stored for each material. This would require assumptions about the function being approximated, however, and here we choose to emphasize generality over efficiency by using the greedy algorithm.

A. Evaluation

To evaluate the BRDF representation in (12), we perform comparisons to both parametric BRDF models and to a non-linear basis (the isotropic Lafortune model [44].) The models are fit to synthetic images of a sphere, and their accuracy is measured by their ability to predict the appearance of the sphere under novel conditions. (Other representations, such as wavelets and the Matusik bases are excluded from this comparison because they require dense, uniform samples.)




Fig. 7. Accuracy of the RBF representation as the number of centers is increased using a greedy algorithm. The input is 10 images of a sphere synthesized using the metallic-blue BRDF measured by Matusik et al. [15]. This is compared to the isotropic Lafortune representation [44] with an increasing number of lobes. Less than 1000 centers are sufficient to represent the available reflectance information using RBFs, whereas the limited flexibility of the Lafortune basis and the existence of local minima in the non-linear fitting process limit the accuracy of the Lafortune representation. (See text for details.)

The input images simulate data from image-based BRDF measurement systems like those in [10], [14], [15]. They are orthographic, directional-illumination images with a resolution of 100×100 , and are generated such that θ_d is uniformly sampled in $[0, \frac{\pi}{2}]$. The accuracy of the recovered models is measured by the relative RMS radiance error over 21 images—also uniformly spaced in θ_d —that are not used as input. For these simulations, we use both specular and diffuse reflectance, one drawn from measured data (the metallic-blue BRDF, courtesy of Matusik et al. [15]), and the other generated using the physics-based Oren-Nayar model [9].

Figure 7 shows the accuracy of increasingly complex RBF and Lafortune representations fit to ten input images. The complexity of the RBF representation is measured by the number of centers selected by the greedy algorithm, and that of the Lafortune model is measured by the number of generalized cosine lobes. An unusually large number of lobes are shown (two or three lobes is typical) so that the resulting Lafortune and RBF representations have comparable degrees of freedom. It is important to note, however, that the size of each representation is different for equivalent complexities; an N -lobe isotropic Lafortune model requires $3N + 1$ parameters, while an N -center RBF interpolant requires $4N + 4$.

Since the basis functions of the Lafortune model are designed for representing BRDFs and are therefore embedded with knowledge of general reflectance behavior, they provide a reasonably good fit with a small number of lobes. For example, a 6-lobe Lafortune model (19 parameters) yields the same RMS error as a 300-center RBF model (1204 parameters.) In addition to being




Fig. 8. *Top*: Error in the estimated BRDF for an increasing number of images of a metallic-blue sphere. As the number of images increases, the RBF representation (with 1000 centers) approaches the true BRDF, whereas the isotropic Ward model [45] and the Lafortune representation are too restrictive to provide an accurate fit. *Bottom*: Synthesized images using the three BRDF representations estimated from 12 input images, where the angle between the source and view directions is 140° .

compact, the Lafortune model has the advantage of being more suitable for direct rendering [5]. But the accuracy of this representation is fundamentally limited; the lack of flexibility and the existence of local minima in the required non-linear fitting process prevent the Lafortune model from accurately representing the reflectance information available in the input images.

In contrast, RBFs provide a general linear basis, and with a sufficient number of centers, they can represent any ‘smooth’ function with arbitrary accuracy (e.g., [36]). Here, the RBF representation converges with less than 1000 centers, suggesting that only a small fraction of the available centers are needed to summarize the reflectance information in the ten input images.

Similar conclusions are drawn from a second experiment in which we investigate the accuracy of these and other representations with a fixed level of complexity and an increasing number of input images. Results are shown in Figs. 8 and 9 for predominantly specular and diffuse reflectance. (Here, six lobes are used in the Lafortune representation since the results do not change significantly with additional lobes.) Since RMS error is often not an accurate perceptual metric, these figures also include synthetic spheres rendered with the recovered models. This experiment demonstrates the flexibility of the RBF representation, which captures both the Fresnel effects in Fig. 8 and the retro-reflection in Fig. 9. Parametric models do not typically afford this flexibility—while it may be possible to find a parametric model that fits a specific

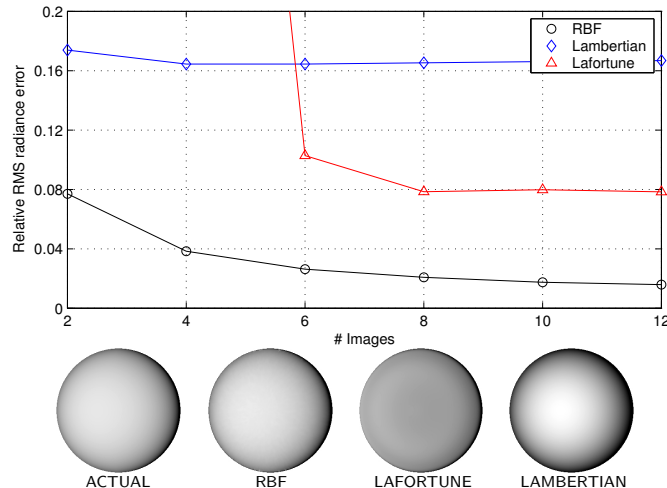


Fig. 9. *Top*: Error in the estimated BRDF for an increasing number of input images of a diffuse Oren-Nayar sphere. Again, the 1000-center RBF representation approaches the true BRDF, whereas the Lafortune and Lambertian BRDF models are too restrictive to accurately represent the data. *Bottom*: Synthesized images comparing the three BRDF representations estimated from 12 input images, where the angle between the source and view directions is 10° .

BRDF quite well, it is very difficult to find a model that accurately fits general BRDFs.

VII. INHOMOGENEOUS SURFACES

The previous section suggests that RBFs can provide a useful representation for homogeneous BRDFs. In this section, we show that essentially the same representation can be used to handle spatially-varying reflectance as well. It enables the exchange between spatial and angular resolution (as described in Sect. III), and drastically reduces the number of required input images. We begin by assuming that the 5D SBRDF varies slowly in the spatial dimensions, and in the next section, we show how this can be generalized to handle rapid spatial variation in terms of a multiplicative albedo or texture.

In the homogeneous case, the BRDF is a function of three dimensions, and the linear RBF $\psi(r) = r$ yields a unique interpolant that minimizes a generalization of the thin-plate energy. Although optimality cannot be proved, this RBF has shown to be useful in higher dimensions as well, since it provides a unique interpolant in any dimension for any n [35]. In the spatially-varying case, the SBRDF is a function of five dimensions, and we let $\vec{q} = (x, y, u, v, w)$ be a point in its domain. Using the linear RBF with $n = 1$, the SBRDF is given by

$$\tilde{f}(\vec{q}) = p(\vec{q}) + \sum_{i=1}^N \lambda_i \|\vec{q} - \vec{q}_i\|, \quad (13)$$

where $p(\vec{q}) = c_1 + c_2x + c_3y + c_4u + c_5v + c_6w$.

We can use any parameterization of the surface \vec{s} , and there has been significant recent work on determining good parameterizations for general surfaces (e.g., [46], [47]). The ideal surface parameterization is one that preserves distance, meaning that $\|\vec{x}_1 - \vec{x}_2\|$ is equivalent to the geodesic distance between $\vec{s}(\vec{x}_1)$ and $\vec{s}(\vec{x}_2)$. For simplicity, here we treat the surface as the graph of a function, with $\vec{s}(x, y) = (x, y, s(x, y))$, $(x, y) \in [0, 1] \times [0, 1]$.

The procedure for recovering the parameters in (13) is almost exactly the same as in the homogeneous case. The coefficients of \tilde{f} are found by solving (10) using a subset of the input SBRDF samples, and this subset is chosen using a greedy algorithm. Radially asymmetric basis functions are realized using $M = \text{diag}(m_{xy}, m_{xy}, 1, 1, m_w)$, where m_{xy} controls the exchange between spatial and angular reflectance information. When $m_{xy} \ll 1$, the basis functions are elongated in the spatial dimensions, and the recovered reflectance function approaches a single BRDF (i.e., a homogeneous representation) with rapid angular variation. When $m_{xy} \gg 1$, we recover a near-Lambertian representation in which the BRDF at each point approaches a constant function of $\vec{\theta}$. Appropriate values of m_{xy} depend on the choice of surface parameterization, and we found typical values to be between 0.2 and 0.4 for the examples in this paper.

A. Evaluation

The SBRDF representation of (13) can be evaluated using experiments similar to those for the homogeneous case. Here, spatial variation is simulated using images of a hemisphere with a Cook-Torrance BRDF [8] with a linearly varying roughness parameter. Five images are shown in Fig. 10, and they demonstrate how the highlight sharpens from left to right across the surface.

The graph in Fig. 10 shows the accuracy of the recovered SBRDF as a function of the number of RBF centers when it is fit to images of the hemisphere under ten uniformly distributed illumination directions. The error is computed over 40 images that are not used as input. Fewer than 2000 centers are needed to accurately represent the spatial reflectance information available in the input images, which is a reasonably compact representation requiring roughly 12,000 parameters. For comparison, an SBRDF representation for a 10,000-vertex surface consisting of two unique Lafortune lobes at each vertex is roughly five times as large.

Figure 11 contrasts reflectance sharing with conventional methods that interpolate only in the angular dimensions, estimating a separate BRDF at each point. This ‘no sharing’ approach is




Fig. 10. *Top*: Accuracy of the SBRDF recovered by reflectance sharing using the RBF representation in (13) as the number of centers is increased using a greedy algorithm. The input is 10 synthetic images of a hemisphere (five of which are shown) with linearly varying roughness.

used by Matusik et al. [13], and is similar in spirit to Wood et al. [11], who also estimate a unique view-dependent function at each point. (In this discussion, angular interpolation in the BRDF domain is assumed to require known geometry, which is different from lighting interpolation (e.g., [12]) that does not. For the particular example in Fig. 11, however, the ‘no sharing’ result can be obtained without geometry, since it is a special case of fixed viewpoint.)

For both the reflectance sharing and ‘no sharing’ cases, the SBRDF is estimated from images with fixed viewpoint and uniformly distributed illumination directions such as those in Fig. 10, and the resulting SBRDF is used to predict the appearance of the surface under novel lighting. The top frame of Fig. 11 shows the actual appearance of the hemisphere under five novel conditions, and the lower frames show the reflectance sharing and ‘no sharing’ results obtained from increasing numbers of input images. Note that many other methods—most notably that of Lensch et al. [16]—are excluded from this comparison because they require the selection of a specific parametric model and therefore suffer from the limitations discussed in Sect. VI.

In this example, reflectance sharing reduces the number of required input images by more than an order of magnitude. Five images are required for good visual results using the RBF representation, whereas at least 150 are needed if one does not exploit spatial coherence. Figure 11 also shows how differently the approaches degrade with sparse input. Reflectance sharing provides a smooth SBRDF whose accuracy gradually decreases away from the convex




Fig. 11. Estimating the spatially-varying reflectance function from a sparse set of images. *Top frame*: five images of a hemisphere under illumination conditions not used as input. *Middle frame*: appearance predicted by reflectance sharing with two and five input images. (The input images are shown in Fig. 10; the two left-most images are used for the two-image case.) *Bottom frame*: appearance predicted by interpolating only in the angular dimensions with 5, 50 and 150 input images. At least 150 images are required to obtain a result comparable to the five-image reflectance sharing result.

hull of input samples. (For example, the sharp specularity on the right of the surface is not accurately recovered when only two input images are used.) In contrast, when interpolating only in the angular dimensions, a small number of images provides only a small number of reflectance samples at each point; and as a result, severe ‘ghosting’ occurs when the surface is illuminated by high-frequency environments like the directional illumination shown here. This is easy to understand in terms of the theoretical analysis in section III. In effect, angular interpolation uses a filter of the form shown in Fig. 5(a) to reconstruct a spectrum of the form of Fig. 5(b), which leads to considerable aliasing.

Even when the input images are captured from a single viewpoint, our method recovers a full SBRDF, and as shown in Fig. 12, view-dependent effects can be predicted. This is made possible by spatial sharing (since each surface point is observed from a unique view in its local coordinate frame) and by reciprocity (since we effectively have observations in which the view and light directions are exchanged).




Fig. 12. Actual and predicted appearance of the hemisphere under fixed illumination and changing view. Given five input images from a single view (Fig. 10), the reflectance sharing method recovers a full SBRDF, including view-dependent effects.

VIII. GENERALIZED SPATIAL VARIATION

This section considers generalizations of the radial basis function SBRDF model by softening the requirement for spatial smoothness, and uses it to model the appearance of a human face. Since our estimate of spatially-varying reflectance allows us to synthesize images from any lighting direction and view, an approach such as this may eventually enable lighting- and pose-insensitive recognition that requires only a few input images for enrollment.

Rapid spatial variation can be handled using a multiplicative albedo or texture as in

$$f(\vec{x}, \vec{\theta}) = a(\vec{x})d(\vec{x}, \vec{\theta}),$$

where $a(\vec{x})$ is an albedo map for the surface and $d(\vec{x}, \vec{\theta})$ is a smooth function of five dimensions. As an example, consider the human face in Fig. 13(a). The function $a(\vec{x})$ accounts for rapid spatial variation due to pigment changes, while $d(\vec{x}, \vec{\theta})$ models the smooth spatial variation that occurs as we transition from a region where skin hangs loosely (e.g., the cheek) to where it is taut (e.g., the nose.)

In some cases, it is advantageous to express the SBRDF as a linear combination of 5D functions. For example, Sato et al. [6] and many others use the dichromatic model of reflectance [48] in which the BRDF is written as the sum of an RGB diffuse component and a scalar specular component that multiplies the source color. We employ the dichromatic model here, and compute the emitted radiance using

$$I_k(\vec{x}, \vec{\theta}) = s_k \left(a_k(\vec{x})d_k(\vec{x}, \vec{\theta}) + g(\vec{x}, \vec{\theta}) \right) \cos \theta_i, \quad (14)$$

where $\hat{s} = \{s_k\}_{k=RGB}$ is an RGB unit vector that describes the color of the light source. In (14), a single function g is used to model the specular reflectance component, while each color channel of the diffuse component is modeled separately. This is significantly more general than




Fig. 13. (a,b) Specular and diffuse components of a single input view. (c) Geometry for SBRDF recovery and image synthesis.

the usual assumption of a Lambertian diffuse component, and it can account for changes in diffuse color as a function of $\vec{\theta}$, such as the desaturation of the diffuse component of skin at large grazing angles witnessed by Debevec et al. [12].

Finally, although not used in our examples, more general spatial variation can be modeled by dividing the surface into a finite number of regions, where each region has spatial reflectance as described above. This technique is used, for example, in [16], [20].

A. Data Acquisition and SBRDF Recovery

For real surfaces, we require geometry and a set of images taken from known viewpoint and directional illumination. In addition, to estimate the separate diffuse and specular reflection components in (14), the input images must be similarly decomposed. Specular/diffuse separation can be performed in many ways (e.g., [6], [49]), one of which uses linear polarizers on both the camera and light source. Two exposures are captured for each view/lighting configuration, one with the polarizers aligned (to observe the sum of specular and diffuse components), and one with the source polarizer rotated by 90° (to observe the diffuse component only.) The specular component is then given by the difference between these two exposures. (See, e.g., [50].)

Geometry can also be recovered in a number of different ways. One possibility is photometric stereo, since under the right conditions, it provides the precise surface normals required for reflectometry. Figure 13 shows an example of a decomposed image along with the corresponding geometry, which is recovered by applying Lambertian photometric stereo to a set of diffuse (ideally Lambertian) images like that shown in the left of the figure.

Given the geometry and a set of decomposed images, the representation in (14) can be fit as follows. First, the effects of shadows and shading are computed, shadowed pixels are discarded, and shading effects are removed by dividing by $\cos \theta_i$. The RGB albedo $a(\vec{x})$ in (14) is estimated

as the median of the diffuse samples at each surface point, and normalized diffuse reflectance samples are computed by dividing by $a(\vec{x})$. The resulting normalized diffuse samples are used to estimate the three functions $d_k(\vec{x}, \vec{\theta})$ in (14) using the RBF techniques described in Sect. V-A. The samples from the specular images are similarly used to compute g .

B. Image Synthesis

In order to synthesize images under arbitrary view and illumination, the SBRDF coordinates \vec{q} at each surface point are determined by the spatial coordinates \vec{x} , the surface normal, and the view and lighting directions. The radiance emitted from that point toward the camera is then given by (14). Because (14) involves sums over a large number of RBF centers for each pixel, image synthesis can be slow. This process can be accelerated, however, using computer graphics techniques, including programmable graphics hardware and image precomputation.

Hardware Rendering Using the GPU. Equation (14) is well suited to implementation in graphics hardware because the same calculations are done at each pixel. For example, a vertex program can compute each \vec{q} and these can be interpolated as texture coordinates for each pixel. A fragment program can then perform the computation in (13), which is simply a sequence of distance calculations. Implementing the sum in (14) is straightforward, since it is simply a modulation by the albedo map and source color.

For the results in this section, we use one rendering pass for each RBF center, accumulating their contributions. On a GeForce FX 5900, rendering a 512×512 image with 2000 centers (and 2000 rendering passes) is reasonably efficient, taking approximately 30s. Further optimizations are possible, such as considering the contributions of multiple RBF centers in each pass.

Real-Time Rendering with Precomputed Images. To predict appearance under complex illumination in real-time, one can trade accuracy for rendering efficiency by using images synthesized with the RBF representation to compute more “rendering-friendly” representations. Suitable representations include that of Meseth *et al.* [51], who precompute lower-dimensional functions for a discrete number of viewpoints and linearly interpolate between these functions at run time; and the double-PCA model of Nayar *et al.* [52], who compress fixed-view data using an algorithm that is conceptually similar to clustered PCA methods [53], [54]. Both of these approaches allow real-time relighting of specular objects with complex illumination and shadows.




Fig. 14. Actual and synthetic images for a novel illumination direction. The image on the right was rendered using the reflectance representation in (14) fit to four input images. Left inset shows a polar plot of the input (+) and output (o) lighting directions, with concentric circles representing angular distances of 10° and 20° from the viewing direction.

We emphasize that despite the gains in efficiency discussed in this section, the RBF representation does not compete with parametric representations for synthesis. Instead, in its current form, it can be viewed as a useful intermediate representation between acquisition and rendering.

C. Results

As a demonstration, the representation of (14) was used to model a human face, which exhibits diffuse texture in addition to smooth spatial variation in its specular component.

The spatially-varying reflectance function estimated from four camera/source configurations is shown in Figs. 14–16. For these results, two polarized exposures were captured in each configuration, and the viewpoint remained fixed throughout. For simplicity, the subject’s eyes remained closed. (Accurately representing the spatial discontinuity at the boundary of the eyeball would require the surface to be segmented as mentioned in Sect. VII.) The average angular separation of the light directions is 21° , spanning a large area of frontal illumination. (See Fig. 14.) This angular sampling rate is considerably less dense than in previous non-parametric work; approximately 150 source directions would be required to cover the sphere at this rate compared to over 2000 source directions used by Debevec et al. [12].

In the recovered SBRDF, 2000 centers were used for each diffuse color channel, and 5000 centers were used for the specular component. (Figure 16 shows scatter plots of the specular component as the number of RBF centers is increased.) Each diffuse channel requires the storage of 2006 coefficients—the weights for 2000 centers and six polynomial coefficients—and 2000

sample points \vec{q}_i , each with five components. This could be reduced, for example, by using the same centers locations for all three color channels. The specular component requires 5006 coefficients and 5000 centers, so the total size is 66,024 single precision floating-point numbers, or 258kB. This is a very compact representation of both view and lighting effects.

Figure 14 shows real and synthetic images of the surface under novel lighting conditions, and shows how a smooth SBRDF is recovered despite the extreme sparsity of the input images. Most importantly, the disturbing aliasing effects observed in the ‘no sharing’ results of Fig. 11 are avoided. Figure 15 shows that the recovered SBRDF is indeed spatially-varying. The graph in the right of this figure is a (θ_h, ϕ_d) scatter plot of the specular SBRDF on the tip of the nose (in transparent blue) and on the cheek (in red), and it shows that the recovered specular lobe on the nose is substantially larger.

At first glance, the fact that this spatial variation is recovered from such a small set of images seems to contradict the theoretical analysis of Sect. III-C. Indeed, that section demonstrates that the maximum recovered spatial frequency is inversely related to the spacing between the directional light sources used for acquisition, which in this case is extremely large. The important difference, however, is that the theoretical analysis of Sect. III-C considers only a single specular highlight on a convex surface. The face, on the other hand, contains a number of concavities, so a single image provides multiple observations of specular highlights at different spatial locations.

While this synthetic result is plausible, careful examination of Fig. 14 reveals deviations from the actual image (the relative RMS difference is 9.5%). For example, the spatial discontinuity in the specular component at the boundary of the lips is smoothed over due to the assumption of slow spatial variation; and more generally, with such a limited input data, the representation is sensitive to noise caused by extreme interreflection and subsurface scattering, motion of the subject during acquisition, calibration errors in the source positions and relative strengths, and errors in the geometry. The accuracy could be improved, for example, by using a high speed acquisition system such as that of Debevec *et al.* [12]; by employing more sophisticated approaches for recovering surface geometry (e.g., [55], [56]); and by identifying spatial discontinuities in the SBRDF, perhaps using clustering techniques that use diffuse color as a cue for segmentation.

We emphasize, however, that only four input images are used, and it would be difficult to improve the results without further assumptions. Even a parametric method like that of Lensch *et al.* [16] may perform poorly in this case, since little more than a Lambertian albedo value

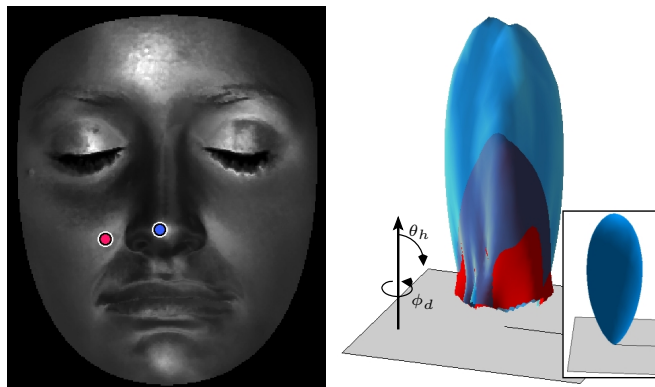


Fig. 15. Spatial variation in the estimated specular reflectance function. *Left*: synthesized specular component used to generate the image in the right of Fig. 14. *Right*: magnitude of the estimated specular SBRDF at two surface points. Plots are the SBRDF as a function of (θ_h, ϕ_d) for $\theta_d = 5^\circ$, with red and transparent-blue plots representing the indicated points on the cheek and nose. (Large values of θ_h near the origin are outside the convex hull of input samples and are not displayed.) For comparison, the inset shows a Cook-Torrance lobe fit to the reflectance of the nose.




Fig. 16. Estimated SBRDF on the cheek (red) and nose (blue) as the number of RBF centers is increased using the greedy algorithm. The 5000-center plots are the same as those on the right of Fig. 15.

could be fit reliably from the four (or less) reflectance samples available at each point.

Finally, Fig. 17 shows synthetic images with a novel viewpoint, again demonstrating that a full SBRDF is recovered despite the fact that only one viewpoint is used as input.

D. A Special Case: One Input Image

The RBF representation can also be adapted to the extreme case when only one input image is available. In one (orthographic, directional illumination) image all reflectance samples lie on a hyperplane of constant w , reducing the dimension of the SBRDF by one.

To exploit the reduced dimension in the case when a specular/diffuse separation is appropriate,




Fig. 17. Synthesized images for two novel viewpoints. Even though the input images are captured from a single viewpoint, a complete SBRDF is recovered, including view-dependent effects.

we use a simplified SBRDF representation and compute the surface radiance according to

$$I_k(\vec{q}) = s_k \left(a_k(\vec{x}) + \sum_{i=1}^N \lambda_i \|\vec{q} - \vec{q}_i\| \right) \cos\theta_i, \quad (15)$$

where $\vec{q} = (x, y, u, v)$. Here, the diffuse component is modeled as Lambertian, and the albedo $a(\vec{x})$ is estimated directly from the reflectance samples in the diffuse component of the input image after shading and shadows are removed. The specular component is estimated from the specular reflectance samples using the same fitting procedure as described in the previous section.

Figure 18 shows the results of fitting the model in 15 using $N = 2000$. The model is fit to the specular and diffuse components of a single view/lighting condition (computed using two polarized exposures), and this model is used to predict the appearance of the face under natural lighting from environment maps. The results in Fig. 18 were rendered using precomputation [52] as discussed in Sect. VIII-B, which allows real-time synthesis under complex lighting.

Since a single view/lighting condition is used as input, only a 2D subset of the angular variation is recovered, and Fresnel effects are ignored. (As done by Debevec et al. [12], this representation could be enhanced to approximate Fresnel effects by using a data-driven microfacet model with an assumed index of refraction.) Also, by using a complete environment map, we necessarily extrapolate the reflectance function beyond the convex hull of input samples, where it is known to be less accurate. Despite these limitations, the method obtains reasonable results, and they would be difficult to improve without assuming a specific parametric BRDF model (as in, e.g., [57]).

IX. CONCLUSIONS AND FUTURE WORK

This paper presents a method for exploiting spatial coherence to estimate a non-parametric, spatially-varying reflectance function from a sparse set of images of known geometry. Reflectance




Fig. 18. Images synthesized using the SBRDF representation in (15) estimated from the single (decomposed) image shown in Fig. 13. These were rendered in real-time using the methods discussed in Sect. VIII-B.

estimation is framed as a scattered-data interpolation problem in a joint spatial/angular domain, an approach that allows the exchange of spatial resolution for an increase in angular resolution of the reflectance function.

The paper introduces a theoretical framework to describe the exchange of spatial and angular reflectance resolution. It also presents a flexible representation of reflectance based on radial basis functions (RBFs) and shows how this representation can be adapted to handle: i) homogeneous BRDF data, ii) smooth spatially-varying reflectance from multiple images, iii) spatial variation with texture, and iv) a single input image. When using this representation, the recovered reflectance model is shown to degrade gracefully as the number of input images decreases.

The most immediate practical issue for future work involves computational efficiency. As presented, the RBF representation is a useful intermediate representation of spatially-varying reflectance since it can be used in combination with current rendering techniques based on precomputation. To improve this, it may be possible to develop real-time synthesis techniques directly from the RBF representation using, for example, fast multipole methods that reduce the evaluation of (9) from $O(N^2)$ to $O(N \log N)$ [42].

Advancing computational power and imaging technology are enabling the use of complex and accurate appearance models for image analysis. In this context, methods for efficiently acquiring and representing these models must be developed. This paper takes a step in this direction by providing a method for using sparse, image-based datasets to recover spatial reflectance.

ACKNOWLEDGMENT

Brian Curless engaged in many helpful discussions towards shaping the theoretical analysis in section III, and Aner Ben-Artzi contributed to hardware rendering using the GPU. This work was funded in part by the National Science Foundation. S. Enrique and R. Ramamoorthi were supported under grants CCF-03-05322, IIS-04-30258 and CCF-04-46916. T. Zickler and P. Belhumeur were supported under grants IIS-00-85864 and IIS-03-08185.

REFERENCES

- [1] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis, "Geometric considerations and nomenclature for reflectance," Monograph 160, National Bureau of Standards (US), 1977.
- [2] A. Georghiadis, "Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo," in *Proc. Int. Conf. on Computer Vision*, 2003, pp. 816–823.
- [3] V. Blanz and T. Vetter, "Face recognition based on fitting a 3D morphable model," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 9, september 2003.
- [4] K. I. K. Hara, K. Nishino, "Light source position and reflectance estimation from a single view without the distant illumination assumption," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 4, pp. 493–505, 2005.
- [5] D. K. McAllister, A. Lastra, and W. Heidrich, "Efficient rendering of spatial bi-directional reflectance distribution functions," in *Graphics Hardware 2002 (Proc. Eurographics/ACM SIGGRAPH Hardware Workshop)*, 2002, pp. 79–88.
- [6] Y. Sato, M. D. Wheeler, and K. Ikeuchi, "Object shape and reflectance modeling from observation," in *Proc. ACM SIGGRAPH*, 1997, pp. 379–387.
- [7] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse global illumination: recovering reflectance models of real scenes from photographs," in *Proc. ACM SIGGRAPH*, 1999, pp. 215–224.
- [8] R. Cook and K. Torrance, "A reflectance model for computer graphics," *Computer Graphics (Proc. ACM SIGGRAPH)*, vol. 15, no. 3, pp. 307–316, 1981.
- [9] M. Oren and S. Nayar, "Generalization of the Lambertian model and implications for machine vision," *Int. Journal of Computer Vision*, vol. 14, pp. 227–251, 1996.
- [10] R. Lu, J. Koenderink, and A. Kappers, "Optical properties (bidirectional reflection distribution functions) of velvet," *Applied Optics*, vol. 37, no. 25, pp. 5974–5984, 1998.
- [11] D. Wood, D. Azuma, K. Aldinger, B. Curless, T. Duchamp, D. Salesin, and W. Stuetzle, "Surface light fields for 3D photography," in *Proc. ACM SIGGRAPH*, 2000, pp. 287–296.
- [12] P. Debevec, T. Hawkins, C. Tchou, H. Duiker, W. Sarokin, and M. Sagar, "Acquiring the reflectance field of a human face," in *Proc. ACM SIGGRAPH*, 2000, pp. 145–156.
- [13] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "Image-based 3D photography using opacity hulls," *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 21, no. 3, pp. 427–437, 2002.
- [14] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg, "Image-based BRDF measurement including human skin," in *Rendering Techniques '99 (Proc. Eurographics Rendering Workshop)*, 1999, pp. 139–152.
- [15] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 22, no. 3, pp. 759–769, 2003.

- [16] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel, "Image-based reconstruction of spatially varying materials," in *Rendering Techniques 2001 (Proc. Eurographics Rendering Workshop)*, 2001.
- [17] K. Karner, H. Mayer, and M. Gervautz, "An image based measurement system for anisotropic reflectance," *Computer Graphics Forum*, vol. 15, no. 3, pp. 119–128, 1996.
- [18] A. Hertzmann and S. Seitz, "Shape and material by example: A photometric stereo approach," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2003.
- [19] M. McCool, J. Ang, and A. Ahmad, "Homomorphic factorization of BRDFs for high-performance rendering," in *Proc. ACM SIGGRAPH*, 2001, pp. 171–178.
- [20] R. Jaroszkiwicz and M. McCool, "Fast extraction of BRDFs and material maps from images," in *Proc. Graphics Interface*, 2003, pp. 1–10.
- [21] B. Cabral, M. Olano, and P. Nemeč, "Reflection space image based rendering," in *Proc. ACM SIGGRAPH*, 1999, p. 165.
- [22] R. Ramamoorthi and P. Hanrahan, "Frequency space environment map rendering," *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 21, no. 3, pp. 517–526, 2002.
- [23] T. Zickler, S. Enrique, R. Ramamoorthi, and P. Belhumeur, "Reflectance sharing: Image-based rendering from a sparse set of images," in *Proc. Eurographics Symposium on Rendering*, 2005.
- [24] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen, "The lumigraph," in *Proc. ACM SIGGRAPH*, 1996, pp. 43–54.
- [25] M. Levoy and P. Hanrahan, "Light field rendering," in *Proc. ACM SIGGRAPH*, 1996, pp. 31–42.
- [26] C. Christoudias, L. Morency, and T. Darrell, "Light field appearance manifolds," in *Proc. European Conf. on Computer Vision*, vol. 4, 2004, pp. 481–493.
- [27] A. Georghiades, P. Belhumeur, and D. Kriegman, "From few to many: illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 643–660, 2001.
- [28] S. Rusinkiewicz, "A new change of variables for efficient BRDF representation," in *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, 1998, pp. 11–22.
- [29] S. Nayar, K. Ikeuchi, and T. Kanade, "Surface reflection: Physical and geometric perspectives," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 611–634, July 1991.
- [30] J. Chai, S. Chan, H. Shum, and X. Tong, "Plenoptic sampling," in *Proc. ACM SIGGRAPH*, 2000, pp. 307–318.
- [31] F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. Sillion, "A frequency analysis of light transport," *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 24, no. 3, pp. 1115–1126, 2005.
- [32] R. Bracewell, K. Chang, A. Jha, and Y. Wang, "Affine theorem for two-dimensional fourier transform," *Electronics Letters*, vol. 29, no. 3, p. 304, 1993.
- [33] S. Marschner, "Inverse rendering for computer graphics," Ph.D. dissertation, Cornell University, 1998.
- [34] D. Terzopoulos, "Multilevel computational processes for visual surface reconstruction," *Computer Vision, Graphics and Image Processing*, vol. 24, pp. 52–96, 1983.
- [35] M. Powell, "The theory of radial basis function approximation in 1990," in *Advances in Numerical Analysis, Vol. II*, W. Light, Ed. Oxford Science Publications, 1992, pp. 105–210.
- [36] M. Buhmann, *Radial basis functions*. Cambridge University Press, 2003.
- [37] J. Duchon, "Splines minimizing rotation-invariant semi-norms in Sobolev spaces," in *Constructive theory of functions of several variables*, W. Schempp and K. Zeller, Eds. Springer-Verlag, 1977, pp. 85–100.
- [38] C. Micchelli, "Interpolation of scattered data: distance matrices and conditionally positive definite functions," *Constructive Approximation*, vol. 1, pp. 11–22, 1986.

- [39] G. Wahba, *Spline models for observational data*, ser. CBMS-NSF regional conference series in applied mathematics. SIAM, 1990, no. 59.
- [40] H. Dinh, G. Turk, and G. Slabaugh, “Reconstructing surfaces using anisotropic basis functions,” in *Proc. Int. Conf. on Computer Vision*, 2001, pp. 606–613.
- [41] G. Beatson, R. Goodsell and M. Powell, “On multigrid techniques for thin plate spline interpolation in two dimensions,” *Lectures in Applied Mathematics*, vol. 32, pp. 77–97, 1995.
- [42] R. Beatson and G. Newsam, “Fast evaluation of radial basis functions: I,” *Comput. Math. Appl.*, vol. 24, pp. 7–19, 1992.
- [43] J. Carr, R. Beatson, J. Cherie, T. Mitchell, W. Fright, B. McCallum, and T. Evans, “Reconstruction and representation of 3D objects with radial basis functions,” in *Proc. ACM SIGGRAPH*, 2001, pp. 67–76.
- [44] E. Lafortune, S. Foo, K. Torrance, and D. Greenberg, “Non-linear approximation of reflectance functions,” in *Proc. ACM SIGGRAPH*, 1997, pp. 117–126.
- [45] G. J. Ward, “Measuring and modeling anisotropic reflection,” *Computer Graphics (Proc. ACM SIGGRAPH)*, vol. 26, no. 2, pp. 265–272, 1992.
- [46] A. Lee, W. Sweldens, P. Schroder, L. Cowsar, and D. Dobkin, “MAPS: Multiresolution adaptive parameterization of surfaces,” in *Proc. ACM SIGGRAPH*, 1998, pp. 95–104.
- [47] X. Gu, S. Gortler, and H. Hoppe, “Geometry images,” *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 21, no. 3, pp. 355–361, 2002.
- [48] S. Shafer, “Using color to separate reflection components,” *COLOR Res. App.*, vol. 10, no. 4, pp. 210–218, 1985.
- [49] S. Nayar, X. Fang, and T. Boult, “Separation of reflection components using color and polarization,” *Int. Journal of Computer Vision*, vol. 21, no. 3, pp. 163–186, 1997.
- [50] S. Mersch, “Polarized lighting for machine vision applications,” in *Proc. RI/SME Third Annual Applied Machine Vision Conf.* Schaumburg, IL, Feb. 1984, pp. 40–54.
- [51] J. Meseth, G. Müller, and R. Klein, “Reflectance field based real-time, high-quality rendering of bidirectional texture functions,” *Computers & Graphics*, vol. 28, no. 1, pp. 105–112, 2004.
- [52] S. Nayar, P. Belhumeur, and T. Boult, “Lighting sensitive display,” *ACM Trans. Graphics*, vol. 23, no. 4, p. 963, 2004.
- [53] W. Chen, J. Bouguet, M. Chu, and R. Grzeszczuk, “Light field mapping: Efficient representation and hardware rendering of surface light fields,” *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 21, no. 3, pp. 447–456, 2002.
- [54] P. Sloan, J. Hall, J. Hart, and J. Snyder, “Clustered principal components for precomputed radiance transfer,” *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 22, no. 3, pp. 382–391, 2003.
- [55] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi, “Efficiently combining positions and normals for precise 3d geometry,” *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, vol. 24, no. 3, 2005.
- [56] T. Zickler, P. Belhumeur, and D. Kriegman, “Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction,” *Int. Journal of Computer Vision*, vol. 49, no. 2/3, pp. 215–227, September 2002.
- [57] S. Boivin and A. Gagalowicz, “Image-based rendering of diffuse, specular and glossy surfaces from a single image,” in *Proc. ACM SIGGRAPH*, 2001, pp. 107–116.



Todd Zickler is an assistant professor of electrical engineering in the Division of Engineering and Applied Sciences at Harvard University. He received a B.Eng. in honours electrical engineering from McGill University in 1996, and an M.S. in electrical engineering from Yale University in 2001. He received a Ph.D. in electrical engineering from Yale in 2004, at which point he joined Harvard University as an assistant professor. His research interests span computer vision, image processing and computer graphics, and he is currently focused on the recovery of shape and reflectance information from image data as well as efficient representations for modeling visual appearance. Todd's research is funded by the National Science Foundation.



Ravi Ramamoorthi received the BS in engineering and applied science and MS degrees in computer science and physics from the California Institute of Technology in 1998. He received his PhD degree in computer science from the Stanford University Computer Graphics Laboratory in 2002. He is currently an assistant professor of Computer Science at Columbia University. His research interests cover many aspects of computer vision and graphics, including mathematical foundations, real-time photorealistic rendering, image-based and inverse rendering, and lighting and appearance in computer vision. He has published papers on all of these topics in leading graphics and vision conferences and journals. In 2005, he was awarded a Sloan research fellowship and a National Science Foundation Career award.



Sebastian Enrique received a Licenciante in Computer Science degree with honors from University of Buenos Aires, Argentina. Then, he received the MS degree and started studies towards a PhD in computer science at Columbia University, New York, under the supervision of Ravi Ramamoorthi, where he worked on the contribution to this paper. His main interests are real-time rendering techniques. After interning in DirectX group at Microsoft, he decided to take a leave and he is currently working at Electronic Arts Canada.



Peter N. Belhumeur is a professor in the Department of Computer Science at Columbia University and the director of the Laboratory for the Study of Visual Appearance. He received an Sc.B. in information sciences from Brown University in 1985, and a Ph.D. in engineering sciences from Harvard University in 1993. He was a postdoctoral fellow at the University of Cambridge's Isaac Newton Institute for Mathematical Sciences in 1994. After that he was assistant, associate, and professor of electrical engineering at Yale University. He joined Columbia University as a professor of computer science in 2002. His research focuses on illumination, reflectance, and shape and their relation to visual appearance. Within these areas, he concentrates on two subproblems: the representation and recognition of objects under variable illumination and the estimation of the geometry of objects from low-level cues like image brightness, binocular stereopsis, and motion. Applications include face and object recognition, image-based rendering, computer graphics, content-based image and video compression, and human computer interfaces. Peter is a recipient of the Presidential Early Career Award for Scientists and Engineers (PECASE) and the National Science Foundation Career Award. Peter won both the Siemens Best Paper Award at the 1996 IEEE Conference on Computer Vision and Pattern Recognition and the Olympus Prize at the 1998 European Conference of Computer Vision. His research in computational vision is funded by the National Science Foundation, the National Institutes of Health, the Army Research Office, and DARPA.