Eurographics Symposium on Rendering (2005)
Kavita Bala, Philip Dutré (Editors)

Reflectance Sharing: Image-based Rendering from a Sparse Set of Images

Todd Zickler[1], Sebastian Enrique[2], Ravi Ramamoorthi[2], and Peter Belhumeur[2]

[1]Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA
[2]Computer Science Department, Columbia University, New York, NY, USA

Abstract

When the shape of an object is known, its appearance is determined by the spatially-varying reflectance function defined on its surface. Image-based rendering methods that use geometry seek to estimate this function from image data. Most existing methods recover a unique angular reflectance function (e.g., BRDF) at each surface point and provide reflectance estimates with high spatial resolution. Their angular accuracy is limited by the number of available images, and as a result, most of these methods focus on capturing parametric or low-frequency angular reflectance effects, or allowing only one of lighting or viewpoint variation. We present an alternative approach that enables an increase in the angular accuracy of a spatially-varying reflectance function in exchange for a decrease in spatial resolution. By framing the problem as scattered-data interpolation in a mixed spatial and angular domain, reflectance information is shared across the surface, exploiting the high spatial resolution that images provide to fill the holes between sparsely observed view and lighting directions. Since the BRDF typically varies slowly from point to point over much of an object's surface, this method enables image-based rendering from a sparse set of images without assuming a parametric reflectance model. In fact, the method can even be applied in the limiting case of a single input image.

Categories and Subject Descriptors (according to ACM CCS): I.2.10 [Artificial Intelligence]: Vision and Scene Understanding I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

Given a set of images of a scene, image-based rendering (IBR) methods strive to build a representation for synthesizing new images of that scene under arbitrary illumination and viewpoint. One effective IBR representation consists of the scene geometry coupled with a reflectance function defined on that geometry. At a given point under given illumination conditions, the reflectance function assigns a radiance value to each exitant ray, so once the geometry and reflectance function are known, realistic images can be synthesized under arbitrary viewpoint and (possibly complex and near-field) illumination.

This approach to IBR involves two stages: recovery of both geometry and reflectance. Yet, while great strides have been made at recovering object shape (e.g., laser-scanners and computer vision techniques), less progress has been made at recovering reflectance properties. Recovering reflectance is difficult because that is where the high-dimensionality of IBR is rooted. 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). The BRDF generally changes spatially over an object's surface, and recovering this spatially-varying BRDF (or 6D SBRDF) without further assumptions generally 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 a near exhaustive sampling of images of the scene from all viewpoints and lighting directions, which can be tens-of-thousands of images or more.

In previous work, recovering spatial reflectance has been made tractable in two different ways. The first is to approximate reflectance using an analytic BRDF model, thereby simplifying the problem from that of recovering a 4D function at each point to that of estimating of a handful of parameters (e.g., [MLH02, SWI97, YDMH99, Geo03]). Since they only recover a few parameters at each point, these methods are able to provide reflectance estimates from a small number of images. They require the selection of a specific parametric BRDF model *a priori*, however, which limits their accuracy and generality.

The second category of methods avoids the restrictions of parametric BRDF models by: i) using a much larger number of input images, and ii) recovering only a subset of the reflectance function. For example, Wood et al. [WAA+00] use over 600 images of an object under fixed (complex) illumination to estimate the 2D view-dependent reflectance variation, and Debevec et al. [DHT+00] use 2048 images to measure the 2D lighting-dependent variation with fixed viewpoint. A full 4D (view and illumination) reflectance function is measured by Matusik et al. [MPBM02], who use more than 12,000 images of an object with known visual hull; but even this large number of images provides only a sparse sampling of the appearance variation at each point, and as a result, images of the object can be synthesized using only low-frequency illumination environments.

This paper presents an alternative approach to estimating spatial reflectance—one that combines the benefits of both parametric methods (i.e., sparse images) and non-parametric methods (i.e., arbitrary reflectance functions.) In developing this approach, it makes the following technical contributions.

- SBRDF estimation is posed as a scattered-data interpolation problem, with images providing dense 2D slices of data embedded in the mixed spatial and angular domain.
- This interpolation problem is solved by introducing: i) a new parameterization of the BRDF domain, and ii) a non-parametric representation of reflectance based on radial basis functions (RBFs).
- This representation is easily adapted to handle: i) homogeneous BRDF data, ii) spatially-varying reflectance from multiple images, and iii) spatially-varying reflectance from a single input image.

Since it is non-parametric, the proposed method does not assume *a priori* knowledge of the SBRDF and is flexible enough to represent arbitrary reflectance functions, including those with high-frequency specular effects. This approach is also very different from previous non-parametric techniques (e.g., [WAA+00, DHT+00, MPBM02]) 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 typically varies slowly from point to point over much of an object's surface, this means that we can often obtain visually pleasing results from a drastically reduced set of images. Additionally, the method degrades gracefully as the number of input images is reduced (see Fig. 6), and as shown in Fig. 13, it can even be applied in the extreme case of a single input image.

2. 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. [YDMH99] and Georghiades [Geo03], who assume that specular BRDF parameters are constant across a surface. Similarly, Sato et al. [SWI97] estimate the specular parameters at only a small set of points, later interpolating these parameters across the surface. Lensch et al. [LKG+01] 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 [LKK98, MWL+99, MPBM03] 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. (For near-planar surfaces where curvature is not available, more angular reflectance information can be obtained using near-field illumination and perspective views [KMG96].)

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 [HS03], efficient BRDF acquisition [MPBM03], efficient rendering of BRDFs [MAA01, JM03], and efficient evaluation of environment maps [CON99, RH02]. Here, we exploit compressibility by assuming that the BRDF typically varies rapidly only in certain dimensions, such as the half-angle.

The three ideas of this section have been developed in very different contexts, and this paper combines and expands them to solve a novel problem: estimating non-parametric SBRDFs from sparse images. The fusion of these ideas is enabled by the BRDF parameterization of Sect. 3, and an interpolation approach that unifies the treatment of spatial and angular dimensions (Sects. 4–6).

2.1. Assumptions

We exploit scene geometry to reduce the number of input images required to accurately represent appearance. Thus, unlike pure light field techniques [GGSC96,LH96], the method requires a set of images of an object with known geometry, viewpoint, and either point-source or directional illumination. A number of suitable acquisition systems have been presented (e.g., [SWI97, DHT+00, MPBM02].)

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. 7.) In this case, our use of the term *SBRDF* is synonymous with the *non-local reflectance field* defined by Debevec et al. [DHT+00].

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. 3 and is left for future work.

3. Notation and BRDF Parameterization

At the core of our approach is the interpolation of scattered data in multiple (3-6) dimensions, the success of which depends on how the SBRDF is parameterized. This section introduces some notation and presents one possible parameterization. Based on this parameterization, our interpolation technique is discussed in Sects. 5 and 6.

The SBRDF is a function of six dimensions and is written $f(\vec{x},\vec{\theta})$, where $\vec{x} = (x,y) \subset \mathbb{R}^2$ 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. When the BRDF is isotropic, the angular variation reduces to a function of three dimensions, commonly parameterized by $(\theta_i,\theta_o,\phi_o-\phi_i)$. In this work, we restrict ourselves to this isotropic case 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.

The BRDF domain can be parameterized in a number of ways, and as discussed below, one good choice is Rusinkiewicz's halfway/difference parameterization [Rus98], shown in Fig. 1(a). Using this parameterization in the isotropic case, the BRDF domain is $\vec{\theta} = (\theta_h,\phi_d,\theta_d) \subset [0,\frac{\pi}{2}) \times [0,\pi) \times [0,\frac{\pi}{2})$. Note that ϕ_d is restricted to $[0,\pi)$ since $\phi_d \longmapsto \phi_d + \pi$ by reciprocity.

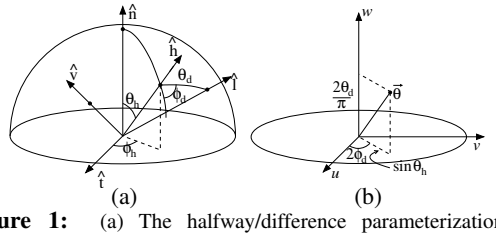

Figure 1: (a) The halfway/difference parameterization of Rusinkiewicz. In the isotropic case, the BRDF domain is parametrized by $(\theta_h,\phi_d,\theta_d)$ [Rus98]. (b) The mapping defined by Eq. (1) that creates a parameterization suitable for interpolation.

The existence of singularities at $\theta_h = 0$ and $\theta_d = 0$ and the required periodicity $(\phi_d \longmapsto \phi_d + \pi)$ make the standard halfway/difference parameterization unsuitable for most interpolation techniques. Instead, we define the mapping $(\theta_h,\phi_d,\theta_d) \longmapsto (u,v,w)$, as

$$(u,v,w) = \left(\sin\theta_h\cos 2\phi_d, \sin\theta_h\sin 2\phi_d, \frac{2\theta_d}{\pi}\right). \quad (1)$$

This mapping is shown in Fig. 1(b). 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 $\theta_d = 0$ (i.e., where the light and view directions are equivalent). This configuration is difficult to create in practice, making it unlikely to occur during acquisition. During synthesis, it must be handled with care.

3.1. Considerations for Image-based Acquisition

The halfway/difference parameterization increases compression rates since common features such as specular and retro-reflective peaks are aligned with the coordinate axes [Rus98]. The modified parameterization of Eq. (1) 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 in IBR for an additional reason: for image-based data, they separate the sparsely- and densely-sampled dimensions of the BRDF. (Marschner's [Mar98] parameterization 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. The image represents a nearly continuous sampling of θ_h and ϕ_d in this plane. Thus, a *set* of images provides dense sampling of (θ_h,ϕ_d) but only as many samples of θ_d as there are images. (The orthographic/directional case is considered for illustrative purposes; it is not required by the method.)

Conveniently, the 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. 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.

4. 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 the previous section, 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 simply 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. [MPBM03], an RBF representation does not require a local preprocessing step to resample the input data at regular intervals.

Additional properties of this method include: i) the cost of computing an RBF interpolant is dimension-independent, and ii) the size of the representation does not grow substantially as the dimension increases. This is in direct contrast to piecewise polynomial splines (e.g., [Ter83]) and local methods like polynomial regression (e.g., [MWL+99]) and the *push/pull* algorithm of Gortler et al. [GGSC96]. These other 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 [JM03] handle this by approximating the high-dimensional SBRDF by a product of 2D functions, each of which is triangulated independently.)

4.1. Radial Basis Functions

To briefly review RBF interpolation (see, e.g., [Pow92, Buh03]), 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 as a sum of a low-order polynomial and a set of scaled, radially symmetric basis functions centered at the sample points;

$$g(\vec{x}) \approx \tilde{g}(\vec{x}) = p(\vec{x}) + \sum_{i=1}^{N} \lambda_i \psi(\|\vec{x}-\vec{x}_i\|), \qquad (2)$$

where $p(\vec{x})$ is a polynomial of order n or less, $\psi : \mathbb{R}^+ \to \mathbb{R}$ 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 determined as the solution of the linear system

$$\begin{bmatrix} \Psi & P \\ P^{\mathsf{T}} & 0 \end{bmatrix} \begin{bmatrix} \vec{\lambda} \\ \vec{c} \end{bmatrix} = \begin{bmatrix} \vec{g} \\ 0 \end{bmatrix}, \qquad (3)$$

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 [Duc77, Mic86].

In many cases we can benefit from using *radially asymmetric* basis functions (which are stretched in certain directions), and here we use them to manage the irregularity in our sampling pattern. (Recall from Fig. 1 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. [DTS01], an asymmetric radial function is created by scaling the Euclidean distance in Eq. (2) so that the basis functions become

$$\psi(\|M(\vec{x}-\vec{x}_i)\|), \qquad (4)$$

where $M \in \mathbb{R}^{d\times d}$. In our case we choose $M = \mathrm{diag}(1,1,m_w)$. 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 have found that typical values for m_w are between 0.1 and 0.5.

When the number of samples is large (i.e., $N > 10,000$), solving Eq. (3) requires care and can be difficult (or impossible) using direct methods. This limitation has been addressed quite recently, and iterative fitting methods [BP95], and fast multipole methods (FMMs) for efficient evaluation [BN92] have been developed for many choices of ψ in many dimensions. In some cases, solutions for systems with over half a million centers have been reported [CBC+01]. 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.

5. Homogeneous Surfaces

In this section, we apply RBF interpolation to homogeneous surfaces, where we seek to estimate a global BRDF that is not spatially-varying. The resulting BRDF representation may be useful for interpolating image-based BRDF data (e.g., [LKK98, MWL+99, MPBM03]).

As discussed in Sect. 3, in the case of 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 Eq. (3) exists for any non-coplanar data, is unique, minimizes a generalization of the thin-plate energy, and is therefore the smoothest in some sense [Duc77, CBC+01]. The BRDF is expressed as

$$\tilde{f}(\vec{\theta}) = c_1 + c_2 u + c_3 v + c_4 w + \sum_{i=1}^{N} \lambda_i \|\vec{\theta} - \vec{\theta}_i\|, \qquad (5)$$

where $\vec{\theta}_i = (u_i, v_i, w_i)$ represents a BRDF sample point from the input images, and $\vec{\lambda}$ and $\vec{c}$ are found by solving Eq. (3).

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. [CBC+01] 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.

5.1. Evaluation

To evaluate the BRDF representation in Eq. (5), we perform an experiment in which we compare it to both parametric BRDF models and to a non-linear basis (the isotropic Lafortune model [LFTG97].) 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. (Many other representations, such as wavelets and the Matusik bases are excluded from this comparison because they require dense, uniform samples.)

The input images simulate data from image-based BRDF measurement systems like those in Refs. [LKK98, MWL+99, MPBM03]. They are orthographic, directional-illumination images with a resolution of 100×100, and are generated such that θ_d is uniformly distributed 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. [MPBM03]), and the other generated using the physics-based Oren-Nayar model [ON94].

Figure 2 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 $4N+4$.

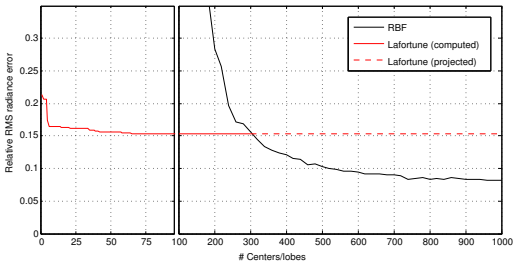

Figure 2: Accuracy of the RBF representation as the number of centers is increased using a greedy algorithm. The input is 10 images of a sphere synthesized using the metallic-blue BRDF measured by Matusik et al. This is compared to the isotropic Lafortune representation with an increasing number of lobes. 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.)

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 compact, the Lafortune model has the advantage of being more suitable for direct rendering [MLH02]. 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 given a sufficient number of centers, they can represent any 'smooth' function with arbitrary accuracy (see, e.g., [Buh03].) In this example, the RBF representation converges with less than 1000 centers, suggesting that only a small fraction of the available centers are required 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 increas-

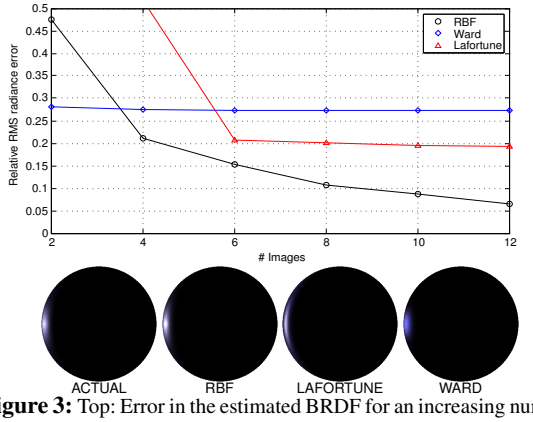

Figure 3: Top: Error in the estimated BRDF for an increasing number of images of a metallic-blue sphere. As the number of images increases, the RBF representation (with 1000 centers) approaches the true BRDF, whereas the isotropic Ward model [War92] and the Lafortune representation are too restrictive to provide an accurate fit. Bottom: Synthesized images using the three BRDF representations estimated from 12 input images. The angle between the source and view directions is 140°.

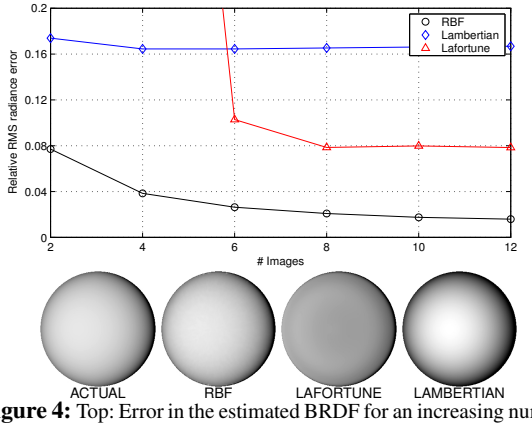

Figure 4: Top: Error in the estimated BRDF for an increasing number of input images of a diffuse Oren-Nayar sphere. Again, the 1000-center RBF representation approaches the true BRDF, whereas the Lafortune and Lambertian BRDF models are too restrictive to accurately represent the data. Bottom: Synthesized images comparing the three BRDF representations estimated from 12 input images. The angle between the source and view directions is 10°.

ing number of input images. Results are shown in Figs. 3 and 4 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 reflection in Fig. 3 and the retro-reflection in Fig. 4. Parametric models do not typically afford this flexibility—while it may be possible to find a parametric model that fits a specific BRDF quite well, it is very difficult to find a model that accurately fits general BRDFs.

6. Inhomogeneous Surfaces

The previous section suggests that RBFs can provide a useful representation for homogeneous BRDFs. In this section, we show that this same representation can be adapted to handle spatially-varying reflectance as well. In this case, it enables an exchange between spatial and angular resolution (a process we refer to as *reflectance sharing*), and it can drastically reduce the number of required input images. We begin by assuming that the 5D SBRDF varies smoothly 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 [Pow92]. 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

$$\tilde{f}(\vec{q}) = p(\vec{q}) + \sum_{i=1}^{N} \lambda_i \|\vec{q} - \vec{q}_i\|, \qquad (6)$$

where $p(\vec{q}) = c_1 + c_2 x + c_3 y + c_4 u + c_5 v + c_6 w$.

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., [LSS+98, GGH02]). 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, so that $\vec{s}(x, y) = (x, y, s(x, y)),\ (x, y) \subset [0, 1] \times [0, 1]$.

The procedure for recovering the parameters in Eq. (6) is almost exactly the same as in the homogeneous case. The coefficients of $\tilde{f}$ are found by solving Eq. (3) using a subset of the SBRDF samples from the input images, and this subset is chosen using a greedy algorithm. Radially asymmetric basis functions are realized using $M = \mathrm{diag}(m_{xy}, m_{xy}, 1, 1, m_w)$, 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.

6.1. Evaluation

The SBRDF representation of Eq. (6) can be evaluated using experiments similar to those for the homogeneous case.

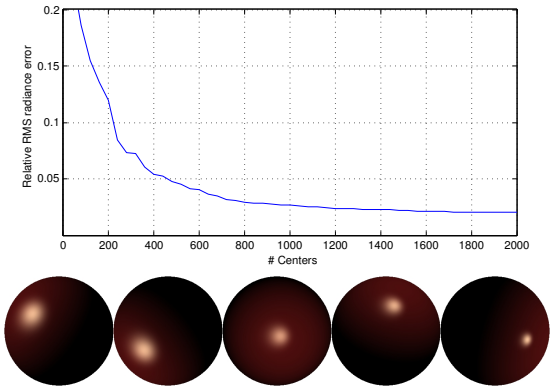

Figure 5: Top: Accuracy of the SBRDF recovered by reflectance sharing using the RBF representation in Eq. (6) as the number of centers is increased using a greedy algorithm. The input is 10 synthetic images of a hemisphere (five of which are shown) with linearly varying roughness.

Here, spatial variation is simulated using images of a hemisphere with a Cook-Torrance BRDF [CT81] with a linearly varying roughness parameter. Five images of the hemisphere are shown in Fig. 5, and they demonstrate how the highlight sharpens from left to right across the surface.

The graph in Fig. 5 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. This 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 6 contrasts reflectance sharing with conventional methods that interpolate only in the angular dimensions, estimating a separate BRDF at each point. This 'no sharing' technique is used by Matusik et al. [MPBM02], and is similar in spirit to Wood et al. [WAA+00], 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., [DHT+00]) that does not. For the particular example in Fig. 6, 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. 5, and it is used to predict the appearance of the surface under novel lighting. The top frame of Fig. 6 shows the actual appearance of the hemisphere under five novel conditions, and the lower frames show the reflectance

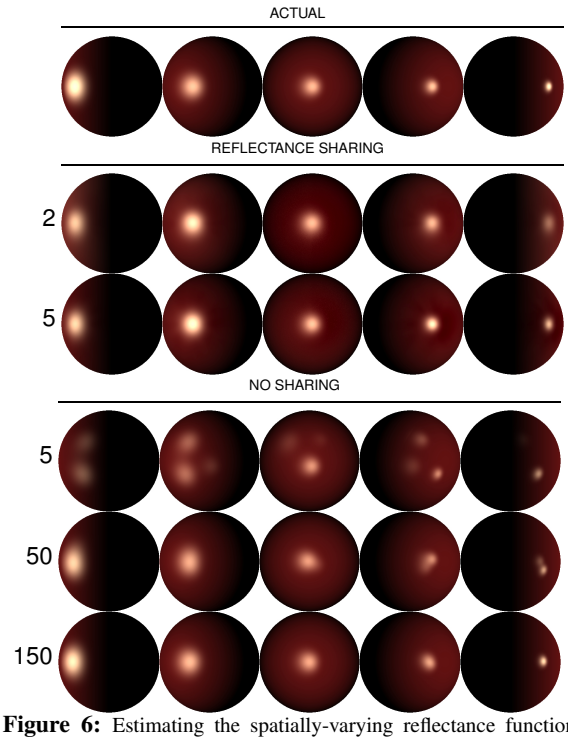

Figure 6: Estimating the spatially-varying reflectance function from a sparse set of images. Top frame: five images of a hemisphere under illumination conditions not used as input. Middle frame: appearance predicted by reflectance sharing with two and five input images. (The input is shown in Fig. 5; the two left-most images are used for 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.

sharing and 'no sharing' results obtained from increasing numbers of input images. Note that many other methods—most notably that of Lensch et al. [LKG+01]—are excluded from this comparison because they require the selection of a specific parametric model and therefore suffer from the limitations discussed in Sect. 5.

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 6 also shows how differently the two approaches degrade with sparse input. Reflectance sharing provides a smooth SBRDF whose accuracy gradually decreases away from the convex hull of input samples. (For example, the sharp specularity on the right side 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 aliasing or 'ghosting' occurs when the surface is illuminated by high-frequency environments like the directional illumination shown here.

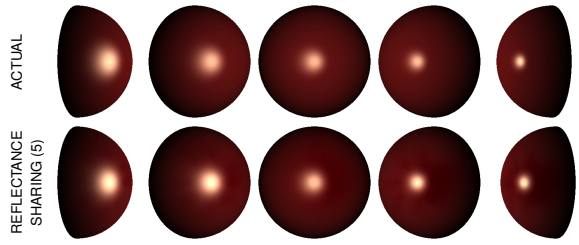

Figure 7: Actual and predicted appearance of the hemisphere under fixed illumination and changing view. Given five input images from a single view (bottom Fig. 5), the reflectance sharing method recovers a full SBRDF, including view-dependent effects.

Even when the input images are captured from a single viewpoint, our method recovers a full SBRDF, and as shown in Fig. 7, 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.)

7. Generalized Spatial Variation

This section considers generalizations of the radial basis function SBRDF model by softening the requirement for spatial smoothness, and applies the model to image-based rendering of a human face.

Rapid spatial variation can be handled using a multiplicative albedo or texture by writing the SBRDF as

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

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. 8(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. [SWI97] and many others use the dichromatic model of reflectance [Sha85] 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

$$I_k(\vec{x},\vec{\theta}) = s_k\left(a_k(\vec{x})d_k(\vec{x},\vec{\theta}) + g(\vec{x},\vec{\theta})\right)\cos\theta_i, \qquad (7)$$

where $\hat{s} = \{s_k\}_{k=RGB}$ is an RGB unit vector that describes the color of the light source. In Eq. (7), 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 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. [DHT+00].

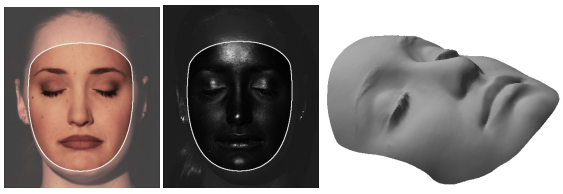

(a) (b) (c)

Figure 8: (a,b) Specular and diffuse components of a single input image. (c) Geometry used for SBRDF recovery and rendering.

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 Refs. [LKG+01, JM03].

7.1. 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, in order to estimate the separate diffuse and specular reflection components in Eq. (7), the input images must be similarly decomposed. Specular/diffuse separation can be performed in many ways (e.g., [SWI97, NFB97]), one of which uses linear polarizers on both the camera and light source and exploits the fact that the specular component preserves the linear polarization of the incident radiance. 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., [DHT+00].)

Geometry can also be recovered in a number of different ways, and one possibility is photometric stereo, since it provides the precise surface normals required for reflectometry. Figure 8 shows an example of a decomposed image along with the corresponding geometry, which is recovered using a variant of photometric stereo.

Given the geometry and a set of decomposed images, the representation in Eq. (7) 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 Eq. (7) 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 Eq. (7) using the RBF techniques described in Sect. 4.1. The samples from the specular images are similarly used to compute g.

7.2. Rendering

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 Eq. (7).

Because Eq. (7) involves sums over a large number of RBF centers for each pixel, image synthesis can be slow. This process can be accelerated, however, using programmable graphics hardware and precomputation.

Hardware Rendering Using the GPU. Equation (7) 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 Eq. (6), which is simply a sequence of distance calculations. Implementing the sum in Eq. (7) is straightforward, since it is simply a modulation by the albedo map and source color.

For the results in this section, we take this approach and 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 enable real-time rendering with complex illumination, images synthesized using the RBF representation can be used as input to methods based on precomputed images. One example is the double-PCA image relighting technique of Nayar et al. [NBB04]. This algorithm is conceptually similar to clustered PCA methods like those in Refs. [CBCG02, SHHS03], and allows real-time relighting of specular objects with complex illumination and shadows (obtained in the input images using shadow mapping in graphics hardware.)

We emphasize that despite these gains in efficiency, the RBF representation does not compete with parametric representations for rendering purposes. Instead, it should be viewed as a useful intermediate representation between acquisition and rendering.

7.3. Results

As a demonstration, the representation of Eq. (7) 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. 9–11. 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. 6.) The average angular separation of the light directions is 21°, spanning a large area of frontal illumination. (See Fig. 9.) This angular sampling rate is considerably less dense than in previous 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. [DHT+00].

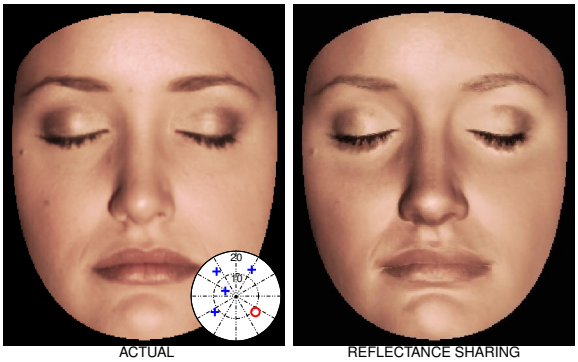

Figure 9: Actual and synthetic images for a novel illumination direction. The image on the right was rendered using the reflectance representation in Eq. (7) fit to four input images. Left inset shows a polar plot of the input (+) and output (○) lighting directions, with concentric circles representing angular distances of 10° and 20° from the viewing direction.

In the recovered SBRDF, 2000 centers were used for each diffuse color channel, and 5000 centers were used for the specular component. (Figure 11 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 9 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 ghosting effects observed in the 'no sharing' results of Fig. 6 are avoided. (The accompanying video includes animations of view and lighting variation.) Figure 10 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.

While this synthetic result is plausible, careful examination of Fig. 9 reveals that it deviates 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 smooth spatial variation; and more generally, with such a limited number of input samples, 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

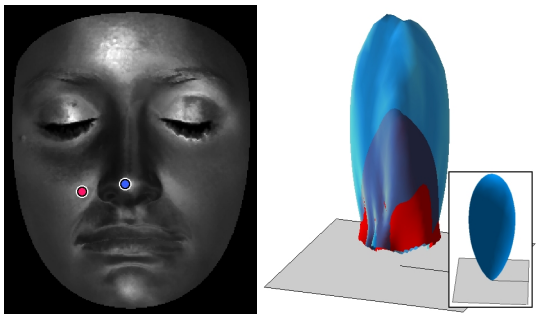

Figure 10: Spatial variation in the estimated specular reflectance function. Left: synthesized specular component used to generate the image in the right of Fig. 9. Right: magnitude of the estimated specular SBRDF at two surface points. Plots are the SBRDF as a function of (θ_h, ϕ_d) for $\theta_d = 5°$, 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.

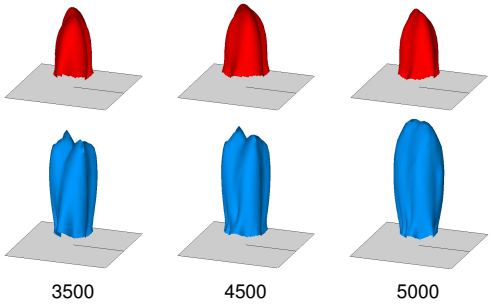

Figure 11: Estimated SBRDF on the cheek (red) and nose (blue) as the number of RBF centers is increased using the greedy algorithm. The 5000-center plots are the same as those on the right of Fig. 10.

example, by using a high speed acquisition system such as that of Debevec et al. [DHT+00], and by identifying spatial discontinuities (perhaps using clustering techniques or by using 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. [LKG+01] may perform poorly in this case, since little more than a Lambertian albedo value could be fit reliably from the four (or less) reflectance samples available at each point.

Finally, Fig. 12 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.

7.4. A Special Case: One Input Image

Due to its dimension-independence, 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. Thus, we can use a simplified SBRDF representation, computing the surface radiance according to

$$I_k(\vec{q}) = s_k \left(a_k(\vec{x}) + \sum_{i=1}^{N} \lambda_i \|\vec{q} - \vec{q}_i\| \right) cos\theta_i, \qquad (8)$$

where $\vec{q} = (x, y, u, v)$.

In this case, 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 the multi-image case. Figure 13 shows an example of a 2000-center SBRDF under natural lighting from environment maps. These were rendered using precomputation [NBB04] as discussed in Sect. 7.2, and the accompanying video demonstrates real-time manipulation of complex lighting.

Since a single image 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. [DHT+00], 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., Ref. [BG01].)

8. 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 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.

This paper 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 degrades gracefully as the number of input images decreases.

The most immediate practical issue for future work involves computational efficiency. We have demonstrated that 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 precomputed information. To improve this, it may be possible to develop real-time rendering techniques directly from

Figure 12: 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.

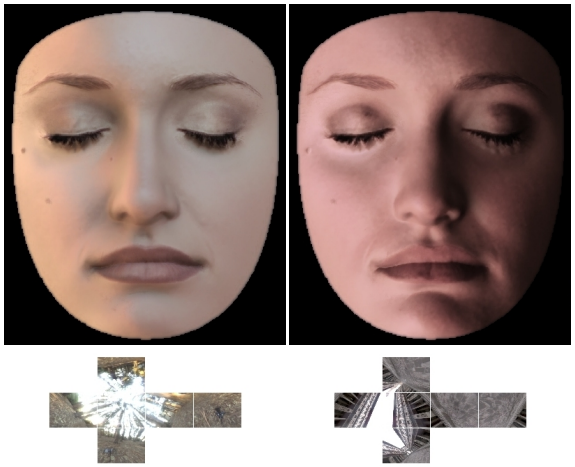

Figure 13: Images synthesized using the SBRDF representation in Eq. (8) estimated from the single (decomposed) image shown in Fig. 8. These were rendered in real-time using the methods discussed in Sect. 7.2.

the RBF representation. For example, fast multipole methods can be used to reduce the evaluation of Eq. (2) from $O(N^2)$ to $O(N \log N)$ [BN92]. This may be a viable alternative to using factored forms of BRDFs [MAA01, JM03] and may provide a practical approach for real-time rendering of surfaces with spatially-varying, non-parametric reflectance.

The increasing use of measured reflectance data in computer graphics requires efficient methods to acquire and represent such data. In this context, appropriate parameterizations, representations and signal-processing techniques are likely to be crucial. This paper presents a step in this direction by providing a method for using sparse, image-based datasets to accurately recover reflectance.

Acknowledgments

We thank Aner Ben-Artzi for his helpful comments and for his work on 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

[BG01] S. Boivin, A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *Proc. ACM SIGGRAPH*, pp. 107–116, 2001. 10

[BN92] R.K. Beatson and G.N. Newsam. Fast evaluation of radial basis functions: I. *Comput. Math. Appl.*, 24:7–19, 1992. 4, 11

[BP95] G. Beatson, R. Goodsell and M. Powell. On multigrid techniques for thin plate spline interpolation in two dimensions. *Lect. Appl. Mathematics*, 32:77–97, 1995. 4

[Buh03] M.D. Buhmann. *Radial basis functions*. Cambridge University Press, 2003. 4, 5

[CBC+01] J. Carr, R. Beatson, J. Cherrie, T. Mitchell, W. Fright, B. McCallum, T. Evans. Reconstruction and representation of 3D objects with radial basis functions. In *Proc. ACM SIGGRAPH*, pp. 67–76, 2001. 4, 5

[CBCG02] 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)*, 21(3):447–456, 2002. 9

[CON99] B. Cabral, M. Olano, and P. Nemec. Reflection space image based rendering. In *Proc. ACM SIGGRAPH*, pp. 165–170, 1999. 2

[CT81] R. Cook and K. Torrance. A reflectance model for computer graphics. *Computer Graphics (Proc. ACM SIGGRAPH)*, 15(3):307–316, 1981. 7

[DHT+00] P. Debevec, T. Hawkins, C. Tchou, H.P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Proc. ACM SIGGRAPH*, pp. 145–156, 2000. 2, 3, 7, 8, 9, 10

[DTS01] H. Dinh, G. Turk, and G. Slabaugh. Reconstructing surfaces using anisotropic basis functions. In *Proc. IEEE Int. Conf. Computer Vision*, pp. 606–613, 2001. 4

[Duc77] J. Duchon. Splines minimizing rotation-invariant semi-norms in Sobelev spaces. In W. Schempp and K. Zeller, eds., *Constructive theory of functions of several variables*, pp. 85–100. Springer-Verlag, 1977. 4

[Geo03] A.S. Georghiades. Recovering 3-D shape and reflectance from a small number of photographs. In *Rendering Techniques 2003 (Proc. Eurographics Symposium on Rendering)*, pp. 230–240, 2003. 1, 2

[GGH02] Xianfeng Gu, Steven Gortler, and Hugues Hoppe. Geometry images. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, 21(3):355–361, 2002. 6

[GGSC96] S.J. Gortler, R. Grzeszczuk, R. Szeliski, and M.F. Cohen. The lumigraph. In *Proc. ACM SIGGRAPH*, pp. 43–54, 1996. 3, 4

[HS03] A. Hertzmann and S. Seitz. Shape and material by example: A photometric stereo approach. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2003. 2

[JM03] R. Jaroszkiewicz and M. McCool. Fast extraction of BRDFs and material maps from images. In *Proc. Graphics Interface*, pp. 1–10, 2003. 2, 4, 8, 11

[KMG96] K. Karner, H. Mayer, M. Gervautz. An image based measurement system for anisotropic reflectance. *Computer Graphics Forum*, 15(3):119–128, 1996. 2

[LFTG97] E. Lafortune, S. Foo, K. Torrance, D. Greenberg. Non-linear approximation of reflectance functions. In *Proc. ACM SIGGRAPH*, pp. 117–126, 1997. 5

[LH96] M. Levoy and P. Hanrahan. Light field rendering. In *Proc. ACM SIGGRAPH*, pp. 31–42, 1996. 3

[LKG+01] 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. 2, 7, 8, 10

[LKK98] R. Lu, J. Koenderink, A. Kappers. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics*, 37:5974–5984, 1998. 2, 4, 5

[LSS+98] Aaron Lee, W. Sweldens, P. Schroder, L. Cowsar, and D. Dobkin. MAPS: Multiresolution adaptive parameterization of surfaces. In *Proc. ACM SIGGRAPH*, pp. 95–104, 1998. 6

[MAA01] M. McCool, J. Ang, and A. Ahmad. Homomorphic factorization of BRDFs for high-performance rendering. In *Proc. ACM SIGGRAPH*, p. 171–178, 2001. 2, 11

[Mar98] S. Marschner. *Inverse rendering for computer graphics*. PhD thesis, Cornell University, 1998. 3

[Mic86] C. A. Micchelli. Interpolation of scattered data: distance matrices and conditionally positive definite functions. *Constructive Approximation*, 1:11–22, 1986. 4

[MLH02] 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)*, pp. 79–88, 2002. 1, 5

[MPBM02] W. Matusik, H. Pfister, M. Brand, and L. McMillan. Image-based 3D photography using opacity hulls. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, 21(3):427–437, 2002. 2, 3, 7

[MPBM03] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, 22(3):759–769, 2003. 2, 4, 5

[MWL+99] 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)*, pp. 139–152, 1999. 2, 4, 5

[NBB04] S.K. Nayar, P.N. Belhumeur, and T.E. Boult. Lighting sensitive display. *ACM Trans. Graphics*, 23(4):963–979, 2004. 9, 10

[NFB97] S.K. Nayar, X. Fang, and T. Boult. Separation of reflection components using color and polarization. *Int. Journal of Computer Vision*, 21(3):163–186, 1997. 8

[ON94] M. Oren and S. Nayar. Generalization of Lambert's reflectance model. In *Proc. ACM SIGGRAPH*, pp. 239–246, 1994. 5

[Pow92] M.J.D. Powell. The theory of radial basis function approximation in 1990. In W. Light, editor, *Advances in Numerical Analysis, Vol. II*, pp. 105–210. Oxford Science Publications, 1992. 4, 6

[RH02] R. Ramamoorthi, P. Hanrahan. Frequency space environment map rendering. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, 21(3):517–526, 2002. 2

[Rus98] S. Rusinkiewicz. A new change of variables for efficient BRDF representation. In *Rendering Techniques '98 (Proc. Eurographics Rendering Workshop)*, pp. 11–22, 1998. 3

[Sha85] S. Shafer. Using color to separate reflection components. *COLOR res. appl.*, 10(4):210–218, 1985. 8

[SHHS03] P. Sloan, J. Hall, J. Hart, and J. Snyder. Clustered principal components for precomputed radiance transfer. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, 22(3):382–391, 2003. 9

[SWI97] Y. Sato, M. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Proc. ACM SIGGRAPH*, pp. 379–387, 1997. 1, 2, 3, 8

[Ter83] D. Terzopoulos. Multilevel computational processes for visual surface reconstruction. *Computer Vision, Graphics and Image Processing*, 24:52–96, 1983. 4

[WAA+00] 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*, pp. 287–296, 2000. 2, 7

[War92] Gregory J. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics (Proc. ACM SIGGRAPH)*, 26(2):265–272, 1992. 6

[YDMH99] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *Proc. ACM SIGGRAPH*, pp. 215–224, 1999. 1, 2