Passive Reflectometry

Fabiano Romeiro, Yuriy Vasilyev, and Todd Zickler

School of Engineering and Applied Sciences,
Harvard University, Cambridge MA 02138, USA
romeiro@fas.harvard.edu

Abstract. Different materials reflect light in different ways, so reflectance is a useful surface descriptor. Existing systems for measuring reflectance are cumbersome, however, and although the process can be streamlined using cameras, projectors and clever catadioptrics, it generally requires complex infrastructure. In this paper we propose a simpler method for inferring reflectance from images, one that eliminates the need for active lighting and exploits natural illumination instead. The method's distinguishing property is its ability to handle a broad class of isotropic reflectance functions, including those that are neither radially-symmetric nor well-represented by low-parameter reflectance models. The key to the approach is a bi-variate representation of isotropic reflectance that enables a tractable inference algorithm while maintaining generality. The resulting method requires only a camera, a light probe, and as little as one HDR image of a known, curved, homogeneous surface.

1 Introduction

Different surfaces modulate light in different ways, and this leads to distinctive lightness, gloss, sheen, haze and so on. Thus, like shape and color, surface reflectance can play a significant role in characterizing objects.

Computationally, surface reflectance is represented by the bi-directional reflectance distribution function, or BRDF; and the task of inferring the reflectance of a surface is formulated as that of inferring a BRDF from radiometric measurements. According to conventional methods, measuring surface reflectance requires the use of controlled, active lighting to sample the double-hemisphere of input and output directions that constitute the BRDF domain. These approaches demand complex infrastructure, including mechanical rotation and translation stages, digital cameras and projectors, and custom catadioptrics.

Perceptual studies suggest that humans can also infer reflectance information from image data, but that they do so in a very different manner. While the vast majority of machine measurement systems rely on illumination by a single moving point source, humans rely on images captured under complex, natural lighting [1]. The human approach has clear practical advantages: it is a passive technique that eliminates the need for controlled lighting, and it substantially reduces the measurement burden.

In this paper we present a passive system for inferring bi-directional surface reflectance that also exploits natural lighting. The approach is general in that,

D. Forsyth, P. Torr, and A. Zisserman (Eds.): ECCV 2008, Part IV, LNCS 5305, pp. 859–872, 2008.

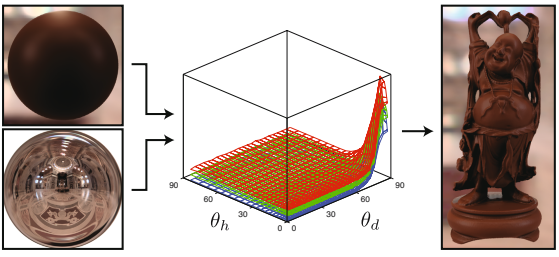

Fig. 1. Reflectometry using only a camera and a light probe (*bottom left*). Using a bivariate representation of reflectance, the constraints induced by a single HDR image (*top left*) of a known shape are sufficient to recover a non-parametric BRDF (*middle*). The recovered BRDF summarizes the object's reflectance properties and is an important material descriptor. Here, its accuracy is demonstrated through its use in rendering a synthetic image of a novel shape (*right*).

besides assuming isotropy, it can handle a rather unrestricted class of materials. It eliminates the need for active illumination and requires only a camera, a light probe, and as little as one HDR image of a known curved shape (Fig. 1).

The success of the approach hinges on its representation of surface reflectance. We avoid the restrictions of low-parameter BRDF models (Lambertian, Lafortune, Ward, Cook-Torrance, etc.) by using general bivariate functions. This approach is motivated by recent empirical studies [2,3], and our evaluations show that when used appropriately, this new representation provides a fine balance between tractability and generality. It enables reflectance to be reliably inferred from as little as one image, and at the same time, it is able to capture important phenomena such as specular and off-specular reflection, retro-reflection, and Fresnel effects. In particular, it is not constrained by any low-parameter analytic BRDF model; and unlike approaches based on de-convolution, it is not limited to radially-symmetric (i.e., one-dimensional) BRDFs.

1.1 Background and Related Work

The BRDF describes the manner in which incident radiant flux is modulated by a uniform surface patch. It is a positive function of four angular dimensions and can be written $f(\mathbf{u}, \mathbf{v})$, where $\mathbf{u}$ and $\mathbf{v}$ are unit vectors on the hemisphere centered about the patch normal. These are the directions of incident and reflected flux, respectively, and they are often expressed in spherical coordinates: (θ_u, ϕ_u) and (θ_v, ϕ_v).

One can measure the BRDF of a planar material by sampling the double hemisphere of input and output directions with a gonioreflectometer. Since this is extremely slow, and since a slight loss of accuracy is often acceptable, a number of camera-based alternatives have been proposed. When a camera is used with a curved mirror [4] or a curved material sample [5], one image provides a dense sampling of a 2D slice of the BRDF. To recover the complete BRDF domain, these can be combined with a moving light source (e.g., [5]) or a projector [6].

These camera-based systems significantly reduce measurement time, but they also require special-purpose hardware and precise lighting control.

Passive methods for reflectometry that require only natural lighting provide an attractive alternative. In the computer graphics community, the inference of reflectance from natural images has been studied under the banner of 'inverse rendering'. Ramamoorthi et al. [7] derive an elegant framework for inverse rendering by interpreting the rendering equation as a convolution. This yields an important theoretical tool that, among other things, enables the recovery of reflectance through de-convolution. Unfortunately, this approach can only yield general isotropic BRDFs when the full 4D output light field is observed. More typically, one has access to a small number of images; and when this is the case, de-convolution can only yield radially-symmetric BRDFs[1], which are incapable of representing off-specular peaks and important grazing-angle effects [2,3].

Inverse rendering can also be formulated, as it is here, directly in the angular domain. Many approaches exist, and almost all of them rely on low-parameter BRDF models (Phong, Cook-Torrance, etc.) to make the problem tractable. Low-parameter BRDF models impose strong constraints on reflectance, and as a result, one can exploit them to recover more than just reflectance information from a set of input images. For example, there are methods for handling global illumination effects and anisotropic reflectance [8,9], spatial reflectance variation, and the simultaneous recovery of illumination and/or shape (e.g., [10,11]). (Patow et al. [12] provide a review.) Every parametric approach suffers from limited accuracy, however, because the expressiveness of existing low-parameter BRDF models is quite restricted [2,3]. This situation is unlikely to improve in the short term. Given the diversity of the world's materials, designing 'general purpose' low-parameter models that are simultaneously accurate, flexible and amenable to tractable analysis has proven to be a very difficult problem.

Unlike these existing approaches, our goal is to recover *general* reflectance information without the restrictions of radial symmetry or low-parameter models. By avoiding these restrictions, we can handle a broader class of materials. To maintain this generality, we find it necessary to assume isotropic reflectance, ignore global illumination effects, and require that shape and illumination be known *a priori*. While the tools we develop can likely be applied to other inverse rendering problems (see discussion in Sect. 5), we leave this for future work.

2 A Bivariate BRDF for Reflectometry

Passive reflectometry is not well-posed without some constraints on the BRDF. Indeed, a BRDF is a function of four (angular) dimensions, while an input image is a function of two. What we require is a way to constrain the BRDF without surrendering our ability to represent important phenomena. Here, we present an approach based on a bivariate representation for isotropic surface reflectance.

[1] A radially-symmetric BRDF is one that, like the Phong model, is radially symmetric about the reflection vector. It's angular domain has dimension one.

For many materials, the dimension of the BRDF domain can be reduced without incurring a significant loss of detail. The domain can be folded in half, for example, because reciprocity ensures that BRDFs are symmetric about the directions of incidence and reflection: $f(\mathbf{u},\mathbf{v}) = f(\mathbf{v},\mathbf{u})$. In many cases, the domain $(\theta_u, \phi_u, \theta_v, \phi_v)$ can be further 'projected' onto the 3D domain $(\theta_u, \theta_v, \phi_u - \phi_v)$ and then folded onto $(\theta_u, \theta_v, |\phi_u - \phi_v|)$. The projection is acceptable whenever a BRDF exhibits little change for rotations of the input and output directions (as a fixed pair) about the surface normal; and additional folding is acceptable whenever there is little change when reflecting the output direction about the incident plane. Materials that satisfy these two criteria—for some definition of 'little change'—are said to satisfy *isotropy* and *bilateral symmetry*, respectively. (It is also common to use the term *isotropy* to mean both.)

It is convenient to parameterize the BRDF domain in terms of halfway and difference angles [13]. Accordingly, the complete 4D domain is written in terms of the spherical coordinates of the halfway vector $\mathbf{h} = (\mathbf{u}+\mathbf{v})/||\mathbf{u}-\mathbf{v}||$ and those of the input direction with respect to the halfway vector: $(\theta_h, \phi_h, \theta_d, \phi_d)$. See Fig. 2. In this parameterization, the folding due to reciprocity corresponds to $\phi_d \rightarrow \phi_d + \pi$, and the projection due to isotropy (without bilateral symmetry) is one onto $(\theta_h, \theta_d, \phi_d)$ [13]. While undocumented in the literature, it is straightforward to show that bilateral symmetry enables the additional folding $\phi_d \rightarrow \phi_d + \pi/2$ which gives the 3D domain $(\theta_h, \theta_d, \phi_d) \subset [0, \pi/2]^3$.

Here, we consider an additional projection of the BRDF domain, one that reduces it from three dimensions down to two. In particular, we project $(\theta_h, \theta_d, \phi_d) \subset [0, \pi/2]^3$ to $(\theta_h, \theta_d) \in [0, \pi/2]^2$. A physical interpretation is depicted in Fig. 2, from which it is clear that the projection is acceptable whenever a BRDF exhibits little change for rotations of the input and output directions (as a fixed pair) about the halfway vector. This is a direct generalization of isotropy, bilateral symmetry and reciprocity, which already restrict the BRDF to be $\frac{\pi}{2}$-periodic for the same rotations. We refer to materials that satisfy this requirement (again, for some definition of 'little change') as being *bivariate*. The accuracy of bivariate representations of the materials in the MERL BRDF

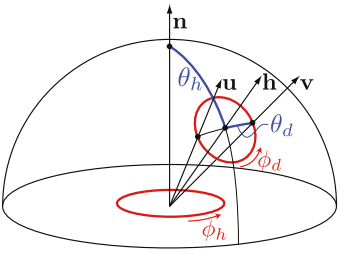

Fig. 2. Domain reduction for reciprocal, isotropic, bilaterally-symmetric, and bivariate BRDFs. Isotropic BRDFs are unchanged by rotations about the surface normal (i.e., changes in ϕ_h), while reciprocity and bilateral symmetry impose periodicity for rotations about the halfway vector (i.e., changes in ϕ_d). Here we consider *bivariate* BRDFs, which are constant functions of ϕ_d.

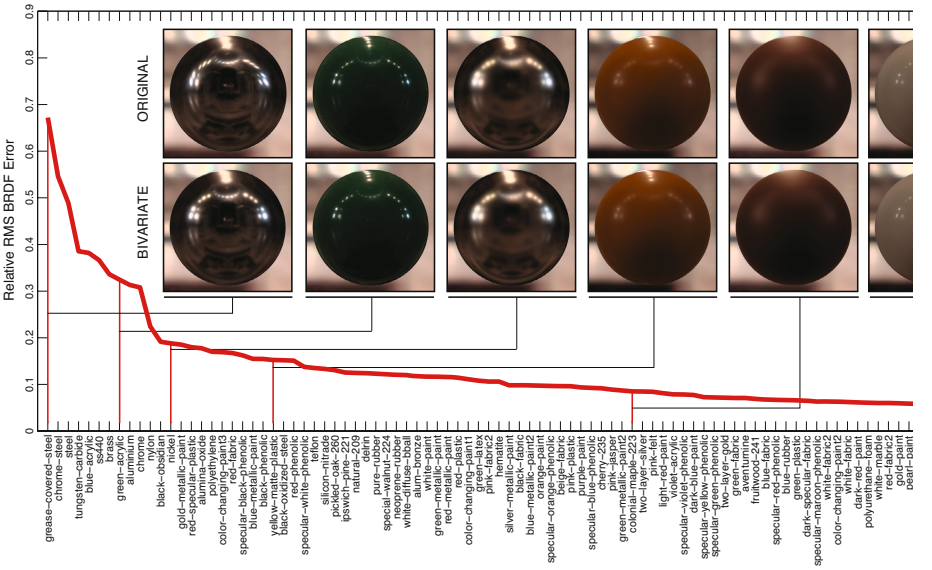

Fig. 3. Accuracy of bivariate representations of materials in the MERL BRDF database. Materials are in order of increasing accuracy, and representative renderings are shown for comparison. Most materials in the database are well-represented by a bivariate function. (Images embedded at high resolution; please zoom in.)

database [14] are shown in Fig. 3, where they are sorted by relative RMS BRDF error:

$$E_{\text{rms}} = \left(\sum_{\theta_h, \theta_d, \phi_d} \frac{(f(\theta_h, \theta_d, \phi_d) - \bar{f}(\theta_h, \theta_d))^2}{(f(\theta_h, \theta_d, \phi_d))^2} \right)^{\frac{1}{2}}, \tag{1}$$

with

$$\bar{f}(\theta_h, \theta_d) = \frac{1}{|\Phi(\theta_h, \theta_d)|} \sum_{\Phi(\theta_h, \theta_d)} f(\theta_h, \theta_d, \phi_d).$$

Here, $\Phi(\theta_h, \theta_d)$ is the set of valid ϕ_d values given fixed values of θ_h and θ_d. The figure also shows synthetic images of materials that are more and less well-represented by a bivariate BRDF. Overall, our tests suggest that the overwhelming majority of the materials in the database are reasonably well-represented by bivariate functions. We even find that the bivariate reduction has positive effects in some cases. For example, the original green-acrylic BRDF has lens flare artifacts embedded in its measurements[2], and these are removed by the bivariate reduction (see Fig. 3).

Motivation for a bivariate representation is provided by the work of Stark et al. [2] who show empirically that a carefully-selected 2D domain is often sufficient

[2] W. Matusik, personal communication.

for capturing (off-)specular reflections, retro-reflections, and important Fresnel effects. The 2D domain (θ_h, θ_d) that is introduced above is homeomorphic to that of Stark et al., which is why it posesses these same properties. Stark et al. propose the '$\alpha\sigma$-parameterization' for two-dimensional BRDFs, and this is related to (θ_h, θ_d) by

$$\alpha = \sin^2 \theta_d, \quad \sigma = \frac{1}{2}\left(1 + \cos 2\theta_d\right) \sin^2 \theta_h.$$

For this reason, Figs. 2 and 3 can be seen as providing a new interpretation and validation for their model. (The original paper examined Cornell BRDF data [15], which is arguably more accurate but also quite sparse.)

One important advantage of our (θ_h, θ_d) parameterization is that it provides an intuitive means for controlling how the 2D domain is sampled. This is explored next, where we use it for reflectometry.

3 Passive Reflectometry

We assume that we are given one or more images of a known curved surface, and that these images are acquired under known distant lighting, such as that measured by an illumination probe. In this case, each pixel in the images provides a linear constraint on the BRDF, and our goal is to infer the reflectance function from these constraints. While the constraints from a single image are not sufficient to recover a general 3D isotropic BRDF [7], we show that they often *are* sufficient to recover plausible bivariate reflectance.

To efficiently represent specular highlights, retro-reflections and Fresnel effects, we can benefit from a non-uniform sampling of the 2D domain. While 'good' sampling patterns can be learned from training data [16], this approach may limit our ability to generalize to new materials. Instead, we choose to manually design a sampling scheme that is informed by common observations of reflectance phenomena. This is implemented by defining continuous functions $s(\theta_h, \theta_d)$ and $t(\theta_h, \theta_d)$ and sampling uniformly in (s, t). Here we use $s = 2\theta_d/\pi$, $t = \sqrt{2\theta_h/\pi}$ which increases the sampling density near specular reflections $(\theta_h \approx 0)$. With this in mind, we write the rendering equation as

$$I(\mathbf{v}, \mathbf{n}) = \int_{\Omega} L(R_n^{-1}\mathbf{u}) f(s(\mathbf{u}, R_n\mathbf{v}), t(\mathbf{u}, R_n\mathbf{v})) \cos\theta_u d\mathbf{u}, \tag{2}$$

where $\mathbf{v}$ is the view direction, $\mathbf{n}$ is the surface normal corresponding to a given pixel, and R_n is the rotation that sends the surface normal to the z-axis and the view direction to the xz-plane. We use overloaded notation for s and t, which depend on the incident and reflected directions indirectly through (θ_h, θ_d).

At each pixel, this integral is computed over the visible hemisphere of light directions Ω. Our use of a bivariate BRDF induces a 'folding' of this hemisphere because light directions $\mathbf{u}$ and $\mathbf{u}'$ that are symmetric about the view/normal plane correspond to the same point in our 2D BRDF domain. When the lighting and surface shape are known, we obtain a constraint from each pixel, and each

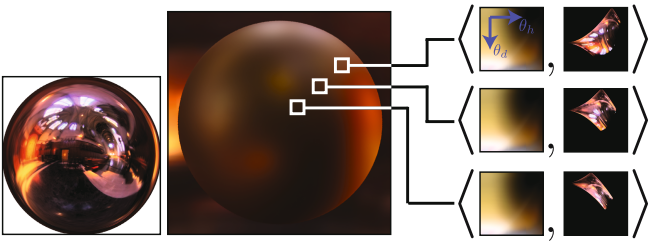

Fig. 4. Constraints on bivariate reflectance from natural lighting. Each pixel of an input image (*middle*) captured under distant illumination (*left*) gives a linear constraint that can be interpreted as an inner product of the 2D BRDF (*right*, first argument) and a visible hemisphere of lighting that is weighted, warped and folded across the local view/normal plane (*right*, second argument).

constraint can be interpreted as an inner product between the unknown BRDF and a hemisphere of illumination that is weighted by $\cos\theta_u$, folded across the local view/normal plane, and warped onto the *st*-plane. See Fig. 4.

To infer the BRDF from these constraints, we create a uniform grid $\mathcal{S} = \{(s_i, t_i)\}$ in the BRDF domain and approximate the rendering equation by a sum over a discrete set Ω_d of lighting directions on the hemisphere:

$$I(\mathbf{v}, \mathbf{n}) \approx \frac{2\pi}{|\Omega_d|} \sum_{\mathbf{u}_k \in \Omega_d} \left(\sum_{s_i, t_j \in N_k} \alpha_{i,j}^k L(R_n^{-1}\mathbf{u}_k) f(s_i, t_j) \right) \cos\theta_{u_k}, \tag{3}$$

where N_k is the set of the four BRDF grid points that are closest to $s(\mathbf{u}_k, R_n\mathbf{v})$, $t(\mathbf{u}_k, R_n\mathbf{v})$, and $\alpha_{i,j}^k$ is the coefficient of the bilinear interpolation associated with these coordinates and s_i, t_j. (We find a piecewise linear approximation of the BRDF to be adequate.) This equation can be rewritten as

$$I(\mathbf{v}, \mathbf{n}) \approx \frac{2\pi}{|\Omega_d|} \sum_{(s_i, t_j) \in \mathcal{S}} f(s_i, t_j) \sum_{\mathbf{u}_k \in \mathrm{bin}_{ij}} \alpha_{i,j}^k L(R_n^{-1}\mathbf{u}_k) \cos\theta_{u_k}, \tag{4}$$

to emphasize its interpretation as an inner product.

Observations of distinct normals $\mathbf{n}_1 \ldots \mathbf{n}_N$ obtained from one or more images provide constraints that are combined into a system of equations

$$I = Lf \tag{5}$$

where $I = [I(\mathbf{v}, \mathbf{n}_1), \ldots, I(\mathbf{v}, \mathbf{n}_N)]$ and L is a lighting matrix whose rows are given by the non-BRDF terms in Eq. 4. The goal is then to find f such that these constraints are satisfied. While this may work well in the noiseless case, in practice we require regularization to handle noise caused by the sensor, the bivariate approximation, the discretization of the rendering equation, and errors in the assumed surface shape.

As with general 4D BRDFs, bivariate BRDFs vary slowly over much of their domain. Regularization can therefore be implemented in the form of a smoothness constraint in the *st*-plane. There are many choices here, and we have found spatially-varying Tikhonov-like regularization to be especially effective. According to this design choice, the optimization becomes

$$\begin{aligned} &\underset{f}{\operatorname{argmin}} \|I - Lf\|_2^2 + \alpha \left(\left\| \Lambda_s^{-1} D_s f \right\|_2^2 + \left\| \Lambda_t^{-1} D_t f \right\|_2^2 \right) \\ &\text{subject to } f \geq 0, \end{aligned} \tag{6}$$

where D_s and D_t are $|\mathcal{S}| \times |\mathcal{S}|$ derivative matrices, and α is a tunable scalar regularization parameter. The matrices Λ_s and Λ_t are diagonal $|\mathcal{S}| \times |\mathcal{S}|$ matrices that affect non-uniform regularization in the bivariate BRDF domain. Their diagonal entries are learned from the MERL database by setting each to the variance of the partial derivative at the corresponding *st* domain point, where the variance is computed across all materials in the database. Probabilistically, this approach can be interpreted as seeking the MAP estimate with independent, zero-mean Gaussian priors on the bivariate BRDF's partial derivatives.

There are many possible alternatives for regularization. For example, one could learn a joint distribution over the entire bivariate domain, perhaps by characterizing this distribution in terms of a small number of modes of variation. However, we have found that the simple approach in Eq. 6 provides reasonable results, does not severely 'over-fit' the MERL database, and is computationally quite efficient (it is a constrained linear least squares problem).

3.1 Adequate Illumination

There is a question of when an environment is adequate for reflectometry to be well-posed and well-conditioned. An algebraic condition is readily available; we simply require the rank of the illumination matrix L to be sufficiently large (i.e., to approach $|\mathcal{S}|$). More intuitively, we require sufficient observations of all portions of the BRDF domain, with regions corresponding to specular reflections ($\theta_h \approx 0$), retro-reflections ($\theta_d \approx 0$), and grazing angles ($\theta_d \approx \pi/2$) being particularly important. In particular, we do not expect good results from simple environments composed of a small number of isolated point sources. This is in agreement with perceptual studies showing that humans are also unable to infer reflectance under such simple and 'unrealistic' conditions [1].

It is interesting to compare our approach to the convolution framework of Ramamoorthi et al. [7]. That approach enables a frequency domain analysis and provides very clear conditions for adequacy. For radially-symmetric BRDFs, for example, we know that an environment is adequate only if its band-limit exceeds that of the BRDF [7]. A frequency domain analysis is difficult to apply in the present case, however, because Eq. 5 does not represent a convolution. While an analysis of the conditions for adequate illumination in the bivariate case may be for worthwhile direction of future work, we focus instead on an empirical investigation here. We show that while the quality of the result depends on the environment, accurate reflectometry is achievable in many cases.

4 Evaluation and Results

We begin with an evaluation that uses images synthesized with tabulated BRDF data from the MERL database [14], measured illumination[3], and a physically based renderer[4]. Using these tools, we can render images for input to our algorithm as well as images with the recovered BRDFs for direct comparison to ground truth. In all cases, we use complete 3D isotropic BRDF data to create the images for input and ground-truth comparison, since this is closest to a real-world setting. Also, we focus our attention on the minimal case of a single input image; with additional images, the performance can only improve. It is worth emphasizing that this data is not free of noise. Sources of error include the fact that the input image is rendered with a 3D BRDF as opposed to a bivariate one, that normals are computed from a mesh and are stored at single precision, and that a discrete approximation to the rendering equation is used.

Given a rendered input image of a defined shape (we use a sphere for simplicity), we harvest observations from 8,000 normals uniformly sampled on the visible hemisphere to create an observation vector I of length 8,000. We discard normals that are at an angle of more than 80° from the viewing direction, since the signal to noise ratio is very low at these points. The bivariate BRDF domain is represented using a regular 32×32 grid on the st-plane, and our observation matrix L is therefore $M \times 1024$, where M is the number of useable normals. The entries in L are computed using Eq. 4 with 32,000 points uniformly distributed on the illumination hemisphere. With I and L determined, we can solve for the unknown BRDF as described in the previous sections.

We find it beneficial to use a small variant of the optimization in Eq. 6: we solve the problem twice using two separate pairs of diagonal weight matrices (Λ_s, Λ_t). One pair gives preference to diffuse reflectance, while the other gives preference to gloss. This provides two solutions, and we choose the one with lowest residual. Using this procedure, we were able to use the same weight matrices and regularization parameter (α) for all results in this paper. In every case, the optimizations were initialized with a Lambertian BRDF.

Results are shown in Fig. 5. The two left columns show results using a single input image synthesized with the Grace Cathedral environment. The recovered bivariate BRDFs are compared to the (3D) ground truth by synthesizing images in another setting (St. Peter's Basilica). Close inspection reveals very little noticeable difference between the two images, and the recovered BRDF is visually quite accurate. There are numerical differences, however, and these have been scaled by 100 for visualization. Note that some of this error is simply due to the bivariate approximation (see Fig. 6). The next two columns similarly show the recovery of the yellow-matte-plastic and green-acrylic materials, this time using the Cafe environment and the St. Peter's Basilica environment (alternately) for input and comparison to ground truth.

[3] Light probe image gallery: `http://www.debevec.org/Probes/`
[4] PBRT: `http://www.pbrt.org/`

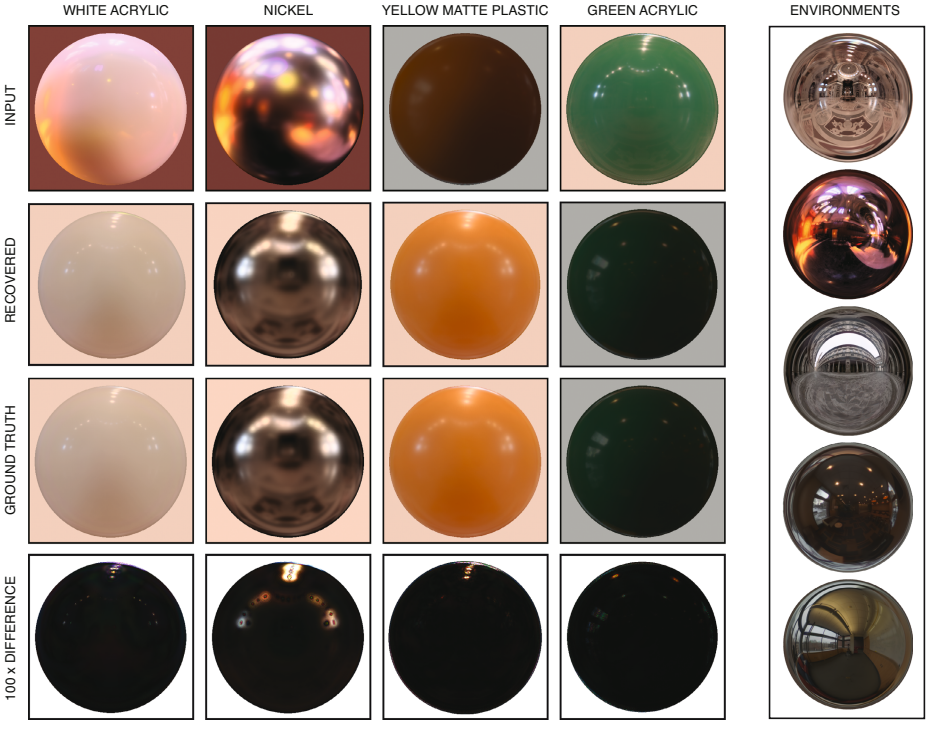

Fig. 5. Visual evaluation with MERL BRDF data. A bivariate BRDF is estimated from a single input image (*top*), and this estimate is used to render a new image under novel lighting (*second row*). Ground truth images for the novel environments are shown for comparison, along with difference images scaled by 100. Few noticeable differences exist. *Far right*: Environment maps used in the paper, top to bottom: St. Peter's Basilica, Grace Cathedral, Uffizi Gallery, Cafe and Corner Office.

In addition to these visual comparisons, we can also evaluate the recovered BRDFs quantitatively using scatter plots and RMS errors. The top of Fig. 6 shows incident-plane scatter plots for the red channels of three recovered BRDFs from Fig. 5, as well as the recovered colonial-maple BRDF from Fig. 1. While the scatter plots reveal clear deviations from ground truth, they suggest that the approach provides reasonable approximations for a variety of materials. This is true even though just a single image is used as input—many fewer than the 300 images that were used to collect the original data [14].

The bottom of the figure displays relative RMS errors for these four recovered BRDFs, along with corresponding results for all materials in the BRDF database. Shown is the accuracy (Eq. 1) of the bivariate BRDF for each material as estimated from one input image. This is done twice—once each using the Grace Cathedral and St. Peter's environments—and the curves are superimposed on the graph from Fig. 3, which shows the accuracy of the 'ground truth' bivariate reduction. (Note that the materials have been re-sorted for display

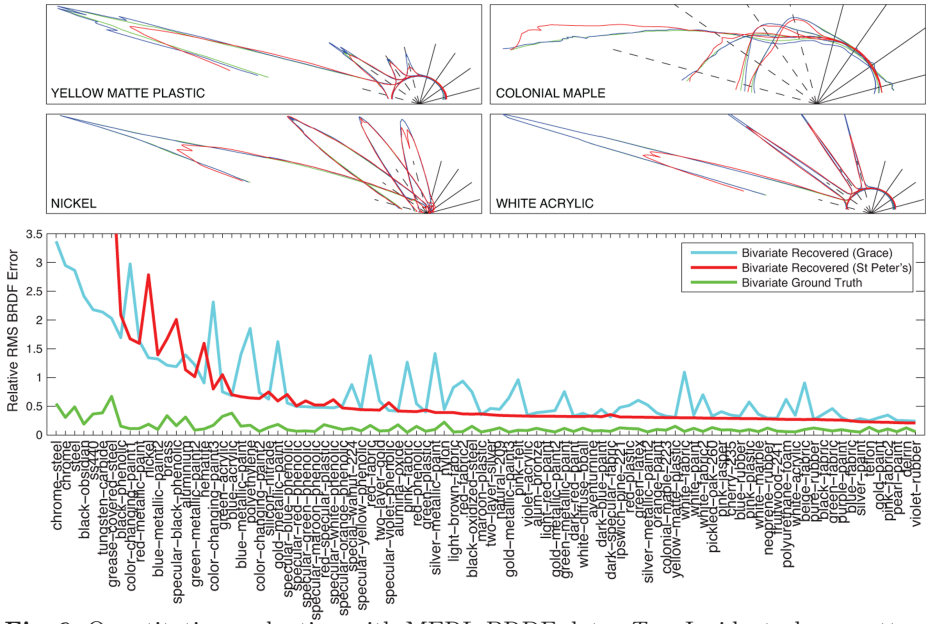

Fig. 6. Quantitative evaluation with MERL BRDF data. *Top*: Incident plane scatterplots for the four materials in Fig. 5, each showing: original 3D BRDF (*blue*); 'ground truth' bivariate BRDF (*green*); and BRDF recovered from one input image (*red*). *Bottom*: Relative RMS BRDF errors for all materials in the MERL database when each is recovered using a single image under the Grace Cathedral or St. Peter's environments. Vertical red lines match the scatterplots above.

purposes). The discrepancy between the results for the two different environments is expected in light of the discussion from Sect. 3.1. To further emphasize this environment-dependence, Fig. 7 compares estimates of yellow-matte-plastic using two different input images. The Uffizi Gallery environment (top left) does not provide strong observations of grazing angle effects, so this portion of the BRDF is not accurately estimated. This leads to noticeable artifacts near grazing angles when the recovered BRDF is used for rendering, and it is clearly visible in a scatter plot. When the Cafe environment is used as input, however, more accurate behavior near grazing angles is obtained.

4.1 Captured Data

The procedure outlined above was applied without change to captured data. Figure 8 shows the results for a number of materials. As before, each BRDF is recovered from a single input image (left), and the recovered BRDFs are used to render synthetic images of the same object from a novel viewpoint. The synthetic images are directly compared to real images captured in the same novel positions.

Captured data contains at least three significant sources of noise in addition to what exists in the rendered data above: 1) errors in the assumed surface geometry;

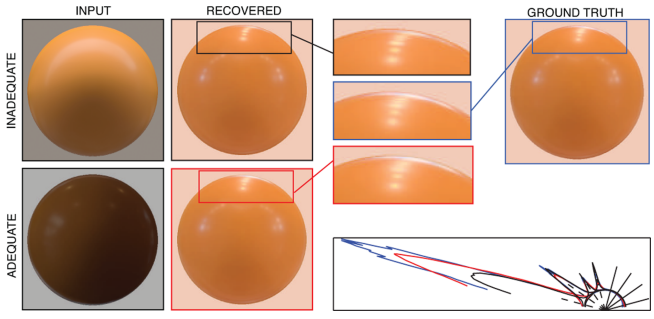

Fig. 7. Dependence on environment used for capture. An input image under the Uffizi Gallery environment (*top left*) does not contain strong observations of grazing angle effects, and as a result, the recovered BRDF is inaccurate. This is visible in a scatter plot (*bottom right*, black curves) and causes noticeable artifacts when used to render in a novel setting. If a different environment is used as input (*bottom left*) these artifacts are largely avoided.

Fig. 8. Results using captured data. A BRDF is estimated from a single input image (*top*) under a known environment. This recovered BRDF is used to render a synthetic image for novel view within the same environment (*middle*). An actual image for the same novel position is shown for comparison (*bottom*). Despite the existence of non-idealities such as surface mesostructure and spatial inhomogeneity, plausible BRDFs are recovered.

2) surface mesostructure (e.g., the green sphere); and 3) spatial reflectance variations (e.g., the grey sphere). Presently, surface shape is computed by assuming the camera to be orthographic and estimating the center and radius of the sphere in the camera's coordinate system. Errors in this process, coupled with errors in

the alignment with the illumination probe, lead to structured measurement noise. Despite this, our results suggest that plausible BRDFs can be recovered for a diversity of materials.

5 Discussion

This paper presents a technique for 'light-weight' reflectometry that eliminates the need for active illumination and requires minimal infrastructure for acquisition. This is enabled by reducing the domain of isotropic bi-directional reflectance functions from three dimensions to two. We provide an empirical evaluation of this reduced representation that compliments recent work [2].

The proposed approach has clear advantages over existing inverse rendering techniques that recover reflectance from 2D images using de-convolution or low-parameter BRDF models. These existing methods recover reflectance functions that are one-dimensional (radially-symmetric) or zero-dimensional (parametric), respectively. In contrast, the method presented here recovers a two-dimensional reflectance function, and thereby matches the dimension of the output with that of the input. For this reason, it can be applied to a much broader class of surfaces.

One of the important things we give up in exchange for generality is the intuition provided by the convolution framework. It becomes difficult to characterize the necessary conditions for adequate illumination, and this suggests a direction for future work. In particular, it may be possible to clarify the role that 'environment foldings' (Fig. 4) play in reducing redundancy in L and 'enhancing the adequacy' of an environment.

There are a number of additional directions for future work. We presented one of many possible regularization schemes, and it is possible that others are more suitable. In exploring this possibility, one must be wary of 'overfitting' existing BRDF databases, since these may provide descriptions of only a fraction of the world's interesting materials. We have largely avoided this in our approach, but even so, we expect our method to be less successful for highly retro-reflective surfaces, which are not well represented in the MERL database.

Our focus in this work is the recovery of *general* reflectance functions, meaning those that are not necessarily well-represented by low-parameter models and those that are not radially-symmetric. For this reason, we considered the case in which the surface is homogeneous, its shape is known, and the illumination environment is also known. Relaxing these conditions is perhaps the most interesting direction for future work, and it is quite likely that the tools presented here will prove useful elsewhere (see [17] for a reconstruction application).

In this vein, the proposed framework provides an opportunity to explore the joint recovery of reflectance and illumination (f and L in Eq. 6), or at least the recovery of reflectance when lighting is unknown. Using our framework, this essentially becomes a blind de-convolution problem. It is possible that this line of research may eventually yield computational systems that can match the human ability to infer reflectance in uncontrolled conditions [1].

Acknowledgements

We thank Wojciech Matusik for helpful discussions regarding the MERL database. Support comes from an NSF CAREER award and a Sloan Foundation fellowship.

References

1. Fleming, R., Dror, R.O., Adelson, E.H.: Real-world illumination and the perception of surface reflectance properties. Journal of Vision 3 (2003)
2. Stark, M., Arvo, J., Smits, B.: Barycentric parameterizations for isotropic BRDFs. IEEE Transactions on Visualization and Computer Graphics 11, 126–138 (2005)
3. Ngan, A., Durand, F., Matusik, W.: Experimental analysis of brdf models. In: Eurographics Symposium on Rendering, pp. 117–126 (2005)
4. Ward, G.: Measuring and modeling anisotropic reflection. Computer Graphics (Proc. ACM SIGGRAPH) (1992)
5. Marschner, S., Westin, S., Lafortune, E., Torrance, K., Greenberg, D.: Image-based BRDF measurement including human skin. In: Proc. Eurographics Symposium on Rendering, pp. 139–152 (1999)
6. Ghosh, A., Achutha, S., Heidrich, W., O'Toole, M.: BRDF acquisition with basis illumination. In: Proc. IEEE Int. Conf. Computer Vision (2007)
7. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. In: Proceedings of ACM SIGGRAPH, pp. 117–128 (2001)
8. Boivin, S., Gagalowicz, A.: Image-based rendering of diffuse, specular and glossy surfaces from a single image. In: Proceedings of ACM SIGGRAPH (2001)
9. Yu, Y., Debevec, P., Malik, J., Hawkins, T.: Inverse global illumination: recovering reflectance models of real scenes from photographs. In: Proceedings of ACM SIGGRAPH (1999)
10. Georghiades, A.: Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In: Proc. IEEE Int. Conf. Computer Vision, pp. 816–823 (2003)
11. Hara, K., Nishino, K., Ikeuchi, K.: Mixture of spherical distributions for single-view relighting. IEEE Trans. Pattern Analysis and Machine Intelligence 30, 25–35 (2008)
12. Patow, G., Pueyo, X.: A Survey of Inverse Rendering Problems. Computer Graphics Forum 22, 663–687 (2003)
13. Rusinkiewicz, S.: A new change of variables for efficient BRDF representation. In: Eurographics Rendering Workshop, vol. 98, pp. 11–22 (1998)
14. Matusik, W., Pfister, H., Brand, M., McMillan, L.: A data-driven reflectance model. ACM Transactions on Graphics (Proc. ACM SIGGRAPH) (2003)
15. Westin, S.: Measurement data, Cornell University Program of Computer Graphics (2003), `http://www.graphics.cornell.edu/online/measurements/`
16. Matusik, W., Pfister, H., Brand, M., McMillan, L.: Efficient isotropic BRDF measurement. In: Proc. Eurographics Workshop on Rendering, pp. 241–247 (2003)
17. Alldrin, N., Zickler, T., Kriegman, D.: Photometric stereo with non-parametric and spatially-varying reflectance. In: Proc. CVPR (2008)