Sparse Sampling for Image-Based SVBRDF Acquisition

Jiyang Yu, Zexiang Xu, Matteo Mannino, Henrik Wann Jensen and Ravi Ramamoorthi
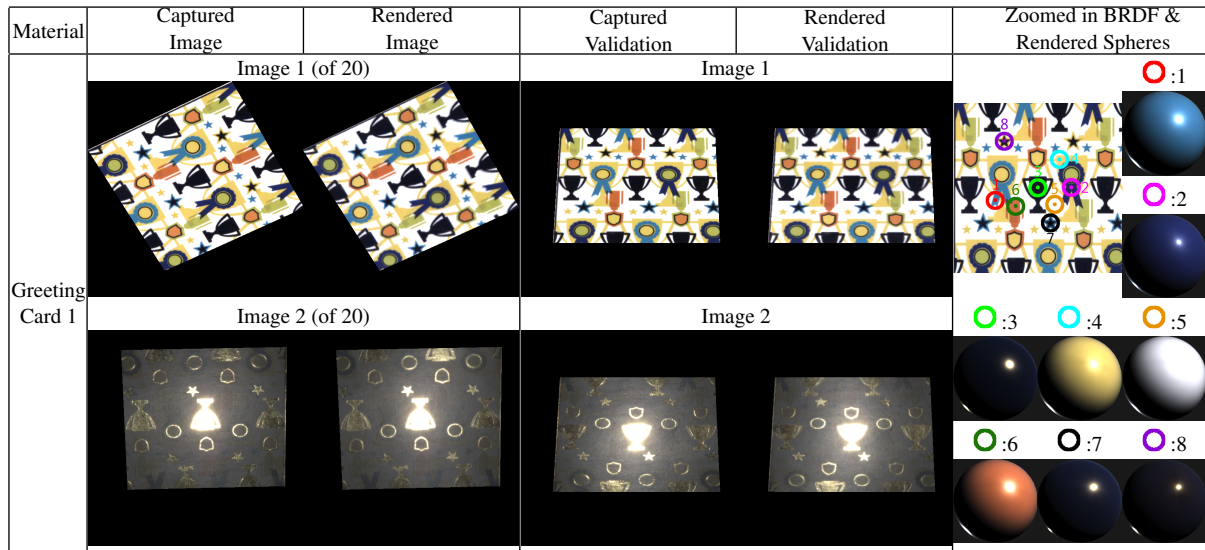
University of California, San Diego

Figure 1: We acquire the measured SVBRDF of a flat material sample using only 20 input images captured on a spherical gantry, with optimized light-view directions. As can be seen, the rendered images match the captured data, and new validation views not used as input at all. We recover an independent BRDF for each point on the surface, visualized by the rendered spheres.

Abstract

We acquire the data-driven spatially-varying (SV)BRDF of a flat sample from only a small number of images (typically 20). We generalize the homogenous BRDF acquisition work of Nielsen et al., who derived an optimal minmal set of lighting/view directions, treating a 4 degree-of-freedom spherical gantry as a gonioreflectometer. In contrast, we benefit from using the full 2D camera image from the gantry to enable SVBRDF acquisition. Like Nielsen et al, our method is data-driven, based on the MERL database of isotropic BRDFs, and finds the optimal directions by minimizing the condition number of the acquisition matrix. We extend their approach to SVBRDFs by modifying the optimal incident/outgoing directions to avoid grazing angles that reduce resolution and make alignment of different views difficult. Another key practical issue is aligning multiple viewpoints, and correcting for near-field effects. We demonstrate our method on SVBRDF measurements of new flat materials, showing that full data-driven SVBRDF acquisition is now possible from a sparse set of only about 20 light-view pairs.

1. Introduction

Parametric BRDFs such as the Phong or Ward models have been widely used, but greater fidelity is provided by mea-

sured data-driven reflectance, and databases such as the MERL BRDFs [MPBM03]. However, each istotropic 3D BRDF in this database contains over a million samples, and

extension to spatially-varying BRDF acquisition is impractical (general SVBRDFs are 6D; we consider isotropic materials, hence 5D). Even when the final SVBRDF representation is factored and compressed [LBAD*06], thousands of images are initially required (we focus only on initial acquisition; such techniques can be applied later if desired).

This paper describes an image-based SVBRDF acquisition system for flat material samples, using a 4 degree-of-freedom spherical gantry to position the light source and camera. We usually need only 20 images (optimized light-view direction pairs), and we provide these directions for other researchers (Fig. 3). As opposed to other methods for efficient SVBRDF acquisition [AWL13, TFG*13], ours is a direct method, without needing to project patterns or apply frequency-space analysis or deconvolution. We also recover an independent BRDF at each point/pixel (Fig. 1), not just a spatial modulation or texture-like quantity [AWL15].

Our paper builds closely on [NJR15], which developed a logarithmic mapping of the MERL database, enabling a simple linear BRDF subspace. Their method also uses a spherical gantry, but treats it as a gonioreflectometer, simply averaging the camera image across the sample. In contrast, we consider the full 2D camera image to acquire SVBRDFs, essentially applying [NJR15] independently to each pixel/point on the material.

The optimal directions in [NJR15] include several outgoing directions close to grazing angles to measure Fresnel effects. However, this causes a loss in spatial resolution, and makes precise alignment between different images/views of a spatially-varying material difficult. We optimize for a new set of directions specialized for acquisition of SVBRDFs. We also address other practical concerns, such as alignment of different views and correction for near-field camera effects over the material sample (Fig. 2). We use simulations on MERL BRDFs to validate the stability of results using our optimal directions (Fig. 4). Practical results show we can reconstruct accurate SVBRDFs of flat samples of unknown real materials from only a few input images (Figs. 1, 5).

2. Background

We base our reconstruction and optimization on the MERL isotropic BRDF database [MPBM03]. This database includes 100 materials, each having 3 color-channels that are treated as independent materials. A total of $m = 300$ BRDFs are considered. Each material is represented using $p = 1,458,000$ measurements in the $(\theta_h, \theta_d, \phi_d)$ parameterization [Rus98] space, with resolution of $90 \times 90 \times 180$. We apply the logarithmic mapping of [NJR15], which transforms BRDF ρ to $\ln[(\rho \cdot \cos_{weight} + \varepsilon)/(\rho_{ref} \cdot \cos_{wieght} + \varepsilon)]$, where $\varepsilon = 0.001$ avoids dividing by zero, and $\cos_{weight} = \max\{\cos(\mathbf{n} \cdot \omega_{\mathbf{i}})\cos(\mathbf{n} \cdot \omega_{\mathbf{o}}), \varepsilon\}$, where n, ω_0, ω_i are the normal, incident and outgoing direction vectors respectively.

Let $X \in \mathbb{R}^{m \times p}$ be the full matrix of all MERL BRDF


Figure 2: Showing near-field effects and view alignment, which we address for practical SVBRDF reconstruction.

observations, where the rows are mapped BRDFs and the columns are sampling directions. We apply principal component analysis on X and get the principal components $Q \in \mathbb{R}^{p \times k}$. In our case we choose the number of samples/images we take as k. Therefore, an arbitrary BRDF x can be obtained as a linear combination of the basis [NJR15],

$$x = Qc + \mu, \qquad (1)$$

where $c \in \mathbb{R}^{k \times 1}$ is a vector of coefficients. $\mu \in \mathbb{R}^{p \times 1}$ is the mean BRDF. In practice, we observe x at some set of n chosen directions, from which we can estimate c as follows:

$$\tilde{x} - \tilde{\mu} = \tilde{Q}c, \qquad (2)$$

where the tildes indicate that we have a reduced set of observations at n directions, with $\tilde{\mu}$ and $\tilde{x} \in \mathbb{R}^{n \times 1}$. $\tilde{Q} \in \mathbb{R}^{n \times k}$ is the set of rows in Q corresponding to the set of observed directions. The above equation can be solved using ridge regression (I is the identity matrix. We use $\eta = 40$.)

$$c = \operatorname{argmin}|(\tilde{x} - \tilde{\mu}) - \tilde{Q}c|^2 + \eta|c|^2 = (\tilde{Q}^T\tilde{Q} + \eta I)^{-1}\tilde{Q}^T(\tilde{x} - \tilde{\mu}). \qquad (3)$$

Therefore given some captured (homogeneous) BRDF samples $\tilde{x}$, the coefficients c can be solved. Then, the full BRDF is recovered by using equation 1.

3. Optimal Directions and Reconstruction

In this paper, we focus on measuring the SVBRDFs of flat planar materials. We assume that the reflectance of any point on the material's surface is an isotropic BRDF. In principle, it is straightforward to apply the BRDF acquisition method above separately to each point on the surface. However, there are some non-trivial practical issues, which we address next.

3.1. Near-Field Effects and View Alignment

Unlike for homogeneous BRDFs, the sample now occupies a finite area within the field of view for the camera (about $4°$ in our gantry setup). Therefore, in a perspective camera image,

each point/pixel is at a slightly different direction with respect to the camera, as shown in Fig. 2. This means that each pixel will have a slightly different $\tilde{Q}$ matrix, but equation 3 can still be applied to reconstruct the BRDF separately at all positions on the material's surface. (We assume the light is distant, but near-field lighting effects are also easy to correct for, since they also correspond simply to modifying $\tilde{Q}$).

Another important issue is alignment of multiple camera views, which must be put in correspondence in order to avoid blurring in SVBRDF reconstruction, as shown in Fig. 2. We tried a number of solutions involving standard homographies, but found that care must be taken to avoid inaccuracies and blurring/ghosting. As discussed in Sec. 4, we pick a reference view, and for each pixel in it, project to the other views using intrinsic and extrinsic camera calibration.

3.2. Optimal Directions

[NJR15] provide optimal directions for homogeneous BRDF acquisition by minimizing the condition number of $\tilde{Q}$. However, we found that simply using these directions and doing a pixel-by-pixel reconstruction on spatially varying materials yields poor results. One reason is that [NJR15] include many viewing directions close to grazing angles. For SVBRDFs, this will result in a significant loss of resolution at oblique angles, where the sample covers only a small field of view. In addition, a precise image alignment to the reference view is difficult to establish. Hence, we re-optimize the sampling directions with an additional constraint that avoids grazing views. We start from 1 sample direction and iteratively do optimization and add one new direction until we get the desired number of n directions. After adding one new direction, similar to [NJR15], we optimize the directions by alternatively moving each direction along the gradient of $\tilde{Q}$, but we only allow changes if the viewing angle is $< 65°$.

Figure 3 shows the optimized 20 sampling directions for SVBRDF reconstruction in Rusinkiewicz coordinates, as well as thumbnails of the corresponding input images. Note that these camera directions apply to the center of the sample; we correct for near-field effects at other pixels, so most points on the material are sampled using directions close, but not exactly equal, to the optimal directions in Fig. 3.

3.3. Verification by Simulation

In order to validate our new directions, we first recover BRDFs in simulation on the 100-material MERL database. This simulation is done on homogenous BRDFs, but assuming a finite sample and field of view, as in Fig. 2. We use 90 materials as training data to calculate the basis matrix $\tilde{Q}$ and use this basis to recover the other 10 materials. We account for the camera field of view, by considering 100 uniformly sampled spatial positions on the flat sample.

It is shown in Fig. 4 that sampling from our close to

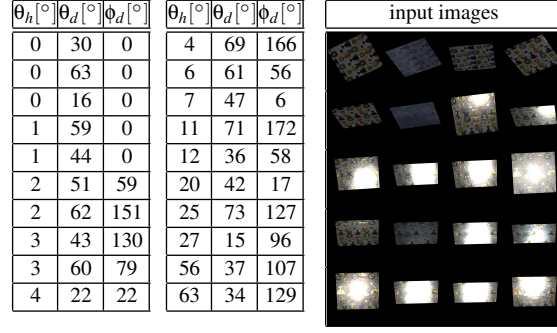| $\theta_h[°]$	$\theta_d[°]$	$\phi_d[°]$	$\theta_h[°]$	$\theta_d[°]$	$\phi_d[°]$
0	30	0	4	69	166
0	63	0	6	61	56
0	16	0	7	47	6
1	59	0	11	71	172
1	44	0	12	36	58
2	51	59	20	42	17
2	62	151	25	73	127
3	43	130	27	15	96
3	60	79	56	37	107
4	22	22	63	34	129

Figure 3: Optimized sampling directions for SVBRDF reconstruction in Rusinkiewicz coordinates. We also show thumbnails of the input images for one of our samples.


Figure 4: Average reconstruction error for our new optimized directions is comparable to previous work, and does not degrade significantly for practical near-field effects.

optimal directions generates comparable error to using the same number of optimal directions from [NJR15]. Moreover, when using the $4°$ camera field of view typical for our spherical gantry setup and acquisition samples, the error does not degrade significantly, even though most points on the material are using slightly different, sub-optimal directions. However, for a larger $10°$ field of view, the errors are more significant, since points at the edge of the sample will have viewing directions far from optimal.

4. Results: Real SVBRDF Measurements

The acquisition device is a high-precision spherical gantry with two robotic arms. A light source and a camera are attached to the two arms, which can be accurately positioned anywhere in the hemisphere around the sample.

We make the common pin-hole camera approximation, and use the method proposed in [Zha99] to compute camera intrinsics using 39 images of a checkerboard at different viewing directions. The checkerboard images for BRDF sampling directions in Fig. 3 are also captured for estimating extrinsic parameters. We also capture an additional checkerboard image by placing the camera at $(\theta, \phi) = (0°, 0°)$. We will call this view the reference view for simplicity. We

Figure 5: Unknown SVBRDFs recovered by our method. We render our reconstructed results for several views with different lighting conditions, and compare them with captured images. These validation views were not used in reconstruction at all.

base our SVBRDFs' resolution on the image of this reference view. Every final recovered SVBRDF contains a set of isotropic BRDFs, each of which corresponds to one small region on the material's surface intersected by a one pixel cone from the reference view. For actual measurements, we use multiple exposures to obtain high-dynamic range images of the flat sample from each light-view sampling direction in Fig. 3. The same procedure is also performed on a spectralon board to calibrate the light intensity.

We find pixel-wise correspondence between sampled images, by transforming the pixels' positions in the reference view to each view using calibrated information of all views. For each spatial position (pixel) in the reference view, we have now measured BRDF samples and corresponding directions from all captured views. We then use the reconstruction framework described in Sec. 2 to recover every position's isotropic BRDF, and a final SVBRDF is generated.

To demonstrate the practicality and efficiency of our method, we measured three different materials' SVBRDFs. These are all greeting cards with many different reflectances for the letters, decorative items and background. The appearance can vary dramatically between diffuse and specular directions, and is not usually well modeled by parametric reflectance models. It is shown in Figs. 1 and 5 that our recovered SVBRDFs look very close to ground truth in both input views and validation views that were not used at all when reconstructing the SVBRDF. We also show spheres rendered with the BRDF recovered at several different pixels; we see that BRDFs at each pixel are independent and accurate, with no artifacts. Close inspection indicates the major limitations are from not accounting for fine-scale normal variation.

References

[AWL13] AITTALA M., WEYRICH T., LEHTINEN J.: Practical SVBRDF capture in the frequency domain. *ACM Transactions on Graphics (TOG) 32*, 4 (2013), 110.

[AWL15] AITTALA M., WEYRICH T., LEHTINEN J.: Two-shot SVBRDF capture for stationary materials. *ACM Transactions on Graphics (TOG) 34*, 4 (2015).

[LBAD*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHI R., RUSINKIEWICZ S.: Inverse Shade Trees for Non-Parametric Material Representation and Editing. *ACM Transactions on Graphics (Proc. SIGGRAPH 06) 25*, 3 (2006), 735–745.

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. In *ACM Transactions on Graphics (TOG)* (2003), pp. 759–769.

[NJR15] NIELSEN J. B., JENSEN H., RAMAMOORTHI R.: On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG) 34*, 6 (2015).

[Rus98] RUSINKIEWICZ S.: A new change of variables for efficient BRDF representation. In *Eurographics Rendering Workshop* (1998), pp. 11–22.

[TFG*13] TUNWATTANAPONG B., FYFFE G., GRAHAM P., BUSCH J., YU X., GHOSH A., DEBEVEC P.: Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Transactions on Graphics (TOG) 32*, 4 (2013).

[Zha99] ZHANG Z.: Flexible camera calibration by viewing a plane from unknown orientations. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on* (1999), vol. 1, IEEE, pp. 666–673.