

Minimal BRDF Sampling for Two-Shot Near-Field Reflectance Acquisition

Zexiang Xu¹ Jannik Boll Nielsen² Jiyang Yu¹
¹University of California, San Diego

Henrik Wann Jensen¹ Ravi Ramamoorthi^{1*}
²Technical University of Denmark

Abstract

We develop a method to acquire the BRDF of a homogeneous flat sample from only two images, taken by a near-field perspective camera, and lit by a directional light source. Our method uses the MERL BRDF database to determine the optimal set of light-view pairs for data-driven reflectance acquisition. We develop a mathematical framework to estimate error from a given set of measurements, including the use of multiple measurements in an image simultaneously, as needed for acquisition from near-field setups. The novel error metric is essential in the near-field case, where we show that using the condition-number alone performs poorly. We demonstrate practical near-field acquisition of BRDFs from only one or two input images. Our framework generalizes to configurations like a fixed camera setup, where we also develop a simple extension to spatially-varying BRDFs by clustering the materials.

Keywords: rendering, reflectance, BRDF, MERL, reconstruction

Concepts: •Computing methodologies → Reflectance modeling;

1 Introduction

Accurate BRDF models are critical for realistic image synthesis. Many analytic BRDF models have been proposed [Torrance and Sparrow 1967; Ward 1992]. However, the greatest fidelity is obtained by data-driven reflectance, such as the MERL BRDF database of 100 real materials [Matusik et al. 2003a].

In this paper, we focus on the canonical problem of measuring the 3D isotropic BRDF of a flat sample of homogeneous material. The conventional approach is to use a gonireflectometer, laboriously sampling illumination-view pairs [Foo 1997]. However, fully sampling a 3D isotropic BRDF domain can require thousands or millions of samples, making this approach very expensive. Mirror-based imaging setups [Ward 1992] can reduce some dimensions, but still require a large number of samples. They also need more complex setups, and can be difficult to calibrate.

Recently, [Nielsen et al. 2015] presented a significant reduction in the number of samples needed to about 20, assuming the BRDF lies approximately in the subspace of the MERL BRDF database. They leverage a logarithmic mapping (originally proposed for BRDF factorization by [McCool et al. 2001]). They then optimize sampling directions to minimize the condition number of the related acquisition matrix (an approach first proposed by [Matusik et al. 2003b]).

*email {zexiangxu, jiy173, henrik, ravir }@eng.ucsd.edu, jbol@dtu.dk
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM.
SA '16 Technical Papers., December 05 - 08, 2016, , Macao
ISBN: 978-1-4503-4514-9/16/12
DOI: <http://dx.doi.org/10.1145/2980179.2982396>




Figure 1: By utilizing the field of view of a camera (A), many radiance observations can be acquired in parallel, to enable efficient BRDF measurement from a homogeneous flat material sample (we assume a near-field camera and distant light source). We demonstrate that as little as 2 view/light configurations enable accurate reconstruction for most materials. The two input images are shown in (B); our acquisition uses data from the circular center of the sample outlined, which may appear elliptical at oblique angles. Qualitatively, the first image captures the overall shape and intensity of the specular highlight and diffuse color. The second image captures grazing angles and Fresnel effects, also refining diffuse shading. From these inputs, we reconstruct a full measured BRDF, which can be visualized on a sphere or used for rendering (C); input material samples are shown in the insets and in results in Fig. 11.

However, there are several limitations of [Nielsen et al. 2015]. First, they use a gantry-based system, where each “measurement” is actually a full 2D image seen by the camera. However, this additional information is not used in their work, providing only a single observation. We seek to exploit the additional degrees of freedom by acquiring multiple BRDF measurements simultaneously, using a near-field camera, so each point on the sample corresponds to a slightly different viewing direction. In this paper, we demonstrate an improved minimal BRDF sampling method for near-field acquisition. Indeed, accuracy comparable to the 20 measurements in [Nielsen et al. 2015] is achievable with only two near-field images, and high-fidelity results are sometimes achieved with a single image, with field-of-view only about 25° , as shown in Figs. 1, 2.

A major technical challenge is finding the optimal light-view directions. The conventional condition number metric is not adequate, since it can increase dramatically (or even go to infinity) for a set of closely-related near-field measurements. While the measurements are no longer completely independent, they do provide additional information. Even for the gonimetric case in [Nielsen et al. 2015], we show that condition number does not fully model the error.

We therefore develop an entirely new framework to accurately estimate the error in BRDF acquisition from a set of samples, considering both deviation from the ideal (noise, BRDF not fitting MERL data), and accuracy of reconstruction based on where samples are located (Fig. 3). The condition number only approximates error from the first term (noise), while we usually want to minimize the latter term (reconstruction error). Our framework enables sig-

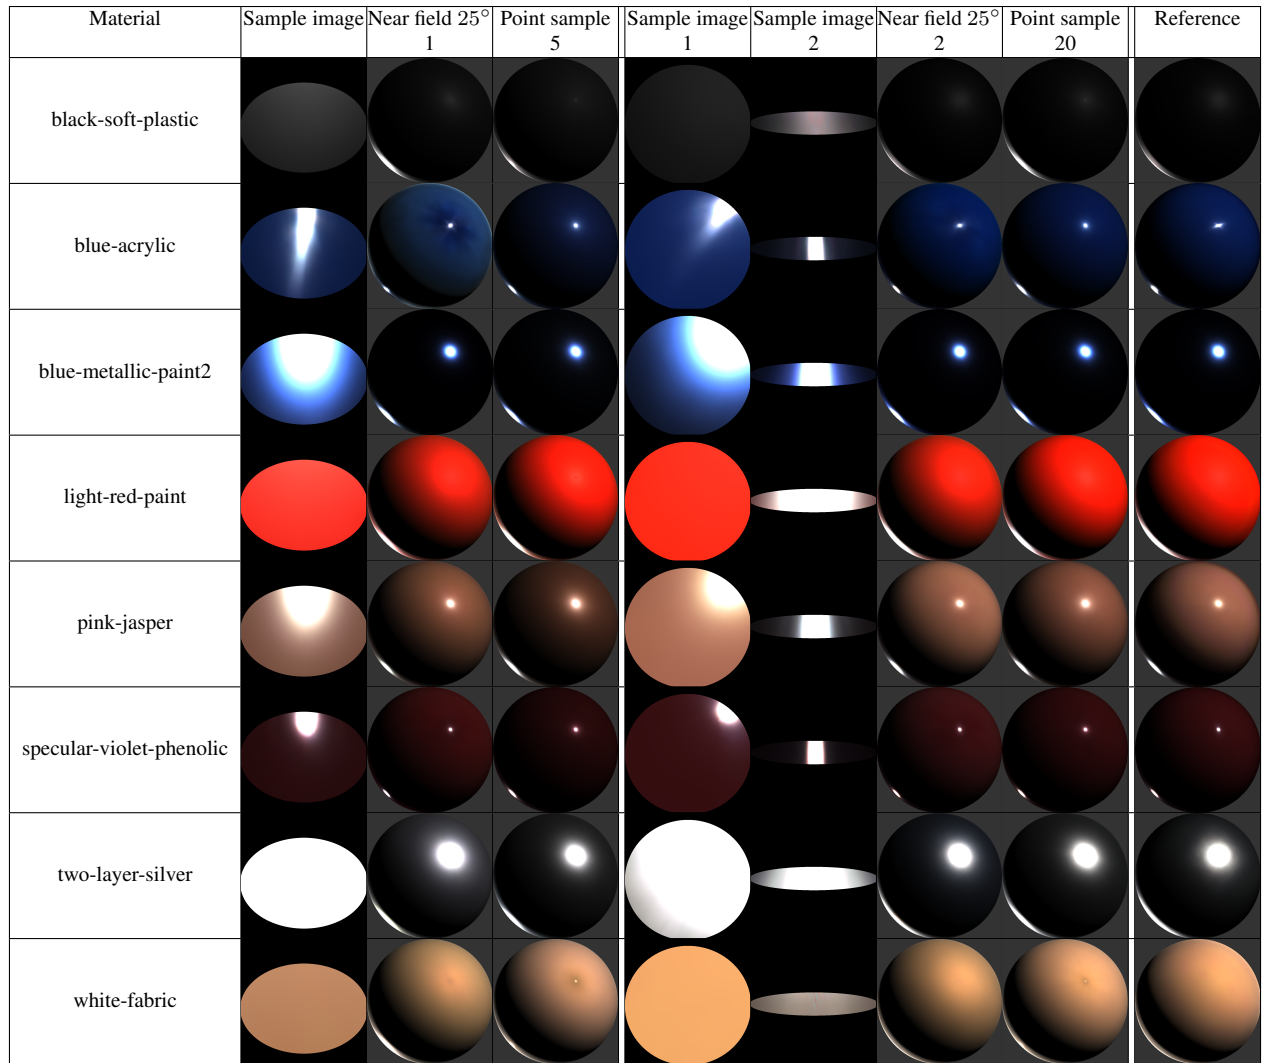


Figure 2: Simulations on MERL BRDF database for one and two-shot near-field BRDF measurement. BRDFs shown are not used at all in analysis/choosing the optimal directions. The sample subtends an angle of 25° when camera is at the zenith (we consider a circular region, sample images elliptical when viewed at an oblique angle; see Figs. 4, 9). After sample label, first 3 columns show single image near-field acquisition, with input, our method and previous work with 5 point samples. The next 4 columns show two image acquisition with inputs, our method, and previous work with 20 point samples. The reference image is shown rightmost. Single-shot measurement is comparable to 5 samples in [Nielsen et al. 2015]. Near-reference quality images, comparable to 20 samples in [Nielsen et al. 2015], are obtained with two-shot measurement. For very broad specularities like light red paint, we do not fully observe the highlights in the input images, and therefore slightly underestimate their width. The supplementary material has results for simulations on all of the MERL BRDF materials.

nificantly better reconstruction error for the multiple-measurement near-field setup (Fig. 5). We make the following contributions:

Novel Theoretical Analysis: We develop a new framework to predict the error in BRDF acquisition from sampling a particular set of directions (Sec. 4), that extends naturally to multiple measurements and allows for optimization (Sec. 5).

Optimal BRDF Sampling Directions: We provide optimal light-view directions for the near-field case (Sec. 6), showing how multiple simultaneous BRDF measurements can dramatically reduce the number of images (see Fig. 2 for simulations on MERL database). The new error framework also provides improved accuracy even for point-sampled measurement (Appendix B).

Practical BRDF Acquisition: We develop BRDF measurement using only two near-field images (two-shot), and we demonstrate results on several real examples (Sec. 7). We also briefly explore

extensions to a fixed camera setup with two light directions (Sec. 8), including a simple first step towards acquiring spatially-varying BRDFs by clustering materials.

2 Related Work

We focus on measured data-driven BRDFs. A few recent works have evaluated and developed improved parametric fits to the MERL BRDF database [Ngan et al. 2005; Low et al. 2012; Brady et al. 2014]. Earlier, Lensch et al. [2003] proposed adaptive sampling of a BRDF, based on the uncertainty of the parametric fit. Fuchs et al. [2007] also proposed an adaptive sampling approach, but first require a relatively dense grid of samples. Our approach actually requires less data (only one or two images), is more accurate, and requires only linear solution, rather than non-linear optimization. Our method is also conceptually related to other areas, such as the use of key points for animation [Meyer and Anderson 2007].

We briefly review related work in efficient BRDF acquisition below.

Basis and Environment Lighting: One approach to speed up BRDF acquisition is to use basis functions for complex illumination [Ghosh et al. 2007; Aittala et al. 2013; Tunwattanapong et al. 2013]. Our setup is simpler and requires fewer measurements, enabling direct BRDF measurement from a directional source, without any deconvolution. Other work has dealt with (uncontrolled) environment lighting [Romeiro et al. 2008; Romeiro and Zickler 2010], but requires non-linear regularization and priors.

Image-Based BRDF Acquisition: The MERL database was acquired by Matusik et al. using image-based BRDF measurement [Marschner et al. 2000] on spheres. Our approach is conceptually similar, making use of images rather than point measurements, but works with conventional flat surfaces (not all BRDFs are easily found as, or can be wrapped on, spheres). Earlier, Karner et al. [1996] fit anisotropic BRDFs to images of a flat sample, but required a parametric model. A variety of other optical setups acquire multiple samples simultaneously, including [Ward 1992; Dana and Wang 2004; Noll et al. 2013]. These methods usually require complex imaging setups, which are hard to calibrate.

One and Two Shot Approaches: Most recently, Aittala et al. [2015] proposed a two-shot spatially-varying BRDF capture setup, but it is aimed at reproducing the “texture” of material samples, rather than a complex measured BRDF. Earlier, [Hertzmann and Seitz 2003] used reference BRDFs to recover shape and spatially-varying reflectance. Ren et al. [2011] developed a pocket reflectometry method, comparing to reference tiles, using a hand-held light source and fixed camera. In contrast, we do not require a physical reference since we can leverage the MERL database, which also has a much broader set of reference materials.

Industry Material Standards: Beyond computer graphics, the materials industry has developed a number of standards for measuring and characterizing reflectance. [Hunter and Judd 1939] proposes a single measurement at 60° perfect reflection. The ASTM Standard D523 for measuring gloss adds near-normal and grazing angle measurements at 85° and 20° [Hunter 1987]. We extend this approach by considering near-field measurements. Our optimal sets of one and two measurements produce results that are close to the above observations, but are based on rigorously minimizing the expected error. Moreover, we can recover a full accurate data-driven BRDF, since we consider multiple measurements over the entire sample. Additional works in graphics include the five measurement directions in [Westlund and Meyer 2001], which are improved on by [Nielsen et al. 2015] and further refined by our method.

3 Background

In this section, we briefly discuss necessary background on using the MERL database [Matusik et al. 2003a] for BRDF measurement, following [Nielsen et al. 2015]. We conclude by providing intuition for why the condition number metric is not ideal, especially in the near-field lighting case, a result that may also impact other problems involving sparse sampling and reconstruction in graphics.

BRDF Database and Processing: The database consists of 100 materials. Each material is represented using $p = 1,458,000$ exhaustive measurements of the 3D isotropic BRDF volume in the $(\theta_h, \theta_d, \phi_d)$ parameterization [Rusinkiewicz 1998] (Fig. 4), with resolution $90 \times 90 \times 180$ degrees. Following [Nielsen et al. 2015], we treat each color channel separately, effectively obtaining $m = 300$ database BRDFs. We also apply their log-relative mapping (inspired by the logarithmic transform proposed in [McCool et al. 2001]). BRDF ρ is transformed to $\ln[(\rho w + \epsilon)/(\rho_{\text{ref}} w + \epsilon)]$, where $\epsilon = 0.001$ avoids division by zero, ρ_{ref} is the reference or

(per-observation) median BRDF, and the weight w is simply the maximum of the cosine of incident and outgoing angles. In this paper, we deal only with these log-mapped BRDFs. Inverse mapping is done at the end to obtain the final measured BRDF.

BRDF Principal Components: Let $X \in \mathbb{R}^{m \times p}$ be the full matrix of all MERL BRDF observations, where the rows are mapped BRDFs and the columns are a particular direction. We use the principal components Q , obtained by performing a singular-valued decomposition (SVD) after subtracting the mean BRDF,

$$X - \hat{\mu} = U \Sigma V^T \quad Q = V \Sigma, \quad (1)$$

where it is convenient to include the singular values in $Q \in \mathbb{R}^{p \times k}$. k is the number of principal components we consider (in our case, $k = m$, but one could use fewer components). The columns of Q are the scaled eigenvectors of the covariance, and correspond to a basis for the space of BRDFs. A particular BRDF x may be obtained as a linear combination of the basis,

$$x = Qc + \mu, \quad (2)$$

where $c \in \mathbb{R}^{k \times 1}$ is a vector of coefficients. $\mu \in \mathbb{R}^{p \times 1}$ is the mean BRDF, while $\hat{\mu} \in \mathbb{R}^{m \times p}$ is a matrix, repeating μ^T over m rows.

Solving for the Measured BRDF: In practice, we observe x at some sample observations, from which we seek to estimate c ,

$$\tilde{x} - \tilde{\mu} = \tilde{Q}c, \quad (3)$$

where the tildes indicate that we have a reduced set of observations at n samples, 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 reduced observations. It is also convenient to define $y = x - \mu$, with $\tilde{y} = \tilde{x} - \tilde{\mu}$ and $\tilde{Q}c = \tilde{y}$. Finally, let $S \in \mathbb{R}^{n \times p}$ be a selection matrix that is zero everywhere, except that $S_{ij} = 1$ in row i iff j is the direction corresponding to observation i . We can now define the reduced $\tilde{Q} = SQ$ and $\tilde{y} = Sy$, which will be useful for the error analysis in Sec. 4.

In [Nielsen et al. 2015], $n \ll p, m$, and typically $n \sim 20$. In our case, for near-field imaging, we have fewer image captures (typically only one or two), but we have several observations at each image, since we make use of the full 2D image. n can now be larger and, in some cases, could even be greater than m . However, the near-field samples are correlated, having the same light and similar view directions, so conceptually we still have a reduced matrix.

The above equation can be solved for the coefficients using Tikhonov regularization (I is the identity matrix. We set $\eta = 40$; we find results are not sensitive to this regularization parameter),

$$c = \arg\min |(\tilde{x} - \tilde{\mu}) - \tilde{Q}c|^2 + \eta|c|^2 = (\tilde{Q}^T \tilde{Q} + \eta I)^{-1} \tilde{Q}^T \tilde{y}. \quad (4)$$

A useful intuition is to consider a closed-form expression for the regularized inverse. Assuming the full SVD of $\tilde{Q} = A \Lambda B^T$,

$$\tilde{Q}_\eta^+ = (\tilde{Q}^T \tilde{Q} + \eta I)^{-1} \tilde{Q}^T = B \Lambda_\eta^+ A^T, \quad (5)$$

where Λ_η^+ is a diagonal matrix with the same shape as Λ , but with a modified set of singular values: $\sigma \rightarrow \sigma/(\sigma^2 + \eta)$. Note that the pseudo-inverse \tilde{Q}^+ and Λ^+ are obtained by setting $\eta = 0$, in which case $\tilde{Q}^+ = B \Lambda^+ A^T$ as expected. In essence, the regularization term creates a η -modified pseudo-inverse where the inversion of small singular values does not blow up.

Optimizing Sampling Directions: In [Matusik et al. 2003b; Nielsen et al. 2015], the optimal sampling directions (the rows of \tilde{Q}

chosen, or equivalently the selection matrix S with $\tilde{Q} = SQ$ are found by optimizing (minimizing) the condition number,

$$\kappa(\tilde{Q}) = \frac{\sigma_{\max}(\tilde{Q})}{\sigma_{\min}(\tilde{Q})}, \quad (6)$$

where σ_{\max} and σ_{\min} are the maximum and minimum singular values of \tilde{Q} . The condition number is a standard numerical tool, and reducing it minimizes the sensitivity to noise and related errors.

Formally, consider a matrix equation such as equation 3, with $\tilde{Q}c = \tilde{y}$ (with $\tilde{y} = \tilde{x} - \tilde{\mu}$ as usual). The condition number is the worst case (upper bound) estimate of the ratio of fractional error δc in output to fractional error/noise $\delta \tilde{y}$ in input,

$$\frac{|\delta c|/|c|}{|\delta \tilde{y}|/|\tilde{y}|} \leq \kappa(\tilde{Q}). \quad (7)$$

3.1 Limitations of Using the Condition Number

The condition number gives good results for point-sampled BRDF measurement [Matusik et al. 2003b; Nielsen et al. 2015]. It can be considered a measure of correlation between samples, and minimizing it chooses sampling directions that discriminate between distinct BRDFs. However, we found that it did not easily extend to near-field measurements, where a large number of related observations are made (Fig. 5). The observation matrix \tilde{Q} is now often rank deficient or nearly so, and close-by observations can drive the condition number very large or even to infinity, reducing its ability to discriminate and choose optimal directions. This leads to the paradox where fewer observations are preferred. In the next section, we formally derive the expected error, considering both reconstruction error and noise. For near-field BRDF measurement, we achieve a dramatic improvement; one to two near-field images is adequate.

There are also many technical limitations of condition number. First, there are two terms related to error: noise or other imperfections (deviations from MERL data); and reconstruction error caused by having too few samples (even in the presence of zero noise or deviation). $\kappa(\tilde{Q})$ only bounds the first term (noise/deviations), but the major component of the error is actually reconstruction error from having fewer observations than principal components. Second, condition number considers fractional error, assuming the error is proportional to the signal. However, the accuracy of measurements from real cameras is determined by a number of factors (shot noise, read noise, dark current), which are constant or proportional to the square root of intensity, and not the intensity itself. Indeed, well lit pixels have less relative noise, and in this paper we more accurately model the noise as a constant magnitude, independent of the signal. Third, κ only provides a worst-case bound, while we are often interested in the average error, say over all of the materials in the MERL BRDF database. Hence, our optimal sampling directions improve somewhat on [Nielsen et al. 2015] even for point-sampling.

4 Sampling Error Analysis

In this section, we conduct a novel analysis of the BRDF reconstruction error from a sparse set of samples. This error can be minimized to find the optimal set of sampling directions, for both conventional point-wise BRDF acquisition, and near-field image-based measurement. For completeness, we consider three sources of error: deviation from the MERL database, sparse sampling, and noise in measurement. In practice, deviation error from the MERL database is not easy to predict, nor is the real noise level easy to evaluate. Therefore, our practical algorithm will focus on minimizing the reconstruction error from sparse sampling, which is the main factor in choosing suitable directions for BRDF acquisition.

ACM Trans. Graph., Vol. 35, No. 6, Article 188, Publication Date: November 2016

Deviation from BRDF Model: We assume the BRDF being measured lies in the subspace spanned by the MERL database (and encapsulated in Q). If this is not the case, we can only find the best projection of the MERL BRDF data. This error is present even when we have all observations. Using pseudo-inverse Q^+ of Q ,

$$c = Q^+(x - \mu) = Q^+y = \Sigma^{-1}V^Ty, \quad (8)$$

where we expand $Q = V\Sigma$. The resulting deviation error is,

$$E_{\text{deviation}} = |Qc - y| = |(VV^T - I)y|. \quad (9)$$

Note that $V \in \mathbb{R}^{p \times k}$ is an orthogonal matrix with $V^TV = I$, but since $k < p$, $VV^T \in \mathbb{R}^{p \times p}$ is not the identity. However, if y is in the MERL BRDF database, it is given as a column of $Y^T = (X - \hat{\mu})^T = V\Sigma U^T$. Using the SVD decomposition, it is easy to see that $(VV^T - I)V\Sigma U^T = 0$, since $V^TV = I$.

Therefore, $E_{\text{deviation}} = 0$ if the material is in the subspace Q spanned by the MERL database, but will be nonzero if it lies outside this subspace. This is an intrinsic property of the material, and independent of the sampling directions chosen.

Projection to Sampling Directions: Choosing a sparse set of n sampling directions corresponds mathematically to choosing a particular selection matrix $S \in \mathbb{R}^{n \times p}$. Noting that $\tilde{Q} = SQ$ and $\tilde{y} = Sy$ by definition, so that $SQc = Sy$, we have

$$\bar{c} = (SQ)_{\eta}^+(Sy), \quad (10)$$

where in the last line we consider the regularized inverse of SQ , as per equation 5, and Sy are the observations we actually make with a camera or a gonioreflectometer. We use the bar on top of c to denote the recovered coefficients, with error

$$c - \bar{c} = \left(Q^+ - (SQ)_{\eta}^+ S\right)y. \quad (11)$$

Finally, the reconstruction error is given by

$$E_{\text{recon}} = \left|Q \left(Q^+ - (SQ)_{\eta}^+ S\right)y\right|. \quad (12)$$

This is the critical error we need to minimize, by choosing sampling directions (and hence S) optimally. It provides the error in reconstruction by measuring only a sparse set of samples, and applies equally whether those are point samples or multiple simultaneous image-based measurements. Note that this error exists even for noise-free measurements, coming purely from reconstruction error when using a sparse set of samples. (By using log-mapped BRDFs, we also limit the ability of intense specularities to unduly influence reconstruction error.) The condition number does not consider this term directly, but only sensitivity to noise. Nevertheless, we show in appendix A that minimizing condition number does adjust SQ to reduce (but not minimize) E_{recon} .

Figure 3 shows both deviation and reconstruction errors for the BRDFs in Fig. 2 (using a different set of 90 MERL materials as our data/training set). As expected, reconstruction error E_{recon} dominates in all cases. Blue acrylic has high $E_{\text{deviation}}$ since the star-shaped highlight deviates significantly from the database.

Noise Error: If we do have noisy data, the image observations y will be corrupted, and we will measure $\tilde{y} = y + \Delta$, where Δ is the noise or error at each pixel. The resulting error in the coefficients is given from equation 10 by $(SQ)_{\eta}^+(S\Delta)$. Therefore,

$$E_{\text{noise}} = \left|Q (SQ)_{\eta}^+ S \Delta\right|. \quad (13)$$




Figure 3: Comparison of E_{recon} and $E_{\text{deviation}}$ for all materials in Fig. 2. Reconstruction error E_{recon} is dominant for most BRDFs.

Conceptually, the condition number seeks to minimize this term. However, condition number provides only a worst-case bound, assuming the noise is proportional to the signal, which is not a correct assumption for cameras, where noise levels are relatively independent of image intensity. Moreover, our main focus is on reconstruction error from sampling (equation 12) rather than noise; one typically acquires high-dynamic range images from high-end cameras where noise is not the most significant challenge. Note that the condition number analysis also does not consider the full process, including the η -regularization. Finally, our focus is on near-field capture where we have several, but closely-related observations. This can lead to a very large condition number, while in fact the additional observations help in reducing the error.

The total error is written simply as (the less than sign comes from the triangle equality, since each error term considers the norm),

$$E_{\text{total}} \leq E_{\text{deviation}} + E_{\text{recon}} + E_{\text{noise}}. \quad (14)$$

Final Error Metric: For simplicity, we do not explicitly consider the deviation error, but just include it as part of the noise/error Δ . A final issue is choosing y in equation 12 and Δ in equation 13, since these quantities depend on the measurements, camera noise, and are not known a-priori. For y , we minimize over all of the m materials in the MERL BRDF, essentially finding the sampling directions that best reconstruct the MERL materials. Define $y_i = x_i - \mu$ where $x_i \in \mathbb{R}^{p \times 1}$ is a vector corresponding to observations of BRDF i in the MERL database. For the noise, we assume a constant user-defined parameter β , corresponding to the noise/error level, $\beta = |\Delta|$, and use a noise vector $O \in \mathbb{R}^{p \times 1}$, where each element is simply 1. This can be seen as the expected magnitude of gaussian-distributed noise, where β controls the magnitude. Putting this together, our final expected error is,

$$E(S) = \left(\frac{1}{m} \sum_{i=1}^m \left| Q(Q^+ - (SQ)_\eta^+ S) y_i \right| \right) + \beta \left| Q(SQ)_\eta^+ SO \right|. \quad (15)$$

Note that we add errors from reconstruction and noise. Each term on the right-hand side is a $p \times 1$ vector, and we take its norm. We also make explicit the dependence of E on selection matrix S .

In this paper, we focus mainly on minimizing the reconstruction error E_{recon} by choosing sampling directions. Therefore, we typically take $\beta = 0$, but we also demonstrate nonzero noise β in supplementary material. Finally, we emphasize that we have so far only defined error; the next section discusses how to choose the sampling directions, corresponding to the selection matrix S , to minimize this expected error. In essence, we seek $S = \text{argmin}_S E(S)$.

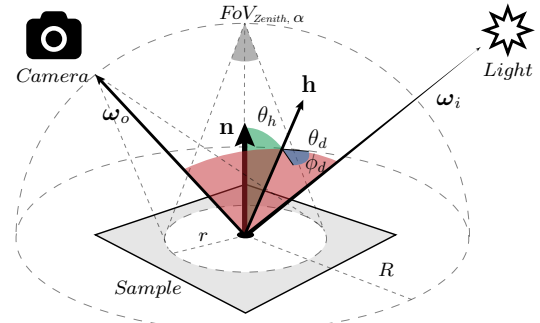


Figure 4: Schematic of near-field reflectance acquisition. The half-diff angles $\theta_h, \theta_d, \phi_d$ are with respect to the center of the sample.

5 Optimal Sampling Directions

We first describe selection of the optimal sampling directions for measuring individual light-view pairs in a BRDF, as in [Nielsen et al. 2015]. We refer to this as point sampling, to distinguish from the near-field image-based BRDF measurement of our method.

5.1 Point-Sampled BRDF Directions

We consider the whole space of valid directions $\mathcal{D} = \{\theta_h, \theta_d, \phi_d\}$ in the MERL database. Our goal is to find the optimal subset \mathcal{D}^n with n directions, and form the corresponding n -row selection matrix S^n . For point samples, each row in S^n is simply one direction in \mathcal{D}^n . In other words, $S_{ij}^n = 1$ iff $\mathcal{D}_i^n = j$. The optimal \mathcal{D}_n (and S_n) must be chosen to achieve the minimal error in equation 15.

We solve the optimization based on a numerical gradient descent framework analogous to [Nielsen et al. 2015], which is shown in that work to be more efficient and higher-quality than the greedy method of [Matusik et al. 2003b]. (Standard numerical optimizers do not work well, given the discrete BRDF space \mathcal{D} and integer steps needed, as well as invalid BRDF regions.) However, we replace the condition number with the accurate error in equation 15, and extend the optimization framework to near-field measurements in Sec. 5.2. We start with an empty set \mathcal{D}^0 with no directions in it. Then we iteratively extend \mathcal{D}^n to \mathcal{D}^{n+1} as follows:

1. Randomly pick t candidate directions from $\mathcal{D} - \mathcal{D}^n$ (Typically we use $t = 500$). For each candidate direction d , we form a selection matrix S_d of $\mathcal{D}^n \cup d$, and evaluate the expected error $E(S_d)$ from equation 15. An initial \mathcal{D}^{n+1} is created with that $\mathcal{D}^n \cup d$ which has a minimal $E(S_d)$.
2. Randomly choose one of the $n + 1$ directions in \mathcal{D}^{n+1} , which we denote as $(\theta_h, \theta_d, \phi_d)$. Numerically estimate the gradient of the error metric $\nabla E(S^{n+1}) = (\frac{\delta E(S^{n+1})}{\delta \theta_h}, \frac{\delta E(S^{n+1})}{\delta \theta_d}, \frac{\delta E(S^{n+1})}{\delta \phi_d})$. Move the chosen direction along ∇E with one step-length (initial step-length is 3 cells) if the destination is a valid location in \mathcal{D} . Repeat until convergence (finding a new direction each time).
3. Reduce step-length and repeat step 2 until convergence with step of 1 cell. Then the final \mathcal{D}^{n+1} and S^{n+1} are formed.

5.2 Near-Field BRDF Directions

We now take advantage of sampling directions for all pixels in an image, instead of only the center of the image for point sampling. In general, the optimal directions depend on the camera's projection matrix and the size of the planar sample. To develop a general framework, we assume the image of interest is a circle on the plane

with radius r . We assume the camera moves on a hemisphere a distance R from the center of the circle, and is always pointed towards the center (i.e., the center pixel corresponds to the center of the sample). We also assume the image always sees the full sample (circle of interest). The key variable is the ratio $\nu = r/R$, which determines the linear field of view when the camera is at the zenith. The angular field of view $\alpha = 2 \tan^{-1} \nu$, which is the angle we use to denote our near-field setup. A schematic of the setup is shown in Figs. 1 and 4. Note that the optimization framework is general, and can also apply to many other configurations. We discuss one-camera multiple light and one-light multiple view cases in Sec. 8.

The goal is still to find an optimal subset of camera directions \mathcal{D}^n . In this case, each direction represents the direction to the camera with respect to the center pixel. However, the corresponding selection matrix is no longer a n -row matrix. We replace S^n with \tilde{S}^n in near field acquisition. One direction in \mathcal{D}^n forms a set of rows in \tilde{S}^n , each of which corresponds to one pixel sample in an image. In general, there will be many more rows than for point-sampling, but many of the directions will be very closely related. Our error metric addresses this directly, and equation 15 still accurately predicts reconstruction error. We can now directly use \tilde{S}^n instead of S^n , and iteratively add directions from \mathcal{D}^0 to \mathcal{D}^n as before. To validate the convergence of our method, we repeated the optimization 50 times with different random conditions, and fields of view. The results all converge well. The supplementary material shows convergence results for $n = 2$ near-field sampling with 25° field of view.

6 Evaluation with Simulations for Near-Field

We now evaluate the minimization of our error metric for near-field image-based BRDF measurement, using simulations with the MERL BRDF database. As shown in Fig. 5, our new error metric has significant advantages over using the condition number for the near-field case. (Visual results on rendered spheres are consistent with these numerical errors; results from minimizing condition number are often even worse than point-sampled measurements, far off from ground truth). Moreover, as seen in Fig. 6, our new image-based method is much more efficient than point-sampling; both methods capture similar images of a flat sample, but we make use of the full 2D image. Section 8 applies the framework to other configurations like fixed camera with multiple lights, or vice versa.

The setup is shown in Fig. 4. For simplicity, we assume a distant light source, and a near-field camera. In a single image, we capture a 2D slice of the BRDF (we consider a circular patch). Since we are assuming a flat sample with distant lighting, the illumination direction is the same everywhere, but the viewing direction varies at each pixel, enabling us to capture multiple observations simultaneously. It is clearly better to have a wider field of view to capture greater view variation, but this may require large samples and close cameras. In fact, we show that a relatively narrow field of view of about 25° suffices for two-shot BRDF acquisition.

We minimize equation 15, choosing the optimal light-view directions, as explained in Sec. 5.2. To evaluate the reconstruction error on the MERL BRDF, we use a slightly different set of directions using 90 training BRDF samples, testing on the 10 other materials not used at all in computing optimal directions. (We use the same training/test materials as [Nielsen et al. 2015] to enable direct comparisons to their approach.)

Figure 5 compares our average normalized reconstruction RMS errors for the unknown materials for several fields of view, as a function of the number of images, and to optimizing condition number alone. As shown in appendix B, condition number is actually a reasonable heuristic for point-sampled BRDF measurement [Nielsen et al. 2015], although our error metric performs somewhat better

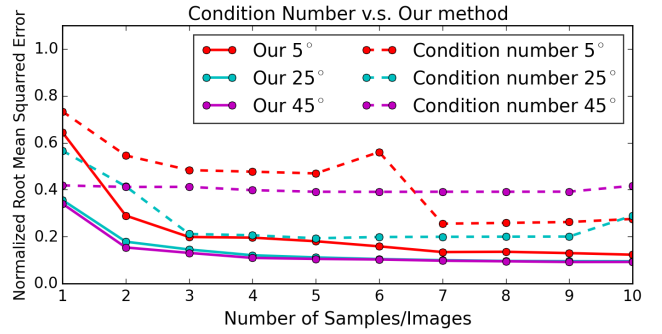


Figure 5: Comparison of errors on unknown samples from our method, and from minimizing condition number, for near-field reflectance acquisition with different fields of view. It is clear that we produce better results for near-field reflectance acquisition.




Figure 6: Average RMS error over unknown samples for near-field reflectance acquisition. We plot the results for a number of different field of view angles. These results clearly show the benefits of our method, often requiring an order of magnitude fewer samples than point-sampled BRDF measurement.

even in that case. However, it breaks down for near-field acquisition as seen in Fig. 5. With several correlated view directions, condition number becomes very large, losing the ability to discriminate between different sets of light-view directions. In some cases, it oscillates or does not decrease with increasing samples, while our method always performs well. *The new error metric is essential for determining optimal light-view directions in the near-field case.*

In Fig. 6, we compare RMS errors for several fields of view, and to point-sampled BRDF measurement (the top red curve is from [Nielsen et al. 2015] while the improved orange curve is using our error metric for the point-sampled case). We see that near-field reflectance acquisition requires almost an order of magnitude fewer images than point-sampled BRDF measurement. Also note that near-field acquisition converges quickly with increasing field of view; while larger fields of view help, 25° is already nearly best (supplementary shows similar curves even for extreme 85° and 175° fields of view). In fact, errors are comparable to standard spherical image-based BRDF measurement [Marschner et al. 2000] (with optimal directions chosen by our error metric; see Fig. 19 in appendix B). However, our approach applies more generally, to flat samples that cannot be obtained or wrapped on a sphere.

Figure 7 shows how errors decrease as more training materials are added to the database (in random order), showing a steady decrease




Figure 7: Average RMS error versus number of materials in database for 2 shot near-field sampling with 25° field of view.

n	$\theta_h [^\circ]$	$\theta_d [^\circ]$	$\phi_d [^\circ]$
1	3	42	1
2	0	60	0
	34	36	32

(a) 15°

n	$\theta_h [^\circ]$	$\theta_d [^\circ]$	$\phi_d [^\circ]$
1	3	52	0
2	4	40	0
	8	81	82

(c) 35°

n	$\theta_h [^\circ]$	$\theta_d [^\circ]$	$\phi_d [^\circ]$
1	3	50	9
2	6	23	41
	16	79	87

(b) 25°

n	$\theta_h [^\circ]$	$\theta_d [^\circ]$	$\phi_d [^\circ]$
1	3	60	4
2	0	81	0
	5	40	4

(d) 45°

Figure 8: Tabulation of one and two near-field acquisition directions for fields of view ranging from 15° to 45°. Note that directions correspond closely to mirror reflection, imaging the highlight shape, and more grazing angles for Fresnel and other effects.





Figure 9: Illustration of one- and two-shot camera/light configurations for a linear field of view of 25°.

with more example BRDFs. Note that the MERL subspace depends on the specific BRDFs used, and the error curve can therefore increase slightly with the addition of a new material.

Note that with two-shot imaging, we can obtain essentially the same accuracy as 20 samples for point-sampled BRDF measurement, and even single-shot near-field acquisition achieves similar accuracy as about 5 samples for point-sampled BRDF measurement. These comparisons, simulated on the MERL data, are shown in Fig. 2. (We omit green fabric and silver metallic paint, whose results are very similar to black soft plastic and two layer silver respectively). We see that two-shot acquisition is adequate in nearly all cases. In some examples like pink jasper, specular violet-phenolic and white fabric, a single near-field image is comparable to 20 point-sampled images. In a few cases, like blue acrylic, two near-field images are required to achieve reasonable results; severe ringing is present in reconstruction from a single image. For very broad specularities like light red paint, we do not observe the full extent of the highlight in any single image, and slightly underestimate highlight width.



(a) Setup A

(b) Setup B

Figure 10: Photograph of our acquisition setup A and B. In setup A, a 6-axis industrial robot precisely positions the camera, and an illumination-arc positioned at $\phi = 0^\circ$ illuminates the sample with halogen lights in 7.5° θ_d intervals. In setup B, a high-angular-resolution spherical gantry positions the light. A DSLR camera is positioned by utilizing the two arms of the gantry and a mirror.

Finally, in Fig. 8, we tabulate our optimal 1,2 directions for near-field angles of 15°, 25°, 35° and 45°. Note that these directions are with respect to the center of the sample; the local view direction will vary at each pixel. For one image, we capture a slightly off-specular direction ($\theta_h = 3^\circ$) at an angle of incidence of about 50°. Similarly, for two images, the first direction for 15° field of view is an exact specular reflection at 60°, although this varies somewhat with field of view. This is as expected, imaging the details of the specular highlight, around the center of the sample, and also accords well with measurements previously used in the appearance industry [Hunter and Judd 1939]. For most materials, this measurement also captures the overall diffuse color. The second direction usually varies somewhat from the specular (more for small fields of view, less for larger fields of view where diffuse and specular reflection are often both available in the same image). Intuitively, the second direction measures Fresnel effects at grazing angles (large θ_d for fields of view 25° and higher). It can also help refine the diffuse shading, especially for materials with broad specular lobes that cover all of the first image. Figure 9 illustrates the one- and two-shot light-view pairs for field of view 25°.

7 Results: Near-Field BRDF Measurements

In this section, we apply the reconstruction method, and optimal sampling directions (Sec. 6, Fig. 8), to image-based near-field BRDF acquisitions of several real samples captured at two different laboratories located in different continents (UCSD in USA and DTU in Denmark). The two laboratory setups deviate slightly and will be described next. We used both approaches to demonstrate the robustness of our method with a variety of simple capture scenarios, which do not require exact configuration or precise alignment between views. We used a portion of the input sample with field of view of 25°, since that achieves near-optimal results (Fig. 6).

Setup A (DTU): In this setup, we utilize a high-precision robot-vision system to precisely position the camera relative to a material sample (Fig. 10 left). The angular error of this positioning is less than 1 degree. The camera used is a Point Grey Grasshopper 3, industrial CCD camera, mounted with a Kowa LM16SC 16mm lens. The light-source consists of arc holding halogen light-bulbs, evenly distributed from 0° to 90° in 7.5° steps.

Setup B (UCSD): We use a (distant) Dolan-Jenner DC 950 light source mounted on one arm of a spherical gantry; the gantry con-

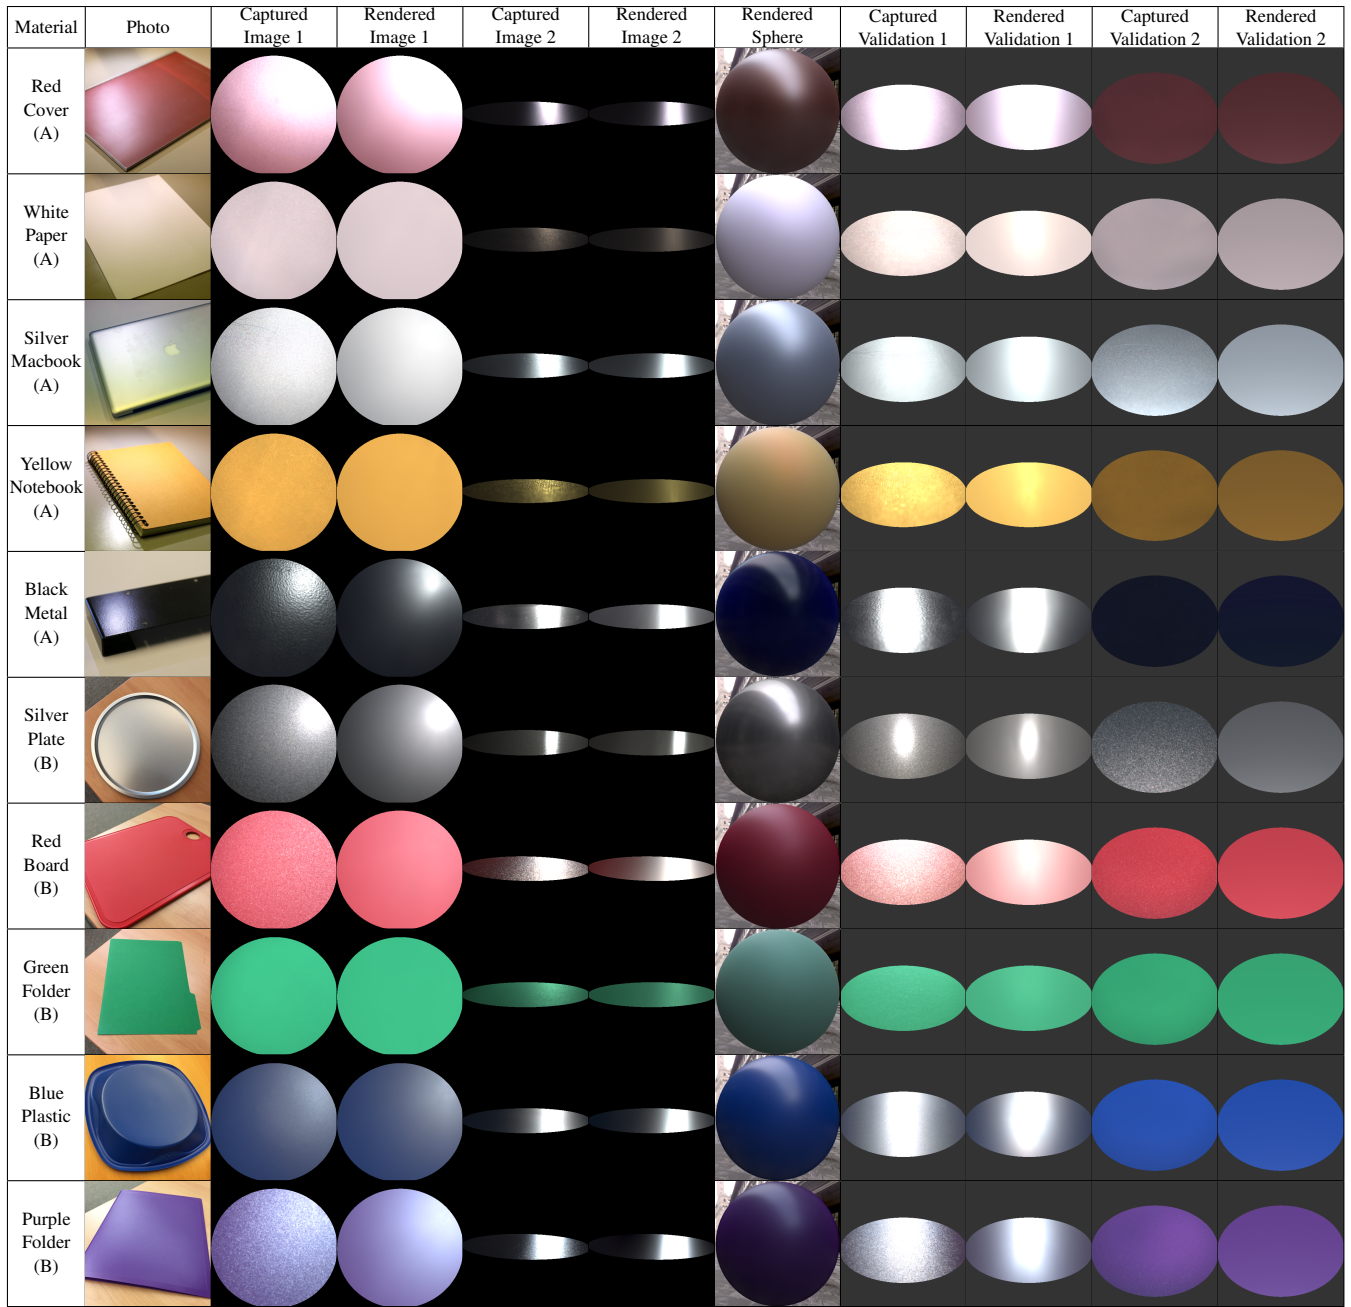


Figure 11: Results for near-field BRDF acquisition on real materials. Materials labeled with (A) and (B) are captured by setups A and B. The material is shown first, followed by comparisons of input photographs and renderings with the measured BRDFs. These include both original views, as well as two new lightings and views not used as input. All renderings use the BRDF reconstructed from the two captured images. Good accuracy is obtained for all materials. We also visualize the full BRDF by rendering a sphere lit by an environment map. Note that we use optimal directions in Fig. 8 from the full MERL data, which differ slightly from those using 90 materials in Fig. 2.

trols two high-precision arms having an angular resolution of 0.1° . However, the gantry's viewing arm/camera is too far away to obtain near-field images directly. We instead manually position a Canon EOS 5D Mark III camera mounted on a tripod close to the input sample (Fig. 10 right). In order to correctly position the camera, we place the gantry's two arms at the mirrored direction of the desired viewing location, and adjust the camera until it points towards both arms' center through a mirror. The final position of the camera is obtained through camera-calibration using a checkerboard.

The setups presented above have different limitations, in that the

light-source confines setup A to a limited set of view/illumination configurations, whereas the manual positioning of the camera limits the precision of setup B. In both cases, we find the configuration that best matches the optimum directions in Figs. 8, 9. We thus demonstrate that our method is robust towards small variations in view/light configurations, while still obtaining very good results.

For acquisition, we capture multiple exposures to produce high-dynamic range images; each exposure is averaged over multiple images to reduce noise. The resulting values are then log-mapped, since our framework works with log-mapped BRDFs. Light inten-

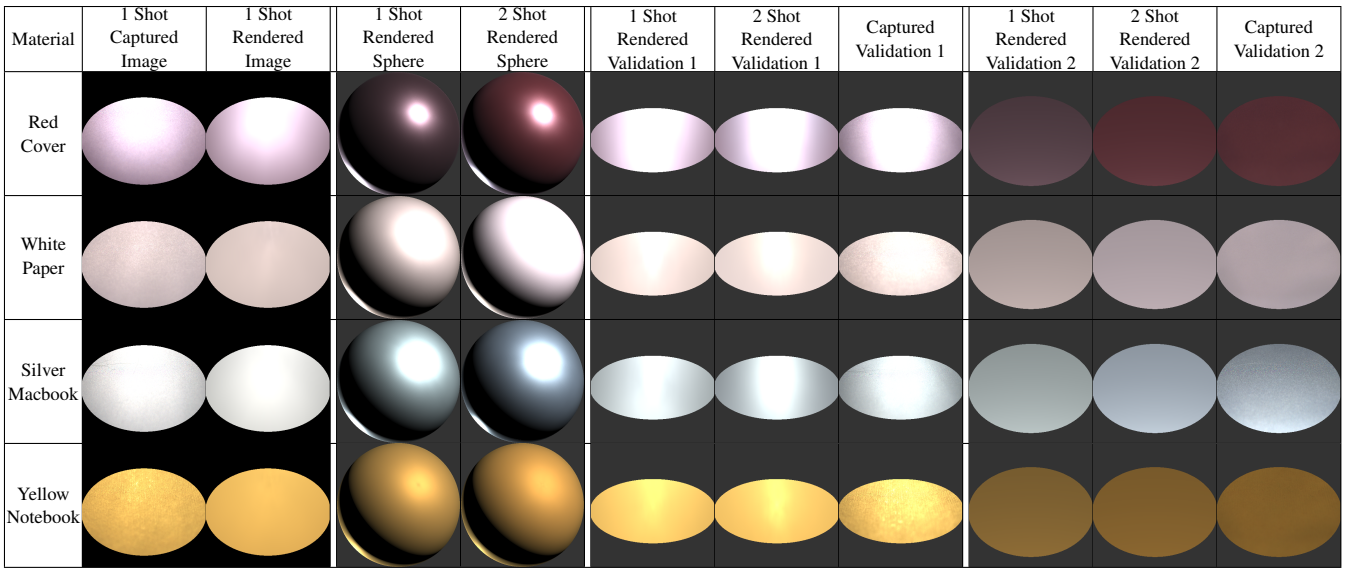


Figure 12: Results for single image near-field BRDF acquisition on real materials, comparing one and two-shot reconstructions.

sity is calibrated using a Spectralon material sample. By proper modeling of the intrinsic and extrinsic scene-parameters, all observed pixels (BRDF values) can be combined to form \tilde{x} or $\tilde{y} = \tilde{x} - \mu$, using their corresponding view/light coordinates. With this, BRDF coefficients c can be solved from equation 4, and the full BRDF x recovered using equation 2, with a final step involving undoing the logarithmic mapping. We emphasize that all results in this section were obtained from only two image configurations of a standard flat planar sample, and in some cases only a single image.

Figure 11 shows several real samples, whose BRDFs we measured using two-shot near-field acquisition (analogous to Fig. 2 for MERL simulations). The first 5 rows are captured using setup A, and the last 5 rows with setup B. The colors differ slightly between the inset photograph and the comparisons, because of the off-white color of the actual light source. We compare the captured image (considering only the circular region of interest, per Fig. 4) with the rendering for both light-view input configurations. We also show two additional light-view configurations for validation, which were not used at all as an input. The validation configurations are chosen with $(\theta_h, \theta_d, \phi_d) = (0^\circ, 60^\circ, 0^\circ); (22.5^\circ, 22.5^\circ, 0^\circ)$ to verify both specular and diffuse appearance. For visualization, we also show a rendered sphere with the corresponding BRDF lit by an environment map. The accompanying video shows the red cover, white paper and silver macbook under changing viewing directions with two illumination directions, comparing real and rendered results, including fading out the illumination to observe highlights without saturation. It can be seen that the real and rendered images match well. Even when the actual material has some noise or a slightly bumpy surface, we recover a smooth BRDF that is an accurate representation, for both diffuse and more glossy materials.

Finally, Fig. 12 shows what can be achieved with a single-shot capture (using setup A). In many cases, one input image is adequate to achieve reasonable results. However, the second input image does help refine the specular reflection somewhat, when comparing the rendered spheres. For example, white paper and yellow notebook are largely diffuse in the first image, and accurate specular and Fresnel information is only achieved at grazing angles in the second view. Moreover, in some cases, the diffuse color and shading can be somewhat refined by using both images.

We briefly discuss some limitations. As with all reconstruction methods based on the MERL data, we are ultimately limited by

the subspace spanned by that data. Our simulations and experiments indicate excellent agreement with reference measurements, but there is unavoidable error when the material deviates from the MERL subspace. Moreover, for small field of view and BRDFs with broad highlights, our measurements may not capture the full range of the highlight in a single image, leading to under-estimating its width in reconstruction (light red paint in Fig. 2). For very dark materials, the noise can be over-fit, causing a blue tint for the black metal in the fifth row of Fig. 11. Finally, we do not account for surface imperfections or normal maps, which also contribute to the “noise” above. Nevertheless, as seen in Fig. 11, we produce accurate smooth BRDFs consistent with the input data.

8 Extension: Fixed Camera Setup

Our error analysis framework and optimization method for sampling directions is general, and could be applied in future to many different configurations. In this section, we consider a one-camera multiple-light case, where we use a single near-field viewing direction, while enabling multiple lighting directions. As before, we show that good results are achieved with two-shot acquisition with two lights. Note that we optimize for lights, and the single view direction, but do so while constraining the camera view to be the same (fixed) for all lighting directions. Using a fixed camera view may also enable simpler acquisition hardware in future. We also briefly discuss the symmetric case of a single fixed light direction, with multiple views. We consider a field of view of 25° .

Figure 13 shows the error of fixed camera, multiple light, and fixed light, multiple view, as a function of the number of images, also comparing to our point-sampling and near-field results. As before, our error analysis framework is essential for finding optimal directions, and condition number does not yield meaningful results.

The errors for one or two images are significantly lower than for point-sampling, and only somewhat more than the unconstrained near-field case considered previously (note that one shot acquisition is the same for fixed or free camera setup). However, the lack of flexibility when fixing light or view, leads to a slower decrease in error for more images. Figure 14(a) indicates the optimal two-image configuration for fixed camera and changing light. The camera is at a 64° angle to the surface, with light sources positioned to enable observation of both diffuse (light close to zenith) and specu-




Figure 13: Average RMS error over unknown samples for fixed camera/multiple lights, and fixed lighting direction/multiple views.

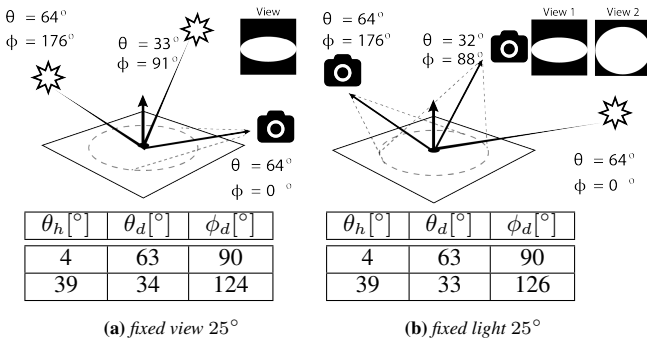


Figure 14: Optimal configurations and angles for single view, multiple light and single light, multiple view cases, analogous to Fig. 9. Angles are shown as standard in-out and Rusinkiewicz coordinates.

lar reflectance (light close to mirror direction). Having the camera at an angle to the surface enables capture of some Fresnel information, but a fixed camera setup will make it harder to fully reproduce grazing angles. We also show the analogous configuration for fixed light with multiple views in Fig. 14(b). Since fixed light/fixed view configurations are symmetric with similar error, we focus on the fixed camera setup, with only a single viewing direction, and therefore simpler calibration and alignment.

Figure 15 shows some synthetic MERL materials reconstructed with fixed view and two images. Note that the first input image is mostly specular while the second is mostly diffuse (dark for specular materials like metallic paint). The results are generally good in most cases. However, some ringing can be observed on the mostly diffuse white fabric. This corresponds to the higher error in Fig. 13, compared to the near-field case where both light and camera can move. Figure 16 shows comparable results for two real materials captured with setup B. We also show a validation view (specular with light/camera at 45°) not used as input. Good results are achieved with the two-shot fixed camera setup, although there is minor variance in the shape of the specular highlight.

Simple Extension to Spatially-Varying BRDFs: So far, we have not considered spatial variation. We take a first step with a simple extension for specific objects, which have two or more materials that have good coverage over the field of view (such as stars spread out on a background). The fixed-camera setup is ideal for this purpose, since no alignment/calibration between different near-field views is required. Note that this is an initial effort, and further work is required to extend the method to general SVBRDFs.

If we can cluster which pixels correspond to which material, we

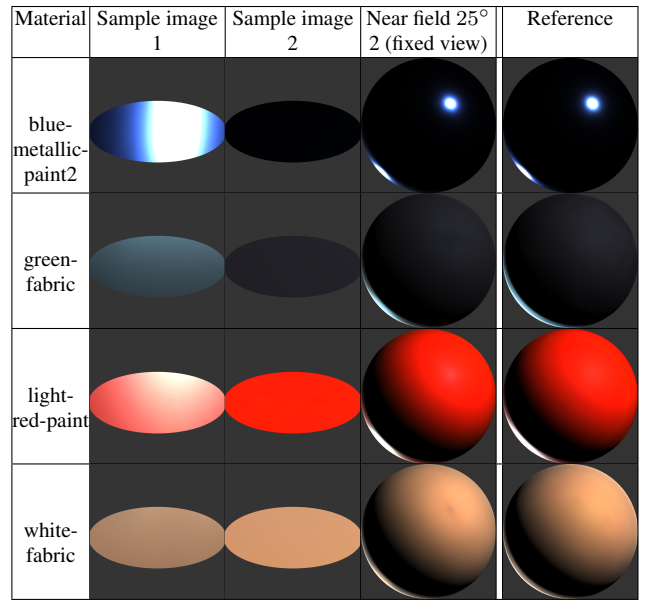


Figure 15: MERL BRDFs reconstructed with fixed view and two images. This configuration works well for most materials.

can separately estimate the BRDFs of the materials, using only the subset of pixels for that BRDF. The key requirement is *coverage* over the field of view, to enable one to see the full range of viewing angles. Using only a subset of pixels does not significantly increase error, especially since a 2D image already contains thousands to millions of observations. In practice, we cluster based on color observed in the second (diffuse) captured image. BRDFs are then estimated separately for each cluster. Figure 17 shows results for two greeting cards with spatial variation, acquired using setup B. In this example, we consider the full field of view, rather than only a circular region. As seen in Fig. 17, we cluster into three materials, and recover full BRDFs for all three materials. The rendered images are close to the captured, with the expected smoothing of surface roughness. (Microstructure and normal variations in the real object cause glints and rough specularities, which increase the apparent size of the highlight for the real object). The validation view, not used as input, also matches well.

9 Conclusions and Future Work

We have developed a method for acquisition of a full measured isotropic 3D BRDF from only two perspective images of a flat sample, lit with a directional light source. This is at least an order of magnitude reduction in effort over previous comparable techniques to measure a full BRDF, and requires only a standard flat homogeneous material sample. Our method is based on using the full 2D image information from a near-field view, and finds the best lighting and viewing directions by minimizing an estimate of the reconstruction error. We provide tables of these directions for different fields of view of the sample, which can directly be used by other researchers. Our major technical contribution is a formal derivation of reconstruction error, which provides a framework for minimization for both point-sampled and near-field BRDF acquisition, producing better results than the previous condition number heuristic.

In future work, we would like to explore other implications of our method. The new reconstruction error framework could have broad impact in problems like many light methods or computation of light transport matrices, where one seeks to reconstruct from a sparse set of samples. Finally, the one or two-shot nature of our method opens the possibility of designing new simple hardware, with light sources and camera in fixed position.

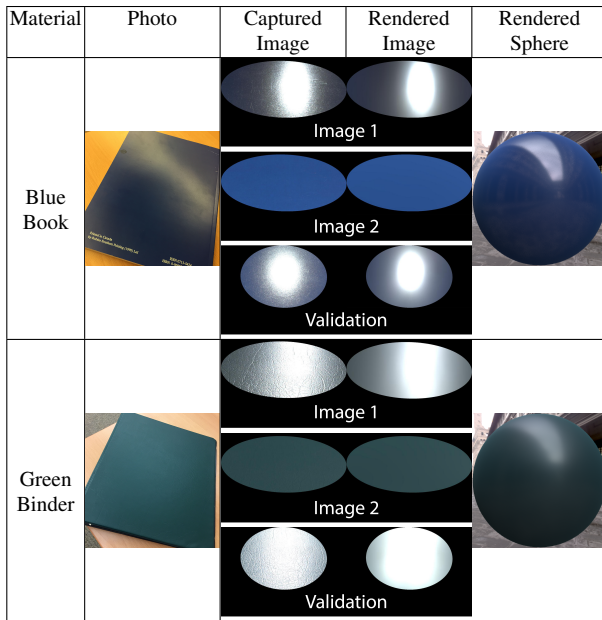


Figure 16: Acquisition of real BRDFs from two images with fixed camera. We also show a validation view not used as input.

Acknowledgements

We thank reviewers for their extensive comments. This work was supported by NSF grants 1451828, 1451830, and the UC San Diego Center for Visual Computing.

References

- AITTALA, M., WEYRICH, T., AND LEHTINEN, J. 2013. Practical SVBRDF capture in the frequency domain. *ACM Transactions on Graphics (TOG)* 32, 4, 110.
- AITTALA, M., WEYRICH, T., AND LEHTINEN, J. 2015. Two-shot SVBRDF capture for stationary materials. *ACM Transactions on Graphics (TOG)* 34, 4.
- BRADY, A., LAWRENCE, J., PEERS, P., AND WEIMER, W. 2014. genBRDF: discovering new analytic BRDFs with genetic programming. *ACM Transactions on Graphics (TOG)* 33, 4.
- DANA, K. J., AND WANG, J. 2004. Device for convenient measurement of spatially varying bidirectional reflectance. *J. Opt. Soc. Am. A* 21, 1, 1–12.
- FOO, S. C. 1997. *A gonireflectometer for measuring the bidirectional reflectance of material for use in illumination computation*. Master's thesis, Cornell University.
- FUCHS, M., BLANZ, V., LENSCH, H. P., AND SEIDEL, H.-P. 2007. Adaptive sampling of reflectance fields. *ACM Transactions on Graphics (TOG)* 26, 2, 10.
- GHOSH, A., ACHUTHA, S., HEIDRICH, W., AND O'TOOLE, M. 2007. BRDF acquisition with basis illumination. In *International Conference on Computer Vision*, 1–8.
- HERTZMANN, A., AND SEITZ, S. M. 2003. Shape and materials by example: A photometric stereo approach. In *Computer Vision and Pattern Recognition*, vol. 1, IEEE, 533–540.
- HUNTER, R. S., AND JUDD, D. B. 1939. Development of a method of classifying paints according to gloss. *ASTM Bulletin*, 97, 11–18.
- HUNTER, R. S. 1987. *The measurement of appearance*. John Wiley & Sons.
- KARNER, K., MAYER, H., AND GERVAUTZ, M. 1996. An image-

based measurement system for anisotropic reflection. *Computer Graphics Forum (EUROGRAPHICS 96)* 15, 3, 119–128.

- LENSCH, H. P. A., LANG, J., S, A. M., AND PETER SEIDEL, H. 2003. Planned sampling of spatially varying BRDFs. *Computer Graphics Forum* 22, 3, 473–482.
- LOW, J., KRONANDER, J., YNNERMAN, A., AND UNGER, J. 2012. BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG)* 31, 1.
- MARSCHNER, S., WESTIN, S., LAFORTUNE, E., TORRANCE, K., AND GREENBERG, D. 2000. Image-Based BRDF Measurement including Human Skin. In *Eurographics Rendering Workshop*, 139–152.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics (TOG)* 22, 3, 759–769.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. Efficient isotropic BRDF measurement. In *Eurographics Rendering Workshop*, 241–247.
- MCCOOL, M., ANG, J., AND AHMAD, A. 2001. Homomorphic Factorization of BRDFs for High-Performance Rendering. In *SIGGRAPH 01*, 171–178.
- MEYER, M., AND ANDERSON, J. 2007. Key point subspace acceleration and soft caching. *ACM Transactions on Graphics (TOG)* 26, 3.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Eurographics Symposium on Rendering*, 117–126.
- NIELSEN, J. B., JENSEN, H., AND RAMAMOORTHY, R. 2015. On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG)* 34, 6.
- NOLL, T., STRICKER, D., KOHLER, J., AND REIS, G. 2013. A full-spherical device for simultaneous geometry and reflectance acquisition. In *Proceedings of the on Applications of Computer Vision*, IEEE Computer Society, WACV, 355–362.
- REN, P., WANG, J., SNYDER, J., TONG, X., AND GUO, B. 2011. Pocket reflectometry. *ACM Transactions on Graphics* 30, 4.
- ROMEIRO, F., AND ZICKLER, T. 2010. Blind reflectometry. In *European Conference on Computer Vision*, 45–58.
- ROMEIRO, F., VASILYEV, Y., AND ZICKLER, T. 2008. Passive reflectometry. In *European Conf. Computer Vision*, 859–872.
- RUSINKIEWICZ, S. 1998. A new change of variables for efficient BRDF representation. In *Eurographics Rendering Workshop*, 11–22.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *JOSA* 57, 9, 1105–1112.
- TUNWATTANAPONG, B., FYFFE, G., GRAHAM, P., BUSCH, J., YU, X., GHOSH, A., AND DEBEVEC, P. 2013. Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Transactions on Graphics (TOG)* 32, 4, 1–12.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92*, 265–272.
- WESTLUND, H. B., AND MEYER, G. W. 2001. Applying appearance standards to light reflection models. In *SIGGRAPH 01*, 501–510.

Appendix A: Details of Reconstruction Error

We analyze equation 12 in more detail, also relating it to the condition number metric. It is convenient to denote $y = Qc$, where c is the accurate coefficient vector to reconstruct the BRDF. In this case, noting $Q^+Q = I$,

$$E_{\text{recon}} = \left| Q \left(Q^+ - (SQ)_\eta^+ S \right) Qc \right| = \left| Q \left(I - (SQ)_\eta^+ SQ \right) c \right|. \quad (16)$$

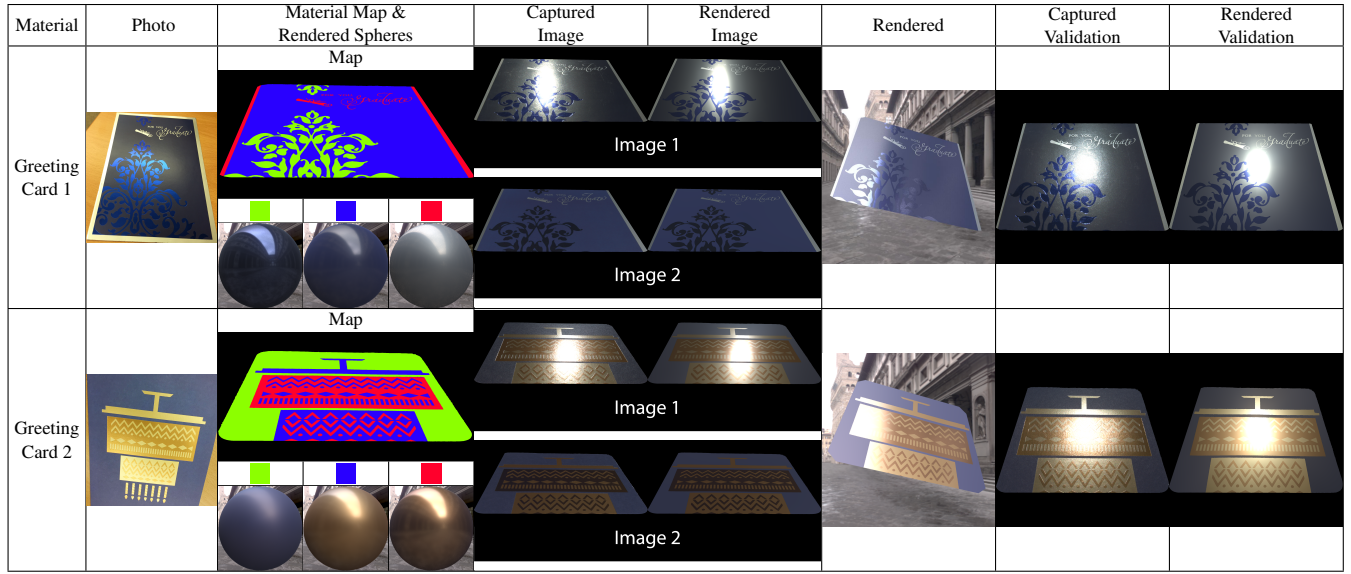


Figure 17: Acquisition of spatially-varying BRDFs from two images with fixed camera. Note the close match of captured and rendered images, including in the validation view, not used as input (rightmost column). Rendered images are under environment lighting. We also show the 3 material clusters used in each case, and spheres rendered with the full BRDFs recovered for each of the 3 materials.

Now, let us denote the SVD of $SQ = \tilde{Q}$ as $A\Lambda B^T$. From equation 5, $(SQ)_\eta^+ = \tilde{Q}_\eta^+ = B\Lambda_\eta^+ A^T$. Now,

$$(SQ)_\eta^+ (SQ) = B\Lambda_\eta^+ A^T A\Lambda B^T = B\Lambda' B^T, \quad (17)$$

where $\Lambda' = \Lambda_\eta^+ \Lambda$ is a diagonal matrix. If the singular values in Λ are σ (and those in Λ_η^+ are $\sigma/(\sigma^2 + \eta)$), then the singular values in Λ' are $\sigma^2/(\sigma^2 + \eta)$. Further simplifying,

$$E_{\text{recon}} = \left| Q \left(I - B\Lambda' B^T \right) c \right| = \left| Q \left(B\Gamma B^T \right) c \right|, \quad (18)$$

where Γ is also a diagonal matrix with singular values $1 - \frac{\sigma^2}{\sigma^2 + \eta} = \eta/(\sigma^2 + \eta)$. To understand E_{recon} , we care about the singular values in Γ . The largest singular value is given by the minimum σ_{\min} , with value $\eta/(\sigma_{\min}^2 + \eta)$. In general, we will reduce E_{recon} if we avoid small σ . Indeed, the condition number optimization affects the σ values and tries to make σ_{\min} larger to reduce the condition number. However, it is not explicitly minimizing the above expression. In contrast, our approach explicitly considers the end-to-end system, as well as the effect of Q , the MERL BRDF materials encoded in the coefficient vector c , and the full spectrum of singular values, to fully minimize the error E_{recon} .

Appendix B: Point-Sampled BRDF Measurement

The main text discusses near-field image-based BRDF measurement. Here, we show that the new error metric also somewhat improves point-sampled BRDF acquisition. We compare our results to [Nielsen et al. 2015] with 5 directions in Fig. 18. (The dotted black curve at the bottom is the lower bound when using all of the input directions, essentially the unavoidable error $E_{\text{deviation}}$.) Note that this evaluation is identical to Fig. 8 in their paper, using the same graphs for their method, as well as parametric fits and the industry-standard 5 directions in [Westlund and Meyer 2001]. It is clear that we have somewhat lower error. This is not surprising since these results are computed assuming the observations are accurate without noise, while the condition number metric measures only sensitivity to noise, not reconstruction error. The supplementary material shows similar results for an example with 20 measurements and 2% noise. Nevertheless, minimizing the condition number is a reasonable heuristic for this setup.

We can also plot the average error over the unknown samples in the MERL BRDF database vs. the number of measurements n in Fig. 19. For both standard point-sampled acquisition, and the image-based spherical acquisition method of [Marschner et al. 2000] (extended to use optimal directions computed with either our error metric or using condition number), our method




Figure 18: Comparison of reconstruction with our new optimized 5 directions, and those from [Nielsen et al. 2015], parametric fits, and industry-standard directions. Our method (green curve) produces lower error than previous work (blue curve) on each BRDF.




Figure 19: Reconstruction error versus number of measurements. We obtain a smooth graph, strictly lower error than previous work.

gives somewhat lower error. Another important observation is the shape of the curves. The result from [Nielsen et al. 2015] oscillates somewhat, since the condition number metric is not directly tied to (or always monotonic with) the actual error. By minimizing the actual expected reconstruction error, we obtain a smooth graph. The supplementary material provides our improved point-sampling directions, and comparisons for a few materials from the MERL database. In some cases we do qualitatively better, while there is a minor improvement in other cases. In general, our 5 directions is comparable to 20 samples using the previous condition number metric.