On Optimal, Minimal BRDF Sampling for Reflectance Acquisition

Jannik Boll Nielsen[*]   Henrik Wann Jensen[†]   Ravi Ramamoorthi[‡]

[*]Technical University of Denmark   [† ‡]University of California, San Diego

Abstract

The bidirectional reflectance distribution function (BRDF) is critical for rendering, and accurate material representation requires data-driven reflectance models. However, isotropic BRDFs are 3D functions, and measuring the reflectance of a flat sample can require a million incident and outgoing direction pairs, making the use of measured BRDFs impractical. In this paper, we address the problem of reconstructing a measured BRDF from a limited number of samples. We present a novel mapping of the BRDF space, allowing for extraction of descriptive principal components from measured databases, such as the MERL BRDF database. We optimize for the best sampling directions, and explicitly provide the optimal set of incident and outgoing directions in the Rusinkiewicz parameterization for $n = \{1, 2, 5, 10, 20\}$ samples. Based on the principal components, we describe a method for accurately reconstructing BRDF data from these limited sets of samples. We validate our results on the MERL BRDF database, including favorable comparisons to previous sets of industry-standard sampling directions, as well as with BRDF measurements of new flat material samples acquired with a gantry system. As an extension, we also demonstrate how this method can be used to find optimal sampling directions when imaging a sphere of a homogeneous material; in this case, only two images are often adequate for high accuracy.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture;

Keywords: reflectance, BRDF, MERL, reconstruction

1 Introduction

The bi-directional reflectance distribution function or BRDF [Nicodemus et al. 1977] characterizes material appearance, representing the unique reflectance of paints, metals, plastics, or velvet. The BRDF is a 4D quantity, depending on incident and outgoing directions, and can be reduced to 3D under the common assumption of isotropic reflectance. Historically, parametric BRDF models such as Phong [1975], Cook-Torrance [1982] or Lafortune [1997] have been used. In the past two decades, as the need for more realistic appearance and reproduction of real-world materials has increased, measured or data-driven

---

[*]jbol@dtu.dk
[†]henrik@cs.ucsd.edu
[‡]ravir@cs.ucsd.edu

Figure 1: *Rendering of isotropic BRDFs of unknown materials, reconstructed with our method from only 20 measurements captured at optimal sample locations (see Fig 10 for error plots).*

BRDF models have become more popular, and databases such as the MERL BRDF database [Matusik et al. 2003a] have been developed. In principle, it is easy to measure the BRDF using a gonioreflectometer (we achieve this in practice using a gantry system with a flat material sample). However, for a dense sampling of even 100 samples in each direction (critical for capturing sharp highlights), one would need a million samples or images over the full 3D isotropic BRDF. This is prohibitive in most applications, and considerably reduces the practicality of measured BRDFs.

In this paper, we address the fundamental problem of where to sample an unknown BRDF when only a very limited number of samples can be acquired. Our goal is to make measured reflectance practical in computer graphics for product design, virtual reality, appearance fabrication and entertainment, as well as in industrial quality control, where appearance must often be monitored in real-time. The question is how best to orient n light-sources and n cameras, such that the best impression of the appearance (BRDF) is obtained. We base our analysis on the MERL isotropic BRDF database [Matusik et al. 2003a], assuming that this covers the majority of real world isotropic material variation. We demonstrate that we can reconstruct a data-driven BRDF from a very small set of (fewer than 20) measurements. Our specific contributions are as follows:

New BRDF Mapping: We introduce a new mapping of the BRDF (Sec. 3) that enables a linear approach, with principal components from the MERL BRDF database that are highly descriptive of reflectance phenomena like diffuse and specular reflection, Fresnel effects and retroreflection (Fig. 3).

Optimization for Sampling Directions: We identify regions of importance based on the condition number of subsets of rows from the principal component matrix of the data [Ipsen and Wentworth 2014] (Sec. 4). The condition number directly relates to the quality at which one can expect to reconstruct a BRDF, with a given set of samples known from it. Based on this, we obtain a pri-

Figure 2: *Reconstructions of unknown samples (10 MERL BRDF samples not used at all for computing principal components (PCs) and sample directions). The BRDFs are rendered as spheres, illuminated by a front-light at a direction of $[1, 1, 1]$, and a back-light causing grazing angle reflections at $[-1, -1, -3]$. Reconstructions are made with $n = \{1, 2, 3, 5, 10, 20\}$ sampling points. In addition, reconstructions using all possible sampling points are shown. This corresponds to a projection of the data into PC-space and shows the best possible reconstruction. Finally, the far right column shows reference renderings of the true BRDFs. We see that generally 3-10 measurements are sufficient to capture the true appearance of a material. For some diffuse materials, more samples are needed to avoid ringing at the highlights. Alternatively, as mentioned in Sec. 8, a separate "soft" set of principal components can be used to avoid this ringing.*

oritized list of light/view direction pairs, which indicates the best directions to sample from in a point-sampling setup, when only n samples are to be acquired. The list does not, by far, cover all variability observable in BRDFs, but it does tell us where the strategically best places to sample are, ensuring that the maximum possible amount of unique information is captured with every sample. We differ from [Matusik et al. 2003b] in considering a very small number of measurements (10-20 instead of 800-1000), and we develop a novel optimization algorithm suited for our purpose. We provide the optimal set of incident and ougoing directions in the Rusinkiewicz [1998] half-difference parameterization for $n = \{1, 2, 5, 10, 20\}$ samples (Table 1 in Sec. 6), which can directly be used by practitioners.

BRDF Reconstruction and Validation: Inspired by the statistical reconstruction methods of [Blanz et al. 2004], we propose a similar ridge-regression based method to reconstruct full BRDFs from this sparse set of samples (Sec. 5). This method penalizes solutions that statistically deviate from the variability learned by the MERL database. We validate our directions and reconstruction method in two ways. First, we (randomly) leave out 10 materials in the MERL BRDF database, obtaining a slightly different set of principal components and sampling directions from the remaining 90. We then use these 10 materials as a validation set. Figure 2 shows the accuracy with increasing numbers of samples, demonstrating that 5-10 samples is often adequate, and 20 BRDF samples provides accurate results in all cases. Next, we consider flat samples of new real-world materials. We use $n = 20$ samples from five unknown materials, and compare the reconstructions with the densely measured in-plane reflectances of the materials (Sec. 7).

Extension to Image-Based BRDF Measurement: Our main contribution is for BRDFs acquired from a sparse set of measurements with a gonioreflectometer (in our case, a gantry setup with a flat material sample). We also demonstrate an extension to the popular image-based BRDF measurement method (Sec. 9), that takes 2D images of spheres of homogeneous materials rather than single measurements [Marschner et al. 2000] (and the approach used to acquire the database [Matusik et al. 2003a]). We demonstrate that two images are often adequate in this case (Table 2 and Fig. 12).

2 Related Work

The canonical approach to BRDF acquisition, considered here, is to sample individual light/view directions for a flat sample using a gonioreflectometer as done in e.g. [White et al. 1998]. As noted above, this is slow, and a variety of alternative setups and sampling strategies leveraging parallel acquisition of multiple BRDF samples with a CCD camera in conjunction with known curved geometry, mirror setups, and adaptive sampling schemes have been proposed.

Image-Based BRDF Measurement: Marschner et al. [2000] proposed a method where a sphere of homogeneous material is imaged (2D set of samples) from a 1D set of lighting drections, to obtain an isotropic 3D BRDF. Other, more complex setups, with curved mirrors etc. have also been proposed [Ward 1992; Dana and Wang 2004; Ghosh et al. 2007; Noll et al. 2013; Schwartz et al. 2014]. Indeed, the approach of Marschner et al. [2000] was used to acquire the MERL database [Matusik et al. 2003a]. However, this method is limited to spheres (or samples that can be painted/wrapped on a sphere). In this paper, we mainly address the canonical case of gonioreflectometric measurement for a flat sample, showing how a very sparse set of measurements suffices. We also demonstrate an extension to image-based BRDF measurement where two images is adequate, rather than full 1D lighting variation.

Adaptive Sampling: Fuchs et al. [2007] suggest an adaptive sampling scheme where the observed properties of the BRDF being measured are taken into account. Here, regions needing to be refined are detected by evaluating how well a sample is modeled by its neighboring samples. A major limitation of this method is that a somewhat dense grid of samples must be acquired before the refinement procedure converges well. In contrast, our method is non-adaptive, and we use a very sparse fixed set of (precomputed and tabulated) sampling directions for any material. Lensch et. al. [2003] suggest a different approach, where the next sampling direction is estimated based on an uncertainty measure of the fitted BRDF parameters. This approach allows for full control of the number of samples used, but in contrast to our method, it relies on fitting a parametric model, with the resulting limitations.

Complex Environment Lighting: Some recent work has explored acquisition of simpler 2D BRDF models under environment lighting [Romeiro et al. 2008] or even using a controlled environment [Ghosh et al. 2007; Tunwattanapong et al. 2013; Aittala et al. 2013]. These methods have shown convincing results, where different variants of convolutional theory are used to obtain homogeneous or spatially varying BRDFs. However, they require solving non-linear systems and regularization, while we focus on more direct measurement of BRDFs from individual samples. Alternative approaches where the environment is unknown but geometry is known have also been suggested [Romeiro and Zickler 2010]; however, the ambiguity between environment and BRDF can be difficult to resolve, requiring various heuristics for regularization.

Rapid Reflectometry: Ren et al. [2011] propose pocket reflectometry, where a reference checkerboard with known reflectance tiles is used in conjunction with a handheld lightsource and a fixed camera. Older work simply compares reference BRDFs to the target-BRDF [Hertzmann and Seitz 2003]. These are somewhat related to our proposed method in that they reconstruct appearance based on reference BRDFs, with the difference that they are limited by the relatively small number of reference BRDFs used and by the need of physically placing the references in the scene.

Efficient Sampling: Closest to our work is that of Matusik et al. [2003b], who mention that for $n = 800$ samples, an unknown BRDF can be modeled by a linear combination of other BRDF models from their dataset. Note that this is for measuring individual BRDF samples, as in our paper, rather than for image-based measurement. Their method is closely related to the optimum sampling directions in our work. However, they do not explicitly provide which directions these 800 samples correspond to, nor investigate how this quality converges. Weyrich et al. [2006] utilizes this method to reconstruct spatially varying BRDFs of human skin, and [Hullin et al. 2010] uses it to interpolate fluorescence measurements for BRRDFs.

We extend the work of Matusik et al. [2003a; 2003b] by introducing a novel BRDF mapping that makes linear modeling possible from fewer samples. Our mapping addresses the issues of non-physical and odd-looking reconstructions, including those having negative reflectance values. Critically, we focus on optimal *minimal* sampling with $n < 20$, rather than $n = 800$, showing that this much smaller set suffices. We develop an optimization method that quickly yields optimal sampling locations for small values of n, compared to the greedy approach [Matusik et al. 2003b]. We provide tables of the $n = 20$ optimal directions, that can be used directly. We also suggest a reconstruction approach that takes into account the statistical variability of BRDFs based on PCA, rather than linear combinations of raw measured BRDFs. Finally we discuss the extension to capturing images of spheres, rather than classical gonioreflectometric point-sampling.

Materials Standards: In the materials industry, various standards have been proposed, and are widely used, to characterize reflectance. For specular or glossy materials, the simplest is the method of Hunter and Judd [1939], who simply measured the reflectance at 60° perfect reflection. An extension to this is adding the near-normal and grazing-angle behavior by measuring the reflectance at 85° and 20° perfect reflection as well [Hunter 1987]. This is also known as the ASTM standard D523. Within the more complex types of materials, such as car-paint and pearlescent colors, additional information is required. Westlund and Meyer [2001] describe that the sets of aspecular angles $\{15°, 45°, 110°\}$ and $\{25°, 45°, 75°\}$, both with 45° incident light, have been suggested for characterizing these more complex materials. Westlund et al. themselves utilize all five angles ($\{15°, 25°, 45°, 75°, 110°\}$) to obtain maximum information about the materials. We adapt these five directions in our comparison in Sec. 6, showing that our method provides significantly more accurate results with our five optimized directions.

Parametric Fits: Finally, considerable effort has been devoted to fit parametric BRDF models to real-world observations. These include both empirical ([Phong 1975; Blinn 1977; Ward 1992; Lafortune et al. 1997]) and physically-based ([Torrance and Sparrow 1967; Cook and Torrance 1982]) reflectance models, and recently also more advanced parametric BRDFs have been proposed, aiming at reproducing the behaviors observed in MERL ([Löw et al. 2012; Brady et al. 2014]). As all models are simplifications of the true behavior, one has to determine what regions of a BRDF the models should prioritize, i.e., what objective function the optimizers should minimize. Using the L_1-norm to emphasize the base of a specular peak, rather than the extremum has been suggested, and various transformations of the observed data have also been proposed such as square or cubic root, logarithmic, and cosine weighting [Ngan et al. 2005]. We leverage some of these observations, including logarithmic mapping and cosine-weighting in our BRDF remapping, but also normalize by the average or a reference BRDF. Crucially, we focus on data-driven reflectance, rather than parametric BRDF models. It is also important to note that although parametric models are thought of as compact, even the most simple models, such as the Phong model, have at least 7 tunable parameters (three if ignoring color). This calls for at least the same number of samples, but ideally many more to get a robust fit. Given our sparse data, even the number of parameters in the most simple analytical models is often comparable to, or exceeds, the number of samples we use. In addition, our proposed method is based on solving a linear system, thus eliminating the need for non-linear optimization and the challenge of local minima.

3 Data and Mapping

We base our analysis on the MERL isotropic BRDF database [Matusik et al. 2003a]. This database covers 100 materials of varying reflectance, from soft diffuse materials like rubber, to hard specular materials like chrome. The reflectance is represented by dense BRDF measurements in a 3D volume using Rusinkiewicz half-difference angle coordinates $(\theta_h, \theta_d, \phi_d)$ [Rusinkiewicz 1998]. The resolution for each RGB color is $(90 \times 90 \times 180)$, or $1,458,000$ measurements. We seek to dramatically reduce the number of measurement samples needed to 20 or fewer. Principal Component Analysis (PCA) has been performed on this dataset before [Matusik et al. 2003a; Ngan et al. 2006], with the obervation that a linear transformation on the raw BRDFs, such as PCA, is inferior to non-linear transformations. Furthermore, Matusik et al. [2003a] observed that linear combinations of raw BRDFs could give rise to unrealistic results such as negative reflectance values and "holes".

Log-Relative Mapping: The poor performance of PCA on raw BRDF values is closely related to the high dynamic range within the BRDF (which is several orders of magnitude), combined with the variation clearly not following a normal distribution. This fact motivates the use of a mapping of data that preserves inter-BRDF variation while dampening intra-BRDF variation. This mapping should at the same time also make the data more normally distributed. We propose a novel log-relative mapping for this purpose:

$$\rho \mapsto \ln\left(\frac{\rho \ \cos_{\text{weight}} + \epsilon}{\rho_{\text{ref}} \ \cos_{\text{weight}} + \epsilon} \right), \tag{1}$$

where ρ is the BRDF, ϵ is a small constant that avoids division by zero and sensitivity to camera noise (we use $\epsilon = 0.001$), and ρ_{ref} is a *reference* BRDF, relative to which the mapping is applied. In our experiments, we choose the reference BRDF to be the median value for each $(\theta_h, \theta_d, \phi_d)$, over the entire dataset of BRDFs. We choose the median rather than the mean, to make the mapping more robust towards outliers in the dataset. A cosine-weight is applied to compensate for extreme grazing-angle values, as in previous work,

$$\cos_{\text{weight}} = \max\left\{ \cos(\mathbf{n} \cdot \boldsymbol{\omega}_i) \cos(\mathbf{n} \cdot \boldsymbol{\omega}_o) , \ \epsilon \right\}, \tag{2}$$

where $\mathbf{n}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o$ are the normal, illumination-direction, and view-direction vectors respectively, obtained from converting a BRDF location $(\theta_h, \theta_d, \phi_d)$ to the Cartesian coordinate frame. The reference BRDF, ρ_{ref}, is shown in Fig. 3.

Note that every $(\theta_h, \theta_d, \phi_d)$ BRDF value is essentially compared individually to a reference BRDF, and the natural logarithm is then used to weight smaller and greater values equally (i.e., a value *half* the size of the reference value should have the same magnitude as a value twice as big as the reference).

RGB Colors as Independent BRDF Samples: It is apparent that many (most) BRDF behaviors are covered by the database, with materials ranging from smooth and soft, to hard and specular, with varying Fresnel effect (refractive indices), and even retro-reflection. It is however also apparent that not all color-variation has been covered, e.g., there is a pink specular material, but not a pink soft material. Therefore, we choose to work in grey-scale and treat each of the 3 color channels of the BRDFs as individual observations. This allows us to abstract away the color, and effectively have three times as many observations (material samples), at the cost of the samples being a little correlated. Our algorithm then simply reconstructs each color channel of the BRDF separately; this does not require any additional measurements.

With 100 materials, each having 3 color-channels that are being treated as individual grey materials, a total of $m = 300$ BRDFs are available in the database. Each material can be vectorized as a point in a $p = 90 \cdot 90 \cdot 180 = 1,458,000$-dimensional space. This gives an observation matrix $\mathbf{Y} \in \mathbb{R}^{m \times p}$. As mentioned above, we use the median of the dataset as the reference BRDF, $\boldsymbol{\rho}_{\text{ref}} \in \mathbb{R}^p$:

$$\rho_{\text{ref},i} = \text{median}\left(Y_{[1,i]}, Y_{[2,i]}, \ldots, Y_{[m,i]} \right) \tag{3}$$

Using the reference BRDF, all observations (rows) in $\mathbf{Y}$ can be mapped using the mapping in Equation 1, resulting in the mapped observation matrix $\mathbf{X}$:

$$X_{j,i} = \text{mapping}(Y_{[j,i]}, \rho_{\text{ref},i}). \tag{4}$$

Principal Components: Up to $k = m = 300$ principal components may be extracted from $\mathbf{X}$ by performing singular value decomposition (SVD) of the mean-subtracted mapped observation matrix:

$$(\mathbf{X} - \hat{\boldsymbol{\mu}}) = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T, \tag{5}$$
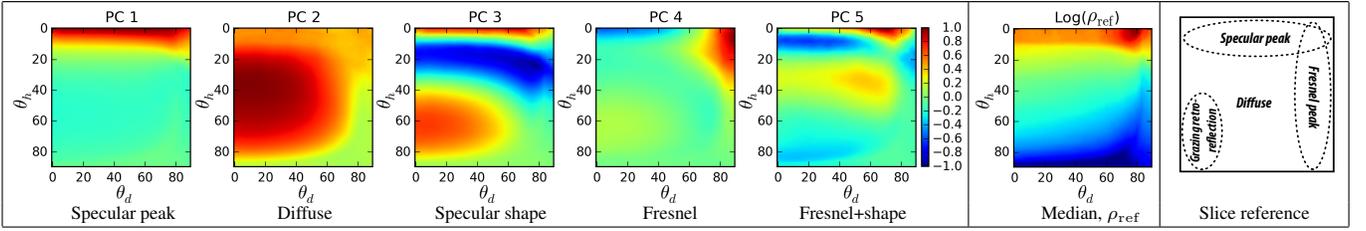
Figure 3: *BRDF-slices of the first 5 principal components of the MERL database, using the mapping from Equation 1. Slices are made at $\phi_d = 90°$, and an interpretation of them is shown on the far right, adapted from [Burley 2012] . Notice how the primary variation (the first principal component) is in the specular peak, and the second principal component primarily contributes to the diffuse region of the BRDF. The third component contracts or widens the specular peak. The fourth component adds Fresnel-effect and finally the fifth both contributes to Fresnel, shaping the specular peak, and a bit of retro-reflection.*
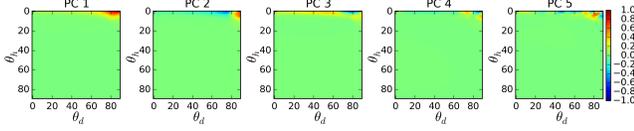
Figure 4: *BRDF-slices of the first 5 principal components of the MERL database without using any mapping. Slices are made at $\phi_d = 90°$. Due to the large dynamic variation in the raw BRDFs, the specular direction is dominating all variation.*

where $\boldsymbol{\mu} \in \mathbb{R}^{1 \times p}$ is the mean over m rows of $\mathbf{X}$, and $\hat{\boldsymbol{\mu}} \in \mathbb{R}^{m \times p}$ simply repeats $\boldsymbol{\mu}$ for each row. The columns of $\boldsymbol{V} \in \mathbb{R}^{p \times k}$ are eigenvectors of the covariance $(\boldsymbol{X} - \hat{\boldsymbol{\mu}})^T(\boldsymbol{X} - \hat{\boldsymbol{\mu}})$, and correspond to the principal components of the data. The diagonal elements, σ_i^2, of $\boldsymbol{\Sigma} \in \mathbb{R}^{k \times k}$ correspond to the variance explained by the i^{th} principal component.

In Fig. 3, the first 5 principal components are visualized as BRDF-slices [Burley 2012]. An interpretation reference is shown to the far right in the figure. The slices reveal very interesting structures of the principal components: the first principal component largely models the intensity of the specular peak, by affecting only the region where θ_h is close to zero. The second principal component models diffuse contribution allowing for either raising or lowering the overall reflectance for all non-grazing angles. The blue band in the third component allows for widening or contraction of the specular peak. Finally, the fourth and fifth components model various combined effects; most notably the Fresnel effect contribution at θ_d close to $90°$ and the additional shaping of the specular peak in PC5, but also a small contribution to the retro-reflective behaviour in the lower left corner. These informative principal components are a result of the log-mapping introduced in Equation 1. Without it, the majority of all numerical variation is concentrated in the specular peak, drowning all off-peak variation of the BRDFs. This is shown in Fig. 4, where slices of the first 5 unmapped principal components are illustrated. Although these components numerically describe the data well, we see qualitatively that they hold little information about the overall nature of a BRDF.

Finally, as it will later become convenient, a matrix of scaled principal components, $\boldsymbol{Q} \in \mathbb{R}^{p \times k}$, can be obtained by:

$$\boldsymbol{Q} = \boldsymbol{V}\boldsymbol{\Sigma}. \qquad (6)$$

Here, the length of each principal component has been scaled by the amount of variance it covers.

4 Optimization of Sampling Directions

Before introducing the optimization scheme used to determine optimal sampling directions, we give the general idea behind recon-struction of BRDFs, as this is the primary driver behind the objective function used to choose the sampling directions.

Using the scaled principal components $\boldsymbol{Q} \in \mathbb{R}^{p \times k}$ obtained from Equation 6, a new BRDF $\boldsymbol{x}$ may be synthesized using a linear combination $\boldsymbol{c}$ of the principal components ($\boldsymbol{c} \in \mathbb{R}^{k \times 1}$ is a vector):

$$\boldsymbol{x} = \boldsymbol{Q}\boldsymbol{c} + \boldsymbol{\mu}, \qquad (7)$$

and likewise, had $\boldsymbol{x}$ been known, the linear combination of principal components best modeling it could be found by solving the linear system for $\boldsymbol{c}$. The expected sensitivity to errors in this modeling can be estimated by the condition number κ of the matrix $\boldsymbol{Q}$:
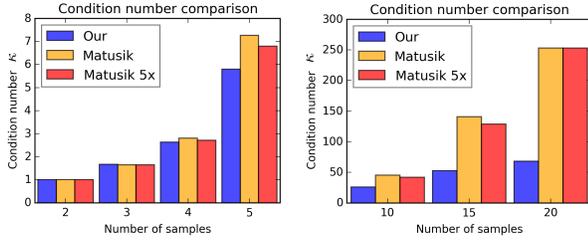
$$\kappa(\boldsymbol{Q}) = \frac{\sigma_{\max}(\boldsymbol{Q})}{\sigma_{\min}(\boldsymbol{Q})} = \frac{\sigma_1}{\sigma_k}, \qquad (8)$$

which is the ratio between the maximum and minimum singular values of $\boldsymbol{Q}$. The lower the condition number, the less the sensitivity to noise and numerical errors [Ipsen and Wentworth 2014].

Recall that every $(\theta_h, \theta_d, \phi_d)$ location in the BRDF volume corresponds to a specific light/view-direction, and that every location additionally corresponds to a row in $\boldsymbol{Q}$. Therefore, the problem of determining the n best directions for sampling becomes a problem of determining the reduced matrix $\tilde{\boldsymbol{Q}} \in \mathbb{R}^{n \times k}$ of n rows from $\boldsymbol{Q}$, that minimizes the condition number $\kappa(\tilde{\boldsymbol{Q}})$.

Minimizing Condition Number: Matusik et al. [2003a] use a greedy algorithm to evaluate when the condition number of $\tilde{\boldsymbol{Q}}$ stabilizes for increasing n. The strategy is to initially pick n random rows from $\boldsymbol{Q}$ and then in a random order try to replace rows in $\tilde{\boldsymbol{Q}}$ with random rows from $\boldsymbol{Q}$, while only keeping the swaps that reduce $\kappa(\tilde{\boldsymbol{Q}})$. Our experience with this approach is that it converges very slowly for small n (see Fig. 5).

The rows of $\boldsymbol{Q}$ are formed by a vectorization of a discrete three-dimensional volume holding values that exhibit continuous variation. This means that the rows of $\boldsymbol{Q}$ are not uncorrelated, and more importantly, we can estimate the gradient of a row by looking up its neighbors in the volumetric representation. Rather than treating the rows in $\boldsymbol{Q}$ as independent and without structure, as Matusik et al.'s method does, we thus propose utilizing gradients to more effectively minimize condition number, $\kappa(\tilde{\boldsymbol{Q}})$. We found that standard numerical optimizers have difficulties in solving this minimization. This is partly due to the integer steps required when moving through the BRDF volume, and partly due to the invalid regions existing in the BRDF volume (views below horizon). We therefore develop our own simple algorithm. We start with a random initialization of $r \leq n$ sampling points (we used $r = 1$). We then randomly pick one of the points, and numerically evaluate the gradient $\nabla\kappa = \left(\frac{\delta\kappa(\tilde{\boldsymbol{Q}})}{\delta\theta_h}, \frac{\delta\kappa(\tilde{\boldsymbol{Q}})}{\delta\theta_d}, \frac{\delta\kappa(\tilde{\boldsymbol{Q}})}{\delta\phi_d} \right)$, moving along this

Figure 5: *Comparison of condition number for our and Matusik et al.'s method. Blue: Condition number reached by our method at convergence. Orange: condition number reached by Matusik et al.'s method using the same computation time as our method. Red: Condition number for Matusik et al.'s method using 5 times the computation time as our method.*

gradient for a pre-determined step-length. This is repeated until convergence. After convergence if $r < n$ we add a point and repeat optimization. The method was not very sensitive to step-length. We chose an initial length of $3°$ (3 cells), and reduced this to $1°$ (1 cell) when the former step-length had converged.

Algorithm: A summary of the full algorithm to minimize condition number to choose optimal sampling directions is as follows:

1. Pick $r \leq n$ random sampling locations in $(\theta_h, \theta_d, \phi_d)$. For a more robust initial guess, repeat this randomization multiple times and choose the guess that had the lowest $\kappa(\tilde{Q})$.

2. Randomly choose one of the r points. Estimate $\nabla \kappa$ and move the point one step-length in this direction if the destination is a valid location in the BRDF volume. Otherwise, clip $\nabla \kappa$ accordingly. Repeat until convergence.

3. Optionally reduce step-length and repeat (2).

4. If $r < n$, add a new point (r++) and repeat (2).

(In the special case of $n = 1$, the condition number cannot be estimated and instead we utilize the leverage or 2-norm of the rows in $\tilde{Q}$ as suggested by [Ipsen and Wentworth 2014].)

Validation: This approach works very well for small values of n, whereas it becomes comparable with the method of [Matusik et al. 2003a] for very large values (beyond the scope of this paper). Within our goals of minimal sampling, $n < 20$, the gains are significant as shown in Fig. 5. We use our method to find the optimum set of n rows that minimizes $\kappa(\tilde{Q})$ and plot this (blue bar). Afterwards we run Matusik et al.'s method for the same amount of computation time and plot the obtained condition number (orange bar). Finally we let Matusik et al.'s method run for 5 times the computation time we used and plot the obtained condition number (red bar); as can be seen, the convergence of Matusik's method is slow, so additional computation time does not significantly change his results. Our condition numbers are significantly lower for $5 \leq n \leq 20$.

To evaluate if our method does in fact converge to a global minimum, we first found the ground truth global minimum by brute-force computations for $n = 2$, and then repeated the algorithm 50 times with different random initial conditions, plotting the resulting sampling locations, the minimum error sampling over all runs (the final result of our algorithm), and the ground-truth, as shown in Fig. 6. Red points indicate the first sampling direction and pink points indicate the second. The blue circles indicate the best solution found and the white stars (on top of them) indicate the global minimum found by brute force computations. The red points are
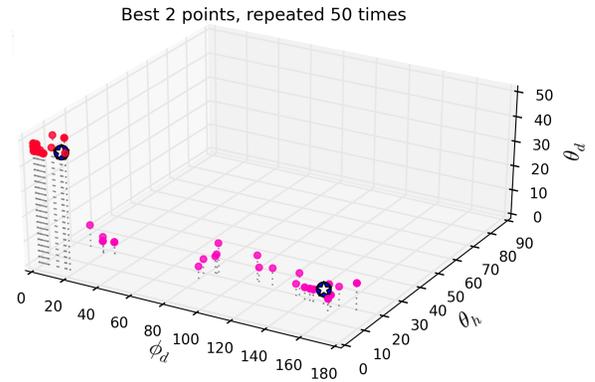
Figure 6: *Optimization repeated 50 times for $n = 2$ sampling directions. Red corresponds to first sampling direction, purple to second. The blue circles indicates the best constellation found. The two white stars indicate the ground truth global minimum found through brute-force evaluation of all combinations ($\kappa = 1.00008$). Our method is seen to correctly find the global minimum. Notice that all red points are fairly well clustered around the global minimum, indicating that our method converges well every time. A scattering is seen over ϕ_d for the purple points; this is however related to the fact that ϕ_d becomes ambiguous for $\theta_d \to 0°$.*

seen to be well clustered around the global minimum. For the purple points we observe a lot of scattering over ϕ_d; this is however related to the fact that ϕ_d becomes ambiguous for $\theta_d \to 0°$ as is the case here. The blue circles and white stars align, showing that we do in fact find the global minimum $\kappa = 1.0008$.

5 Reconstruction

Given a small number of samples n from an unknown BRDF, we wish to reconstruct the missing elements of the BRDF. We do this by projecting the known information into the PCA space, and by using the information of the principal components in Q to reconstruct the remaining information.

Let $\tilde{x} \in \mathbb{R}^n$ be the vector of known values of a BRDF, let $\tilde{\mu} \in \mathbb{R}^n$ be the vector of corresponding mean values from Equation 5, and let $\tilde{Q} \in \mathbb{R}^{n \times k}$ be the corresponding rows of the principal components in Q. The linear combination of principal components, c, that best models the observed data is then obtained by:

$$
\begin{aligned}
(\tilde{x} - \tilde{\mu}) &= \tilde{Q}c \\
c &= \operatorname*{argmin}_{c} \|(\tilde{x} - \tilde{\mu}) - \tilde{Q}c\|^2 \\
&= \left(\tilde{Q}^T \tilde{Q}\right)^{-1} \tilde{Q}^T (\tilde{x} - \tilde{\mu}), \quad (9)
\end{aligned}
$$

and the full BRDF, $x \in \mathbb{R}^p$, is then reconstructed by using the full principal components:

$$
x = Qc + \mu. \quad (10)
$$

Finally, we can apply the inverse mapping from Equation 1 to obtain the original unmapped BRDF.

Although the least squares solution above is unbiased, it usually results in severely over-fitted results, deviating significantly from ground truth. Blanz et al. [2004] addresses this issue using ridge regression and we adapt this approach for BRDFs. The length of the scaled principal components, Q, is proportional to the amount of variance they explain. Hence, the magnitude of the elements in c directly links to how much a fit deviates from the mean behavior,
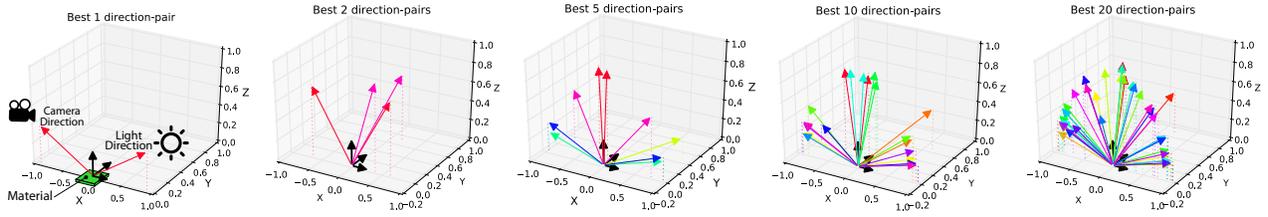
Figure 7: *Optimal sampling directions for $n = \{1, 2, 5, 10, 20\}$ samples. Each color pair corresponds to a view/illumination-combination. Notice that for a single observation, the $60°$ constellation is the optimum. When extending to two observations, an off-peak sampling should be performed, and continuing to 5, 10 and 20 we see that the majority of measurements should lie within the plane of incidence.*

i.e., $\|c\|^2$ is proportional to the unlikelihood of a reconstruction. Introducing the hyper-parameter η in conjuction with the 2-norm of c, it is possible to favour reconstructions closer to the observed distribution of BRDFs:

$$c = \underset{c}{\operatorname{argmin}}\|(\tilde{x} - \tilde{\mu}) - \tilde{Q}c\|^2 + \eta\|c\|^2. \qquad (11)$$

This expression has the closed form solution:

$$c = \left(\tilde{Q}^T\tilde{Q} + \eta I\right)^{-1}\tilde{Q}^T(\tilde{x} - \tilde{\mu}), \qquad (12)$$

where I is the identity matrix. As may be seen, a value of $\eta = 0$ corresponds to the least squares solution in Equation 9, and for $\eta \to \infty$, c goes to zero and the reconstruction moves towards the mean solution μ. We have found that the method is not very sensitive to the value of η; we use $\eta = 40$. For completeness we evaluated the use of L_1 penalized minimization (Lasso, [Tibshirani 1996]) to obtain a sparser set of PCs, and did not observe better performance.

Finally, recall that as mentioned in Sec. 3 color information has been removed from the training data. This means that in order to reconstruct a colored BRDF, each color channel must be reconstructed separately by Equations 10 and 12.

6 Results

Based on the optimization approach described in Sec. 4, we were able to obtain a prioritized list of optimal sampling directions for $n = \{1, 2, 5, 10, 20\}$ samples, and validate the accuracy of reconstruction using these directions for both the MERL database materials, as well as for new measured BRDFs (Sec. 7).

Optimal Sampling Directions: The optimization method was repeated 50 times in order to verify repeatability, and to enable us to pick the minimum condition-number solution. In Table 1, these directions are listed for different values of n, and they are visualized by colored pairs of arrows for view-illumination combinations in Fig. 7. We emphasize that these directions can be used directly for minimal sampling of any BRDF within a gonioreflectometer or gantry-based setup for a flat sample; the implementer need not re-run our optimization. We will also make the source code for reconstruction, principal components and mean values available online upon publication, to directly allow application of Equations 10 and 12 for BRDF reconstruction from the sparse measured samples.

Note that the computed sampling directions lie primarily in the plane of incidence and, for some measurements, resemble the as-pecular directions $15°, 25°, 45°, 75°$, and $100°$ reported by Westlund and Meyer [2001] to be industry standards for measuring goniochromatic surfaces. Intuitively, it makes sense that the single most important measurement is at the perfect reflection around $60°$ c.f. Fig. 7, in order to determine the intensity of the specular peak. Likewise, for $n = 2$, we see that while the first measurement should

n	$\theta_h[°]$	$\theta_d[°]$	$\phi_d[°]$	n	$\theta_h[°]$	$\theta_d[°]$	$\phi_d[°]$
1	0	66	0		9	12	69
2	2	40	24		66	35	103
	29	21	52		4	76	150
5	1	4	36		3	79	107
	77	13	86		33	63	105
	8	79	85		1	7	174
	3	74	145		21	8	180
	13	52	80		7	64	180
10	3	12	28		6	73	180
	63	19	89		2	63	95
	5	77	77	20	11	11	124
	2	60	180		6	74	123
	15	4	130		1	62	129
	1	6	37		33	49	24
	2	79	110		14	71	55
	39	76	89		8	35	180
	0	71	104		31	77	91
	5	75	180		2	72	130
					2	47	113
					10	43	68

Table 1: *Optimium sampling directions in Rusinkiewicz coordinates [Rusinkiewicz 1998], when n BRDF samples can be acquired.*

capture the specular peak, the second measurement moves far off-peak and measures the diffuse component of the material. It is important to emphasize, as stated in the introduction, that the reported sampling directions do not cover the full variability of BRDFs. To fully capture a BRDF, generally thousands of samples are required [Lensch et al. 2003]. What these results provide are the strategically best locations to sample from in a minimal sampling setup, in order to capture as much unique information as possible per sample. Note also that these directions hold for all BRDFs, and no complex adaptive acquisition scheme is required.

Reconstruction Quality on MERL Database: To assess the reconstruction quality, we first describe experiments on the MERL BRDF database itself; the next section discusses validation on new samples. To avoid overfitting, the data was randomly split into two groups: 90 materials for training, and 10 materials for testing. The optimization was performed again on the reduced training-dataset and although it did not provide exactly the same directions, these were close to what was observed for the full dataset. Hence, the 10 testing materials are completely separate from the training phase.

We first evaluate the performance of the sampling directions, by comparing the reconstruction quality using the five industry standard directions reported by Westlund and Meyer [2001], with the reconstruction quality using the five optimal directions we have obtained. Using ours and Westlund and Meyer's [2001] five directions for sampling—which correspond to five known rows in x and Q—the testing materials were reconstructed and the results compared to the ground truth. Fig. 8 shows the root-mean-squared error of mapped values between reconstruction and ground truth, normal-
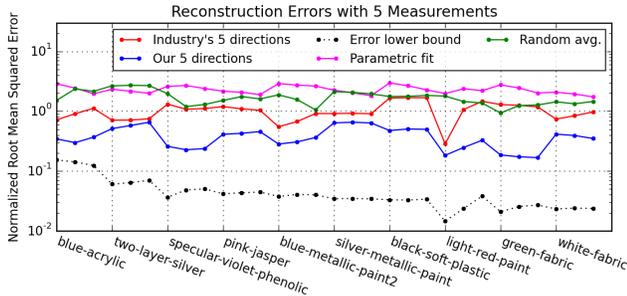
Figure 8: *RMS error of reconstruction of unknown mapped BRDFs normalized by mean mapped BRDF value, using (blue) our 5 best points of sampling, (red) 5 common industry directions [Westlund and Meyer 2001], and (dashed) all values of the BRDF. For comparison we also fitted the Ward parametric model to our 5 samples (purple), and also evaluated the average of 20 reconstructions using 5 random directions (green). Each material is represented by 3 datapoints for its R, G, and B errors.*
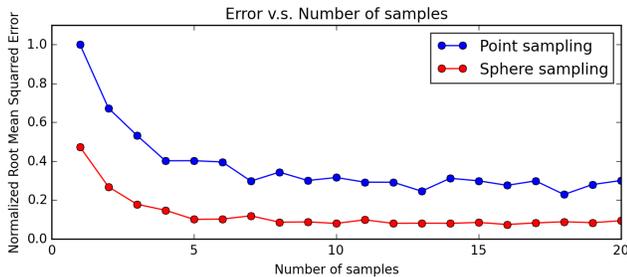
Figure 9: *Normalized average reconstruction error of unknown samples versus number of sampling directions used.*

ized by the mean mapped BRDF value, for three different reconstructions: using our five directions; using the five directions of Westlund and Meyer [2001]; using all directions ($\tilde{\boldsymbol{x}} = \boldsymbol{x}$), which is a projection of the data into PCA space and is the lower limit for the error. In addition, Fig. 8 also shows the (much higher) error of the parametric Ward model fitted to our five sampling directions (purple), as well as the average error of 20 runs where five purely random sampling directions were used. It is apparent that our proposed five directions outperform the industry sampling directions in reconstruction quality by up to an order of magnitude, and have very low error.

An evaluation of reconstruction error for an increasing number of sampling points is shown in the blue curve in Fig. 9 (the red curve for sphere sampling corresponds to Sec. 9). The figure shows the average normalized root-mean-squared error of reconstruction for the 10 testing materials, using a range of 1 to 20 sampling directions. Note that already at 4-5 samples the average error is starting to stabilize. The long tail following indicates that whereas the first samples yield great improvements to the reconstruction, latter samples only improve it slowly.

In Fig. 2, reconstructions of the 10 testing materials for an increasing number samples, $n = \{1, 2, 3, 5, 10, 20\}$, are shown. The renderings of spheres with the respective materials are done using a front light at a direction of $[1, 1, 1]$ and a back-light causing grazing angle reflections from a direction of $[-1, -1, -3]$. In addition to the low-number sampling directions, projections into PC space are also shown. These are made by fitting the principal components to the full BRDFs and illustrate the best possible reconstructions obtainable by the PCs. Finally, in the last column, reference renderings of the materials are shown. In accordance with Fig. 9, it

is observed that the reconstructions stabilize after $n = 5$ samples, and at $n = 20$ there are no noticeable visual differences between the reconstruction and the reference. A single outlier is the "white-fabric" material, where the method has difficulties capturing the diffuse appearance of the material due to the strong specular prior in the data. This is addressed in Sec. 8 and a better reconstruction is shown in the last row of the figure.

In the supplementary material, we also report results on all 100 materials (without separating training and testing data) for completeness, with comparable results.

7 Validation on New Materials

To validate the reconstruction method with real data, beyond what is found in the MERL database, the in-plane BRDF of five flat samples of new materials was densely sampled using a spherical gantry. The materials evaluated were: a glossy blue book, a brown-red notebook with a smooth highlight, a specular binder-cover, a diffuse piece of green cloth, and a diffuse piece of yellow paper. We considered the BRDF at a single spatial location, the center of the sample. Note that diffuse materials are not purely Lambertian. Reference images of the five samples are shown in the lower right corners in Fig. 10. The in-plane BRDF profiles were densely sampled for all materials with a $1°$ resolution using a $45°$ incident light. In addition, the BRDF at the best 20 sampling directions, listed in Table 1, were also acquired, in order to apply our method to reconstruct the full 3D isotropic BRDF.

Using our proposed reconstruction method from Equation 12, we reconstruct the full BRDFs, using the best $n = 20$ samples. In the left column of Fig. 10, the measured in-plane BRDF values are plotted as solid curves for red, green, and blue channels. The in-plane reflectances, extracted from the full reconstructed BRDFs are plotted as dashed curves. Note that these plots are 1-D curves extracted from the full 3D BRDF volumes and as such are only very small fractions of all the data that has been reconstructed. In general, there is a very good match, indicating accurate BRDF reproduction on real samples.

The most significant deviations are observed in the 3rd reconstruction, the binder-cover, and the 4th reconstruction, the green cloth. For the binder-cover, the diffuse component remains constant much longer than what has been learned from the data as being "natural". This is most likely caused by a very isotropic subsurface scattering in the material. For the green cloth, a retroreflective behavior is observed at the incident light direction. Although retroreflective behavior is represented by a few samples in the MERL database, it is not enough to match the retroreflectivity of the cloth using the first 20 principal components. Note that overall appearance of materials is reproduced well, with very few BRDF measurements.

Renderings of the materials using the same rendering setup as in Fig. 2 are shown in the top right corners of Fig. 10. In addition, renderings of the materials using more complex geometry (Killeroo model) and high dynamic range environment lighting (Grace Cathedral, [Debevec 1998]) are in Fig. 1. Color variations are due to the different colors in the environment lighting.

8 Refinement

During reconstructions we noticed that for diffuse materials, a ringing sometimes appeared around the specular peak. This ringing is caused by the bias towards specular materials in the MERL database. This bias is not only due to a predominance of specular materials, but also due to the numerical magnitude that specular peaks have. A simple way of addressing this ringing, if needed, is
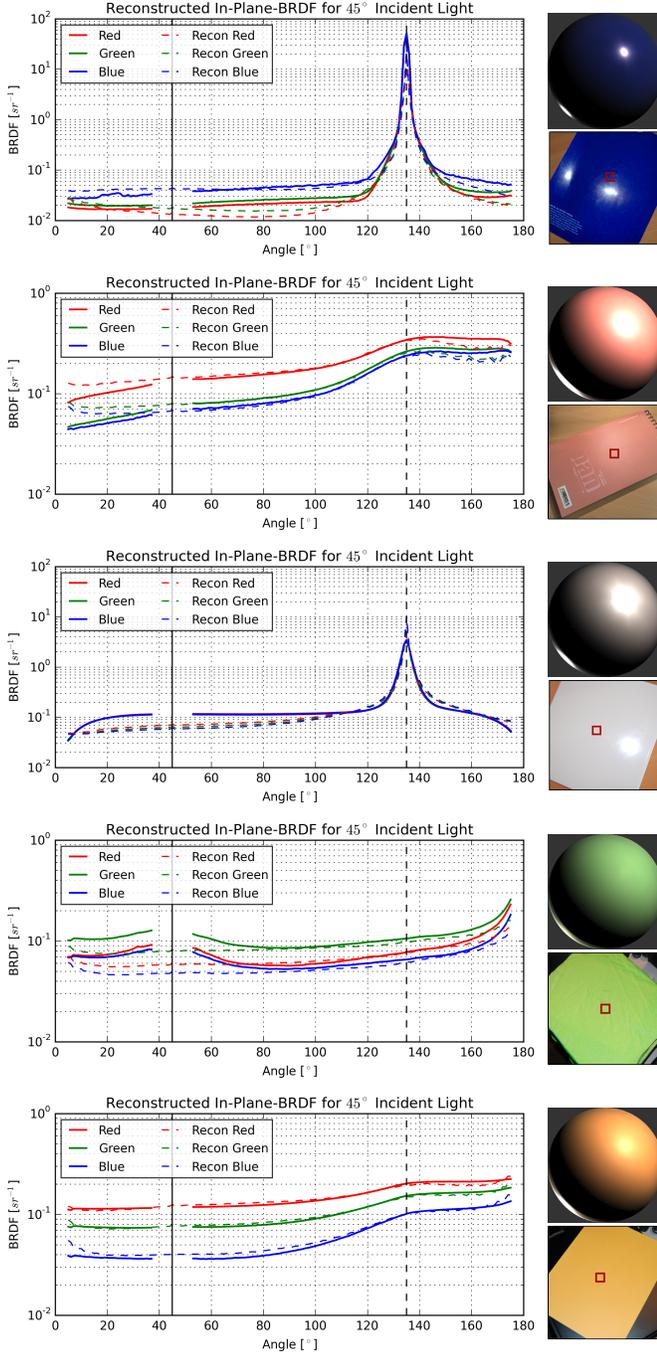
Figure 10: *Reconstruction of BRDF from five unknown materials: A glossy dark-blue book, a soft-specular notebook, a binder-cover, a piece of green cloth, and a piece of yellow paper. Lower right: reference photos of the measured materials. Upper right: renderings of reconstructed BRDFs with a front light at $[1,1,1]$ and a back light at $[-1,-1,-3]$. Left: Comparison between the measured in-plane BRDF values (solid) and the reconstructed in-plane BRDF values (dashed). Incident light is at $45°$ (marked as solid vertical line) and perfect reflection at $135°$ (marked with dashed vertical line). Note that these plots represent a single curve extracted from the fully reconstructed 3D BRDF volume.*

n	1	2		5				
$\theta_d[°]$	5	4	70	4	10	31	68	75

Table 2: *Optimal sampling directions in Rusinkiewicz coordinates [Rusinkiewicz 1998], when n images are acquired by imaging a sphere. For each acquired image, all combinations of (θ_h, ϕ_d) are captured. Thus, only θ_d needs to be varied.*

splitting up the database into "soft" and "specular" materials, and using the respective principal components $\boldsymbol{Q}_{\mathrm{soft}}$, and $\boldsymbol{Q}_{\mathrm{specular}}$ to reconstruct a material. Determining if a material is specular is easily done by inspecting the magnitude of the ratio between an in-peak sample and an out-of-peak sample. This refinement procedure is not strictly needed, but does in some cases improve reconstructions. An example of this is shown in the two bottom rows of Fig. 2. Here, the reconstruction of the diffuse "white-fabric" BRDF has introduced a ringing around the specular highlight, and even at 20 samples the ringing persists. By using "soft" PCs the artifact is effectively removed, producing convincing results with as few as 3 BRDF measurements.

## 9 Extension to Image-Based BRDF Capture

An effective way of capturing multiple BRDF samples per image is from spheres of a homogeneous material [Marschner et al. 2000]. This is how the extensive MERL database was captured [Matusik et al. 2003b], and is an approach that is often used today. This of course puts a constraint on the types of BRDFs possible to capture, as not all materials can be cut or molded into perfect spheres.

For any angle $\theta_d \leq 90°$ between camera and lightsource, an image of a perfect sphere covers all surface normal orientations in the positive hemisphere. An interesting observation is that an image of an illuminated sphere (with $\theta_d \leq 90°$) actually corresponds to a 2D slice of the 3D Rusinkiewicz coordinate frame (θ_h, ϕ_d) at θ_d. Hence, to capture the full 3D isotropic BRDF, only a sweep over θ_d must be made. This leads to a natural extension of our work, namely in determining the best n slices through the BRDF-volume, corresponding to the best n angles between camera and illumination when capturing a BRDF from an image of a spherical material.

We modify the optimization algorithm in Sec. 4. Where a measurement in the point-sampling setup corresponds to a single location in the BRDF volume and a single row in $\boldsymbol{Q}$, a measurement in a spherical-sampling setup corresponds to many rows in $\boldsymbol{Q}$. For a given angle θ_d, let $\mathcal{L}(\theta_d)$ be the set of BRDF locations, $(\theta_h, \theta_d, \phi_d)$, visible on the sphere. $\tilde{\boldsymbol{Q}}$ will now include the set of rows in $\boldsymbol{Q}$ corresponding to $\mathcal{L}(\theta_d)$ for each of the n measurements. In this case, the condition number κ depends only on the θ_d values, so that $\nabla\kappa = \frac{\delta\kappa(\tilde{\boldsymbol{Q}})}{\delta\theta_d}$. Using the new $\tilde{\boldsymbol{Q}}$ and $\nabla\kappa$, the optimization in Sec. 4 is again used to find the optimal sampling directions θ_d.

Our results are presented in Table 2 and visualized in Fig. 11. It is seen that the single most important sample is of the oblique angle reflection of the material (low θ_d). The second most important sample is the grazing angle reflection (high θ_d). For additional samples, we observe a spreading over all angles, but with a predominance at low and high values of θ_d.

As with the point-sampling setup, we qualitatively evaluate the reconstructions of the 10 test materials. The results are shown in Fig. 12 for $n = \{1, 2, 5\}$ sampling directions. After only two images, appearance is accurately captured for all materials. A quantitative analysis (red error graph in Fig. 9) confirms this, and shows that almost optimal reconstruction is achieved with 5 images. In agreement with Table 2, the first image captures the correct oblique angle appearance (first column of Fig. 12) and the second image captures the correct grazing angle appearance (second column). We
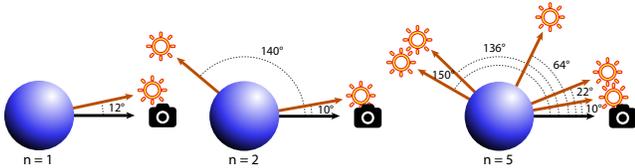
Figure 11: *Optimal sampling directions for spheres with $n = \{1, 2, 5\}$ images. Black arrows denote the direction to the camera. Orange arrows denote the direction to the light-source (due to reciprocity these may be interchanged). Notice that θ_d is the angle to the half-vector; the angle between view and light is $2\theta_d$.*

also performed a comparison to randomly chosen sampling directions. Intuitively, with more samples available from an image, the difference in error should be lower than for point sampling. We observe that random sampling results in roughly 35% higher errors for up to $n = 5$, after which it converges, approaching our error, again indicating again that 5 images are sufficient for reconstruction.

For a final validation we used the Cornell Reflectance database [Foo 1997] to simulate additional θ_d slices completely uncorrelated with the MERL data. These are shown in the last 3 rows.

10 Conclusions and Future Work

In this paper, we have developed a method for optimal, minimal sampling of BRDFs. Perhaps surprisingly, we show that $n = 20$ individual measurements is adequate in most cases for accurate isotropic BRDF reconstuction, and $n = 2$ images suffices for image-based BRDF measurements of spherical samples.

Our method leverages the MERL database [Matusik et al. 2003a], and proposes a novel mapping of BRDFs, allowing extraction of very descriptive principal components. A reconstruction approach based on ridge-regression, that utilizes the learned principal components, is described. We believe we are the first to present a BRDF reconstruction method that utilizes the statistical likelihood of a synthesized BRDF through the magnitude of the eigenvalue-scaled principal components. This approach yields better reconstructions than previous methods, and may have broader applicability.

We also develop a method for determining the regions of greatest importance, to sample the BRDFs. We provide an explicit table of the strategically best $n = \{1, 2, 5, 10, 20\}$ directions for sampling any unknown BRDF, and validate our results against previous industry standard sets of directions, as well as with measurements on new BRDFs not in the MERL database. We also show how the approach can be extended to sampling spheres via image-based BRDF measurement. In that case, two images often suffice.

In the future, the approach could also be extended to other acquisition geometries and customized for a particular gantry setup and near-field views of a homogeneous flat sample, where multiple light-view directions are available in a single image. Another interesting extension is to allow for capturing spatially varying BRDFs, since most gonioreflectometers are equipped with cameras. Our prioritized lists could also be used for importance sampling general BRDFs for rendering. The descriptive principal components may also be useful in editing and synthesizing novel BRDFs. In summary, we believe our results take an important step towards making rapid acquisition of data-driven reflectance models more practical in many applications.

Learned principal components, lists for $n = [1, 50]$ sampling directions, sample code, and reconstructed BRDFs, can be downloaded from our website: http://brdf.compute.dtu.dk.

Material	$n = 1$	$n = 2$	$n = 5$	Reference
black-soft-plastic				
blue-acrylic				
blue-metallic-paint2				
green-fabric				
light-red-paint				
pink-jasper				
silver-metallic-paint				
specular-violet-phenolic				
two-layer-silver				
white-fabric				
cayman [Cornell]				
garnet-red [Cornell]				
krylon-blue [Cornell]				

Figure 12: *Reconstructions of test samples, simulating BRDF capture using a sphere [Marschner et al. 2000]. The BRDFs are rendered as spheres, illuminated by a front light at a direction of $[1, 1, 1]$, and a back light causing grazing angle reflections at $[-1, -1, -3]$. Reconstructions are made with $n = \{1, 2, 5\}$ sampling directions. The far right column shows reference renderings of the true BRDFs. We see that two measurements (images) are sufficient to capture the true appearance of a material. The bottom 3 rows show materials from the Cornell Database [Foo 1997].*

Acknowledgements

References

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

BLANZ, V., MEHL, A., VETTER, T., AND PETER SEIDEL, H. 2004. A Statistical Method for Robust 3D Surface Reconstruction from Sparse Data. In *3D Data Processing Visualization and Transmission*, 293–300.

BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. In *SIGGRAPH 77*, 192–198.

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, 114.

BURLEY, B. 2012. Physically based shading at Disney. In *Practical physically-based shading in film and game production*, S. Hill and S. McAuley, Eds., ACM SIGGRAPH Courses. Article 10.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG) 1*, 1, 7–24.

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.

DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH 98*, 189–198.

FOO, S. C. 1997. *A gonioreflectometer 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.

HULLIN, M. B., HANIKA, J., AJDIN, B., SEIDEL, H.-P., KAUTZ, J., AND LENSCH, H. 2010. Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Transactions on Graphics (TOG) 29*, 4, 97.

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.

IPSEN, I. C. F., AND WENTWORTH, T. 2014. The effect of coherence on sampling from matrices with orthonormal columns, and preconditioned least squares problems. *SIAM J. Matrix Analysis Applications 35*, 4, 1490–1520.

LAFORTUNE, E., FOO, S., TORRANCE, K., AND GREENBERG, D. 1997. Non-Linear Approximation of Reflectance Functions. In *SIGGRAPH 97*, 117–126.

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.

LÖW, 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, 9.

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. In *ACM Transactions on Graphics (TOG)*, 759–769.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. Efficient isotropic BRDF measurement. In *Eurographics Rendering Workshop*, 241–247.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Eurographics Symposium on Rendering*, 117–126.

NGAN, A., DURAND, F., AND MATUSIK, W. 2006. Image-driven navigation of analytical BRDF models. In *Eurographics Symposium on Rendering*, 399–407.

NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. *Geometric Considerations and Nomenclature for Reflectance (NBS Monograph 160)*. National Bureau of Standards (US).

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.

PHONG, B. T. 1975. Illumination for computer generated pictures. *Communications of the ACM 18*, 6, 311–317.

REN, P., WANG, J., SNYDER, J., TONG, X., AND GUO, B. 2011. Pocket reflectometry. *ACM Transactions on Graphics (TOG) 30*, 4, 45.

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. on Computer Vision*, 859–872.

RUSINKIEWICZ, S. 1998. A new change of variables for efficient BRDF representation. In *Eurographics Rendering Workshop*, 11–22.

SCHWARTZ, C., SARLETTE, R., WEINMANN, M., RUMP, M., AND KLEIN, R. 2014. Design and implementation of practical btf measurement devices focusing on the developments at the university of Bonn. *Sensors 14*, 5, 7753–7819.

TIBSHIRANI, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological) 58*, 267–288.

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.

WEYRICH, T., MATUSIK, W., PFISTER, H., BICKEL, B., DONNER, C., TU, C., MCANDLESS, J., LEE, J., NGAN, A., JENSEN, H. W., AND GROSS, M. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics (TOG) 25*, 3, 1013–1024.

WHITE, D., SAUNDERS, P., BONSEY, S. J., VAN DE VEN, J., AND EDGAR, H. 1998. Reflectometer for measuring the bidirectional reflectance of rough surfaces. *Applied optics 37*, 16, 3450–3454.