


# SVBRDF Bootstrapping

Jiaping Wang\* Yue Dong†\* Xin Tong\* John Snyder‡ Moshe Ben-Ezra\* Yanxiang Lan† Baining Guo\*†  
\*Microsoft Research Asia †Tsinghua Univeristy ‡Microsoft Research



**Figure 1:** Renderings of high resolution SVBRDFs reconstructed by our method from low dimensional data captured and processed within minutes. (a,b) Anisotropic brushed metal. (c,d) Satin with complex needlework.

## Abstract

SVBRDF bootstrapping is a new method for data-driven modeling of real-world, spatially-varying reflectance that provides a high resolution result in both the spatial and angular domains. It decomposes reflectance measurement into two phases. The first acquires *representatives* of high angular dimension but sampled sparsely over the surface, while the second acquires *keys* of low angular dimension but sampled densely over the surface. Because reflectance is spatially coherent, keys can be used as feature vectors to “unlock” or interpolate between representatives, using a linear combination of a small number of representatives ( $k=15$ ) to match each key.

We develop a hand-held, high-speed BRDF capturing device for measurements in the first phase. A condenser-based optical setup collects rays emanating from a single point on the target sample over a dense hemisphere of outgoing directions. The device includes 6 LED light sources and captures the sample’s response to each light direction in sequence, yielding 10 BRDF point measurements per second. We then scan this device over the target to capture representatives at different, indiscriminate points. Sparse light measurements are amplified to a high-resolution 4D BRDF using a general microfacet model.

The second phase captures images of the entire sample from a fixed view. Lighting is varied by scanning a point source over the sample, yielding  $N=20-60$  images from different light directions. Arbitrary background lighting in the environment may also be present and is measured and accounted for. The key measurement thus represents the material’s response to  $N$  key lighting conditions at a dense set of key points. We show that the resulting,  $N$ -dimensional response captures much of the distance information in the original BRDF space. At each key point, the same linear weighting computed to match the measured key response is applied to the full 4D representatives, yielding a high-resolution SVBRDF. A few minutes of capture on a simple device yields sharp and anisotropic specularly and rich spatial detail.

## 1 Introduction

Real-world materials exhibit rich and detailed reflectance variation. Modeling or acquiring it is essential for realistic CG imagery. Surface reflectance is represented by the six-dimensional spatially

varying bidirectional reflectance distribution function (SVBRDF)  $\rho(\mathbf{x}, \mathbf{i}, \mathbf{o})$  [Nicodemus et al. 1977], describing how much radiance reflects at each surface point  $\mathbf{x}$  when viewed from direction  $\mathbf{o}$  and lit from direction  $\mathbf{i}$ . The parameter  $\mathbf{x}$  represents the *spatial* domain component of the SVBRDF while  $\mathbf{i} \times \mathbf{o}$  represents its *angular* domain component. Capturing a SVBRDF having high resolution in both domain components is difficult, and generally requires hours of measurement and processing, and use of large, specialized and expensive hardware rigs. As a result, only a few materials have been measured. Our goal is to make SVBRDF measurement available to more applications, by speeding up acquisition and doing it on simple, inexpensive devices.

Reducing the dimensionality of data capture is crucial to this goal. Our approach assumes that reflectance is spatially coherent; in particular, that its variation lies in a low-dimensional subspace over a small set of representatives. We therefore separate acquisition into two, lower-dimensional phases. The first measures a spatially sparse set of BRDF *representatives* of high angular dimension, and ignores their spatial location. The second densely maps over the surface a set of reflectance measurements of low angular dimension, called *keys*. The overall structure of the BRDF manifold for a given material is captured in the first phase, while the second phase determines the position on that manifold occupied at each  $\mathbf{x}$ .

Brute-force 4D capture is a time-consuming task even for a single representative BRDF, and still more so for a set of them. We obtain high-resolution data from a single camera position using a compact dual condenser setup that collects rays exiting in all directions from a single point on the material sample. We can quickly capture hundreds of representative points in a few minutes, by scanning this device over the sample. Reflectance is measured with respect to 6 different lighting directions. We amplify this sparse lighting measurement by using the BRDF microfacet model to infer a normal distribution function (NDF) [Ngan et al. 2005; Wang et al. 2008]. For anisotropic BRDFs, we further enlarge the set of representatives by a discrete series of azimuthal rotations on the derived NDFs. Matching against this enlarged set allows us to later recover an unknown local rotation of the target surface.

The challenge of the second phase is to minimize the number of key measurements that still discriminate adequately between high-resolution BRDFs. We scan an LED light source over the sample and capture  $N$  images of it from a single view. This produces an  $N$ -dimensional *key vector* at each point  $\mathbf{x}$ , each component of

\*email: {jiapw,xtong,johnsny,mosheb,bainguo}@microsoft.com

which represents the material’s response to a different *key lighting measurement*. Surprisingly, we find that a fairly small  $N$  (20-60) preserves distances quite well from the much higher-dimensional BRDF space. In other words, representatives having a similar response to the low-dimensional key lighting measurements are also close in the full BRDF space, while those having a different response are more distant in BRDF space.

To *bootstrap*, for each spatial location  $\mathbf{x}$ , we find a linear combination of *neighbor* representatives whose response matches the measured key vector at  $\mathbf{x}$ . The simulated response of each representative to the key lighting measurements is called its  $N$ -dimensional *matching vector*. Here the difficulty is in choosing neighbors and their weights so as to avoid synthesizing a physically unrealizable and visually implausible BRDF [Matusik et al. 2003]. Distance preservation in our key measurements implies that neighbors in the space of matching vectors are also neighbors in the original BRDF space and thus safe to interpolate. By choosing only a few (e.g., 15) neighbors and computing weights based on local linear embedding [Roweis and Saul 2000]), we ensure the linear combination produces a plausible result.

Our method is the first to capture high resolution, spatially-varying surface reflectance from sparse, two-phase captured data. Previous work [Lensch et al. 2003; Lawrence et al. 2006] has also exploited the spatial coherence of reflectance, typically by fitting special-purpose parametric models. Our work handles more general materials by pushing this exploitation all the way into the data acquisition, and using the ideas of local linear embedding and data bootstrapping. This greatly speeds up the acquisition process and makes measurement of new materials quick and easy. We demonstrate compelling visual accuracy for acquired materials.


## 2 Related Work

**SVBRDF Direct Acquisition** Dana et al. directly measured SVBRDFs using a spatial gonioreflectometer [1999]. Systems have since been developed to capture SVBRDFs [McAllister et al. 2002; Gu et al. 2006], and related BTFs [Dana 2001; Muller et al. 2005] and reflectance fields [Garg et al. 2006]. These methods densely sample the angular domain of view and light directions as well as the spatial domain. In all cases, special rigs are needed and the 6D datasets are huge and require hours to collect and process.


**SVBRDF Data-Driven Modeling** A number of approaches capture surface reflectance from a sparse set of views [Lensch et al. 2003; Gardner et al. 2003] and fit parametric BRDF models at each point. These methods capture spatial variation in reflectance, but the simple parametric models on which they rely do not accurately capture angular variation of real-world materials. SVBRDFs can also be captured from sparse data by constraining the type of material, such as to human skin [Debevec et al. 2000].

Another class of methods exploits the spatial redundancy of reflectance by merging data from surface points that are nearby [Zickler et al. 2005] or have similar reflectance properties (similar partial NDFs) [Wang et al. 2008]. We use the idea in [Wang et al. 2008] to represent BRDFs using a general microfacet model for the data amplification in phase one. But no previous work has yet exploited our general bootstrapping idea of capturing an SVBRDF in two separate phases: one to capture the structure of the BRDF manifold with high angular resolution, and another to map this structure over the sample with high spatial resolution.

Our approach handles general materials with sparse measurements. It avoids “hallucinating” data without verifying it by some measurement, such as merging partial NDFs from other surface points that happen to agree in their area of overlap. Most importantly, our method produces SVBRDFs that have much higher resolution in both the spatial and angular domains than previous data-driven modeling approaches. Yet the inner loop over each spatial point does only simple processing (finding a few neighbors



**Figure 2:** Single point BRDF measurement device (phase 1): (a) optical design, (b/c) prototype from side/bottom view.



**Figure 3:** Device setup for capturing reflectance maps (phase 2): (a) diagram, (b) photograph.

in matching vector space and solving a small linear system) rather than expensive model fitting using nonlinear optimization. The result is that total capture and processing takes only about 20 minutes for each new material, compared to up to 21 hours in [Wang et al. 2008].


## 3 SVBRDF Data Acquisition

Our approach captures two data from a flat sample of the target material. Typical sample dimension is  $10\text{cm} \times 10\text{cm}$ . Since our measuring devices are portable and we can handle natural background lighting, materials can be captured on-site and need not be moved to a dedicated capture room.

### 3.1 Acquiring Representatives: BRDF Samples

We have developed a portable device for capturing a hemispherical field of reflected rays emanating from a single point on the material sample using a single camera position. Data is acquired by illuminating the surface point using  $n_l = 6$  lighting directions and capturing its resulting reflectance. We scan the sample to acquire about 1000 BRDFs sampled randomly over its surface. A high-resolution BRDF is then derived from this captured data at each point.

**Device Setup** Figure 2 shows the design of our single-point BRDF measurement device. Our setup includes a pair of Anchor Optics 47mm condenser lenses with 21mm focus length, a  $200\mu\text{m}$  pinhole and a Firefly<sup>(R)</sup> MV camera from Point Grey Research. These components are mounted along the same optical axis using a lens tube from Thorlabs. We use six high brightness LEDs as light sources; each is attached to a carbon fiber tube to generate a light beam. One (top light beam) is mounted between the two condenser lenses and illuminates the capturing point at roughly a 10 degree bias from the vertical direction. The other five (side light beams) are mounted around the optical axis between the field condenser lens and the target surface, and illuminate the capturing point at exactly 20 degrees from vertical. A flat sample is placed at the focal plane of the field condenser lens,  $f_f$ . The pinhole is placed at the focal plane of the ocular condenser lens,  $f_o$ , and images the light field at a single point on the target surface onto a video camera. The camera communicates with a laptop via an IEEE1394 cable, which also supplies power for the LEDs and their control unit. A housing ensures the device is at the correct distance from the target sample.



**Figure 4:** NDF reconstruction: (a) 2D BRDF slice captured using top light beam, (b) using side light beam, (c) reconstructed NDF, (d) covered region in the reconstructed NDF. BRDF slices from these six lighting directions cover most of the NDF domain.

**Calibration** The lens tube ensures optic alignment of the lenses, pinhole and the camera. Distances between them are manually adjusted. The LED for the top light beam is held by an acrylic disc; its position is calibrated by measuring a mirror. Positions of the side LEDs are calibrated in manufacture. We calibrate the color and intensity of each LED by measuring a color checker pattern. Radial distortion of the dual condenser system is analytically calculated based on the specification from Anchor Optics, and determines the view direction at each pixel in the captured image.

**Capturing** The device is basically a cylinder 50mm in diameter and 150mm tall, and weighs around 500g. We scan the device over the sample to collect BRDF samples at different locations. For each position, we acquire six images lit by each LED and two images per light for exposure bracketing. The camera captures images of resolution  $320 \times 240$  at 135Hz, which finally captures 10 BRDFs per second. In a postprocess, each exposure pair is merged into an HDR image [Debevec and Malik 1997], and the resulting 6 images of  $240 \times 240$  used to derive a high-resolution BRDF. Figure 4a shows an example.

The top light LED component occludes a 3mm diameter hole in the captured image. Since the top light beam is biased from the optic axis, this hole does not occlude the peak of the specular lobe. We obtain the occlusion mask when calibrating with the color checker. If the hole contains no high frequency features, we fill it with harmonic interpolation [Schuster 2001]. We detect this by querying the intensity range of pixels surrounding the hole and testing whether the max/min ratio exceeds 2. In that case, we discard the sample.

Reflectance samples are then computed from the six HDR images by removing the clamped cosine factor


$$\rho(\mathbf{o}(\mathbf{u}), \mathbf{i}_l) = \frac{G_l(\mathbf{u})}{(\mathbf{n} \cdot \mathbf{i}_l) L_l} \quad (1)$$

where  $\mathbf{u}$  is the pixel position in the image corresponding to the view direction  $\mathbf{o}(\mathbf{u})$ , and  $\mathbf{i}_l$  and  $L_l$  are the direction and intensity of the  $l$ -th LED. These quantities are all determined in calibration.

**Reconstruction** To reconstruct a high resolution 4D BRDF from this captured data, we use a strategy similar to example-based microfacet synthesis [Wang et al. 2008]. There are two main differences. In our case, the view direction varies densely rather than the lighting direction. And we reconstruct the full NDF from partial NDFs inferred using a sparse set of  $n_l$  lighting directions, rather than using data from a single view but at other surface points. We represent the recovered NDF by a  $400 \times 400$  square image using the spherical parameterization in [Shirley and Chiu 1997]. Figure 4 summarizes the process and shows an example. The reconstruction algorithm is fully described in the appendix.

We sample each representative at  $N_b = N_o \times N_l$  samples,  $N_o$  for viewing direction and  $N_l$  for lighting direction, and pack the result into a *representative vector* denoted  $\mathbf{b}_p^*$ . The resulting set of  $M$  representatives is indexed by  $p$  and denoted  $\mathcal{B}^* = \{\mathbf{b}_p^*(\mathbf{o}, \mathbf{i}) \mid p = 1, 2, \dots, M\}$ .

**Representative Set Enlargement** To handle anisotropic materials, we rotate the derived NDF around the normal direction by



**Figure 5:** SVBRDF bootstrapping. A key vector at each spatial position  $\mathbf{x}$  is projected to the space of matching vectors to determine a local linear embedding. The linear weights and neighborhood indices are then applied the full-resolution representatives to reconstruct the BRDF at  $\mathbf{x}$ .

a discrete set of 360 azimuthal angles and add the corresponding BRDFs to the example set. This set is thus enlarged 360 times. The enlargement procedure also allows us to recover the local orientation at each pixel, by applying the local linear interpolation weights  $w_p$  in equation (5) to the local tangent orientation of each neighborhood representative (see Section 4.1).

### 3.2 Acquiring Keys: Reflectance Maps

Keys are based on reflectance maps captured from a single view and lit by  $N$  different lighting configurations, as illustrated in Figure 3. These lighting configurations can include fixed or variable area/environmental sources and moving point sources. This allows on-site capture of samples since our rig is portable and we may not be able to control the background lighting. A mirror ball is used to probe the lighting applied. A Canon 30D camera with EF-100 2.8 lens is placed over and 2.0m away from the center of the material sample. Image resolution is  $3504 \times 2336$ .

Before capturing, we calibrate the camera’s position and orientation using the method in [Zhang 2000]. For each lighting change, we record an HDR image including the material and the mirror ball using exposure bracketing as in [Debevec and Malik 1997]. In our prototype system, we simply move a single lamp around and at various distances from the sample by hand.

The process is finished after capturing  $N$  images, resulting in the material’s reflectance responses,  $r_j(\mathbf{x})$ , and reconstructed source radiance fields,  $L_j(\mathbf{i})$ , for  $j \in 1, 2, \dots, N$ . The point light source is far enough away to reasonably assume that the radiance field is constant over the entire material sample. We also compute the viewing direction  $\mathbf{o}^*$  at the sample center and assume it is constant over all  $\mathbf{x}$  as well. We finally assemble all  $N$  reflectance responses at each key point  $\mathbf{x}$  into an  $N$ -dimensional *key vector*,  $\mathbf{r}_\mathbf{x} = (r_1(\mathbf{x}), r_2(\mathbf{x}), \dots, r_N(\mathbf{x}))^T$ .

## 4 SVBRDF Bootstrapping

Given the  $M$  high-resolution representatives BRDF  $\mathbf{b}_p^*$  and the  $N$ -dimensional keys  $\mathbf{r}_\mathbf{x}$ , both acquired from the same material sample, we bootstrap to reconstruct a high resolution SVBRDF, using a linear combination of representatives to match each key (Figure 5).

Arbitrary linear combinations of different BRDFs can produce physically implausible reflectance [Matusik et al. 2003]. Though the space of physical BRDFs is not globally linear, we assume it forms a smooth manifold and is therefore locally linear. Our reconstruction employs a convex linear combination over a *representative neighborhood* for each surface point  $\mathbf{x}$ , containing a small number  $k$  of representatives that are all close to each other.

We choose  $k$  in the range 15-20 based on experiments in [Matusik et al. 2003], which observed a local dimensionality for BRDFs of around 16.

### 4.1 Basic Idea

As described in Section 3.2, keys are captured at each point  $\mathbf{x}$  using measurements from a single view  $\mathbf{o}^*$  and  $N$  different source radiance fields  $L_j$ . These provide constraints on integrals of the BRDFs



$\mathbf{b}_x(\mathbf{i}, \mathbf{o})$ , via

$$\mathbf{r}_j(\mathbf{x}) = \int_{\Omega_x(\mathbf{n})} \mathbf{b}_x(\mathbf{i}, \mathbf{o}) (\mathbf{n} \cdot \mathbf{i}) L_j(\mathbf{i}) d\mathbf{i}. \quad (2)$$

This equation can be represented in matrix form as

$$\mathbf{r}_x = \mathbf{R} \mathbf{b}_x, \quad (3)$$

where  $\mathbf{r}_x$  is the measured key vector, and  $\mathbf{b}_x$  is its corresponding and unknown BRDF vector. The  $N \times N_b$  key measurement matrix  $\mathbf{R}$  is given by

$$\mathbf{R}_{jk} = \begin{cases} (\mathbf{n} \cdot \mathbf{i}_{k_i}) L_j(\mathbf{i}_{k_i}), & k_{\mathbf{o}} = k^* \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

Here,  $k_i = k \bmod N_i$  indexes the lighting direction,  $k_{\mathbf{o}} = (k - k_i) / N_i$  indexes the view direction in the packed representative vector  $\mathbf{b}^*$ . Recall that  $N_i$  is the number of lighting directions used in the representative vector; see the last paragraph of Section 3.1.  $k^*$  indexes the particular view direction  $\mathbf{o}^*$  used for all key measurements.

We represent the BRDF  $\mathbf{b}_x$  as a convex linear combination over the representative neighborhood,  $\mathbf{b}_x^*$ ,  $p \in \delta(\mathbf{x})$ :

$$\mathbf{b}_x \approx \sum_{p \in \delta(\mathbf{x})} w_p(\mathbf{x}) \mathbf{b}_p^*, \quad \sum_{p \in \delta(\mathbf{x})} w_p(\mathbf{x}) = 1, \quad (5)$$

Substituting this BRDF approximation into (3), we obtain a constrained linear equation on the weights  $w_p(\mathbf{x})$ :

$$\mathbf{r}_x = \sum_{p \in \delta(\mathbf{x})} w_p(\mathbf{x}) \mathbf{r}_p^*, \quad \sum_{p \in \delta(\mathbf{x})} w_p(\mathbf{x}) = 1 \quad (6)$$

where

$$\mathbf{r}_p^* = \mathbf{R} \mathbf{b}_p^*. \quad (7)$$

The projection in (7) numerically applies the source lighting we captured in phase two to the representative BRDFs we captured in phase one, and also evaluates at the single key viewing direction  $\mathbf{o}^*$ . We call the projected representatives,  $\mathbf{r}_p^*$ , *matching vectors*. In other words, equations (3) and (5) imply that the measured key vector  $\mathbf{r}_x$  be represented as a linear combination of matching representative vectors,  $\mathbf{r}_p^*$ ,  $p \in \delta(\mathbf{x})$ .

## 4.2 Uniform Measurement Scaling

Use of area or environmental light sources in key measurement produces a non-orthogonal key measurement matrix. This leads to ellipsoidal rather than spherical neighborhoods in key space, and so complicates the selection of neighbors. We orthogonalize the projection by applying the singular value decomposition (SVD) to  $\mathbf{R}$ , yielding

$$\mathbf{R} = \mathbf{U} \Sigma \mathbf{V} \quad (8)$$

where  $\mathbf{U}$  is a  $N \times N$  orthogonal matrix of left-hand eigenvectors,  $\Sigma$  is a diagonal matrix of eigenvalues, and  $\mathbf{V}$  is an  $N \times N_b$  orthogonal matrix of transposes of the right-hand eigenvectors. We assume  $\Sigma$  contains no zero or very small elements. If it does, then we are measuring redundant (i.e., linearly dependent) lighting configurations, which add no new information to the key.

To remove non-uniform scaling in our key measurements, we apply the SVD in (8) to obtain the *uniform key vector*

$$\hat{\mathbf{r}}_x = \Sigma^{-1} \mathbf{U}^T \mathbf{r}_x. \quad (9)$$

We also define the *uniform matching vector*  $\hat{\mathbf{r}}_p^*$  as

$$\hat{\mathbf{r}}_p^* = \mathbf{V} \mathbf{b}_p^*. \quad (10)$$

Neighbors can now be found in the uniform key space using a simple distance threshold over these  $N$ -dimensional vectors.

## 4.3 Neighborhood Selection

The representative neighborhood  $\delta(\mathbf{x})$  is determined by finding the  $k$ -nearest uniform matching vectors  $\hat{\mathbf{r}}_p^*$  to the uniform key  $\hat{\mathbf{r}}_x$ . If the matrix  $\mathbf{V}$  doesn't collapse meaningful dimensions of the BRDFs, then neighborhoods in the space of uniform matching vectors  $\hat{\mathbf{r}}_p^*$  will also be neighborhoods (i.e., close to each other) in the full BRDF space  $\mathbf{b}_p^*$ . This ensures their linear combination will produce physically plausible BRDFs. We will validate this assumption in the next section.

We use approximated nearest neighbor (ANN) search [Mount and Arya 1997] to accelerate finding the  $k$ -nearest neighbors. We also remove outliers having distance more than 5 times of the average distance over all neighborhoods.

## 4.4 Local Linear Combination

We determine the linear weights,  $w_p(\mathbf{x})$ , based on the distance metric in each local neighborhood [Roweis and Saul 2000], via:

$$w_p = \sum_{q \in \delta(\mathbf{x})} \mathbf{C}_{pq}^{-1} (\hat{\mathbf{r}}_x \cdot \hat{\mathbf{r}}_q^* + \lambda), \quad (11)$$

$$\lambda = \frac{1 - \sum_{p,q \in \delta(\mathbf{x})} \mathbf{C}_{pq}^{-1} (\hat{\mathbf{r}}_x \cdot \hat{\mathbf{r}}_q^*)}{\sum_{p,q \in \delta(\mathbf{x})} \mathbf{C}_{pq}^{-1}}. \quad (12)$$

$\mathbf{C}_{pq} = \hat{\mathbf{r}}_p^* \cdot \hat{\mathbf{r}}_q^*$  denotes the covariance matrix of the neighborhood and  $\mathbf{C}^{-1}$  is its inverse. Occasionally, in the case of a nearly singular  $\mathbf{C}$ , we remove the most distant neighbor and solve (11) again.

## 4.5 Key Measurement Sufficiency

Key measurements amount to BRDF projections via (3) and (4). The orthogonalization of our key light configuration, represented by the matrix  $\mathbf{V}$ , must not collapse meaningful dimensions of the representative BRDFs  $\{\mathbf{b}_p^*\}$ . Intuitively, increasing  $N$  by adding more linearly independent lighting configurations should help.

We can quantify the discriminative power of our key measurements by considering how well distance measurements agree between the space of BRDF representatives and their corresponding  $N$ -dimensional matching vectors. We measure the overall distance preservation factor  $\tau$  over a neighborhood of representatives of radius  $r$ ,  $\delta(p, r) = \{q \mid \|\mathbf{b}_p^* - \mathbf{b}_q^*\| < r\}$ , via:

$$\tau(\delta) = \frac{\sum_{i,j \in \delta} \|\hat{\mathbf{r}}_i^* - \hat{\mathbf{r}}_j^*\|}{\sum_{i,j \in \delta} \|\mathbf{b}_i^* - \mathbf{b}_j^*\|}. \quad (13)$$

In the uniformly-scaled space, we have  $0 \leq \|\hat{\mathbf{r}}_i^* - \hat{\mathbf{r}}_j^*\| \leq \|\mathbf{b}_i^* - \mathbf{b}_j^*\|$ . The closer  $\tau$  is to 1, the better our key measurement is at discriminating between representatives in the neighborhood. We can then examine the average local distance preservation for a given radius:

$$\bar{\tau}(r) = 1/M \sum_{p=1}^M \tau(\delta(p, r)). \quad (14)$$

Finally, we define the *global distance preservation factor*,  $\tau_g = \bar{\tau}(\infty)$ ; i.e., by calculating distance over all pairs of representatives. We also define the *local distance preservation factor*,  $\tau_l = \bar{\tau}(\bar{r})$  where  $\bar{r}$  is the average radius of the  $k$  nearest neighbors over all representative BRDFs.

To investigate the sufficiency of key light conditions, we captured 100 lighting conditions and randomly selected  $N$  as input to generate the matching vector space of 1200 BRDFs sampled from the example in Figure 7a. Figure 6 plots the results as a function of  $N$ . Both global and local distance preservation factors converge quickly as  $N$  increases. In our experiments, convincing results are obtained with  $\tau_g > 0.9$  and  $\tau_l > 0.85$ . Evaluating  $\tau_l$  and  $\tau_g$  at a few hundred representatives takes little time and indicates whether our lighting configurations and value for  $N$  are sufficient. So it is useful for guiding the on-site capturing procedure in phase 2.




Figure 6: Distance preservation of the matching vector space.

Material Sample	Reflectance Maps		Representative BRDFs	
	Res.	Num.	Num.	$\tau_g / \tau_l$
wrapping paper	1000×1000	50	120	0.90/0.87
old copper	2000×2000	80	1200	0.93/0.85
aluminum pan	2000×2000	80	10 (3600*)	0.99/0.85
satin	1500×1500	80	30 (10800*)	0.93/0.85

Table 1: Statistics for our examples. The starred column denotes the size of the enlarged representative set that includes tilt/pan rotated NDF versions.

## 5 Experimental Results

We implemented our SVBRDF bootstrapping algorithm on a Intel Core™2 Duo 2.13G PC with 2G memory. Capturing the BRDF representatives and reflectance map keys takes 10-20 minutes. Subsequent data processing takes less than 10 minutes. Table 1 lists the parameters used in capturing. We infer 2D NDFs of resolution  $400 \times 400$ , yielding 4D BRDFs with more than 10 million angular samples in viewing and lighting direction. The spatial resolution ranges from 1 to 4 million samples.

To validate our method, we tested our bootstrapping scheme on fully sampled SVBRDF data (greeting card from [Lawrence et al. 2006]). We selected 1000 BRDFs from random positions as the representatives. We simulated capture of the reflectance map by applying the Grace Cathedral environment map [Debevec and Malik 1997] along with a point light source at a controlled direction. We then measured reconstruction error of our bootstrapping method as a function of the number  $N$  of different light directions. For each  $N$ , we average over many randomly generated sets of light directions. Figure 7a shows average reconstruction error, which falls quickly as  $N$  increases. The right two columns of the figure compare rendered results between the original data (b) and our reconstructed SVBRDF (c), at a view not used for phase 2 capture. An accurate match is obtained.

Figure 8 shows results for isotropic material samples with different kinds of reflectance. We compare the rendered result of our reconstructed SVBRDFs to photographs of the captured sample with the same lighting conditions. Materials with progressive (a/c) and sharp (b/d) spatial variation are both handled well.

Rendered results with captured SVBRDFs are shown in Figures 1, 9, and 10. Circular brushed metal is reconstructed with rich spatial detail in the brushing tracks. Satin with needlework also demonstrates spatially varying anisotropic reflectance.




Figure 7: Validation example. (a) Reconstruction error as a function of number of lighting measurements,  $N$ . (b) Rendering with original SVBRDF. (c) Rendering with reconstructed SVBRDF.




Figure 8: Captured isotropic SVBRDF examples. Old copper: (a) original, (c) reconstructed. Wrapping paper: (b) original, (d) reconstructed.


## 6 Conclusion

SVBRDF bootstrapping simplifies and accelerates the capture of complex reflectance. It does this by decomposing data acquisition into two phases: one to capture the overall BRDF manifold and another to find the location on this manifold for each point on the surface. Both phases make only sparse measurements of the overall 6D SVBRDF. The first is dense over view direction, but sparse in light direction and spatial position, while the second is dense over spatial position but sparse in lighting and uses only a single view. We propose a new compact device based on a pair of condenser lenses to capture BRDF point samples in the first phase. Using local linear embedding and representative set enlargement, we are able to produce SVBRDFs of high resolution in both the spatial and angular domains from this data. Materials captured with our technique exhibit convincing realism, isotropic and anisotropic specularity, and spatial detail.


The bootstrapping approach is general and can accommodate different methods for acquiring representative BRDFs in phase 1 and reflectance map keys in phase 2. Our hand-held BRDF sampling device measures only a few light directions, and so requires amplification via the (single-bounce) microfacet model. Though this model has wide applicability [Ashikhmin et al. 2000; Wang et al. 2008], it does prohibit some anomalous materials such as retro-reflective ones. Our method for acquiring reflectance maps is limited to flat surfaces without significant self-shadowing and self-masking. We would like to improve these drawbacks in future work. We are also interested in combining reflectance maps captured from multiple views to further improve the discriminative capacity of phase 2 measurements.

## References

- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Siggraph 2000, Computer Graphics Proceedings*, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 65–74.
- COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1, 7–24.



**Figure 9:** Rendered results with isotropic SVBRDFs. Image pairs show two different views and rotation of the sample with respect to the lighting environment.



**Figure 10:** Rendered results with anisotropic SVBRDFs.

- DANA, K. J., NAYAR, S. K., VAN GINNEKEN, B., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics* 18, 1, 1–34.
- DANA, K. J. 2001. BRDF/BTF measurement device. In *Proceedings of eighth IEEE international conference on computer vision (ICCV)*, vol. 2, 460–466.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH*, 369–378.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proc. SIGGRAPH 2000*, 145–156.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *ACM Trans. Graph.* 22, 3, 749–758.
- GARG, G., TALVALA, E.-V., LEVOY, M., AND LENSCH, H. P. A. 2006. Symmetric photography: exploiting data-sparseness in reflectance fields. In *Eurographics Workshop/Symposium on Rendering*, Eurographics Association, Nicosia, Cyprus, 251–262.
- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 43–54.
- GU, J., TU, C.-I., RAMAMOORTHY, R., BELHUMEUR, P., MATUSIK, W., AND NAYAR, S. 2006. Time-varying surface appearance: acquisition, modeling and rendering. *ACM Trans. Graph.* 25, 3, 762–771.
- LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (July).
- LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Transaction on Graphics* 22, 2 (Apr.), 234–257.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. Graph.* 22, 3, 759–769.
- MCALLISTER, D. K., LASTRA, A. A., AND HEIDRICH, W. 2002. Efficient rendering of spatial bi-directional reflectance distribution functions. In *Proceedings of the 17th Eurographics/SIGGRAPH workshop on graphics hardware (EGGH-02)*, ACM Press, New York, S. N. Spencer, Ed., 79–88.
- MOUNT, D., AND ARYA, S. 1997. Ann: A library for approximate nearest neighbor searching. In *CGC 2nd Annual Fall Workshop on Computational Geometry*.
- MÜLLER, G., MESETH, J., SÄTTLER, M., SARLETTE, R., AND KLEIN, R. 2005. Acquisition, synthesis, and rendering of bidirectional texture functions. *Computer Graphics Forum* 24, 1, 83–109.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. *Eurographics Symposium on Rendering 2005*.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. Geometric considerations and nomenclature for reflectance. *Monograph 161, National Bureau of Standards (US)*.
- PRESS, W. H., ET AL. 1992. Numerical recipes in C (second edition). *Cambridge University Press*.
- ROWEIS, S. T., AND SAUL, L. K. 2000. Nonlinear dimensionality reduction by locally linear embedding. In *Science*, 2323–2326.
- SCHUSTER, W. 2001. Harmonische interpolation. In *Math. Semesterber*, Springer-Verlag, 1–27.
- SHIRLEY, P., AND CHIU, K. 1997. A low distortion map between disk and square. *J. Graph. Tools* 2, 3, 45–52.

WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, ACM, New York, NY, USA, 1–9.

ZHANG, Z. 2000. A flexible new technique for camera calibration. In *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 22, 1330–1334.

ZICKLER, T., ENRIQUE, S., RAMAMOORTHY, R., AND BELHUMEUR, P. 2005. Reflectance sharing: image-based rendering from a sparse set of images. In *Eurographics Symposium on Rendering*, Eurographics Association, Konstanz, Germany, K. Bala and P. Dutré, Eds., 253–264.

## Appendix: BRDF Reconstruction using the Microfacet Model

This appendix describes how a high-resolution 4D BRDF can be derived from measurements over a sparse set of light directions and a dense set of views.

Reflectance is decomposed into diffuse and specular components. We determine the diffuse component  $\rho_d$  using a simple minimum filter on the samples  $\rho$  in (1), via

$$\rho_d = \frac{\sum_l \min_{\mathbf{u}} \{\rho(\mathbf{o}(\mathbf{u}), \mathbf{i}_l)\}}{n_l}. \quad (15)$$

The specular component is the residual after subtracting this diffuse term:

$$\rho_s(\mathbf{o}(\mathbf{u}), \mathbf{i}_l) = \rho(\mathbf{o}(\mathbf{u}), \mathbf{i}_l) - \rho_d. \quad (16)$$

We then represent the specular component with a general microfacet model [Ashikhmin et al. 2000] as

$$\rho_s(\mathbf{o}, \mathbf{i}) = c_s \frac{D(\mathbf{h}) S(\mathbf{i}) S(\mathbf{o}) F(\mathbf{o}, \mathbf{i})}{\pi(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})}, \quad (17)$$

This model is defined in terms of five factors: a microfacet normal distribution function (NDF)  $D$ , its shadowing factor  $S$ , a Fresnel reflection factor  $F$ , and the scalar specular coefficient,  $c_s$ . The NDF is defined over the hemisphere of half angle vectors  $\mathbf{h} = (\mathbf{o} + \mathbf{i}) / \|\mathbf{o} + \mathbf{i}\|$ . For a flat material sample, the surface normal is aligned to the  $z$  axis:  $\mathbf{n} = \mathbf{z} = (0, 0, 1)$ .

Since  $D$  dominates the other factors in determining the high-frequency characteristics of surface reflectance, we follow [Ashikhmin et al. 2000; Debevec et al. 2000; Wang et al. 2008] and tabulate it as a square image using the spherical parameterization in [Shirley and Chiu 1997]. The shadow factor  $S$  is derived from  $D$ , and the Fresnel factor  $F$  reduced to a single scalar parameter, as will be described later. We reconstruct the microfacet BRDF in (17) from the measured specular data  $\rho_s$  in (16) using two steps.

The first step fits  $D$ . To simplify the computation, we exclude reflectance data near the grazing angle and assume  $F(\mathbf{o}, \mathbf{i}) = 1$ . The remaining scattered BRDF measurements  $\rho_s(\mathbf{o}(\mathbf{u}), \mathbf{i}_l)$  for each lighting direction  $\mathbf{i}_l$  are interpolated with the push-pull method [Gortler et al. 1996], as in [Lawrence et al. 2006]. This yields  $n_l$  2D slices of the specular component  $\rho_s(\mathbf{o}, \mathbf{i}_l)$ , each uniformly sampled over the hemicycle of viewing directions  $\mathbf{o}$ . We then fit the partial NDF for each lighting direction  $\mathbf{i}_l$  using an iterative algorithm.

Starting from an initial shadow factor  $S(\mathbf{i}) = S(\mathbf{o}) = 1$ , the partial NDF is computed as:

$$D_l(\mathbf{h}) = \begin{cases} \frac{\pi \rho_s(\mathbf{o}_l(\mathbf{h}), \mathbf{i}_l) (\mathbf{o}_l(\mathbf{h}) \cdot \mathbf{n}) (\mathbf{i}_l \cdot \mathbf{n})}{S(\mathbf{o}_l(\mathbf{h})) S(\mathbf{i}_l)}, & \mathbf{h} \in \Omega_l \\ 0, & \mathbf{h} \notin \Omega_l \end{cases} \quad (18)$$

where  $\mathbf{o}_l(\mathbf{h}) = \text{refl}(\mathbf{i}_l, \mathbf{h}) = 2(\mathbf{i}_l \cdot \mathbf{h})\mathbf{h} - \mathbf{i}_l$ , and  $\Omega_l$  is the region of the hemisphere covered by half vectors between  $\mathbf{o}$  and  $\mathbf{i}_l$ . It is defined

as

$$\Omega_l = \left\{ \frac{\mathbf{o} + \mathbf{i}_l}{\|\mathbf{o} + \mathbf{i}_l\|} \mid \mathbf{o} \in \Omega_+(\mathbf{n}) \right\} = \{\mathbf{h} \mid \mathbf{o}_l(\mathbf{h}) \in \Omega_+(\mathbf{n})\}, \quad (19)$$

where the positive hemisphere around a vector  $\mathbf{v}$  is denoted  $\Omega_+(\mathbf{v}) = \{\mathbf{u} \mid \mathbf{u} \cdot \mathbf{v} \geq 0\}$ .

We then merge the six partial NDFs  $D_l$  and average them in overlapping regions to generate the completed NDF  $D$ . Values not covered by any  $\Omega_l$  are extrapolated via push-pull.

The shadowing factor,  $S(\mathbf{v})$ , can be derived from the completed NDF  $D$  via

$$S(\mathbf{v}) = \frac{(\mathbf{v} \cdot \mathbf{n})}{\int_{\Omega_+(\mathbf{v}) \cap \Omega_+(\mathbf{n})} (\mathbf{h} \cdot \mathbf{v}) D(\mathbf{h}) d\mathbf{h}}. \quad (20)$$

Given this shadow factor, we iteratively update the partial NDFs  $D_l$  using (18), merge them, and then derive  $S$  again using (20) until convergence.

The second step computes the remaining microfacet model parameters, given the completed NDF  $D$ . The specular coefficient is derived via  $c_s = \int_{\Omega_+(\mathbf{n})} (\mathbf{n} \cdot \mathbf{h}) D(\mathbf{h}) d\mathbf{h}$ . We then normalize the NDF by  $c_s$  and use it to compute the shadowing factor via (20). The Fresnel factor is based on [Cook and Torrance 1982], and defined in terms of a single parameter,  $\eta$ , representing the relative index of refraction:

$$F(\mathbf{o}, \mathbf{i}, \eta) = \frac{(g-c)^2}{2(g+c)^2} \left( 1 + \frac{(c(g+c)-1)^2}{(c(g-c)+1)^2} \right), \quad (21)$$

where  $g^2 = \eta^2 + c^2 - 1$  and  $c = |\mathbf{i} \cdot \mathbf{h}|$ . To fit this model, we minimize

$$E(\eta) = \int_{\Omega_+(\mathbf{n})} \|F(\mathbf{o}, \mathbf{i}, \eta) - F_m(\mathbf{o}, \mathbf{i})\|^2 d\omega. \quad (22)$$

This integration is over  $\omega$  representing the 6 lighting directions, and the image pixels representing view directions captured by the camera. using the Levenberg-Marquardt algorithm [Press et al. 1992].  $F_m(\mathbf{o}, \mathbf{i})$  is computed from the measured  $\rho_s(\mathbf{o}, \mathbf{i})$  as

$$F_m(\mathbf{o}, \mathbf{i}) = \frac{\pi \rho_s(\mathbf{o}, \mathbf{i}) (\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})}{S(\mathbf{i}) S(\mathbf{o}) c_s D(\mathbf{h})}. \quad (23)$$

The final representative BRDF combines the diffuse and specular components via  $\rho(\mathbf{o}, \mathbf{i}) = \rho_d + \rho_s(\mathbf{o}, \mathbf{i})$  where  $\rho_s$  is computed via (17).