

Modeling Anisotropic Surface Reflectance with Example-Based Microfacet Synthesis

Jiaping Wang* Shuang Zhao† Xin Tong* John Snyder‡ Baining Guo*

*Microsoft Research Asia

†Shanghai Jiaotong University

‡Microsoft Research

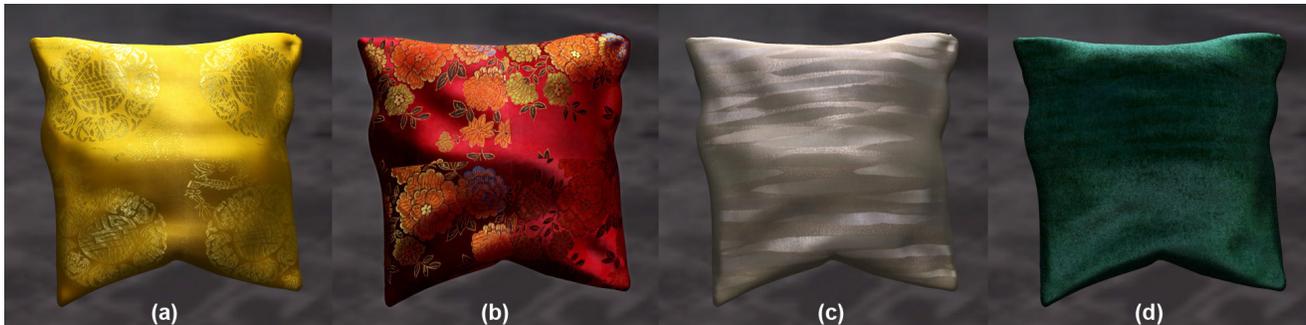


Figure 1: Fabric SVBRDFs from our algorithm mapped onto a pillow: (a) yellow satin, (b) red satin with colorful needlework, (c) wallpaper, (d) velvet.

Abstract

We present a new technique for the visual modeling of spatially-varying anisotropic reflectance using data captured from a single view. Reflectance is represented using a microfacet-based BRDF which tabulates the facets’ normal distribution (NDF) as a function of surface location. Data from a single view provides a 2D slice of the 4D BRDF at each surface point from which we fit a partial NDF. The fitted NDF is partial because the single view direction coupled with the set of light directions covers only a portion of the “half-angle” hemisphere. We complete the NDF at each point by applying a novel variant of texture synthesis using similar, overlapping partial NDFs from other points. Our similarity measure allows azimuthal rotation of partial NDFs, under the assumption that reflectance is spatially redundant but the local frame may be arbitrarily oriented. Our system includes a simple acquisition device that collects images over a 2D set of light directions by scanning a linear array of LEDs over a flat sample. Results demonstrate that our approach preserves spatial and directional BRDF details and generates a visually compelling match to measured materials.

1 Introduction

The physical world contains many complex materials whose reflectance properties must be modeled or captured to produce realistic CG imagery. Spatial variation and anisotropy are crucial and particularly challenging to reproduce. Such effects can be 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 light is reflected at each surface point

\mathbf{x} as the radiance ratio measured when viewed from direction \mathbf{o} and lit from direction \mathbf{i} . Current techniques have proven inadequate for obtaining visually compelling SVBRDFs.

The traditional approach represents the BRDF at each surface point using parametric models [Cook and Torrance 1981; Ward 1992; Lafortune et al. 1997; Rusinkiewicz 1998]. It is difficult to create realistic appearance by building such models *ab initio*. Simple models lack power to capture real-world materials. More complicated, multi-lobe models represent any BRDF to arbitrary accuracy but their many parameters are too hard to specify and control. This has led to extensive research work in BRDF measurement.

A number of systems measure 6D SVBRDFs directly from surface samples [Dana et al. 1999; McAllister et al. 2002; Lawrence et al. 2006]. However, densely scanning over light and view directions necessitates lengthy capture on expensive devices. It also requires precise calibration of multiple or moving cameras, adding to the cost and fragility of these systems. As a result, few materials have been measured in this way and little SVBRDF data is available. Another measurement approach fits parametric BRDF models at each point, given data from a sparse set of views [Lensch et al. 2003a; Gardner et al. 2003]. This simplifies the acquisition process, but fails to capture detailed, anisotropic reflectance. Simple parametric models do not capture details in the measured data, while multi-lobe models overfit the sparse view data and generalize poorly to uncaptured views.

We present a new technique for modeling SVBRDFs from images of a surface sample taken from a single view. Our approach represents anisotropic specularities at each point with a general microfacet BRDF model [Ashikmin et al. 2000], and tabulates the microfacets’ 2D normal distribution function (NDF). The result is a spatially-varying NDF (SVNDF). Data from a single view determines the NDF over only a partial region of the hemisphere. We generate the full NDF at each surface point by iteratively growing these partial NDFs using partial NDFs from other points. Each iteration searches candidates that overlap the current point’s partial NDF and are similar within the overlapping region. Since this procedure is similar to example-based texture synthesis [Efros and Freeman 2001; Kwatra et al. 2003], we call our method *Example-Based Microfacet Synthesis*.

Our basic assumption is that for any surface point we can find others having a similar but rotated microstructure (i.e., a different local frame). A single view thus yields different slices of the BRDF at points sharing similar reflectance, equivalent to partial NDFs de-

*Email: {jpwang, xtong, johnsny, bainguo}@microsoft.com

†Shuang Zhao was a visiting student at Microsoft Research Asia.

fined over different subregions of the hemisphere. By converting slices of 4D BRDFs to 2D partial NDFs, we are able to detect which surface points share rotated microstructure and merge them to complete their NDFs. Although we do not produce the original sample’s SVBRDF exactly, we do capture its statistical variation and rich appearance, as shown in Figure 1.

Our method represents a new application of ideas from texture synthesis to SVBRDF acquisition in order to exploit spatial redundancy. It is also the first sparse-view (actually single-view) acquisition method that preserves spatial and angular details. The sparsity of input data required by our method greatly simplifies acquisition and reduces processing costs. Results demonstrate good visual accuracy using data acquired by a simple and inexpensive device.

2 Related Work

Anisotropic BRDF models were introduced by Kajijiya[1985]. A number of empirical BRDF models have since been proposed, including Gaussian [Ward 1992] and generalized Phong [Lafortune et al. 1997; Ashikhmin and Shirley 2000]. While these models approximately capture anisotropic reflectance, they ignore the underlying microstructure and can miss details in many real-world materials [Ngan et al. 2005].

Another strategy is to explicitly model geometric microstructure. Physically-based anisotropic BRDF models have been developed for specialized materials such as finished wood [Marschner et al. 2005], hair [Marschner et al. 2003], and cloth [Yasuda et al. 1992; Irawan and Marschner 2006], but it is difficult to extend these models to other materials.

Models can also be developed to represent more general microstructure. Early work by Poulin et al.[1990] used a set of small cylinders. Westin et al.[1992] computed a BRDF by simulating the optical scattering of specified micro-geometry. Ashikhmin et al.[2000] proposed a model based on an NDF for surface microfacets. Recently Ngan et al.[2005] showed that this microfacet-based model captures real-world anisotropic materials better than traditional parametric models. They developed a method to fit the NDF from captured 4D BRDF data.

We also model anisotropic BRDFs using tabulated normal distribution functions as in [Ashikhmin et al. 2000; Ngan et al. 2005]. But we deal with the challenge of inferring *spatially varying* BRDFs from sparse measured data.

Dense view acquisition using a spatial gonioreflectometer was proposed by Dana et al.[1999] to directly measure SVBRDFs and bi-directional texture functions (BTFs) of real world surfaces. Different systems have since been developed to capture SVBRDFs [McAllister et al. 2002; Gu et al. 2006], BTFs [Dana 2001; Muller et al. 2005], and reflectance fields [Garg et al. 2006]. These methods all require a dense sampling of view and light directions over the hemisphere. This is achieved by moving any two of the three components (camera, object, and light source) or by using a dome containing multiple light sources and cameras. Although several techniques have been proposed to simplify the capturing work with special devices [Dana 2001; Han and Perlin 2003], or by exploiting intrinsic properties of reflectance data [Garg et al. 2006; Lensch et al. 2003b], acquisition still takes a long time or needs an expensive hardware setup. Precise image registration is also required to detect surface correspondences, a difficult task that leads to blurring if not done carefully [Weyrich et al. 2007].

Sparse view acquisition has been employed by several researchers to solve the camera calibration problem and make BRDF acquisition cheaper and easier. In [Gardner et al. 2003], the surface is scanned with a linear light source and captured from a fixed view. An isotropic Ward model is then fit to the captured data at each point. Lensch et al.[2003a] reconstructed SVBRDFs of a real object of known geometry. BRDFs are grouped into a small set

each fit using a Lafortune model basis, and reflectance at every point is represented as a linear combination over this basis. Goldman et al.[2005] use the same linear combination idea but with an isotropic Ward model as the BRDF basis, to reconstruct both an object’s shape and its SVBRDF from sparse reflectance data. Although these methods preserve detailed spatial variation of surface reflectance, the simple parametric models on which they rely do not accurately capture details or anisotropy of angular variation.

Debevec et al.[2000] apply a generalized microfacet model to human skin. A face is acquired from a dense set of illumination directions, at a single view from each side. Given specific assumptions about the skin’s BRDF (i.e., the refraction index is known and spatially constant, and the shadowing is based on the V-shaped groove model) the face can then be rendered under new viewing and lighting conditions. We capture surface reflectance for arbitrary and unknown materials from a single view. We are also able to reconstruct reflectance for view/light directions whose half-angle was not actually acquired. We do this by explicitly inferring partial NDFs and completing them by merging from other surface points.

Recently, Zickler et al.[2005] presented a method to infer a BRDF at one point by exploiting data from neighboring surface points. Their method models the SVBRDF using six-dimensional radial basis functions which fit sparse reflectance data in each local region. Angular resolution is enhanced at the cost of spatial resolution. Moreover, they assume that reflectance is isotropic in direction and varies smoothly in space. Our representation makes no assumptions about the material distribution over the surface – it models the BRDF as a tabular microfacet normal distribution function that is spatially variant. By synthesizing the partial NDF fitted at different surface points, our approach captures angular variation without decreasing spatial resolution.

3 Microfacet-based SVBRDF Model

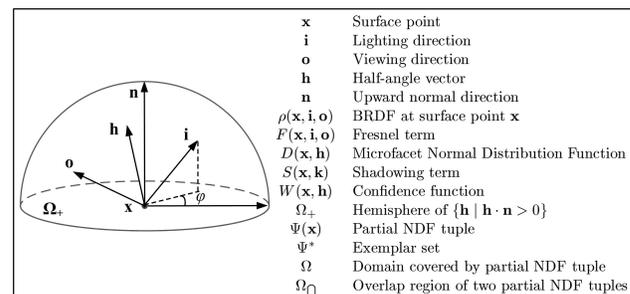


Figure 2: Symbols used in BRDF definition.

Microfacet theory represents surface microstructure with a large number of tiny mirror facets, each purely reflective and having the same refraction index [Cook and Torrance 1981]. The BRDF $\rho(\mathbf{x}, \mathbf{i}, \mathbf{o})$ at a surface point \mathbf{x} can be decomposed into two parts:

$$\rho(\mathbf{x}, \mathbf{i}, \mathbf{o}) = \rho_d(\mathbf{x}, \mathbf{i}, \mathbf{o}) + k_s(\mathbf{x}) \rho_s(\mathbf{x}, \mathbf{i}, \mathbf{o}) \quad (1)$$

where $\rho_s(\mathbf{x}, \mathbf{i}, \mathbf{o})$ is the specular part due to single-bounce reflections, $\rho_d(\mathbf{x}, \mathbf{i}, \mathbf{o})$ is the diffuse part resulting from subsurface scattering and multiple-bounce reflections, and $k_s(\mathbf{x})$ controls the diffuse-to-specular ratio. We model the diffuse part with the Lambertian model $\rho_d(\mathbf{x}, \mathbf{i}, \mathbf{o}) = k_d(\mathbf{x})/\pi$.

Assuming the microfacets form a height field, the specular term is given by [Cook and Torrance 1981]:

$$\rho_s(\mathbf{x}, \mathbf{o}, \mathbf{i}) = \frac{D(\mathbf{x}, \mathbf{h}) G(\mathbf{x}, \mathbf{o}, \mathbf{i}) F(\mathbf{x}, \mathbf{o}, \mathbf{i})}{4(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})} \quad (2)$$

where $\mathbf{h} = (\mathbf{o} + \mathbf{i})/|\mathbf{o} + \mathbf{i}|$ is the half-angle vector (shown in Figure 2), and \mathbf{n} is the upward surface normal at \mathbf{x} . Note that for a flat sample, $\mathbf{n} = (0, 0, 1)$. The model consists of three spatially-varying

components: the Fresnel term F , the microfacet Normal Distribution Function (NDF) D , and the shadowing-masking term G .

The Fresnel term $F(\mathbf{x}, \mathbf{o}, \mathbf{i})$ affects the specular shape near grazing angles. Following [Cook and Torrance 1981], it is given by

$$F(\mathbf{x}, \mathbf{o}, \mathbf{i}) = \frac{(g-c)^2}{2(g+c)^2} \left(1 + \frac{(c(g+c)-1)^2}{(c(g-c)+1)^2} \right) \quad (3)$$

for unpolarized light. Here $g^2 = \eta(\mathbf{x})^2 + c^2 - 1$ and $c = |\mathbf{i} \cdot \mathbf{h}|$. $\eta(\mathbf{x})$ denotes the relative refraction index at surface location \mathbf{x} .

The normal distribution function (NDF) $D(\mathbf{x}, \mathbf{h})$ describes the distribution of microfacet orientations on the surface. It satisfies $D(\mathbf{x}, \mathbf{h}) \geq 0$ and $\int_{\Omega_+} (\mathbf{n} \cdot \mathbf{h}) D(\mathbf{x}, \mathbf{h}) d\omega_{\mathbf{h}} = 1$ where the integration is done in half-angle space over the hemisphere

$$\Omega_+ = \Omega_+(\mathbf{n}) = \left\{ \mathbf{h} \mid \mathbf{h} \cdot \mathbf{n} > 0 \right\}.$$

High-frequency characteristics of surface reflectance such as glossiness and anisotropy are dominated by $D(\mathbf{x}, \mathbf{h})$, so we make no assumptions about it and represent it in tabular form.

$G(\mathbf{x}, \mathbf{o}, \mathbf{i})$ represents the shadowing and masking effects of microfacets on both incoming and outgoing light. It can be approximately decomposed into the product of two directional shadowing terms S as in [Beckmann 1965; Smith 1967; Ashikmin et al. 2000]:

$$G(\mathbf{x}, \mathbf{o}, \mathbf{i}) = S(\mathbf{x}, \mathbf{i}) S(\mathbf{x}, \mathbf{o}). \quad (4)$$

We follow the derivation in [Ashikmin et al. 2000] to compute the shadowing term $S(\mathbf{x}, \mathbf{k})$ for a given direction \mathbf{k} from the NDF $D(\mathbf{x}, \mathbf{h})$ via

$$S(\mathbf{x}, \mathbf{k}) = \frac{(\mathbf{k} \cdot \mathbf{n})}{\int_{\Omega_+(\mathbf{k}) \cap \Omega_+(\mathbf{n})} (\mathbf{h} \cdot \mathbf{k}) D(\mathbf{x}, \mathbf{h}) d\omega_{\mathbf{h}}}. \quad (5)$$

The shadowing term is smooth [Ashikmin et al. 2000; Ngan et al. 2005], a fact we exploit in partial NDF recovery to constrain the shadowing term (Section 5.1).

The full SVBRDF model is therefore

$$\rho(\mathbf{x}, \mathbf{o}, \mathbf{i}) = \frac{k_d(\mathbf{x})}{\pi} + k_s(\mathbf{x}) \frac{D(\mathbf{x}, \mathbf{h}) S(\mathbf{x}, \mathbf{i}) S(\mathbf{x}, \mathbf{o}) F(\mathbf{x}, \mathbf{o}, \mathbf{i})}{4(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})} \quad (6)$$

and is represented by four quantities: an NDF D , a scalar Fresnel refraction index η , and scalar specular and diffuse coefficients, k_d and k_s . Each is a function of surface location \mathbf{x} .

Given the SVBRDF, outgoing radiance $L(\mathbf{x}, \mathbf{o})$ at each point \mathbf{x} can be computed by the rendering equation as

$$L(\mathbf{x}, \mathbf{o}) = \int_{\Omega_+} \rho(\mathbf{x}, \mathbf{i}, \mathbf{o}) (\mathbf{i} \cdot \mathbf{n}) L(\mathbf{x}, \mathbf{i}) d\omega_{\mathbf{i}} \quad (7)$$

where $L(\mathbf{x}, \mathbf{i})$ is the incoming radiance from direction \mathbf{i} .

Our implementation parameterizes direction space with a hemicycle, typically subdivided into 32×32 cells on the top face and 32×16 cells on each side. For more specular materials such as aluminum, we subdivide the hemicycle more finely. Here ‘‘direction space’’ refers to half-angle directions \mathbf{h} for the NDF, and light directions \mathbf{i} for the 2D BRDF slice for a single view direction, \mathbf{o} .

4 Reflectance Data Acquisition

Data is acquired by illuminating the surface from different directions and capturing its reflectance from one view. We adapt the device introduced in [Gardner et al. 2003]. As shown in Figure 3, the setup includes a camera and a linear light source mounted above a flat sample and controlled by a stepping motor. Typical dimensions of a surface sample are $10\text{cm} \times 10\text{cm}$. We replace the linear light

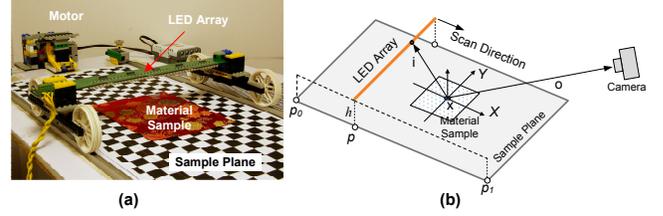


Figure 3: Our data acquisition device: (a) photograph, (b) diagram.

source used in [Gardner et al. 2003] with a linear array of 40 LEDs. The distance between neighboring LEDs in the array is 0.6cm. A Canon 30D camera with EF-100 2.8 lens is placed about 1.0m away and 0.6m above the center of the surface sample, making an angle of roughly $\theta=60^\circ$ from vertical. Image resolution is 3504×2336 .

Before capturing, we calibrate the camera’s position and orientation with the method in [Zhang 2000]. A colored checker pattern is used to calibrate the color and intensity of each LED. After geometric and photometric calibration, we put the flat material sample onto the XY plane, centered at the origin. Based on the size and thickness of the material sample, we can adjust the height h of the LED array above the sample plane, but the typical value is 4cm. The LED array is then placed at the starting position p_0 , as far as possible away from the sample, and translated over it. The step length is equal to the LED separation size. For each LED position, images of different exposures are taken to reconstruct an HDR image as in [Debevec and Malik 1997]. The scan is finished when the LED array reaches p_1 , where it starts to occlude the sample from the camera’s view. The result is a sequence of images $I_q(\mathbf{x})$ where the subscript $q \in \{1, 2, \dots, n\}$ indicates a different point light source (i.e. different LED with different Y translation). We manually measure p_0 and p_1 so that the light source position can be recovered.

After capturing, the HDR images are reconstructed and rectified. A BRDF sample at each surface point \mathbf{x} can then be computed from the image sequence I_q , as

$$\rho(\mathbf{x}, \mathbf{i}_q(\mathbf{x}), \mathbf{o}(\mathbf{x})) = \frac{I_q(\mathbf{x})}{(\mathbf{n} \cdot \mathbf{i}_q(\mathbf{x})) L_q \|P_q - \mathbf{x}\|^2} \quad (8)$$

where P_q is the position of the light source in image q and L_q is its intensity. The lighting direction \mathbf{i} is computed as $\mathbf{i}_q(\mathbf{x}) = (P_q - \mathbf{x}) / \|P_q - \mathbf{x}\|$. Since the camera is relatively far from the sample, we use the viewing direction at the sample center for all surface points, $\mathbf{o} = V / \|V\|$, where V is the calibrated camera position.

We then separate the diffuse from specular components using

$$\rho_d(\mathbf{x}) = \min_q \{ \rho(\mathbf{x}, \mathbf{i}_q(\mathbf{x}), \mathbf{o}) \mid I_q(\mathbf{x}) > 0.05 I_{avg} \} \quad (9)$$

where $\rho(\mathbf{x}, \mathbf{i}_q(\mathbf{x}), \mathbf{o}) = I_q(\mathbf{x}) / (I_q(\mathbf{x}) \cdot \mathbf{n})$. Intensity measurements below five percent of average I_{avg} tend to be noisy and unreliable and so are removed. This method is naive compared with existing approaches [Debevec et al. 2000; Gardner et al. 2003], but it produces reasonable results nonetheless. After separation we obtain the scaled specular measurement $\tilde{\rho}_s(\mathbf{x}, \mathbf{i}_q(\mathbf{x}), \mathbf{o}) = \rho(\mathbf{x}, \mathbf{i}_q(\mathbf{x}), \mathbf{o}) - \rho_d(\mathbf{x}) = k_s(\mathbf{x}) \rho_s(\mathbf{x}, \mathbf{i}_q(\mathbf{x}), \mathbf{o})$ at densely-sampled lighting directions. Finally, these scattered measurements are interpolated with the push-pull method [Gortler et al. 1996], as in [Lawrence et al. 2006]. This yields a 2D slice of the specular component $\tilde{\rho}_s(\mathbf{x}, \mathbf{i}, \mathbf{o})$ that is uniformly sampled over the hemicycle of incident lighting directions \mathbf{i} . The next phase uses $\tilde{\rho}_s$ to fit the NDF and the BRDF model’s other parameters.

It is possible that a single view can fail to provide enough exemplars for NDF synthesis, if there is insufficient spatial redundancy or local frame variation over the material sample. An example is brushed metal whose brush direction is constant over the entire sample. Such cases can easily be addressed by capturing image sequences for different rotations of the sample. BRDF slices from

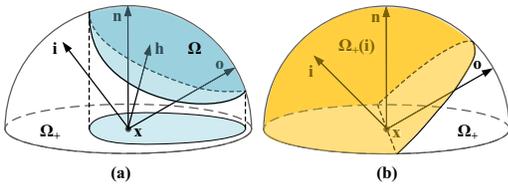


Figure 4: Domains for the partial NDF and its shadowing term. (a) The partial NDF’s domain Ω for a given view direction, \mathbf{o} , is shaded in blue. (b) Computing the shadow term S for a given light direction \mathbf{i} needs D over the domain shaded in yellow. Only part of it is covered in (a).

these different rotations can then be put together as if they were captured from a larger sample in a single pass. None of our results have required additional rotations.

5 Example-Based Microfacet Synthesis

With the captured 2D BRDF data at each surface point, the surface’s SVBRDF is modeled in three steps. We first construct the partial NDF of the microfacets at each point. The NDF is partial because the single view direction generates half-angle reflected directions covering only a portion of the hemisphere. Microfacet synthesis then completes the partial NDF by “stitching” partial NDFs of similar regions together. The shadowing and Fresnel terms are derived from the completed NDF in the final step. Each step is detailed in the following subsections.

5.1 Partial NDF Fitting

Given a 2D slice of specular BRDF data at each surface point, we fit its NDF in this step. Each surface point’s NDF is computed independently. To avoid effects from the Fresnel term that will be derived later, we exclude reflectance data near the grazing angle and assume $F(\mathbf{x}, \mathbf{i}, \mathbf{o}) = 1$. As shown in Figure 4a, with a fixed view direction \mathbf{o} and light directions \mathbf{i} over the whole hemisphere, the sampled half-angle reflectance directions cover only a sub-region Ω of the hemisphere:

$$\Omega = \left\{ \frac{\mathbf{i} + \mathbf{o}}{\|\mathbf{i} + \mathbf{o}\|} \mid \mathbf{i} \in \Omega_+ \right\} = \left\{ \mathbf{h} \mid \text{refl}(\mathbf{o}, \mathbf{h}) \cdot \mathbf{n} > 0 \right\} \quad (10)$$

where $\text{refl}(\mathbf{o}, \mathbf{h}) = 2(\mathbf{h} \cdot \mathbf{o})\mathbf{h} - \mathbf{o}$ reflects the unit vector \mathbf{o} about \mathbf{h} . So the captured data specifies the NDF only partially, over Ω .

A straightforward solution for fitting this partial normal distribution function $D(\mathbf{x}, \mathbf{h})$ ¹ adapts the method in [Ngan et al. 2005]. Starting from constant shadow term $S(\mathbf{x}, \mathbf{k})=1$, the partial NDF is computed as:

$$D(\mathbf{x}, \mathbf{h}) = \begin{cases} \frac{4\pi \tilde{\rho}_s(\mathbf{x}, \mathbf{o}, \mathbf{i}(\mathbf{h})) (\mathbf{i}(\mathbf{h}) \cdot \mathbf{n}) (\mathbf{o} \cdot \mathbf{n})}{S(\mathbf{x}, \mathbf{i}(\mathbf{h})) S(\mathbf{x}, \mathbf{o})}, & \mathbf{h} \in \Omega \\ 0, & \mathbf{h} \notin \Omega \end{cases} \quad (11)$$

where $\mathbf{i}(\mathbf{h}) = \text{refl}(\mathbf{o}, \mathbf{h})$. The shadowing term is then derived from the NDF using (5). This leads to a scheme that iteratively updates first D and then S until convergence. Although this method fits the captured data well, it is biased as illustrated in Figure 5b, and does not accurately predict specular response at other views.

The problem arises because single-view data determines a partial NDF over Ω (Figure 4a), while the shadowing term requires the full NDF over Ω_+ (Figure 4b). If we assume that $D = 0$ outside Ω and use it to estimate the shadowing term in (5), we end up overestimating S . After iterating, the NDF becomes biased.

To solve this problem, we constrain S to its minimum value over all azimuthal angles via $S(\mathbf{x}, \mathbf{i}) = \min_{\mathbf{i}'} \{S(\mathbf{x}, \mathbf{i}') \mid \mathbf{i}' \cdot \mathbf{n} = \mathbf{i} \cdot \mathbf{n}\}$. With

¹Though the NDF D and its domain Ω are inferred in terms of half-angle space, \mathbf{h} , they determine a probability distribution over the hemisphere of microfacet normal directions. Nevertheless, we will continue to denote the function’s argument as \mathbf{h} .

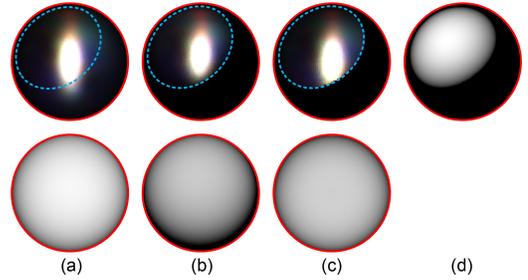


Figure 5: Recovering the partial NDF from a single-view BRDF slice: (a) ground truth, (b) reconstruction using [Ngan et al. 2005] (domain Ω is inside the dashed blue curve), (c) reconstruction computed by our method, (d) confidence function W . The top row (a-c) shows the NDF D while the bottom row (a-c) shows its shadowing term S (log of value). Both are visualized using an orthographic projection.

this isotropically constrained S , we iteratively update the NDF (11) and its shadowing term (5) until convergence, as in the straightforward method. Figure 5c shows that the resulting partial NDF is closer to the ground truth in Figure 5a. We have experimented with other schemes to extend the domain of D outside Ω in order to define S , including projection to low-order spherical harmonics and push/pull. None has worked as well as this simple min-azimuthal method. Note that the constrained shadowing term is only used to fit the partial NDF; the final S is derived from the completed NDF in the last step.

To represent the partial NDF’s domain Ω , we define a confidence function via

$$W(\mathbf{x}, \mathbf{h}) = \left[(\text{refl}(\mathbf{o}, \mathbf{h}) \cdot \mathbf{n})^{\frac{1}{2}} \right]_+ \quad (12)$$

where $[\cdot]_+$ denotes $\max\{\cdot, 0\}$. W is positive for \mathbf{h} within Ω , gradually decreases to 0 at its boundary, and remains 0 outside it. Low confidence is assigned near the boundary of Ω where BRDF values tend to be extrapolated and dominated by the Fresnel term. The partial NDF $D(\mathbf{x}, \mathbf{h})$ and its confidence function $W(\mathbf{x}, \mathbf{h})$ are represented as a hemicube over \mathbf{h} .

Note that it is possible to explicitly represent spatial variation in the set of light directions $\mathbf{i}_q(\mathbf{x})$ and view direction $\mathbf{o}(\mathbf{x})$, by using a spatially varying NDF domain $\Omega(\mathbf{x})$. This accounts for the finite lighting span and camera distance. Unfortunately, it also complicates our synthesis algorithm and we leave the idea for future work.

5.2 NDF Completion using Synthesis

After recovering the partial NDF $D(\mathbf{x}, \mathbf{h})$ at each surface point, we complete it using partial NDFs at other points. The key observation is that on a sample surface, we can always find surface points having similar NDFs but different local frames; i.e., different *azimuthal* rotations around the normal. So the captured data actually reveals different parts of the common NDF at such points. As shown in Figure 6, we exploit this idea by finding rotated partial NDFs at other surface points that are similar within their shared half-space domain and merging them.

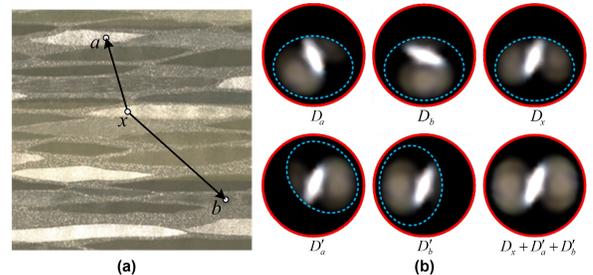


Figure 6: Merging partial NDFs from different points. (a) Partial NDF of point x is similar to points a and b . (b) D_a and D_b are aligned to D_x by azimuthal rotation. Aligned NDFs D'_a and D'_b are merged to complete D_x .

Given the partial NDF tuple $\Psi(\mathbf{x}) = \langle D(\mathbf{x}, \mathbf{h}), W(\mathbf{x}, \mathbf{h}) \rangle$ defined on each surface point, the exemplar set Ψ^* is generated by rotating $\Psi(\mathbf{x})$ around the normal \mathbf{n} by all possible azimuthal angles φ :

$$\begin{aligned} \Psi^* &= \{\Psi(\mathbf{x}, \varphi)\} = \{\langle D^\varphi(\mathbf{x}, \mathbf{h}), W^\varphi(\mathbf{x}, \mathbf{h}) \rangle\} \\ &= \left\{ \langle D(\mathbf{x}, R(\varphi)\mathbf{h}), W(\mathbf{x}, R(\varphi)\mathbf{h}) \rangle \mid \varphi \in [0, 2\pi] \right\} \end{aligned} \quad (13)$$

where $R(\varphi)$ is an azimuthal rotation of φ .

```

NDFSynthesis( $\Psi(\mathbf{x}), \Psi^*$ )
For each surface point  $\mathbf{x}$ 
   $\Psi_0(\mathbf{x}) = \Psi(\mathbf{x})$ 
  While ( $\Omega_i(\mathbf{x}) < \Omega_+$ )
     $\Psi(\mathbf{x}', \varphi') = \text{BestMatch}(\Psi_i(\mathbf{x}), \Psi^*)$ 
     $\Psi_{i+1}(\mathbf{x}) = \text{Merge}(\Psi_i(\mathbf{x}), \Psi(\mathbf{x}', \varphi'))$ 

```

We complete each point's partial NDF with iterative synthesis using the pseudo-code shown above. For each surface point \mathbf{x} , starting from the initial NDF tuple $\Psi_0(\mathbf{x}) = \Psi(\mathbf{x}) = \langle D(\mathbf{x}, \mathbf{h}), W(\mathbf{x}, \mathbf{h}) \rangle$, we update the NDF and its confidence function iteratively, yielding $\Psi_i(\mathbf{x}) = \langle D_i(\mathbf{x}, \mathbf{h}), W_i(\mathbf{x}, \mathbf{h}) \rangle$ at each iteration i . Each iteration extends the NDF domain incrementally in counter-clockwise order around the azimuth, as we will further explain in the next section.

The update searches the exemplar set Ψ^* to find a match of the current NDF tuple $\Psi_i(\mathbf{x})$, $\Psi(\mathbf{x}', \varphi')$, via

$$(\mathbf{x}', \varphi') = \underset{\tilde{\mathbf{x}}, \tilde{\varphi}}{\text{argmin}} \{ \text{dist}(\Psi_i(\mathbf{x}), \Psi(\tilde{\mathbf{x}}, \tilde{\varphi})) \}. \quad (14)$$

Distance between two partial NDF tuples $\Psi_1 = \langle D_1, W_1 \rangle$ and $\Psi_2 = \langle D_2, W_2 \rangle$ is defined using the integral

$$\text{dist}(\Psi_1, \Psi_2) = \frac{\int_{\Omega_\cap} \sqrt{W_1(\mathbf{h})W_2(\mathbf{h})} \left\| \frac{D_1(\mathbf{h})}{s_1} - \frac{D_2(\mathbf{h})}{s_2} \right\|^2 d\omega_h}{\int_{\Omega_\cap} \sqrt{W_1(\mathbf{h})W_2(\mathbf{h})} d\omega_h} \quad (15)$$

over their overlapping region

$$\Omega_\cap = \{\Omega_1 \cap \Omega_2\} = \{\mathbf{h} \mid W_1(\mathbf{h})W_2(\mathbf{h}) > 0\}.$$

The scalars s_1 and s_2 normalize the partial NDFs over the overlapped region and are defined as $s_k = \int_{\Omega_\cap} D_k(\mathbf{h}) d\omega_h$ for $k=1,2$.

To maintain synthesis quality, we exclude candidates from the search set Ψ^* whose partial domains overlap insufficiently with $\Psi_i(\mathbf{x})$. We also exclude candidates that insufficiently extend $\Psi_i(\mathbf{x})$. We adopt the simple rule that candidates must have 50 – 85% overlap with $\Psi_i(\mathbf{x})$.

After determining the best match, the new NDF tuple $\Psi_{i+1}(\mathbf{x})$ is updated by merging in the optimal extension $\Psi(\mathbf{x}', \varphi')$ via

$$D_{i+1}(\mathbf{x}, \mathbf{h}) = \frac{W_i(\mathbf{x}, \mathbf{h})D_i(\mathbf{x}, \mathbf{h}) + \frac{s_1}{s_2}W^{\varphi'}(\mathbf{x}', \mathbf{h})D^{\varphi'}(\mathbf{x}', \mathbf{h})}{W_i(\mathbf{x}, \mathbf{h}) + W^{\varphi'}(\mathbf{x}', \mathbf{h})}, \quad (16)$$

$$W_{i+1}(\mathbf{x}, \mathbf{h}) = \max \left\{ W_i(\mathbf{x}, \mathbf{h}), W^{\varphi'}(\mathbf{x}', \mathbf{h}) \right\}. \quad (17)$$

Continuing with the new partial NDF tuples $\Psi_{i+1}(\mathbf{x})$ at each surface point, we repeat the above steps until each point's NDF covers the hemisphere. The NDF coverage $\Omega_i(\mathbf{x})$ after each iteration can be reconstructed from $W_i(\mathbf{x}, \mathbf{h})$ via $\Omega_i(\mathbf{x}) = \{\mathbf{h} \mid W_i(\mathbf{x}, \mathbf{h}) > 0\}$. There may be small uncovered regions left in some NDFs after synthesis, which can be interpolated with the push-pull method. We find that NDFs are suitably completed after five to seven iterations.

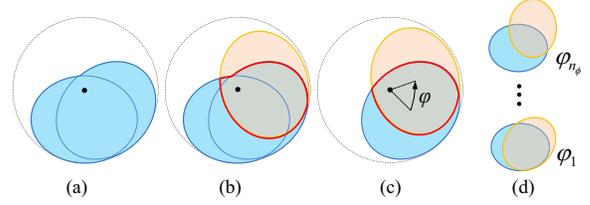


Figure 7: Overlap region approximation in search pruning. (a) Current $\Psi_i(\mathbf{x})$, formed by an earlier merge of two partial NDFs. (b) Overlap region Ω_\cap (enclosed by red line) of $\Psi_i(\mathbf{x})$ and a candidate (orange). (c) The overlap region can be approximated by the intersection of Ω with a rotated version of itself. (d) A range of rotations, φ_i , are uniformly sampled to get a discrete set of overlap regions used for search pruning.

5.3 Synthesis Acceleration

A naive synthesis implementation is prohibitively slow due to the expensive distance calculation in the candidate search and the large number of surface points. Two ideas speed up synthesis.

Search Pruning We accelerate the search by pruning the set of candidates. At first sight, the region $\Psi_i(\mathbf{x})$ is extended differently at each point and each iteration, depending on which candidate is selected for the merge. The overlap region therefore cannot be pre-determined. However as shown in Figure 7b, as we extend the NDF domain azimuthally, the overlap region is mostly determined by the candidate we merged last and the current candidate. Therefore, the overlap region can be represented as the intersection of Ω with a rotated version of itself, $\Omega(\varphi) = \Omega \cap R(\varphi)\Omega$ as in Figure 7c.

Since the candidates have a 50-85% overlap with $\Psi_i(\mathbf{x})$, φ need only be chosen in a limited range which we uniformly sample using $n_\phi=20$ angles. The result is a set of regions $\{\Omega(\varphi_i) \mid i=1, 2, \dots, n_\phi\}$ as shown in Figure 7d. Within each $\Omega(\varphi_i)$, we compute the histogram of $D(\mathbf{x}, \mathbf{h})$ at each surface point \mathbf{x} using $m=32$ buckets over D 's range $[0,1]$. We then use the resulting 32D vector as a search key to find merge candidates that are similar within the region of overlap. To accelerate this search, we precompute an ANN tree of the histogram vectors before synthesis [Mount and Arya 1997]. A separate tree is built for each φ_i .

During synthesis, for each φ_i , we compute the histogram of the current region $\Psi_i(\mathbf{x})$ over $\Omega(\varphi_i)$. We use the corresponding ANN tree to quickly find the five best matches in terms of L^2 distance in histogram space. For each match, gradient descent obtains the optimal φ' by minimizing (14) using the full-blown distance computation in (15). The minimal error match over all φ is then selected to merge. In the final iteration, we must consider overlaps between the first as well as the last merged domain because azimuthal rotation is periodic. We simply do the ANN search over both these domains to form a larger candidate set before applying the φ optimization.

NDF Clustering The second acceleration idea reduces both the number of NDFs that must be synthesized as well as searched. Since many surface points have similar reflectance and can be approximated well by a linear combination of a few representatives [Matusik et al. 2003; Lawrence et al. 2006], we can perform expensive NDF completion on a smaller set of representatives, and obtain a high-resolution result by interpolating their completions.

To find the set of representatives, we apply k -means clustering to the partial NDFs of all surface points. A representative in each cluster is selected that is closest to the cluster center where distance is computed according to (15). All examples in the paper set the number of representatives to be 1% of the number of surface points. A large number of clusters ensures each contains only samples that are similar.

We then find interpolation weights on the partial data. Each (non-representative) partial NDF, $D(\mathbf{x}_i, \mathbf{h})$, is approximated by lin-

ear interpolation on a set of neighbor representatives, $D(\mathbf{x}_j^*, \mathbf{h})$, via

$$D(\mathbf{x}_i, \mathbf{h}) = \sum_{j \in N_i} w_{ij} D(\mathbf{x}_j^*, \mathbf{h}), \quad (18)$$

where $j \in N_i$ indexes one of the neighbor representatives of \mathbf{x}_i , and w_{ij} are the interpolation coefficients. To determine the neighbor set N_i for \mathbf{x}_i , we first collect its $k=16$ nearest representatives, excluding those whose distance is larger than 5λ where λ is the smallest distance between two representative NDFs. We then solve for the weights w_{ij} in the linear equation system represented by (18), augmented by the single equation $\sum_{j \in N_i} w_{ij} = 1$. We set $k=16$ based on the observation that the intrinsic dimension of the BRDF space is roughly 16 [Matusik et al. 2003].

5.4 Estimating the Remaining Parameters

We compute the specular coefficient given the completed NDF D via $k_s(\mathbf{x}) = \int_{\Omega_+} (\mathbf{n} \cdot \mathbf{h}) D(\mathbf{x}, \mathbf{h}) d\omega_h$. We then normalize the NDF by $D(\mathbf{x}, \mathbf{h})/k_s(\mathbf{x})$ and use it to compute the shadowing term from (5).

Finally, the relative refraction index $\eta(\mathbf{x})$ at each surface point is estimated by minimizing

$$E(\eta(\mathbf{x})) = \int_{\Omega_+} |F_c(\mathbf{i}, \mathbf{o}, \eta(\mathbf{x})) - F_m(\mathbf{x}, \mathbf{i}, \mathbf{o})|^2 d\omega \quad (19)$$

where $F_c(\mathbf{i}, \mathbf{o}, \eta(\mathbf{x}))$ is computed according to (3), and $F_m(\mathbf{x}, \mathbf{i}, \mathbf{o})$ is computed from the measured $\rho_s(\mathbf{x}, \mathbf{o}, \mathbf{i})$ as

$$F_m(\mathbf{x}, \mathbf{i}, \mathbf{o}) = \frac{4\pi\rho_s(\mathbf{x}, \mathbf{o}, \mathbf{i})(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})}{S(\mathbf{x}, \mathbf{i})S(\mathbf{x}, \mathbf{o})k_s(\mathbf{x})D(\mathbf{x}, \mathbf{h})}. \quad (20)$$

We solve for $\eta(\mathbf{x})$ with the Levenberg-Marquardt algorithm [Press et al. 1992].

6 Experimental Results

We tested our method on publicly available, full 6D SVBRDF data as well as 4D, fixed-view data slices captured with our simple device. Table 1 summarizes statistics for the latter, including spatial resolution of acquired images, light resolution (product of the number of LEDs used in capture and the number of translation steps), resolution of the top hemicube face of the resulting NDFs, and the camera viewing angles. Viewing direction is measured in terms of its elevation angle (angle with vertical axis Z) θ and azimuthal angle (angle in XY plane) φ .

We implemented our microfacet synthesis algorithm on a PC with Intel Core™2 Quad 2.13GHz CPU and 4GB memory. Capturing takes about 1 hour using single-exposure acquisition (for less specular materials like velvet) and 5-10 hours using multiple-exposure acquisition (for highly specular materials like aluminum). The angular sampling density for lighting is manually chosen: for relatively diffuse materials, we can disable half the LEDs to speed up the capturing process. Image data processing (including calibration, HDR reconstruction, diffuse separation, and resampling) takes about 2-4 hours, and is dominated by disk I/O. Partial NDF reconstruction takes about an hour, synthesis 2-3 hours, and estimation of the remaining BRDF parameters 3-4 hours. Rendering results shown in the paper use ray tracing. Only direct lighting effects are included; inter-reflection between surfaces is ignored.

6.1 Validation with Dense-View Data

We applied our synthesis algorithm on two SVBRDF datasets from [Lawrence et al. 2006] and shared by the Princeton Graphics Group. This data was captured by a spherical gantry whose movable camera/lamp samples hundreds of lighting and tens of viewing directions. We applied our algorithm to a single view slice of the data, reserving the remaining views as a validation reference. The comparison is shown in Figure 8. On the wallpaper example (a,b,c,d),

| Sample | Image Res. | Light Res. | NDF Res. | View (θ, φ) |
|--------------|------------|------------|----------|----------------------------|
| red satin | 850×850 | 20×24 | 32×32 | (57.6°, 0.6°) |
| yellow satin | 750×750 | 20×24 | 32×32 | (58.3°, -0.2°) |
| wallpaper | 800×800 | 20×20 | 32×32 | (63.6°, -0.8°) |
| velvet | 600×500 | 20×20 | 32×32 | (61.2°, 2.1°) |
| rose wood | 600×600 | 40×65 | 32×32 | (53.3°, 4.6°) |
| oak wood | 800×800 | 40×65 | 32×32 | (48.6°, -0.3°) |
| aluminium | 250×400 | 40×65 | 128×128 | (40.8°, 4.9°) |
| copper | 800×800 | 40×50 | 32×32 | (51.0°, 1.9°) |

Table 1: Acquisition parameters for various samples.

the anisotropic reflectance and the local frame variation over fan-like regions is reproduced well by our technique. The greeting card result (e,f,g,h) preserves the reflectance properties (in this case, isotropic) and their spatial variation.

6.2 Results with Single-View Data

We also experimented with data from our simple acquisition device. The following results were all obtained from data captured on this device and synthesized using our technique.

Figure 9 compares real images to rendered results from our method, on three material samples. The reconstruction uses a different view than the one captured. Although a good match is achieved, subtle differences could be due to errors in photometric calibration and white balance correction in the camera. A slight reduction in sharpness on the aluminum highlights in (b) probably means that our NDF hemicube resolution could be usefully increased. Results for other material samples can be found in the accompanying video.

Figure 10 compares our synthesized microfacet model with a fitted Ward model. Given data from our device, the Ward model parameters are fit at each surface point using the Levenberg-Marquardt algorithm [Press et al. 1992], as in [Goldman et al. 2005]. This process took about 6 hours for the anisotropic version of the model and so is comparable to the processing cost of our synthesis algorithm. Though the anisotropic Ward model matches better than the isotropic one, neither captures the visual characteristics of this red satin example as well as our technique. A comparison of 6D SVBRDFs is difficult to convey in a single image; please see the accompanying video for more extensive comparisons.

Figure 1 shows pillows decorated with four different materials and rendered under an HDR lighting environment. Satin (a,b) exhibits strong anisotropy due to the consistent fiber orientation. Fine details in the needlework (b) are reproduced by our model. The wallpaper example (c) shows local frame rotation of the plastic coating microstructure. A result for velvet is shown in (d).

Figure 11ab shows dishes mapped with weathered copper and brushed aluminum. Our model captures spatial variation to pro-

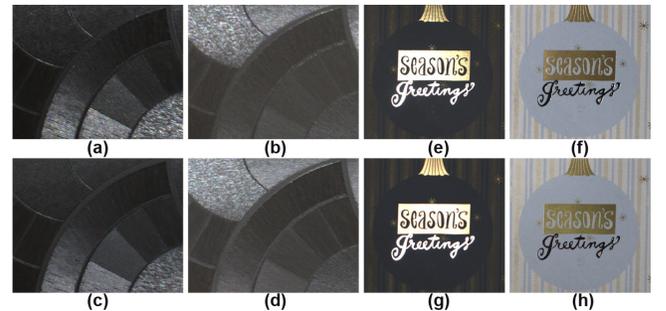


Figure 8: Validation comparison. The first row shows original measurements from the database; the second shows our synthesis results. The leftmost two columns are from wallpaper data [a,c = view#1/light#125; b,d = view#2/light#124]. The rightmost two columns are from greeting card data [e,g = view#0/light#188; f,h = view#2/light#12]. View #3 (wallpaper) and #0 (greeting card) were used as the input slice for our synthesis.

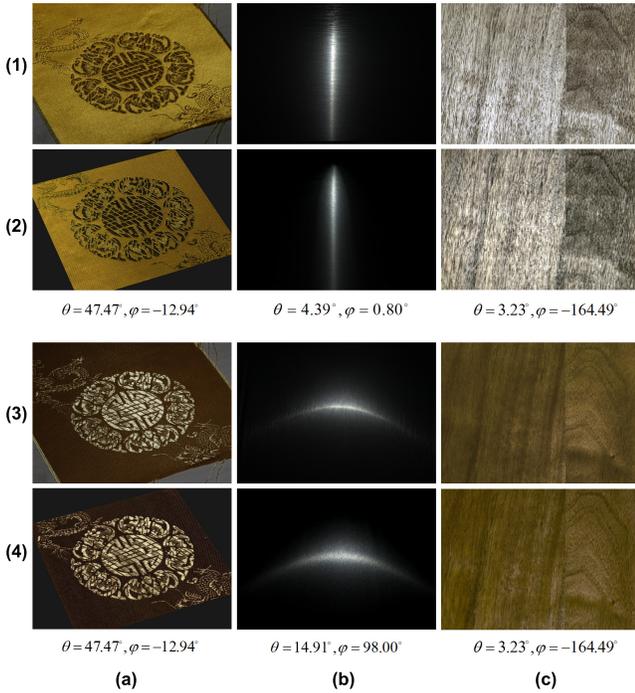


Figure 9: Comparison using single-view data from our simple device: (a) yellow satin, (b) brushed aluminum, (c) oak. We compare two different lighting directions (consecutive pairs of rows, direction is specified below the row pair), at a view ($\theta=50.9^\circ$, $\phi=-15.3^\circ$) different from the one captured. Rows (1) and (3) are images of the original sample; (2) and (4) are rendered from our synthesized model.

vide a realistic weathered appearance in (a), which includes both isotropic and anisotropic reflectance. The brushed aluminum in (b) is mapped onto the dish model using texture coordinates which simulate a circular brushing. Convincing, fan-shaped anisotropic highlights are generated.

Figure 11cd shows two kinds of wood mapped on a simple board model. Our method visually captures the anisotropic reflectance of wood and preserves subtle differences in appearance between the coarser oak strands vs. the much finer rosewood fibers. Our method models both surface and (thin layer) subsurface reflections [Marschner et al. 2005], without requiring their explicit separation. The accompanying video shows more light/view configurations.

7 Conclusion

Our microfacet synthesis technique generates anisotropic, spatially varying surface reflectance consistent with the appearance of real

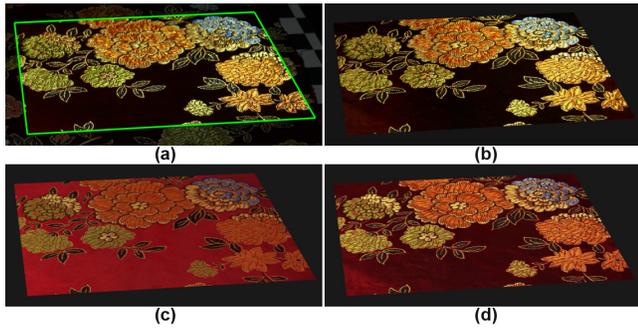


Figure 10: Comparison with parametric models: (a) real measured appearance, (b) our result, (c) isotropic Ward, (d) anisotropic Ward. The viewing direction reconstructed here is ($\theta=50.9^\circ$, $\phi=-15.3^\circ$).

samples. We have demonstrated successful results for a variety of materials from data captured from a single view. Our method avoids image registration and greatly simplifies data acquisition and processing. We hope this leads to a proliferation of measured SVBRDF data for many materials and applications.

Our approach is also subject to a number of limitations. Since the NDF at each point is synthesized from partial NDFs at other points rather than from 4D BRDF data, the resulting BRDF is not captured exactly. Our micro-facet model assumes only direct reflection, and so does not capture unusual phenomena dominated by multiple light bounces, such as retro-reflection. The intent of our method is to visually convey anisotropic reflectance of common materials from a small amount of easily captured data. Acquisition from a single view assumes enough sample points with shared but rotated reflectance. As we state in Section 4, more reflectance data can be obtained simply by rotating the sample.

In future work, we are interested in further optimizing the performance of the data capture, synthesis, and parameter estimation algorithms. Acquisition and processing might be sped up by exploiting sparsity still further, by capturing smaller domains over light direction and completing them by synthesis. We would also like to handle samples that are not flat. Both these ideas would require extending our synthesis algorithm to account for spatial variation in the NDF domain $\Omega(\mathbf{x})$. Backtracking during the NDF merging and capturing multiple views with (perhaps limited) surface correspondence information might improve synthesis accuracy. Generalizing the microfacet model to capture translucent objects and multiple-bounce effects is another extension. Finally, we want to explore ways to accelerate rendering of our SVBRDF model.

Acknowledgements

The LED array was designed and produced by Le Ma. Qiang Dai helped us to capture the raw data. The authors thank Steve Lin for discussions on BRDF modeling and for proofreading this paper, and Dong Xu for discussions on NDF clustering. We also thank the anonymous reviewers for their helpful suggestions and comments.

References

- ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic phong BRDF model. *Journal of Graphics Tools* 5, 2, 25–32.
- ASHIKHMIN, M., PREMOŽE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000, Computer Graphics Proceeding, Annual Conference Series*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 65–74.
- BECKMANN, P. 1965. Shadowing of random rough surfaces. *IEEE Transactions on Antennas and Propagation*, 13, 384–388.
- COOK, R. L., AND TORRANCE, K. E. 1981. A reflectance model for computer graphics. In *Computer Graphics (Proceedings of SIGGRAPH 81)*, 307–316.
- 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*, vol. 2, 460–466.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proceedings of ACM SIGGRAPH 1997, Computer Graphics Proceeding, Annual Conference Series*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 369–378.

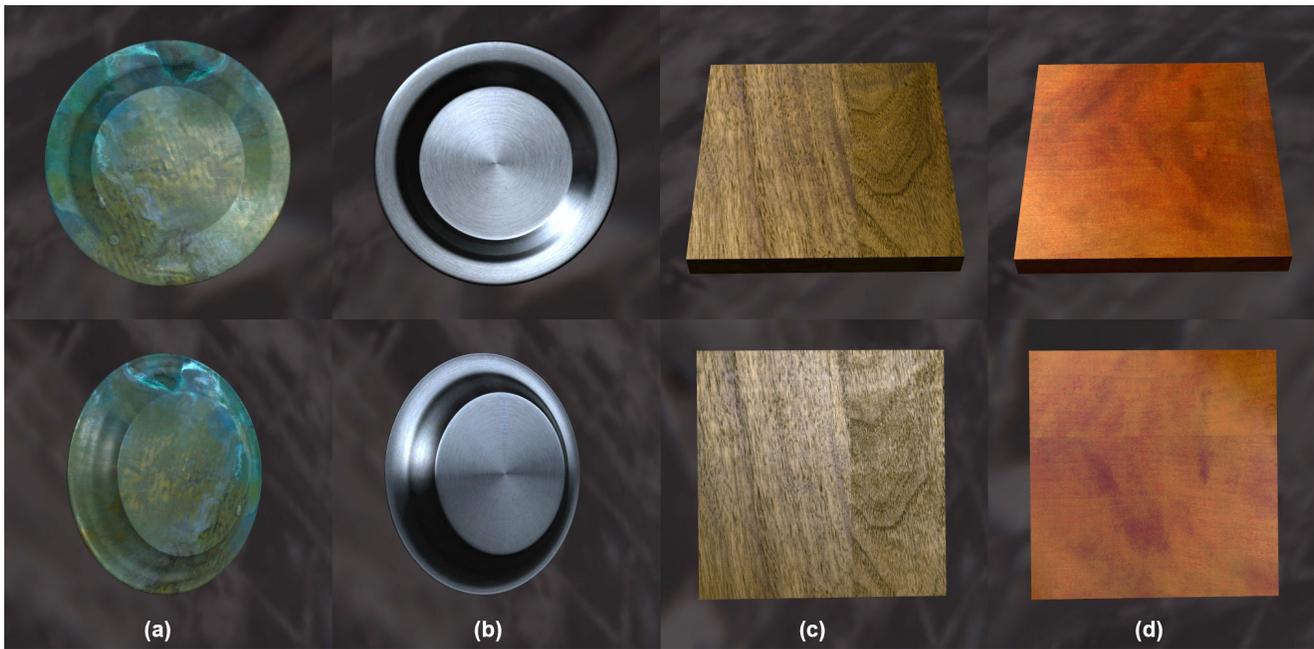


Figure 11: More results: (a) weathered copper, (b) brushed aluminum, (c) oak wood, (d) rose wood.

- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of ACM SIGGRAPH 2000, Computer Graphics Proceeding, Annual Conference Series*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 145–156.
- EFROS, A. A., AND FREEMAN, W. T. 2001. Image quilting for texture synthesis and transfer. In *Proceedings of ACM SIGGRAPH 2001, Computer Graphics Proceeding, Annual Conference Series*, ACM, New York, NY, USA, 341–346.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *ACM Transactions on Graphics* 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.
- GOLDMAN, D. B., CURLESS, B., HERTZMANN, A., AND SEITZ, S. M. 2005. Shape and spatially-varying BRDFs from photometric stereo. In *International Conference on Computer Vision*, I: 341–348.
- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In *Proceedings of ACM SIGGRAPH 1996, Computer Graphics Proceeding, Annual Conference Series*, 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 Transactions on Graphics* 25, 3, 762–771.
- HAN, J. Y., AND PERLIN, K. 2003. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics* 22, 3, 741–748.
- IRAWAN, P., AND MARSCHNER, S. R. 2006. A simple, accurate texture model for woven cotton cloth. Technical report PCG-06-01, Program of Computer Graphics, Cornell University, June.
- KAJIYA, J. T. 1985. Anisotropic reflection models. In *Computer Graphics (Proceedings of ACM SIGGRAPH 85)*, ACM, New York, NY, USA, 15–21.
- KWATRA, V., SCHÖDL, A., ESSA, I., TURK, G., AND BOBICK, A. 2003. Graphcut textures: Image and video synthesis using graph cuts. *ACM Transactions on Graphics* 22, 3 (July), 277–286.
- LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *Proceedings of ACM SIGGRAPH 1997, Computer Graphics Proceeding, Annual Conference Series*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 117–126.
- 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* 25, 3, 735–745.
- 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.
- LENSCH, H. P. A., LANG, J., SÁ, A. M., AND SEIDEL, H.-P. 2003. Planned sampling of spatially varying BRDFs. *Computer Graphics Forum* 22, 3, 473–482.
- MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. In *Proceedings of ACM SIGGRAPH 2003*, J. Hodgins and J. C. Hart, Eds., vol. 22(3) of *ACM Transactions on Graphics*, 780–791.

- MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Transactions on Graphics* 24, 3 (July), 727–734.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics* 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*.
- MULLER, G., MESETH, J., SATTLER, 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. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, 117–126.
- 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)*.
- POULIN, P., AND FOURNIER, A. 1990. A model for anisotropic reflection. In *Computer Graphics (Proceedings of ACM SIGGRAPH 90)*, F. Baskett, Ed., vol. 24, 273–282.
- PRESS, W. H., ET AL. 1992. Numerical recipes in C (second edition). *Cambridge University Press*.
- RUSINKIEWICZ, S. M. 1998. A new change of variables for efficient BRDF representation. In *Rendering Techniques (Proceeding of Eurographics Workshop on Rendering)*, 11–22.
- SMITH, B. G. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation*, 15, 668 – 671.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of ACM SIGGRAPH 92)*, ACM Press, New York, NY, USA, 265–272.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Computer Graphics (Proceedings of ACM SIGGRAPH 92)*, ACM Press, New York, NY, USA, 255–264.
- WEYRICH, T., LAWRENCE, J., LENSCH, H. P. A., RUSINKIEWICZ, S., AND ZICKLER, T. 2007. Principles of appearance acquisition and representation. In *Short Course ICCV 2007*.
- YASUDA, T., YOKOI, S., ICHIRO TORIWAKI, J., AND INAGAKI, K. 1992. A shading model for cloth objects. *IEEE Computer Graphics and Applications* 12, 6 (Nov.), 15–24.
- ZHANG, Z. 2000. A flexible new technique for camera calibration. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, 1330–1334.
- ZICKLER, T., ENRIQUE, S., RAMAMOORTHI, 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.