

Pocket Reflectometry

Peiran Ren ^{* †}

Jiaping Wang [†]

John Snyder [‡]

Xin Tong [†]

Baining Guo ^{† *}

^{*} Tsinghua University

[†] Microsoft Research Asia

[‡] Microsoft Research

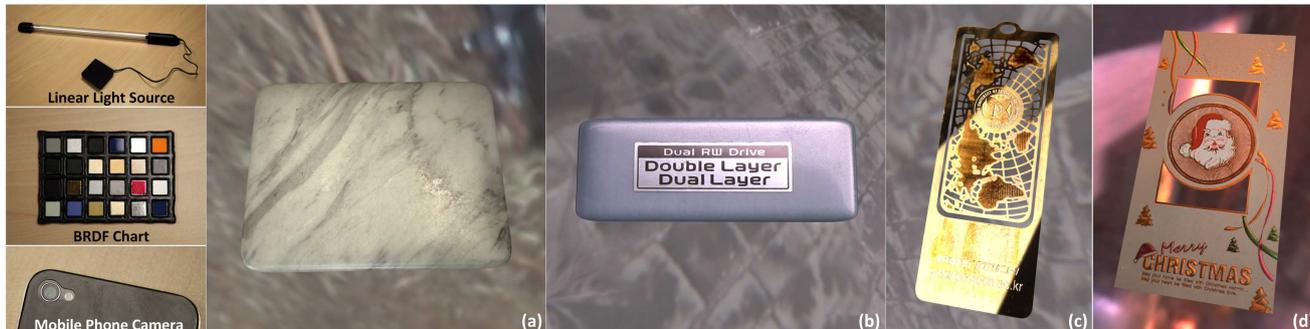


Figure 1: We capture spatially-varying, isotropic reflectance in about half a minute of casual scanning using three simple tools shown on the far left. Rendered results from four captured examples are shown on the right.

Abstract

We present a simple, fast solution for reflectance acquisition using tools that fit into a pocket. Our method captures video of a flat target surface from a fixed video camera lit by a hand-held, moving, linear light source. After processing, we obtain an SVBRDF.

We introduce a *BRDF chart*, analogous to a color “checker” chart, which arranges a set of known-BRDF reference tiles over a small card. A sequence of light responses from the chart tiles as well as from points on the target is captured and matched to reconstruct the target’s appearance.

We develop a new algorithm for BRDF reconstruction which works directly on these LDR responses, without knowing the light or camera position, or acquiring HDR lighting. It compensates for spatial variation caused by the local (finite distance) camera and light position by warping responses over time to align them to a specular reference. After alignment, we find an optimal linear combination of the Lambertian and purely specular reference responses to match each target point’s response. The same weights are then applied to the corresponding (known) reference BRDFs to reconstruct the target point’s BRDF. We extend the basic algorithm to also recover varying surface normals by adding two spherical caps for diffuse and specular references to the BRDF chart.

We demonstrate convincing results obtained after less than 30 seconds of data capture, using commercial mobile phone cameras in a casual environment.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture;

Keywords: BRDF chart, dynamic time warping (DTW), local linear embedding, reflectance sequence/response, spatially varying BRDF (SVBRDF)

Links: [DL](#) [PDF](#)

1 Introduction

Even neglecting wavelength dependence, an object’s spatially-varying reflectance is a complex, 6D function: its SVBRDF. Realis-

tic reflectance is critical for convincing CG rendering. Capturing it from real world targets remains a challenging problem that requires expensive hardware and slow scanning and processing.

Our goal is to make reflectance acquisition easy for almost anyone. More ubiquitous reflectometry engenders applications that customize virtual environments, with materials captured from each user’s own home, workplace, or places he might visit. Examples include user design of personalized car body finishes and decals in a racing game, or scanning of fabric and upholstery samples by individual clothing and furniture makers for e-commerce preview. Essentially, we seek a more accessible SVBRDF design type, which can be chosen and tuned with little more difficulty than textured region fills in a 2D drawing program.

Our method takes a video of the target, along with a reference BRDF chart, under a moving light. We use a *linear* light source [Gardner et al. 2003] to adequately sample highlights on most targets via a simple 1D movement from periphery to overhead. This measurement yields a 1D (per rgb channel) reflectance response over time for each chart tile, called a *representative*, and for each target point. At each target point, we match over a *neighborhood* or set of similar representative responses using a distance metric that performs temporal warping to compensate for the variation of view and light directions over an extended target. We then compute an overall diffuse and specular coefficient as well as an optimal blending of specular components over this neighborhood. Our BRDF chart is designed for generality by condensing a large measured database but could also be specialized to smaller domains such as textiles, fabrics, building materials, etc.

Our contributions include a new method for reflectance reconstruction that works directly on LDR reflectance responses without measuring HDR lighting or knowing the camera’s position. We develop a new method to match responses between target points and chart references based on dynamic time warping [Sakoe and Chiba 1978] which compensates for spatial variation in view and light direction over the target and handles saturated pixels. We introduce the idea of a BRDF chart and demonstrate one method for its manufacture. We leverage *local reconstruction* from manifold bootstrapping [Dong et al. 2010], but increase the fitting power of a small set of representatives by performing the local reconstruction over purely specular components and allowing arbitrary per-point variation of diffuse and specular coefficients. We also extend our basic approach to recover variation in surface normal.

Our method captures isotropic SVBRDFs and surface normal variation, making reflectance acquisition easy with inexpensive tools that can be carried anywhere in a pocket.

2 Related Work

2.1 BRDF Acquisition

A gonioreflectometer directly measures the BRDF at a single surface point by densely sampling the angular domain of light and view directions. A single camera and light source can be moved [Dana et al. 1999] or multiple cameras and light sources mounted over a spherical dome [Moshe et al. 2008]. Other solutions leverage a curved mirror [Ward 1992; Dana 2001; Ghosh et al. 2007; Mukaigawa et al. 2007] or a condenser lens [Dong et al. 2010]. Directly measure datas have been increasingly employed in recent rendering and editing works [Lawrence et al. 2006; Xu et al. 2008; Wang et al. 2009; An and Pellacini 2008; Xu et al. 2009].

Image-based methods capture a single BRDF from a homogeneous curved surface of known geometry by varying directional lighting [Lu et al. 1998; Marschner et al. 1999; Ngan et al. 2005]. Non-directional lighting has also been applied. Ramamoorthi and Hanrahan [2001] infer BRDF and illumination from sparse views of a homogeneous sphere. Romeiro et al. [2008] model an isotropic BRDF as a general bivariate function of two angles (between normal and halfway vectors, and light and halfway vectors) to acquire a BRDF from a single HDR image of a homogeneous sphere and known environmental lighting captured by a light probe. Most recently, Romeiro and Zickler [2010] exploit the statistics of real-world illumination to estimate a BRDF from a single image of a homogeneous sphere under unknown lighting. Holroyd et al. [2010] acquire both reflectance and geometry of a homogeneous curved surface using a specialized coaxial optical scanner with spatially modulated light source.

These approaches all capture the BRDF of a homogeneous target but cannot be easily extended to spatially varying reflectance.

2.2 SVBRDF Acquisition

Spatial gonioreflectometers directly measure SVBRDFs [McAllister et al. 2002; Lawrence et al. 2006], BTFs [Dana et al. 1999; Han and Perlin 2003; Muller et al. 2005], and surface reflectance fields [Debevec et al. 2000; Garg et al. 2006]. These methods require a dedicated device setup and careful calibration. The scanning is data-intensive and lengthy.

Polarization-based methods capture surface reflectance using polarized light to separate specular from diffuse reflectance and estimate each component’s albedo. Ma et al. [2007] estimate normal

maps from images under a set of polarized, spherical gradient lighting configurations. Ghosh et al. [2009] further estimate specular roughness and anisotropy using a set of second-order spherical gradient illuminations. Recently, Ghosh et al. [2010] apply a single illumination measurement using a uniform, circularly-polarized, spherical light to estimate specular roughness and Fresnel coefficients. These methods require a specialized light dome and accurate orthogonal orientation of the camera’s polarizing filters.

Image-based methods fit BRDF model parameters at each surface point using images captured over different view and light directions. Some work [Lensch et al. 2003; Goldman et al. 2005; Holroyd et al. 2010] applies linear combinations over a small basis set derived from parametric models. This approach regularizes sparse and noisy data but often fits real targets poorly. Alldrin et al. [2008] apply a more general model based on the two-angle bivariate function to reconstruct the SVBRDF and surface normal from ~ 100 images which vary the light direction but fix a single view. Zickler et al. [2005] use a radial BRDF basis to reconstructing smoothly-varying surface reflectance from samples in each local region. These methods all require accurate knowledge of view and light direction and careful movement of the light source to adequately sample each point’s specular lobe.

Another class of methods applies a dense lighting scan. Gardner et al. [2003] capture a flat target by scanning a linear light source and imaging from a fixed camera. Given the camera and light source position at each frame, an isotropic Ward model is fit to each pixel’s 1D reflectance response. Geometric details are recovered by laser strip scanning combined with two scans of the light source in a diagonal orientation. Wang et al. [2008] extend to anisotropic reflectance using a similar setup that replaces the linear light source by an LED array, performs a 2D scan over lighting direction, and merges similar BRDF data from different pixels.

Like [Gardner et al. 2003], our method is single-view and performs a dense 1D scan of a linear light source. However, it manipulates each pixel’s 1D reflectance response directly without prior knowledge of light or camera position, to simplify scanning. We also model reflectance using real-world representatives that provide more realistic angular details than does a parametric model.

Example-based methods recover reflectance and normal variation with the help of reference samples or objects. Hertzmann and Seitz [2003] recover surface normals by capturing reference objects of homogeneous reflectance and known geometry together with the target. Each target point’s BRDF is represented as a linear combination of reference object BRDFs. Points on reference objects having the same normal are grouped and matched to each target point to determine its normal. Treuille et al. [2004] extend this method with voxel coloring to reconstruct full geometric models from multi-view data. These methods need no knowledge of camera or lighting pose, but require that reference BRDFs be similar to the target’s to resolve ambiguity between the normal and the reflectance variation. Our method captures more general SVBRDFs using a greater number of reflectance representatives, but on a flat target. It also handles local viewing and lighting.

Matusik et al. [2003] represent an isotropic BRDF as a linear combination of measured representatives. A realistic BRDF spanned by their database can be reconstructed from 800 measurements. Weyrich et al. [2006] apply this model to reconstructing human skin SVBRDFs from dense measurements. Dong et al. [2010] present a general, two-pass method which reconstructs each target point as a local linear combination of (nearby) representatives acquired in a separate scanning pass. Our BRDF chart representatives are generic and limited in number and so yield a less accurate fit than the spe-

cialized BRDF scanning of the actual target. The benefit of our method is a much faster and simpler capture session that need only take a short video of uncalibrated light motion.

3 Pocket Solution Overview

This section summarizes our pocket reflectometry approach. Later sections then detail our algorithms for reflectance sequence alignment and SVBRDF reconstruction.

3.1 Device Setup

As shown in Figure 2(c), our device setup consists of three compact and portable components.

A *hand-held linear light source* is a 40cm fluorescent tube powered by a 12V battery. During capture, the tube is manually scanned over the surface. We also experimented with a flexible LED strip, but the resulting light was not uniform enough for capturing highly specular materials like brass.

A *BRDF chart*, shown in Figure 2(b), comprises a palette of known BRDF samples used as representatives for reconstruction. Our BRDF chart contains 24 flat tiles, each about 0.5mm thick and glued onto a 4mm thick plastic base. The chart measures 8.5cm×5.5cm, little bigger than a credit card. It is placed alongside and imaged with the target surface.

The chart contains two material types. One is a nearly ideal Lambertian material that provides the diffuse reference and also calibrates the camera’s exposure and white balance. The tile is a 1cm² square of the Spectralon diffuse standard (albedo = 80%) from Labsphere, Inc. The rest of the tiles, also 1cm², consist of different specular materials whose choice is discussed in Section 3.3.

A *video camera* captures reflected light from the target and chart as the light is moved. We use a mobile phone held in a car mount with flexible stand. Careful adjustment of camera settings is not required; uncontrolled auto-exposure and white-balance defaults on mobile phone cameras work well. Of course, a high resolution camera is desirable to obtain high spatial detail in the resulting SVBRDF. Our experiments capture 1280 × 720 video at 30fps.

3.2 BRDF Model

Let the 3D unit vectors \mathbf{i} and \mathbf{o} denote the light and view direction respectively. We represent a BRDF ρ as a linear combination of a normalized Lambertian BRDF α and a normalized specular BRDF β shared by all three color channels, via

$$\rho(\mathbf{i}, \mathbf{o}) = d\alpha(\mathbf{i}, \mathbf{o}) + s\beta(\mathbf{i}, \mathbf{o}), \quad (1)$$

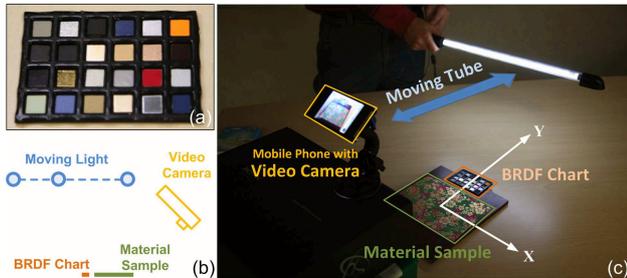


Figure 2: Device setup: (a) BRDF chart, (b) schematic, (c) photo.

$$\alpha(\mathbf{i}, \mathbf{o}) = \frac{1}{4\pi^2}, \quad \int_{\Omega^+} \int_{\Omega^+} \beta(\mathbf{i}, \mathbf{o}) d\mathbf{i} d\mathbf{o} = 1, \quad (2)$$

where $\int_{\Omega^+} \rho(\mathbf{i}, \mathbf{o}) d\mathbf{i} \leq 1$, and $d \geq 0$ and $s \geq 0$ are three-channel diffuse and specular coefficients, respectively. While the normalized Lambertian BRDF α is constant, the normalized specular BRDF β is complex and affected by several factors, including surface roughness, Fresnel reflection, and shadowing, that cause nonlinear variation over the target.

We assume this specular variation forms a low-dimensional manifold which can be reconstructed by local linear embedding using a limited number of representatives. Denote the set of BRDF chart samples as $\{\rho_i\}$, with corresponding normalized specular components $\Phi = \{\beta_i\}$. Our model represents the BRDF $\rho(\mathbf{x})$ at each surface point \mathbf{x} using overall diffuse and specular coefficients and a local linear combination over Φ , via

$$\rho(\mathbf{x}) = d(\mathbf{x})\alpha + s(\mathbf{x}) \sum_{j=1}^k u_j(\mathbf{x})\beta_j, \quad \beta_j \in \Phi(\beta), \quad (3)$$

where $\Phi(\beta)$ denotes the k -nearest neighbors in Φ of the normalized specular part of $\rho(\mathbf{x})$. The weights $u_j(\mathbf{x})$ are non-negative and sum to 1, for each \mathbf{x} . In our current implementation, $k=8$.

3.3 BRDF Chart Design

Representative BRDFs on the chart must be carefully chosen to permit the reconstruction of different targets. We develop an algorithm that computes a small “spanning” set of representatives from a given BRDF database of interest. The computation is based only on normalized specular BRDF components and ignores variation in diffuse and specular coefficient, which is handled separately.

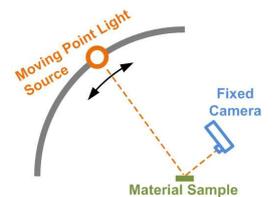
Given a set of normalized specular BRDFs $\{\beta_i\}$, we choose a subset Ψ for which local linear embedding can accurately reconstruct any member of the database. We first construct the k -nearest neighbor graph for the database, and find the most distant pair of samples $\{\beta_1, \beta_2\}$ based on greatest geodesic distance in the graph. We then select $2k$ additional samples by iteratively picking the sample having the maximum geodesic distance to all samples accumulated so far. Beyond these, each further sample is chosen that has maximum reconstruction error from Equation 3. The error is measured by the weighted L2-norm:

$$E^2 = \frac{\int_{\Omega^+} \int_{\Omega^+} (\mathbf{i} \cdot \mathbf{n})^2 (\beta^*(\mathbf{i}, \mathbf{o}) - \beta(\mathbf{i}, \mathbf{o}))^2 d\mathbf{i} d\mathbf{o}}{\int_{\Omega^+} \int_{\Omega^+} (\mathbf{i} \cdot \mathbf{n})^2 d\mathbf{i} d\mathbf{o}}, \quad (4)$$

where \mathbf{n} is the surface normal. We then place (physical) material tiles for the selected representatives on our BRDF chart. An alternative approach [Matusik et al. 2009] is to print the selected BRDFs as regular tiles.

Our experiments are based on a BRDF database of 100 homogeneous isotropic materials. We performed the above computation using an error threshold of 2.5% to obtain 24 BRDF chart representatives. Tiles are arranged as shown in Figure 3. Refer to supplementary material for more information about the materials selected for the BRDF chart.

These materials were measured with a gonioreflectometer similar to [Murry-Coleman 1990] as shown in the right figure. The gonioreflectometer densely samples the 1D polar angle of lighting direction with a fixed camera. The light source is 2 meters away from the material and the camera is about 1 meter away. The camera is tilted at about 45° and



plaster	sliver paint	rubber	polished acrylic	aluminium	fluorescent paint
matte tape	black paper	polished resin	bronze	bronze metallic paint	acrylic
plastic	brass	coated metallic paint	polyethylene	red metallic paint	alumina
80% Spectralon	leather	matte golden paint	alum-bronze	tinfoil	lactoprene

Figure 3: BRDF chart tile arrangement. Further details about these materials are included in supplementary material.

its precise position is calibrated before data capturing based on [Zhang 2000]. Positions of the moving light source are determined by precision mechanical control. We uniformly sample 90 images while the light source moves from top ($\theta = 0^\circ$) to bottom ($\theta = 90^\circ$), which takes about ten minutes. The captured data was then fit using the general microfacet BRDF model [Ashikhmin et al. 2000] with a tabulated 1D normal distribution function. The specular component was separated by subtracting the minimum reflectance value from the measured BRDF.

3.4 Data Capture and Preprocessing

Figure 2(c) illustrates our capture setup. We denote the plane of the flat target as XY with its center at the origin. A typical target measures $10\text{cm} \times 10\text{cm}$. The camera is placed 50cm away from the origin along Y and 40cm above it, making an angle of roughly 40 degrees with the Z axis. The chart adjoins the target in the X direction. The user holds the light 100cm above the sample, with its length aligned to X, and slowly moves it back and forth from $Y=-100\text{cm}$ to $Y=20\text{cm}$. The above measurements are approximate, and light movement need not be controlled accurately. Roughly 30 seconds of video is captured, generating 900 images.

We process this image sequence by first calibrating intensity based on a pre-computed gamma (power law) curve for the camera’s response. Reflected light from the background environment is removed by subtracting a frame with the light source off. This yields the *partially-calibrated sequence* $\check{r}(t)$. Reflectance responses depend on the target location \mathbf{x} ; our notation makes that dependence implicit.

We then divide the partially-calibrated sequence by the Lambertian tile’s sequence, $\check{r}_*(t)$, at each time sample t :

$$\tilde{r}(t) = \check{r}(t) / \check{r}_*(t). \quad (5)$$

This compensates for time-varying camera exposure, lighting distance variation, and approximate cosine factor in the radiance integral. We assume the Lambertian response does not saturate and is independent of local viewing and lighting effects. We call the resulting $\tilde{r}(t)$ the *calibrated sequence*.

4 Reflectance Sequence Alignment

Casual capture presents several difficulties for matching reflectance responses captured on the target and chart. Limited camera zoom and light brightness require that both devices be placed close to the target. View and light positions are thus local and their corresponding directions vary spatially over the target. The light source is moved manually; its speed and orientation may vary over the scan. Camera and light positions are both unknown. The target BRDF is also unknown and varies spatially. Finally, we must handle saturated pixels in the LDR measurement.

We solve these problems by introducing a robust alignment algorithm for reflectance responses based on dynamic time warping. The key observation is that variations caused by local view and light position can be approximately matched by shifting and possibly scaling the 1D reflectance sequence in time. For robust match-

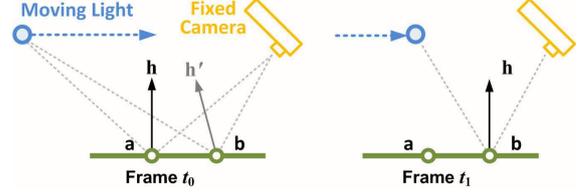


Figure 4: Directional variation as time shifting. The reflection half-vector \mathbf{h} at surface point a and frame time t_0 does not equal the one at surface point b due to viewing and lighting locality. Instead, b reaches that half-vector at a slightly different time t_1 as the light is moved.

ing, response sequences are normalized and their saturated samples repaired before the alignment.

4.1 Directional Variation as Time Shifting

Our setup implies variation of light and view direction over the target. Responses captured at two different points thus differ even when their reflectance is identical. We match sequences by assuming that reflectance is isotropic and mainly determined by the *half-vector*, \mathbf{h} , midway between the light and view directions. This assumption is suggested by the microfacet model after neglecting smooth factors based on shadowing (S) and Fresnel effects (F):

$$\rho(\mathbf{i}, \mathbf{o}) = \frac{S(\mathbf{i}, \mathbf{o}) F(\mathbf{i}, \mathbf{o}) N(\mathbf{h})}{4(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})} \quad (6)$$

where $\mathbf{h} = (\mathbf{i} + \mathbf{o}) / \|\mathbf{i} + \mathbf{o}\|$ and \mathbf{n} is the surface normal. An isotropic BRDF further restricts the normal distribution function (NDF), N , to be a 1D function of angle, θ , between \mathbf{h} and \mathbf{n} .

Because a linear light source is held parallel over a flat acquisition target, each time point in the moving light scan yields just a small range for (well-lit) θ , at each target point \mathbf{x} . But this dominant angle varies as a function of \mathbf{x} because of light/view locality.

Figure 4 illustrates the situation with two target points a and b . At time t_0 , a exhibits a highlight since its half-vector \mathbf{h} aligns with the (upward) normal. Because of view/light locality, b has a different half-vector and no highlight. As the light source moves, b attains this same half-vector and highlight at a different time, t_1 . This motivates a temporal warping approach to compensate for locality and non-uniformity in a manual light scan.

Detailed processing steps are explained in the following sections, and illustrated in Figure 5.

4.2 Normalization

We normalize the calibrated responses to eliminate the effect of a spatially-varying albedo, via

$$\hat{r}(t) = \frac{\tilde{r}(t) - \tilde{r}_{\min}}{\tilde{r}_{\max} - \tilde{r}_{\min}}, \quad (7)$$

where \tilde{r}_{\min} and \tilde{r}_{\max} are the minimal and maximal values of $\tilde{r}(t)$. For robustness, minimal/maximal values are determined by averaging over the 10 values of least/greatest magnitude. The resulting *normalized sequence*, $\hat{r}(t)$, is used as input for alignment.

We replace saturated values with an estimate generated by fitting a 1D Gaussian to neighboring, non-saturated values. The fitting uses 10 to 20 samples on each side. This is done before the above normalization and only for matching. It has no further effect on BRDF reconstruction, which blends previously-measured representatives.

4.3 Alignment by Dynamic Time Warping

Dynamic Time Warping (DTW) [Sakoe and Chiba 1978] is widely used in speech recognition and other time-series analysis for measuring similarity or aligning two discrete sequences. Given a reference sequence $p(t)$ of length n (to warp to) and a source sequence $q(t)$ of length m (to warp from), the algorithm seeks a pair of corresponding warp functions, $(w_p(t), w_q(t))$, that minimize the sum of element-wise differences between them:

$$\min_{w_p, w_q} \sum_{t=1}^h D((p(w_p(t))), q(w_q(t))), \quad (8)$$

where h is the common length of the two sequences after warping. As shown in the inset, this amounts to finding a minimal path from $(1, 1)$ to (n, m) in the $n \times m$ discrete 2D domain where each grid point (i, j) represents distance between $p(i)$ and $q(j)$. The problem can be solved by dynamic programming.

We use a distance function tailored to our application of matching reflectance responses:

$$D(p(i), q(j)) = (p(i)^\gamma - q(j)^\gamma)^2 + \lambda C_m, \quad (9)$$

where p and q are two normalized 1D reflectance sequences. The first term attempts to align specular peaks. A power function is used to suppress non-peak values; we set $\gamma = 3$. The second term C_m penalizes temporal shrinking or expansion (scaling). We use a small

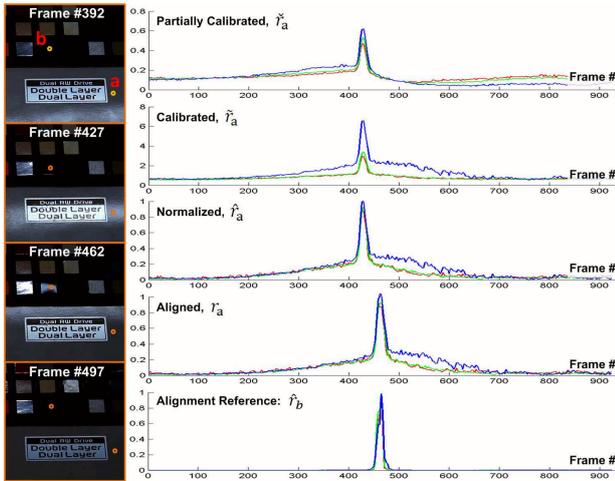


Figure 5: Processing for reflectance sequences. The source sequence at a is processed and aligned to the reference sequence at b .

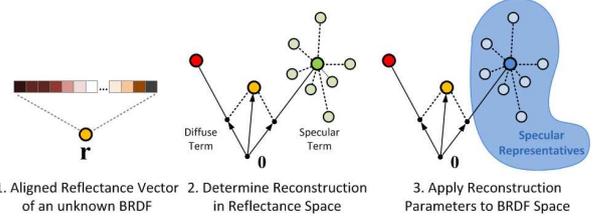


Figure 6: BRDF Reconstruction based on aligned reflectance vectors.

weight λ , equal to 1% of the average intensity, which avoids undue scaling in noisy but low-amplitude regions. To preserve the falloff shape of the source sequence's specular peak, we permit no stretching in regions where the intensity exceeds 50% of its maximum. Detailed algorithm pseudocode is included in the Appendix.

We then apply the resulting warps to align q to p , via

$$q'(t) = q(w_{q \rightarrow p}(t)), \quad w_{q \rightarrow p}(t) = w_q(w_p^{-1}(t)). \quad (10)$$

The aligned sequence is finally resampled to the frame rate of the original video.

To obtain a canonical peak layout for BRDF reconstruction, we align reflectance sequences of all pixels to a single specular tile on the BRDF chart. We choose the one having the highest unsaturated response peak (Figure 5, bottom).

The resulting *aligned sequence* is denoted $r(t)$ (without tilde or hat) and the corresponding time-sampled vector as \mathbf{r} . Essentially, DTW provides a robust method of recentering each sequence around the peak of the canonical reference sequence chosen above.

5 BRDF Reconstruction

After alignment and resampling, the resulting reflectance sequences $r(t)$ are effectively responses from the same, infinitely-distant but time-varying environmental lighting $\{L_t(\mathbf{i}), t = 1, 2, \dots, n\}$ and a constant view direction, \mathbf{o} . The resulting sequence thus characterizes the BRDF at each target point and is called a *reflectance vector*. Reflectance vectors comprise h components, where h is about 900.

The reflectance function $r(t)$ at some location is given by

$$r(t) = \int_{\Omega^+} L_t(\mathbf{i}) [d \alpha(\mathbf{i}, \mathbf{o}) + s \beta(\mathbf{i}, \mathbf{o})] (\mathbf{n} \cdot \mathbf{i}) \, d\mathbf{i} \quad (11)$$

We separate diffuse and specular terms to yield

$$\mathbf{r} = d\mathbf{a} + s\mathbf{b}, \quad (12)$$

$$a(t) = \int_{\Omega^+} L_t(\mathbf{i}) \alpha(\mathbf{i}, \mathbf{o}) (\mathbf{n} \cdot \mathbf{i}) \, d\mathbf{i}, \quad b(t) = \int_{\Omega^+} L_t(\mathbf{i}) \beta(\mathbf{i}, \mathbf{o}) (\mathbf{n} \cdot \mathbf{i}) \, d\mathbf{i},$$

where reflectance vectors \mathbf{a} and \mathbf{b} represent purely diffuse and specular responses, respectively.

The normalized reference diffuse response is given by $\mathbf{a} = \mathbf{r}_*/d_*$ where \mathbf{r}_* is the response of the Lambertian reference tile and d_* is its albedo. The normalized specular response for reference material j is then given by

$$\mathbf{b}_j = \frac{\mathbf{r}_j - d_j \mathbf{a}}{s_j}, \quad (13)$$

where \mathbf{r}_j is the reference's aligned response and d_j and s_j are its diffuse and specular coefficients.

To reconstruct an unknown BRDF ρ with its measured reflectance vector \mathbf{r} , we minimize the reconstruction error as shown in Figure 6, via

$$\min_{u_0, u_1, \dots, u_k} \left\| \mathbf{r} - u_0 \mathbf{a} - \sum_{j=1}^k u_j \mathbf{b}_j \right\|, \quad u_j \geq 0, \quad \mathbf{b}_j \in \Phi(\mathbf{r}). \quad (14)$$

Given the set of \mathbf{r} 's nearby reference responses, $\Phi(\mathbf{r})$, the above is a well-conditioned least squares problem with $k+1$ unknowns and h equations ($h \gg k$). To ensure non-negativity in the weighting, we solve (14) by quadratic programming with non-negative parameter constraints [Lawson and Hanson 1974] if negative values appear in the unconstrained least-squares solution.

Since the neighborhood is unknown, we try all possible representative neighborhoods. These are precomputed by randomly generating convex linear combinations of reference specular responses and finding their k -nearest neighborhood sets. We try 10,000 combinations, resulting in about 20 unique neighborhoods for $k=8$.

Saturated values in reflectance vectors cause problems when solving Equation 14. Since the equations are already highly constrained, we can just eliminate any saturated components in the vectors \mathbf{r} and \mathbf{b}_j . Typically, only 1-3% of the roughly 900 components are removed.

Equation 14 can be applied separately to the reflectance sequences of each color channel. This result yields significant color noise. We obtain a more robust color estimate by first computing an RGB diffuse and specular coefficient and then solving for the linear combination of references. This shares the same specular falloff shape among all three channels.

We average over the 10 values of least brightness in the RGB sequence to obtain the diffuse component (d'_r, d'_g, d'_b) , and over the 10 greatest brightness samples to get the specular component (s'_r, s'_g, s'_b) . The diffuse coefficient is then normalized: $(d_r, d_g, d_b) = (d'_r, d'_g, d'_b) / \|(d'_r, d'_g, d'_b)\|$. The specular coefficient is likewise normalized via

$$(s_r, s_g, s_b) = \frac{(s'_r, s'_g, s'_b) - (d'_r, d'_g, d'_b)}{\|(s'_r, s'_g, s'_b) - (d'_r, d'_g, d'_b)\|}. \quad (15)$$

We then substitute these estimated RGB coefficients into Equation 14 to solve

$$\min_{u_0, u_1, \dots, u_k} \left\| \begin{bmatrix} \mathbf{r}_r \\ \mathbf{r}_g \\ \mathbf{r}_b \end{bmatrix} - u_0 \begin{bmatrix} d_r \\ d_g \\ d_b \end{bmatrix} \mathbf{a} - \sum_{j=1}^k u_j \mathbf{b}_j \right\|, \quad (16)$$

$$u_j \geq 0, \quad \mathbf{b}_j \in \Phi(\mathbf{r}).$$

6 Bumpy Surface Extension

To recover normal variation, we extend the BRDF chart with two *spherical cap references* shown in Figure 7(a). A spherical cap covers a portion of the sphere that is less than a hemisphere. One Lambertian and one specular reference provide the corresponding responses for a range of surface orientations. Response at every target point is then matched with those from different points on the spherical cap references; the best match determines the normal.

Unlike [Hertzmann and Seitz 2003], our method matches based on temporal warping to obtain a metric that is largely material-invariant, and so does not rely on references whose reflectance most closely match the target's. This is because pixels on the spherical cap have identical reflectance (specular lobe shape) and differ only in their responses' peak location, while DTW provides a robust way of aligning each sample's pixel response, of arbitrary reflectance, to these references. Other work in normal estimation for specular

surfaces [Chen et al. 2006; Holroyd et al. 2008] applies directional lighting and requires a known view and light direction.

For robust normal recovery, we need an extra pass of light scanning in the orthogonal (X) direction as well as the original (Y), as shown in Figure 7(c). Two-axis light scanning is required to distinguish what is an intrinsically 2D (unit) surface normal.

Each point's surface normal is initialized to the best-matching Lambertian reference. We apply a distance metric based on an inner product:

$$\max_j \frac{\check{\mathbf{r}} \cdot \check{\mathbf{r}}_j}{\|\check{\mathbf{r}}\| \|\check{\mathbf{r}}_j\|}, \quad (17)$$

where $\check{\mathbf{r}}$ is the reflectance sequence to be matched and $\check{\mathbf{r}}_j$ are the candidate Lambertian reference sequences at different normals. Note that partially calibrated sequences $\check{\mathbf{r}}$ are used here instead of (fully) calibrated ones since the surface orientation is unknown. This simple method suffices to recover normals at nearly diffuse target points. Points that are more specular must be processed further as described in the following. We deem a point specular if its maximum value is 10% larger than its average, after dividing by the diffuse reference sequence with the best-matching normal.

Like BRDF reconstruction, normal recovery for a specular point is based on its calibrated sequence $\check{\mathbf{r}}$ as described in Section 3.4. But in this case, it should be normalized by dividing it by the sequence from the diffuse spherical cap reference *having the same normal*. Since that normal is unknown, we apply an iterative procedure. We perform normalization based on the previous step's normal. We then match this normalized response against the specular cap reference using a material-invariant metric from the DTW algorithm in Section 4. It measures overall temporal distortion from the target sequence $p(t)$ to the reference $q(t)$ via

$$\frac{1}{n} \sum_{t=1}^n \left(\frac{t-1}{n-1} - \frac{w_{q \rightarrow p}(t)}{n-1} \right)^2, \quad (18)$$

where $w_{q \rightarrow p}$ is the warp function from source to reference sequence defined in Equation 10. This metric cares only about the position of the specular peak and ignores its height and falloff. The chosen normal corresponds to the reference having minimum temporal distortion. These two steps (first, intensity-normalize by the Lambertian response of the previous step's normal, and second, select the specular reference normal having minimum distortion) are iterated until convergence, usually in 2-5 steps.

The above implies a binary decision: a point is either diffuse or specular. To suppress spatial artifacts from a hard boundary, we instead blend between the two recovered normals, \mathbf{n}_d and \mathbf{n}_s , for mostly diffuse points that have maximum value 5-15% larger than the average. The blend is based on estimated specularity η of the material, via

$$\mathbf{n} = (1 - \eta) \mathbf{n}_d + \eta \mathbf{n}_s, \quad \eta = P(\lambda(\bar{r}_{\max}/\bar{r}_{\text{avg}} - 1.1)), \quad (19)$$

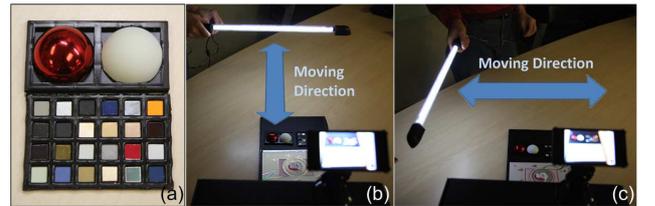


Figure 7: Extensions for normal recovery: (a) BRDF chart with two added spherical cap references, (b,c) two-axis light movement.

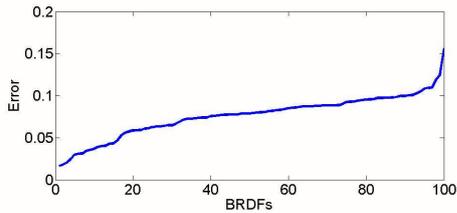


Figure 8: Validation of our method on 100 measured BRDFs. RMS error (y axis) is plotted against BRDF index (x axis). BRDFs are sorted in order of increasing reconstruction error.

where P is the sigmoid function switching from 0 to 1, and \tilde{r}_{avg} denotes the average value of the sequence $\tilde{r}(t)$. The parameter $\lambda = 100$ determines the sharpness of the transition from diffuse to specular.

After recovering the normal, each target point’s normalized sequence $\hat{r}(t)$ (based on its recovered normal) is then aligned with the reference sequence from the BRDF chart and used to reconstruct a BRDF, as described in Sections 4.3 and 5.

To eliminate light/view locality problems in the above processing, we assume locality effects are smooth (low-frequency) over the target while normal variation over a flat target lacks any low-frequency component. After recovering the normal map as described above, we therefore estimate and remove its low-frequency variation. This is done by fitting a second-order b-spline separately to the map’s x and y components. The z component of the smooth normal is then obtained by $z^2 = 1 - x^2 - y^2$. At each surface point, a compensation rotation matrix is determined by the original normal \mathbf{n} and its b-spline fit \mathbf{n}_b by rotating around $\mathbf{n}_b \times \mathbf{n}$ by an angle of $-\arccos(\mathbf{n} \cdot \mathbf{n}_b)$. The final normal then applies this rotation to \mathbf{n} at every surface point.

7 Results

Our experiments use the video camera on Apple’s iPhone 4 16G. Capture takes around 30 seconds and produces 900 images at resolution 1280×720 . Light source movement should be slow and

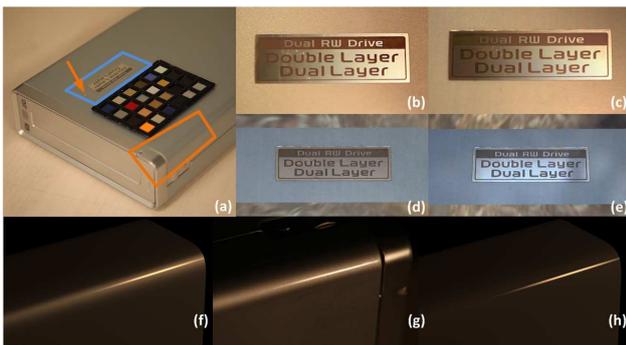


Figure 9: Results with a DVD burner example, shown in (a). The top right two columns compare a real photo (b) to our reconstruction (c). The middle row (d,e) shows renderings of our reconstruction with natural environmental lighting applied. The bottom row compares our reconstruction (f) with one based on Cook-Torrance model fitting (h), and a real photo in (g). For both (f) and (h), a geometric model for the entire case was rendered using BRDFs reconstructed at the single point marked with the orange arrow in (a). The blue box shows the area of the target displayed in (b,c,d,e). The orange box shows the region of the DVD case rendered in (f,g,h).



Figure 10: Reflectance sequence alignment in reconstruction: (a) ground truth, (b) reconstruction with alignment, (c) reconstruction without alignment. Alignment is necessary to obtain a good match.

steady to ensure that specular materials are adequately sampled. A movement pass of about 30 seconds suffices.

We manually localize the BRDF chart and the pre-modeled geometry of spherical cap references in the captured video frame. The output SVBRDF is also rectified manually. This manual work could be avoided by adding markers on the BRDF chart and spherical cap card to allow their automatic detection and estimate the camera’s projection.

We then perform reconstruction on a PC with two Intel Core™2 Quad CPUs and 8GB memory. Total data processing time per example is about 20 minutes and is fully automatic. This includes tens of seconds for preprocessing, 15 minutes for reflectance sequence alignment, and 5 minutes for final BRDF reconstruction.

To validate our method, we tested it on synthetic data generated by the database of 100 BRDFs acquired from real materials in Section 3.3. We simulated the geometric configuration exactly as described in Section 3.1. Images of a $10\text{cm} \times 10\text{cm}$ virtual target comprising all 100 tiles were rendered as lit by a moving linear light source, and then processed by our method. Results showing error between the reconstructed result and the original BRDF are shown in Figure 8. Error is computed using the metric in [Ngan et al. 2005], similar to Equation 4 but weighted by a solid angle correction factor and followed by an overall square root. Errors across materials are normalized by the maximum albedo of each BRDF. Our reconstruction matches well, with average RMS error (across all 100 samples) of 8% and worst case RMS error of 16%.

Figure 12 shows results for three different materials: a shiny wallpaper (top), a weathered cold-electroplated copper plate (middle),

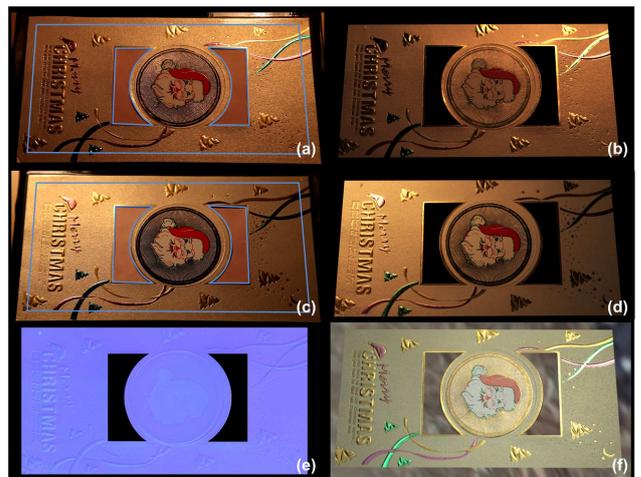


Figure 11: Results for a bumpy example: (a/c) real photo, (b/d) rendered result of reconstructed SVBRDF with normal map in (e), (f) rendered result under natural environmental lighting.

and a block of marble (bottom). The figure's left two columns compare appearance between the original surface (left) and our reconstruction (right), at a new view direction (nearly top-view) and lit by an incandescent bulb. Our results's appearance closely matches the original in all three examples. Just for these experiments, we captured the lighting environment with a light probe so that we could compare renderings of our reconstruction with photos of the original material. This data is not used in our reconstruction. The view used in this comparison is nearly overhead while in capturing it was about 40 degrees away from overhead.

Figure 12's right two columns show rendered results in a natural (measured) lighting environment, using two different orientations of the surface with respect to the lighting. Our method accurately captures detailed spatial variation in these materials, with behavior that convincingly responds to changing light and view. Notice the spatially-varying oxidation in the copper plate and pattern of specular veining, scratches, and smudges on the marble.

Figure 9 shows results for a homogeneous plastic DVD burner case with a small metallic label, photographed in (a). The top row compares an original photo (b) with a synthetic rendering using our reconstruction (c). The middle row (d,e) shows two different renderings in measured environmental lighting. The bottom row compares our result with a fit based on the Cook-Torrance model [Cook and Torrance 1982]. The parameter fitting is done as described in [Ngan et al. 2005]. We sampled the BRDF at one point on the case (marked by the orange arrow), and then compared the two reconstructions, assuming homogeneous reflectance of the whole case geometry. For this experiment, we built a geometric model of the DVD case by hand. The best-fitting parametric model exhibits an unduly sharp highlight (h) that poorly matches the original (g); our method (f) is based on real-world references and better preserves this glossy highlight.

Figure 10 compares reconstruction with and without temporal alignment of reflectance responses. We expect temporal warping to be crucial for highly specular materials. Even for the less shiny material used in this experiment, specular coefficients are significantly underestimated without alignment in (c).

Figure 11 shows an example captured with our bumpy surface extension from Section 6. It is a greeting card with colorful and raised metallic glazing. After initialization with the diffuse spherical cap reference, 23% of surface points required further refinement by matching with the specular spherical cap reference. The whole computation took about 4 hours. We compare real photos (a/c) with the rendered result of our method (b/d) at two different view and light configurations. A synthetic rendering in measured natural lighting is shown in (f).

Please refer to the accompanying video for animations under different viewing and lighting conditions.

8 Conclusions and Future Work

We have presented a new method for high-quality SVBRDF capture that works quickly, in a casual setting, using inexpensive and portable tools.

Our technique is subject to a number of limitations. Processing times are large, but we have not yet significantly optimized the computation. A fixed view is sufficient for reflectance reconstruction, but hampers normal recovery for bumps oriented away from the camera. Specular reflectance is complex, and a limited number of tiles can span only a limited part of it. Specializing the BRDF chart to a particular domain better samples variations in the specular lobe's precise shape and provides a better reconstruction. Our method for selecting representatives supports such spe-

cialized databases. Our approach can not easily be extended to handle anisotropic reflectance and tangent rotations. We observe that anisotropic reflectance is a much larger and more complex space than isotropic reflectance. Although our method succeeds in visually accurate SVBRDF reconstruction under casual conditions, it does not completely eliminate problems from local viewing and lighting. A local view point in our setup causes time scaling in the pixel sequence \check{r} , ultimately yielding about 5% average and 13% maximal reconstruction error over a target of size 10cm \times 10cm. We expect that further research can reduce or eliminate this error.

Acknowledgements

The authors would like to thank Steve Lin for helpful discussions, Xiaoming Fu for helping data acquisition, and Matt Scott for proofreading the paper and dubbing the video. We appreciate helpful suggestions and comments from anonymous reviewers. We also recommend the freeware iFunbox (<http://www.i-funbox.com/>) for efficiently transferring files from the iPhone to PC.

References

- ALLDRIN, N., ZICKLER, T. E., AND KRIEGMAN, D. 2008. Photometric stereo with non-parametric and spatially-varying reflectance. In *Conf. on Comp. Vision and Pattern Recognition (CVPR)*, 1–8.
- AN, X., AND PELLACINI, F. 2008. Approp: all-pairs appearance-space edit propagation. *ACM Transactions on Graphics* 27, 40:1–40:9.
- 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.
- CHEN, T., GOESELE, M., AND SEIDEL, H.-P. 2006. Mesostructure from specularity. In *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, CVPR '06, 1825–1832.
- COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1, 7–24.
- 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., 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.
- DONG, Y., WANG, J., TONG, X., SNYDER, J., LAN, Y., BEN-EZRA, M., AND GUO, B. 2010. Manifold bootstrapping for SVBRDF capture. *ACM Trans. Graph.* 29, 4, 98:1–98:10.
- 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.

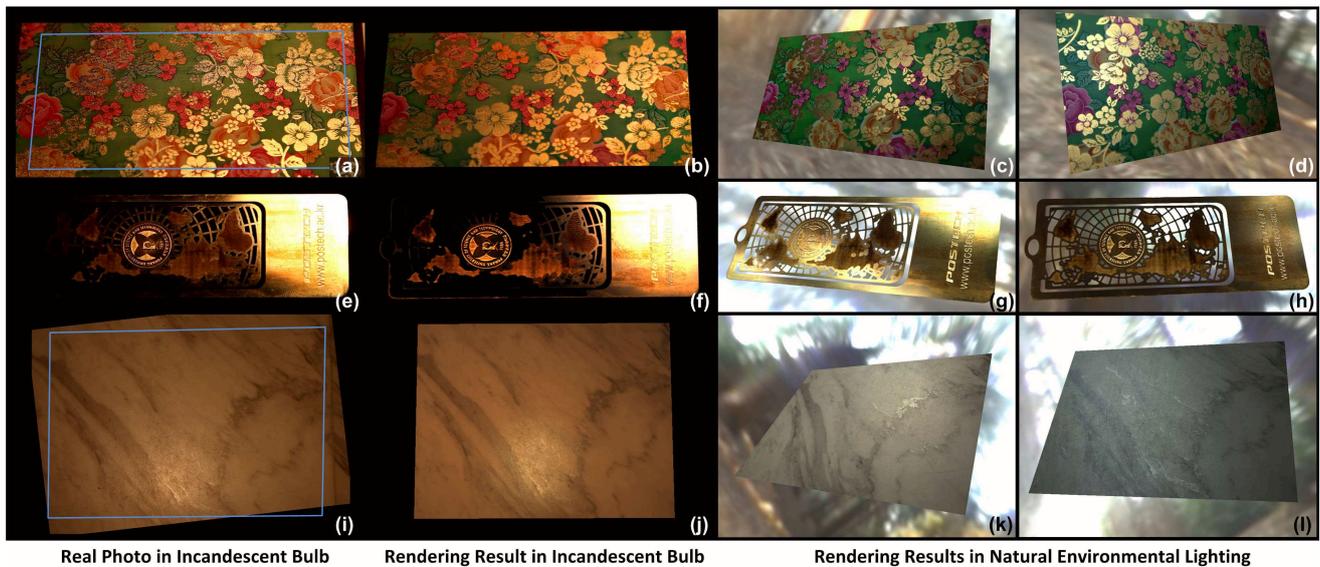


Figure 12: Results with three materials: wallpaper, electroplated copper, and marble. The left two columns compare original photos with a synthetically rendered result of our method. The right two show two different renderings using measured environmental lighting.

- GHOSH, A., ACHUTHA, S., HEIDRICH, W., AND O'TOOLE, M. 2007. BRDF acquisition with basis illumination. *Computer Vision, IEEE International Conference on*, 1–8.
- GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. E. 2009. Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Computer Graphics Forum* 28, 4, 1161–1170.
- GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. 2010. Circularly polarized spherical illumination reflectometry. *ACM Trans. Graph.* 29 (December), 162:1–162:12.
- 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.
- HAN, J. Y., AND PERLIN, K. 2003. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. Graph.* 22, 3, 741–748.
- HERTZMANN, A., AND SEITZ, S. M. 2003. Shape and materials by example: A photometric stereo approach. In *IN PROC. IEEE CONF. COMPUTER VISION AND PATTERN RECOGNITION*, 533–540.
- HOLROYD, M., LAWRENCE, J., HUMPHREYS, G., AND ZICKLER, T. 2008. A photometric approach for estimating normals and tangents. *ACM Transactions on Graphics* 27, 5, 1–9.
- HOLROYD, M., LAWRENCE, J., AND ZICKLER, T. 2010. A coaxial optical scanner for synchronous acquisition of 3d geometry and surface reflectance. *ACM Trans. Graph.* 29 (July), 99:1–99:12.
- 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).
- LAWSON, C., AND HANSON, R. 1974. Solving least squares problems. *Prentice-Hall Chapter* 23, 161.
- 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.
- LU, R., KOENDERINK, J. J., AND KAPPERS, A. M. L. 1998. Optical properties (bidirectional reflectance distribution functions) of velvet. *Applied Optics* 37, 25 (Sept.), 5974–5984.
- MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C. F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. *Rendering Techniques*, 183–194.
- MARSCHNER, S., WESTIN, S., LAFORTUNE, E., TORRANCE, K., AND GREENBERG, D. 1999. Image-based BRDF measurement including human skin. In *10th Eurographics Rendering Workshop*.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. Efficient isotropic BRDF measurement. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 241–247.
- MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)* 28, 5 (Dec.).
- 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.
- MOSHE, B.-E., WANG, J., BENNETT, W., LI, X., AND MA, L. 2008. An LED-only BRDF measurement device. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, 1–8.
- MUKAIGAWA, Y., SUMINO, K., AND YAGI, Y. 2007. High-speed measurement of BRDF using an ellipsoidal mirror and a projec-

tor. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, 1–8.

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.

MURRSY-COLEMAN, J.F., A. S. 1990. The automated measurement of BRDFs and their application to luminaire modeling. *Journal of the Illuminating Engineering Society, Winter*.

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

RAMAMOORTHY, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '01, 117–128.

ROMEIRO, F., AND ZICKLER, T. 2010. Blind reflectometry. In *Proceedings of the 11th European conference on Computer vision: Part I*, Springer-Verlag, Berlin, Heidelberg, ECCV'10, 45–58.

ROMEIRO, F., VASILYEV, Y., AND ZICKLER, T. 2008. Passive reflectometry. In *Proceedings of the 10th European Conference on Computer Vision: Part IV*, Springer-Verlag, Berlin, Heidelberg, 859–872.

SAKOE, H., AND CHIBA, S. 1978. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26, 43–49.

TREUILLE, A., HERTZMANN, A., AND SEITZ, S. M. 2004. Example-based stereo with general BRDFs. In *European Conference on Computer Vision*, 457–469.

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.

WANG, J., REN, P., GONG, M., SNYDER, J., AND GUO, B. 2009. All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Transactions on Graphics* 28, 5, 133:1–133:10.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, ACM Press, New York, NY, USA, 265–272.

WEYRICH, T. 2006. Acquisition of human faces using a measurement-based skin reflectance model. In *PhD thesis, Department of Computer Science, ETH Zurich*, vol. 16741.

XU, K., JIA, Y.-T., FU, H., HU, S.-M., AND TAI, C.-L. 2008. Spherical piecewise constant basis functions for all-frequency precomputed radiance transfer. *IEEE Transactions on Visualization and Computer Graphics* 14, 2, 454–467.

XU, K., LI, Y., JU, T., HU, S.-M., AND LIU, T.-Q. 2009. Efficient affinity-based edit propagation using k-d tree. *ACM Transactions on Graphics* 28, 5, 118:1–118:6.

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: Adapted DTW

We apply dynamic time warping to align a source sequence $q(t)$ to a reference sequence $p(t)$. The distance metric is given by:

$$D(p(i), q(j)) = b(p(i)^\gamma - a(j)^\gamma)^2 + \lambda \max(C_I[i, j] - 1, C_D[i, j] - 1, 0)^2. \quad (20)$$

C_I/C_D denote the number of insertion/deletion operations on the minimum path to (i, j) , and thus penalize temporal expansion or shrinking. To preserve the falloff shape of source peaks, we prohibit temporal scaling (insertion or deletion) when the intensity of the source sequence exceeds τ .

```

DTW_Align( p[1..n], q[1..m]; Dsum, wp, wq )
// workspace: minimum summed difference along all possible path
E[0..n, 0..m] ← ∞,   E[0, 0] ← 0
// workspace: counters of sustained insertion & deletion
CI[0..n, 0..m] ← 0,   CD[0..n, 0..m] ← 0
// search path with minimum summed difference
for i in [1..n]
for j in [1..m]
  if p[j] > τ // enforce align by match
    CI[i, j] ← 0,   CD[i, j] ← 0
    E[i, j] ← D(p[i], q[j]) + E[i - 1, j - 1])
  else
    Epath ← min(E[i - 1, j], E[i, j - 1], E[i - 1, j - 1])
    if E[i - 1, j] is the minimum // align by insertion
      CI[i, j] ← CI[i - 1, j] + 1,   CD[i, j] ← 0
    if E[i, j - 1] is the minimum // align by deletion
      CD[i, j] ← CD[i, j - 1] + 1,   CI[i, j] ← 0
    if E[i - 1, j - 1] is the minimum // align by match
      CI[i, j] ← 0,   CD[i, j] ← 0
    E[i, j] ← D(p[i], q[j]) + Epath
// summed difference on the minimum path
Dsum ← E[n, m]
// back trace the minimum path
path ← [(n, m)] // the path
while path[0] ≠ (0, 0)
  (i, j) ← path[0]
  λ ← argmin(iλ, jλ) ∈ Λ E[iλ, jλ];   Λ = {(i - 1, j), (i - 1, j - 1), (i, j - 1)}
  path ← [λ] ∪ path // concatenate the path
// construct backward warping
for k in [1..|path|]
  (wp[k], wq[k]) ← path[k]

```

Figure 13: Pseudo-code for dynamic time warping.