Eurographics Symposium on Rendering 2009
Hendrik Lensch and Peter-Pike Sloan
(Guest Editors)

Volume 28 (2009), Number 4

Estimating Specular Roughness and Anisotropy from Second Order Spherical Gradient Illumination

Abhijeet Ghosh Tongbo Chen Pieter Peers Cyrus A. Wilson Paul Debevec

University of Southern California,
Institute for Creative Technologies

Abstract

This paper presents a novel method for estimating specular roughness and tangent vectors, per surface point, from polarized second order spherical gradient illumination patterns. We demonstrate that for isotropic BRDFs, only three second order spherical gradients are sufficient to robustly estimate spatially varying specular roughness. For anisotropic BRDFs, an additional two measurements yield specular roughness and tangent vectors per surface point. We verify our approach with different illumination configurations which project both discrete and continuous fields of gradient illumination. Our technique provides a direct estimate of the per-pixel specular roughness and thus does not require off-line numerical optimization that is typical for the measure-and-fit approach to classical BRDF modeling.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism—

1. Introduction

Measuring the appearance of real materials is an active research area in computer graphics. Often the appearance is described by the *bidirectional reflectance distribution function (BRDF)* [NRH*77], a 4D function that relates the ratio of reflectance between the incident and outgoing directions for a single surface point. Usually these BRDF models depend on a sparse set of non-linear parameters that roughly correspond to albedo, specular roughness, surface normal, and tangent directions. Measuring and fitting these parameters for a particular material model often requires a dense sampling of incident and outgoing lighting directions. In many cases fitting the parameters of an a-priori chosen BRDF model to the observed measurements relies on complex fragile non-linear optimization procedures.

Recently a number of methods have been proposed that estimate fundamental parameters of appearance, such as normal direction and albedo, without assuming an a-priori material model. Instead they rely on general properties shared by many physical materials such as symmetry [ZBK02, AK07, MHP*07, AZK08, HLHZ08]. The work presented here is most related to [MHP*07], where it is shown that the first order spherical statistics of the reflectance under distant illumination correspond to the normal and reflection vector for diffuse and specular materials respectively. Furthermore, Ma et al. [MHP*07] demonstrate that the first order statistics can be efficiently measured using linear gradient illumination conditions. In this work, we build upon Ma et al.'s work to estimate a *per-pixel specular roughness* using polarized *second order* spherical gradients that provide measurements of the *variance* about the mean (i.e., reflection vector). We demonstrate that for isotropic BRDFs, only three additional axis-aligned second order spherical gradient illumination patterns are sufficient for a robust estimate of per pixel specular roughness (Figure 1). We also demonstrate that by using the five second order spherical harmonics, related to the second order spherical gradient illumination patterns, reliable estimates of the specular roughness and tangent directions of general anisotropic BRDFs are possible. Using a lookup table, these estimated high order statistics can then be directly translated to the parameters of any BRDF model of choice.

An example of a direct application of our method is the estimation of spatially varying reflectance parameters of arbitrary objects. Furthermore, since the proposed method relies on only up to nine distinct illumination conditions with

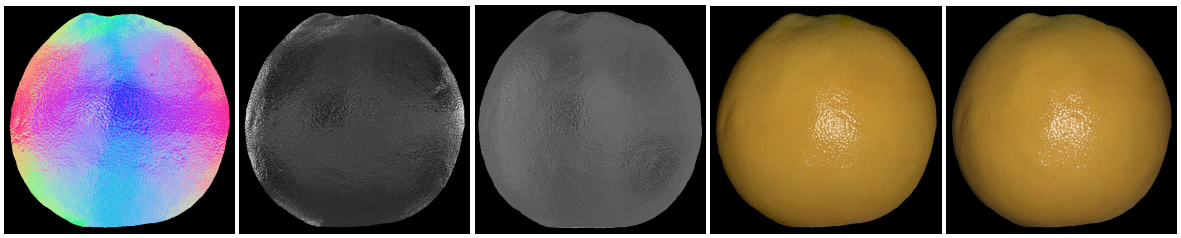

(a) specular normals (b) specular albedo (c) specular roughness (d) rendering (e) photograph

Figure 1: *Specular reflectance properties ((a)-(c)) of a plastic orange estimated using polarized second order spherical gradient illumination conditions. The estimated specular roughness map (c) is used as the per-pixel distribution for a Torrance-Sparrow BRDF to create a rendering (d) that closely matches the validation photograph (e).*

minimal capture time, it is amenable to capturing per-surface point roughness parameters of human subjects.

2. Related Work

A large body of work covers the representation and acquisition of the appearance of physical materials and objects. We focus this discussion on reflectance modeling and related surface properties that influence the appearance of a material or an object (e.g., surface normal and tangent directions). We can broadly classify prior work in two categories: analytical BRDF-based methods, and non-parametric and photometric methods.

Analytical BRDF-based Methods: In the past few decades a large number of analytical material representation models have been developed. Each of these BRDF models are designed with certain goals in mind, such as physical accuracy [TS67, HTSG91], facilitating rendering [APS00], and versatility and flexibility in approximating physical materials [Ash06, LFTG97, War92].

Numerous methods have been proposed that capture the spatially varying appearance, based on the above analytical BRDF models, to recreate complete digital copies of existing real world objects (e.g., [GTHD03, Geo03, GCHS05, LGK*01, LKG*03, Mar98]). Each of these methods requires either a special acquisition device or an acquisition scheme tailored towards a particular prefixed analytical material model, and often sacrifice spatial variation for angular variation to reduce the total acquisition time. However, the choice of BRDF model can impact the quality of appearance reproduction significantly [NDM05]. Furthermore, fitting the model-specific parameters is often a complicated and ill-conditioned process due to the non-linear nature of the parameters and the presence of measurement noise. In many cases it only becomes clear after attempting to fit the measured data that the choice of material model may be suboptimal. Switching to a better suited BRDF model after the fact is often difficult due to the inherent reliance of the acquisition setup/scheme of these methods on a particular model.

Our method in contrast enables the acquisition of general appearance statistics per surface point without prior knowledge of the object geometry and without assuming a pre-selected BRDF model. The acquisition scheme is fast and optimized to capture these general statistics, which can then be subsequently mapped to most analytical BRDF parameter spaces.

Non-parametric and Photometric Techniques: Classical photometric stereo [Woo78] estimates surface normals by assuming an underlying Lambertian material, and using a small set of fixed viewpoint observations under point lighting. However, most materials are not purely Lambertian, and thus an inaccurate surface normal is estimated. As a result photometric stereo has been extended to non-Lambertian materials (e.g., [Geo03,GCHS05]). While these methods can handle a wider range of material types, they still rely on (mostly isotropic) analytical BRDF models that limit their generality.

To overcome this limitation, a number of techniques have been proposed that avoid using parametric BRDF models. Mallick et al. [MZKB05] reduce a general material to Lambertian by removing the specular "component", and subsequently apply traditional photometric stereo. Hertzmann and Seitz [HS03] estimate surface normals using a reference object with known shape and similar material properties as the target object. While this method does not rely on parametric appearance models, it requires a reference object which is not always available.

Another class of methods exploit general properties of surface reflectance to infer surface statistics. A common assumption is that the maximum reflectance is observed when the halfway vector coincides with the surface normal (e.g., [FCM*08]). Various forms of symmetry constitute another property that has been exploited extensively. Zickler et al. [ZBK02] exploit the Helmholtz reciprocity to recover depth and normal directions. Alldrin and Kriegman [AK07] exploit the symmetry about the view-normal plane in case of isotropic BRDFs. Ma et al. [MHP*07] assume symmetry about the mean vector of a BRDF observed from a fixed viewpoint. They show that for Lambertian reflections the mean corresponds to the surface normal, and for specular reflections to the reflected direction. Both can be efficiently and directly measured by using polarized spherical gradi-

ent illumination. Holroyd et al. [HLHZ08] assume a similar type of symmetry, but use a dense sampling to resolve both the normal direction as well as tangent vectors per surface point.

While these methods do not rely on a parametric model, they do not provide complete information regarding the surface reflectance. Lawrence et al. [LBAD*06] used inverse shade trees and an optimization scheme coined ACLS to decompose the spatially varying material properties of planar samples from dense hemispherical samplings in a collection of 1D curves and 2D textures. Alldrin et al. [AZK08] also employ ACLS to compute bivariate representations of isotropic surface appearances, including surface normals. Finally, Zickler et al. [ZREB06] share reflectance information from different surface points to create a dense non-parametric reconstruction of appearance.

All of the above methods either require a dense sampling of the lighting directions, integrate information over multiple surface points, or deliver incomplete appearance information. The method proposed in this work extends [MHP*07] to capture the second order statistics of surface reflection independently per surface point. We show that these second order statistics correspond to specular roughness and the tangent vectors for specular reflections. These statistics can be readily captured by extending the linear spherical gradients of Ma et al. with second order spherical gradients. As such it requires only a few observations, while delivering statistics suitable for creating complete appearance descriptions.

3. Theoretical Background

In this section, we introduce the necessary notations and definitions that we need to infer specular roughness and anisotropy from second order gradient illumination. We first recap the definitions of moments on general 1D functions, which are then extended to a spherical domain. Armed with these definitions, we then show how these moments can be used to infer specular roughness. Finally, we show the connection between spherical harmonics and spherical moments, and how they can be used to measure and infer roughness and anisotropy.

0th, 1st, and 2nd Moments In statistics the zeroth, first and second moments of a general 1D function $\mathbf{f}(x)$ correspond to the total energy α, mean μ, and variance σ^2 of that function. These moments can be directly computed from the inner-products of the function $\mathbf{f}(x)$ and a constant function ($\mathbf{g}(x) = 1$), linear gradient ($\mathbf{g}(x) = x$), and quadratic function ($\mathbf{g}(x) = x^2$), denoted by L_0, L_1, and L_2 respectively:

$$\begin{aligned} \alpha &= \int \mathbf{f}(x)\,dx, \\ &= L_0, \end{aligned} \tag{1}$$

$$\begin{aligned} \mu &= \int x\frac{\mathbf{f}(x)}{\alpha}\,dx, \\ &= \frac{1}{\alpha}\int x\mathbf{f}(x)\,dx, \\ &= \frac{L_1}{L_0}, \end{aligned} \tag{2}$$

$$\begin{aligned} \sigma^2 &= \int (x-\mu)^2\frac{\mathbf{f}(x)}{\alpha}\,dx, \\ &= \frac{1}{\alpha}\left(\int x^2\mathbf{f}(x)\,dx - 2\mu\int x\mathbf{f}(x)\,dx + \mu^2\int \mathbf{f}(x)\,dx\right), \\ &= \frac{1}{\alpha}\left(L_2 - 2\mu L_1 + \mu^2 L_0\right), \\ &= \frac{L_2}{L_0} - \frac{L_1^2}{L_0^2}. \end{aligned} \tag{3}$$

0th, 1st, and 2nd Spherical Moments We can extend these moments to the spherical domain by redefining L_0, L_1, and L_2 on the sphere Ω. This can be compactly denoted using the following vector notation:

$$L_0 = \int_\Omega \mathbf{f}(\vec{\omega})\,d\vec{\omega}, \tag{4}$$

$$L_1 = \int_\Omega \vec{\omega}\,\mathbf{f}(\vec{\omega})\,d\vec{\omega}, \tag{5}$$

$$L_2 = \int_\Omega \vec{\omega}\vec{\omega}^T\,\mathbf{f}(\vec{\omega})\,d\vec{\omega}, \tag{6}$$

where $\vec{\omega} = [\omega_x, \omega_y, \omega_z] \in \Omega$, and each integration is with respect to solid angle. Thus: $\omega_x^2 + \omega_y^2 + \omega_z^2 = 1$. Note that $\vec{\omega}\vec{\omega}^T$ is the generalization of the quadratic function x^2 to the spherical domain, and is a symmetric 3×3 matrix. Furthermore L_0 is a scalar, L_1 is a vector of length 3, and L_2 is a 3×3 symmetric matrix. Applying these to Equations (1), (2), and (3), yield the 0th, 1st, and 2nd order spherical moments. The three linear spherical gradient of Ma et al. [MHP*07] correspond to $\vec{\omega}$. Likewise, $\vec{\omega}\vec{\omega}^T$ defines six new second order spherical gradients. On the diagonal of this matrix are the quadratic gradients ω_x^2, ω_y^2, and ω_z^2. On the off-diagonal elements are the mixed linear gradients: $\omega_x\omega_y$, $\omega_x\omega_z$, and $\omega_y\omega_z$.

Specular Roughness Ma et al. [MHP*07] demonstrated that for specular reflections, the zeroth and first order moments correspond to the specular albedo and reflection vector. In this work, we argue that the second moment under specular reflections is directly proportional to specular roughness. Note that this "specular roughness" is independent of a chosen analytical BRDF model. While the exact relation to a particular BRDF model's specular parameters is highly model dependent, it is still instructive to verify the correlations of the 2nd moment and the specular parameters of an analytical model. For example, consider the Ward BRDF [War92]:

$$\mathbf{R}(\vec{h}) = c \cdot e^{\left(\frac{h_x}{h_z\sigma_x}\right)^2 + \left(\frac{h_y}{h_z\sigma_y}\right)^2}, \tag{7}$$

where c is a normalization constant, $\vec{h} = (h_x, h_y, h_z)$ is the halfway vector between incident $\vec{\omega}_i$ and outgoing $\vec{\omega}_o$ di-

rections, and σ_x and σ_y are the anisotropic specular roughness parameters. By definition, $\mathbf{R}$ is normalized, and consequently $\alpha = 1$. For simplicity, lets assume that the local shading frame (reflected direction + tangent directions) are known, and the spherical gradients are defined in this frame (this can be accomplished by an appropriate rotation such that the reflection direction is aligned with $[0,0,1]$). Then, $\mu = [0,0,1]$. In this case Equation (3) simplifies to $\sigma^2 = L_2$, and:

$$L_2 = \int_\Omega \vec{\omega}\vec{\omega}^T \mathbf{R}(\vec{\omega})\, d\vec{\omega} \sim \begin{pmatrix} \sigma_x^2 & 0 & 0 \\ 0 & \sigma_y^2 & 0 \\ 0 & 0 & \sigma_z^2 \end{pmatrix}. \quad (8)$$

Intuitively, Equation (7) closely resembles a normal distribution for which the 2nd moment corresponds to the variance (the diagonal elements). The off-diagonal element in L_2 are zero because the gradients are axis aligned. Note that σ_z^2 depends on the values of σ_x^2 and σ_y^2. The exact value of σ_z^2 does not have a direct physical meaning, hence we will ignore this value. Practically, the above states that the observed radiance is proportional to the specular roughness, when illuminating a surface point with shading frame aligned quadratic gradients ω_x^2 and ω_y^2.

2nd Order Spherical Harmonics The above discussion assumes that the gradients are aligned with the local shading frame. However, this shading frame is most likely not known beforehand. Furthermore, every surface point has a potentially different shading frame orientation. Additionally, the six second order spherical gradients do not form an orthogonal spherical basis, and as such are not optimal in terms of number of patterns. Ideally, we would like to capture the responses under some optimal (i.e., orthogonal) canonical spherical basis illumination, from which the responses of gradients aligned with the local shading frame can be computed during processing for each surface point.

A well known set of orthogonal spherical basis functions are *spherical harmonics*. Spherical harmonics can be thought of as the spherical generalization of Fourier functions [SKS02]. An interesting property of spherical harmonics is that they are invariant under rotation, similar to how Fourier functions are invariant under translation. While the first order spherical gradients correspond exactly to the first order spherical harmonics [MHP*07], second order spherical gradient do *not* correspond exactly to the second order spherical gradients (i.e., there are six second order spherical gradients, but only five second order spherical harmonics). However, there is some overlap, in particular, the second order spherical harmonic that only depends on the azimuthal angle corresponds exactly to ω_z^2 (up to a scale factor, and assuming that the azimuthal angle is defined with respect to the $\omega_z = 1$ axis). Due to the rotation invariance, rotations of spherical harmonics are just linear combinations of the same functions [SKS02]. Thus, ω_z can be aligned to any axis using only rotations, and as a result, any second order spherical gradient (including ω_x^2 and ω_y^2) can be expressed using second order spherical harmonics only.

Practically, this implies that by capturing the response of a surface point under second order spherical harmonic illumination, we can compute what the response of that surface point would be under any second order spherical gradient due to linearity of light transport. Furthermore, the rotation from world coordinates to shading frame is solely determined by the reflected direction and the tangent directions. As shown by Ma et al. [MHP*07], the reflected direction corresponds to the first moment (acquired using linear spherical gradients). The tangent directions are then defined as the other two principal vectors of the specular lobe. Let the reflected vector correspond to the Z axis of some rotated frame with arbitrary orthogonal X and Y axes. It is unlikely that the tangent directions will exactly correspond to these arbitrary X and Y axes. However, we can still compute the responses of the second order spherical gradients ω_x^2 and ω_y^2 in this frame. Furthermore, computing the response under $\omega_x\omega_y$ allows us to create a 2×2 covariance matrix. Computing the principal vectors of this covariance matrix yields the tangent directions. The magnitude of these principal vectors correspond to σ_x^2 and σ_y^2.

If the underlying material is isotropic, then the exact orientation of the tangent vectors does not matter, and $\sigma_x = \sigma_y$. In this case, any rotated frame (with Z corresponding to the reflected direction) is sufficient to compute the specular roughness, and thus we only need to compute ω_x^2 for an arbitrary X axis orthogonal to Z. This allows us to capture the specular roughness using a subset of just three second order spherical harmonics. These are the spherical harmonics that are proportional to Z^2, XZ and $X^2 - Y^2$ (i.e., $m = \{0, -2, +2\}$), omitting the spherical harmonics proportional to XY and YZ (i.e., $m = \{+1, -1\}$). We can then rotate the spherical harmonics frame such that the Z axis is perpendicular to the reflected direction. The computed response of this rotated Z^2 spherical harmonics corresponds to the isotropic specular roughness.

The above analysis ignores the effects of Fresnel reflections, masking and shadowing, foreshortening, and albedo. As in Ma et al. [MHP*07], we will assume that foreshortening varies little over the specular lobe, and thus can be considered a constant scale factor. The effects of the albedo can be easily compensated by dividing by the zeroth moment (i.e., albedo). Note that unless special precautions are taken, the Fresnel effects will be "baked" into the zeroth moment, and thus a division by the zeroth moment will also compensate for the Fresnel effects on the second moment. Masking and shadowing will be "baked" into the second moment as well.

4. Error Analysis

In this section, we analyze the accuracy and limits of our proposed technique for estimating specular roughness and

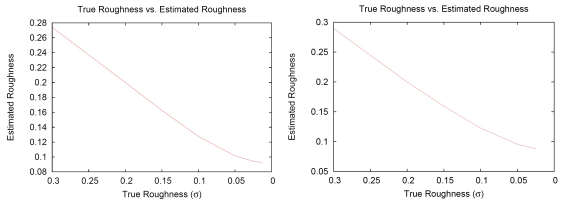

(a) Ward model (b) Ashikhmin model

Figure 2: *Simulation of estimated specular roughness using second order spherical gradient illumination for various roughness parameters. As seen in the plots, the estimates are accurate up to roughness parameter $\sigma = 0.1$, beyond which the roughness parameter is overestimated.*

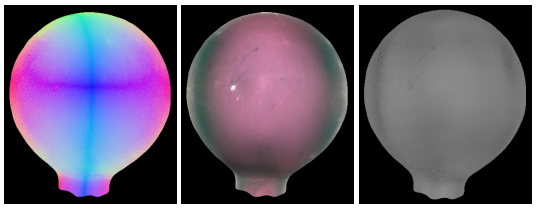

(a) spec. normals (b) spec. albedo (c) spec. roughness

Figure 3: *Viewpoint independence of the estimated specular roughness for a glossy ornament. The division by the albedo removes any Fresnel effects at grazing angles.*

anisotropy. We also analyze the effects of Fresnel reflectance and cosine factors in the estimation of specular roughness. For this purpose, we carried out simulations of the effects of second order spherical gradient illumination on analytic BRDFs (Ward and Ashikhmin) with various roughness parameters. Figure 2, (a) demonstrates that the roughness estimation for the Ward BRDF model is indeed linearly proportional as predicted by Equation (8), and is accurate up to roughness $\sigma = 0.1$, beyond which the method overestimates the computed roughness due to the finite resolution of the simulation. For highly specular BRDFs, there is very little difference in the computed integrals over the specular lobe for the first and second order spherical gradients which impacts the precision of the estimate. Figure 2, (b) shows a similar plot for the Ashikhmin BRDF model with Fresnel reflectance $F_0 = 0.1$, demonstrating that the division by the zeroth term removes any Fresnel effects from the specular roughness estimation. This is also seen in real acquired data in Figure 3, which illustrates the view independence of the estimated roughness parameters. Our simulations also found the cosine factor to not have any significant effect on the roughness estimate. Finally, the recovery of the tangents of an anisotropic BRDF is also demonstrated in Figure 4.

5. Measurement and Analysis

In this section we apply the theory derived in Section 3 to the measurement of specular roughness (and tangent vectors) of physical objects. We first discuss the different measurement setups and the required calibration. Next, we describe the algorithms for computing the roughness and tangent vectors. Finally, an analysis on the required lighting resolution of the measurement setups is discussed before presenting the results of several acquired objects in Section 6.

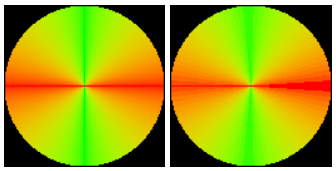

(a) true tangent (b) estimated

Figure 4: *Simulation of tangent frame estimation for a flat anisotropic disk with rotating tangent frame. The simulation uses an anisotropic Ward BRDF with roughness parameters $\sigma_x = 0.3$ and $\sigma_y = 0.1$.*

5.1. Measurement Setup and Calibration

We employ three different measurement setups to illustrate the effectiveness our method. Using each of the three devices we illuminate an object using the zeroth (i.e., constant), first (i.e., linear), and second order spherical gradients. Depending whether or not the target object only contains isotropic reflections, we just emit 3 or all 5 second-order spherical gradients. This yields a total of 7 and 9 illumination conditions respectively.

The first setup consists of an LED sphere with approximately 150 individually controllable lights. Each light is covered with a linear polarizer in the pattern of [MHP*07]. We observe the object placed at the center of the sphere through a polarized camera placed outside the sphere. We record the effects of the gradient illumination patterns on the object under both cross and parallel polarization conditions in order to separate the diffuse and specular reflections.

The second setup uses an LCD monitor as a controllable light source. The object of interest is placed at a small distance from the LCD screen, aimed such that all reflections observed through the camera from the object see the LCD screen. Because the monitor covers only a small part of the full sphere of incident directions, we need to carefully calibrate the exact mapping from LCD pixel to incident lighting direction on the object. This is achieved by placing a mirrored sphere at the location of the object, and computing a one-to-one correspondence between pixels on the sphere (i.e., directions) and pixels on the LCD. Even though the LCD monitor does not cover the full sphere of incident directions, and thus the spherical illumination patterns are not completely emitted onto the object, a correct estimate of the second order statistics can still be made as long as the solid angle extended by the specular reflection from a surface point is completely included in the projected solid angle of the LCD monitor seen from the respective surface point. In such a case the observed radiance is the result of illumination incident from the extended solid angle, and thus

from a portion of the LCD monitor. Any illumination or lack thereof outside this extended solid angle does not influence the observed radiance for that surface point. Due to the limited angular support of specular BRDFs, this condition can be easily met by restricting the normal directions for which roughness parameters can be estimated. Diffuse and specular reflection can be separated, by exploiting the linear polarization of the LCD screen, and capturing the effects of each gradient pattern under a large set of (camera) polarization orientations. The maximum and minimum observed (per-pixel) intensities over the different polarization orientations correspond to the parallel and cross polarized observations, respectively.

The third setup is a rough specular hemisphere similar to [PHD06]. The object is placed near the center of the hemisphere, next to a projector equipped with a fish-eye lens. The object is observed through a hole in the apex of the hemisphere. Light emitted by the projector is reflected by the hemisphere onto the object, and subsequently observed by the camera. A similar geometric calibration as above is performed to ensure we have a one-to-one mapping between directions and projector pixels. We do not separate diffuse and specular reflections using this setup, and restrict the objects placed in this setup to ones exhibiting specular reflections only.

The first setup has the advantage that it covers the full sphere of incident lighting directions, and thus can deal with objects of any geometry. The second setup is the most restrictive in this case, and is mostly suited for planar surfaces. In terms of separating the diffuse from the specular reflections, the first device is by far the easiest (only two photographs per lighting pattern are needed), followed by the second device. Due to the limited sampling density of the first device, it is only suited for objects with materials exhibiting rough specular reflections. The other two devices have a very dense sampling, and can deal with almost any type of specular reflection.

5.2. Practical Methodology

We first discuss how we compute specular roughness and tangent vectors from measurements in the various setups described above. For isotropic materials, in principle we only require measurements under 7 illumination conditions. This assumes we can steer the first and second order spherical gradients to the appropriate directions as discussed in Section 3. However, in practice this may not be possible depending upon the measurement setup. In our case, the LCD panel setup and the reflective dome do not cover the entire sphere of directions and hence are not suited for spherical harmonics based rotations. For such setups we approximate the specular roughness as the magnitude of the roughness measured along the orthogonal X and Y directions:

$$\sigma^2 = \| \sigma_x^2 + \sigma_y^2 \|, \tag{9}$$

where the reflection vector is aligned with the Z direction. This is a reasonable approximation for surfaces that are mostly flat. In order to enforce symmetry on the above approximation, we measure both the X and Y aligned first order gradients and the corresponding reverse gradients. The second order X and Y spherical gradients are symmetric, by definition, about these axes. This still leads to a total of 7 illumination conditions for such measurement setups. When computing the specular roughness according to Equation 3 with these measurements, we simply choose between the first order gradients and the reverse gradients depending upon the orientation of the reflection vector.

Measurements with the LED sphere setup, on the other hand, allow us to take advantage of the spherical harmonic rotations to compute the specular roughness more accurately according to the procedure discussed in Section 3. For anisotropic materials, in order to recover the tangent vectors in practice, we simply search for the local X and Y orientations with the largest anisotropy ($\frac{\sigma_x}{\sigma_y}$).

To relate the estimated model-independent roughness parameter to a particular model specific parameter, we follow one of the the following two strategies. For model specific parameters that are linearly dependent on the variance, we precompute the corresponding scale factor. For non-linear dependencies (e.g., the sharpness parameter for the Blinn-Phong BRDF), one could tabulate the non-linear mapping function, and perform a look-up (with interpolation) to find the corresponding roughness parameters.

5.3. Lighting Resolution Analysis

Before presenting some acquisition results in Section 6, we analyze the effects of the discrete lighting resolution of the employed LED sphere on the specular roughness estimation. Given that the employed LED sphere has 20 lights around the equator, it can resolve spherical harmonic frequencies of up to order 10 in the limit (Nyquist frequency). From the frequency space analysis of Ramamoorthi and Hanrahan [RH02], we know that the number of spherical harmonic frequencies F is related to the lobe width s of a Phong BRDF as $F \approx \sqrt{s}$, or to the width of a microfacet lobe as $F \approx 1/\sigma$. Given that a BRDF is defined over a hemisphere, the frequency F equals 5 in our case. Hence we obtain $s \approx 25$, or $\sigma \approx 0.2$ to be the limit of the specular roughness (around the reflection direction) that can be resolved by such a lighting apparatus.

Alternatively, one can also determine the limit of specular roughness through simulation of such discrete lighting on BRDFs with varying roughness parameters. For our simulations, we rendered a sphere with a Ward BRDF model with spatially uniform specular roughness under the various second order spherical gradient illumination conditions as represented on the LED sphere. Figure 5 presents the accuracy in terms of spatial variation in the estimated roughness

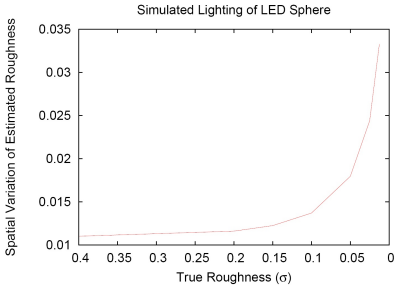

Figure 5: *Spatial variation in the estimated specular roughness as a function of decreasing roughness value (σ) with the simulated discrete lighting resolution of the LED sphere with* 18° *light spacing. As seen, the estimate is not reliable for very specular materials.*

under such discrete lighting. As can be seen, the estimate is stable up to a specular roughness value of around 0.2, beyond which the error in the estimate increases dramatically due to the discrete nature of the lighting. Thus, we can accurately measure roughness of relatively broad specular lobes with such a setup.

6. Results and Discussion

In this section, we present the results of estimated spatially varying specular roughness parameters of real objects captured with second order spherical gradient illumination. For anisotropic samples, we also recover the local tangent vectors. The recovered reflectance parameters are then used with several commonly used BRDF models to demonstrate the generality of the approach. For validation, we present comparisons of renderings with the estimated parameters to photographs under distant point lighting conditions. We generate our renderings using the hybrid normal map rendering technique of Ma et al. [MHP*07].

Figure 1 presents the various specular reflectance parameters recovered for an isotropic material, in this case a plastic orange, with the presented technique. The reflectance properties of the plastic orange were measured with 7 cross and parallel polarized spherical gradient illumination condition using the LED sphere setup described in Section 5.1. The image (d) shows a rendering of the synthetic orange with the recovered specular (and diffuse) reflectance properties under a point lighting condition that is a close match to a photograph (e) of the object under similar illumination. Here, we employ the recovered specular roughness parameters in a Torrance-Sparrow BRDF model while assuming Lambertian diffuse reflection for generating the rendering.

Figure 6 shows results of specular reflectance parameter estimation on two anisotropic samples: a red necktie (top-row), and a red satin pouch (bottom-row). Here, we employ all the first 9 spherical harmonic illumination conditions (i.e., 0th, 1st, and 2nd order) under both polarization states to recover the various specular reflectance parameters (a)-(c), as well as the local tangent (d) and bitangent vectors (e) according to the procedure discussed in Section 5.2. Note how the local surface and tangent orientations of such samples can be reliably estimated using only 9 illumination patterns using our technique, which is a significant improvement over previous techniques. We employ the recovered anisotropic reflectance parameters for the mildly anisotropic tie to a Ward BRDF model (top-row), and the Ashikhmin BRDF model for the satin pouch exhibiting significant anisotropy. The renderings of these objects with the recovered anisotropic parameters are again a close match to the validation photographs. The main differences between the renderings and the reference photographs are because the diffuse BRDF is most likely not Lambertian, and that the highlights are slightly broader due to the limited lighting resolution of the employed LED sphere.

Figure 7 presents more results for various spatially varying isotropic BRDFs captured in the various measurement setups discussed in Section 5.1. These include a female subject (top-row) and a wax candle (second-row) measured in the LED sphere setup, a flat spatially varying BRDF sample measured using an LCD panel as an emitter, and a (dark) specular object, an Obsidian sculpture, measured using a reflective hemispherical dome. The recovered reflectance parameters are applied to a Torrance-Sparrow BRDF model to generate the renderings (d) that are a close match to the corresponding photographs (e). Note that the validation photographs are under illumination from the frontal direction (top two examples) as well as from the side (bottom two examples). We are able to reliably recover spatially varying specular reflectance properties for a live subject using our technique due to the small number of required photographs, taken in just 5 seconds. The estimated roughness parameters in different regions of the face correspond well to those reported in the literature for faces [WMP*06].

There are slight mismatches in the renderings with the recovered parameters and the corresponding photographs. For example, the wax candle has a very smooth surface and its specular roughness is slightly overestimated in some regions due to the limited lighting resolution of the LED sphere. On the other hand, some mismatches on the flat spatially varying BRDF sample (third row) can be attributed to slight inaccuracies in the estimated surface normals with the limited angular extent of the illumination from the LCD panel, and slight errors in diffuse-specular separation due the discrete sampling of the polarizer orientations in front of the camera.

As predicted from simulations, our technique overestimates the specular roughness of a very specular object such as the Obsidian sculpture. Note that the reflective hemispherical setup does not allow for polarization based diffuse-specular separation. Hence, currently we are limited to objects with minimal diffuse reflectance when using this setup for specular roughness estimation.

Finally, although the second order gradient illumination conditions have a global spherical support, occlusions do

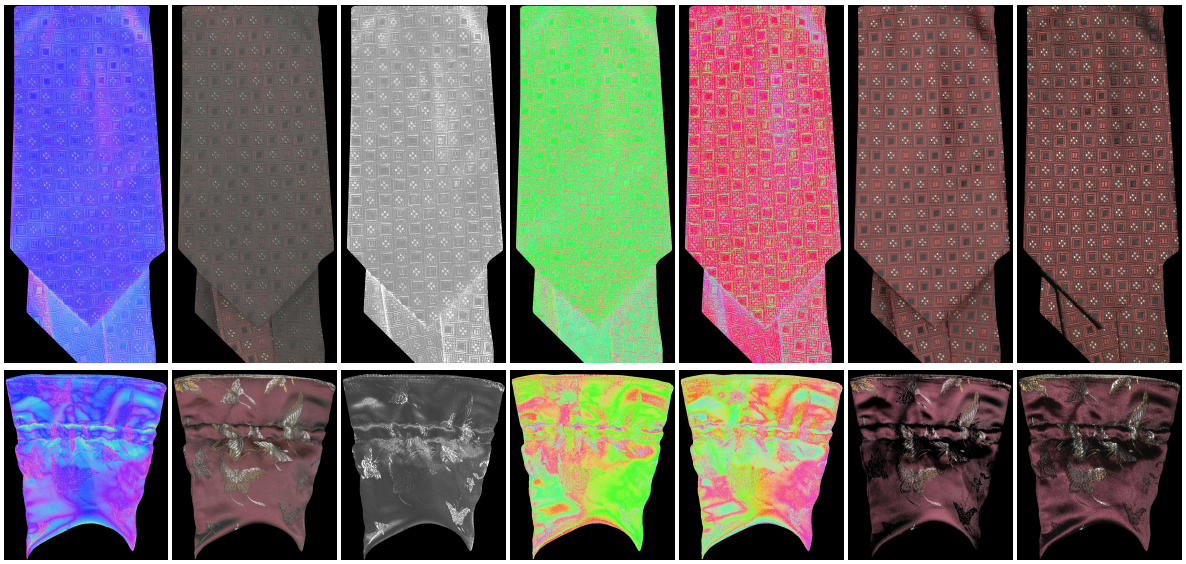

(a) spec. normals (b) spec. albedo (c) anisotropy ($\frac{\sigma_x}{\sigma_y}$) (d) tangent (e) bitangent (f) rendering (g) photograph

Figure 6: *Anisotropic specular reflectance properties ((a)-(c)) estimated using polarized second order spherical gradient illumination conditions. The estimated roughness anisotropy factor (c), and local tangent (d) and bitangent (e) orientations are applied to anisotropic BRDF models to create renderings (f) that closely matches the validation photographs (g). Top-row: a red necktie rendered with a Ward BRDF. Bottom-row: a red satin pouch rendered with an Ashikhmin BRDF.*

not affect the roughness and anisotropy estimation much, unlike the estimation of normals from diffuse reflections as in [MHP*07]. While the diffuse BRDF has a global angular support, specular BRDFs only have a very limited angular support. This reduces the effects of occlusions to an (almost) binary function. Either the mirror direction is occluded, or not. Note that even methods that rely on exhaustive sampling of all incident lighting directions do not yield reliable results when the mirror direction is occluded. For rough specular BRDFs, however, occlusions do affect the estimates. In such a case, part of the lobe is occluded, and thus a lower variance is reported (i.e., the estimated BRDF will be more specular).

7. Conclusion and Future Work

This work introduces higher order statistics in order to measure the appearance properties of a spatially varying material sample. We demonstrate that per-surface point specular roughness and anisotropy can be accurately estimated using second order spherical gradient illumination for glossy to moderately specular materials. We examine the limits of the proposed technique using simulations and demonstrate practical acquisition results with various measurement setups. Using this technique, we are able to obtain a more complete material BRDF information with fewer measurements compared to previous work.

We believe this technique can be very useful for even traditional BRDF fitting approaches as the obtained roughness estimates can serve as the starting point for any off-line numerical optimization. The obtained roughness map can also be used to factor out the view dependent Fresnel effects from the albedo map. Note that we assume the BRDF to be symmetric about the mean direction in this work. Higher order gradients would be required to measure asymmetries such as off-specular peaks in the BRDF. The available lighting resolution would become even more critical in order to reliably measure such statistics in practice.

Acknowledgments

The authors would like to thank Saskia Mordijck, Jay Busch, Monica Nichelson, Bill Swartout, Randy Hill, and Randolph Hall for their generous support to this project. Finally, we thank the anonymous reviewers for their helpful suggestions and comments. This work was supported by the U.S. Army Research, Development, and Engineering Command (RDECOM) and the University of Southern California Office of the Provost. The content of the information does not necessarily reflect the position or the policy of the U.S. Government, and no official endorsement should be inferred.

References

[AK07] ALLDRIN N., KRIEGMAN D.: Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2007), pp. 1–8.

[APS00] ASHIKHMIN M., PREMOZE S., SHIRLEY P. S.: A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000* (2000), Computer Graphics Proceedings, Annual Conference Series, pp. 65–74.

[Ash06] ASHIKHMIN M.: Distribution-based BRDFs, 2006. http://jesper.kalliope.org/blog/library/dbrdfs.pdf.

[AZK08] ALLDRIN N., ZICKLER T., KRIEGMAN D.: Photometric stereo with non-parametric and spatially-varying reflectance. In *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR)* (2008).

[FCM*08] FRANCKEN Y., CUYPERS T., MERTENS T., GIELIS J., BEKAERT P.: High quality mesostructure acquisition using specularities. *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on* (2008), 1–7.

[GCHS05] GOLDMAN D. B., CURLESS B., HERTZMANN A., SEITZ S. M.: Shape and spatially-varying brdfs from photometric stereo. In *ICCV '05: Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1* (2005), pp. 341–348.

[Geo03] GEORGHIADES A.: Recovering 3-D shape and reflectance from a small number of photographs. In *Rendering Techniques* (2003), pp. 230Ű–240.

[GTHD03] GARDNER A., TCHOU C., HAWKINS T., DEBEVEC P.: Linear light source reflectometry. In *ACM SIGGRAPH 2003* (2003), pp. 749–758.

[HLHZ08] HOLROYD M., LAWRENCE J., HUMPHREYS G., ZICKLER T.: A photometric approach for estimating normals and tangents. In *SIGGRAPH Asia '08: ACM SIGGRAPH Asia 2008 papers* (2008), pp. 1–9.

[HS03] HERTZMANN A., SEITZ S.: Shape and materials by example: a photometric stereo approach. *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on 1* (2003), 533–540, vol.1.

[HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. *SIGGRAPH Comput. Graph. 25*, 4 (1991), 175–186.

[LBAD*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHI R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics 25*, 3 (2006), 735–745.

[LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), pp. 117–126.

[LGK*01] LENSCH H. P. A., GOESELE M., KAUTZ J., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatially varying materials. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques* (2001), pp. 103–114.

[LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics 22*, 2 (2003), 234–257.

[Mar98] MARSCHNER S.: *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, 1998.

[MHP*07] MA W.-C., HAWKINS T., PEERS P., CHABERT C.-F., WEISS M., DEBEVEC P.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Rendering Techniques* (2007), pp. 183–194.

[MZKB05] MALLICK S. P., ZICKLER T. E., KRIEGMAN D. J., BELHUMEUR P. N.: Beyond lambert: Reconstructing specular surfaces using color. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition* (2005).

[NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of brdf models. In *Proceedings of the Eurographics Symposium on Rendering* (2005), pp. 117–226.

[NRH*77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERIS T.: Geometric considerations and nomenclature for reflectance. *National Bureau of Standards Monograph 160* (1977).

[PHD06] PEERS P., HAWKINS T., DEBEVEC P.: *A Reflective Light Stage*. Tech. Rep. ICT Technical Report ICT-TR-04.2006, ICT-USC, 2006.

[RH02] RAMAMOORTHI R., HANRAHAN P.: Frequency space environment map rendering. In *Proc. of ACM SIGGRAPH '02* (2002), pp. 517–526.

[SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques* (2002), pp. 527–536.

[TS67] TORRANCE K. E., SPARROW E. M.: Theory of off-specular reflection from roughened surfaces. *J. Opt. Soc. Am. 57* (1967), 1104–1114.

[War92] WARD G. J.: Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph. 26*, 2 (1992), 265–272.

[WMP*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics 25*, 3 (2006), 1013–1024.

[Woo78] WOODHAM R. J.: Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proc. SPIE's 22nd Annual Technical Symposium* (1978), vol. 155.

[ZBK02] ZICKLER T. E., BELHUMEUR P. N., KRIEGMAN D. J.: Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *Int. J. Comput. Vision 49*, 2-3 (2002), 215–227.

[ZREB06] ZICKLER T., RAMAMOORTHI R., ENRIQUE S., BELHUMEUR P. N.: Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *IEEE Trans. Pattern Anal. Mach. Intell. 28*, 8 (2006), 1287–1302.

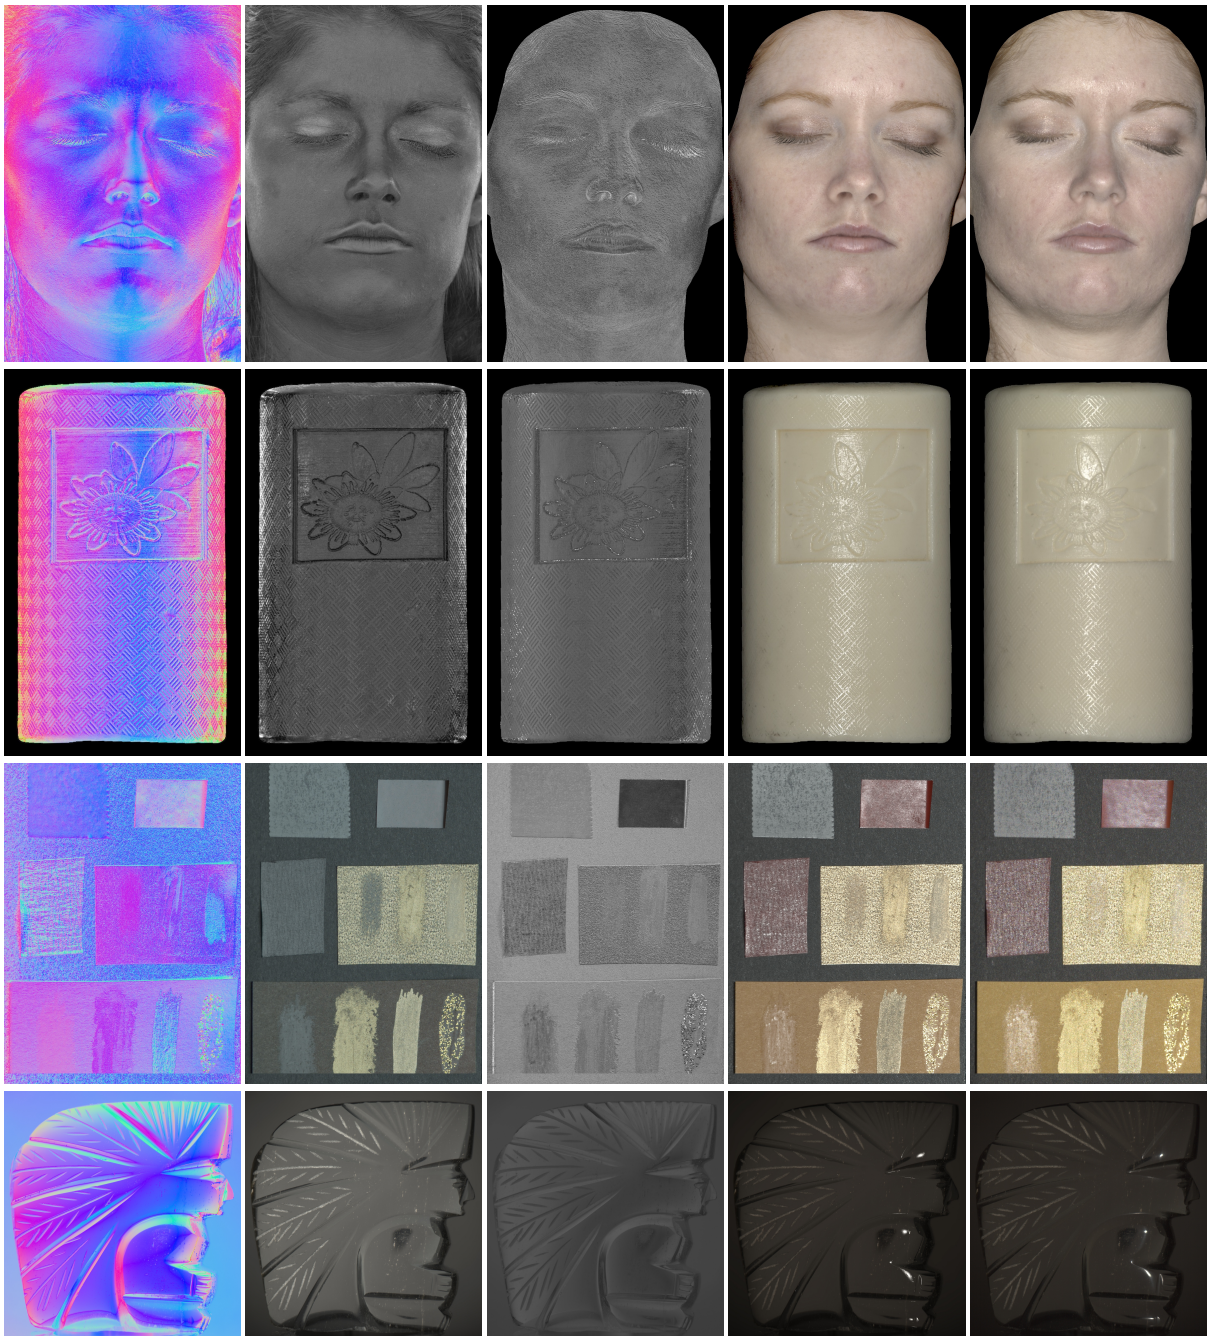

(a) specular normals (b) specular albedo (c) specular roughness (d) rendering (e) photograph

Figure 7: *Specular reflectance properties ((a)-(c)) estimated using second order spherical gradient illumination conditions. The estimated specular roughness map (c) is used as the per pixel distribution for a Torrance-Sparrow BRDF to create renderings (d) that closely matches the validation photographs (e) with illumination from the front (top two rows), as well as from the side (bottom two rows). Top-row: a female subject. Second-row: a wax candle. Third row: a flat spatially varying BRDF sample measured using an LCD monitor as emitter. Bottom row: a specular Obsidian sculpture measured using a reflective hemisphere.*