

Acquiring Reflectance and Shape from Continuous Spherical Harmonic Illumination

Borom Tunwattanapong¹ Graham Fyffe¹ Paul Graham¹ Jay Busch¹ Xueming Yu¹ Abhijeet Ghosh^{1,2} Paul Debevec¹

USC Institute for Creative Technologies¹

Imperial College London²

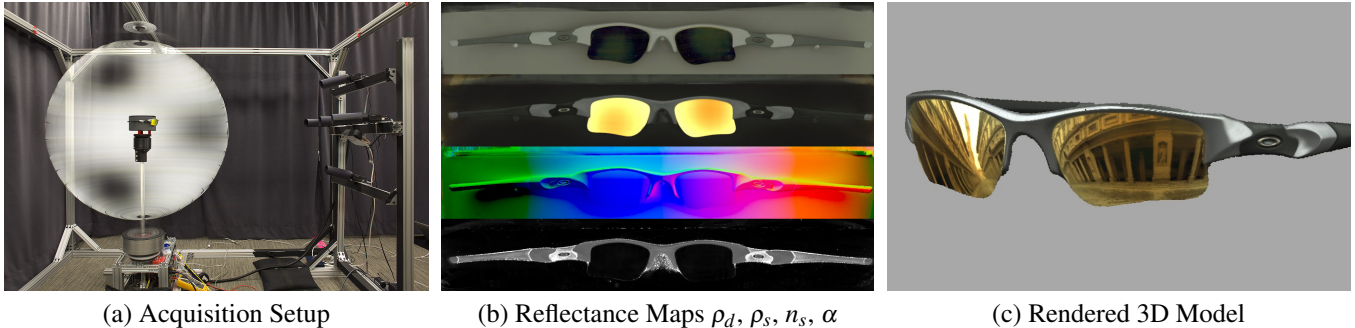


Figure 1: (a) A pair of sunglasses lit by continuous spherical harmonic illumination (harmonic $y_3^{-2}(\omega)$) in our capture setup. (b) Recovered diffuse, specular, specular normal, and specular roughness maps. (c) A rendering of the sunglasses with geometry and reflectance derived from the SH illumination and multiview reconstruction.

Abstract

We present a novel technique for acquiring the geometry and spatially-varying reflectance properties of 3D objects by observing them under continuous spherical harmonic illumination conditions. The technique is general enough to characterize either entirely specular or entirely diffuse materials, or any varying combination across the surface of the object. We employ a novel computational illumination setup consisting of a rotating arc of controllable LEDs which sweep out programmable spheres of incident illumination during 1-second exposures. We illuminate the object with a succession of spherical harmonic illumination conditions, as well as photographed environmental lighting for validation. From the response of the object to the harmonics, we can separate diffuse and specular reflections, estimate world-space diffuse and specular normals, and compute anisotropic roughness parameters for each view of the object. We then use the maps of both diffuse and specular reflectance to form correspondences in a multiview stereo algorithm, which allows even highly specular surfaces to be corresponded across views. The algorithm yields a complete 3D model and a set of merged reflectance maps. We use this technique to digitize the shape and reflectance of a variety of objects difficult to acquire with other techniques and present validation renderings which match well to photographs in similar lighting.

Keywords: specular scanning, spherical illumination, spherical harmonics

1 Introduction

Digitally recording realistic models of real-world objects is a long-standing problem in computer graphics and vision, with applications in online commerce, industrial design, visual effects, and interactive entertainment. Some of the most successful techniques use a combination of 3D scanning and photography under different lighting conditions to acquire models of an object’s shape and reflectance. When both of these characteristics are measured, the models can be used to render how the object would look from any viewpoint, reflecting the light of any environment, allowing the digital model to represent the object faithfully in a virtual world.

Acquiring the shape of an object independent of its reflectance characteristics is a largely solved problem. Active illumination techniques such as structured light scanning can record the shape of an object as long as it has a diffuse component, and passive stereo can obtain robust results as long as the object has surface texture. Unfortunately, these techniques alone do not record the reflectance properties of the surfaces being scanned, and cannot work on smooth shiny objects unless the reflectance is modified through dulling spray or powder coating.

Acquiring the reflectance properties of an object means measure its spatially-varying BRDF (SVBRDF) [McAllister 2002], or Bidirectional Texture Function (BTF) [Dana et al. 1999]. The problem with directly acquiring an SVBRDF is that a very large number of photographs is required to observe the object from all possible angles and lighting directions. Practical simplifications of the problem such as Sato et al. [1997] and Lensch et al. [2003] observe the object from a tractably sparse set of viewpoints and lighting conditions, and assume that BRDFs change smoothly across regions to extrapolate the measurements to full SVBRDFs. However, such techniques fail to record the BRDF independently at each surface point, meaning that interesting reflectance information could be missed at various places, and that two scans of the same object could yield significantly different models of the surface reflectance.

The largest challenge to efficient SVBRDF acquisition is that many materials exhibit sharp specular reflections which require very many observations to characterize using point light reflectometry techniques. If a point on an object is shiny, the specular highlight will only be observed close to the reflection vector, so many such angles must be observed. Extended light sources (e.g. [Ikeuchi 1981; Nayar et al. 1990; Gardner et al. 2003]) can be used effectively to excite more specular reflection directions at once, and they also reduce the brightness disparity between diffuse and specular reflections. Drawbacks remain, however, in that these techniques typically acquire less information about the shape of the specular lobe, or are limited to recording reflectance for a limited range of surface normals, or still require hundreds of images to observe the specular characteristics. In addition, shiny surfaces which lack a diffuse component, such as polished metal or tinted glass, pose significant problems for 3D reconstruction with either structured light or stereo correspondence techniques.

In this work, we present an improved practical acquisition and

analysis approach for digitizing the shape and reflectance of objects with spatially-varying BRDFs, including completely diffuse or specular surfaces, including mirror-like specularity. We do this with a tractable number of images by using continuous spherical harmonic illumination patterns generated by a semicircular arm of LEDs which traces out spherical lighting patterns during each exposure. We show that spherical harmonic lighting can be used to separate diffuse and specular reflections, to measure diffuse and specular albedo and surface orientation, and that higher orders responses can be used to characterize the shape of both isotropic and anisotropic specular reflections independently at each point. Furthermore, we show that the surface orientation measurements allow specular surfaces to be used in the context of multiview stereo geometry reconstruction in the same manner as diffuse surfaces, allowing the shape and reflectance of objects such as sunglasses, metallic parts, and glossy figurines to be digitized in the same straightforward manner.

In summary, the principal contributions of this work are:

- A novel computational illumination system for illuminating an object with high-resolution continuous spherical lighting conditions.
- The use of spherical harmonic illumination conditions for diffuse and specular reflectance component separation.
- A method of using higher-order spherical harmonics to measure the albedo, reflection vector, roughness, and anisotropy parameters of a specular reflectance lobe.
- A multi-view stereo algorithm which uses both diffuse and specular albedo and surface orientation measurements for high-quality geometry reconstruction.

2 Related work

An extensive body of work in the graphics and vision literature addresses the acquisition of geometry and reflectance from images under controlled and uncontrolled lighting conditions. Two recent overviews are Weyrich et al. [2009], which covers a wide variety of techniques for acquiring and representing BRDFs and SVBRDFs over object surfaces, and Ihrke et al. [2010], which focused on the acquisition of purely specular and transparent objects. In the following, we highlight some of the most relevant work in capturing opaque objects with spatially-varying diffuse and specular reflectance components.

Spatially Varying BRDF Capture SVBRDFs can be captured exhaustively using point light sources (e.g. [Dana et al. 1999; Mcallister 2002]), but this requires a large number of high-dynamic range photographs to capture every possible combination of incident and radiant angle of light. Similar to our approach, many techniques (e.g. [Debevec et al. 2000; Gardner et al. 2003; Holroyd et al. 2008; Ren et al. 2011]) look instead at BRDF slices of spatially-varying materials observed from a single viewpoint to infer parameters of a reflectance model, which can be used to extrapolate reflectance to novel viewpoints. Other approaches [Sato et al. 1997; Lensch et al. 2003; Zickler et al. 2006] use sparse sets of viewpoint and lighting directions and extrapolate BRDFs per surface point assuming that the reflectance varies smoothly over the object. This approach is also used by Dong et al. [2010], which employs a dedicated BRDF measurement system to sample representative surface BRDFs, which are extrapolated to the surface of the entire object based on its appearance under a moderate number of environmental lighting conditions. None of these techniques, however, produces independent measurements of diffuse and specular reflectance parameters for each observed surface point, and thus may miss important surface reflectance detail. Holroyd et al. [2010] describes a complete system for high-precision shape

and reflectance measurement of 3D objects using a pair of co-axial camera-projector units. Their setup uses phase-shifted structured light leveraging Helmholtz reciprocity [Zickler et al. 2002] for high-quality geometry estimation of nonconvex objects, and a clustering technique to derive SVBRDFs across object surfaces from a relatively sparse sampling of viewpoints. While their system can produce high-quality results for many objects, it would likely have trouble estimating geometry and reflectance where sharp specular reflections are dominant, such as sunglasses lenses.


Using Extended Light Sources Ikeuchi [1981] extended the original photometric stereo approach of Woodham [1980] to specular surfaces, using a set of angularly varying area light sources to estimate the specular surface orientation. Nayar et al. [1990] used an extended light source technique to measure orientations of hybrid surfaces with both diffuse and specular reflectance, but they did not characterize the BRDF of the specular component. Gardner et al. [2003] employed a moving linear light source to derive BRDF models of spatially-varying materials, including highly specular materials, but still required hundreds of images of the moving light to record sharp reflections. Hawkins et al. [2005] recorded diffuse and specular reflectance behavior of objects with high angular resolution using a surrounding spherical dome and a laser to excite the various surface BRDFs through Helmholtz reciprocity, but achieved limited spatial resolution and required a high-powered laser equipment. Recently, Wang et al. [2011] used step-edge illumination to estimate dual scale reflectance properties of highly glossy surfaces, but did not estimate per-pixel BRDFs.

Reflectance from Spherical Illumination Ma et al. [2007] used spherical gradient illumination representing the 0th and 1st order spherical harmonics in an LED sphere to perform view-independent photometric stereo for diffuse and/or specular objects, and used polarization difference imaging to independently model diffuse and specular reflections of faces. Ghosh et al. [2009] extended this approach by adding 2nd order spherical harmonics to estimate spatially varying specular roughness and anisotropy at each pixel. Unfortunately, the use of an LED sphere with limited resolution made the reflectance analysis applicable only to relatively rough specular materials such as human skin, and the use of polarization for component separation becomes complicated for metallic surfaces and near the Brewster angle. To avoid using polarization for reflectance component separation, Lamond et al. [2009] modulated gradient illumination patterns with phase-shifted high-frequency patterns to separate diffuse and specular reflections and measure surface normals of 3D objects. Our work reformulates and generalizes this frequency-based component separation approach to increasing orders of spherical harmonic illumination. Noting that BRDFs can usefully be represented by spherical harmonic functions (e.g. [Westin et al. 1992]), Ghosh et al. [2010] used spherical harmonic illumination projected to a zone of a hemisphere for reflectance measurement, but only for single BRDFs from flat samples. Their approach also did not separate diffuse and specular reflectance using the measurements and required very high orders of zonal basis function to record sharp specular reflectance. Instead, in this work we propose employing up to 5th order spherical harmonics for both diffuse-specular separation as well as estimating reflectance statistics such as specular roughness and anisotropy.


Geometry from Specular Reflection Highly specular objects have long posed problems for image-based shape reconstruction [Blake and Brelstaff 1992]. Bontford and Strum [2003] proposed a multi-view voxel carving technique based on reflected observations of a calibrated world pattern, but achieved very low resolution results. Tarini et al. [2005] proposed an alternate shape from distortion approach for high resolution reconstruction. Using a CRT screen as an extended light source, they illuminate a specular object

with several stripe patterns to obtain first a matte and then iteratively solve for the depth-normal ambiguity. Chen et al. [2006] propose measuring mesostructure from specularly using a hand held light source that is waved around while a camera observes the moving specular highlights on the sample to estimate its surface normals. Francken et al. [2008] propose measuring surface mesostructure instead by using a set of calibrated gray codes projected from an LCD panel in order to localize the surface normal of each surface point. These techniques work well in principle but are limited to scanning small flat objects that are covered by the illumination from an extended source. Adato et al. [2007] have proposed an alternate approach for shape reconstruction by formulating a set of coupled PDEs based on observed specular flow at a specific surface point. They also derive a simple analytic formulation for a special case of camera rotation about the view axis. While having the advantage of not requiring control over the incident illumination, in practice the method yields only very simple shape reconstructions. In contrast, our technique combines cues from both diffuse and specular reflectance information to derive high-fidelity geometric models for many common types of objects.

3 Setup and Acquisition



(a) Light arc, object, and cameras



(b) Exposure for one rotation (c) Object lit by sphere of light

Figure 2: *Spinning Spherical Reflectance Acquisition Apparatus*

Our lighting apparatus is designed to illuminate an object at its center with any series of continuous spherical incident illumination conditions. The light is produced by a 1m diameter semi-circular arc (Fig. 2, a) of 105 white LEDs (Luxeon Rebels) which rotates about its central vertical axis using a motion control motor. As seen in cross-section in 3, each LED is focused toward the center with a clear plastic optical element which is aimed through two spaced-apart layers of diffusion. The diffusion allows the LEDs to form a smooth arc of light when they are all on, but baffles between the optics help each LED have only a local effect on the arc (top graphs in Fig. 3). Since the arc spins through space more slowly near its top and bottom than at the equator, it would naturally produce more light per solid angle near the poles than from the equator. To counteract this, a curved aperture slit is applied to the arc which is

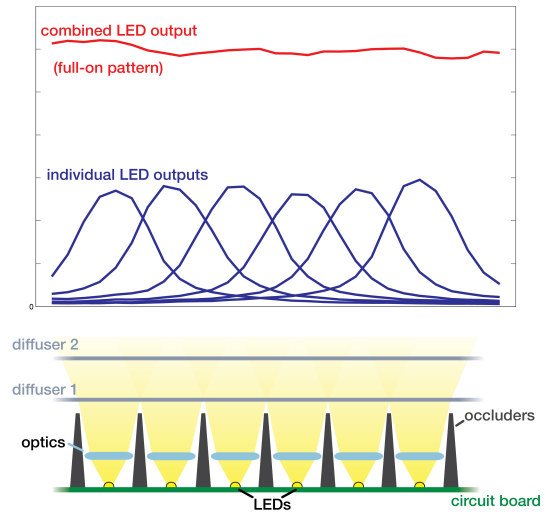


Figure 3: *Cross-section of the LED arm and plots of measured vertical intensity profiles for six of the 105 LEDs (blue curves) and the nearly constant intensity achieved along the arm with all LEDs driven to equal intensity (red curve).*

1cm wide at the equator and tapers toward 0cm wide at the poles proportional to the cosine of the angle to the center.

The object to be scanned sits on a small platform at the center of the arc. The platform is motorized to rotate around the vertical axis yielding additional views of the object. Typically, the object is rotated to eight positions, 45 degrees apart. One version of the platform is a dark cylinder which can light up from LEDs mounted inside it; this version can measure an additional transparency map for objects such as eyeglasses.

In front of the object and outside the arm is an array of five machine vision cameras (PointGrey Grasshopper 2.0)(Fig. 2, c) arranged in a plus sign configuration, with each camera being spaced about fifteen degrees apart from its neighbor(s). Each camera has a narrow field of view lens framed and focused on the object.

In this work, we spin the arm at one revolution per second during which the intensities of the 105 LEDs are modulated to trace out arbitrary spherical illumination environments. Differences in LED intensity due to manufacturing are compensated for by calibrating the intensity of each LED as reflected in a chrome ball placed in the center of the device. We use pulse width modulation to achieve 256 levels of intensity with 400 divisions around the equator, allowing a resolution of 400×105 pixel lighting environments to be produced. We expose each of the cameras for the full second of each rotation to record a full sphere of incident illumination as the arc rotates (Fig. 2, b). Additionally, we do one rotation pass where the object is only illuminated from the back and dark on the front in order to obtain masks for visual hull for subsequent stereo processing. The use of a spinning arm largely eliminates problematic interreflections which would occur if one were to surround the object with a projection surface.

Our spherical lighting patterns minimize the brightness disparity between diffuse and specular reflections compared to techniques with more concentrated illumination sources: a perfectly specular sphere and perfectly diffuse sphere of the same albedo appear the same brightness under uniform light. Nonetheless, to optimally estimate reflectance properties of both high albedo and low albedo surfaces with low noise, we employ high dynamic range photography with 3 exposures per lighting condition, one and a half stops apart.

The motion of the arc creates a vertical sliver of reflection occlusion when it passes in front of each camera (see the dark line down the center of the sphere in Fig. 2, d). In processing, we estimate this missing reflectance information from data in the neighboring views. Typical datasets in this work are captured in approximately ten minutes.

4 Reflectometry from Spherical Harmonics

This section describes our algorithm to estimate per-pixel reflectance parameters for a single viewpoint of an object observed under SH illumination functions $y_l^m(\omega) = y_l^m(x, y, z)$ up to 5th order where ω is the unit vector (x, y, z) (Fig. 4). We assume a traditional reflectance model consisting of a Lambertian diffuse lobe D and a specular lobe S with roughness and anisotropy. Holding the view vector fixed, we parameterize the reflectance functions $D(\omega)$ and $S(\omega)$ by the unit vector ω indicating the incident illumination direction. Of course, we do not observe $D(\omega)$ and $S(\omega)$ directly, but rather the responses of their sum $f(\omega) = D(\omega) + S(\omega)$ to the SH illumination functions. We denote these responses $f_l^m = \int_{\Omega} f(\omega) y_l^m(\omega) d\omega$.

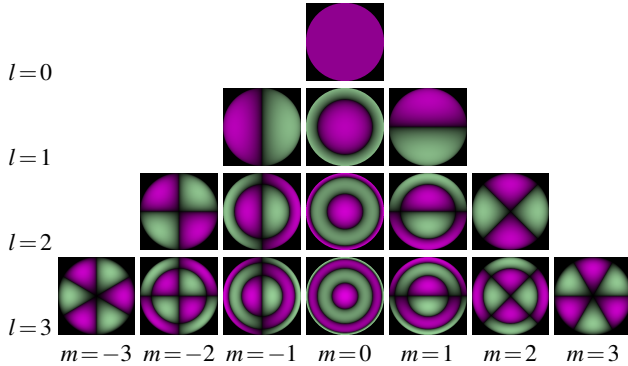


Figure 4: The spherical harmonic functions $y_l^m(\omega)$ up to 3rd order, seen from the top as reflected in a mirrored sphere, with $x = y = 0$ and $z = 1$ in the center. Positive values are shown in magenta and negative values are shown in green. Each harmonic locally resembles the harmonic above it around $x = y = 0$.




Figure 5: Responses of diffuse D , specular S , and mixed $D + S$ reflectance lobes to the first six zonal harmonics $y_l^0(\omega)$; negative values are shown in absolute value. (Non-zonal responses are 0). Above second order, the diffuse response is small and the higher order responses to the mixed lobe closely match the response to the specular lobe on its own.

A key to our reflectometry technique is the observation by Ramamoorthi and Hanrahan [2001] that a Lambertian diffuse lobe exhibits the vast majority of its energy in only the 0th, 1st, and 2nd-order spherical harmonic bands. We observe that this implies that the 3rd order SH coefficients and above respond only to the specular lobe $S(\omega)$, so that $S_l^m \approx f_l^m$ for $l \geq 3$, as seen in Fig. 5. We then es-

timate the specular lobe’s albedo, reflection vector, roughness, and anisotropy parameters from the higher order responses by comparing to the higher-order responses of lobes from a reflectance model such as [Ward 1992]. From the specular reflectance parameters, we can estimate the responses S_0^0, S_1^m of the specular lobe to the lower-order harmonics, and subtract this response from the observations to estimate the response of just the diffuse lobe $D(\omega)$ to the 0th and 1st order harmonics D_0^0, D_1^m . From those, we can estimate the diffuse albedo and diffuse surface normal as in [Ma et al. 2007], yielding a complete model of diffuse and specular reflectance per pixel from a small number of observations. Specifically, our reflectance measurement process is as follows:

4.1 Acquiring Spherical Harmonic Responses

We use our acquisition setup to acquire the responses of the object to the thirty-six SH illumination conditions up to the 5th order. Since the device cannot produce negative light, we offset and scale the SH functions above 0th order to produce two lighting patterns, one with pixel values between 0 and 255 and a complementary condition with pixel values from 255 to 0. The difference of these two images yields the response to the spherical harmonic. One could acquire fewer images by using the harmonics scaled 0 to 255, but our approach distributes camera noise more evenly throughout the range of intensities.

4.2 Building the Reflectance Table

We compute the response of our chosen reflectance model’s specular lobe to the SH illumination basis over its range of valid roughness values. In this work, we arbitrarily choose the Ward [1992] model’s specular lobe f_{α_1, α_2}^s and choose anisotropic roughness parameters $\alpha_1 \geq \alpha_2$ ranging from 0 (perfectly sharp) to 0.35 (very rough) in increments of 0.005. We view the surface at normal incidence along the z-axis, choose a unit specular albedo $\rho_s = 1$, and align the axis of anisotropy to 0 degrees along the x-axis. We then numerically integrate the lobes against the SH basis to determine the coefficient table $R_l^m(\alpha_1, \alpha_2)$ for each order l across the range of α .




Figure 6: Plots of the responses $R_l^m(\alpha, \alpha)$ (with $l = 0$ in black) of an isotropic specular lobe to the first six zonal harmonics y_l^0 as roughness increases from $\alpha = 0$ to $\alpha = 0.35$. The dotted line plots the ratio R_3^0/R_5^0 which is used in determining specular roughness.

The response of the lobes to the SH basis has useful properties for reflectance measurement. Since isotropic lobes are radially symmetric, only the zonal SH basis functions y_l^0 yield a nonzero response when $\alpha_1 = \alpha_2$. Even when the lobe is anisotropic, it is symmetrical across both the x and y axes in that $f_{\alpha_1, \alpha_2}^s(x, y, z) = f_{\alpha_1, \alpha_2}^s(\pm x, \pm y, z)$ so its response to the basis functions $y_l^{\pm 1}$ (for $l \geq 1$) are both zero since $y_l^{-1}(x, y, z) = -y_l^{-1}(-x, y, z)$ and $y_l^1(x, y, z) = -y_l^1(x, -y, z)$ (seen in Fig. 4). Furthermore, the response to the basis functions y_l^{-2} (for $l \geq 2$) are zero since they also have the




Figure 7: The rendered reflectance tables which map measured responses $(u, v) = (f_5^0/f_3^0, f_3^2/f_3^0)$ back to anisotropic roughness values α_1 (left) and α_2 (right).

property $y_l^{-2}(x, y, z) = -y_l^{-2}(-x, y, z)$. The responses to y_l^2 will be nonzero, however, when the lobe is anisotropic, since for small x and y , $y_l^2(x, y, z)$ is positive when $|x| < |y|$ and negative when $|x| > |y|$, so lobes stretched more along x will have a negative response and lobes stretched more along y will have a positive response to y_l^2 . We will use this response to measure the anisotropy of the observed specular lobes.

Fig. 6 shows the reflectance data $R_l^0(\alpha_1, \alpha_2)$ for $0 \leq l \leq 5$ for isotropic lobes where $\alpha_1 = \alpha_2$. In practice, the elements of interest in our reflectance table R are the zonal responses R_3^0 and R_5^0 (for roughness) and the tessoral response R_3^2 (for anisotropy). When measured from a real surface, they will all be scaled by the unknown specular albedo ρ_s , so we divide each of them by R_3^0 to obtain the two independent measurements $u = R_5^0/R_3^0$ and $v = R_3^2/R_3^0$ which generally correlate with specular roughness and anisotropy, respectively, as will be discussed further in Sec. 4.4 and Sec. 4.5.

Our calculations tell us values (u, v) for given values (α_1, α_2) . When it is time for reflectance measurement, we will need to evaluate the inverse of this mapping. Fortunately, the mapping is smooth and monotonic over our ranges of roughness, which allows us to construct a fast inverse table lookup. For each of α_1 and α_2 , we scan convert the mesh of the values they take on into the range of (u, v) (Fig. 7). Then, when we measure values u and v , we can quickly look up the anisotropic roughness parameters α_1 and α_2 to which they correspond.

4.3 Estimating the Specular Surface Normal

The specular lobe of our pixel’s reflectance function will be centered around a reflection vector r , implying a specular surface normal n_s halfway between r and the view vector. We search for this specular peak to be located at the maximum of the $l = 3$ order SH reconstruction of the function:

$$r = \arg \max_{\omega} \sum_{m=-l}^l f_m^3 y_m^3(\omega) \quad (1)$$

Since we assume the specular lobe is relatively narrow, we note that rotating the lobe to align with the zonal $+z$ axis will maximize its response to the zonal harmonic y_3^0 , which assumes its global maximum along the $+z$ axis. Around this same location, all other SH functions of the third order are close to zero. Thus, we observe that a narrow specular lobe’s projection into the 3rd order SH basis will resemble a rotated version of the 3rd order zonal harmonic itself, and thus will have a clear global maximum.

Since finding where the SH reconstruction of a function attains its maximum is complicated to perform analytically, we use the quick-to-converge hill climbing approach of Sloan [2008] to find r starting from six possible initial estimates. From r , we can calculate the world-space surface normal n_s , and can rotate the other SH responses f_l^m using the $(2l + 1) \times (2l + 1)$ SH rotation matrices so

that the reconstructed specular lobe aligns with the zonal $+z$ axis, yielding a set of rotated SH coefficients \hat{f}_l^m .

4.4 Estimating the Angle of Anisotropy

The responses \hat{f}_l^m now measure the lobe as if it were aligned with the zonal axis, but if the lobe is anisotropic, the angle of anisotropy ψ could be anywhere. We would like to determine this angle, and further rotate the harmonics around the zonal axis to align the direction of anisotropy with the x axis so that it will match with our reflectance table. To measure the angle of anisotropy of a distribution $g(x, y)$ centered about the origin, one typically computes the covariances $s = \int g(x, y)(x^2 - y^2)dx dy$ and $t = \int g(x, y)(2xy)dx dy$, where s characterizes the spread of values along the x -axis versus the y -axis and t characterizes the spread of values along versus the $x = y$ diagonal versus the $x = -y$ diagonal. Then, the angle of anisotropy can be found as $\frac{1}{2} \tan^{-1}(s/t)$.

Conveniently, we observe that the \hat{f}_3^{-2} and \hat{f}_3^2 harmonic responses essentially perform these integrations for us (visualized in Fig. 8), up to a scaling factor, since from the spherical harmonic formulae in the neighborhood of the zonal peak $z = 1$ are:

$$y_3^{-2}(x, y, z) = \frac{1}{4} \sqrt{\frac{105}{\pi}} 2xyz \approx \frac{1}{4} \sqrt{\frac{105}{\pi}} 2xy \quad (2)$$

and

$$y_3^2(x, y, z) = \frac{1}{4} \sqrt{\frac{105}{\pi}} (x^2 - y^2)z \approx \frac{1}{4} \sqrt{\frac{105}{\pi}} (x^2 - y^2) \quad (3)$$

Thus, the angle of anisotropy can be computed as $\psi = \frac{1}{2} \tan^{-1}(\hat{f}_3^2/\hat{f}_3^{-2})$. If we were to rotate the harmonic responses around the z axis by ψ to form \hat{f}_l^m , we note that the \hat{f}_3^{-2} response becomes zero and the \hat{f}_3^2 response becomes $\sqrt{(\hat{f}_3^{-2})^2 + (\hat{f}_3^2)^2}$. We have simply calculated the angle and magnitude of the vector $(\hat{f}_3^{-2}, \hat{f}_3^2)$. This magnitude is our indication of how anisotropic the specular lobe is relative to its roughness.




Figure 8: Responses of an anisotropic lobe S to the rotated SH functions \hat{y}_l^m whose responses together determine estimates of the anisotropic angle ψ , roughnesses α_1 and α_2 , and the specular albedo ρ_s .

4.5 Estimating Roughness and Anisotropy

Ideally, to measure the roughness of the rotated specular lobe $\hat{S}(x, y, z)$, we would compute its second moment as the following integral:

$$\int \hat{S}(x, y, z)(x^2 + y^2)dx dy \quad (4)$$

Notably, the zonal harmonics for $l > 1$ perform something similar to this integration, visualized in the right images of Fig. 8. For example, for $l = 3$, near the apex $z = 1$:

$$y_3^0(x, y, z) = \frac{1}{4} \sqrt{\frac{7}{\pi}} z(2 - (x^2 + y^2)) \approx \frac{1}{2} \sqrt{\frac{7}{\pi}} - \frac{1}{4} \sqrt{\frac{7}{\pi}} (x^2 + y^2) \quad (5)$$

Thus the zonal harmonic approximation has an $(x^2 + y^2)$ term but also a constant term. If we knew the response of the specular lobe S to the 0th order harmonic y_0^0 , we could subtract off the response to the constant term. Unfortunately, we only know the response of y_0^0 to $D + S$. To resolve this problem, we can look to a higher zonal harmonic we capture such as y_5^0 :

$$y_5^0(x, y, z) = \frac{1}{16} \sqrt{\frac{11}{\pi}} z (8 - 56(x^2 + y^2) + 63(x^4 + 2x^2y^2 + y^4)) \quad (6)$$

$$\approx \frac{1}{2} \sqrt{\frac{11}{\pi}} - \frac{7}{2} \sqrt{\frac{11}{\pi}} (x^2 + y^2) \quad (7)$$

To make this approximation, we set the 4th order terms to zero (and drop $z \approx 1$). This approximation also contains a constant term and an $(x^2 + y^2)$ term, with linearly independent coefficients from the y_3^0 , approximation, meaning that one can theoretically determine the integral of the specular lobe S against the constant term and the $(x^2 + y^2)$ term from the responses \hat{f}_3^0 and \hat{f}_5^0 , yielding measures of the specular lobe's albedo and roughness.

However, since these polynomial expansions are approximate, and since our reflectance model's parameters may or may not relate closely to these same measures of roughness and anisotropy, we employ a lookup table to determine which specular model parameters will produce the same responses as we find in our captured data. But the derivations above explain why the measurements we seek are contained in the available data.

We return to our reflectance function's rotated SH responses \hat{f}_3^0 , \hat{f}_5^0 , and \hat{f}_3^2 . As discussed earlier in Sec. 4.2, these will all be proportionately scaled by the specular albedo, so we normalize them by dividing by the 3rd order zonal response \hat{f}_3^0 to obtain $(u, v) = (\hat{f}_5^0/\hat{f}_3^0, \hat{f}_3^2/\hat{f}_3^0)$. From the precomputed reflectance table of Sec. 4.2, look up which anisotropic roughness parameters α_1 and α_2 yield a specular lobe in our reflectance model which has the same response to the rotated harmonics.

Note that we could in principle similarly use the 4th order zonal harmonic y_4^0 instead of the 5th order harmonic. In practice, we prefer the 5th order harmonic because of it provides a greater contrast to the 3rd order measurement and hence better conditioned measurements. Also, the 4th order measurement, being an even order harmonic, still has some diffuse response in the signal making it less suitable for this purpose than the the 5th order measurement (Fig. 5).

4.6 Estimating Specular Albedo

As shown above, the responses \hat{f}_3^0 and \hat{f}_5^0 of our specular lobe to the rotated zonal harmonics contain information about both specular roughness and albedo. With the roughness values α_1 and α_2 determined, we refer to our tabulation R (Fig. 6) to determine the 3rd order zonal response to a unit specular albedo lobe with our lobe's roughness parameters. Dividing our lobe's 3rd order zonal response yields our estimate of the specular albedo $\rho_s = \hat{f}_3^0/R_3^0(\alpha_1, \alpha_2)$.

4.7 Estimating Diffuse Albedo and Normal

We have now fully characterized our specular lobe with specular albedo ρ_s , normal n_s , angle of anisotropy ψ , and anisotropic roughness parameters α_1 and α_2 . From ρ_s , α_1 , and α_2 , we can determine the 0th and 1st order responses to our lobe from the table $R_l^m(\alpha_1, \alpha_2)$. We use spherical harmonic rotation to rotate the 1st

order responses by the angle ψ around the z axis and then rotate to align the lobe's center to the reflection vector from the view vector to n_s . We subtract these rotated responses $\rho_s \hat{R}_l^m(\alpha_1, \alpha_2)$ from our observations f_l^m (which include the response to both the diffuse and specular lobes) to estimate the response to the diffuse lobe on its own. From the 0th and 1st order responses D_l^m to the diffuse lobe, it is straightforward to estimate the diffuse normal n_d as (D_{-1}^0, D_1^0, D_1^1) (which should be normalized) and albedo as $\frac{1}{\pi} D_0^0/y_0^0$, where the $\frac{1}{\pi}$ factor divides out the integral of a unit Lambertian lobe over the sphere and y_0^0 is the constant value of the 0th order harmonic $\frac{1}{2} \sqrt{\frac{1}{\pi}}$.

4.8 Reflectometry Discussion

We do not consider the effect of Fresnel gain, which makes grazing lobes considerably brighter. This could be accounted for using the Fresnel terms from a physically based reflectance model and estimates of the index of refraction of the materials. In our setup, we capture enough viewpoints to view most surface normals at an angle where Fresnel gain is minor, and our map merging process assigns frontal reflectance data the highest weight.

Our technique for measuring roughness and anisotropy is different from Ghosh et al. [2009], which uses only the 2nd order harmonics for reflectance analysis. They require diffuse reflection to be eliminated through polarization difference imaging, and integrate against the equatorial region of the zonal harmonic rather than the zonal peaks. We believe our technique is more general since it can be applied to higher-order harmonics (which exclude diffuse reflections without polarization) and can obtain better-conditioned estimates of sharp specular behavior.

5 Geometry Reconstruction from Diffuse and Specular Reflections

Once the reflectance maps are acquired and processed for each of the viewpoints, we have estimates of the spatially varying diffuse and specular albedo and diffuse and specular normal maps (which correspond to diffuse and specular reflected *directions* on the illumination sphere). Our acquisition setup has five color cameras in a "plus sign" arrangement with fifteen degrees between views, and a motorized platform rotates the object about the vertical axis at 45° increments to provide views all around the object. We estimate a 3D surface model for the object using the maps from all of these views using a geometry reconstruction algorithm that leverages multi-view stereo correspondence and the surface normal estimates from the reflectance analysis. We first describe our multiview camera calibration procedure before describing the stereo reconstruction in Section 5.1.

Camera calibration We employ a single-shot camera calibration process using a 15cm diameter checkerboard cylinder (Fig. 9). The cylinder has three colored dots to consistently identify the cylinder's orientation in each view. The checker size provides a known distance to establish scale. The checker corners are detected using a corner detector [Harris and Stephens 1988] and refined to sub-pixel coordinates. The pixel coordinates of the detected corners are matched to the corresponding 3-D points on an ideal cylinder.



Figure 9: Top, left, front, right, and bottom views of the calibration cylinder from the camera array.

We provide initial estimates of the camera poses and focal lengths to a least squares minimization [Moré et al. 1984] of the distance between the detected corner points and their 3D reprojections in the five camera models. We first optimize only the cameras’ extrinsic parameters, and then also optimize the cameras’ intrinsics (including two terms of radial distortion), and finally also optimize the 3D positions of the points on the cylinder so that it need not be constructed with great precision. The solver converges in about one minute yielding average point re-projection errors of 0.15 to 0.35 pixels.

5.1 Stereo Reconstruction

We solve for the object’s shape beginning with an initial approximate base mesh of the geometry obtained from either sparse stereo correspondence and Poisson reconstruction (e.g. [Furukawa and Ponce 2009]), shape-from-silhouettes, or simply a cylinder at the approximate location of the object being scanned. We define a displacement map over the vertices of this base mesh, refining its estimated shape by displacing the base vertices along their normals. We compute displacement values to minimize the following energy function, defined over the set of base mesh vertices \mathcal{V} and edges \mathcal{E} :

$$E = \omega(P, V) + \sum_{i \in \mathcal{V}} \rho(p_i, n_i, v_i) + \sum_{i, j \in \mathcal{E}} \eta(p_i, p_j, n_i, n_j) + \zeta(p_i, p_j, n_i, n_j, v_i, v_j), \quad (8)$$

where $P = \{p_i\}_{i \in \mathcal{V}}$ are the estimated vertex positions (with $p_i = p_i^0 + d_i n_i^0$), p_i^0 and n_i^0 are the base mesh vertices and normals, respectively, $D = \{d_i\}_{i \in \mathcal{V}}$ are the displacement values, $N = \{n_i\}_{i \in \mathcal{V}}$ are the estimated surface normals, $V = \{v_i\}_{i \in \mathcal{V}}$ are the estimated sets of views in which each vertex is visible, ρ is a photoconsistency term, ω is a visibility term, η is a surface normal consistency term, and ζ is a photometric curvature term.

To simplify optimization, we take an iterative approach, first fixing P and optimizing V and N , and then vice versa. For efficiency we employ a multi-resolution scheme, where we vary the number of iterations and the mesh vertex density. We perform 8 iterations at $\frac{1}{8}$ vertex density, 4 iterations at $\frac{1}{4}$ vertex density, 2 iterations at $\frac{1}{2}$ vertex density and 1 iteration at full vertex density. The result of each pass initializes the next higher resolution pass using simple bilinear upsampling. The entire reconstruction takes four minutes on an 8-core Intel Xeon E5620 system with hyperthreading enabled. In detail, we initialize $D = \{0\}$ and iterate the following three stages:

Optimizing V Each set v_i indicates the views in which the vertex i is visible. We define $\omega = 0$ if every v_i matches the true visibility for vertex i given the mesh vertex positions P , and ∞ otherwise. With P fixed, ω dominates all other terms and we simply compute visibility from the current mesh estimate, omitting any vertices outside of the visual hull defined by the data masks captured in Section 3. We say that a vertex is outside the visual hull if the following holds:

$$\frac{\sum_{k \in \text{views}} w_{k,i} m_{k,i}}{\sum_{k \in \text{views}} w_{k,i}} < 0.75, \quad (9)$$

where $w_{k,i} = \max(0, n_i^g \cdot l_{k,i})$, n_i^g is a geometric surface normal computed from the vertex positions neighboring p_i , $l_{k,i}$ is the view vector $(c_k - p_i)/|c_k - p_i|$, c_k is the position of the camera for view k , and $m_{k,i}$ is the data mask value for view k at projected position p_i . Besides the possibility of a vertex being occluded by other parts of the mesh, we also consider back-facing vertices to be occluded.

Optimizing N With P fixed and V already computed, we optimize N considering only η and ζ . (ρ also influences N , but less so.)

The surface normal consistency term η prefers the surface tangent vector $(p_i - p_j)/|p_i - p_j|$ to be perpendicular to the surface normals:

$$\eta = \lambda(n_i \cdot n_j)^\alpha \left(\frac{(p_i - p_j) \cdot (n_i + n_j)}{|p_i - p_j||n_i + n_j|} \right)^2. \quad (10)$$

Normalizing the tangent vector keeps the solution consistent regardless of mesh vertex density. λ is a global smoothing weight (0.01 in our work) and α is an anisotropic smoothing weight modulation exponent (64 in our work) to reduce the smoothing weight across high-curvature boundaries.

We replace (10) with the following approximate solution to yield a least-squares linear problem:

$$\eta \approx \lambda(n_i^p \cdot n_j^p)^\alpha \frac{1}{2} (|n_i - n_i^g|^2 + |n_j - n_j^g|^2), \quad (11)$$

where n_i^p is the photometric normal estimate at vertex i :

$$n_i^p = \sum_{k \in v_i} w_{k,i} \left(\frac{1}{10} a_{k,i}^d \max(0, n_{k,i}^d \cdot l_{k,i}) n_{k,i}^d + a_{k,i}^s \max(0, n_{k,i}^s \cdot l_{k,i}) n_{k,i}^s \right), \quad (12)$$

normalized. This combines the diffuse normal estimate n^d and the specular normal estimate n^s , each weighted by their respective albedos a^d and a^s , sampled over all visible views, and weighted based on similarity to the view vector. The diffuse normal is weighted one tenth as much as the specular normal, because it is typically softened by scattering.

The diffuse and specular normals are derived from the measured diffuse and specular directions as follows:

$$n_{k,i}^d = \beta r_{k,i}^d - p_i, \quad (13)$$

normalized, where β is the radius of the illumination sphere and $r_{k,i}^d$ is the photometric diffuse direction (averaged over the three color channels) for view k at projected position p_i , and:

$$n_{k,i}^s = (\beta r_{k,i}^s - p_i) / |\beta r_{k,i}^s - p_i| + l_{k,i}, \quad (14)$$

normalized, where $r_{k,i}^s$ is the photometric specular direction (averaged over the three color channels) for view k at projected position p_i . (Albedos are also averaged over the three color channels.)

The curvature term ζ prefers the spatial change in estimated surface normal to agree with the change in photometric surface normal:

$$\zeta = \frac{\lambda(n_i^p \cdot n_j^p)^\alpha}{|p_i - p_j|^2} |(n_i - n_j) - (n_i^p - n_j^p)|^2. \quad (15)$$

As in η , normalizing by $|p_i - p_j|^2$ keeps the solution consistent regardless of mesh vertex density.

For any vertex that is outside the visual hull (as determined when optimizing V), we let $\eta = 0$ and $\zeta = \epsilon |n_i - n_j|^2$ (with some small ϵ) to provide a smooth interpolation between the other vertices.

The sum of η and ζ over edges $(i, j) \in \mathcal{E}$ is a sparse least-squares linear problem in terms of the x , y and z components of the normals N , which we solve by invoking Gaussian TRW-S message passing [Kolmogorov 2006] three times (for x , y , and z). Finally, we normalize the resulting surface normals to unit length.

Optimizing P With V and N fixed, we optimize P considering only ρ and η (as ζ has only second-order effects on P).

We use the full form (10) of the surface normal term η , but we replace $|p_i - p_j|$ in the denominator with the values from the previous iteration to yield a least-squares linear system in terms of displacement values D . This is an acceptable approximation in our iterative scheme as the denominator changes slowly with respect to displacement.

Our photoconsistency term ρ employs a novel blend of matching costs between multiple sets of calibrated cameras, when the calibration between sets is imprecise. In our case, we have 5 cameras mounted rigidly to our structure, and thus the calibration between these 5 cameras is precise. However, we have multiple copies of these cameras with different scan object rotations, which are not precise enough for stereo matching. Thus we compute matching costs within sets, and compute a weighted average of the costs, weighted by $w_{k,i}$ for the center camera k of each set.

The matching cost for each set is normalized cross correlation (NCC) with a 3×3 window aligned in space to the surface normal n_i . The NCC cost is a weighted average over the cameras (weighted by $w_{k,i}$) and scaled by the window variance in a primary view to avoid undue influence from noisy low-intensity pixel values. The primary view is the view within the set that is most facing n_i .

We average the NCC cost over 10 data channels (diffuse albedo RGB, mean specular albedo, mean diffuse normal XYZ, and mean specular normal XYZ), and truncate the cost to 1 for robustness to outliers. Vertex positions outside of the visual hull are given a constant cost of 1.

The landscape of ρ is highly irregular, so we employ the data-driven mean-shift scheme (DDMS) [Park et al. 2010], which fits a smooth approximation to the cost function over a weighted region of interest determined by the solution from the previous iteration of the outer optimization loop. We fit a weighted quadratic cost approximation over the region of interest based on discrete samples at 0.1mm intervals, yielding a sparse least-squares linear system in terms of the displacement values D .

We invoke Gaussian TRW-S message passing to compute D that minimizes the sum of (approximate) ρ and η . While the original DDMS scheme requires a double loop, we perform only the inner Gaussian loop since it is already nested in our V, N, P outer loop.

Reflectance Blending After the geometry reconstruction is complete, we blend the various channels of reflectance data using similar weights as the photometric normals, modulated by the similarity to the final estimated normals:

$$\bar{c}_i = \frac{\sum_{k \in v_i} w_{k,i} w_{k,i}^n c_{k,i}}{\sum_{k \in v_i} w_{k,i} w_{k,i}^n}, \quad (16)$$

where \bar{c}_i is the blended reflectance, $c_{k,i}$ is the reflectance data for view k at projected position p_i , and $w_{k,i}^n = \frac{1}{10} a_{k,i}^d \max(0, n_{k,i}^d \cdot l_{k,i}) \max(0, n_{k,i}^d \cdot n_i)^2 + a_{k,i}^s \max(0, n_{k,i}^s \cdot l_{k,i}) \max(0, n_{k,i}^s \cdot n_i)^2$.

6 Results

We now present some results of scanning objects with complex varying reflectance using our technique. All viewpoints were recorded under SH lighting up to fifth order, for 2×36 photographs per camera per view (though the 2×14 images from the 2nd and 4th orders are not used).

Figure 10 shows recovered maps for the shiny red plastic ball seen in Fig. 2 under uniform illumination. The maps show good diffuse/specular separation and the normals trace out the spherical shape, except in areas of reflection occlusion near the bottom. The roughness is low and consistent across the ball, but becomes especially low around the periphery presumably due to Fresnel gain, which can be observed in the specular map.

Fig. 11 shows maps recovered for an arrangement of five slightly bent anisotropic brushed metal petals at different angles. The maps exhibit four different angles of anisotropy. The major-axis roughness α_1 is consistent for all five petals, and the minor-axis roughness α_2 is sharper overall but especially low for the vertical petal. We

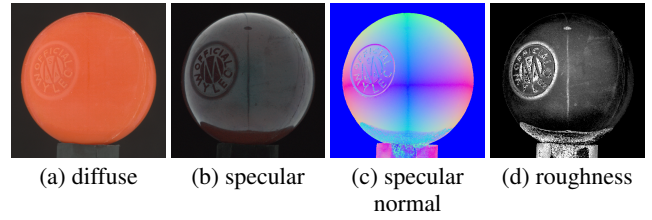


Figure 10: Estimated reflectance maps for a plastic ball.

believe this inconsistency is due to the LED arm’s higher resolution and lack of blur around the equator compared to along the length of the arm. The rendering and validation photograph, both under point-source illumination, are still reasonably consistent, though the tails of the specular lobes are wider in the photograph. This is likely because a the tendency of the Ward model lobe to fall to zero too quickly; a different reflectance model might better represent this reflectance.

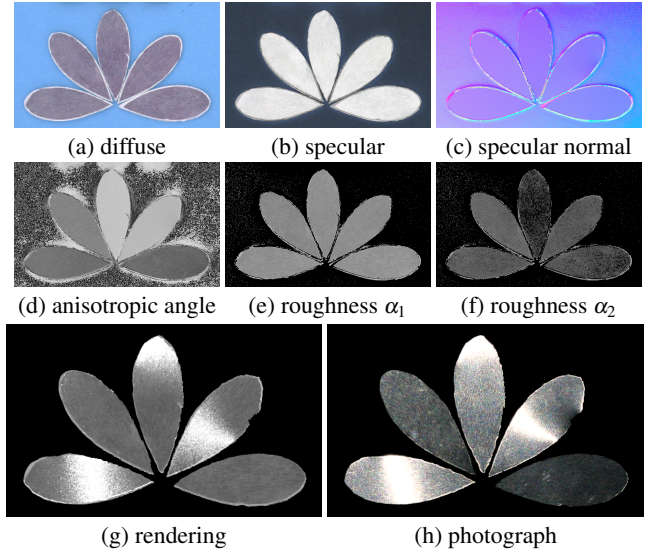


Figure 11: Maps, rendering, and photo for brushed metal.

Fig. 1 shows the scanning process for a pair of sunglasses with several materials of varying roughness on the frames and strongly hued lenses with a mirror-like reflection. The reflectance maps correctly show very little diffuse reflection on the lenses and spatially-varying roughness on the ear pieces. The glasses were scanned in five poses for twenty-five viewpoints total. The diffuse and specular albedo and diffuse and specular normals were used in the multiview shape reconstruction shown in Fig. 12, providing merged maps in a cylindrical texture space from all 25 views. The geometry successfully reconstructs both the diffuse ear pieces and the mirror-like lenses in a single process. Fig. 12 also provides a ground truth comparison of the rendered glasses to the real pair of glasses lit by environmental illumination, and an animation of the glasses is shown in the accompanying video.

Fig. 13 shows maps and geometry for a digital camera with several different colors and roughnesses of metal and plastic and an anisotropic brushed metal bezel around the lens. The maps successfully differentiate the materials and allow the renderings to give a faithful impression of the original device.

Finally, Fig. 14 presents an error analysis for shape and reflectance measurement for different types of spheres with the presented acquisition setup. As can be seen, our technique is able to correctly separate diffuse and specular reflectance for a mirror sphere, a red

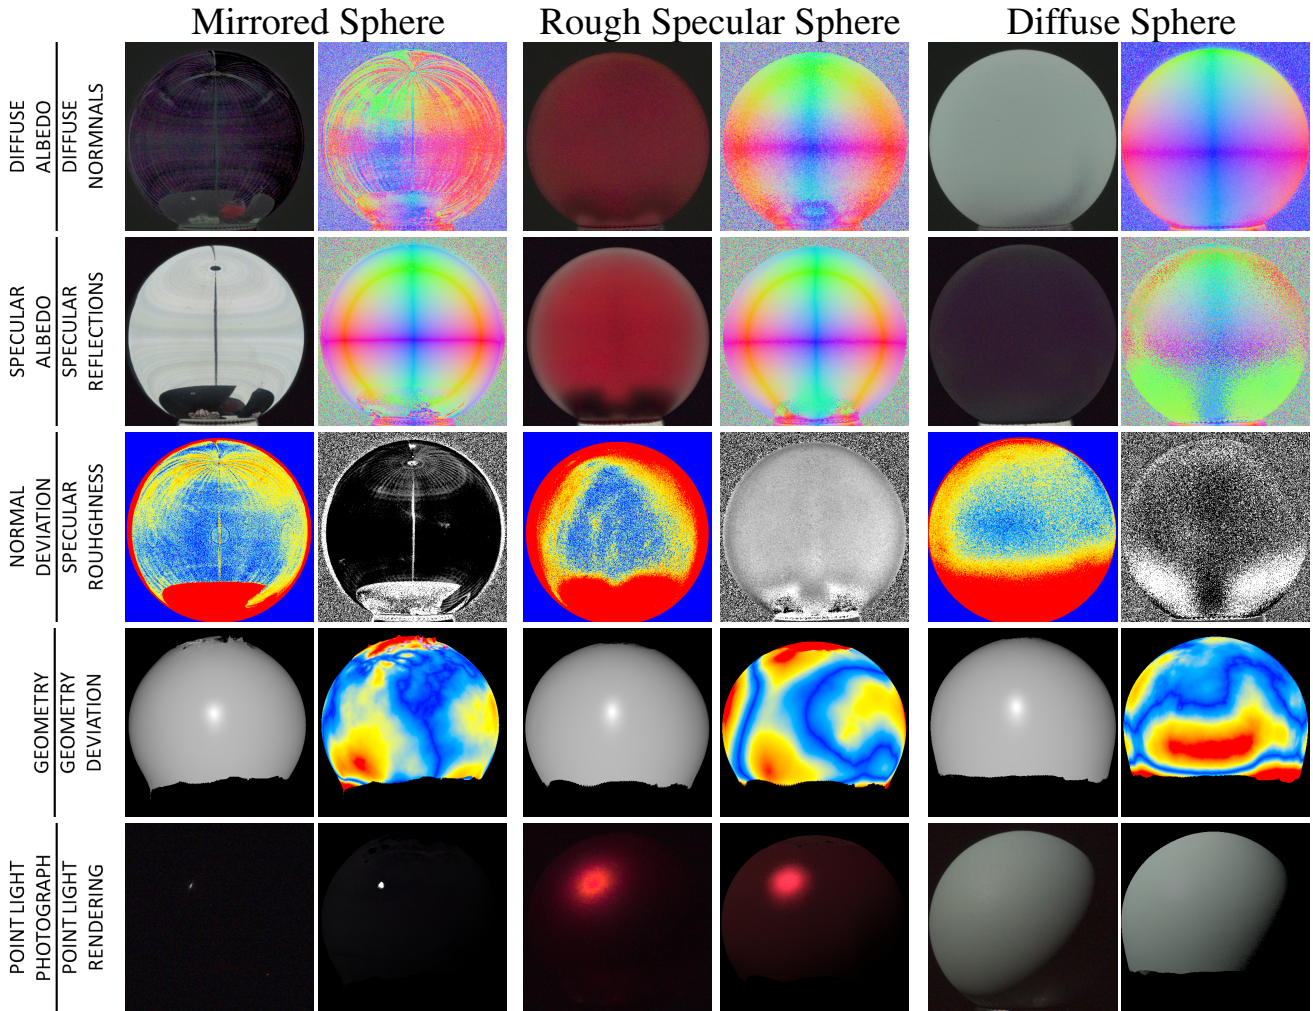


Figure 14: Analysis of shape and reflectance measurement for different types of spheres. Left two columns: *Mirrored Sphere*. Middle two columns: *Rough Specular Red Sphere*. Right two columns: *Diffuse Sphere*. Top row:(left) diffuse albedo (right) diffuse normal. Second row:(left) specular albedo (right) specular reflection vector. Third row:(left) specular reflection vector deviation from ideal sphere (blue = 0° , yellow = 5° , red $\geq 10^\circ$) (right) specular roughness (dark is sharp specular, light is broad specular). Fourth row:(left) reconstructed geometry (right) geometry deviation from ideal sphere (blue = 0mm, yellow = 0.25mm, red ≥ 0.5 mm; sphere diameter ranges from 7cm to 8cm). Fifth row:(left) validation photograph (right) point light rendering.

metallic rough specular sphere, and a diffuse sphere. Fig. 14 also presents plots of deviation in measured surface orientation and surface geometry compared to an ideal sphere, as well as rendering under point light illumination compared to validation photographs. As can be seen, the reconstruction error is low near the equatorial regions of the spheres that are unoccluded and well sampled by the camera viewpoints and higher near the poles due to occultations (bottom) and insufficient views for stereo (top). The reconstruction error near the poles could be reduced with data acquisition from additional viewpoints.

7 Discussion and Future Work

The principal advantage of our approach is the ability to estimate BRDF parameters, including specular roughness and anisotropy, over a complete set of object surface normals with a relatively small number of measurements. Recording the 0th, 1st, 3rd, and 5th bands of harmonics requires just 44 photographs (positive and negative), which can be acquired in minutes from multiple viewpoints. Since our LED arm creates images of incident illumination with 400×105 pixel resolution, it would take thousands of pho-

tographs to record the reflectance information one direction at a time, and hundreds of photographs using linear light source reflectometry. We believe our technique is the first to record diffuse and arbitrarily sharp specular SVBRDF behavior for 3D objects from a small number of photographs.

The proposed technique suggests several avenues for improvement. Currently, we estimate reflectance model parameters per pixel using images captured from a single viewpoint, and then we merge the reflectance parameter maps per view into maps covering the entire 3D object. This fails to make full use of the multiple viewpoints which may be available for a given surface point for BRDF fitting. We also do not consider self-shadowing or interreflections, assuming that each surface point receives light from the entire hemisphere around its surface normal. While the technique appears to degrade gracefully for small amounts of occlusion, reflective objects with significant concavities will not reconstruct well with this approach. We also note that the range of observable BRDFs exceeds the ones which can be expressed with a single diffuse/specular lobe reflectance model. For more interesting materials which must be modeled faithfully, using a more complex reflectance model fitted to higher-order SH responses may be required.




Figure 12: Two views of the 3D geometry for sunglasses (top), with a validation rendering (lower middle) and photo (bottom) with environmental illumination created by the LED arm.

8 Conclusion

We have presented a new technique for measuring the shape and reflectance of objects with arbitrary diffuse and specular reflectance properties at each surface point using spherical harmonic lighting based on a new continuous spherical illumination device. Unlike previous work which uses spherical illumination patterns, we avoid the problems associated with polarization-based reflectance separation and can measure the full range of specular materials from entirely diffuse to perfectly sharp specular. Furthermore, we leverage both the diffuse and specular reflectance maps to form surface correspondences in the geometry reconstruction process, allowing even textureless specular surfaces such as the lenses of sunglasses to be reconstructed accurately. While the technique is less applicable to translucent materials and geometrically concave shapes, it can reconstruct accurate models of a wide range of man-made objects usually deemed to be failure cases for existing 3D scanning and reflectance measurement approaches.

9 Acknowledgments

We thank Jonathan Coon, Sean Forsgren, Adam Gravois, Dominic Jones, Richard Adams, Callum Rex Reed, Randal Hill, Clarke Lethin, Bill Swartout, Randolph Hall, Ankur Agarwal, Kathleen Haase, Valerie Dauphin, Dava Cassoni and Santa Datta, for their support and assistance with this work. We also thank our anonymous reviewers for their helpful suggestions and comments. This work was sponsored by GLASSES.COM, the University of Southern California Office of the Provost, the U.S. Army Research, Development, and Engineering Command (RDECOM), NSF grant IIS-1016703, and a Royal Society Wolfson Research Merit Award. The content of the information does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.




Figure 13: Reflectance maps for a 3D model of a digital camera and a rendering and photograph with environment lighting.

References

- ADATO, Y., VASILYEV, Y., BEN-SHAHAR, O., AND ZICKLER, T. 2007. Toward a theory of shape from specular flow. In *Proc. IEEE International Conference on Computer Vision*, 1–8.
- BLAKE, A., AND BRELSTAFF, G. 1992. In *Physics-Based Vision, Principles and Practice: Shape Recovery*, L. B. Wolff, S. A. Shafer, and G. E. Healey, Eds. Jones and Bartlett Publishers, Inc., USA, ch. Geometry from specularities, 277–286.
- BONFORT, T., AND STURM, P. 2003. Voxel carving for specular surfaces. In *Proc. IEEE International Conference on Computer Vision*, 591–596.
- CHEN, T., GOESELE, M., AND SEIDEL, H. P. 2006. Mesostructure from specularities. In *CVPR*, 1825–1832.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1 (Jan.), 1–34.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of ACM SIGGRAPH 2000*, 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 (July), 98:1–98:10.
- FRANCKEN, Y., CUYPERS, T., MERTENS, T., GIELIS, J., AND BEKAERT, P. 2008. High quality mesostructure acquisition using specularities. *CVPR*, 1–7.
- FURUKAWA, Y., AND PONCE, J. 2009. Dense 3D motion capture for human faces. In *Proc. of CVPR 09*.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. In *ACM TOG*, 749–758.

- 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. *Comput. Graph. Forum* 28, 4, 1161–1170.
- GHOSH, A., HEIDRICH, W., ACHUTHA, S., AND O'TOOLE, M. 2010. A basis illumination approach to brdf measurement. *Int. J. Comput. Vision* 90, 2 (Nov.), 183–197.
- HARRIS, C., AND STEPHENS, M. 1988. A combined corner and edge detector. In *Proc. of Fourth Alvey Vision Conference*, 147–151.
- HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. 2005. A dual light stage. In *Proc. EGSR*, 91–98.
- HOLROYD, M., LAWRENCE, J., HUMPHREYS, G., AND ZICKLER, T. 2008. A photometric approach for estimating normals and tangents. *ACM Trans. Graph.* 27, 5 (Dec.), 133:1–133:9.
- 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, 4 (July), 99:1–99:12.
- IHRKE, I., KUTULAKOS, K. N., LENSCH, H. P. A., MAGNOR, M., AND HEIDRICH, W. 2010. Transparent and specular object reconstruction. *Computer Graphics Forum* 29, 8, 2400–2426.
- IKEUCHI, K. 1981. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Trans. Pattern Anal. Mach. Intell.* 3, 6 (June), 661–669.
- KOLMOGOROV, V. 2006. Convergent tree-reweighted message passing for energy minimization. *IEEE Trans. Pattern Anal. Mach. Intell.* 28 (October), 1568–1583.
- LAMOND, B., PEERS, P., GHOSH, A., AND DEBEVEC, P. 2009. Image-based separation of diffuse and specular reflections using environmental structured illumination. In *Proc. IEEE International Conf. Computational Photography*.
- 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 TOG* 22, 2, 234–257.
- 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. In *Rendering Techniques*, 183–194.
- MCALLISTER, D. K. 2002. *A generalized surface appearance representation for computer graphics*. PhD thesis, The University of North Carolina at Chapel Hill. AAI3061704.
- MORÉ, J. J., SORENSEN, D. C., HILLSTROM, K. E., AND GARBOW, B. S. 1984. The MINPACK project. In *Sources and Development of Mathematical Software*, 88–111.
- NAYAR, S., IKEUCHI, K., AND KANADE, T. 1990. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Trans. Robotics and Automation* 6, 4, 418–431.
- PARK, M., KASHYAP, S., COLLINS, R., AND LIU, Y. 2010. Data driven mean-shift belief propagation for non-gaussian mrfs. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, 3547–3554.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *Proc. of ACM SIGGRAPH '01*, 497–500.
- REN, P., WANG, J., SNYDER, J., TONG, X., AND GUO, B. 2011. Pocket reflectometry. *ACM Trans. Graph.* 30, 4 (July), 45:1–45:10.
- SATO, Y., WHEELER, M. D., AND IKEUCHI, K. 1997. Object shape and reflectance modeling from observation. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '97, 379–387.
- SLOAN, P.-P., 2008. Stupid spherical harmonics (sh) tricks. Game Developer's Conference, Feb. <http://www.ppsloan.org/publications/>.
- TARINI, M., LENSCH, H. P., GOESELE, M., AND SEIDEL, H.-P. 2005. 3D acquisition of mirroring objects using striped patterns. *Graphical Models* 67, 4, 233–259.
- WANG, C.-P., SNAVELY, N., AND MARSCHNER, S. 2011. Estimating dual-scale properties of glossy surfaces from step-edge lighting. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 30, 6.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.* 26, 2, 265–272.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. *SIGGRAPH Comput. Graph.* 26, 2 (July), 255–264.
- WEYRICH, T., LAWRENCE, J., LENSCH, H. P. A., RUSINKIEWICZ, S., AND , T. 2009. Principles of appearance acquisition and representation. *Found. Trends. Comput. Graph. Vis.* 4, 2 (Feb.), 75–191.
- WOODHAM, R. J. 1980. Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1, 139–144.
- ZICKLER, T. E., BELHUMEUR, P. N., AND KRIEGMAN, D. J. 2002. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *Int. J. Comput. Vision* 49, 2-3, 215–227.
- ZICKLER, T., RAMAMOORTHY, R., ENRIQUE, S., AND BELHUMEUR, P. N. 2006. Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *PAMI* 28, 8, 1287–1302.