

Simultaneous Acquisition of Microscale Reflectance and Normals

Giljoo Nam[†] Joo Ho Lee[†] Hongzhi Wu[§] Diego Gutierrez* Min H. Kim[†]

[†]KAIST

[§]Zhejiang University, State Key Lab of CAD & CG

*Universidad de Zaragoza, I3A




Figure 1: Overview of our system and results. (a) Our custom-built microscopic imager equipped with electrically-controlled LED lights over a fabricated dome, for simultaneous microscale capture of reflectance and normals. (b) A \$20 dollar bill captured by our system. We show the zoomed-in letter T as captured by Johnson et al. [2011], and compared with our geometric result (note that we had to use a different bill so an exact match is not possible). (c) Different from Johnson et al., we can simultaneously acquire reflectance information. We show a comparison of microscopic photographs of two details of the bill, and our reconstructed results from factorized basis BRDFs.

Abstract

Acquiring microscale reflectance and normals is useful for digital documentation and identification of real-world materials. However, its simultaneous acquisition has rarely been explored due to the difficulties of combining both sources of information at such small scale. In this paper, we capture both spatially-varying material appearance (diffuse, specular and roughness) and normals *simultaneously* at the *microscale* resolution. We design and build a microscopic light dome with 374 LED lights over the hemisphere, specifically tailored to the characteristics of microscopic imaging. This allows us to achieve the highest resolution for such combined information among current state-of-the-art acquisition systems. We thoroughly test and characterize our system, and provide microscopic appearance measurements of a wide range of common materials, as well as renderings of novel views to validate the applicability of our captured data. Additional applications such as bi-scale material editing from real-world samples are also demonstrated.

Keywords: material appearance, SVBRDF, microscopic imaging, photometric stereo

Concepts: •Computing methodologies → 3D imaging; Appearance and texture representations;

[†]Corresponding author e-mail: minhkim@kaist.ac.kr

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 ACM.

SA '16 Technical Papers, December 05-08, 2016, Macao
ISBN: 978-1-4503-4514-9/16/12

DOI: <http://dx.doi.org/10.1145/2980179.2980220>

ACM Reference Format

Nam, G., Lee, J., Wu, H., Gutierrez, D., Kim, M. 2016. Simultaneous Acquisition of Microscale Reflectance and Normals. ACM Trans. Graph. 35, 6, Article 185 (November 2016), 11 pages.
DOI = 10.1145/2980179.2980220 <http://doi.acm.org/10.1145/2980179.2980220>.

1 Introduction

Accurate depiction of material appearance has seen great progress in the last years [Dorsey et al. 2007; Weyrich et al. 2008; Kim et al. 2012; Kim et al. 2014]. It is well-known that the object-scale reflectance properties of materials are highly dependent on their underlying microscale geometry. This observation gave rise to microfacet theory [Torrance and Sparrow 1967; Cook and Torrance 1982], which explains surface reflection phenomena according to random irregularities at the *microscale* geometry level. Due to the extremely small size of microfacets, their distributions have been modeled statistically in existing works, instead of measuring them directly.

Related work in graphics has focused on image-based measurements of microscale geometry [Levoy et al. 2006; Johnson et al. 2011]. Dong et al. [2015] used a profilometer to measure microscopic geometry of metallic surfaces, from which reflectance information was *predicted*. Wang and Dana [2006] measured mesoscale appearance with a spatial resolution of 0.1 mm. Recently, appearance information from macro-scale photographs has been combined with acquired micro-geometry for high quality cloth modeling [Zhao et al. 2011]. In this paper, we aim to bridge the gap between microscopic imaging and material appearance acquisition: We present a system that allows for the *simultaneous acquisition* of both reflectance and geometry at the microscale, which had not yet been achieved due to the difficulties in combining both sources of information at such scale. To achieve our goal, the following **challenges** need to be overcome:

- Due to optics, the distance between the lens and the object needs to be very small (smaller as the magnification ratio increases). This constraint on the form factor poses a strong limit when trying to accommodate additional space to support any illumination structure. Moreover, one cannot directly apply traditional calibration algorithms involving a camera, a target, and a light source.
- Traditional reflectance modeling based on 4D reflectometry is not applicable to microscopic data. Since the depth of field

(DoF) of microscopic imaging is extremely shallow at micron scales, the microscopic imager should be oriented perpendicular to the surface to achieve in-focus images, effectively freezing the view vector.

- Imperceptible mechanical vibrations during the capture process, including vibrations from the shutter in a DSLR camera, will hinder accurate measurements. This is particularly problematic in the microscopic system, since each pixel captures information at sub-micron resolution, which is close to the limit for reflectance acquisition by itself. Moreover, we rely on high-dynamic-range (HDR) imaging to accurately obtain specular information, so vibrations will lead to registration issues in the captured images.

To tackle these challenges, we design and introduce a custom-built microscopic acquisition setup, which allows us to capture *spatially-varying* bidirectional reflection and normal maps at microscale *simultaneously*, at the highest resolution (620 nm per pixel) among existing systems. Note that although a commercial microscopic imager (like the Canon macro lens we use) can acquire images at micron resolution, corresponding geometry cannot be simultaneously captured at such scale. Our pipeline consists of two steps: microscale image acquisition, followed by microscale reflectance and normal estimation. Our custom-built hemispherical light dome, accommodating a dense array of LED lights, enables us to produce the structural light patterns of the spherical harmonics (SH) bases, used to capture microgeometry, as well as acting as point lights for measuring bidirectional reflectance. We first capture HDR images using spherical harmonic patterns and point-wise lighting. Next, we estimate SVBRDF and normals, in a joint optimization framework. To model our captured micro-scale data, we adapt the traditional Torrance-Sparrow model with a non-parametric normal distribution function (NDF) term. Our results can then be used to render novel views under different illuminations.

In summary, our main **contributions** are:

- A custom-built microscopic system for capturing microscale SVBRDF and normals simultaneously. Reflectance information is captured at the highest possible resolution, with geometric information matching this resolution.
- A microscale material appearance representation using our custom-built acquisition system. Moreover, the high complexity of the captured data is reduced to a manageable form suitable for editing, so that novel material appearances can be easily created from real-world data.
- An experimental dataset of real-world microscopic material appearance of a wide range of common materials, compared to microscopic photographs (to be released with the paper).

Nonetheless, our system is not free from limitations. At the micron scale, light transport is often dominated by subsurface scattering (SSS) in the thin outer layer. To allow for spatially-varying SSS measurements at this scale, structured light patterns of high frequency need to be used [Nayar et al. 2006]. To the best of our knowledge, there is no microscale projector that can support high-frequency patterns at micro-resolution within an area of approximately 1.5×1.5 mm. We therefore narrow our range of measured materials to relatively opaque, and let our diffuse term indirectly account for subsurface scattering in our reflectance model. Despite not explicitly capturing or modeling spatially-varying SSS, we show that our reflectance model represents our observations well, as demonstrated in the results and video.

The configuration of our system is described, along with the main design decisions made, in Section 3. We then describe our algorithmic workflow to obtain microscale reflectance and normals through alternating optimization in Section 4. The system is thoroughly

evaluated in Section 5, along with the presentation of our results. We hope our work inspires future work and helps to gain a deeper understanding of microscale material appearance, as well as to develop new multiscale models and novel acquisition strategies.

2 Previous Work

2.1 Reflectance Acquisition Systems

Multiple Views with Multiple Lights These works attempt to capture *all* combinations of the view and light directions [Lensch et al. 2003; Marschner et al. 2005; Lawrence et al. 2006; Holroyd et al. 2010] using a four-dimensional mechanical gantry. This approach allows to measure spatially varying reflectance of isotropic or anisotropic materials, although it requires both an elaborate hardware design and a long acquisition time. In contrast, the viewpoint in microscopic imaging must remain perpendicular to the surface, since the depth of field is extremely shallow. Moreover, since accurate acquisition of specular highlights requires HDR information, we need to avoid any vibration that would hinder microscale capture.

Fixed View with Multiple Lights Other methods measure reflectance and normals simultaneously by fixing the viewpoint and capturing shading under multiple light directions, e.g., [Holroyd et al. 2008; Ghosh et al. 2008; Alldrin et al. 2008]. These techniques jointly optimize surface normals and the parameters of a reflectance model iteratively. Recently, Aittala et al. [2015] propose a more efficient method for the particular case of texture-like materials, using just a flash/no-flash approach, while Dong et al. [2014] rely on motion information to reproduce appearance at object scale. Other systems use a linear light source to produce structured patterns, allowing for more efficient acquisition of isotropic [Ren et al. 2011] and anisotropic materials [Chen et al. 2014; Wang et al. 2008]. Structured lighting with varying frequency has also been proposed by Aittala et al. [2013] and Tunwattanapong et al. [2013]. Ghosh et al. [2009a] estimate per-pixel specular roughness using polarized second-order spherical gradients, using three different measurement setups. In later work, object-scale reflectance information is captured using four different polarizing filters [Ghosh et al. 2010]. Last, gray codes have been employed for acquiring gloss and normals of mesostructures [Francken et al. 2009].

Our microscopic acquisition system relies on a fixed viewpoint and electrically controlled point lights over a custom-built hemisphere. Different from previous works, designed for object-scale reflectance acquisition, we obtain both spatially varying reflectance and normals at the *microscale* level. In particular, we utilize both structured patterns and multiple-light reflection to acquire more accurate normals and to factorize images into the effects due to spatially-varying reflectance and normals.

Discrete Illumination Design Using discrete lights at fixed positions to measure reflectance has been attempted before [Marschner et al. 2005; Gu et al. 2006; Ghosh et al. 2008; Ghosh et al. 2009b]. As these works focus on *object-scale* acquisition of reflectance and normals, fundamental system designs and algorithmic solutions are different from ours. Ghosh et al. [2008] captured normals and a single BRDF per-region at object scale. Gu et al. [2006] and Marschner et al. [2005] measure object-scale SVBRDFs, without measuring geometry. Ghosh et al. [2009b] capture SVBRDF and normals simultaneously, but their object-scale approach is not directly applicable to microscopic imaging. Recently, Graham et al. [2013] introduced a novel microscale acquisition setup for human faces, based on photometric stereo, and fitting a parametric




Figure 2: (a) Schematic overview of our system design. (b) The object holder (a small drawer) on its three-axis micro-stage, firmly screwed to the optical table with vibration control. (c) The light dome is attached to the fixed structure of the object holder to avoid the need to recalibrate the system when changing a sample. (d) The wired dome with 374 white LED lights (as seen from the bottom). (e) The macro lens is positioned by the supporting structure on top of the light dome. (f) The structure tightly screwed to the table with vibration isolation.

BRDF model for each region of the face. As the authors found, their method is insufficient to capture microscale SVBRDFs of general materials, and their reconstruction method is tailored to the particular characteristics of skin.

2.2 Microscale Acquisition Systems

At the microscale, existing works either capture geometry or reflectance. Different from these, our work captures both microscale spatially-varying reflectance and normals *simultaneously*.

Microscale Geometry The first works on microscale acquisition focused on shape-only approaches, for instance using shape-from-focus or shape-from-specularity [Chen et al. 2006; Wang and Dana 2006]. Levoy et al. [2006] captured microscale shape using a light field lens array. Li and Li [2011] used photometric stereo with an optical microscope to obtain micron scale surface structure. However, their approach assumes Lambertian reflectance at the microscale. Recently, Johnson et al. [2011] also relied on photometric stereo, and used an elastomeric sensor with a reflective skin made of silver powder, to remove the influence of the optical properties of the material. The system is therefore incapable of capturing colors. Zhao et al. [2011] introduced micro CT imaging to capture the micro geometry of fabric. They successfully captured the high frequency details of fabric, which can then be used for rendering using a volumetric appearance model. However, since the system’s spatial resolution is higher than the frequency of visible light, it is impossible to capture color information from this setup. Dong et al. [2015] used a profilometer to measure the microscopic geometry of metallic surfaces, from which the reflectance of the surface is predicted, but not measured.

Microscale Reflectance Very few works exist that attempt to capture microscale material appearance. In conservation science, Kotoula and Kyranoudi [2013] presented a preliminary attempt to combine reflectance transformation imaging (RTI) with microscopy, yielding polynomial texture maps (PTM). The zenith angle is only coarsely sampled, and the dependence on the light direction is encoded as a low-order polynomial to a fixed viewpoint [Malzbender et al. 2001], which is fundamentally different from our goal of SVBRDF capture.

3 Microscopic Acquisition System

3.1 Design Decisions and System Description

Given the small scale of our acquisition setup, unique challenges arise during its design. A key aspect is *vibration control*, since we

aim to capture an area smaller than a micron per pixel. Thus the first decision is to avoid mechanical operations as much as possible, inevitable in existing four-dimensional gantries when moving light and sensor arms [Lensch et al. 2003; Marschner et al. 2005; Lawrence et al. 2006; Holroyd et al. 2010]; instead, we choose a design based on a light dome structure, where lights are controlled electronically. Moreover, since we need to capture multiple exposures to build an HDR image with different light vectors, capturing three thousand snapshots per object, we cannot use a high-resolution DSLR camera for the imaging sensor, since its mechanical shutter would introduce severe vibration; we choose to use a color machine vision camera instead. Last, to further stabilize our prototype, we build it on a custom-built block aluminum structure to which all its components are tightly screwed, while placing the entire system on a pneumatic vibration isolation table to eliminate even slight ground vibrations.

Given these key design decisions, our final system is made up of a three-axis micro-translation stage with a custom-built sample holder, a custom-built LED light dome, a machine vision camera equipped with a macro lens, and the global supporting structure. Figure 2 shows a schematic view (a), as well as close-ups of the different components, and a view of the entire system. We provide detailed descriptions of each component in the following paragraphs, along with other specific challenges that need to be overcome, and a discussion of the tradeoffs of our design.

XYZ Micro-Translation Stage As shown in Figure 2(b), we assembled a custom XYZ micro-translation stage from two off-the-shelf, micro-translation stages from Edmund Optics: one for XY translation, and the other for Z translation controlling the focusing distance, with micron resolution. The target sample is placed on a 3D-printed object holder, which slides in place and whose structure is firmly screwed to the translation stage. This allows to change samples without having to recalibrate the system (Section 3.2).

Light Dome As we have seen, our decision to design a micro light dome is largely motivated by the need to eliminate all mechanical vibrations during capture. This in turn forces us to accommodate its design within the restricted space that the working distance of the macro lens imposes, while packing as many lights as possible near the zenith angle without casting shadows on the object being captured. Our light dome consists of a 3D-printed hemispherical structure with a radius of 40 mm, where we accommodate 374 white LEDs (each with a 3 mm diameter) in dense 7-degree intervals, as Figures 2(c) and (d) show. Their intensity is controlled by an Arduino micro-controller and multiple TLC5940 chips, each providing 16 channels of 12-bit, pulse-width-modulation output. This allows for 4095 luminance levels for each LED, and discrete spher-




Figure 3: (a) An HDR capture of a chrome ball ($\phi=1$ mm) for calibrating the position of LED lights (all lights on). The yellow circle indicates the outer boundary of the ball. (b) Schematic diagram for calibrating the optical axis of the camera w.r.t. the stage coordinates. (c) Estimated 3D positions of all 374 LEDs on the hemisphere. White points show estimated positions and red crosses indicate the reference positions from the 3D model data of the hemisphere. Average angle difference is 0.68 degrees (std. dev.: 0.36).

ical harmonic illumination up to level 3 which will be used to capture geometry. Each LED is coupled with a plano-concave lens, to approximate ideal point lights. Accordingly, we will assume each LED to be a point light source when formulating light transport.

Camera and Macro Lens We use a trichromatic machine vision camera, Pointgrey Grasshopper3 (GS3-U3-120S6C-C), with a resolution of 4240×2824 pixels (12MP, pixel pitch: $3.1 \mu\text{m}$). This camera provides the highest resolution among cameras equipped with an electronic global shutter, thus avoiding mechanical vibrations. The camera is attached to a macro lens via a C-mount adapter. We employ a commercial DSLR macro lens (Canon MP-E 65 mm f/2.8), shown in Figure 2(e), with a magnification ratio up to 5:1. The pixel resolution of our microscopic system is 620 nm per pixel. Since the depth-of-field of our system is extremely shallow (less than $100 \mu\text{m}$, see Table 1), we do not control the focus distance directly with the lens; instead, we use the micro-stage (Z axis) as explained above. Note that the working distance of the lens at the 5:1 magnification ratio is 41 mm, which imposes additional constraints on the design of the light dome.

Supporting Structure As shown in Figure 2(f), the camera unit is attached to a supporting structure made of black aluminum profiles. The overall dimensions are $80 \times 40 \times 100$ cm. The structure is firmly screwed on a commercial optical table (Daeil DVIO-B-2410M-200T), equipped with pneumatic vibration isolation.

Discussion Some of the design decisions present some inevitable tradeoffs. First, the shallow depth of field forces the camera and the lens to remain perpendicular to the sample holder (the actual axis is obtained through calibration). Second, the light dome only allows for discrete incident light angles in 7-degree intervals. Last, given the relatively large form factor of the lens (diameter: 60 mm) and its working distance, this results in an inevitable hole at the top of the lighting dome, of ~ 20 degrees around the zenith axis, to allow the camera to capture images. We overcome these hardware tradeoffs algorithmically, as described in Section 4.

3.2 Calibration and Registration

Geometric Calibration Our goal here is to calibrate the light positions and the actual camera axis with respect to the surface of the object holder. However, the unavailability of suitable microscopic planar patterns (such as checkered boards) renders the traditional Zhang’s method [2000] invalid in our context. Our calibration follows three steps: We first calibrate the 374 lights positions with respect to the camera; second, we calibrate the stage’s normal vector, also with respect to the camera. Finally, we convert the camera




Figure 4: Even a sample of Spectralon of 99% diffuse reflectance presents clear specular peaks when captured at the microscale. Moreover, the diffuse and specular patterns that arise do not correspond to what would be expected when simply rendering the captured geometry.

view and all light vectors to the stage’s local coordinates. Each pixel shares a common view vector, assuming orthographic projection since the captured area is very small (1.5×1.5 mm), using a lens of 65 mm focal length.

For the first step, we employ a small, 1.0 mm diameter chrome ball as a light probe with all lights on [see Figure 3(a)]. We capture HDR images of the chrome ball under each light source (one light on, all the rest off) and estimate each initial light vector $\hat{\mathbf{i}}_k$ according to the mirror-reflection equation $\hat{\mathbf{i}}_k = 2(\hat{\mathbf{n}}_k \cdot \hat{\mathbf{o}})\hat{\mathbf{n}}_k - \hat{\mathbf{o}}$, where $\hat{\mathbf{o}}$ is the initial view vector $[0, 0, 1]^T$ in local camera coordinates, and $\hat{\mathbf{n}}_k$ is the corresponding surface normal of the chrome ball at the brightest pixel in camera coordinates, oriented toward the view vector $\hat{\mathbf{o}}$. In the second step, we find out the orientation $\hat{\mathbf{n}}$ of the object holder in camera coordinates from averaged shading information of each light, as Figure 3(b) shows. We then rely on photometric stereo to measure surface normals of a flat diffuse white surface placed on the holder stage. In the last step, we convert the camera coordinates to object stage coordinates with respect to an orthogonal vector frame with the same z-axis component of $\hat{\mathbf{n}}$. We transform the initial view vector $\hat{\mathbf{o}}$ in the camera coordinates to the actual view vector \mathbf{o} in the object holder coordinates by applying a rotation transformation on the vector frame. This completes our system calibration [see Figure 3(c) for light calibration results].

Registration During acquisition of the HDR images, any slight vibration may cause problems during registration, despite the pneumatic vibration isolation system. In particular, we found a long-term registration problem of up to 15 pixels (roughly $10 \mu\text{m}$). Since we capture specular highlights from various angles, most image registration algorithms would fail searching correspondences. To overcome this, we capture images with all the lights on at regular intervals, and use them for searching correspondences for rigid body registration.

4 Microscale Normals and SVBRDFs

Our system explores the co-design of novel hardware configurations and algorithmic solutions, to overcome some of the intrinsic tradeoffs of combining a light dome approach with microscopic imaging conditions. We show here how to factorize normals and SVBRDF from our captured data. Note that we use point-wise illumination only to capture reflectance, while we use both SH lighting and point-wise illumination to capture the surface normals. Please refer to the supplemental material for complete implementation details of all the contents in this section.

4.1 Microscale Normals

We first need to estimate per-pixel surface normals. As shown in Figure 4, many specular highlights on the outer layer can be seen at our capture resolution. Therefore common photometric stereo (PS) approaches that estimate *shape-from-shading* are inappropriate, since they would require separating the diffuse and specular information at microscale resolution. Ma et al. [2007] apply PS with polarization; however, this is impractical in our case, given our hardware configuration including densely packed LED lights in our small dome. Johnson et al. [2011] also apply PS using a flexible silver layer, although their goal is to capture geometry only, since the coated layer removes reflectance information. Recently, Nam and Kim [2014] introduced a hyperspectral photometric stereo method that effectively removes interreflection while estimating normals from diffuse reflection.

We therefore choose a *shape-from-specular* (SfS) approach [Chen et al. 2006; Tunwattanapong et al. 2013], given the large number of light sources in our dome. First, the initial surface normal at each pixel is computed from measurements under SH illumination (level 3) [see Figure 7(c) for an example]. This allows us to obtain normal cues from first-bounce specular reflections on the outer layer. However, our light dome design includes a hole of 20 degrees around the zenith axis to allow for imaging; this means that our setup misses discrete SH illumination patterns in that area, and as a consequence the corresponding obtained values tend to be clamped. To solve this, we first interpolate mirror-like reflection vectors illuminated by point lights circling the edge of the area where information is missing, and then update the clamped values and outlier artifacts with interpolated normals from the SfS method.

4.2 Microscale SVBRDFs

We now describe our algorithm for reconstructing SVBRDFs from the initially estimated normals and the captured radiance. Each BRDF is represented as a diffuse term and a specular term with a non-parametric, tabulated 1D normal distribution function [Torrance and Sparrow 1967]. We then factorize reflectance into basis BRDFs and corresponding blending coefficients, via alternating optimization.

Given our fixed-view configuration, we estimate spatially-varying BRDFs that share basis reflectance functions. We adapt a non-parametric factorization approach [Lawrence et al. 2006; Alldrin et al. 2008] and represent reflectance as a linear combination of shared basis BRDFs and corresponding weighting coefficients. By assuming that surface normals include enough angle variation, we can overcome the lack of light sources near the zenith axis and obtain enough mirror-angle observations. To take into account random irregularities in specular reflection at the microscopic resolution, we choose a non-parametric NDF [Ren et al. 2011], instead of relying on existing analytic functions such as Beckmann [Cook and Torrance 1982] or GGX [Walter et al. 2007], commonly used for object and mesoscale BRDF models.

4.2.1 Reflectance Model

The reflected radiance L at each point \mathbf{x} along the view direction \mathbf{o} under one directional light from \mathbf{i} can be obtained as

$$L(\mathbf{x}, \mathbf{o}) = R(\mathbf{x}, \mathbf{o}, \mathbf{i})(\mathbf{n} \cdot \mathbf{i})L(\mathbf{x}, \mathbf{i}), \quad (1)$$

where \mathbf{n} is the normal at \mathbf{x} (our normal map has μm resolution). We can express the reflectance term R to the diffuse and the specular terms, where the basis of our specular term is originated from the

Torrance-Sparrow model [1967] as follows:

$$R(\mathbf{x}, \mathbf{o}, \mathbf{i}) = \frac{1}{\pi} \rho_d(\mathbf{x}) + \rho_s(\mathbf{x}) \frac{D(\mathbf{x}, \mathbf{h})G(\mathbf{n}, \mathbf{o}, \mathbf{i})F(\mathbf{x}, \mathbf{h}, \mathbf{i})}{4(\mathbf{n} \cdot \mathbf{o})(\mathbf{n} \cdot \mathbf{i})}, \quad (2)$$

where ρ_d and ρ_s are diffuse and specular albedos at microfacet scale, $\mathbf{h} = (\mathbf{o} + \mathbf{i})/|\mathbf{o} + \mathbf{i}|$ is the half angle, D is the NDF term for specularity, G is the geometric term, and F is the Fresnel term.

NDF Our NDF is represented as a non-parametric tabulated function of 90 coefficients. We factorize the specular lobe as a single non-parametric NDF D , with the monotonicity constraint only, following Ren et al. [2011]. We found that this simple approach works well with our captured data.

Shadowing/Masking We formulate the shadowing/masking effects on both the light and view directions, following Ashikhmin's formulation [2000]. Since this formulation relates the shadowing/masking G and the distribution D functions, we apply an alternating optimization for the factorization of both terms, after we obtain the initial G based on V-grooves [Cook and Torrance 1982].

Fresnel To determine the Fresnel term, we require prior knowledge about the material properties, such as whether it is dielectric or metallic. Many recent works set the $F(0)$ term manually [Holroyd et al. 2008; Aittala et al. 2013; Ngan et al. 2005]. Given our form-factor constraints in our lighting dome, we do not capture grazing angles in general. Thus, to reduce complexity during the optimization, we simplify $F(\mathbf{x}, \mathbf{h}, \mathbf{i})$ to a constant color vector F per BRDF basis. In our formulation, the specular coefficient is a monochromatic scalar, and only the Fresnel term accounts for color.

4.2.2 Alternating Optimization

We initially estimate surface normals as described above, and optimize the basis BRDFs and its coefficients iteratively. Through 10-fold cross validation, we then repeat the optimization iteration until the test RMSE error starts to increase, following machine learning convention. The number of iterations varies depending on samples from one to three (usually two are sufficient). Instead of following Lawrence's non-parametric factorization of shading [2006], we combine an alternative photometric algorithm [Alldrin et al. 2008] with a non-parametric and isotropic NDF [Ren et al. 2011]. This configuration helps reduce overfitting in estimating both reflectance and normals in our experiments. We use a commercial sparse convex quadratic programming solver (e04nkc) [NAG 2015]. In addition, linear constraints are set to impose non-negativity on basis BRDFs, and the monotonicity of the distribution D . We found that no additional regularization terms for smoothness are required.

5 Evaluation and Results

5.1 System Evaluation

Radiometric Characterization We characterize the spectral power distribution of our white LEDs, shown in Figure 5(a). Since we use multiple LEDs, their intensity and color may vary slightly; we also build a light intensity profile for each LED by measuring the reflected light from a Spectralon tile for each light, taking into account the cosine law for each light direction. We then map device-dependent HDR RGB values of the trichromatic camera to device-independent values in the standardized linear sRGB color space, using an HDR color characterization method [Kim and Kautz 2008]. Figure 5(b) compares the colors captured by our microscopic imager with reference colors measured by a spectroradi-

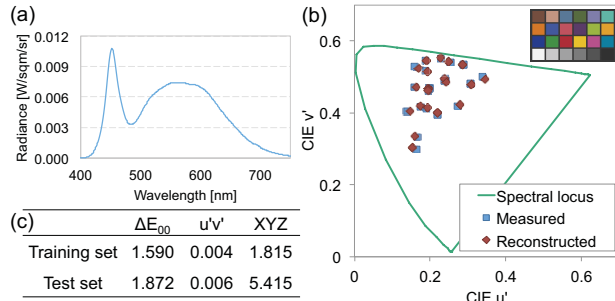


Figure 5: (a) Spectral power distribution of the LED lights, measured by a spectroradiometer. (b) & (c) Results of HDR color characterization. (b) compares the color difference on 24 color patches between our captured colors and reference measurements from a spectroradiometer, in the CIE $u'v'$ chromaticity diagram. (c) presents color differences of the training and the test datasets in CIE ΔE_{00} , $u'v'$ chromaticity, and the CIE XYZ color space (Y is normalized to 100).




Figure 6: (a) Short-term repeatability of the LEDs. We illuminate the measured surface 100 times in 300 seconds. The coefficient of variance of the measured luminance is 0.73%. (b) Long-term stability of the LEDs. The LED stays stable after approximately 30 minutes of warming up.

diometer (Jeti Specbos 1200¹) in the CIE $u'v'$ chromaticity diagram. In addition, we measured 24 standard color patches from the X-rite ColorChecker Passport to train the model, and tested it with eight different test colors [Figure 5(c)].

LED Light Characterization We have evaluated short-/long-term repeatability of the LED lights for validating stability. Figure 6(a) shows short-term stability when illuminating a surface 100 times in 300 seconds after 30 minutes of warming up. The coefficient of variation (CV) of the measured luminance levels is 0.73%. Figure 6(b) shows long-term stability. After approximately 30 minutes of warming up, luminance levels remain stable at 93% of the maximum luminance.

Stray Light Scattering Our dome structure is painted with black diffuse color to minimize the effect of interreflections; however, some scattering due to stray light is inevitable. To analyze its impact, we measured the ratio of the total scattered energy in the dome to the energy of one light, by capturing HDR images of our test chrome ball lit from different angles. Figure 7(a) shows some of these images, while Figure 7(b) shows the result. The amount of stray light depends on the direction of the incident light: from around 10 % when perpendicular, gradually decreasing to 2.7 % for grazing angles. When we capture a light probe under a point light source in an ordinary darkroom, the ratio ranges similarly from 9 to 5 %. The amount of stray light in our setup is therefore similar to object-scale darkroom setups.

¹Luminance accuracy: ± 0.05 at 1000 cd/m²; xy chromaticity repeatability: ± 0.0005 [Morgenstern et al. 2004]

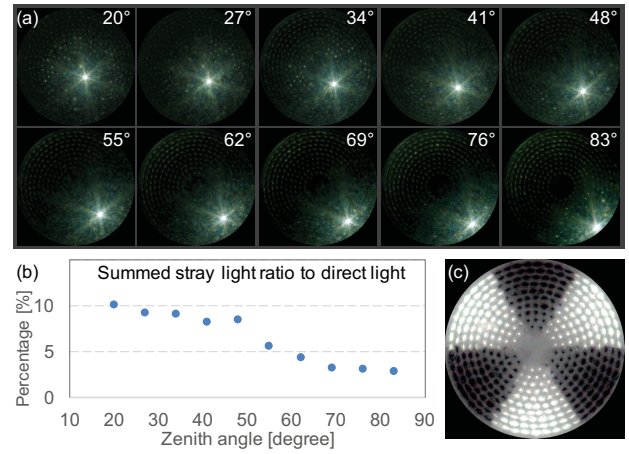


Figure 7: (a) HDR captures of a chrome ball lit from varying zenith angles. Note that the star-shaped spikes are typical diffraction artifacts produced by the aperture wings in the lens. (b) Measured ratio between the direct and stray interreflection light for varying latitudes. (c) An example of SH patterns (Y_3^{-3} of level 3) for capturing microscale geometry.

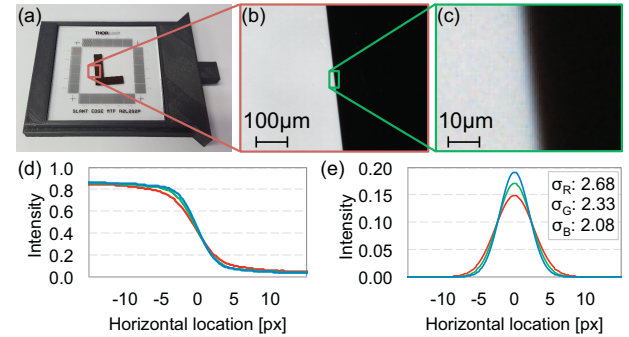


Figure 8: Point spread function with a 5:1 magnification ratio and $f/5.6$. (a) Reference MTF target. (b) & (c) Captured slanted edge. (d) & (e) Estimated edge and point spread functions for each color.

Diffraction Limit Given our large magnification ratio, Airy disks due to diffraction will inevitably appear, whose size is wavelength-dependent and inversely proportional to the numerical aperture size. The diffraction limit of our system can be evaluated by the size of these Airy disks as point spread functions (PSF). To measure it and analyze its impact, we capture an etched slanted edge of a standard MTF target (ThorLabs R2L2S2P), shown in Figures 8(a) – (c), from object scale to μm scale. Following Burns’s method [2000], we obtain the first derivatives of the edge spread functions (ESF) which lead to discrete PSFs for each channel; these are then fitted to Gaussian functions [see Figures 8(d) and (e)], finally applying non-blind deconvolution using the Richardson-Lucy method. Table 1 shows the measured resolving powers (and the corresponding DOFs) for each configuration we tested. The resolving power is calculated according to the Sparrow limit [Minin and Minin 2016], which is widely used for evaluating microscopic imaging performance. We highlight the best configurations in green, which yield an optical resolution of about 2–3 μm . We found that the heights of the micro-structure of the objects that we capture span approximately 50–100 μm . We therefore choose the camera configuration for each magnification ratio considering not only resolving limit, but also depth of field. The diffraction limit of this configuration is evaluated as PSFs in Figure 8.

(a) Resolving limit (object-side) [μm]					(b) Depth-of-field [μm]			
	f/2.8	f/4.0	f/5.6	f/8.0	f/2.8	f/4.0	f/5.6	f/8.0
1x	6.71	6.32	7.49	9.56	396	560	792	1,120
2x	4.65	4.72	4.58	7.06	148	210	297	420
3x	2.56	2.97	3.68	5.20	88	124	176	249
4x	2.82	3.05	3.45	4.79	62	88	124	175
5x	1.98	2.60	3.10	4.27	48	67	95	134

Table 1: Measured optical resolving limit (a) and corresponding depth of field (b) for available lens configurations in our system. The green cells indicate good configurations providing around $3\mu\text{m}$ optical resolution, accounting for diffraction.




Figure 9: We measured four standard surface roughness samples (Flexbar, Model 16008), described in (a), and compare the normal deviations measured by our system and a using a Nanofocus micro-surface profiler as ground truth, achieving a very good correlation.

Normal Accuracy We employ a standard surface roughness target [Flexbar, Model No. 16008; see Figure 9(a) for data of the four samples used]. We measured the four samples using our system and a micro-surface profiler (Nanofocus μsurf) as ground truth. Since the two systems work at different resolutions, direct comparison of measured normals per pixel is impossible. Instead, we compare measured surface roughness in terms of normal deviation, shown in Figure 9(b). It can be seen how both sets of measurements show a strong correlation ($R^2 = 0.9989$).

Reflectance Accuracy We first evaluate the microscale reflectance accuracy of our measurements on eight different materials. We conducted leave-one-out cross validation (LOOCV) instead, by randomly sampling 10 % of the 374 microscopic captures (one for each light). Table 2 summarizes the resulting factorization accuracy for our eight appearance datasets, measured as RMSE, PSNR, and SSIM. From this, we can see how our approach yields a very close reproduction of microscale appearance.

We qualitatively evaluate the accuracy of our system. We create synthetic images as input, using a reference surface [sample #3 in Figure 9(a), whose average height difference from the mean height (R_a) is $0.8\mu\text{m}$], and rendered with all 100 MERL BRDFs [Matusik et al. 2003]. From these rendered images, we apply our reconstruction algorithm, and compare the estimated reflectance with the ground truth BRDFs. Figure 10 shows two examples of specular and diffuse BRDFs, *steel* (a) and *green latex* (b), where we plot the reflection around the mirror reflection angle (135 degrees, red line). Overall our measurements present good agreement with the ground truth. Nevertheless, there are some noticeable differences between our estimated BRDFs and the ground truth, in particular near the mirror reflection in (a), and the near grazing angle in (b). This is due to our sampling strategy around the mirror reflection, which is limited by the zenithal hole and the density of our LEDs. Moreover, in our reconstruction algorithm we choose a single non-parametric representation for NDFs, following Ren et al. [2011]; while this allows us to factorize SVBRDFs and normals simultaneously, this could limit the accuracy in estimating sharp

Samples	Training set			Test set		
	RMSE	PSNR	SSIM	RMSE	PSNR	SSIM
Copper coin	1.4507	45.3860	0.9678	1.2557	45.0217	0.9742
Dollar bill (Tr)	1.9863	49.0532	0.9928	1.9107	46.3339	0.9886
Dollar bill (eagle)	0.3300	34.0659	0.7922	0.3244	34.3071	0.8047
Halftone printout	0.3791	34.4255	0.8434	0.3659	34.1288	0.8405
Leaf	0.2893	36.4879	0.8667	0.2940	36.3763	0.8653
Blue notebook	1.1011	46.1157	0.9916	0.9386	44.8029	0.9893
Leather	0.0231	42.2071	0.9660	0.0182	41.7205	0.9641
Textile	1.3310	38.0945	0.9196	1.3730	36.8318	0.9051

Table 2: Leave-one-out cross validation for accuracy test of our reconstruction (RMSE/PSNR/SSIM), on eight different materials.




Figure 10: We compare the measurements of our system with all 100 MERL BRDFs (green channel only), rendered over the standard surface sample [S_3 , shown in Figure 9(a)]. We show results for steel (a) and green latex (b). Our measurements show good agreement with the ground truth, with some noticeable errors around the mirror reflection and near grazing angles (see text for details).

specular reflections. Recently, Bagher et al. [2016] presented an effective non-parametric representation for homogeneous BRDFs, using *three* individual tabulated factors for the specular term: the NDF, the shadowing/masking, and the Fresnel term, respectively. This could be a good alternative to improve performance.

5.2 Captured Results

We now show results for our eight datasets shown in Table 2. The image acquisition process takes ~ 9 hours for each dataset; the decomposition process requires also ~ 9 hours to converge per iteration with an Intel i7 3770K processor and 32GB RAM. Our non-parametric, data-driven rendering is implemented on the LuxRender framework with a high-fidelity HDR color reproduction method [Kim et al. 2009; Kim 2010]. Our supplemental material includes animations, rendering novel views under dynamic lighting.

Microscale SVBRDFs and Geometry For our first example, we captured a \$20 dollar bill. Figure 1(b) shows our measured geometry of a $\sim 1\text{mm}^2$ area around the letter *T* in the bill, compared to the state-of-the-art geometry capture of Johnson et al. [2011]. Our captured data is factorized to microscale normals, which are converted to the geometric shape. This is done by solving an objective function that minimizes the difference between the partial derivatives of the geometry and the captured normals [Basri et al. 2007]. Even though the bill itself is different from the one that Johnson captured, we clearly see a very strong structural similarity and level of detail. In addition to capturing microscale geometry at state-of-the-art resolution, we can also capture microscale material

appearance simultaneously; Figure 1(c) shows a comparison with a microscopic photograph of the same region. Note that Johnson’s work cannot capture reflectance information.

Figure 11(a) shows a microscopic photograph of the eagle head in the bill, next to our reconstruction and a normal map by factorizing our data into four basis BRDFs and blending coefficients. Note how we can clearly separate the striped pattern on the background into the two ivory and yellowish lines, as well as the metallic flakes and greenish pigments of the head. Note also how the metallic specular function (material #4) shows a very sparse, sharp distribution.

Figure 11(b) shows a grayscale logo image, printed with a standard, high-resolution CMYK laser printer. The microscale appearance of a gray halftone printout is significantly different from object-scale appearance. We show here a $460 \times 460 \mu\text{m}^2$ area, photographed and reconstructed with our approach. The microscopic photograph clearly shows three different main pigments of cyan, magenta, yellow, as well as the white paper. Figure 12 shows the impact of the number of basis BRDFs for this example; when the number of basis BRDFs in our factorization matches the number of different materials in the actual sample, the reconstruction converges.

Figure 11(c) shows an example of a red leaf with green dots. Our reconstruction result is close to the microscopic photograph, while our captured normals clearly reveal the microscale surface of the leaf. Our SVBRDFs represent red, green and beige color components of the leaf. Note that we chose a particularly thick leaf to minimize the effects of surface scattering. Last, Figures 11(d) and (e) show a notebook cover and a coin, respectively. From these examples, we can see how our microscale reconstructions appear virtually identical to the original surfaces. Moreover, as shown in Figure 13, our data allow for novel rendered views under dynamic lighting, which indicates that our factorization does not suffer from overfitting (see video).

Comparison with Other Models We compare our results with the recent capture method by Aittala et al. [2015] (Figure 14). In particular, we capture and reconstruct a leather sample with Aittala’s (a) and our method (b). Note that since Aittala’s method is not designed for microscale capture, we apply it to object scale reflectance and normals for a fair comparison of scale. While their method provides a much simpler acquisition setup using a mobile device and yields plausible results at mesoscale, our proposed method captures much finer geometric and reflectance detail, yielding an accurate representation at microscale, shown in a comparison of our normals with the geometric measurement using a Nanofocus confocal 3D scanner. Both systems are therefore complementary, in the sense that they tackle the acquisition problem at different scales.

Second, we compare our SVBRDF to the existing microfacet reflectance models of Cook-Torrance and Oren-Nayar, on our textile dataset (Figure 15). Since these algorithms are originally designed for BRDF only, we share our captured weight map and normals for the side-by-side comparison. The parameters for each BRDF model are fit to the measured appearance using Levenberg-Marquardt non-linear optimization. Our method can capture better microscale appearance, with more details in the specular reflection.

5.3 Editing Microscale Materials

One additional advantage of our SVBRDF decomposition at such small scale is that it enables novel appearance edits from real-world data. Each captured microscale BRDF can be edited independently, from which object scale appearance changes emerge. This was shown in bi-scale material design for the case of *synthesized* data [Wu et al. 2011], focusing on efficient rendering of the edited

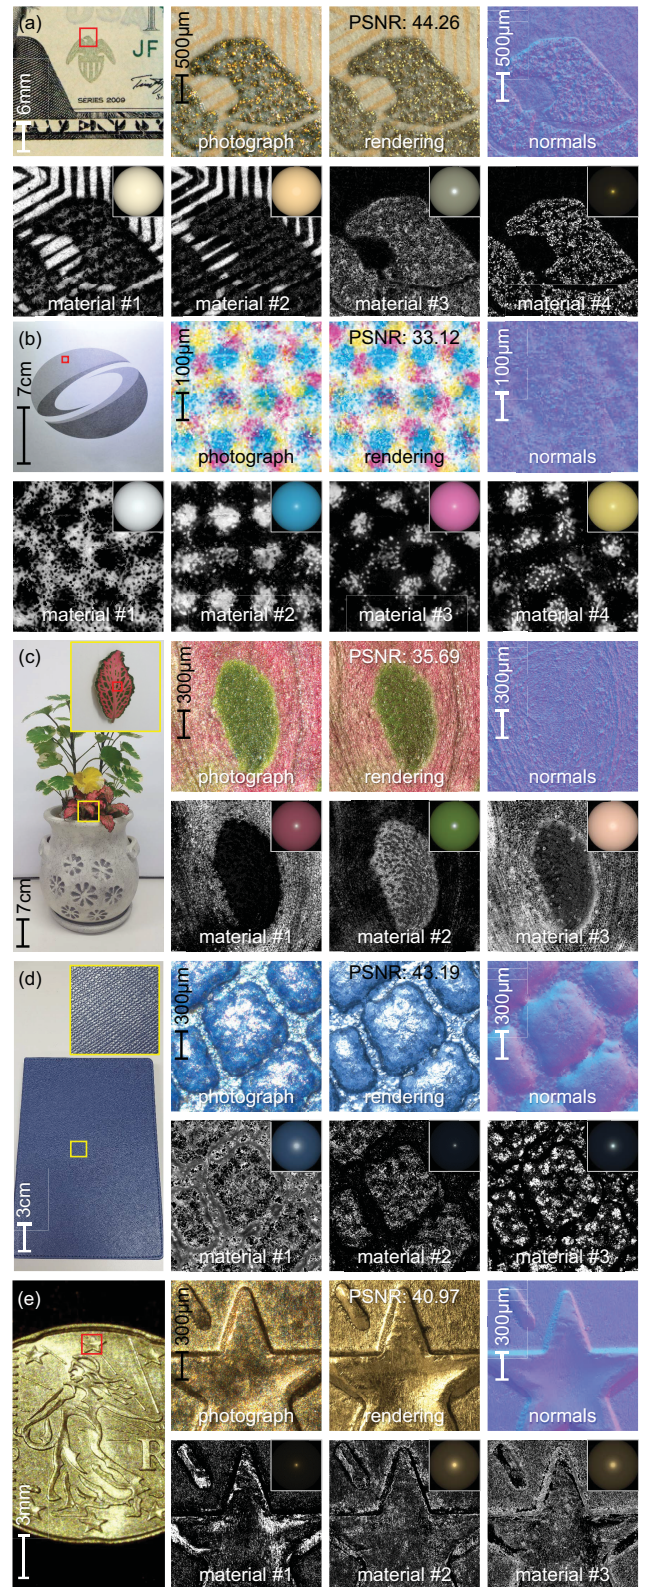


Figure 11: Results of our acquisition. For each material, we present an object-scale photograph, a micro-scale photograph, our reconstruction result, in addition to measured normals and decomposed basis SVBRDFs with corresponding weight maps for five daily life objects. More datasets are also shown in Figures 1, 14, and 15. Refer to the supplemental material for more results.

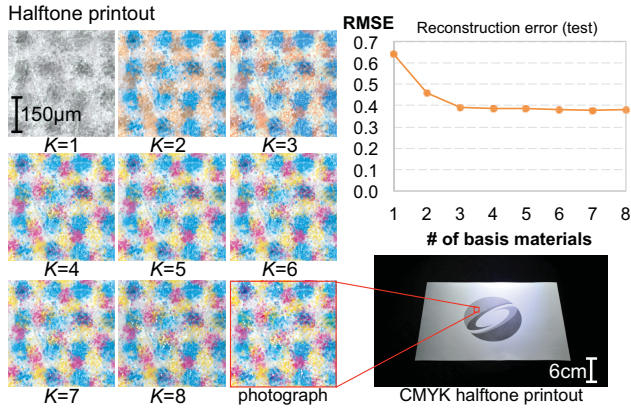


Figure 12: We analyze the impact of the number of basis BRDFs on the material shown in Figure 11(b). When the number of basis functions reaches the number of different materials in the actual sample, our reconstruction converges both numerically and visually.




Figure 13: Novel-view renderings under different view point and lighting (refer to the video in the supplemental material).




Figure 14: A leather sample captured for comparison. The first column shows object-scale and mesoscale appearance. (a) presents the captured SVBRDF using the method by Aittala et al. [2015]. Their method can reproduce plausible appearance and normals at mesoscale. (b) shows our reconstructed image and estimated normals. (c) provides the reference microscopic photograph and separate geometry from a confocal 3D micro scanner. The last two columns show that both geometric detail and material appearance at microscale are very well captured with our system.




Figure 15: Comparison of our BRDF reconstruction method, and other existing microfacet reflectance models.

objects. To illustrate this, we render object-scale appearance using a simple tiling synthesis from the measured SVBRDF, and apply supersampling to our synthesized weight maps and normal maps at rendering time. Figure 16 shows results using two different objects. Apart from generating strikingly different color appearances, we can also create an anisotropic appearance by altering our microscopic normals.

6 Discussion

There are many exciting opportunities for future work. The microscopic photographs of metallic surfaces reveal many irregular, complex, and multi-colored patterns, in addition to diffraction. Such irregular color patterns can be observed in Figures 17(a) and (b), and are due to interference effects (according to the Huygens-Fresnel principle). The physical computation of such interference patterns is challenging in our setup; we therefore simplify this effect using a Fresnel color constant in our reflectance model [Figures 17(c) and (d)], which does not reproduce such subtle color changes. Highly translucent materials are also challenging, as stated in the introduction and shown in Figures 17(e) through (h) for the case of artificial plastic fibers. Solving interreflection by applying hyperspectral imaging [Nam and Kim 2014] remains as our future work.

In summary, we have presented a novel microscopic system to simultaneously acquire microscale reflectance and normals, by combining microscopic imaging with reflectometry. This is the first time that this information is simultaneously captured at such small scale. We have shown examples of reconstruction and editing on many different materials, as well as favorable comparisons with state-of-the-art acquisition techniques. Several other potential applications can benefit from this work, such as data-driven, bi-scale material editing for novel material fabrication, digital forensics for detecting counterfeit notes or passports, etc. To foster further research on appearance acquisition, reconstruction, and editing, we have made our captured datasets available at <http://vclab.kaist.ac.kr/siggraphasia2016p2/>.

Acknowledgements

We thank the anonymous reviewers and Holly Rushmeier for their insights, as well as Adrian Jarabo, Daniel S. Jeon, Young Beum Lee for proof-reading and help. Min H. Kim gratefully acknowledges Korea NRF grants (2016R1A2B2013031, 2013R1A1A1010165 and 2013M3A6A6073718) and additional support by an ICT R&D program of MSIP/IITP (10041313). Diego Gutierrez acknowledges an ERC Consolidator Grant (682080), as well as the national MINECO grant (TIN2013-41857-P). Hongzhi Wu acknowledges NSF China (61303135) and the National Key Research & Development Plan of China (2016YFB1001403).

References

- AITTALA, M., WEYRICH, T., AND LEHTINEN, J. 2013. Practical SVBRDF capture in the frequency domain. *ACM Trans. Graph.* 32, 4, 110:1–12.
- AITTALA, M., WEYRICH, T., AND LEHTINEN, J. 2015. Two-shot SVBRDF capture for stationary materials. *ACM Trans. Graph.* 34, 4, 110:1–13.
- ALLDRIN, N., ZICKLER, T., AND KRIEGMAN, D. 2008. Photometric stereo with non-parametric and spatially-varying reflectance. In *Proc. IEEE CVPR 2008*, 1–8.

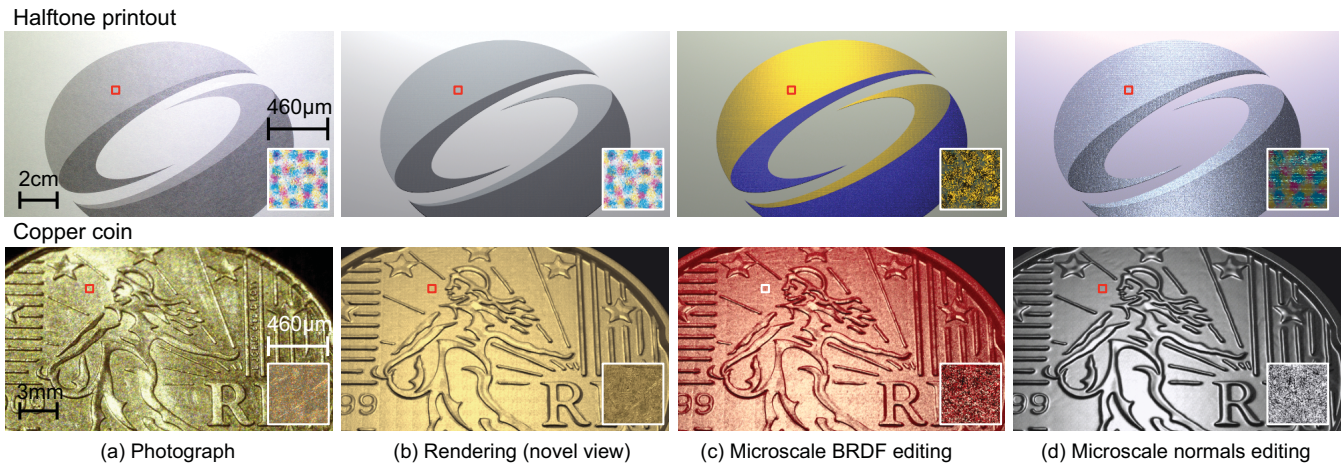


Figure 16: Example of microscale material editing. Column (a) shows object-scale photographs and microscale insets of $460 \times 460 \mu\text{m}^2$. (b) presents multi-scale rendering results using our captured SVBRDF, both at object- and micro scale, changing the lighting. (c) shows examples of edited appearance by swapping microscale BRDFs, including not only diffuse and specular albedo, but also microfacet distributions. (d) shows additional examples by editing the microscale normal maps into anisotropic (top) and flat (bottom) normals. Note that even though we use isotropic facet distribution functions, we can create anisotropic appearance by using microscopic normals.




Figure 17: (a) and (b) show color changes over microfacets depending on the light angle. (c) and (d) present our reconstruction using two different basis materials. Results remain as a single color consistently. We trace back that this color change occurs due to the Fresnel effect by mirror-like interreflection on microfacets, our reflectance model is incapable of capturing as we exclude interreflection of facets like other microfacet models.

ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Proc. ACM SIGGRAPH 2000*, 65–74.

BAGHER, M. M., SNYDER, J., AND NOWROUZSAHRAI, D. 2016. A non-parametric factor microfacet model for isotropic brdfs. *ACM Transactions on Graphics (TOG)* 35, 5, 159.

BASRI, R., JACOBS, D. W., AND KEMELMACHER, I. 2007. Photometric stereo with general, unknown lighting. *Int. J. Comput. Vision* 72, 3.

BURNS, P. D. 2000. Slanted-edge MTF for digital camera and scanner analysis. In *Proc. the Conference on Image Processing, Image Quality, Image Capture Systems (PICS-00)*, 135–138.

CHEN, T., GOESELE, M., AND SEIDEL, H.-P. 2006. Mesoscale structure from specularity. In *Proc. IEEE CVPR 2006*, 1825–1832.

CHEN, G., DONG, Y., PEERS, P., ZHANG, J., AND TONG, X. 2014. Reflectance scanning: estimating shading frame and

BRDF with generalized linear light sources. *ACM Trans. Graph.* 33, 4, 117:1–11.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1, 7–24.

DONG, Y., CHEN, G., PEERS, P., ZHANG, J., AND TONG, X. 2014. Appearance-from-motion: recovering spatially varying surface reflectance under unknown lighting. *ACM Trans. Graph.* 33, 6, 193.

DONG, Z., WALTER, B., MARSHNER, S., AND GREENBERG, D. P. 2015. Predicting appearance from measured microgeometry of metal surfaces. *ACM Trans. Graph.* 35, 1 (Dec.), 9:1–9:13.

DORSEY, J., RUSHMEIER, H., AND SILLION, F. X. 2007. *Digital Modeling of Material Appearance*. Morgan Kaufmann, Burlington, MA, USA.

FRANCKEN, Y., CUYPERS, T., MERTENS, T., AND BEKAERT, P. 2009. Gloss and normal map acquisition of mesostructures using gray codes. In *ISVC (2)*, Springer, vol. 5876 of *Lecture Notes in Computer Science*, 788–798.

GHOSH, A., HAWKINS, T., PEERS, P., FREDERIKSEN, S., AND DEBEVEC, P. 2008. Practical modeling and acquisition of layered facial reflectance. *ACM Trans. Graph.* 27, 5, 139.

GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. 2009. Estimating Specular Roughness and Anisotropy from Second Order Spherical Gradient Illumination. *Computer Graphics Forum* 28.

GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. 2009. Estimating specular roughness and anisotropy from second order spherical gradient illumination. In *Computer Graphics Forum*, vol. 28, 1161–1170.

GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. 2010. Circularly Polarized Spherical Illumination Reflectometry. *ACM Trans. Graph.* (Dec.).

GRAHAM, P., TUNWATTANAPONG, B., BUSCH, J., YU, X., JONES, A., DEBEVEC, P., AND GHOSH, A. 2013. Measurement-based synthesis of facial microgeometry. *Computer Graphics Forum* 32, 335–344.

- GU, J., TU, C.-I., RAMAMOORTHY, R., BELHUMEUR, P., MATUSIK, W., AND NAYAR, S. 2006. Time-varying surface appearance: acquisition, modeling and rendering. *ACM Trans. Graph.* 25, 3, 762–771.
- HOLROYD, M., LAWRENCE, J., HUMPHREYS, G., AND ZICKLER, T. 2008. A photometric approach for estimating normals and tangents. *ACM Trans. Graph.* 27, 5, 133.
- 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, 99.
- JOHNSON, M. K., COLE, F., RAJ, A., AND ADELSON, E. H. 2011. Microgeometry capture using an elastomeric sensor. *ACM Trans. Graph.* 30, 4, 46.
- KIM, M. H., AND KAUTZ, J. 2008. Characterization for high dynamic range imaging. *Computer Graphics Forum (Proc. EUROGRAPHICS 2008)* 27, 2, 691–697.
- KIM, M. H., WEYRICH, T., AND KAUTZ, J. 2009. Modeling human color perception under extended luminance levels. *ACM Trans. Graph. (Proc. SIGGRAPH 2009)* 28, 3, 27:1–9.
- KIM, M. H., HARVEY, T. A., KITTLE, D. S., RUSHMEIER, H., DORSEY, J., PRUM, R. O., AND BRADY, D. J. 2012. 3D imaging spectroscopy for measuring hyperspectral patterns on solid objects. *ACM Trans. Graph. (Proc. SIGGRAPH 2012)* 31, 4, 38:1–11.
- KIM, M. H., RUSHMEIER, H., FRENCH, J., PASSERI, I., AND TIDMARSH, D. 2014. Hyper3d: 3d graphics software for examining cultural artifacts. *ACM Journal on Computing and Cultural Heritage* 7, 3, 1:1–19.
- KIM, M. H. 2010. *High-Fidelity Colour Reproduction for High-Dynamic-Range Imaging*. Ph.D. Thesis, University College London.
- KOTOULA, E., AND KYRANOUDI, M. 2013. Study of ancient greek and roman coins using reflectance transformation imaging. *E-Conservation Magazine* 25, 74–88.
- LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.* 25, 3, 735–745.
- LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.* 22, 2, 234–257.
- LEVOY, M., NG, R., ADAMS, A., FOOTER, M., AND HOROWITZ, M. 2006. Light field microscopy. *ACM Trans. Graph.* 25, 3, 924–934.
- LI, Z., AND LI, Y. 2011. Microscopic photometric stereo: A dense microstructure 3d measurement method. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, IEEE, 6009–6014.
- 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 *Proc. Eurographics Conference on Rendering Techniques*, Eurographics Association, 183–194.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. J. 2001. Polynomial texture maps. In *Proc. ACM SIGGRAPH 2001*, 519–528.
- MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Trans. Graph.* 24, 3, 727–734.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July), 759–769.
- MININ, I., AND MININ, O. 2016. *Diffraction Optics and Nanophotonics: Resolution Below the Diffraction Limit*. Springer, New York.
- MORGENSTERN, T., BORNHOEF, G., AND GOERLICH, S. 2004. Miniaturized spectroradiometer. Light 2004, JETI Technische Instrumente GmbH, Brno, Jun.
- NAG, 2015. The NAG Library, Numerical Algorithms Group. <http://www.nag.com/>.
- NAM, G., AND KIM, M. H. 2014. Multispectral photometric stereo for acquiring high-fidelity surface normals. *IEEE Computer Graphics and Applications* 34, 6, 57–68.
- NAYAR, S., KRISHNAN, G., GROSSBERG, M. D., AND RASKAR, R. 2006. Fast Separation of Direct and Global Components of a Scene using High Frequency Illumination. *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH)* 25, 3, 935–944.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental Analysis of BRDF Models. *Rendering Techniques 2005*, 16.
- REN, P., WANG, J., SNYDER, J., TONG, X., AND GUO, B. 2011. Pocket reflectometry. *ACM Trans. Graph.* 30, 4, 45:1–10.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *JOSA* 57, 9, 1105–1112.
- TUNWATTANAPONG, B., FYFFE, G., GRAHAM, P., BUSCH, J., YU, X., GHOSH, A., AND DEBEVEC, P. 2013. Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Trans. Graph.* 32, 4, 109:1–12.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Proc. Eurographics*, 195–206.
- WANG, J., AND DANA, K. J. 2006. Relief texture from specularities. *IEEE Trans. PAMI* 28, 3, 446–457.
- WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph.* 27, 3, 41:1–10.
- WEYRICH, T., LAWRENCE, J., LENSCH, H., RUSINKIEWICZ, S., AND ZICKLER, T. 2008. Principles of appearance acquisition and representation. *Foundations and Trends in Computer Graphics and Vision* 4, 2, 75–191.
- WU, H., DORSEY, J., AND RUSHMEIER, H. 2011. Physically-based interactive bi-scale material design. *ACM Trans. Graph.* 30 (Dec.), 145:1–10.
- ZHANG, Z. Y. 2000. A flexible new technique for camera calibration. *IEEE Trans. Pattern Analysis and Machine Intelligence* 22, 11 (Nov.), 1330–1334.
- ZHAO, S., JAKOB, W., MARSCHNER, S., AND BALA, K. 2011. Building volumetric appearance models of fabric using micro CT imaging. *ACM Trans. Graph.* 30, 4, 44:1–10.