

Polarization imaging reflectometry in the wild

JÉRÉMY RIVIERE, Imperial College London
ILYA RESHETOUSKI, Imperial College London
LUKA FILIPI, Imperial College London
ABHIJEET GHOSH, Imperial College London




Fig. 1. **Polarization imaging reflectometry in the wild.** High resolution reflectance and normal maps of planar samples obtained by exploiting multiple polarization observations under uncontrolled outdoor illumination. The acquisition involves commodity photography equipment - a tripod-mounted DSLR camera (a,b,c) (see Figure 4), as well as hand-held acquisition using a mobile phone (d).

We present a novel approach for on-site acquisition of surface reflectance for planar, spatially varying, isotropic samples in uncontrolled outdoor environments. Our method exploits the naturally occurring linear polarization of incident and reflected illumination for this purpose. By rotating a linear polarizing filter in front of a camera at three different orientations, we measure the polarization reflected off the sample and combine this information with multi-view analysis and inverse rendering in order to recover per-pixel, high resolution reflectance and surface normal maps. Specifically, we employ polarization imaging from two near orthogonal views close to the Brewster angle of incidence in order to maximize polarization cues for surface reflectance estimation. To the best of our knowledge, our method is the first to successfully extract a complete set of reflectance parameters with passive capture in completely uncontrolled outdoor settings. To this end, we analyze our approach under the general, but previously unstudied, case of incident partial linear polarization (due to the sky) in order to identify the strengths and weaknesses of the method under various outdoor conditions. We provide practical guidelines for on-site acquisition based on our analysis, and demonstrate high quality results with an entry level DSLR as well as a mobile phone.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling; Computational photography**;

Additional Key Words and Phrases: Surface reflectometry, linear polarization, Stokes parameters, index of refraction

ACM Reference format:

Jérémy Riviere, Ilya Reshetouski, Luka Filipi, and Abhijeet Ghosh. 2017. Polarization imaging reflectometry in the wild. *ACM Trans. Graph.* 36, 6, Article 206 (November 2017), 14 pages. <https://doi.org/10.1145/3130800.3130894>

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.

© 2017 Association for Computing Machinery.
0730-0301/2017/11-ART206 \$15.00
<https://doi.org/10.1145/3130800.3130894>

1 INTRODUCTION

Accurately reproducing the appearance of real-world materials has been an active area of research in graphics and vision with a wide range of applications including visual effects, games, virtual reality, cultural heritage, advertising, design, analysis/recognition, etc. Advances in digital imaging over the last two decades has resulted in image based measurements becoming an integral component of appearance modeling. A recent and emerging trend here has been towards making acquisition and modeling more practical, employing commodity off-the-shelf equipment. However, despite these advances, the acquisition has for the most part still been limited to controlled indoor settings such as dark rooms.

In this work, we aim to extend surface reflectometry for the first time to completely uncontrolled outdoor settings, under natural daylight illumination. We propose a novel passive acquisition approach for estimating a complete set of reflectance maps for spatially varying planar surfaces exhibiting isotropic BRDFs using linear polarization imaging. We limit the data required for reflectometry by fusing observations of the polarization state of light reflected off the surface from up to three views – one near normal incidence, and two observations from orthogonal viewpoints near the Brewster angle of incidence. We motivate our approach of employing polarization cues for reflectometry with the following observations:

- Open sky is strongly linearly polarized due to single scattering of light by molecules in the atmosphere [Strutt 1871].
- Light from overcast skies tends to be unpolarized due to multiple scattering events in the clouds. However, light gets partially polarized upon reflection, with strong polarization achieved near the Brewster angle.

While some work exists on employing similar polarization imaging for reflectance estimation, these have been restricted to settings with controlled illumination, and assumptions of unpolarized [Miyazaki et al. 2003] or circularly polarized [Ghosh et al. 2010] incident illumination. In this work, we extend such measurements

to completely uncontrolled outdoor settings and address the challenges of reflectometry under general partially polarized incident illumination. A practical advantage of our approach is that it only requires standard photography equipment often employed on-site for image based lighting measurements [Debevec 1998]: a camera, a linear polarizing filter, a calibration target and a mirror ball (Figure 4). Besides achieving high quality results with a tripod mounted DSLR, we also present a few results acquired with hand-held imaging with a mobile phone (see Fig. 1).

In summary, the principal contributions of this work are:

- Practical method for passive acquisition of spatially varying surface reflectance with a few measurements under uncontrolled outdoor illumination using commodity equipment.
- Novel analysis of polarization imaging under general partial linearly polarized incident illumination and its application to on-site reflectometry in the wild.

2 RELATED WORK

There exist a significant body of prior work on reflectance capture (e.g., see [Guarnera et al. 2016; Weyrich et al. 2009]), with a primary focus on accuracy of measurements and reduction of the time-complexity of the acquisition process. Traditionally, reflectometry setups have been complex and suited only for laboratory-like settings. More recent work has however investigated development of practical acquisition techniques employing off-the-shelf equipment, as well as appearance estimation outside of laboratory environments. In the following, we discuss these latter approaches in more detail before discussing some related work on polarization imaging.

2.1 Commodity hardware

Most prior work on reflectometry using commodity hardware is designed around a camera-light source pair. Linear light sources are generally preferred for spatially varying samples as they help reduce the amount of measurements. Gardner et al. [2003] employed a linear light source mounted on a translation gantry in order to record per-pixel reflectance traces of planar samples. This design was later modified by Ren et al. [2011] for portability, and Chen et al. [2014] to allow measurement of anisotropic BRDFs.

An alternative to linear light sources, to help further reduce the amount of data and time spent in measurements, is to employ an LCD panel as an extended source of illumination. Ghosh et al. [2009] proposed to project second order spherical gradient patterns from an LCD screen to estimate per-pixel, spatially varying reflectance maps of planar samples and exploited the inherent polarization of LCD illumination for diffuse-specular separation. Francken et al. [2009] instead proposed employing Gray codes for this purpose, and more recently Aittala et al. [2013] have proposed to capture a sample's response to band-limited illumination patterns in the frequency domain using 2D Fourier patterns. For homogeneous samples, Wang et al. [2011] have proposed employing step edge illumination for a dual-scale statistical modeling of surface appearance.

Advances in mobile technology have recently given rise to more compact and portable designs for reflectance measurements. Wu & Zhou [2015] have proposed an integrated system for hand-held acquisition of shape and reflectance of objects with a Kinect sensor.

Aittala et al. [2015] have proposed a two-shot method for acquisition of stationary materials using a mobile phone. They employ a pair of flash-no flash observations (in general indoor environment) of the sample coupled with statistical analysis to extract reflectance maps. The method has been extended to a single flash image for stationary materials using neural synthesis [Aittala et al. 2016]. Riviere et al. [2016] have also recently proposed two mobile acquisition setups (one involving free-form acquisition with mobile camera-flash pair, and the other employing a tablet LCD screen) for acquisition of more general spatially varying planar surfaces with isotropic BRDFs. However, none of these above approaches can be exploited in general outdoor lighting conditions as they rely on active illumination.

2.2 Uncontrolled environment

Reflectometry in uncontrolled and/or unknown lighting environments is a very challenging problem that has recently attracted attention in the vision and graphics community. Romeiro et al. [2008] proposed an image-based method for passive reflectometry of a homogeneous curved object under known but uncontrolled lighting, such that each pixel of the image provides a linear constraint on the BRDF. They later extended the approach for reflectometry under unknown illumination that leverages a bi-linear constraint and the statistics of real-world illumination [Romeiro and Zickler 2010]. Glencross et al. [2008] proposed a depth hallucination method for diffuse textured surfaces. They require only two observations of a sample under diffuse outdoors lighting with and without flash illumination to recover albedo and surface shape. [Hauagge et al. 2014] also assume a Lambertian image formation model and a model of outdoor illumination to recover per-pixel albedo. Lombardi & Nishino [2012] have investigated reflectance estimation under unknown lighting for objects with known shape and homogeneous BRDF. They employ an expectation maximization approach with appropriate priors on both the BRDF (directional statistics) and illumination (natural image statistics) for this purpose. A similar approach has also been employed for simultaneous estimation of object shape and homogeneous BRDF under known illumination [Oxholm and Nishino 2012]. Dong et al. [2014] have further recovered spatially varying isotropic reflectance from a video of a rotating object (with known shape) under unknown lighting. Their approach alternatively estimates reflectance and lighting in an iterative process. This approach has just recently been further extended for simultaneous estimation of reflectance, object shape and unknown illumination [Xia et al. 2016]. While extremely powerful as a framework for reflectance estimation, appearance from motion does require a calibrated turntable setup.

In our work, we further push the limits of reflectometry to general uncontrolled outdoor settings by designing a portable solution that does not require active illumination or any translation/rotation gantry, while recovering highly detailed spatially varying surface reflectance. We do so by taking advantage of the inherent polarization of natural illumination and polarization from reflection.

2.3 Exploiting polarization

Polarization has been extensively studied in both vision and graphics, but mainly in strictly controlled environments where the polarization state of the incident light can be fine tuned by an operator. It

has proved to be a useful channel of information for shape estimation, material classification and reflectance components separation. The vast majority of previous work have studied the polarization resulting from reflection under unpolarized or linearly polarized incident light. Two notable exceptions are [Koshikawa 1992] and [Ghosh et al. 2010] who employ circularly polarized illumination to recover surface orientation and surface reflectance, respectively.

2.3.1 Reflectance separation/estimation. Appearance modeling methods strongly rely on the accurate separation of surface reflectance into its diffuse and specular components. Here, researchers have looked at polarization imaging, both exclusively [Debevec et al. 2000; Ghosh et al. 2008; Ma et al. 2007; Müller 1995; Wolff and Boulton 1991] as well as in conjunction with color space methods [Nayar et al. 1997; Umeyama and Godin 2004], for diffuse-specular separation. These methods all exploit the fact that diffuse reflection tends to depolarize incident polarized illumination due to multiple subsurface scattering, while specular reflection preserves incident polarization due to single bounce.

Ma et al. [2007] proposed combining polarization with spherical gradient illumination (using an LED sphere) to obtain high quality diffuse and specular albedo and normal maps. Their view dependent polarization solution was later extended for multi-view capture [Ghosh et al. 2011]. Closer to our work, Ghosh et al. [2010] have proposed measurement of the complete Stokes parameters of reflected circularly polarized spherical illumination to recover detailed reflectance parameters including index of refraction and specular roughness. Also related to our work is that of Miyazaki et al. [2003], who employ polarization imaging under unpolarized illumination coupled with inverse rendering in order to estimate shape, albedo and specular roughness of a homogeneous convex object. Their measurements are however made under controlled conditions with a few point light sources. In this work, we aim to extend such reflectometry using polarization imaging outside the laboratory to uncontrolled outdoor environments and resolve a more complete set of spatially varying reflectance parameters.

2.3.2 Surface normals estimation. Shape from polarization has been previously studied in the vision literature, under the restriction of uniform unpolarized incident illumination. In such conditions, the angle of polarization determines the direction perpendicular to the plane of incidence, which is the plane that contains the incident direction, the normal to the surface and the view direction. Two strategies are then typically employed to infer orientation.

The first approach relies on the degree of polarization and inverting the Fresnel equations. Here, most prior work has focused on shape from specular reflection and on solving the ambiguity in azimuth angle due to the degree of polarization reaching an extremum at Brewster angle [Guarnera et al. 2012; Saito et al. 1999; Thilak et al. 2007]. However, Atkinson et al. [2006] measure the degree of polarization due to diffuse reflection for shape estimation. Kadambi et al. [2015] have recently proposed a method to enhance coarse depth maps by fusing shape from polarization cues with the output of a depth sensor. They follow the unpolarized world assumption to estimate the zenith angle and employ the coarse 3D geometry to resolve the azimuthal ambiguity in polarization normals. Instead,

Smith et al. [2016] have recently proposed direct inference of surface depth instead of normals by combining specular and diffuse polarization cues with a linear depth constraint formulation. They demonstrate depth recovery under uncalibrated (unpolarized) point source as well as low order spherical harmonic illumination.

A second approach consists of observing the sample through multiple polarized cameras [Miyazaki et al. 2003; Rahmann and Canterakis 2001; Sadjadiz and Sadjadi 2007; Wolff 1989]. The key idea is then that one view constrains the surface normal to one plane and in theory only one additional view (and at most two [Wolff 1989]) are necessary to fully determine the normal to the surface. The advantage over the previous method is that no ambiguity exists in the determination of surface normal, at the cost of requiring carefully calibrated cameras in order to obtain per-pixel correspondences for each view. In this work, we follow a multi-view polarization approach and demonstrate its successful application under general outdoor sky polarization.

3 OVERVIEW AND BACKGROUND

In this work, we propose a passive method for reflectometry under uncontrolled outdoors lighting conditions. We combine observations of the polarization state of light reflected off a surface from multiple vantage points with inverse rendering to estimate high quality reflectance maps of the sample under consideration. The rest of the paper is organized as follows: we first briefly present some background on polarization and Mueller calculus in Section 3.1. We then develop our on-site reflectometry method in Section 4 by first deriving the Mueller calculus for polarization measurements in the wild (Section 4.1), which informs our measurement protocol (Section 4.2). Finally, we describe our data analysis pipeline in Section 4.3, before discussing results and limitations in Section 5. We refer to Section 1 of the supplemental material for a list of symbols and abbreviations used throughout the paper.

3.1 Polarization and Mueller Calculus

The polarization state of light can be formalized by Stokes parameters, expressed as a 4-vector $s = [s_0, s_1, s_2, s_3]^T$, where s_0 is the power of the incident beam, s_1 and s_2 respectively the power of 0° and $+45^\circ$ linear polarization and s_3 the power of right circular polarization. Upon reflection off a surface, the incident polarization state of light is altered according to Mueller calculus [Collett 2005]:

$$s_{ref} = M_{rot}(\phi)M_{ref}(\theta_i; \delta; \vec{n})M_{rot}(\psi_i)s_i \quad (1)$$

Here, $M_{rot}(\psi_i)$ rotates the incident Stokes vector s_i into a canonical frame of reference (plane of incidence), and $M_{rot}(\phi)$ rotates the reflected light into the camera's reference frame. $M_{ref}(\theta_i; \delta; \vec{n})$ is the concatenation of the Mueller matrices of a linear diattenuator, also referred to as Mueller reflection matrix, and a linear retarder of phase δ (refer to the supplemental material for more details).

Light in outdoor environments due to the sky is either unpolarized (on a cloudy day) or partially linearly polarized in the general case. Hence for such illumination, we only need to measure the first three components of the Stokes vector. This can be done by rotating a linear polarizer in front of a camera at three or more different orientations, e.g., 0° , 45° and 90° .




Fig. 2. *Plane of incidence*: Plane spanned by the normal \vec{n} to the surface, incident direction $\vec{\omega}_i$ and outgoing direction $\vec{\omega}_o$.

Input coordinate frame: $\vec{\omega}_i$, E_{\parallel} and E_{\perp} form a right-handed basis. ψ_i is the angle made by the polarization ellipse's fast-axis with E_{\perp} .

Output coordinate frame: $\vec{\omega}_o$, x_o and y_o form a right-handed frame. $\phi_{\vec{n}}$ is the angle made by the projection of the plane of incidence on the image plane with x_o .

Note that Eq. 1 holds for pure specular reflection only. Most real world materials also exhibit diffuse and rough specular reflection. We can account for rough specular reflections by modeling the surface with a microfacet BRDF where each microfacet behaves as per Eq. 1. The resulting Stokes vector can then be computed as:

$$s_o(\vec{\omega}_o) = \int_{\Omega} \left(\frac{\rho_d}{\pi} + f_r(\sigma; \eta; \vec{\omega}_i, \vec{\omega}_o) s_{ref}(\vec{\omega}_i) \right) L_i(\vec{\omega}_i) (\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i \quad (2)$$

Equation (2) is our complete image formation model, where we model $f_r(\sigma; \eta; \vec{\omega}_i, \vec{\omega}_o)$ as a Cook-Torrance microfacet BRDF [1982] with a GGX distribution [Walter et al. 2007]. The specular BRDF $f_r(\sigma; \eta; \vec{\omega}_i, \vec{\omega}_o)$ forms a narrow lobe around the reflection vector, within which the incident polarization can be assumed to be uniform as the polarization field varies smoothly over the sky [Können 1985]. Our goal is to recover the four parameters of diffuse albedo (ρ_d), index of refraction (η), surface normal (\vec{n}) and specular roughness (σ) from observations of $s_o(\vec{\omega}_o)$ under natural outdoors illumination.

4 POLARIZATION REFLECTOMETRY IN THE WILD

Outdoors illumination is known to be unpolarized in the special case of very overcast conditions, but partially linearly polarized under clearer sky conditions. While the latter case is more general and common for outdoor illumination, it has not been previously studied in the context of reflectometry. Hence, in the following we examine the more complex Mueller calculus of partially linearly polarized incident illumination.

4.1 Partially linearly polarized light

Partially linearly polarized light is the superposition of an unpolarized beam and a purely polarized beam [Collett 2005], where the contribution of the latter is mitigated by \mathcal{P}_i , the degree of linear polarization (DOLP) of the light. For light incident from the sky, we can therefore express its Stokes vector as:

$$s_i = L_i(\vec{\omega}_i) \begin{bmatrix} 1 \\ \mathcal{P}_i \cos 2\psi_i \\ \mathcal{P}_i \sin 2\psi_i \\ 0 \end{bmatrix} \quad (3)$$




Fig. 3. The intensity profile through a linear polarizer has the form of phase-shifted sinusoid of phase ϕ , which can be measured with only three measurements (red lines) at 0° , 45° and 90° .

In this form, we can conveniently describe the incident polarization direction in terms of the local coordinates defined by the plane of incidence. In this configuration, Eq. 3 is directly expressed in the coordinates of the reflecting element (see Fig. 2), and can be substituted in Eq. 1 to obtain the reflected Stokes parameters:

$$s_{ref} = M_{rot}(\phi) M_{ref}(\theta_i; \delta; \vec{n}) s_i$$

$$= L_i(\vec{\omega}_i) \begin{bmatrix} \frac{R_{\perp} + R_{\parallel}}{2} + \mathcal{P}_i \frac{R_{\perp} - R_{\parallel}}{2} \cos 2\psi_i \\ \frac{R_{\perp} - R_{\parallel}}{2} \cos 2\phi + \mathcal{P}_i * A \\ \frac{R_{\perp} - R_{\parallel}}{2} \sin 2\phi + \mathcal{P}_i * B \\ -\mathcal{P}_i \sqrt{R_{\perp} R_{\parallel}} \sin 2\psi_i \sin \delta \end{bmatrix} \quad (4)$$

where

$$A = \frac{R_{\perp} + R_{\parallel}}{2} \cos 2\phi \cos 2\psi_i - \sqrt{R_{\perp} R_{\parallel}} \sin 2\phi \sin 2\psi_i \cos \delta$$

$$B = \frac{R_{\perp} + R_{\parallel}}{2} \sin 2\phi \cos 2\psi_i + \sqrt{R_{\perp} R_{\parallel}} \cos 2\phi \sin 2\psi_i \cos \delta$$

Note that the rotation $M_{rot}(\phi)$ brings the resulting Stokes vector into the camera's coordinate system (Figure 2). In this general form, it is interesting to note the differences between Eq. 4 and corresponding simpler expressions found in the literature for incident unpolarized [Guarnera et al. 2012; Wolff 1989] and circularly polarized illumination [Ghosh et al. 2010].

- (1) The reflected radiance component $s_{ref,0}$ is affected by the polarization state of the incident illumination. Therefore, accurate inverse rendering might require in principle that we recover not only the intensity $L_i(\vec{\omega}_i)$ of the incident light but also its polarization state. We will discuss the challenges as well as some practical solutions for this in Section 4.6.
- (2) Unlike in previous work, it is not possible to estimate the azimuth $\phi_{\vec{n}}$ of the surface normal directly from the linear components $s_{ref,1}$ and $s_{ref,2}$ as they also depend on the incident polarization. We therefore describe another way of estimating $\phi_{\vec{n}}$ for obtaining cues for surface normal estimation in the following.

4.1.1 Transmitted Radiance Sinusoid (TRS). A known method for shape from polarization (unpolarized illumination assumption) is to consider the intensity profile of reflected light passing through a linear polarizer, which has the form of a phase-shifted sinusoid

(Figure 3) (consider for now that there is no diffuse component). Previous work [Huynh et al. 2010; Kadambi et al. 2015; Miyazaki et al. 2012] has shown the phase ϕ and the I_{min} and I_{max} amplitudes of the TRS can be recovered with just three measurements, e.g., with the polarizer at 0° , 45° and 90° orientations. Under unpolarized incident illumination assumption, the phase is directly related to the azimuth $\phi_{\vec{n}}$ of the surface normal as $\phi_{\vec{n}} = \phi + \pi/2$. We however found that this does not necessarily hold true for partially linearly polarized illumination, which in the general case has a more complicated expression for TRS (we refer to the supplemental material for details on the modified TRS expression).

However, from the general expression of the TRS under partial linear polarization, we derive three special cases to connect it back to the well-known expression under unpolarized incident illumination:

- (1) At the Brewster angle, dielectrics completely transmit the component of light parallel to the plane of incidence, i.e. $R_{\parallel} = 0$. This simplifies the TRS expression and we obtain measurements at the Brewster angle that behave as if the incident illumination were unpolarized.
- (2) Under horizontal or vertical polarization, i.e. $\psi_i = 0^\circ$ (resp. 90°), the TRS expression also simplifies back to an expression of the same form as unpolarized illumination.

In these particular cases, the expression for the TRS including the diffuse component (I_d) is given by:

$$I(\phi_o) = \frac{I_d}{2} + \frac{I_{\perp} + I_{\parallel}}{2} + \frac{I_{\perp} - I_{\parallel}}{2} \cos(2(\phi_o - \phi))$$

$$\text{where } I_{\perp} = L_i(\vec{\omega}_i) \frac{(1 + \mathcal{P}_i \cos 2\psi_i) R_{\perp}}{2} \quad (5)$$

$$I_{\parallel} = L_i(\vec{\omega}_i) \frac{(1 - \mathcal{P}_i \cos 2\psi_i) R_{\parallel}}{2}$$

Note that eq. (5) is a generalization which still holds under the ideal case of unpolarized incident polarization when $\mathcal{P}_i = 0$.

4.1.2 Implications. These observations inform us on how to develop our acquisition. The first point suggests that we make our measurements around oblique views close to the Brewster angle of incidence (Section 4.2). We will later show in Section 5 that finding the exact Brewster angle is not that important in practice and it suffices to be within a 15° window around the Brewster angle (which is easily judged visually) to obtain good qualitative and quantitative results. The second point informs us on when to best perform the acquisition: near mid-day when the sky is horizontally polarized [Können 1985], the polarization is mostly perpendicular to the incidence plane (for upwards facing samples) maximizing specular signal near the Brewster angle. We will later also discuss how our measurements are affected near sunrise/sunset when the sky is strongly vertically polarized in the direction orthogonal to the sun.

4.2 Acquisition

In light of the analysis presented in the previous sections, we propose a simple method for reflectometry in general outdoor conditions. Our primary measurement setup is composed of standard




Fig. 4. Our primary measurement setup composed of a DSLR camera, a linear polarizer, a mirror ball, and a color chart.

photography equipment often employed for image based lighting measurements (Figure 4):

- An entry level DSLR (Canon EOS 650D, 18 Megapixel), with an attached 18-55 mm lens. The camera sits on a tripod for stability during measurements.
- A glass linear polarizer which we mount on the camera lens. We mark the 0, 45 and 90 degrees orientations on the polarizer for precise rotations during measurements.
- A stainless steel mirror ball placed next to the planar sample during the capture process to record the incoming light.
- An X-Rite ColorChecker chart also placed flat next to the sample for white balancing and radiometric calibration.

We will later describe an even lighter weight hand-held capture process using a mobile phone in Section 5 that allows estimation of a useful subset of reflectance information. The measurement process typically proceeds as follows (Figure 5): We start by imaging the planar sample close to normal incidence to have a canonical view of reference. While it would be possible to simply image this view without a polarizer, we found it tedious to have to unmount and remount the polarizer between views. We then proceed to the measurement of linear Stokes parameters for two roughly orthogonal oblique views of the sample close to the Brewster angle of incidence to maximize the strength of the polarization signal and to be in the conditions where Equation (5) holds. While in principle Equation (5) is valid only at the Brewster angle, we found that in practice it suffices to just be near the Brewster angle (see Section 5.3 for an in-depth analysis).

Each set of oblique angle photographs consists of nine images: three different exposures (using auto-exposure bracketing on the camera) for HDR imaging for each of the three orientations of the polarizing filter. These shots are then combined to produce linear sRGB HDR images for all three polarization filter orientations using pfstools [Mantiuk et al. 2007] and subsequently employed for reflectance estimation. We also apply radiometric calibration to the acquired data by scaling the observed intensity of the white square of the color chart to match its sRGB value ($[.9, .9, .9]$) for each view. Note that this step also calibrates the brightness of the corresponding light probes. Typical measurements took us around 5 minutes per sample when uninterrupted, and up to 20 minutes when interrupted by passers by in some busy walk ways.

We further require camera pose estimation to recover surface normals from multi-view shape from polarization (Section 4.5), and for the estimation of specular roughness using inverse rendering




Fig. 5. Our method typically requires two sets of linear Stokes measurements close to the Brewster angle of incidence from near orthogonal view-points, and one observation near normal incidence for registration.

(Section 4.6). For this purpose, we employ VisualSfM [Wu 2011] which is a GUI-based software for structure from motion (SfM). Given the sparse set of views, in some cases VisualSfM has difficulties in finding matching features across or around the sample surface (e.g., for highly specular samples), resulting in 4-12 pixels in reprojection error. For such cases, we manually specify four corners of the sample for homography across the views and provide the homography based feature matches as input to VisualSfM for obtaining the camera poses with subpixel accuracy. This semi-automatic process for pose estimation also turns out to be very useful later for hand-held mobile acquisition (Section 5).

4.3 Reflectance extraction

We now describe our method for extracting high resolution reflectance maps from the acquired data. Recall that we took two sets of polarization measurements close to the Brewster angle, by sampling the intensity through a linear polarizer at $\phi_o = 0^\circ, 45^\circ$ and 90° orientation. From those observations, it is possible to fit a sinusoid as defined by Eq. 5 to obtain the three unknowns I_\perp, I_\parallel and ϕ . We start by rewriting Equation 5 as:

$$I(\phi_o) = \begin{bmatrix} 1 & \cos 2\phi_o & \sin 2\phi_o \end{bmatrix} \begin{bmatrix} \frac{I_d}{2} + \frac{I_\perp + I_\parallel}{2} \\ \frac{I_\perp - I_\parallel}{2} \cos 2\phi \\ \frac{I_\perp - I_\parallel}{2} \sin 2\phi \end{bmatrix} \quad (6)$$

The fitting problem becomes linear of the form $Ax = b$ which can be solved very efficiently using SVD. The intermediate result to this linear problem, $\hat{x} = [x_1, x_2, x_3]^T$ can then be exploited to obtain the three unknowns related to the TRS:

$$\hat{I}_\perp = \frac{I_d}{2} + I_\perp = x_1 + \sqrt{x_2^2 + x_3^2}, \hat{I}_\parallel = \frac{I_d}{2} + I_\parallel = x_1 - \sqrt{x_2^2 + x_3^2} \quad (7)$$

$$\phi = \frac{1}{2} \arctan \frac{x_3}{x_2}$$

4.4 Diffuse and specular albedo

4.4.1 Diffuse albedo. From Equation 7, it is clear that the diffuse component of reflection can be recovered from \hat{I}_\parallel , due to cancellation of specular signal near the Brewster angle. \hat{I}_\parallel is then related to the surface's diffuse albedo ρ_d as:

$$\hat{I}_\parallel = \frac{I_d}{2} = \frac{\rho_d}{2\pi} \int_{\Omega} \underbrace{(\vec{n} \cdot \vec{\omega}_i) L_i(\vec{\omega}_i)}_{\pi} d\vec{\omega}_i = \frac{\rho_d}{2} \quad (8)$$

Note that the integral under natural illumination is simplified because of the radiometric calibration step of Section 4.2. While any of the two views can in principle be employed for estimation of ρ_d , in practice we found slightly better specular cancellation in one of the two views in some cases and chose the minimum of the two results as the diffuse albedo.

4.4.2 Specular reflectance. By subtracting \hat{I}_\parallel from \hat{I}_\perp obtained from the same view point as that used to estimate ρ_d , we obtain a diffuse-free image that encodes R_\perp up to a scale factor that depends on the polarization state of incident light (Equation 5). Without knowledge of the latter, finding the specular albedo is thus an ill-posed problem. To overcome this, we propose a template based approach, where we employ the specular response of the black plastic casing around the color chart as our template. The rationale for this is because our samples are mainly planar, they are subject to the same incident illumination as the specular reflection on the planar casing around the color chart.

Given that the chart casing is made of plastic with a known index of refraction $\eta_{chart} = 1.46$, we pre-compute its perpendicular reflection coefficient at the Brewster angle under uniform spherical illumination, $R_{\perp,real}$. We then compute the scale factor between the plastic casing's measured diffuse subtracted maximum intensity and the pre-computed $R_{\perp,real}$. The same scaling factor is then applied to the sample's diffuse subtracted maximum intensity to obtain an estimate of R_\perp for the sample. We finally apply a method akin to that of [Ghosh et al. 2010] to estimate a per-pixel index of refraction:

$$\eta^2 = \frac{(1 + \sqrt{R_\perp(\theta_B)})}{(1 - \sqrt{R_\perp(\theta_B)})} \quad (9)$$

In our case, Equation (9) is a simplification of the more general equation given that $R_\parallel = 0$ at the Brewster angle. From the estimated η , we can then compute the reflectance at normal incidence as used in Schlick's approximation [Schlick 1994] as $R_0 = \frac{(n-1)^2}{(n+1)^2}$. Either parameter (η or R_0) can then be used with their corresponding equations to model Fresnel effects in renderings.

4.5 Surface normal estimation

We formulate our normal estimation in a multi-view shape from polarization framework, akin to [Miyazaki et al. 2012]. Consider the tangent vector $\vec{b} = [\cos(\phi), \sin(\phi), 0]^T$ orthogonal to the surface normal \vec{n} . Essentially, the knowledge of \vec{b} constrains the surface normal to lie in the plane of incidence. Therefore, by combining two or more observations from different viewing directions it is possible to fully resolve the normal to the surface.

Given two views close to Brewster angle, whose camera rotations in world coordinate are define as R_1 and R_2 and the observation of \vec{b}_1 and \vec{b}_2 , the surface normal can be estimated by solving the following linear system of equations:




Fig. 6. Change in reflected radiance $s_{ref,0}$ due to incident polarized illumination ($DOP = 80\%$) is very similar for stainless steel (solid plots) and a dielectric ($\eta = 1.5$, dotted plots) around the Brewster angle of incidence. The colors indicate three different angles of polarization w.r.t. plane of incidence (Red: $\psi_i = 0^\circ$, Green: $\psi_i = 45^\circ$, Blue: $\psi_i = 90^\circ$).

$$\begin{bmatrix} R_1^T b_1 \\ R_2^T b_2 \\ 0 \end{bmatrix} \begin{bmatrix} n_x \\ n_y \\ n_z \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad (10)$$

The above problem can be solved in closed-form for two viewpoints as the intersection of the two planes of reflection, as well as using SVD decomposition. We employ SVD since it is a more general solution extensible to more than two viewpoints.

4.6 Specular roughness

We formulate our specular roughness estimation as a least squares problem within an inverse rendering framework. Given a set of observations of a sample under natural illumination $s_{0,i}, i \in [0, 2]$, we compute $\hat{s}_{0,i}$ using Equation 2, where ρ_d, \vec{n}, η are estimated from the previous sections, leaving only σ as an unknown. We thus solve for the σ that minimizes the residual (in the least squares sense) between our rendered images $\hat{s}_{0,i}$ and photographs $s_{0,i}$:

$$\min_{\sigma} \sum_i \frac{1}{2} \|s_{0,i} - \hat{s}_{0,i}(\sigma)\|_2 \quad (11)$$

This can be efficiently solved using Levenberg-Marquardt (LM) non-linear optimizer. We employed the dlib [King 2009] implementation of LM for this purpose.

It is worth pointing out here that in our inverse rendering step, we did not explicitly account for the polarization of the incident illumination which, as shown in Equation 4, has an effect on the reflected radiance. This is because there are practical challenges with doing so in our measurement setup: we employ a stainless steel mirror ball as a light probe in the scene. Even though we record measurements of the linear Stokes parameters of reflection on the mirror ball, Equation 4 shows that under the general case of incident partial linear polarization, the reflected light becomes somewhat elliptically polarized. In fact, ellipticity of reflected light has previously been employed for metal vs. dielectric classification [Chen and Wolff 1998]. To accurately recover the incident polarization using the mirror ball would require an additional measurement with a circular polarizer to recover the ellipticity which would add measurement complexity (e.g. employing a filter wheel to switch between different polarizers).

Instead, we propose a practical solution based on the observation that the incident radiance recorded using the mirror ball already encodes the modulation of intensity in the s_0 component due to the incident partial linear polarization. This is why light probes captured in outdoor environments exhibit darker and brighter sections in

the sky due to polarization effects. Figure 6 shows that this change in intensity of reflected light observed on a stainless steel mirror ball is very similar to that on a dielectric around the Brewster angle of incidence. Hence, we employ view dependent light probes for the inverse rendering step, where we render images $\hat{s}_{0,i}$ using the corresponding view's light probe to compute the error with respect to the corresponding photograph $s_{0,i}$. We found that this first order approximation of incident polarization with view dependent light probes gave satisfying results in practice (see Section 5).

5 RESULTS AND ANALYSIS

5.1 DSLR acquisition

We first present some results obtained with our primary acquisition setup involving a tripod mounted DSLR. To evaluate our approach, we measured a few planar samples outdoors, under various illumination conditions and different locations. Figure 7 shows the reflectance maps estimated from the method presented in Section 4 for seven exemplar datasets. We chose these datasets as they exhibit a wide variety of reflectance properties and were captured in different lighting conditions. It can be seen from Figure 7 that our method is fairly agnostic to the incident illumination and is able to produce good qualitative results under various conditions (Figure 8). Note that the majority of these datasets cannot be acquired in a controlled setting as they are permanent on-site structures.

Our dielectric reflectance assumption is not strictly true for the “specular drain cover” (Figures 7 and 8, 1st row). However, we obtain good qualitative results for this metal-dielectric composite (cast-iron) material. Being an outdoor surface, the metal is fairly oxidized adding to the dielectric like behaviour of the polarized reflectance. The data was acquired on a busy street sidewalk with surrounding buildings, showing that our method is robust to cluttered environments. The “red bricks” (Figures 7 and 8, 2nd row) sample is diffuse dominated and was captured in an early afternoon with thin clouds in the sky. The “stone pavement” (Figures 7 and 8, 3rd row) data was acquired in the late afternoon (4pm) under a very sunny sky condition. Our method is able to effectively resolve strong specular reflections (glints) on the surface and mesostructure of the stone slabs (see accompanying video). The “garden pavement” (Figures 7 and 8, 4th row) was acquired under overcast (unpolarized) conditions in a back-garden surrounded by buildings. Our method is successful in estimating reflectance maps even for this mostly diffuse surface. The surface normals however are a bit noisy due to drop in signal-to-noise ratio (SNR) of the specular signal. The “red book” (Figures 7 and 8, 5th row) demonstrates the results for a hard book cover acquired under partially overcast conditions. The book surface has prominent bumps which form a meandering pattern which is faithfully represented in the estimated surface normal map and in the spatial variation in the specular roughness map. The “canvas print” (Figures 7 and 8, 6th row) shows significant spatial variation in specularity, with stronger specularity towards the bottom due to printing over white canvas. Finally, “wooden bench” (Figures 7 and 8, bottom row) presents the results for two planar panels of a wooden bench which are connected by a supporting beam from below. While mostly diffuse, the panels do exhibit some




Fig. 7. Reflectance maps [(a)-(d)] estimated from two views of the sample close to the Brewster angle of incidence, under natural outdoors illumination (e). Our method is agnostic to the incident illumination and robust to changes in illumination during capture.

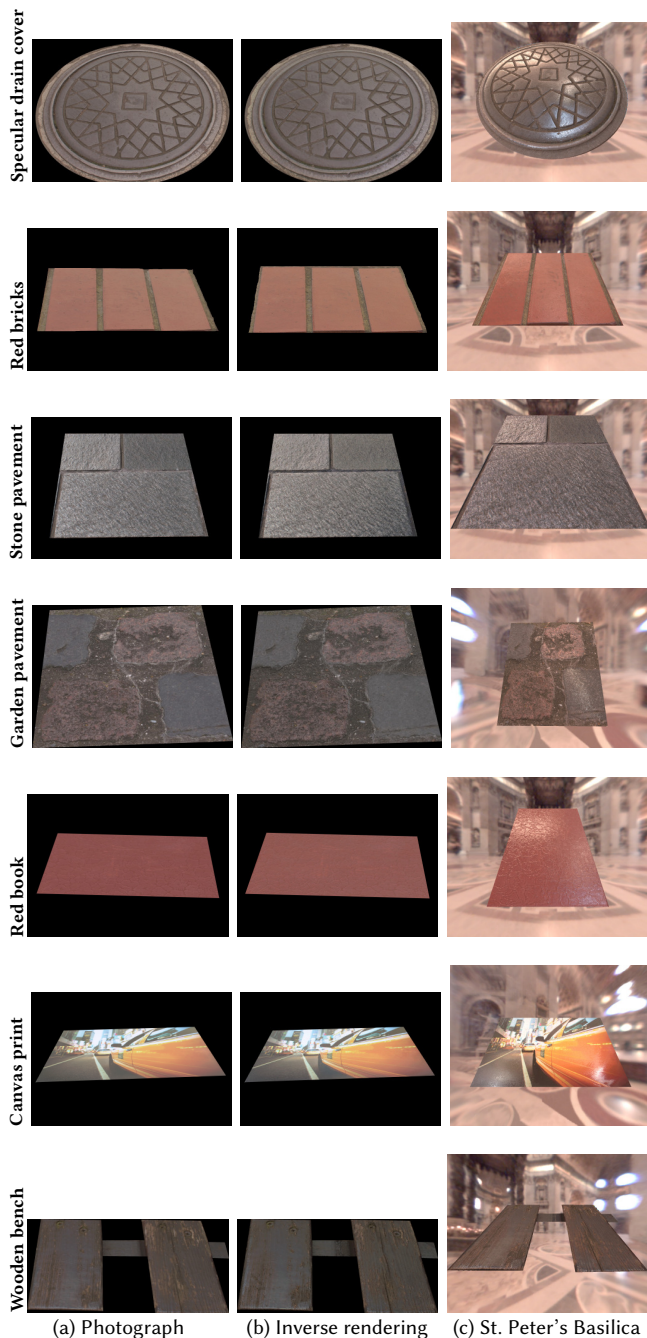


Fig. 8. Comparisons of sample photographs (a) taken under the illumination conditions shown in Figure 7 with inverse rendering results (b) under the same incident illumination. Here, we also show rendering under a novel lighting environment (c).

irregular specularity on the surface due to worn out layer of varnish (see accompanying video). Some occlusion artifacts can also be seen on the support beam which is to be expected. Here, the data was captured in an outdoor patio with surrounding buildings. We

provide additional rendering comparisons under novel viewpoint and lighting in the supplemental document.

5.2 Mobile acquisition



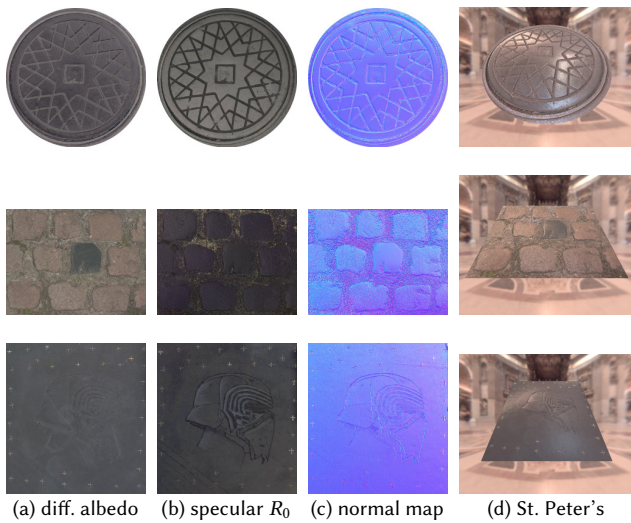
Fig. 9. Polarization measurements with hand-held acquisition with a mobile phone. A static polarizer is mounted in front of the lens and the entire device is rotated in-plane to the desired 0° , 45° , and 90° orientations to acquire the data.

We additionally present a few results that we acquired using light weight hand-held acquisition using a mobile device. Here, we replaced the DSLR (and tripod) with an Android smartphone (Samsung Galaxy S4) for the acquisition. The primary (back) camera on the S4 has 13 megapixel resolution, and a lens with 31mm fixed focal length and $f/2.2$ aperture. Unlike the DSLR where we were able to screw a glass linear polarizer to the lens while still allowing for polarizer rotation, such a system was not available off-the-shelf for a phone camera¹. Instead, we statically mounted a laminated plastic sheet of linear polarizing film (Edmund Optics) in front of the lens, while aligning the linear polarizing axis with the phone held in landscape mode. Our hand-held capture process then proceeds as follows: we actually rotate the entire phone (and its camera) sequentially in-plane from landscape mode (0°) all the way to portrait mode (90°), while orienting the phone diagonally in between to record the scene with a 45° polarizer orientation (Figure 9). In order to do this hand-held rotation precisely, we rely on the internal accelerometer and gyroscope readings of the device and display the current phone orientation (as estimated by the sensors) on the screen.

Note that in practice, we were not able to take the three photographs at precisely the desired 0 , 45 and 90° orientations. However, such precise rotation is not really required by the method as the TRS equation can be solved in principle with any combination of three well spaced orientations. Hence we just allowed the user to get as close as possible to the target orientations and record the images while storing the actual recording orientations for TRS. Also, given the hand-held nature of the acquisition, we limited ourselves to single exposure capture in order to minimize motion blur. We ensured that the camera was well exposed for the reflectance observed on the sample and the color chart within a single exposure. However, this resulted in the mirror ball being recorded over-exposed for the sky illumination with our chosen exposure.

Compared to DSLR acquisition, an important difference that arises with the above acquisition process is that the entire scene is recorded at different orientations for the three polarizer orientations. Hence, before alignment of data across viewpoints, we need to first align data acquired from a single viewpoint but with different device orientations. We found our semi-automatic tool for homography

¹Clip-on circular polarizing lenses are currently available for phone cameras, but not linear. <https://photojojo.com/awesomeness/cell-phone-lenses>



(a) diff. albedo (b) specular R_0 (c) normal map (d) St. Peter's
 Fig. 10. Reflectance maps [(a)-(c)] estimated using hand-held acquisition with a mobile phone. Here, in-plane hand-held rotation of the mobile device is used to record the polarization data. Rendering in a novel lighting environment (d) with the acquired data.

based alignment to be very useful for this purpose for aligning the three camera orientations of a viewpoint. We then employed VisualSfM as before for alignment of data across the three primary viewpoints as well as to obtain camera parameters for these views.

Figure 10 presents the recovered reflectance parameters, including diffuse albedo (ρ_d), specular reflectance (R_0) and surface normals (\vec{n}), of a few planar samples with the above hand-held mobile capture process. This includes the specular drain cover that was previously also acquired with the DSLR setup, a brick pavement, and a sketch book cover. As can be seen, the maps of the specular drain cover are qualitatively very similar to the results obtained with the DSLR setup. Our method is able to estimate detailed surface normals and spatially varying specular reflectance for the diffuse dominated brick pavement dataset. Finally, the sketch book cover has a very interesting emboss pattern depicting Star Wars character “Kylo Ren” in the specular reflectance and the normal map. This embossing is hardly visible in the corresponding diffuse albedo and correctly attributed to the specular reflectance.

Note that due to the saturation of the light probes with single exposure acquisition and difficulty in alignment of the probes acquired at various camera orientations, we were not able to obtain good recordings of the incident illumination with this capture process for inverse rendering. Hence, for these datasets, we did not compute the specular roughness parameter but manually selected a suitable value ($\sigma = 0.2$) for the pavement and sketch book, while borrowing from the DSLR dataset for the drain cover. As seen in the renderings under environmental illumination, the recovered diffuse and specular reflectance and the surface normals go a long way towards producing compelling rendering results for the acquired samples.

5.3 Discussion and error analysis

The main assumption for our method to work is that surfaces are planar, which is consistent with many recent SVBRDF measurement

methods (we refer to Section 4 of the supplemental material for details). We also assume measurements can be made around the Brewster angle of incidence. While in principle finding the exact Brewster angle is challenging outside of a laboratory setup, we found that being “close to” Brewster angle suffices to produce good results. Our proposed Brewster angle measurements for reflectometry are also consistent with the recent study of Nielsen et al. [2015] who found measurements near the Brewster angle to be nearly optimal for reflectance estimation (best 1-direction pair in their work) although they did not consider polarization.

5.3.1 Index of refraction estimation. To validate our refractive index estimation, we compared our estimates for the canvas print surface to known values of various types of inks [Peiponen et al. 2008], which are as follows: black 1.65 (ref. 1.6), blue & yellow 1.49 (ref. 1.53 – 1.54). We refer to the supplemental material for details.

5.3.2 Normal estimation. Figure 11, 2nd row, shows plots of the TRS for different angles of incidence θ_i , and angles of polarization ψ_i for a beam of light linearly polarized at 80%. We chose 80% as the degree of polarization as that is the maximum DOP predicted by Rayleigh sky model. When $\theta_i = \theta_B$, the maximum of the TRS correctly predicts the azimuth of the surface normal. Within a 10° window around the Brewster angle, it can be observed that the maximum of the TRS is always within a small window of the true azimuth with a small mean error in azimuth estimation to be around 3 to 4°. The latter is an easy condition to fulfill when making measurements: in all our measurements, we subjectively selected oblique views that we judged close to the Brewster angle by rotating the polarizer and testing the specular cancellation. Furthermore, the error is small in magnitude for even up to 15° away from θ_B . This implies that the method allows for a planar surface with up to 30° cone of surface normal variation.

5.3.3 Diffuse-specular separation. One added benefit of making measurements close to the Brewster angle is that the component of light polarized parallel to the plane of incidence is completely transmitted for a dielectric material. Mathematically, I_{\perp} is then always greater than I_{\parallel} and both can thus be equivalently referred to as I_{max} and I_{min} respectively. We refer to Section 3 of the supplemental for more details. This then provides us with a mechanism for diffuse-specular separation and estimation of diffuse albedo and specular reflectance R_0 from the observed minima and maxima respectively of the TRS near Brewster angle. We also note that while the Brewster angle can vary over large surfaces, we only found a 5–6° variations over our largest sample (canvas print/~0.75 m wide), which is still within the acceptable range as per Figure 11.

Figure 12 presents a comparison between reflectance maps of the red book cover acquired with our approach under uncontrolled outdoor illumination (Figure 7) and those acquired with controlled LCD panel illumination using the approach of [Ghosh et al. 2009]. Overall, the maps acquired with our proposed approach are qualitatively very similar to those obtained with controlled measurements. There is however some noticeable difference in the specular reflectance computed with both the methods with some color visible in the specular of the polarized LCD illumination based separation which is not seen in our result. This is in fact due to the red book surface

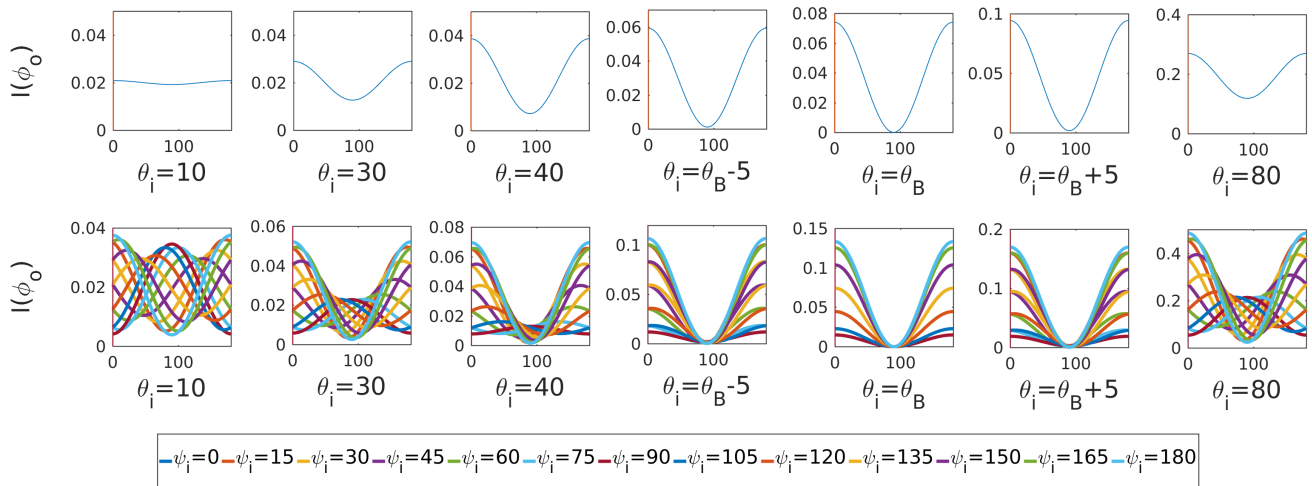


Fig. 11. Simulated TRS for a glass material (Index of Refraction $\eta = 1.5$) oriented at an azimuth $\phi_{\vec{n}} = 90^\circ$. First row: Simulation under unpolarized incident illumination - the minimum of the TRS is found at $\phi_o = 90^\circ$, as expected, for any angle of incidence θ_i . Second row: Simulation under partially linearly polarized illumination with a DOP of 80%. The different curves represent different angles of polarization ψ_i . Unlike under unpolarized incident illumination, the maximum of the TRS is not always located at $\psi_{pol} = 0^\circ$. Instead, the maximum is shifted depending on the incident angle of polarization ψ_i and angle of incidence θ_i . However, behaviour similar to that observed under unpolarized illumination can be observed again at Brewster angle.




Fig. 12. Comparison between the reflectance and normal maps of red book estimated with our proposed approach (top row) vs. those obtain with controlled measurements using an LCD panel [Ghosh et al. 2009] (center row). Error is visualized in false color (bottom row).

being coated with a slightly metallic paint which is not modelled with our approach which assumes pure dielectric reflectance. We also computed quantitative errors (RMSE) for the various maps as follows: diffuse $\rho_d = ([0.0456, 0.0227, 0.0187])$, specular $R_0(0.0428)$, specular roughness $\sigma(0.05)$ and surface normals (6.8° angular).

Figure 13 presents a qualitative comparison between the albedo and normal maps of the red book cover acquired with our approach and those obtained using controlled flash illumination measurements with the recent two-shot method of [Aittala et al. 2015]. Overall, there are both qualitative similarities as well as differences between the maps acquired with the two approaches. The maps estimated using the method of Aittala et al. are of higher resolution than our results due to close-up acquisition with flash illumination. However, we also note that the estimated spatial variation on the book cover using the method of Aittala et al. is not an exact reproduction of the true surface variation but rather its statistical reproduction

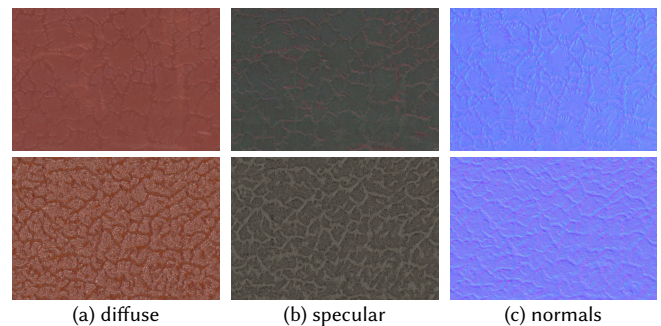


Fig. 13. Comparison between the albedo and normal maps of red book estimated with our proposed approach (top row) vs. those obtained with controlled flash illumination using the two-shot technique of [Aittala et al. 2015] (bottom row).

(using tile based synthesis). This can be seen most clearly in the normal map where the meandering veins on the surface estimated with the two shot method do not exactly match in location with those estimated with our proposed method.

Given the absence of pixel-level correspondence between our result and that of [Aittala et al. 2015], we conducted a statistical comparison by means of the mean and standard deviation values for the reflectance maps in Figure 13. The statistics (mean, std. dev.) are as follows: our method - diffuse ($[0.2925, 0.0718, 0.0555]$, $[0.0169, 0.0086, 0.0085]$), specular ($0.0647, 0.0108$), surface normal ($[0, 0, 1], 3.80^\circ$); [Aittala et al. 2015] - diffuse ($[0.2741, 0.0969, 0.068]$, $[0.1141, 0.073, 0.0629]$), specular ($0.0806, 0.0305$), surface normal ($[0, 0, 1], 5.13^\circ$). As can be seen, our result has a fair degree of consistency with that of Aittala et al. (albeit with a slightly lower standard deviation) for this near-stochastic sample.

We also conducted a statistical analysis of the surface normals estimated for the red book under various outdoor illumination conditions including cloudy sky, and sunny sky conditions at various

		Ours		[Ghosh et al. 2009]		[Aittala et al. 2015]	
Surf. normals	Time Of Day	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean
	Cloudy (mid-day)	3.80°	[0, 0, 1]	5.32°	[0, 0, 1]	5.13°	[0, 0, 1]
	Sunny (10-10:30am)	8.91°					
	Sunny (3-3:30pm)	5.97°					
	Sunny (6-6:30pm)	8.97°					

Table 1. Statistical variation in surface normals of red book under different lighting conditions (left column), compared to two measurement methods employing controlled illumination.




Fig. 14. Differences between the reflectance and normal maps estimated with the primary tripod mounted DSLR setup and hand-held mobile acquisition, visualized in false color.




Fig. 15. Differences between the reflectance and normal maps estimated with hand-held mobile acquisition vs. controlled measurements with LCD panel, visualized in false color.

times of day (Table 1). Here, we report the standard deviation of the surface normal distribution around the mean (up vector) $[0, 0, 1]$ compared against controlled measurements with the techniques of [Ghosh et al. 2009] and [Aittala et al. 2015] respectively. As can be seen, the surface normal distribution is closest to the reference distributions under a cloudy sky (unpolarized incident illumination), and for a near mid-day² measurement at 3pm on a sunny day (horizontal partial polarization). The statistical difference increases slightly for the sunny sky earlier in the morning (10am) when the sky polarization has a more arbitrary orientation, as well as later in the evening (6pm) when the horizon is more vertically polarized towards the direction orthogonal to the sun. However, the estimated surface normals are fairly consistent even in these non-ideal conditions.

Figure 14 presents a comparison between the maps of the drain cover acquired with our primary DSLR camera setup and our proposed hand-held acquisition with mobile phone. As can be seen, the maps acquired with mobile acquisition are numerically very close to those acquired with a tripod mounted DSLR highlighting the practicality of the approach for on-site acquisition. Visually, we do observe some differences in the color tone of the diffuse albedo estimated with mobile acquisition. We believe this may be due to the automatic white balancing employed by the mobile camera during acquisition as well as any differences in transmission characteristics of the glass linear polarizer on the DSLR camera vs. the laminated plastic polarizing sheet employed for the mobile camera. The surface

²The mid-day sun was at 1pm at this geographic location.




Fig. 16. Diffuse-specular separation with polarization measurements near normal incidence. The sky near sunset is highly polarized from above (a), which allows good quality of diffuse-specular separation on both the color chart and the red book.

normal map has expectedly higher error inside the grooves of the drain cover where self-occlusions are a contributing factor. Figure 15 presents a comparison of maps of the sketch book acquired with mobile acquisition compared to controlled measurements with LCD panel illumination (also see supplemental document). As seen, the acquired maps have low error for this planar dielectric surface.

While most of our reflectometry data is acquired at oblique angles of incidence, Figure 16 presents a special case of acquisition near sunset (after 6pm) where the sun is very low in the sky near the horizon. In this case, the sky zenith is strongly linearly polarized as can be seen in the mirror ball (a). In this case, we can obtain good diffuse-specular separation with polarization imaging when also viewing the sample from above because most of the incident light is strongly polarized near normal incidence. Note that some color can be seen in specular reflectance of the color chart. This is likely due to polarization preserving single scattering [Ghosh et al. 2008] under polarized incident illumination.

We also note that we do not explicitly consider diffuse polarization [Atkinson and Hancock 2006; Wolff and Boulton 1991] in this work because it has been shown to have a minor contribution ($\sim 9\%$) to polarized reflectance near the Brewster angle. Furthermore, given our day-time measurements of upwards facing planar samples, any resulting diffuse polarization will be mostly parallel-polarized (due to in-plane transmission) and hence contribute to the diffuse estimate, but not the specular reflectance estimates with our measurements.

5.4 Limitations

Our method currently models all types of materials as dielectrics. While this works well for many dielectric-metal composites and oxidized metallic surfaces commonly found outdoors, it certainly is not accurate for highly metallic surfaces exhibiting a complex index of refraction. Figure 17 presents one such example of an ornate book cover with sections of dielectric (reddish) grooves and metallic (golden) paint. Our method does not accurately estimate the diffuse-specular reflectance over the golden surface of the sample, treating it like a dielectric. This also affects the estimated specular roughness of the golden surface. And while the estimated surface normals appear to have rich qualitative surface details, they suffer from higher absolute error (RMSE 28.4° angular) due to combination of a metallic surface and grooves around the dielectric parts.

Our result is slightly sensitive to variation in incident illumination over the spatial extent of a sample due to employed calibration with

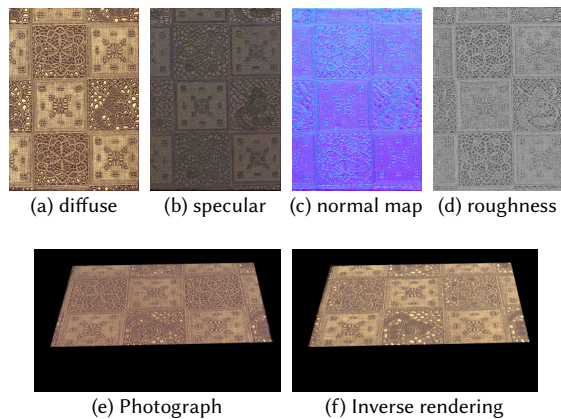


Fig. 17. Estimated reflectance and normal maps for an ornate book cover with metallic gold paint. Our method has problems with accurate estimation of surface reflectance over the metallic surface.

respect to the color chart. This can affect the result slightly under very strong directional lighting of direct sunlight on a clear day as can be seen in the maps of the stone pavement where there is slight variation in the estimated specular albedo and roughness over the sample surface. Such illumination variation can be a significant issue for highly specular materials.

Our method also assumes unoccluded incident illumination over most of the upper hemisphere and the results can suffer from bias when the hemisphere is partially occluded. Figure 18 presents one such partial success/failure result for a bas-relief. This relief is mounted on an interior wall of an enclosed entrance to an open courtyard. As seen in the corresponding light probe, the sample receives direct illumination from the sky only from the top-right direction (w.r.t. camera), while receiving indirect bounce illumination from other directions. The surface is also not perfectly planar and has self-occlusions (around the nose) and high curvature profiles. Even in this difficult case, our method achieves very reasonable results for the diffuse and specular albedo and surface normals. The recovered normals however suffer from some bias with reduced variation along the x-direction due to lack of direct illumination from the front and left sides. We are also not able to reliably estimate spatially varying specular roughness over the sample due to a narrow cone of direct illumination and instead estimate an average specular roughness value ($\sigma = 0.15$) over the entire sample for rendering. Despite these limitations, the recovered maps are of sufficient quality for producing compelling visualizations of the bas-relief in novel lighting (see accompanying video).

We currently rely on VisualSfM to provide us with multi-view correspondence for surface normal computation. While this works well for many surfaces in the real world, this can potentially fail for some texture less highly specular surfaces. Our mobile acquisition currently relies on our semi-automatic alignment procedure. Going forward, mobile AR/VR platforms such as Tango could be exploited for automatic tracking based alignment. Our inverse rendering step employs a first order approximation for modeling polarized incident illumination through the usage of view dependent light probes. While this has worked well enough in practice for our application,

a more accurate simulation of the incident polarization might be required for other applications. Finally, very sunny conditions can create issues with dynamic range for specular surfaces and may require extended HDR acquisition which can be a bit cumbersome.

6 CONCLUSION

We have presented a novel approach for passive reflectometry of planar surfaces in completely uncontrolled outdoor environments using a combination of (linear) polarization imaging, multi-view acquisition and inverse rendering. We demonstrate high quality estimation of spatially varying diffuse and specular albedo, index of refraction, surface normals and specular roughness for a wide variety of planar real world materials ranging from diffuse dominated brick and stone surfaces to very specular metal-dielectric composite surfaces. To the best of our knowledge, we are the first to apply polarization-based reflectometry in such complex and completely uncontrolled outdoor environments including busy urban settings. Unlike previous work on polarization based shape/reflectance analysis which has assumed unpolarized or circularly polarized illumination, we take into account the potential partial linear polarization of outdoor illumination and propose steps to mitigate the effect of such incident polarization in our reflectance acquisition and analysis. While understandably not quite as accurate as completely controlled measurements, our method achieves sufficient accuracy for realistic rendering applications and is particularly suited for surfaces that are permanent on-site structures that cannot be brought indoors. In future work, we would like to extend our analysis for truly metallic surfaces exhibiting a complex index of refraction and also investigate polarization effects on anisotropic surface reflectance.

ACKNOWLEDGEMENTS

We take this opportunity to thank Raoul Guiazon and Ching Pang Ho for useful discussions, Yiming Lin for help with some data acquisition, and Antoine Toisoul for assistance with rendering. We also thank our anonymous reviewers for suggestions on improving this manuscript. This work was supported by an EPSRC Early Career Fellowship (EP/N006259/1) and a Royal Society Wolfson Research Merit Award.

REFERENCES

- Miika Aittala, Timo Aila, and Jaakko Lehtinen. 2016. Reflectance Modeling by Neural Texture Synthesis. *ACM Trans. Graph.* 35, 4, Article 65 (July 2016), 13 pages.
- Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. 2013. Practical SVBRDF Capture In The Frequency Domain. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 32, 4 (2013).
- Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. 2015. Two-shot svbrdf capture for stationary materials. *ACM Transactions on Graphics* 34, 4 (2015), 110.
- Gary Atkinson and Edwin R Hancock. 2006. Recovery of surface orientation from diffuse polarization. *Image Processing, IEEE Transactions on* 15, 6 (2006), 1653–1664.
- Guojun Chen, Yue Dong, Pieter Peers, Jiawan Zhang, and Xin Tong. 2014. Reflectance Scanning: Estimating Shading Frame and BRDF with Generalized Linear Light Sources. *ACM Trans. Graph.* 33, 4, Article 117 (July 2014), 11 pages.
- Hua Chen and Lawrence B. Wolff. 1998. Polarization Phase-Based Method For Material Classification In Computer Vision. *IJCV* 28, 1 (1998), 73–83.
- E. Collett. 2005. *Field Guide to Polarization*, SPIE Field Guides vol. FG05. SPIE.
- R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM TOG* 1, 1 (1982), 7–24.
- Paul Debevec. 1998. Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography. In *Proceedings of ACM SIGGRAPH 98*.
- Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. 2000. Acquiring the Reflectance Field of a Human Face. In *Proceedings of ACM SIGGRAPH 2000*. 145–156.




Fig. 18. Partial result: Bas-relief on an interior wall of an enclosed entrance to a courtyard acquired with our method.

- Yue Dong, Guojun Chen, Pieter Peers, Jiawan Zhang, and Xin Tong. 2014. Appearance-from-motion: recovering spatially varying surface reflectance under unknown lighting. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 193.
- Yannick Francken, Tom Cuyppers, Tom Mertens, and Philippe Bekaert. 2009. Gloss and Normal Map Acquisition of Mesostructures Using Gray Codes. In *Proceedings of the 5th International Symposium on Advances in Visual Computing: Part II (ISVC '09)*. Springer-Verlag, Berlin, Heidelberg, 788–798.
- Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. 2003. Linear light source reflectometry. *ACM Trans. Graph. (Proc. SIGGRAPH)* 22, 3 (2003), 749–758.
- Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul Debevec. 2010. Circularly polarized spherical illumination reflectometry. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 29 (December 2010), 162:1–162:12. Issue 6.
- Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul E. Debevec. 2009. Estimating Specular Roughness and Anisotropy from Second Order Spherical Gradient Illumination. *Comput. Graph. Forum* 28, 4 (2009), 1161–1170.
- Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. 2011. Multiview face capture using polarized spherical gradient illumination. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, Article 129 (2011), 10 pages.
- Abhijeet Ghosh, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul Debevec. 2008. Practical Modeling and Acquisition of Layered Facial Reflectance. *ACM Transactions on Graphics* 27, 5 (Dec. 2008), 139:1–139:10.
- Mashhuda Glencross, Gregory J. Ward, Francho Melendez, Caroline Jay, Jun Liu, and Roger Hubbard. 2008. A Perceptually Validated Model for Surface Depth Hallucination. *ACM Trans. Graph.* 27, 3, Article 59 (Aug. 2008), 8 pages.
- D. Guarnera, G.C. Guarnera, A. Ghosh, C. Denk, and M. Glencross. 2016. BRDF Representation and Acquisition. *Computer Graphics Forum* 35, 2 (2016), 625–650.
- Giuseppe Claudio Guarnera, Pieter Peers, Paul E. Debevec, and Abhijeet Ghosh. 2012. Estimating Surface Normals from Spherical Stokes Reflectance Fields.. In *ECCV Workshop on Color and Photometry in Computer Vision*. 340–349.
- Daniel Hauagge, Scott Wehrwein, Paul Upchurch, Kavita Bala, and Noah Snavely. 2014. Reasoning about photo collections using models of outdoor illumination. In *Proceedings of the British Machine Vision Conference. BMVA Press*.
- Cong Phuoc Huynh, Antonio Robles-Kelly, and Edwin Hancock. 2010. Shape and refractive index recovery from single-view polarisation images. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 1229–1236.
- Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. 2015. Polarized 3D: High-Quality Depth Sensing with Polarization Cues. In *Proceedings of the IEEE International Conference on Computer Vision*. 3370–3378.
- Davis E. King. 2009. Dlib-ml: A Machine Learning Toolkit. *Journal of Machine Learning Research* 10 (2009), 1755–1758.
- G.P. Können. 1985. *Polarized Light in Nature*. Cambridge University Press.
- Kazutada Koshikawa. 1992. A polarimetric approach to shape understanding of glossy objects. (1992), 190–192.
- Stephen Lombardi and Ko Nishino. 2012. Reflectance and Natural Illumination from a Single Image. In *Proc. European Conference on Computer Vision (ECCV'12)*. 582–595.
- Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. 2007. Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination. In *Rendering Techniques*. 183–194.
- Rafal Mantiuk, Grzegorz Krawczyk, Radoslaw Mantiuk, and Hans-Peter Seidel. 2007. High dynamic range imaging pipeline: Perception-motivated representation of visual content. In *Electronic Imaging 2007*. International Society for Optics and Photonics, 649212–649212.
- Daisuke Miyazaki, Masataka Kagesawa, and Katsushi Ikeuchi. 2003. Polarization-based Transparent Surface Modeling from Two Views. In *ICCV*. 1381–1386.
- Daisuke Miyazaki, Takuya Shigetomi, Masashi Baba, Ryo Furukawa, Shinsaku Hiura, and Naoki Asada. 2012. Polarization-based surface normal estimation of black specular objects from multiple viewpoints. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2012 Second International Conference on*. IEEE, 104–111.
- Daisuke Miyazaki, Robby T. Tan, Kenji Hara, and Katsushi Ikeuchi. 2003. Polarization-based Inverse Rendering from a Single View. In *ICCV*. 982–987.
- Volker Müller. 1995. Polarization-Based Separation of Diffuse and Specular Surface Reflection. In *Mustererkennung 1995*. Springer, 202–209.
- Shree K. Nayar, Xi-Sheng Fang, and Terrance Boulton. 1997. Separation of Reflection Components Using Color and Polarization. *IJCV* 21, 3 (1997).
- Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. On Optimal, Minimal BRDF Sampling for Reflectance Acquisition. *ACM Trans. Graph.* 34, 6, Article 186 (Oct. 2015), 11 pages. <https://doi.org/10.1145/2816795.2818085>
- Geoffrey Oxholm and Ko Nishino. 2012. Shape and Reflectance from Natural Illumination. In *Proc. European Conference on Computer Vision (ECCV'12)*. 528–541.
- KE Peiponen, V Kohturi, I Niskanen, M Juuti, J Rätty, H Koivula, and M Toivakka. 2008. On estimation of complex refractive index and colour of dry black and cyan offset inks by a multi-function spectrophotometer. *Measurement Science and Technology* 19, 11 (2008), 115601.
- Stefan Rahmann and Nikos Canterakis. 2001. Reconstruction of Specular Surfaces using Polarization Imaging. *CVPR* 1 (2001), 149.
- Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo. 2011. Pocket reflectometry. *ACM Trans. Graph.* 30, 4 (July 2011), 45:1–45:10.
- J. Riviere, P. Peers, and A. Ghosh. 2016. Mobile Surface Reflectometry. *Computer Graphics Forum* 35, 1 (2016), 191–202.
- Fabiano Romeiro, Yuriy Vasilyev, and Todd Zickler. 2008. Passive reflectometry. In *Computer Vision—ECCV 2008*. Springer, 859–872.
- Fabiano Romeiro and Todd Zickler. 2010. Blind reflectometry. In *Proc. of ECCV*. 45–58.
- Firooz Sadjadiz and Farzad Sadjadi. 2007. Extraction of surface normal and index of refraction using a pair of passive infrared polarimetric sensors. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1–5.
- Megumi Saito, Yoichi Sato, Katsushi Ikeuchi, and Hiroshi Kashiwagi. 1999. Measurement of surface orientations of transparent objects by use of polarization in highlight. *J. Opt. Soc. Am. A* 16, 9 (1999), 2286–2293.
- C. Schlick. 1994. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246.
- W.A.P. Smith, R. Ramamoorthi, and S. Tozza. 2016. Linear Depth Estimation from an Uncalibrated, Monocular Polarisation Image. In *Proc. of European Conference on Computer Vision (ECCV)*. 517–526.
- John W Strutt. 1871. XV. On the light from the sky, its polarization and colour. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 41, 271 (1871), 107–120.
- Vimal Thilak, David G. Voelz, and Charles D. Creusere. 2007. Polarization-based index of refraction and reflection angle estimation for remote sensing applications. *Appl. Opt.* 46, 30 (2007), 7527–7536.
- Shinji Umeyama and Guy Godin. 2004. Separation of Diffuse and Specular Components of Surface Reflection by Use of Polarization and Statistical Analysis of Images. *PAMI* 26, 5 (2004).
- Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*. Eurographics Association, 195–206.
- Chun-Po Wang, Noah Snavely, and Steve Marschner. 2011. Estimating Dual-scale Properties of Glossy Surfaces from Step-edge Lighting. *ACM Trans. Graph.* 30, 6, Article 172 (Dec. 2011), 12 pages.
- Tim Weyrich, Jason Lawrence, Hendrik P. A. Lensch, Szymon Rusinkiewicz, and Todd Zickler. 2009. Principles of Appearance Acquisition and Representation. *Found. Trends. Comput. Graph. Vis.* 4, 2 (Feb. 2009), 75–191.
- L. B. Wolff. 1989. Surface orientation from two camera stereo with polarizers. In *Proc. SPIE Conf. Optics, Illumination and Image Sensing for Machine Vision IV*, Vol. 1194. 287–297.
- Lawrence B. Wolff and Terrance E. Boulton. 1991. Constraining Object Features Using a Polarization Reflectance Model. *PAMI* 13, 7 (1991), 635–657.
- Changchang Wu. 2011. VisualSFM: A visual structure from motion system. (2011).
- Hongzhi Wu and Kun Zhou. 2015. AppFusion: Interactive Appearance Acquisition Using a Kinect Sensor. *Computer Graphics Forum* 34, 6 (2015), 289–298.
- Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. 2016. Recovering Shape and Spatially-varying Surface Reflectance Under Unknown Illumination. *ACM Trans. Graph.* 35, 6, Article 187 (Nov. 2016), 12 pages.