

Polarization imaging reflectometry in the wild*

Jérémy Riviere, Ilya Reshetouski, Abhijeet Ghosh
Imperial College London




Figure 1: Polarization imaging reflectometry in the wild. Reflectance maps (a,c) obtained using commodity photography equipment (fig. 3), by exploiting multiple observations of the sample under sky illumination. Specifically, we recover diffuse ((a,c) top left) and specular ((a,c) top right) albedos, surface normals ((a,c) bottom left) and specular roughness ((a,c) bottom right). These can then be used to render the objects captured in new lighting conditions (b,d)

Abstract

We present a novel approach for on-site acquisition of surface reflectance for planar, spatially varying, isotropic materials in uncontrolled outdoor environments. Our method exploits the naturally occurring linear polarization of incident illumination: by rotating a linear polarizing filter in front of a camera at 3 different orientations, we measure the linear polarization reflected off the sample and combine this information with multiview analysis and inverse rendering in order to recover per-pixel, high resolution reflectance maps. We exploit polarization both for diffuse/specular separation and surface normals estimation by combining polarization measurements from at least two near orthogonal views close to Brewster angle of incidence. We then use our estimates of surface normals and albedos in an inverse rendering framework to recover specular roughness. 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 environments.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Polarization I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Reflectometry;

Keywords: Surface reflectometry, partial linear polarization, Stokes parameters, index of refraction

1 Introduction

Accurately reproducing the appearance of physical objects plays an important role in domains such as visual effects for movies or video games, virtual and augmented reality, cultural heritage preservation, etc. as it helps create a more believable and immersive virtual environment for the end user. While many techniques for image based reflectometry have been developed over the past 20 years, they gen-

erally require dark room conditions which can only be achieved inside a laboratory, thus limiting the scope of measurements.

In this paper, we propose a novel image-based acquisition system for the full measurement of reflectance maps for spatially varying planar surfaces commonly found outdoors, exhibiting isotropic BRDFs. We limit the amount of data needed for reflectometry by fusing observations of the polarization state of light reflected off the surface from up to 3 views. We motivate our approach of using polarization cues by 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 Brewster angle, as predicted by Fresnel equations.

Our approach only requires standard photography equipment often used on-site for image based lighting measurements [Debevec 1998]: a DSLR camera, a linear polarizing filter, a calibration target and a mirror ball (fig. 3).

The principal contributions of this work are:

- Practical method for passive acquisition of surface reflectance with a few measurements under hemispherical uncontrolled outdoor illumination, using commodity photography equipment.
- Novel combination of polarization imaging with multiview analysis and inverse rendering for estimating detailed reflectance maps of planar surfaces.
- Novel analysis of the intensity profile observed through a linear polarizer applied to reflectometry under partially linearly polarized incident illumination.

*Please cite this technical report as "Jérémy Riviere, Ilya Reshetouski and Abhijeet Ghosh. Polarization imaging reflectometry in the wild. Technical Report 2016/8, Department of Computing, Imperial College London, ISSN 1469-4174, May 2016."

2 Related work

Reflectance measurement generally requires very complex set-ups in laboratory-like environments, thus limiting their accessibility to non-expert users. However, in the past few years, the advances in Computer Vision and Pervasive Computing have given rise to a new body of research in reflectance acquisition that seeks to develop off-the-shelf, easy to use solutions for creative people based on commodity hardware, that can work outside of a laboratory environment.

2.1 Commodity hardware

Most prior work on reflectometry based on commodity hardware is designed around a camera-light source pair, where a typical light source can be a point or linear light source. Linear light sources are generally preferred as they help reduce the amount of data required in measurements. [Gardner et al. 2003] estimated spatially varying BRDF of planar surfaces by translating a linear light source with a gantry over the sample, while keeping the camera static. This design was later modified by [Ren et al. 2011] for portability and [Chen et al. 2014] to allow the 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 use an LCD panel as an extended source of illumination. [Ghosh et al. 2009] proposed to project spherical gradient illumination patterns from an LCD screen to estimate per-pixel, spatially varying reflectance maps of planar samples. [Aittala et al. 2013] later proposed a setup to capture high resolution surface reflectance maps by measuring the sample's response to band-limited illumination patterns in the frequency domain.

More recently, the advances in mobile technology have given rise to more compact and portable designs for reflectance measurements. [Aittala et al. 2015] proposed a design for reflectance measurement of specular stationary materials based on a mobile device. They use a pair of flash-no flash observations of the sample and statistical analysis tools to extract the reflectance maps in general office room lighting. [Riviere et al. 2015] have also recently proposed two setups for the acquisition of spatially varying planar surfaces in general office room lighting, based on a mobile device. Their first setup is targetted at rough specular samples and relies on a hand-held dense sampling of the BRDF in the back-scattering direction. Their second setup is used to resolve highly specular material by projecting second order spherical gradient illumination patterns from the device's screen, similar to [Ghosh et al. 2009].

None of these approaches can be exploited in general lighting conditions, and especially not outdoors, as they rely on active illumination that would likely be dominated by the ambient lighting.

2.1.1 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 an homogeneous curved object under known but uncontrolled lighting, such that each pixel of the image provide a linear constraint on the BRDF. Their goal is then to find the BRDF function that best fits those constraints. Two years later, the same authors proposed a method for reflectometry under unknown illumination for homogeneous curved objects that leverages the statistics of real-world illumination. [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 in

overcast conditions and flash lighting to recover shape and diffuse shading. [Hauagge et al. 2014] also assume a lambertian image formation model and a model of outdoor illumination to recover per-pixel albedo. Closer in spirit to ours is the work of [Dong et al. 2014] on Appearance from Motion. In their work, Dong et al. recover spatially varying isotropic reflectance from a video of a rotating object under unknown lighting, by alternatively estimating for reflectance and lighting in an iterative process.

In our work paper, we further push the limits of reflectometry to general outdoors lighting conditions by designing a portable solution that does not require active illumination. We do so by taking advantage of the inherent polarization of natural lighting and polarization from reflection to provide cues for albedo and normals estimation.

2.2 Polarization

Polarization has been extensively studied in both Computer Vision and Computer 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 on polarization has been interested in studying the polarization resulting from reflection under unpolarized incident light. Two notable exceptions are [Koshikawa 1992] who use circularly polarized illumination to recover surface orientation and [Ghosh et al. 2010] who also use circularly polarised incident lighting for reflectometry.

2.2.1 Surface normals estimation

Shape from Polarization has been extensively studied in Computer Vision, under the restriction that the incident illumination comes from an unpolarized, uniform source. In such conditions, the angle of polarization determines the direction perpendicular to the plane of incidence, which is the plane that goes through the incident direction, normal to the surface and view direction. Two strategies are then typically used to infer orientation from polarization.

The first approach relies on the degree of polarization and inverting the Fresnel equations. Most prior work has focused on shape from specular reflection, and solving the ambiguity in azimuth angle due to the degree of polarization reaching an extremum at Brewster angle [Thilak et al. 2007; Saito et al. 1999; Guarnera et al. 2012]. Unlike those previous methods, [Atkinson et al. 2006] measure the degree of polarization due to diffuse reflection for shape estimation.

A second approach consists in observing the sample through different polarized cameras [Wolff 1989; Rahmann and Canterakis 2001; Miyazaki et al. 2003a; Sadjadiz and Sadjadi 2007]. 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 2 [Wolff 1989]) is 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 that careful camera calibration has to be obtained in order to get per-pixel correspondences for each view. In this paper, we follow a multiview polarization approach and show its successful application under partially polarized outdoor illumination.

Recently, [Kadambi et al. 2015] proposed a method to enhance coarse depth maps by fusing Shape from Polarization cues with the output of a depth sensor (Microsoft Kinect or 3D scanner). They follow the unpolarized world assumption to get estimate the zenith angle and use the coarse 3D geometry to resolve the azimuthal ambiguity in polarization normals.

2.2.2 Reflectance separation

Appearance modeling methods strongly rely on the accurate separation of surface reflectance into its diffuse and specular component. Most existing methods therefore either rely on color-space heuristics [Lee et al. 2006; Mallick et al. 2006], polarization [Wolff and Boult 1991; Müller 1995; Ma et al. 2007; Ghosh et al. 2010] or a combination of both [Nayar et al. 1997b; Umeyama and Godin 2004] for diffuse-specular separation. For the purpose of this paper, we will limit our discussion to polarization-based methods.

These methods all exploit the fact that diffuse reflection tends to depolarize light due to multiple subsurface scattering events happening within the material, while specular reflection preserves the polarization state of light. [Ma et al. 2007] exploit polarisation for the separation of diffuse and specular reflection, which they combine with spherical gradient illumination to obtain high resolution diffuse and specular normal maps. [Ghosh et al. 2010] later proposed a method where they measure the full Stokes parameters of reflected circularly polarized illumination to recover detailed reflectance parameters including index of refraction and specular roughness. In this work, we aim to extend such reflectometry using polarization imaging to completely uncontrolled outdoor environments.

3 Overview

In this paper, we propose a novel method for passive reflectometry applied under uncontrolled outdoors lighting conditions. We combine multiple observations of the polarization state of light reflected off an object at different vantage points and inverse rendering to estimate high resolution reflectance maps of the sample under consideration. We start our discussion by giving an overview of polarization and Mueller calculus that defines the notations used in the remainder of the paper (section 4). We then develop our method in section 5 by first deriving the theory needed for polarization measurements in the wild (section 5.1), from which we derive our measurement protocol (section 5.3). Finally, we describe our data analysis pipeline in section 5.4, before discussing results and limitations in section 6.

We propose a three-step pipeline for the extraction of reflectance maps from the measured data. We start by separating the reflection into its diffuse and specular components from which we derive a per-pixel diffuse albedo (ρ_d) and index of refraction (η) (section 5.5). We then estimate a per surface point normal map in a multi-view Shape from Polarization framework, by combining polarization measurements from two near orthogonal views close to Brewster angle of incidence (section 5.6). Finally, we combine our estimates of albedo, index of refraction and normals in an inverse rendering framework to estimate specular roughness throughout the surface (section 5.7).

4 Background - Polarization

4.1 Mueller calculus

Before we get to describing our method, we will give a brief overview of the necessary background in polarization and Mueller calculus. The polarization state of light can be described in vector form by the 4-element Stokes 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 polarization state of light can be expressed as a linear transformation, according to Mueller calculus (assume for now that there is no dif-

fuse component):

$$s' = C(\phi)M_{ref}(\eta; \theta_i; \delta; \vec{n})C(-\phi)s \quad (1)$$

where $L_i(\vec{\omega}_i)$ is the incoming radiance along a ray incident at $\vec{\omega}_i = (\theta_i, \phi_i)$. $C(\phi)M_{ref}(\eta; \theta_i; \delta; \vec{n})C(-\phi)$ is the aggregate transformation resulting from a reflecting element rotated at an angle ϕ from a canonical frame of reference [Collett 2005]. The Mueller matrix for rotation, $C(\phi)$ is defined as:

$$C(\phi) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\phi & -\sin 2\phi & 0 \\ 0 & \sin 2\phi & \cos 2\phi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2)$$

where ϕ provides a cue to constrain the azimuth of the surface normal to the plane of incidence, defined as the plane containing the exitant direction $\vec{\omega}_o$, the incident direction $\vec{\omega}_i$ and the surface normal.

$M(\eta; \theta_i; \delta; \vec{n})$ is the concatenation of the Mueller matrices of a linear diattenuator, also referred to as Mueller reflection matrix in previous work, and a linear retarder of phase δ :

$$M_{ref} = \begin{bmatrix} \frac{R_{\parallel} + R_{\perp}}{2} & \frac{R_{\perp} - R_{\parallel}}{2} & 0 & 0 \\ \frac{R_{\perp} - R_{\parallel}}{2} & \frac{R_{\parallel} + R_{\perp}}{2} & 0 & 0 \\ 0 & 0 & \sqrt{\frac{R_{\parallel}R_{\perp}}{2}} \cos \delta & \sqrt{\frac{R_{\parallel}R_{\perp}}{2}} \sin \delta \\ 0 & 0 & -\sqrt{\frac{R_{\parallel}R_{\perp}}{2}} \sin \delta & \sqrt{\frac{R_{\parallel}R_{\perp}}{2}} \cos \delta \end{bmatrix} \quad (3)$$

Light being a transverse electromagnetic wave, its electric field can be decomposed in an orthonormal basis spanned by two components respectively parallel (p-polarized) and perpendicular (s-polarized) to the plane of incidence. In eq. (3), R_{\perp} and R_{\parallel} represent the relative amounts of reflected s-polarized (resp. p-polarized) light as predicted by Fresnel equations. δ is the relative phase between the s- and p-polarized components of light. In particular, for dielectric materials, $\delta = 180^\circ$ for any angle of incidence before Brewster angle (θ_B) and 0° after. While that is strictly true only for dielectric materials, we have found in our experiments that some metallic surfaces found outdoors (see fig. 6, specular drain cover) exhibit dielectric-like behaviour due to oxidation.

Finally, the Mueller matrix for a linear polarizer rotated at an angle ψ_{pol} with respect to the observation's coordinate system is defined as:

$$M_{pol} = \frac{1}{2} \begin{bmatrix} 1 & \cos 2\psi_{pol} & \sin 2\psi_{pol} & 0 \\ \cos 2\psi_{pol} & \cos^2 2\psi_{pol} & \cos 2\psi_{pol} \sin 2\psi_{pol} & 0 \\ \sin 2\psi_{pol} & \cos 2\psi_{pol} \sin 2\psi_{pol} & \sin^2 2\psi_{pol} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (4)$$

To image the polarization state of light, a common method is to rotate a linear polarizer in front of a camera at 3 or more different orientations of ψ_{pol} . In particular, orienting the polarizer at 0° , 45° and 90° (referred to as P_0 , P_{45} and P_{90} in the remainder of the paper), allows to compute the 3 first elements of the Stokes vector of reflected light:

$$\begin{aligned} s'_0 &= P_0 + P_{90} \\ s'_1 &= P_0 - P_{90} \\ s'_2 &= 2 * P_{45} - s'_0 \end{aligned} \quad (5)$$

Equation (1) holds for pure specular reflection only, which does not generally correspond to how real world materials reflect light.

Indeed, most real world materials are rough to some extent. It is however possible to account for rough specular reflections by modeling the surface with a microfaceted BRDF where each microfacet behaves as per eq. (1). The resulting Stokes vector can then be computed as:

$$s'_{rough}(\vec{\omega}_o) = \int_{\Omega} \left(\frac{\rho_d}{\pi} + f_r(\vec{n}; \sigma; \eta; \vec{\omega}_i, \vec{\omega}_o) s'(\eta; \vec{\omega}_i) \right) L_i(\vec{\omega}_i) (\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i \text{ where} \quad (6)$$

Equation (6) is our complete image formation model with both diffuse and rough specular reflection, where $f_r(\vec{n}; \sigma; \eta; \vec{\omega}_i, \vec{\omega}_o)$ we model as a Cook-Torrance microfacet BRDF [Cook and Torrance 1982] with a GGX distribution [Walter et al. 2007]. Our goal is to recover all 4 of diffuse albedo (ρ_d), real index of refraction (η), surface normal (\vec{n}) and specular roughness (σ) from observations of $s'_{rough}(\vec{\omega}_o)$ under natural outdoors illumination.

5 Polarization reflectometry in the wild

Polarization has been extensively employed either for shape estimation [Miyazaki et al. 2012; Guarnera et al. 2012; Kadambi et al. 2015] or reflectance estimation [Ghosh et al. 2010] but rarely for both, with the exception of [Miyazaki et al. 2003b] who recover diffuse and specular albedos, surface normals and illumination from a single view. These approaches, however, rely heavily on the incident illumination to be either unpolarized ([Miyazaki et al. 2012; Guarnera et al. 2012; Kadambi et al. 2015]) or circularly polarized ([Ghosh et al. 2010; Guarnera et al. 2012]). While the latter case cannot be found naturally, outdoors illumination is known to be unpolarized in overcast conditions, and partially linearly polarized in open sky conditions. The latter case being more general, we will derive our theory for incident light that is partially linearly polarized.

5.1 Partially linearly polarized light

In this section, we will only be interested in the linear components of the reflected Stokes vector. Therefore, Mueller matrices reduce to 3x3 matrices and Stokes vectors to a 3-element vector. It can be shown that 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 the degree of linear polarization (DOLP) of the light. For light incident from the sky, we can therefore express its Stokes vector as:

$$s_{plp} = [1 \quad \mathcal{P}_i \cos 2\psi_i \quad \mathcal{P}_i \sin 2\psi_i]^T \quad (7)$$

where \mathcal{P}_i is the degree of linear polarization (DOLP) and ψ_i the angle of polarization of the incident ray. Note that if $\mathcal{P}_i = 0$, eq. (7) expresses the Stokes vector for unpolarized light.

We obtain the reflected Stokes vector by applying eq. (1) to eq. (7):

$$s'_{plp} = \begin{bmatrix} \frac{R_{\perp} + R_{\parallel}}{2} + \mathcal{P}_i \frac{R_{\perp} - R_{\parallel}}{2} \cos(2(\phi - \psi_i)) \\ \frac{R_{\perp} - R_{\parallel}}{2} \cos 2\phi + \mathcal{P}_i * A \\ \frac{R_{\perp} - R_{\parallel}}{2} \sin 2\phi + \mathcal{P}_i * B \end{bmatrix}$$

$$A = \left(\frac{R_{\perp} + R_{\parallel}}{2} \cos^2 2\phi + \sqrt{R_{\perp} R_{\parallel}} \sin^2 2\phi \cos \delta \right) \cos 2\psi_i$$

$$+ \left(\frac{R_{\perp} + R_{\parallel}}{2} - \sqrt{R_{\perp} R_{\parallel}} \cos \delta \right) \cos 2\phi \sin 2\phi \sin 2\psi_i$$

$$B = \left(\frac{R_{\perp} + R_{\parallel}}{2} - \sqrt{R_{\perp} R_{\parallel}} \cos \delta \right) \cos 2\phi \sin 2\phi \cos 2\psi_i$$

$$+ \left(\frac{R_{\perp} + R_{\parallel}}{2} \sin^2 2\phi + \sqrt{R_{\perp} R_{\parallel}} \cos^2 2\phi \cos \delta \right) \sin 2\psi_i \quad (8)$$

It is interesting to note the differences that exist with the expressions found in the literature for incident unpolarized illumination [Guarnera et al. 2012] and circularly polarized incident illumination [Ghosh et al. 2010].

1. The reflected intensity component $s'_{0,plp}$ depends on the polarization state of the incident illumination. It is therefore important that we recover not only the intensity $L_i(\vec{\omega}_i)$ of the incident light but also its polarization properties for inverse rendering (see section 5.7).
2. Unlike in previous work [Wolff 1989; Guarnera et al. 2012], it is not possible to estimate ϕ directly from the linear components $s_{1,plp}$ and $s_{2,plp}$ as they depend on the incident polarization as well. We must therefore find another way of estimating ϕ to obtain cues for surface normal estimation.

In the following section, we will describe our approach to recover the angle ϕ .

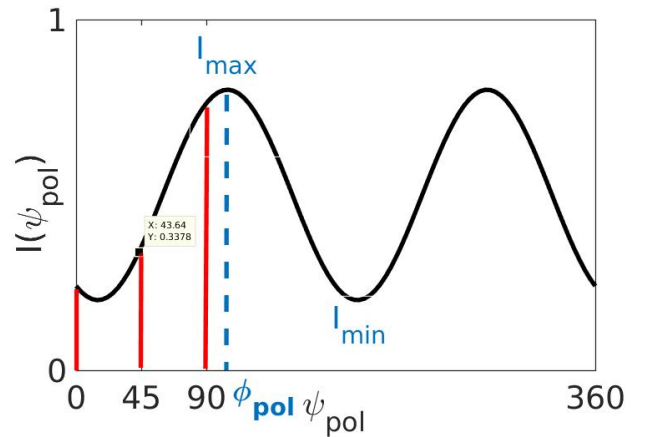


Figure 2: The intensity profile through a linear polarizer has the form of phase-shifted sinusoid of phase ϕ_{pol} , which can be measured with only three measurements (red lines) at 0, 45, 90°




Figure 3: Our measurement apparatus, composed only of commodity photography equipment often used for image-based measurements.

5.2 Transmitted Radiance Sinusoid

A common method for Shape from Polarization is to consider the intensity profile of reflected light passing through a linear polarizer, which has the form of a phase-shifted sinusoid of phase ϕ_{pol} , with minimum and maximum amplitudes denoted as I_{min} and I_{max} respectively (fig. 2):

$$I(\psi_{pol}) = \frac{I_d}{2} + \frac{I_{max} + I_{min}}{2} + \frac{I_{max} - I_{min}}{2} \cos(2(\psi_{pol} - \phi_{pol}))$$

where I_d accounts for diffusely reflected light

(9)

In previous work [Kadambi et al. 2015; Huynh et al. 2010; Miyazaki et al. 2012], the phase ϕ_{pol} of the Transmitted Radiance Sinusoid has been shown to be directly related to the azimuth of the surface normal ϕ as $\phi = \phi_{pol} + \frac{\pi}{2}$. We found however that this holds true only when the incident illumination has no linear polarization component. Under partially linearly polarized light, ϕ_{pol} is related to the azimuth of the surface normal ϕ only when observing the sample at Brewster angle (fig. 5). We will thus develop our method around oblique views close to Brewster angle of incidence section 5.3. We will later show in section 6 that finding the exact Brewster angle is not important and that it is sufficient to be within a 10° window around Brewster angle, which is easily judged visually, to obtain good qualitative and quantitative results.

5.3 Acquisition setup

In light of the analysis performed in the previous section, we propose a simple method for reflectometry in general outdoor conditions. Our measurement apparatus is composed only of standard equipment often used on-site for image based measurements (fig. 3):

- A Canon EOS 650D, 18 Megapixel DSLR camera, to which we have attached a 18-55 mm zoom lens. The camera sits on a heavy duty tripod so as to make the camera stable during measurements.
- A linear polarizer onto which we have marked the 0, 45 and 90 degrees orientations to make measurements of reflected




Figure 4: Our method typically requires two sets of linear Stokes measurement close to Brewster angle of incidence taken at orthogonal directions and one near normal incidence for registration.

Stokes parameters.

- A stainless steel mirror ball placed next to our sample during the capture process to record the incoming light.
- An X-rite color checker chart also placed flat next to the sample for white balancing and radiometric calibration.

Our measurement process typically goes as follows (fig. 4): Once the scene is set up, we start by imaging the reflected intensity close to normal incidence to have a canonical view of reference. While it would be possible to simply image that view without the polarizer mounted, we found it tedious to have to unmount and remount the polarizer between views. We then proceed to the measurement of Stokes parameters for two roughly orthogonal oblique views of the sample close to Brewster angle of incidence, to maximize the strength of the polarization signal and to be in the conditions where eq. (9) holds. While in theory eq. (9) is valid only at Brewster angle, we found that in practice it suffices to be close to Brewster angle (see section 6.2 for an in-depth analysis).

Each set of photographs close to Brewster angle consists of 9 images: 3 different exposure level images for each of the 3 different orientations of the polarizing filter. These shots are then combine to produce linear sRGB HDR images for all three polarization filter orientations using pfstools[Mantiuk et al. 2007] which we subsequently process to estimate the linear Stokes parameters of the captured reflected light by applying eq. (5) per-pixel. We complete our radiometric calibration by scaling the observed intensity of the white square of the color chart to match its sRGB value ([.9.9.9]) for each view.

We further require camera pose estimation to recover surface normals from multi-view Shape from Polarization section 5.6 and for the estimation of specular roughness through inverse rendering section 5.7. For this purpose, we use VisualSfM [Wu 2011], a gui-based software for Structure from Motion.

5.4 Reflectance extraction

In the following section, we will describe our method for extracting high resolution reflectance maps from data captured as per the protocol described in section 5.3. Recall that we took two sets of polarization measurements close to Brewster angle, by sampling the intensity through a linear polarizer at $\psi_{pol} = 0, 45$ and 90 degrees. From those observations, it is possible to fit a sinusoid as defined by eq. (9) to obtain the three unknowns I_{max} , I_{min} and ϕ . This could be done using any non linear least-squares fitting algorithm. However, this can be highly computationally expensive as we need

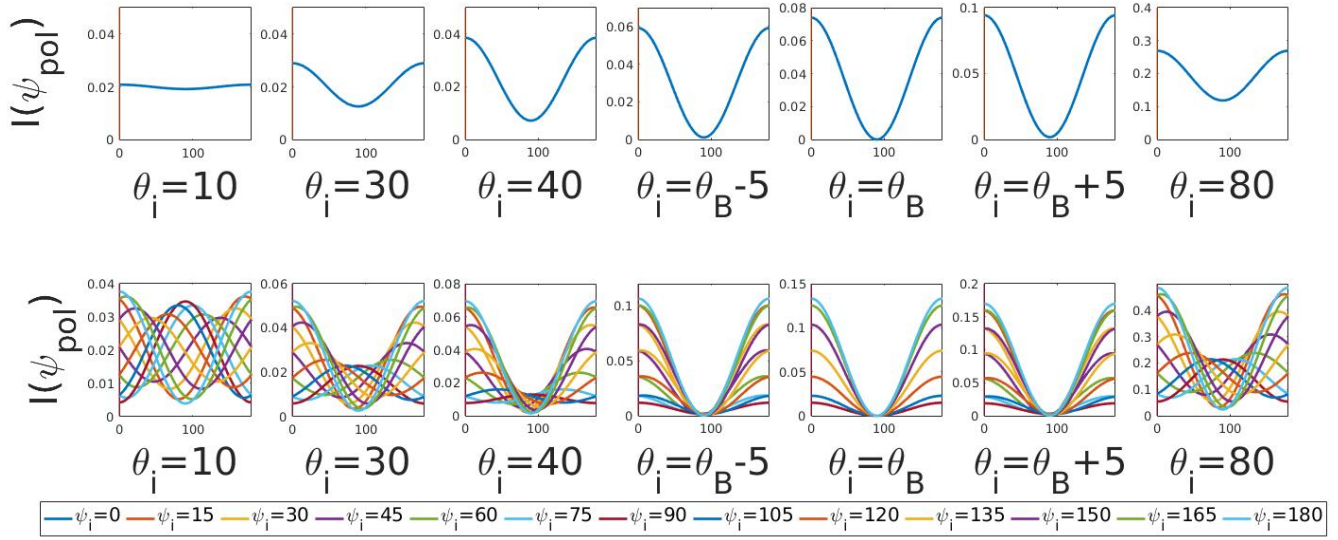


Figure 5: Simulated TRS for a glass material ($\eta = 1.5$) oriented at an azimuth $\phi = \frac{\pi}{2}$. First row: Simulation under unpolarized incident illumination - the maximum of the TRS is found at $\psi_{pol} = 0^\circ$ as expected for any angle of incidence θ_i . Second row: Simulation under 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 . The behaviour observed under unpolarized illumination can be observed again at Brewster angle.

to fit a sinusoid for each of our 18 millions of pixel. Instead, we follow an approach akin to that of [Nayar et al. 1997a]. By rewriting eq. (9) as:

$$I(\psi_{pol}) = \begin{bmatrix} 1 & \cos 2\psi_{pol} & \sin 2\psi_{pol} \end{bmatrix} \begin{bmatrix} \frac{I_{max} + I_{min}}{2} \\ \frac{I_{max} - I_{min}}{2} \cos 2\phi_{pol} \\ \frac{I_{max} - I_{min}}{2} \sin 2\phi_{pol} \end{bmatrix} \quad (10)$$

the problem becomes a linear problem of the form $Ax = b$ which can be solved for very efficiently using Single Value Decomposition. 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:

$$\begin{aligned} I_{max} &= x_1 + \sqrt{x_2^2 + x_3^2} \\ I_{min} &= x_1 - \sqrt{x_2^2 + x_3^2} \\ \phi_{pol} &= \frac{1}{2} \arctan \frac{x_3}{x_2} \end{aligned} \quad (11)$$

5.5 Diffuse specular separation

Diffuse albedo (fig. 6,a)

From eq. (9), it is clear that the diffuse component of reflection is

observed at the minimum of the TRS, when $\psi_{pol} = \phi + \frac{\pi}{2}$:

$$\begin{aligned} I_{min} &= \frac{I_d}{2} \\ &= \frac{\rho_d}{2\pi} \underbrace{\int_{\Omega} (\vec{n} \cdot \vec{\omega}_i) L_i(\vec{\omega}_i) d\vec{\omega}_i}_{\pi} \\ &= \frac{\rho_d}{2} \end{aligned} \quad (12)$$

where the integral can be simplified because of the radiometric calibration step of section 5.3. In most of our experiments, we have found that any of the two oblique views could be used to extract the diffuse albedo, except for the "specular drain cover" dataset (see fig. 6). In that case, we picked the view with the most diffuse-like appearance overall.

Reflectance at normal incidence (fig. 6,b)

By subtracting the diffuse albedo to I_{max} 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 (eq. (9)). 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 use the specular response of the color chart as our template. The key idea is that because our samples are mainly planar, they are subject to the same incident illumination as the color chart.

Moreover, because the color chart is made of plastic, we assume its index of refraction η_{chart} to be uniform and equal to 1.46 and compute its perpendicular reflection coefficient at Brewster angle

under uniform spherical illumination, $R_{\perp,real}$. We then compute the scale factor between its measured diffuse subtracted maximum intensity and $R_{\perp,real}$. The same scaling factor is then applied to the sample’s diffuse subtracted maximum intensity to obtain an estimate of R_{\perp} . We finally apply a method akin to that of [Ghosh et al. 2010] to estimate a per-pixel index of refraction:

$$\eta^2 = \frac{(1 + \sqrt{R_{\perp}(\theta_B)})}{(1 - \sqrt{R_{\perp}(\theta_B)})} \quad (13)$$

From 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 form (η or R_0) can then be used with their corresponding equations to model Fresnel effects, depending on the accuracy vs computation time trade-off.

5.6 Surface normals estimation (fig. 6,c)

We formulate our normal estimation as a multi-view Shape from Polarization framework, akin to the work of [Miyazaki et al. 2012]. Let us consider the vector $\vec{b} = [\cos(\phi_{pol}) \sin(\phi_{pol}) 0]^T$. This vector is orthogonal to the surface normal \vec{n} which means that the dot product (\vec{b}, \vec{n}) is 0. Essentially, the knowledge of \vec{b} constrains the surface normal to lie in the plane of incidence. Therefore, by combining 2 or more observations at 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:

$$\begin{bmatrix} R_1^T \vec{b}_1 \\ R_2^T \vec{b}_2 \\ 0 \end{bmatrix} \begin{bmatrix} n_x \\ n_y \\ n_z \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad (14)$$

While the above problem may easily be solved in closed-form for two views by writing the equations of the two planes of reflection for each view and intersecting them, we decided to follow the same approach as [Miyazaki et al. 2012] using SVD decomposition as it is more easily extendable to a scenario where more views are available.

5.7 Specular roughness (fig. 6,d)

We formulate our specular roughness estimation as a least squares problem. Given a set of observations of a sample under natural illumination $s'_{0,i}, i \in [0, 2]$, we compute $s_{\hat{0},i'}$ using eq. (6), where $\rho\mathbf{a}, \vec{n}, \eta$ we estimated from the previous sections, leaving only σ as an unknown. We thus solve for the σ that minimizes the residual between our rendered images $s_{\hat{0},i'}$ and photos $s'_{0,i}$ in the least squares sense:

$$\min_{\sigma} \sum_i \frac{1}{2} \|s_{0,i} - s_{\hat{0},i}(\sigma)\|_2 \quad (15)$$

This can be efficiently solved using a Levenberg-Marquardt (LM) non-linear optimizer. In our work, we chose to use the implementation of LM implemented in dlib [King 2009].

Furthermore, in our inverse rendering implementation, we did not account for the polarization of the incident illumination as is predicted by eq. (8) because this requires that we invert the Mueller matrix of the mirror ball in order to recover the incident polarization which is very challenging to do for metals. Instead, we simply assumed the incident illumination to be unpolarized (as is often the case in Computer Graphics) and modeled the Fresnel term of the BRDF with a simple Schlick’s approximation which we found to give satisfying results (see section 6). It would be interesting however to study how factoring the incident polarization in the inverse rendering would change the end results, as future work.

6 Results and discussion

6.1 Results

To validate our approach, we measured different objects found outdoors under various illumination conditions. Figure 6 shows the reflectance maps estimated from the method presented in section 5 for four exemplar datasets, namely the ”specular drain cover”, ”red bricks”, ”garden pavement” and ”red book”. We chose those datasets as they show a wide variety of reflectance properties and were captured in drastically different lighting conditions.

It is clear from fig. 6 that our method is agnostic to the incident illumination and is able to produce good qualitative results in very challenging conditions.

Specular drain cover (figs. 6 and 7, first row). Despite our assumption that all our samples are dielectrics, which is not strictly accurate for the specular drain cover, we obtain good results with our method. We attribute those results to two main factors:

1. The drain cover is actually made of a composite material which can be explained by a high, real index of refraction.
2. Surfaces outdoors tend to rust because of humidity, which in turn results in a dielectric-like behavior in reflectance.

This example also shows that our method is robust to cluttered environments as can be found in city centers.

Red bricks (figs. 6 and 7, second row) This sample was captured an early afternoon on a rather clear day, which hints for strongly linearly polarized incident illumination. Because we are taking measurements close to Brewster angle, we are able to accurately extract high-resolution reflectance maps.

Garden pavement (figs. 6 and 7, third row) In this example, we can see that our method is also well-suited to cloudy conditions where the illumination from the sky tends to be depolarized due to multiple scattering in the clouds. Furthermore, both this sample and the red bricks are largely diffuse-dominant, to which our method is robust.

Red book (figs. 6 and 7), fourth row demonstrates the robustness of our method to conditions where the sky is neither overcast nor open. In such conditions, the incident illumination is not only partially polarized but also has gradients in intensity.

Overall, our proposed method is able to estimate high quality reflectance maps fig. 6 under various lighting conditions and for different types of materials, from diffuse-dominated rocks and bricks to highly specular metallic surfaces. These maps can then be plugged into any rendering system to produce photorealistic renderings (fig. 7,d) that closely match reality (fig. 7, b-c).




Figure 6: Reflectance maps ((a)-(d)) estimated from two views of the same scene taken close to Brewster angle of incidence, under outdoors natural illumination ((e)). Our method is agnostic to the incident illumination and robust to changes in illumination during capture.

6.2 Discussion

The main assumption for our method to work is that measurements should be made at 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.

Normal estimation Figure 5, second row, shows plots of the TRS for different angles of incidence θ_i , 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. Furthermore, within a 10° window around Brewster angle, it can be observed that the maximum of the TRS is always within a small window of the true azimuth. We calculated the mean error in azimuth estimation to be around 3 to 4° when within 10° of Brewster angle. The latter is an easy condition to fulfill when making measurements: in all our measurements, we only subjectively picked directions that we

judged close to Brewster angle, and found no problem processing the measured data.

Diffuse-specular separation

One added benefit of making measurements close to Brewster angle is that the component of light polarized parallel to the plane of incidence is completely transmitted at Brewster angle. It follows that at Brewster angle, when the polarizing filter's axis is parallel to the plane of incidence, the s-polarized component is completely blocked while the diffuse and p-polarized component are allowed through the polarizer. The p-polarized component being 0, only the diffuse is left thus providing a means for diffuse-specular separation.

6.3 Limitations

Our method currently models all types of materials as dielectrics in polarization analysis. While this works well for many dielectric-metal composites and metallic surfaces commonly found in outdoor environments, it certainly is not accurate for highly metallic sur-




Figure 7: Comparisons of the scene photographs (b) taken under the illumination conditions shown in (a) with renderings (c) in the same conditions. Here we also show renderings in the Pisa environment (d). Notice how corresponding renderings are matching real photographs well.

faces exhibiting a complex index of refraction. We currently rely on VisualSFM to provide us with multiview 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 inverse rendering step currently does not model the incident polarization of the environmental illumination. While this has worked well enough for specular roughness estimation under the illumination conditions in which we made our measurements, there could certainly be conditions (e.g., clear sunny sky) where the illumination can be more strongly polarized which can lead to errors in reflectance estimation (under or over) with our current approach.

7 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, multiview 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 metallic 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

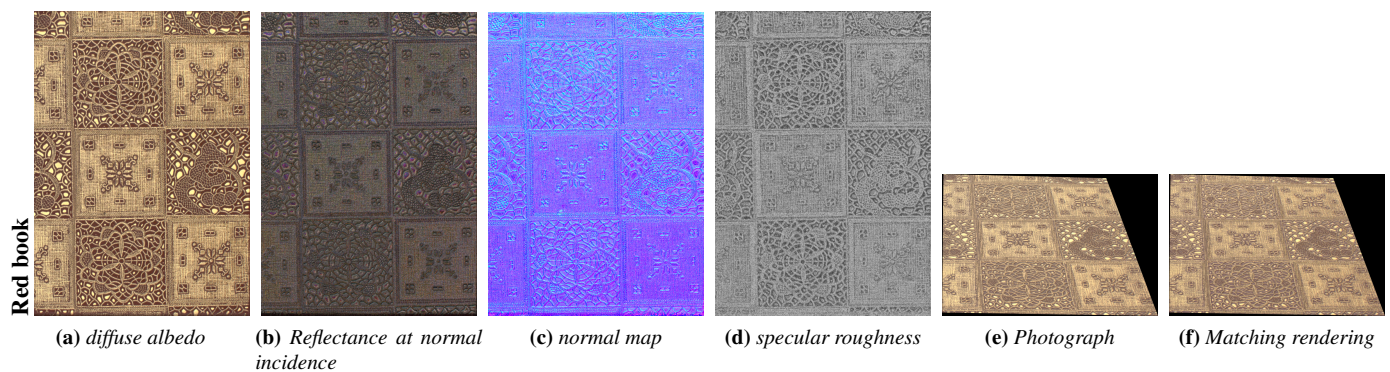


Figure 8: Additional result: Book cover with small-scale features, captured in the same conditions as the red book

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. We would like to more completely address this issue in future work by incorporating the incident polarization in our inverse rendering step for specular roughness estimation and also extend our analysis more accurately for truly metallic surfaces exhibiting a complex index of refraction.

References

- AITALA, M., WEYRICH, T., AND LEHTINEN, J. 2013. Practical SVBRDF capture in the frequency domain. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 32, 4.
- AITALA, M., WEYRICH, T., AND LEHTINEN, J. 2015. Two-shot svbrdf capture for stationary materials. *ACM Transactions on Graphics* 34, 4, 110.
- ATKINSON, G., HANCOCK, E. R., ET AL. 2006. Recovery of surface orientation from diffuse polarization. *Image Processing, IEEE Transactions on* 15, 6, 1653–1664.
- 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 (July), 117:1–117:11.
- COLLETT, E. 2005. *Field Guide to Polarization, SPIE Field Guides vol. FG05*. SPIE.
- COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM TOG* 1, 1, 7–24.
- DEBEVEC, P. 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*.
- DONG, Y., CHEN, G., PEERS, P., ZHANG, J., AND TONG, X. 2014. Appearance-from-motion: recovering spatially varying surface reflectance under unknown lighting. *ACM Transactions on Graphics (TOG)* 33, 6, 193.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *ACM Trans. Graph. (Proc. SIGGRAPH)* 22, 3, 749–758.
- GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. E. 2009. Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Comput. Graph. Forum* 28, 4, 1161–1170.
- GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. 2010. Circularly polarized spherical illumination reflectometry. *ACM Trans. Graph.* 29 (December), 162:1–162:12.
- GLENCROSS, M., WARD, G. J., MELENDEZ, F., JAY, C., LIU, J., AND HUBBOLD, R. 2008. A perceptually validated model for surface depth hallucination. *ACM Trans. Graph.* 27, 3 (Aug.), 59:1–59:8.
- GUARNERA, G. C., PEERS, P., DEBEVEC, P. E., AND GHOSH, A. 2012. Estimating surface normals from spherical stokes reflectance fields. In *ECCV Workshop on Color and Photometry in Computer Vision*, 340–349.
- HAUAGGE, D., WEHRWEIN, S., UPCHURCH, P., BALA, K., AND SNAVELY, N. 2014. Reasoning about photo collections using models of outdoor illumination. In *Proceedings of the British Machine Vision Conference. BMVA Press*.
- HUYNH, C. P., ROBLES-KELLY, A., AND HANCOCK, E. 2010. Shape and refractive index recovery from single-view polarization images. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, IEEE, 1229–1236.
- KADAMBI, A., TAAMAZYAN, V., SHI, B., AND RASKAR, R. 2015. Polarized 3d: High-quality depth sensing with polarization cues. In *Proceedings of the IEEE International Conference on Computer Vision*, 3370–3378.
- KING, D. E. 2009. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research* 10, 1755–1758.
- KOSHIKAWA, K. 1992. A polarimetric approach to shape understanding of glossy objects. 190–192.
- LEE, S., KOO, H., CHO, N., AND PARK, J. 2006. Stochastic approach to separate diffuse and specular reflections. In *ICIP*.
- MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Rendering Techniques*, 183–194.
- MALLICK, S. P., ZICKLER, T., BELHUMEUR, P. N., AND KRIEGSMAN, D. J. 2006. Specularity removal in images and videos: A pde approach. In *ECCV*.
- MANTIUK, R., KRAWCZYK, G., MANTIUK, R., AND SEIDEL, H.-P. 2007. High dynamic range imaging pipeline:

- Perception-motivated representation of visual content. In *Electronic Imaging 2007*, International Society for Optics and Photonics, 649212–649212.
- MIYAZAKI, D., KAGESAWA, M., AND IKEUCHI, K. 2003. Polarization-based transparent surface modeling from two views. In *ICCV*, 1381.
- MIYAZAKI, D., TAN, R. T., HARA, K., AND IKEUCHI, K. 2003. Polarization-based inverse rendering from a single view. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, IEEE, 982–987.
- MIYAZAKI, D., SHIGETOMI, T., BABA, M., FURUKAWA, R., HIURA, S., AND ASADA, N. 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.
- MÜLLER, V. 1995. Polarization-based separation of diffuse and specular surface-reflection. In *Mustererkennung 1995*. Springer, 202–209.
- NAYAR, S., FANG, X., AND BOULT, T. 1997. Separation of reflection components using color and polarization. *IJCV* 21, 3, 163–186.
- NAYAR, S. K., FANG, X.-S., AND BOULT, T. 1997. Separation of reflection components using color and polarization. *IJCV* 21, 3.
- RAHMANN, S., AND CANTERAKIS, N. 2001. Reconstruction of specular surfaces using polarization imaging. *CVRP* 1, 149.
- REN, P., WANG, J., SNYDER, J., TONG, X., AND GUO, B. 2011. Pocket reflectometry. *ACM Trans. Graph.* 30, 4 (July), 45:1–45:10.
- RIVIERE, J., PEERS, P., AND GHOSH, A. 2015. Mobile surface reflectometry. *Computer Graphics Forum*, n/a–n/a.
- ROMEIRO, F., VASILYEV, Y., AND ZICKLER, T. 2008. Passive reflectometry. In *Computer Vision–ECCV 2008*. Springer, 859–872.
- SADJADIZ, F., AND SADJADI, F. 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.
- SAITO, H., OMATA, K., AND OZAWA, S. 1999. Recovery of shape and surface reflectance of specular object from rotation of light source. In *Proceedings of the 2nd international conference on 3-D digital imaging and modeling*, IEEE Computer Society, Washington, DC, USA, 3DIM'99, 526–535.
- SCHLICK, C. 1994. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 3, 233–246.
- STRUTT, J. W. 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, 107–120.
- THILAK, V., VOELZ, D. G., AND CREUSERE, C. D. 2007. Polarization-based index of refraction and reflection angle estimation for remote sensing applications. *Appl. Opt.* 46, 30, 7527–7536.
- UMEYAMA, S., AND GODIN, G. 2004. Separation of diffuse and specular components of surface reflection by use of polarization and statistical analysis of images. *PAMI* 26, 5.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, Eurographics Association, 195–206.
- WOLFF, L. B., AND BOULT, T. E. 1991. Constraining object features using a polarization reflectance model. *PAMI* 13, 7, 635–657.
- WOLFF, L. B. 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.
- WU, C. 2011. Visualsfm: A visual structure from motion system.