

Image-Based Acquisition and Modeling of Polarimetric Reflectance

SEUNG-HWAN BAEK, KAIST
TIZIAN ZELTNER, EPFL
HYUN JIN KU and INSEUNG HWANG, KAIST
XIN TONG, Microsoft Research Asia
WENZEL JAKOB, EPFL
MIN H. KIM, KAIST




Fig. 1. We have acquired the first polarimetric BRDF dataset of real-world materials that provides coverage of arbitrary viewing and lighting configurations using our novel acquisition setup (shown on the left top). Acquired pBRDFs can be used in physically based simulations that correctly account for the change in polarization state during multiple scattering. We visualize the polarimetric information using the degree, azimuth, chirality, and type of polarization at the wavelength of 550 nm, following Wilkie and Weidlich [2010].

Realistic modeling of the bidirectional reflectance distribution function (BRDF) of scene objects is a vital prerequisite for any type of physically based rendering. In the last decades, the availability of databases containing real-world material measurements has fueled considerable innovation in the development of such models. However, previous work in this area was mainly focused on increasing the visual realism of images, and hence ignored the effect of scattering on the polarization state of light, which is normally imperceptible to the human eye. Existing databases thus only

capture scattered flux, or polarimetric BRDF datasets are too directionally sparse (e.g., in-plane) to be usable for simulation.

While subtle to human observers, polarization is easily perceived by any optical sensor (e.g., using polarizing filters), providing a wealth of additional information about shape and material properties of the object under observation. Given the increasing application of rendering in the solution of inverse problems via analysis-by-synthesis and differentiation, the ability to realistically model polarized radiative transport is thus highly desirable.

Polarization depends on the wavelength of the spectrum, and thus we provide the first polarimetric BRDF (pBRDF) dataset that captures the polarimetric properties of real-world materials over the full angular domain, and at multiple wavelengths. Acquisition of such reflectance data is challenging due to the extremely large space of angular, spectral, and polarimetric configurations that must be observed, and we propose a scheme combining image-based acquisition with spectroscopic ellipsometry to perform measurements in a realistic amount of time. This process yields raw Mueller matrices, which we subsequently transform into Rusinkiewicz-parameterized pBRDFs that can be used for rendering.

Our dataset provides 25 isotropic pBRDFs spanning a wide range of appearances: diffuse/specular, metallic/dielectric, rough/smooth, and different

Authors' addresses: Seung-Hwan Baek, KAIST, School of Computing, Daejeon, South Korea, 34141; Tizian Zeltner, EPFL, , Hyun Jin Ku; Inseung Hwang, KAIST, School of Computing, Daejeon, South Korea, 34141; Xin Tong, Microsoft Research Asia, Wenzel Jakob, EPFL, , Min H. Kim, KAIST, School of Computing, Daejeon, South Korea, 34141, minhkim@kaist.ac.kr.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).
0730-0301/2020/7-ART139

<https://doi.org/10.1145/3386569.3392387>

color albedos, captured in five wavelength ranges covering the visible spectrum. We demonstrate usage of our data-driven pBRDF model in a physically based renderer that accounts for polarized interreflection, and we investigate the relationship of polarization and material appearance, providing insights into the behavior of characteristic real-world pBRDFs.

CCS Concepts: • **Computing methodologies** → *Image and video acquisition*; **Reflectance modeling**.

Additional Key Words and Phrases: polarimetric imaging, polarimetric bidirectional reflectance distribution function

ACM Reference Format:

Seung-Hwan Baek, Tizian Zeltner, Hyun Jin Ku, Inseung Hwang, Xin Tong, Wenzel Jakob, and Min H. Kim. 2020. Image-Based Acquisition and Modeling of Polarimetric Reflectance. *ACM Trans. Graph.* 39, 4, Article 139 (July 2020), 14 pages. <https://doi.org/10.1145/3386569.3392387>

1 INTRODUCTION

Realistic modeling of material appearance via the bidirectional reflectance distribution function (BRDF) is an essential ingredient of physically-based light transport simulation. In the last decades, the availability of data-driven BRDF models [Matusik 2003] based on goniophotometric measurements has led to significant improvements in the understanding of real-world scattering, enabling comparative analysis and inspiring the development of novel parametric BRDFs [Ashikhmin and Premoze 2007; Bagher et al. 2012; Burley 2012; Löw et al. 2012; Ngan et al. 2005] and applications [Nam et al. 2018, 2016].

Existing datasets mainly focus on the scalar or RGB intensity of scattered light, neglecting other physical properties of light such as its spectrum and polarization state [Baek et al. 2018; Kadambi et al. 2015; Saragadam and Sankaranarayanan 2019]. While neither is strictly necessary to produce images that are meant for human consumption, they play an important role in applications that require a particularly high level of accuracy, such as predictive rendering. Rendering algorithms are also increasingly used in an inverse sense, for instance by combining a differentiable renderer with a gradient-based optimization algorithm. Such methods can reconstruct shape and material properties from empirical observations, in which case the polarization state and spectral distribution provide important extra cues that can aid the inversion process.

In such tasks, a repository of empirical material measurements characterizing intensity, spectrum, and polarization would be a powerful asset, but motorized acquisition of such information at sufficiently high angular resolution remains prohibitive due to the extremely large set of configurations that must be observed. Existing methods have thus been limited to in-plane acquisition [Hyde IV et al. 2009; Riviere et al. 2012], which is too sparse to subsequently reproduce the measured material in a physically based simulation.

In this work, we propose an image-based acquisition scheme for isotropic *polarimetric BRDFs* (pBRDFs) along with a data-driven interpolant that enables usage of the captured data in light transport simulations. Polarization is known to depend on the wavelength of the spectrum [Huynh et al. 2013; Hyde IV et al. 2009], hence our goal is to acquire not only polarimetric but also multispectral BRDFs of real-world objects. To this end, we devise an efficient acquisition setup, where we combine an image-based BRDF acquisition

method [Marschner et al. 2000; Matusik 2003] from computer graphics with spectroscopic ellipsometry [Azzam 1978] developed in the field of optics. Combining these two different methods enables us to efficiently sample the large and high-dimensional space underlying multispectral polarized BRDFs.

Using the proposed acquisition setup, we captured 25 spherical objects at five different wavelengths spanning the visible spectrum. The measurements are then converted into complete pBRDFs representation based the Rusinkiewicz parameterization and discretized into a 6D tensor that can be used for rendering. Our pBRDF database includes a wide variety of different properties: diffuse/specular, metallic/dielectric, rough/smooth, and different color albedos.

Lastly, we analyze the acquired pBRDFs to investigate the relation of pBRDFs with appearance parameters of surface normals, roughness, colors, and dielectricity, unveiling hidden links between pBRDFs and material appearance. We then present polarimetric rendering applications using the acquired data-driven pBRDF models. We compare our measurements of a common material with the measurements by an in-plane spectroscopic ellipsometry system, which we built to acquire the ground-truth measurement in the partial, directional space of pBRDF for the sake of the validation of our image-based measurements. Using extensive comparisons and simulations to alternative analytic polarimetric BRDF models, we demonstrate that our data-driven pBRDF model outperforms the state-of-the-art methods. To stimulate further research in this area and ensure reproducibility, we will release our code and pBRDF measurements. Figure 1 shows the captured materials, a polarimetric image rendered with our models, and polarimetric visualization.

2 RELATED WORK

Polarimetric BRDF models. Building on microfacet theory [Cook and Torrance 1982; Torrance and Sparrow 1967], several analytic polarimetric BRDF models have been proposed in prior work. They replace the unpolarized Fresnel term of the original model with a more accurate expression [Collett 2005] that accounts for changes in the polarization state [Hyde IV et al. 2009; Priest and Gerner 2000]. Recent work has furthermore added a polarized diffuse term [Baek et al. 2018]. Like microfacet models, the parameterizations of such pBRDFs span an intuitive low-dimensional space of material appearances that replicates the most important characteristics of single-layered metals and dielectrics. However, most real-world materials are significantly more complex—they consist of multiple layers with different properties and a multitude of effects on polarization, and their directional behavior is poorly approximated by standard distributions such as the Beckmann and GGX distributions.

We provide the first comprehensive database of multispectral polarimetric reflectance of real-world objects together with a data-driven pBRDF implementation. We believe that this database will lead to a better understanding of polarimetric reflectance, enabling the design of improved analytic pBRDF models.

Polarimetric BRDF acquisition. Ellipsometry is an optical measurement technique used to characterize how interactions with a material change the polarization state of incident light. For instance,

linearly polarized light that obliquely reflects from a surface typically becomes elliptically polarized. A standard problem in ellipsometry, then, entails determining parameters (eccentricity, azimuth, etc.) of such an ellipse, which provides a wealth of information about the material under observation. Ellipsometry is used to classify minerals or infer the properties of thin films, and has been studied for decades in the field of optics [Azzam 2016; Fujiwara 2007].

Different parameterizations can be used to represent polarized radiance and reflectance, the most general being *Stokes vectors* and *Mueller matrices*. The former is a 4D vector space, whose components encode total radiance, horizontal and diagonal linear polarization, and circular polarization. The latter are 4×4 matrices that describe how a Stokes vector must be modified following reflection.

The most common approach used to measure Mueller matrices is the *dual-rotating-retarder method* by Azzam [1978], which we also adopt in this article. Briefly, the idea is to illuminate and observe the material through optical paths that contain rotating *retarders*, which are optical elements—typically very thin slices of a birefringent material like quartz—that change the polarization state of light depending on angle. A sequence of measurements then yields a sufficient amount of information to reconstruct all 16 entries of the Mueller matrix.

In prior work, this technique has also been used for multispectral ellipsometry [Riviere et al. 2012; Wellems et al. 2000]. These previous pBRDF acquisition systems only capture angular configurations where the view/light vectors and the surface normal all lie in the same plane, which dramatically reduces the measurement time. On the flipside, the resulting data is far too sparse for direct use in a light transport simulation. Indirect use by fitting an analytic model is in principle feasible—however, this is also made difficult due to the limited availability of suitable parametric pBRDFs.

Recently, Boher et al. [2017] applied Fourier optics to the problem of efficient pBRDFs acquisition. Their approach relies on a beamsplitter and illumination with a fixed polarization state, which not only prevents reconstruction of the full Mueller matrix but potentially also introduces inaccuracies, since interactions with the beamsplitter at oblique angles will generally alter the polarization state. Our method similarly targets efficient pBRDF acquisition, but with the goal of capturing the full Mueller matrix on its high-dimensional domain by combining image-based BRDF acquisition and spectroscopic ellipsometry [Dupuy and Jakob 2018; Marschner et al. 2000; Matusik 2003].

Polarimetric rendering. Several works in the field of computer graphics have proposed rendering algorithms that track the polarization state of light via the Stokes/Mueller calculus. Wilkie and Weidlich [2012] were the first to propose a unidirectional method, and two later works build on the path integral formulation [Jarabo and Arellano 2018; Mojzik et al. 2016] to enable bidirectional connection strategies [Veach 1998]. Recently, Nimier-David et al. [2019] have proposed an open-source renderer, Mitsuba 2, that supports forward and inverse rendering with polarimetric light simulation. Our data-driven pBRDF provides access to realistic polarized material models that are compatible with any of these methods.




Fig. 2. (a) A Stokes vector can be represented in a coordinate system, where the z -axis is aligned with the propagation direction. The x -/ y -axes are in the tangent plane forming a local basis. (b) Polarization ellipse describes the polarization property of light by describing the tip of the electric field vector forms an ellipse in the tangent plane. The orientation angle of the ellipse ψ and the ellipticity angle χ determine the shape of the ellipse.

3 BACKGROUND ON POLARIZATION

This article relies on Stokes vector and Mueller matrix-based parameterizations of polarized illumination and reflectance; hence, we begin with a brief review of these representations. A very comprehensive discussion is provided by Collet et al. [2005].

Stokes vector and Mueller matrix. The Stokes vector is a four-dimensional quantity that fully describes the polarization state of light traveling along a ray. Its component representation $\mathbf{s} = [s_0, s_1, s_2, s_3]$ is defined in a specific coordinate system, where the z -axis aligns with the direction of propagation. The x, y axes can be chosen arbitrarily to form a right-handed coordinate system (Figure 2(a)); hence, the properties of any given ray can be described by an infinite set of Stokes vectors with different axes. The elements of the Stokes vector encode the following information: s_0 denotes the total amount of radiance, s_1 and s_2 is the amount of radiance with linear polarization at 0° and 45° , respectively, and s_3 denotes (right) circular polarization. Any other polarization state can be created using positive or negative superpositions of these components. The total radiance must satisfy the inequality $s_0^2 \geq s_1^2 + s_2^2 + s_3^2$. Light is fully polarized when both sides are equal—otherwise, it is only partially polarized. For examples, a Stoke vector of unpolarized light with intensity normalization is $[1, 0, 0, 0]$, a Stoke vector of 0° linearly polarized light is $[1, 1, 0, 0]$, and a Stoke vector of circularly polarized light is $[1, 0, 0, 1]$.

A property of Stokes vectors, particularly important in the context of rendering, is that they form a vector matching physical behavior, meaning that the addition of Stokes vectors replicates the additive combination of two beams of polarized light (this assumes that the beams are not mutually coherent).

The *polarization ellipse* (Figure 2(b)) provides another way of parameterizing the polarization state that is particularly convenient for drawing intuitive visualizations. Its parameters are the light intensity L , the degree of polarization ρ , the orientation angle ψ , and the ellipticity angle χ , which satisfy the following relation to Stokes vectors: $\mathbf{s} = L(1, \rho \cos 2\psi \cos 2\chi, \rho \sin 2\psi \cos 2\chi, \rho \sin 2\chi)$.

Finally, the Mueller matrix $\mathbf{M} \in \mathbb{R}^{4 \times 4}$ describes the response of a reflective surface subject to light with arbitrary polarization

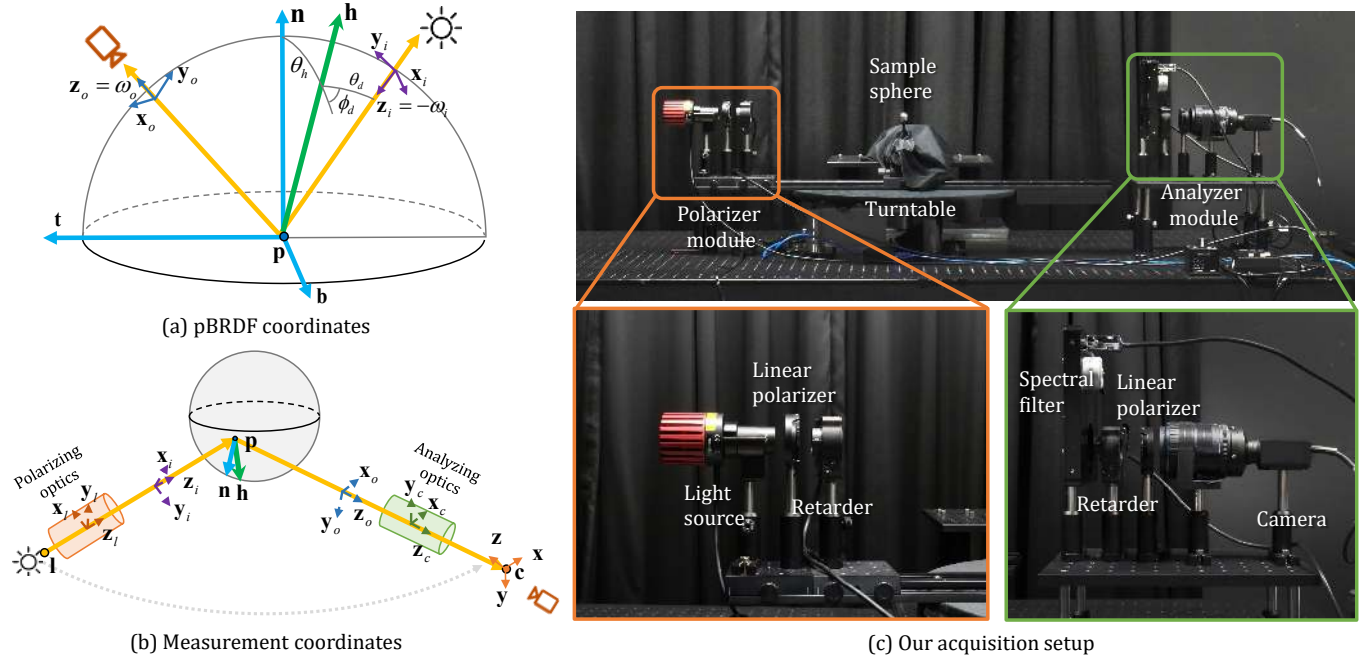


Fig. 3. Our coordinate system of pBRDFs and our acquisition system. (a) We parameterize our BRDF using half/difference angles θ_h , θ_d , and ϕ_d following Rusinkiewicz [1998]. The vertical axis of the incident/outgoing Stokes parameterization lies in the plane containing \mathbf{p} , ω_i , and ω_o . (b) We illuminate the sample using light with different polarization states that is analyzed following reflection. Both sides contain optics with their own respective coordinate systems: $\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i$ and $\mathbf{x}_c, \mathbf{y}_c, \mathbf{z}_c$, respectively. The world coordinates of the measurement setup are based on the camera coordinates $\mathbf{x}, \mathbf{y}, \mathbf{z}$. (c) Photograph of our acquisition setup. An isotropic sample sphere is placed at the center of the turntable. A polarizer module attached to the arm of the turntable rotates, illuminating the sphere from different directions. The polarizer module contains a linear polarizer and a rotating achromatic quarter wave plate to produce variety of different polarization states. The analyzer module then captures a series of sample images with the analyzing optics: spectral filters, a rotating achromatic quarter wave plate, and a linear polarizer.

states. Following reflection, the incident (\mathbf{s}_i) and outgoing (\mathbf{s}_o) Stokes vectors satisfy the relation $\mathbf{s}_o = \mathbf{M}\mathbf{s}_i$.

Polarimetric visualization. The linear and circular components of the Stokes vector are generally imperceptible to the human eye. To visualize polarized light in this article, we thus rely on a visual encoding proposed by Wilkie and Weidlich [2010]. Figure 1 is an example: it depicts the *degree of polarization* (DoP), i.e., the fraction of light that is fully polarized. The *angle of linear polarization* (AoLP) expresses the azimuth of the polarization ellipse and corresponds to the ψ parameter in Figure 2(b). The *chirality of polarization* (CoP) distinguishes between left- (yellow) and right- (blue) handedness of circular polarization, corresponding to the ellipticity angle χ . Finally, the *type of polarization* (ToP) denotes the relative degree of linear (cyan) versus circular (yellow) polarization and also corresponds to the ellipticity angle χ . We refer to the original article [Wilkie and Weidlich 2010] for a thorough discussion of this type of visualization.

4 POLARIMETRIC BRDF

Given a pair of incident (ω_i) and outgoing (ω_o) directions, a polarimetric BRDF model [Baek et al. 2018; Hyde IV et al. 2009; Priest and Gerner 2000] normally returns Mueller matrices characterizing the material's effect on polarized light, as opposed to scalar reflectance

values. In our case, the pBRDF additionally models the effect of polarization with respect to different wavelengths and thus takes the form $\mathbf{M}(\lambda, \omega_i, \omega_o)$, relating incident and outgoing spectral radiance:

$$\mathbf{s}_o^\lambda(\omega_o) = \int_{S^2} \mathbf{M}(\lambda, \omega_i, \omega_o) \mathbf{s}_i^\lambda(\omega_i) \cos\theta_i d\omega_i, \quad (1)$$

where \mathbf{s}_i^λ and \mathbf{s}_o^λ are now wavelength-dependent quantities, and $\cos\theta_i$ models foreshortening.

Coordinate convention. The coordinate system of Stokes vectors (and hence also that of Mueller matrices) is arbitrary, hence we clarify the convention used by our method. As shown in Figures 3(a) and (b), the axis of propagation of Stokes vectors follows the propagation of light, i.e., $\mathbf{z} = -\omega_i$ or $\mathbf{z} = \omega_o$ for the incident and outgoing directions, respectively. The y-axis is set to $\omega_i - \omega_o$ or $\omega_o - \omega_i$ for the incident and outgoing directions, respectively, and then orthonormalized with respect to \mathbf{z} via the Gram-Schmidt method. The last axis is given by $\mathbf{x} = \mathbf{y} \times \mathbf{z}$.

Parameterization. Like the MERL materials [Matusik 2003], our discretization of the pBRDF relies on the Rusinkiewicz [1998] parameterization, which leads to a significantly better-behaved interpolant. In the parameterization, directional data is represented with three parameters: the half- and difference vector elevation

$\theta_h, \theta_d \in [0, \pi/2]$, and azimuth difference $\phi_d \in [-\pi, \pi]$, resulting in the pBRDF $\mathbf{M}(\lambda, \theta_h, \theta_d, \phi_d)$.

Reciprocity. Scalar BRDFs satisfy the Helmholtz reciprocity condition, meaning that the BRDF is invariant to reversal of the incident and outgoing directions: $f(\omega_i, \omega_o) = f(\omega_o, \omega_i)$. Isotropic BRDFs furthermore satisfy a bilateral symmetry, meaning that the azimuth difference ϕ_d is invariant to a rotation by π radians. In the MERL database, this was, e.g., used to reduce the range of ϕ_d by half. While Helmholtz reciprocity approximately holds in our dataset, we found that large discrepancies occur when assuming bilateral symmetry during polarized measurements. Similar observations were made in prior studies of polarized reflectance [Clarke and Parry 1985; Leroy 2001] and diffuse-specular pBRDFs [Baek et al. 2018]. For this reason, our pBRDF parameterization uses $\phi_d \in [-\pi, \pi]$ in contrast to previous BRDF datasets.

5 POLARIMETRIC BRDF ACQUISITION

5.1 System Design

We now discuss our acquisition setup, which is motivated by the following desiderata to produce pBRDFs that can be used in computer graphics, computer vision, and for optical experimentation. In particular, our goals are to

- (1) capture the full 4×4 Mueller matrix including effects on linear and circular polarization.
- (2) achieve the directional coverage that is needed for rendering.
- (3) capture spectral variation in the visible spectrum, and
- (4) accomplish the above within a reasonable amount of time per material (on the order of days rather than months/years).

Prior work on measuring polarized reflectance has focused on individual items of this list, but no comprehensive technique exists to date. One noteworthy facet of (1) is that the Mueller matrix encodes the effects of *chirality*, i.e., it distinguishes clockwise versus counter-clockwise rotation in the polarization ellipse. This places additional constraints on our measurement setup because the chirality of light is invisible to simple sensing techniques that are only based on polarizing filters.

System configuration. Our approach tackles these challenges by combining an image-based acquisition system [Matusik 2003] with a spectroscopic ellipsometry acquisition system [Azzam 1978]. Our system, shown in Figure 3(c), captures pBRDFs of isotropic materials with complete angular polarimetric information (linear/circular, and chirality) on five spectral channels that regularly sample the visible range. Our setup consists of a machine vision camera, an apochromatic lens, a LED light source, an achromatic doublet, a turntable, a motorized spectral filter wheel with bandpass filters (450, 500, 550, 600, 650 nm), two motorized rotary stages equipped with achromatic quarter-wave retarders, and two manual rotary stages with linear polarizers. A detailed list of components is provided in the supplemental material.

Capturing procedure. This system is able to acquire high-dynamic range images of an isotropic spherical sample with variations in illumination angles, spectral bands, and polarization states. The light source is located at the end of a motorized rotating arm, while

the camera and sample remain static, the latter being fixed in a holder at the center of the device.

We designed our capture procedure, considering the rotation speed of the mechanical arm, the spectral filter wheel, and the motorized rotary stage for the retarders. The main thread of the capture program controls all motorized devices to capture HDR images. Using our acquisition system, we captured various samples of 25 materials (Figure 1(a)), using approximately 2.5 days per measurement. We will release the resulting material database and associated data-driven pBRDF model to the community to facilitate modeling of polarized reflectance. We now discuss handling of the various pBRDF dimensions in turn.

5.2 Directional Measurement

The pBRDF depends on incident/exitant light vectors, which together span a large four-dimensional space. We restrict ourselves to isotropic materials that are invariant under rotation around the surface normal, which removes one dimension. However, the remaining 3D space is still dauntingly large, particularly in combination with the other effects that must be captured (wavelength, polarization).

Following previous image-based acquisition systems for isotropic samples [Marschner et al. 2000; Matusik 2003], our system captures photographic images of spherical samples. This significantly reduces the overall acquisition time because many viewing/lighting combinations can be captured in parallel. Our device consists of a static analyzer module and a rotating polarizer module placed on a motorized arm. The former contains a camera (FLIR GS3-U3-91S6M), and the latter contains a broadband LED light source (Thorlabs MCWHLP1). Details on the specific components used are provided in the supplemental material. A single HDR photograph taken by the camera captures a full 2D slice of the BRDF, and rotation of the illumination direction provides coverage in the remaining dimension.

5.3 Spectroscopic Measurement

The refractive index of many materials depends on wavelengths, which manifests itself in a spectral dependence on polarization [Huyhn et al. 2013]. A wavelength dependent Mueller matrix can also arise in multi-layered materials containing layers that selectively absorb in some wavelengths. For this reason, we perform full spectro-polarimetric measurements with our acquisition system inspired by Azzam's system [1978] rather than assuming multiplicative separability of these effects.

We use a set of hard-coated narrow-band bandpass filters (Edmund Optics #88-299) to sample the wavelength domain. The coatings of such bandpass filters consist of hundreds of thin layers of dielectric material, whose thickness and index of refraction are optimized to achieve a boxcar function-shaped transmission profile. In contrast to other types of bandpass filters, such as liquid crystal tunable filters (LCTFs), they have no effect on polarization for light with perpendicular incidence (as is the case in our setup). We use filters with a bandwidth of 10 nm and mean wavelengths of 450, 500, 550, 600, and 650 nm. These are then placed into a motorized filter wheel holder located in front of the camera, as shown in Figure 3(c).




Fig. 4. Schematic diagram of our dual-rotating retarder setup.

We only use a single filter wheel and thus ignore fluorescence, which would require another filter in front of the light source.

5.4 Polarization Measurement

As stated previously, we wish to capture the full Mueller matrix including effects of scattering on the chirality of light, which rules out simpler methods that only rely on polarizing filters. We build on a classic ellipsometry approach known as the *dual-rotating-retarder* (DRR) method, which was pioneered by Azzam [1978]. It surrounds the sample with polarization optics and analyzer optics, which each consist of a polarizing filter (Edmund Optics #47-215) and a *quarter wave plate* (QWP) in a rotation mount (Thorlabs AQWP10M-580). A QWP is a transparent optical element that is part of the larger class of *retarders*, typically a circular slice of a birefringent (i.e., optically anisotropic) crystal such as quartz, which *retards*, i.e., delays, the phase of light that oscillates along a fixed direction in the crystal by $1/4$ wavelength relative to oscillations along the orthogonal axis. Its effect on the polarization state of light therefore depends on its rotation relative to the optical axis, and the Mueller matrix of a QWP is given by

$$\mathbf{Q}(\theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos^2 2\theta & \sin 2\theta \cos 2\theta & -\sin 2\theta \\ 0 & \sin 2\theta \cos 2\theta & \sin^2 2\theta & \cos 2\theta \\ 0 & \sin 2\theta & -\cos 2\theta & 0 \end{bmatrix}, \quad (2)$$

which (depending on θ) converts linear polarization to left or right circular polarization and a 1D space of elliptical states in between.

Using the combination of a polarizer and a rotating QWP (angle θ), the DRR method thus “probes” the sample with various forms of polarized light. The reflection is measured after passing through symmetric analyzer optics consisting of a second rotating QWP (angle θ') and a second polarizer as shown in Figure 4.

Measuring many combinations of θ and θ' would be burdensome, and Azzam showed that fixing $\theta' = 5\theta$ produces an intensity $f(\theta)$ that, when expressed as a 12-term Fourier series

$$f(\theta) = a_0 + \sum_{k=1}^{12} [a_k \cos(2k\theta) + b_k \sin(2k\theta)], \quad (3)$$

leads to an analytic solution of the Mueller matrix in terms of the series coefficients. For instance, $\mathbf{M}_{11} = a_0 - a_2 + a_8 - a_{10} + a_{12}$, $\mathbf{M}_{12} = 2a_2 - 2a_8 - 2a_{12}$, etc. While this approach can work very well with sensors that are continuously read out during a rotation, it is prohibitive for sensors with an intrinsic exposure time, such

as CCD/CMOS cameras. In this case, a very large number of photographs would be required to perform a sufficiently high quality continuous Fourier transform into coefficients a_k, b_k ($k = 1, \dots, 12$).

Here, a more resourceful way of obtaining Mueller matrices entails the solution of a linear least squares problem

$$\mathbf{M}_{\text{meas}} = \underset{\mathbf{M}'}{\operatorname{argmin}} \sum_{k=1}^K \left(f(\theta_k) - \left[\mathbf{A}(\theta'_k) \mathbf{M}' \mathbf{P}(\theta_k) \right]_{11} \right)^2, \quad (4)$$

for a much smaller number of measurements $f(\theta)$ taken at angles $\theta_1, \dots, \theta_K$. This term inside the sum constitutes a *forward model* of the entire optical system: $\mathbf{P}(\theta)$ denotes the combined Mueller matrix of the polarization optics on the light source side, \mathbf{M}' is the hypothesized Mueller matrix of the sample, $\mathbf{A}(\theta')$ is the combined Mueller matrix of the analyzer optics, and $[\dots]_{11}$ extracts the (1, 1) entry corresponding to overall radiance measured by the camera. We then optimize over \mathbf{M}' so that it becomes consistent with the actual observation $f(\theta)$, which can be accomplished via the least-squares solution of an overdetermined linear system with K equations and 16 unknowns. As before, we set the second QWP rotation to a fixed multiple $\theta' = 5\theta$, which yields a well-conditioned linear system [Smith 2002], though other combinations are also possible. The definitions of the polarizer and analyzer terms are

$$\mathbf{P}(\theta) = \mathbf{R}(\theta) \mathbf{L}(0) \mathbf{I} \quad \text{and} \quad \mathbf{A}(\theta') = \mathbf{L}(0) \mathbf{R}(\theta'), \quad (5)$$

where $\mathbf{L}(0)$ is the Mueller matrix of a linear polarizer at horizontal orientation, and \mathbf{I} denotes the unpolarized intensity of the light source.

Our measurement technique exploits the property that the polarimetric transformations $\mathbf{P}(\theta)$ and $\mathbf{A}(\theta')$ can be shared by many simultaneous measurements when Azzam’s scalar sensor is replaced by a high-resolution camera. Our system then takes $K = 36$ measurements that regularly sample $\theta \in [0, \pi/2]$ and applies the above optimization independently to each pixel.

Finding \mathbf{M}_{meas} via optimization has two additional benefits: even achromatic QWPs like the ones used in our setup deviate from ideal $1/4\lambda$ retardation across different wavelengths. Such non-ideal retardation breaks the analytic solution of the Fourier approach, but is easy to account for in the least-squares solver. Another benefit is that regularization and priors on expected structures of Mueller matrices can be imposed straightforwardly. We currently do not use priors or regularization but account for spectrally varying retardation.

One important point to note is that the resulting pixel-matrices \mathbf{M}_{meas} are not directly usable in a pBRDF interpolant $\mathbf{M}(\dots)$ and require further conversion. In particular, the above optimization produces matrices defined in the polarizer/analyzer coordinate system, and which are unaware of the varying tangential planes observed through different pixels. We thus perform an additional rotation from the measurement frame to the convention discussed at the beginning of Section 3:

$$\mathbf{M} = \mathbf{C}_{\text{out}} \mathbf{M}_{\text{meas}} \mathbf{C}_{\text{in}}. \quad (6)$$

To acquire the pBRDF of a new material, our acquisition setup captures 26,460 HDR images of a sphere with 147 different light

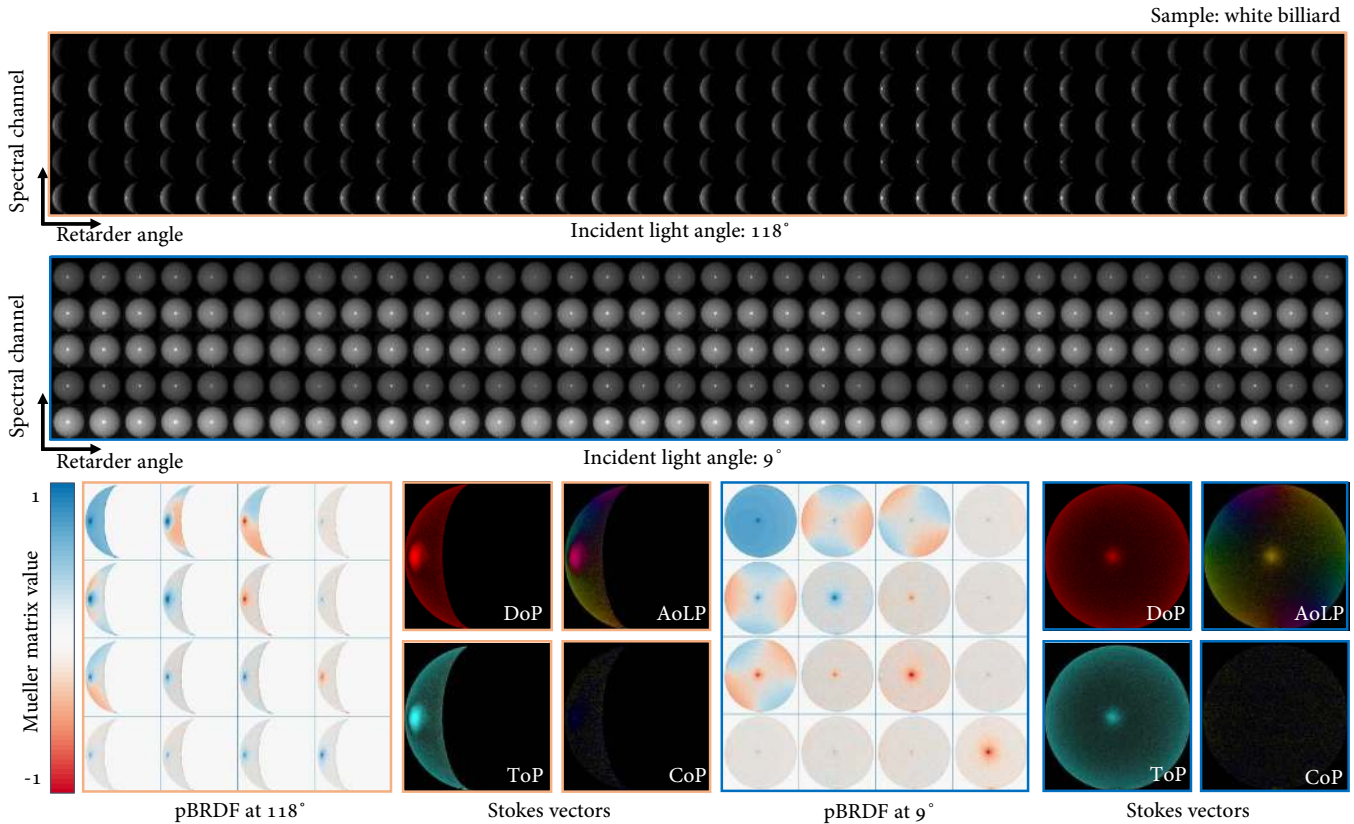


Fig. 5. We captured input images polarized with 147 different longitude angles of the light source from the camera from 9° to 173° degrees about the center of the sphere. For each light position, five spectral bands from 450 nm to 650 nm with 50 nm intervals, and 36 different rotation angles of the retarders are captured. The first and the second rows show the captured images of the white billiard ball for two different light angles. For each light angle and spectral channel, we reconstruct a per-pixel pBRDF in a form of Mueller matrix. These matrices are visualized with gamma correction ($\gamma = 5$). Also, we visualize the Stokes vector of the reflected light at the wavelength of 550 nm under an unpolarized incident light: DoP, AoLP, ToP, and CoP.

positions, 5 spectral bands, and 36 retarder angles. The specifics of the angular sampling are discussed in the next section.

Depending on roughness and reflectivity of the captured material, we acquire either four or eight photographs with different exposure times per configuration and merge them into an HDR image. Figure 5 shows the captured data, Mueller matrices, and a polarimetric visualization of a spherical sample (a white billiard ball).

5.5 Polarimetric BRDF Tabulation

The optimization technique discussed in the previous section assigns a Mueller matrix to every pixel on the spherical target. In the next phase, we transform these measurements into the Rusinkiewicz parameterization and discretize them into a 6D tensor. Its dimensions correspond to the half and difference angles $\theta_h, \theta_d, \phi_d$, wavelength λ , and the row and column of the Mueller matrix. For the former three, we use a discretization with resolution [91, 91, 361] following Matusik [2003], also adopting their nonlinear mapping in the θ_h parameter. This yields a total of 14, 947, 205 Mueller matrices and thus roughly 912 MiB of storage per material in single precision.

Backward mapping. Based on the geometry of our acquisition setup, we construct a function that maps from entries of this 6D

tensor to positions in the input photographs. To populate the tensor with the measured data, we then simply iterate through its entries and perform lookups through this backward mapping. During this process, it is important to note that the input photographs contain measurement noise, and that tensor entries often map to large regions on the spherical sample—a single lookup may therefore be unnecessarily noisy. To reduce measurement noise during tabulation, we perform 10 jittered lookups and record the median value.

Hole filling. As in other devices for measuring reflectance, there are certain directions that cannot be measured due to occlusion. In our case, this happens for retro-reflective directions with $0^\circ \leq \theta_d \leq 3^\circ$. We fill these missing values in our tabulation by extrapolating the pBRDF values from nearby entries. In particular, we solve a Poisson problem via Gauss-Seidel iterations to find the smoothest interpolant along the angular dimensions (i.e., not mixing the spectral or polarization-related dimensions). To clearly distinguish measurement from extrapolated data in future analysis of our pBRDFs, we will also release a raw version of our data with holes.

6 CALIBRATION

Measurements using our system require careful geometric, radiometric, and polarimetric calibration. Most calibration steps should only be performed once, while others are sample-dependent. We briefly discuss them in turn.

Geometric calibration. We use the Bougouet calibration toolbox based on the method of Zhang [2000] to obtain the camera’s intrinsic parameters from checkerboard images, specifically: focal length, principal points, and distortion parameters. We observed an average reprojection error of 0.47 pixel for our camera that has a pixel pitch of 3.69 μm .

The light source is installed at the end of the robotic arm on a motorized turntable. To ensure that the plane of its orbit is perpendicular to the camera’s vertical axis, we balanced the center of mass mechanically with a high-precision spirit level and a set of metal weights. In addition, the turntable body, the camera axis and the optical tables are also individually calibrated to be parallel.

We approximate the light source as a point light, whose exact position on the arm must be calibrated via photographs. In our setup, the light emitted from the LED light source passes through an achromatic doublet to weakly collimate incident light on the sphere. Therefore, the *virtual* light source position lies at a greater distance compared to its physical position. To determine its effective position, we align the light source, the sample holder and the camera on a straight line. We then measure the projected diameters of the light illumination progressively at five different distances from the location of the sphere holder to the light source. This enables us to calibrate the effective distance between the virtual light source and the sphere holder as a constant using the least-squares method.

Finally, once the obtained light positions are determined with respect to the sphere sample, we convert them to the camera coordinates with the known 3D position of the sphere sample holder.

To validate our light source’s position and angle, we evaluate the geometric calibration accuracy by projecting the virtual highlights (computed by our geometry model) to captured specular highlights on a chrome sphere. We observed a mean reprojection error of 1.79 pixel across all 147 light source angles.

Radiometric calibration. To correct for spectrally varying emission from the LED and transmission variations in the analyzer and polarizer modules, we capture a spherical Spectralon target with flat 99% reflectance (Labsphere SRM-990) using the five different filter wheel configurations. We determine a constant scaling factor per spectral band from the resulting measurement to ensure correct color rendition and energy conservation of measured pBRDFs.

Polarimetric calibration. Accurate polarimetric reconstruction requires that all optical elements in the analyzer and polarizer have the correct rotational offset around the optical axis. For instance, the fast axis of the QWPs should be perfectly horizontal at the beginning of each measurement—a drift of 1-2 degrees can introduce significant errors. To ensure that this is the case, we initially remove all optical elements and configure them in an extinction configuration. We begin by installing a single linear polarizer, using its printed indicator to align it with the camera’s horizontal axis. We then add the second polarizer and rotate it into a cross-polarized configuration that can

easily be tested by directly observing the light source through both filters when all elements are aligned in a straight line. We then add the QWPs individually with their (highlighted) fast axis oriented horizontally. In each case, we perform a fine alignment to ensure that the extinction configuration is preserved. Finally, we rotate the second polarizer by 90° to achieve maximum transmission.

Sample-specific calibration steps. The radius of each of our spherical samples is slightly different and ranges from 15.03 mm to 25.4 mm. To correct for this variation, we calibrate 3D surface points on the sphere of each sample. We first measure the radius of each sphere using a caliper and also capture the sphere under ambient illumination to capture the overall shape of the sphere in the screen space. After applying the camera intrinsic calibration to the sphere shape captured in the screen space, we estimate the position of the sphere’s center with respect to the camera coordinate system through trigonometry by matching the sphere with the measured radius and the backprojected sphere in 3D.

Refraction calibration. The retarders and spectral filter have a non-zero thickness, and we found that this causes a small but noticeable offset in photographs that are taken through the retarder. This offset furthermore depends on the angular configuration, which impedes extraction of the final pBRDF. We therefore calculate homography matrices using a small checkerboard target that is visible in each photograph and use them to correct for this offset.

7 POLARIMETRIC RENDERING

Rendered images shown in this paper were produced using a polarized Monte Carlo renderer based on Jarabo and Arellano’s [2018] path integral formulation of light transport implemented on Mitsubishi 2 [Nimier-David et al. 2019]. In particular, the Stokes vector of a rendered pixel j is computed using an unbiased estimate of the path integral

$$\mathbf{s}^{(j)} = \int_{\mathcal{P}} \mathbf{f}^{(j)}(\mathbf{x}) d\bar{\mathbf{x}}, \quad (7)$$

over a path space with elements $\bar{\mathbf{x}} = (\lambda, \mathbf{x}_1, \dots, \mathbf{x}_k) \in \mathcal{P}$ endowed with an additional wavelength dimension. The vector-valued path contribution function \mathbf{f} pixel j is given by

$$\mathbf{f}^{(j)}(\lambda, \mathbf{x}_1, \dots, \mathbf{x}_k) = \mathbf{W}_e^{(j)}(\mathbf{x}_0 \leftarrow \mathbf{x}_1) \left[\prod_{i=2}^{k-1} \mathbf{M}(\mathbf{x}_{i-1} \leftarrow \mathbf{x}_i \leftarrow \mathbf{x}_{i+1}) \right] \left[\prod_{i=1}^{k-1} G(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1}) \right] \mathbf{L}_e(\mathbf{x}_{k-1} \leftarrow \mathbf{x}_k), \quad (8)$$

where W_e is the importance, \mathbf{M} denotes the Mueller matrices of scattering interactions, G is the (scalar) geometry term including visibility, and \mathbf{L}_e is the polarized emission profile at position \mathbf{x}_k located on a light source. All functions except for G possibly depend on wavelength λ , which is omitted for readability.

Compared to typical implementations of rendering algorithms, the product of Mueller matrices in the above expression requires an additional degree of scrutiny: polarized BRDFs are generally non-reciprocal, and the matrix-matrix products furthermore do not commute. It is thus imperative that the multiplication is carried out in the right order, and with the pBRDF’s directional arguments

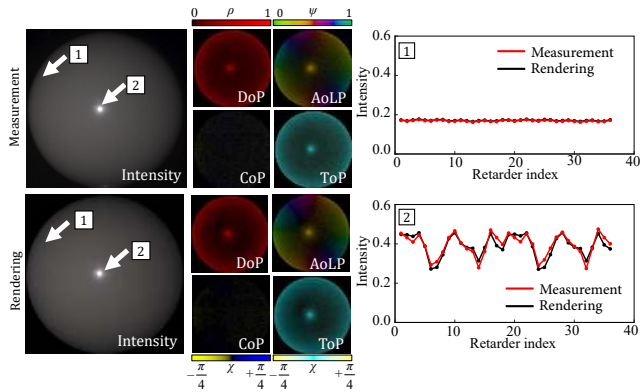


Fig. 6. Photograph vs. rendering. The first row shows the pBRDF measurement of the intensity, DoP, AoLP, CoP, and ToP of the white billiard ball at the wavelength of 550 nm under unpolarized illumination. The second row presents polarimetric rendering images with our pBRDF model in terms of the same polarization properties. These two measurement and rendering present a strong agreement in terms of polarization states.

matching the propagation of light. Another crucial detail is that pBRDFs return Mueller matrices in a local coordinate frame that depends on ω_i and ω_o . Simulating interreflection involving multiple scattering events requires intermediate rotations that reconcile frame incompatibilities along the way.

Spectral rendering. Our renderer integrates over continuous wavelengths spanning the visible range, while the captured pBRDF data provides five discrete spectral samples at 50 nm intervals. To reconcile this difference, our data-driven implementation of \mathbf{M} performs linear interpolation to evaluate the model at intermediate wavelengths. Lookups outside of the captured range are clamped (in other words, we use constant extrapolation). Finally, the importance function W_e performs a product integral of the spectral path throughput against the CIE color matching functions, which results in an RGB Stokes vector per pixel in the rendered image.

Importance sampling. Similar to the MERL BRDF [Matusik 2003], the half/difference-direction transformation of the data unfortunately prevents direct importance sampling of scattered directions. We instead approximate the overall scattered radiance (\mathbf{M}_{11} entry) with a linear combination of a Lambertian and GGX microfacet model with the weights of 0.1 and 0.9, using the resulting function as proxy for sample generation. We set roughness value of the GGX model manually per material to guide specular samples to high-valued BRDF directions. Following sample generation, we evaluate the data-driven model to determine associated BRDF values. We note that more sophisticated fitting strategies, such as the approach of Dupuy et al. [2015], may also be useful in this context. That said, we find that our simple scheme yields excellent noise characteristics in renderings (Figures 1 and 8) for the entire range of the captured materials.

Next-event estimation. One important sanity check of our system is that we are able to synthetically regenerate the photographs of the spherical sample in our Monte Carlo renderer (Figure 6). One challenge here is that polarizing and analyzing optics are ordinarily

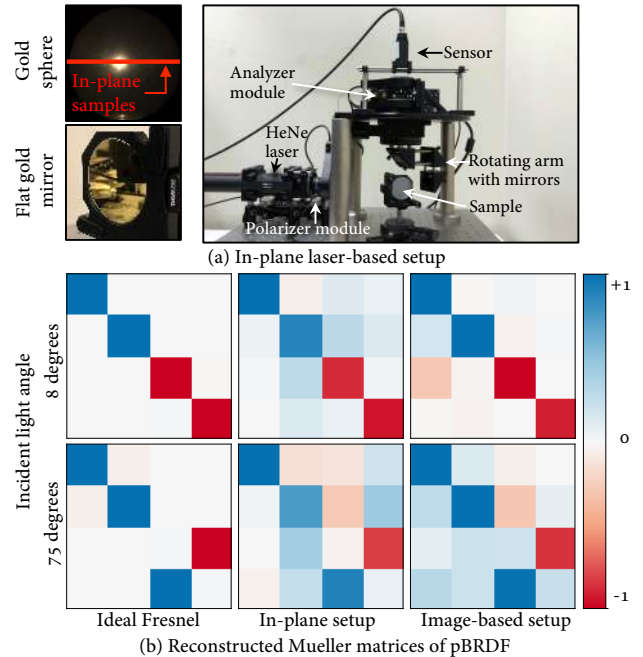


Fig. 7. Validation setup using an in-plane measurement. (a) We built a laser-based ellipsometry system that can accurately measure in-plane pBRDFs of a flat sample. We captured a gold plane that is made of the same gold material with the gold sphere. Note that this in-plane setup can only capture partial angle combinations indicated by the red line while our image-based setup can capture the complete angle combinations. (b) Assuming the sample as a pure gold material, we compare the obtained pBRDFs by the laser-based setup and our image-based setup with the ideal Fresnel reflection at two incident angles. Our image-based system accurately captures pBRDFs compared to the laser-based one, which well matches with the ideal gold material of perfect Fresnel reflection.

treated as occluders, which effectively disable direct illumination sampling strategies and results in renderings that are contaminated by severe variance. To address this issue, we implemented a connection strategy that is able to sample positions on light sources observed through an arbitrary number of polarizing filters and QWPs.

8 RESULTS

We now turn to validation experiments and analysis that are enabled by our pBRDF database.

8.1 Comparison with In-plane Measurement

Tracking the polarization state of light introduces a disproportionate number of pitfalls that must be considered during both measurement and simulation. These include the convention of how polarization ellipse is parameterized (e.g., azimuths counter-clockwise as seen from the sensor), frame conversions from the spherical sample to a reference frame, and further frame conversions between different components in a renderer. The lack of reference data (e.g., a “polarized Cornell box”) and our inability to see polarization further exacerbate these difficulties. As previously discussed, one important sanity check is the ability to re-render images of the spherical sample




Fig. 8. Polarimetric rendering. Our pBRDF dataset enables physically-meaningful polarimetric rendering with various real-world materials. For the top row, we render an object under an unpolarized environmental light map. The below rows visualize the rendered images of the same scene when we place polarization filters in front of the camera: horizontal linear polarizer (+LPH), vertical linear polarizer (+LPV), and right-hand circular polarizer (+CPR) from the second row to the last row. To only focus on polarimetric differences and not overall brightness change, we show absolute error compared a rendering with an (unpolarizing) neutral density filter that removes the same average intensity.

using a simulated reproduction of our measurement setup. To exclude the chance of systematic coordinate inconsistencies that might not be visible in re-rendered images, we built a second polarized light measurement device that provides in-plane measurements.

This device, shown in Figure 7, also builds the dual rotating retarder method but relies on Azzam’s frequency-space variant with a dense set of retarder configurations (see Equation (3)) rather than solving a least squares problem based on a few observations. We illuminate a *flat* sample with a stabilized HeNe laser (633 nm) being linearly polarized for the dual-rotating-retarder setup and use QWPs that are optimized for this wavelength. We avoid de-stabilizing back-reflections from optical elements into the laser using a Faraday isolator. Light striking the sample is then reflected by a sequence of 3 mirrors that guide light to a statically mounted analyzer module at the top. The sample and the three mirrors can be rotated independently, which enables measurements for arbitrary values of θ_i , θ_o . However, azimuths are constrained to $\phi_d = \pm 180^\circ$. We note that the sequence of mirrors change the polarization state of reflected light, and we therefore perform reference measurement without a sample to capture the Mueller matrix of the device itself, whose inverse is applied to subsequent measurements. One important design aspect of this device that the flat sample and planar configuration remove many of the coordinate conversions that would ordinarily be required.

We then measured the Mueller matrix of an elemental gold surface, specifically: a gold-plated metal sphere on a Carbon steel base, and an unprotected gold mirror (Thorlabs PF20-03-M03) in both devices. We also computed the ideal Mueller matrix predicted by the complex index of refraction of gold at 633 nm. The resulting matrices are shown in Figure 7 and provide an addition confirmation of the correctness of the various coordinate conversions even though there is still discrepancy between the results coming from compositional difference of the spherical sample and the planar sample, in addition to potential systematic errors.

8.2 Polarimetric Rendering

Following acquisition of our pBRDF dataset, we are able to render photorealistic images that are accurate not only in terms of spectral intensity but also polarization state. To validate our data-driven pBRDF model, we initially created a digital replica of our measurement setup in the renderer to simulate the measurement procedure of capturing images with different retarder angles, light source positions, and spectral filters. Figure 6 shows a comparison between real and re-rendered spherical samples; an animated comparison is shown in the supplemental video.

Figure 8 shows renderings with eight different materials that exhibit a range of different polarimetric effects, and Figure 1 visualizes polarization parameters in a more complex setting.




Fig. 9. Impact of surface normals on polarization. These pBRDF measurements visualize the intensities, AoLPs, DoPs, and CoPs of a chrome ball and a white billiard ball at the wavelength of 550 nm. Under unpolarized illumination, the AoLP and DoP exhibit a clear relation with surface normals for the dielectric material. For the metallic material, the normal dependency becomes weaker for the AoLP and DoP. However, for the chrome sample, we can observe another normal dependency on the CoP.

8.3 Material Appearance vs. Polarization

We now use the captured data to investigate the interplay of material appearance and its effect on polarization. We believe that observations of such relationships will motivate the development of analytic pBRDF models in the future.

Surface normals. One of the most popular uses of polarized light is the estimation of surface normals by analyzing the degree and angle of linear polarization. Previous methods depend on simplified analytic models that assume mirror-like specular or perfect Lambertian reflectance. Our dataset enables a new data-driven analysis of the relationship between surface normals and polarization.

Figure 9 shows the captured AoLP, DoP, and CoP for a metallic chrome ball and a dielectric white billiard ball under unpolarized point light illumination. In the case of the white billiard ball, the captured AoLP and DoP show a clear normal dependency due to *diffuse* reflection, as predicted by previous methods. We observe that this normal dependency changes in the regions, where both diffuse and specular reflection exist. Furthermore, there are strong differences in the nature of the normal dependency when comparing the billiard and chrome ball: the DoP of the metallic surface is almost constant, while the CoP presents a slight normal dependency that is also present in other materials in our dataset. See Figure 11 for examples. This effect on chirality has been rarely discussed in previous studies [Guarnera et al. 2012]. Developing shape-from-polarization methods with a more accurate model of the normal dependency of polarization could make them more robust in real-world scenarios. In this regard, our dataset could also be used to create realistic benchmarks for various shape-from-polarization methods.

Surface roughness. Surface roughness has traditionally been an important parameter to control the appearance of parametric BRDF models [Cook and Torrance 1982]. Previous pBRDF models assumed that the diffuse polarization and surface roughness are unrelated [Baek et al. 2018; Hyde IV et al. 2009; Priest and Gerner




Fig. 10. Roughness impact on polarization. These pBRDF measurements show the intensities, AoLPs, DoPs, and CoPs of the spherical Spectralon (fluoropolymer) and a ceramic ball (zirconium dioxide, ZrO₂) at the wavelength of 550 nm under unpolarized illumination. High surface roughness tends to reduce the normal dependency of the AoLP and DoP.




Fig. 11. Dielectric vs. metallic materials. These pBRDF measurements show the intensities, AoLPs, DoPs, and CoPs of the fake gold and real gold ball at the wavelength of 550 nm under unpolarized illumination. Dielectric “fake” gold ball presents similar appearance to the metallic gold material in terms of intensity. However, the chirality of polarization reveals that they have the *opposite handedness of polarization* (see the arrows).

2000], which motivated us to investigate the relation between surface roughness and polarization in our dataset. Figure 10 shows the AoLP and DoP for a Spectralon sphere made of fluoropolymer and a ceramic ball made of zirconium dioxide (ZrO₂). Our spectralon sample has a rough sanded surface, while the ceramic ball shows a much smoother finish. We observe that the ceramic material presents strong AoLP and DoP values close to the specular reflection, while the little amount of specular reflection of the Spectralon sphere is mostly unpolarized. We also see that the normal dependency of the AoLP and DoP becomes more evident for the ceramic ball with its low surface roughness. We hypothesize that multiple scattering both above and below the rough surface boundary induces a stronger depolarization effect in the case of Spectralon. Such relationships between polarization and surface roughness could be used as a cue to directly estimate roughness from polarization in inverse rendering applications.

Metallic vs. dielectric materials. Fresnel reflection is the fundamental optical mechanism that distinguishes dielectric from conducting surfaces. While dielectric surfaces can often appear metallic, e.g., via multi-layer coatings involving interference pigments, they




Fig. 12. Color and polarization. These pBRDF measurements show the intensities, AoLPs, and DoPs of five differently colored billiard balls at the wavelength of 550 nm under unpolarized illumination. We found that, at the wavelengths of weak diffuse reflection, the specular polarization dominates the diffuse polarization even near the outer boundary regions in the captured sphere image.

display fundamentally different polarimetric behavior that can be used to recognize them. We analyze two materials that present overall similar golden appearance: a gold-plated steel sphere with a thin coating of elemental gold, and a *dielectric* “fake” gold sphere. Figure 11 shows that the chirality of the latter material has the opposite handedness compared to the true gold surface. The effect could be used to identify the fundamental material class in inverse rendering applications.

Color and wavelength. Investigating the wavelength dependency of polarization is challenging, and many previous studies [Kadambi et al. 2015] have assumed a multiplicative separability of these dimensions. Here we investigate the wavelength dependency for differently pigmented objects made of the same material. Figure 12 shows the AoLP and DoP of different-colored billiard balls (white,

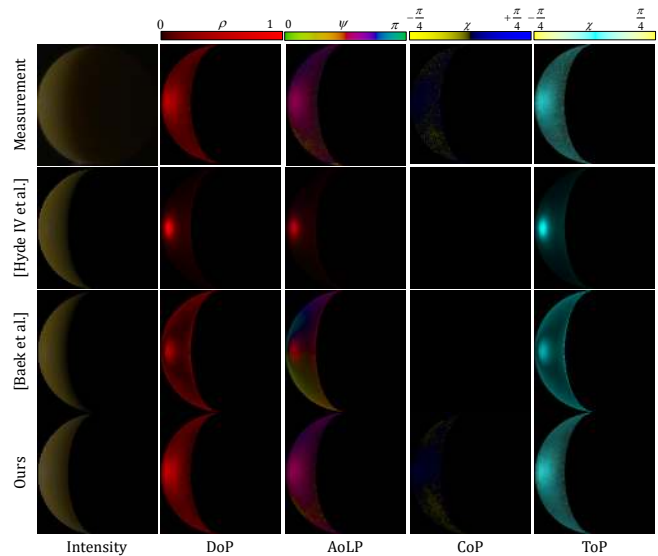


Fig. 13. Comparison to parametric pBRDFs. We fit two parametric pBRDF models [Baek et al. 2018; Hyde IV et al. 2009] to the yellow silicon material and observe significant disagreement in both intensity and polarimetric properties. The plot shows a lighting angle of 118° at 550 nm.

black, red, green, and blue) captured at 5 bands spanning 450–650 nm. For the white billiard ball that is spectrally uniform, we found that the observed AoLP and DoP are also similar. As indicated by previous approaches [Baek et al. 2018; Hyde IV et al. 2009; Priest and Gerner 2000], this could change depending on how much the refractive index of the material changes with respect to wavelength. In addition, we found that AoLP and DoP are low for wavelengths with a high diffuse albedo, which appears to reduce the effects of both diffuse and specular polarization. In contrast, AoLP and DoP are high for the black billiard ball, where specular polarization is dominant. For the red, green, and blue balls, this dependency on diffuse polarization is clearly visible across different wavelengths. This is an example of a wavelength-dependent polarization change, which also has been discussed in previous studies [Huynh et al. 2013; Riviere et al. 2012; Wellemes et al. 2000]. In addition to the specular billiard balls, we also acquired a set of diffuse silicon samples, which could be an interesting data source for further investigation of the spectral dependence of pBRDFs with varying roughness.

8.4 Analytic pBRDF Models

We conducted experiments to see whether existing analytic pBRDF models based on microfacet theory are able to reproduce the characteristics observed in our pBRDF database. Figure 13 shows such an example where we fit two models proposed in prior work to the captured yellow silicon material. The model of Hyde et al. [2009] only represents specular polarization and fails to reproduce polarization of the diffuse component. In contrast, the model of Baek et al. [2018] does include a diffuse polarization term, but we find that there remains significant disagreement between the fit and captured material. Naturally, the pBRDF model (last row) matches the measurement since it is an interpolant of the measured data. Details on the fitting procedure are provided in the supplemental material.

9 LIMITATIONS

Acquisition setup. Although we carefully aligned and calibrated all hardware components, the measurement setup is still not perfect. Acquiring samples at grazing angles is challenging due to the sphere-to-plane projection in small image regions. Interreflection between components of our measurement setup affects the measurement of extremely smooth samples with strong specular highlights (e.g., the black billiard ball) particularly by the spectral filters. Our placement of the spectral filters following the camera retarder aims to reduce this effect.

On the illumination path, the use of an achromatic lens to collimate the light source into a parallel beam in principle requires a point emission source. Although a pinhole could be used to improve the directional quality of the beam, ideal collimation cannot not be achieved with a broadband light LED when one must at the same time gather enough radiance for a CCD or CMOS sensor.

Also, our method requires samples with an ideal spherical shape, and geometric deviations introduce angular errors in the recorded pBRDF values. The top of the sample holder often reflects a small amount of light toward the bottom side of the sphere, producing a reflection. We remove associated regions and reconstruct them via the interpolation similar to other configurations that cannot be measured. Lastly, sensor noise is critical especially for accurate analysis of polarimetric measurements, which we tried to mitigate by HDR imaging with multi-exposure captures. These and other practical limitations leave room for improvement. We believe that the pBRDF dataset and our detailed discussion of technique, calibration, and limitations will stimulate future work on polarized reflectance acquisition.

Validity of captured data. Not all Stokes vectors correspond to valid physical configurations: s must have positive energy and DoP, i.e., $s_0 \geq 0$ and $s_0^2 - (s_1^2 + s_2^2 + s_3^2) \geq 0$. Based on these constraints, a Mueller matrix is considered valid if interactions with arbitrary valid Stokes vectors once more yield a valid configuration. This leads to three necessary and sufficient conditions for a Mueller matrix \mathbf{M} [Givens and Kostinski 1993]: $\text{Im}(\sum_k)_{k=1..4} = 0$, $S_{\sigma_1}^T \mathbf{G} S_{\sigma_1} \geq 0$, and $s_{0,\sigma_1} \geq 0$, where $[S, \sum_k] = \text{eig}(\mathbf{G}^T \mathbf{M} \mathbf{G})$, $\mathbf{G} = \text{diag}(1, -1, -1, -1)$, $\text{Im}()$ returns the imaginary component of a given complex number, σ is the eigenvalue, and S is the eigenvector. We checked these conditions on three of our pBRDF measurements: Spectralon, white the billiard ball, and brass. For the three samples, we have 93.79/87.06/43.31% of table bins which are valid according to the above criterion. This trend is related to the signal-to-ratio (SNR) of the measurements: for the brass sample, scattering away from the specular peak rapidly falls off, and the entries of Mueller matrices are therefore close to the noise floor. We believe that further post-processing of our measurements e.g., employing efficient convex, optimization techniques, could impose priors to ensure that the above spectral condition is satisfied even in low-SNR configurations.

Interpolation of missing configurations. As previously discussed, certain configurations, particularly those close to retro-reflection, cannot be measured and require interpolation. We solve a Poisson problem to find the smoothest interpolant per Mueller matrix entry,

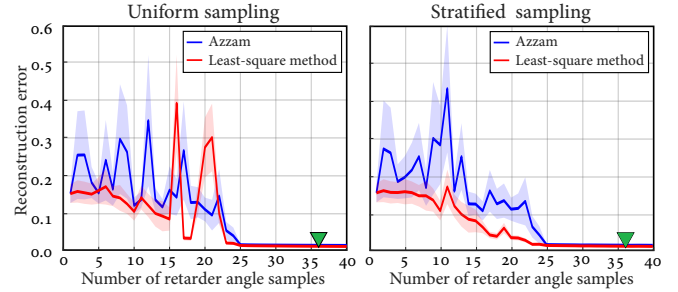


Fig. 14. Comparison of the least-square reconstruction approach used in this paper to Azzam’s frequency method [1978]. The plots show the reconstruction error (Frobenius norm divided by 16, the number of matrix elements) of 10,000 randomly simulated physically-valid Mueller matrices with additive Gaussian noise of standard deviation 0.01. The left plot shows the results for varying number of retarder samples with uniform sampling. While the least-square method mostly outperforms Azzam’s method, there are degraded cases at 16, 20, and 21 by falling into local minima in optimization. This problem can be resolved by adopting stratified sampling: jittering the uniformly-sampled retarder angles with the standard deviation 1° . The right plot shows that the least-square method with stratified sampling outperforms the Azzam’s method for all number of retarder angles. The light blue and red regions around the plotted curves indicate the standard deviation of reconstructed errors for the simulated Mueller matrices.

which is a rather naïve way of recovering this information (it is, e.g., unaware of the previously mentioned validity criterion). Similar to spherical linear interpolation (“slerp”) in the context of rotations, future work could lead to the development of more sophisticated polarimetric interpolants.

Polarimetric rendering. Our polarized rendering system currently importance samples roughly proportionally to the intensity of scattered light, using a density function that is reciprocal, enabling standard reverse path tracing starting at the camera. In the future, it could be interesting to develop more advanced conditional sampling schemes that refine the sampling pattern conditioned on the polarimetric properties of the current subpath.

Sampling resolution of retarder angles. In theory, even an extremely very coarse sampling suffices to accurately reconstruct Mueller matrix via our least-squares approach, assuming that there is no noise, and that the number of observations exceeds the number of free parameters. In practice, actual measurements include sensor noise and small instrument-specific biases, in which cases increasing the sampling rate of retarder angles improves the quality of the reconstruction. Figure 14 shows that a least-square approach to reconstruction in most generally provides more accurate Mueller matrices than Azzam’s method [1978], particularly when the sampling rate is low. For varying number of samples, we tested two different sampling schemes: uniform sampling and stratified sampling. The uniform sampling causes the least-square method to fail due to local minima at a specific number of angle samples, while stratified sampling could resolve this issue, resulting in a consistent performance gain for the least-square method over Azzam’s method [1978]. Our measurements use 36 retarder angle samples, which leads to an over-determined least squares problem that is

more resilient to noise (indicated by the green symbol in Figure 14).

10 CONCLUSION

We introduced an efficient acquisition technique to measure the large and high-dimensional space underlying multispectral polarized BRDFs, addressing issues with limited coverage in either spectral, angular, or polarimetric dimensions in previous datasets. Our method draws upon both spectroscopic ellipsometry and image-based BRDF acquisition to make such extensive measurements practical. This enabled us to create the first polarimetric BRDF dataset containing isotropic pBRDFs of 25 real-world objects. Our data-driven model can use these measurements to create renderings with an unprecedented level of realism in the polarimetric reflectance of materials. We also use the captured data to investigate polarization and material appearance and observe relationships that could be leveraged in future parametric pBRDFs with further analysis on the dataset. We expect that our work will enable future progress in areas including shape-from-polarization, appearance-from-polarization, and polarimetric rendering.

ACKNOWLEDGMENTS

Min H. Kim acknowledges Korea NRF grants (2019R1A2C3007229, 2013M3A6A6073718), Samsung Research, the CT Research & Development Program of KOCCA in MCST of Korea, MSIT/IITP of Korea, MSRA, and Cross-Ministry Giga KOREA (GK17P0200).

REFERENCES

- Michael Ashikhmin and Simon Premoze. 2007. Distribution-based BRDFs. *Technical Report, University of Utah* (2007).
- RMA Azzam. 1978. Photopolarimetric measurement of the Mueller matrix by Fourier analysis of a single detected signal. *Optics Letters* 2, 6 (1978), 148–150.
- RMA Azzam. 2016. Stokes-vector and Mueller-matrix polarimetry. *JOSA A* 33, 7 (2016), 1396–1408.
- Seung-Hwan Baek, Daniel S Jeon, Xin Tong, and Min H Kim. 2018. Simultaneous acquisition of polarimetric SVBRDF and normals. *ACM Trans. Graph.* 37, 6 (2018), 268–1.
- M. M. Bagher, C. Soler, and N. Holzschuch. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Comput. Graph. Forum* 31, 4 (2012), 1509–1518.
- Pierre Boher, Thierry Leroux, Ludvine Cave, Thibault Bignon, and Véronique Collom-Patton. 2017. Polarimetric multispectral bidirectional reflectance distribution function measurements using a Fourier transform instrument. *Electronic Imaging* 2017, 8 (2017), 19–24.
- Brent Burley. 2012. Physically-based shading at Disney. In *ACM SIGGRAPH 2012 Courses (SIGGRAPH '12)*. ACM, New York, NY, USA.
- FJJ Clarke and DJ Parry. 1985. Helmholtz reciprocity: Its validity and application to reflectometry. *Lighting Research & Technology* 17, 1 (1985), 1–11.
- Edward Collett. 2005. Field guide to polarization. SPIE Bellingham, WA.
- Robert L. Cook and Kenneth E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Transactions on Graphics (TOG)* 1, 1 (1982), 7–24. <https://doi.org/10.1145/357290.357293>
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, and Victor Ostromoukhov. 2015. Extracting microfacet-based brdf parameters from arbitrary materials with power iterations. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 21–30.
- Jonathan Dupuy and Wenzel Jakob. 2018. An adaptive parameterization for efficient material acquisition and rendering. *ACM Transactions on graphics (TOG)* 37, 6 (2018), 1–14.
- Hiroyuki Fujiwara. 2007. *Spectroscopic ellipsometry: principles and applications*. John Wiley & Sons.
- Clark R Givens and Alexander B Kostinski. 1993. A simple necessary and sufficient condition on physically realizable Mueller matrices. *Journal of Modern Optics* 40, 3 (1993), 471–481.
- Giuseppe Claudio Guarnera, Pieter Peers, Paul Debevec, and Abhijeet Ghosh. 2012. Estimating surface normals from spherical stokes reflectance fields. In *European Conference on Computer Vision*. Springer, 340–349.
- Cong Phuoc Huynh, Antonio Robles-Kelly, and Edwin R Hancock. 2013. Shape and refractive index from single-view spectro-polarimetric images. *International journal of computer vision* 101, 1 (2013), 64–94.
- MW Hyde IV, JD Schmidt, and MJ Havrilla. 2009. A geometrical optics polarimetric bidirectional reflectance distribution function for dielectric and metallic surfaces. *Optics express* 17, 24 (2009), 22138–22153.
- Adrian Jarabo and Victor Arellano. 2018. Bidirectional rendering of vector light transport. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 96–105.
- Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. 2015. Polarized 3D: High-Quality Depth Sensing with Polarization Cues. In *Proc. ICCV*. IEEE Computer Society, 3370–3378.
- Marc Leroy. 2001. Deviation from reciprocity in bidirectional reflectance. *Journal of Geophysical Research: Atmospheres* 106, D11 (2001), 11917–11923.
- Joakim Löw, Joel Kronander, Anders Ynnerman, and Jonas Unger. 2012. BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph.* 31, 1 (2012), 9:1–14.
- Stephen R Marschner, Stephen H Westin, Eric PF Lafortune, and Kenneth E Torrance. 2000. Image-based bidirectional reflectance distribution function measurement. *Applied Optics* 39, 16 (2000), 2592–2600.
- Wojciech Matusik. 2003. *A data-driven reflectance model*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- Michal Mojzík, Tomas Skrivan, Alexander Wilkie, and Jaroslav Krivanek. 2016. Bi-Directional Polarised Light Transport. In *Eurographics Symposium on Rendering*.
- Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H. Kim. 2018. Practical SVBRDF Acquisition of 3D Objects with Unstructured Flash Photography. *ACM Transactions on Graphics (TOG)* 37, 6 (2018).
- Giljoo Nam, Joo Ho Lee, Hongzhi Wu, Diego Gutierrez, and Min H. Kim. 2016. Simultaneous Acquisition of Microscale Reflectance and Normals. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2016)* 35, 6 (2016). <https://doi.org/10.1145/2980179.2980220>
- Addy Ngan, Frédo Durand, and Wojciech Matusik. 2005. Experimental Analysis of BRDF Models. In *Proceedings of the Eurographics Symposium on Rendering*. Eurographics Association, 117–226.
- Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: a retargetable forward and inverse renderer. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 203.
- Richard G Priest and Thomas A Gerner. 2000. *Polarimetric BRDF in the microfacet model: Theory and measurements*. Technical Report. NAVAL RESEARCH LAB WASHINGTON DC.
- Nicolas Riviere, Romain Ceoloto, and Laurent Hespel. 2012. Multispectral polarized BRDF: design of a highly resolved reflectometer and development of a data inversion method. *Optica applicata* 42, 1 (2012).
- Szymon Rusinkiewicz. 1998. A New Change of Variables for Efficient BRDF Representation. In *Proc. Eurographics Rendering Workshop*. 11–22.
- Vishwanath Saragadam and Aswin C Sankaranarayanan. 2019. KRISM—krylov subspace-based optical computing of hyperspectral images. *ACM Transactions on Graphics (TOG)* 38, 5 (2019), 1–14.
- Matthew H Smith. 2002. Optimization of a dual-rotating-retarder Mueller matrix polarimeter. *Applied optics* 41, 13 (2002), 2488–2493.
- Kenneth E Torrance and Ephraim M Sparrow. 1967. Theory for off-specular reflection from roughened surfaces. *JOSA* 57, 9 (1967), 1105–1112.
- Eric Veach. 1998. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph.D. Dissertation. Stanford, CA, USA. Advisor(s) Guibas, Leonidas J. AAI9837162.
- D Welles, M Serna, SH Sposato, MP Fetrow, KP Bishop, SA Arko, and TR Caudill. 2000. Spectral polarimetric BRDF model and comparison to measurements from isotropic roughened glass. In *Workshop on Multi/Hyperspectral Sensors, Measurements, Modeling and Simulation*. US Army Aviation and Missile Command Huntsville, AL.
- Alexander Wilkie and Andrea Weidlich. 2010. A standardised polarisation visualisation for images. In *Proceedings of the 26th Spring Conference on Computer Graphics*. ACM, 43–50.
- Alexander Wilkie and Andrea Weidlich. 2012. Polarised light in computer graphics. In *SIGGRAPH Asia 2012 Courses*. ACM, 8.
- Zhengyou Zhang. 2000. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence* 22, 11 (2000), 1330–1334.