

Hyperspectral Polarimetric BRDFs of Real-world Materials

YUNSEONG MOON, POSTECH, South Korea

RYOTA MAEDA, POSTECH, South Korea and University of Hyogo, Japan

SUHYUN SHIN, POSTECH, South Korea

INSEUNG HWANG, KAIST, South Korea

YOUNGCHAN KIM, POSTECH, South Korea

MIN H. KIM, KAIST, South Korea

SEUNG-HWAN BAEK, POSTECH, South Korea

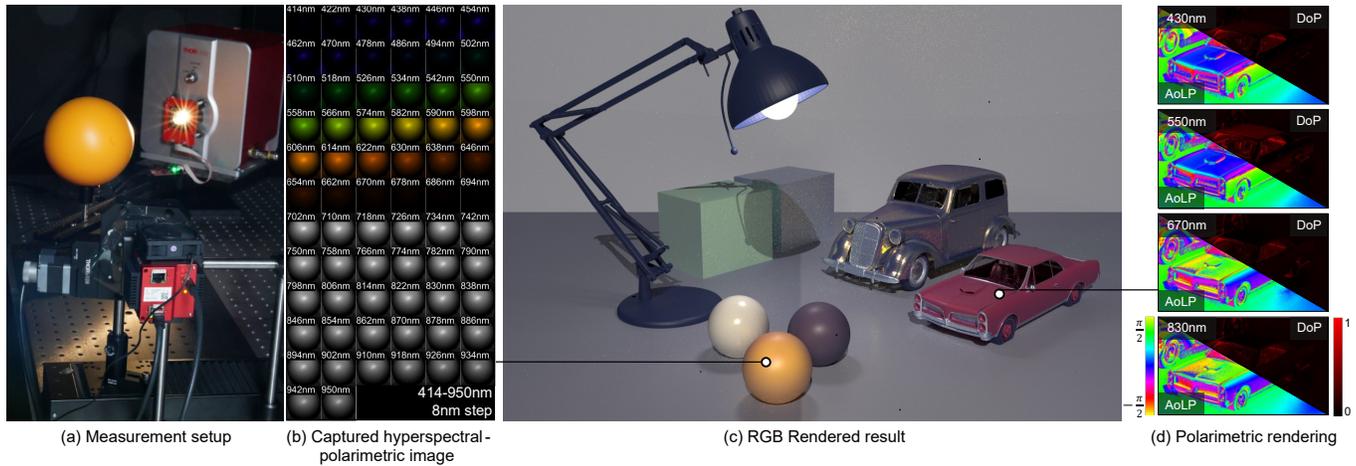


Fig. 1. We present the first dataset of hyperspectral polarimetric BRDFs of real-world materials acquired by (a) our imaging setup, providing full Mueller-matrix BRDF measurements at (b) 68 spectral channels from visible to near-infrared (414–950 nm). This enables rendering of not only (c) intensity but also (d) polarization states from visible to NIR spectral bands.

Acquiring bidirectional reflectance distribution functions (BRDFs) is essential for simulating light transport and analytically modeling material properties. Over the past two decades, numerous intensity-only BRDF datasets in the visible spectrum have been introduced, primarily for RGB image rendering applications. However, in scientific and engineering domains, there remains an unmet need to model light transport with polarization—a fundamental wave property of light—across hyperspectral bands. To address this gap, we present the first hyperspectral-polarimetric BRDF (hpBRDF) dataset of real-world materials, spanning wavelengths from 414 to 950 nm and densely sampled at 68 spectral bands. This dataset covers both the visible and near-infrared (NIR) spectra, enabling detailed material analysis and light reflection simulations that incorporate polarization at each narrow spectral band. We develop an efficient hpBRDF acquisition system that captures high-dimensional hpBRDFs within a feasible acquisition time. Using this system, we demonstrate hyperspectral-polarimetric rendering using the acquired hpBRDFs. To provide insights on hpBRDF, we analyze the hpBRDFs with respect to their dependencies on wavelength, polarization

state, material type, and illumination/viewing geometry. Also, we propose compact representations through principal component analysis and implicit neural hpBRDF modeling. Dataset is available on our project page¹.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling; Image-based rendering.**

Additional Key Words and Phrases: BRDF, polarization imaging, hyperspectral imaging, rendering

ACM Reference Format:

Yunseong Moon, Ryota Maeda, Suhyun Shin, Inseung Hwang, Youngchan Kim, Min H. Kim, and Seung-Hwan Baek. 2025. Hyperspectral Polarimetric BRDFs of Real-world Materials. 1, 1 (September 2025), 11 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Modeling how light reflects from materials is a fundamental problem across science and engineering, extending beyond computer graphics. The bidirectional reflectance distribution function (BRDF) models surface reflection as a 4D function over the 2D incident and 2D outgoing directions. BRDFs are central to physically based rendering for simulating light transport, and to inverse rendering for recovering geometry and material properties [Pharr et al. 2023].

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

© 2025 Association for Computing Machinery.

XXXX-XXXX/2025/9-ART \$15.00

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

¹<https://yunseong0518.github.io/projects/hpBRDF/>

Measured BRDF datasets have significantly advanced BRDF research over the past two decades. Early datasets focused on RGB intensity, facilitating realistic RGB rendering [Marschner et al. 2000; Matusik et al. 2003]. With the increased importance of RGB-near-infrared (NIR) imaging and display applications, Dupuy and Jakob [2018] expanded BRDF measurements from 360 nm to 1,000 nm, enabling rendering across visible and NIR spectra.

On the other hand, the emergence of polarimetric imaging has driven the development of polarimetric BRDF (pBRDF) datasets. Baek et al. [2020] introduced a pBRDF dataset at five visible wavelengths, enabling polarimetric rendering at the five-selected spectra. However, five discrete wavelengths are insufficient to capture the spectral variability of polarization even within the visible spectrum.

Spectral and polarimetric approaches are also interconnected. Polarimetric inverse rendering leverages wavelength-dependent polarization cues to enhance shape reconstruction and reflectance decomposition [Baek et al. 2018; Han et al. 2024; Hwang et al. 2022; Ichikawa et al. 2023; Kondo et al. 2020; Li et al. 2024]. Several studies have combined spectral-polarimetric modeling [Ha et al. 2024; Huynh et al. 2013; Kim et al. 2023].

Despite this progress, the lack of hyperspectral-polarimetric BRDF (hpBRDF) datasets limits the accuracy of modeling and simulation for applications that rely on spectral and polarimetric data.

To address this gap, we introduce a hpBRDF dataset, spanning 68 spectral bands from 414 nm to 950 nm, with a full width at half maximum (FWHM) of 10 nm and a step size of 8 nm. The dataset covers the full angular domain of surface reflection and is represented using Mueller matrices, enabling accurate modeling and simulation of arbitrary polarization states at each spectral band. This enables the rendering and analysis of complex wavelength-dependent polarization phenomena in real-world materials.

The primary challenge in acquiring an hpBRDF dataset is its high dimensionality, necessitating prolonged capture times using traditional scanning approaches. To efficiently acquire this high-dimensional data, we have developed an hpBRDF acquisition system integrating image-based BRDF acquisition [Matusik 2003], single-shot hyperspectral imaging, and broadband visible-to-NIR optical ellipsometry. Our system employs a hyperspectral light-field camera for single-shot visible-to-NIR hyperspectral imaging, significantly reducing acquisition times. We also implemented optical ellipsometry with accurate spectro-polarimetric calibration for broadband visible-to-NIR operation.

Using our newly acquired real-world hpBRDF dataset comprising 14 materials, we demonstrate hyperspectral-polarimetric rendering. Furthermore, we analyze the dataset to explore dependencies on wavelength, polarization state, material characteristics, and illumination/viewing geometry. We also perform principal component analysis (PCA) and introduce an implicit neural representation of the hpBRDF, providing compact data representations. The neural approach additionally supports interpolation and continuous spectral-angular modeling.

In summary, our key contributions are:

- A visible-to-NIR hpBRDF dataset for 14 real-world materials in Mueller-matrix form.

- An efficient hpBRDF acquisition system, including the acquisition procedure, calibration, and processing methods.
- An analysis of hpBRDF dependencies on wavelength, polarization state, and material characteristics, complemented by compact data representations using PCA and implicit neural representations.

2 RELATED WORK

2.1 Acquisition Systems

BRDF Acquisition. BRDF measurements involve capturing reflected light intensity across combinations of incident and outgoing directions. Various optomechanical systems have been developed to facilitate these measurements, each balancing trade-offs in speed, accuracy, and flexibility. The gonio-reflectometer is a classical device for BRDF measurement [White et al. 1998]. Conventional gonio-reflectometers measure light intensities in various light/view directions by rotating a light source and a sensor across a hemispherical range around a planar target. This approach has been extensively used for capturing BRDFs in both spectral [Dupuy and Jakob 2018; Filip and Vavra 2014; Li et al. 2006; Tsuchida et al. 2005] and polarization domains [Kondo et al. 2020]. However, achieving dense angular sampling with this approach is inherently time-consuming. To overcome this limitation, an image-based BRDF measurement method has been proposed [Marschner et al. 2000; Matusik 2003; Ngan et al. 2005]. This method captures convex objects, such as spheres or cylinders, instead of planar targets. A camera can capture dense angular combinations in a single shot, owing to the geometric variation of the objects, significantly reducing acquisition time. This strategy has been used for both spectral [Kim et al. 2010] and polarization [Baek et al. 2020] measurements. In this work, we also adopt an image-based approach, utilizing a sphere object to efficiently acquire hpBRDF within a practical time frame.

Spectral Acquisition. Early work on BRDF measurements often relied on a conventional RGB camera, which provided only limited spectral information. In contrast, multi- or hyperspectral imaging techniques acquire reflectance data in more than three bands. Spectrometers can measure thousands of spectral bands and have been integrated with gonio-reflectometers [Dupuy and Jakob 2018]. However, the spectrometer only measures a single point and is unsuitable for image-based BRDF acquisition. To achieve spectral sampling in an image-based setup, previous approaches have employed multiple band-pass filters mounted on a rotating stage [Baek et al. 2020; Li et al. 2006; Tsuchida et al. 2005] or a liquid crystal tunable filter [Kim et al. 2010]. However, these time-multiplexing approaches require capturing additional images for each spectral band, substantially increasing measurement time. In contrast, our approach utilizes a hyperspectral light field camera capable of capturing 68 spectral bands in a single shot. This reduces the total acquisition time, making hpBRDF measurement feasible.

Polarimetric Acquisition. Polarimetric reflectance can be described by a Mueller matrix, which characterizes how a surface changes the polarization state of incident light. Such measurements are commonly performed using an ellipsometer [Collett 2005], which systematically modulates the polarization states on both the light source

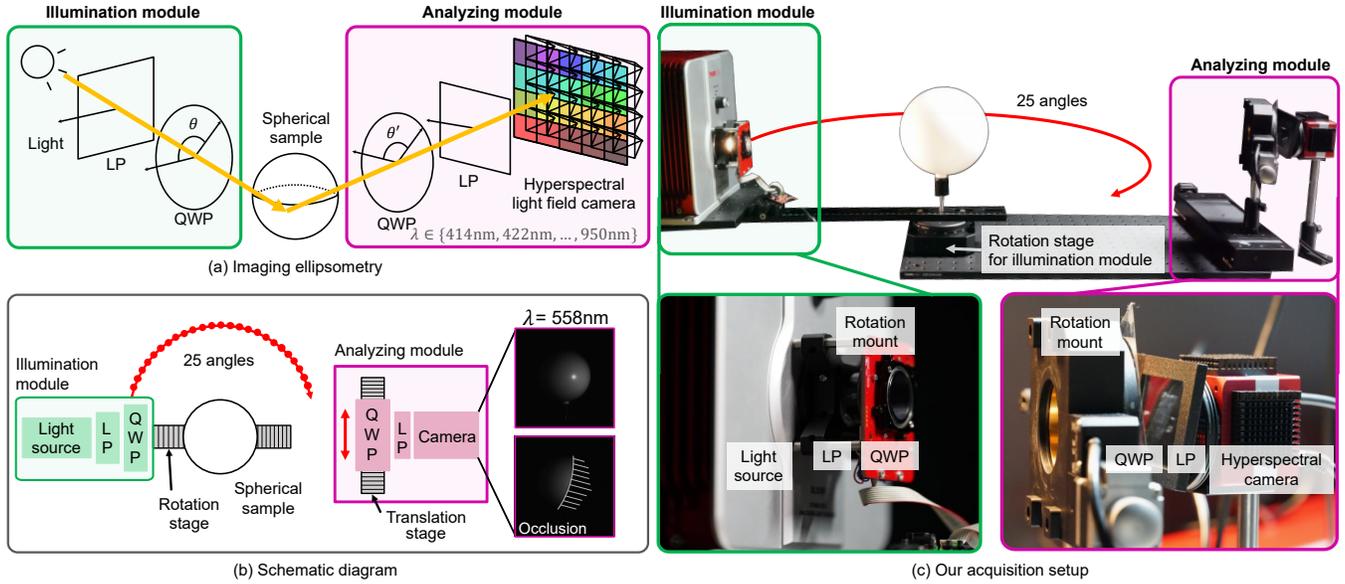


Fig. 2. **hpBRDF acquisition system.** (a) Our system consists of an illumination module and an analyzing module. The illumination module emits broadband polarized light, which is reflected at a sample sphere. The analyzing module captures the reflected light using a hyperspectral light field camera, which captures different spectral images from multiple viewpoints. Both the illumination and analyzing modules have achromatic LP and QWP to generate and analyze the polarization state. (b) The illumination module is rotated around the sample using a motorized rotation stage. To address occlusion, we translate the QWP and acquire two separate image sets. (c) Photographs of our acquisition setup.

and detector sides. Kondo et al. [2020] measured a pBRDF using rotating linear polarizers, thereby obtaining a 3×3 Mueller matrix limited to linear polarization components. Baek et al. [2020] constructed a pBRDF dataset that captures the full 4×4 Mueller matrix, encompassing both linear and circular polarization effects. Their measurement employed a dual-rotating-retarder (DRR) method [Azam 1978], which rotates quarter-wave plates (QWPs) to sample a comprehensive range of polarization states. Following this approach, our method also adopts the DRR technique to obtain a complete 4×4 Mueller matrix for visible-to-NIR wavelengths.

2.2 BRDF Datasets

The MERL dataset [Matusik 2003] and the UTIA dataset [Filip and Vávra 2014] each contain BRDF measurements of approximately one hundred materials, but only at RGB intensity. Dupuy and Jakob [2018] introduced a hyperspectral BRDF dataset spanning the visible to NIR domain. Baek et al. [2020] introduced a pBRDF dataset that includes polarization properties at five discrete wavelengths in the visible spectrum. Our hpBRDF dataset offers both hyperspectral and polarimetric measurements.

2.3 Applications of Spectral-polarimetric Imaging

Our hpBRDF dataset provides an opportunity for exploring the relationships between spectral and polarimetric visual information. Parametric pBRDF models [Baek et al. 2018; Ha et al. 2024; Hwang et al. 2022; Hyde IV et al. 2009; Ichikawa et al. 2023; Kondo et al. 2020; Priest and Germer 2000] are fundamental to not only forward rendering, but also inverse rendering for geometry and appearance estimation [Baek et al. 2018; Dave et al. 2022; Ha et al. 2024; Han et al.

2024; Hwang et al. 2022; Kondo et al. 2020; Li et al. 2024]. Yet, most existing pBRDF models neglect spectral dependencies on pBRDF. On the other hand, incorporating spectral-polarimetric dependencies has been found to be useful for accurately capturing various optical phenomena. For example, wavelength-dependent variations in the refractive index have been leveraged for shape recovery from a spectral-polarimetric image [Huynh et al. 2013] and for glass segmentation in RGB-polarization images [Mei et al. 2022]. Moreover, large-scale spectral-polarimetric datasets have begun to uncover complex interrelationships between spectral and polarization characteristics in natural images [Jeon et al. 2023]. Our hpBRDF dataset paves the way for developing more physically grounded hpBRDF models and spectro-polarimetric imaging applications.

3 hpBRDF IMAGING SYSTEM

To efficiently acquire hpBRDFs of real-world materials, we design a hyperspectral-polarimetric reflectance imaging system that combines spectral and polarization sensitivity with angular sampling efficiency. Figure 2 illustrates our hpBRDF imaging system. It captures the hpBRDF of a spherical object, represented as a Mueller matrix \mathbf{M} , within a practical acquisition time. To this end, we extend image-based BRDF imaging [Matusik 2003] with single-shot hyperspectral imaging and visible-to-NIR optical ellipsometry.

3.1 Polarized VIS-NIR Illumination

The illumination module of our imaging system uses a stabilized quartz tungsten-halogen lamp (*Thorlabs SLS301L*), which produces a consistent visible-to-NIR spectral output denoted by L^λ . Since the emitted light is unpolarized, its Stokes vector is: $s^\lambda = [L^\lambda, 0, 0, 0]^\top$.

The unpolarized light passes through two polarization-modulating components: an ultra-broadband linear polarizer (LP, *Thorlabs WP25M-UB*) characterized by Mueller matrix \mathbf{L} , and an achromatic quarter-wave plate (QWP, *Thorlabs AQWP10M-580*) mounted on a high-speed motorized rotation stage (*Thorlabs ELL14K*). By rotating the QWP to an angle θ , multiple polarization states can be generated. The resulting Stokes vector of the light exiting the illumination module at wavelength λ is given by:

$$\mathbf{s}_{\text{emitted}}^\lambda(\theta) = \mathbf{R}^\lambda(\theta)\mathbf{L}\mathbf{s}^\lambda, \quad (1)$$

where $\mathbf{R}^\lambda(\theta)$ is the Mueller matrix of the QWP rotated to angle θ . This configuration enables rapid and flexible control over the polarization state of broadband light across the visible to near-infrared spectrum. The entire illumination module is mounted on a heavy-duty motorized rotation stage (*Thorlabs HDR50/M*), which rotates around the target sample as shown in Figure 2(b).

3.2 Reflection on Spherical Samples

We capture homogeneous spherical samples as in previous image-based BRDF imaging [Matusik 2003] as shown in Figure 2(c). The continuous variation of surface normals across the sphere, combined with changing incident angles from the rotating light source, enables dense sampling over a wide angular domain. For a surface point on the sphere, the incident light direction is denoted by ω_i and the direction of the reflected light by ω_o . The hpBRDF at that point is defined as a Mueller matrix $\mathbf{M}(\omega_i, \omega_o)$, describing the polarization transformation of incident light. The reflected Stokes vector is:

$$\mathbf{s}_{\text{reflected}}^\lambda(\theta, \omega_o) = \mathbf{M}(\omega_i, \omega_o)\mathbf{s}_{\text{incident}}^\lambda(\theta, \omega_i). \quad (2)$$

As shown in Figure 3(c), we apply a coordinate transformation to align the emitted-light Stokes vector and the incident-light Stokes vector as $\mathbf{s}_{\text{incident}}^\lambda(\theta, \omega_i) = \mathbf{C}_{e \rightarrow i}\mathbf{s}_{\text{emitted}}^\lambda(\theta)$.

3.3 Hyperspectral Polarimetric Imaging

Visible-to-NIR Polarimetric Modulation. Our analyzing module measures the Stokes vector $\mathbf{s}_{\text{reflected}}^\lambda(\theta, \omega_o)$ of the reflected light, as shown in Figure 2(a). The light first passes through a 2-inch achromatic polymer QWP (*Edmund Optics WP140HE*), mounted on a motorized rotation stage (*Thorlabs HDR50/M*) and rotated to angle θ' . It then passes through a broadband 2-inch linear polarizer (*Thorlabs WP50L-UB*) mounted on a custom 3D-printed holder. To mitigate ghosting artifacts caused by the internal reflections between the QWP and the LP, both components are tilted by approximately 15° with respect to the optical axis. The selected tilt angles are within the specified operational range of each component. The resulting polarization-modulated Stokes vector is:

$$\mathbf{s}_{\text{analyzed}}^\lambda(\theta', \omega_o) = \mathbf{L}\mathbf{R}^\lambda(\theta')\mathbf{s}_{\text{captured}}^\lambda(\theta, \omega_o), \quad (3)$$

where, as shown in Figure 3(c), coordinate conversion is applied to the reflected-light Stokes vector: $\mathbf{s}_{\text{captured}}^\lambda(\theta, \omega_i) = \mathbf{C}_{r \rightarrow c}\mathbf{s}_{\text{reflected}}^\lambda(\theta)$.

Single-shot Hyperspectral Imaging. The light is then captured by a hyperspectral light-field camera (*Cubert Ultris X20*), which uses an array of narrow bandpass filters integrated into its microlens array to enable single-shot acquisition of spectral data at each different viewpoint, as shown in Figure 2(a). This configuration minimizes

acquisition time—critical for practical hpBRDF measurement. The hyperspectral camera provides a spatial resolution of 410×410 pixels with 164 spectral bands by default. We select 68 channels at 8 nm intervals from 414 nm to 950 nm to balance spectral resolution and data size, considering FWHM of the camera is 10 nm. The recorded intensity at wavelength λ under the QWP angles of the illumination module and analyzing module θ and θ' , respectively, is:

$$f(\lambda, \theta, \theta') = [\mathbf{s}_{\text{analyzed}}^\lambda(\theta', \omega_o)]_0, \quad (4)$$

where $[\]_0$ extracts the first entry of the Stokes vector, representing the intensity of light. Note that the outgoing direction ω_o varies per wavelength due to the multi-view geometry of the light-field camera. Because each spectral image corresponds to a micro-lens view direction, spatial misalignment between spectral bands is inherent. This necessitates geometric calibration across the spectral bands to correctly associate spatial and angular information in the hyperspectral-polarimetric image domain.

Occlusion Handling. The 2-inch QWP in the analyzer module limits the effective field of view of the hyperspectral light-field camera. To overcome this, we perform two acquisitions with the QWP positioned at distinct lateral offsets, using a linear translation stage (*Thorlabs NRT150*) as shown in Figure 2(b). In the figure, we show the raw captured image for two different QWP positions for wavelength 558 nm. For each spectral channel, we select an occlusion-free mask to ensure complete Mueller matrix reconstruction.

3.4 Calibration

We perform geometric, radiometric, and polarimetric calibration of our hpBRDF imaging system. Geometrically, our hyperspectral light-field camera is modeled as an array of 68 multi-view cameras, each capturing distinct narrow spectral bands. We calibrate their intrinsic and extrinsic parameters using multiple checkerboard captures. We also estimate the positions of the sample sphere via sphere fitting, and accurately determine light-source positions by minimizing the differences between simulated and captured specular highlights. Radiometrically, we calibrate the camera response by imaging a Spectralon reference with known diffuse reflectance across visible-to-NIR wavelengths, facilitating stable and linear HDR hyperspectral imaging. Polarimetrically, we precisely align LPs and QWPs using a polarimeter, allowing controlled polarization states through rotations of QWPs in both illumination and analyzing modules. For more details, please refer to the Supplementary Document.

4 hpBRDF DATASET

Using our hpBRDF imaging system, we have acquired the hpBRDFs of 14 real-world materials as shown in Figure 4. We represent hpBRDF \mathbf{M} in a 4D tabulation form: $\mathbf{P}(\lambda, \phi_d, \theta_d, \theta_h) \in \mathbb{R}^{4 \times 4}$, where we use Rusinkiewicz coordinate [1998] of $\phi_d, \theta_d, \theta_h$ for the incoming and outgoing light directions discretized into 361, 91 and 91 bins respectively. For the spectral axis, the hpBRDF table has 68 bins from 414 nm to 950 nm with 8 nm interval. For each table element at the index of $\lambda, \phi_d, \theta_d, \theta_h$, the table stores a 4×4 Mueller matrix in a single-precision float, leading to 13 GB per single hpBRDF table.

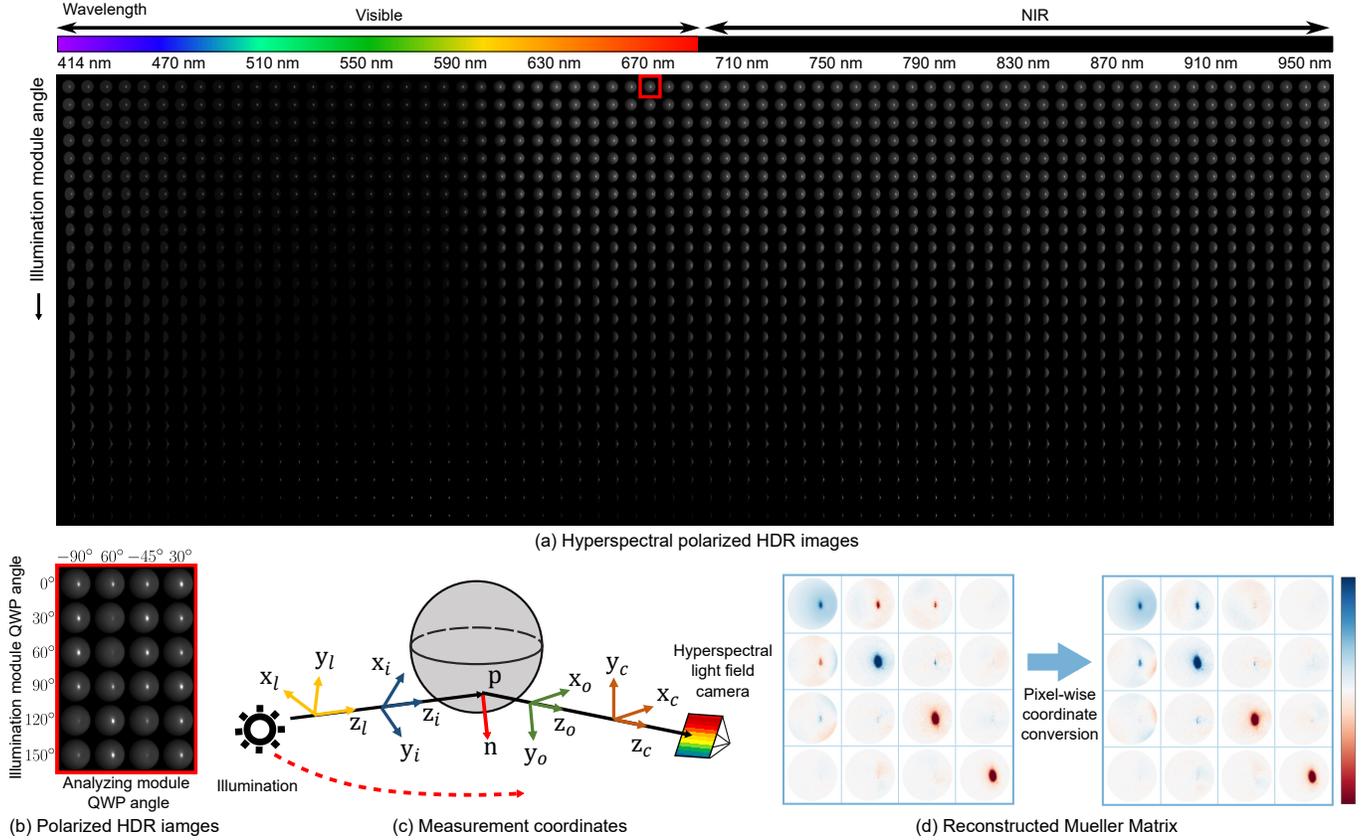


Fig. 3. **hpBRDF acquisition.** (a) Hyperspectral HDR images captured under different rotation angles of the illumination module. (b) Captured images with varying QWP angles corresponding to a specific configuration marked by the red rectangle in (a). (c) Given the coordinate systems of the illumination and analyzing modules, (x_l, y_l, z_l) , (x_c, y_c, z_c) , respectively, the hpBRDF coordinate is defined with respect to (x_i, y_i, z_i) and (x_o, y_o, z_o) that dependent on the TBN (tangent, bitangent, normal) space of the surface. (d) The reconstructed Mueller matrix is transformed to the hpBRDF coordinate.

Index	Material	
1	Aluminium	Metallic
2	SUJ2	
3	Fake gold	
4	Black glass	
5	White billiard	
6	White smooth plastic	Dielectric
7	White rough plastic	
8	Black rough plastic	
9	Silver rough plastic	
10	Red rough plastic	
11	Yellow rough plastic	
12	Plum rough plastic	
13	Green rough silicone	
14	Gray rough silicone	

Fig. 4. **Material samples.** Our dataset provides the measured hpBRDFs of 14 real-world spherical samples.

Image Acquisition. To acquire hpBRDF table, we first capture images covering all combinations of system parameters, including illumination-module angle, analyzing-module QWP translation,

QWP-angle configurations, and exposure levels. Specifically, for angular sampling, the illumination module is rotated through 25 discrete angles, ranging from 40° to 160° at intervals of 5° . At each illumination angle, two positions of the analyzing-module QWP are used using a translation stage to mitigate occlusion caused by the rotation mount, as shown in Figure 2. Ellipsometric measurements require rotating QWPs in both the illumination and analyzing modules to predefined angular configurations: specifically, angles $\Theta \in \{30^\circ, -45^\circ, 60^\circ, -90^\circ\}$ for the illumination module, and angles $\Theta' \in \{0^\circ, 30^\circ, 60^\circ, 90^\circ, 120^\circ, 150^\circ\}$ for the analyzing module. For each QWP and illumination configuration, we capture an HDR image by merging multiple exposure-bracketed images. The number of exposures ranges from 3 to 11, depending on the reflectance properties of the material. Consequently, the total number of hyperspectral images acquired per material ranges between 3,744 and 12,480. The total acquisition time per material ranges from 13 to 35 hours.

hpBRDF Reconstruction. We then reconstruct hpBRDF M for each pixel, illumination angle, and wavelength channel. Using the image formation models of Equations (1), (2), (3), (4), this is formulated as

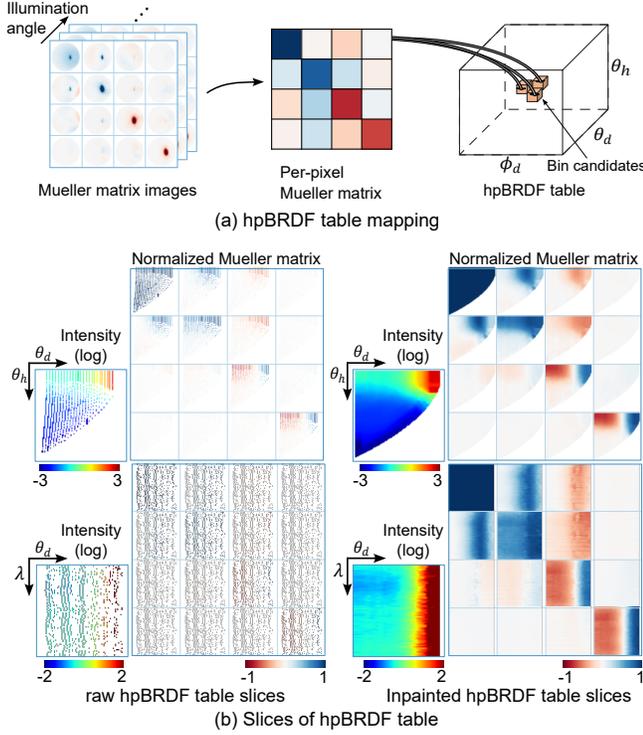


Fig. 5. **hpBRDF table construction** (a) We map per-pixel Mueller matrices to the neighbour bins in the hpBRDF table. (b) Representative slices of the hpBRDF table for raw and inpainted versions.

a least-squares problem:

$$\underset{\theta, \theta'}{\text{minimize}} \sum_M \left(f(\lambda, \theta, \theta') - [\mathbf{L}\mathbf{R}^\lambda(\theta')\mathbf{C}_{r \rightarrow c}\mathbf{M}\mathbf{C}_{e \rightarrow i}\mathbf{R}^\lambda(\theta)\mathbf{L}\mathbf{s}^\lambda]_0 \right)^2. \quad (5)$$

Figure 3(d) shows the reconstructed Mueller matrix image.

hpBRDF Table Construction. To construct the hpBRDF table, we calculate per-pixel Rusinkiewicz coordinates for each Mueller matrix. As shown in Figure 5(a), we map a single pixel’s Mueller matrix to the target hpBRDF table bin including the neighbor bins. Due to mechanical constraints in our imaging system, some entries are missing in the hpBRDF table especially for low θ_d . We perform numerical inpainting of the missing entries. We use 3D Gaussian convolution in the $(\phi_d, \theta_d, \theta_h)$ space for each Mueller matrix element at each wavelength for inpainting. Figure 5(b) shows the slices of the constructed hpBRDF table before and after applying the inpainting. For constructing the hpBRDF table, it takes about 10 hours per material.

5 RESULTS

5.1 Validation

We perform three validation experiments. First, we verify the accuracy of our reconstructed Mueller matrices by comparing them with ground-truth measurements for air, linear polarizer, and quarter-wave plate as shown in Figure 10(a). Results demonstrate close alignment with the ground truth. Second, we assess the physical

validity of the reconstructed Mueller matrices. A physically valid Mueller matrix maps any physically admissible Stokes vector to another admissible Stokes vector. A Stokes vector $\mathbf{s} = [s_0, s_1, s_2, s_3]^T$ is considered admissible if it satisfies: $s_0 \geq 0$, $s_0^2 - (s_1^2 + s_2^2 + s_3^2) \geq 0$. We use Givens–Kostinski’s method [1993] to verify this condition for three selected materials from our raw hpBRDF dataset. The results indicate that 94.50% of the reconstructed Mueller matrices satisfy the physical validity criteria across all hpBRDF bins. A complete analysis of the physical validity results for all 14 materials are provided in the supplementary document. Third, we qualitatively evaluate the fidelity of our measurements by comparing the hpBRDF data with similar visual appearances named “white billiard” provided by Baek et al. [2020]. Figure 10(b) shows qualitative agreement between our dataset and the reference measurements at the five wavelengths.

5.2 Rendering Polarimetric Visible-NIR Images

Using our hpBRDF dataset on Mitsuba 3 [Jakob et al. 2022] enables the simulation of polarized light transport for arbitrary scenes across the visible-NIR spectrum. Since our hpBRDF table follows the same format as Baek’s pBRDF table [2020], this allows for rendering with minimal code modifications. Figure 6 shows an example of a nighttime driving scenario, exhibiting different polarization characteristics of the RGB and NIR spectra due to the use of polarized NIR illumination. Figure 1 shows another example of hyperspectral polarimetric rendering.

5.3 hpBRDF Analysis

To analyze the polarization characteristics of our hpBRDF, we employ the Lu-Chipman decomposition [1996], which factorizes each Mueller matrix into three physically interpretable matrices: a depolarization matrix, a retardance matrix, and a diattenuation matrix. We can also derive four scalar quantities: diattenuation, polarizance, polarization preservation, and retardance. For detailed mathematical formulations and definitions, refer to the Supplementary document.

Wavelength Dependency. We analyze the wavelength-dependent behavior of hpBRDF. Figure 7 shows diattenuation and preservation images of three plastic spheres with distinct colors: white, red, and yellow. The white plastic sphere exhibits uniformly high spectral intensity across the whole wavelength range, with minimal variation. In contrast, the red and yellow plastic spheres show pronounced spectral intensity changes within the visible wavelength range. Notably, preservation tends to increase at wavelengths with lower reflectance, whereas diattenuation remains relatively stable despite changes in spectral intensity.

Surface Roughness. Figure 12 shows that smooth surfaces exhibit significantly stronger polarization-preserving behavior than rough surfaces consistently over the entire measured spectrum. This indicates that the polarization information of illumination is not preserved for rough surfaces, regardless of the light spectrum, potentially due to multiple scattering between microfacets.

Dielectric vs. Metallic. Figure 13 shows the comparison of the dielectric and metallic material using retardance Mueller matrix images \mathbf{M}_R . While both materials exhibit a similar diagonal component,

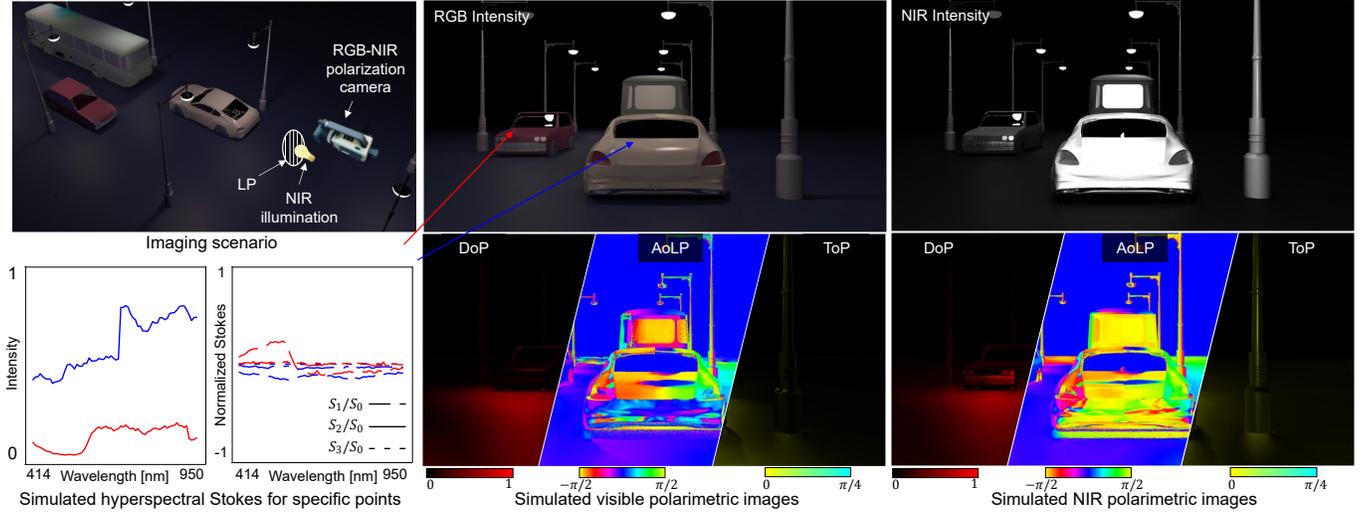


Fig. 6. **Hyperspectral polarimetric rendering.** We simulated a driving scenario under nighttime illumination using our hpBRDF dataset. The scene includes multiple vehicles and halogen-based street lights. We use the imaging setup with an RGB-NIR polarization camera and an active linearly polarized NIR LED illumination, showing different polarization characteristics for the visible and NIR spectrum.

clear differences appear in non-diagonal components, particularly $M_{R,13}$, $M_{R,23}$, $M_{R,31}$, $M_{R,32}$. These components show sign inversions between metallic and dielectric materials, indicating distinct phase shift behaviors in reflection since the retardance matrix represents the phase changes.

5.4 hpBRDF Representations

Implicit Neural Representation. We introduce an implicit neural representation for hpBRDF that enables continuous interpolation and extremely compact storage of hpBRDF table. Whereas recent neural BRDF models map incident and outgoing directions to RGB radiance [Dou et al. 2024; Rainer et al. 2019; Sztajman et al. 2021; Zeltner* et al. 2024], our formulation extends the input domain with wavelength λ [Jeon et al. 2023; Kim et al. 2023] and predicts the full 4×4 Mueller matrix, thereby jointly modeling spectral and polarization variation. Figure 9 shows the trained result with (a) different network sizes and (b) visualization of the trained hpBRDF. The predicted Mueller matrices closely resemble the original inpainted table, well representing the whole hpBRDF. In the rendering with a table, nearest-neighbor lookups introduce discretization artifacts, resulting in noticeable discontinuities at specular highlights. In contrast, the neural hpBRDF provides a continuous prediction of the Mueller matrix across both directional and spectral domains, enabling smooth rendering. Furthermore, the neural hpBRDF achieves a substantial reduction in storage size. For instance, a network with four hidden layers of 256 neurons in a single-precision floating-point format occupies 146 kB—about 10^5 times smaller than the original 13 GB tabulation—while preserving visual fidelity.

Principal Component Analysis. Inspired by previous works [Baek and Heide 2021; Matusik 2003; Ngan et al. 2006; Nielsen et al. 2015], we analyze low-rank structures of hpBRDF using PCA over spectral-polarimetric-angular domains. Figure 11(a) shows the dimensionality of the normalized hpBRDF. We also visualize the first two

principal components across selected angular and spectral slices, highlighting characteristic patterns in the data in Figure 11(b). These visualizations reveal distinct structural patterns within each slice, reflecting how the data varies across angular and spectral dimensions. Notably, the slices along the angular dimensions, particularly (θ_d, θ_h) , exhibit stronger and more complex variations, indicating a higher degree of angular dependency. In contrast, the variations in the spectral domain (θ_h, λ) appear more gradual and less spatially structured, showing that the spectral properties are relatively smoother compared to the angular variations.

Analytical Models. We evaluate an analytical parametric pBRDF model [Hwang et al. 2022] on our hpBRDF data. Figure 8 shows the fitted result on the yellow rough plastic material. While the rendered sRGB intensity appears visually similar to our measurement, notable differences are observed in the Lu-Chipman decomposition results. In the depolarization matrix M_Δ , the model overestimates near the edge of the sphere. Moreover, the model fails to reconstruct the non-diagonal components in the retardance matrix M_R , which we identified in Section 5.3 as important indicators of material metallicity. These components are either underestimated or absent in the fitted result.

6 DISCUSSION

Image resolution. Since hyperspectral light field cameras spatially divide a single image sensor to capture images at multiple wavelengths, the resulting spatial resolution per wavelength is limited to 410×410 pixels, which is lower than that of conventional image sensors. This resolution limits the dense angular sampling on a target sphere surface.

Tilted QWP. To mitigate inter-reflections between the QWP and the LP, we slightly tilted the QWP mounted on the camera side.

We carefully selected the tilt angle to minimize reflections while ensuring that no perceptible optical artifacts were introduced.

Inpainting. To address the missing Mueller matrices due to inevitable geometric constraints of our imaging setup, we interpolated the missing entries by finding the smoothest values in the angular domain. Although this inpainting produces visually plausible results, it does not guarantee physical validity. To improve the overall quality of the hpBRDF table, future work could incorporate physics-informed or data-driven methods that jointly model both the angular and polarization domains. Such approaches potentially produce more accurate and physically consistent estimations of the missing Mueller matrix entries.

7 CONCLUSION

In this work, we present the first hpBRDF dataset of real-world materials, leveraging our reflectance acquisition system. This dataset addresses the longstanding gap in modeling spectral and polarimetric BRDF across the visible to NIR wavelengths. It enables rendering of spectro-polarimetric images, facilitating realistic simulation of diverse imaging scenarios.

We further analyze how hpBRDF properties vary with wavelength, polarization state, material type, and geometric conditions, and propose compact representations through PCA and implicit neural representation. We anticipate that our hpBRDF dataset will contribute to spectro-polarimetric imaging and display research and inverse rendering that exploit combined spectral and polarimetric features.

Limitations. Although our acquisition setup benefits from single-shot hyperspectral imaging, it also faces limitations such as reduced spatial resolution per wavelength due to the sensor’s division into multiple spectral channels. Also, anisotropic samples cannot be captured in our imaging system. To handle missing Mueller matrix entries arising from imaging-system geometry, we employ inpainting, which, although plausible, may not exactly match the physical ground truth. Overcoming the aforementioned challenges will be an interesting area for future work.

ACKNOWLEDGMENTS

This work was partly supported by the National Research Foundation (NRF) grant (RS-2024-00438532, RS-2023-00211658, RS-2024-00357548), and the Institute of Information & Communications Technology Planning & Evaluation (IITP)-ITRC (Information Technology Research Center) grant (IITP-2025-RS-2024-00437866), funded by the Korean government (MSIT).

REFERENCES

RMA Azzam. 1978. Photopolarimetric measurement of the Mueller matrix by Fourier analysis of a single detected signal. *Optics Letters* 2, 6 (1978), 148–150.

Seung-Hwan Baek and Felix Heide. 2021. Polarimetric spatio-temporal light transport probing. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–18.

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.

Seung-Hwan Baek, Tizian Zeltner, Hyunjin 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 (2020), 139.

E. Collett. 2005. *Field Guide to Polarization*. SPIE Press. <https://books.google.co.kr/books?id=5JjwcSLbLsC>

Akshat Dave, Yongyi Zhao, and Ashok Veeraraghavan. 2022. Pandora: Polarization-aided neural decomposition of radiance. In *European Conference on Computer Vision*. Springer, 538–556.

Yishun Dou, Zhong Zheng, Qiaoqiao Jin, Bingbing Ni, Yugang Chen, and Junxiang Ke. 2024. Real-Time Neural BRDF with Spherically Distributed Primitives. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4337–4346.

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.

Jiri Filip and Radomír Vávra. 2014. Template-based sampling of anisotropic BRDFs. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 91–99.

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.

Hyunho Ha, Inseung Hwang, Nestor Monzon, Jaemin Cho, Donggun Kim, Seung-Hwan Baek, Adolfo Muñoz, Diego Gutierrez, and Min H Kim. 2024. Polarimetric BSSRDF Acquisition of Dynamic Faces. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–11.

Yufei Han, Heng Guo, Koki Fukai, Hiroaki Santo, Boxin Shi, Fumio Okura, Zhanyu Ma, and Yungpeng Jia. 2024. NeRSP: Neural 3D Reconstruction for Reflective Objects with Sparse Polarized Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11821–11830.

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.

Inseung Hwang, Daniel S Jeon, Adolfo Munoz, Diego Gutierrez, Xin Tong, and Min H Kim. 2022. Sparse ellipsometry: portable acquisition of polarimetric SVBRDF and shape with unstructured flash photography. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–14.

Milo W Hyde IV, Jason D Schmidt, and Michael J Havrilla. 2009. A geometrical optics polarimetric bidirectional reflectance distribution function for dielectric and metallic surfaces. *Optics express* 17, 24 (2009), 22138–22153.

Tomoki Ichikawa, Yoshiki Fukao, Shohei Nobuhara, and Ko Nishino. 2023. Fresnel Microfacet BRDF: Unification of Polari-Radiometric Surface-Body Reflection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16489–16497.

Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. Mitsuba 3 renderer. (2022). <https://mitsuba-renderer.org>.

Yujin Jeon, Eunsue Choi, Youngchan Kim, Yunseong Moon, Khalid Omer, Felix Heide, and Seung-Hwan Baek. 2023. Spectral and Polarization Vision: Spectro-polarimetric Real-world Dataset. *arXiv preprint arXiv:2311.17396* (2023).

Duck Bong Kim, Myoung Kook Seo, Kang Yeon Kim, and Kwan H Lee. 2010. Acquisition and representation of pearlescent paints using an image-based gonioscophotometer. *Optical engineering* 49, 4 (2010), 043604–043604.

Youngchan Kim, Wonjoon Jin, Sunghyun Cho, and Seung-Hwan Baek. 2023. Neural spectro-polarimetric fields. In *SIGGRAPH Asia 2023 Conference Papers*. 1–11.

Yuhi Kondo, Taishi Ono, Legong Sun, Yasutaka Hirasawa, and Jun Murayama. 2020. Accurate polarimetric BRDF for real polarization scene rendering. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX* 16. Springer, 220–236.

Chenhao Li, Taishi Ono, Takeshi Uemori, Hajime Mihara, Alexander Gatto, Hajime Nagahara, and Yusuke Moriuchi. 2024. NeISF: Neural Incident Stokes Field for Geometry and Material Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21434–21445.

Hongsong Li, Sing-Choong Foo, Kenneth E Torrance, and Stephen H Westin. 2006. Automated three-axis gonioreflectometer for computer graphics applications. *Optical Engineering* 45, 4 (2006), 043605–043605.

Shih-Yau Lu and Russell A Chipman. 1996. Interpretation of Mueller matrices based on polar decomposition. *JOSA A* 13, 5 (1996), 1106–1113.

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.

Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2003. A Data-Driven Reflectance Model. *ACM Transactions on Graphics* 22, 3 (July 2003), 759–769.

Haiyang Mei, Bo Dong, Wen Dong, Jiayi Yang, Seung-Hwan Baek, Felix Heide, Pieter Peers, Xiaopeng Wei, and Xin Yang. 2022. Glass segmentation using intensity and spectral polarization cues. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12622–12631.

Addy Ngan, Frédo Durand, and Wojciech Matusik. 2005. Experimental Analysis of BRDF Models. *Rendering Techniques* 2005, 16th (2005), 2.

Addy Ngan, Frédo Durand, and Wojciech Matusik. 2006. Image-driven Navigation of Analytical BRDF Models. *Rendering Techniques* 2006 (2006), 399–407.

- Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–11.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation*. MIT Press.
- Richard G Priest and Thomas A Germer. 2000. Polarimetric BRDF in the microfacet model: Theory and measurements. In *Proceedings of the 2000 Meeting of the Military Sensing Symposia Specialty Group on Passive Sensors*, Vol. 1. Infrared Information Analysis Center, 169–181.
- Gilles Rainer, Wenzel Jakob, Abhijeet Ghosh, and Tim Weyrich. 2019. Neural BTF compression and interpolation. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 235–244.
- Szymon M Rusinkiewicz. 1998. A new change of variables for efficient BRDF representation. In *Rendering Techniques' 98: Proceedings of the Eurographics Workshop in Vienna, Austria, June 29–July 1, 1998* 9. Springer, 11–22.
- Alejandro Sztrajman, Gilles Rainer, Tobias Ritschel, and Tim Weyrich. 2021. Neural brdf representation and importance sampling. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 332–346.
- Masaru Tsuchida, Hiroyuki Arai, Masami Nishiko, Yoshiyuki Sakaguchi, T Uchiyama, Masahiro Yamaguchi, Hideaki Haneishi, and Nagaaki Ohyama. 2005. Development of BRDF and BTF measurement and computer-aided design systems based on multispectral imaging. *Proc. AIC Colour* 5 (2005), 129–132.
- D Rod White, Peter Saunders, Stuart J Bonsey, John van de Ven, and Hamish Edgar. 1998. Reflectometer for measuring the bidirectional reflectance of rough surfaces. *Applied optics* 37, 16 (1998), 3450–3454.
- Tizian Zeltner*, Fabrice Rousselle*, Andrea Weidlich*, Petrik Clarberg*, Jan Novák*, Benedikt Bitterli*, Alex Evans, Tomáš Davidovič, Simon Kallweit, and Aaron Lefohn. 2024. Real-time neural appearance models. *ACM Transactions on Graphics* 43, 3 (2024), 1–17.

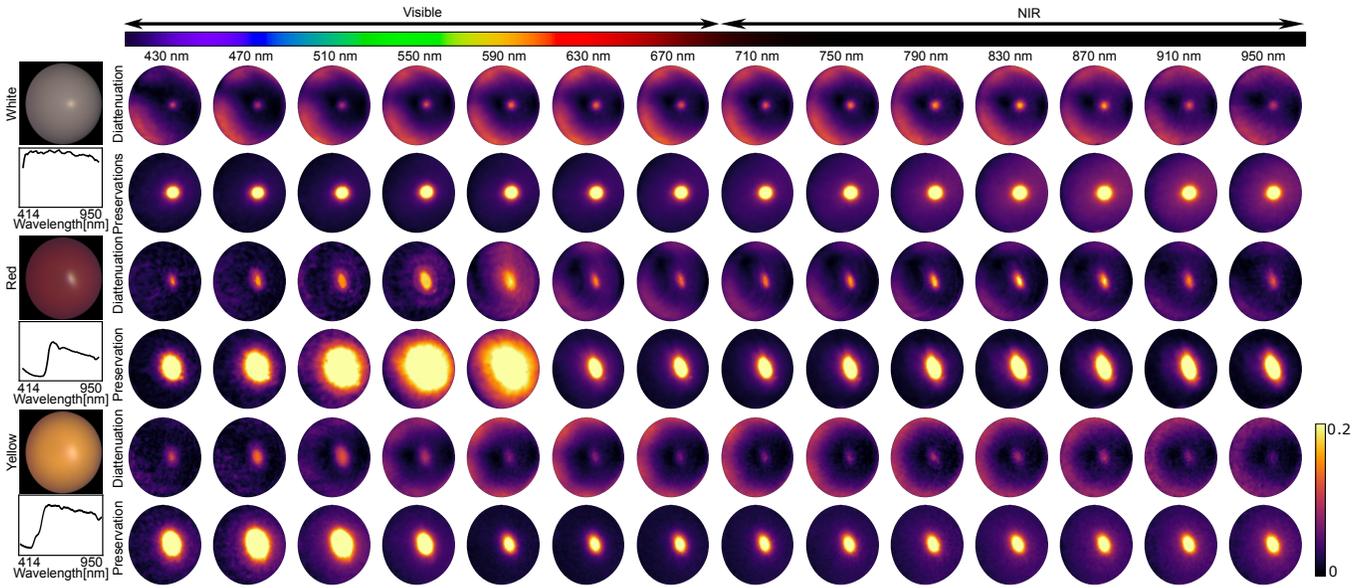


Fig. 7. **Wavelength dependency of hpBRDF.** We visualize three colored rough plastic samples (white, red, yellow) and their diattenuation and polarization preservation across wavelengths. The result shows wavelength-dependent characteristics for each material.

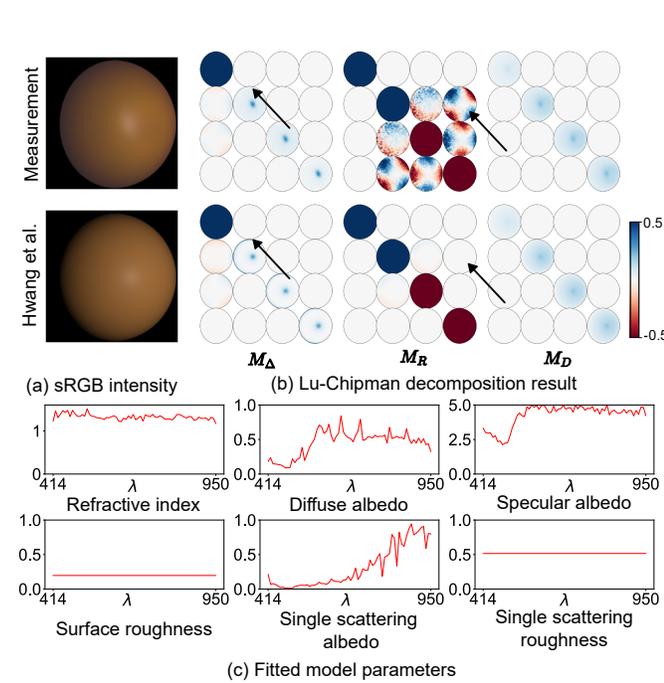


Fig. 8. **Evaluation of analytical pBRDF model on our hpBRDF.** We fit analytic parametric pBRDF model [Hwang et al. 2022] to our yellow rough plastic hpBRDF table. We show the (a) rendered sRGB intensity and (b) Lu-Chipman decomposition result of specific wavelength (550 nm) using our hpBRDF table and the fitted model. In the decomposition matrix M_{Δ} , we observe overestimation near the edges of the sphere. In the retardance matrix M_R , the analytical model fails to represent the non-diagonal component patterns. (c) Fitted model parameters.

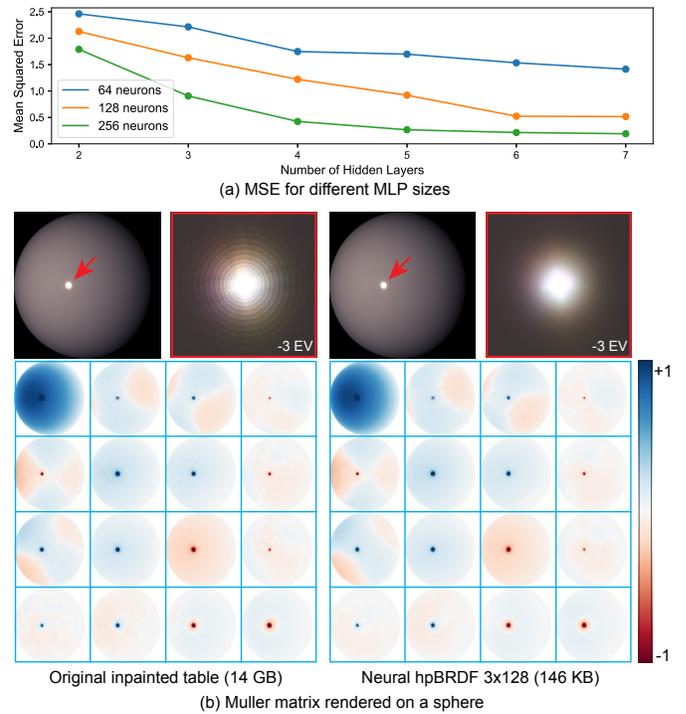


Fig. 9. **Implicit neural representation for hpBRDF.** (a) Mean squared error (MSE) of the predicted Mueller matrices for "white smooth plastic", evaluated across different MLP configurations. (b) Visualization of the original inpainted table and the neural hpBRDF corresponds to incident angles of 30° and a wavelength of $\lambda = 670$ nm. The neural hpBRDF closely matches the original Mueller matrix with continuous interpolation, while significantly reducing storage requirements. The rendered results are visualized with gamma correction ($\gamma = 2.2$).

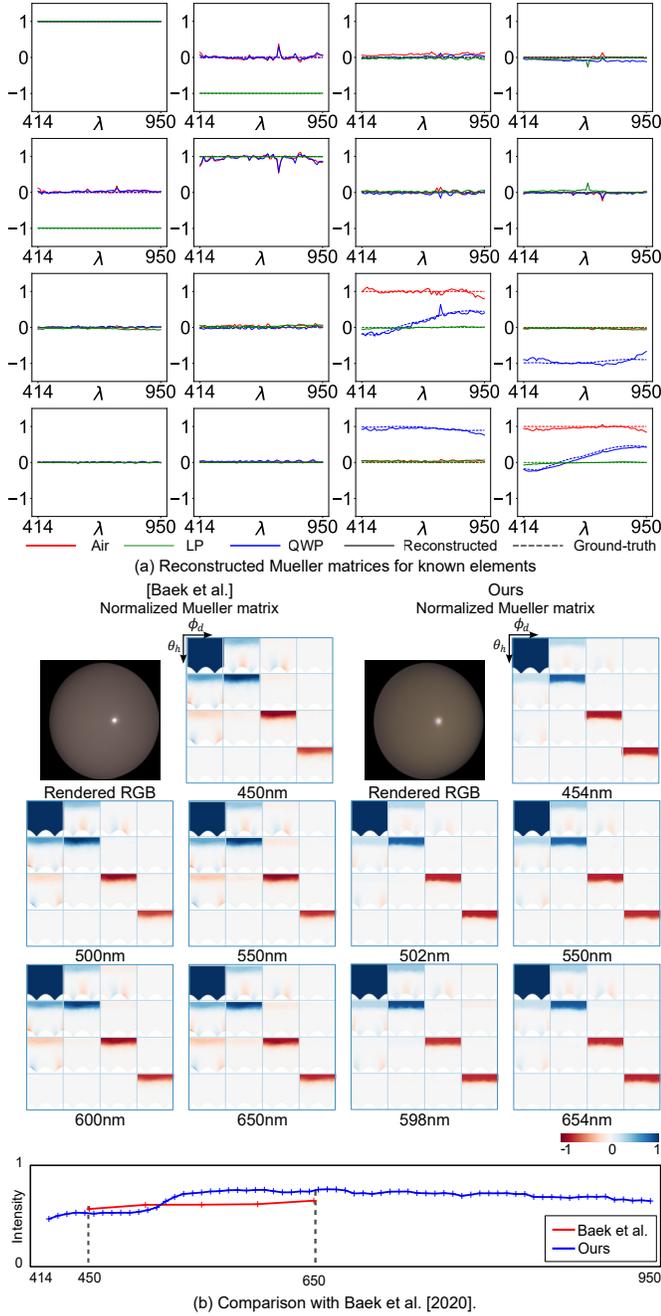


Fig. 10. **Validation.** (a) We validate our Mueller matrix reconstruction pipeline by capturing elements with known Mueller matrices. For QWP, the captured QWP has wavelength dependent retardance, so we calculated ground truth with known retardance. Reconstructed Mueller matrices are closely similar to the ground truth across the entire wavelength range. (b) We compare our hpBRDF table with the existing pBRDF table [2020], using the “white billiard” where it has a similar visual appearance. The pBRDF tables from both datasets show similar Mueller matrix, qualitatively validating the accuracy of our measurements. Note, our dataset offers denser and broader spectral sampling, in contrast to the five discrete wavelengths in [Baek et al. 2020].

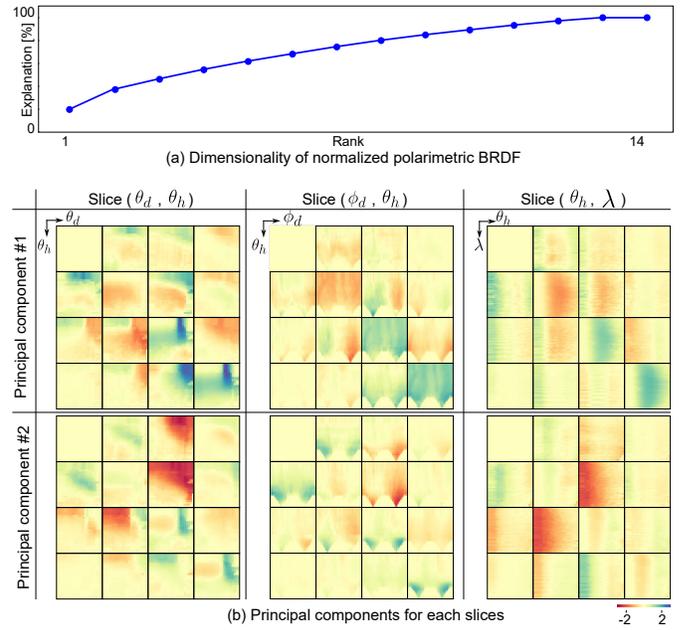


Fig. 11. **PCA analysis of hpBRDFs.** (a) The dimensionality of hpBRDFs using PCA. (b) Estimated first two principal components.

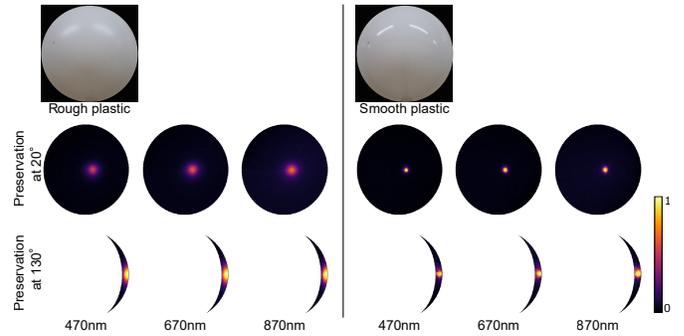


Fig. 12. **Surface roughness and hpBRDF.** For smooth material, the preservation value remains consistently close to 1 in specular highlight region, regardless of θ_d . In contrast, the preservation value of rough material shows a sharp transition around a specific θ_d in specular highlight region. At high θ_d values, preservation value approaches 1, similar to smooth material, while at low θ_d , it stabilizes a lower preservation value.

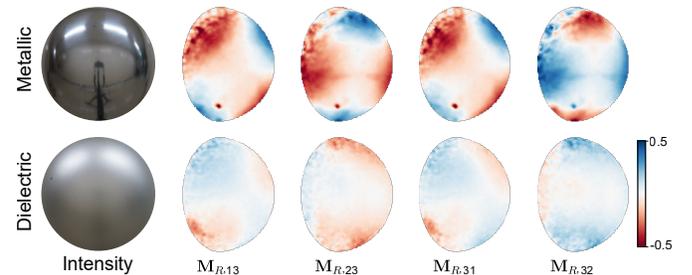


Fig. 13. **Metallic vs. dielectric.** The retardance matrices (M_R) of the metallic and dielectric materials at 710 nm. In their $M_{R,13}$, $M_{R,23}$, $M_{R,31}$, and $M_{R,32}$ elements, the metallic material shows pattern with opposite sign compared to the dielectric material.