

Polarimetric Spatio-Temporal Light Transport Probing

SEUNG-HWAN BAEK, Princeton University

FELIX HEIDE, Princeton University

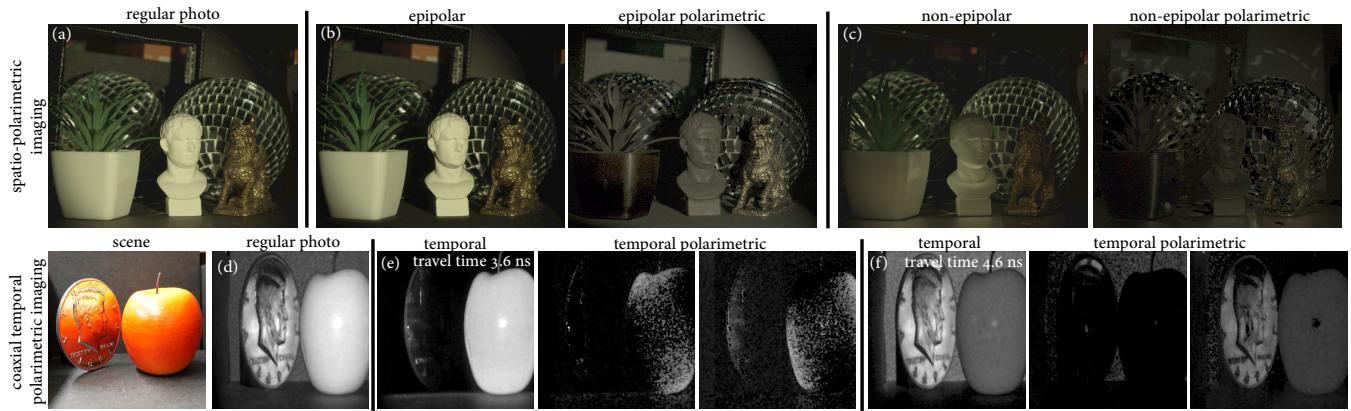


Fig. 1. We propose a computational light transport probing method that decomposes transport into full polarization, spatial and temporal dimensions. We model this multi-dimensional light transport as a tensor and analyze low-rank structure in the polarization domain which is exploited by our polarimetric probing method. We instantiate our approach with two imaging systems for spatio-polarimetric and coaxial temporal-polarimetric capture. (a)&(d) Conventional intensity imagers integrate incident light intensity over space and time independently of the polarization states of light, losing geometric and material information encoded in the polarimetric transport. Capturing polarization-resolved spatial transport components of (b) epipolar and (c) non-epipolar dimensions enable fine-grained decomposition of light transport. Combining temporal and polarimetric dimensions, we separate (e) geometry-dependent reflections and (f) direct/indirect reflections that cannot be resolved in the temporal-only measurements.

Light emitted from a source into a scene can undergo complex interactions with multiple scene surfaces of different material types before being reflected towards a detector. During this transport, every surface reflection and propagation is encoded in the properties of the photons that ultimately reach the detector, including travel time, direction, intensity, wavelength and polarization. Conventional imaging systems capture intensity by integrating over all other dimensions of the incident light into a single quantity, hiding this rich scene information in these aggregate measurements. Existing methods are capable of untangling these measurements into their spatial and temporal dimensions, fueling geometric scene understanding tasks. However, examining polarimetric material properties jointly with geometric properties is an open challenge that could enable unprecedented capabilities beyond geometric scene understanding, allowing for material-dependent scene understanding and imaging through complex transport, such as macroscopic scattering.

In this work, we close this gap, and propose a computational light transport imaging method that captures the spatially- and temporally-resolved complete polarimetric response of a scene, which encodes rich material properties. Our method hinges on a novel 7D tensor theory of light transport. We discover low-rank structure in the polarimetric tensor dimension and propose a data-driven rotating ellipsometry method that learns to exploit redundancy of polarimetric structure. We instantiate our theory with two imaging prototypes: spatio-polarimetric imaging and coaxial temporal-polarimetric imaging. This allows us, for the first time, to decompose scene light transport into temporal, spatial, and complete polarimetric dimensions that unveil scene properties hidden to conventional methods. We validate the applicability of our method on diverse tasks, including shape reconstruction with subsurface scattering, seeing through scattering media, untangling

multi-bounce light transport, breaking metamerism with polarization, and spatio-polarimetric decomposition of crystals.

1 INTRODUCTION

Light transport theory has proven itself to be a powerful tool across fields that rely on how light flows from a source through a scene to a detector. As photons travel through the scene, they interact with matter on surfaces or subsurfaces. This interaction modifies the photon properties (light direction, travel time, intensity, wavelength, and polarization) and, as such, encodes information about the scene. These properties have been formalized as plenoptic light dimensions augmented with wave properties [Adelson et al. 1991; Wetzstein et al. 2011]. Conventional detectors integrate over all dimensions of this high-dimensional plenoptic function. While this integration is beneficial in terms of photon efficiency, the individual plenoptic dimensions that encode scene information are lost during the conventional measurement process.

To recover this hidden scene information, a large body of work in computer graphics and vision has explored light transport approaches that aim to decompose the integral into a subset of the plenoptic dimensions. Specifically, considerable progress has been made in untangling spatial and temporal dimensions of light transport. Spatially controlling emitted and detected light allows for path-dependent geometric scene understanding [O’Toole et al. 2014b, 2012]. Similarly, advances in fast illumination and detection devices have allowed for temporal light transport capture [Heide et al. 2013; O’Toole et al. 2014a; Velten et al. 2013]. Joint spatio-temporal decomposition has also been explored recently [Kotwal

Authors’ addresses: Seung-Hwan Baek, Princeton University, Computer Science; Felix Heide, Princeton University, Computer Science.

et al. 2020; O’Toole et al. 2014a]. While these existing light transport decomposition methods have been successful for geometric scene understanding, analyzing material properties has been an open challenge. Real-world materials are complex and shape transport in the light-matter interaction. Specifically, we observe complex material-dependent scattering effects such as complex reflection, subsurface scattering, backscattering, multiple reflection between surfaces, and refractive-index-dependent reflection. Decomposing material-dependent light transport effects could allow for scene understanding and reconstruction tasks that exploit material properties, potentially fuelling broad applications across disciplines, including localization in VR/AR, object digitization, seeing through scattering media, and material analysis.

Polarization is an often overlooked dimension that may allow for analysis of material-dependent scattering across application domains. Polarization is a wave property of light describing how the electric field of light oscillates along the propagation direction in time. Previous polarization imaging methods in graphics focus on cross-polarization configurations that use either linear or circular polarization filters in front of a camera and a light source with perpendicular orientation [Ghosh et al. 2010, 2011; Gupta et al. 2008; Ma et al. 2007; Treibitz and Schechner 2008]. This approach, however, fundamentally prevents measurement of the full polarimetric response of the light transport, excluding polarization states of light such as elliptically polarized light and partially polarized light. Using linear polarizers at multiple angles as in polarization cameras has been also extensively used in computer vision [Atkinson and Hancock 2006; Ba et al. 2020; Kadambi et al. 2015; Tanaka et al. 2020], however, this configuration cannot acquire circular and elliptical polarization, and does not utilize the polarization of the emitted light. While Baek et al. [2018] proposed using a pair of rotating linear polarizers for reflectance acquisition, their approach also cannot model elliptically and circularly polarized light. More recently, goniometry systems have allowed researchers to acquire isotropic polarimetric BRDFs with full polarimetric resolution [Baek et al. 2020]. Orthogonal to the proposed light transport acquisition method, this method captures spherical homogenous objects and hence is not applicable to light transport analysis in a complex scene. In addition, it cannot capture the space and time dimensions.

In this work, we analyze material- and geometry-dependent scene transport effects by proposing the first spatio-temporal light transport imaging methods with full polarimetric resolution. The proposed acquisition method allows us to decompose light into components that have a specific travel time, path geometry and polarization state – allowing for a 7-dimensional representation of light transport covering 2D camera space, 2D illumination space, 1D travel time, and 2D polarization change. To this end, we propose an imaging method that combines ellipsometry, a technique in optics for characterizing polarimetric response, and spatial and temporal light transport acquisition from graphics and vision. We also introduce a novel tensor-based light transport theory that maintains natural backward compatibility to existing matrix-based light transport models in order to analyze and decompose multi-dimensional light transport. Given the high dimensionality of the full 7D light transport, directly acquiring the tensor is impractical with existing time-resolved sensors, and for the experimental setup shown

in this work would require *6 years per scene*. Instead, we analyze the formalized 7D light transport and identify low-dimensional embeddings of the polarimetric light transport that we exploit in our acquisition method. However, in contrast to existing methods for space and wavelength [O’Toole and Kutulakos 2010; Saragadam and Sankaranarayanan 2019] that *directly* acquire low-rank transport, *it is infeasible to directly acquire low-rank polarimetric light transport* as we have to solve an additional reconstruction problem, converting the intensity measurements to the Stokes-Mueller domain [Azzam 1978]. To overcome this challenge, we propose a novel data-driven rotating ellipsometry method that learns to rely on the discovered low-dimensional structure in the polarimetric dimension. We then instantiate our approach with a projector-camera ellipsometric imager and a coaxial SPAD-laser ellipsometric imager, that allow for spatio-polarimetric and coaxial temporal-polarimetric light transport imaging, respectively. We validate our method in simulation and experimentally using these prototype systems. We demonstrate novel applications of the recovered spatio-polarimetric and temporal-polarimetric light transport decompositions, including shape reconstruction with subsurface scattering, seeing through scattering media, untangling multi-bounce light transport, spatio-polarimetric decomposition of crystals, and breaking metamerism with polarization.

Specifically, we make the following contributions in this work

- We propose a computational light transport imaging method for shape and material analysis that jointly captures full polarimetric information along with spatial and temporal transport dimensions, spanning 7-dimensional light transport.
- We introduce a novel tensor-based light transport theory that is naturally backward-compatible with existing matrix models.
- We analyze the low-rank structure of polarimetric light transport and propose a learned ellipsometry method which implicitly exploits this low-dimensional polarimetric embedding.
- We instantiate two experimental prototypes of the proposed method for spatio-polarimetric imaging and coaxial temporal-polarimetric imaging.
- We validate the method with five applications, including shape reconstruction with subsurface scattering, seeing through scattering media, multi-bounce light transport decomposition, breaking metamerism with polarization, and spatio-polarimetric decomposition of crystals, where the proposed method outperforms existing approaches in *all* experiments.

Overview of Limitations. While the proposed imaging method makes it possible, for the first time, to acquire the complete polarization information together with the spatial and temporal dimensions, each of the two prototype systems only captures partial dimensions of the 7D spatio-temporal polarimetric tensor we analyze in this work. Although both systems could be combined with a high-resolution SPAD array, as proposed in Morimoto et al. [2020], pixel counts of existing commodity sensors are limited to a few kilopixels. We leave the substantial engineering efforts of such a joint

high-resolution temporal acquisition approach for future research. Also, the capture time of the proposed experimental prototypes is currently restricted to sequential acquisition on the order of hours per scene, limiting the proposed method to static scenes.

2 RELATED WORK

Computational Light Transport. Acquiring and decomposing light transport has been extensively studied in computer graphics and vision. Early approaches exploit statistics of global and direct reflections to decompose light transport based on the number of light-matter bounces [Gupta et al. 2011; Nayar et al. 2006]. Geometric constraints such as stereoscopic epipolar geometry have been also exploited to selectively capture light paths that satisfy the certain geometric constraints [Kubo et al. 2018; O’Toole et al. 2015; O’Toole and Kutulakos 2010; O’Toole et al. 2014b, 2012]. These spatial decomposition methods are typically implemented with a camera and a projector where the projection is masked during the acquisition process. Another line of research untangles the temporal dimension of light transport [Jarabo et al. 2017; O’Toole et al. 2017, 2014a; Velten et al. 2013; Wu et al. 2014]. Combined with a fast pulsed laser, direct temporal sampling approaches rely on a streak camera [Velten et al. 2013] or a single-photon avalanche diode (SPAD) [O’Toole et al. 2017] to directly measure such temporal scene responses. An alternative line of work relies on a correlation sensor and amplitude-synchronized continuous illumination to acquire temporal light transport [Gupta et al. 2015; Heide et al. 2013; Kadambi et al. 2013; Peters et al. 2015]. Interferometric approaches have demonstrated micron-scale temporal light transport [Gkioulekas et al. 2015]. Combining these two threads of spatial and temporal dimensions makes it possible to reason over rich decompositions that encode how light travels at a given time and position. Such joint spatio-temporal analysis has been implemented using frequency-aware probing methods with a correlation sensor-illumination pair [O’Toole et al. 2014a] or, recently, using an interferometric setup with programmable amplitude and phase modulators [Kotwal et al. 2020]. The spatio-temporal transport structure enables direct-indirect separation for multipath interference removal [Agresti and Zanuttigh 2018; Naik et al. 2015; Whyte et al. 2015]. Also, existing works have explored joint analysis of temporal and partial polarimetric dimensions for light field imaging [Manakov et al. 2013], surface reconstruction [Dashpute et al. 2018; Maeda et al. 2018] and imaging through scattering media [Wu et al. 2018, 2016].

None of these existing methods are capable of capturing full polarization states (linear, circular, ellipsoidal with varying degrees) jointly with spatial and temporal dimensions. The only polarimetric sampling demonstrated before captures partial linear polarization [Gkioulekas et al. 2015]. In this work, we propose a novel spatially- and temporally-resolved ellipsometry methods that allow us to capture a complete polarimetric light transport of any polarization state. As such, for the first time, we demonstrate full spatio-polarimetric and temporal-polarimetric captures of the 7-dimensional light transport (2D spatial, 2D detector, 2D polarimetric, 1D temporal dimension).

Polarimetric Imaging in Graphics and Vision. The polarization state of light has been exploited by a broad range of methods in computer graphics and vision to encode material-dependent scene

properties or source-dependent transport. Cross-polarization imaging, one of the most common polarimetric acquisition approaches, uses either linear or circular polarization filters in front of a camera and a light source with perpendicular orientation which separates captured light into two components depending on whether light maintains the polarization of the light source or not. This approach to source separation has been adopted in various applications, most notably seeing through scattering media [Gupta et al. 2008; Tanaka et al. 2013; Treibitz and Schechner 2008] and its application for shape estimation [Chen et al. 2007], also diffuse-specular separation [Ghosh et al. 2010, 2008; Ma et al. 2007; Nagano et al. 2015; Riviere et al. 2017] for facial capture and geometric scene understanding. This relies on unique polarization encoding of material-dependent reflectance properties. Polarimetric imaging also has been used for estimating depth using birefringent materials [Baek et al. 2016; Meuleman et al. 2020] and surface normals [Atkinson and Hancock 2006; Baek et al. 2018; Deschaintre et al. 2021; Ichikawa et al. 2021; Kadambi et al. 2015] in addition to metallic surface characterization [Berger et al. 2011]. While substantial progress has been made in exploiting polarization for these diverse applications, the information of *all* polarization states jointly, along with spatial and temporal resolutions, has not been exploited in graphics and vision. Indeed, only very recently researchers proposed data-driven polarimetric BRDF models [Baek et al. 2020], however, without considering the spatial and temporal dimensions. In this work, we exploit rich structure in light polarization along with space and time dimensions, and propose a complete spatio- and temporal-polarimetric imaging and decomposition method that allows us to analyze scene properties hidden to existing light transport imaging methods.

Optical Ellipsometry. Consider a material sample observed by a sensor and illuminated by a light source. The polarimetric response of the sample encodes material information including its refractive index, roughness, composition, crystalline structure, and electrical conductivity. Ellipsometry describes a method of obtaining such material properties by capturing the complete polarimetric transport that the material sample is subject to, often by synchronously modulating the polarization states of emitted and captured light. Ellipsometric methods have been studied during the past few decades in optics [Azzam 2016], material science [Jellison Jr et al. 2018; Ramsey and Ludema 1994], and biology [Arwin 2011; Ghosh and Vitkin 2011]. In this work, we go beyond conventional ellipsometry by jointly capturing spatial, temporal, and polarimetric dimensions. As a result, the proposed spatio- and temporal-polarimetric imaging allows us to decompose light transport into unprecedented detail, revealing joint polarimetric spatial and temporal information that appears “hidden” to existing methods. In addition, we analyze low-rank structure of real-world polarimetric light transport and develop novel data-driven ellipsometry methods which rotate polarizing optics at optimal angles to achieve high-fidelity acquisition of polarimetric light transport with a reduced number of captures.

3 BACKGROUND

Before introducing the proposed spatio- and temporal-polarimetric light transport imaging method, we review existing models for

dimension	(a) space	(b) time	(c) polarization	(d) polarimetric spatio-temporal (ours)
diagram				
model	$\begin{bmatrix} i_0^s \\ i_1^s \end{bmatrix} = \mathbf{T}_s \times \begin{bmatrix} p_0^s \\ p_1^s \end{bmatrix}$	$\begin{bmatrix} i_0^t \\ i_1^t \end{bmatrix} = \mathbf{T}_t \times p^t$	$\begin{bmatrix} i_0^p \\ i_1^p \\ i_2^p \\ i_3^p \end{bmatrix} = \mathbf{T}_p \times \begin{bmatrix} p_0^p \\ p_1^p \\ p_2^p \\ p_3^p \end{bmatrix}$	$\mathcal{I}_{4D} = \mathcal{T}_{7D} \otimes \mathcal{P}_{3D}$

Fig. 2. Previous light transport models for (a) space [O’Toole et al. 2012], (b) time [Velten et al. 2013], and (c) polarization [Collett 2005] are based on linear modeling as a matrix-vector multiplication. (d) We propose a tensor-based light transport model which simultaneously comprises space, time, and polarization. We model light transport as a 7D tensor \mathcal{T} covering 2D camera space, 2D illumination space, 1D travel time, and 2D polarization change. This 7D transport is then multiplied with a 3D illumination tensor \mathcal{P} covering 2D illumination space and 1D polarization, resulting in a captured 4D tensor \mathcal{I} covering the 2D camera space, 1D polarization, and 1D travel time.

spatial and temporal light transport in graphics and vision, and polarimetric light transport in optics.

Spatial Light Transport. Given a 2D projector and a 2D camera, the spatial light transport between them can be modeled as a matrix-vector multiplication [Ngan et al. 2005; Sen et al. 2005]

$$\mathbf{i}_s = \mathbf{T}_s \mathbf{p}_s, \quad (1)$$

where $\mathbf{i}_s, \mathbf{p}_s \in \mathbb{R}^{S \times 1}$ are the vectorized 2D images captured and emitted by the camera and the projector, respectively, and S being the number of pixels. The matrix $\mathbf{T}_s \in \mathbb{R}^{S \times S}$ is the spatial light transport matrix that describes how light intensity, emitted at a specific angle from a projector pixel, changes as it travels through a scene and is captured at a specific pixel on the camera. The matrix \mathbf{T}_s , therefore, decomposes the *spatial dimension* of light transport indexed as pixels on the projector and camera, respectively, see Figure 2(a).

Temporal Light Transport. Spatial light transport assumes a steady-state equilibrium of global light transport. However, in reality, emitted light reaches the detector along potentially complex paths with varying lengths, and hence may arrive at the sensor at different travel times. This temporal light transport, before reaching a steady state, can be modeled for a *single detector pixel* as

$$\mathbf{i}_t = \mathbf{T}_t \mathbf{p}^t, \quad (2)$$

where $\mathbf{i}_t \in \mathbb{R}^{\Gamma \times 1}$ is the vector of detected intensities of travel time t discretized in Γ different temporal bins. The scalar p^t describes the emitted light intensity at the time index zero. The matrix $\mathbf{T}_t \in \mathbb{R}^{\Gamma \times 1}$ is *temporal light transport matrix* for a given detector pixel, see Figure 2(b). This can be further generalized to shift-invariant light transport with the convolutional model

$$\mathbf{i}_t = \mathbf{T}_t * \mathbf{p}_t, \quad (3)$$

where $*$ is a matrix-vector convolution [O’Toole et al. 2014a] across the temporal dimension t and $\mathbf{p}_t \in \mathbb{R}^{\Gamma \times 1}$ are the intensities vectorized in time bins.

Polarimetric Light Transport.

Polarization is an intrinsic wave property of light describing the geometric oscillation pattern of the electromagnetic field, i.e., light, in time. The Stokes-Mueller formalism provides a representation for polarization, describing any polarization state of light and its change [Collett 2005]. Specifically, the Stokes vector is a four-dimensional vector describing the polarization state of light: $\mathbf{s} = [s_0, s_1, s_2, s_3]$, where each element describes the radiance of light, the linear-polarization components at 0° and 45° , and left-right circular-polarization component.

To provide intuition for the Stokes vector formalism, Poincaré spheres offer an illustrative visualization of the complete polarization state of light based on the Stokes vector by plotting the last three Stokes elements normalized by the total power s_0 : $[s_1/s_0, s_2/s_0, s_3/s_0]$, see Figure 3. A Stokes vector is then represented as a vector inside the Poincaré sphere. The sphere’s center represents unpolarized light and light becomes more polarized as we are approaching the surface of the Poincaré sphere. The direction of the Stokes vector in the Poincaré sphere describes the type of polarization such as horizontal linear or right circular polarization.

To characterize polarimetric transport, we use the Mueller matrix which describes the polarimetric response of a scene as the change of a Stokes vector

$$\mathbf{i}_p = \mathbf{T}_p \mathbf{p}_p, \quad (4)$$




Fig. 3. Poincaré sphere.

where $\mathbf{i}_p, \mathbf{p}_p$ are the Stokes vectors of the detected and emitted light, and \mathbf{T}_p is the polarimetric light transport matrix in the form of a Mueller matrix for a single detector pixel, see Figure 2(c).

4 TENSOR LIGHT TRANSPORT

Existing spatial [O’Toole et al. 2012], temporal [Velten et al. 2013], and polarimetric [Collett 2005] light transport models rely on matrix-vector representations. While these linear models offer convenient links to linear algebra for each plenoptic dimension, they do not treat light transport jointly across space, time, and polarization. As a generalization to these existing models, we *propose a tensor-based model of light transport*. For a primer on tensor algebra, we refer the reader to Itskov [2007].

We start by modeling the polarization state and intensity of emitted pulsed light as a 3D tensor $\mathcal{P}(s, p)$, parameterized by the 2D spatial direction s and the 1D Stokes vector dimension p . Light modeled by the tensor \mathcal{P} then travels through a scene, e.g., undergoing potentially complex scattering effects, and encodes material and geometric properties of the scene. When light returns to the detector, we model it as a 4D tensor $\mathcal{I}(s, p, t)$. The temporal dimension t (1D) denotes the travel time, s is the camera pixel location (2D) and p is the Stokes vector dimension (1D). With the two tensors \mathcal{P} and \mathcal{I} in hand, we model the spatio-temporal polarimetric light transport as a 7D tensor $\mathcal{T}(s, s', p, p', t)$, where s and s' are the 2D spatial sample locations at the detector and the light source, p and p' are the 1D Stokes vectors of the detected and the emitted light, and t is the 1D travel time. Refer to Figure 2(d). We write the light transport equation using tensor contraction [Itskov 2007] as

$$\mathcal{I}(s, p, t) = \sum_{s', p'} \mathcal{T}(s, s', p, p', t) \mathcal{P}(s', p'). \quad (5)$$

For shift-invariant light transport, this can be formulated as a convolutional integral

$$\mathcal{I}(s, p, t) = \sum_{s', p'} \int_{-\infty}^{\infty} \mathcal{T}(s, s', p, p', t - t') \mathcal{P}(s', p', t') dt', \quad (6)$$

where t' is the time dimension of the time-resolved illumination tensor \mathcal{P} .

This 7D spatio-temporal polarimetric light transport tensor \mathcal{T} describes the complete polarimetric change of light for every combination of camera and projector pixels with high temporal resolution of the travel time. *Probing* operations can be performed by slicing the tensor along one or two indices, reducing the tensor into a vector or a matrix. This means that our tensor-based transport model is backward-compatible to the existing matrix representations of light transport in Equations (1), (2), (4) as follows

$$\begin{aligned} \mathbf{T}_t &= \sum_{s, s'} \mathcal{T}(s, s', 0, 0, :), \\ \mathbf{T}_p &= \sum_{s, s', t} \mathcal{T}(s, s', :, :, t). \end{aligned} \quad (7)$$

4.1 Path Tracer for 7D Light Transport Simulation

To provide further intuition on the properties of the high-dimensional 7D light transport tensor \mathcal{T} , we developed a path tracer and we simulate a synthetic capture system consisting of a high-speed 2D

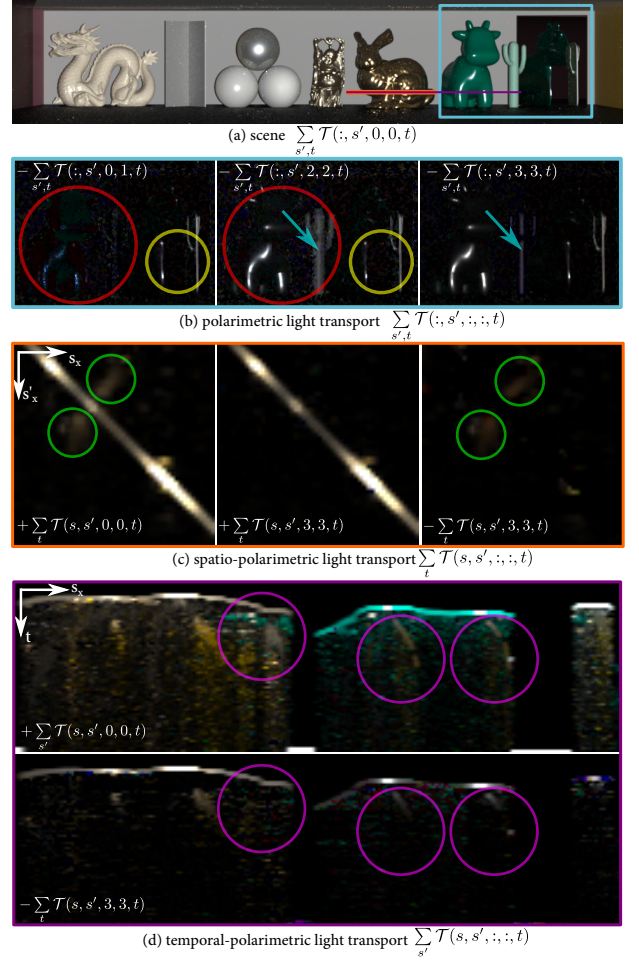


Fig. 4. Rendered simulation of a complete 7D light transport \mathcal{T} tensor using a multi-dimensional path tracer. (a) A conventional camera integrates light intensity emitted from all projector pixels s' at temporal steady-state t , capturing $\sum_{s', t} \mathcal{T}(:, s', 0, 0, t)$. We analyze the tensor \mathcal{T} by slicing different dimensions of space, time, and polarization. (b) The polarimetric light transport $\sum_{s', t} \mathcal{T}(:, s', :, :, t)$ reveals material-dependent scene responses, for instance, untangling reflections from diffuse surface points (red circles) and mirrors (yellow circles), and allows us to separate single-bounce specular reflection from scattered specular reflection (cyan arrows). (c) An additional spatial dimension, spatio-polarimetric light transport $\sum_t \mathcal{T}(s, s', :, :, t)$ separates single-bounce specular reflection from multi-bounce specular reflection (green circles) visualized for the orange scanline shown in (a). (d) Capturing the temporal-polarimetric light transport $\sum_{s'} \mathcal{T}(s, s', :, :, t)$ makes it possible to decompose time-varying global illumination into diffuse and specular components (magenta circles) visualized for the purple scanline shown in (a).

projector, a high-speed 2D camera with picosecond temporal resolution, and polarizing optics. We note, that although such high fidelity transient projectors and sensors do not exist today. We validate the proposed method using two practical imaging methods in the remainder of this work. We implement our renderer on top of the Mitsuba2 renderer [Nimier-David et al. 2019]. Specifically, we set




Fig. 5. PCA analysis of polarimetric light transport via real-world materials. (a) Polarimetric BRDFs with intensity normalization exhibit low-rank structure. (b) To perform PCA analysis, we propose the arctangent mapping that handles both negative and positive values of polarimetric light transport, in contrast to conventional positive-only log mappings.

up spatial acquisition through masking operations for a projector-camera pair. For the temporal dimension, we keep track of the time-of-flight for each light ray throughout the rendering process [Jarabo et al. 2014; Pediredla et al. 2019]. For polarization, we use a real-world polarimetric BRDF dataset [Baek et al. 2020] and implement a dual-rotating-retarder setup [Azzam 1978]. We render trichromatic imagery at wavelengths of 450, 550, and 660 nm. Figure 4 shows the rendered 7D light transport, revealing geometry- and material-dependent reflections. We observe that by separating polarization, space, and time dimensions, it is possible to distinguish mirror and scene reflections, acquire mirror-like components isolated from cluttered specular reflection, decompose specular reflections into single and multi-bounces, and analyze global illumination as mixed diffuse and specular transport.

4.2 Low-dimensional Polarimetric Transport Embeddings

We further analyze the light transport tensor, in particular its polarization slice, which has been overlooked in graphics and vision compared to the other well-understood dimensions, such as space and time. Specifically, we analyze low-dimensional polarimetric embeddings of the high-dimensional light transport tensor \mathcal{T} . Such lower-dimensional embeddings offer statistical priors that may allow for efficient acquisition and provide insight into the transport properties. Low-dimensional transport structure has also been used in the spatial dimension [O’Toole and Kutulakos 2010], spectral dimension [Saragadam and Sankaranarayanan 2019], and temporal dimension [Heide et al. 2019; O’Toole et al. 2018].

To this end, we perform principal component analysis (PCA) on 25 real-world polarimetric BRDFs [Baek et al. 2020]. This facilitates efficient polarimetric analysis, in contrast to naively analyzing a large collection of rendered polarimetric images using costly polarization-aware path tracing. Furthermore, polarimetric light transport can be represented as an intensity-scaled polarimetric BRDF by integrating over indirect transport components.

Inspired by PCA analysis on intensity BRDFs [Nielsen et al. 2015], we develop our analysis method with two major distinctions. First, we normalize the polarimetric BRDF in order to decouple polarization from intensity via division by the intensity component: $\mathbf{T}_p \mapsto \mathbf{T}_p / \mathbf{T}_p(0, 0)$. This allows us to separately investigate the polarimetric variation independently of intensity scale. Second, to distribute the polarimetric energy across a high dynamic range of inputs, we propose to apply the arctangent mapping: $\mathbf{T}_p \mapsto$

$\arctan(c\mathbf{T}_p)$, where c is a scalar which is set to 8 in our experiment. This handles positive and negative inputs in contrast to the positive-only log mapping [Nielsen et al. 2015]. We arrange the normalized and compressed polarimetric BRDFs to form an observation matrix $\mathbf{X} \in \mathbb{R}^{N \times 16}$, where N is the number of valid angular bins for all dataset materials. We extract the principal components of \mathbf{X} using singular value decomposition (SVD)

$$\mathbf{X} - \mu = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (8)$$

where μ is the mean matrix of \mathbf{X} over the valid angular bins of different materials. \mathbf{U} and \mathbf{V} contain the eigenvectors and $\mathbf{\Sigma}$ has the eigenvalues. Then, we obtain the scaled principal components $\mathbf{V}\mathbf{\Sigma}$ shown in Figure 5. These visualizations reveal the low-rank structure in the captured light transport. We refer to the Supplemental Document for a low-rank analysis of the angular-polarimetric dimensions.

5 LEARNING ROTATING ELLIPSOMETRY

To acquire the complete polarimetric light transport, conventional optical ellipsometry rotates polarizing optics in front of a light source and a sensor while acquiring polarization-modulated intensities from which polarimetric light transport is reconstructed. For instance, dual-rotating-retarders (DRR) [Azzam 1978] uses two fixed linear polarizers and two rotating retarders at hand-crafted angles, resulting in many captures, typically 36. The previous section illustrates low-dimensional polarimetric embeddings in the transport tensor which opens a potential avenue to a more efficient capture procedure. A similar principle has been applied to the plenoptic dimensions defined on intensity by iteratively acquiring low-rank approximations of respective light transport slices [O’Toole and Kutulakos 2010; Saragadam and Sankaranarayanan 2019]. However, *directly capturing low-dimensional polarimetric structure is infeasible* as polarimetric light transport is based on Stokes-Mueller formalism, not intensity measurements, therefore an additional ellipsometric reconstruction step is required. To overcome this challenge, we propose a data-driven ellipsometry method that implicitly exploits this polarimetric low-rank structure for efficient capture. We dub our method *learned quad-rotating polarizer retarder* (L-QRPR). It uses rotating linear polarizers and retarders at learned angles to provide accurate reconstructions for real-world polarimetric light transport tensors.

To learn the set of optimal angles, we first propose a differentiable image formation model of a rotating ellipsometry

$$I^{\theta_k, \theta_k^3, \theta_k^2, \theta_k^1} = \left[\mathbf{L}_k^{\theta_k} \mathbf{Q}_k^{\theta_k^3} \mathbf{T}_p \mathbf{Q}_k^{\theta_k^2} \mathbf{L}_k^{\theta_k^1} \mathbf{p} \right]_0 + \eta, \quad (9)$$

where k is the rotation index out of the total number of captures K . $\mathbf{p} = [1, 0, 0, 0]^T$ is the unit Stokes vector and $[\cdot]_0$ is the first element of the vector. \mathbf{Q}^θ and \mathbf{L}^θ are the Mueller matrices of a quarter-wave plate (QWP) and a linear polarizer (LP) at an angle θ . For the definitions of the Mueller matrices, we refer to the Supplemental Document. $I^{\theta_k, \theta_k^3, \theta_k^2, \theta_k^1}$ is the simulated intensity and η is Gaussian noise with standard deviation 5×10^{-4} . Note that the matrix-vector multiplications based on Stokes-Mueller formalism make this image formation naturally differentiable with respect to the angles of the polarizing optics, $\theta_k^{i \in \{1, 2, 3, 4\}} \in \{1, \dots, K\}$.

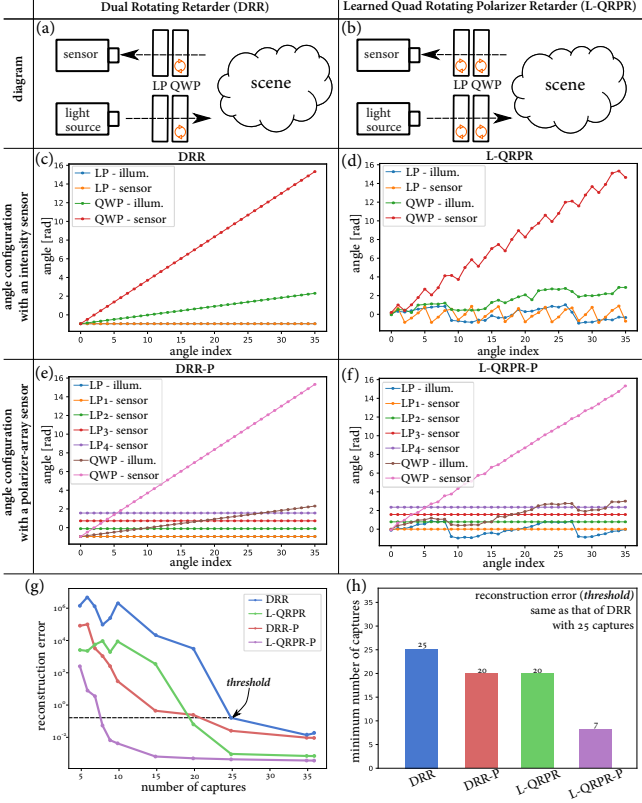


Fig. 6. (a) Conventional DRR ellipsometry rotates the two quarter-wave plates at (c) manually-chosen angle sets. (b) In contrast, we propose to rotate the two linear polarizers and the two retarders (d) using a smaller set of angles that learns to exploit the low-rank structure of natural polarimetric light transport. (f) We also show how the proposed method can be extended to an emerging polarizer-array sensor. To evaluate the proposed learned ellipsometry method, we conduct 5-fold cross-validation. (g) Our learned ellipsometry consistently outperforms the conventional DRR method both for the intensity sensor and the polarizer-array sensor. (h) Our method substantially reduces the number of required captures by more than 3× without sacrificing reconstruction accuracy.

Given the simulated intensities, we reconstruct the polarimetric light transport by solving an optimization problem

$$\underset{\hat{\mathbf{T}}_p}{\text{minimize}} \sum_{k=1}^K \left(I^{\theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1} - \left[\mathbf{L}^{\theta_k^4} \mathbf{Q}^{\theta_k^3} \hat{\mathbf{T}}_p \mathbf{Q}^{\theta_k^2} \mathbf{L}^{\theta_k^1} \mathbf{p} \right]_0 \right)^2, \quad (10)$$

where $\hat{\mathbf{T}}_p$ is the estimate of the polarimetric light transport. Note that the solution to Equation (10) can be analytically obtained as the least-square solution [Baek et al. 2020], and is therefore differentiable.

Last, we compute the loss with the ground truth \mathbf{T}_p , and optimize the angle set to minimize this loss

$$\underset{\Theta_k^4, \Theta_k^3, \Theta_k^2, \Theta_k^1}{\text{minimize}} \mathcal{L}(\hat{\mathbf{T}}_p, \mathbf{T}_p), \quad (11)$$

where \mathcal{L} is the mean-squared-error loss function and $\Theta_k^i = \{\theta_k^1, \dots, \theta_k^i\}$ is an angle set for each optical element. We use the Adam optimizer [Paszke et al. 2017] and initialize the angles of polarizing optics following the conventional DRR setting [Azzam 1978] as $\theta_k^1 =$

$0^\circ, \theta_k^2 = 5(k-1)^\circ, \theta_k^3 = 25(k-1)^\circ, \theta_k^4 = 0^\circ$ with $k \in \{1, 2, \dots, K\}$. As ground truth data, we use the real-world polarimetric BRDF dataset from Baek et al. [2020]. Note that data-driven optimal capture has also been studied for reflectance acquisition [Nielsen et al. 2015; Xu et al. 2016].

Learning Rotation Angles with a Polarizer-array Sensor. The proposed method allows for learning optimal angles not only for the conventional CMOS intensity sensors, but also for emerging polarizer-array CMOS sensors, e.g., Sony IEDM2016. These emerging sensors increase the sampling rate by four times by using on-sensor linear polarizers with four different orientations ($0^\circ, 45^\circ, 90^\circ, 135^\circ$). We incorporate these hardcoded polarization directions into the differentiable image formation as

$$I_{q \in \{1, 2, 3, 4\}}^{\theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1} = \left[\mathbf{L}_{q \in \{1, 2, 3, 4\}} \mathbf{Q}^{\theta_k^3} \mathbf{T}_p \mathbf{Q}^{\theta_k^2} \mathbf{L}^{\theta_k^1} \mathbf{p} \right]_0 + \eta, \quad (12)$$

where q is the index of the sensor linear polarizers. Inserting this into Equation (10) and Equation (11) allows us to optimize for the angle sets of $\theta_k^3, \theta_k^2, \theta_k^1$ while the polarizer angles for the sensor are fixed.

Figure 6 shows that our ellipsometry methods, learned quad-rotating polarizer and retarder (L-QRPR) and its variant with a polarization camera (L-QRPR-P), outperform the conventional dual rotating retarder (DRR) method in terms of capture efficiency. For synthetic validation, we conduct a 5-fold cross-validation experiment. L-QRPR and L-QRPR-P both achieve the same accuracy as conventional DRR using 25 captures while using 20% and 72% fewer captures respectively. As such, we validate that the proposed data-driven rotating ellipsometry effectively exploits the low-rank structure in natural polarimetric light transport. For experimental validation, we use the angles learned from the pBRDF dataset [Baek et al. 2020] and show that our method generalizes to unseen materials, see Figures 10 and 11. We refer to the Supplemental Document for additional assessment.

6 SPATIO-TEMPORAL POLARIMETRIC IMAGING

The full 7D light transport tensor \mathcal{T} is challenging to acquire due to its high dimensionality of space, time, and polarization, even with our learned ellipsometry. While methods have been proposed before that acquire transport for each individual dimension, the joint acquisition and analysis of the full 7D light transport has been a challenge, especially for the polarization dimension. We propose polarimetric acquisition approaches that allow us to capture space, time and polarization dimensions of light transport jointly, while being realizable as practical imaging systems. The absence of high-resolution, time-tagged sensors and pulsed projectors poses a practical restriction on capture time. To address this challenge, we propose two novel computational imaging systems that capture spatio-polarimetric light transport \mathcal{T}_{sp} (6D) and coaxial temporal-polarimetric light transport \mathcal{T}_{pt} (5D) respectively

$$\begin{aligned} \mathcal{T}_{sp} &= \sum_t \mathcal{T}(:, :, :, :, t), \\ \mathcal{T}_{pt} &= \mathcal{T}(s, s, :, :, :), \forall s. \end{aligned} \quad (13)$$

In the following, we describe these two instantiations of our generalized polarimetric transport acquisition. We refer the reader

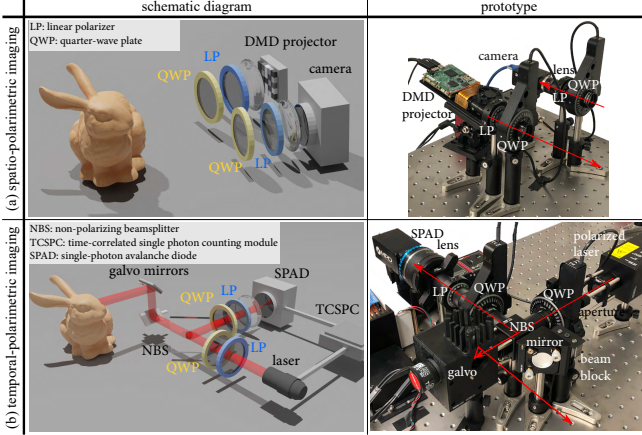


Fig. 7. We propose two novel imaging systems that capture spatio-polarimetric light transport \mathcal{T}_{sp} and coaxial temporal-polarimetric light transport \mathcal{T}_{pt} . The spatio-polarimetric system acquires the complete polarimetric scene response for every combination of camera and projector pixels over trichromatic spectral bands, using camera-projector masks with rotating polarization optics. The temporal-polarimetric setup captures the coaxial spatial dimension at picosecond temporal resolution for a narrow spectral band, also with the full polarimetric resolution. We scan the transport using a 2-axis galvo mirror and rotating QWPs.

to the Supplemental Document for comparison with other systems and calibration details.

6.1 Spatio-Polarimetric Imaging

We present a computational imaging method that captures the 6D spatio-polarimetric light transport \mathcal{T}_{sp} using a high-resolution intensity sensor. The 2D intensity sensor allows us to efficiently resolve spatial dimensions of the 2D illumination s' and camera s . Figure 7(a) illustrates this spatially-multiplexed acquisition system.

Illumination. To spatially multiplex illumination patterns, we employ a DMD projector that emits unpolarized light with a programmable intensity from a certain projector pixel s' . Light polarization is then modulated by passing through a LP and a QWP. The LP and the QWP are mounted on rotation stages at angles θ_k^1 and θ_k^2 , respectively at the k -th capture index. This provides the emitted light tensor $\mathcal{P}^{\theta_k^2, \theta_k^1}(s', :)$ from a projector pixel s'

$$\mathcal{P}^{\theta_k^2, \theta_k^1}(s', :) = \mathcal{Q}^{\theta_k^2} \mathbf{L}^{\theta_k^1} \mathbf{p}^{s'}, \quad (14)$$

where $\mathbf{p}^{s'} = [p^{s'}, 0, 0, 0]^T$ is the Stokes vector of the unpolarized DMD light and $p^{s'}$ is the intensity for that spatial direction s' .

Spatio-Polarimetric Light Transport. The emitted light interacts with a scene described by the spatio-polarimetric light transport \mathcal{T}_{sp} that is $\mathcal{R}^{\theta_k^2, \theta_k^1}(s, :) = \mathcal{T}_{sp}(s, s', :; :)\mathcal{P}^{\theta_k^2, \theta_k^1}(s', :)$, where s is the spatial position on the detection module. The spatio-polarimetric light transport \mathcal{T}_{sp} assumes steady-state equilibrium of light transport, integrating the full tensor \mathcal{T} over the temporal dimension t .

Detection. We capture the light that returns to the camera using an intensity array sensor (CMOS sensor) equipped with a QWP and a LP in front of the camera. The QWP and the LP are mounted

on rotation stages at angles θ_k^3 and θ_k^4 , resulting in the light tensor arriving at the pixel s on the CMOS sensor

$$\mathcal{I}^{s', \theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1}(s, :) = \mathbf{L}^{\theta_k^4} \mathcal{Q}^{\theta_k^3} \mathcal{R}^{\theta_k^2, \theta_k^1}(s, :). \quad (15)$$

For a 2D CMOS sensor pixel s , our spatio-polarimetric measurement f_{sp} is then formed as the total recorded radiance of incident light, that is

$$f_{sp}(\theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1, s, s') = \mathcal{I}^{s', \theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1}(s, 0). \quad (16)$$

Rotating Polarimetric Modulation. The rotation angles of the two LPs and the two QWPs are crucial for diverse polarimetric sampling. We evaluate two configurations including the established DRR scheme implemented with the fixed LPs and rotating QWPs at manually-chosen angles [Azzam 1978] and our learned QRPR method with rotating LPs and QWPs at learned angles, see Figure 6 for the specific angles we used.

Polarizer-array Sensor. By replacing the CMOS intensity camera with a polarizer-array camera, we can sample multiple linear polarization states simultaneously, further reducing the number of captures as described in Section 5. We experimentally implemented this polarizer-array configuration shown in Figure 7.

6.2 Coaxial Temporal-Polarimetric Imaging

Next, we describe the second instance of the proposed light transport theory. Specifically, to acquire temporal-polarimetric light transport \mathcal{T}_{pt} , we present a novel coaxial ellipsometric SPAD-laser imaging method. The proposed system consists of an illumination module and a detection module; a pulse of light is emitted to a scene and detected by a high-speed detector SPAD where both emission and detection paths involve polarization-modulating optics. See Figure 7(b) for the schematic diagram and prototype photograph.

Illumination. We use a picosecond pulsed laser to emit a linearly-polarized pulse of light at time zero at the linear polarization angle of zero $\theta_k^1 = 0^\circ$. The polarization state of the emitted pulsed light is then modulated by a QWP mounted on a rotation stage at an angle θ_k^2 . Light passes through a non-polarizing beam splitter in transmission mode and is reflected by a 2D galvo mirror specified by a spatial orientation s , exiting the imaging system towards a scene. Then, for each QWP angle θ_k^2 and spatial orientation s , the light emitted to a scene is modeled as a tensor

$$\mathcal{P}^{\theta_k^2, \theta_k^1}(s, :) = \mathbf{G}^s \mathbf{B}_t \mathcal{Q}^{\theta_k^2} \mathbf{p}, \quad (17)$$

where \mathbf{p} is the Stokes vector of the linearly polarized laser pulse at the polarization angle of $\theta_k^1 = 0^\circ$, \mathbf{B}_t is the beamsplitter Mueller matrix in transmission mode, \mathbf{G}^s is the galvo mirror Mueller matrix oriented to the spatial direction s .

Temporal-Polarimetric Light Transport. The emitted light travels through a scene and arrives at a detection module after potentially complex light-matter interactions. This temporal-polarimetric light transport is modeled as a tensor \mathcal{T}_{pt} and is applied to the illumination tensor $\mathcal{P}^{\theta_k^2}$, resulting in the reflected tensor $\mathcal{R}^{\theta_k^2}$ as

$$\mathcal{R}^{\theta_k^2, \theta_k^1}(s, :, t) = \mathcal{T}_{pt}(s, s, :, :; t) \mathcal{P}^{\theta_k^2, \theta_k^1}(s, :). \quad (18)$$




Fig. 8. A conventional photo integrates light intensity reflected from a scene over temporal and spatial dimensions. Specifically, conventional photography captures $\sum_{t,s'} \mathcal{T}(s, s', 0, 0, t)$, while losing geometric and material information encoded in the scene transport. Spatially resolving the light transport allows for (a) epipolar and (f) non-epipolar imaging that approximates direct and indirect light transport components respectively as $\sum_t \mathcal{T}(s, s'_e, 0, 0, t)$ and $\sum_t \mathcal{T}(s, s'_n, 0, 0, t)$. Our spatio-polarimetric imaging approach acquires the complete polarimetric scene response which reveals strongly material-dependent structure in the transport. We observe the specular components that maintain horizontal linear polarization in (b) the approximated direct and (g) indirect reflection by epipolar and non-epipolar imaging. Note that the scattering inside of the crystal can be observed in the non-epipolar polarimetric component (g). (c) Analyzing the epipolar circular-polarization component, we can examine the surface of the transparent crystal. We further decompose the spatio-polarimetric light transport to analyze (d) & (i) polarizance and (e) & (j) retardance of the crystal. This reveals the birefringence properties of crystals indicated by red arrows. The two indicated crystal objects are made of calcite (left) and quartz (right) respectively.

To improve signal-to-noise ratio (SNR), we design our imaging system to be coaxial with the non-polarizing beamsplitter, sharing the same optical path for detection and emission. Therefore, we use the same spatial sample s for the illumination and the detection.

Detection. The light returning from the scene first passes through the 2D galvo mirror. It then enters the non-polarizing beamsplitter with reflection mode followed by two polarizing optics, a QWP and a LP. The QWP and the LP are mounted on rotation stages at angles θ_k^3 and θ_k^4 . The SPAD detector captures the temporal response of the scene at a specific spatial orientation s as

$$I^{\theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1}(s, :, t) = \mathbf{L}^{\theta_k^4} \mathbf{Q}^{\theta_k^3} \mathbf{B}_r \mathbf{G}^s \mathcal{R}^{\theta_k^2, \theta_k^1}(s, :, t), \quad (19)$$

where \mathbf{B}_r is the Mueller matrix of the beamsplitter in reflection mode. Then, the SPAD sensor measures the total radiance and we write this temporal-polarimetric image formation f_{pt} as

$$f_{pt}(\theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1, s, t) = I^{\theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1}(s, 0, t). \quad (20)$$

For our temporal-polarimetric experiments, we fix the LP angles as zero. We refer to the Supplemental Document for discussions on extending our quad-rotating ellipsometry to this temporal setup.

7 POLARIMETRIC RECONSTRUCTION

We record spatial and temporal information as *intensity* observations using a 2D CMOS camera and a SPAD sensor combined with a time-correlated single photon counting (TCSPC) device. As these conventional sensors cannot directly measure the polarization state of light, we computationally recover polarimetric light transport in the form of Mueller matrix reconstruction from intensities observed

with varying angles of polarizing optics. Note that the Mueller matrix corresponds to the polarimetric slicing of the full light transport tensor.

To recover polarimetric information, we record a set of intensities $I^{\Theta_k}(s, t)$ and $I^{s', \Theta_k}(s)$ with varying angle of the polarizing optics, where Θ_k describes the angle configuration as $\{\theta_k^4, \theta_k^3, \theta_k^2, \theta_k^1\}$. From the polarization-modulated observations, we reconstruct the polarimetric light transport \mathcal{T}_p , a slice of the full light transport tensor \mathcal{T} , by solving a least-squares optimization problem. For spatio-polarimetric imaging, this is

$$\mathcal{T}_{sp}(s, s', :, :) = \arg \min_{\mathcal{T}_p} \sum_{k=1}^K \left(I^{\Theta_k}(s, s') - f_{sp}(\Theta_k, s, s') \right)^2, \quad (21)$$

and for temporal-polarimetric imaging we solve

$$\mathcal{T}_{pt}(s, s, :, t) = \arg \min_{\mathcal{T}_p} \sum_{k=1}^K \left(I^{\Theta_k}(s, t) - f_{pt}(\Theta_k, s, t) \right)^2. \quad (22)$$

Specifically, we solve these problems per pixel with the least-squares solver inspired by Baek et al. [2020] but with a different image formation model, see Supplemental Document. The polarimetric recovery completes the proposed acquisition of the spatio-polarimetric transport tensor \mathcal{T}_{sp} and the temporal-polarimetric transport tensor \mathcal{T}_{pt} . See Figure 8 and Figure 9 for the reconstruction results.

7.1 Polarimetric Decomposition

We further analyze the polarimetric light transport for the basic polarization properties of a material: depolarization, retardance, and diattenuation. These three polarimetric properties have been traditionally studied in optics [Lu and Chipman 1996] and other application domains, notably biology [Ghosh and Vitkin 2011]. We bring

these concepts to graphics and vision with an application to light transport analysis. Depolarization describes how much light loses its polarization state, turning into unpolarized light for both linear polarization and circular polarization. As such, this quantity can be used to characterize properties of scattering media such as density and particle size. Retardance describes the phase delay between two orthogonal polarization states including vertical/horizontal linear polarization and right/left circular polarization, therefore revealing the difference of the polarization-dependent refractive index of a material. Retardance has been used to examine plastic stretch, defects on thin films, and birefringent materials. As a generalized concept of dichroism, diattenuation describes the differential attenuation of light's orthogonal polarization states including linear and circular polarization. Representing the three basic properties of depolarization, retardance, and diattenuation as individual Mueller matrices \mathbf{M}_Δ , \mathbf{M}_R , and \mathbf{M}_D , we decompose the polarimetric light transport \mathbf{T}_p into a product of three Mueller matrices following the optical Lu-Chipman decomposition [1996]

$$\mathbf{T}_p = \mathbf{M}_\Delta \mathbf{M}_R \mathbf{M}_D. \quad (23)$$

Once decomposed, we analyze the low-dimensional projections of the Mueller matrices as three scalar quantities of polarizance, retardance, and diattenuation for intuitive analysis [Ghosh and Vitkin 2011]. Polarizance P is the degree of polarization of exiting light after light-matter interactions of an unpolarized incident light in a scene as

$$P = \frac{1}{\mathbf{T}_p(0,0)} \sqrt{\mathbf{T}_p^2(1,0) + \mathbf{T}_p^2(2,0) + \mathbf{T}_p^2(3,0)}. \quad (24)$$

Retardance R is the phase shift between the orthogonal polarization states of light

$$R = \cos^{-1} \left(\frac{\text{tr}(\mathbf{M}_R)}{2} - 1 \right), \quad (25)$$

where $\text{tr}(\cdot)$ is the trace operator. Diattenuation D characterizes how incident light polarization affects the transmission of light

$$D = \frac{1}{\mathbf{T}_p(0,0)} \sqrt{\mathbf{T}_p^2(0,1) + \mathbf{T}_p^2(0,2) + \mathbf{T}_p^2(0,3)}. \quad (26)$$

See Figure 8 for decomposition.

8 EXPERIMENTAL PROTOTYPE

We describe our two experimental prototypes for spatio-polarimetric and coaxial temporal-polarimetric acquisition. See Figure 7 for schematic diagrams and system photographs. For details, including calibration procedures, we refer to the Supplemental Document.

8.1 Spatio-polarimetric Setup

Spatial Imaging. To acquire spatially-resolved light transport, we use a DMD projector (Lightcraft LC3000-G2-Pro) and a machine-vision camera (FLIR GS3-U3-51S5C) with an objective lens of focal length 25 mm (Fujinon HF25HA-1S 25mm f/1.4). The DMD projector emits unpolarized light with a programmable projection pattern $M_p(s')$ and the camera has a programmable region-of-interest feature to implement the imaging pattern $M_c(s)$. While arbitrary projection and imaging patterns can be implemented, we demonstrate two effective mask configurations for epipolar and non-epipolar imaging proposed in O'Toole et al. [2012]. To this end, we geometrically

calibrate the projector and the camera by capturing checkerboard images at different poses, obtaining the intrinsic and distortion parameters of the devices, and the fundamental matrix between them in the undistorted coordinates. This allows us to compute epipolar/non-epipolar patterns in the distorted coordinates that can be used as $M_p(s')$ and $M_c(s)$.

Polarimetric Imaging. We use a pair of achromatic QWPs (Thorlabs AQWP10M-580) and LPs (Newport 10LP-VIS-B) in front of the projector and the camera. Unpolarized light emitted from the DMD projector passes through a LP mounted on a manual rotary stage (Newport RM25A) along a vertical axis for the DRR ellipsometry. For our learned rotating ellipsometry, the LP is mounted on a motorized rotary stage (Thorlabs KPRM1E). The linearly polarized light is then modulated by a QWP mounted on another motorized rotary stage so that the fast axis is aligned with the polarization axis of the LP at the zero base. The second module of a QWP and a LP is installed in front of the camera, modulating the polarization state of the returning light from a scene. For DRR ellipsometry, we rotate the QWPs in front of the projector and the camera 36 times at angles $\theta_k = 5k$ and $\theta'_k = 25k$, where $k \in \{0, \dots, 35\}$. For our learned ellipsometry, we use the optimized angles of the polarizing optics as shown in Figure 6. We obtain all the results with the DRR setting except for the evaluation experiments for the learned ellipsometry (Figures 10 and 11). For each angle index k , we obtain epipolar and non-epipolar images with the precomputed spatial masks M_p and M_c .

Acquisition Parameters. We set the acquisition parameters to obtain accurate reconstructions of light transport tensors at acceptable acquisition time. Our camera is set to have 600 ms exposure and captures 16bit raw images. The projector is configured to emit a 8bit intensity image and we use zero or 255 depending on whether the projector mask is turned off or on. For each polarizing-optics angle set $\Theta_k = \{\theta_k^1, \theta_k^2, \theta_k^3, \theta_k^4\}$, we project a sequence of the precomputed masks for epipolar and non-epipolar imaging. To accelerate the capture speed, we reconfigure the region of interest for the camera depending on the camera mask. In addition, the capture script is written in Python with multi-threading programming to reduce any capture-to-save delay. This results in an acquisition time of one hour for the epipolar and non-epipolar polarimetric captures with DRR ellipsometry.

8.2 Coaxial Temporal-polarimetric Setup

Temporal Dimension. To achieve high temporal resolution, we use a picosecond pulsed laser unit (Edinburgh instruments EPL-635) with an optical wavelength of 635 nm. The laser pulse is linearly polarized and passes through a non-polarizing beamsplitter (Thorlabs BS013) and is directed to a scene by using a mirror galvanometer (Thorlabs GVS012). As our system is coaxial, we detect light coming from a scene using the same path by the non-polarizing beamsplitter. A single-photon avalanche diode (MPD series SPAD) is used with a time-correlated single photon counter (PicoQuant TimeHarp 260 PICO) for controlling the timings of the SPAD and the laser. Overall, our system has a temporal resolution of 25 picoseconds.




Fig. 9. (a) Coaxial monochromatic imagers integrates temporal dimension without polarimetric resolution, capturing only intensity of light $\sum_t \mathcal{T}(s, s, 0, 0, t)$. (b) Transient imaging disentangles the time integration enabling light-in-flight imaging $\mathcal{T}(s, s, 0, 0, t)$. (c) Our coaxial temporal-polarimetric imaging method captures not only incoming intensity, but the complete polarization change at the picosecond temporal resolution as $\mathcal{T}(s, s, :, :, t)$. We visualize the 4×4 temporal-polarimetric light transport of the travel time 3.65 ns. This temporal-polarimetric decomposition enables us to see (c) strong specular reflection that preserves the circular polarization state of incident light as $\mathcal{T}(s, s, :, :, t)$ and (d) specular components that invert the circular polarization of light similar to mirror as $-\mathcal{T}(s, s, :, :, t)$. This reveals the material difference of the hand-painted patterns on the doll as indicated by the red arrows.

Polarimetric Dimension. Similar to the spatio-polarimetric setup, we employ two QWPs and a LP. Note that the laser light is linearly polarized to the vertical orientation and passes through the QWP mounted on a motorized rotation stage before entering the beam-splitter. The same configuration of the QWP and the LP is installed for the detection path after the beamsplitter reflection. We use the DRR configuration for all temporal-polarimetric experiments.

Acquisition Parameters. We use an exposure time of 10 ms for each galvo mirror position, which we sample at 128×128 points. We repeat this for every angle configuration of the polarizing optics $\Theta_{k \in \{1, \dots, 36\}}$. For a single scene, the total acquisition time is 7 hours.

9 ASSESSMENT

We experimentally validate the proposed method by acquiring and analyzing spatio-temporal polarimetric transport tensors. For the first time, we demonstrate complete polarimetric light transport capture combined with temporal and spatial transport imaging. The resulting transport tensors provide detailed information about a scene both in terms of geometry and material composition, which are challenging to acquire for existing temporal and spatial transport methods. We also experimentally validate our learned rotating ellipsometry and its polarizer-array variant.

9.1 Spatio-Polarimetric Light Transport Analysis

Spatial light transport methods are capable of isolating light paths satisfying a certain transport geometry connecting a projector pixel and a detector pixel. Existing methods untangle light transport into different path geometries by masking transport components in the

illumination path, using a projector, and on the sensor path, with per-pixel masks on a camera. Commonly used mask patterns include high-frequency checkerboards for direct-global separation [Nayar et al. 2006] and epipolar/non-epipolar imaging configurations [O’Toole et al. 2012]. Although the proposed projector-camera ellipsometry method allows for arbitrary mask patterns, we demonstrate epipolar and non-epipolar imaging with complete polarimetric sampling. As such, for the first time, we demonstrate the complete polarimetric decomposition of epipolar and non-epipolar light transport. Figure 8 shows the decomposed spatio-polarimetric light transport of crystals that reveal direct and indirect transport further resolved in polarization state. This allows us to acquire specular reflections that maintain the horizontal linear polarization of incident light for direct and indirect reflections in $\sum_t \mathcal{T}(s, s'_e, 1, 1, t)$ and $\sum_t \mathcal{T}(s, s'_n, 1, 1, t)$. Direct surface reflection on the crystals are revealed in the circular-polarization inverting component $-\sum_t \mathcal{T}(s, s'_e, 3, 3, t)$ with epipolar scanning that could allow for examining the surface profiles of the transparent crystal. Further, we decomposed the epipolar/non-epipolar polarimetric light transport to obtain diattenuation, polarizance, and retardance. This reveals the intrinsic properties of various crystals, notably the birefringence of calcite and quartz in the retardance estimates. Note that complete Mueller matrix acquisition with spatial probing is essential for this spatio-polarimetric decomposition and its application to estimating crystal properties.

9.2 Temporal-Polarimetric Light Transport Analysis

Existing transient imaging methods capture light propagating through a scene at high temporal resolution on the order of picoseconds [Heide et al. 2013; O’Toole et al. 2017; Velten et al. 2013]. This acquisition

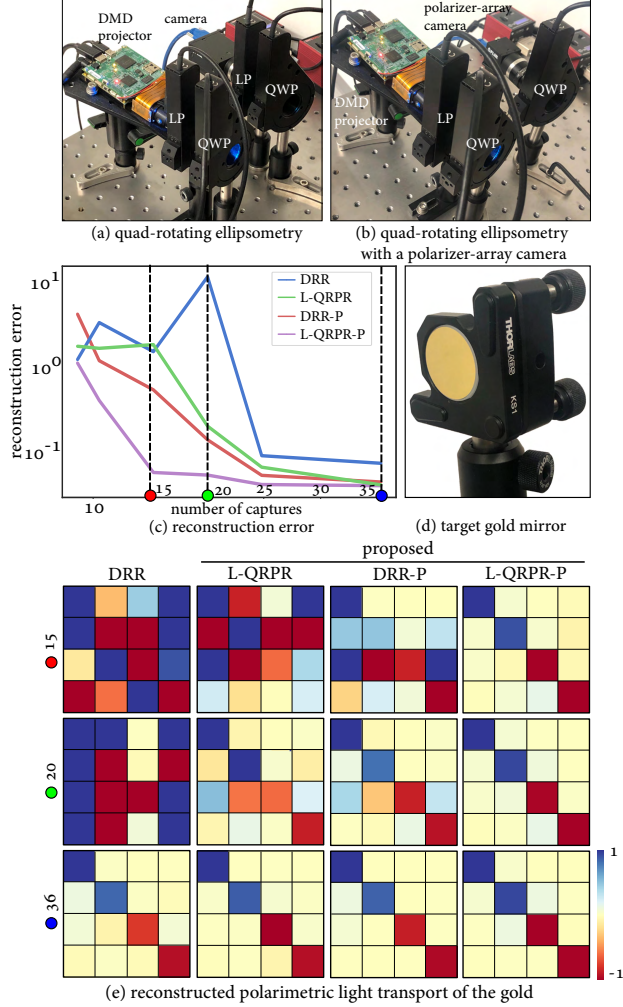


Fig. 10. We modify our spatio-polarimetric imager to demonstrate the proposed learned ellipsometry method of (a) the quad-rotating ellipsometry and (b) its extension with the polarizer-array sensor. (d) We capture the polarimetric light transport of an unprotected gold mirror. (c) & (e) Our method outperforms the conventional DRR method, achieving same reconstruction accuracy with fewer captures. We use acronyms L-QRPR for the learned quad-rotating polarizer retarder, DRR-P for the dual-rotating retarder with a polarizer sensor, and L-QRPR-P for the learned quad-rotating polarizer retarder with a polarizer sensor. Note that compared to the conventional DRR method, we reduce the number of captures from 36 to 15 for the L-QRPR-P method.

approach reveals the scene geometry both in the direct sight [Heide et al. 2018] and outside the direct line of sight [Velten et al. 2012]. However, material-dependent scene understanding has been challenging for transient imaging due to temporal blur and temporal metamerism that light rays undergo since different material-dependent scattering events have similar travel distances [Wu et al. 2014]. For the first time, we demonstrate that polarimetric transient imaging improves the decomposition capability of material-dependent light transport. Figure 9 shows the decomposed coaxial temporal-polarimetric light transport. For each time bin t , we obtain a complete polarimetric decomposition. This allows us to

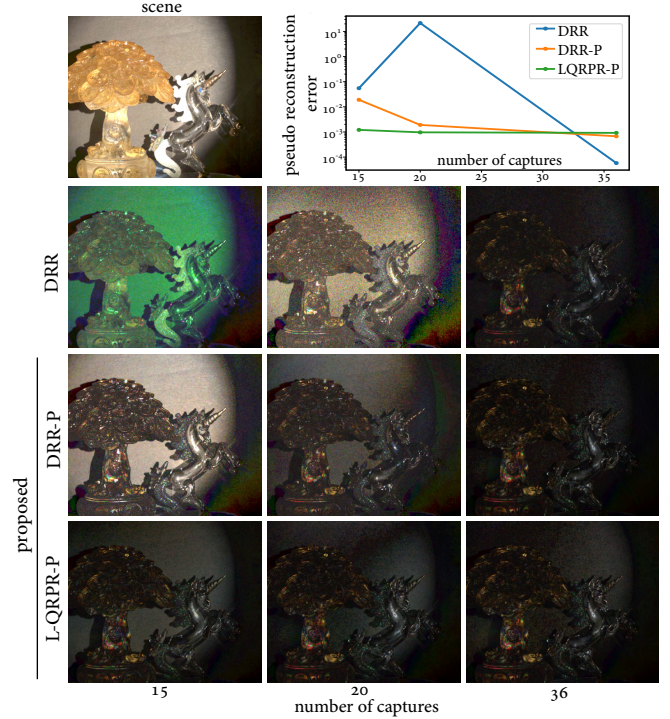


Fig. 11. The proposed learned ellipsometry method efficiently acquires complex polarimetric light transport with fewer captures than existing methods. For quantitative evaluation, we use pseudo ground-truth data from 50 DRR captures [Azzam 1978]. The proposed L-QRPR-P method provides accurate reconstructions despite only using 15 captures. We visualize the (2,3)-th component of the polarimetric transport $-\sum_{s',t} \mathcal{T}(s, s', 2, 3, t)$.

separate specular reflections that maintaining circular polarization of incident light. Furthermore, we distinguish the specular components that invert or maintain the incident circular polarization with $\mathcal{T}(s, s, :, t)$ and $-\mathcal{T}(s, s, :, t)$.

9.3 Learned Ellipsometry and Polarizer-array Sensor

In addition to the synthetic assessment in Figure 6, we experimentally evaluate our proposed ellipsometry methods compared to the conventional DRR method [Azzam 1978]. To this end, we modify our spatio-polarimetric setup by mounting the linear polarizers on the motorized rotation stages to implement the quad-rotating ellipsometry. For the polarizer-array extension, we replace the CMOS camera with the polarizer-array camera (FLIR Blackfly S BFS-U3-51S5PC). For details on geometric and polarimetric calibrations, we refer to the Supplemental Document. We capture an unprotected gold mirror (Thorlabs PF10-03-M03) which has known polarimetric light transport as shown in Figure 10. We confirm that our proposed ellipsometry reduces the number of captures required for acquiring polarimetric light transport compared to the conventional DRR technique [Azzam 1978]. We also evaluate the ellipsometric methods (DRR, DRR-P, L-QRPR-P) on a scene with unseen materials, see Figure 11. To this end, we use the polarizer-array setup and obtain pseudo ground-truth data using 50 DRR captures [Azzam 1978]. Note that DRR inputs can be acquired by sampling the on-sensor polarizers oriented at zero degree, resulting in pixel-aligned inputs




Fig. 12. (a) We reconstruct the shape of a dragon statue with strong subsurface scattering and interreflections. (g) & (h) Spatial light transport of non-epipolar and epipolar imaging reveals the approximate indirect and direct components. (c) Conventional structured light (SL) methods, such as (b) graycode probing, assume diffuse appearance and struggle with strong subsurface scattering. (d) Shape-from-polarization (SfP) methods estimate surface normals with ambiguity in azimuthal angle and refractive index of the material [Atkinson and Hancock 2006]. (e) Combining structured light with the shape-from-polarization method resolves the polarimetric-normal ambiguity but the interreflections and subsurface scattering corrupt estimated normals. See red arrows. (k) Exploiting both geometric cues with epipolar capture eliminates the ambiguity in existing methods and allows us to reconstruct fine geometry detail. (f) We sprayed anti-reflective diffuse particles to the dragon statue, obtaining (l) the pseudo ground-truth geometry by using structured light scanning. We compute the reconstruction errors of surface normals with this registered ground truth, reported in the subfigure caption headings, revealing that our method also quantitatively outperforms existing methods. We refer to the Supplemental Document for further results.

to the DRR-P and the L-QRPR-P methods. We again observe that L-QRPR-P provides reliable reconstructions of polarimetric transport with only 15 captures while DRR fails. We refer to the Supplemental Document for qualitative evaluation of all ellipsometry methods and details on the spatio-polarimetric acquisition with the proposed ellipsometry method.

10 APPLICATIONS

We demonstrate applications of the proposed method for shape reconstruction with subsurface scattering, seeing through scattering media, multi-bounce reflection analysis, and resolving metamerism.

10.1 Shape Reconstruction with Subsurface Scattering

Reconstructing the geometry of an object with strong subsurface scattering poses a challenge for existing 3D imaging systems. The proposed spatio-polarimetric imaging method solves this problem by jointly exploiting polarimetric and spatial information that encodes information about a scene geometry. Figure 12 shows the reconstructed geometry for a dragon object with severe subsurface scattering and object interreflections. We note that conventional structured-light methods, such as graycode patterns, assume diffuse surface properties, and, as such, suffer in presence of strong subsurface scattering and interreflections on the object, failing to resolve fine surface structure. Epipolar structured light methods approximately capture direct reflections in order to mitigate such problems. However, epipolar structured light approaches *fail with strong subsurface scattering* that does not fulfill the epipolar constraint. Shape from polarization [Atkinson and Hancock 2006] provides geometric cues of the surface with π azimuthal ambiguity. Although these cues allow for the recovery of surface normals for convex objects, interreflections can still cause severe errors for normal estimation. Even combining structured light and shape from polarization, which

resolves the polarization-normal ambiguity, suffers from such inter-reflection artifacts [Baek et al. 2018]. In contrast, we obtain robust geometry from structured-light normals and polarization normals using epipolar-polarimetric captures. As such, the proposed method is capable of exploiting the intertwined relationship between the epipolar spatial domain and polarization domain. We observe improved geometric reconstruction compared to [Baek et al. 2018] at the neck and torso of the dragon indicated by the red arrows in Figure 12. To quantitatively assess reconstruction quality, we have acquired the pseudo ground-truth geometry using structured light after spraying diffuse particles onto the object. In Figure 12, we report the reconstruction errors of surface normals of each method with this ground-truth measurement after registration. The proposed epipolar structured polarimetric scanning qualitatively and quantitatively outperforms existing methods as validated by the reported normal estimation error. For details on the reconstruction method and additional results, we refer to the Supplemental Document.

10.2 Seeing through Scattering Media

Removing the effect of volumetric scattering is an open challenge in imaging and computer vision. We demonstrate imaging through scattering media to see objects of interest through scattering media such as fog and milky water in macroscopic scenes. In the presence of scattering in Figure 13(a), light transport \mathcal{T} from a source to a detector is modeled as a sum of the light \mathcal{T}_b back-scattered by the scattering and the ballistic object-scattered light \mathcal{T}_o reflected on the object surfaces: $\mathcal{T} = \mathcal{T}_b + \mathcal{T}_o$. Descattering then refers to the process of recovering the object-scattered component \mathcal{T}_o from the integrated transport \mathcal{T} . Existing methods have shown that the structure in polarimetric light transport can aid the recovery in this ill-posed problem. A key intuition of existing approaches is that

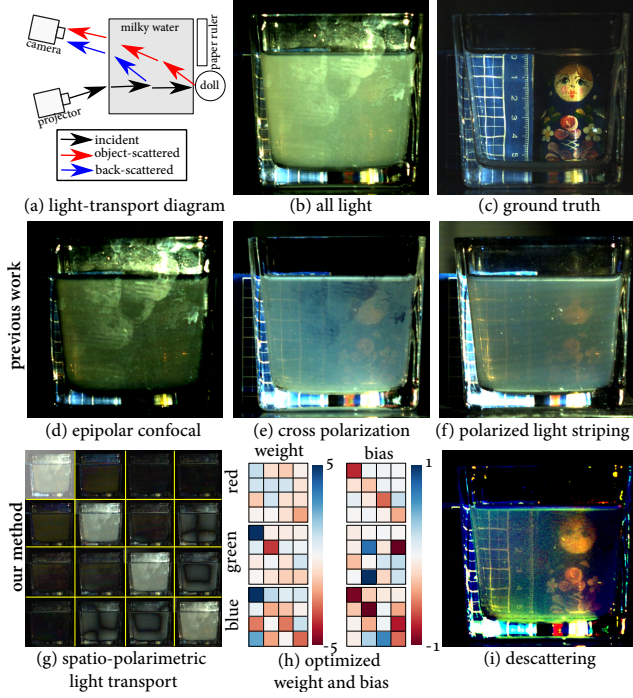


Fig. 13. Experimental seeing through scattering media. (a) Experimental configuration where milky water in a glass pot is in between the target objects and the imaging setup. (b) With severe backscatter, conventional imaging methods that integrate backscattered light and ballistic photons have limited visibility of (c) light interacting with the object. (d) Epipolar confocal imaging [O’Toole et al. 2012] removes scattering outside of the epipolar plane while epipolar scattering is ignored. (e) Cross-polarization imaging [Treibitz and Schechner 2008] relies on the orthogonal polarization states of linear/circular components for back-scattered and object-scattered light. (f) Polarized light striping [Gupta et al. 2008] combines line scanning and degree-of-polarization analysis on the linear polarization component, further improving descattering results. (g) We take a step forward by capturing the complete epipolar polarimetric light transport $\sum_t \mathcal{T}(s, s_e, :, t)$. (h) Our optimized weights and biases for polarimetric components reveal the contributions of all spatio-polarimetric information for descattering, leading to (i) the improved visibility of the hidden scene.

the back-scattered component \mathcal{T}_b tends to maintain the polarization of incident light, while the object-scattered component \mathcal{T}_o is depolarized [Gupta et al. 2008; Treibitz and Schechner 2008]. As such, capturing two images with parallel and orthogonal states of the linear polarizers allows us to filter out linearly polarized light which is assumed to be the backscattered light. While this approach reduces some backscatter, it relies on incomplete polarization analysis; only one linearly-polarized angle is investigated. Other linear polarization angles, circular, and elliptical polarization are ignored. Existing methods also have explored spatial probing [Fuchs et al. 2008; Gupta et al. 2008; O’Toole et al. 2012] for descattering but independently of the polarimetric analysis.

We overcome the limitations of existing approaches by imaging through scattering media with a variant of the proposed spatio-polarimetric imaging system. Our system allows us to capture the complete polarimetric change in scattering media, in contrast to

the existing incomplete polarization-based descattering methods, while simultaneously performing epipolar probing. To recover the undisturbed latent image (without the scattering media present), we solve an optimization problem to estimate the optimal weight and bias of the spatio-polarimetric light transport for descattering, that is

$$\begin{aligned} & \text{minimize}_{W, b} \|S_{\text{syn}} - S_{\text{gt}}\|_2, \\ S_{\text{syn}}^{k \in \{R, G, B\}} &= \sum_{i=0}^3 \sum_{j=0}^3 \left\{ W_{ijk} \odot \left(\sum_t \mathcal{T}_k(s, s_e, i, j, t) + b_{ijk} \right) \right\}, \quad (27) \end{aligned}$$

where W and b are the 4×4 matrices for trichromatic channels globally applied to all pixels, \mathcal{T}_k is the transport tensor at k -th color channel, and S_{gt} is the ground-truth image without scattering. We solve this optimization problem with a quasi-Newton L-BFGS optimizer [Liu and Nocedal 1989]. Figure 13 shows that our method provides improved visibility of the doll and the ruler compared to existing approaches. It is interesting to note that the optimized weight and bias perform linear operations on all polarimetric elements including linear, circular, and elliptical components across the 4×4 polarimetric Mueller matrix. This shows that complete polarization sensing combined with the spatial epipolar scanning enables effective descattering.

10.3 Untangling Multi-Bounce Light Transport

The light transport in a complex scene featuring diverse materials can include multiple interreflection between individual scene surface points. In conventional intensity images, the superposition of this multi-bounce transport results in complex diffusion and specular caustics that are challenging to analyze with conventional intensity images. We demonstrate that adding spatio-polarimetric and temporal-polarimetric separation resolves this ambiguity. Figure 14 shows that we can decouple interreflections between the specular coins, diffuse floor, specular disco ball, and colored diffuse planes. By exploiting the complicated spatio-polarimetric transport response, we separate the specular interreflections between the coins and diffuse interreflection between the colored diffuse planes with the disco ball. This can be explained by the fact that specular interreflection often flips the circular polarization, while diffuse interreflection is encoded in the linear polarization. Also, we separate reflections from the statue/lion and the mirror with specular and diffuse decomposition. This fine-grained decomposition of light transport is enabled by our joint analysis of spatial and polarimetric light transport.

10.4 Breaking Metamerism with Polarization

The proposed polarized light transport analysis in the temporal and spatial dimensions allows us to resolve metamerism. Conventional cameras integrate over large regions of the spectral irradiance, which can make objects with different reflectance appear identical as a result of this integration. Acquiring polarimetric light transport allows us to resolve this ambiguity without having to acquire hyperspectral data with high spectral resolution. Figure 15 validates that the proposed polarimetric analysis allows us to identify different materials for various objects including fake/real fruits and

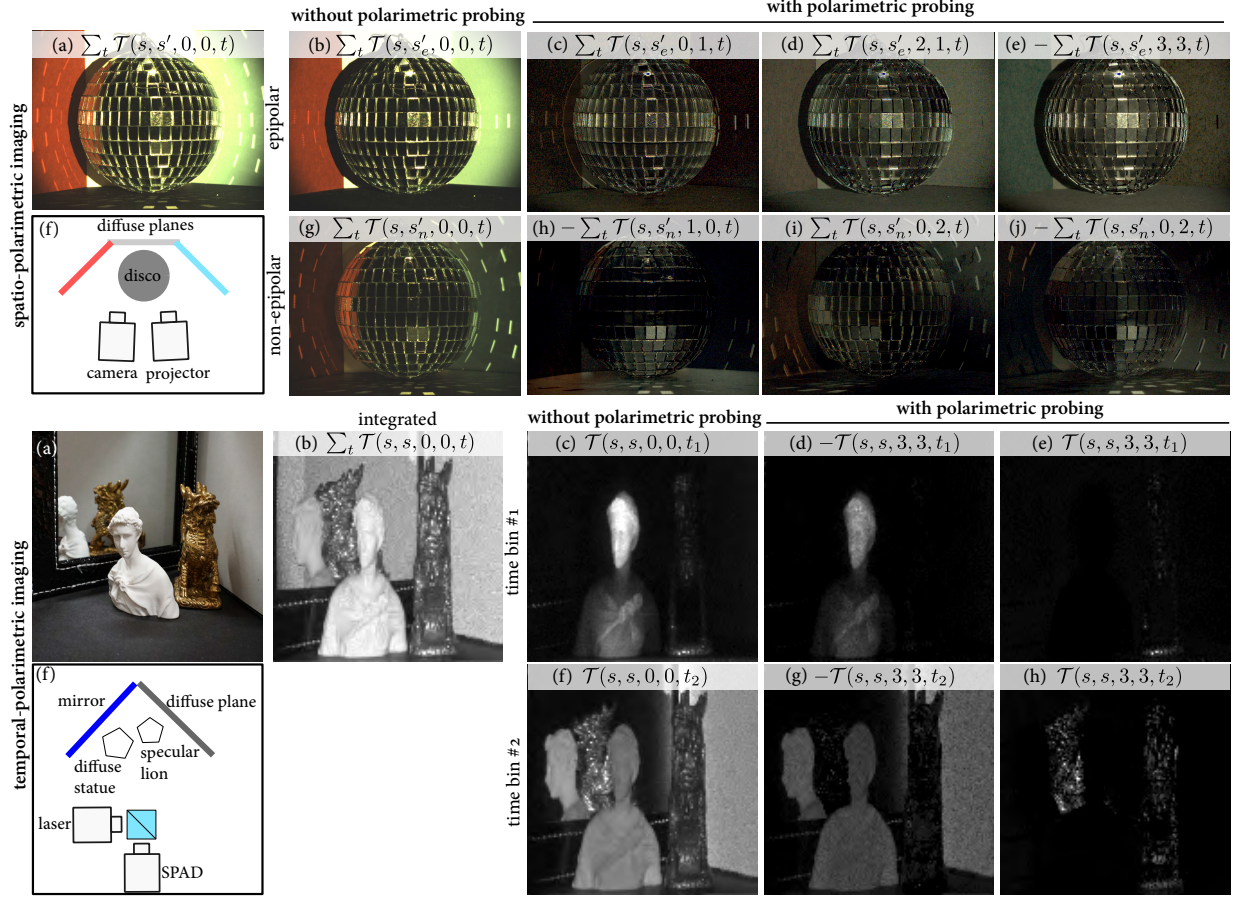


Fig. 14. Spatio-polarimetric Imaging: (a) & (f) We untangle complex reflections between a specular disco ball and diffuse planes using our spatio-polarimetric imaging approach. (b) Epipolar and (g) non-epipolar spatial imaging partially separates direct and indirect reflections between the objects [O’Toole et al. 2012]. Our epipolar polarimetric probing (c) isolates the remaining specular reflection on the diffuse planes ($\sum_t \mathcal{T}(s, s'_e, 0, 1, t)$), (d) removes the specular reflection ($\sum_t \mathcal{T}(s, s'_e, 2, 1, t)$), (e) and captures the diffuse direct interreflections between the red and blue planes as revealed by the changed colors of the planes ($-\sum_t \mathcal{T}(s, s'_e, 3, 3, t)$). Our polarimetric non-epipolar imaging provides further decomposition of the indirect reflections, allowing us to untangle (h) the reflection from the planes and the disco ball only to the floor ($-\sum_t \mathcal{T}(s, s'_n, 1, 0, t)$), (i) & (j) the reflections with different directions from the disco ball to the planes ($\sum_t \mathcal{T}(s, s'_n, 0, 2, t)$ and $-\sum_t \mathcal{T}(s, s'_n, 0, 2, t)$). Temporal-polarimetric imaging: (a) & (b) Conventional photography integrates along the time dimension, capturing the steady-state equilibrium of light transport. (c) & (f) Transient imaging captures the temporal dimension, revealing light in flight. Our temporal-polarimetric probing further decomposes this temporal transport depending on material properties, effectively distinguishing (d) & (g) the diffuse transient reflections as $\mathcal{T}(s, s, 3, 3, t_1)$ and (e) & (h) specular transient reflections from metallic surfaces as $\mathcal{T}(s, s, 3, 3, t_2)$.

flowers by discovering scattering that are “hidden” in conventional and spatial captures. Here we utilize the epipolar spatial dimension with 45 degree linear-polarization maintaining component, $-\sum_t \mathcal{T}(s, s'_e, 2, 2, t)$, where s and s'_e are the pairs of pixels in the camera and the projector within the same epipolar plane. In this tensor slice, the plastic fake objects of fruits and flowers have low intensity because the captured light returns from the direct surface reflection from the surface paint. In contrast, the real organic materials present rich subsurface scattering revealing its unique organic material structure.

11 DISCUSSION

In this work, we develop a method for capturing a light transport tensor decomposed into its spatial, temporal, and polarimetric dimensions. This comes at the cost of sequential acquisition and reduced light efficiency. We partly mitigated the light efficiency issue in our spatio-polarimetric imaging prototype by spatially binning 2×2 pixels, resulting in an effective pixel pitch of $12.9 \mu\text{m}$. Departing from sequential acquisition and developing photon-efficient capture are exciting areas of future research. In particular, one potential direction that we envision is to combine the proposed method with liquid crystal retarders which could substantially reduce capture time, possible even below a minute. Our current implementation of the polarimetric reconstruction requires 12 sec. and 0.3 sec. to process the spatio-polarimetric and the temporal-polarimetric data,




Fig. 15. The proposed spatio-polarimetric imaging approach allows us to resolve metamerism observed in (a) real and fake fruits and flowers. (b) & (c) Conventional spatial epipolar and non-epipolar imaging methods cannot identify the subtle difference between the real and the fake. (d) Our spatio-polarimetric imaging reveals that the fake plastic objects show minimal subsurface scattering. The right apple, right pear, and the bottom flowers are fake objects. These objects contain low-intensity polarimetric scattering effects as shown in the epipolar 45-degree linear-polarization component $-\sum_t \mathcal{T}(s, s'_e, 2, 2, t)$.

respectively. Accelerating the per-pixel reconstruction which is separable on a modern GPU could unlock real-time capture and reconstruction in the future. Combining the learned principal components of polarimetric transport and polarimetric reconstruction is another exciting direction which could facilitate robust reconstruction with fewer captures. Energy-efficient spatial probing [O'Toole et al. 2015] is further exciting direction for future research.

12 CONCLUSION

We propose a computational light transport imaging method that allows, for the first time, to capture full polarization information along with temporal and spatial resolution. Our approach hinges on a novel tensor-based light transport theory that jointly models temporal, spatial and polarimetric dimensions. We analyze the low-rank polarimetric embeddings and propose a data-driven ellipsometry that learns to exploit this low-rank structure. We instantiate our approach with two imaging systems for coaxial temporal-polarimetric imaging and spatio-polarimetric imaging. The proposed polarimetric probing method provides in-depth analysis on material and geometry characteristics of a scene, enabling unprecedented decomposition of temporal-polarimetric and spatio-polarimetric light transport. We validate the proposed method with five applications: reconstruction of shape in the presence of subsurface scattering, seeing through scattering media, analyzing multiple light-bounces, untangling metamerism, and analyzing crystal birefringence. We outperform existing methods for *all* applications. We envision that

the proposed method will inspire broad applications in computer graphics, vision, and beyond.

ACKNOWLEDGMENTS

We thank our reviewers for their invaluable comments. We also thank Julian Knodt for helping the path-tracer implementation and Ethan Tseng for proofreading. Felix Heide was supported by an NSF CAREER Award (2047359) and a Sony Faculty Innovation Award.

REFERENCES

- Edward H Adelson, James R Bergen, et al. 1991. *The plenoptic function and the elements of early vision*. Vol. 2. MIT.
- Gianluca Agresti and Pietro Zanuttigh. 2018. Combination of spatially-modulated ToF and structured light for MPI-free depth estimation. In *European Conference on Computer Vision (ECCV) Workshops*. 0–0.
- Hans Arwin. 2011. Application of ellipsometry techniques to biological materials. *Thin Solid Films* 519, 9 (2011), 2589–2592.
- Gary A Atkinson and Edwin R Hancock. 2006. Recovery of surface orientation from diffuse polarization. *IEEE Transactions on Image Processing (TIP)* 15, 6 (2006), 1653–1664.
- 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. *Journal of the Optical Society of America A (JOSA A)* 33, 7 (2016), 1396–1408.
- Yunhao Ba, Alex Gilbert, Franklin Wang, Jinfa Yang, and Rui Chen. 2020. Deep Shape from Polarization. In *European Conference on Computer Vision (ECCV)*.
- Seung-Hwan Baek, Diego Gutierrez, and Min H Kim. 2016. Birefractive stereo imaging for single-shot depth acquisition. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 194.
- Seung-Hwan Baek, Daniel S Jeon, Xin Tong, and Min H Kim. 2018. Simultaneous acquisition of polarimetric SVBRDF and normals. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 268–1.
- 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 Transactions on Graphics (TOG)* 39, 4 (2020).
- Kai Berger, Andrea Weidlich, Alexander Wilkie, and Marcus Magnor. 2011. Modeling and verifying the polarizing reflectance of real-world metallic surfaces. *IEEE Computer Graphics and Applications (CGA)* 32, 2 (2011), 24–33.
- Tongbo Chen, Hendrik PA Lensch, Christian Fuchs, and Hans-Peter Seidel. 2007. Polarization and phase-shifting for 3D scanning of translucent objects. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1–8.
- Edward Collett. 2005. Field guide to polarization. SPIE Bellingham, WA.
- Aniket Dashpute, Chandani Anand, and Mukul Sarkar. 2018. Depth resolution enhancement in time-of-flight cameras using polarization state of the reflected light. *IEEE Transactions on Instrumentation and Measurement* 68, 1 (2018), 160–168.
- Valentin Deschaintre, Yiming Lin, and Abhijeet Ghosh. 2021. Deep Polarization Imaging for 3D Shape and SVBRDF Acquisition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 15567–15576.
- Christian Fuchs, Michael Heinz, Marc Levoy, Hans-Peter Seidel, and Hendrik PA Lensch. 2008. Combining confocal imaging and descattering. In *Computer Graphics Forum (CGF)*, Vol. 27. Wiley Online Library, 1245–1253.
- Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A Wilson, and Paul Debevec. 2010. Circularly polarized spherical illumination reflectometry. *ACM Transactions on Graphics (TOG)* 29, 6 (2010), 162.
- Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. 2011. Multiview face capture using polarized spherical gradient illumination. *ACM Transactions on Graphics (TOG)* 30, 6 (2011), 129.
- Abhijeet Ghosh, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul Debevec. 2008. Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics (TOG)* 27, 5 (2008), 139.
- Nirmalya Ghosh and Alex I Vitkin. 2011. Tissue polarimetry: concepts, challenges, applications, and outlook. *Journal of Biomedical Optics* 16, 11 (2011), 110801.
- Ioannis Gkioulekas, Anat Levin, Frédo Durand, and Todd Zickler. 2015. Micron-scale light transport decomposition using interferometry. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–14.
- Mohit Gupta, Amit Agrawal, Ashok Veeraraghavan, and Srinivasa G Narasimhan. 2011. Structured light 3D scanning in the presence of global illumination. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 713–720.
- Mohit Gupta, Srinivasa G Narasimhan, and Yoav Y Schechner. 2008. On controlling light transport in poor visibility environments. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1–8.
- Mohit Gupta, Shree K Nayar, Matthias B Hullin, and Jaime Martin. 2015. Phasor imaging: A generalization of correlation-based time-of-flight imaging. *ACM Transactions on Graphics (TOG)* 34, 5 (2015), 1–18.

- Felix Heide, Steven Diamond, David B Lindell, and Gordon Wetzstein. 2018. Sub-picosecond photon-efficient 3D imaging using single-photon sensors. *Scientific Reports* 8, 1 (2018), 1–8.
- Felix Heide, Matthias B Hullin, James Gregson, and Wolfgang Heidrich. 2013. Low-budget transient imaging using photonic mixer devices. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.
- Felix Heide, Matthew O’Toole, Kai Zang, David B Lindell, Steven Diamond, and Gordon Wetzstein. 2019. Non-line-of-sight imaging with partial occluders and surface normals. *ACM Transactions on Graphics (TOG)* 38, 3 (2019), 1–10.
- Tomoki Ichikawa, Matthew Purri, Ryo Kawahara, Shohei Nobuhara, Kristin Dana, and Ko Nishino. 2021. Shape From Sky: Polarimetric Normal Recovery Under the Sky. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 14832–14841.
- Mikhail Itskov. 2007. *Tensor algebra and tensor analysis for engineers*. Springer.
- Adrian Jarabo, Julio Marco, Adolfo Munoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A framework for transient rendering. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 1–10.
- Adrian Jarabo, Belen Masia, Julio Marco, and Diego Gutierrez. 2017. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* 1, 1 (2017), 65–79.
- Gerald Earle Jellison Jr, Donovan N Leonard, Lawrence M Anovitz, Chad M Parish, Eliot D Specht, and TM Rosseel. 2018. Crystallographic orientation of uniaxial calcite and dolomite determined using reflection generalized ellipsometry. *Journal of Applied Physics* 124, 22 (2018), 223102.
- Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. 2015. Polarized 3D: High-Quality Depth Sensing with Polarization Cues. In *IEEE International Conference on Computer Vision (ICCV)*. 3370–3378.
- Achuta Kadambi, Refael Whyte, Ayush Bhandari, Lee Streeter, Christopher Barsi, Adrian Dorrington, and Ramesh Raskar. 2013. Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–10.
- Alankar Kotwal, Anat Levin, and Ioannis Gkioulekas. 2020. Interferometric transmission probing with coded mutual intensity. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 74–1.
- Hirofumi Kubo, Suren Jayasuriya, Takafumi Iwaguchi, Takuya Funatomi, Yasuhiro Mukaigawa, and Srinivasa G Narasimhan. 2018. Acquiring and characterizing plane-to-ray indirect light transport. In *IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–10.
- Dong C Liu and Jorge Nocedal. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical programming* 45, 1 (1989), 503–528.
- Shih-Yau Lu and Russell A Chipman. 1996. Interpretation of Mueller matrices based on polar decomposition. *Journal of the Optical Society of America A (JOSA A)* 13, 5 (1996), 1106–1113.
- Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Eurographics Conference on Rendering Techniques*. 183–194.
- Tomohiro Maeda, Achuta Kadambi, Yoav Y Schechner, and Ramesh Raskar. 2018. Dynamic heterodyne interferometry. In *IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–11.
- Alkhazur Manakov, John Restrepo, Oliver Klehm, Ramon Hegedus, Elmar Eisemann, Hans-Peter Seidel, and Ivo Ihrke. 2013. A reconfigurable camera add-on for high dynamic range, multispectral, polarization, and light-field imaging. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 47–1.
- Andreas Meuleman, Seung-Hwan Baek, Felix Heide, and Min H Kim. 2020. Single-Shot Monocular RGB-D Imaging Using Uneven Double Refraction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2465–2474.
- Kazuhiro Morimoto, Andrei Ardelean, Ming-Lo Wu, Arin Can Ulku, Ivan Michel Antolovic, Claudio Bruschini, and Edoardo Charbon. 2020. Megapixel time-gated SPAD image sensor for 2D and 3D imaging applications. *Optica* 7, 4 (Apr 2020), 346–354.
- Koki Nagano, Graham Fyfe, Oleg Alexander, Jernej Barbic, Hao Li, Abhijeet Ghosh, and Paul E Debevec. 2015. Skin microstructure deformation with displacement map convolution. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 109–1.
- Nikhil Naik, Achuta Kadambi, Christoph Rhemann, Shahram Izadi, Ramesh Raskar, and Sing Bing Kang. 2015. A light transport model for mitigating multipath interference in time-of-flight sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 73–81.
- Shree K Nayar, Gurunandan Krishnan, Michael D Grossberg, and Ramesh Raskar. 2006. Fast separation of direct and global components of a scene using high frequency illumination. In *ACM Transactions on Graphics (TOG)*. 935–944.
- 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.
- 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.
- 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), 1–17.
- Matthew O’Toole, Supreeth Achar, Srinivasa G Narasimhan, and Kiriakos N Kutulakos. 2015. Homogeneous codes for energy-efficient illumination and imaging. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–13.
- Matthew O’Toole, Felix Heide, David B Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. 2017. Reconstructing transient images from single-photon sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1539–1547.
- Matthew O’Toole, Felix Heide, Lei Xiao, Matthias B Hullin, Wolfgang Heidrich, and Kiriakos N Kutulakos. 2014a. Temporal frequency probing for 5D transient analysis of global light transport. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.
- Matthew O’Toole and Kiriakos N Kutulakos. 2010. Optical computing for fast light transport analysis. *ACM Transactions on Graphics (TOG)* 29, 6 (2010), 1–164.
- Matthew O’Toole, David B Lindell, and Gordon Wetzstein. 2018. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature* 555, 7696 (2018), 338–341.
- Matthew O’Toole, John Mather, and Kiriakos N Kutulakos. 2014b. 3d shape and indirect appearance by structured light transport. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3246–3253.
- Matthew O’Toole, Ramesh Raskar, and Kiriakos N Kutulakos. 2012. Primal-dual coding to probe light transport. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 39–1.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. (2017).
- Adithya Pediredla, Ashok Veeraraghavan, and Ioannis Gkioulekas. 2019. Ellipsoidal path connections for time-gated rendering. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.
- Christoph Peters, Jonathan Klein, Matthias B Hullin, and Reinhard Klein. 2015. Solving trigonometric moment problems for fast transient imaging. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–11.
- David A Ramsey and Kenneth C Ludema. 1994. The influences of roughness on film thickness measurements by Mueller matrix ellipsometry. *Review of Scientific Instruments* 65, 9 (1994), 2874–2881.
- Jérémy Riviere, Ilya Reshetouski, Luka Filipi, and Abhijeet Ghosh. 2017. Polarization Imaging Reflectometry in the Wild. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 206:1–206:14.
- 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.
- Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R Marschner, Mark Horowitz, Marc Levoy, and Hendrik PA Lensch. 2005. Dual photography. *ACM Transactions on Graphics (TOG)* (2005), 745–755.
- Kenichiro Tanaka, Yasuhiro Mukaigawa, and Achuta Kadambi. 2020. Polarized non-line-of-sight imaging. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2136–2145.
- Kenichiro Tanaka, Yasuhiro Mukaigawa, Yasuyuki Matsushita, and Yasushi Yagi. 2013. Descattering of transmissive observation using parallel high-frequency illumination. In *IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–8.
- Tali Treibitz and Yoav Y Schechner. 2008. Active polarization descattering. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 31, 3 (2008), 385–399.
- Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mouni G Bawendi, and Ramesh Raskar. 2012. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature communications* 3, 1 (2012), 1–8.
- Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mouni Bawendi, Diego Gutierrez, and Ramesh Raskar. 2013. Femtophotography: capturing and visualizing the propagation of light. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–8.
- Gordon Wetzstein, Ivo Ihrke, Douglas Lanman, and Wolfgang Heidrich. 2011. Computational plenoptic imaging. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 2397–2426.
- Refael Whyte, Lee Streeter, Michael J Cree, and Adrian A Dorrington. 2015. Resolving multiple propagation paths in time of flight range cameras using direct and global separation methods. *Optical Engineering* 54, 11 (2015), 113109.
- Di Wu, Andreas Velten, Matthew O’Toole, Belen Masia, Amit Agrawal, Qionghai Dai, and Ramesh Raskar. 2014. Decomposing global light transport using time of flight imaging. *International Journal of Computer Vision (IJCV)* 107, 2 (2014), 123–138.
- Rihui Wu, Adrian Jarabo, Jinli Suo, Feng Dai, Yongdong Zhang, Qionghai Dai, and Diego Gutierrez. 2018. Adaptive polarization-difference transient imaging for depth estimation in scattering media. *Optics Letters* 43, 6 (2018), 1299–1302.
- Rihui Wu, Jinli Suo, Feng Dai, Yongdong Zhang, and Qionghai Dai. 2016. Scattering robust 3D reconstruction via polarized transient imaging. *Optics letters* 41, 17 (2016), 3948–3951.
- Zexiang Xu, Jannik Boll Nielsen, Jiyang Yu, Henrik Wann Jensen, and Ravi Ramamoorthi. 2016. Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM*

Transactions on Graphics (TOG) 35, 6 (2016), 1–12.