

All-photon Polarimetric Time-of-Flight Imaging

Seung-Hwan Baek^{1,2}
Princeton University¹

Felix Heide¹
POSTECH²

Abstract

Time-of-flight (ToF) sensors provide an image modality fueling diverse applications, including LiDAR in autonomous driving, robotics, and augmented reality. Conventional ToF imaging methods estimate depth by sending pulses of light into a scene and measuring the ToF of the first-arriving photons directly reflected from a scene surface without any temporal delay. As such, all photons following this first response are typically considered as unwanted noise. In this paper, we depart from the principle of using first-arriving photons and propose an all-photon ToF imaging method that relies on the temporal-polarimetric analysis of first- and late-arriving photons which encode rich scene information in terms of geometry and material. To this end, we propose a novel temporal-polarimetric reflectance model, an efficient capture method, and a reconstruction method that exploits the temporal-polarimetric changes of light reflected by the surface and sub-surface reflection. The proposed all-photon polarimetric ToF imaging method allows us to acquire depth, surface normals, and material parameters of a scene by utilizing all photons captured by the system, whereas conventional ToF imaging only obtains coarse depth from the first-arriving photons. We validate our method in simulation and experimentally with a prototype system.

1. Introduction

Depth cameras have equipped us with the ability to acquire scene geometry that is essential for 3D vision both *directly* as a geometric input to a perception or reconstruction method, and *indirectly*, as training data for supervised learning in rich 3D vision datasets [9, 11, 42]. As such, depth sensor have fueled diverse applications across domains, including autonomous vehicles, virtual and augmented reality, mobile photography, and robotics. A large body of prior work explores different approaches to extract depth, including parallax [37], defocus [23], double refraction [4], and correlation imaging [22]. In particular, *pulsed* ToF cameras [32] have been rapidly adopted across application domains for their high depth precision, low baseline, long

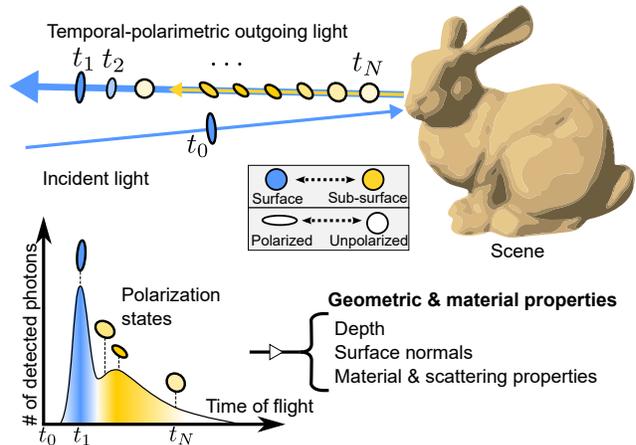


Figure 1. All-photon polarimetric ToF imaging. In contrast to conventional ToF imaging methods that analyze the first-reflected light, we propose to capture and exploit both early and late photons and untangle them via polarization. We analyze temporal-polarimetric characteristics of surface and sub-surface reflections that enable us to infer scene parameters (depth, normals, and material parameters) from *all* captured photons.

range and robustness to scene texture.

ToF cameras directly measure the round-trip return time of light emitted into a scene. This requires high-speed sensors and electronics that can time-tag photon arrival events, such as single-photon avalanche diode (SPAD) and time-correlated single-photon counting (TSCPC) electronics. By repeating pulses and recording arrival events hundreds of thousands of times per measurement, these pulsed ToF cameras are capable of recording an entire temporal histogram for each pixel, reporting the number of photons detected by the sensor for each travel time bin.

Photons that arrive at different times at a single pixel can be broadly categorized into two temporally distinct components that sum up to the total measured photons. The first group of photons originates from the light directly reflected off from an object surface. Therefore, this surface reflection allows for estimating the object depth directly from its ToF measurement. Surface reflection often has higher intensity in comparison to the other component. As a result, one common practice for ToF depth imaging is to estimate

the ToF through peak-finding over the photon counts. In contrast to the surface reflection, the second group of photons comes from the light that penetrates the object surface, undergoes sub-surface scattering, and leaves the surface towards the detector. This sub-surface scattered light arrives at the sensor with a longer travel time than the light originating from surface reflections. *At first glance, sub-surface scattering appears to harm accurate depth estimation* as it disturbs the true ToF by an additional unknown scattering time. Existing ToF methods typically discard this sub-surface observation and only rely on the surface reflection.

In this work, we treat this sub-surface scattering component as a signal instead of noise in estimating geometric and material properties of a scene. The proposed *all-photon ToF imaging method combines ToF imaging with polarization analysis*, exploiting the unique temporal-polarimetric signature of surface and sub-surface reflections, as shown in Figure 1. The proposed method hinges on a novel temporal-polarimetric bidirectional reflectance distribution function (BRDF) model that describes both surface and sub-surface reflections in temporal and polarimetric domains. We implement this BRDF model with differentiable operators allowing us to pose *scene reconstruction as an inverse rendering problem*. To efficiently acquire high-dimensional polarimetric-temporal data, we also propose a light-efficient temporal-polarimetric acquisition approach that captures the temporal-polarimetric signals of a scene combined with a learned optical ellipsometry. We demonstrate that all-photon polarimetric ToF imaging can record diverse scene information, including depth, normals, and material parameters from temporal-polarimetric measurements – both in simulation and experimentally. Specifically, we make the following contributions.

- We devise an all-photon polarimetric ToF imaging method that combines ToF imaging and polarization analysis to jointly exploit surface and sub-surface reflected photons for scene reconstruction.
- We present the first temporal-polarimetric BRDF model which characterizes surface and sub-surface reflection based on micro-facet theory and the Stokes-Mueller formalism.
- We propose an efficient temporal-polarimetric acquisition approach tailored for polarimetric imaging with polarized-laser illumination.
- We validate our method in simulation and with an experimental prototype for capturing depth, normals, and material parameters.

Limitations As a result of the mechanical rotation of the polarizing optics and sequential scanning, the acquisition of the temporal-polarimetric data takes five hours for

each scene. We note that the limited number of temporal-polarimetric photons and low sensor quantum efficiency makes the acquisition challenging. In the future, accelerating the polarimetric capture with electronically-controllable liquid crystal modulators instead of mechanically rotating polarizers, or using a polarization array filters placed on photon-efficient SPAD sensor arrays, may eliminate the need for sequential acquisition. Our experimental prototype employs a laser module with low average power of 0.13 mW available to us. We note that higher-power laser sources are commercially available which could increase the signal-to-noise ratio of the measurements for a given exposure time budget.

2. Related Work

ToF Imaging Modern ToF imaging methods can be broadly categorized into correlation and pulsed approaches. Correlation ToF methods emit intensity-modulated light into a scene and *indirectly* estimate the round-trip time of the light returned back to the sensor by performing intensity demodulation [22, 49]. In contrast, pulsed ToF methods *directly* measure ToF of light in an emitted pulse. These methods combine picosecond-pulsed illumination with time-resolved sensors, such as a SPAD, synchronized with the laser using TSCPC electronics [15, 26]. Not only have pulsed ToF approaches been deployed broadly for LiDAR systems in autonomous vehicles [40], robotics [21], and space missions [1], but recently in consumer mobile phones such as the iPhone 12 Pro. Typical LiDAR sensors extract depth by reporting a detection event for the *first-arriving* photons among all photons reflected from a scene surface. To this end, existing methods implement peak finding via analog thresholding [40], spatial coherence [31], temporal coherence [41], and learned priors [24]. This acquisition principle also means that *photons arriving later are ignored*, along with rich material and geometric information that these photons carry about the scene. Researchers have attempted to utilize this hidden information by modeling a more comprehensive *temporal* response of a scene with analytic temporal BRDF models [35, 36, 48]. However, temporal-only analysis restricts them to only extracting scene depth from the *fitted* first-arriving photons, while the later arriving photons are only used for obtaining a better fit of the first-arriving photons. In contrast, we analyze *all* detected photons by jointly analyzing their polarization and temporal characteristics to extract geometric and material properties of a scene.

Polarimetric Imaging A rich body of work has explored using polarization of light for various visual-computing tasks. Early works rely on cross-polarization, which equips illumination and detector with orthogonal polarization filters. This allows for separating polarized reflection from

unpolarized reflection, or vice versa [28, 38]. This principle has led to many applications such as seeing through scattering media [44] and diffuse-specular separation [13, 25]. Recently, polarization cameras with spatially multiplexed linear polarizers have been used for the acquisition of surface normals [2, 19]. These methods explicitly or implicitly rely on polarimetric BRDFs that couple local geometry to a polarization response [6, 7]. Combining polarization and ToF imaging has largely been an unexplored field. Callenberg et al. [8] investigated polarization-difference imaging by configuring a two-illumination correlation ToF camera in a cross-polarization mode. Most relevant to our work, Baek et al. [5] combine polarization imaging and ToF imaging to capture the temporal-polarimetric response of a scene. Their method is *unable to extract geometry and reflectance* from the measurements due to the absence of a temporal-polarimetric BRDF model that links the scene parameters to the observation. We close this gap in this work and introduce *the first temporal-polarimetric BRDF model*, which makes it possible to estimate scene parameters directly from temporal-polarimetric observations.

Polarization LiDAR Interestingly, combining polarization with ToF imaging has been studied for decades in geoscience [34]. Notably, researchers in cloud science have applied polarization analysis to LiDAR systems in order to examine cloud properties and aerosol concentration by sending out polarized laser pulses to the cloud and measure how much of returned light, at a given travel time, is depolarized [14, 33, 47]. This principle, first studied by Schotland et al. [39], has led to remarkable success, not only in cloud analysis but also in other fields, including biology [17] and ocean science [45]. The proposed method shares this motivation in that we combine polarization and ToF imaging, however, we tackle the problem of scene reconstruction.

3. Temporal-polarimetric Light Transport

The proposed method hinges on a novel analytical model for temporal-polarimetric light transport, which we introduce in this section. Our model is the first to describe both surface and sub-surface reflection based on the Stokes-Mueller formalism. We represent polarized light and reflectance with a Stokes vector and a Mueller matrix [10].

3.1. Temporal-polarimetric Reflectance

We introduce a temporal-polarimetric reflectance model $\mathbf{M}(\tau, \omega_i, \omega_o)$. This model describes how light polarization and intensity change when impinging on a surface with given incident and outgoing direction of light (ω_i and ω_o), and with temporal delay upon the arrival of incident light (τ). Specifically, we model reflectance \mathbf{M} as a sum of sur-

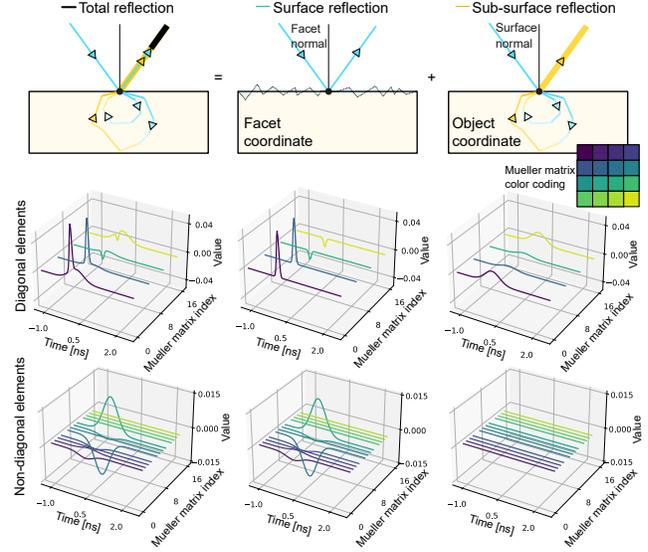


Figure 2. We introduce the first temporal-polarimetric reflectance model which explains the total reflection (left) as a sum of the temporal polarization change of surface (middle) and sub-surface reflections (right). Here, we visualize an example reflectance function in the form of a temporal-polarimetric Mueller matrix. Time-varying Mueller matrix elements are shown in the second and third rows for diagonal and non-diagonal elements, respectively.

face and sub-surface reflection (\mathbf{M}_s and \mathbf{M}_{ss})

$$\mathbf{M}(\tau, \omega_i, \omega_o) = \mathbf{M}_s(\tau, \omega_i, \omega_o) + \mathbf{M}_{ss}(\tau, \omega_i, \omega_o). \quad (1)$$

We describe both components next, see also Figure 2.

Surface Reflection Some of the incident photons are immediately scattered in a surface reflection, including both retro-reflection and specular reflection [29]. Following microfacet theory [43], we can model a surface as a collection of microfacets. Surface reflection originates from the microfacets oriented orthogonal to the halfway vector $\mathbf{h} = \frac{\omega_i + \omega_o}{2}$. To describe the polarimetric change induced by this microfacet reflection, existing work has relied on the Fresnel-reflection Mueller matrix \mathbf{F}_R . This is a function of the refractive index of the material η [6, 10, 18]. In our work, we also consider time in addition to polarization, which can induce a slight temporal delay of the surface reflection, e.g., originating from reflection between near surfaces or microfacets. This results in the depolarization of light, which we model with a time-varying attenuation Mueller matrix $\mathbf{D}^s(\tau)$ using polarization-dependent Gaussian distributions, that is $\mathbf{D}_i^s(\tau) = a_i^s \exp(-\frac{(\tau - \mu_i^s)^2}{2(\sigma_i^s)^2})$, where \mathbf{D}_i^s is the i -th diagonal element, a_i^s is the amplitude, and μ_i^s and $(\sigma_i^s)^2$ are the mean and variance, respectively. The matrix \mathbf{D}^s can attenuate the intensity and polarization components of incident light at different degrees. Combining this with Smith’s shading and masking function [16] and

the GGX normal distribution function [46], we arrive at the proposed temporal-polarimetric surface-reflection model

$$\mathbf{M}_s(\tau, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \frac{D(\theta_h; \sigma)G(\theta_i, \theta_o; m)}{4 \cos \theta_i \cos \theta_o} \mathbf{D}^{ss}(\tau) \mathbf{F}_R, \quad (2)$$

where $\theta_h = \cos^{-1}(\mathbf{h} \cdot \mathbf{n})$ is the halfway angle, \mathbf{n} is the normal, m is the roughness, and $\theta_i = \cos^{-1} \mathbf{n} \cdot \boldsymbol{\omega}_i$ and $\theta_o = \cos^{-1} \mathbf{n} \cdot \boldsymbol{\omega}_o$ are the incident/outgoing angles.

Sub-surface Reflection The second type of reflection originates from the light transmitted into the object. The transmitted light undergoes sub-surface scattering and eventually leaves the object. The Fresnel-transmission Mueller matrix \mathbf{F}_T [10] describes the polarization change of light passing through the interface between the two mediums (air and bulk material). Baek et al. [6] combined this Fresnel matrix with a perfect depolarization Mueller matrix to describe sub-surface reflection with an assumption that light becomes completely unpolarized after sub-surface scattering. However, in reality, depolarization by sub-surface scattering is highly correlated with the path length of light; sub-surface scattering with longer path lengths results in more depolarization [12]. Modeling this behavior is essential for the proposed temporal-polarimetric imaging method because path length determines the temporal delay. To this end, we model sub-surface scattering with time-varying depolarization. We use time-varying Gaussian distributions as the diagonal elements of a depolarization Mueller matrix $\mathbf{D}^{ss}(\tau)$ similar to the surface-reflection model with the parameters a_i^{ss} , μ_i^{ss} , σ_i^{ss} . For sub-surface scattering, we also consider the coordinates of Stokes vectors. As a Stokes vector of light can only be defined in a basis coordinate system [10], we use the halfway coordinate approach [7]. For more details on the coordinate conversion, see the Supplemental Document.

With all components described above, we define our temporal-polarimetric sub-surface reflection model

$$\mathbf{M}_{ss}(\tau, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \mathbf{C}_{n \rightarrow o} \mathbf{F}_T^o \mathbf{D}^{ss}(\tau) \mathbf{F}_T^i \mathbf{C}_{i \rightarrow n}, \quad (3)$$

where $\mathbf{C}_{i \rightarrow n}$ and $\mathbf{C}_{n \rightarrow o}$ are the coordinate-conversion Mueller matrices from the halfway coordinates of incident and outgoing light to the surface normal coordinate [10]. Here, \mathbf{F}_T^i and \mathbf{F}_T^o are the Fresnel transmission Mueller matrices for the incident and outgoing light which both depend on the refractive index of the material η . Note that our surface-reflection model does not require coordinate conversion, as surface reflection is already defined in halfway coordinates following microfacet theory [43].

3.2. Temporal-polarimetric Rendering

The proposed temporal-polarimetric reflectance model returns a Mueller matrix $\mathbf{M}(\tau, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ at a temporal delay

τ and incident and outgoing directions, $\boldsymbol{\omega}_i$ and $\boldsymbol{\omega}_o$. Incorporating this Mueller matrix into the rendering equation [20] describes the Stokes vector of the outgoing light at each temporal delay as

$$\begin{aligned} \mathbf{s}_o(\boldsymbol{\omega}_o, t) &= \int_{S^2} \int_0^t \cos \theta_i \mathbf{M}(\tau, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \mathbf{s}_i(\boldsymbol{\omega}_i, t - \tau) d\tau d\boldsymbol{\omega}_i \\ &= \int_{S^2} \int_0^t \mathbf{H}(\tau, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \mathbf{s}_i(\boldsymbol{\omega}_i, t - \tau) d\tau d\boldsymbol{\omega}_i, \end{aligned} \quad (4)$$

where \mathbf{H} is the Mueller matrix scaled by the cosine foreshortening. As such, the Stokes vector of the outgoing light (\mathbf{s}_i) is the function of time and propagation direction.

4. Polarimetric ToF Imaging

We now turn to developing a polarimetric ToF imaging system to capture a temporal-polarimetric scene response.

Coaxial ToF Imaging Although the proposed method is not restricted to a specific time-resolved sensing system, we opt to use a pulsed laser illumination and a synchronized SPAD sensor in a coaxial configuration typical to LiDAR systems [40]. This means that the illumination and detection share the same optical path to optically sample direct illumination with minimal indirect illumination by inter-reflection. This allows us to simplify the rendering equation (4) by removing the outer integral over direction, that is

$$\begin{aligned} I(t, \boldsymbol{\omega}) &= \left[\int_0^{t'} \mathbf{H}(\tau, \boldsymbol{\omega}, \boldsymbol{\omega}) \mathbf{s}_{\text{illum}}(\boldsymbol{\omega}, t' - \tau) d\tau \right]_0 \\ &= [\mathbf{H}(t', \boldsymbol{\omega}, \boldsymbol{\omega}) \mathbf{s}_{\text{illum}}(\boldsymbol{\omega}, 0)]_0, \end{aligned} \quad (5)$$

where $I(t, \boldsymbol{\omega})$ is the captured intensity of light with ToF t along the direction $\boldsymbol{\omega}$. The term $[\mathbf{x}]_0$ denotes the first element of a Stokes vector \mathbf{x} , c is the speed of light, and $\mathbf{s}_{\text{illum}}$ is the Stokes vector of the laser illumination. Here, t' is the shifted time by the ToF ($t' = t - \frac{2d}{c}$), where d is the Euclidean distance between the collocated laser at position $\mathbf{0}$, set to the origin without loss of generality, and a scene point $\mathbf{p} = \mathbf{0} + d\boldsymbol{\omega}$. In addition, the second integral over the temporal dimension is simplified as a result of using pulsed laser illumination.

Rotating Ellipsometry for Polarized Illumination

Starting with the coaxial ToF imaging systems described above, we acquire the complete polarization change induced by light-matter interaction using rotating optical ellipsometry [10]. Rotating optical ellipsometry commonly incorporates pairs of linear polarizers (LP) and quarter-wave plates (QWP) in front of an *unpolarized* illumination source and a detector [3, 5]. Multiple intensity measurements $I_{i \in \{1, \dots, N\}}$ are then captured by rotating

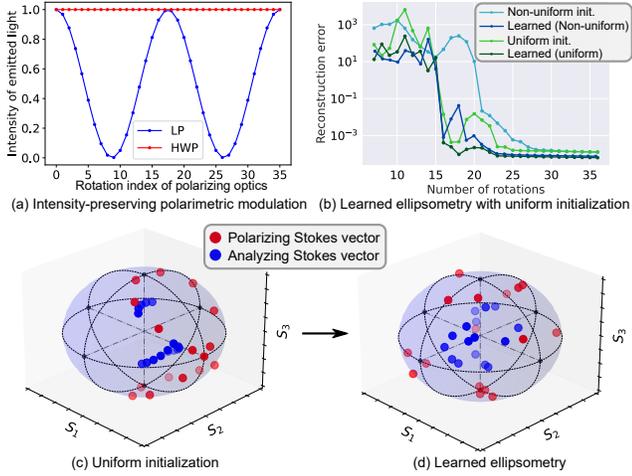


Figure 3. Energy-efficient Polarimetric ToF Imaging. (a) For a polarimetric capture, we rotate a HWP instead of a LP to preserve the light intensity of the polarized source. (b) We learn the rotation angles of the polarizing optics with uniform initialization (see text) which allows us to reconstruct polarimetric reflectance with fewer captures than conventional ellipsometry methods. We visualize the Poincaré sphere of the polarizing (c) and analyzing (d) Stokes vectors ($[\mathbf{P}_i]_{\forall,0}$, $[\mathbf{A}_i]_{0,\forall}$) before and after learning.

the quarter-wave plates [3], or together with rotating linear polarizers [5] to reduce the required number of captures.

Unfortunately, rotating linear polarizers in our ToF imaging setup results in *significant* loss of light energy, because our laser illumination itself is linearly polarized; the polarized laser illumination loses energy as it passes through the linear polarizer. Instead of a linear polarizer, we present a simple solution for this light-inefficient polarization modulation by placing a half-wave plate (HWP) in front of the laser illumination. This allows us to rotate the polarization axis of the laser illumination without sacrificing illumination intensity as shown in Figure 3(a).

As a result, our coaxial polarimetric ToF imaging system consists of a pulsed laser, HWP, and QWP, i.e., the illumination module, and QWP, LP, and SPAD forming an analyzer module. We model the image formation of the proposed method as

$$I_i(t, \omega) = [\mathbf{A}_i \mathbf{H}(t', \omega, \omega) \mathbf{P}_i \mathbf{s}_{\text{illum}}(\omega, 0)]_0, \quad (6)$$

where \mathbf{A}_i and \mathbf{P}_i are the i -th Mueller matrices of the analyzing optics and the polarizing optics defined as $\mathbf{A}_i = \mathbf{L}(\theta_i^4) \mathbf{Q}(\theta_i^3)$ and $\mathbf{P}_i = \mathbf{Q}(\theta_i^2) \mathbf{W}(\theta_i^1)$, with $\theta_i^{\{1,2,3,4\}}$ as the rotation angles of the polarizing module HWP and QWP and the analyzing module QWP and LP. Here, \mathbf{W} , \mathbf{Q} , and \mathbf{L} are the Mueller matrices of HWP, QWP, and LP [10]. Equation (6) formulates the temporal-polarimetric transport in our system from laser illumination via polarizing optics, scene transport, and analyzing optics, to SPAD detection. The polarimetric response of the galvo-mirror is cal-

ibrated by capturing an uncoated gold mirror with known response [7].

Learned Ellipsometry with Uniform Initialization

Next, we investigate how few intensity measurements I_i can be captured to faithfully reconstruct the Mueller matrix \mathbf{H} . This is critical for the proposed polarimetric ToF imaging method, as it performs high-dimensional probing via sequential galvo scanning and rotation of polarization optics, resulting in hour-long capture time. Recently, Baek and Heide [5] presented a learned ellipsometry which optimizes the rotation angles of polarizing optics using first-order data-driven optimization. In the same spirit, we learn the optimal rotation angles using a differentiable implementation of Equation (6). One modification we make is to use a different initialization scheme for the angles, which is essential to overcome local minima. Specifically, we initialize the rotation angles of the polarization optics to produce diverse polarization states, uniformly distributed over the Poincaré sphere [10]. Obtained with the proposed initialization scheme, the experiments in Figure 3(b) validate that the polarization states provide more reliable constructions than using the non-uniform initialization [5].

5. All-photon Scene Reconstruction

The proposed temporal-polarimetric image formation model explains how geometric and material parameters of a scene relate to our sensor measurements in the presence of surface and sub-surface reflections. This allows us to solve the corresponding inverse problem: Given the temporal-polarimetric sensor measurements I_i , we estimate a set of scene parameters Θ , including travel distance (d), surface normals (\mathbf{n}), refractive index (η), surface roughness (m), and scattering parameters ($a_i^{s/ss}$, $\mu_i^{s/ss}$, $\sigma_i^{s/ss}$): $\Theta = \{d, \mathbf{n}, \eta, m, a_i^{s/ss}, \mu_i^{s/ss}, \sigma_i^{s/ss}\}$. Note that conventional ToF imaging methods estimate travel distance by measuring the round-trip time of the first-returning photons with digital or analog peak-finding methods. In contrast, we do not discard photons that are arriving later and exploit the temporal-polarimetric information in the surface and sub-surface reflections to reconstruct a scene.

Ellipsometric Reconstruction We first convert the temporal measurements I_i for all the rotation angles of the polarizing optics into a time-varying Mueller matrix \mathbf{H}_{meas} . To this end, we invert the image formation model of Equation (6) in a least-squares manner following Baek et al. [7], solving the following optimization problem for each spatio-temporal pixel

$$\underset{\mathbf{H}_{\text{meas}}}{\text{minimize}} \sum_{i=1}^N (I_i - [\mathbf{A}_i \mathbf{H}_{\text{meas}} \mathbf{P}_i \mathbf{s}_{\text{illum}}]_0)^2. \quad (7)$$

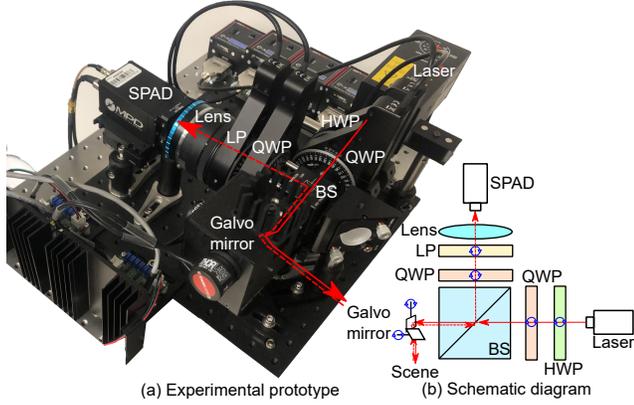


Figure 4. Experimental Setup. We built (a) an experimental prototype for polarimetric ToF imaging, allowing us to efficiently capture temporal-polarimetric images. Its optical path is shown in (b).

We implement the corresponding solver efficiently by parallelizing the pseudo-inverse matrix-vector product, see details in the Supplemental Document.

Scene Reconstruction Once we have obtained the Mueller matrix \mathbf{H}_{meas} for each temporal bin t' and coaxial incident/outgoing direction ω , we estimate the scene parameters Θ best explaining the Mueller matrix \mathbf{H}_{meas} , that is

$$\underset{\Theta}{\text{minimize}} \underbrace{\left\| \dot{\mathbf{W}}_p \odot (f(\Theta) - \mathbf{H}_{\text{meas}}) \right\|_1}_{\text{Data term}} + \underbrace{\left\| \dot{\mathbf{W}}_d \odot \nabla \mathbf{n} \right\|_1}_{\text{Regularization term}}, \quad (8)$$

where $\dot{\mathbf{x}}$ is the stacked vector of pixel quantities of variable \mathbf{x} for each temporal bin, spatial direction, and Mueller matrix element. Here, \odot is the Hadamard product, and ∇ is the spatial gradient operator. The data term evaluates the predicted Mueller matrix using the forward rendering function $f(\Theta) = \mathbf{H}(\tau, \omega, \omega; \Theta)$. The matrix \mathbf{W}_p weights the diagonal/non-diagonal Mueller matrix elements due to their different scales, see Supplemental Document for details. The regularization term penalizes steep surface normal changes in the spatial dimension. The matrix \mathbf{W}_d is a weighting matrix that penalizes the occlusion edges. With the objective in Equation (8) being differentiable with respect to the scene parameters Θ , we solve it using the first-order Adam optimization [30].

To circumvent the ill-posedness of Equation (8), we propose the three following modifications. First, we take a regional approach commonly used in inverse rendering methods [6, 27]. That is, while surface normals and travel distance are per-pixel quantities, we define refractive index, surface roughness, and scattering parameters as per-material quantities specified by a given cluster map. This reduces the total parameter count. We compute the cluster map based on the per-pixel average intensity using k -means

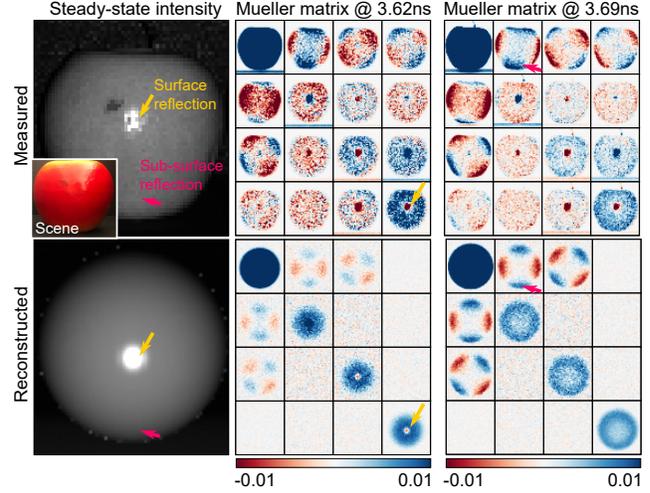


Figure 5. Temporal-Polarimetric BRDF Analysis. The proposed BRDF model is a parametric representation that is capable of describing temporal-polarimetric reflections for surface (orange arrows) and sub-surface scattering (pink arrows). For validation of the proposed BRDF model, we measured the temporal-polarimetric response of a spherical object (top) using our prototype. Our parametric BRDF model (bottom) accurately describes the measurements in terms of surface reflection (middle) and sub-surface reflection (right) at different time stamps.

clustering. Second, we incorporate the unit-norm and range constraints of each scene parameter to reduce the parameter search space with vector normalization and scaled sigmoid functions, see Supplemental Document for details. Third, we initialize the travel distance d and surface normals \mathbf{n} with conventional peak-finding outputs from the unpolarized ToF measurements corresponding to the first element of the Mueller matrix $[\mathbf{H}_{\text{meas}}]_{00}$.

6. Experimental Setup

Figure 4 shows our experimental prototype. We use a picosecond pulsed laser of optical wavelength 635 nm (Edinburgh Instruments EPL-635) and a SPAD sensor (MPD Series) with 25 ps temporal resolution. A time-correlation circuit (PicoQuant TimeHarp 260 PICO) synchronizes the laser and the SPAD. Light emitted from the laser passes through a half-wave plate (Thorlabs AHWP10M-600) and a quarter-wave plate (Thorlabs AQWP10M-580). A galvo mirror (Thorlabs GVS012) redirects the beam to a scene. We use a non-polarizing beamsplitter (Thorlabs BS013) to implement a coaxial configuration. Another quarter-wave plate and a linear polarizer (Newport 10LP-VIS-B) modulate the returning light. We use a 50 mm objective lens for the SPAD sensor, and mount the polarizing optics on motorized rotary stages (Thorlabs KPRM1E). We implement capture and reconstruction in Matlab and Pytorch.

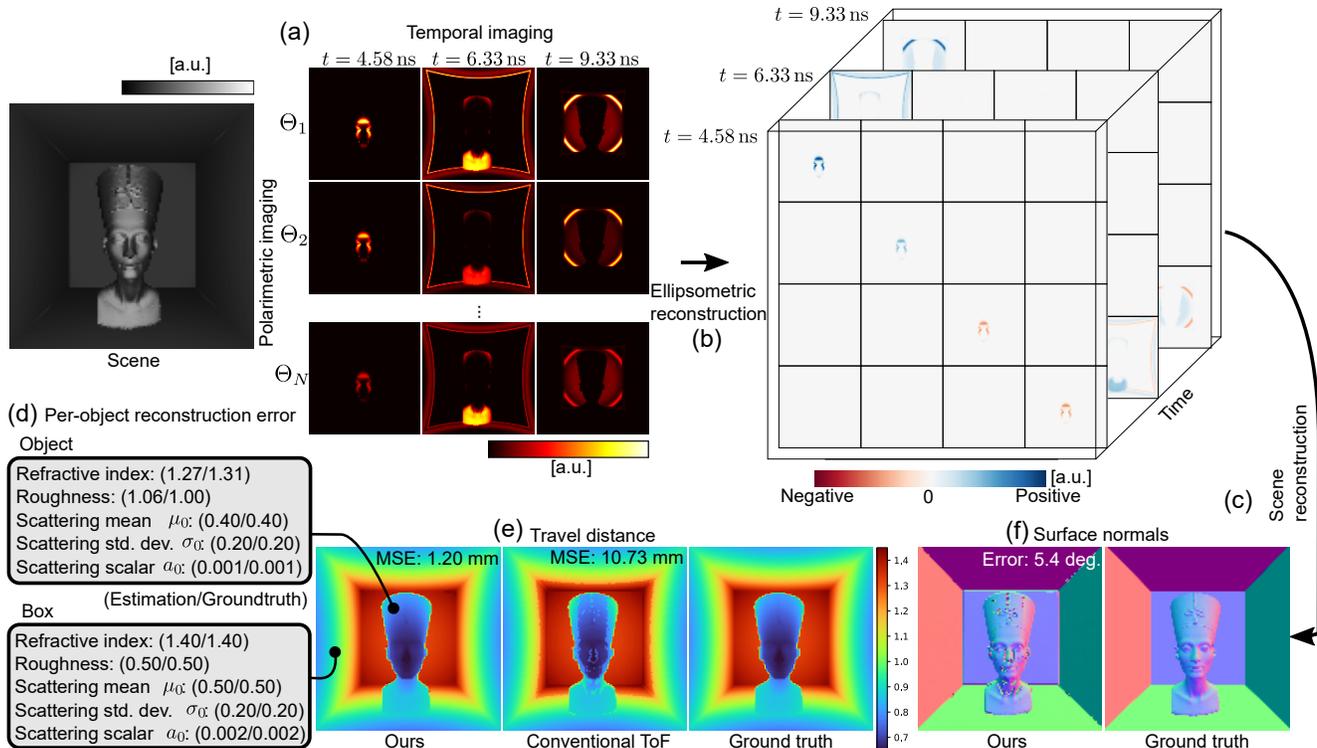


Figure 6. Synthetic Validation. We evaluate our all-photon ToF imaging on synthetic data simulating the entire process of (a) temporal-polarimetric capture, (b) ellipsometric reconstruction, and (c) scene reconstruction. The proposed method accurately estimates (c) per-object material parameters, (d) travel distance, and (e) surface normals.

7. Assessment

Temporal-polarimetric BRDF To evaluate the representation capacity of our temporal-polarimetric BRDF model, we capture a spherical object and reconstruct its corresponding Mueller matrix for each time bin: $\mathbf{H}_{\text{meas}}(t, \omega, \omega)$. We then find the scene parameters Θ of a spherical object, which best explains the observed temporal Mueller matrices, resulting in the reconstructed Mueller matrix $\mathbf{H}(t, \omega, \omega; \Theta)$. We further convert \mathbf{H} to another Mueller matrix \mathbf{H}' so that it follows similar noise statistics of the measured Mueller matrix \mathbf{H}_{meas} by simulating the sensor measurements I'_i with the reconstructed Mueller matrix \mathbf{H} and again reconstructing this Mueller matrix. The experiment in Figure 5 validates that the measured temporal-polarimetric data from our prototype matches simulated data from the proposed BRDF model. In particular, our model predicts the normal-dependent temporal-polarimetric structure of sub-surface reflection at the boundaries of the spherical object in $\mathbf{H}'_{\{i,0\},\{0,i\},\forall i}$. We also see the sign changes of the circular polarimetric sub-surface and surface reflection in $\mathbf{H}'_{3,3}$.

Synthetic Evaluation We evaluate the proposed method on three synthetic scenes shown in Figure 6. As ground-

truth scene parameters Θ for the synthetic scenes are known, we simulate the synthetic sensor measurements I_i using the image formation model from Equation (6). Here, we add Gaussian noise with a standard deviation of 10^{-4} to simulate measurement noise. With the simulated observations in hand, we solve Equation (8) by performing the proposed all-photon scene reconstruction, that is ellipsometric reconstruction of the temporal-polarimetric Mueller matrix \mathbf{H}_{meas} and scene reconstruction. Figure 6 validates that our reconstruction approach estimates geometric and material parameters accurately, both qualitatively and quantitatively. For additional results, we refer to the Supplemental Document.

Experimental Evaluation Next, we validate the proposed method on the additional real-world scenes shown in Figures 7 and 8. We reconstruct the temporal-polarimetric Mueller matrix \mathbf{H}_{meas} from experimental measurements I_i followed by finding the optimal scene parameters Θ . Our all-photon polarimetric ToF imaging reconstructs travel distance estimates that correspond to per-point depth and surface normals. We also acquire scattering parameters of the scene. Together with the obtained geometry, the proposed method accurately reconstructs the original measurements.

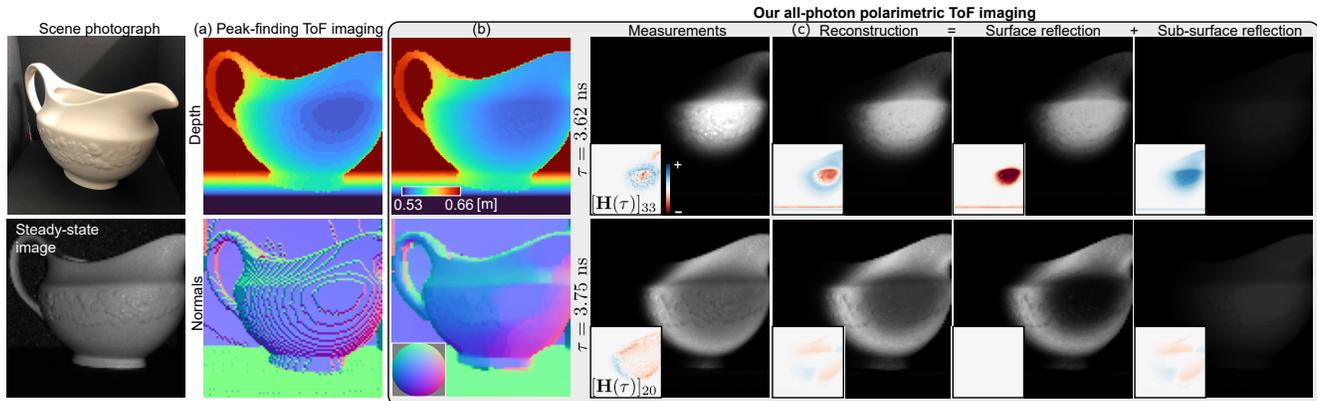


Figure 7. Experimental Validation. We validate our all-photon polarimetric ToF imaging on real-world scenes. (a) Conventional peak-finding ToF imaging methods acquire scene depth from first-reflected photons only, resulting in inaccurate geometry reflected in imperfect depth and surface normals estimates. In contrast, the proposed method relies on all photons, enabling (b) accurate geometry reconstruction and the recovery of material parameters by (c) decomposing light transport into surface and sub-surface components.

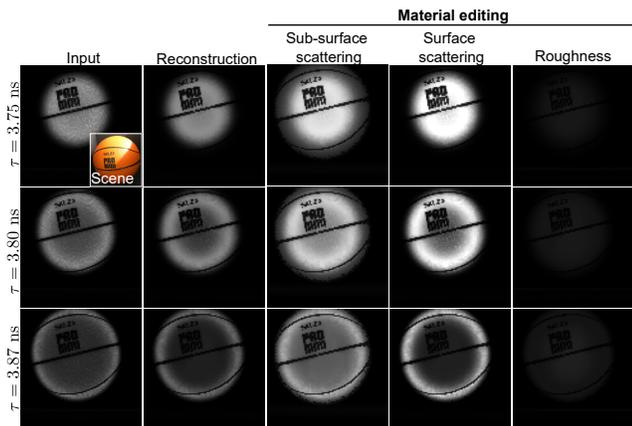


Figure 8. Material Editing. We edit the material parameters reconstructed from our polarimetric ToF imaging. Specifically, we change the scattering parameters of sub-surface and surface reflection as well as the surface roughness. We refer to the text for details.

Figure 7 shows the rendered measurements with our estimated parameters Θ and the corresponding sensor measurements. Our scene reconstructions and the proposed reflectance model also facilitate material editing, which we demonstrate with three editing examples in each column of Figure 8. Specifically, we first change the sub-surface scattering parameters by boosting $a_i \leftarrow 3a_i$, and changing its polarization-dependent temporal shift to $\mu \leftarrow [0.1\mu_0, 2\mu_1, 2\mu_2, 2\mu_3]$. This results in stronger sub-surface scattering. We then edit the surface scattering parameters in the same fashion, but only doubling $a_i \leftarrow 2a_i$. This mimics a material with stronger surface reflection. Lastly, increasing the surface roughness provides the appearance of rough surface with a dominating sub-surface reflection.

8. Conclusion

We introduce an all-photon polarimetric ToF imaging method, which extracts information from both surface and sub-surface reflection for acquiring geometry and scattering parameters. To this end, we propose the first temporal-polarimetric BRDF model, which relates sensor measurements to geometric and scattering parameters of a scene. The proposed learned ellipsometric ToF imaging system permits the acquisition of temporal-polarimetric data with energy-efficient polarimetric modulation. We reconstruct scene parameters from all photons instead of the first reflected ones. We validate the proposed method in simulation and experimentally. Although we validate the utility of temporal-polarimetric measurements, our experimental setup is restricted to sequential scanning by the mechanical rotation of the polarization optics available to us. Exciting future directions include the extension of the proposed system with electronically-controllable liquid crystal modulators, or combining SPAD array sensors with polarization filter arrays, which may pave the way towards single-shot temporal-polarimetric acquisition. Aside from making a step towards high-dimensional computational imaging, the approach may also pave the way for practical implementation such as gated imaging with polarization cues as a low-cost automotive imaging modality.

Acknowledgements

Felix Heide was supported by an NSF CAREER Award (2047359), a Sony Young Faculty Award, and a Project X Innovation Award.

References

- [1] Waleed Abdalati, H Jay Zwally, Robert Bindenschadler, Bea Csatho, Sinead Louise Farrell, Helen Amanda Fricker, David Harding, Ronald Kwok, Michael Lefsky, Thorsten Markus, et al. The icesat-2 laser altimetry mission. *Proceedings of the IEEE*, 98(5):735–751, 2010. 2
- [2] Gary A Atkinson and Edwin R Hancock. Recovery of surface orientation from diffuse polarization. *IEEE Transactions on Image Processing (TIP)*, 15(6):1653–1664, 2006. 3
- [3] RMA Azzam. Photopolarimetric measurement of the mueller matrix by fourier analysis of a single detected signal. *Optics Letters*, 2(6):148–150, 1978. 4, 5
- [4] Seung-Hwan Baek, Diego Gutierrez, and Min H Kim. Birefractive stereo imaging for single-shot depth acquisition. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016. 1
- [5] Seung-Hwan Baek and Felix Heide. Polarimetric spatio-temporal light transport probing. *ACM Transactions on Graphics (TOG)*, 2021. 3, 4, 5
- [6] Seung-Hwan Baek, Daniel S Jeon, Xin Tong, and Min H Kim. Simultaneous acquisition of polarimetric svbrdf and normals. *ACM Transactions on Graphics (TOG)*, 37(6):268–1, 2018. 3, 4, 6
- [7] Seung-Hwan Baek, Tizian Zeltner, Hyunjin Ku, Inseung Hwang, Xin Tong, Wenzel Jakob, and Min H Kim. Image-based acquisition and modeling of polarimetric reflectance. *ACM Transactions on Graphics (TOG)*, 39(4):139, 2020. 3, 4, 5
- [8] Clara Callenberg, Felix Heide, Gordon Wetzstein, and Matthias Hullin. Snapshot difference imaging using time-of-flight sensors. *ACM Transactions on Graphics (TOG)*, 2017. 3
- [9] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. 1
- [10] Edward Collett. Field guide to polarization. Spie Bellingham, WA, 2005. 3, 4, 5
- [11] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5828–5839, 2017. 1
- [12] Thomas A Germer. Evolution of transmitted depolarization in diffusely scattering media. *Journal of the Optical Society of America (JOSA)*, 37(6):980–987, 2020. 4
- [13] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A Wilson, and Paul Debevec. Circularly polarized spherical illumination reflectometry. pages 1–12, 2010. 3
- [14] Dennis H Goldstein. *Polarized light*. CRC press, 2017. 3
- [15] Felix Heide, Steven Diamond, David B Lindell, and Gordon Wetzstein. Sub-picosecond photon-efficient 3d imaging using single-photon sensors. *Scientific Reports*, 8(1):1–8, 2018. 2
- [16] Eric Heitz. Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques*, 3(2):32–91, 2014. 3
- [17] J Alex Huffman, Anne E Perring, Nicole J Savage, Bernard Clot, Benoît Croczy, Fiona Tummon, Ofir Shoshanim, Brian Damit, Johannes Schneider, Vasanthi Sivaprakasam, et al. Real-time sensing of bioaerosols: Review and current perspectives. *Aerosol Science and Technology*, 54(5):465–495, 2020. 3
- [18] Milo W Hyde IV, Jason D Schmidt, and Michael J Havrilla. A geometrical optics polarimetric bidirectional reflectance distribution function for dielectric and metallic surfaces. *Optics Express*, 17(24):22138–22153, 2009. 3
- [19] Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. Polarized 3d: High-quality depth sensing with polarization cues. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3370–3378, 2015. 3
- [20] James T Kajiya. The rendering equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*, pages 143–150, 1986. 4
- [21] Inki Kim, Renato Juliano Martins, Jaehyuck Jang, Trevon Badloe, Samira Khadir, Ho-Youl Jung, Hyeongdo Kim, Jongun Kim, Patrice Genevet, and Junsuk Rho. Nanophotonics for light detection and ranging technology. *Nature Nanotechnology*, 16(5):508–524, 2021. 2
- [22] Robert Lange and Peter Seitz. Solid-state time-of-flight range camera. *Journal of Quantum Electronics*, 37(3):390–397, 2001. 1, 2
- [23] Anat Levin, Rob Fergus, Frédo Durand, and William T Freeman. Image and depth from a conventional camera with a coded aperture. *ACM Transactions on Graphics (TOG)*, 26(3):70–es, 2007. 1
- [24] David B Lindell, Matthew O’Toole, and Gordon Wetzstein. Single-photon 3d imaging with deep sensor fusion. *ACM Transactions on Graphics (TOG)*, 37(4):113–1, 2018. 2
- [25] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, Paul E Debevec, et al. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. *Eurographics Symposium on Rendering (EGSR)*, 2007(9):10, 2007. 3
- [26] Aongus McCarthy, Robert J Collins, Nils J Krichel, Verónica Fernández, Andrew M Wallace, and Gerald S Buller. Long-range time-of-flight scanning sensor based on high-speed time-correlated single-photon counting. *Applied Optics*, 48(32):6241–6251, 2009. 2
- [27] Abhimitra Meka, Mohammad Shafiei, Michael Zollhöfer, Christian Richardt, and Christian Theobalt. Real-time global illumination decomposition of videos. *ACM Transactions on Graphics (TOG)*, 40(3):1–16, 2021. 6
- [28] Shree K Nayar, Xi-Sheng Fang, and Terrance Boult. Separation of reflection components using color and polarization. *Springer International Journal of Computer Vision (IJCV)*, 21(3):163–186, 1997. 3
- [29] Michael Oren and Shree K Nayar. Generalization of lambert’s reflectance model. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, pages 239–246, 1994. 3

- [30] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. [6](#)
- [31] Joshua Rapp and Vivek K Goyal. A few photons among many: Unmixing signal and noise for photon-efficient active imaging. *IEEE Transactions on Computational Imaging (TCI)*, 3(3):445–459, 2017. [2](#)
- [32] Dieter Renker. Geiger-mode avalanche photodiodes, history, properties and problems. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 567(1):48–56, 2006. [1](#)
- [33] Kenneth Sassen. The polarization lidar technique for cloud research: A review and current assessment. *Bulletin of the American Meteorological Society*, 72(12):1848–1866, 1991. [3](#)
- [34] Kenneth Sassen. Polarization in lidar. In *Lidar*, pages 19–42. Springer, 2005. [3](#)
- [35] Guy Satat, Barmak Heshmat, Dan Raviv, and Ramesh Raskar. All photons imaging through volumetric scattering. *Scientific Reports*, 6(1):1–8, 2016. [2](#)
- [36] Guy Satat, Matthew Tancik, and Ramesh Raskar. Towards photography through realistic fog. In *Int. Conf. Comput. Photog. (ICCP)*, pages 1–10. IEEE, 2018. [2](#)
- [37] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Springer International Journal of Computer Vision (IJCV)*, 47(1):7–42, 2002. [1](#)
- [38] Yoav Y Schechner, Srinivasa G Narasimhan, and Shree K Nayar. Instant dehazing of images using polarization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages I–I. IEEE, 2001. [3](#)
- [39] Richard M Schotland, Kenneth Sassen, and Richard Stone. Observations by lidar of linear depolarization ratios for hydrometeors. *Journal of Applied Meteorology and Climatology*, 10(5):1011–1017, 1971. [3](#)
- [40] Brent Schwarz. Mapping the world in 3d. *Nature Photonics*, 4(7):429–430, 2010. [2](#), [4](#)
- [41] Dongeek Shin, Feihu Xu, Dheera Venkatraman, Rudi Lussana, Federica Villa, Franco Zappa, Vivek K Goyal, Franco NC Wong, and Jeffrey H Shapiro. Photon-efficient imaging with a single-photon camera. *Nature Communications*, 7(1):1–8, 2016. [2](#)
- [42] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wilmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. [1](#)
- [43] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America (JOSA)*, 57(9):1105–1114, 1967. [3](#), [4](#)
- [44] Tali Treibitz and Yoav Y Schechner. Active polarization descattering. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 31(3):385–399, 2008. [3](#)
- [45] Alexander P Vasilkov, Yury A Goldin, Boris A Gureev, Frank E Hoge, Robert N Swift, and C Wayne Wright. Airborne polarized lidar detection of scattering layers in the ocean. *Applied Optics*, 40(24):4353–4364, 2001. [3](#)
- [46] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. *Rendering techniques*, 2007:18th, 2007. [4](#)
- [47] David M Winker, Mark A Vaughan, Ali Omar, Yongxiang Hu, Kathleen A Powell, Zhaoyan Liu, William H Hunt, and Stuart A Young. Overview of the calipso mission and calipso data processing algorithms. *Journal of Atmospheric and Oceanic Technology*, 26(11):2310–2323, 2009. [3](#)
- [48] Di Wu, Andreas Velten, Matthew O’Toole, Belen Masia, Amit Agrawal, Qionghai Dai, and Ramesh Raskar. Decomposing global light transport using time of flight imaging. *Springer International Journal of Computer Vision (IJCV)*, 107(2):123–138, 2014. [2](#)
- [49] Zhengyou Zhang. Microsoft kinect sensor and its effect. *IEEE Multimedia*, 19(2):4–10, 2012. [2](#)