

Differentiable Polarized Path Tracing

Pramod Rao^{1,2,3}, J  r  my Riviere⁴, Xilong Zhou^{1,2}, Abhijeet Ghosh⁴,
Abhimitra Meka⁴, Thabo Beeler⁴, Marc Habermann^{1,2},
Christian Theobalt^{1,2,3}, and Delio Vicini⁴

¹ Max Planck Institute for Informatics

² VIA Research Center

³ Saarland Informatics Campus

⁴ Google




Fig. 1: Polarization plays an important role in differentiable inverse rendering. **(a)** Parameters that control polarization effects such as a linear polarizer rotation θ can produce significant gradient contributions, but existing differentiable path tracing methods struggle to compute these gradients efficiently and accurately. We introduce a caching mechanism for polarized path replay that avoids inverting rank-deficient Mueller operators. **(b)** Our method produces reverse-mode gradients with respect to scene parameters that match conventional automatic differentiation (AD), while using substantially less memory. Our algorithm improves runtime relative to existing unbiased baselines and yields substantially lower gradient error.

Abstract. Physically based differentiable rendering has proven to be a powerful tool for inverse rendering problems (e.g., 3D reconstruction, reflectance estimation, lighting estimation). However, most existing methods operate solely on radiometric intensity, discarding valuable polarization cues that constrain scene geometry and material properties. While forward simulation of polarized light is well-defined via Mueller-Stokes calculus, extending reverse-mode differentiation to this domain presents significant challenges. The rank-deficient nature of common polarimetric operators, such as linear polarizers and diffuse reflections, violates the invertibility assumptions of standard gradient estimators like path replay backpropagation and results in numerical instability. We address this by proposing a robust, polarization-aware differentiable path tracing method. Our approach estimates unbiased gradients through a combination of path replay and local caching. This formulation enables efficient

and stable optimization of material and lighting parameters in complex scenes, broadening the applicability of physically based inverse rendering. Project page: <https://vcai.mpi-inf.mpg.de/projects/DPPT/>

Keywords: Differentiable rendering, Polarization, Path Replay Backpropagation, Mueller-Stokes Calculus

1 Introduction

Polarimetric imaging offers valuable cues for inverse problems such as 3D reconstruction [1, 3, 17, 20, 24], appearance estimation [4, 9, 27, 35], and intrinsic decomposition [41, 42]. Capturing information about the polarization state of light helps constrain reconstructions to disambiguate material or surface properties that may appear similar in intensity-only images [41].

Recently, *inverse rendering* approaches have gained popularity to tackle 3D reconstruction problems [22, 28, 32, 39]. Specifically, *physically based differentiable rendering* (PBDR) methods optimize scene parameters while accounting for the intricate interaction of light with scene geometry and materials [12, 25, 32, 48]. Using reverse-mode differentiation, these methods can jointly optimize very high-dimensional parameter vectors (e.g., textures or mesh vertices). The practical application of these methods can be challenging due to local minima and ambiguities. In this work, we combine polarimetric measurements with PBDR. By using these richer observations, we enable better identification of material parameters in inverse rendering and improve robustness.

The combination of PBDR with polarization information is non-trivial: While PBDR methods offer a powerful and general framework, simply applying reverse-mode automatic differentiation [13] to path tracing [21] quickly exhausts the available system memory [31, 37]. *Path replay backpropagation* (PRB) [37] addresses this by leveraging local numerical invertibility to efficiently compute gradients using a constant amount of memory and linear complexity. However, simulating polarized light transport typically uses *Mueller-Stokes calculus* [8]. This means scalar radiance values are replaced by 4D Stokes vectors, and interactions of light with surfaces require multiplication with 4×4 Mueller matrices. Even for simple scenes, the Mueller matrices are often singular and thus violate PRB’s local invertibility assumptions. Therefore, a direct application of PRB leads to unstable or severely biased gradient estimators.

We first identify these numerically problematic cases and the scenarios where standard PRB fails. We then propose a general, efficient, and easy-to-implement algorithm that produces unbiased gradients (Fig. 1). The key idea is to extend PRB by selectively storing a small amount of additional information in local memory. This is significantly less memory-intensive than conventional automatic differentiation while resolving numerical issues, which enables fully differentiable polarized path tracing. When applying this method to material reconstruction tasks, the additional cues from polarimetric imaging effectively reduce problematic local minima that hinder unpolarized inverse rendering. In this paper, we

focus on the continuous *interior* integral of the differentiable light transport, and leave the handling of geometric discontinuities and shape reconstruction to future work. In summary, our contributions are:

- An analysis of why standard PRB does not directly generalize to polarized light transport, identifying rank-deficient Mueller operators as the source of ill-defined or unstable replay cancellation.
- To our knowledge, the first practical framework for differentiable polarized path tracing along long light paths, combining cached suffix replay with a hybrid variant to avoid unreliable Mueller-matrix inversion.
- An experimental validation demonstrating that polarized PBDR improves material, texture, and normal recovery compared to unpolarized baselines, while producing gradients consistent with conventional AD.

2 Related Work

2.1 Polarimetric Reconstruction

Polarization has a long history in computer vision and graphics, dating back to the 1970s when polarization was established as a useful signal for shape understanding [7, 23, 43], material classification [40, 42], and the separation of diffuse and specular reflection [41]. The seminal work of Boult & Wolff [7] establishes the theoretical framework for Shape from Polarization (SfP) for specular reflection via the Transmitted Radiance Sinusoid (TRS), measured by rotating a linear polarizer in front of a camera at 3–4 angles⁵. The phase of the TRS is related to the surface normal up to an ambiguity of 180° in azimuth. Subsequent works address this ambiguity using multi-view constraints [1, 30], fusion with coarse geometry [20], or deep learning approaches [3, 24]. While most SfP methods focus on dielectric specular materials, extensions to transparent objects [29] and diffuse surfaces [2] have also been proposed. Another interesting application of polarization for scene understanding is diffuse-specular separation, where specular surfaces preserve the polarization properties of incident polarized illumination while diffuse surfaces depolarize light. Ma *et al.* [27] introduce cross and parallel polarized spherical gradient illumination to recover diffuse and specular albedo as well as surface normals. Ghosh *et al.* [11] further extend this setup to multi-view polarization capture. Researchers also investigate methods to reduce capture complexity for polarimetric reflectance acquisition. Baek *et al.* [4] propose a near co-axial projector-camera ellipsometry setup for acquiring shape and spatially varying surface reflectance, combining inverse rendering with a polarimetric BRDF (p-BRDF) modeling both specular and diffuse polarization effects. Baek *et al.* [5] further extend this to spectral measurements and provide a database of isotropic polarimetric BRDFs. Hwang *et al.* [18] simplify capture further by using sparse ellipsometry with a hand-held polarized flash

⁵ Typically 0° , $\pm 45^\circ$, and 90° . A polarization camera may also be used, where the Bayer pattern is replaced by micro-polarizers at different orientations per pixel.

camera for shape and SVBRDF reconstruction. Deschaintre *et al.* [9] leverage deep learning frameworks for 3D shape and material reconstruction. Riviere *et al.* [35] propose a method to acquire material reflectance properties through multi-view polarized images, which was later extended by [33] to single-view material acquisition. Polarization techniques are also widely used for complex facial appearance capture [15, 26].

Polarization has also been integrated with neural fields to improve geometry reconstruction and challenging material acquisition. Han *et al.* [16] address the specific problem of reconstructing reflective objects from sparse views. By incorporating a p-BRDF [4] and enforcing normal consistency derived from the Angle of Polarization (AoP) cues, NeRSP significantly reduces shape-radiance ambiguity for sparse captures. Moving towards higher physical fidelity, Wanaset *et al.* [38] propose a method that fits a polarized neural field directly to raw, mosaiced sensor data. More recently, Han *et al.* [17] explicitly incorporate polarimetric constraints into 3D Gaussian Splatting [22]. They utilize a Gaussian Surfel representation to better approximate surface geometry and employ a p-BRDF [4] module to separate specular and diffuse components, allowing high-fidelity reconstruction of shiny objects from sparse inputs. Concurrently, Guo *et al.* [14] focus on ambiguity-free reconstruction. They leverage the Degree of Linear Polarization (DoLP) to detect photometrically ambiguous regions (e.g., textureless or perfectly specular areas) and apply a polarization-guided geometric correction to align Gaussian normals with physical constraints.

While these methods demonstrate promising reconstruction results, they primarily use polarization cues as additional supervision signals during optimization, without addressing the underlying challenge of differentiating through the polarized light transport governed by Mueller-Stokes calculus. In contrast, our work addresses the fundamental rank-deficiency of polarimetric operators and extends PRB to support polarization-aware differentiable path tracing, enabling efficient and stable inverse rendering.

2.2 Physically Based Inverse Rendering

Physically based inverse rendering methods differentiate (unbiased) Monte Carlo path tracing [12, 25, 32, 39, 48]. Many works in this space focus on differentiating visibility discontinuities, as required for geometry optimization [6, 25, 44, 48, 49]. On the other hand, Stam [36], Nimier-David *et al.* [31] and Vicini *et al.* [37] present approaches to reduce the memory use for differentiating long light paths. Specifically, we extend path replay backpropagation [37] to support polarization. Handling visibility discontinuities is orthogonal, and we leave this for future work. To the best of our knowledge, Yang *et al.* [45] is the only prior work demonstrating differentiable polarization rendering across multiple interactions. The authors demonstrate the use of a differentiable polarized light transport simulation for waveguide design. The implementation is specialized to waveguide design, and uses up to 70 GB of memory across four GPUs. Our method is aimed at general scenes, and carefully avoids storing large gradient checkpoints.

3 Background

3.1 Inverse Rendering

Problem Definition. Our goal is to solve inverse rendering problems of the form:

$$\pi^* = \arg \min_{\pi} g(I(\pi)), \quad (1)$$

where I is a rendered image as a function of a vector of scene parameters π , and g is a differentiable objective function. To solve this problem using gradient descent, we need to efficiently compute derivatives of the rendering function $I(\pi)$. Without loss of generality, all the following derivations will be written in terms of a single pixel and a single parameter π .

Rendering Equation. Physically based rendering algorithms estimate the pixel color by applying Monte Carlo integration to the *rendering equation* [21]:

$$L_o(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \int_{S^2} f_s(\mathbf{x}, \omega_o, \omega_i) L_i(\mathbf{x}, \omega_i) d\omega_i^\perp, \quad (2)$$

where $\mathbf{x}$ is a position, ω_o the outgoing direction, and ω_i the incident direction. L_o , L_i , and L_e denote outgoing, incident, and emitted radiance, respectively. The bidirectional scattering distribution function (BSDF) f_s differentially relates incident to outgoing radiance. In unpolarized rendering, radiance is either modeled spectrally, or simply as RGB values. The standard representation to simulate polarization uses Stokes vectors and Mueller matrices (Sec. 3.3).

3.2 Path Replay Backpropagation

Before expanding on polarized light transport, we review path replay backpropagation (PRB) [37]. The path replay algorithm allows computing (unbiased) derivatives of arbitrarily long light paths. For this, it exploits the local invertibility of the path tracer’s light path contribution calculation. Our method is an extension of this algorithm to polarized light transport.

```

1 def sample_path_adjoint(ray, L, δL):
2     β = 1
3     for i in range(N):
4         L -= β * L_e(...)
5         ω_i, bsdf_val, bsdf_pdf = sample_bsdf(...)
6         δπ += backward(δL * L / bsdf_val * bsdf_val)
7         β *= bsdf_val / bsdf_pdf
8     return δπ

```

Algorithm 1: Pseudocode for path replay backpropagation.

PRB computes gradients in three phases:

1. An initial *primal* phase estimates the image pixels and is used to compute the derivative of the loss function with respect to the pixel intensities.
2. A decorrelated rendering pass traces new light paths and retains the estimated radiance per path.
3. Finally, the *adjoint* pass reuses the random seed from step 2, *replaying* the same light paths. At each interaction, it differentiably evaluates the local contribution of the BSDF and emitters, and accumulates the gradients due to that particular interaction.

Notably, automatic differentiation is only used within the loop body and not across the iteration boundaries. This avoids memory-intensive checkpoints that retain the full loop state, as would be needed to differentiate a loop using automatic differentiation [13].

Algorithm 1 provides pseudocode for the adjoint phase (step 3) of the algorithm. The pseudocode computes derivatives of the BSDF parameters. The inputs of the algorithm are the current ray, the resulting radiance estimate from step 2, and the gradient of the loss function with respect to the path radiance δL . This adjoint pass evaluates a vector-Jacobian product between the Jacobian $\partial L / \partial \boldsymbol{\pi}$ and the vector δL . The term β denotes the *path throughput* and L_e is the emitted radiance. The `sample_bsdf` function samples an outgoing direction proportional to the BSDF. Line 6 evaluates the local derivative of the current BSDF parameters. The function `backward(x)` “backpropagates” gradients of the computation `x`. PRB computes gradients selectively with respect to only some of the quantities within the loop body. For clarity, we therefore indicate the differentiated quantity in blue color, with any non-differentiated weights in black. We refer to [37] for an in-depth introduction to PRB.

3.3 Polarized Light Transport

We now define notation and terminology related to polarized rendering. For a comprehensive theory, please refer to Collet *et al.* [8], and Zeltner *et al.* [47].

Mueller-Stokes Calculus. The standard approach to rendering with polarization is to use *Mueller-Stokes calculus*. In this framework, the scalar radiance is replaced by 4D *Stokes vectors*. A Stokes vector $\mathbf{s} \in \mathbb{R}^4$ has components s_0 for intensity, s_1 distinguishes horizontal and vertical linear polarization, s_2 accounts for diagonal polarization, and s_3 for circular polarization. For unpolarized light, components s_1 , s_2 and s_3 are zero. A *Mueller matrix* is a 4×4 matrix that characterizes a linear operation on Stokes vectors. In polarized light transport, a BSDF evaluation returns a Mueller matrix. Various optical filters (e.g., a linear polarizer) can be modeled using their Mueller matrix.

Coordinate Conventions. Polarized quantities are always defined with respect to a 3D coordinate system. Any operation (e.g., addition of two Stokes vectors) requires both quantities to be expressed with respect to the same coordinate frame. We implement our algorithms on top of Mitsuba 3 [19], which implements all the necessary coordinate conversions.

Path Contribution. Implementing the Mueller-Stokes calculus requires modifying Algorithm 1. For simplicity, a practical implementation can represent all quantities using Mueller matrices [32]. For that, Stokes vectors are represented as a Mueller matrix, where the first column contains the Stokes vector, and the other entries are zeros. Therefore, in a polarized path tracer, both the accumulated path contribution and throughput become 4×4 matrices. A key assumption of PRB is that `bsdf_val` is non-zero, and hence invertible. In polarized rendering, `bsdf_val` is a 4×4 Mueller matrix. For many polarization operators (e.g., optical filters), the Mueller matrix is singular (or close to it), as it maps the 4D Stokes space to the lower-dimensional subspace of the transmitted/reflected polarization state. Mathematically, such a lossy operation does not have a well-defined inverse. Therefore, the scalar inversion used by PRB does not directly carry over to the polarized setting.

4 Method

4.1 Problem Statement

We discuss the generalization of PRB to support polarized rendering. Our goal is to enable the computation of gradients with an arbitrary path depth while avoiding approximations and bias.

Pseudocode. Algorithm 2 shows a basic implementation of PRB applied to polarization. The key difference from Algorithm 1 is that all quantities are now Mueller matrices, and M^{-1} denotes matrix inversion.

Limitations of Algorithm 2. In practice, the matrix M is most often rank-deficient and non-invertible. For example, an ideal diffuse BSDF’s Mueller matrix M has only a single non-zero component M_{00} (the albedo). All other entries are zero, and the matrix is therefore not invertible. Intuitively, a diffuse interaction loses information about the incident light’s polarization state. Similarly, a linear polarizer filters out circular polarization, thus losing information about the last component of the Stokes parameters. On the other hand, some optical elements do produce invertible matrices. For example, a quarter waveplate simply rotates the incident polarization state by 90° around its axis, and does not lose any information.

```

1 def sample_polarized_basic(ray, L,  $\delta L$ ):
2      $\beta = I_4$ 
3     for i in range(N):
4          $L -= \beta * L_e(\dots)$ 
5          $\omega_i, M, \text{bsdf\_pdf} = \text{sample\_bsdf}(\dots)$ 
6          $L = M^{-1} * L$  # Remove the BSDF contribution from L
7          $\delta_\pi += \text{backward}(\delta L * M * L)$ 
8          $\beta *= M / \text{bsdf\_pdf}$ 
9     return  $\delta_\pi$ 

```

Algorithm 2: Pseudocode for polarized path replay backpropagation.




Fig. 2: Example for Cached Suffix Replay on a light path of depth 2. The primal pass accumulates the full path contribution, tracks the detached BxDF contributions β_i , and constructs the cache of suffix contributions via a backward fold. The adjoint pass then directly re-loads the incident polarization state and multiplies it by the differentiated BxDF weight. This enables unbiased gradient estimation without any matrix inversion.

Analysis. The examples above reveal that there are three different cases our algorithm ultimately needs to handle:

1. *Invertible elements:* Any BxDF with an invertible Mueller matrix can be handled by directly inverting M .
2. *Ideal depolarizers:* BxDFs that fully depolarize incident illumination (e.g., diffuse BxDFs) can be handled by standard PRB. Their derivatives do not require any information about the incident polarization state of light.
3. *Non-invertible polarization-aware elements:* Elements that consider polarization, but produce a non-invertible M are not supported by the baseline algorithm.

The first two cases can already be handled within the framework of the basic Algorithm 2, but the last case requires additional considerations.

4.2 Regularization Using Noise

Vicini *et al.* [37] propose to handle a similar non-invertibility for *attached* sampling by adding i.i.d. noise to the matrix diagonal. The key idea here is to add $u \cdot I_4$ to the matrix M , where $u \sim \mathcal{N}(0, \sigma)$ is randomly sampled. By using the same noise for both the initial path tracing as well as the adjoint phase, this technically does not introduce bias into the gradients [37]. This regularization can help in some cases, but is not without issues: First, there is no guarantee that the resulting matrix is well-conditioned. Second, adding too much noise will deteriorate the quality of the derivatives. This method alone cannot produce sufficiently high-quality gradients, regardless of the choice of the noise magnitude σ , as shown in Sec. 5.1.

4.3 Cached Suffix Replay

The basic polarized extension of PRB (Algorithm 2) follows the same replay logic as in the scalar setting. During the adjoint pass, it “removes” the detached BxDF

contribution by multiplying with the inverse of the sampled transport factor. In polarized light transport, this factor is a Mueller matrix M , which is frequently non-invertible, and thus results in incorrect derivatives. Our alternative avoids inversion by caching a compact summary of the *suffix* of each sampled path. We illustrate the core idea using an example light path in Fig. 2.

```

1  def sample_polarized_primal(ray, ...):
2      L = [04, ...], β = [I4, ...] # Size N+1
3      for i in range(N):
4          L[i] = Le(...)
5          ωi, M, bsdf_pdf = sample_bsdf(...)
6          β[i] = M / bsdf_pdf
7
8      # Loop backwards: reconstruct suffix radiance
9      for j in range(N-2, -1, -1):
10         L[j] += β[j] * L[j + 1]
11
12     return L[0], L
13
14 def sample_polarized_adjoint(ray, L, δL):
15     β = I4
16     for i in range(N):
17         ωi, M, bsdf_pdf = sample_bsdf(...)
18
19         Li = L[i + 1] # <-- Suffix radiance from cache (no inversion needed)
20         δπ += backward(δL * β * M * Li)
21
22         β *= M / bsdf_pdf
23
24     return δπ

```

Algorithm 3: Pseudocode: polarized path replay backpropagation with caching.

Algorithm. We implement this cached path replay with two passes (Algorithm 3). The primal pass stores per-vertex emission and throughput factors in local arrays and reconstructs suffix radiance by a backward fold. During the adjoint replay, instead of canceling the detached contribution via M^{-1} , we directly look up the cached suffix radiance $L_i = L[i + 1]$. We then compute the local gradient term using the replayed throughput and a differentiable BSDF evaluation. This removes the need for inverting rank-deficient Mueller operators while producing unbiased gradients. Unlike conventional automatic differentiation, we do not store the complete loop state, but instead retain path replay to recompute surface interaction structures, BSDF values and next-event estimation. This offers a more favorable memory-compute tradeoff, and scales to scenes of high complexity, as demonstrated in the Sec. 5.

4.4 Hybrid Cached Replay

Cached Suffix Replay stores suffix radiance at every path interaction. This is efficient for the path depths used in most of our inverse-rendering experiments,

but the local cache size still grows with the maximum path depth. To reduce this storage for deeper paths, we introduce a hybrid variant that combines block-wise caching with recursive recomputation.

The hybrid method stores cached suffix values only at checkpoint intervals. Let k denote the checkpoint interval. When $k = 1$, the method reduces to the fully cached variant described above. Larger values of k store fewer cached states and therefore reduce memory usage, but require additional work during replay for vertices that lie between two checkpoints. At such vertices, the adjoint pass cannot directly load the suffix radiance and must recover it either by local inversion or by recursively tracing until the next checkpoint.

The decision is made using the detached block throughput $\bar{\beta}$, which summarizes the polarized transport accumulated since the last checkpoint. We first evaluate the depolarization index $DI(\bar{\beta})$. If $DI(\bar{\beta}) < \gamma$, the accumulated transport is treated as effectively depolarizing and we use the scalar PRB update. Otherwise, we retain the polarized update and test whether local inversion is numerically admissible. The high-level replay rule is:

$$\text{suffix recovery} = \begin{cases} \text{scalar PRB update,} & DI(\bar{\beta}) < \gamma, \\ \bar{\beta}^{-1}L, & DI(\bar{\beta}) \geq \gamma \text{ and } |\det(\bar{\beta})| > \varepsilon, \\ \text{recursive recomputation,} & \text{otherwise,} \end{cases} \quad (3)$$

where L is the cached suffix at the next checkpoint. The threshold γ controls when polarized replay is retained, while ε avoids unstable inversions of near-singular Mueller operators. This yields a memory-reduced variant of our method for larger path depths. Full pseudocode and extended benchmarks are provided in the supplemental material.

5 Results

In this section, we report correctness tests, practical applications of polarization, and performance evaluations of our method. We ran all evaluations on a workstation equipped with an NVIDIA GeForce RTX 4090 (24 GiB). To our knowledge, aside from Mitsuba’s conventional automatic differentiation, our method is the first to support full differentiable polarized path tracing. We therefore compare against polarized extensions of PRB [37] and RB [31], which we denote P-PRB and P-RB, respectively. While these extensions are not part of the original works, we discussed the implementation details with the authors and validated their correctness.

5.1 Gradient Analysis

We first validate gradient correctness in Fig. 3 by comparing our method against Mitsuba’s conventional automatic differentiation (Conv. AD) and polarized variants of RB and PRB. The figure reports reverse-mode gradients (first three rows) and forward-mode gradients (last row) across multiple scenes and materials.




Fig. 3: Gradient Analysis. We visualize gradient images and compare against the Conv. AD reference. Columns report P-PRB, P-RB, and our method, with each gradient tile annotated by its *relative error (RE)* w.r.t. Conv. AD (preferred over RMSE to better account for small gradient magnitudes). Rows correspond to different scenes with their associated parameters listed on the right.

We use conventional AD as the reference. Since our method and P-RB are unbiased, both produce gradients that closely match the AD baseline. In contrast, the polarized extension of PRB based on noise-regularized matrix inversion (P-PRB) often yields incorrect gradients. The discrepancy becomes more pronounced when incident polarization strongly affects the objective, such as when optimizing surface normals or roughness textures (rows 2 and 3 in Fig. 3).

To further isolate this limitation, we evaluate the gradient in the Cornell Box scene with respect to the linear polarizer rotation parameter θ . Here, the sampled BSDF contribution is a rank-deficient Mueller matrix (linear polarizer). The baseline P-PRB relies on inverting the sampled BSDF’s Mueller matrix, which is not well-defined. As a result, P-PRB breaks down in this setting, making it ineffective for differentiating through a linear polarizer, which is a canonical optical element in polarimetric scenes.

5.2 Impact of Polarization on Inverse Rendering

While the previous section establishes gradient correctness, we now examine the role of polarization in inverse rendering tasks. We first demonstrate the variation




Fig. 4: Qualitative comparison of reverse-mode gradients under different polarization configurations. We render a virtual light stage setup with right-circular illumination and accumulate gradients over a multi-view camera setup. Columns correspond to polarizer configurations (0° , 45° , 90° , and left-circular), and rows show gradients with respect to different scene parameters (surface normals, roughness, and specular reflectance). The resulting gradients exhibit distinct spatial structures across polarization states, illustrating that polarization provides complementary information when optimizing different parameters.

of gradients under different polarization configurations. Inspired by the multi-view capture setups of [10, 34, 50], we construct a virtual light stage with right-circular illumination and a centrally placed object observed from multiple views. The object is rendered under four polarization configurations: 0° , 45° , 90° and left-circular.

Fig. 4 shows reverse-mode gradients with respect to surface normals, roughness, and specular reflectance under these configurations. Each polarization state produces distinct gradient structures, indicating that polarization provides complementary information for parameter estimation. Having verified both correctness and diversity of gradients, we next evaluate the practical impact of polarization in optimization tasks (Fig. 5). We compare unpolarized PRB [37] against our polarization-aware method. In the first task (*Kitchen* scene), we aim to remove glare in the cooking area (highlighted in red). We introduce a linear polarizer in front of the camera and optimize its orientation. Since PRB cannot differentiate parameters controlling the polarization state, it fails to suppress the glare. In contrast, our method successfully optimizes the configuration and reduces the specular reflection.

In the second task (*Veatch* scene, row 2), we jointly optimize the roughness and diffuse texture of a plate. PRB struggles to disentangle these parameters, leading to incorrect texture recovery (see inset). By leveraging cross and parallel




Fig. 5: Polarized inverse rendering. We evaluate four inverse rendering tasks (rows) and compare conventional automatic differentiation (Conv. AD), unpolarized PRB [37], and our polarization-aware method against the ground truth reference. From left to right, we show the initial state, the optimized result using Conv. AD, PRB, our method, and the reference. The tasks include optimizing the rotation angle of a linear polarizer (*Kitchen*), jointly optimizing diffuse reflectance and roughness (*Veach*), optimizing diffuse reflectance (*Living Room*), and recovering a normal map (*Marble Bust*). Conv. AD fails to scale to several of these settings and runs out of memory (OOM). Compared to PRB, our method achieves lower reconstruction error (RMSE overlays) and produces results that more closely match the reference across all tasks.

polarized observations, our method better separates specular and diffuse components and recovers parameters closer to the reference. While this is a known application of polarization for reflectance separation [11, 27, 35], it is important to note here that we do not explicitly pre-separate diffuse and specular observations to model with diffuse and specular BSDFs separately. Details of experimental design are in supplemental material.

We observe similar behavior in the *Living Room* scene when reconstructing the diffuse floor texture. Finally, we jointly optimize surface normals and diffuse reflectance of a marble bust from a multi-view dataset consisting of 12 images to recover a 1K-resolution normal map. PRB produces inaccurate normals, particularly in highly detailed regions such as the statue’s hair. Polarization provides additional cues about surface orientation, enabling improved normal recovery.

Across these experiments, polarization enhances parameter identifiability, and our contribution is to enable such tasks with an efficient polarization-compatible integrator. While our results match Conv. AD in correctness, AD does not scale to scenes with many bounces, large parameter spaces, or high resolutions; in several cases it runs out of memory (gray entries in Fig. 5).




Fig. 6: Single-view 3D reconstruction. Ours (*Pol.*) reconstructed geometry is sharper than the unpolarized baseline (*Unpol.*), shown from test views.

5.3 Geometry Reconstruction with Polarization

Classical shape-from-polarization methods exploit the fact that the polarization state of reflected light carries surface-orientation information. To demonstrate this, we further extend the projective sampling integrator [49] to polarized light transport with suffix caching, enabling visibility-discontinuity gradients while retaining our inversion-free polarized replay. We evaluate single-view 3D reconstruction over seven meshes, comparing our polarized variant against the corresponding unpolarized projective-sampling baseline under the same initialization and optimization schedule. As shown in Fig. 6, polarized observations recover sharper geometric details and more faithful silhouettes that are smoothed or distorted in the unpolarized reconstructions. Quantitatively, polarization reduces Chamfer distance from **0.1874** to **0.0904** and normal error from **65.61°** to **46.72°**. These results indicate that polarized light transport provides crucial shape cues beyond intensity alone, and that our replay formulation can be combined with visibility-aware shape-gradient estimators. We analyze scalability and performance in the next section.

5.4 Performance Evaluation

Mitsuba represents polarized quantities using per-channel Mueller matrices. Instead of a simple 3D vector, each radiance quantity becomes a $4 \times 4 \times 3$ tensor. This increased representation cost makes differentiable rendering in polarized mode significantly more demanding in both memory and computation time, motivating the need for an efficient gradient estimator.

We benchmark reverse-mode gradient computation in terms of runtime and memory consumption (Fig. 7). Thanks to the use of local memory, our method exhibits nearly constant memory usage up to a path depth of 32, with only a slight increase beyond that point. It always remains substantially more memory-efficient than conventional AD. In the supplemental material, we discuss a variant of our method that combines local caching with recursive recomputation based on polarization heuristics to further reduce memory consumption. We also evaluate the hybrid variant introduced in Sec. 4.4. For the path depths used in our




Fig. 7: Performance evaluation on a polarized scene rendered at 1280×720 resolution with 1 sample per pixel. (a) We use the *Modern Hall* scene and introduce a stack of 85 linear polarizers to simulate long light paths. (b) We illustrate the effect of increasing the maximum path depth on the rendered appearance. We then compute gradients with respect to the floor albedo using conventional automatic differentiation (Conv. AD), P-RB, P-PRB, and our method. (c) Conv. AD exhausts GPU memory of a NVIDIA GeForce RTX 5090, whereas our method maintains an approximately constant memory up to depth 32 and increases only slightly thereafter, while remaining substantially below Conv. AD. (d) P-PRB attains low runtime but produces inaccurate gradients. P-RB produces unbiased gradients, but its runtime is higher; in this experiment it is approximately $2\times$ slower than our method.

main applications, the fully cached method remains the most efficient choice. The hybrid variant becomes useful at larger depths, where storing suffix values at every interaction increases local memory use. Detailed plots are included in the supplemental material.

Although P-PRB also demonstrates constant memory behavior, its gradient inaccuracies limit its applicability in polarized settings. Conventional AD, while producing correct gradients, incurs substantially higher memory and runtime costs. In comparison, our method achieves correct gradients with improved efficiency, and is approximately twice as fast as P-RB in our experiments, making it a practical approach for polarized differentiable rendering.

6 Conclusion and Future Work

In this paper, we introduced a memory-efficient differentiable path tracing method for polarized light transport. We showed that standard PRB does not directly generalize to polarimetric rendering, since common Mueller operators can be rank-deficient and make replay-based inversion ill-defined. Our method avoids this issue using cached suffix replay, and further introduces a hybrid cache–recompute variant for larger path depths. Together, these algorithms enable stable and efficient gradient computation for polarized scenes. We validated the method through gradient comparisons against conventional AD and demonstrated that polarimetric measurements improve material, texture, and normal recovery compared to unpolarized baselines.

Future work includes extending the framework to subsurface scattering. Since multiple scattering tends to depolarize light while the first surface interactions can retain polarization cues, differentiable polarized subsurface transport could help separate surface specularities from internal scattering in translucent materials such as marble, jade, or skin. It would also be interesting to apply our method to computational design problems such as waveguide design [46].

References

1. Atkinson, G.A., Hancock, E.R.: Multi-view surface reconstruction using polarization. In: IEEE International Conference on Computer Vision (ICCV). vol. 1, pp. 309–316. IEEE (2005)
2. Atkinson, G.A., Hancock, E.R.: Recovery of surface orientation from diffuse polarization. *IEEE transactions on image processing* **15**(6), 1653–1664 (2006)
3. Ba, Y., Gilbert, A., Wang, F., Yang, J., Chen, R., Wang, Y., Yan, L., Shi, B., Kadambi, A.: Deep shape from polarization. In: European Conference on Computer Vision (ECCV). pp. 554–571. Springer (2020)
4. Baek, S.H., Jeon, D.S., Tong, X., Kim, M.H.: Simultaneous acquisition of polarimetric svbrdf and normals. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **36**(6), 268:1–15 (2018). <https://doi.org/10.1145/3272127.3275018>, <http://dx.doi.org/10.1145/3272127.3275018>
5. Baek, S.H., Zeltner, T., Ku, H.J., Hwang, I., Tong, X., Jakob, W., Kim, M.H.: Image-based acquisition and modeling of polarimetric reflectance. *ACM Trans. Graph. (Proc. SIGGRAPH)* **39**(4) (2020)
6. Bangaru, S., Li, T.M., Durand, F.: Unbiased warped-area sampling for differentiable rendering. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **39**(6), 245:1–245:18 (2020). <https://doi.org/10.1145/3414685.3417833>
7. Boulton, T.E., Wolff, L.B.: Physically-based edge labeling. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 656–662 (1991)
8. Collet, E.: Field Guide to Polarization. SPIE Publications (2005)
9. Deschaintre, V., Lin, Y., Ghosh, A.: Deep polarization imaging for 3d shape and svbrdf acquisition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2021)
10. Ghosh, A., Chen, T., Peers, P., Wilson, C.A., Debevec, P.: Circularly polarized spherical illumination reflectometry. *ACM Transactions on Graphics (TOG)* **29**(6), 1–12 (2010)
11. Ghosh, A., Fyffe, G., Tunwattanapong, B., Busch, J., Yu, X., Debevec, P.: Multiview face capture using polarized spherical gradient illumination. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **30**(6), 1–10 (Dec 2011). <https://doi.org/10.1145/2070781.2024163>
12. Gkioulekas, I., Levin, A., Zickler, T.: An evaluation of computational imaging techniques for heterogeneous inverse scattering. In: European Conference on Computer Vision (ECCV). pp. 685–701. Springer International Publishing (2016). https://doi.org/10.1007/978-3-319-46487-9_42
13. Griewank, A., Walther, A.: Evaluating derivatives: principles and techniques of algorithmic differentiation, vol. 105. SIAM (2008)
14. Guo, B., Wen, S., Zhao, Y., Li, J., Zheng, Z.: Polargs: Polarimetric cues for ambiguity-free gaussian splatting with accurate geometry recovery. *arXiv preprint arXiv:2512.00794* (2025)
15. Ha, H., Hwang, I., Monzon, N., Cho, J., Kim, D., Baek, S.H., Muñoz, A., Gutierrez, D., Kim, M.H.: Polarimetric bssrdf acquisition of dynamic faces. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2024)* **43**(6) (2024)
16. Han, Y., Guo, H., Fukai, K., Santo, H., Shi, B., Okura, F., Ma, Z., Jia, Y.: Nersp: Neural 3d reconstruction for reflective objects with sparse polarized images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11821–11830 (2024)

17. Han, Y., Tie, B., Guo, H., Lyu, Y., Li, S., Shi, B., Jia, Y., Ma, Z.: Polgs: Polarimetric gaussian splatting for fast reflective surface reconstruction. In: IEEE International Conference on Computer Vision (ICCV). pp. 28073–28082 (2025)
18. Hwang, I., Jeon, D.S., Muñoz, A., Gutierrez, D., Tong, X., Kim, M.H.: Sparse ellipsometry: Portable acquisition of polarimetric svbrdf and shape with unstructured flash photography. *ACM Trans. Graph. (Proc. SIGGRAPH)* **41**(4) (2022)
19. Jakob, W., Speierer, S., Roussel, N., Nimier-David, M., Vicini, D., Zeltner, T., Nicolet, B., Crespo, M., Leroy, V., Zhang, Z.: Mitsuba 3 renderer (2022), <https://mitsuba-renderer.org>
20. Kadambi, A., Taamazyan, V., Shi, B., Raskar, R.: Polarized 3d: High-quality depth sensing with polarization cues. In: IEEE International Conference on Computer Vision (ICCV). pp. 3370–3378 (2015)
21. Kajiya, J.T.: The rendering equation. *Computer Graphics (Proc. SIGGRAPH)* **20**(4), 143–150 (Aug 1986). <https://doi.org/10/cvf53j>
22. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph. (Proc. SIGGRAPH)* **42**(4), 139–1 (2023)
23. Koshikawa, K.: A polarimetric approach to shape understanding of glossy objects. In: Proceedings of the 6th international joint conference on Artificial intelligence—Volume 1. pp. 493–495 (1979)
24. Lei, C., Qi, C., Xie, J., Fan, N., Koltun, V., Chen, Q.: Shape from polarization for complex scenes in the wild. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12632–12641 (2022)
25. Li, T.M., Aittala, M., Durand, F., Lehtinen, J.: Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **37**(6), 222:1–222:11 (2018). <https://doi.org/10.1145/3272127.3275109>
26. Lin, A., Ghosh, A.: Single-shot facial capture using polarized rgb sinusoidal illumination. In: Proceedings of the SIGGRAPH Asia 2025 Conference Papers. SA Conference Papers '25, Association for Computing Machinery, New York, NY, USA (2025). <https://doi.org/10.1145/3757377.3763942>, <https://doi.org/10.1145/3757377.3763942>
27. Ma, W.C., Hawkins, T., Peers, P., Chabert, C.F., Weiss, M., Debevec, P.: Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination. In: Kautz, J., Pattanaik, S. (eds.) *Rendering Techniques (Proc. EG Symposium on Rendering)*. Eurographics Association (2007). <https://doi.org/10.2312/EGWR/EGSR07/183-194>
28. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision (ECCV) (2020). <https://doi.org/10.1145/3503250>
29. Miyazaki, Kagesawa, Ikeuchi: Polarization-based transparent surface modeling from two views. In: IEEE International Conference on Computer Vision (ICCV). pp. 1381–1386. IEEE (2003)
30. Miyazaki, D., Shigetomi, T., Baba, M., Furukawa, R., Hiura, S., Asada, N.: Polarization-based surface normal estimation of black specular objects from multiple viewpoints. In: 2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission. pp. 104–111. IEEE (2012)
31. Nimier-David, M., Speierer, S., Ruiz, B., Jakob, W.: Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. *ACM Trans. Graph. (Proc. SIGGRAPH)* **39**(4) (Jul 2020). <https://doi.org/10.1145/3386569.3392406>

32. Nimier-David, M., Vicini, D., Zeltner, T., Jakob, W.: Mitsuba 2: A retargetable forward and inverse renderer. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **38**(6) (Nov 2019). <https://doi.org/10.1145/3355089.3356498>
33. Nogue, E., Lin, Y., Ghosh, A.: Polarization-imaging Surface Reflectometry using Near-field Display. In: Ghosh, A., Wei, L.Y. (eds.) *Comp. Graph. Forum (Proc. EGSR)*. The Eurographics Association (2022). <https://doi.org/10.2312/sr.20221154>
34. Rao, P., Meka, A., Zhou, X., Fox, G., B R, M., Zhan, F., Weyrich, T., Bickel, B., Pfister, H., Matusik, W., Beeler, T., Elgharib, M., Habermann, M., Theobalt, C.: 3dpr: Single image 3d portrait relighting with generative priors. In: *ACM SIGGRAPH ASIA 2025 Conference Proceedings* (2025)
35. Riviere, J., Reshetouski, I., Filipi, L., Ghosh, A.: Polarization imaging reflectometry in the wild. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **36**(6) (Nov 2017). <https://doi.org/10.1145/3130800.3130894>, <https://doi.org/10.1145/3130800.3130894>
36. Stam, J.: Computing light transport gradients using the adjoint method. *CoRR abs/2006.15059* (2020), <https://arxiv.org/abs/2006.15059>
37. Vicini, D., Speierer, S., Jakob, W.: Path replay backpropagation: Differentiating light paths using constant memory and linear time. *ACM Trans. Graph. (Proc. SIGGRAPH)* **40**(4) (Aug 2021). <https://doi.org/10.1145/3450626.3459804>
38. Wanaset, R., Guarnera, C., Smith, W.A.P.: Neural field multi-view shape-from-polarisation. In: *Comp. Graph. Forum*. vol. 44. Wiley-Blackwell (2025)
39. Weier, P., Riviere, J., Guseinov, R., Garbin, S., Slusallek, P., Bickel, B., Beeler, T., Vicini, D.: Practical inverse rendering of textured and translucent appearance. *ACM Trans. Graph. (Proc. SIGGRAPH)* **44**(4) (Aug 2025). <https://doi.org/10.1145/3730855>
40. Wolff, L.B.: Material classification and separation of reflection components using polarization/radiometric information. In: *Proceedings of a workshop on Image understanding*. pp. 232–244 (1989)
41. Wolff, L.B.: Using polarization to separate reflection components. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 363–369. IEEE (1989)
42. Wolff, L.B.: Polarization-based material classification from specular reflection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(11), 1059–1071 (1990)
43. Wolff, L.B.: Surface orientation from two camera stereo with polarizers. In: *Optics, Illumination, and Image Sensing for Machine Vision IV*. vol. 1194, pp. 287–297. SPIE (1990)
44. Xu, P., Bangaru, S., Li, T.M., Zhao, S.: Markov-chain monte carlo sampling of visibility boundaries for differentiable rendering. In: *SIGGRAPH Asia Conference Papers* (2024)
45. Yang, X., Liu, Z., Nie, Z., Fan, Q., Shi, Z., Bonar, J., Heidrich, W.: End-to-end differentiable design of geometric waveguide displays (2026)
46. Yang, X., Liu, Z., Nie, Z., Fan, Q., Shi, Z., Bonar, J., Heidrich, W.: End-to-end differentiable design of geometric waveguide displays. *arXiv preprint arXiv:2601.04370* (2026)
47. Zeltner, T.: Mitsuba 3 documentation: Polarization. https://mitsuba.readthedocs.io/en/stable/src/key_topics/polarization.html (2021), [Online; accessed 16-January-2026]

48. Zhang, C., Wu, L., Zheng, C., Gkioulekas, I., Ramamoorthi, R., Zhao, S.: A differential theory of radiative transfer. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **38**(6) (Nov 2019). <https://doi.org/10.1145/3355089.3356522>
49. Zhang, Z., Roussel, N., Jakob, W.: Projective sampling for differentiable rendering of geometry. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* **42**(6) (Dec 2023). <https://doi.org/10.1145/3618385>
50. Zhou, X., Chen, J., Rao, P., Teufel, T., Lyu, L., Minasian, T., Sotnychenko, O., Long, X.X., Habermann, M., Theobalt, C.: Olatverse: A large-scale real-world object dataset with precise lighting control. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 28848–28859 (2026)