

Stokes-Informed Diffusion for Robust Linear Polarization Estimation

Yidong Luo^{1,2*}, Chenggong Li^{3*}, Yuchao Feng⁴, Boxin Shi⁵,
Junchao Zhang^{3†}, and Xin Yuan^{2†}

¹ Zhejiang University, Hangzhou, China

² School of Engineering, Westlake University, Hangzhou, China
{luoyidong, xyuan}@westlake.edu.cn

³ School of Automation, Central South University, Changsha, China
{244603040, junchaozhang}@csu.edu.cn

⁴ College of Computer Science and Technology, Zhejiang University of Technology,
Hangzhou, China

⁵ School of Computer Science, Peking University, Beijing, China

Abstract. Polarization cues benefit applications such as material detection and de-reflection, yet acquiring them typically requires dedicated hardware. This motivates us to estimate the linear polarization from a single RGB image. However, the task is inherently ill-posed, with the Angle of Polarization (AoP) becoming particularly unstable in weak-polarization regions, where the polarimetric signal is overwhelmed by noise, leading to erratic angle estimates. To address these limitations, we propose GenPolar, a Stokes-informed diffusion framework grounded in the Mueller formalism from an intensity observation. Specifically, GenPolar predicts channel-wise linear Stokes components (S_1, S_2) from intensity S_0 , from which degree of linear polarization (DoLP) and AoP are analytically derived; AoP is further supervised with an observability-aware loss. In addition, to enable efficient and high-fidelity inference, we adopt a two-stage training strategy. Firstly, a multi-step conditional diffusion model is trained with a physics-based loss. Subsequently, we distill it into a one-step generator, which further supports stable Low-Rank Adaptation (LoRA) of the VAE encoder to mitigate domain-specific autoencoding bias. Extensive experiments across rotating-polarizer, division-of-focal-plane, and hybrid datasets demonstrate that GenPolar achieves state-of-the-art performance in both DoLP fidelity and AoP stability. Crucially, these improvements translate to significant and consistent gains in downstream applications, including material detection and de-reflection.

Keywords: Linear polarization estimation · Data-physics dual-driven learning · Polarimetric imaging

1 Introduction

Polarization reveals geometric and material cues that are largely imperceptible to standard RGB imaging, supporting applications such as material detection

* Equal contribution. † Corresponding authors.




Fig. 1: (a) DoFP/RLP capture requires specialized hardware. (b) Direct DoLP/AoP regression is unstable in weakly polarized regions. (c) Predicting (S_1, S_2) from S_0 offers a Mueller-grounded alternative. (d) GenPolar further uses AoP observability-aware supervision and LoRA adaptation for more coherent cues.

and reflection separation [6]. However, capturing polarization typically depends on dedicated hardware, i.e., division-of-focal-plane polarimeters DoFP [41], or multi-shot capture with a rotating linear polarizer (RLP; Fig. 1(a)), incurring additional cost and constraining deployment. These limitations motivate single-image polarization estimation from an intensity observation, with the goal of recovering cues that are not only visually consistent but also physically grounded and reliable for downstream reasoning.

Notably, when the observation is an RGB intensity image, polarization is not a single scalar property but exhibits wavelength-dependent behavior [2]. In many common formation processes, the degree of linear polarization (DoLP) varies with wavelength due to wavelength-dependent diattenuation and dispersion, whereas angle of polarization (AoP) is often less sensitive to wavelength under typical surface-reflection settings [3]. Consequently, for an RGB observation, the polarization parameters should be channel-wise estimated to preserve physically meaningful spectral differences in polarization strength.

However, an RGB intensity observation collapses the underlying polarization formation into a severely incomplete measurement, making single-image polarization estimation fundamentally ill-posed: the scene-dependent polarization transfer is unknown and varies spatially with material, geometry, and illumination. In addition, AoP is *not uniformly observable* [4, 36]: when DoLP is small, AoP becomes ill-conditioned and even minor perturbations can induce large angular fluctuations, which explains the speckled AoP maps produced by direct DoLP/AoP regression in Fig. 1(b). A seemingly safer alternative is to predict Stokes components and derive DoLP/AoP analytically, yet optimizing (S_1, S_2) in latent diffusion model [32] with generic reconstruction losses can still yield unstable cues (Fig. 1(c)), partly because Stokes fields follow different statistics from natural RGB images and off-the-shelf latent VAEs introduce structured autoencoding bias. These observations motivate us to seek a physically grounded and spectrally faithful $S_0 \rightarrow (S_1, S_2)$ formulation, together with observability-aware AoP supervision and stable latent-domain adaptation.

In this paper, we propose **GenPolar**, a Stokes-informed diffusion framework grounded in the Mueller formalism [3] for single-image linear polarization estimation. We cast the task as channel-wise $S_{0,\lambda} \rightarrow (S_{1,\lambda}, S_{2,\lambda})$, preserving the spectral dependence of polarization strength, and derive DoLP/AoP as standard cues. To avoid fitting ill-conditioned angles, we supervise AoP only on a DoLP-thresholded observability mask. Finally, we adopt a two-stage training strategy: a physics-supervised multi-step diffusion base model, followed by one-step distillation [40] that enables stable LoRA adaptation [7] of the VAE encoder to reduce polarization-domain autoencoding bias. Experiments on RLP, DoFP, and hybrid datasets show consistent gains in DoLP fidelity and AoP stability, and we further verify their utility on material detection and polarization de-reflection. Our contributions are summarized as:

- A Mueller-grounded, RGB channel-wise $S_0 \rightarrow (S_1, S_2)$ formulation with observability-aware AoP supervision.
- A two-stage diffusion training with one-step distillation enabling stable encoder LoRA adaptation.
- Cross-modality evaluation, with downstream validation on material detection and polarization de-reflection.

2 Related Work

Hardware for polarization acquisition. Polarization is commonly represented by Stokes parameters, from which DoLP and AoP are computed for geometry and material reasoning, often focusing on linear polarization (S_1, S_2) . Measuring polarization requires optical modulation that maps polarization states to intensity observations. Typical systems include DoFP polarimeters with micro-polarizer mosaics (single-shot but spatially multiplexed, requiring polarimetric demosaicking) and rotating-analyzer capture with a rotating linear polarizer (high fidelity but multi-shot and motion-sensitive) [4]. Other single-shot designs include division-of-amplitude [1] and division-of-aperture polarimeters [28] that split the beam/sub-aperture into multiple analyzer channels, as well as liquid-crystal modulation [5] for programmable polarization analysis. These designs trade off optical complexity, calibration burden and reconstruction difficulty, motivating polarization estimation from standard RGB observations.

Polarization restoration and downstream tasks. Since many acquisition pipelines multiplex polarization in space or time, a substantial body of work studies polarization restoration and reconstruction, including DoFP demosaicking [14, 15, 22, 45], denoising [17, 18, 22] and so on, often exploiting correlations across analyzer angles and Stokes-domain constraints. Polarization cues are then used in downstream pipelines such as material detection [10, 24, 25, 34, 37], de-reflection [12, 23, 39], normal estimation [13, 16], low-light enhancement [8, 20, 44] and dehazing [33, 38]. In such applications, AoP can be highly informative yet unstable under weak polarization, so practical systems benefit from explicitly handling angle uncertainty rather than assuming uniformly reliable AoP.

Polarization estimation from RGB. Recent work formulates polarization estimation as predicting polarization information directly from an RGB input. The benchmark study RGB-to-Polarization Estimation [19] establishes RGB-to-polarization as a task and evaluates diverse model families, treating the RGB image as the total intensity (S_0) and estimating polarization components in a Stokes parameterization. In parallel, PolarAnything [42] explores diffusion-based polarimetric synthesis from a single RGB image, and shows that representing polarization with compact cues (e.g., DoLP/AoP) can facilitate generative modeling. Overall, these studies highlight the promise of strong priors for plausible polarization synthesis, while leaving open how to obtain AoP/DoLP that are consistently usable for downstream inference, especially under weak polarization. In contrast, we adopt an observability-aware, physics-grounded Stokes estimation paradigm that targets more stable AoP/DoLP, and assess its practical utility through downstream stress tests.

3 Methodology

3.1 Physical plausibility of $S_{0,\lambda} \rightarrow (S_{1,\lambda}, S_{2,\lambda})$

We represent polarization by the Stokes vector $\mathbf{s} = [S_0, S_1, S_2, S_3]^\top$, where S_0 is the total intensity and (S_1, S_2) encode linear polarization; we assume negligible circular polarization ($S_3 \approx 0$). Given four analyzer-angle measurements $I_{0^\circ}, I_{45^\circ}, I_{90^\circ}, I_{135^\circ}$, the Stokes components are computed as

$$S_0 = \frac{I_{0^\circ} + I_{45^\circ} + I_{90^\circ} + I_{135^\circ}}{2}, \quad S_1 = I_{0^\circ} - I_{90^\circ}, \quad S_2 = I_{45^\circ} - I_{135^\circ}. \quad (1)$$

DoLP and AoP are then computed as standard cues:

$$\text{DoLP} = \frac{\sqrt{S_1^2 + S_2^2}}{S_0}, \quad \text{AoP} = \frac{1}{2} \text{atan2}(S_2, S_1). \quad (2)$$

The task definition is summarized in Fig. 2. The physical plausibility of inferring nonzero (S_1, S_2) from intensity-only input can be viewed through Mueller transport [3]. In natural scenes, the incident illumination can be approximated as unpolarized, since it is a mixture of many polarization states with random orientations. For incident light $\mathbf{s}_{in} = [I, 0, 0, 0]^\top$, the outgoing Stokes $\mathbf{s}_{out}$ satisfies

$$\mathbf{s}_{out} = \begin{bmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{bmatrix} = \mathbf{M} \cdot \mathbf{s}_{in} = \begin{bmatrix} m_{0,0} & m_{0,1} & m_{0,2} & m_{0,3} \\ m_{1,0} & m_{1,1} & m_{1,2} & m_{1,3} \\ m_{2,0} & m_{2,1} & m_{2,2} & m_{2,3} \\ m_{3,0} & m_{3,1} & m_{3,2} & m_{3,3} \end{bmatrix} \begin{bmatrix} I \\ 0 \\ 0 \\ 0 \end{bmatrix} = I \begin{bmatrix} m_{0,0} \\ m_{1,0} \\ m_{2,0} \\ m_{3,0} \end{bmatrix}. \quad (3)$$

Here $\mathbf{M}$ is the scene Mueller matrix, and $S_3 \approx 0$ in linear polarization. $m_{0,0}$ measures the intensity throughput, while $m_{1,0}$ and $m_{2,0}$ quantify how unpolarized intensity is converted into linear polarization along the $0^\circ/90^\circ$ and $45^\circ/135^\circ$ bases, respectively, indicating that nonzero $m_{1,0}$ or $m_{2,0}$ can produce linear polarization. Since $S_0 = m_{0,0}I$, $S_1 = m_{1,0}I$, and $S_2 = m_{2,0}I$ in Eq. (3), the linear polarization components are coupled to the intensity via the Mueller parameters.




Fig. 2: Physical motivation for the Stokes-estimation formulation. (a) Under an unpolarized illumination approximation, Mueller transport implies an effective $S_0 \rightarrow (S_1, S_2)$ relation, making intensity-to-linear-polarization estimation physically plausible. (b) Polarization strength is spectral, motivating channel-wise estimation.

We therefore define the task as estimating (S_1, S_2) from S_0 . Under the standard assumption of unpolarized incident illumination, Mueller transport implies

$$S_1 = \frac{m_{1,0}}{m_{0,0}} S_0, \quad S_2 = \frac{m_{2,0}}{m_{0,0}} S_0, \quad (4)$$

where the normalized ratios $\frac{m_{1,0}}{m_{0,0}}$ and $\frac{m_{2,0}}{m_{0,0}}$ can be viewed as components of an effective diattenuation, which varies with local material, geometry, and illumination. Importantly, Eq. (4) is a physics-derived structural constraint from Mueller transport under unpolarized illumination: it shows that (S_1, S_2) can be written as a spatially varying linear transfer of S_0 , with coefficients determined by scene polarimetric properties. It provides a physically plausible parameterization that our model learns under physics supervision. Together with S_0 , the predicted (S_1, S_2) allows us to compute DoLP and AoP via Eq. (2). In scenarios with partially polarized illumination or complex propagation (e.g., skylight), Eq. (4) should be understood in an effective sense.

A key formation property motivating our RGB channel-wise prediction is that polarization strength depends on wavelength and incidence geometry. For a broad class of diattenuation/Fresnel effects [36], linear Stokes components admit the amplitude-orientation form

$$[S_{1,\lambda}, S_{2,\lambda}] = \rho(\lambda, \theta, n(\lambda)) S_{0,\lambda} [\cos 2\psi, \sin 2\psi], \quad (5)$$

where $\rho(\lambda, \theta, n(\lambda))$ captures the polarization strength and depends on wavelength λ , the local incidence angle θ , and the refractive index $n(\lambda)$ of the material, which accounts for wavelength-dependent diattenuation and dispersion. The angle of polarization ψ reflects the orientation of the plane of incidence, primarily determined by macroscopic geometry, such as surface normals and viewing directions.




Fig. 3: AoP reliability analysis across public datasets. The line plot shows the decreasing trend of cross-channel error with increasing DoLP. The pie chart illustrates the pixel-wise DoLP distribution across all images in public datasets. True-color visualizations of DoLP and AoP clearly reveal pronounced noise in weakly polarized regions.

This decomposition yields $\text{DoLP}_\lambda = |\rho(\lambda, \theta, n(\lambda))|$ and $\text{AoP}_\lambda = \psi$. Thus DoLP is inherently wavelength-dependent, motivating estimation of $(S_{1,\lambda}, S_{2,\lambda})$ for all RGB channels rather than collapsing polarization to grayscale.

A seemingly simpler alternative is to regress $(\text{DoLP}_\lambda, \cos(2\text{AoP}_\lambda), \sin(2\text{AoP}_\lambda))$ directly from $S_{0,\lambda}$. Yet this discards the explicit Stokes-level structure in Eq. (4) and does not resolve the instability of AoP under weak polarization. We quantify this instability on public datasets by cross-channel AoP coherence, defined as the mean angular deviation of one channel relative to the other two, which serves as a practical reliability proxy. As shown in Fig. 3, this deviation is much larger at low DoLP and decreases rapidly as DoLP increases. These observations motivate an observability-aware objective and evaluation that explicitly account for where polarization cues are reliable.

Overall, we implement $S_{0,\lambda} \rightarrow (S_{1,\lambda}, S_{2,\lambda})$ using a pretrained stable diffusion model and fine-tune it on polarization data, letting the generative prior regularize this ill-posed inference while learning the spatially varying polarization transfer implied by Eq. (4). Moreover, the pretrained model provides a rich representation: its features encode appearance cues correlated with geometry, material, and illumination factors that drive the transfer in Eq. (4).

3.2 Stokes-Informed Stable Diffusion

We propose a Stokes-Informed Stable Diffusion framework for generating polarization components $(S_{1,\lambda}, S_{2,\lambda})$ from RGB intensity proxy proportional to $S_{0,\lambda}$. The overall pipeline is shown in Fig. 4. Our method operates in a clear inference process, followed by two training stages. At inference time, the model receives an RGB intensity proxy c proportional to $S_{0,\lambda}$ and Gaussian noise ϵ_0 . The one-step generator $\mathcal{G}_\theta$ predicts the clean latent $z_0^{\text{pred}} = \mathcal{G}_\theta(\epsilon_0, c, t^*)$, where t^* is the fixed timestep corresponding to the maximum noise level. This latent z_0^{pred} is then decoded by $\mathcal{D}_{\text{vae}}$ to obtain $(\hat{S}_{1,\lambda}, \hat{S}_{2,\lambda})$.

Directly encoding and decoding polarization targets with the pretrained VAE yields noticeable errors (Fig. 5), indicating a mismatch between polarization maps and the pre-trained latent representation. We therefore insert LoRA adapters [7] into the VAE encoder $\mathcal{E}_{\text{vae}}$ to mitigate this autoencoding bias, while




Fig. 4: Overview of GenPolar. Given RGB intensity proxy S_0 , GenPolar predicts (S_1, S_2) and derives DoLP/AoP. Stage I trains a multi-step diffusion model with physics supervision, and Stage II distills it into a one-step generator with VAE-encoder LoRA.

keeping the decoder fixed to maintain a consistent output parameterization. However, stably adapting $\mathcal{E}_{vae}$ requires end-to-end optimization, which is impractical to backpropagate through long multi-step diffusion trajectories. This motivates our two-stage training: (i) Stage I trains the multi-step diffusion backbone μ_θ and ControlNet [43] C_θ while keeping the VAE frozen. (ii) Stage II distills the multi-step diffusion process into a one-step generator $\mathcal{G}_\theta$ and adapts the $\mathcal{E}_{vae}$ via LoRA. As evidenced in Fig. 5, the LoRA-adapted encoder substantially reduces the reconstruction error, validating our two-stage training scheme.

Training stage I: In Stage I, we start with a pretrained VAE encoder that produces a 4-channel latent representation. Since $(S_{1,\lambda}, S_{2,\lambda})$ form a 6-channel target (denoting as $x = [S_{1,\lambda}, S_{2,\lambda}] \in \mathbb{R}^{6 \times H \times W}$), we encode the 3-channel $S_{1,\lambda}$ and $S_{2,\lambda}$ separately using the same encoder, then concatenate the resulting latents into an 8-channel tensor $z_0 \in \mathbb{R}^{8 \times \frac{H}{8} \times \frac{W}{8}}$. The noisy latent z_t is given by

$$z_t = \alpha_t z_0 + \sigma_t \epsilon, \quad (6)$$

where α_t and σ_t are determined by the noise schedule, and $\epsilon \sim \mathcal{N}(0, I)$ is Gaussian noise. The diffusion UNet μ_θ is adapted to handle the 8-channel latents, and the C_θ is constructed to inject the RGB structural guidance, while keeping the VAE frozen. Let $x = [S_{1,\lambda}, S_{2,\lambda}]$ denote the polarization targets and c the RGB conditioning image proportional to $S_{0,\lambda}$. We adopt an z_0 -prediction parameterization to train μ_θ and C_θ with the clean latent prediction loss

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{x,c,\epsilon,t} \left[\left\| \mu_\theta(z_t, c, t) - z_0 \right\|_2^2 \right]. \quad (7)$$

To enforce physical consistency, we decode the predicted latent $\hat{z}_0 = \mu_\theta(z_t, c, t)$ to obtain $(\hat{S}_{1,\lambda}, \hat{S}_{2,\lambda})$, and compute $\widehat{\text{DoLP}}_\lambda$ and $\widehat{\text{AoP}}_\lambda$ via Eq. (2). Let Ω denote all pixels and $\Omega_\tau = \{\mathbf{y} \in \Omega \mid \text{DoLP}_\lambda(\mathbf{y}) > \tau\}$ be the observability mask defined




Fig. 5: The error map is significantly reduced after LoRA adaptation, demonstrating its effectiveness in improving the reconstruction of polarization targets.

by a DoLP threshold τ . We supervise linear Stokes components globally with

$$\mathcal{L}_s = \frac{1}{|\Omega|} \sum_{\mathbf{y} \in \Omega} \left(|\hat{S}_{1,\lambda}(\mathbf{y}) - S_{1,\lambda}(\mathbf{y})| + |\hat{S}_{2,\lambda}(\mathbf{y}) - S_{2,\lambda}(\mathbf{y})| \right), \quad (8)$$

and constrain AoP only on Ω_τ , where Ω_τ is computed from GT DoLP. This follows our observation in Fig. 3 that AoP becomes ill-conditioned under weak polarization, while high-DoLP regions exhibit consistent AoP cues across channels. We define an angular distance $d_{\text{ang}}(\cdot, \cdot)$ that is π -periodic, and write

$$\mathcal{L}_{\text{AoP}} = \frac{1}{|\Omega_\tau|} \sum_{\mathbf{y} \in \Omega_\tau} d_{\text{ang}}\left(\widehat{\text{AoP}}_\lambda(\mathbf{y}), \text{AoP}_\lambda(\mathbf{y})\right) \quad (9)$$

as AoP observability-aware loss. Thus, the physics loss is $\mathcal{L}_{\text{phys}} = \mathcal{L}_s + \gamma_{\text{AoP}} \mathcal{L}_{\text{AoP}}$, and the overall training objective for the first stage is

$$\mathcal{L}_{\text{stg}_1} = \mathcal{L}_{\text{diff}} + \gamma_p \mathcal{L}_{\text{phys}}. \quad (10)$$

Training stage II: In Stage II, we convert the multi-step diffusion process learned in Stage I into a one-step generator $\mathcal{G}_\theta$. This stage accelerates inference and, more importantly, enables stable end-to-end adaptation of the VAE encoder via LoRA: backpropagating through a long stochastic denoising trajectory is computationally expensive and often yields high-variance gradients for the encoder, which can lead to unstable updates. By distilling the diffusion model into a one-step mapping, we shorten the backpropagation path to a single denoising step, yielding a well-conditioned end-to-end generator that supports reliable gradient propagation while allowing task-specific physical supervision during distillation. We adopt distribution match distillation (DMD) [40] for this conversion, as it provides a distribution-level distillation objective by matching real/fake score functions, leading to more stable one-step training under our physics supervision. We equip the $\mathcal{E}_{vae}$ with LoRA adapters and denote the resulting encoder as $\mathcal{E}_\phi$. For each training sample, we obtain the clean latent and apply forward noising:

$$\tilde{z}_{t^*} = \alpha_{t^*} \mathcal{E}_\phi(x) + \sigma_{t^*} \epsilon, \quad (11)$$

Algorithm 1: Training Procedure for Stokes-Informed Stable Diffusion

Input: Paired dataset $\{(x, c)\}$ with $x = [S_{1,\lambda}, S_{2,\lambda}]$ and $c \propto S_{0,\lambda}$;

Output: One-step generator $\mathcal{G}_\theta$, and LoRA parameters ϕ

- 1 **Stage I: Train multi-step diffusion (freeze VAE);**
 - 2 Freeze VAE parameters (encoder/decoder), train UNet μ_θ and ControlNet $\mathcal{C}_\theta$;
 - 3 **while Stage I training do**
 - 4 Sample batch $(x, c)^B$;
 - 5 $z_0 \leftarrow \mathcal{E}_{vae}(x)$ ($\mathbb{R}^{B \times 8 \times \frac{H}{8} \times \frac{W}{8}} \leftarrow \mathbb{R}^{2B \times 3 \times H \times W}$);
 - 6 Sample t and $\epsilon \sim \mathcal{N}(0, I)$; obtain z_t using Eq. (6);
 - 7 $\hat{z}_0 \leftarrow \mu_\theta(z_t, c, t)$;
 - 8 $(\hat{S}_{1,\lambda}, \hat{S}_{2,\lambda}) \leftarrow \mathcal{D}_{vae}(\hat{z}_0)$ ($\mathbb{R}^{2B \times 3 \times H \times W} \leftarrow \mathbb{R}^{B \times 8 \times \frac{H}{8} \times \frac{W}{8}}$);
 - 9 Compute $\mathcal{L}_{\text{diff}}$ using Eq. (7), $\mathcal{L}_s$ using Eq. (8) and $\mathcal{L}_{\text{AoP}}$ using Eq. (9);
 - 10 Form $\mathcal{L}_{\text{phys}} = \mathcal{L}_s + \gamma_{\text{AoP}} \mathcal{L}_{\text{AoP}}$ and $\mathcal{L}_{\text{stg1}}$ using Eq. (10);
 - 11 Update μ_θ and $\mathcal{C}_\theta$ by minimizing $\mathcal{L}_{\text{stg1}}$;
 - 12 **Stage II: One-step distillation + LoRA fine-tuning;**
 - 13 Initialize $\mathcal{G}_\theta$ from μ_θ ; Freeze μ_θ , $\mathcal{C}_\theta$, and $\mathcal{D}_{vae}$; Insert LoRA into the $\mathcal{E}_{vae}$;
 - 14 **while Stage II training do**
 - 15 Sample batch $(x, c)^B$;
 - 16 $\tilde{z}_0 \leftarrow \mathcal{E}_\phi(x)$ ($\mathbb{R}^{B \times 8 \times \frac{H}{8} \times \frac{W}{8}} \leftarrow \mathbb{R}^{2B \times 3 \times H \times W}$);
 - 17 Sample t^* and $\epsilon \sim \mathcal{N}(0, I)$; obtain $\tilde{z}_{t^*}$ using Eq. (11);
 - 18 $\hat{\tilde{z}}_0 \leftarrow \mathcal{G}_\theta(\tilde{z}_{t^*}, c, t^*)$;
 - 19 Compute g_{dmd} using Eq. (12);
 - 20 $(\hat{S}_{1,\lambda}, \hat{S}_{2,\lambda}) \leftarrow \mathcal{D}_{vae}(\hat{\tilde{z}}_0)$ ($\mathbb{R}^{2B \times 3 \times H \times W} \leftarrow \mathbb{R}^{B \times 8 \times \frac{H}{8} \times \frac{W}{8}}$);
 - 21 Compute $\mathcal{L}_{\text{phys}}$ using Eq. (8)–Eq. (9); Compute $\mathcal{L}_{\text{kl}}$ using Eq. (14);
 - 22 Update θ (in $\mathcal{G}_\theta$) and ϕ (LoRA in $\mathcal{E}_\phi$) using Eq. (15);
-

where $(\alpha_{t^*}, \sigma_{t^*})$ are taken from the Stage-I noise schedule at a fixed timestep t^* . We distill the Stage-I multi-step diffusion $\mu_{\text{real}} = \mu_\theta$ into a one-step generator $\mathcal{G}_\theta$ using DMD [40], which minimizes an approximate KL divergence $\text{KL}(p_{\text{fake}} \| p_{\text{real}})$ via score differences. Let $\hat{\tilde{z}}_0 = \mathcal{G}_\theta(\tilde{z}_t, t^*, c)$ be the one-step denoised latent and $\hat{\tilde{z}}_t = \alpha_t \hat{\tilde{z}}_0 + \sigma_t \epsilon$ its re-noised version under the same schedule; the KL gradient is estimated by the real/fake scores computed from the frozen μ_{real} and an online-trained auxiliary generator μ_{fake} :

$$g_{\text{dmd}} = \mathbb{E}_{t,c,\epsilon} \left[\left(s_{\text{fake}}(\hat{\tilde{z}}_t, t, c) - s_{\text{real}}(\hat{\tilde{z}}_t, t, c) \right) \cdot \frac{\partial \hat{\tilde{z}}_t}{\partial \theta} \right]. \quad (12)$$

where $s_{\text{real}}(\hat{\tilde{z}}_t, t, c)$ and $s_{\text{fake}}(\hat{\tilde{z}}_t, t, c)$ are computed by the μ_{real} and μ_{fake} using:

$$s_{\text{real}}(\hat{\tilde{z}}_t, t, c) = -\frac{\hat{\tilde{z}}_t - \alpha_t \mu_{\text{real}}(\hat{\tilde{z}}_t, t, c)}{\sigma_t^2}, \quad s_{\text{fake}}(\hat{\tilde{z}}_t, t, c) = -\frac{\hat{\tilde{z}}_t - \alpha_t \mu_{\text{fake}}(\hat{\tilde{z}}_t, t, c)}{\sigma_t^2}. \quad (13)$$

We keep μ_{real} frozen and train μ_{fake} online Eq. (7), while updating θ using Eq. (12). We reuse the physics loss $\mathcal{L}_{\text{phys}}$ defined in Stage I after decoding the predicted latent, and noting that $\mathcal{L}_s$ also serves as a regression loss for distillation. Since LoRA adapters are inserted into the encoder $\mathcal{E}_\phi$, we regularize its

posterior to ensure that the learned latent z distribution aligns with the standard Gaussian prior $p(z) = \mathcal{N}(0, I)$, as in standard VAEs. Specifically, the encoder defines $q_\phi(z|x) = \mathcal{N}(m_\phi(x), \text{diag}(v_\phi^2(x)))$ with prior $p(z) = \mathcal{N}(0, I)$, where m_ϕ/v_ϕ denote mean/standard deviation, and the KL term of LoRA is

$$\mathcal{L}_{\text{kl}} = \mathbb{E}_x [\text{KL}(q_\phi(z|x) || p(z))]. \quad (14)$$

In Stage II, we update θ and ϕ with learning rate η_θ and η_ϕ as follows:

$$\theta \leftarrow \theta - \eta_\theta (g_{\text{dmd}} + \nabla_\theta \mathcal{L}_{\text{phys}}), \quad \phi \leftarrow \phi - \eta_\phi (\nabla_\phi \mathcal{L}_{\text{phys}} + \nabla_\phi \mathcal{L}_{\text{kl}}), \quad (15)$$

During Stage II, we jointly optimize the one-step generator parameters θ and the LoRA parameters ϕ inserted in the encoder ($\mathcal{D}_{\text{vae}}$ fixed). In practice, g_{dmd} distills the multi-step denoising behavior into $\mathcal{G}_\theta$, $\mathcal{L}_{\text{kl}}$ regularizes the LoRA-adapted encoder posterior, and $\mathcal{L}_{\text{phys}}$ provides physical supervision that is applied for both θ and ϕ to improve the stability of DoLP/AoP estimation. The overall training is summarized in Algorithm 1.

4 Experiments

4.1 Experimental Setup

Data preparation. We train and evaluate on multiple polarization datasets spanning two real capture modalities and one hybrid set: (i) Rotating linear polarizer group. This group includes Monno [26, 31], Qiu [30], PIDSr [45], and Wen [35], which share the acquisition protocol of RLP captures (363 samples). We additionally capture 235 low-noise scenes using the same rotating-polarizer setup, yielding a total of 598 samples in the RLP group. (ii) Division-of-focal-plane polarimeter group. This group includes PolarFree [39], CPIF [21], and Jeon [9], all collected by DoFP cameras, totaling 2188 samples. (iii) RSP group. RSP [11] is treated as a separate group containing a mixture of simulated and real measurements, totaling 2000 samples. Across all datasets, we compute Stokes components following Eq. (1) (or the equivalent processing provided by the dataset when the native measurement format differs), and use $c \propto S_{0,\lambda}$ as the conditioning input and $x = [S_{1,\lambda}, S_{2,\lambda}]$ as supervision. Furthermore, we split each group independently into training and testing sets. For each group, we randomly select 50 samples for testing; the rest are used for training.

Metrics. We evaluate polarization cues computed from the predicted Stokes components. Given predicted $(\hat{S}_{1,\lambda}, \hat{S}_{2,\lambda})$ and the conditional input $S_{0,\lambda}$, we compute $\widehat{\text{DoLP}}_\lambda$ and $\widehat{\text{AoP}}_\lambda$ via Eq. (2). We report PSNR/SSIM on $\widehat{\text{DoLP}}_\lambda$ and MeanAE on $\widehat{\text{AoP}}_\lambda$. Metrics are computed per color channel and averaged over $\lambda \in \{R, G, B\}$. AoP is π -periodic and unreliable under weak polarization; hence, we compute AoP-MeanAE only on $\Omega_{\tau_{\text{eval}}} = \{\mathbf{y} \mid \text{DoLP}_\lambda(\mathbf{y}) > \tau_{\text{eval}}\}$ (Sec. 3) using a π -periodic angular distance. For fair comparison, $\Omega_{\tau_{\text{eval}}}$ is constructed from GT DoLP and is shared across all methods. Following the analysis in Fig. 2,




Fig. 6: Visual comparisons among baselines and our GenPolar.

we fix $\tau_{eval} = 0.05$ for all experiments as a practical trade-off between AoP reliability and spatial coverage: it yields a reasonable cross-channel AoP consistency while retaining $\sim 70\%$ of pixels in $\Omega_{\tau_{eval}}$.

Baselines. We retrain all baselines on the same union training set and evaluate them on identical group-wise test splits: PolarAnything [42], Restormer, and MAE from RGB-to-Polarization [19]. All methods use the same RGB intensity proxy $c \propto S_0$ as input. Since PolarAnything predicts grayscale DoLP/AoP cues whereas GenPolar predicts channel-wise RGB Stokes components, we preserve each method’s original output representation during training and convert all outputs to a common DoLP/AoP space for evaluation.

Implementation details. We use SD1.5 as the base model and optimize all trainable modules with AdamW, using betas (0.9, 0.999) and weight decay 10^{-3} . Stage I trains $(\mu_\theta, \mathcal{C}_\theta)$ for 100 epochs with a learning rate of 4×10^{-5} , while Stage II jointly trains the one-step generator $\mathcal{G}_\theta$ and encoder LoRA parameters ϕ for 30 epochs with a learning rate of 10^{-5} , keeping μ_θ and $\mathcal{D}_{vae}$ fixed. We set $\gamma_p = \gamma_{AoP} = 0.5$, LoRA rank to 4, patch size to 512, and batch size to 1. S_1 and S_2 are directly used in $[-1, 1]$, while $S_0 \in [0, 2]$ is scaled to $[-1, 1]$. ControlNet is initialized from the diffusion UNet. Stage-I inference uses 20 steps, and training is performed on 4 A40 GPUs.

4.2 Comparison with Prior Methods

Qualitative results. Fig. 6 compares pseudo-colored DoLP and AoP predictions across methods. As PolarAnything only generates grayscale DoLP and AoP, pseudo-coloring is used for visualization consistency, and true-color results are provided in the supplementary material (SM). The first column shows the RGB input (S_0 proxy), followed by DoLP/AoP visualizations for all methods. GenPolar produces more coherent DoLP fields and AoP maps that better match the ground truth, with fewer artifacts and angular outliers. In contrast, PolarAnything tends to deviate from the GT in the zoomed DoFP and RSP

Table 1: Main comparison on polarization cues across test groups.

Method	RLP test			DoFP test			RSP test		
	PSNR $\uparrow$	SSIM $\uparrow$	MeanAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	MeanAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	MeanAE $\downarrow$
Restormer	20.64	0.540	35.35	16.82	0.532	32.72	15.15	0.588	35.17
MAE	26.98	0.638	28.23	21.40	0.598	29.82	16.67	0.616	33.51
PolarAnything	19.77	0.563	33.56	17.85	0.572	29.33	15.47	0.604	30.45
GenPolar (Ours)	29.54	0.741	20.15	24.01	0.666	18.77	21.29	0.676	27.60

regions, likely because direct DoLP/AoP prediction is more sensitive to AoP ambiguity in weakly polarized areas, while its channel-collapsed output provides less flexibility for wavelength-dependent polarization cues. Among deterministic baselines, Restormer improves upon MAE in some cases but is less robust to sensor artifacts, while MAE is stable but produces blurred DoLP boundaries and inconsistent AoP patterns. GenPolar’s performance benefits from two key factors: (i) an AoP observability-aware objective that stabilizes angles in weak-polarization areas, and (ii) LoRA-based VAE encoder adaptation, which reduces polarization-domain bias and stabilizes downstream DoLP/AoP.

Quantitative results across modalities. Tab. 1 reports DoLP-PSNR/SSIM and AoP-MeanAE on the three test groups. GenPolar achieves the best overall performance across capture modalities, with a clear advantage in AoP stability (lower MeanAE), especially on the DoFP and RSP groups where sensor artifacts and domain shifts make AoP reconstruction particularly challenging.

4.3 Ablation Studies



We ablate GenPolar on the RLP testset, which provides a controlled setting with consistent acquisition. As shown in Tab. 2, Fig. 7 and 8, the quantitative metrics and visual examples are displayed. Detailed analysis is as follows.

Stage II enables stable VAE adaptation. We report four ablations. (i) The Stage-I multi-step model with a frozen VAE already yields reasonable polarization cues. (ii) Pre-adapting $\mathcal{E}_{vae}$ with LoRA via reconstruction losses [32] on our polarization data does not help either; using this adapted $\mathcal{E}_{vae}$ in Stage-I training further degrades the multi-step model, indicating that encoder adaptation is impractical under long diffusion trajectories. (iii) Directly inserting LoRA into $\mathcal{E}_{vae}$ during Stage-I multi-step training is unstable: DoLP remains plausible but AoP collapses. (iv) Distilling to a one-step generator resolves this optimization issue, but without LoRA the distilled model still inherits a latent-domain mismatch, leading to larger AoP error than the full model. Overall, one-step distillation is necessary to enable stable encoder adaptation, and LoRA is further needed to correct the latent mismatch for accurate AoP.

Diffusion vs. E2E. For E2E + $\mathcal{L}_{phys}$, we replace diffusion sampling with a single feed-forward pass, using the diffusion UNet to directly map the RGB latent to the (S_1, S_2) latent and training it end-to-end with $\mathcal{L}_{phys}$. This baseline underperforms substantially, suggesting that the multi-step base diffusion provides useful distributional regularization beyond a single-pass regression network.

Table 2: Quantitative ablations on the RLP testset.

Setting	PSNR $\uparrow$	SSIM $\uparrow$	MeanAE $\downarrow$
Only Stage I	28.84	0.684	22.61
Pre-LoRA	27.92	0.674	23.49
LoRA in Stage I	19.78	0.460	50.03
w/o LoRA	28.49	0.681	23.12
E2E + $\mathcal{L}_{\text{phys}}$	24.15	0.553	30.73
w/o $\mathcal{L}_{\text{AoP}}$	27.21	0.668	23.96
$\tau = 0$	29.47	0.724	21.89
Gray Stokes	27.43	0.621	23.56
$S_0 \rightarrow [\text{DoLP}, \text{AoP}]$	22.84	0.502	35.40
Ours	29.54	0.741	20.15

**Fig. 7:** Ablation in τ_{train} of $\mathcal{L}_{\text{AoP}}$.**Fig. 8:** Visualization for ablation.

Observability-aware AoP supervision. Removing $\mathcal{L}_{\text{AoP}}$ ($\tau_{train} = 1$) primarily harms AoP. The optimal AoP stability is achieved at $\tau_{train} = 0.05$ according to Fig. 7, where the loss focuses on observable regions, preventing AoP from being enforced in weakly polarized areas.

Channel-wise matters for color inputs. Gray-Stokes collapses ($S_{1,\lambda}, S_{2,\lambda}$) by replacing ($S_{1,R}, S_{1,G}, S_{1,B}$) with their channel mean (similarly for S_2), which degrades polarization cues and increases angular inconsistency. This supports our channel-wise Stokes formulation, which respects wavelength-dependent polarization strength and avoids the blurring effect introduced by channel averaging.

Direct DoLP/AoP regression is unstable. Obvious degradation occurs for direct DoLP, $\sin(2\text{AoP})$, and $\cos(2\text{AoP})$ regression, since small $\sin/\cos$ errors can translate into large angular deviations after inversion and this representation does not preserve Stokes-level coupling between magnitude and orientation.


Overall, the ablations show that robust linear polarization estimation benefits from jointly addressing representation, observability, and latent adaptation. Predicting per-channel ($S_{1,\lambda}, S_{2,\lambda}$) avoids channel averaging that blurs wavelength-dependent polarization strength, while observability-aware AoP supervision improves angular stability by avoiding ill-conditioned regions. Finally, the one-step generation training stage is essential for stable LoRA adaptation of the VAE encoder, yielding the strongest DoLP fidelity and AoP stability.

4.4 Downstream Applications


We further evaluate whether the estimated polarization cues are usable for downstream pipelines, rather than only improving cue-level metrics. We consider two

Table 3: Material accuracy (dielectric vs. metallic).

Method	Acc. (%) $\uparrow$
Captured	93.75
MAE	82.75
Restormer	75.12
PolarA.	45.36
Ours	90.12

**Fig. 9:** Qualitative material detection results (input DoLP/AoP and predicted masks).**Table 4:** De-reflection comparison in PSNR and LPIPS.

Method	PSNR $\uparrow$	LPIPS $\downarrow$
Captured	22.44	0.133
MAE	22.04	0.147
Restormer	21.87	0.169
PolarA.	20.08	0.190
Ours	22.41	0.127

**Fig. 10:** Visual input and de-reflected output.

representative tasks with different output types: (i) material detection as a non-regression task (discrete categories), and (ii) polarization de-reflection as a regression task (continuous image restoration). In both tasks, we keep the downstream model fixed and only replace its polarization inputs (DoLP/AoP) with those computed from different methods, including polarization derived from measured captures and from four estimation baselines.

Material detection. We adopt SfPUEL [24] to segment two material types (dielectric and metallic), and we evaluate detection accuracy on its real test set. For each method, we feed SfPUEL the corresponding DoLP/AoP cues, and obtain a material mask. Fig. 9 visualizes the input DoLP/AoP and the resulting predictions. We observe that material discrimination is sensitive to both DoLP and AoP quality: specular reflections can corrupt polarization cues and lead to incomplete or fragmented masks. In particular, PolarAnything often confuses material categories and produces incorrect masks, while other baselines and even measured cues may suffer from reflection-induced artifacts. In contrast, GenPolar performs best among estimation methods and is close to the captured-input reference in both accuracy and visual quality, indicating that more coherent polarization cues improve downstream material reasoning.

Polarization de-reflection. We evaluate polarization-based de-reflection using PolarFree [39]. Following the same protocol, we keep PolarFree fixed and only replace its polarization inputs with those from different methods. We report PSNR and LPIPS. Tab. 4 and Fig. 10 shows that GenPolar consistently outperforms estimation baselines and remains close to the captured-input reference. Notably, PolarAnything is less catastrophic in this regression setting than

in material detection, which is consistent with the visual results: it can partially suppress reflections but often leaves residual artifacts and loses fine details.

Across a non-regression material detection task and a regression de-reflection task, improved polarization cues translate to better downstream performance. These results serve as empirical evidence that observability-aware Stokes estimation can provide more usable DoLP/AoP cues for practical pipelines.

4.5 Limitations

GenPolar focuses on linear polarization, estimating (S_1, S_2) and the derived DoLP/AoP, but does not recover the circular Stokes component S_3 ; full-Stokes estimation from RGB remains future work. Our effective unpolarized-illumination Mueller/Fresnel approximation may be less reliable under polarized illumination, skylight, transparent or multiple-scattering paths, and severe domain shift. These cases do not necessarily lead to failure, but they can weaken the RGB-to-polarization constraint; once DoLP is underestimated, AoP becomes less observable and its estimation accuracy inevitably degrades. Handling such coupled DoLP–AoP uncertainty remains an open problem. Representative failure cases are provided in the SM.

5 Conclusion

We presented GenPolar, a Stokes-informed stable diffusion framework for robust linear polarization estimation from a single RGB intensity proxy proportional to S_0 . Our key insight is to predict channel-wise (S_1, S_2) and treat AoP as observable only when DoLP is sufficiently high, incorporating this observability into both training and evaluation. To improve efficiency and reduce latent-domain mismatch, we adopt a two-stage pipeline that first trains a conditional diffusion model with physics-driven polarization supervision and then distills it into a one-step generator with lightweight encoder LoRA adaptation. Experiments across RLP, DoFP, and hybrid datasets show improved DoLP fidelity and AoP stability over retrained baselines, with consistent gains in downstream material detection and polarization-based de-reflection. Code is available here.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (grant number U2541205, 62271414), National Key R&D Program of China (2024YFF0505603), “Pioneer” and “Leading Goose” R&D Program of Zhejiang (grant number 2024SDXHDX0006, 2024C03182), the 2023 International Sci-tech Cooperation Projects under the purview of the “Innovation Yongjiang 2035” Key R&D Program (grant number 2024Z126), the National Natural Science Foundation of China (grant number 62105372) and Hunan Provincial Research and Development Project (grant number 2025QK3019).

References

1. Azzam, R.M.A.: Division-of-amplitude photopolarimeter (DOAP) for the simultaneous measurement of all four Stokes parameters of light. *Optica Acta: International Journal of Optics* **29**(5), 685–689 (1982)
2. Azzam, R.M.A., Bashara, N.M.: *Ellipsometry and Polarized Light*. North-Holland, Amsterdam (1977)
3. Chenault, D.B., Chipman, R.A.: Measurements of linear diattenuation and linear retardance spectra with a rotating sample spectropolarimeter. *Applied Optics* **32**(19), 3513–3519 (1993)
4. Chipman, R., Lam, W.S.T., Young, G.: *Polarized Light and Optical Systems*. CRC Press (2018)
5. Coleman, D.A., Fernsler, J., Chattham, N., Nakata, M., Takanishi, Y., Korblova, E., Link, D.R., Shao, R.F., Jang, W.G., MacLennan, J.E., et al.: Polarization-modulated smectic liquid crystal phases. *Science* **301**(5637), 1204–1211 (2003)
6. Collett, E.: *Polarized light: Fundamentals and applications*. Marcel Dekker, New York (1993)
7. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: LoRA: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021)
8. Hu, H., Lin, Y., Li, X., Qi, P., Liu, T.: IPLNet: a neural network for intensity-polarization imaging in low light. *Optics Letters* **45**(22), 6162–6165 (2020)
9. Jeon, Y., Choi, E., Kim, Y., Moon, Y., Omer, K., Heide, F., Baek, S.H.: Spectral and polarization vision: Spectro-polarimetric real-world dataset. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 22098–22108 (2024)
10. Kalra, A., Taamazyan, V., Rao, S.K., Venkataraman, K., Raskar, R., Kadambi, A.: Deep polarization cues for transparent object segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 8602–8611 (2020)
11. Kurita, T., Kondo, Y., Sun, L., Moriuchi, Y.: Simultaneous acquisition of high quality RGB image and polarization information using a sparse polarization sensor. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 178–188 (2023)
12. Lei, C., Huang, X., Zhang, M., Yan, Q., Sun, W., Chen, Q.: Polarized reflection removal with perfect alignment in the wild. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1750–1758 (2020)
13. Lei, C., Qi, C., Xie, J., Fan, N., Koltun, V., Chen, Q.: Shape from polarization for complex scenes in the wild. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 12632–12641 (2022)
14. Li, C., Luo, Y., Wu, C., Zhang, J., Yang, D., Zhao, D.: Demosaicking customized diffusion model for snapshot polarization imaging. *Optics and Laser Technology* **188**, 112868 (2025)
15. Li, C., Luo, Y., Zhang, J., Yang, D.: Polarization uncertainty-guided diffusion model for color polarization image demosaicking. In: *AAAI*. vol. 40, pp. 6028–6036 (2026)
16. Li, C., Ono, T., Uemori, T., Mihara, H., Gatto, A., Nagahara, H., Moriuchi, Y.: NeISF: Neural incident stokes field for geometry and material estimation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 21434–21445 (2024)
17. Li, N., Wang, B., Goudail, F., Zhao, Y., Pan, Q.: Joint denoising-demosaicking network for long-wave infrared division-of-focal-plane polarization images with mixed noise level estimation. *IEEE Trans. Image Process.* **32**, 5961–5976 (2023)

18. Li, Z., Jiang, H., Cao, M., Zheng, Y.: Polarized color image denoising. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 9873–9882 (2023)
19. Lin, B., Yuan, Z., Chen, T.: RGB-to-polarization estimation: A new task and benchmark study. In: Adv. Neural Inform. Process. Syst. (2025)
20. Lu, Y., Tian, J., Su, Y., Liu, J., Hao, C., Yue, C., Yang, X.: Polarization-aware low-light image enhancement for nighttime intelligent vehicles. *IEEE Transactions on Intelligent Vehicles* **10**(3), 1944–1958 (2025)
21. Luo, Y., Zhang, J., Li, C.: CPIFuse: Toward realistic color and enhanced textures in color polarization image fusion. *Information Fusion* **120**, 103111 (2025)
22. Luo, Y., Zhang, J., Shao, J., Tian, J., Ma, J.: Learning a non-locally regularized convolutional sparse representation for joint chromatic and polarimetric demosaicking. *IEEE Trans. Image Process.* **33**, 5029–5044 (2024)
23. Lyu, Y., Cui, Z., Li, S., Pollefeys, M., Shi, B.: Reflection separation using a pair of unpolarized and polarized images. In: Adv. Neural Inform. Process. Syst. vol. 32 (2019)
24. Lyu, Y., Guo, H., Zhang, K., Li, S., Shi, B.: SfPUEL: Shape from polarization under unknown environment light. In: Adv. Neural Inform. Process. Syst. vol. 37, pp. 97184–97202 (2024)
25. Mei, H., Dong, B., Dong, W., Yang, J., Baek, S.H., Heide, F., Peers, P., Wei, X., Yang, X.: Glass segmentation using intensity and spectral polarization cues. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 12622–12631 (2022)
26. Morimatsu, M., Monno, Y., Tanaka, M., Okutomi, M.: Monochrome and color polarization demosaicking based on intensity-guided residual interpolation. *IEEE Sensors Journal* **21**(23), 26985–26996 (2021)
27. Parazzoli, C., Gregor, R., Li, K., Koltenbah, B., Tanielian, M.: Experimental verification and simulation of negative index of refraction using snell’s law. *Physical review letters* **90**(10), 107401 (2003)
28. Pezzaniti, J.L., Chenault, D.B.: A division of aperture MWIR imaging polarimeter. In: *Polarization Science and Remote Sensing II*. vol. 5888, pp. 239–250. SPIE (2005)
29. Pust, N.J., Shaw, J.A.: Digital all-sky polarization imaging of partly cloudy skies. *Applied optics* **47**(34), H190–H198 (2008)
30. Qiu, S., Fu, Q., Wang, C., Heidrich, W.: Linear polarization demosaicking for monochrome and colour polarization focal plane arrays. In: *Comput. Graph. Forum*. vol. 40, pp. 77–89 (2021)
31. Rahman, M.D.A.B.A., Monno, Y., Tanaka, M., Okutomi, M.: Polarization denoising and demosaicking: Dataset and baseline method. In: IEEE Int. Conf. Image Process. pp. 2724–2729 (2025)
32. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10684–10695 (2022)
33. Sun, R., Chen, L., Liao, T., Fan, Z.: Self-supervised polarization image dehazing method via frequency domain generative adversarial networks. *Pattern Recognition* **165**, 111615 (2025)
34. Terrier, P., Devlaminck, V., Charbois, J.M.: Segmentation of rough surfaces using a polarization imaging system. *Journal of the Optical Society of America A* **25**(2), 423–430 (2008)
35. Wen, S., Zheng, Y., Lu, F.: A sparse representation based joint demosaicing method for single-chip polarized color sensor. *IEEE Trans. Image Process.* **30**, 4171–4182 (2021)
36. Wolff, L.B., Boulton, T.E.: Constraining object features using a polarization reflectance model. *IEEE Trans. Pattern Anal. Mach. Intell.* **13**(7), 635–657 (1991)

37. Xiang, K., Yang, K., Wang, K.: Polarization-driven semantic segmentation via efficient attention-bridged fusion. *Optics Express* **29**(4), 4802–4820 (2021)
38. Yang, Z., Hao, C., Lu, Y., Su, Y., Zhang, Y., Zhang, J., Tian, J.: A novel dehazing approach: Recovery of color and polarization information using polarized characteristics. *IEEE Trans. Multimedia* **27**, 8679–8693 (2025)
39. Yao, M., Wang, M., Tam, K.M., Li, L., Xue, T., Gu, J.: Polarfree: Polarization-based reflection-free imaging. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 10890–10899 (2025)
40. Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T.: One-step diffusion with distribution matching distillation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 6613–6623 (2024)
41. Zhang, J., Luo, H., Hui, B., Chang, Z.: Image interpolation for division of focal plane polarimeters with intensity correlation. *Optics Express* **24**(18), 20799–20807 (2016)
42. Zhang, K., Lyu, Y., Guo, H., Li, S., Ma, Z., Shi, B.: Polaranything: Diffusion-based polarimetric image synthesis. In: *Int. Conf. Comput. Vis.* pp. 26466–26476 (2025)
43. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Int. Conf. Comput. Vis.* pp. 3836–3847 (2023)
44. Zhou, C., Teng, M., Lyu, Y., Li, S., Xu, C., Shi, B.: Polarization-aware low-light image enhancement. In: *AAAI*. vol. 37, pp. 3742–3750 (2023)
45. Zhou, S., Zhou, C., Lyu, Y., Guo, H., Ma, Z., Shi, B., Sato, I.: PIDSR: Complementary polarized image demosaicing and super-resolution. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 16081–16090 (2025)

A Details on DMD




Fig. S1: Overview of distribution matching distillation in our training stage II.

Fig. S1 illustrates the Stage-II DMD mechanism. The only component not detailed in the main paper is the online update of the fake score model. Starting from the one-step output

$$\hat{\tilde{z}}_0 = \mathcal{G}_\theta(\tilde{z}_{t^*}, c, t^*), \quad (\text{S1})$$

we sample a random diffusion step t and re-noise it as

$$\hat{\tilde{z}}_t = \alpha_t \hat{\tilde{z}}_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, I). \quad (\text{S2})$$

The frozen Stage-I diffusion model serves as the real score model ($\mu_{\text{real}} = \mu_\theta$), while an auxiliary multi-step diffusion model μ_{fake} is maintained online to estimate the score of the evolving generated distribution. Following the DMD formulation, the two scores are

$$s_{\text{real}}(\hat{\tilde{z}}_t, t, c) = -\frac{\hat{\tilde{z}}_t - \alpha_t \mu_{\text{real}}(\hat{\tilde{z}}_t, t, c)}{\sigma_t^2}, \quad s_{\text{fake}}(\hat{\tilde{z}}_t, t, c) = -\frac{\hat{\tilde{z}}_t - \alpha_t \mu_{\text{fake}}(\hat{\tilde{z}}_t, t, c)}{\sigma_t^2}. \quad (\text{S3})$$

To keep μ_{fake} synchronized with the current generator distribution, we update it online using the generated clean latent $\hat{\tilde{z}}_0$ as the denoising target:

$$\mathcal{L}_{\text{fake}} = \mathbb{E}_{t, \epsilon} [\|\mu_{\text{fake}}(\hat{\tilde{z}}_t, t, c) - \hat{\tilde{z}}_0\|_2^2]. \quad (\text{S4})$$

Unlike the main one-step generator, μ_{fake} is updated only in the latent domain and does not reuse the image-domain physics losses. This simplification avoids an additional decoding path for the auxiliary model and thus reduces the memory and computational overhead of Stage II training.

The gradient flow in Stage II is separated as follows:

- the DMD score-difference gradient updates only the one-step generator $\mathcal{G}_\theta$;
- the auxiliary latent diffusion loss $\mathcal{L}_{\text{fake}}$ updates only μ_{fake} ;
- the physics loss $\mathcal{L}_{\text{phys}}$ updates both $\mathcal{G}_\theta$ and the LoRA parameters ϕ in the encoder;
- the posterior regularization term $\mathcal{L}_{\text{kl}}$ updates only ϕ ;
- the real score model μ_{real} and the VAE decoder remain fixed throughout Stage II.

B Additional analysis on LoRA

To further analyze the role of LoRA, we compare three VAE adaptation strategies on the RLP test set: (1) the original pretrained VAE without LoRA (NO-LoRA), (2) encoder LoRA pre-adapted only by plain autoencoding on the polarization training set (PRE-LoRA), and (3) our Stage-II end-to-end LoRA adaptation (OURS). For each setting, we examine two aspects: (i) the direct VAE reconstruction behavior obtained by encoding and decoding the polarization target, and (ii) the final DoLP/AoP predictions after diffusion inference. This comparison distinguishes improving autoencoding alone from improving the latent representation actually used by the diffusion model.

As shown in Fig. S2, the pretrained RGB VAE produces noticeable structured reconstruction errors on polarization targets, which propagate to the diffusion output and lead to relatively large DoLP/AoP errors. The PRE-LoRA setting reduces the direct VAE reconstruction artifacts to some extent, indicating that encoder-only adaptation can partially fit the polarization dataset in an autoencoding sense. However, this improvement does not translate to better DoLP/AoP after diffusion inference: the predicted cues remain similar to, or no better than, the w/o-LoRA case. This shows that plain VAE pre-adaptation does not resolve the mismatch between the latent distribution expected by the diffusion UNet and the latent representation induced by the adapted encoder.

In contrast, our Stage-II end-to-end LoRA adaptation reduces the direct VAE reconstruction error and consistently improves the final diffusion outputs, producing more accurate DoLP/AoP maps with fewer artifacts. This supports our main claim that LoRA is effective only when it is optimized jointly with the one-step generator under physics supervision, so that the adapted encoder remains aligned with the latent distribution used by the diffusion model.

A further observation is that the direct VAE reconstruction errors are much smaller for S_1 , S_2 , and DoLP than for angular quantities such as $\sin(2\text{AoP})$ and $\cos(2\text{AoP})$. This is expected because AoP is obtained through a nonlinear angular transform of (S_1, S_2) , and small Stokes-domain perturbations can induce disproportionately large changes in angle-related representations, especially under weak polarization. Therefore, the visible AoP degradation does not imply that the Stage-I multi-step diffusion is fundamentally unreliable; rather, it indicates that the remaining latent-domain reconstruction bias is sufficient to destabilize angular cues, which motivates the additional error reduction enabled by Stage-II end-to-end LoRA adaptation.




Fig. S2: Error maps and DoLP/AoP outputs after diffusion across the three pipeline related to LoRA.

Tab. S1 reveals a clear mismatch between direct VAE reconstruction quality and final diffusion performance. Although PRE-LoRA substantially reduces the direct reconstruction error of S_1 , S_2 , DoLP, and the AoP-related proxies, it does not improve the final diffusion output, yielding even lower DoLP PSNR and higher AoP MeanAE than w/o LoRA. In contrast, our end-to-end LoRA improves both the direct VAE reconstruction and the final polarization estimation, confirming that standalone VAE pre-adaptation does not align the latent representation with the diffusion UNet, whereas joint optimization in Stage II does.

C Inference efficiency analysis

We further analyze the inference efficiency of four methods: Restormer, MAE, PolarAnything, and our method. Table S2 reports both runtime and peak GPU memory usage under the same input resolution and hardware setting. The four methods differ not only in architecture, but also in prediction target and inference protocol.

Restormer. Restormer is a direct discriminative regressor that takes the 3-channel S_0 proxy as input and jointly outputs the 6-channel color Stokes components (S_1, S_2). Its single-pass design leads to both very fast inference and the lowest peak memory among the compared methods.

Table S1: LoRA adaptation comparison on the RLP test set. PRE-LoRA improves direct VAE reconstruction but does not reduce the final DoLP/AoP error after diffusion inference, whereas our end-to-end LoRA improves both.

Setting	VAE Recon. Mean Abs. Err. ↓					Diffusion Output	
	S_1	S_2	DoLP	$\sin(2\text{AoP})$	$\cos(2\text{AoP})$	$\text{PSNR}_{\text{DoLP}} \uparrow$	$\text{MeanAE}_{\text{AoP}} \downarrow$
w/o LoRA	0.020	0.019	0.081	0.476	0.503	28.49	23.12
PRE-LoRA	0.010	0.012	0.035	0.307	0.375	27.92	23.49
OURS	0.005	0.008	0.018	0.232	0.268	29.54	20.15

Table S2: Inference efficiency comparison. “Time” denotes the end-to-end inference time per sample, and “Peak Memory” denotes the maximum GPU memory recorded during inference. For MAE, S_1 and S_2 are obtained by two separate generative checkpoints; hence the reported runtime includes two sequential forward passes, while peak memory is measured per pass rather than summed across both passes.

Method	Output representation	Time (s) ↓	Peak Memory (MB) ↓
Restormer	color (S_1, S_2) (3→6 channels)	0.1081	3733
MAE	color (S_1, S_2), separate generative prediction of S_1 and S_2	2.3378×2	7400
PolarAnything	grayscale DoLP, $\cos(2\text{AoP})$, $\sin(2\text{AoP})$	2.2623	7330
Ours	color (S_1, S_2) in one step	0.1181	7434

MAE. For efficiency analysis, we characterize MAE by its actual prediction protocol. In our implementation, MAE serves as a conditioning prior for the latent generative backbone, playing a role functionally similar to the guidance branch in ControlNet [43]. Specifically, instead of using a shallow CNN-based condition encoder, it uses an MAE-based encoder to extract structural features from S_0 and inject them into the generative model. Importantly, the final S_1 and S_2 are generated by two separate checkpoints rather than predicted jointly in a single pass. This separate-generation protocol is not an implementation-specific workaround, but follows the original generative formulation adopted in RGB-to-Polarization [19], where the polarization components are modeled as separate generation targets. Therefore, two sequential forward passes are naturally required to obtain the complete color Stokes output, while peak memory is measured per pass rather than summed across both passes.

PolarAnything. PolarAnything feeds S_0 into a ControlNet-style generative backbone and predicts single-channel DoLP, $\cos(2\text{AoP})$, and $\sin(2\text{AoP})$. As a result, it estimates polarization from color input, but does not estimate color polarization: the output representation itself is grayscale and discards wavelength-specific polarization information. Its inference time and memory are both dominated by the diffusion/ControlNet backbone.

Our GenPolar. Our method directly predicts color Stokes components (S_1, S_2) in a single step. Compared with MAE and PolarAnything, one-step distillation substantially reduces inference time while preserving a color-aware Stokes representation. Although our peak memory remains relatively high due to the Stable-Diffusion-based backbone and VAE decoder, the overall runtime is the lowest among the generative methods and remains close to the direct-regression Restormer baseline.

Overall, the comparison highlights two points. First, one-step distillation primarily improves inference speed rather than eliminating model-residency memory. Second, compared with existing generative baselines, our method achieves a more favorable trade-off: it is substantially faster than MAE and PolarAnything, while jointly predicting full color Stokes components instead of either generating S_1 and S_2 separately or collapsing polarization into grayscale cues.

D Special Fresnel cases and failure analysis

In the main paper, we explain why DoLP is naturally wavelength-dependent through the factor $\rho(\lambda, \theta, n(\lambda))$. A further question is why AoP is often observed to be much less wavelength-sensitive in ordinary cases, and under what conditions this approximation breaks down. Here we clarify these special cases and analyze several representative failures of GenPolar.

D.1 Why AoP is usually weakly wavelength-dependent

Recall the formulation in the main paper:

$$\begin{bmatrix} S_{1,\lambda} \\ S_{2,\lambda} \end{bmatrix} = \rho(\lambda, \theta, n(\lambda)) S_{0,\lambda} \begin{bmatrix} \cos 2\psi \\ \sin 2\psi \end{bmatrix}. \quad (\text{S5})$$

Under ordinary single-bounce dielectric reflection, $\rho(\lambda, \theta, n(\lambda))$ carries the wavelength dependence through the Fresnel amplitudes and refractive index, while the angle ψ is mainly determined by the plane of incidence, i.e., by local geometry and viewing direction. Therefore,

$$\text{DoLP}_\lambda = |\rho(\lambda, \theta, n(\lambda))|, \quad \text{AoP}_\lambda = \psi, \quad (\text{S6})$$

which implies that DoLP is spectrally variant, whereas AoP is often approximately wavelength-independent in common reflective scenes. In other words, wavelength mainly changes the *strength* of linear polarization, but not its dominant *orientation*.

D.2 Special cases: when AoP becomes wavelength-dependent

The above approximation can break down when the effective polarization direction is no longer governed by a single reflection plane. This is especially relevant for transparent media and sky regions.

Transparent objects. For transparent or semi-transparent materials, the outgoing signal is often a mixture of reflection and transmission/refraction components. In this case, the relevant direction inside the medium is governed by Snell’s law [27]:

$$n_1 \sin \theta_i = n_2(\lambda) \sin \theta_t(\lambda), \quad (\text{S7})$$

which gives

$$\theta_t(\lambda) = \arcsin\left(\frac{n_1}{n_2(\lambda)} \sin \theta_i\right). \quad (\text{S8})$$




Fig. S3: Special cases beyond the basic Fresnel interpretation. GenPolar remains effective on the transparent object, but the sky case is visibly more challenging.

Importantly, the external incident angle θ_i does not change with wavelength, but the refracted angle $\theta_t(\lambda)$ does because the refractive index is dispersive. As a result, the internal propagation path, the relative contribution of different interfaces, and the effective polarization plane can all vary across RGB channels. Consequently, the observed AoP may also exhibit noticeable wavelength dependence. This explains why transparent objects, such as the transparent phone case in our example (Fig. S3), can show channel-dependent AoP even though the simple single-reflection Fresnel picture would suggest $\text{AoP}_\lambda \approx \psi$.

Sky regions. For an ideal clear sky under the single-scattering Rayleigh model, the angle of polarization is perpendicular to the scattering plane and is theoretically independent of wavelength, while the degree of polarization varies spectrally [6]. However, real captured sky images are rarely ideal Rayleigh-only observations, and are often affected by aerosols, clouds, and multiple scattering [29]. They often contain wavelength-dependent mixtures of molecular scattering, aerosols, multiple scattering, refraction, and even surface-reflected environmental light, all of which can perturb the effective polarization direction. Multiple scattering is known to modify the simple single-scattering polarization pattern, especially where single-scattering polarization is weak [6]. Therefore, the sky example in our figure should not be interpreted as a contradiction to Rayleigh theory. Rather, it suggests that the captured sky region is not an ideal single-scattering case, and that additional effects in real scenes can cause the observed AoP to differ across channels.

Overall, these special cases suggest that the approximation $\text{AoP}_\lambda \approx \psi$ is accurate for many ordinary reflective surfaces, but becomes less reliable when dispersive transmission, multi-interface mixing, or multiple scattering significantly contributes to image formation. This also explains why our channel-wise formulation is still necessary even though AoP is often less spectral than DoLP.

D.3 Failure cases of GenPolar

We also present several failure cases of GenPolar in Fig. S4. In the red-boxed regions, the GT indicates that these pixels belong to relatively reliable high-




Fig. S4: Representative failure cases of GenPolar. The red boxes mark regions where the GT indicates relatively reliable high-polarization areas. However, GenPolar underestimates the DoLP magnitude in these regions, and the weakened predicted linear polarization further causes large AoP deviations.

polarization areas, where AoP should in principle be easier to estimate. Nevertheless, our predictions still deviate noticeably from the GT.

A consistent observation is that the generated DoLP in these regions is lower than the GT. This directly affects AoP estimation, because AoP is computed from the predicted Stokes components as

$$\widehat{\text{AoP}}_{\lambda} = \frac{1}{2} \text{atan2}(\hat{S}_{2,\lambda}, \hat{S}_{1,\lambda}), \quad (\text{S9})$$

and becomes highly sensitive to perturbations once the linear polarization magnitude is underestimated. Although our training mask is constructed from GT DoLP, this only determines *where supervision is applied*; it does not guarantee that the model prediction itself remains in a high-confidence regime. Therefore, when the model underestimates the local polarization strength, the predicted DoLP drops below the GT level, and the corresponding AoP error increases accordingly. These examples reveal a practical gap between *GT observability* and *predicted observability*. Our observability-aware loss prevents AoP from being enforced in clearly unreliable GT regions, but it cannot completely eliminate errors caused by insufficiently predicted Stokes magnitude in challenging cases. This is consistent with the main-paper analysis: once the predicted polarization strength becomes weak, AoP is intrinsically more unstable.

Qualitatively, we find that GenPolar still behaves reasonably well on many transparent objects, where local appearance cues remain informative for inferring the effective polarization transfer. In contrast, sky reconstruction is more challenging, likely because its polarization formation depends on large-scale atmospheric scattering and spectral mixing that are only weakly constrained by local RGB appearance.

E Additional visual results

Additional true-color results. The main paper mainly presents pseudo-color visualizations of DoLP and AoP for clearer structural comparison. Here we additionally provide true-color visualizations for the same three scenes in Fig. S5, together with two extra examples for each of the three test groups in Figs. S6, S7, and S8. Compared with pseudo-color rendering, true-color visualization better reflects the channel-wise spectral characteristics of polarization. Pseudo-color maps are useful for highlighting local structures and artifacts, but the visualization mapping may sometimes make a prediction look less consistent with the GT than it actually is. By contrast, true-color rendering more directly reflects whether the estimated polarization cues are spectrally consistent with the GT.

Additional downstream results. We further provide extra qualitative results for the two downstream tasks used in the main paper, namely material detection and polarization de-reflection.

Fig. S9 shows additional material-detection examples. A notable observation is that the predicted AoP is not always visually perfect, yet the final material masks can still be highly accurate. This suggests that, in our setting, material discrimination depends more strongly on whether the polarization *strength* is correctly captured than on whether the polarization *angle* is perfectly reconstructed. Intuitively, the distinction between dielectric and metallic regions is more directly reflected by how strongly the surface polarizes light, whereas AoP is more sensitive to local geometry and thus can tolerate moderate errors without necessarily changing the final category prediction. This explains why imperfect AoP estimation does not always harm the final detection result.

Fig. S10 shows additional de-reflection examples. Compared with material detection, this task is more sensitive to the overall quality of the estimated polarization cues. When DoLP and AoP are jointly reconstructed more coherently, the downstream method can better separate reflection from transmission, leading to cleaner reflection suppression and better preservation of fine image details. The additional examples further support the same conclusion as in the main paper: improved polarization estimation translates to more reliable downstream restoration.




Fig. S5: Additional true-color visualizations for the same three scenes shown in the main paper.




Fig. S6: Additional qualitative results on the RLP test group. Two extra examples are shown, each with both pseudo-color and true-color visualizations.




Fig. S7: Additional qualitative results on the DoFP test group. Two extra examples are shown, each with both pseudo-color and true-color visualizations.




Fig. S8: Additional qualitative results on the RSP test group. Two extra examples are shown, each with both pseudo-color and true-color visualizations.



Fig. S9: Additional qualitative results for material detection. Although the predicted AoP is not always visually perfect, the final material masks remain highly accurate, indicating that this task is more strongly driven by polarization magnitude cues.




Fig. S10: Additional qualitative results for polarization de-reflection. More accurate polarization estimation leads to cleaner reflection removal and better preservation of image details.