

PolarAPP: Beyond Polarization Demosaicking for Polarimetric Applications

Yidong Luo^{1,2*}, Chenggong Li^{3*}, Yunfeng Song^{1,2}, Ping Wang²,
Boxin Shi⁴, 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 School of Computer Science, Peking University, Beijing, China

Abstract. Polarimetric imaging enables advanced vision applications such as normal estimation and de-reflection by capturing unique surface-material interactions. However, existing applications (alternatively called downstream tasks) rely on datasets constructed by naïvely regrouping raw measurements from division-of-focal-plane sensors—where pixels of the same polarization angle are extracted and aligned into sparse images without proper demosaicking. This reconstruction strategy results in suboptimal, incomplete targets that limit downstream performance. Moreover, current demosaicking methods are task-agnostic, optimizing only for photometric fidelity rather than utility in downstream tasks. Towards this end, we propose PolarAPP, the first framework to jointly optimize demosaicking and its downstream tasks. PolarAPP introduces a feature alignment mechanism that semantically aligns the representations of demosaicking and downstream networks via meta-learning, guiding the reconstruction to be task-aware. It further employs an equivalent imaging constraint for demosaicking training, enabling direct regression to physically meaningful outputs without relying on rearranged data. Finally, a task-refinement stage fine-tunes the task network using the stable demosaicking front-end to further enhance accuracy. Extensive experimental results demonstrate that PolarAPP outperforms existing methods in both demosaicking quality and downstream performance. Code is available here.

Keywords: Polarimetric vision · DoFP imaging · Task-aware learning

1 Introduction

Polarimetric imaging captures physical properties beyond RGB, such as surface normal and subsurface scattering, enabling advanced vision tasks in target detection [28], early cancer diagnosis [47] and other high-level tasks. However, off-the-shelf polarization sensors (e.g., Sony IMX250MZR) output mosaic images,

* Equal contribution. † Corresponding authors.




Fig. 1: (a) Comparison between two baselines and our **joint PolarAPP**, which directly performs polarization-based tasks from a single CPFA raw image. (b) The current SOTA pipeline, SfPUEL [32], suffers from artifacts in the reconstructed normal map due to noisy AoP estimation. (c) A naïve combination of the SOTA demosaicker PIDSr [56] with SfPUEL still fails to improve normal estimation. (d) In contrast, our PolarAPP yields sharper AoP and a more accurate normal map through customized joint learning with equivalent imaging transformation (EIT) and feature alignment.

where each pixel records only one polarization angle, necessitating demosaicking to reconstruct linear polarization images for downstream applications.

Despite the growing interest in polarization vision, existing algorithm designs do not take division-of-focal-plane (DoFP) imaging into consideration. Existing methods for Shape from Polarization (SfP) [1, 20, 32], De-reflection from Polarization (DfP) [19, 31, 50] and other polarization-based tasks [57] are trained on datasets constructed by naively rearranging DoFP measurements into four polarization channels. Concretely, a CPFA (color polarization filter array) sensor produces a single raw mosaic (e.g., 2048×2448) that interleaves an RGB Bayer pattern with a 2×2 DoFP micro-polarizer pattern. Prior pipelines typically (i) regroup pixels within each 2×2 DoFP block into four polarization views $\{0^\circ, 45^\circ, 90^\circ, 135^\circ\}$, yielding a half-resolution tensor of size $1024 \times 1224 \times 4$, and then (ii) apply Bayer demosaicking each view to obtain color polarization images of size $1024 \times 1224 \times 3 \times 4$. While this simplification provides convenient supervision, it bypasses polarization-aware demosaicking, leaving interpolation artifacts and polarization-angle misalignment uncorrected; moreover, the 2×2 regrouping inevitably reduces spatial resolution. As shown in Fig. 1(b), noisy AoP computed from the regrouped measurements can propagate to downstream predictions (e.g., inaccurate normals), motivating an explicit demosaicking stage that reconstructs physically consistent polarization cues for task learning.

Meanwhile, current deep learning-based demosaickers [14, 21, 36, 46, 56] are primarily designed for image reconstruction, optimizing for pixel-level or perceptual fidelity metrics such as LPIPS [53]. Although these objectives enhance visual quality, they overlook the semantic requirements of downstream tasks. Consequently, the demosaicked images may appear visually correct yet remain not task-ready. Moreover, when demosaicker and task models are trained separately, the demosaicker receives no task-level feedback and thus cannot adapt its reconstruction toward downstream utility. As shown in Fig. 1(c), the gap between fidelity-oriented demosaicking and SfP prevents further improvement in

task performance, underscoring the need for a joint optimization paradigm that makes polarization reconstruction both physically accurate and task-aware.

We argue that these performance gaps stem from two fundamental flaws: (i) performing polarimetric task without essential demosaicking, and (ii) treating demosaicking and downstream tasks as isolated processes. To this end, we propose **PolarAPP**, *the first framework to jointly optimize demosaicking and its downstream tasks*. PolarAPP incorporates a feature alignment mechanism optimized via meta-learning, which aligns the feature spaces of the demosaicking and task networks to produce task-adaptive demosaicking reconstructions and accurate task results. Then, inspired by equivalent imaging theory [3–5], PolarAPP adopts an equivalent imaging transformation (EIT) prior that explicitly enables the network to focus on learning the intrinsic reconstruction mapping, rather than overfitting to dataset-specific artifacts. After joint optimization, we further refine the downstream model with a dedicated fine-tuning stage, leveraging the stable demosaicking front-end to further improve downstream accuracy under high-quality inputs. PolarAPP enables end-to-end learning from mosaic input to task-specific outputs by propagating task-level supervision throughout the entire pipeline. As shown in Fig. 1(d), we validate the feasibility of jointly optimizing SfP and demosaicking, and further extend this approach to other tasks like DfP. Our contributions are summarized as follows:

- We introduce PolarAPP, the first framework that jointly optimizes reconstruction and polarimetric applications, bridging low-level recovery with downstream polarization-based analysis tasks.
- We design a feature alignment mechanism that semantically links demosaicking and downstream task features through a customized meta-learning-based pipeline, guiding the reconstruction with task-aware feedback.
- We employ an EIT prior that decouples demosaicking from low-quality data artifacts, enabling physics-consistent and faithful polarization reconstruction under the constrained task-specific datasets.
- Extensive experiments demonstrate that PolarAPP outperforms conventional two-stage methods on demosaicking, SfP and DfP, validating the effectiveness of our joint, task-aligned design.

2 Related Work

2.1 Polarization Image Demosaicking

Early polarization demosaicking methods rely on interpolation [24, 26, 27, 33, 49, 52] or optimization [9, 29, 30, 38, 48] with hand-crafted priors (e.g., edge preservation, chrominance consistency), generally suffering from artifacts in complex scenes. Recent deep learning approaches leverage CNNs [23, 36, 46], GANs [14], or Transformers [21, 56] to learn end-to-end mappings, achieving superior reconstruction quality. However, these methods typically treat demosaicking as an isolated task, and thus ignore its intrinsic coupling with downstream tasks (e.g. SfP and DfP), limiting the overall performance.

2.2 Polarimetric Applications in SfP and DfP

Polarimetric imaging supports a range of vision tasks such as SfP and DfP. SfP exploits the Fresnel reflection effect, where the degree and angle of polarization correlate with surface normals, enabling geometry recovery from a single polarized view. DfP utilizes the difference in polarization states between reflected and transmitted components, particularly near the Brewster angle [10, 34], to separate specular reflections from diffuse layers. Both tasks have evolved from physics-based analytical formulations [6, 8, 10, 17, 25, 34, 35, 41, 44, 45] to deep-learning approaches [1, 7, 13, 15, 16, 19, 20, 22, 31, 32, 42, 50], which achieve improved robustness and reconstruction quality. Despite this progress, most learning-based approaches rely on polarization data constructed by rearranging DoFP sensor pixels into separate channels, discarding high-frequency details and bypassing the necessary demosaicking stage. This unrealistic preprocessing causes models to overlook mosaic-induced artifacts, leading to degraded accuracy in normal estimation and reflection removal. These limitations motivate a unified framework that couples polarization demosaicking with downstream polarimetric tasks, forming the basis of our proposed PolarAPP.

2.3 Meta-Learning in Vision

Meta-learning develops algorithms to automatically fine-tune model parameters for specific tasks, demonstrating strong adaptability across domains. MAML [11, 12, 37] learns task-agnostic initialization parameters that enable fast adaptation with limited data. Other approaches [39, 43] focus on sample weighting, and they use a small validation set to identify informative samples under noise. Concurrently, some methods [2, 54] explore learning parametric loss functions to better align training objectives with downstream tasks. However, no existing work applies meta-learning to polarimetric imaging, where sensor-level degradation critically impacts performance of downstream tasks. This gap motivates our joint PolarAPP, which leverages meta-learning to bridge the gap between low-level reconstruction and downstream shape and reflection estimation.

3 Proposed Method

3.1 Overview

The overall framework of PolarAPP is illustrated in Fig. 2. During inference (Fig. 2(a)), the pipeline takes a single CPFA raw image $\mathbf{R}$ as input and outputs a task prediction $\mathbf{Y}$. We first apply Bayer demosaicking to obtain half-resolution color polarization images $\mathbf{I}^{\downarrow 2}$, which are fed into our color polarization demosaicker $\mathcal{D}$ to reconstruct full-resolution polarization images $\mathbf{I} = \{\mathbf{I}_{0^\circ}, \mathbf{I}_{45^\circ}, \mathbf{I}_{90^\circ}, \mathbf{I}_{135^\circ}\}$. From $\mathbf{I}$, we compute physically interpretable polarization features [55], including Stokes components $\mathbf{S}$ ($\mathbf{S}_0 = \frac{\mathbf{I}_{0^\circ} + \mathbf{I}_{45^\circ} + \mathbf{I}_{90^\circ} + \mathbf{I}_{135^\circ}}{2}$, $\mathbf{S}_1 = \mathbf{I}_{0^\circ} - \mathbf{I}_{90^\circ}$, $\mathbf{S}_2 = \mathbf{I}_{45^\circ} - \mathbf{I}_{135^\circ}$), DoLP = $\frac{\sqrt{\mathbf{S}_1^2 + \mathbf{S}_2^2}}{\mathbf{S}_0}$, AoP = $\frac{1}{2} \text{atan2}(\mathbf{S}_2, \mathbf{S}_1)$, along with




Fig. 2: The pipeline of PolarAPP. (a) Inference of PolarAPP. (b) Overview of the meta-learning-based iteration for joint training. (c) Inner update for updating demosaicking and task network under one gradient descent. (d) Outer update for updating the feature transform modules. (e) Joint demosaicking and task learning.

normalized spatial coordinates. These cues are concatenated channel-wise and fed into a task-specific network $\mathcal{T}$ to produce Y .

Although inference resembles a standard two-stage pipeline, the key observation behind PolarAPP is that demosaicking and downstream performance are coupled. The task network $\mathcal{T}$ is driven by physically derived polarization cues (DoLP/AoP), whose fidelity is directly determined by the demosaicked polarization images produced by $\mathcal{D}$. Consequently, a task-agnostic demosaicker optimized solely for reconstruction fidelity may preserve appearance yet distort subtle polarization structures that are critical to SfP/DfP, motivating joint learning to make $\mathcal{D}$ explicitly task-aware. However, naïve end-to-end joint training with demosaicking loss $\mathcal{L}_d$ and task loss $\mathcal{L}_t$ is prone to objective mismatch, where reconstruction fidelity and task objectives may induce conflicting gradients and the interaction between the two networks is mediated only through intermediate images. PolarAPP instead introduces feature alignment and meta-learns an alignment metric via a bi-level optimization, so that the alignment objective is updated only when it improves the end objectives measured by $\mathcal{L}_d$ and $\mathcal{L}_t$. Importantly, the alignment modules are used only during training and are discarded at inference, introducing zero additional inference overhead.

3.2 Loss Function

Before detailing the meta-learning training strategy, we define the loss functions used in the PolarAPP. These include standard reconstruction and task-specific losses, as well as a feature alignment loss that enables task-aware optimization.

Demosaicking loss. $\mathcal{D}$ is trained to reconstruct full-resolution polarization images, from which the S , DoLP, and AoP are derived. To ensure high-fidelity

reconstruction, we adopt a composite loss function $\mathcal{L}_d$, which jointly constrains $\mathbf{I}$, $\mathbf{S}$ and DoLP through ℓ_1 and gradient losses inspired by PIDSr [56]. AoP is supervised with ℓ_1 loss on wrapped angular differences to respect periodicity.

Task-Specific loss. $\mathcal{T}$ is trained with a task-specific loss $\mathcal{L}_t$, designed according to the downstream application. For SfP, we use cosine similarity and mean angular error losses to supervise surface normal estimation. For DfP, we follow the training protocol of PolarFree [50], incorporating pixel, VGG, total variation, phase, and diffusion prior losses.

Feature Alignment loss. To connect demosaicking and downstream task for mutual promotion, we introduce a feature alignment loss $\mathcal{L}_{fa}$, which operates through a meta-learned feature metric. Unlike $\mathcal{L}_d$ and $\mathcal{L}_t$, which provide direct supervision, $\mathcal{L}_{fa}$ encourages the demosaicking features to be task-ready and the task features to be reconstruction-aware. Specifically, in the meta-training loop, two feature transform modules (FT , Each FT consists of three CNN blocks) extract features $\mathbf{F}_{dj}$ and $\mathbf{F}_{tj}$ (where $j \in \{1, 2, 3\}$ indexes the feature hierarchy) from the output of $\mathcal{D}$ and $\mathcal{T}$ for calculating $\mathcal{L}_{fa}$. $\mathcal{L}_{fa}$ is defined as $\mathcal{L}_{fa} = 3 - \sum_{j=1}^3 \frac{\langle FT_1(\mathbf{F}_{dj}), FT_2(\mathbf{F}_{tj}) \rangle}{\|FT_1(\mathbf{F}_{dj})\|_2 \cdot \|FT_2(\mathbf{F}_{tj})\|_2}$, where the FT modules are meta-trained in the outer loop to make $\mathcal{L}_{fa}$ a meaningful signal for joint optimization. The total loss $\mathcal{L}$ is denoted as $\mathcal{L}_d + \lambda_t \mathcal{L}_t + \lambda_{fa} \mathcal{L}_{fa}$, where λ_t and λ_{fa} balance the trade-off.

3.3 Meta-Learning-based Feature Alignment

As discussed in Sec. 3.1, naïve joint training of $\mathcal{D}$ and $\mathcal{T}$ with $\mathcal{L}_d + \mathcal{L}_t$ can suffer from objective mismatch and limited cross-network communication. We therefore introduce two FT modules $\{FT_1, FT_2\}$ to align intermediate representations and enable feature-level information exchange in a latent metric space. However, learning a useful shared metric space is non-trivial: directly constraining the original feature spaces provides weak or misleading gradients due to representation discrepancy, and a fixed alignment loss can admit degenerate solutions (e.g., collapsed projections that minimize $\mathcal{L}_{fa}$ without improving $\mathcal{L}_d$ and $\mathcal{L}_t$). To address this, we meta-learn the alignment metric via a bi-level scheme (Fig. 2(c-d)): the inner update uses $\mathcal{L}_{fa}$ to obtain tentative parameters ($\mathcal{D}'$, $\mathcal{T}'$), and the outer update optimizes $\{FT_1, FT_2\}$ only through the resulting decreases in $\mathcal{L}_d$ and $\mathcal{L}_t$ on a meta-test split. This ties the alignment space to task and reconstruction objectives, yielding a goal-directed $\mathcal{L}_{fa}$ that stabilizes joint learning of $\mathcal{D}$ and $\mathcal{T}$.

Inner Update. Given the meta-train set $\{\mathbf{I}_{mtr}^{ref}, \mathbf{Y}_{mtr}^{gt}\}$, we form the input to $\mathcal{D}$ by applying the imaging operator $\mathcal{A}$ to the reference color polarization image $\mathbf{I}_{mtr}^{ref}$, where $\mathcal{A}$ denotes CPFA down-sampling and Bayer interpolation. Here, $\mathbf{I}_{mtr}^{ref}$ denotes the color polarization image obtained by regrouping and Bayer demosaicking into half resolution (relative to real-captured sensor resolution), which serves as the supervision target of $\mathcal{L}_d$ in the meta-learning loop. In the inner step, we use $\mathcal{L}_{fa}$ to perform a probe update that assesses the quality of the current alignment metric. Specifically, both $\mathcal{D}$ and $\mathcal{T}$ take one gradient descent




Fig. 3: Grad-CAM visual comparison of $\mathcal{L}_t$ and $\mathcal{L}_{fa}$ on SfP and DfP. It can be seen that $\mathcal{L}_{fa}$ yields more distributed, structure-aware responses for both $\mathcal{D}$ and $\mathcal{T}$.

step guided by $\mathcal{L}_{fa}$ computed from the meta-features $\mathbf{F}_{dj}^m$ and $\mathbf{F}_{tj}^m$:

$$\theta_{\mathcal{D}'} = \theta_{\mathcal{D}} - \beta_{\mathcal{D}'} \sum_{j=1}^3 \frac{\partial \mathcal{L}_{fa}(FT_1(\mathbf{F}_{dj}^m), FT_2(\mathbf{F}_{tj}^m))}{\partial \theta_{\mathcal{D}}}, \quad (1)$$

$$\theta_{\mathcal{T}'} = \theta_{\mathcal{T}} - \beta_{\mathcal{T}'} \sum_{j=1}^3 \frac{\partial \mathcal{L}_{fa}(FT_1(\mathbf{F}_{dj}^m), FT_2(\mathbf{F}_{tj}^m))}{\partial \theta_{\mathcal{T}}}, \quad (2)$$

where $\beta_{\mathcal{D}'}$ & $\beta_{\mathcal{T}'}$ are the inner-step learning rates and $\mathbf{I}_{mtr}^{out}$ the inner-step demosaicked output. The original parameters ($\theta_{\mathcal{D}}$, $\theta_{\mathcal{T}}$) of ($\mathcal{D}$, $\mathcal{T}$) remain unchanged, and the temporary ($\mathcal{D}'$, $\mathcal{T}'$) are created only to preserve the computational graphs required for optimizing the $\{FT_1, FT_2\}$ in the outer update.

Outer Update. Using the meta-test set $\{\mathbf{I}_{mts}^{ref}, \mathbf{Y}_{mts}^{gt}\}$, we update FT parameters θ_n ($n=1, 2$) under supervision from the ($\mathcal{L}_d$, $\mathcal{L}_t$) computed with ($\mathcal{D}'$, $\mathcal{T}'$):

$$\begin{aligned} \theta_n &\leftarrow \theta_n - \beta_n \frac{\partial \mathcal{L}_d(\mathcal{D}'(\mathcal{A}(\mathbf{I}_{mts}^{ref}); \theta_{\mathcal{D}'}), \mathbf{I}_{mts}^{ref})}{\partial \theta_n} - \beta_n \frac{\partial \mathcal{L}_t(\mathcal{T}'(\mathbf{I}_{mts}; \theta_{\mathcal{T}'}), \mathbf{Y}_{mts}^{gt})}{\partial \theta_n} \\ &= \theta_n + \beta_n \cdot \beta_{\mathcal{D}'} \frac{\partial \mathcal{L}_d(\mathcal{D}'(\mathcal{A}(\mathbf{I}_{mts}^{ref}); \theta_{\mathcal{D}'}), \mathbf{I}_{mts}^{ref})}{\partial \theta_{\mathcal{D}'}} \cdot \frac{\partial^2 \mathcal{L}_{fa}}{\partial \theta_n \partial \theta_{\mathcal{D}}} \\ &\quad + \beta_n \cdot \beta_{\mathcal{T}'} \frac{\partial \mathcal{L}_t(\mathcal{T}'(\mathbf{I}_{mts}; \theta_{\mathcal{T}'}), \mathbf{Y}_{mts}^{gt})}{\partial \theta_{\mathcal{T}'}} \cdot \frac{\partial^2 \mathcal{L}_{fa}}{\partial \theta_n \partial \theta_{\mathcal{T}}}, \end{aligned} \quad (3)$$

where β_n is the outer-step learning rate and $\mathbf{I}_{mts}^{out}$ is the outer-step demosaicked output. Backpropagating through Eq. (1)–(2) yields second-order gradients that adjust $\{FT_1, FT_2\}$ so that an $\mathcal{L}_{fa}$ -induced inner update leads to lower $\mathcal{L}_d$ and $\mathcal{L}_t$ on the meta-test split.

The bi-level optimization makes $\mathcal{L}_{fa}$ goal-directed: the alignment metric is updated only when it improves the end objectives measured by $\mathcal{L}_d$ and $\mathcal{L}_t$. This prevents degenerate alignment and yields an alignment signal that can be reused in the subsequent joint learning stage. In addition, we visualize gradient-based attribution for $\mathcal{D}$ and $\mathcal{T}$ under $\mathcal{L}_t$ and $\mathcal{L}_{fa}$ in Fig. 3. For SfP, $\mathcal{L}_{fa}$ emphasizes low-intensity regions in $\mathbf{S}_0$, complementing task supervision. For DfP, $\mathcal{L}_{fa}$ highlights object in strong reflection regions on $\mathcal{D}$ compared to $\mathcal{L}_t$; correspondingly, $\mathcal{L}_t$ drives $\mathcal{T}$'s gradients on dominant reflection regions, whereas $\mathcal{L}_{fa}$ distributes them to broader structural areas, encouraging structure-preserving predictions. Overall, the visualization indicates that the meta-learned alignment provides complementary training-time regularization with zero inference overhead.

3.4 Joint Learning

Using the meta-learned alignment metric, we perform joint learning of the $\mathcal{D}$ and $\mathcal{T}$, as illustrated in Fig. 2(e). Concretely, we adopt an epoch-wise alternat-




Fig. 4: Learning demosaicking based on equivalent imaging transformation prior.

ing schedule: within each epoch, we first update $\{FT_1, FT_2\}$ via the bi-level inner/outer steps, and then freeze $\{FT_1, FT_2\}$ and update $\mathcal{D}$ and $\mathcal{T}$ using standard end-to-end backpropagation. This design ensures that meta-learning is responsible only for learning the alignment metric, while joint learning performs the actual task-aware optimization of $\mathcal{D}$ and $\mathcal{T}$.

During joint learning, we compute $\mathcal{L}_{fa}$ on the current feature hierarchies $\{\mathbf{F}_{dj}\}$ and $\{\mathbf{F}_{tj}\}$ extracted from $\mathcal{D}$ and $\mathcal{T}$ (with $j \in \{1, 2, 3\}$), and backpropagate it jointly with $\mathcal{L}_d$ and $\mathcal{L}_t$. The update of $\mathcal{D}$ is:

$$\theta_{\mathcal{D}} \leftarrow \theta_{\mathcal{D}} - \beta_{\mathcal{D}} \left(\frac{\partial \mathcal{L}_d(\mathbf{I}, \mathbf{I}^{ref}) + \lambda_t \partial \mathcal{L}_t(\mathbf{Y}, \mathbf{Y}^{gt}) + \lambda_{fa} \sum_{j=1}^3 \partial \mathcal{L}_{fa}(FT_1(\mathbf{F}_{dj}), FT_2(\mathbf{F}_{tj}))}{\partial \theta_{\mathcal{D}}} \right), \quad (4)$$

and the $\mathcal{T}$ is updated as:

$$\theta_{\mathcal{T}} \leftarrow \theta_{\mathcal{T}} - \beta_{\mathcal{T}} \left(\frac{\lambda_t \partial \mathcal{L}_t(\mathbf{Y}, \mathbf{Y}^{gt}) + \lambda_{fa} \sum_{j=1}^3 \partial \mathcal{L}_{fa}(FT_1(\mathbf{F}_{dj}), FT_2(\mathbf{F}_{tj}))}{\partial \theta_{\mathcal{T}}} \right). \quad (5)$$

Here, gradients of $\mathcal{L}_{fa}$ flow through the fixed transforms $\{FT_1, FT_2\}$ into both networks, enabling feature-level communication while keeping inference unchanged (the transforms are used only for training). Overall, this stage uses the meta-learned alignment objective to bias demosaicking toward task-relevant structures and to provide the task network with reconstruction-aware context, so that improvements in $\mathcal{D}$ translate more directly into downstream performance.

3.5 Learning Demosaicking via EIT prior

Standard demosaicking datasets lack corresponding task GT, while polarization-task datasets (e.g., Sfp or Dfp) contain four polarization angles but often suffer from low image quality due to extraction from single-frame CPFA captures. Although we introduce $\mathcal{A}$ to construct paired data on task-specific datasets for training, its inherent low quality prevent reconstructions from high performance when $\mathcal{D}$ is only supervised by $\mathcal{L}_d$. To overcome this limitation, we first introduce the EIT prior for demosaicking, a self-supervised constraint that incorporates the imaging process. Since imaging operator is fully known and deterministic, EIT enables learning imaging-aware representation from limited observation by




Fig. 5: Overview of the task-refinement stage for SFP and DfP. $\mathcal{D}$ remains fixed while $\mathcal{T}$ is optimized with clean polarization inputs to enhance task accuracy.

leveraging transformation consistency, thereby overcoming the performance bottleneck of supervised learning on constrained task-specific datasets.

The training of demosaicking is shown in Fig. 4, where $\mathcal{D}$ is a Restormer-inspired UNet [40, 51], and $\mathcal{G}$ denotes a set of geometric transformations T_g . In practice, T_g includes one random non-zero integer translation, one 10° rotation, and one horizontal flip, all applied to the reconstructed full-resolution polarization image before re-sampling by $\mathcal{A}$. The equivariant demosaicking hypothesis asserts that, for any $T_g \in \mathcal{G}$:

$$\mathcal{D}(\mathcal{A}(T_g \mathbf{I}^{ref})) = T_g \mathcal{D}(\mathcal{A}(\mathbf{I}^{ref})). \quad (6)$$

This means that applying a geometric transformation in the reference polarization-image domain and then demosaicking should be consistent with demosaicking first and then transforming the reconstruction. This equivariance constraint provides a self-consistency condition that can be enforced without additional labeled data. Specifically, given a reference polarization image $\mathbf{I}^{ref}$, we first obtain its reconstruction:

$$\mathbf{I}^{rec} = \mathcal{D}(\mathcal{A}(\mathbf{I}^{ref})). \quad (7)$$

We then apply T_g in the reconstructed polarization-image domain and re-sample it using the fixed imaging operator $\mathcal{A}$ to synthesize a valid CPFA-consistent observation. The EIT loss is defined as:

$$\mathcal{L}_{eit} = \mathbb{E}_{\mathbf{I}^{ref}, g} \left[\mathcal{L}_d \left(\mathcal{D}(\mathcal{A}(T_g \mathbf{I}^{rec})), T_g \mathcal{D}(\mathcal{A}(\mathbf{I}^{ref})) \right) \right]. \quad (8)$$

In this way, transformations are applied before re-sampling by $\mathcal{A}$, so the CPFA or Bayer sampling layout remains valid and the generated inputs are still consistent with the known imaging process. $\mathcal{L}_{eit}$ regularizes $\mathcal{D}$ to produce consistent polarization reconstructions under different transformed observations. The demosaicking objective becomes $\mathcal{L}_d \leftarrow \mathcal{L}_d + \mathcal{L}_{eit}$.

Algorithm 1: PolarAPP Training Procedure

Input: Training set $\{(\mathbf{I}^{ref}, \mathbf{Y}^{gt})\}$
Output: Learned parameters $\theta_{\mathcal{D}}, \theta_{\mathcal{T}}, \theta_1, \theta_2$

```

1 Initialize  $\theta_{\mathcal{D}}, \theta_{\mathcal{T}}, \theta_1, \theta_2$ ;
2 while PolarAPP training do
3   // Meta-learning of the feature alignment metric;
4   while feature-alignment meta-learning do
5     Sample  $(\mathbf{I}_{mtr}^{ref}, \mathbf{Y}_{mtr}^{gt})$ ;
6     Obtain  $(\mathbf{F}_d^m, \mathbf{F}_t^m)$  and update  $\theta_{\mathcal{D}'}, \theta_{\mathcal{T}'}$  by Eq. (1)–Eq. (2);
7     Sample  $(\mathbf{I}_{mts}^{ref}, \mathbf{Y}_{mts}^{gt})$ ;
8     Obtain  $(\mathbf{I}_{mts}, \mathbf{Y}_{mts})$  and update  $\theta_1, \theta_2$  by Eq. (3);
9   // Joint learning of demosaicking and downstream task;
10  while joint demosaicking-task learning do
11    Sample  $(\mathbf{I}^{ref}, \mathbf{Y}^{gt})$ ;
12    Obtain  $(\mathbf{I}, \mathbf{Y}, \mathbf{F}_d, \mathbf{F}_t)$  and update  $\theta_{\mathcal{D}}, \theta_{\mathcal{T}}$  by Eq. (4)–Eq. (5);
13  // Task-network refinement with frozen demosaicker;
14  while task-network refinement do
15    Sample  $(\mathbf{I}^{ref}, \mathbf{Y}^{gt})$  without applying  $\mathcal{A}$ ;
16    Obtain  $(\mathbf{I}^{\uparrow 2}, \mathbf{Y}^{\uparrow 2}, \mathbf{F}_d^{\uparrow 2}, \mathbf{F}_t^{\uparrow 2})$ ;
17    Obtain  $(\mathbf{I}_{ref}^{\uparrow 2}, \mathbf{Y}_{gt}^{\uparrow 2})$  and update  $\theta_{\mathcal{T}}$  by Eq. (5);
  
```


3.6 Task Refinement and Training

Task Refinement: As mentioned earlier, joint training begins by generating low-quality inputs paired with reference images through $\mathcal{A}$. Although $\mathcal{D}$ performs well with the aid of the EIT prior, the compression introduced by $\mathcal{A}$ limits $\mathcal{D}$ ’s output quality during training, thereby hindering the improvement of $\mathcal{T}$. Therefore, we perform a refinement stage after the joint learning, where $\mathcal{A}$ is removed to fully unleash $\mathcal{D}$ ’s restoration capability. Notably, $\mathcal{D}$ is frozen in this stage, and $\mathcal{T}$ is optimized using $\mathcal{L}_t$ and $\mathcal{L}_{fa}$, as shown in Fig. 5. Removing $\mathcal{A}$ increases the task output resolution, thus requiring upsampling of the corresponding task GT for supervision. Specially, we use bilinear interpolation to upsample the SfP normals GT purely for scale alignment. The $2\times$ reflection-free supervision for DfP is obtained by applying the *frozen* $\mathcal{D}$ to generate full-resolution pseudo references (used only for supervision, without updating $\mathcal{D}$). This refinement brings two benefits: (i) $\mathcal{D}$ processes full-resolution inputs, yielding higher-quality demosaicking images for downstream training; and (ii) aligns training with inference conditions, reducing the resolution gap. Overall, it consolidates joint optimization and improves task accuracy under realistic inputs.

Training: The complete training procedure is summarized in Alg. 1. Each epoch alternates among (i) meta-learning (inner/outer updates to refine FT_1, FT_2), (ii) joint learning (updating $\mathcal{D}, \mathcal{T}$ using $\mathcal{L}_d, \mathcal{L}_t, \mathcal{L}_{fa}$), and (iii) task refinement. For the SfP task, $\mathcal{T}$ adopts an architecture similar to $\mathcal{D}$; for DfP, we integrate the PolarFree framework to provide a diffusion prior.

Table 1: Metrics comparisons on demosaicking.

Methods	DCPM	PIDSR	Ours
	S_0		
PSNR(dB)	42.98	43.07	43.12
SSIM	0.976	0.973	0.982
	DoLP		
PSNR(dB)	37.92	37.86	38.10
SSIM	0.918	0.925	0.928
	AoP		
MAE($^{\circ}$)	6.456	5.587	5.444

**Fig. 6:** Visual comparison in demosaicking.

4 Experiment

4.1 Experimental Setting

In training, the number of epochs L and the meta-learning iterations M are set to 100 and 200, respectively. The joint learning iterations N depends on the training set size. The batch size is fixed to 2, and all models in different training stages are optimized using Adam with a learning rate of 5×10^{-5} . The hyperparameters λ_t are set 20 for balanced scaling, while λ_{fa} is chosen based on ablation study. All experiments are conducted on servers equipped with NVIDIA RTX 4090 GPUs.

Dataset Preparation. PolarAPP is evaluated on polarization tasks including SfP and DfP. For SfP, we use the SfPUeL dataset [32], where 19,800 of 20,000 image groups are used for training and the remainder, together with several real-captured scenes, are used for testing. For DfP, 6,312 of 6,500 image groups from the PolarRGB dataset [50] are used for training and the rest for testing. In addition, the Qiu dataset [38] and several DoFP raw images are included as supplementary benchmarks for evaluating polarization demosaicking quality.

Comparison methods. For demosaicking evaluation, we compare PolarAPP with SOTA methods DCPM [21] and PIDSR [56], and they are re-trained on the same dataset as PolarAPP’s. The S_0 and DoLP are evaluated using PSNR and SSIM, while the AoP is measured by Mean Angular Error (MAE). For SfP comparison, SfPW [20] and SfPUeL [32] are applied either to demosaicked images or to direct inputs, forming 4×3 result combinations. Since SfP primarily concerns normal-direction accuracy rather than image-level appearance, the predicted normal maps are evaluated by angular-accuracy thresholds (11.25° , 22.5° , 30.0°) and angle error metrics (mean, median, RMSE). For DfP comparison, PolarFree [50] is paired with different upstream modules, producing another 4×2 set of results. Because DfP jointly requires effective reflection removal and detail preservation, the de-reflected outputs are assessed using PSNR, SSIM, LPIPS [53] and MUSIQ [18] metrics.

Table 2: Quantitative task comparisons. PIDS_R+ denotes retraining the corresponding task networks, SfPUEL for SfP and PolarFree for DfP, on outputs from a fixed PIDS_R demosaicker.

SfP	Accuracy $\uparrow$ (%)			Error $\downarrow$ ($^\circ$)			DfP	w Ref.		w/o Ref.	
	11.25 $^\circ$	22.5 $^\circ$	30 $^\circ$	Mean	Med.	RMSE		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MUSIQ $\uparrow$
SfPUEL	70.11	91.15	95.02	10.58	7.59	15.30	PolarFree	22.44	0.868	0.132	60.31
PIDS _R +	78.12	90.73	95.55	6.78	4.88	11.45	PIDS _R +	22.41	0.859	0.124	60.79
Ours	92.37	97.45	98.48	3.57	2.71	8.64	Ours	22.90	0.871	0.129	61.29




Fig. 7: Combination landscape in SfP (left) and DfP (right). Each column denotes combination of one demosaicker and one task method, serving as a practical reference for module compatibility. Complete PolarAPP shows the best metrics on both tasks.

4.2 Demosaicking Comparison

The quantitative results on the Qiu dataset are summarized in Tab. 1. PolarAPP delivers demosaicking performance that is competitive with SOTA methods while maintaining the stability required for downstream tasks—a balance that purely fidelity-driven demosaickers do not achieve. Visual comparisons on both Qiu and real-world raw data (Fig. 6) show that PolarAPP preserves edge structure and polarization patterns more reliably than DCPM and PIDS_R, which often produce blurred or unstable DoLP/AoP maps. In real scenes, PolarAPP reconstructs polarization cues that remain visually coherent and physically plausible, highlighting its stronger generalization under unconstrained conditions. Additional qualitative results are provided in the supplementary material (SM).

4.3 Task Comparison

Our task comparisons include two settings. First, in simulation, we retain the imaging operator $\mathcal{A}$ to ensure compatibility with task labels from mosaic rearranged datasets. Second, in real-world experiments, we remove $\mathcal{A}$ to match the inference setup with practical deployment conditions.

The quantitative results are shown in Tab. 2, with the broader combination landscape in Fig. 7. In Tab. 2, PIDS_R+ denotes the two-stage baseline where




Fig. 8: Visual results in SfP (top) and DfP (bottom). Top: Our complete PolarAPP avoids artifacts in synthetic image (left) and ensure the normal consistency in real scene (right). Bottom: Our complete PolarAPP better removes reflection while retaining necessary details in reflective area. (**Zooming in for better comparison**)

the corresponding task network is retrained on outputs from a fixed PIDSr demosaicker, i.e., SfPUEL for SfP and PolarFree for DfP. PolarAPP achieves the best performance on all SfP metrics and the main DfP metrics, showing that the gain is not merely from using a stronger demosaicker or adapting the task network to PIDSr outputs. Fig. 7 further evaluates different demosaicker-task pairings as practical two-stage references. The results show that both modules of PolarAPP are individually effective: $\mathcal{D}$ provides task-effective polarization cues, and $\mathcal{T}$ remains robust to different upstream inputs. However, the best performance is achieved when both modules are trained together, confirming the benefit of task-aware coupling.

Qualitative comparisons in Fig. 8 support these findings. In SfP scenes, alternative combinations produce distorted normals and inconsistent textures, while PolarAPP recovers correct orientations, sharper edges, and fewer shading artifacts. In DfP, PolarFree with other demosaickers suppresses reflections, and $\mathcal{T}$ paired with other demosaickers retains details, but only full PolarAPP removes reflections cleanly without losing fine structure. Real DoFP imaging results without $\mathcal{A}$ are shown in Fig. 9. For direct task baselines such as SfPUEL and PolarFree, their predictions are bilinearly upsampled only for resolution alignment; in the PIDSr-SfPUEL and the PIDSr-PolarFree baseline, $\mathcal{A}$ is removed for consistency with the real setting. PolarAPP closely matches the GT on both SfP and DfP, whereas the other pipelines exhibit artifacts and texture degradation. Additional results are provided in the SM.




Fig. 9: Visual comparisons follow real DoFP imaging (without $\mathcal{A}$).




Fig. 10: Ablation studies of λ_{fa} on SfP (left) and DfP (right).

Table 3: Quantitative results on ablation study.

Metrics	DoLP $\uparrow$		AoP $\downarrow$	Normal Accuracy (%) $\uparrow$			Normal Error ($^\circ$) $\downarrow$			De-reflection $\uparrow$	
	PSNR	SSIM	MAE	11.25 $^\circ$	22.5 $^\circ$	30 $^\circ$	MAE	MdAE	RMSE	PSNR	SSIM
w/o $\mathcal{D}$ in training	/	/	/	78.14	91.11	95.97	6.01	5.61	13.40	22.26	0.851
Task-agnostic $\mathcal{D}$	38.13	0.927	5.502	76.91	89.10	94.55	5.87	5.43	12.98	22.40	0.846
w/o $\mathcal{L}_{eit}$	36.75	0.885	8.011	86.12	92.43	95.75	5.73	4.96	12.57	22.46	0.869
Naïve Joint	38.43	0.924	5.658	82.44	94.62	97.43	6.09	4.43	11.86	22.37	0.858
w/o Refinement	38.05	0.920	5.789	89.01	93.55	96.33	4.68	4.50	10.13	22.59	0.869
Ours	38.10	0.928	5.444	92.37	97.45	98.48	3.57	2.71	8.64	22.90	0.871

4.4 Ablation study

To evaluate the contribution of each component in PolarAPP, we conduct ablation studies on demosaicking and downstream tasks (SfP and DfP) under a unified setting that retains the imaging operator $\mathcal{A}$ during joint learning. Results are summarized in Tab. 3.

***w/o* $\mathcal{D}$ in training.** We remove the demosaicker and train the task network directly on CPFA-processed inputs, which matches the task-only pipeline used by SfPUeL in our comparisons. This variant leads to significant performance drop on both SfP and DfP, indicating that a task network alone cannot reliably compensate for the lack of explicit polarization demosaicking.

Task-agnostic $\mathcal{D}$. We further include a two-stage baseline where $\mathcal{D}$ is trained alone with $\mathcal{L}_d$ (including $\mathcal{L}_{eit}$) and then frozen while $\mathcal{T}$ is trained with $\mathcal{L}_t$ only. Compared with PolarAPP, this task-agnostic pipeline remains inferior, indicating that a fidelity-oriented $\mathcal{D}$ can be suboptimal for downstream utility even when $\mathcal{T}$ is trained on its fixed outputs. This supports the need for task-aware coupling between demosaicking and downstream prediction.

***w/o* $\mathcal{L}_{eit}$.** Removing the equivariant imaging transformation loss $\mathcal{L}_{eit}$ significantly reduces DoLP/AoP reconstruction quality and in turn harms downstream accuracy, confirming that $\mathcal{L}_{eit}$ provides effective self-supervision for stable polarization demosaicking under limited paired data.

Naïve joint training (*w/o* meta-learned alignment). We ablate the proposed meta-learning-based feature alignment by disabling it and performing standard end-to-end joint optimization of $\mathcal{D}$ and $\mathcal{T}$ using only $\mathcal{L}_d + \mathcal{L}_t$. This naïve baseline is consistently inferior to PolarAPP, demonstrating the value of goal-directed feature alignment.

Effect of λ_{fa} . We sweep the alignment weight λ_{fa} in Fig. S2. Performance peaks at $\lambda_{fa} = 100$ for SfP and $\lambda_{fa} = 10$ for DfP (for radar visualization, angular-accuracy thresholds are converted to angular-error thresholds so that lower is better). The different optima reflect task-dependent sensitivity: SfP benefits more from strong global geometric consistency, whereas DfP is more sensitive to fine appearance details, and overly strong alignment can interfere with task-specific objectives. In both cases, too small λ_{fa} weakens coupling, while too large λ_{fa} over-regularizes optimization.

w/o refinement. Finally, the benefits of refinement are also reflected in improved metrics, indicating that this stage further unlocks the potential of the joint model. More discussion is provided in the SM, and we additionally evaluate a non-regression downstream task (material detection) in the SM.

5 Conclusion

In this work, we introduce PolarAPP, the first framework to jointly optimize polarization demosaicking and downstream tasks through task-specific joint learning. The proposed feature alignment mechanism, optimized through meta-learning, effectively bridges the objectives of the demosaicking and task networks and stabilizes their joint training. The demosaicking network is also guided by the equivalent imaging constraint, enabling effective regression learning without relying on low-quality datasets. Both innovations lead to superior performance in demosaicking and downstream tasks, including normal estimation and de-reflection. Experimental results demonstrate that PolarAPP outperforms existing methods, providing a flexible and robust solution for polarimetric imaging.

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. Ba, Y., Gilbert, A., Wang, F., Yang, J., Chen, R., Wang, Y., Yan, L., Shi, B., Kadambi, A.: Deep shape from polarization. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 554–571. Springer (2020)
2. Bai, H., Zhang, J., Zhao, Z., Wu, Y., Deng, L., Cui, Y., Feng, T., Xu, S.: Task-driven image fusion with learnable fusion loss. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7457–7468 (2025)

3. Chen, D., Davies, M., Ehrhardt, M.J., Schönlieb, C.B., Sherry, F., Tachella, J.: Imaging with equivariant deep learning: From unrolled network design to fully unsupervised learning. *IEEE Signal Processing Magazine* **40**(1), 134–147 (2023)
4. Chen, D., Tachella, J., Davies, M.E.: Equivariant imaging: Learning beyond the range space. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 4379–4388 (2021)
5. Chen, D., Tachella, J., Davies, M.E.: Robust equivariant imaging: a fully unsupervised framework for learning to image from noisy and partial measurements. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5647–5656 (2022)
6. Chen, G., He, L., Guan, Y., Zhang, H.: Perspective phase angle model for polarimetric 3d reconstruction. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 398–414. Springer (2022)
7. Dave, A., Zhao, Y., Veeraraghavan, A.: Pandora: Polarization-aided neural decomposition of radiance. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 538–556. Springer (2022)
8. Deschaintre, V., Lin, Y., Ghosh, A.: Deep polarization imaging for 3d shape and svbrdf acquisition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 15567–15576 (2021)
9. Dumoulin, R., Lapray, P.J., Thomas, J.B., Farup, I.: Impact of training data on lmmse demosaicing for colour-polarization filter array. In: *2022 16th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*. pp. 275–280. IEEE (2022)
10. Farid, H., Adelson, E.H.: Separating reflections and lighting using independent components analysis. In: *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*. vol. 1, pp. 262–267. IEEE (1999)
11. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: *International Conference on Machine Learning*. pp. 1126–1135. PMLR (2017)
12. Finn, C., Rajeswaran, A., Kakade, S., Levine, S.: Online meta-learning. In: *International Conference on Machine Learning*. pp. 1920–1930. PMLR (2019)
13. Fukao, Y., Kawahara, R., Nobuhara, S., Nishino, K.: Polarimetric normal stereo. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 682–690 (2021)
14. Guo, Y., Dai, X., Wang, S., Jin, G., Zhang, X.: Attention-based progressive discrimination generative adversarial networks for polarimetric image demosaicing. *IEEE Transactions on Computational Imaging* **10**, 713–725 (2024)
15. 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: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 11821–11830 (2024)
16. Hwang, I., Jeon, D.S., Munoz, A., Gutierrez, D., Tong, X., Kim, M.H.: Sparse ellipsometry: portable acquisition of polarimetric svbrdf and shape with unstructured flash photography. *ACM Transactions on Graphics (TOG)* **41**(4), 1–14 (2022)
17. Ichikawa, T., Purri, M., Kawahara, R., Nobuhara, S., Dana, K., Nishino, K.: Shape from sky: Polarimetric normal recovery under the sky. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 14832–14841 (2021)

18. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: Musiq: Multi-scale image quality transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5148–5157 (2021)
19. Lei, C., Huang, X., Zhang, M., Yan, Q., Sun, W., Chen, Q.: Polarized reflection removal with perfect alignment in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1750–1758 (2020)
20. Lei, C., Qi, C., Xie, J., Fan, N., Koltun, V., Chen, Q.: Shape from polarization for complex scenes in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12632–12641 (2022)
21. 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)
22. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21434–21445 (June 2024)
23. Li, J., Hao, J., Tong, G., Karim, S., Sun, X., Yu, Y.: Unsupervised demosaicking network using the recurrent renovation and the pixel-wise guidance. *Optics Letters* **47**(16), 4008–4011 (2022)
24. Liu, S., Chen, J., Xun, Y., Zhao, X., Chang, C.H.: A new polarization image demosaicking algorithm by exploiting inter-channel correlations with guided filtering. *IEEE Transactions on Image Processing* **29**, 7076–7089 (2020)
25. Logothetis, F., Mecca, R., Sgallari, F., Cipolla, R.: A differential approach to shape from polarisation: A level-set characterisation. *International Journal of Computer Vision* **127**(11), 1680–1693 (2019)
26. Lu, Y., Ren, W., Su, Y., Zhang, Z., Zhang, J., Tian, J.: Polarization image demosaicking based on homogeneity space. *Optics and Lasers in Engineering* **178**, 108179 (2024)
27. Lu, Y., Tian, J., Su, Y., Luo, Y., Zhang, J., Hao, C.: A hybrid polarization image demosaicking algorithm based on inter-channel correlation. *IEEE Transactions on Computational Imaging* **10**, 1400–1413 (2024)
28. Luo, Y., Zhang, J., Li, C.: Cpifuse: Toward realistic color and enhanced textures in color polarization image fusion. *Information Fusion* **120**, 103111 (2025)
29. 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 Transactions on Image Processing* **33**, 5029–5044 (2024)
30. Luo, Y., Zhang, J., Tian, D.: Sparse representation-based demosaicking method for joint chromatic and polarimetric imagery. *Optics and Lasers in Engineering* **164**, 107526 (2023)
31. Lyu, Y., Cui, Z., Li, S., Pollefeys, M., Shi, B.: Reflection separation using a pair of unpolarized and polarized images. *Advances in Neural Information Processing Systems* **32** (2019)
32. Lyu, Y., Guo, H., Zhang, K., Li, S., Shi, B.: Sfpuel: Shape from polarization under unknown environment light. In: *Advances in Neural Information Processing Systems*. vol. 37, pp. 97184–97202 (2024)
33. 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)
34. Nayar, S.K., Fang, X.S., Boulton, T.: Separation of reflection components using color and polarization. *International Journal of Computer Vision* **21**(3), 163–186 (1997)

35. Ngo, T.T., Nagahara, H., Taniguchi, R.i.: Surface normals and light directions from shading and polarization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(9), 5618–5630 (2021)
36. Nguyen, V., Tanaka, M., Monno, Y., Okutomi, M.: Two-step color-polarization demosaicking network. In: *IEEE International Conference on Image Processing (ICIP)*. pp. 1011–1015 (2022)
37. Qin, X., Quan, Y., Pang, T., Ji, H.: Ground-truth free meta-learning for deep compressive sampling. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9947–9956 (2023)
38. Qiu, S., Fu, Q., Wang, C., Heidrich, W.: Linear polarization demosaicking for monochrome and colour polarization focal plane arrays. In: *Computer Graphics Forum*. vol. 40, pp. 77–89. Wiley Online Library (2021)
39. Ren, M., Zeng, W., Yang, B., Urtasun, R.: Learning to reweight examples for robust deep learning. In: *International Conference on Machine Learning*. pp. 4334–4343. PMLR (2018)
40. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 234–241. Springer (2015)
41. Schechner, Y.Y., Shamir, J., Kiryati, N.: Polarization and statistical analysis of scenes containing a semireflector. *Journal of the Optical Society of America A* **17**(2), 276–284 (2000)
42. Shao, M., Xia, C., Yang, Z., Huang, J., Wang, X.: Transparent shape from a single view polarization image. In: *IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 9243–9252 (2023)
43. Shu, J., Xie, Q., Yi, L., Zhao, Q., Zhou, S., Xu, Z., Meng, D.: Meta-weight-net: Learning an explicit mapping for sample weighting. In: *Advances in Neural Information Processing Systems*. vol. 32 (2019)
44. Smith, W.A., Ramamoorthi, R., Tozza, S.: Linear depth estimation from an uncalibrated, monocular polarisation image. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 109–125. Springer (2016)
45. Smith, W.A., Ramamoorthi, R., Tozza, S.: Height-from-polarisation with unknown lighting or albedo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(12), 2875–2888 (2018)
46. Sun, Y., Zhang, J., Liang, R.: Color polarization demosaicking by convolutional neural network. *Optics Letters* **46**(17), 4338–4341 (2021)
47. Tuniyazi, A., Mu, T., Jiang, X., Lang, X., Li, Q., Yu, H., Han, F., Ban, J., Qin, B.: Snapshot polarized light scattering spectrometric fiberscopy for early cancer detection. *Laser & Photonics Reviews* **18**(12), 2400486 (2024)
48. Wen, S., Zheng, Y., Lu, F.: A sparse representation based joint demosaicing method for single-chip polarized color sensor. *IEEE Transactions on Image Processing* **30**, 4171–4182 (2021)
49. Xin, J., Li, Z., Wu, S., Wang, S.: Demosaicking dofp images using edge compensation method based on correlation. *Optics Express* **31**(9), 13536–13551 (2023)
50. Yao, M., Wang, M., Tam, K.M., Li, L., Xue, T., Gu, J.: Polarfree: Polarization-based reflection-free imaging. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2025)
51. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5728–5739 (2022)

52. Zhang, J., Chen, J., Yu, H., Yang, D., Liang, B., Xing, M.: Polarization image demosaicking via nonlocal sparse tensor factorization. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–10 (2021)
53. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 586–595 (2018)
54. Zhao, W., Xie, S., Zhao, F., He, Y., Lu, H.: Metafusion: Infrared and visible image fusion via meta-feature embedding from object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 13955–13965 (2023)
55. Zhao, Y., Qiao, X., Li, N., Pan, Q.: Polarization vision. *SCIENTIA SINICA Informationis* **54**(7), 1620–1645 (2024)
56. Zhou, S., Zhou, C., Lyu, Y., Guo, H., Ma, Z., Shi, B., Sato, I.: Pidsr: Complementary polarized image demosaicing and super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2025)
57. Zhu, Z., Li, X., Zhai, J., Hu, H.: Podb: A learning-based polarimetric object detection benchmark for road scenes in adverse weather conditions. *Information Fusion* p. 102385 (2024)

A Theoretical Analysis of FT Update

To better understand the learning behavior of the feature-alignment module, we analyze the optimization of FT_1, FT_2 . During the outer update, their parameters θ_n ($n \in 1, 2$) are optimized using losses computed on the meta-testing set, while the resulting gradients must account for how the demosaicking network $\mathcal{D}$ and task network $\mathcal{T}$ responded to the meta-training set in the inner update. This results in a bi-level gradient consisting of two terms, one from each network.

Gradient through Demosaicking Loss. Let $\mathcal{L}_d^{mts} = \mathcal{L}_d(\mathcal{D}'(\mathcal{A}(\mathbf{I}_{mts}^{ref})), \mathbf{I}_{mts}^{ref})$ denote the demosaicking loss computed using the demosaicker $\mathcal{D}'$. The FT parameters appear in this loss indirectly because: (i) θ_n affects $\mathcal{L}_{fa}$ in the inner update, (ii) $\mathcal{L}_{fa}$ affects the update of $\theta_{\mathcal{D}'}$ and (iii) $\theta_{\mathcal{D}'}$ affects the demosaicking loss on meta-testing. This produces the following gradient pathway:

$$\frac{\partial \mathcal{L}_d^{mts}}{\partial \theta_n} = \frac{\partial \mathcal{L}_d^{mts}}{\partial \theta_{\mathcal{D}'}} \cdot \frac{\partial \theta_{\mathcal{D}'}}{\partial \theta_n}. \quad (\text{S1})$$

Using inner-update rule $\theta_{\mathcal{D}'} = \theta_{\mathcal{D}} - \beta_{\mathcal{D}'} \frac{\partial \mathcal{L}_{fa}^{mtr}}{\partial \theta_{\mathcal{D}'}}$, we obtain:

$$\frac{\partial \theta_{\mathcal{D}'}}{\partial \theta_n} = -\beta_{\mathcal{D}'} \frac{\partial^2 \mathcal{L}_{fa}^{mtr}}{\partial \theta_n \partial \theta_{\mathcal{D}'}}. \quad (\text{S2})$$

Substituting into Eq. (S1) yields the demosaicking-related contribution:

$$-\beta_n \frac{\partial \mathcal{L}_d^{mts}}{\partial \theta_n} = \beta_n \beta_{\mathcal{D}'} \frac{\partial \mathcal{L}_d^{mts}}{\partial \theta_{\mathcal{D}'}} \cdot \frac{\partial^2 \mathcal{L}_{fa}^{mtr}}{\partial \theta_n \partial \theta_{\mathcal{D}'}}. \quad (\text{S3})$$

Gradient through Task Loss. Similarly, defining the task loss on the meta-testing set as $\mathcal{L}_t^{mts} = \mathcal{L}_t(\mathcal{T}'(\mathbf{I}_{mts}), \mathbf{Y}_{mts}^{gt})$. Since θ_n influences $\theta_{\mathcal{T}'}$ through the inner-update loss $\mathcal{L}_{fa}$, this gives:

$$\frac{\partial \mathcal{L}_t^{mts}}{\partial \theta_n} = \frac{\partial \mathcal{L}_t^{mts}}{\partial \theta_{\mathcal{T}'}} \cdot \frac{\partial \theta_{\mathcal{T}'}}{\partial \theta_n}. \quad (\text{S4})$$

With inner-update rule $\theta_{\mathcal{T}'} = \theta_{\mathcal{T}} - \beta_{\mathcal{T}'} \frac{\partial \mathcal{L}_{fa}^{mtr}}{\partial \theta_{\mathcal{T}'}}$, we obtain:

$$\frac{\partial \theta_{\mathcal{T}'}}{\partial \theta_n} = -\beta_{\mathcal{T}'} \frac{\partial^2 \mathcal{L}_{fa}^{mtr}}{\partial \theta_n \partial \theta_{\mathcal{T}'}} \quad (\text{S5})$$

and thus yielding the task-loss contribution:

$$-\beta_n \frac{\partial \mathcal{L}_t^{mts}}{\partial \theta_n} = \beta_n \beta_{\mathcal{T}'} \frac{\partial \mathcal{L}_t^{mts}}{\partial \theta_{\mathcal{T}'}} \cdot \frac{\partial^2 \mathcal{L}_{fa}^{mtr}}{\partial \theta_n \partial \theta_{\mathcal{T}'}}. \quad (\text{S6})$$

Summing both pathways yields exactly the update rule in Eq. (3) (main text). This decomposition makes the learning roles clear: (i) the meta-testing losses $\mathcal{L}_d$ and $\mathcal{L}_t$ determine the update direction for the FT modules, and (ii) the meta-training loss $\mathcal{L}_{fa}$ controls how strongly the FT modules influence the adaptations of $\mathcal{D}$ and $\mathcal{T}$ through second-order gradients. Together, these interactions enable the FT modules to learn alignment behaviors that improve joint optimization stability and downstream task performance.

Intuitively, the FT modules observe how changes in feature alignment, induced by $\mathcal{L}_{fa}$ on the meta-train set, modify the behavior of $\mathcal{D}$ and $\mathcal{T}$, and are then updated according to how well these adapted networks perform on the meta-test set. In this way, the FT modules gradually learn alignment strategies that make the demosaicking and task networks more compatible within the joint training framework.

B Additional Implementation Details

Resolution clarification of PolarAPP. Fig. S1 summarizes the resolution flow and data availability in PolarAPP. A DoFP CPFA sensor captures a raw mosaic at $2h \times 2w$, whereas the corresponding full-resolution color polarization images ($2h \times 2w \times 3 \times 4$) are unknown in current public datasets. Instead, available “reference” polarization images are typically obtained by DoFP regrouping followed by Bayer demosaicking, and are provided at a unified spatial resolution of $h \times w$ ($h \times w \times 3 \times 4$); the downstream task labels are also defined at $h \times w \times 3$ of both SfP and DfP.

To construct paired supervision from task datasets, we introduce a deterministic imaging degradation operator $\mathcal{A}$ to synthesize DoFP-consistent low quality inputs from the reference images, thereby forming supervised pairs for training $\mathcal{D}$. Built on this paired construction, we further impose the EIT prior to enforce transformation-consistent demosaicking, which helps the model learn the underlying demosaicking process rather than overfitting dataset-specific statistics. In the refinement and inference stages, we remove $\mathcal{A}$ and feed the reference images directly; notably, $\mathcal{D}$ is frozen during refinement to bridge the train–test resolution gap. Inference ultimately produces full-resolution task outputs.

Training time. Different downstream tasks are trained separately due to their distinct objectives and training strategies. All experiments are conducted on 4 NVIDIA RTX 4090 GPUs with batch size 2 per GPU (global batch size 8) and input resolution 256×256 . In each epoch, the meta-learning stage performs 200 fixed inner–outer cycles (~ 5 min), followed by joint learning and refinement over the full training set. One full training run requires roughly six/three days for SfP/DfP.

Efficiency. We compares PolarAPP with other methods in terms of parameter count, FLOPs and runtime, and all the methods are evaluated on an NVIDIA RTX 4090 GPU. All demosaickers take inputs of size 512×512 , while the task models operate on 1024×1024 inputs with a batch size of 1 throughout. For runtime evaluation, we perform 100 inference runs and report the average latency of the last 99 iterations to exclude the warm-up overhead of the first run. As shown in Tab. S1, our demosaicker $\mathcal{D}$ achieves lightweight design—significantly reducing parameters, FLOPs, and latency compared to existing methods. In the SfP task, our task model $\mathcal{T}$ incurs only marginally higher computational cost than the underperforming SfPW baseline, yet uses fewer parameters. Thanks to the meta-learning-based training framework, PolarAPP maintains remark-




Fig. S1: Resolution clarification of PolarAPP.

Table S1: Model efficiency comparison. (Our $\mathcal{T}$ and PolarFree share the same network architecture in DfP, resulting in identical computational complexity).

	Metrics	Params (M)	GFLoPs	Runtime (s)
Dem	DCPM	573.75	17936.02	0.5688
	PIDSR	7.45	230.92	0.1361
	Ours ($\mathcal{D}$)	4.40	117.83	0.0921
SfP	SfPW	42.48	187.33	0.1036
	SfPUeL	138.33	5303.67	2.5441
	Ours ($\mathcal{T}$)	4.55	352.86	0.2218
DfP	Ours ($\mathcal{T}$)	22.07	460.21	0.5077

ably low computational complexity while delivering SOTA performance across demosaicking, SfP, and DfP tasks.

C Additional Ablations Study

In Section 4.4, we analyze the contributions of individual modules in PolarAPP through controlled ablations. Although the refinement stage appears to have limited impact in simulation-based experiments, it provides clear benefits in real DoFP imaging. To highlight its practical value, we include an additional ablation study that directly follows the real imaging mechanism.

Qualitative comparisons with and without refinement are shown in the second column of Fig. S2. For SfP, both synthetic and real DoFP inputs reveal that removing refinement causes noticeable deviations from the upsampled ground truth. This reflects the strong sensitivity of normal estimation to resolution changes, since SfP depends on consistent angular cues across neighboring pixels. For DfP, the influence of refinement is reflected in the quality of reflection removal. Without refinement, PolarAPP suppresses specular reflections only partially, leaving visible color inconsistencies in reflective regions. These observations indicate that refinement plays an important role in restoring geometric accuracy for SfP and improving photometric consistency for DfP by compensating for resolution-related distortions and residual artifacts.

The remaining results in Fig. S2 illustrate the effect of removing other components under real DoFP imaging. Replacing the demosaicker with PIDSR or substituting the task networks with SfPUeL/PolarFree leads to inferior perfor-




Fig. S2: Visual results follow real DoFP imaging for ablation study.




Fig. S3: Visual results for ablation study of $\mathcal{L}_{eit}$.

mance. Similar degradation appears when feature alignment is removed or excessively strengthened, or when the task network is trained independently. Although removing the EIT prior affects demosaicking more strongly than the downstream tasks, it still degrades normal estimation and de-reflection. As shown in Fig. S3, retaining the EIT prior preserves more fine-scale details and provides cleaner polarization reconstructions.

D Additional Visual Comparisons

This section presents additional visual results referenced in the main experiments, including comparisons for demosaicking, SfP, and DfP. Fig. S4 presents the SfP and DfP results when the imaging operator $\mathcal{A}$, is retained. Fig. S5 shows additional demosaicking results on the Qiu dataset and real-captured images.




Fig. S4: Additional visual results in SfP and DfP.

Fig. S6 compares PolarAPP with two other baseline methods on a real-captured dataset that lacks normal GT references. It demonstrates that our clean, low-noise DoLP and AoP are critical factors for generating higher-quality normal maps. Fig. S7 displays the corresponding results under the standard DoFP imaging pipeline without using $\mathcal{A}$.

E Task generalization

Since SfP and DfP are both regression-based tasks, we further evaluate the generality of PolarAPP on a *non-regression* downstream task, namely material detection, which classifies dielectric and metallic materials from polarization cues. We use the dataset released with SfPUeL, where 19,800 out of 20,000 image groups are used for training and the remaining 200 for testing, and adopt




Fig. S5: Additional visual results in demosaicing.




Fig. S6: Visual results on Real-captured data without normal GT. Compared with DoLP&AoP of low-quality ones and PIDSr, our clean DoLP&AoP lead better normal estimation.

the same task network $\mathcal{T}$ as in SfP. For supervision, we use the standard cross-entropy loss for two-class prediction, and set $\lambda_{fa} = 100$.




Fig. S7: Additional visual results in SfP and DfP following the standard DoFP imaging.

For fair comparison, we evaluate five representative pipelines: (i) *SfPUEL*. The task-only baseline that inputs CPFA directly. (ii) *PIDSR+SfPUEL*. Task-agnostic demosaicking followed by a separately trained task network. (iii) *D+SfPUEL*. Our demosaicker combined with the external task network. (iv) *PIDSR+T*. External demosaicker combined with our task network. (v) Our complete PolarAPP. Notably, the material labels are discrete binary maps, so their $2\times$ supervision can be obtained by lossless upsampling without introducing interpolation ambiguity, allowing direct quantitative evaluation at the refined resolution. As shown in Fig. S8, the full PolarAPP achieves the highest material




Fig. S8: Visual comparisons on material detection. The detection accuracy of each pipeline is listed below the results, and PolarAPP still achieves the highest.

detection accuracy among all compared settings, indicating that the proposed joint structure generalizes beyond regression tasks and remains beneficial for discrete polarimetric perception.

F Discussion

PolarAPP provides a unified framework for polarization demosaicking and downstream tasks, but several limitations remain. First, although PolarAPP achieves consistent gains across all tasks, its ultimate performance is influenced by the strength of the underlying demosaicking and task architectures. PolarAPP improves their interaction through feature alignment and meta-learning, but further advances in backbone design or physically grounded architectures may unlock additional improvements. Second, publicly available polarimetric datasets are limited in quality and diversity; many rely on regrouped mosaics or provide incomplete ground truth, which restricts the attainable accuracy of both reconstruction and downstream predictions. Third, this work trains task-specific $\mathcal{D}$ - $\mathcal{T}$ pairs for SfP and DfP because their task losses and preferred polarization cues differ. A shared demosaicker with multiple task heads can be incorporated into the same framework by aggregating task losses, but balancing potentially conflicting task gradients is left for future work. Fourth, our current study isolates mosaic-induced degradation and does not explicitly model shot noise or readout noise in CPFA measurements. Extending PolarAPP to noisy raw observations is an important future direction.

Looking forward, the increasing availability of compact DoFP sensors and the growing interest in polarization-based vision suggest opportunities for more comprehensive datasets and physics-aware architectures. We expect future research

to explore a broader spectrum of polarimetric applications, including underwater imaging, dehazing, transparent and specular scene perception, low-visibility target sensing, remote sensing, biomedical imaging, and industrial inspection. These directions further highlight the importance of polarization cues for robust visual understanding in challenging real-world environments, as well as the need for unified and task-optimized polarimetric frameworks.