





PISR: Polarimetric Neural Implicit Surface Reconstruction for Textureless and Specular Objects

Guangcheng Chen^{1,2}, Yicheng He^{1,2}, Li He^{1,2}, and Hong Zhang^{*1,2}

¹ Southern University of Science and Technology, Shenzhen, China

² Shenzhen Key Laboratory of Robotics and Computer Vision

{chengc2023, 12332136}@mail.sustech.edu.cn, {he1, hzhang}@sustech.edu.cn

Abstract. Neural implicit surface reconstruction has achieved remarkable progress recently. Despite resorting to complex radiance modeling, state-of-the-art methods still struggle with textureless and specular surfaces. Different from RGB images, polarization images can provide direct constraints on the azimuth angles of the surface normals. In this paper, we present PISR, a novel method that utilizes a geometrically accurate polarimetric loss to refine shape independently of appearance. In addition, PISR smooths surface normals in image space to eliminate severe shape distortions and leverages the hash-grid-based neural signed distance function to accelerate the reconstruction. Experimental results demonstrate that PISR achieves higher accuracy and robustness, with an L1 Chamfer distance of 0.5 mm and an F-score of 99.5% at 1 mm, while converging $4 \sim 30\times$ faster than previous polarimetric surface reconstruction methods.

Keywords: Polarization · Surface reconstruction · Multi-view reconstruction · Neural implicit surface · Polarimetric constraint

1 Introduction

Surface reconstruction from multi-view images is a fundamental problem in computer vision, with various applications in graphics, robotics, and more. Recently, neural implicit surface has shown powerful flexibility for representing 3D shape and appearance. Combined with differentiable volume rendering [23], neural implicit surface reconstruction is able to estimate 3D shapes from multi-view images accurately [42, 47]. However, these methods struggle with textureless and specular objects due to the shape-radiance ambiguity [52]. Recent methods resort to diffuse and specular separation modeling [40], specular reflection detection [9] and inverse rendering techniques [18, 51]. While these methods show promising results, they may still encounter difficulties since their optimizations rely on image reconstruction losses that tightly couple shape and appearance.

* Corresponding author.

The source code is available at <https://github.com/GCChen97/PISR>

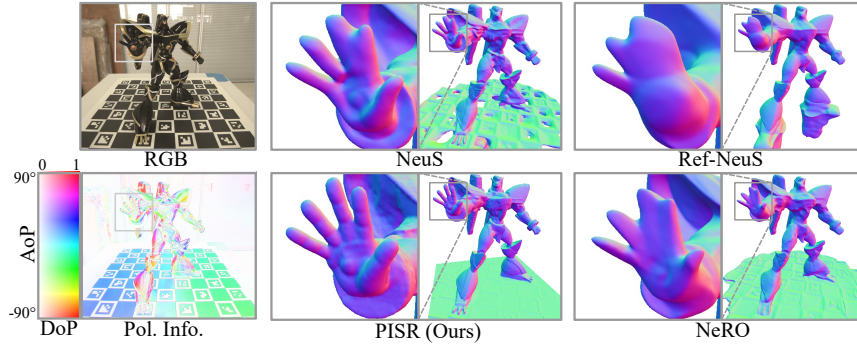


Fig. 1: Normal maps of reconstructed surfaces. By leveraging polarization information, our method is able to reconstruct textureless and specular surfaces with fine-grained details.

Polarization, as another property of light, provides direct cues on object shape. The angle of polarization of light from object surfaces is independent of light intensity and mainly determined by the surface shapes. From the perspective of a camera, the angle of polarization of a pixel is equal to the azimuth of the surface normal, i.e., the 2D projection direction of the normal onto the image plane, up to a $\pi/2$ -ambiguity and a π -ambiguity [5]. This geometric property has been extensively explored in shape from polarization [1, 2, 25, 49], photo-polarimetric stereo [7, 8, 27, 38], polarimetric multi-view 3D reconstruction [5, 54, 55] and dense mapping [34, 46].

In this work, we introduce PISR (polarimetric neural implicit surface reconstruction), a novel method that combines the merits of polarimetric 3D reconstruction and neural implicit surface reconstruction for textureless and specular objects, such as ceramics and plastics. Different from recent works [6, 16] that utilize the polarization cues through polarimetric volume rendering, we propose a novel polarimetric loss to refine an object shape directly. Thanks to the integration of the perspective polarimetric constraint [4] in the proposed loss, PISR eliminates artifacts on the reconstructed surfaces, resulting in a 30% reduction in L1 Chamfer distance. Besides, instead of treating all possible normals equally [54], our loss adaptively adjusts according to the degree of polarization to address specular-reflection-dominant surfaces, such as ceramics with dark color.

To accelerate the optimization process, PISR uses a multi-resolution hash grid and multi-layer perceptrons (MLP) to represent an object, achieving $4 \sim 30\times$ speed up as shown in Tab. 3. This representation is also crucial for accurately reconstructing highly reflective surfaces, enabling PISR to rectify the shape topology during optimization, a processing step not available for the mesh-based method, PMVIR [54]. However, the discrete nature of a hash grid also reduces the smoothness of the representation and can lead to the reconstruction

results being trapped in local minima. To mitigate the problem, PISR smooths rendered normals for better robustness against poor shape initialization.

Since there is no publicly available dataset for quantitative evaluation of polarimetric 3D reconstruction, we collect a dataset containing textureless and specular objects with a color polarization camera [36]. Using the dataset, we show that PISR is able to improve the surface reconstruction accuracy over the SOTA methods, with an L1 Chamfer distance of 0.5 mm and an F-score of 99.5% at 1 mm. Shown in Fig. 1 are recovered shapes by different methods and our result shows an accurate shape and fine-grained details.

In summary, our contributions are as follows:

- PISR, a novel method that leverages polarization cues and hash-grid-based neural SDF for robust and efficient reconstruction of textureless and specular objects.
- A novel polarimetric loss that utilizes an accurate formulation of the polarimetric constraint for higher surface reconstruction accuracy.
- A polarization image dataset of real textureless and specular objects with ground-truth meshes for quantitative evaluation.

2 Related Works

Neural Implicit Surface Reconstruction: The success of neural radiance field reconstruction [23] in novel view synthesis greatly inspires the development of neural implicit surface [28, 30, 48]. Combining with differentiable volume rendering, neural implicit surface reconstruction [29, 42, 47] is able to reconstruct implicit surfaces from dense multi-view images while showing excellent performance in reconstructing fine-grained details. Since then, several works have improved it in reconstruction quality [50], sparse-view performance [19] and optimization speed [43].

Yet textureless and specular surfaces are still challenging for neural implicit surface reconstruction. To address the problem, NeRO [18] incorporates inverse rendering to alleviate the shape-radiance ambiguity [52] induced by specular reflection. Ref-NeuS [9] borrows the idea of anomaly detection to decrease the weights of specular pixels during optimization and conditions the radiance on reflection direction in reflective scenes as in Ref-NeRF [39]. UniSDF [40] models diffuse and specular radiances in two MLPs separately.

Estimating a 3D radiance field and an SDF from 2D images is a highly ill-posed problem. During optimization, these methods may distort shapes to fit the reflection colors, causing severe reconstruction errors, or, as is commonly referred to, the shape-radiance ambiguity [52]. In contrast, polarization images provide linear constraints on surface normal by the angle of polarization. Therefore, to avoid the shape-radiance ambiguity, our method utilizes the geometry cues from polarization to regularize object shapes independently of appearance, keeping the shape from distortions induced by the visually distorted appearance.

Polarimetric 3D Reconstruction: Shape from polarization (SfP) [1,25,31,49] and photo-polarimetric stereo [7,8,27,38] are single-view model-based methods that utilize polarization images for normal and depth estimation. These methods rely on assumptions of either spotlighting or distant lighting to explicitly model the relationship between the degree of polarization and the zenith angle of a surface normal, making them difficult to work under natural lighting conditions. Several deep learning-based methods [2,21,35] leverage data priors to enhance the robustness to more complex indoor lighting conditions. SfPW [15] and DPS-Net [37] learn to estimate scene-level normal maps and depth maps. Yet their generalization is still limited by the relatively small amount of training datasets of polarization images.

Polarimetric multi-view 3D reconstruction leverages multi-view polarimetric constraints to adapt to natural lighting conditions. The polarimetric constraint is established on the fundamental fact that the azimuth angle of a surface normal is equal to the angle of polarization up to a $\pi/2$ -ambiguity and a π -ambiguity under natural lighting [5,54], which is relatively independent of lighting [5]. Miyazaki *et al.* [24] uses multi-view polarimetric constraints to reconstruct black specular objects. PolMVS [5] proposes to propagate the depths of sparse points and refine the depth maps with this constraint. Inspired by PolMVS, PDMS [46] and PMDM [34] integrate this constraint into SLAM for dense mapping. PMVIR [54] proposes a framework to simultaneously refine coarse meshes and estimate illuminations and surface albedos. PolarPMS [55] combines this constraint with photo consistency to estimate depth maps from multi-view polarization images.

More closely related to our approach is PMVIR [54]. PMVIR represents shapes using meshes and thereby has limitations in handling incorrect shape topology. In contrast, by leveraging neural SDF, our method is able to rectify incorrect shape topology during the optimization. More importantly, our method utilizes the perspective polarimetric constraint [4] to consider the camera ray directions in the polarimetric loss and thus achieves further performance improvement.

Polarimetric Inverse Rendering: The development of polarimetric bidirectional reflectance distribution function (pBRDF) [3,11,12,14] has enabled polarimetric inverse rendering. Based on mesh representation, sparse ellipsometry [11] leverages polarimetric inverse rendering to simultaneously estimate object meshes and materials using polarization images captured under flashlight illumination.

With the emergence of neural fields in 3D reconstruction, a few recent works [6,16] have integrated polarization into neural implicit surfaces through polarization image rendering. Concurrent work NeISF [16] introduces the incident Stokes field to represent multi-path polarized light and reparameterizes pBRDF using MLPs, achieving highly accurate geometry reconstruction. PANDORA [6] aims to decompose diffuse and specular reflections using polarization as auxiliary cues, relying on the Fresnel equations to calculate the diffuse and specular components of radiance.

Although polarimetric inverse rendering excels at accurate material and illumination estimation, it requires proper geometry initialization to decouple shape and appearance. As a result, it may still suffer from the shape-radiance ambiguity, especially when dealing with textureless and specular surfaces. To avoid this problem, our method utilizes the polarimetric constraint to regularize shape directly, independent of appearance. The surface reconstruction results of our method can be used to initialize polarimetric inverse rendering.

3 Preliminary

Polarimetric Constraint: Assuming four synchronized and aligned images captured at polarizer angles of 0, 45, 90, 135 degrees [36], denoted as \mathbf{I}_0 , \mathbf{I}_{45} , \mathbf{I}_{90} and \mathbf{I}_{135} . For each pixel, the angle of polarization (AoP) $\varphi \in [-\pi/2, \pi/2]$ and degree of polarization (DoP) $\rho \in [0, 1]$ are calculated as:

$$\varphi = \frac{1}{2} \arctan2(s_2, s_1), \quad \rho = \frac{\sqrt{s_1^2 + s_2^2}}{s_0} \quad (1)$$

where $\mathbf{s} = [s_0, s_1, s_2]^T = [\frac{1}{2} \sum I_i, I_0 - I_{90}, I_{45} - I_{135}]^T$ is the Stokes vector for representing the polarization state of light. Given the estimated normal $\hat{\mathbf{n}}$ in the camera frame and the AoP φ of a pixel, the orthographic polarimetric constraint on the azimuth angle of the normal is calculated as:

$$[\sin(\varphi + \Delta), \cos(\varphi + \Delta), 0] \cdot \hat{\mathbf{n}} = 0, \quad (2)$$

where the disambiguation scalar $\Delta \in \{0, \pi/2\}$ is used for solving the $\pi/2$ -ambiguity [5]. While this constraint is commonly used in previous works [5, 34, 38, 46, 54], there is a perspective polarimetric constraint which is geometrically more accurate than it by considering the perspective effect of the lens [4]:

$$[v_z \sin \varphi', v_z \cos \varphi', -(v_y \cos \varphi' + v_x \sin \varphi')] \cdot \hat{\mathbf{n}} = 0, \quad \varphi' = \varphi + \Delta \quad (3)$$

where $\mathbf{v} = [v_x, v_y, v_z]^T$ is a normalized camera ray in the camera frame.

$\pi/2$ ambiguity: The $\pi/2$ -ambiguity arises from the differences between the AoPs of specular reflection and diffuse reflection [5]. According to the Fresnel equations, reflected and transmitted light have distinct energy distributions in the parallel and perpendicular directions to the plane of incident, resulting in a $\pi/2$ difference between their AoPs, i.e., $\Delta = 0$ for diffuse reflection and $\Delta = \pi/2$ for specular reflection. Under natural lighting conditions, the diffuse reflections and specular reflections are usually mixed, making it difficult to distinguish the dominant component, and thus resulting in the $\pi/2$ -ambiguity. Therefore, the $\pi/2$ -ambiguity poses a challenge in terms of applying the polarimetric constraints.

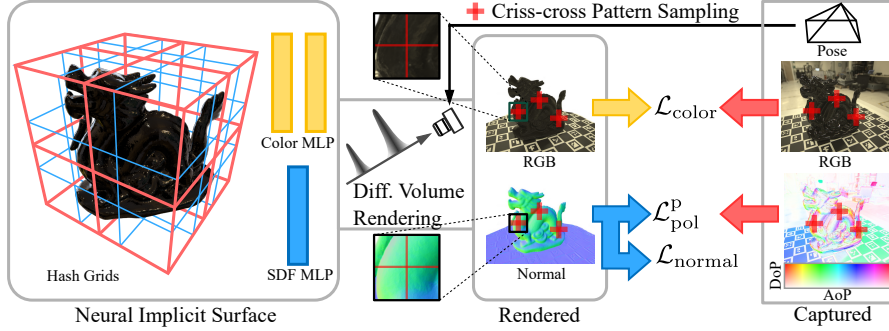


Fig. 2: Overview of PISR. Pixels for the optimization are sampled in criss-cross patterns at the middle stage of the optimization. The polarimetric loss $\mathcal{L}_{\text{pol}}^p$ and the normal loss $\mathcal{L}_{\text{normal}}$ are used for regularizing the shape independently of appearance.

4 Method

PISR aims to recover object shapes from multi-view polarization images and camera poses. Fig. 2 shows an overview of our method. Object shape is represented as a hash-grid-based neural SDF and rendered to images for computing losses (Sec. 4.1). During the reconstruction, the coarse shape of an object is estimated with the photometric loss $\mathcal{L}_{\text{color}}$. Then the polarimetric loss $\mathcal{L}_{\text{pol}}^p$ (Sec. 4.2) and the normal loss $\mathcal{L}_{\text{normal}}$ with criss-cross pattern sampling (Sec. 4.3) are jointly used to rectify the shape from distortions. Finally, the shape is refined with $\mathcal{L}_{\text{color}}$ and $\mathcal{L}_{\text{pol}}^p$. The optimization scheme is introduced in Sec. 4.4.

4.1 Surface Representation and Image Rendering

Since Instant-NGP [26], combining a multi-resolution hash grid and MLPs has been commonly used for faster convergence in neural field reconstruction. Following recent works [17, 41, 43], we represent the shape and appearance of an object as an L -level multi-resolution hash grid $\mathcal{G} = \{\mathcal{G}_l\}_{l=1}^L$ with two MLPs Φ_s and Φ_c for decoding SDF values s and color \mathbf{c} respectively. For each 3D space point \mathbf{x} , the SDF value s and color \mathbf{c} are computed as

$$(s, \mathbf{f}) = \Phi_s(\mathcal{G}(\mathbf{x})), \quad \mathbf{c} = \Phi_c(s, \mathbf{f}, \mathbf{v}, \mathbf{n}) \quad (4)$$

with the camera ray \mathbf{v} and surface normal \mathbf{n} . The normal \mathbf{n} can be computed by normalizing the gradient of SDF w.r.t. to \mathbf{x} or by finite difference [17]. Besides the foreground object, we use Instant-NGP [26] to represent the background.

To optimize the shape with images, we use differentiable volume rendering [22] to render images given the surface representation. The pixel color $\hat{\mathbf{c}}$ and normal $\hat{\mathbf{n}}$ are computed as the weighted sum of M samples $\{\mathbf{x}_i = \mathbf{o} + t_i \mathbf{v}\}_{i=1}^M$ on a cast-off pixel ray with the camera position \mathbf{o} and different depth t_i :

$$\hat{\mathbf{c}} = \sum_{i=1}^M w_i \mathbf{c}_i, \quad \hat{\mathbf{n}} = \sum_{i=1}^M w_i \mathbf{n}_i, \quad (5)$$

where w_i is the blending weight converted from SDF values using the unbiased weight function of NeuS [42], and \mathbf{c}_i and \mathbf{n}_i are the color and normal of the sample \mathbf{x}_i .

4.2 Polarimetric and Photometric Losses

Polarimetric Loss: Previous works [5, 34, 38, 46, 54] leverage polarization cues with the following loss function h^o that is derived from the orthographic polarimetric constraint Eq. (2):

$$h^o(\varphi, \Delta, \mathbf{n}) = \left(\mathbf{a}^o(\varphi, \Delta) \cdot \mathbf{n} \right)^2, \quad \mathbf{a}^o(\varphi, \Delta) = [\sin(\varphi + \Delta), \cos(\varphi + \Delta), 0], \quad (6)$$

where the coefficient vector \mathbf{a}^o is calculated with the AoP φ and the disambiguation scalar Δ . However, as discussed in [4], Eq. (2) neglects the perspective effect of the lens and may introduce distortions in the reconstruction results.

Therefore, we instead adopt the perspective polarimetric constraint Eq. (3) [4] and propose a novel loss function h^p :

$$h^p(\varphi, \Delta, \mathbf{v}, \mathbf{n}) = \left(\frac{\mathbf{a}^p(\varphi, \Delta, \mathbf{v}) \cdot \mathbf{n}}{\|\mathbf{a}^p(\varphi, \Delta, \mathbf{v})\|} \right)^2 \quad (7)$$

with \mathbf{a}^p calculated as:

$$\mathbf{a}^p(\varphi, \Delta, \mathbf{v}) = [v_z \sin \varphi', v_z \cos \varphi', -(v_y \cos \varphi' + v_x \sin \varphi')], \quad \varphi' = \varphi + \Delta. \quad (8)$$

The loss function h^p additionally considers the camera ray direction \mathbf{v} to model the perspective effect. In Sec. 5.3, we will show that h^p effectively improves the reconstruction accuracy.

Besides, the $\pi/2$ -ambiguity must be addressed when applying the polarimetric constraint. We combine the merits of PolMVS [5] and PMVIR [54], proposing the segmented loss function f^p as follows:

$$f^p(\varphi, \mathbf{v}, \mathbf{n}, \rho) = \begin{cases} h^p(\varphi, 0, \mathbf{v}, \mathbf{n}) \cdot h^p(\varphi, \pi/2, \mathbf{v}, \mathbf{n}) & \text{if } \rho < \theta, \\ h^p(\varphi, \pi/2, \mathbf{v}, \mathbf{n}) & \text{if } \rho \geq \theta, \end{cases} \quad (9)$$

where ρ is the DoP and θ is a preset threshold. This design is based on two observations: 1) the reflected light of an object is usually a mixture of the diffuse and specular components, and thus the DoP is relatively low [5], 2) only the DoP of specular reflection of dielectric objects can exceed 0.4 according to Fresnel equations. We use multiplication instead of addition when $\rho < \theta$ because only one of the two needs to be minimized.

Finally, we calculate the polarimetric loss for a batch of sampled pixels \mathcal{S} as follows:

$$\mathcal{L}_{\text{pol}}^p = \frac{1}{|\mathcal{S}|} \sum_{u \in \mathcal{S}} f^p(\varphi_u, \mathbf{v}_u, \hat{\mathbf{n}}_u, \rho_u). \quad (10)$$

Note that the camera ray \mathbf{v}_u and normal $\hat{\mathbf{n}}_u$ are transformed into the camera frame according to the camera pose. In Sec. 5.3, the experiments will demonstrate that calculating the loss $\mathcal{L}_{\text{pol}}^p$ with h^p instead of h^o can reduce the L1 Chamfer distance by 30%.

Photometric Loss: Since the polarimetric loss alone is insufficient for recovering object shapes due to the $\pi/2$ ambiguity and image noises, we used the photometric loss for initializing and maintaining the basic shape:

$$\mathcal{L}_{\text{color}} = \frac{1}{|\mathcal{S}|} \sum_{u \in \mathcal{S}} |\hat{\mathbf{c}}_u - \mathbf{c}_u|. \quad (11)$$

4.3 Normal Regularization

The discrete nature of the hash grid makes the estimated SDF less smooth than those purely MLP-based, which can introduce reconstruction artifacts such as holes or cracks. Directly smoothing the normals in 3D space as in [17, 53] can lead to a bumpy surface because the sampled normals are solely constrained by the smoothness loss. Inspired by the belief propagation used in traditional MVS methods [33, 45, 56], we smooth the rendered normals so that all the samples are also constrained by both the photometric loss and the polarimetric loss. We therefore utilize the following normal loss:

$$\mathcal{L}_{\text{normal}} = -\frac{1}{|\mathcal{S}_c|} \sum_{u \in \mathcal{S}_c} \frac{1}{|\mathcal{N}_u|} \sum_{\hat{\mathbf{n}}_j \in \mathcal{N}_u} \text{SG}(\hat{\mathbf{n}}_u)^T \hat{\mathbf{n}}_j \quad (12)$$

where \mathcal{S}_c is a set of pixels sampled among pixels of all the images, \mathcal{N}_u is a set of pixel normals sampled around $\hat{\mathbf{n}}_u$ in the image space. The function SG (“Stop Gradient”) is used to stop propagating the gradients of the losses w.r.t. $\hat{\mathbf{n}}_u$ during backpropagation, i.e., fix the values of $\hat{\mathbf{n}}_u$ so that it can be propagated to neighbor pixels.

Now the pixel sample set \mathcal{S} of each iteration during the optimization is the union of the set of center pixels \mathcal{S}_c and the sets of neighbor pixels \mathcal{N}_u :

$$\mathcal{S} = \mathcal{S}_c \cup \left\{ \bigcup_{u \in \mathcal{S}_c} \mathcal{N}_u \right\} \quad (13)$$

Empirically, larger \mathcal{S}_c can lead to better results because it helps capture global structures, which are important for optimizing a neural field [44], while larger \mathcal{N}_u is beneficial to smooth the surface. Therefore, we sample \mathcal{N}_u in a criss-cross pattern instead of a whole patch around $\hat{\mathbf{n}}_u$ to balance $|\mathcal{S}_c|$ and $|\mathcal{N}_u|$ with fixed $|\mathcal{S}|$. In principle, \mathcal{N}_u can be sampled in other relatively sparse patterns. Here we use the criss-cross one for its simple implementation in code.

4.4 Optimization Strategy

The overall loss is

$$\mathcal{L}_{\text{all}} = \mathcal{L}_{\text{color}} + \lambda_p \mathcal{L}_{\text{pol}}^p + \lambda_n \mathcal{L}_{\text{normal}} + \lambda_e \mathcal{L}_{\text{eikonal}} \quad (14)$$

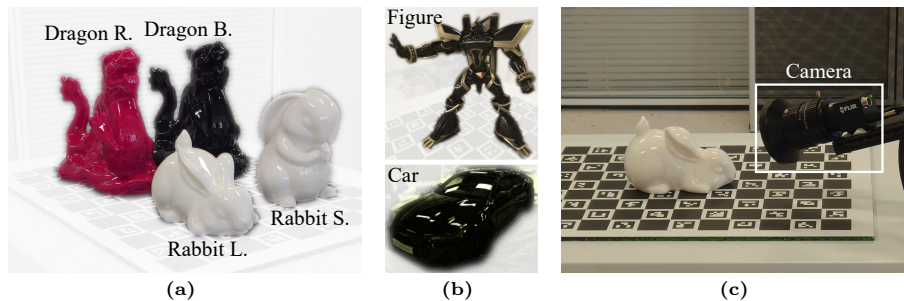


Fig. 3: Dataset and collection. (a) Objects with ground-truth meshes. (b) Objects without ground-truth meshes. (c) Capture setup.

The extra Eikonal loss $\mathcal{L}_{\text{eikonal}} = \frac{1}{M|S|} \sum_{k=1}^{M|S|} (\|\nabla \hat{s}(\mathbf{x}_k)\| - 1)^2$ is used for encouraging the neural field to approximate a signed distance field [10]. During the optimization, we apply progressive training [17] on the multi-resolution hash grid. To initialize a coarse shape, we set $\lambda_p = \lambda_n = 0$ and use the photometric and eikonal losses at the early stage. Then we gradually increase λ_p and λ_n to rectify the distorted shape. Finally, we gradually decrease λ_n to 0 to avoid over smoothing the details.

5 Experiments

5.1 Experimental Setup

Dataset: Since there is no public polarization image dataset of textureless and specular objects with ground-truth meshes, we collect a real-world dataset, including four objects (*Black Dragon*, *Red Dragon*, *Standing Rabbit*, *Lying Rabbit*, *Figure* and *Car*) as shown in Fig. 3. The reasons for capturing real object images, rather than synthesizing using a rendering engine, such as [13], are two-fold: 1) the natural ambient lighting conditions are complex and therefore difficult to synthesize, and 2) the shadows and occlusions of the camera can be realistically generated.

Each object’s polarization images are captured indoors under natural lighting conditions from 40 (60 for *Figure* for its relatively complex shape) different viewpoints around the object with a color polarization camera [36]. Color images are obtained by first decomposing polarization images into images of four polarizer angles and then averaging and demosaicing. Camera poses are estimated from color images using SuperGlue [32] and COLMAP [33] given pre-calibrated camera intrinsics. We use a checkerboard to help pose estimation. To obtain ground-truth meshes, the objects are coated by 3D scanning spray to eliminate specular reflection and scanned by a structure light-based scanner except for *Figure* and *Car*. Please refer to our supplementary material for more details.




Fig. 4: Signed error maps. Red and blue mean surface swelling and shrinking respectively. Regions with higher color saturation indicate longer distances between the reconstruction result and the ground truth. Errors are truncated to within $\pm 2\text{mm}$ for better visualization.

Table 1: Summary of Methods in Experiments

Methods	Image	Shape Repre.	Appearance Repre.
Ref-NeuS [9]	RGB	Neural SDF	Neural radiance
NeRO [18]	RGB	Neural SDF	Neural BRDF
NeuS [42]	RGB	Neural SDF	Neural radiance
PANDORA [6]	Pol. RGB	Neural SDF	Neural BRDF
PMVIR [54]	Pol. RGB	Mesh	Spherical harmonics
PISR (Ours)	Pol. RGB	Grids+Neural SDF	Neural radiance

Performance Metrics: Since we focus on geometry reconstruction, we only evaluate the surface reconstruction quality using L1 Chamfer distance and F-score. We remove the background meshes before calculating the metrics.

Implementation Details: The SDF MLP Φ_s has one hidden layer with a size of 64, and the color MLP Φ_c has two hidden layers with sizes of 64. The multi-resolution grid contains 16 levels of resolution, and the coarsest and finest resolutions are 32 and 2700, respectively. The hash table size of each level is 2^{19} . The optimization of the grid starts at the fourth-coarsest resolutions and progressively extends one level every 1000 iterations. Throughout the whole optimization process, the threshold of DoP θ is set to 0.3 and λ_{eikonal} is set to 0.1. After the first 2.5k iterations, λ_{pol} , λ_{normal} and $|\mathcal{N}_p|$ are linearly increased from

Table 2: Reconstruction accuracy in Chamfer distance (CD↓, mm) and F-score (FS↑, %). FS is calculated at a threshold of 1 mm. Note that the values in the table are rounded to one decimal place. The CD of PMVIR and PISR on Rabbit L. are 0.584 mm and 0.578 mm respectively. **Bold**, underline, and double underline mean the best, second-best and third-best performance, respectively.

Methods		Dragon B.		Dragon R.		Rabbit S.		Rabbit L.		Avg.	
		CD	FS	CD	FS	CD	FS	CD	FS	CD	FS
RGB	Ref-NeuS	2.1	65.3	2.1	66.1	1.2	84.5	1.0	90.5	1.6	76.6
	NeRO	2.4	59.2	1.8	66.3	1.3	80.5	1.3	85.0	1.7	72.8
	NeuS	<u>1.9</u>	<u>67.1</u>	<u>1.2</u>	<u>86.0</u>	<u>1.1</u>	<u>86.8</u>	<u>0.9</u>	<u>93.0</u>	<u>1.3</u>	<u>83.2</u>
P. RGB	PANDORA	2.3	55.3	1.8	64.4	1.4	78.3	1.5	77.1	1.8	68.8
	PMVIR	<u>1.8</u>	<u>76.5</u>	<u>1.0</u>	<u>90.5</u>	<u>0.6</u>	<u>97.5</u>	<u>0.6</u>	<u>97.3</u>	<u>1.0</u>	<u>90.5</u>
	PISR (Ours)	0.5	99.9	0.6	99.1	0.5	99.7	0.6	99.4	0.5	99.5

0 to 2.0, 1.0 and 28 in 2.5k iterations, respectively. Then λ_{normal} and $|\mathcal{N}_p|$ are linearly decreased to 0 in another 2.5k iterations. The maximum number of iterations is 20k. The final SDF is converted to meshes using Marching Cubes [20].

Baselines: We compare our method with five methods, including PMVIR [54], PANDORA [6], NeRO [18], Ref-NeuS [9] and NeuS [42]. Tab. 1 summarizes the type of input images, geometry representation and color representation of all the six methods. Since PMVIR is a refinement method based on an initial shape, we provide the coarse initial meshes reconstructed by PISR-C (see the ablation study in Sec. 5.3 for the details of PISR-C).

5.2 Results

Quantitative Results: The Chamfer distances and F-scores of the reconstruction results are reported in Tab. 2. Our method achieves the best reconstruction quality on all four objects shown in Fig. 3a and significantly outperforms the other five methods. The average Chamfer distance and F-score of our method are 0.5 mm and 99.5%, respectively. As the second-best method in the comparison, PMVIR achieves twice the Chamfer distance and a 9% lower F-score compared to ours. Both PMVIR and our method utilize polarimetric constraints, and the results show that the geometry cue of polarization is effective for shape estimation. Although PANDORA also uses polarization images as input, it implicitly utilizes polarization cues through polarimetric rendering, i.e., optimizing shapes using an image reconstruction loss, which could suffer from the shape-radiance ambiguity as RGB-based methods. Among RGB-based methods, NeuS achieves better results than Ref-NeuS and NeRO which are designed for specular objects.

We visualize the error distributions in Fig. 4. For methods that rely solely on image reconstruction losses (Ref-NeuS, NeRO, NeuS and PANDORA), their reconstruction errors in non-convex regions are notably higher compared to




Fig. 5: Normal maps of Dragon B. before and after optimizations. Using neural SDFs to represent shapes allows changes in the shape topology during the optimization.

Table 3: Average optimization time of different methods.

Methods	Ref-NeuS	NeRO	NeuS	PANDORA	PMVIR	PISR (Ours)
Iterations	200k	200k	200k	100k	4	20k
Time (h)	20.0	20.5	10.8	15.0	2.0	0.5

PMVIR and our method. While the inherent continuity of MLPs ensures the watertightness of their reconstructed shapes, their accuracy cannot be guaranteed due to the shape-radiance ambiguity introduced by specular reflections and the lack of textures. Additionally, for NeRO and PANDORA, the lighting shadows cast by the camera further exacerbate the difficulty of estimating surface material parameters. The real HDR ambient lighting also poses challenges for PANDORA and Ref-NeuS to accurately estimate object shapes. In contrast, PMVIR and our method can recover non-convex regions using the polarimetric constraints. Despite the polarization cues, representing the shapes as neural SDF helps our method better reconstruct regions that are poorly initialized. As shown in Fig. 5, neural SDF allows continuous changes of the shape topology during optimization, which is crucial for evolving surfaces from holes.

Optimization Time: We report the average optimization time in Tab. 3. The experiments are conducted on a machine equipped with an Intel Core i9-12900K and an NVIDIA RTX 3090 GPU. We set the maximum number of iterations to 200k. If the shape converges, we stop the optimization earlier. Our method represents shape as a multi-resolution hash grid and two shallow MLPs, significantly reducing the time complexity for rendering a pixel. Therefore, our method converges much faster, achieving up to $4 \sim 41\times$ convergence speed while maintaining the highest reconstruction accuracy.

5.3 Ablation Study

We evaluate the effectiveness of the polarimetric loss and normal regularization, as well as the influence of the perspective effect of the lens. We compare PISR(-

Table 4: Results of the ablation study on different combinations of losses. The L1 CD↓ (mm) and FSs↑ (%) at 1.0 mm and 0.5 are the averages over four objects in Fig.3a. The ranks are the averages over the three metrics. $\mathcal{L}_{\text{pol}}^{\text{o}}$ differs from $\mathcal{L}_{\text{pol}}^{\text{p}}$ by its use of the orthographic constraint Eq. 2 instead of the perspective one Eq. 3. Note that $\mathcal{L}_{\text{eikonal}}$ is used in all experiments and is omitted it in this table.

Variants	Used Losses				Avg.			Rank
	$\mathcal{L}_{\text{color}}$	$\mathcal{L}_{\text{pol}}^{\text{o}}$	$\mathcal{L}_{\text{pol}}^{\text{p}}$	$\mathcal{L}_{\text{normal}}$	CD	FS(1.0)	FS(0.5)	
PISR-C	✓				1.10	88.10	70.30	6
PISR-N	✓			✓	1.00	88.20	71.10	5
PISR-O	✓	✓			0.80	<u>96.60</u>	<u>75.54</u>	3
PISR-ON	✓	✓		✓	<u>0.77</u>	96.10	75.50	4
PISR-P	✓		✓		0.51	99.76	92.79	1
PISR(-PN)	✓		✓	✓	<u>0.54</u>	<u>99.54</u>	<u>91.75</u>	2

PN) with its five variants, referred to as *PISR-P*, *PISR-ON*, *PISR-O*, *PISR-N* and *PISR-C*, based on different combinations of losses as shown in Tab. 4. The quantitative experiment is conducted on objects shown in Fig. 3a and the qualitative experiment is conducted on objects shown in Fig. 3b.

As shown in Tab. 4, the top four performing variants are all incorporated with the polarimetric losses, and fine-grained details can be recovered without complicated appearance modeling. As for the normal regularization, although the quantitative comparison between PISR-PN and PISR-P shows that it slightly reduces the accuracy, the qualitative results in Fig. 6 show its effectiveness for recovering surfaces from severe distortions.

Among the top four performing variants, PISR-PN and PISR-P achieve higher accuracy by considering the perspective effect of the lens using Eq. (3) instead of Eq. (2). As the example error maps of Rabbit L. shown in Fig. 7, the perspective effect leads to surface swelling and shrinking, causing severe cracks on the surface. This phenomenon does not appear in the results of PMVIR and PANDORA. We speculate the reason to be that the hash-grid-based neural SDF we use is too flexible in topology compared to the mesh-based one used in PMVIR, and lacks the smoothness of purely MLP-based one used in PANDORA. Besides, the results in Tab. 4 show that PISR-PN outperforms PISR-ON by a 30% reduction in L1 Chamfer distance and a 17% improvement in F-score at 0.5 mm, which suggests the necessity of considering the perspective effect to improve the accuracy of polarimetric 3D reconstruction.

6 Conclusion

We have proposed PISR, an accurate and efficient polarimetric neural implicit surface reconstruction method for textureless and specular objects. The key idea of PISR is to explicitly regularize the SDF using the polarimetric constraints to avoid the shape-radiance ambiguity. More importantly, PISR achieves high accu-




Fig. 6: Normal maps of Figure and Car. The polarimetric loss and normal loss jointly recover the shape from distortions induced by strong reflections.




Fig. 7: Signed error maps of Rabbit L. PISR-PN leverages the perspective polarimetric constraint to eliminate surface cracks encountered by PISR-ON. Additional error maps can be found in our supplementary material.

racy by leveraging the perspective polarimetric constraint, eliminating the surface cracks caused by the traditional orthographic one. Besides, PISR regularizes surface normals in image space for robustness and integrates the multi-resolution hash grid for efficiency.

Limitations: Despite the high reconstruction accuracy in the experiments, PISR’s performance degenerates with noisy AoP maps, limiting its application in reconstructing rough, well-textured or mirror-like objects. Recent works [16, 18] have shown solutions to mitigate this problem. In addition, PISR also requires dense input views, adding shape constraints derived from normal and depth estimation models would be a promising solution.

References

1. Atkinson, G.A., Hancock, E.R.: Recovery of surface orientation from diffuse polarization. *IEEE transactions on image processing* **15**(6), 1653–1664 (2006)
2. Ba, Y., Gilbert, A., Wang, F., Yang, J., Chen, R., Wang, Y., Yan, L., Shi, B., Kadambi, A.: Deep shape from polarization. In: *European Conference on Computer Vision*. pp. 554–571. Springer (2020)
3. Baek, S.H., Zeltner, T., Ku, H., Hwang, I., Tong, X., Jakob, W., Kim, M.H.: Image-based acquisition and modeling of polarimetric reflectance. *ACM Trans. Graph.* **39**(4), 139 (2020)
4. Chen, G., He, L., Guan, Y., Zhang, H.: Perspective phase angle model for polarimetric 3d reconstruction. In: *European Conference on Computer Vision*. pp. 398–414. Springer (2022)
5. Cui, Z., Gu, J., Shi, B., Tan, P., Kautz, J.: Polarimetric multi-view stereo. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1558–1567 (2017)
6. Dave, A., Zhao, Y., Veeraraghavan, A.: Pandora: Polarization-aided neural decomposition of radiance. In: *European Conference on Computer Vision*. pp. 538–556. Springer (2022)
7. Ding, Y., Ji, Y., Ye, J.: Polar-photometric stereo under natural illumination. In: *2022 International Conference on 3D Vision (3DV)*. pp. 690–699. IEEE (2022)
8. Ding, Y., Ji, Y., Zhou, M., Kang, S.B., Ye, J.: Polarimetric helmholtz stereopsis. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5037–5046 (2021)
9. Ge, W., Hu, T., Zhao, H., Liu, S., Chen, Y.C.: Ref-neus: Ambiguity-reduced neural implicit surface learning for multi-view reconstruction with reflection. *arXiv preprint arXiv:2303.10840* (2023)
10. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. In: *Proceedings of the 37th International Conference on Machine Learning*. pp. 3789–3799 (2020)
11. 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)
12. Ichikawa, T., Fukao, Y., Nobuhara, S., Nishino, K.: Fresnel microfacet brdf: Unification of polari-radiometric surface-body reflection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16489–16497 (2023)
13. Jakob, W., Speierer, S., Roussel, N., Nimier-David, M., Vicini, D., Zeltner, T., Nicolet, B., Crespo, M., Leroy, V., Zhang, Z.: Mitsuba 3 renderer (2022), <https://mitsuba-renderer.org>
14. Kondo, Y., Ono, T., Sun, L., Hirasawa, Y., Murayama, J.: Accurate polarimetric brdf for real polarization scene rendering. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX* **16**. pp. 220–236. Springer (2020)
15. 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 (June 2022)
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. *arXiv preprint arXiv:2311.13187* (2023)

17. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8456–8465 (2023)
18. Liu, Y., Wang, P., Lin, C., Long, X., Wang, J., Liu, L., Komura, T., Wang, W.: Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. In: SIGGRAPH (2023)
19. Long, X., Lin, C., Wang, P., Komura, T., Wang, W.: Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In: European Conference on Computer Vision. pp. 210–227. Springer (2022)
20. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH Computer Graphics* **21**(4), 163–169 (1987)
21. Lyu, Y., Zhao, L., Li, S., Shi, B.: Shape from polarization with distant lighting estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(11), 13991–14004 (2023). <https://doi.org/10.1109/TPAMI.2023.3298376>
22. Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7210–7219 (2021)
23. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
24. Miyazaki, D., Shigetomi, T., Baba, M., Furukawa, R., Hiura, S., Asada, N.: Polarization-based surface normal estimation of black specular objects from multiple viewpoints. In: 2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission. pp. 104–111. IEEE (2012)
25. Miyazaki, D., Tan, R.T., Hara, K., Ikeuchi, K.: Polarization-based inverse rendering from a single view. In: Computer Vision, IEEE International Conference on. vol. 3, pp. 982–982. IEEE Computer Society (2003)
26. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* **41**(4), 1–15 (2022)
27. Ngo Thanh, T., Nagahara, H., Taniguchi, R.i.: Shape and light directions from shading and polarization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2310–2318. IEEE (2015)
28. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2020)
29. Oechsle, M., Peng, S., Geiger, A.: Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In: International Conference on Computer Vision (ICCV) (2021)
30. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019)
31. Rahmann, S., Canterakis, N.: Reconstruction of specular surfaces using polarization imaging. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001. vol. 1, pp. I–I. IEEE (2001)
32. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: SuperGlue: Learning feature matching with graph neural networks. In: CVPR (2020)

33. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
34. Shakeri, M., Loo, Y., Zhang, H.: Polarimetric monocular dense mapping using relative deep depth prior [j]. 2021. Shakeri M Loo Y Zhang H Polarimetric Monocular Dense Mapping Using Relative Deep Depth Prior (2021)
35. Shao, M., Xia, C., Yang, Z., Huang, J., Wang, X.: Transparent shape from a single view polarization image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9277–9286 (October 2023)
36. Sony: Polarization image sensor with four-directional on-chip polarizer and global shutter function, <https://www.sony-semicon.co.jp/e/products/IS/industry/product/polarization.html>, accessed 2024-02-04
37. Tian, C., Pan, W., Wang, Z., Mao, M., Zhang, G., Bao, H., Tan, P., Cui, Z.: Dps-net: Deep polarimetric stereo depth estimation. In: The IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
38. Tozza, S., Zhu, D., Smith, W., Ramamoorthi, R., Hancock, E.: Uncalibrated, two source photo-polarimetric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
39. Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T., Srinivasan, P.P.: Ref-nerf: Structured view-dependent appearance for neural radiance fields. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5481–5490. IEEE (2022)
40. Wang, F., Rakotosaona, M.J., Niemeyer, M., Szeliski, R., Pollefeys, M., Tombari, F.: Unisdf: Unifying neural representations for high-fidelity 3d reconstruction of complex scenes with reflections. In: arXiv (2023)
41. Wang, H., Wang, J., Agapito, L.: Co-slam: Joint coordinate and sparse parametric encodings for neural real-time slam. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13293–13302 (2023)
42. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems* **34**, 27171–27183 (2021)
43. Wang, Y., Han, Q., Habermann, M., Daniilidis, K., Theobalt, C., Liu, L.: Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3295–3306 (2023)
44. Xie, Z., Yang, X., Yang, Y., Sun, Q., Jiang, Y., Wang, H., Cai, Y., Sun, M.: S3im: Stochastic structural similarity and its unreasonable effectiveness for neural fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18024–18034 (2023)
45. Xu, Q., Tao, W.: Multi-scale geometric consistency guided multi-view stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
46. Yang, L., Tan, F., Li, A., Cui, Z., Furukawa, Y., Tan, P.: Polarimetric dense monocular slam. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3857–3866 (2018)
47. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. In: Thirty-Fifth Conference on Neural Information Processing Systems (2021)
48. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* **33** (2020)

49. Yu, Y., Zhu, D., Smith, W.A.: Shape-from-polarisation: a nonlinear least squares approach. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 2969–2976 (2017)
50. Yu, Z., Peng, S., Niemeyer, M., Sattler, T., Geiger, A.: Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems* **35**, 25018–25032 (2022)
51. Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: PhySG: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In: The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
52. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492* (2020)
53. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)* **40**(6), 1–18 (2021)
54. Zhao, J., Monno, Y., Okutomi, M.: Polarimetric multi-view inverse rendering. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022)
55. Zhao, J., Oishi, J., Monno, Y., Okutomi, M.: Polarimetric patchmatch multi-view stereo. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3476–3484 (2024)
56. Zheng, E., Dunn, E., Jovic, V., Frahm, J.M.: Patchmatch based joint view selection and depthmap estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2014)