

# PhyGaP: Physically-Grounded Gaussians with Polarization Cues

Jiale Wu<sup>2,1,†,‡</sup> Xiaoyang Bai<sup>1,\*</sup> Zongqi He<sup>1</sup> Weiwei Xu<sup>2,†</sup> Yifan Peng<sup>1,\*</sup>

<sup>1</sup>The University of Hong Kong <sup>2</sup>Zhejiang University

jialewu2022@zju.edu.cn xybai@hku.hk zongqi\_he@connect.hku.hk

xww@cad.zju.edu.cn evanpeng@hku.hk

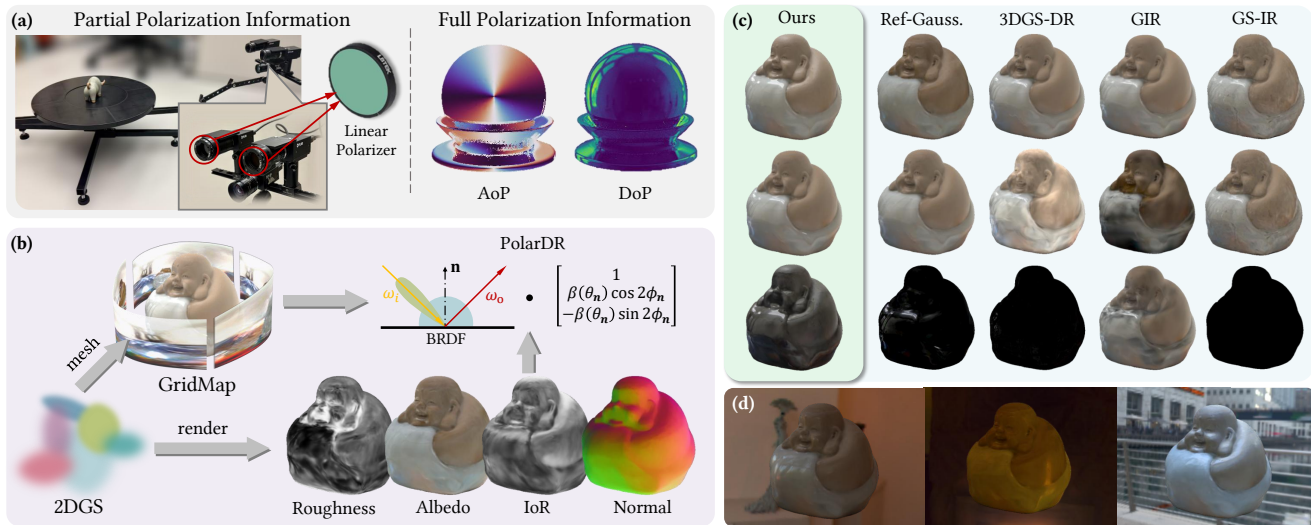


Figure 1. We propose **PhyGaP**, a physically-grounded 3DGS method that (a) takes full or partial polarization information as input, (b) accurately reconstructs the shape and physical attributes of glossy object, (c) achieves decomposed rendering of object appearance (*top*), diffuse reflection (*mid*), and specular reflection (*bottom*), as well as (d) enables robust and realistic relighting with natural reflection. Results in this visualization are from our captured *buddha* scene.

## Abstract

Recent advances in 3D Gaussian Splatting (3DGS) have demonstrated great success in modeling reflective 3D objects and their interaction with the environment via **deferred rendering (DR)**. However, existing methods often struggle with correctly reconstructing physical attributes such as albedo and reflectance, and therefore they do not support high-fidelity relighting. Observing that this limitation stems from the lack of **shape and material** information in RGB images, we present *PhyGaP*, a physically-grounded 3DGS method that leverages polarization cues to facilitate precise reflection decomposition and visually consistent relighting of reconstructed objects. Specifically, we design a polarimetric deferred rendering (*PolarDR*) process to model polarization by reflection, and a self-occlusion-aware environment map building technique (*GridMap*) to resolve in-

direct lighting of non-convex objects. We validate on multiple synthetic and real-world scenes, including those featuring only partial polarization cues, that *PhyGaP* not only excels in reconstructing the appearance and surface normal of reflective 3D objects ( $\sim 2$  dB in PSNR and 45.7% in Cosine Distance better than existing RGB-based methods on average), but also achieves state-of-the-art inverse rendering and relighting capability. Our code is available at <https://github.com/Kelvar00/PhyGaP>.

## 1. Introduction

Modeling reflective objects ranks among the most challenging task in 3D object reconstruction. For 3DGS [28] variants particularly, the lack of explicit geometry representation and the inability of the splatting pipeline to simulate secondary light transport limit their capability of representing glossy surfaces [5, 15]. Only recently have several works emerged to tackle this bottleneck. To regularize surface smoothness, 2DGS [23] and Gaussian Surfel [11] equip Gaussian primitives with explicit normal attributes, while

\*Corresponding authors.

†Affiliated with the State Key Lab of CAD&CG, Zhejiang University.

‡Work conducted during Jiale Wu’s visit to HKU.

GaussTR [26] and Feat2GS [8] use foundation models such as Metric3D V2 [22] to provide surface normal supervision. On the other hand, the introduction of **deferred rendering (DR)** into GS training [29, 32] enables the modeling of reflection. Combining both branches together, state-of-the-art methods such as Ref-Gaussian [47] are able to accurately synthesize novel views of reflective objects.

However, objects reconstructed by those pipelines are prone to color shifts, unrealistic shading, or even surface discontinuities when placed under a different lighting condition, indicating incorrect decomposition between object albedo (inherent color) and reflected light (diffuse and specular) (Fig. 1(c)). This is because ray tracing-based DR relies on precise estimation of surface normal, reflectance, and roughness, which ordinary RGB images do not encode. Inspired by prior works that leverage *polarization information* to learn those physical attributes [12, 19], this work presents **PhyGaP**, a physically-grounded GS pipeline using polarization cues to enable *reliable reflection decomposition* and supports *coherent relighting* for glossy objects.

The PhyGaP framework aims to bridge the gap between Gaussian representations and the physical world via accurate reflection modeling. Specifically, we jointly model Gaussian appearance, shape, index of refraction (IoR), and roughness, as well as a learnable environment map. We further derive a polarimetric deferred rendering (**PolarDR**) process that computes pixel-wise Stokes values from this physically-grounded scene representation using the polarimetric BRDF (pBRDF) model, thus enabling direct scene optimization with either full or partial polarization information. Additionally, we present a self-occlusion-aware environment map building technique (**GridMap**) that addresses complex shading of non-convex objects without learning scene-specific indirect lighting, an improvement crucial for robustly relighting objects with arbitrary shapes.

We validate PhyGaP on both synthetic and real-world datasets of multiview polarization images, and evaluate its performance across different aspects: novel view synthesis (NVS) quality, normal accuracy, reflection decomposition, and relighting. Experimental results validate that its reconstruction of object appearance and normal is comparable to state-of-the-art reflection-aware methods, regardless of the availability of polarization cues, while it better estimates the albedo, diffuse, and specular reflection components of target scenes, demonstrating superior decomposition capability and leading to more realistic relighting results.

In summary, this work’s contributions are threefold:

- Our PhyGaP pipeline implements polarization-based GS optimization with a physically-grounded PolarDR process, facilitating accurate normal reconstruction and precise reflection decomposition of glossy objects.
- Our GridMap technique resolves indirect lighting and complex inter-reflection without learning and querying

scene-specific parameters, expanding the relighting capacity of PhyGaP to non-convex objects.

- We validate our claim experimentally on a wide range of synthetic and real-world data. We also demonstrate the effectiveness of PhyGaP when only partial polarization information is available with real-world captures.

## 2. Related Work

**Reflective Object Reconstruction and Inverse Rendering.** As a long-standing topic in the graphics and vision communities, existing 3D reconstruction methods are mostly designed for opaque surfaces with the Lambertian reflectance model [9, 30]. To model view-dependent visual features, Neural Radiance Field (NeRF) [36] takes viewing direction as an argument of its scene function, while 3D Gaussian Splatting (3DGS) [28] leverages spherical harmonics (SH) for color representation. However, these tricks are not physically-grounded and thus still struggle with specular reflection [39].

Attempts to extend 3D reconstruction methods to reflective objects started early in the evolution of NeRF. Works like Mirror-NeRF [49], NeRFReN [18], MS-NeRF [48], and NeRF-MD [40] focus exclusively on mirror and metal surfaces. Ref-NeRF [41] models view-dependent appearances with MLPs which requires a large number of parameters. Spec-NeRF [35], NDE [45], NeRF-Casting [42], and NeRO [34] incorporate physics-based rendering with neural fields to achieve realistic reflection. Yet all these methods suffer from the problem of long runtime and heavy computation that all NeRF variants share.

The mainstream solution to modeling reflection with 3DGS is through *deferred rendering*, which firstly rasterizes per-view physical attribute maps (albedo, roughness, metallic, etc.) through  $\alpha$ -blending, and then computes reflection from evaluating the rendering equation. To obtain accurate and consistent surface normals to estimate incident light direction, GShader [27], GS-IR [32], 3DGS-DR [29], and GIR [37] directly use the shortest principal axis of each Gaussian; Relightable 3DGS (R3DG) [17] instead designates normal as learnable attributes; and Ref-Gaussian [47] bases itself on 2DGS [23]. Despite their effectiveness, many of these methods do not correctly decompose albedo from reflection, and therefore can not support object relighting. We attribute this limitation to insufficient observability of per-pixel reflection state in RGB images alone, a gap we seek to bridge by leveraging polarization cues.

**Polarization-Based Vision and 3D Reconstruction.** Polarization describes the orientation in which light wave oscillates. Since circular polarization is rarely observed in nature, this work is only concerned with *linear polarization*. Linear polarization is known to encode rich cues on surface geometry and material, and therefore has been widely ap-

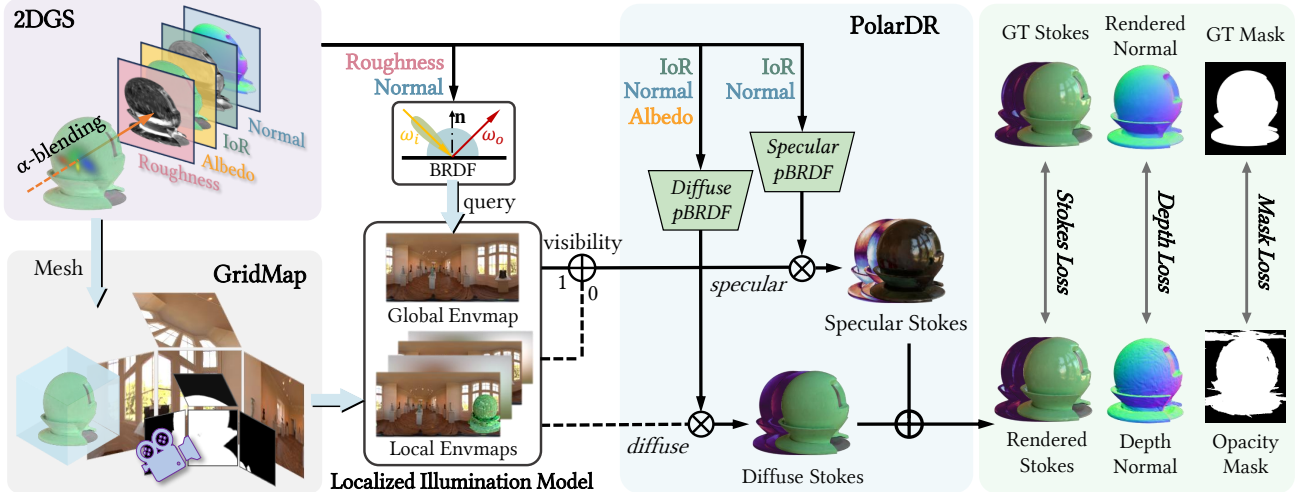


Figure 2. Overview of the PhyGaP pipeline. We represent physically-grounded attributes such as roughness, albedo, IoR and surface normal with 2DGS, and render them into Stokes values via the PolarDR process. Furthermore, we design the GridMap technique to tackle self-occlusion of nonconvex objects. By utilizing polarization cues, we achieve *accurate, explicit* and *disentangled* representation of object albedo, diffuse reflection and specular reflection in PhyGaP.

plied to estimate shape (SfP) [1, 13, 31, 38] and classify material [14, 33]. Its interaction with reflection and refraction further enables applications such as specular reflection removal [43], dehazing [51], underwater localization [4], and time-of-flight imaging [2].

Stemming from existing works on polarization-based multi-view stereo [10, 16, 50], SLAM [46], and depth-from-polarization [6, 16], polarization-based NeRF has recently demonstrated state-of-the-art accuracy in modeling glossy objects. The pioneering PANDORA [12] learns to disentangle diffuse and specular reflection, and renders pixel-wise Stokes vectors via pSVBRDF [3]. NeRSP [19] further utilizes polarization cues to achieve better surface reconstruction. Given the practical challenges of acquiring polarization images, Wu et al. [44] explore shape reconstruction from monocular linearly polarized inputs. The concurrent work PolGS [20] is, to the best of our knowledge, the only attempt to fuse RGB and polarization modalities within 3DGS. However, none of these works supports re-lighting, which is enabled by our proposed PhyGaP.

### 3. Preliminaries: Ref-Gaussian

Our framework is built upon *Ref-Gaussian* [47], which leverages 2D Gaussian Splatting (2DGS) [23] for accurate geometry representation and deferred rendering (DR) to model reflection. We here give a brief overview of its pipeline before diving into PhyGaP itself.

#### 3.1. 2D Gaussian Splatting

2DGS differs from 3DGS in that each Gaussian primitive is two-dimensional, defined by mean position  $\mathbf{p}$ , two tangential vectors  $\mathbf{t}_u$  and  $\mathbf{t}_v$ , and scaling factors  $(s_u, s_v)$ . The

mapping from coordinates  $(u, v)$  on each Gaussian’s local tangent plane to the 3D world frame is thus given by:

$$P(u, v) = \mathbf{p} + s_u \mathbf{t}_u \cdot u + s_v \mathbf{t}_v \cdot v = \mathbf{H} \cdot [u \ v \ 1 \ 1]^T, \quad (1)$$

$$\text{where } \mathbf{H} = \begin{bmatrix} s_u \mathbf{t}_u & s_v \mathbf{t}_v & 0 & \mathbf{p} \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Then, the intersection between the camera ray passing through a given pixel  $(x, y)$  and the tangent plane can be calculated by solving the equation [52]:

$$[xz \ yz \ z \ 1]^T = \mathbf{W}\mathbf{H} \cdot [u \ v \ 1 \ 1]^T, \quad (2)$$

where  $\mathbf{W}$  is the world-to-camera transformation matrix. After  $u, v$ , and  $z$  have been determined, we may then perform standard  $\alpha$ -blending, ordered by  $z$  from front to back, where the influence of each Gaussian is defined as:

$$\mathcal{G}(u, v) = \exp\left(-\frac{u^2 + v^2}{2}\right). \quad (3)$$

Notably, the normal of each Gaussian is easily computed as  $\mathbf{n} = \mathbf{t}_u \times \mathbf{t}_v$ , and a normal map can be subsequently rendered for each viewpoint using  $\alpha$ -blending.

#### 3.2. Deferred Rendering (DR)

*Ref-Gaussian* defines the following physical attributes to each Gaussian: albedo  $\lambda$ , metallic  $m$ , and roughness  $r$ . Instead of directly splatting per-Gaussian colors into an RGB image, it  $\alpha$ -blends these attributes into 2D material maps during rendering, and feeds them into the rendering equation to obtain per-pixel radiance with reflection:

$$L(\omega_o) = \int_{\Omega} L_i(\omega_i) f_r(\omega_i, \omega_o) \langle \omega_i \cdot \mathbf{n} \rangle d\omega_i, \quad (4)$$

$$L_o = L_d(\omega_o) + L_s(\omega_o). \quad (5)$$

Here  $L_d$  and  $L_s$  are the diffuse and specular radiances.  $\omega_o$  and  $\omega_i$  denote the outgoing and incident directions, respectively, determined by the viewpoint direction and surface normal  $\mathbf{n}$ ,  $L_i(\omega_i)$  is the incident radiance obtained by sampling a learnable environment map  $E$ , and  $f_r$  is the BRDF that characterizes how much light is reflected by the object surface. Specifically,  $f_r$  follows the Cook-Torrance model [9] for  $L_s$ :

$$f_r^s(\omega_i, \omega_o) = \frac{DGF}{4\langle\omega_o, \mathbf{n}\rangle\langle\omega_i, \mathbf{n}\rangle}, \quad (6)$$

and the Lambertian model  $f_r^d = \lambda/\pi$  for  $L_d$ . Refer to the **Supplementary Material** for details.

However, this formulation is unable to model the inter-reflection between two object surfaces. To this end, *Ref-Gaussian* simply learns a spherical harmonics (SH) representation of indirect light for each Gaussian and uses it to replace DR-generated color for pixels not directly illuminated by the environment. The whole model is trained with the same set of losses as in vanilla 3DGS.

## 4. Method

Our PhyGaP framework, as illustrated in Fig. 2, features two main improvements upon *Ref-Gaussian* that crucially enable polarization-based optimization and relighting for nonconvex objects. In Sec. 4.1, we detail our PolarDR process and its core pBRDF formulation. In Sec. 4.2, we introduce GridMap, an alternative strategy to model indirect light tailored for relighting tasks.

### 4.1. PolarDR: Polarimetric Deferred Rendering

The polarization state of light is represented by the **Stokes vector**  $\mathbf{s} = [s_0 \ s_1 \ s_2 \ s_3]^\top$ , where  $s_0$  denotes the total intensity,  $s_1$  and  $s_2$  represent the state of linear polarization, and  $s_3$  characterizes circular polarization. Notably, changes to polarization during light-surface interaction can be modeled by a **Mueller matrix**  $\mathbf{M}$  describing the surface’s physical property, which maps the incident Stokes vector  $\mathbf{s}_{in}$  to the outgoing one  $\mathbf{s}_{out} = \mathbf{M}\mathbf{s}_{in}$ .

According to the established pSVBRDF model [3, 12], specular reflection creates strong linear polarization from unpolarized incident light, while diffuse reflection results in much weaker polarization with angle of polarization (AoP) shifted by  $90^\circ$ . Therefore, ground truth polarization cues captured by either polarization camera or regular RGB camera overlaid with linear polarizers (LPs) encode rich information about the object’s reflection property and thus they serve as an ideal modality to facilitate the learning of decomposed physical attributes.

Following the formulation in pSVBRDF, the rendering equation Eq. 4 can be extended to a polarimetric form, *i.e.*,

**pBRDF**, which computes the polarization state of outgoing light after reflection. Assuming unpolarized incident light, specular reflection is modeled as:

$$S_{\omega_o}^s = \begin{bmatrix} 1 \\ \beta_s(\theta_{\mathbf{n}}) \cos 2\phi_{\mathbf{n}} \\ -\beta_s(\theta_{\mathbf{n}}) \sin 2\phi_{\mathbf{n}} \end{bmatrix} L_s(\omega_o), \quad (7)$$

where  $\beta_s(\theta_{\mathbf{n}}) = \frac{R^\perp - R^\parallel}{R^\perp + R^\parallel}$ ,

and diffuse reflection is modeled as:

$$S_{\omega_o}^d = \begin{bmatrix} 1 \\ \beta_d(\theta_{\mathbf{n}}) \cos 2\phi_{\mathbf{n}} \\ -\beta_d(\theta_{\mathbf{n}}) \sin 2\phi_{\mathbf{n}} \end{bmatrix} L_d(\omega_o), \quad (8)$$

where  $\beta_d(\theta_{\mathbf{n}}) = \frac{T^\perp - T^\parallel}{T^\perp + T^\parallel}$ .

Herein,  $L_d(\omega_o)$  and  $L_s(\omega_o)$  denote outgoing radiance as formulated in Eq. 4, and the Fresnel coefficients  $T$  and  $R$  denote the ratios of transmitted and reflected power to the incident light intensity, respectively, with  $\perp$  and  $\parallel$  indicating different orientations of linear polarization. All four Fresnel coefficients can be computed from the index of refraction (IoR)  $\eta$  of the surface and the incident angle  $\langle\mathbf{n}, \omega_i\rangle$ . The zenith angle  $\theta_{\mathbf{n}}$  is defined as  $\theta_{\mathbf{n}} = \cos^{-1}(\mathbf{n} \cdot \omega_o)$ , and the azimuth angle of the polarized light [3, 21]  $\phi_{\mathbf{n}}$  is computed as  $\phi_{\mathbf{n}} = \cos^{-1}(\mathbf{n}_{cam} \cdot \mathbf{y}_{cam})$ , where  $\mathbf{y}_{cam}$  and  $\mathbf{n}_{cam}$  are respectively the  $y$ -axis of the camera coordinate system and the projection of the surface normal  $\mathbf{n}$  onto the image plane (perpendicular to the viewing direction). See Supplementary Material for a detailed formulation of pBRDF.

Similar to *Ref-Gaussian*, our PolarDR process splats the following Gaussian attributes into material maps: albedo  $\lambda$ , IoR  $\eta$ , surface normal  $\mathbf{n}$ , and roughness  $r$ . Specifically,  $r$  is used to determine the mipmap level in which the learnable environment map  $E$  is queried for specular incident radiance  $L_i^s(\omega_i)$  computation. Then we follow Eqs. 7-8 to obtain  $S_{\omega_o}^s$ ,  $S_{\omega_o}^d$ , and their sum  $S_{\omega_o}$  as the overall rendered Stokes vector for each pixel, which is subsequently supervised by ground truth polarization information. Note that we refrain from using SH to represent color since albedo  $\lambda$  should be view-independent by definition.

### 4.2. GridMap: Self-Occlusion-Aware Environment Map Building

We use an environment cube mipmap  $E$  to represent incident light, whose pixel values stand for radiance from a specific direction at infinite distance. During PolarDR,  $L_i^s$  and  $L_i^d$  are obtained by querying  $E$  with the incident direction  $\omega_i$ . However, when the target object is nonconvex, it may cast shadow on or even occlude itself, which results in inconsistency between queried and actual incident light, and creates unrealistic artifacts. This phenomenon is especially pronounced for the diffuse component  $L_i^d$ , since it

integrates  $E$  over a whole hemisphere and is more likely to be affected by self-occlusion.

On the one hand, *Ref-Gaussian* circumvents this issue by learning a standalone SH representation of indirect light  $\mathbf{l}_{\text{ind}}$  for each Gaussian primitive. It then replaces DR-generated radiances with colors computed from  $\mathbf{l}_{\text{ind}}$  for any pixel not “visible” (*i.e.*, the corresponding  $\omega_i$  is blocked by the object itself) from the current viewpoint. Simple as this solution is, it fares poorly in relighting, since  $\mathbf{l}_{\text{ind}}$  simultaneously approximates global illumination and inter-reflection. Yet on the other hand, the classic approach to trace light through multiple reflection bounces till it hits the environment [7, 17] is also undesirable due to its high computation overhead during training.

Our GridMap finds balance between two extremes. As shown in Fig. 3, we divide each face of the object’s bounding box into  $3 \times 3$  grids and place an **anchor camera** at each grid point except for those on the bottom face. For each of the  $N = 52$  cameras, we construct a *local cubemap*  $\tilde{E}_i$  with a resolution  $D^2$  by performing  $6 \times D \times D$  one-step ray tracing, starting from the camera’s location  $\mathbf{c}_i$  and passing through each pixel on the cubemap.

When rendering for the pixel  $(x, y)$ , we perform one PolarDR step using each local cubemap as illumination to obtain a set of diffuse and specular Stokes, namely  $\{\tilde{S}_d^{(1)}, \tilde{S}_d^{(2)}, \dots, \tilde{S}_d^{(N)}\}$  and  $\{\tilde{S}_s^{(1)}, \tilde{S}_s^{(2)}, \dots, \tilde{S}_s^{(N)}\}$ . Then, the *localized diffuse Stokes* is computed as:

$$\tilde{S}_d = \frac{\sum_{i=1}^N \tilde{S}_d^{(i)} / \|\mathbf{p} - \mathbf{c}_i\|_2}{\sum_{i=1}^N 1 / \|\mathbf{p} - \mathbf{c}_i\|_2} \quad (9)$$

where  $\mathbf{p}$  is the surface point corresponding to  $(x, y)$ . The *localized specular Stokes*  $\tilde{S}_s$  is calculated likewise, except that it only applies to “invisible” pixels as defined in prior work [47]. This design is necessary for optimizing  $E$ , since GridMaps are detached from the computation graph and  $E$  cannot receive gradients from Eq. 9 directly.

By blending object color with global illumination, each cubemap  $\tilde{E}_i$  better captures the “environment” that surface points near the anchor  $\mathbf{c}_i$  see, which may include the object itself. The distance-weighted formulation in Eq. 9 further ensures that the object surface has smooth appearance and is free of abrupt color change. Since local cubemaps require no gradient and only need to be updated infrequently, the computation overhead of GridMap is mainly caused by the additional PolarDR steps, which is faster than performing multi-bounce ray tracing and easily parallelizable on GPUs.

### 4.3. Training Strategies

In all, the training loss of our model is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{rgb}} + \lambda_1 \mathcal{L}_{\text{pol}} + \lambda_2 \mathcal{L}_{\text{mask}} + \lambda_3 \mathcal{L}_{\text{depth}} + \lambda_4 \mathcal{L}_{\text{smooth}}. \quad (10)$$




Figure 3. Overview of GridMap. (a) We sample a set of anchor cameras on the object bounding box. (b) For each anchor camera, we ray trace in all directions to blend object color with the global environment map. (c) The resulting local environment map enables more accurate diffuse irradiance computation.

Following 3DGS, the RGB reconstruction loss is defined as  $\mathcal{L}_{\text{rgb}} = 0.8 \cdot \mathcal{L}_1(s_0, \hat{s}_0) + 0.2 \cdot \mathcal{L}_{\text{DSSIM}}(s_0, \hat{s}_0)$ . The polarization reconstruction loss is defined as  $\mathcal{L}_{\text{pol}} = \mathcal{L}_1(s_1, \hat{s}_1) + \mathcal{L}_1(s_2, \hat{s}_2)$ . Here  $\hat{s}_0, \hat{s}_1$  and  $\hat{s}_2$  are the ground truth Stokes values.

Next, as we apply TSDF [47] to construct object mesh and use it for deferred rendering, floating Gaussians can be fatal to the reconstruction process. Therefore, we introduce a mask loss  $\mathcal{L}_{\text{mask}}$  to eliminate out-of-place floaters:  $\mathcal{L}_{\text{mask}} = \mathcal{L}_1(M, \mathcal{O})$ , where  $\mathcal{O}$  denotes the rendered opacity map and  $M$  is the per-frame segmentation mask given by pretrained segmentation models (for captured data) or the renderer (for synthetic data).

In addition, to make sure that 2DGS primitives align consistently with the object surface, we apply the depth-normal consistency loss  $\mathcal{L}_{\text{depth}}$  to enforce the  $\alpha$ -blended Gaussian normal  $\mathbf{n}$  matches the depth normal  $\tilde{\mathbf{n}}$ , calculated from local depth gradients. That is:  $\mathcal{L}_{\text{depth}} = 1 - \tilde{\mathbf{n}}^\top \mathbf{n}$ . Finally,  $\mathcal{L}_{\text{smooth}} = \|\nabla \mathbf{n}\| \exp(-\|\nabla \hat{s}_0\|)$  is an edge-aware normal smoothness loss, aiming to regularize normal variation.

## 5. Experiments

### 5.1. Implementation Details

**Baselines.** We compare PhyGaP against both RGB-only and polarization-assisted models that are able to reconstruct reflective 3D objects. The former category includes GS-IR [32], 3DGS-DR [29], Relightable3DGS (R3DG) [17], GIR[37], and Ref-Gaussian [47]. The latter includes two NeRF-based methods, PANDORA [12] and NeRSP [19], as well as the concurrent work PolGS [20].

We train all GS-based models on an NVIDIA RTX

Table 1. Quantitative results of different methods on novel view synthesis (*left*) and surface normal reconstruction (*right*). Best results are highlighted as **1st**, **2nd**, and **3rd**. This visualization applies to all following quantitative assessments.

	PSNR (dB) $\uparrow$										Cosine Distance (CD) $\downarrow$					
	PANDORA		RMVP		SMVP			Mitsuba3		RMVP		SMVP			Mitsuba3	
	owl	vase	frog	dog	squirrel	snail	david	matpre.	teapot	frog	dog	squirrel	snail	david	matpre.	teapot
R3DG	24.49	25.35	26.56	17.77	19.71	26.97	22.93	25.41	28.30	0.0679	0.1376	0.0916	0.0473	0.1054	0.0795	0.0326
GS-IR	23.76	24.40	31.07	33.33	18.59	27.92	20.47	16.43	21.77	0.0954	0.1523	0.0931	0.0909	0.1005	0.1130	0.0847
GIR	23.81	23.37	28.26	41.10	20.88	23.67	22.78	24.41	27.10	0.0603	0.1674	0.0536	0.0491	0.1307	0.0616	0.0213
3DGS-DR	24.20	26.94	34.68	39.59	20.63	28.46	24.72	24.33	29.07	0.0933	0.1420	0.0484	0.0462	0.1353	0.1069	0.0325
Ref-Gaussian	22.39	26.19	34.13	37.94	17.90	27.08	23.00	23.69	29.67	0.0501	0.1565	0.1122	0.0207	0.1067	0.0362	0.0093
PolGS	24.99	24.85	28.25	28.15	20.92	26.23	27.20	-	-	0.0491	0.1085	0.0343	0.0297	0.0658	-	-
PANDORA*	27.43	30.05	28.42	26.29	27.97	31.56	30.54	-	-	0.0533	0.0863	0.0154	0.0177	0.0663	-	-
NeRSP*	-	-	-	-	23.97	31.56	30.19	-	-	-	-	0.0243	0.0128	0.0545	-	-
<b>Ours</b>	28.14	27.39	32.92	37.82	20.01	28.18	24.62	26.56	29.69	0.0467	0.0910	0.0482	0.0261	0.0597	0.0334	0.0079




Figure 4. Visualization of reconstructed surface normal for synthetic (*snail*) and real-world (*frog*) objects.

4090D GPU for 30k iterations. PANDORA and NeRSP models are trained on an NVIDIA V100 GPU for 150k and 50k iterations, respectively.

**Datasets.** We test all methods on 9 scenes from 4 datasets: two (*owl* and *vase*) from PANDORA [12], two (*frog* and *dog*) from RMVP3D [19], three (*squirrel*, *snail* and *david*) from SMVP3D [19], and two (*matpre* and *teapot*) rendered using Mitsuba3 [25]. Among them, PANDORA and RMVP3D are captured datasets, while SMVP3D and Mitsuba3-rendered scenes are synthetic.

## 5.2. 3D Reconstruction Results

We evaluate 3D reconstruction quality from two aspects: novel view synthesis (NVS, measured by peak signal-to-noise ratio, or PSNR) and surface normal reconstruction (measured by Cosine distance, or CD). Table 1 displays a thorough quantitative comparison among different methods. More metrics such as SSIM and LPIPS for NVS and MAE for surface normal reconstruction are available in the Supplementary Material. We observe that our method exhibits overall better scores than RGB-based baselines and PolGS. While free of the conventional spherical harmonics (SH) representation of Gaussian color, PhyGaP is still able to match state-of-the-art methods in NVS, thanks to its accurate physically-grounded modeling of reflection. Moreover, polarization cues also assist with the learning of surface normal consistency, a trait that is hard to attain with GS

\*Training data for PANDORA and NeRSP models contain test viewpoints, and thereby their scores are only weakly comparable to others.

variants. Our method is even comparable to PANDORA and NeRSP, trained with additional viewpoints for significantly more iterations, on several scenes.

A visual comparison of reconstructed surface normal is available in Fig. 4. For results on other scenes and NVS visualizations, refer to the Supplementary Material.

**Reflection Decomposition.** By leveraging polarization information, PhyGaP achieves precise decomposition of reflection, including each pixel’s albedo, diffuse radiance, specular radiance and the surrounding environment map, as visualized in Fig. 6. Note that not all methods capable of reconstructing glossy objects learn explicit disentanglement of all relevant components. RGB-only methods, as we have predicted, are typically unable to resolve between the object’s albedo and global illumination, while even the polarization-based PolGS is unable to correctly recover the brightness of the environment map.

A more thorough comparison on environment map estimation is available in Fig. 5. More results on *owl* (from the PANDORA dataset) and *frog* (from RMVP) are available in the Supplementary Material.

## 5.3. Relighting Results

Relighting is a crucial task that determines whether the reconstructed 3D model can adapt to changes in scene content and subsequently be used in various downstream applications including VR/AR, gaming and interactive design. Unlike PANDORA that implicitly encodes environment maps or PolGS that does not decompose albedo, PhyGaP stands




Figure 5. Qualitative comparison on estimated environment maps.



Figure 6. Comparison on reflection decomposition. Note that PolGS and Ref-Gaussian do not explicitly model object albedo.

out as the first polarization-based reconstruction method that supports object relighting.

We firstly compare the accuracy of environment map estimation in Table 2 left. Three Mitsuba-rendered scenes where we have access to ground truth environment irradiance scales are used. Then, we test relighting by replacing the reconstructed environment maps and re-rendering the *david* object. Quantitative and qualitative results are available in Table 2 right and Fig. 7, respectively. Thanks to the high-fidelity albedo reconstructed by PolarDR and the occlusion-aware rendering enabled by GridMap, PhyGaP more faithfully preserves object color and simulates shades and highlights across different unseen environment maps.

Table 2. Quantitative evaluation of environment map reconstruction and relighting. *Left*: PSNR of optimized envmaps for different methods on the *Mitsuba-rendered* dataset. *Right*: Polarimetric evaluation of *david* relighted with the *sunset* envmap. Ablation results are also included. “No PDR” is realized by setting  $\lambda_1 = 0$ .

Method	Envmap (PSNR $\uparrow$ )			Relighting		
	Teapot	Matpre.	David	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
R3DG	-	-	-	11.42	0.916	0.0754
GS-IR	6.70	6.70	6.72	16.90	0.957	0.0359
GIR	10.30	10.73	10.01	18.02	0.960	0.0327
3DGS-DR	6.71	8.26	6.71	-	-	-
Ref-Gaussian	7.51	9.86	8.28	-	-	-
Ours (no PDR & GM)	9.97	9.78	10.69	15.56	0.955	0.0369
Ours (no GM)	<b>12.17</b>	15.82	<b>13.00</b>	17.81	0.967	0.0321
Ours	11.50	<b>17.46</b>	12.39	19.18	0.973	0.0255

#### 5.4. Reconstruction with Partial Polarization Cues

To test our method’s applicability in the absence of a dedicated polarization camera, we additionally set up an acquisition system (see Fig. 1(a)) of two FLIR GS3-U3-23S6C-C RGB cameras, each overlaid with an off-the-shelf LBTEK FLP25-VIS-M linear polarizer (LP). Notably, this system only records *partial* polarization information, yet as we can simulate LP by simply multiplying its corresponding Mueller matrix to the Stokes vectors calculated in Eqs. 7–8, PhyGaP is readily trainable under such a setup. Examples on captured data and their reconstruction results are available in Fig. 8. Details of the acquisition prototype and more results are available in the Supplementary Material.

#### 5.5. Ablation Study

We conduct our ablation study on *david* by controlling whether PolarDR and GridMap are used. As shown in Fig. 9, polarization cues effectively exclude specular reflection from the reconstructed albedo, which consequently leads to significantly better environment map quality, while GridMap resolves shadows cast by the sculpture’s non-convex geometry and helps to restore a consistent shade of whiteness across the surface of the object. Therefore, both components are essential for PhyGaP to achieve its established inverse rendering and relighting capabilities. Furthermore, quantitative results on environment map reconstruction and relighting are provided in Table 2.




Figure 7. Relighting results from different methods. The environment maps are presented at the top. Note that R3DG suffers severe overexposure in the rightmost *Brown Photostudio* environment.

## 6. Conclusion

Unlike prior studies on 3DGS-based reconstruction relying solely on a sequence of multi-view RGB inputs, this work explores a robust polarization-based GS optimization pipeline for 3D reconstruction by incorporating physically-grounded PolarDR processing. Furthermore, the presented GridMap mechanism successfully resolves indirect lighting and complex inter-reflection without learning and querying scene-specific parameters, substantially extending the relighting capability of PhyGaP to non-convex geometries and objects. Experiments conducted on extensive synthetic and real-world data have demonstrated state-of-the-art performance in NVS and surface normal reconstruction of our PhyGaP, as well as its exceptional reflection decomposition and relighting capabilities, even when only partial polarization information is available. We envision our proposed PhyGaP pipeline paves the way towards accurate and realistic 3D reconstruction of glossy objects, with broad applicability to VR/AR/MR and interactive designs.

**Limitation and Future Works.** The performance of PhyGaP may degenerate under certain circumstances. Metal surfaces, whose pBRDFs involve complex-valued indices and phase terms, may lead to inaccurate reconstructions. Moreover, although GridMap effectively enhances existing




Figure 8. Examples of our real-world captured data and (from top-right, clockwise) the reconstructed surface normal, specular reflection, diffuse reflection, and final rendered results.

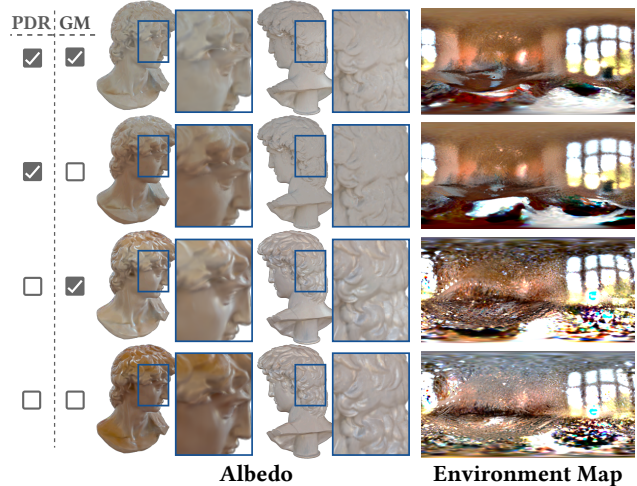


Figure 9. Ablation results with and without PolarDR (PDR) and/or GridMap (GM). We observe visible differences in color correctness, albedo quality, and environment map smoothness.

envmap-based pipelines on global illumination modeling, it is still far from being the ultimate solution and may struggle with highly irregular object shapes or extremely specular surfaces that produce multiple inter-reflection bounces. We have presented failure cases and a comparison between different anchor placement strategies in the Supplementary Material. We also observe that reconstruction results on real-world data can be heavily affected by the disparity between actual 3D environments and our environment map formulation which assumes infinite depth of incident light.

We have envisioned several key steps to tackle these limitations, including: (a) integrating metal surfaces into the PolarDR model; (b) modeling multi-bounce inter-reflection; (c) reducing device requirements of the PhyGaP framework; and (d) leveraging depth-aware environment representations, such as GaussProbe in TransparentGS [24], to further boost the model’s robustness and fidelity on real-world data.

## Acknowledgement

This work was partially supported by the National Key Research and Development Program of China (2024YFE0216600), the National Science Foundation of China (62322217, 62421003), the Innovation and Technology Fund of Hong Kong (MHP/313/24), and the Research Grants Council of Hong Kong (GRF 17208023).

## References

- [1] Yunhao Ba, Alex Gilbert, Franklin Wang, Jinfa Yang, Rui Chen, Yiqin Wang, Lei Yan, Boxin Shi, and Achuta Kadambi. Deep shape from polarization. In *European Conference on Computer Vision*, pages 554–571. Springer, 2020. 3
- [2] Seung-Hwan Baek and Felix Heide. All-photon polarimetric time-of-flight imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17876–17885, 2022. 3
- [3] Seung-Hwan Baek, Daniel S Jeon, Xin Tong, and Min H Kim. Simultaneous acquisition of polarimetric SVBRDF and normals. *ACM Trans. Graph.*, 37(6):268, 2018. 3, 4
- [4] Xiaoyang Bai, Zuodong Liang, Zhongmin Zhu, Alexander Schwing, David Forsyth, and Viktor Gruev. Polarization-based underwater geolocation with deep learning. *Light*, 3(1):15, 2023. 3
- [5] Yanqi Bao, Tianyu Ding, Jing Huo, Yaoli Liu, Yuxin Li, Wenbin Li, Yang Gao, and Jiebo Luo. 3d gaussian splatting: Survey, technologies, challenges, and opportunities. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025. 1
- [6] Marc Blanchon, Désiré Sidibé, Olivier Morel, Ralph Seulin, Daniel Braun, and Fabrice Meriaudeau. P2d: a self-supervised method for depth estimation from polarimetry. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7357–7364. IEEE, 2021. 3
- [7] Hongze Chen, Zehong Lin, and Jun Zhang. Gi-gs: Global illumination decomposition on gaussian splatting for inverse rendering. *arXiv preprint arXiv:2410.02619*, 2024. 5
- [8] Yue Chen, Xingyu Chen, Anpei Chen, Gerard Pons-Moll, and Yuliang Xiu. Feat2gs: Probing visual foundation models with gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6348–6361, 2025. 2
- [9] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG)*, 1(1):7–24, 1982. 2, 4
- [10] Zhaopeng Cui, Jinwei Gu, Boxin Shi, Ping Tan, and Jan Kautz. Polarimetric multi-view stereo. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1558–1567, 2017. 3
- [11] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality surface reconstruction using gaussian surfels. In *ACM SIGGRAPH 2024 conference papers*, pages 1–11, 2024. 1
- [12] Akshat Dave, Yongyi Zhao, and Ashok Veeraraghavan. Pandora: Polarization-aided neural decomposition of radiance. In *European conference on computer vision*, pages 538–556. Springer, 2022. 2, 3, 4, 5, 6
- [13] Valentin Deschaintre, Yiming Lin, and Abhijeet Ghosh. Deep polarization imaging for 3d shape and svbrdf acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15567–15576, 2021. 3
- [14] Wen Dong, Haiyang Mei, Ziqi Wei, Ao Jin, Sen Qiu, Qiang Zhang, and Xin Yang. Exploiting polarized material cues for robust car detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1564–1572, 2024. 3
- [15] Ben Fei, Jingyi Xu, Rui Zhang, Qingyuan Zhou, Weidong Yang, and Ying He. 3d gaussian splatting as new era: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 1
- [16] Yoshiki Fukao, Ryo Kawahara, Shohei Nobuhara, and Ko Nishino. Polarimetric normal stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 682–690, 2021. 3
- [17] Jian Gao, Chun Gu, Youtian Lin, Zhihao Li, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. Relightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing. In *European Conference on Computer Vision*, pages 73–89. Springer, 2024. 2, 5
- [18] Yuan-Chen Guo, Di Kang, Linchao Bao, Yu He, and Song-Hai Zhang. Nerfren: Neural radiance fields with reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18409–18418, 2022. 2
- [19] Yufei Han, Heng Guo, Koki Fukai, Hiroaki Santo, Boxin Shi, Fumio Okura, Zhanyu Ma, and Yunpeng Jia. Nersp: Neural 3d reconstruction for reflective objects with sparse polarized images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11821–11830, 2024. 2, 3, 5, 6
- [20] Yufei Han, Bowen Tie, Heng Guo, Youwei Lyu, Si Li, Boxin Shi, Yunpeng Jia, and Zhanyu Ma. Polgs: Polarimetric gaussian splatting for fast reflective surface reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 28073–28082, 2025. 3, 5
- [21] Eugene Hecht. *Optics*. Pearson Education India, 2012. 4
- [22] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 2
- [23] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In *SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery, 2024. 1, 2, 3
- [24] Letian Huang, Dongwei Ye, Jialin Dan, Chengzhi Tao, Huiwen Liu, Kun Zhou, Bo Ren, Yuanqi Li, Yanwen Guo, and Jie Guo. Transparentgs: Fast inverse rendering of transparent objects with gaussians. *ACM Transactions on Graphics (TOG)*, 44(4):1–17, 2025. 8
- [25] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet,

- Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>. 6
- [26] Haoyi Jiang, Liu Liu, Tianheng Cheng, Xinjie Wang, Tianwei Lin, Zhizhong Su, Wenyu Liu, and Xinggang Wang. Gausstr: Foundation model-aligned gaussian transformer for self-supervised 3d spatial understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 11960–11970, 2025. 2
- [27] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussian-shader: 3d gaussian splatting with shading functions for reflective surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5322–5332, 2024. 2
- [28] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4), 2023. 1, 2
- [29] Ye Keyang, Hou Qiming, and Zhou Kun. 3d gaussian splatting with deferred reflection. 2024. 2, 5
- [30] Johann Heinrich Lambert. *Photometria sive de mensura et gradibus luminis, colorum et umbrae*. sumptibus viduae E. Klett, typis CP Detleffsen, 1760. 2
- [31] Chenyang Lei, Chenyang Qi, Jiaxin Xie, Na Fan, Vladlen Koltun, and Qifeng Chen. Shape from polarization for complex scenes in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12632–12641, 2022. 3
- [32] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. GS-IR: 3D Gaussian Splatting for Inverse Rendering. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21644–21653, 2024. 2, 5
- [33] Yi-Hsin Lin, Hao-Hsin Huang, Yu-Jen Wang, Huai-An Hsieh, and Po-Lun Chen. Image-based polarization detection and material recognition. *Optics Express*, 30(22):39234–39243, 2022. 3
- [34] Yuan Liu, Peng Wang, Cheng Lin, Xiaoxiao Long, Jiepeng Wang, Lingjie Liu, Taku Komura, and Wenping Wang. Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. *ACM Transactions on Graphics (ToG)*, 42(4):1–22, 2023. 2
- [35] Li Ma, Vasu Agrawal, Haithem Turki, Changil Kim, Chen Gao, Pedro Sander, Michael Zollhöfer, and Christian Richardt. Specnerf: Gaussian directional encoding for specular reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21188–21198, 2024. 2
- [36] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [37] Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jian Zhang, Bin Zhou, Errui Ding, and Jingdong Wang. Gir: 3d gaussian inverse rendering for relightable scene factorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2, 5
- [38] Chaoran Tian, Weihong Pan, Zimo Wang, Mao Mao, Guofeng Zhang, Hujun Bao, Ping Tan, and Zhaopeng Cui. Dps-net: Deep polarimetric stereo depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3569–3579, 2023. 3
- [39] Fabio Tosi, Youmin Zhang, Ziren Gong, Erik Sandström, Stefano Mattocchia, Martin R Oswald, and Matteo Poggi. How nerfs and 3d gaussian splatting are reshaping slam: a survey. *arXiv preprint arXiv:2402.13255*, 2024. 2
- [40] Leif Van Holland, Michael Weinmann, Jan U Müller, Patrick Stokto, and Reinhard Klein. Nerfs are mirror detectors: Using structural similarity for multi-view mirror scene reconstruction with 3d surface primitives. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1795–1807. IEEE, 2025. 2
- [41] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 2
- [42] Dor Verbin, Pratul P Srinivasan, Peter Hedman, Ben Mildenhall, Benjamin Attal, Richard Szeliski, and Jonathan T Barron. Nerf-casting: Improved view-dependent appearance with consistent reflections. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–10, 2024. 2
- [43] Sijia Wen, Yinqiang Zheng, and Feng Lu. Polarization guided specular reflection separation. *IEEE Transactions on Image Processing*, 30:7280–7291, 2021. 3
- [44] Bojian Wu, Yifan Peng, Ruizhen Hu, and Xiaowei Zhou. Glossy object reconstruction with cost-effective polarized acquisition. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 422–431, 2025. 3
- [45] Liwen Wu, Sai Bi, Zexiang Xu, Fujun Luan, Kai Zhang, Iliyan Georgiev, Kalyan Sunkavalli, and Ravi Ramamoorthi. Neural directional encoding for efficient and accurate view-dependent appearance modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21157–21166, 2024. 2
- [46] Luwei Yang, Feitong Tan, Ao Li, Zhaopeng Cui, Yasutaka Furukawa, and Ping Tan. Polarimetric dense monocular slam. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3857–3866, 2018. 3
- [47] Yuxuan Yao, Zixuan Zeng, Chun Gu, Xiatian Zhu, and Li Zhang. Reflective Gaussian Splatting. In *ICLR*, 2025. 2, 3, 5
- [48] Ze-Xin Yin, Jiaxiong Qiu, Ming-Ming Cheng, and Bo Ren. Multi-space neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12407–12416, 2023. 2
- [49] Junyi Zeng, Chong Bao, Rui Chen, Zilong Dong, Guofeng Zhang, Hujun Bao, and Zhaopeng Cui. Mirror-nerf: Learning neural radiance fields for mirrors with whitted-style ray tracing. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 4606–4615, 2023. 2
- [50] Jinyu Zhao, Yusuke Monno, and Masatoshi Okutomi. Polarimetric multi-view inverse rendering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7):8798–8812, 2022. 3

- [51] Chu Zhou, Mingui Teng, Yufei Han, Chao Xu, and Boxin Shi. Learning to dehaze with polarization. *Advances in neural information processing systems*, 34:11487–11500, 2021. [3](#)
- [52] M. Zwicker, H. Pfister, J. van Baar, and M. Gross. EWA volume splatting. In *Proceedings Visualization, 2001. VIS '01.*, pages 29–538, 2001. [3](#)