

Robust Depth Enhancement via Polarization Prompt Fusion Tuning

Kei Ikemura^{6*} Yiming Huang^{5*} Felix Heide²
 Zhaoxiang Zhang^{1,4} Qifeng Chen³ Chenyang Lei^{1,2†}

¹CAIR, HKISI-CAS ²Princeton University ³HKUST ⁴CASIA ⁵CUHK ⁶KTH Royal Institute of Technology

Abstract

Existing depth sensors are imperfect and may provide inaccurate depth values in challenging scenarios, such as in the presence of transparent or reflective objects. In this work, we present a general framework that leverages polarization imaging to improve inaccurate depth measurements from various depth sensors. Previous polarization-based depth enhancement methods focus on utilizing pure physics-based formulas for a single sensor. In contrast, our method first adopts a learning-based strategy where a neural network is trained to estimate a dense and complete depth map from polarization data and a sensor depth map from different sensors. To further improve the performance, we propose a Polarization Prompt Fusion Tuning (PPFT) strategy to effectively utilize RGB-based models pre-trained on large-scale datasets, as the size of the polarization dataset is limited to train a strong model from scratch. We conducted extensive experiments on a public dataset, and the results demonstrate that the proposed method performs favorably compared to existing depth enhancement baselines. Code and demos are available at <https://lastbasket.github.io/PPFT/>.

1. Introduction

Estimating accurate depth is crucial for a wide range of applications, from autonomous driving to 3D reconstruction. Over the past decades, various depth sensing methods have been investigated, including Direct [45, 59] and Indirect [16, 20, 49] Time-of-Flight (ToF) cameras, structured light sensors [14, 15], and stereo cameras [13]. Despite wide application, these sensing systems live in a complex trade-off space where each has its own vulnerabilities, including specularities [42], transparency [3], Multi-Path-Interference (MPI) [16, 49], strong ambient light [14, 15], and texturelessness [3] just to name a few.

In this work, we are interested in using polarization data

*Equal contribution

†Corresponding author

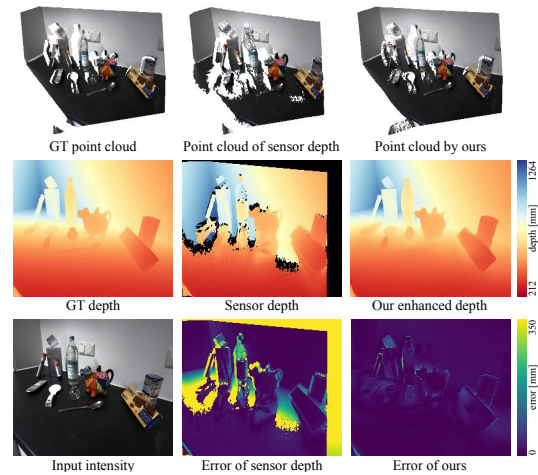


Figure 1. **Visualization of the results of our proposed framework.** The sensor depth shown is from a d-ToF sensor. Our method leverages the dense shape cues from polarization and produces accurate results on challenging depth enhancement problems, including depth in-painting, restoring depth on transparent surfaces, shape correction, etc. For example, see that the depth of the transparent water bottle is restored. Note that on the error map of the sensor depth, we clamp values to the maximum scale used.

as generalizable geometry cues for depth enhancement. The polarization state is closely related to the local surface orientation, providing valuable information for accurate depth estimation. Moreover, as a passive sensor, polarization cameras have the advantage of capturing dense geometry cues without being limited by the depth range [29]. Further, polarization measurement can be used to detect abnormal conditions such as reflection [29, 37] and transparency [25]. This can be significantly reliable for detecting and correcting false depth measurements. Hence, polarization can serve as a highly effective complementary component to existing depth sensors, enhancing their overall performance and expanding their capabilities in various applications.

To the best of our knowledge, no existing work focuses on a general depth enhancement that applies to various depth-sensing technologies using polarization guidance. Some previous studies have explored physics-based methods [22, 62], achieving accurate depth enhancement

Model	RMSE↓	MAE↓	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
Without polarization	43.64	22.84	0.978	0.999	0.999
Naive with polarization	41.17	22.14	0.981	0.999	0.999
Ours with polarization	28.63	14.17	0.992	0.999	1.000

Table 1. **Naively incorporated polarization on HAMMER [21] Dataset.** Results are average on all sensors and "Without polarization" indicates using the baseline [65]. One observes little performance gains by naively passing polarization data into RGB-guided methods directly. In contrast, our approach could effectively utilize the polarization, achieving significantly better performance.

results. However, they are limited by several strict assumptions. For example, Yoshida *et al.* [62] assumed specular reflection and that the target object in interest represents a closed surface. Other studies leveraged deep learning with polarization to address particular depth sensing problems, including texturelessness for stereo-vision [52], and multipath inference for i-ToF sensors [19]. Despite the robust enhancement result, they are limited to specific depth sensing technology. In this work, we aim to extend the capability of deep learning techniques to achieve robust polarization-guided depth enhancement that generalizes across different sensing modalities.

However, we noticed that learning with polarization for general depth enhancement is not trivial. From our primary experiments, we observe little performance gains by directly passing polarization into existing RGB-based depth enhancement models, as illustrated in Table 1. Moreover, the size of existing polarization datasets [21, 55], are not comparable to the large-scale RGB-D datasets [39, 53], resulting in a limited size of the training data available for polarization. On this, the naive approach of collecting as much polarization data as RGB images alleviates the problem but is unfavorable as it is extremely labor-intensive.

To overcome these issues, we propose a cross-modal transfer learning strategy, Polarization Prompt Fusion Tuning (PPFT). First, to prevent overfitting on small polarization datasets, we incorporate pre-trained model weights from large-scale RGB-D datasets into our backbone. Second, we target the modality misalignment by the proposed Polarization Prompt Fusion Block (PPFB), a parallel branch design where a spatial attention-like operation and a channel fovea operation are introduced. Utilizing our PPFB throughout the feature extraction network, the polarization cues can be integrated into the network more effectively.

We conduct extensive experiments on different depth sensor types to validate the effectiveness of our approach. Examples of various depth modalities are shown in Figure 2. We find that our model significantly improves the quality across all types of sensor depths with the help of polarization. While other methods fail in challenging cases such as in the presence of transparent objects, reflective surfaces, etc., our enhanced depth maps are more robust and can restore detailed geometry.




Figure 2. **Visualization of our evaluation data by [21].** One can observe varying types of data degradation in each depth sensor type. For example, in (f) we can see missing measurements at the table surface, due to insufficient texture. Similarly, as the dataset provides all sensor depth mapped to the same camera frame, in (h) we observe a region of unknown depth at the boundary. This arises from the smaller Field-of-View (FoV) of the i-ToF sensor. In this work, we propose a general framework to resolve multiple types of sensor depth degradation.

In summary, we make the following contributions

- We devise a general depth enhancement approach that leverages polarization information as dense shape cues. To the best of our knowledge, our method is the first polarization-based depth enhancement method that can handle different types of sensor depth.
- To tackle data scarcity in existing polarization datasets, we propose a novel cross-modal transfer learning strategy named Polarization Prompt Fusion Tuning (PPFT) for polarization-based vision. PPFT improves the performance significantly on our depth enhancement task. Besides, we also prove its effectiveness on another polarization-based task, Shape-from-Polarization.
- We conduct extensive experiments that validate our method outperforms state-of-the-art methods on different types of depth sensors. We provide ablation studies to understand the effect of polarization and various components of our proposed method.

2. Related Work

Depth Enhancement via Polarization Cues Studies [2, 7, 66] have demonstrated the usefulness of polarization in assisting multi-view stereo depth, particularly in textureless regions. Kadambi *et al.* [22–24] enhanced coarse depth map from structure light cameras by using the shape information from polarization cues. On the other hand, [8, 47, 62, 64] employed polarization to refine incomplete depths from Time-of-Flight sensors, and Jeon *et al.* [19] proposed to address incorrect depth measurement given scattering media by polarization data. In addition, Shakeri *et al.* [50] explored the dense depth map reconstruction with polarization. Recently, [52] extended the polarimetric stereo depth using deep learning method. In contrast, we propose a polarization-based deep learning method for depth enhance-

ment generalized to multi-sensor depths.

Polarization-Guided Vision Task We can also draw insight from works that study alternative vision tasks with polarization. [25, 32, 37] utilized polarization data for object segmentation, fusing the angle of linear polarization, degree of linear polarization, and intensity image to achieve material-specific segmentation, such as on transparent and semi-transparent materials. [4, 18] focus on physics-based polarimetric 3D methods, while [1, 10, 17, 29, 38, 51] explored object-level and scene-level Shape-from-Polarization (SfP) with deep learning. In addition, [9, 26] demonstrate the utilization of polarization in rendering. On low-level vision tasks, [28, 57] leveraged polarization properties of light and deep learning for specular reflection removal. Ono *et al.* [41] proposed to solve color constancy with DoLP, and [27, 31] contributed to novel polarization enhancement approaches. By drawing insights from previous polarimetric vision tasks, we focus on the fusion and complement of polarization and coarse depth information, mitigating ambiguity with deep learning.

RGB-Based Depth Enhancement A major line of research for depth enhancement investigates deep learning with (RGB and/or intensity) images as guidance [35, 40, 43, 45, 48, 56, 59, 60, 63, 65]. Some works have studied a coarse-to-fine strategy with spatial propagation networks (SPN) [36], by learning spatial affinities [5, 6, 43, 60]. Further, geometric constraints are also exploited in the depth enhancement process. Zhang *et al.* [63] used a deep CNN to predict the occlusion boundary and surface normal to regularize the depth enhancement process. Qiu *et al.* [45] and Xu *et al.* [59] proposed end-to-end deep learning frameworks with surface normal as intermediate representations to guide depth enhancement. Following recent advancements in Transformer architectures [11, 54], Rho *et al.* proposed GuideFormer [48], a Transformer-based depth completion model with RGB image guidance. More recently, Zhang *et al.* combined convolution operation and Transformer layer in a unified block [65] for depth completion.

Although there have been numerous studies on depth enhancement using image guidance and deep learning methods, alternative modalities-guided processes have received little attention. In this paper, we propose to use polarization as guidance, offering superior surface geometry information and enabling the detection of irregular conditions.

3. Method

We propose the first end-to-end learning-based general depth enhancement method with polarization guidance. Figure 3 illustrates the proposed method. The method operates on the raw sensor depth \mathbf{d} , as well as the polarization data \mathbf{P} . In the following, we first present the preliminaries

of polarization imaging as it relates to depth enhancement in Section 3.1. We then describe the proposed learning-based polarization-guided depth enhancement method in Section 3.2. Finally, we describe our polarization prompt fusion tuning (PPFT) strategy in Section 3.3.

3.1. Polarization as General Depth Guidance

A division-of-focal-plane (DoFP) polarization camera measures the light intensity \mathbf{I}_{pol} passing through the linear polarizer at polarization angle ϕ_{pol} is given by [1, 29, 55]:

$$\mathbf{I}_{\text{pol}} = \mathbf{I}_{\text{un}} * (1 + \rho * \cos(2\phi - 2\phi_{\text{pol}})), \quad (1)$$

where \mathbf{I}_{un} is the total incident light into the camera, which is unpolarized in general, ρ is the degree of linear polarization (DoLP), and ϕ is the angle of linear polarization (AoLP).

Polarization is affected by the surface geometry of scenes, which can, in turn, provide crucial guidance to depth enhancement. Specifically, assuming a perspective projection, the AoLP ϕ , and corresponding surface normal \mathbf{n} , can be related via [29]:

$$\begin{cases} \Phi = \mathbf{n} \times \mathbf{v} \times [0, 0, 1]^T, & \text{if diffuse dominant} \\ \Phi = \mathbf{n} \times \mathbf{v} \times \mathbf{v} \times [0, 0, 1]^T, & \text{if specular dominant} \end{cases}$$

and

$$\phi = \arctan\left(\frac{\Phi_y}{\Phi_x}\right),$$

where $\Phi = [\Phi_x, \Phi_y, \Phi_z]^T$, and $\mathbf{v} \in \mathbb{R}^3$ is the viewing direction. Though there is a strong relationship between AoLP and surface normal, directly deriving the surface normal vector given polarization data is challenging under the perspective projection assumption [29].

In this work, we rely on the strong correlation between surface normal and polarization data and investigate a learning-based strategy to extract and exploit such guidance for depth enhancement. The polarization guidance is obtained by concatenating the intensity image \mathbf{I} , the angle of polarization image ϕ , the degree of polarization image ρ , and the viewing direction \mathbf{V} as

$$\mathbf{P} = [\mathbf{I}; \phi; \rho; \mathbf{V}], \quad (2)$$

where $[\cdot]$ denote concatenation. A typical DoFP polarization camera captures a raw image \mathbf{I} consisting of four polarization images under different polarizer angles: $0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}$. One can recover the complete linear polarization state ϕ and ρ following Equation 1. As most depth sensors capture scene-level data, the commonly adopted orthographic projection assumption [1, 46, 58] fails for the setting targeted in this work and we assume a perspective projection. Since under such an assumption, the polarization state is closely related to the viewing direction \mathbf{V} , we pass it along with the polarization states. Further details are discussed in the Supplementary Material.

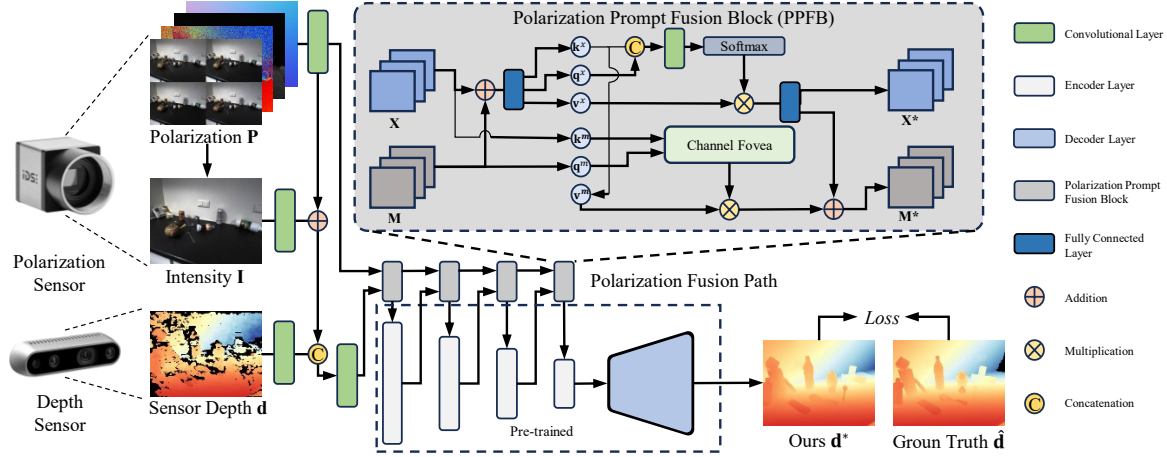


Figure 3. **Polarization Prompt Fusion Tuning (PPFT)**. We fuse polarization embeddings to the features extracted from pre-trained layers sequentially using our proposed Polarization Prompt Fusion Block (PPFB). Specifically, polarization features are passed into our PPFB as prompt M , and features from the pre-trained foundation as the input X . Both are then updated and passed into the next set of pre-trained encoder and our PPFB respectively.




Figure 4. **Channel Fovea Operation in Proposed PPFB**. The summation of the input features is transformed to have doubled channel size with an MLP, and then distilled to a list of weights by global average pooling [33]. These weights are separated into two sets and multiplied with the two inputs, respectively. Output attention weights are computed via a softmax operation on the sum of the re-weighted inputs.

3.2. Learning Polarization Depth Enhancement

Polarization provides strong cues for scene geometry, and existing methods [24, 62] have proposed physics-based approaches for depth enhancement. As such, these methods are designed for specific reflections or materials, which prevent them from being applied to general depth sensors under complex environments. Departing from existing methods, we propose a polarization-guided depth enhancement method that generalizes to diverse types of depth sensors.

Given a depth map d and the corresponding polarization representation P as described above, our method estimates the refined depth with a neural network f as

$$\hat{d} = f(P, d; \theta), \quad (3)$$

where θ are the parameters of the network f .

Following [43], we use a combination of L_1 and L_2 loss

for depth supervision, shown as follows:

$$\mathcal{L}(d^*, \hat{d}) = \frac{1}{|\mathbf{M}_{\text{depth}}|} \sum_{m \in \mathbf{M}_{\text{depth}}} (||d_m^* - \hat{d}_m||_1 + ||d_m^* - \hat{d}_m||_2^2), \quad (4)$$

where d^*, \hat{d} are the ground truth and predicted depth, respectively. $\mathbf{M}_{\text{depth}}$ is the mask for valid depth pixels in d^* .

3.3. Polarization Prompt Fusion Tuning

Pre-Trained Model Weights Existing polarization-depth datasets are relatively small in comparison to RGB-D datasets [39, 53]. This poses a significant challenge to learning a highly generalizable model. This is more severe under our setting since there are diverse types of depth sensors, each having a unique set of measurement degradation characteristics. To overcome this, we transfer knowledge extracted from large-scale RGB-guided depth completion datasets and use it as a foundation for our model generalizability. Specifically, we do this by loading pre-trained model weights to our encoder-decoder backbone, as annotated in Figure 3. We then refine the model by incorporating polarization guidance, through a fusion path consisting of a stack of our proposed Polarization Prompt Fusion Block (PPFB). Our experiments validate that this strategy is necessary to exploit the polarization data. In this work, we use CompletionFormer [65] pre-trained on NYU-Depth V2 dataset [39] as our foundation.

Fusion. As illustrated in Figure 3, our network first encodes each modality via three convolutional encoders. The modality-specific embeddings are then combined as a

multi-modal embedding and passed as an input feature to the first Polarization Prompt Fusion Block (PPFB). Here, we treat the polarization embedding as the prompt feature. A series of PPFB modules then take each previously updated prompt as the input prompt feature, the previous encoder output as the input feature, and produce an updated feature to be passed into the next encoder layer. This operation allows the polarization information to persist and be fused in each layer of the encoder module. Lastly, the encoded features are passed through a decoder block to obtain the final enhanced depth $\hat{\mathbf{d}}$.

It is non-trivial to introduce another modality to a pre-trained foundation model as the polarization modality deviates drastically from the RGB one. To handle this modality misalignment, we devise a Polarization Prompt Fusion Block (PPFB). The PPFB is designed to put emphasis on the polarization input, it updates encoder features sequentially across the encoder block. Inspired by Ye *et al.* [61], the proposed PPFB includes two parallel branches: (1) An input feature update branch in which the polarization prompt \mathbf{M} are utilized to update and produce a fused feature \mathbf{X}^* with a spatial-wise operation; (2) A polarization prompt learning branch where the input features, \mathbf{X} , are used to compute the updated polarization prompt \mathbf{M}^* via a channel-wise fusion strategy. Inspired by the spatial fovea operation introduced by Zhu *et al.* [67], we utilize an attention-like operation for input feature enhancement. In this process, the polarization prompt and the input patch token (C -dimensional) are concatenated and projected by a multi-layer perceptron (MLP) to produce the key \mathbf{k}^x , query \mathbf{q}^x , and value \mathbf{v}^x vectors with learnable parameters $\mathbf{W}_k^x, \mathbf{W}_q^x, \mathbf{W}_v^x \in \mathbb{R}^{2C \times C}$ as

$$[\mathbf{k}^x; \mathbf{q}^x; \mathbf{v}^x] = [\mathbf{M}; \mathbf{X}] \times [\mathbf{W}_k^x; \mathbf{W}_q^x; \mathbf{W}_v^x], \quad (5)$$

$$[\mathbf{X}^*; \mathbf{M}^*] = \text{FC}_d(\lambda \cdot \text{softmax}(\text{FC}([\mathbf{q}^x; \mathbf{k}^x])) \odot \mathbf{v}^x), \quad (6)$$

where \odot is the element-wise multiplication, FC_d is a fully-connected layer with dropout.

With the spatial attention-like output \mathbf{M}^x , we learn a new polarization prompt \mathbf{M}^* with a channel fovea operation, which is illustrated in Figure 4, and as follows:

$$\mathbf{k}^m, \mathbf{q}^m, \mathbf{v}^m = \mathbf{X}, \mathbf{M}, \mathbf{k}^x, \quad (7)$$

$$[s_q; s_k] = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W (\text{FC}(\mathbf{M} + \mathbf{X})), \quad (8)$$

$$\mathbf{M}^* = \mathbf{M}^x + \text{FC}((s_q \mathbf{Q}^m + s_k \mathbf{K}^m) \odot \mathbf{V}^m). \quad (9)$$

This follows the design of Li *et al.* [30], in which we explore the channel-wise relationship for our prompt update by learning channel-wise statistics s_q, s_k for query and key by shrinking their embedding with height and width H, W through spatial dimensions.

With PPFB, we can input our polarization modality into the pre-trained models more efficiently. We randomly ini-

tialize all the parameters of our PPFB. The entire network is then trained and updated with polarization data.

4. Experiments

4.1. Experimental Setup

Dataset We evaluate our method on the public real-world HAMMER dataset [21]. The HAMMER dataset is a polarization-depth dataset consisting of 13 different scenes, with 7200 distinct frames capturing scenes with objects in total. It aids our analysis as it contains both depths from commodity-level depth sensors and paired hyper-quality rendered ground truth depths of the captured scene. Specifically, the dataset includes 3 types of depth sensor measurement: stereo camera depths (D435), d-ToF sensor (L515), and i-ToF sensor depths. This variety of real-world depths allows us to study different depth sensor degradation. We use the official training set for training, which consists of the first nine scenes, giving 5253 frames in total. Further, we split the remaining data into two scenes for validation and two scenes for testing.

Baselines There is no existing polarization-based depth enhancement approach that is tackling the proposed setup. Specifically, given polarization guidance, the task is to enhance commodity-level depth sensor measurement capturing real-world scenes, regardless of the sensing technology used (active stereo, d-ToF, i-ToF, etc.). Thus, we first compare against Polarized 3D [22], an object-level physics-based depth enhancement method. In addition, we adopt 3 state-of-the-art RGB-based depth completion methods as baselines. Namely, we compare with CompletionFormer [65], SemAttNet [40], and DySPN [34], which are learning-based approaches that receive the RGB images as input guidance. We change the network input from RGB images to polarization images and train the models on the same dataset for fair comparison. The fairness is in the sense that polarization information is given to all models. Note that SemAttNet [40] requires additional instance segmentation maps, which are provided in the Hammer dataset.

4.2. Assessment

Table 2 presents the quantitative results of all methods evaluated on the HAMMER [21] dataset. We evaluate our method on each depth sensor type separately to validate the generalizability across various kinds of degradation. Our approach achieves the best performance consistently on all sensor depths and on all metrics. Specifically, compared to CompletionFormer [65], the proposed method obtains a performance gain of 24.4%, 29.7%, 25.3% on the RMSE metric; and 31.8%, 30.6%, 26.5% on the MAE metric, evaluated on stereo, d-ToF, and i-ToF benchmarks respectively. The prediction results on i-ToF sensors are relatively worse

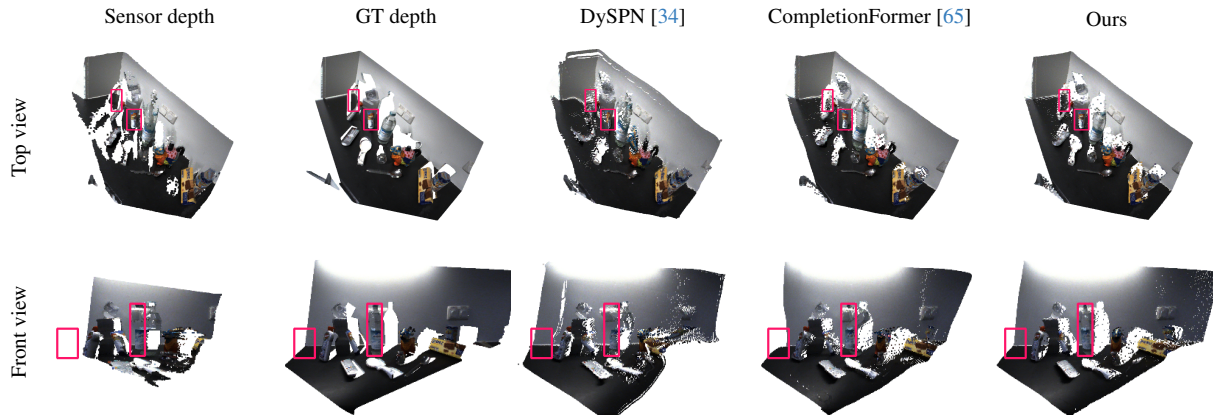


Figure 5. **Comparison between our approach and baselines [34, 65] using point cloud visualization.** In addition to restoring challenging irregularities (e.g. the transparent bottle highlighted in row 2), we also observe a higher degree of regularity (e.g. the left-bottom corner highlighted in row 2, which is misaligned in the sensor depth, resulting in blank point clouds) using our proposed method, showing strong surface geometry information provided by polarization.

Sensor Type	Model	RMSE (mm)↓	MAE (mm)↓	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
Active Stereo	Polarized 3D* [22]	755.64	659.83	0.077	0.164	0.284
	SemAttNet [40]	56.36	36.68	0.975	0.998	0.999
	DySPN [34]	54.56	32.22	0.975	0.997	0.998
	CompletionFormer [65]	32.78	15.64	0.991	0.999	0.999
	CompletionFormer [†] [65]	33.86	15.68	0.991	0.991	0.999
	Ours	24.78	10.66	0.995	0.999	1.000
d-ToF	Polarized 3D* [22]	851.86	777.43	0.034	0.073	0.130
	SemAttNet [40]	54.32	33.92	0.977	0.999	0.999
	DySPN [34]	51.82	31.19	0.977	0.997	0.998
	CompletionFormer [65]	33.93	15.10	0.987	0.999	0.999
	CompletionFormer [†] [65]	36.40	17.22	0.984	0.991	0.999
	Ours	23.84	10.47	0.995	0.999	1.000
i-ToF	Polarized 3D* [22]	709.97	629.98	0.035	0.074	0.123
	SemAttNet [40]	81.35	56.49	0.934	0.997	0.999
	DySPN [34]	69.94	48.11	0.949	0.997	0.999
	CompletionFormer [65]	49.95	29.08	0.974	0.998	0.999
	CompletionFormer [†] [65]	54.94	32.75	0.971	0.998	0.999
	Ours	37.27	21.37	0.987	0.999	1.000
All	Polarized 3D* [22]	772.49	689.08	0.048	0.103	0.179
	SemAttNet [40]	64.01	42.36	0.962	0.998	0.999
	DySPN [34]	58.77	37.17	0.967	0.997	0.998
	CompletionFormer [65]	38.88	19.94	0.984	0.998	0.999
	CompletionFormer [†] [65]	41.73	21.88	0.982	0.993	0.999
	Ours	28.63	14.17	0.992	0.999	1.000

Table 2. **Quantitative comparison on the HAMMER [21] dataset.** To facilitate fair comparisons, we pass polarization data to the RGB-guided models [34, 40, 65]. Sensor depths are obtained from an active stereo camera, d-ToF sensor, and i-ToF sensor. Our method achieves the best performance consistently in all cases. * indicates that directly applying Polarized 3D [22] with author provided code fails in our scenario. † denotes the model is fine-tuned with RGB images. We train on all data and test on different sensor types for analysis.

compared to those on active stereo and d-ToF, this is due to the large area of missing depths that require in-painting since i-ToF sensors have smaller field-of-view (FoV) compared to other devices.

Figure 6 reports the qualitative comparisons of the best three baselines. The proposed method improves the qual-

ity of depth maps from different depth sensors. The results of DySPN [34] are blurry, especially at regions with raw, shape-incorrect depth measurements. Further, the enhanced depth lacks accuracy at detailed regions, this is more apparent in the stereo depth maps. While CompletionFormer [65] produces poor boundaries, our enhanced depths are more

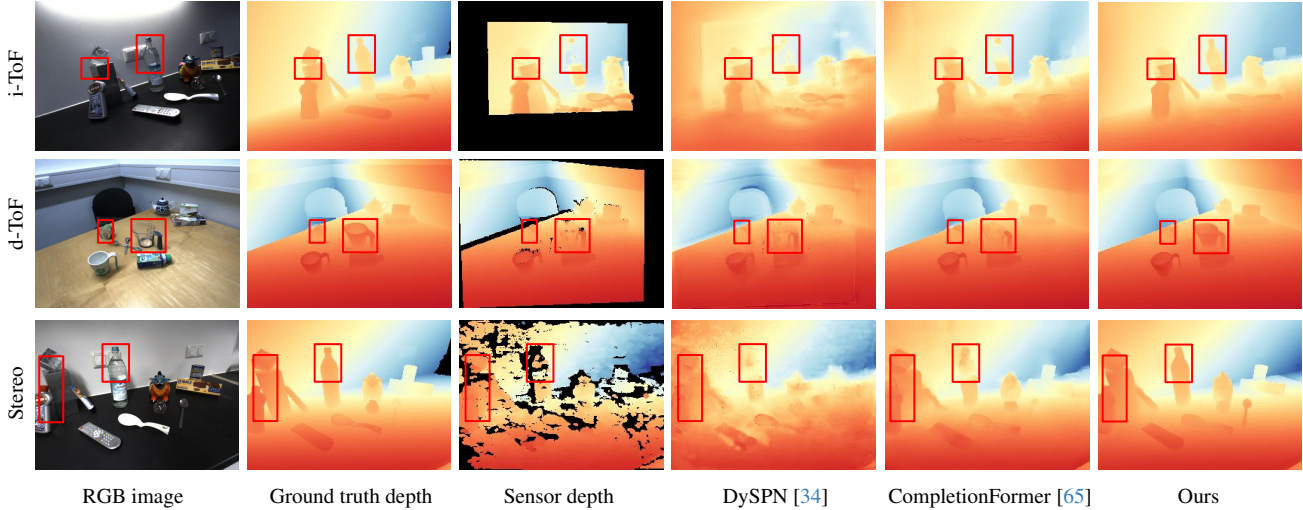


Figure 6. **Qualitative comparison on HAMMER dataset [21].** We present results on different depth sensors. The red boxes highlight regions to emphasize. We can observe significantly better results on transparent surfaces, i.e. the water bottle, and detailed regions, for example, the stack of objects on the left in Row 2.

robust to irregularities that are hard to generalize using RGB guidance, given a small training dataset. Further, the proposed method is able to correct false surface depth in abnormal conditions, such as transparent object surfaces. As the point cloud results in Figure 5 present, our method generates highly regularized depth, for example, the flat walls. Hence we argue that our method can learn the underlying surface normal information from polarization cues. Further, as the stereo (row 1) and d-ToF (row 2) results in Figure 6 illustrate, our method separates foreground and background surfaces adequately (e.g., the teapot and the wall). In the i-ToF case (row 3), our method presents the ability to generalize to the unseen surface, obtaining a smoother transition between the known and unknown regions.

4.3. Analysis

In the following, we conduct ablation experiments on the average of all sensor types to evaluate each component in our method. We adopt the same setting for all experiments and the details are included in the supplementary.

Polarization Guidance We first analyze the importance of polarization information to depth enhancement. In Table 3, we demonstrate the results using our proposed architecture, with and without polarization input during both training and testing. Note that we only substitute the polarization data by the RGB image at the input layer, without modifying the rest of the proposed model structure. We observe that having no polarization guidance results in inferior performance. Using polarization cues improves the performance substantially, validating the proposed method. Figure 7 further illustrates that, with polarization as guid-

PPFT Guidance	RMSE↓	MAE↓	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
RGB	33.06	16.55	0.989	0.999	0.999
Polarization	28.63	14.17	0.992	0.999	1.000

Table 3. **Ablation experiment on the polarization guidance.** By using RGB guidance in PPFT, the performance drops significantly, showing the importance of polarization as geometry guidance.

Method	RMSE (mm)↓	δ RMSE (mm)↓	MAE (mm)↓	δ MAE (mm)↓
CompletionFormer* [65]	313.60	-	241.40	-
Without PPFT	41.17	-	22.14	-
Ours	28.63	-12.54	14.17	-7.97

Table 4. **Ablation experiments for cross-modal transfer learning strategy.** With our proposed Polarization Prompt Fusion Tuning (PPFT) the performance is improved significantly. * denotes that the foundation model is pre-trained on NYU-Depth V2 dataset [39]. δ here indicates the performance gain of our method compared to the case **Without PPFT**.

ance, the proposed method achieves accurate estimation of the transparent regions.

Cross-modal Transfer Learning Our cross-modal transfer learning design is the key to achieving superior performance on real data. The existing polarization-based depth dataset is significantly smaller than RGB-based datasets, preventing most data-driven methods from achieving satisfactory performance. As Table 4 reports, the proposed cross-modal transfer learning module improves the performance significantly compared with training the model on the HAMMER[21] dataset from scratch. Without the proposed Polarization Prompt Fusion Tuning (PPFT) strategy, the model is not able to fit such a complex dense geometry




Figure 7. **Qualitative results for the ablation study on polarization input.** We can observe a significant performance drop given no polarization information, especially in detailed regions.

Fusion Method	RMSE↓	MAE↓	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
Early Fusion [12]	30.33	15.30	0.991	0.999	1.000
Ours (PPFB-shallow)	29.77	14.79	0.991	0.999	0.999
Ours (PPFB-deep)	28.63	14.17	0.992	0.999	1.000

Table 5. **Quantitative comparison on different fusion strategies.** It is vital to inject the PPFB module into different layers of the neural network instead of only the first layer.

prediction problem. In this regard, our proposed method is able to effectively utilize the limited polarization information, adapting well to sensor depth in complex scenes.

Polarization Fusion Module We evaluate our Polarization Prompt Fusion Block (PPFB) with respect to multimodal data fusion. We compare an early fusion strategy [12] with naive polarization concatenation and two variants of our proposed method, (a) a shallow (only a single PPFB at the first encoder layer), and (b) a deep version (i.e. our full model). Quantitative results are presented in Table 5. This result illustrates that with our full-depth parallel branches design, the additional modality can be forwarded into deep layers of the pre-trained foundation more effectively, enabling a thorough fusion with the input features.

4.4. PPFT for Surface Normal Estimation

We highlight the versatility of the proposed PPFT strategy to a wider range of polarization vision tasks, for example, Shape-from-Polarization (SfP) [1, 29]. As Table 6 shows our method can further improve the performance of the SfP method by incorporating our proposed PPFT strategy and a pre-trained RGB-based normal estimation model. Specifically, we apply our method to iDisc [44], a foundation model trained on NYU-Depth V2 dataset [39] that consists of 30k training samples. We then train and evaluate on SPW dataset [29], a scene-level SfP dataset, which in comparison involves 400 training frames. Our method produces high-accuracy surface normals with fine details as shown in Figure 8, emphasizing the potential applications of the proposed PPFT module to more polarization vision tasks.

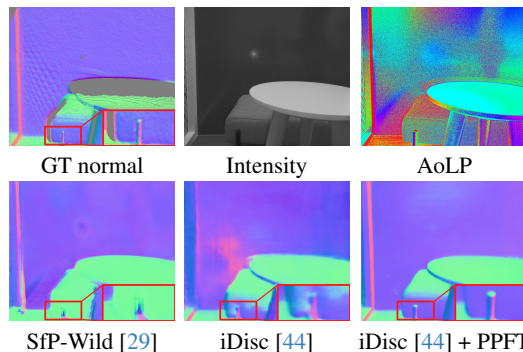


Figure 8. **Qualitative results of shape-from-polarization methods on the SPW dataset [29].** Our proposed PPFT can be combined with other polarization-based tasks to achieve better performance on limited polarization training data. Here SfP-Wild [29] is the state-of-the-art shape from the polarization approach. We use iDisc [44] as the backbone with the same polarization input and achieve better performance. In this example, our proposed PPFT with iDisc [44] produces significantly cleaner results, where the details are clearly visible as shown in the highlight region.

Model	Mean	Med	RMSE	11.5°	22.5°	30°
	Angular Error ↓ Percentage% ↑					
iDisc* [44]	43.56	38.46	51.92	16.61	35.48	43.56
SfP-Wild [29]	18.59	14.46	24.08	45.38	75.86	83.81
iDisc [44]	12.99	10.11	16.85	59.68	87.78	92.48
iDisc [44] + PPFT	10.91	8.92	14.03	67.82	92.12	95.77

Table 6. **Quantitative results on shape-from-polarization.** We achieve the best surface normal prediction performance by applying our PPFT strategy on iDisc [44] pre-trained on NYU Depth V2 [39]. For comparison, we pass polarization data to train iDisc with naive concatenation. *: this model is the RGB-based foundation model that is pre-trained on NYU-Depth V2 dataset [39].

5. Conclusion

In this work, we introduce a unified learning-based method that utilizes polarization information to improve sensor depth quality. To target the data scarcity problem with the existing polarization-depth dataset, we propose a Polarization Prompt Fusion Tuning (PPFT) strategy, which utilizes models pre-trained on large-scale RGB-D datasets. This allows both efficient use of pre-trained weights and effective fusion of polarization data. With PPFT, we achieve generalization to various depth-sensing approaches. Experiments validate substantial depth quality improvements for diverse depth inputs. Furthermore, we showcase the capability of applying the proposed method to other polarization-based vision tasks. As such, we hope that our work can contribute to the advancement of polarization-based depth enhancement and its potential applications in autonomous driving, 3D reconstruction, and beyond.

Acknowledgements This work was supported by the InnoHK program.

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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, pages 554–571. Springer, 2020. 3, 8
- [2] Kai Berger, Randolph Voorhies, and Larry H Matthies. Depth from stereo polarization in specular scenes for urban robotics. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 1966–1973. IEEE, 2017. 2
- [3] Chun-Yu Chai, Yu-Po Wu, and Shiao-Li Tsao. Deep depth fusion for black, transparent, reflective and texture-less objects. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6766–6772. IEEE, 2020. 1
- [4] Guangcheng Chen, Li He, Yisheng Guan, and Hong Zhang. Perspective phase angle model for polarimetric 3d reconstruction. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 398–414. Springer, 2022. 3
- [5] Xinjing Cheng, Peng Wang, and Ruigang Yang. Depth estimation via affinity learned with convolutional spatial propagation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 103–119, 2018. 3
- [6] Xinjing Cheng, Peng Wang, Chenye Guan, and Ruigang Yang. Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10615–10622, 2020. 3
- [7] 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. 2
- [8] Aniket Dashpute, Chandani Anand, and Mukul Sarkar. Depth resolution enhancement in time-of-flight cameras using polarization state of the reflected light. *IEEE Transactions on Instrumentation and Measurement*, 68(1):160–168, 2018. 2
- [9] Akshat Dave, Yongyi Zhao, and Ashok Veeraraghavan. Pandora: Polarization-aided neural decomposition of radiance. *arXiv preprint arXiv:2203.13458*, 2022. 3
- [10] 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
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [12] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 22(3):1341–1360, 2020. 8
- [13] Frederic Garcia, Djamila Aouada, Hashim Kemal Abdella, Thomas Solognac, Bruno Mirbach, and Björn Ottersten. Depth enhancement by fusion for passive and active sensing. In *European Conference on Computer Vision*, pages 506–515. Springer, 2012. 1
- [14] Mohit Gupta, Amit Agrawal, Ashok Veeraraghavan, and Srinivasa G Narasimhan. Structured light 3d scanning in the presence of global illumination. In *CVPR 2011*, pages 713–720. IEEE, 2011. 1
- [15] Mohit Gupta, Qi Yin, and Shree K Nayar. Structured light in sunlight. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 545–552, 2013. 1
- [16] Masaya Horio, Yu Feng, Tomoya Kokado, Taishi Takasawa, Keita Yasutomi, Shoji Kawahito, Takashi Komuro, Hajime Nagahara, and Keiichiro Kagawa. Resolving multi-path interference in compressive time-of-flight depth imaging with a multi-tap macro-pixel computational cmos image sensor. *Sensors*, 22(7):2442, 2022. 1
- [17] Tianyu Huang, Haoang Li, Kejing He, Congying Sui, Bin Li, and Yun-Hui Liu. Learning accurate 3d shape based on stereo polarimetric imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17287–17296, 2023. 3
- [18] Tomoki Ichikawa, Matthew Purri, Ryo Kawahara, Shohei Nobuhara, Kristin Dana, and Ko Nishino. Shape from sky: Polarimetric normal recovery under the sky. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14832–14841, 2021. 3
- [19] Daniel S Jeon, Andréas Meuleman, Seung-Hwan Baek, and Min H Kim. Polarimetric ifof: Measuring high-fidelity depth through scattering media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12353–12362, 2023. 2
- [20] HyunJun Jung, Nikolas Brasch, Aleš Leonardis, Nassir Navab, and Benjamin Busam. Wild tofu: Improving range and quality of indirect time-of-flight depth with rgb fusion in challenging environments. In *2021 International Conference on 3D Vision (3DV)*, pages 239–248. IEEE, 2021. 1
- [21] HyunJun Jung, Patrick Ruhkamp, Guangyao Zhai, Nikolas Brasch, Yitong Li, Yannick Verdie, Jifei Song, Yiren Zhou, Anil Armagan, Slobodan Ilic, Ales Leonardis, and Benjamin Busam. Is my depth ground-truth good enough? hammer – highly accurate multi-modal dataset for dense 3d scene regression, 2022. 2, 5, 6, 7
- [22] Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. Polarized 3d: High-quality depth sensing with polarization cues. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3370–3378, 2015. 1, 2, 5, 6
- [23] Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. Polarized 3d: Synthesis of polarization and depth cues for enhanced 3d sensing. In *SIGGRAPH 2015: Studio*, New York, NY, USA, 2015. Association for Computing Machinery.
- [24] Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. Depth sensing using geometrically constrained po-

- larization normals. *International Journal of Computer Vision*, 125:34–51, 2017. 2, 4
- [25] Agastya Kalra, Vage Taamazyan, Supreeth Krishna Rao, Kartik Venkataraman, Ramesh Raskar, and Achuta Kadambi. Deep polarization cues for transparent object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8602–8611, 2020. 1, 3
- [26] Yuhi Kondo, Taishi Ono, Legong Sun, Yasutaka Hirasawa, and Jun Murayama. 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*, pages 220–236. Springer, 2020. 3
- [27] Teppei Kurita, Yuhi Kondo, Legong Sun, and Yusuke Moriuchi. Simultaneous acquisition of high quality rgb image and polarization information using a sparse polarization sensor. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 178–188, 2023. 3
- [28] Chenyang Lei, Xuhua Huang, Mengdi Zhang, Qiong Yan, Wenxiu Sun, and Qifeng Chen. Polarized reflection removal with perfect alignment in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1750–1758, 2020. 3
- [29] 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. 1, 3, 8
- [30] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5
- [31] Zhuoxiao Li, Haiyang Jiang, Mingdeng Cao, and Yinqiang Zheng. Polarized color image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9873–9882, 2023. 3
- [32] Yupeng Liang, Ryosuke Wakaki, Shohei Nobuhara, and Ko Nishino. Multimodal material segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19800–19808, 2022. 3
- [33] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013. 4
- [34] Yuankai Lin, Tao Cheng, Qi Zhong, Wending Zhou, and Hua Yang. Dynamic spatial propagation network for depth completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1638–1646, 2022. 5, 6, 7
- [35] Lina Liu, Xibin Song, Xiaoyang Lyu, Junwei Diao, Mengmeng Wang, Yong Liu, and Liangjun Zhang. Fcfr-net: Feature fusion based coarse-to-fine residual learning for depth completion. In *proceedings of the AAAI Conference on Artificial Intelligence*, pages 2136–2144, 2021. 3
- [36] Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, and Jan Kautz. Learning affinity via spatial propagation networks, 2017. 3
- [37] Haiyang Mei, Bo Dong, Wen Dong, Jiayi Yang, Seung-Hwan Baek, Felix Heide, Pieter Peers, Xiaopeng Wei, and Xin Yang. Glass segmentation using intensity and spectral polarization cues. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12622–12631, 2022. 1, 3
- [38] Manasi Muglikar, Leonard Bauersfeld, Diederik Paul Moeys, and Davide Scaramuzza. Event-based shape from polarization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1547–1556, 2023. 3
- [39] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 2, 4, 7, 8
- [40] Danish Nazir, Marcus Liwicki, Didier Stricker, and Muhammad Zeshan Afzal. Semattnet: Towards attention-based semantic aware guided depth completion, 2022. 3, 5, 6
- [41] Taishi Ono, Yuhi Kondo, Legong Sun, Teppei Kurita, and Yusuke Moriuchi. Degree-of-linear-polarization-based color constancy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19740–19749, 2022. 3
- [42] Roy Or-El, Rom Hershkovitz, Aaron Wetzler, Guy Rosman, Alfred M Bruckstein, and Ron Kimmel. Real-time depth refinement for specular objects. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4378–4386, 2016. 1
- [43] Jinsun Park, Kyungdon Joo, Zhe Hu, Chi-Kuei Liu, and In So Kweon. Non-local spatial propagation network for depth completion. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 120–136. Springer, 2020. 3, 4
- [44] Luigi Piccinelli, Christos Sakaridis, and Fisher Yu. idisc: Internal discretization for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21477–21487, 2023. 8
- [45] Jiaxiong Qiu, Zhaopeng Cui, Yinda Zhang, Xingdi Zhang, Shuaicheng Liu, Bing Zeng, and Marc Pollefeys. DeepLidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3313–3322, 2019. 1, 3
- [46] Stefan Rahmann and Nikos Canterakis. Reconstruction of specular surfaces using polarization imaging. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, pages I–I. IEEE, 2001. 3
- [47] Mohamed Reda, Yongqiang Zhao, and Jonathan Cheung-Wai Chan. Polarization guided autoregressive model for depth recovery. *IEEE Photonics Journal*, 9(3):1–16, 2017. 2
- [48] Kyeongha Rho, Jinsung Ha, and Youngjung Kim. Guideformer: Transformers for image guided depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6250–6259, 2022. 3
- [49] Michael Schelling, Pedro Hermosilla, and Timo Ropinski. Weakly-supervised optical flow estimation for time-of-flight.

- In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2135–2144, 2023. [1](#)
- [50] Moein Shakeri, Shing Yang Loo, Hong Zhang, and Kangkang Hu. Polarimetric monocular dense mapping using relative deep depth prior. *IEEE Robotics and Automation Letters*, 6(3):4512–4519, 2021. [2](#)
- [51] Mingqi Shao, Chongkun Xia, Zhendong Yang, Junnan Huang, and Xueqian Wang. Transparent shape from a single view polarization image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9277–9286, 2023. [3](#)
- [52] 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 (ICCV)*, pages 3569–3579, 2023. [2](#)
- [53] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns. In *2017 international conference on 3D Vision (3DV)*, pages 11–20. IEEE, 2017. [2](#), [4](#)
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [3](#)
- [55] Yannick Verdié, Jifei Song, Barnabé Mas, Benjamin Busam, Aleš Leonardis, and Steven McDonagh. Cromo: Cross-modal learning for monocular depth estimation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3927–3937. IEEE, 2022. [2](#), [3](#)
- [56] Yufei Wang, Bo Li, Ge Zhang, Qi Liu, Tao Gao, and Yuchao Dai. Lrru: Long-short range recurrent updating networks for depth completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9422–9432, 2023. [3](#)
- [57] Patrick Wieschollek, Orazio Gallo, Jinwei Gu, and Jan Kautz. Separating reflection and transmission images in the wild. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 89–104, 2018. [3](#)
- [58] Lawrence B Wolff. Surface orientation from two camera stereo with polarizers. In *Optics, Illumination, and Image Sensing for Machine Vision IV*, pages 287–297. SPIE, 1990. [3](#)
- [59] Yan Xu, Xinge Zhu, Jianping Shi, Guofeng Zhang, Hujun Bao, and Hongsheng Li. Depth completion from sparse lidar data with depth-normal constraints, 2019. [1](#), [3](#)
- [60] Zheyuan Xu, Hongche Yin, and Jian Yao. Deformable spatial propagation networks for depth completion. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 913–917. IEEE, 2020. [3](#)
- [61] Hanrong Ye and Dan Xu. Taskprompter: Spatial-channel multi-task prompting for dense scene understanding. In *ICLR*, 2023. [5](#)
- [62] Tomonari Yoshida, Vladislav Golyanik, Oliver Wasenmüller, and Didier Stricker. Improving time-of-flight sensor for specular surfaces with shape from polarization. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 1558–1562. IEEE, 2018. [1](#), [2](#), [4](#)
- [63] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 175–185, 2018. [3](#)
- [64] Yixin Zhang, Xia Wang, Yuwei Zhao, and Yujie Fang. Time-of-flight imaging in fog using polarization phasor imaging. *Sensors*, 22(9):3159, 2022. [2](#)
- [65] Youmin Zhang, Xianda Guo, Matteo Poggi, Zheng Zhu, Guan Huang, and Stefano Mattoccia. Completionformer: Depth completion with convolutions and vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18527–18536, 2023. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [66] Dizhong Zhu and William AP Smith. Depth from a polarisation+ rgb stereo pair. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7586–7595, 2019. [2](#)
- [67] Jiawen Zhu, Simiao Lai, Xin Chen, Dong Wang, and Huchuan Lu. Visual prompt multi-modal tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9516–9526, 2023. [5](#)