

DEPTHOR: Depth Enhancement from a Practical Light-Weight dToF Sensor and RGB Image

Jijun Xiang¹, Xuan Zhu¹, Xianqi Wang¹,
Yu Wang², Hong Zhang², Fei Guo², Xin Yang^{1†}

¹ Huazhong University of Science and Technology ² Honor Device Co., Ltd

{jijunx, xuanzhu, xianqi, xinyang2014}@hust.edu.cn

{wangyu24, zhanghong70, guofei2}@honor.com




Figure 1. Depth predictions based on practical dToF data. The results from left to right are: a lightweight PENet [12], PENet with our training strategy, and our method. Our training strategy improves the performance of existing methods on real-world data. Our method, with the proposed training strategy and depth completion network, further refines predictions in challenging regions.

Abstract

Depth enhancement, which uses RGB images as guidance to convert raw signals from dToF into high-precision, dense depth maps, is a critical task in computer vision. Although existing super-resolution-based methods show promising results on public datasets, they often rely on idealized assumptions like accurate region correspondences and reliable dToF inputs, overlooking calibration errors that cause misalignment and anomaly signals inherent to dToF imaging, limiting real-world applicability. To address these challenges, we propose a novel completion-based method, named DEPTHOR, featuring advances in both the training strategy and model architecture. First, we propose a method to simulate real-world dToF data from the accurate ground truth in synthetic datasets to enable noise-robust training. Second, we design a novel network that incorporates monocular depth estimation (MDE), leveraging global depth relationships and contextual information to improve prediction in challenging regions. On the ZJU-L5

dataset, our training strategy significantly enhances depth completion models, achieving results comparable to depth super-resolution methods, while our model achieves state-of-the-art results, improving Rel and RMSE by 27% and 18%, respectively. On a more challenging set of dToF samples we collected, our method outperforms SOTA methods on preliminary stereo-based GT, improving Rel and RMSE by 23% and 22%, respectively. Our Code is available at <https://github.com/ShadowBbBb/Depthor>

1. Introduction

The direct Time-of-Flight (dToF) sensor is a depth sensor that measures depth by calculating the time it takes for emitted light pulses to reflect off objects and return to the sensor. With advantages like miniaturization and low power consumption, dToF sensors are widely deployed on mobile devices for applications like autofocus and obstacle detection

[†]Corresponding author.

[18, 33]. However, the depth information provided by dToF sensors is typically too coarse for high-precision tasks like 3D reconstruction [53, 54] and SLAM [8, 14, 19, 50, 52]. A common solution is using depth enhancement methods to produce dense, high-resolution depth maps from raw sensor data, using RGB images as guidance.

Based on the sensor’s data format, these methods fall into two categories: depth completion and depth super-resolution. In depth completion, the sensor generates a high-resolution depth map where valid depth points are sparsely distributed. The algorithm then propagates these sparse measurements to reconstruct dense depth through geometric reasoning guided by RGB context. Conversely, in depth super-resolution, the sensor returns a low-resolution dense depth map where each element corresponds to a local image region. The algorithm subsequently upsamples the depth map to match RGB resolution by recovering high-frequency details through cross-modal guidance.

Existing dToF enhancement approaches [9, 18, 36] are typically designed for depth super-resolution, which are generally based on two assumptions: (1) *ideal calibration exists between the RGB camera and dToF sensor*, and (2) *the dToF sensor operates reliably*. However, these assumptions often fall short in real-world environments. Through a systematic analysis of the RGB-dToF samples collected from a mobile phone, we found that calibration errors between devices are inevitable and may increase over time, leading to misalignment and conflicts in region correspondences, with a maximum deviation of up to two dToF pixels. In addition, the imaging principle of the dToF sensor results in signal loss or anomalous values in specific regions.

To address these challenges, we project dToF signals into a high-resolution sparse depth map using device parameters, redefining the problem within the scope of depth completion rather than depth super-resolution. The key motivation behind this projection is that, after transformation, both calibration errors and signal anomalies manifest as depth point inconsistencies at global or local scales. This reformulation allows us to focus solely on improving the robustness of the depth completion model against anomalous depth measurements, eliminating the need for complex region correspondences and thereby relaxing the restrictive assumptions of previous methods.

For this depth completion task with noisy input, we propose a noise-robust training strategy consisting of two key aspects: training dataset selection and dToF simulation, whose core idea is to simulate real-world dToF inputs during training while ensuring accurate supervision. While high-precision sensor-acquired datasets are commonly used in previous methods, their ground truth often lacks high-frequency details and exhibits similar anomaly patterns to dToF, limiting their effectiveness for training. In contrast, synthetic datasets offer more precise and detailed supervi-

sion, making them a preferable choice. To further improve alignment with real-world dToF characteristics, we introduce a dToF simulation method that accounts for four key aspects: overall distribution, specific abnormal regions, calibration errors, and random anomalies.

Although our training strategy effectively improves models’ prediction on real-world dToF data, correcting erroneous dToF signals or distinguishing general empty regions from signal loss remains an ill-posed problem for many methods, as they primarily focus on propagating existing depth measurements. Thus, we propose a simple yet effective depth completion network that integrates the monocular depth estimation (MDE) model, leveraging its global depth relationships and contextual information to improve predictions in challenging regions. Specifically, our model consists of two stages: multimodal fusion and refinement. First, we employ an encoder-decoder to extract and aggregate RGB and depth features along with the relative depth map, generating an coarse prediction. Then, we fuse the MDE feature with the decoder feature, compute mixed affinity, and further refine the initial depth map.

Our contributions are as follows:

- We conduct a comprehensive analysis of real-world dToF data and propose a noise-robust training strategy with a novel dToF simulation method on synthetic datasets.
- We design a depth completion network that effectively integrates MDE model at multiple stages to enhance predictions in ill-posed regions inherent to dToF imaging.
- On the ZJU-L5 dataset, our training strategy enhances depth completion models, achieving results comparable to depth super-resolution methods. Meanwhile, our model surpasses all types of state-of-the-art methods, improving Rel and RMSE by 27% and 18%, respectively.
- On a more challenging set of mobile phone-based dToF samples we collected, our model outperforms SOTA models on preliminary GT generated by stereo matching, improving Rel and RMSE by 23% and 22%.

2. Related Work

Depth Completion. Conventional 2D depth completion methods can be broadly categorized into encoder-decoder approaches [12, 55, 56] and affinity propagation approaches [5, 6, 41], with most evaluations conducted on real datasets [1, 31, 38]. Recent methods [30, 37, 47] utilize camera parameters to project depth points into 3D space, enhancing accuracy by leveraging spatial information. Additionally, some methods focus on aspects such as point sparsity [7, 10, 15, 25] and cross-dataset generalization [20, 24, 57] to further improve practicality.

However, assessing the applicability of existing methods to real-world dToF data remains challenging due to two key limitations. First, many methods rely on idealized data simulations. For instance, on the NYUv2 dataset, 500 perfectly

accurate depth points are randomly sampled from GT, causing evaluation metrics to focus on depth propagation rather than robustness to sensor noise in real-world scenarios. Second, the unique characteristics of dToF data are rarely considered. To the best of our knowledge, only methods from the MIPI competition [11, 34, 35] have attempted to simulate the uniform distribution of dToF data using grid sampling. Therefore, many designs that are considered effective in other depth modalities are not suitable for dToF.

Depth Super-resolution. For dToF data, Deltar [18] proposed a dual-branch depth super-resolution network that utilizes PointNet to extract dToF features and employs a transformer-based fusion module to integrate RGB and depth information. Building upon this, CFPNet [9] addressed the limited FoV coverage of dToF sensors by incorporating large convolution kernels and cross-attention mechanisms to enhance predictions in border regions. DVSR [36] specifically addresses depth super-resolution in dToF video sequences, using optical flow and deformable convolutions to aggregate multi-frame information, thereby enhancing prediction consistency.

The dToF simulation in these approaches typically begin by computing the depth histogram within a given region of the ground truth, followed by processing the histogram based on the characteristics of the target sensor (*e.g.* mean, peak, variance, rebin histograms). Among these, DVSR accounts for signal loss in low-reflectance regions by estimating the probability of missing depth measurements based on the mean RGB value.

However, existing methods typically rely on accurate RGB-dToF correspondence. For example, the ZJU-L5 dataset provides the coordinates of the RGB region corresponding to each dToF signal, and both Deltar and CFPNet leverage these coordinates to guide their feature aggregation modules. DVSR assumes that the dToF data is uniformly distributed, dividing a 480×640 image into 30×40 patches and directly simulating the dToF from depth GT corresponding to each patch. When real-world devices fail to provide accurate correspondences, the performance of these methods deteriorates significantly.

Monocular Depth Estimation. These methods predict the depth of each pixel in the input RGB image to obtain a dense depth map. Early methods trained on each dataset predict metric depth, but due to the inherent lack of depth scale information, these methods have poor generalization across different datasets. Ranftl *et al.* [27] introduced an affine-invariant loss, which predicts the inverse depth $1/d$, making the model focus more on relative distance relationships rather than absolute depth values. This mitigates the impact of scale shifts between different datasets and further enhances the model’s generalization.

Recent models [13, 16, 48, 49, 51] have significantly advanced this field with methods like pseudo-label generation,

diffusion model priors, and additional normal supervision. Among these, Ke *et al.* [16] achieves high detail in depth image outputs; however, its reliance on a diffusion model results in long inference times, which conflicts with many dToF application scenarios. Hu *et al.* [13], by introducing a normalized camera model and additional normal supervision, enables the model to output scaled depth with some generalization. Depth Anything [48, 49], on the other hand, outputs inverse depth and, through techniques like pseudo-label generation and teacher-student models, provides both generalization and detail performance. Since dToF sensors already offer depth scale information but lack detail, we select Depth Anything V2 [49] as the pre-trained monocular depth estimation model for our subsequent experiments.

3. The Proposed Method

3.1. Training Strategy with dToF Data Simulation

We collected a set of RGB-dToF samples using an Honor Magic6 Ultra to analyze the distribution characteristics and potential anomalies of real-world dToF data, with resolutions of 912×684 and 40×30 , respectively. Readers interested in dToF imaging are referred supplementary material for more details. Figure 2a shows an RGB-dToF sample obtained under ideal conditions.




Figure 2. Ideal and anomalous samples in real-world dToF data.

Using the intrinsic and extrinsic parameters of both the dToF sensor and the RGB camera, along with calibration transformation matrices, we project dToF data into a high-resolution, uniformly distributed sparse depth map. For this depth map, we design our simulation method based on the following four key aspects:

Overall Distribution. Due to the substantial difference in FoV between dToF sensor and RGB camera, the captured depth data does not cover the entire RGB image. Moreover, the high-resolution sparse depth points are inherently imprecise, as they theoretically correspond to peak values

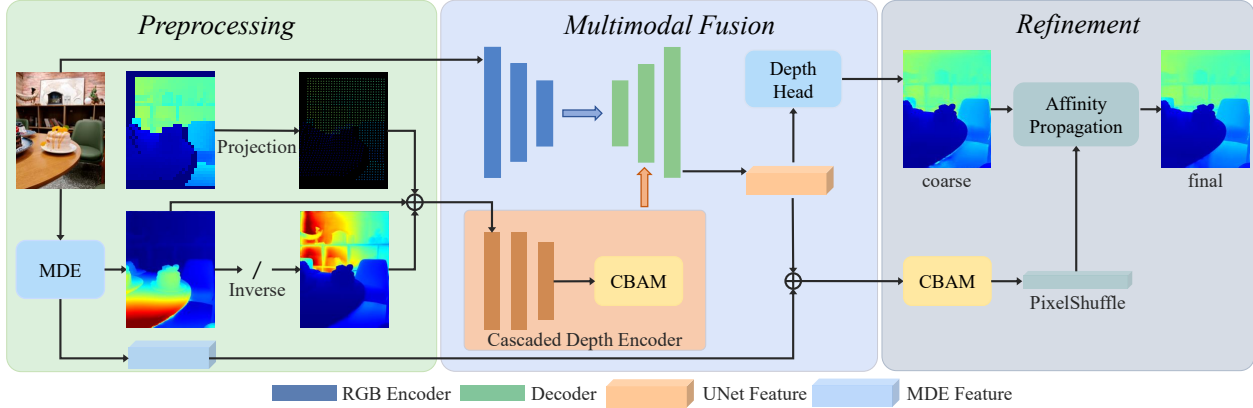


Figure 3. **Overview of our model:** We first project dToF signals into sparse depth points, and use a pre-trained MDE model to generate inverse and relative depth maps. In multimodal fusion, we employ a simple encoder-decoder structure to obtain a coarse estimation. In refinement, we update the initial depth map using mixed affinity propagation.

within a defined iFoV (Fig. 2h). Thus, we perform random translations and rotations on the depth GT within the iFoV range, followed by approximately uniform sampling within the roughly defined FoV.

Specific Abnormal Regions. We also address abnormal conditions in specific areas inherent to dToF imaging.

1. *Non-Lambertian surfaces* (Fig. 2b, Fig. 2c): Photons may pass through objects, leading to signal loss or returning depth values from a farther distance. We use diffuse reflection intensity to identify non-Lambertian regions and randomly determine the type of anomaly.
2. *Low-reflectivity areas* (Fig. 2e): In low-light environments or on dark surfaces, photons are absorbed rather than reflected, leading to signal loss. We convert RGB to HSV space and assign a probability of signal loss to points with low brightness values in the V channel.
3. *Long-distance regions* (Fig. 2f): At greater distances, photons are more susceptible to environmental noise and may be lost entirely if they exceed the device’s maximum reception time. The theoretical maximum detection range of our dToF sensor is 8.1 meters, but signals beyond 6 meters are frequently lost in practical use.

Random Anomalies. To enhance the model’s robustness to random anomalies (Fig. 2g), we introduced approximately 5% noise points and 5% blank points. The depth values of the noise points were randomly assigned within the theoretical detection range.

Calibration Errors. Calibration errors manifest as regional shifts after projection. Specifically, foreground points generally project with high precision, while background points often experience a noticeable shift (Fig. 2d). Therefore, we select a percentile from GT as the threshold, treat points above it as background, and apply a random shift (within 0–2 dToF pixels).

Using this simulation method, we obtain training dToF

data from synthetic datasets that closely resemble real-world distributions. Supervising this input with accurate ground truth encourages the model to learn robust features under imperfect data conditions, thereby better adapting to real-world conditions and partially mitigating issues inherent to dToF imaging, as shown in Fig. 1.

3.2. Depth Completion Model Integrating MDE

We formulate the problem as: given a projected sparse depth map S , the corresponding RGB image I , and the inverse depth map D_{inv} and features F_{mde} output by the MDE model (based on I), the goal of the depth completion model is to predict a dense depth map D .

As shown in Fig. 3, our proposed network consists of two stages: multi-modal fusion and refinement. The first stage outputs an initial depth map at half resolution, while the second stage refines it to produce the precise full-resolution depth prediction using affinity propagation.

Multimodal Fusion. In this stage, we implement an encoder-decoder network. The encoder extracts multi-resolution features from image and depth separately, which are subsequently fused in the decoder. The fused feature F_{unet} is then passed through a depth head to produce an initial depth estimation.

We employed the network from BpNet [37] as the RGB encoder, progressively downsampling the RGB image and generating feature maps F_{img} at resolutions ranging from $1/2$ to $1/32$. We modified its architecture and feature dimensions to reduce computational cost and parameters.

Since distant regions in D_{inv} are numerically compressed toward zero, we introduce the relative depth map $D_{rel} = 1/(D_{inv} + \epsilon)$ to emphasize structural details in these areas. The inputs of depth encoder include $\{D_{rel}, D_{inv}, S\}$, while D_{rel} and D_{inv} are normalized, S remains unnormalized to retain absolute scale information. With these simple design choices, the depth encoder maintains a stable bal-

ance between near and far regions, as well as between relative and absolute depth.

To effectively extract depth features F_{dep} , we first apply a combination of convolution layers, including large-kernel dilated convolution to enhance the perception of scale information in S and small-kernel downsampling convolution to capture high-frequency details in D_{rel} and D_{inv} . Then, we feed the output feature into a CBAM module [43], where spatial and channel attention mechanisms are employed for aggregation.

In the decoder, we progressively fuse RGB and depth features through convolution and upsampling layers, ultimately producing the decoder feature F_{unet} .

Following the design of [2, 18], we input F_{unet} into a depth head. Rather than directly regressing depth values, the depth head generates a set of N non-uniformly normalized depth bins b for each image, along with weighting coefficients k_i for each pixel corresponding to b_i . After restoring the depth bins to metric depth using hyperparameters and computing each bin’s center c_i , the initial depth is computed using the following formula:

$$d = \sum_{i=1}^N k_i c_i \quad (1)$$

Since computing the depth map requires generating N -dimensional features for each pixel, we set N to 128 to balance computational cost and accuracy. Additionally, the initial depth map is predicted at half resolution.

Refinement. The initial depth map often contains various anomalies, such as artifacts in regions without depth measurement coverage and residual erroneous signals in certain areas, which is particularly pronounced when the MDE model produces sharp numerical changes. To address this, we deploy an affinity propagation module, based on CSPN++[6], to further refine the initial depth map.

Unlike previous methods that compute affinity using single-modality features from the decoder, we jointly compute affinity, since the rich semantic information in F_{mde} helps correct errors in F_{unet} caused by inaccurate depth signals. Conversely, incorporating F_{unet} mitigates the resolution discrepancies introduced by the Transformer architecture and the lack of scale information in F_{mde} .

We begin by interpolating F_{mde} to align with the half-resolution of F_{unet} . We then concatenate F_{unet} and F_{mde} , feeding this combined feature into a CBAM module and a PixelShuffle [29] layer to upsample to the full resolution. Using this merged feature F_{cspn} , we calculate mixed affinity ω_k corresponding to kernel [3, 5, 7].

During the propagation, the update process of pixel i under affinity kernel k at the t -th iteration is formulated in (2).

$$\hat{D}_{i,k,t} = \omega_{i,k} \hat{D}_{i,t-1} + \sum_{j \in \mathbb{N}_k(i)} \omega_{j,k} \hat{D}_{j,t-1} \quad (2)$$

Following BPNet[37], we aggregate the outputs across different iterations and affinity kernels using two normalized weights produced by a convolution and softmax layer, as described in (3), where $t \in \{0, T/2, T\}$.

$$D = \sum_{t \in T} \tau_t \sum_{k \in \mathcal{K}} \sigma_k \hat{D}_{k,t} \quad (3)$$

Conventional affinity propagation modules typically embed sparse depth measurements at each iteration, directly assigning the original sparse depth values to the updated depth map. However, since dToF points are not entirely accurate, we remove this setting.

3.3. Implementation Details

We employ a scaled affine-invariant loss for supervision following [18], the expression of which is as follows:

$$L = \alpha \sqrt{\frac{1}{T} \sum_i g_i^2 - \frac{\lambda}{T^2} \left(\sum_i g_i \right)^2} \quad (4)$$

Where $g_i = \log \tilde{d}_i - \log d_i$, \tilde{d}_i , d_i , represent the predicted values and ground truth for valid pixel points, respectively, and in all experiments, $\alpha=10$, $\lambda=0.85$. We calculate the loss only for pixels within the sensor’s theoretical detection range. The detailed training settings are provided in the supplementary material.

4. Experiments

4.1. Datasets & Evaluation Metrics

Hypersim dataset for training. We trained our model on the Hypersim [28] dataset, as it provides precise ground truth and extensive labels, with 59,544 frames for training and 7,386 frames for testing. For different dToF sensors, we modified the simulation method to ensure a similar distribution between the training and testing data. Please refer to the supplementary material for more details.

ZJU-L5 dataset for testing. Deltar[18] employs the ST VL53L5CX (L5) and the Intel RealSense 435i to capture raw dToF data and ground truth. The dToF and image resolutions are 8×8 and 480×640 , respectively. In our method, we utilize the provided iFoV to convert each dToF signal into a sparse depth point at the center of its corresponding region, without using the variance information.

Real-world samples we collected for testing. We collected RGB-dToF samples on an Honor Magic6 Ultra and projected dToF using its internal parameters. We use the results

from an SOTA stereo matching method [4] as a preliminary ground truth for evaluation. The detailed pipeline is introduced in supplementary material.

We adopt a high-quality stereo matching pipeline using the main and ultra-wide cameras on the mobile phone to get ground truth, manually filter failed samples, and mask noisy regions. Lens distortion and baseline mismatch may slightly affect the epipolar geometry, leading to a global shift. While not perfect, we believe the SOTA stereo methods [4, 42, 45, 46] to be more effective than common sensors, especially in complex regions targeted in our work (Fig. 4), the results in the main paper are from MonSter [4].




Figure 4. Ground truth from various sensors and stereo methods

Evaluation Metrics We reported standard metrics including δ_i , Rel, RMSE, \log_{10} . To further evaluate performance at boundaries, we also reported edge-weighted mean absolute error (EWMAE) [21, 34], which assigns greater weight to pixels with larger gradients when calculating MAE. The details are introduced in the supplementary material.

4.2. Effectiveness of Our Training Strategy.

We trained our model and a lightweight PENet (denoted as PENet*) using different simulation methods on Hypersim and evaluated on the ZJU-L5. For PENet*, we retain the original design but reduce the number of layers and channels to accelerate training. As a result, the parameters and FLOPs are reduced from 131M / 592G to 48M / 110G.

As shown in Tab. 1, our method significantly improves the performance of both models. Notably, PENet* even outperforms the SOTA super-resolution method CFPNet on several metrics, demonstrating that our simulation strategy effectively narrows the gap between depth completion and super-resolution, without relying on precise dToF-RGB region correspondences. We further analyze the impact of training datasets in Sec. 4.4.

Model	Simulation	δ_1	δ_2	Rel	RMSE	\log_{10}
CFPNet	Deltar	0.883	0.949	0.103	0.431	0.047
PENet	Deltar	0.807	0.914	0.161	0.498	0.065
PENet*	Deltar	0.815	-	0.152	0.510	-
PENet*	MIPI	0.865	0.929	0.118	0.493	0.061
PENet*	Ours	0.889	0.949	0.093	0.447	0.046
Ours	Deltar	0.804	0.883	0.164	0.562	0.097
Ours	MIPI	0.853	0.909	0.123	0.511	0.089
Ours	Ours	0.933	0.972	0.075	0.350	0.034

Table 1. Performance under different simulation methods. The results in the second row are reported by Deltar [18].

4.3. Comparison with SOTAs

In addition to referencing results from published papers, we conducted additional experiments to ensure a comprehensive comparison. For monocular depth estimation, we separately evaluated the Metric-Large (fine-tuned on Hypersim) and Relative-Small (used in our method) versions of Depth Anything V2 (indicated as DAV2 -ML and DAV2 -RS). The former was tested directly, while for the latter, we fitted its output using dToF measurements. For depth completion, we conducted two types of experiments: (1) retraining existing methods with our strategy, including the PENet*, the SOTA 2D (CompletionFormer) and 3D (BPNet) method on conventional benchmark; (2) testing the latest generalizable model OMNI-DC, which incorporates multiple depth modalities and varying levels of sparsity during training.

Method	Type	Pub	δ_1	δ_2	Rel	RMSE	\log_{10}
BTS[17]	MDE	arXiv19	0.739	0.914	0.174	0.523	0.079
AdaBins[2]	MDE	CVPR21	0.770	0.926	0.160	0.494	0.073
PnP-Depth[39]	DS	ICRA19	0.805	0.904	0.144	0.560	0.068
PrDepth[44]	DS	CVPR20	0.800	0.926	0.151	0.460	0.063
PENet[12]	DC	ICRA21	0.807	0.914	0.161	0.498	0.065
Deltar[12]	DS	ECCV22	0.853	0.941	0.123	0.436	0.051
CFPNet[9]	DS	3DV25	0.883	0.949	0.103	0.431	0.047
PENet*[12]	DC	ICRA21	0.889	0.949	0.093	0.447	0.046
BPNet[37]	DC	CVPR24	-	-	-	0.671	-
CFormer[55]	DC	CVPR23	0.873	0.938	0.103	0.480	0.053
OMNI-DC[57]	DC	arXiv24	0.871	0.933	0.099	0.502	0.053
DAv2 -ML [49]	MDE	Neurips24	0.703	0.905	0.220	0.467	0.083
DAv2 -RS [49]	MDE	Neurips24	0.869	0.937	0.109	0.480	0.063
Ours-Small	DC	-	0.921	0.963	0.080	0.379	0.038
Ours-Large	DC	-	0.933	0.972	0.075	0.350	0.034

Table 2. Quantitative comparison on ZJU-L5. MDE: Monocular Depth Estimation; DS: Depth Super-resolution; DC: Depth Completion. The best and second best are marked with colors, while the best result among existing methods is underlined. CFPNet is the published SOTA method focusing on this dataset.

Method	Deltar	CFPNet	CFormer	OMNI-DC	DAv2-L	DAv2-S	Ours-S	Ours-L
Params (M)	18	20	81	84	335	24	6 + 24	12 + 24
FLOPs (G)	42	46	380	398	674	47	26 + 13	64 + 47

Table 3. Complexity comparison. We separately list the depth completion and the MDE model in our method, with FLOPs calculated at a resolution of 480×640.

We found that due to changes in depth pattern and underlying anomalies, many models and designs that are effective in traditional benchmarks are not well-suited for real-world dToF data. First, since depth points are roughly uniformly distributed yet extremely sparse (0.02% in ZJU-L5 and 0.18% in our data), 3D methods struggle to capture effective spatial interactions, and architectures sensitive to sparsity also suffer from performance degradation. Second, some designs that assume absolute accuracy of depth sensors are sensitive to real-world noise, focusing solely

on preserving and rapidly propagating the sparse measurement. Examples include the point embedding operation in the affinity propagation module and residual connections with the initial sparse depth map.




Figure 5. **Qualitative results on ZJU-L5**, our model further improve anomalies present in the ground truth

In addition, we believe that the metrics from real datasets primarily reflect how well depth enhancement models improve low-cost sensors, bringing them closer to high-precision ones (specifically, L5 and RealScene 435i). However, even the depth sensors used to collect ground truth still exhibit limited performance in the challenging regions targeted in our study. Thus, these metrics may not fully capture the performance of our method. As shown in Fig. 5, our model’s predictions not only outperform existing methods but also further improve anomalies present in the ground truth, which leads to a decrease in metrics. More qualitative results are provided in the supplementary material.

Results on our real-world data. For this type of data, we employed the full dToF simulation method described in Sec. 3.1 and maintained the same evaluation settings. Due to the higher image resolution, we modified some methods to accelerate training.

Model	δ_1	δ_2	RMSE	Rel	EWMAE
BPNet[37]	-	-	0.630	-	-
OMNI-DC[57]	0.593	0.768	0.643	0.292	0.195
DAv2 -RS[49]	0.687	0.833	<u>0.292</u>	0.237	0.141
PENet*[12]	<u>0.740</u>	0.878	0.327	<u>0.202</u>	<u>0.139</u>
CFFormer*[55]	0.732	<u>0.883</u>	0.320	0.206	0.159
Ours	0.790	0.911	0.226	0.155	0.108

* indicates a lightweight version.

Table 4. Quantitative comparison on our real-world samples.

Table 4 presents the quantitative results based on ground truth obtained from stereo matching. Despite some scale

shifts and anomalies in the GT, these metrics still offer a meaningful preliminary evaluation of the model’s performance. Notably, our model achieves the best results.

Figure 6 presents the qualitative comparison. Our method effectively integrates the MDE model, improving performance in fine details and challenging regions. Additionally, we observed that PENet achieves better detail prediction than CFormer, which is consistent with the EWMAE metric in Table 4. More qualitative results are provided in the supplementary material.




Figure 6. Qualitative results on our real-world samples.

4.4. Ablation Studies

Components of Simulation Method. We performed ablation studies on each component of our simulation method using the ZJU-L5 dataset, as shown in Tab. 5. Since calibration errors are not considered, we validate their impact through qualitative results on real-world samples, provided in the supplementary material.

Method	δ_1	δ_2	Rel	RMSE	\log_{10}
Standard	0.933	0.972	0.075	0.350	0.034
w/o RA	0.923	0.966	0.076	0.362	0.037
w/o OD	0.912	0.965	0.091	0.381	0.043
w/o SR	0.773	0.847	0.175	0.566	0.138
w/o (RA + OD)	0.905	0.965	0.102	0.395	0.045
w/o (OD + SR)	0.780	0.855	0.191	0.641	0.168

Table 5. **Ablation studies about simulation method.** OD: Over-all Distribution, SR: Specific Region, RA: Random Anomalies.

Depthor with Different MDE Models. As shown in Tab. 6, we replaced different MDE models in our method. We found that using more powerful MDE models does not significantly improve performance on ZJU-L5 compared to Hypersim, particularly in EWMAE, which further highlights the potential issues in the GT of real datasets.

Model	ZJU-L5			Hypersim		
	RMSE	Rel	EWMAE	RMSE	Rel	EWMAE
DAv2 -RS	0.350	0.075	0.136	0.445	0.068	0.110
DAv2 -RB	0.335	0.071	0.136	0.406	0.061	0.103
DAv2 -RL	0.330	0.070	0.135	0.390	0.059	0.101
DAv2 -MS	0.372	0.095	0.141	0.554	0.102	0.114

Table 6. **Ablation studies on different MDE models.** The results on Hypersim are based on ZJU-L5’s dToF simulation. R: relative; M: metric; S: small; B: base; L: large.

Refinement of Mixed Affinity Propagation. We analyze this module through quantitative metrics from synthetic datasets and qualitative results of real-world samples. As shown in Tab. 7, on the Hypersim dataset, refining the initial depth map in full resolution using affinity propagation improves the model’s overall performance, particularly on boundary-focused metric EWMAE. The qualitative results in Fig. 7 also reveal that this module effectively improves the model’s performance in regions beyond the sensor’s FoV and at foreground-background boundaries, mitigating anomalies while enhancing prediction consistency.

Refine	Input Feature		Hypersim			Params. (M)
	MDE	UNet	RMSE	REL	EWMAE	
/	/	/	0.267	0.039	0.091	-
✓	✓	/	0.269	0.039	0.087	+0.048
✓	/	✓	0.258	0.037	0.081	+0.085
✓	✓	✓	0.248	0.034	0.079	+0.122
+ Point Embedding			0.328	0.038	0.098	+0.122

Table 7. **Ablation studies about refinement.** The results on Hypersim are based on our samples’ dToF simulation

Furthermore, our experimental results demonstrate that due to the lack of scale information and the resolution differences, computing affinity solely based on F_{mde} improves EWMAE but adversely affects scale metrics. However, the contextual information in F_{mde} can still be leveraged to enhance F_{unet} . Additionally, the regional characteristics and anomalies of dToF signals conflict with the assumptions of point embedding, leading to performance degradation.

Complementarity of Our Training Strategy and Model. In Fig. 8, we present our model’s predictions on the NYUv2 dataset under different training strategies, where dToF points are sampled from the inaccurate ground truth: (a) Trained on the NYUv2 dataset, model tends to disregard MDE outputs due to conflicts with the inaccurate ground truth; (b) Trained on the Hypersim dataset without our simulation method, model extracts only contextual information from MDE, to propagate accurate depth points




Figure 7. Refinement of mixed affinity propagation.

while neglecting global depth relationships; (c) Our training strategy enhances performance through both global relationships and local details.




Figure 8. Prediction of our model under different combinations of training datasets and simulation methods.

5. Conclusion

In this paper, we present a comprehensive solution to real-world dToF enhancement. Unlike previous super-resolution methods, we reformulate the problem within completion to streamline the task. We introduce a noise-robust training strategy that improves existing depth completion models, achieving results comparable to depth super-resolution methods. Additionally, we design a novel network that effectively integrates MDE to enhance predictions in challenging regions. Our method achieves state-of-the-art results on both the ZJU-L5 dataset and a challenging set of dToF samples. Not fast enough for real-time inference is the main limitation of our method. Extending our method to other sensors is a promising direction for future research.

Acknowledgement. This research is supported by the National Key R&D Program of China (2024YFE0217700), National Natural Science Foundation of China (62472184), the Fundamental Research Funds for the Central Universities, and the Innovation Project of Optics Valley Laboratory (Grant No. OVL2025YZ005)

References

- [1] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *arXiv preprint arXiv:2111.08897*, 2021. 2
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4009–4018, 2021. 5, 6
- [3] E Charbon. Single-photon imaging in complementary metal oxide semiconductor processes. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2012):20130100, 2014. 1
- [4] Junda Cheng, Longliang Liu, Gangwei Xu, Xianqi Wang, Zhaoxing Zhang, Yong Deng, Jinliang Zang, Yurui Chen, Zhipeng Cai, and Xin Yang. Monster: Marry monodepth to stereo unleashes power. *arXiv preprint arXiv:2501.08643*, 2025. 6
- [5] Xinjing Cheng, Peng Wang, and Ruigang Yang. Learning depth with convolutional spatial propagation network. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2361–2379, 2019. 2
- [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. 2, 5
- [7] Andrea Conti, Matteo Poggi, and Stefano Mattoccia. Sparsity agnostic depth completion. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 5871–5880, 2023. 2
- [8] Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (ToG)*, 36(4): 1, 2017. 2
- [9] Laiyan Ding, Hualie Jiang, Rui Xu, and Rui Huang. Cfpnet: Improving lightweight tof depth completion via cross-zone feature propagation. *arXiv preprint arXiv:2411.04480*, 2024. 2, 3, 6
- [10] Jakub Gregorek and Lazaros Nalpantidis. Steeredmarigold: Steering diffusion towards depth completion of largely incomplete depth maps. *arXiv preprint arXiv:2409.10202*, 2024. 2
- [11] Dewang Hou, Yuanyuan Du, Kai Zhao, and Yang Zhao. Learning an Efficient Multimodal Depth Completion Model. *arXiv preprint arXiv:2208.10771*, 2022. 3
- [12] Mu Hu, Shuling Wang, Bin Li, Shiyu Ning, Li Fan, and Xiaojin Gong. Penet: Towards precise and efficient image guided depth completion. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13656–13662. IEEE, 2021. 1, 2, 6, 7, 3
- [13] 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. *arXiv preprint arXiv:2404.15506*, 2024. 3
- [14] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, et al. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on user interface software and technology*, pages 559–568, 2011. 2
- [15] Jinyoung Jun, Jae-Han Lee, and Chang-Su Kim. Masked spatial propagation network for sparsity-adaptive depth refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19768–19778, 2024. 2
- [16] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. 3
- [17] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019. 6
- [18] Yijin Li, Xinyang Liu, Wenqi Dong, Han Zhou, Hujun Bao, Guofeng Zhang, Yinda Zhang, and Zhaopeng Cui. Deltar: Depth estimation from a light-weight tof sensor and rgb image. In *European conference on computer vision*, pages 619–636. Springer, 2022. 2, 3, 5, 6
- [19] Xinyang Liu, Yijin Li, Yanbin Teng, Hujun Bao, Guofeng Zhang, Yinda Zhang, and Zhaopeng Cui. Multi-modal neural radiance field for monocular dense slam with a light-weight tof sensor. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1–11, 2023. 2
- [20] Zhiheng Liu, Ka Leong Cheng, Qiuyu Wang, Shuzhe Wang, Hao Ouyang, Bin Tan, Kai Zhu, Yujun Shen, Qifeng Chen, and Ping Luo. Depthlab: From partial to complete. *arXiv preprint arXiv:2412.18153*, 2024. 2
- [21] Javier López-Randulfe, César Veiga, Juan J Rodríguez-Andina, and José Farina. A quantitative method for selecting denoising filters, based on a new edge-sensitive metric. In *2017 IEEE International Conference on Industrial Technology (ICIT)*, pages 974–979. IEEE, 2017. 6, 1
- [22] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 1
- [23] Desmond O’Connor. *Time-correlated single photon counting*. Academic press, 2012. 1
- [24] Hyoungseob Park, Anjali Gupta, and Alex Wong. Test-time adaptation for depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20519–20529, 2024. 2
- [25] Jinhung Park, Yu-Jhe Li, and Kris Kitani. Flexible depth completion for sparse and varying point densities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21540–21550, 2024. 2

- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 1
- [27] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020. 3
- [28] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021. 5, 2
- [29] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016. 5
- [30] Yunxiao Shi, Manish Kumar Singh, Hong Cai, and Fatih Porikli. Decotr: Enhancing depth completion with 2d and 3d attentions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10736–10746, 2024. 2
- [31] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*, pages 746–760. Springer, 2012. 2
- [32] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, pages 369–386. SPIE, 2019. 1
- [33] STMicroelectronics. STMicroelectronics Ships 1 Billionth Time-of-Flight Module. https://www.st.com/content/st_com/en/about/media-center/pressitem.html/t4210.html, 2022. Accessed: 19-Jul-2022. 2
- [34] Qianhui Sun, Qingyu Yang, Chongyi Li, Shangchen Zhou, Ruicheng Feng, Yuekun Dai, Wenxiu Sun, Qingpeng Zhu, Chen Change Loy, Jinwei Gu, et al. Mipi 2023 challenge on rgbw remosaic: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2878–2885, 2023. 3, 6, 1
- [35] Wenxiu Sun, Qingpeng Zhu, Chongyi Li, Ruicheng Feng, Shangchen Zhou, Jun Jiang, Qingyu Yang, Chen Change Loy, Jinwei Gu, Dewang Hou, et al. Mipi 2022 challenge on rgb+ tof depth completion: Dataset and report. In *European Conference on Computer Vision*, pages 3–20. Springer, 2022. 3, 1
- [36] Zhanghao Sun, Wei Ye, Jinhui Xiong, Gyeongmin Choe, Jialiang Wang, Shuo Chen Su, and Rakesh Ranjan. Consistent direct time-of-flight video depth super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5075–5085, 2023. 2, 3, 1
- [37] Jie Tang, Fei-Peng Tian, Boshi An, Jian Li, and Ping Tan. Bilateral propagation network for depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9763–9772, 2024. 2, 4, 5, 6, 7
- [38] 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
- [39] Tsun-Hsuan Wang, Fu-En Wang, Juan-Ting Lin, Yi-Hsuan Tsai, Wei-Chen Chiu, and Min Sun. Plug-and-play: Improve depth prediction via sparse data propagation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5880–5886. IEEE, 2019. 6
- [40] Wenshan Wang, DeLong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. Tartanair: A dataset to push the limits of visual slam. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4909–4916. IEEE, 2020. 2
- [41] 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. 2
- [42] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. Foundationstereo: Zero-shot stereo matching. *arXiv preprint arXiv:2501.09898*, 2025. 6
- [43] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 5
- [44] Zhihao Xia, Patrick Sullivan, and Ayan Chakrabarti. Generating and exploiting probabilistic monocular depth estimates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 65–74, 2020. 6
- [45] Gangwei Xu, Yun Wang, Junda Cheng, Jinhui Tang, and Xin Yang. Accurate and efficient stereo matching via attention concatenation volume. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4):2461–2474, 2023. 6
- [46] Gangwei Xu, Xianqi Wang, Zhaoxing Zhang, Junda Cheng, Chunyuan Liao, and Xin Yang. Igev++: iterative multi-range geometry encoding volumes for stereo matching. *arXiv preprint arXiv:2409.00638*, 2024. 6
- [47] Zhiqiang Yan, Yuankai Lin, Kun Wang, Yupeng Zheng, Yufei Wang, Zhenyu Zhang, Jun Li, and Jian Yang. Tri-perspective view decomposition for geometry-aware depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4874–4884, 2024. 2
- [48] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024. [3](#)
- [49] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. [3](#), [6](#), [7](#)
- [50] Xin Yang, Jingyu Chen, Yuanjie Dang, Hongcheng Luo, Yuesheng Tang, Chunyuan Liao, Peng Chen, and Kwang-Ting Cheng. Fast depth prediction and obstacle avoidance on a monocular drone using probabilistic convolutional neural network. *IEEE Transactions on Intelligent Transportation Systems*, 22(1):156–167, 2019. [2](#)
- [51] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023. [3](#)
- [52] Zikang Yuan, Qingjie Wang, Ken Cheng, Tianyu Hao, and Xin Yang. Sdv-loam: Semi-direct visual–lidar odometry and mapping. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):11203–11220, 2023. [2](#)
- [53] Zikang Yuan, Jie Deng, Ruiye Ming, Fengtian Lang, and Xin Yang. Sr-livo: Lidar-inertial-visual odometry and mapping with sweep reconstruction. *IEEE Robotics and Automation Letters*, 2024. [2](#)
- [54] Zikang Yuan, Fengtian Lang, Tianle Xu, and Xin Yang. Sr-livo: Lidar-inertial odometry with sweep reconstruction. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7862–7869. IEEE, 2024. [2](#)
- [55] 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](#), [6](#), [7](#)
- [56] Shanshan Zhao, Mingming Gong, Huan Fu, and Dacheng Tao. Adaptive context-aware multi-modal network for depth completion. *IEEE Transactions on Image Processing*, 30: 5264–5276, 2021. [2](#)
- [57] Yiming Zuo, Willow Yang, Zeyu Ma, and Jia Deng. Omnidc: Highly robust depth completion with multiresolution depth integration. *arXiv preprint arXiv:2411.19278*, 2024. [2](#), [6](#), [7](#)