DEPTH ANYTHING WITH ANY PRIOR

Anonymous authors

000

001 002 003

004

006

008 009

010

011

012

013

014

015

016

017

018

019

021

024

025

027 028 029

031

033

034

037

040

041

042

043

044

045

046

047

048

051

052

Paper under double-blind review

ABSTRACT

This work presents *Prior Depth Anything*, a framework that combines incomplete but precise metric information in depth measurement with relative but complete geometric structures in depth prediction, generating accurate, dense, and detailed metric depth maps for any scene. To this end, we design a coarse-to-fine pipeline to progressively integrate the two complementary depth sources. First, we introduce pixel-level metric alignment and distance-aware weighting to pre-fill diverse metric priors by explicitly using depth prediction. It effectively narrows the domain gap between prior patterns, enhancing generalization across varying scenarios. Second, we develop a conditioned monocular depth estimation (MDE) model to refine the inherent noise of depth priors. By conditioning on the normalized pre-filled prior and prediction, the model further implicitly merges the two complementary depth sources. Our model showcases impressive zero-shot generalization across depth completion, super-resolution, and inpainting over 7 real-world datasets, matching or even surpassing previous task-specific methods. More importantly, it performs well on challenging, unseen mixed priors and enables test-time improvements by switching prediction models, providing a flexible accuracy-efficiency trade-off while evolving with advancements in MDE models.

1 Introduction

Fine-detailed and dense metric depth information is a fundamental pursuit in computer vision and robotics applications (Zhang et al., 2023a; Deng et al., 2022; Chung et al., 2024; Roessle et al., 2022; Esser et al., 2023; Wang et al., 2024; Zhen et al., 2024; Wang et al., 2019; Wofk et al., 2019). Although monocular depth estimation (MDE) models (Ranftl et al., 2020; Yang et al., 2024a;b; Ke et al., 2024; Bochkovskii et al., 2025; He et al., 2025; Hu et al., 2024; Yin et al., 2023) have made significant progress, enabling complete and detailed depth predictions, predicted

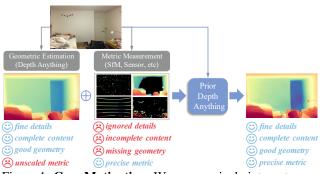


Figure 1: **Core Motivation**. We progressively integrate complementary information from metric measurements (accurate metrics) and relative predictions (completeness and fine details) to produce dense and fine-grained metric depth maps.

depths are relative and lack precise metric information. On the other hand, depth measurement technologies, such as Structure from Motion (SfM) (Schonberger & Frahm, 2016) or depth sensors (Lange & Seitz, 2001), provide precise but often incomplete and coarse metric information.

In this paper, we explore *prior-based monocular depth estimation*, which takes RGB images and measured depth priors as inputs to output detailed and precise metric depth maps. We unify different depth estimation tasks by abstracting various types of depth measurements as depth prior. To clarify the scope, we first outline common types of depth priors and their primary applications:

1) Sparse points (depth completion): Depth from LiDAR or SfM (Schonberger & Frahm, 2016) is typically sparse. Completing sparse depth priors is crucial for applications such as 3D reconstruction (Roessle et al., 2022; Chung et al., 2024) and autonomous driving (Tao et al., 2022; Häne et al.,

Methods	Torget Task	Sparse Point			Low-	Missing Area			Mixed
Methods	Target Task	SfM	LiDAR	Extreme	resolution	Range	Shape	Object	Mixeu
Marigold-DC (Viola et al., 2024)	Depth Completion	1	√	√	Х	×	Х	Х	Х
Omni-DC (Zuo et al., 2024)	Depth Completion	/	✓	✓	✓	X	X	X	X
PromptDA (Lin et al., 2025)	Depth Super-resolution	X	X	X	✓	X	X	X	X
DepthLab (Liu et al., 2024a)	Depth Inpainting	X	X	X	X	X	✓	✓	X
Prior Depth Anything	All-Rounder	/	✓	✓	✓	✓	✓	✓	✓

Table 1: Applicable scenarios of current prior-based monocular depth estimation models. *SfM*: sparse matching points from SfM, *LiDAR*: sparse LiDAR line patterns, *Extreme*: extremely sparse points (100 points), *Range*: missing depth within a specific range, *Shape*: missing regular-shaped areas, *Object*: missing depth of an object.

2017; Carranza-García et al., 2022). **2) Low-resolution (depth super-resolution):** Low-power Time of Flight (ToF) cameras (Lange & Seitz, 2001), commonly used in mobile phones, capture low-resolution depth maps. Depth super-resolution is essential for spatial perception, VR (Rasla & Beyeler, 2022), and AR (Slater & Wilbur, 1997) in portable devices (He et al., 2021a). **3) Missing areas (depth inpainting):** Stereo matching failures (Lowe, 2004; Rublee et al., 2011) or 3D Gaussian splatting edits (Liu et al., 2024b; Yu et al., 2024) may leave large missing areas in depth maps. Filling these gaps is vital for 3D scene generation and editing (Yu et al., 2024). **4) Mixed prior:** In real-world scenarios, different depth priors often coexist. For instance, structured light cameras (Herrera et al., 2012) often generate low-resolution and incomplete depth maps, due to their limited working range. Handling these mixed priors is vital for practical applications.

In Tab. 1, we detail the patterns of each prior. Existing methods mainly focus on specific limited priors, limiting their use in diverse, real-world scenarios. In this work, we propose *Prior Depth Anything*, motivated by the complementary advantage between predicted and measured depth maps, as illustrated in Fig. 1. Technically, we design a coarse-to-fine pipeline to explicitly and progressively combine the depth prediction with measured depth prior, achieving impressive robustness to any image with any prior.

We first introduce coarse metric alignment to pre-fill incomplete depth priors using predicted relative depth maps, which effectively narrows the domain gap between various prior types. Next, we apply the fine structure refinement to rectify misaligned geometric structures in the pre-filled depth priors caused by inherent noises in the depth measurements. Specifically, the pre-filled depth prior (with accurate metric data) and the relative depth prediction (with fine details and structure) are provided as additional inputs to a conditioned MDE model. Guided by the RGB image input, the model can combine the strengths of two complementary depth sources for the final output.

We evaluate our model on 7 datasets with varying depth priors. It achieves zero-shot depth completion, super-resolution, and inpainting within a single model, matching or outperforming previous models that are specialized for only one of these tasks. More importantly, our model achieves significantly better results when different depth priors are mixed, highlighting its effectiveness in more practical and varying scenarios.

Our contribution can be summarized as follows:

- We propose *Prior Depth Anything*, a unified framework to estimate fine-detailed and complete metric depth with any depth priors. Our model can seamlessly handle zero-shot depth completion, super-resolution, inpainting, and adapt to more varied real-world scenarios.
- We introduce coarse metric alignment to pre-fill depth priors, narrowing the domain gap between different types of depth prior and enhancing the model's generalization.
- We design fine structure refinement to alleviate the inherent noise in depth measurements. This involves a conditioned MDE model to granularly merge the pre-filled depth prior and prediction based on image content.
- Our method exhibits superior zero-shot results across various datasets and tasks, surpassing
 even state-of-the-art methods specifically designed for individual tasks.

2 RELATED WORK

2.1 MONOCULAR DEPTH ESTIMATION

Monocular depth estimation (MDE) is a fundamental computer vision task that predicts the depth of each pixel from a single color image (Eigen et al., 2014; Fu et al., 2018; Bhat et al., 2021). Recently, with the success of "foundation models" (Bommasani et al., 2021), some studies (Ranftl et al., 2020; Yang et al., 2024;b; Xu et al., 2025; He et al., 2025; Ke et al., 2024; Bochkovskii et al., 2025; Yin et al., 2023) have attempted to build depth foundation models by scaling up data and using stronger backbones, enabling them to predict detailed geometric structures for any image.

MiDaS (Ranftl et al., 2020) made the pioneering study by training an MDE model on joint datasets to improve generalization. Following this line, Depth Anything v1 (Yang et al., 2024a) scaled training with massive unlabeled image data, while Depth Anything v2 (Yang et al., 2024b) further enhanced its ability to handle fine details, reflections, and transparent objects by incorporating highly precise synthetic data (Wang et al., 2021; 2020; Yao et al., 2020; Cabon et al., 2020; Roberts et al., 2021).

Although these methods have demonstrated high accuracy and robustness, they primarily produce unscaled relative depth maps due to the significant scale differences between indoor and outdoor scenes. While Metric3D (Yin et al., 2023; Hu et al., 2024) and Depth Pro (Bochkovskii et al., 2025) achieve zero-shot metric depth estimation through canonical camera transformations, the precision remains limited compared to measurement technologies.

Our method builds on the strength of existing depth foundation models, which excel at precisely capturing relative geometric structures and fine details in any image. By progressively integrating accurate but incomplete metric information in the depth measurements, our model can generate dense and detailed metric depth maps for any scene.

2.2 PRIOR-BASED MONOCULAR DEPTH ESTIMATION

In practical applications, depth measurement methods like multi-view matching (Cordts et al., 2016) or sensors (Silberman et al., 2012; Geiger et al., 2013) can provide accurate metric information, but due to their inherent nature or cost limitations, these measurements often capture incomplete information. Some recent studies have tried to use this measurement data as prior knowledge in the depth estimation process to achieve dense and accurate metric depth. These methods, however, primarily focus on specific patterns of depth measurement, which can be categorized into three types based on their input patterns:

Depth Completion As noted in (Roessle et al., 2022), SfM reconstructions from 19 images often result in depth maps with only 0.04% valid pixels. Completing the sparse depth maps with observed RGB images is a fundamental computer vision task (Tang et al., 2024; Cheng et al., 2020; 2019; Zuo & Deng, 2024; Zhang et al., 2023b; Park et al., 2024). Recent approaches like Omni-DC (Zuo et al., 2024) and Marigold-DC (Viola et al., 2024) have achieved certain levels of zero-shot generalization across diverse scenes and varying sparsity levels. However, due to the lack of explicit scene geometry guidance, they face challenges in extremely sparse scenarios.

Depth Super-resolution Obtaining high-resolution metric depth maps with depth cameras usually demands significant power. A more efficient alternative is to use low-power sensors to capture low-resolution maps and then enhance them using super-resolution. Early efforts (Zhong et al., 2023; Xian et al., 2020; Zhao et al., 2022; He et al., 2021b), however, show limited generalization to unseen scenes. Recent PromptDA (Lin et al., 2025) achieves effective zero-shot super-resolution by using the low-resolution map as a prompt for depth foundation models (Yang et al., 2024b).

Depth Inpainting As discussed in (Yang et al., 2024b; Ke et al., 2024), due to inherent limitations in stereo matching and depth sensors, even "ground truth" depth data in real-world datasets often have significant missing regions. Additionally, in applications like 3D Gaussian editing and generation (Liu et al., 2024b; Yu et al., 2024; Chung et al., 2023), there is a need to fill holes in depth maps. DepthLab (Liu et al., 2024a) first fills holes using interpolation and then refines the results with a depth-guided diffusion model. However, interpolation errors reduce its effectiveness for large missing areas or incomplete depth ranges.

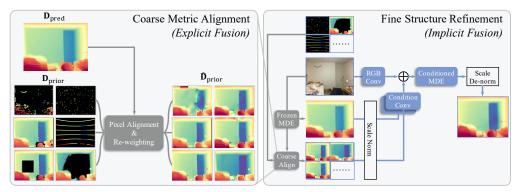


Figure 2: **Prior Depth Anything.** Considering RGB images, any form of depth prior \mathbf{D}_{prior} , and relative prediction \mathbf{D}_{pred} from a frozen MDE model, coarse metric alignment first explicitly combines the metric data in \mathbf{D}_{prior} and geometry structure in \mathbf{D}_{pred} to produce the pre-filled depth map $\hat{\mathbf{D}}_{prior}$. Fine structure refinement implicitly merges the complementary information to produce the final metric depth map.

These previous methods have two main limitations: 1) Poor performance when prior is limited. 2) Difficulty generalizing to unseen prior patterns. Our approach, Prior Depth Anything, tackles these challenges by explicitly using geometric information from depth prediction in a coarse-to-fine process, achieving impressive generalization and accuracy across various patterns of prior input.

3 Prior Depth Anything

Advanced monocular depth estimation models excel in predicting detailed, complete relative depth maps with precise geometric structures for any image. In contrast, depth measurement technologies can provide metric depth maps but suffer from inherent noise and varying patterns of incompleteness. Inspired by the complementary strengths of estimated and measured depth, we introduce *Prior Depth Anything* to progressively and effectively merge the two depth sources. To handle diverse real-world scenarios, we take measurement depth in any form as the metric prior, producing fine-grained and complete metric depth maps for any image with any prior.

3.1 PRELIMINARY

Given an RGB image $\mathbf{I} \in \mathbb{R}^{3 \times H \times W}$ and its corresponding metric depth prior $\mathbf{D}_{\text{prior}} \in \mathbb{R}^{H \times W}$, prior-based monocular depth estimation takes the \mathbf{I} and $\mathbf{D}_{\text{prior}}$ as input, aiming to output the depth map $\mathbf{D}_{\text{output}} \in \mathbb{R}^{H \times W}$ that is detailed, complete, and metrically precise. As discussed in 1, depth priors obtained by different measurement techniques often display various forms of incompleteness. To handle various priors with a unified framework, we uniformly represent the coordinates of valid positions in $\mathbf{D}_{\text{prior}}$ as $\mathbf{P} = \{x_i, y_i\}_{i=0}^N$, which N pixels are valid.

3.2 COARSE METRIC ALIGNMENT

As shown in Fig. 2, different types of depth priors exhibit distinct missing patterns (e.g. sparse points, low-resolution grids, or irregular holes). These differences in sparsity and incompleteness restrict models' ability to generalize across various priors. To tackle this, we propose pre-filling missing regions to transform all priors into a shared intermediate domain.

However, interpolation-based filling, used in previous methods (Liu et al., 2024a; Lin et al., 2025), preserves pixel-level metric information but ignores geometric structure, leading to significant errors in the filled areas. On the other hand, global alignment (Chung et al., 2024; 2023), which scales relative depth predictions to match priors, maintains the fine structure of predictions but loses critical pixel metric details. To address these challenges, we propose pixel-level metric alignment, which aligns geometry predictions and metric priors at pixel level, preserving both predicted structure and original metric information.

$$\hat{\mathbf{D}}_{\text{prior}}(x, y) = \mathbf{D}_{\text{pred}}(x, y), \text{ where } (x, y) \in \mathbf{P}$$
 (1)

For each missing pixel (\hat{x}, \hat{y}) , we first identify its K closest valid points $\{x_k, y_k\}_{k=1}^K$ from the valid pixel set \mathbf{P} using k-nearest neighborhood (kNN). Then, we compute the optimal scale s and shift t parameter that minimizes the least-squares error between the depth values of \mathbf{D}_{pred} and \mathbf{D}_{prior} at the K supporting points:

$$s, t = \underset{s,t}{\operatorname{arg\,min}} \sum_{k=1}^{K} \| s \cdot \mathbf{D}_{\text{pred}}(x_k, y_k) + t - \mathbf{D}_{\text{prior}}(x_k, y_k) \|_2$$
 (2)

With the optimal scale s and shift t, we fill the missing pixels in $\hat{\mathbf{D}}_{prior}$ by linearly aligning the predicted depth value at (\hat{x}, \hat{y}) to metric prior:

$$\hat{\mathbf{D}}_{\text{prior}}(\hat{x}, \hat{y}) = s \cdot \mathbf{D}_{\text{pred}}(\hat{x}, \hat{y}) + t \tag{3}$$

To handle this issue, we further introduce distance-aware weighting for more smooth and accurate alignment. Within the alignment objective of Eq. 2, we re-weight each supporting point based on its distance to the query pixel, modifying Eq. 2 to:

$$s, t = \arg\min_{s,t} \sum_{k=1}^{K} \frac{\|s \cdot \mathbf{D}_{\text{pred}}(x_k, y_k) + t - \mathbf{D}_{\text{prior}}(x_k, y_k)\|_2}{\|(\hat{x}, \hat{y}) - (x_k, y_k)\|_2}$$
(4)

This simple modification ensures smoother transitions between regions and improves robustness by emphasizing geometrically closer measurements.

In summary, by explicitly integrating the accurate metric information from \mathbf{D}_{prior} and the fine geometric structures from \mathbf{D}_{pred} , we cultivate the pre-filled dense prior $\hat{\mathbf{D}}_{prior}$, which offers two main advantages: 1) Similar Pattern: Filling the missing area narrows the differences between various prior types, improving the generalization across different scenarios. 2) Fine Geometry: The filled regions, derived from linear transformations of depth prediction, natively preserve the fine geometric structures, significantly boosting performance when prior information is limited.

3.3 FINE STRUCTURE REFINEMENT

Although the prefilled coarse dense depth is generally accurate in metric, the parameter-free approach is sensitive to noise in depth priors. A single noisy pixel on blurred edges can disrupt all filled regions relying on it as a supporting point. To tackle these errors, we further implicitly leverage the MDE model's ability in capturing precise geometric structures in RGB images, learning to rectify the noise in priors and produce refined depth.

Metric Condition Specifically, we incorporate the pre-filled prior $\hat{\mathbf{D}}_{prior}$ as extra conditions to the pre-trained MDE model. With the guidance of RGB images, the conditioned MDE model is trained to correct the potential noise and error in $\hat{\mathbf{D}}_{prior}$. To this end, we introduce a condition convolutional layer parallel to the RGB input layer, as shown in Fig. 2. By initializing the condition layer to zero, our model can natively inherit the ability of pre-trained MDE model.

Geometry Condition In addition to leveraging the MDE model's inherent ability in capturing geometric structures from RGB input, we also incorporate existing depth predictions as an external geometry condition to help refine the coarse pre-filled prior. The depth prediction \mathbf{D}_{pred} , obtained from the frozen MDE model, is also passed into the condition MDE model through a zero-initialized conv layer.

Scale Normalization Then, we normalize the metric condition $\hat{\mathbf{D}}_{prior}$ and geometry condition \mathbf{D}_{pred} to [0,1] for two key benefits: 1) Better Scene Generalization: different scenes (e.g. indoor vs. outdoor) have significant depth scale differences. Normalization removes this scale variance, improving performance across diverse scenes. 2) Better MDE Model Generalization: predictions from different frozen MDE models also have varying scales. Normalizing \mathbf{D}_{pred} enables test-time model switching, offering flexible accuracy-efficiency trade-offs for diverse demands, and enabling seamless improvements as MDE models advance.

Synthetic Training Data As discussed in (Ke et al., 2024; Yang et al., 2024b), real depth datasets often face issues such as blurred edges and missing values. Therefore, we leverage synthetic datasets, Hypersim (Roberts et al., 2021) and vKITTI (Cabon et al., 2020), with precise GT to

Model	Encoder	NYUv2		5	ScanNet		ETH-3D		DIODE		KITTI		[
Model	Elicodei	S+M	L+M	S+L	S+M	L+M	S+L	S+M	L+M	S+L	S+M	L+M	S+L	S+M	L+M	S+L
DAv2 Depth Pro	ViT-L ViT-L	4.59 4.46	4.95 4.87	5.08 5.41	4.34 4.22	4.55 4.44	4.62 4.51	6.82 6.39	10.99 7.08	7.49 9.29	11.57 12.92	10.23 8.42	12.46 8.91	10.20 6.29	11.12 9.24	10.97 9.18
Omni-DC Marigold-DC DepthLab PromptDA	SDv2 SDv2 ViT-L	2.86 2.26 6.33 17.00	3.26 3.38 3.96 3.76	3.81 3.82 5.80 17.13	2.88 2.19 5.16 15.27	2.76 2.87 3.38 4.21	3.64 3.37 5.24 15.44	2.09 2.15 7.87 18.34	4.17 4.77 4.70 9.01	4.59 5.13 7.45 18.73	4.23 4.98 8.83 18.24	4.80 6.73 6.40 9.97	0.02	٠,٠	8.63 9.67 23.12 54.35	20.70
PriorDA (ours)	DAv2-B+ViT-S DAv2-B+ViT-B Depth Pro+ViT-B	2.09 2.04 2.01	2.88 2.82 2.82	3.17 3.09 3.08	2.14 2.11 <u>2.06</u>	2.56 2.50 2.48	2.94 2.86 <u>2.82</u>	1.65 1.56 1.61	3.96 3.84 3.83	4.16 4.08 4.05	3.76 3.62 <u>3.40</u>	4.43 4.30 <u>4.23</u>	4.89 4.73 4.60	3.97 3.86 <u>3.37</u>	8.38 8.40 8.12	8.53 8.57 8.25

Table 2: **Zero-shot depth estimation with mixed prior.** All results are reported in AbsRel ↓. "S": "Extreme" setting in sparse points, "L": "×16" in low-Resolution, "M": "Shape" (square masks of 160) in missing area. We highlight **Best**, <u>second best</u> results. "Depth Pro+ViT-B" indicates the frozen MDE and conditioned MDE. DAv2-B: Depth Anything v2 ViT-B (Yang et al., 2024b), SDv2: Stable Diffusion v2. (Rombach et al., 2022).

Model	Encoder		NYU	v2		ScanNet		ETH-3D		DIODE			KITTI			
Model	Elicodei	SfM	LiDAR	Extreme	SfM	LiDAR	Extreme	SfM	LiDAR	Extreme	SfM	LiDAR	Extreme	SfM	LiDAR	Extreme
DAv2 Depth Pro	ViT-L ViT-L	5.31 4.80	4.85 4.47	4.77 4.41	5.88 5.27	4.60 4.12	4.68 4.23	5.84 5.68	7.41 5.31	6.61 6.51	12.45 10.53	12.00	14.37 8.98	9.83 6.45	8.86 6.05	9.25 6.19
Omni-DC Marigold-DC DepthLab PromptDA	SDv2 SDv2 ViT-L	2.87 2.65 5.92 18.68	2.12 1.90 4.30 17.59	2.63 2.13 6.30 16.96	6.09 4.32 9.87 18.13	2.02 1.76 3.56 15.99	2.71 2.12 5.09 15.18	2.57 4.65 13.82 25.02	1.88 2.27 6.40 18.86	1.98 2.03 8.01 18.18	4.99 8.41 16.67 25.46	3.96 5.12 7.45 18.58	0.00	-0.71	5.27 6.88 37.17 21.96	4.17 5.62 40.29 21.39
PriorDA (ours)	DAv2-B+ViT-S DAv2-B+ViT-B Depth Pro+ViT-B	2.42 2.38 2.36	2.01 1.95 1.96	2.01 1.97 1.96	3.90 3.85 3.84	2.19 2.15 2.14	2.09 2.04 2.01	3.26 3.10 3.13	1.90 1.82 1.84	1.61 1.50 1.54	5.25 5.16 5.07	3.88 3.77 3.66	3.65 3.64 <u>3.37</u>	3.73 3.74 3.35	4.81 4.73 <u>4.12</u>	3.76 3.76 3.28

Table 3: **Zero-shot depth completion** (e.g. sparse points prior). All results are reported in AbsRel ↓. "SfM": points sampled with SIFT (Lowe, 2004) and ORB (Rublee et al., 2011), "LiDAR": 8 LiDAR lines, "Extreme": 100 random points.

drive our conditioned MDE model to rectify the noise in measurements. From the precise ground truth, we randomly sample sparse points, create square missing areas, or apply downsampling to construct different synthetic priors. To mimic real-world measurement noise, we add outliers and boundary noise to perturb the sampled prior, following (Zuo et al., 2024).

Learning Objective As mentioned earlier, both the metric and geometry conditions are normalized. Thus, we apply the de-normalization transformation to convert the output into the ground truth scale. Following ZoeDepth (Bhat et al., 2023), we use the scale-invariant log loss for pixel-level supervision.

4 Experiment

4.1 EXPERIMENTAL SETTING

Implementation Details During training, we utilize the Depth Anything V2 ViT-B model as the frozen MDE model to produce relative depth predictions. During inference, the frozen MDE model can be swapped with any other pre-trained model. The k-value of the kNN process in Sec. 3.2 is set to 5. We initialize the conditioned MDE Model with two versions of Depth Anything V2: ViT-S and ViT-B. We train the conditioned MDE model for 200K steps with a batch size of 64, using 8 GPUs. The AdamW optimizer with a cosine scheduler is employed, where the base learning rate is set to 5e-6 for the MDE encoder and 5e-5 for the MDE decoder.

Benchmarks Our method aims to provide accurate and complete metric depth maps in a zero-shot manner for any image with any prior. **To cover "any image"**, we evaluate models on 7 unseen real-world datasets, including NYUv2 (Silberman et al., 2012) and ScanNet (Dai et al., 2017) for indoor, ETH3D (Schops et al., 2017) and DIODE (Vasiljevic et al., 2019) for indoor/outdoor, KITTI (Geiger et al., 2012) for outdoor, ARKitScenes (Baruch et al., 2021) and RGB-D-D (He et al., 2021b) for captured low-resolution depth. **To cover "any prior"**, we construct 9 individual patterns: sparse points (SfM, LiDAR, Extremely sparse), low-resolution (captured, x8, x16), and missing areas (Range, Shape, Object). We also mix these patterns to simulate more complex scenarios.

Model	Encoder	ARKit	Scenes	RGB	RGB-D-D		NYUv2		ScanNet		I-3D	DIODE		KITTI	
Model	Elicoder	AbsRel↓	RMSE↓	AbsRel↓	RMSE↓	$8 \times$	16×	8×	16×	8×	16×	8×	16×	8×	16×
DAv2 Depth Pro	ViT-L ViT-L	3.67 3.25	0.0764 0.0654	4.67 4.28	0.1116 0.1030	4.77 4.48	5.13 4.83	4.64 4.17	4.85 4.40	6.27 5.88	7.38 6.79	12.49 8.20	11.20 8.33	9.54 6.76	11.22 9.16
Omni-DC Marigold-DC DepthLab PromptDA	SDv2 SDv2 ViT-L	2.14 2.17 2.10 1.34	0.0435 0.0448 <u>0.0411</u> 0.0347	2.09 2.15 2.13 2.79	0.0685 0.0672 0.0624 0.0708	1.57 1.83 2.60 <u>1.61</u>	3.11 3.32 3.73 1.75	1.29 1.63 1.89 1.87	2.65 2.83 3.19 1.93	1.86 2.33 2.60 1.80	4.09 4.75 4.50 2.56	2.81 4.28 4.42 3.18	4.71 6.60 6.16 3.73	4.05 5.17 17.17 3.92	8.35 9.47 22.90 4.95
PriorDA (ours)	DAv2-B+ViT-S DAv2-B+ViT-B Depth Pro+ViT-B	2.09 1.94 1.95	0.0414 0.0404 0.0408	2.07 2.02 2.02	0.0597 0.0581 0.0581	1.73 1.72 1.72	2.79 2.73 2.74	1.60 1.61 1.58	2.50 2.45 2.43	2.06 2.00 1.99	3.91 3.79 <u>3.77</u>	3.09 3.10 3.01	4.36 4.23 4.15	4.54 4.65 4.44	8.20 8.24 <u>7.99</u>

Table 4: **Zero-shot depth super-resolution** (e.g. low-resolution prior). ARKitScenes and RGB-D-D provide captured low-resolution depth. For other datasets, results are reported in AbsRel \downarrow , with low-resolution maps created by downsampling the GT depths.

Model	Encoder	NYUv2		ScanNet		ETH-3D		DIODE			KITTI					
Wiodei	Elicodei	Range	Shape	Object	Range	Shape	Object	Range	Shape	Object	Range	Shape	Object	Range	Shape	Object
DAv2 Depth Pro	ViT-L ViT-L	17.40 10.89	5.24 9.20	6.56 6.52	16.75 16.76	4.64 15.39	6.74 6.80	68.76 10.37	8.23 34.28	19.22 17.28				31.12 14.51	14.93 16.11	17.94 8.19
Omni-DC Marigold-DC DepthLab PromptDA	SDv2 SDv2 ViT-L	23.24 19.83 23.85 36.67	5.94 2.37 2.66 20.88	13.79 6.18 10.87 23.14	22.89 17.14 21.17 35.86	5.44 <u>1.97</u> 2.08 17.87	8.71 6.66 10.40 21.89	30.61	4.81 2.15 2.75 24.94	17.97 7.72 10.53 27.42	38.83 39.33 41.01 49.50	7.59 6.51	25.43 18.97 17.17 28.29	33.44 40.43	9.21 13.60	15.06 7.72 18.66 38.29
PriorDA (ours)	DAv2-B+ViT-S DAv2-B+ViT-B Depth Pro+ViT-B	16.86 16.61 16.31	2.30 2.30 2.17	5.72 5.49 5.59	14.29 14.48 14.18	2.01 1.99 1.98	5.87 5.73 5.87	21.16 21.90 22.72	1.98 1.76 1.76	6.52 6.09 6.21	36.59 36.64 34.90	5.58 5.94 4.86	10.77 <u>9.72</u> 11.99	30.04 30.79 30.44	6.67 6.29 <u>5.47</u>	7.99 7.52 <u>6.04</u>

Table 5: **Zero-shot depth inpainting** (e.g. missing area prior). All results are reported in AbsRel ↓. Metrics are calculated only on the masked (inpainted) regions. "Range": masks for depth beyond 3m (indoors) and 15m (outdoors), "Shape": average result for square masks of sizes 80, 120, 160, and 200, "Object": object segmentation masks detected by YOLO (Jocher et al., 2023).

4.2 Comparison on Mixed Depth Prior

We quantitatively evaluate the ability to handle challenging unseen mixed priors in Tab 2. In terms of absolute performance, all versions of our model outperform compared baselines. More importantly, our model is less impacted by the additional patterns. For example, compared to the setting that only uses sparse points in Tab 3, adding missing areas or low-resolution results in only minor performance drops $(1.96\rightarrow2.01,\ 3.08$ in NYUv2). In contrast, Omni-DC $(2.63\rightarrow2.86,\ 3.81)$, and Marigold-DC $(2.13\rightarrow2.26,\ 3.82)$ show larger declines. These results highlight the robustness of our method to different prior inputs.

4.3 Comparison on Individual Depth Prior

Zero-shot Depth Completion Tab 3 shows the zero-shot depth completion results with different kinds and sparsity levels of sparse points as priors. Compared to Omni-DC (Zuo et al., 2024) and Marigold-DC (Viola et al., 2024), which are specifically designed for depth completion and rely on sophisticated, time-consuming structures, our approach achieves better overall performance with simpler and more efficient designs.

Zero-shot Depth Super-resolution In Tab 4, we present results for super-resolution depth maps. On benchmarks where low-resolution maps are created through downsampling (e.g. NYUv2 (Silberman et al., 2012), ScanNet (Dai et al., 2017), *etc*), our approach achieves performance comparable to state-of-the-art methods. However, since downsampling tends to include overly specific details from the GT depths, directly replicating noise and blurred boundaries from GT leads to better results instead. Therefore, ARKitScenes (Baruch et al., 2021) and RGB-D-D (He et al., 2021b) are more representative and practical, as they use low-power cameras to capture the low-resolution depths. On these two benchmarks, our method achieves leading performance compared to other zero-shot methods.

Zero-shot Depth Inpainting In Tab 5, we evaluate the performance of inpainting missing regions in depth maps. In the practical and challenging "Range" setting, our method achieves superior results, which is highly meaningful for improving depth sensors with limited effective working ranges. Additionally, it outperforms all alternatives in filling square and object masks, demonstrating its potential for 3D content generation and editing.

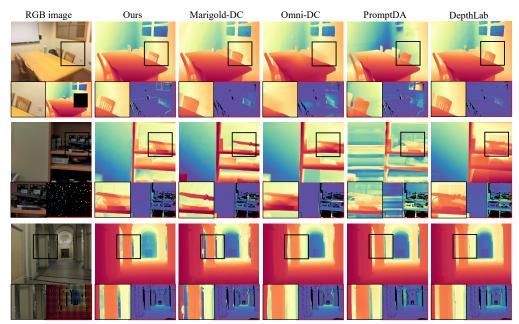


Figure 3: Qualitative comparisons with previous methods. The depth prior or error map is shown below each sample.

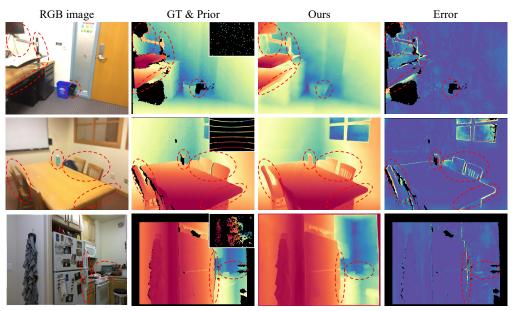


Figure 4: Error analysis on widely used but indeed noisy benchmarks (Silberman et al., 2012; Dai et al., 2017). Red means higher error, while blue indicates lower error.

4.4 QUALITATIVE ANALYSIS

In Fig 3, we provide a qualitative comparison of the outputs from different models. Our model consistently outperforms previous approaches, offering richer details, sharper boundaries, and more accurate metrics.

Fig 4 visualizes the error maps of our model. The errors mainly occur around blurred edges in the "ground truth" of real data. Our method effectively corrects the noise in labels while aligning with the metric information from the prior. These "beyond ground truth" cases highlight the potential of our approach in addressing the inherent noise in depth measurement techniques. More visualizations can be found in the supp.

	S	L	M	S+M	L+M	S+L
Interpolation	7.93	3.88	8.96	8.38	4.55	7.99
Ours (w/o re-weight)	2.92	3.44	6.91	3.22	4.53	4.36
Ours	2.42	3.51	6.70	2.60	4.32	4.40

Table 6: Accuracy of pre-filled depth maps with different strategies. We compare the pre-filled dense prior (i.e., $\hat{\mathbf{D}}_{prior}$) with ground truth.

Model	Encoder	S	L	M	S+M	L	S+L
	ViT-S	2.15	2.77	2.68	2.22	2.87	3.20
Depth	ViT-B	1.97	2.73	2.50	2.02	2.82	3.09
Anything V2	ViT-L	1.92	2.71	2.29	1.97	2.79	3.04
	ViT-G	1.87	2.70	2.22	1.94	2.76	3.02
Depth Pro	ViT-L	1.96	2.74	2.35	2.01	2.82	3.08

Table 8: Effect of using different frozen MDE models. The conditioned MDE model is ViT-B version here.

Metric	Geometry	S	L	M	S+M	L+M	S+L
X	✓	5.46	5.29	5.48	5.36	5.30	5.46
1	×	2.10	2.94	2.58	2.17	3.02	3.31
/	✓	1.96	2.74	2.48	2.01	2.82	3.08

Table 7: Effect of each condition for conditioned MDE models.

	Seen	een Unseen							
	S	L	M	S+M	L+M	S+L			
None	2.50	3.71	46.07	2.50	3.74	3.64			
Interpolation	3.40	2.68	4.28	3.53	2.94	3.56			
Ours (w/o re-weight)	2.13	2.86	2.58	2.19	2.94	3.25			
Ours	1.99	2.82	2.26	2.06	2.90	3.11			

Table 9: Effect of pre-filled strategies on generalization. We train models with various pre-fill strategies using only sparse points and evaluate their ability to generalize to unseen types.

4.5 ABLATION STUDY

We use Depth Anything V2 ViT-B as the frozen MDE and ViT-S as the conditioned MDE for ablation studies by default. All results are evaluated on NYUv2.

Accuracy of different pre-fill strategy As shown in Fig 6, our pre-fill method outperforms simple interpolation across all scenarios by explicitly utilizing the precise geometric structures in depth prediction. Additionally, the re-weight mechanism further enhances performance.

Effectiveness of fine structure refinement Comparing the pre-filled coarse depth maps in Tab 6 with the final output accuracy in Tab 3, 4, 5 and 2, the performance improvements after fine structure refinement (sparse points: $2.42 \rightarrow 2.01$, low-resolution: $3.51 \rightarrow 2.79$, missing areas: $6.70 \rightarrow 2.48$, S+M: $2.60 \rightarrow 2.09$, L+M: $4.32 \rightarrow 2.88$, S+L: $4.40 \rightarrow 3.17$) demonstrate its effectiveness in rectifying misaligned geometric structures in pre-filled depth maps while maintaining its metric information.

Effectiveness of metric and geometry condition We evaluate the impact of metric and geometry guidance for the conditioned MDE model in Tab 7. The results show that combining both conditions achieves the best performance, emphasizing the importance of reinforcing geometric information during the fine structure refinement stage.

Testing-time improvement We investigate the potential of test-time improvements in Tab 8. Our findings reveal that larger and stronger frozen MDE models consistently bring higher accuracy, while smaller models maintain competitive performance and enhance the efficiency of the entire pipeline. These findings underscore the flexibility of our model and its adaptability to various scenarios.

Pre-fill strategy for generalization From Tab 9, we observe that our pixel-level metric alignment helps the model generalize to new prior patterns, and the re-weighting strategy further enhances the robustness by improving the accuracy of the pre-filled depth map.

5 CONCLUSION

In this work, we present *Prior Depth Anything*, a robust and powerful solution for prior-based monocular depth estimation. We propose a coarse-to-fine pipeline to progressively integrate the metric information from incomplete depth measurements and the geometric structure from relative depth predictions. The model offers three key advantages: 1) delivering accurate and fine-grained depth estimation with any type of depth prior; 2) offering flexibility to adapt to extensive applications through test-time module switching; and 3) showing the potential to rectify inherent noise and blurred boundaries in real depth measurements.

REFERENCES

- 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. In *NeurIPS*, 2021.
- Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *CVPR*, 2021.
- Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zeroshot transfer by combining relative and metric depth. *arXiv:2302.12288*, 2023.
- Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. In *ICLR*, 2025.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv:2108.07258*, 2021.
- Yohann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2. arXiv:2001.10773, 2020.
- Manuel Carranza-García, F Javier Galán-Sales, José María Luna-Romera, and José C Riquelme. Object detection using depth completion and camera-lidar fusion for autonomous driving. *Integrated Computer-Aided Engineering*, 2022.
- Xinjing Cheng, Peng Wang, and Ruigang Yang. Learning depth with convolutional spatial propagation network. *TPAMI*, 2019.
- Xinjing Cheng, Peng Wang, Chenye Guan, and Ruigang Yang. Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion. In AAAI, 2020.
- Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv:2311.13384*, 2023.
- Jaeyoung Chung, Jeongtaek Oh, and Kyoung Mu Lee. Depth-regularized optimization for 3d gaussian splatting in few-shot images. In *CVPR*, 2024.
- Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017.
- Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *CVPR*, 2022.
- David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NeurIPS*, 2014.
- Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *ICCV*, 2023.
- Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018.
- Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 2013.
 - Christian Häne, Lionel Heng, Gim Hee Lee, Friedrich Fraundorfer, Paul Furgale, Torsten Sattler, and Marc Pollefeys. 3d visual perception for self-driving cars using a multi-camera system: Calibration, mapping, localization, and obstacle detection. *Image and Vision Computing*, 2017.

- Jing He, Haodong Li, Wei Yin, Yixun Liang, Leheng Li, Kaiqiang Zhou, Hongbo Zhang, Bingbing Liu, and Ying-Cong Chen. Lotus: Diffusion-based visual foundation model for high-quality dense prediction. In *ICLR*, 2025.
- Lingzhi He, Hongguang Zhu, Feng Li, Huihui Bai, Runmin Cong, Chunjie Zhang, Chunyu Lin, Meiqin Liu, and Yao Zhao. Towards fast and accurate real-world depth super-resolution: Benchmark dataset and baseline. In *CVPR*, 2021a.
 - Lingzhi He, Hongguang Zhu, Feng Li, Huihui Bai, Runmin Cong, Chunjie Zhang, Chunyu Lin, Meiqin Liu, and Yao Zhao. Towards fast and accurate real-world depth super-resolution: Benchmark dataset and baseline. In *CVPR*, 2021b.
 - Daniel Herrera, Juho Kannala, and Janne Heikkilä. Joint depth and color camera calibration with distortion correction. *TPAMI*, 2012.
 - 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. *TPAMI*, 2024.
 - Glenn Jocher, Jing Qiu, and Ayush Chaurasia. Ultralytics YOLO, January 2023. URL https://github.com/ultralytics/ultralytics.
 - Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *CVPR*, 2024.
 - Robert Lange and Peter Seitz. Solid-state time-of-flight range camera. *IEEE Journal of quantum electronics*, 2001.
 - Haotong Lin, Sida Peng, Jingxiao Chen, Songyou Peng, Jiaming Sun, Minghuan Liu, Hujun Bao, Jiashi Feng, Xiaowei Zhou, and Bingyi Kang. Prompting depth anything for 4k resolution accurate metric depth estimation. In *CVPR*, 2025.
 - 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:2412.18153*, 2024a.
 - Zhiheng Liu, Hao Ouyang, Qiuyu Wang, Ka Leong Cheng, Jie Xiao, Kai Zhu, Nan Xue, Yu Liu, Yujun Shen, and Yang Cao. Infusion: Inpainting 3d gaussians via learning depth completion from diffusion prior. *arXiv:2404.11613*, 2024b.
 - David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004.
 - Jin-Hwi Park, Chanhwi Jeong, Junoh Lee, and Hae-Gon Jeon. Depth prompting for sensor-agnostic depth estimation. In *CVPR*, 2024.
 - René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *TPAMI*, 2020.
 - Alex Rasla and Michael Beyeler. The relative importance of depth cues and semantic edges for indoor mobility using simulated prosthetic vision in immersive virtual reality. In *Proceedings of the 28th ACM symposium on virtual reality software and technology*, 2022.
 - 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 *ICCV*, 2021.
- Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *CVPR*, 2022.
 - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
 - Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *ICCV*, 2011.

- Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016.
 - Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *CVPR*, 2017.
 - Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012.
- Mel Slater and Sylvia Wilbur. A framework for immersive virtual environments (five): Speculations on the role of presence in virtual environments. *Presence: Teleoperators & Virtual Environments*, 1997.
 - Jie Tang, Fei-Peng Tian, Boshi An, Jian Li, and Ping Tan. Bilateral propagation network for depth completion. In *CVPR*, 2024.
 - Yifu Tao, Marija Popović, Yiduo Wang, Sundara Tejaswi Digumarti, Nived Chebrolu, and Maurice Fallon. 3d lidar reconstruction with probabilistic depth completion for robotic navigation. In *IROS*, 2022.
 - Igor Vasiljevic, Nick Kolkin, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z Dai, Andrea F Daniele, Mohammadreza Mostajabi, Steven Basart, Matthew R Walter, et al. Diode: A dense indoor and outdoor depth dataset. *arXiv:1908.00463*, 2019.
 - Massimiliano Viola, Kevin Qu, Nando Metzger, Bingxin Ke, Alexander Becker, Konrad Schindler, and Anton Obukhov. Marigold-dc: Zero-shot monocular depth completion with guided diffusion. *arXiv*:2412.13389, 2024.
 - Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In CVPR, 2025.
 - Qiang Wang, Shizhen Zheng, Qingsong Yan, Fei Deng, Kaiyong Zhao, and Xiaowen Chu. Irs: A large naturalistic indoor robotics stereo dataset to train deep models for disparity and surface normal estimation. In *ICME*, 2021.
 - Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *CVPR*, 2024.
 - 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 *IROS*, 2020.
 - Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *CVPR*, 2019.
 - Diana Wofk, Fangchang Ma, Tien-Ju Yang, Sertac Karaman, and Vivienne Sze. Fastdepth: Fast monocular depth estimation on embedded systems. In *ICRA*, 2019.
 - Chuhua Xian, Kun Qian, Zitian Zhang, and Charlie CL Wang. Multi-scale progressive fusion learning for depth map super-resolution. *arXiv:2011.11865*, 2020.
 - Guangkai Xu, Yongtao Ge, Mingyu Liu, Chengxiang Fan, Kangyang Xie, Zhiyue Zhao, Hao Chen, and Chunhua Shen. Diffusion models trained with large data are transferable visual models. In *ICLR*, 2025.
- Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 2024a.
 - Lihe Yang, Bingyi Kang, Zilong Huang, Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. In *NeurIPS*, 2024b.
 - Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmys: A large-scale dataset for generalized multi-view stereo networks. In *CVPR*, 2020.

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 ICCV, 2023. Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. Wonderworld: Interactive 3d scene generation from a single image. arXiv:2406.09394, 2024. Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In ICCV, 2023a. Youmin Zhang, Xianda Guo, Matteo Poggi, Zheng Zhu, Guan Huang, and Stefano Mattoccia. Completionformer: Depth completion with convolutions and vision transformers. In CVPR, 2023b. Zixiang Zhao, Jiangshe Zhang, Shuang Xu, Zudi Lin, and Hanspeter Pfister. Discrete cosine trans-form network for guided depth map super-resolution. In CVPR, 2022. Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, and Chuang Gan. 3d-vla: A 3d vision-language-action generative world model. In ICML, 2024. Zhiwei Zhong, Xianming Liu, Junjun Jiang, Debin Zhao, and Xiangyang Ji. Guided depth map super-resolution: A survey. ACM Computing Surveys, 2023. Yiming Zuo and Jia Deng. Ogni-dc: Robust depth completion with optimization-guided neural iterations. In ECCV, 2024. Yiming Zuo, Willow Yang, Zeyu Ma, and Jia Deng. Omni-dc: Highly robust depth completion with multiresolution depth integration. arXiv:2411.19278, 2024.

A APPLICATION

	Monoc	ular Depth Est	imation	Multi-view D	epth Estimation
	NYU	ETH-3D	KITTI	ETH-3D	KITTI
VGGT	3.54 (-)	4.94 (-)	6.56 (-)	2.46 (-)	18.75 (-)
+Omni-DC	4.12 (0.58)	6.08 (1.14)	6.85 (<mark>0.29</mark>)	2.64 (0.18)	18.66 (-0.09)
+Marigold-DC	4.06 (0.52)	5.43 (0.49)	7.63 (1.07)	2.81 (0.35)	18.86 (<mark>0.11</mark>)
+DepthLab	3.56 (0.02)	4.92 (-0.02)	7.97 (1.41)	2.25 (-0.21)	19.47 (<mark>0.72</mark>)
+PromptDA	3.43 (-0.11)	4.97 (0.03)	6.50 (-0.06)	2.48 (0.02)	18.91 (<mark>0.16</mark>)
+PriorDA	3.45 (-0.09)	4.43 (-0.51)	6.39 (-0.17)	1.99 (-0.47)	18.61 (-0.14)

Table 10: Results of refining VGGT depth prediction with different methods. All results are reported as AbsRel.

To demonstrate our model's real-world applicability, we employ prior-based monocular depth estimation models to refine the depth predictions from VGGT Wang et al. (2025), a state-of-the-art 3D reconstruction foundation model. VGGT provides both a depth and confidence map. We take the top 30% most confident pixels as depth prior and apply different prior-based models to obtain finer depth predictions. ¹

Table 10 reports VGGT's performance in monocular and multi-view depth estimation, along with the effectiveness of different prior-based methods as refiners. We observe that only our PriorDA consistently improves VGGT's predictions, primarily due to its ability to adapt to diverse priors. These surprising results highlight PriorDA's broad application potential.

B Inference efficiency analysis

Model	Encoder	Param	Latency(ms)
Omni-DC	-	85M	223
Marigold-DC	SDv2	1,290M	28,607
DepthLab	SDv2	2,080M	9,078
PromptDA	ViT-L	340M	108
PriorDA (ours)	DAv2-B+ViT-S DAv2-B+ViT-B Depth Pro+ViT-B	97M+25M 97M+98M 952M+98M	79 ₄₉₊₉₊₂₁ 107 ₄₉₊₉₊₄₉ 941 ₈₈₃₊₉₊₄₉

Table 11: Analysis of inference efficiency. " $_{x+x+x}$ " represents the latency of the frozen MDE model, coarse metric alignment, and conditioned MDE model, respectively.

In Tab 11, we analyze the inference efficiency of different models on one A100 GPU for an image resolution of 480×640 . Overall, compared to previous approaches, our model variants achieve leading performance while demonstrating certain advantages in parameter number and inference latency. For a more detailed breakdown, we provide the time consumption for each stage of our method. Notably, the coarse metric alignment, which relies on kNN and least squares, introduces negligible overhead. Our method demonstrates significant efficiency advantages over previous methods.

C MORE ABLATION STUDY

Effect of k-value in Coarse Metric Alignment In Tab 12, we analyze the impact of the k-value on the accuracy of the final output $\mathbf{D}_{\text{output}}$. Overall, our method is non-sensitive to the selection of k-value, with most selections yielding strong structural results. This highlights the effectiveness of the nearest neighbor approach in preserving detailed metric information.

¹For models less adept at handling missing pixels (DepthLab, PromptDA), the entire VGGT depth prediction was provided as prior.

L

2.52

2.16

2.31

2.27

756	•		S
757			- 5
758		k=3	2.00
759		k=5	1.97
760		k=10 k=20	2.00
761		κ=20	2.10
762	Table 12: Im	pact of o	differer
763	_		
764			
765	_		

772 773 774

775

776

777

778 779 780

781 782

783

784

785

786 787 788

789 790

791

792

793

794

795

796

797 798 799

800 801

802

803

804 805 806

807 808

809

2.76 ifferent k-value on the accuracy of the final output $\mathbf{D}_{\text{output}}$.

M

2.74

2.73

2.74

S+M

2.10

2.04

2.09

2.16

L+M

2.83

2.82

2.83

2.83

S+L

3.07

3.09

3.12

3.14

	NYU					
	S	L	M	S+M	L+M	S+L
Omni-DC	2.6	3.1	6.6	2.8	3.3	3.8
PromptDA	17.0	1.8	21.4	17.0	3.8	17.1
PriorDA _{SP-only}	2.0	2.8	2.3	2.1	2.9	3.1
PriorDA _{LW-only}	2.2	3.0	3.4	2.4	3.3	3.7
PriorDA _{ALL}	2.0	<u>2.7</u>	2.5	2.0	2.8	3.1

Table 13: Impact of used prior patterns during training.

Trained on Single Prior PriorDA trained with sparse point or low-resolution prior is compared with baselines in Tab. 13, showing advantages over baseline trained with same single prior. This indicates that our method's ability to generalize to any form of prior stems from the coarse metric alignment, rather than the use of multiple patterns during training.

D More Qualitative Results

To further explore the boundaries of our model's capabilities and its potential to rectify the "ground truth" depth, we offer more error analysis with different patterns of priors on the 7 unseen datasets (Figure 5 for RGB-D-D, Figure 6 for ARKitScenes, Figure 7 for NYUv2, Figure 8 for ScanNet, Figure 9 for ETH-3D, Figure 10 for DIODE, Figure 11 for KITTI).

MORE TRAINING DETAILS

For each training sample, we randomly select one of three patterns (e.g., sparse points, lowresolution, or missing areas) with equal probability to sample the depth prior from the ground truth depth map. Specifically: for sparse points, we randomly select 100 to 2000 pixels as valid; for low-resolution, we downsample the GT map by a factor of 8; and for missing areas, we generate a random square mask with a side length of 160 pixels. It is worth mentioning that, we find that using multiple patterns or only using sparse points lead to similar results, as shown in Tab 13. This indicates that our method's ability to generalize to any form of prior stems from the coarse metric alignment, rather than the use of multiple patterns during training.

MORE DISSCUSION

Why PromptDA preforms poor in mixed prior, sparse point and missing area? PromptDA is only trained on low-resolution but complete priors. Therefore, it struggles to handle incomplete priors, even if we already apply DepthLab's pre-processing to fill in missing areas.

USAGE OF LARGE LANGUAGE MODELS

In this work, we employed Large Language Models (LLMs) in the writing phase for sentence-level polishing.

H LIMITATIONS AND FUTURE WORKS

Currently, our largest conditioned MDE model is initialized with Depth Anything v2 ViT-B. Given that larger versions of the Depth Anything v2 model exhibit stronger capabilities, training conditioned MDE models based on larger backbones is an important direction for future work. Additionally, following Depth Anything, all training images are resized to 518×518. In contrast, PromptDA is natively trained at 1440×1920 resolution. Therefore, training at higher resolutions to better handle easily accessible high-resolution RGB images is another crucial direction for our future research.

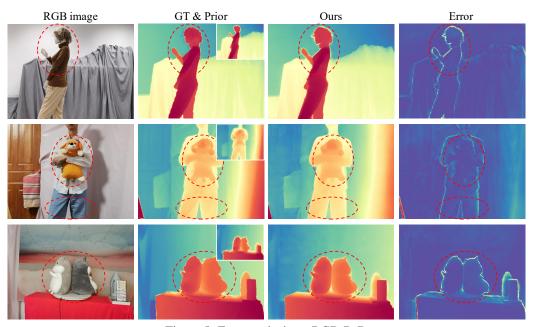


Figure 5: Error analysis on RGB-D-D.

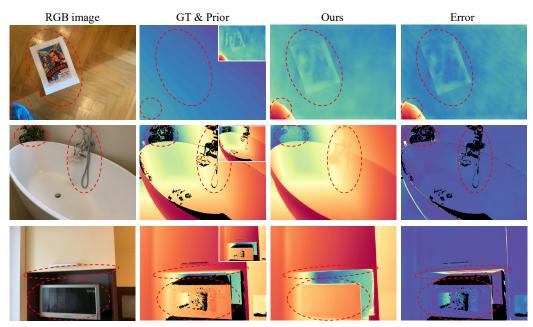


Figure 6: Error analysis on ARKitScenes.

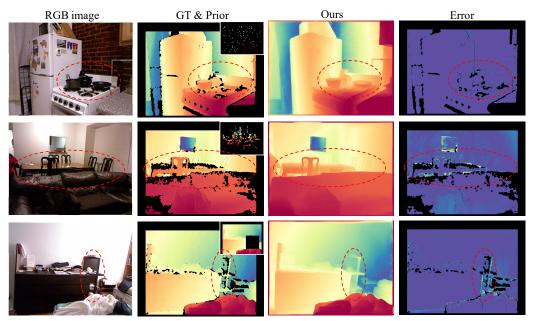


Figure 7: Error analysis on NYUv2.

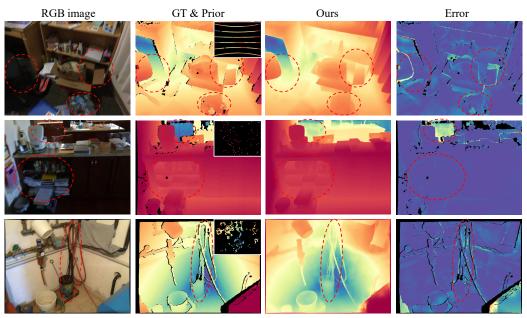


Figure 8: Error analysis on ScanNet.

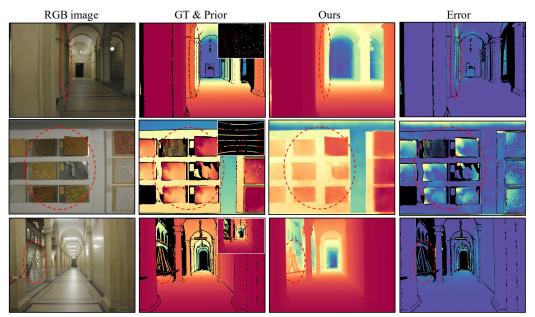


Figure 9: Error analysis on ETH-3D.



Figure 10: Error analysis on DIODE.

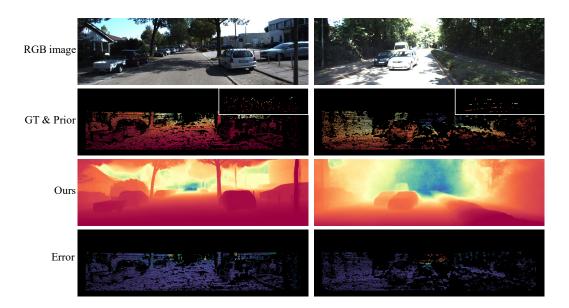


Figure 11: Error analysis on KITTI.