

# One Look is Enough: Seamless Patchwise Refinement for Zero-Shot Monocular Depth Estimation on High-Resolution Images

Byeongjun Kwon  
KAIST\*

byeongjun.kwon@kaist.ac.kr

Munchurl Kim<sup>†</sup>  
KAIST\*

mkimee@kaist.ac.kr

[https://kaist-viclab.github.io/One-Look-is-Enough\\_site](https://kaist-viclab.github.io/One-Look-is-Enough_site)

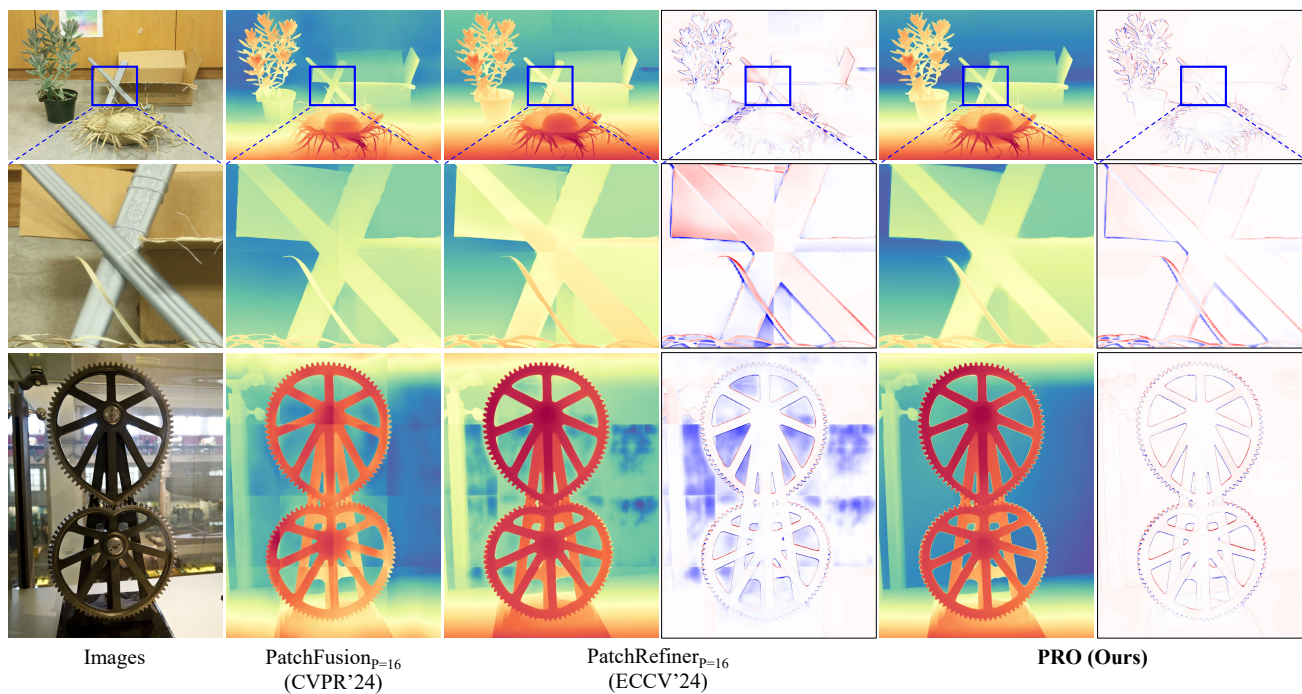


Figure 1. **Qualitative comparison of patch-based depth estimation models (PatchFusion [23], PatchRefiner [24] and PRO (Ours)) for high-resolution images.** PatchRefiner [24] and PRO (Ours) also visualize the residuals to effectively highlight depth discontinuity issues in the 4th and 6th columns. All models perform depth estimations on  $4 \times 4$  patches of input images. Zoom-in can help distinguish the grid lines (depth discontinuities) of the  $4 \times 4$  patches. Our proposed PRO achieves smooth boundary transitions without depth discontinuity artifacts along the grid boundaries.

## Abstract

Zero-shot depth estimation (DE) models exhibit strong generalization performance as they are trained on large-scale datasets. However, existing models struggle with high-resolution images due to the discrepancy in image resolutions of training (with smaller resolutions) and inference (for high resolutions). Processing them at full resolution leads to decreased estimation accuracy on depth with

tremendous memory consumption, while downsampling to the training resolution results in blurred edges in the estimated depth images. Prevailing high-resolution depth estimation methods adopt a patch-based approach, which introduces depth discontinuity issues when reassembling the estimated depth patches, resulting in test-time inefficiency. Additionally, to obtain fine-grained depth details, these methods rely on synthetic datasets due to the real-world sparse ground truth depth, leading to poor generalizability. To tackle these limitations, we propose Patch Refine Once (PRO), an efficient and generalizable tile-based framework.

\*Korea Advanced Institute of Science and Technology.

<sup>†</sup>Corresponding author.

Our PRO consists of two key components: (i) Grouped Patch Consistency Training that enhances test-time efficiency while mitigating the depth discontinuity problem by jointly processing four overlapping patches and enforcing a consistency loss on their overlapping regions within a single backpropagation step, and (ii) Bias Free Masking that prevents the DE models from overfitting to dataset-specific biases, enabling better generalization to real-world datasets even after training on synthetic data. Zero-shot evaluations on Booster, ETH3D, Middlebury 2014, and NuScenes demonstrate that our PRO can be seamlessly integrated into existing depth estimation models. It preserves the performance of original depth estimation models even under grid-based inference on high-resolution images, exhibiting minimal depth discontinuities along patch boundaries. Moreover, our PRO achieves significantly faster inference speed compared to prior patch-based methods.

## 1. Introduction

Monocular Depth Estimation (MDE) [10, 13, 14, 27, 33, 36, 43] has been widely investigated as the demand for 3D information continues to grow in autonomous driving, robotics, and virtual reality applications. After Midas [30] proposed an MDE network trainable with mixed training datasets, numerous zero-shot MDE networks [2, 9, 11, 19, 31, 46, 47] have been studied to predict depth for unseen real-world images. However, due to the architectural limitation or limited image resolutions of training datasets, most zero-shot MDE networks are trained specifically with the images of low resolutions (e.g.  $384 \times 384$ ,  $518 \times 518$ ). Consequently, when high resolution images are fed into these models, the resulting estimated depth images tend to contain disrupted overall structures with low-frequency artifacts, although yielding improved edge details [7, 26].

Patch-based methods [23, 24, 26] have shown promising results by splitting high-resolution images into patches for DE, mitigating memory consumption issues while achieving remarkable performance. However, they suffer from depth discontinuity problems (e.g. boundary artifacts along grid boundaries) which occur when independent DE patches are reassembled to construct a complete (whole) depth map because depth continuity is maintained within each patch but not between patches. Previous methods [23, 24] alleviate this depth discontinuity problem (i) by incorporating a consistency loss during training and (ii) by ensemble averaging at test time which slows down inference speed, making it impractical for real-world applications.

A common strategy for training high-resolution depth estimation is to train models on real-world high-resolution datasets [17, 20, 34, 39, 42] or leverage high-resolution synthetic datasets [18, 38, 40, 41]. Real-world datasets have the advantage of a smaller domain gap compared to synthetic datasets, but their ground truth (GT) depths are of-




Figure 2. Comparison of performance and efficiency on Middlebury 2014 [34]. The area of each circle represents the inference time. Circles of the same color represent the same model with different patch numbers ( $P = 16$  for small-sized yellow and grey circles,  $P = 177$  for large-sized yellow and grey circles). Our model (PRO) achieves the best performance in terms of edge errors ( $D^3R$ ) and depth errors (AbsRel) while maintaining the fastest inference time. Inference time for all models is measured on the RTX 4090.

ten sparse, particularly around edges, where supervision is crucial. Additionally, the difficulty of acquiring dense depth data in real-world scenarios hinders the training of DE models for high-resolution images. On the other hand, synthetic datasets provide dense GT depth with fine details, but models trained on synthetic datasets often struggle with zero-shot inference on real-world data due to the domain gap [21]. Moreover, we observe that transparent objects are inaccurately labeled in the UnrealStereo4K dataset [38], the only synthetic dataset providing 4K depth annotations, as their depth values correspond to the background behind them.

To address the aforementioned problems, we propose Patch Refine Once (PRO), a novel refinement model that enables DE models to generate accurate depth predictions with only a single refinement per patch at inference time (i.e., “One Look”). Our PRO consists of two strategies: Grouped Patch Consistency Training and Bias Free Masking. Fig. 1 shows a qualitative comparison of patch-based depth estimation models for high-resolution images. As shown, our PRO yields depth refinement results with minimal depth discontinuities along grid boundaries, whereas two recent state-of-the-art (SOTA) models still exhibit noticeable boundary artifacts. Fig. 2 illustrates a comparison of performance and efficiency between recent SOTA depth refinement models and our PRO. As observed, our PRO provides the best performance in terms of  $D^3R$ , AbsRel and inference speed. Our contribution is summarized as:

- We propose Grouped Patch Consistency Training (GPCT) strategy that can be well harmonized with existing DE models to mitigate the depth discontinuity problem for patch-wise DE approaches on high-resolution images. The GPCT does not require test-time ensemble, allowing a  $12 \times$  faster inference speed compared to the SOTA patch-based method.
- We introduce Bias Free Masking (BFM) to selectively

apply supervision signals by masking out unreliable regions in images. The unreliable regions are the regions that have inaccurately annotated depth in GT, which is often observed in the UnrealStereo4K dataset. This issue is effectively addressed by our BFM, which prevents the model from overfitting to domain-specific biases.

- Our PRO achieves SOTA performance in zero-shot evaluation on Booster [29], ETH3D [35], Middlebury 2014 [34], and NuScenes [4] across all metrics, and achieves efficient inference ( $12\times$  faster than PatchRefiner with 177 patches), outperforming recent state-of-the-art depth refinement methods on high-resolution images.

## 2. Related work

### 2.1. Zero-shot Monocular Depth Estimation

Early MDE models [1, 10, 25, 33, 49] were trained individually on each dataset [12, 37], achieving high performance on the training data but poor generalization to unseen datasets due to the domain gap. Zero-shot MDE models that perform well on unseen images have been widely studied to improve generalization. MegaDepth [22] and DiverseDepth [48] construct large-scale datasets by collecting Internet images, improving adaptability to images from diverse scenes and conditions. They have inspired further efforts to scale up training datasets to enhance zero-shot performance. MiDaS [30] proposes a scale and shift invariant loss, which enables MDE networks to be trained on a diverse mixture of datasets, making the model robust to unseen datasets. DPT [31], Omnidata [9], and MiDaS v3.1 [2] improve DE performance by replacing CNNs with transformer architectures. DepthAnything [46] introduces a semi-supervised approach that utilizes 62M unlabeled images as pseudo labels for training. This substantial increase in data scale has led to improved generalization performance. DepthAnythingV2 [47] observes that real-world depth datasets contain label noise and lack fine details. It is trained with pseudo GT labels obtained from a teacher model trained on a synthetic dataset, achieving high-quality depth maps.

Instead of scaling up datasets, some studies utilize prior knowledge from diffusion models for DE. Marigold [19] fine-tunes a pretrained Stable Diffusion [32] model on a relatively small synthetic dataset, showing competitive results. Geowizard [11] jointly estimates depth and surface normals by leveraging self-attention across different domains to enhance geometric consistency. However, as most zero-shot MDE models are trained on low-resolution images (e.g.,  $384\times 384$ ,  $518\times 518$ ), their performance degrades when applied to high-resolution images. When high-resolution images are downsampled to the training resolution for inference, fine details are lost. On the other hand, performing inference with a larger resolution input preserves details but degrades depth accuracy.

### 2.2. High-Resolution Depth Estimation

High-resolution DE models aim to predict accurate depth while preserving fine details. SMD-Net [38] utilizes an implicit function [5] to predict a mixture density for precise DE at object boundaries. Dai *et al.* [7] introduce a Poisson fusion-based depth map optimization method by applying guided filtering in a self-supervised learning framework. SDDR [21] proposes a self-distilled depth refinement method that generates pseudo edge labels, addressing local consistency issues and edge deformation noise through a noisy Poisson fusion process. Patch-based high-resolution DE methods select patches and merge patch-wise results to improve depth details. BoostingDepth [26] introduces a patch selection process based on edge density, and employs iterative refinement through multi-resolution depth merging. PatchFusion [23] avoids patch selection by applying a predefined partitioning scheme, and leverages shifted window self-attention to fuse global and local information. PatchRefiner [24] initializes depth estimation with a low-resolution depth map, predicts residuals, and proposes a detail and disentangling (DSD) loss for training on both real-world and synthetic datasets. However, patch-based methods inherently suffer from depth discontinuities at patch boundaries, as they emphasize local information over global consistency, or maintain depth continuity only within individual patches. To address this problem, prior works [23, 24] employ a large number of patches for test-time ensemble, which significantly slows down inference. By contrast, our PRO (Patch Refine Once) model efficiently achieves fine-grained depth refinement with only a single pass per patch during inference, despite employing a tile-based approach.

### 2.3. Training Data for HR Depth Estimation

High-resolution DE requires datasets with fine-grained depth details and minimal label noise. Real-world datasets [4, 12, 15, 22, 37] are often sparse due to the limitations of LiDAR or contain inaccurate depth boundaries when obtained through Structure-from-Motion (SfM). Therefore, previous research [7, 26] has trained depth networks using a small number of relatively dense real-world datasets [20, 34, 44]. Some studies [7, 21] utilize pseudo labels for self-supervision, but these labels often contain noise. Other works [23, 24] leverage synthetic datasets for more accurate depth, but suffer from poor generalization ability due to the domain gap. To take advantage of dense depth annotations in synthetic datasets without compromising generalization performance, we introduce Bias Free Masking (BFM), which selectively identifies the regions where supervision is applied. It mitigates overfitting to biases in synthetic datasets and preserves robustness in zero-shot high-resolution DE.




Figure 3. Labeling error examples in UnrealStereo4K [38] dataset.

### 3. Method

We describe our Patch Refine Once (PRO) model. Section 3.1 introduces the overall pipeline of PRO, followed by the two key contributions: Grouped Patch Consistency Training (GPCT) and Bias Free Masking (BFM), which are detailed in Sections 3.2 and 3.3, respectively.

#### 3.1. Overview of Patch Refine Once (PRO)

Fig. 4 illustrates the overall pipeline of the proposed PRO framework. Following [23], given an original input image  $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ , we estimate a coarse depth map  $\mathbf{D}_c$  by feeding a resized version of  $\mathbf{I}$  into a pretrained zero-shot MDE network  $\Psi$ .  $\mathbf{D}_c$  captures the overall scene structure, which is essential for preserving global consistency.

To complement the local details, we take the following steps. First, we crop the original image  $\mathbf{I}$  into  $N$  grid-based patches  $\{\mathbf{P}^i\}_{i=1}^N$ . Each individual patch  $\mathbf{P}^i$  is then resized and fed into  $\Psi$  to estimate the corresponding fine depth map  $\mathbf{D}_f^i$ . We extract two sets of five-level features from the decoder of  $\Psi$ : (i)  $\mathbf{F}_c = \{\mathbf{f}_{c,j}^i\}_{j=1}^5$  for the resized image, and (ii)  $\mathbf{F}_f = \{\mathbf{f}_{f,j}^i\}_{j=1}^5$  for the patch  $\mathbf{P}^i$ . To align  $\mathbf{F}_c$  with  $\mathbf{F}_f$ , we apply the ROI operation [16] to extract patch-aligned features from  $\mathbf{F}_c$ . The patch-level features  $\mathbf{F}_f^i$  and the corresponding ROI-extracted features are subsequently passed into a fusion module within the residual prediction network  $\theta$ . The fusion module integrates these features using a wavelet transform [8] and shallow convolutional layers, effectively combining global and local information. See *supplementary material* for details of the fusion module. Finally, the refined depth map for  $\mathbf{P}^i$ , denoted as  $\mathbf{D}_{\text{refine}}^i$ , is obtained as:

$$\mathbf{D}_{\text{refine}}^i = \theta(\text{concat}(\mathbf{P}^i, \text{ROI}(\mathbf{D}_c, \mathbf{P}^i), \mathbf{D}_f^i)), \quad (1)$$

where  $\text{ROI}(\mathbf{D}_c, \mathbf{P}^i)$  extracts the region of interest from  $\mathbf{D}_c$  corresponding to the spatial extent of  $\mathbf{P}^i$ , yielding  $\mathbf{D}_c^i$ . The operator  $\text{concat}$  denotes channel-wise concatenation. Unlike previous methods [23, 24] where separate networks are trained to predict  $\mathbf{D}_c$  and  $\mathbf{D}_f$  individually, we utilize the same pretrained zero-shot MDE network  $\Psi$  and train only the residual prediction network as shown in Fig. 4-(a).

#### 3.2. Grouped Patch Consistency Training (GPCT)

To address the boundary artifacts introduced by patch-wise refinement, we propose a simple yet effective GPCT strategy that ensures depth consistency across patch boundaries. While previous approaches [23, 24] focus only on two diagonally adjacent overlapping patches (e.g., (A, D) or (B, C) in Fig. 4-(b)) to enforce consistency constraints, our method employs four overlapping patches (A, B, C, and D in Fig. 4-(b)) simultaneously. As shown in Fig. 4-(b), we divide the training sample into overlapping patches, so that each patch overlaps with its neighbors. After refining each patch independently, we apply a depth consistency loss  $\mathcal{L}_{\text{con}}$  to enforce consistency between the depth refinement results of overlapping patches:

$$\mathcal{L}_{\text{con}} = \sum_{i \neq j} \frac{1}{|\Omega|} \sum_{p \in \Omega} \left( \mathbf{D}_{\text{refine}}^i(p) - \mathbf{D}_{\text{refine}}^j(p) \right)^2, \quad (2)$$

where  $\mathbf{D}_{\text{refine}}^i$  and  $\mathbf{D}_{\text{refine}}^j$  denote the depth predictions from overlapping patches  $i$  and  $j$ , respectively. We denote the overlapping region between patches  $i$  and  $j$  as  $\Omega$ , with  $|\Omega|$  indicating the number of pixels within this region. The merged depth map  $\mathbf{D}_{\text{merged}}$  at position  $p$  is computed as:

$$\mathbf{D}_{\text{merged}}(p) = \frac{1}{N_o(p)} \sum_{i \in \{A, B, C, D\}} \mathbf{D}_{\text{refine}}^i(p), \quad (3)$$

where  $\mathbf{D}_{\text{refine}}^i$  denotes the refined depth prediction from patch  $i$  (i.e.,  $i \in \{A, B, C, D\}$  as shown in Fig. 4-(b)), and  $N_o(p)$  is the number of overlapping patches that contribute a depth estimate at pixel  $p$ . Then, we calculate the loss between merged depth  $\mathbf{D}_{\text{merged}}$  and GT depth  $\mathbf{D}_{\text{gt}}$ .

Since our method computes the loss using all four patches simultaneously, every backward propagation step enforces consistency at all boundaries. In contrast, the prior method [23] only applied a consistency loss on the overlapped region between two patches (e.g., (A, D) or (B, C)) at a time, leading to insufficient boundary supervision. Our approach provides a stronger supervision signal, yielding significantly improved cross-patch consistency and smooth boundary transitions during inference. By training with this GPCT strategy, our PRO model mitigates boundary artifacts without requiring additional refinement during inference, allowing efficient patch-wise depth estimation, as only a single refinement step per patch (i.e., ‘‘One Look’’) is required.

#### 3.3. Bias Free Masking (BFM)

To utilize dense GT depth when training depth refinement models for fine-grained details, we use a synthetic dataset such as the UnrealStereo4K [38]. However, the UnrealStereo4K dataset, which provides 4K resolutions, annotates the depths of transparent objects as the depths of their backgrounds (e.g., window objects in Fig. 3), which introduces a

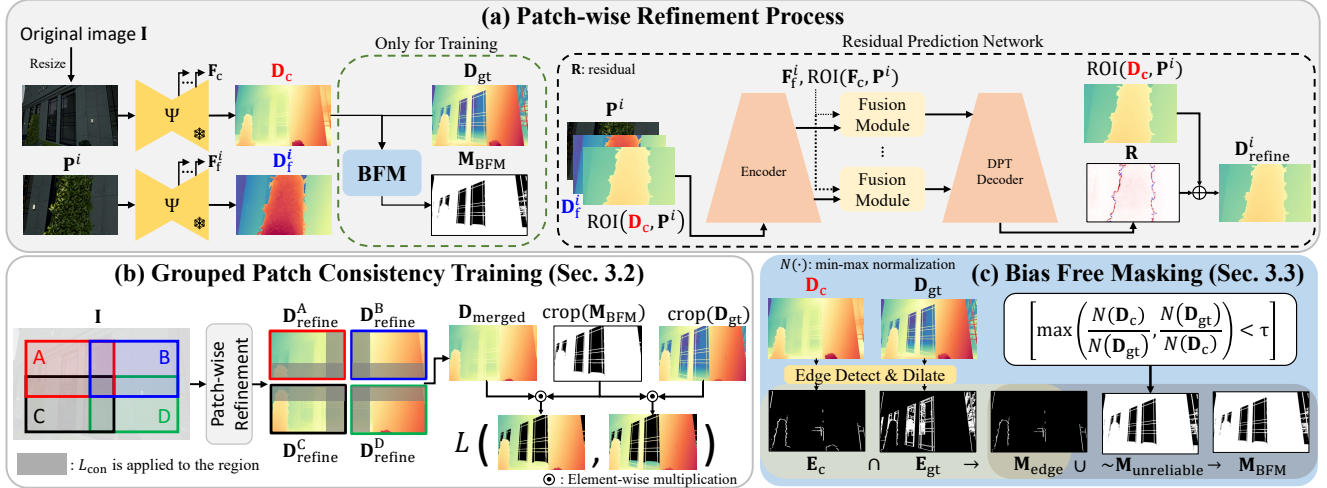


Figure 4. **Overview of the Framework.** (a) **Patch-wise Refinement Process.** First, the pretrained MDE model  $\Psi$  estimates a coarse depth  $\mathbf{D}_c$  and a fine depth  $\mathbf{D}_f^i$  from the resized input image  $\mathbf{I}$  and the  $i$ -th patch  $\mathbf{P}^i$ . Based on  $\mathbf{D}_c$ , the residual prediction network predicts residuals  $\mathbf{R}$ . This process is applied to every patch. Note that Bias Free Masking (BFM) is only required at training. (b) **Grouped Patch Consistency Training.** The training image is cropped with overlapping regions, followed by patch-wise refinement. Subsequently, depth consistency loss  $\mathcal{L}_{\text{con}}$  is applied to the gray-shaded regions to enforce consistency between depth results refined separately. (c) **Bias Free Masking.** We identify the reliable region as  $\mathbf{M}_{\text{BFM}}$  by utilizing  $\mathbf{D}_c$  and  $\mathbf{D}_{\text{gt}}$ .

specific bias in the dataset. To maintain the benefits of dense depth labeling while avoiding overfitting to the dataset-specific biases, we propose Bias Free Masking (BFM) that uses the prior knowledge of the pretrained zero-shot MDE models. That is, BFM aims to exclude the unreliable regions corresponding to the incorrect GT depths in transparent objects in the synthetic dataset (UnrealStereo4K) during training. Since the pretrained zero-shot MDE model  $\Psi$  can be considered to have informative prior knowledge learned from large scale image-depth pairs, the model can be used to identify unreliable regions where  $\mathbf{D}_c$  and  $\mathbf{D}_{\text{gt}}$  significantly deviate from each other, indicating potential biases to the synthetic dataset.

Given a resized version for input image  $\mathbf{I}$ , we obtain coarse depth  $\mathbf{D}_c$  from  $\Psi$  and GT depth  $\mathbf{D}_{\text{gt}}$  from the synthetic dataset. Then, by measuring the relative consistency between  $\mathbf{D}_c$  and  $\mathbf{D}_{\text{gt}}$ , we identify unreliable regions as:

$$\mathbf{M}_{\text{unreliable}} = \left[ \max \left( \frac{N(\mathbf{D}_c)}{N(\mathbf{D}_{\text{gt}})}, \frac{N(\mathbf{D}_{\text{gt}})}{N(\mathbf{D}_c)} \right) > \tau \right], \quad (4)$$

where  $[\cdot]$  represents the Iverson bracket,  $N(\mathbf{D})$  denotes min-max normalization for  $\mathbf{D}$  and  $\tau$  is an empirically chosen threshold (set to 2 in our experiments).  $\mathbf{M}_{\text{unreliable}}$  as an unreliable mask identifies the regions where  $\mathbf{D}_c$  and  $\mathbf{D}_{\text{gt}}$  significantly deviate, indicating potential synthetic dataset biases or inconsistencies, as aforementioned. Since  $\mathbf{D}_c$  lacks sharp edges, complementing  $\mathbf{M}_{\text{unreliable}}$  to obtain the reliable regions is a naive approach, which causes the exclusion of critical edge regions from training, thus resulting in the removal of valuable depth gradients required for refinement. To address this, we incorporate edge information into

the reliable mask generation. However, simply utilizing the edge map from  $\mathbf{D}_{\text{gt}}$  as an edge mask could bring noise in a training sample because the flat transparent region could include the edges due to background with edges as observed in Fig. 4-(c). To handle this problem, we additionally employ the edge maps from  $\mathbf{D}_c$ . Consequently, we first extract edges from both  $\mathbf{D}_c$  and  $\mathbf{D}_{\text{gt}}$ , and dilate them as:

$$\mathbf{E}_c = \text{dilate}(\text{edge}(\mathbf{D}_c)), \quad \mathbf{E}_{\text{gt}} = \text{dilate}(\text{edge}(\mathbf{D}_{\text{gt}})), \quad (5)$$

to guarantee that the final edge mask should not miss the regions that require depth refinements, otherwise they might be excluded simply from complementing  $\mathbf{M}_{\text{unreliable}}$ . We define the edge mask as  $\mathbf{M}_{\text{edge}} = \mathbf{E}_c \cap \mathbf{E}_{\text{gt}}$ . Finally, the reliable mask for BFM is constructed as  $\mathbf{M}_{\text{BFM}} = \mathbf{M}_{\text{edge}} \cup \sim \mathbf{M}_{\text{unreliable}}$  where  $\sim \mathbf{M}_{\text{unreliable}}$  is the complement of  $\mathbf{M}_{\text{unreliable}}$ . This ensures that the unreliable regions caused by large discrepancies between  $\mathbf{D}_c$  and  $\mathbf{D}_{\text{gt}}$  are ignored, while essential edge details are preserved for training. Consequently, supervision is applied only to the reliable regions via a masked loss:

$$\mathcal{L}_{\text{masked}} = \frac{1}{|\mathbf{M}_{\text{BFM}}|} \sum_{p \in \mathbf{M}_{\text{BFM}}} \mathcal{L}(\mathbf{D}_{\text{merged}}(p), \mathbf{D}_{\text{gt}}(p)) \quad (6)$$

where  $\mathcal{L}$  is the loss between merged depth and GT depth, which is the combination of an L1 loss, an L2 loss, and a multi-scale gradient loss with four scale levels [22] using a ratio of 1:1:5. By restricting training to reliable regions, the depth refinement model refines depth only where necessary, avoiding modifications in regions where the pretrained model's prior knowledge is considered more reliable than

Methods	Publications	Runtime (s)	Booster		ETH3D		Middle14			NuScenes		DIS
			AbsRel↓	$\delta_1$ ↑	AbsRel↓	$\delta_1$ ↑	AbsRel↓	$\delta_1$ ↑	D <sup>3</sup> R ↓	AbsRel↓	$\delta_1$ ↑	BR↑
DepthAnythingV2 [47]	NIPS 2024	-	0.0307	0.993	0.0465	0.983	0.0307	0.994	0.0979	0.106	0.882	0.059
BoostingDepth [26]	CVPR 2021	12.7	0.0330	0.993	0.0552	0.974	0.0330	0.995	0.1035	0.115	0.870	0.170
PatchFusion <sub>P=16</sub> [23]	CVPR 2024	3.4	0.0504	0.985	0.0735	0.956	0.0450	0.989	0.1124	0.141	0.831	<b>0.206</b>
PatchFusion <sub>P=177</sub> [23]	CVPR 2024	36.8	0.0496	0.986	0.0723	0.957	0.0448	0.989	0.1050	0.139	0.833	0.189
PatchRefiner <sub>P=16</sub> [24]	ECCV 2024	1.5	0.0348	0.989	0.0435	<b>0.985</b>	0.0292	0.995	0.0830	0.107	0.879	0.151
PatchRefiner <sub>P=177</sub> [24]	ECCV 2024	16.2	0.0336	0.991	0.0430	<b>0.985</b>	0.0292	0.995	0.0805	0.106	0.881	0.141
<b>PRO (Ours)</b>	-	<b>1.4</b>	<b>0.0304</b>	<b>0.994</b>	<b>0.0422</b>	<b>0.985</b>	<b>0.0287</b>	<b>0.996</b>	<b>0.0803</b>	<b>0.104</b>	<b>0.883</b>	0.156

Table 1. Quantitative comparison of depth estimation methods on Booster [29], ETH3D [35], Middlebury 2014 [34], NuScenes [4], and DIS-5K [28] datasets. **Bold** indicates the best performance in each metric.

the (noisy) GT. This helps prevent the model from refining transparent regions incorrectly.

**Final Loss function.** The final training loss consists of masked loss  $\mathcal{L}_{\text{masked}}$  and depth consistency loss  $\mathcal{L}_{\text{con}}$  as:

$$\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{masked}} + \lambda \mathcal{L}_{\text{con}}. \quad (7)$$

where  $\lambda$  is empirically set to 4 in our experiments.

## 4. Experiments

### 4.1. Datasets and Metrics

We train recent SOTA depth refinement models and our PRO model on a synthetic dataset (UnrealStereo4K) to get the advantage of dense GT depth for a fair comparison. At test time, we evaluate its zero-shot performance on four real datasets to demonstrate generalizability.

**UnrealStereo4K.** The UnrealStereo4K dataset [38] is a synthetic dataset comprising stereo images with 4K resolution ( $2,160 \times 3,840$ ) and pixel-wise GT disparity labels. As aforementioned, its some GT labels are inaccurate especially for transparent objects, as shown in Fig. 3.

**Booster.** The Booster dataset [29] has high-resolution ( $3,008 \times 4,112$ ) indoor images with specular and transparent surfaces. We use the training set with GT depth for testing. To ensure diversity, we remove images captured from the same scenes, resulting in a test set of 38 images.

**ETH3D.** The ETH3D dataset [35] contains high-resolution indoor and outdoor images ( $6,048 \times 4,032$ ) with GT depth captured by LiDAR sensors. In the dataset, 34 images were rotated 90 degrees counterclockwise, causing the bottom of the scenes to appear on the right. Before testing, we rotate these images 90 degrees clockwise to ensure that the floor sides are correctly positioned at the bottom sides.

**Middlebury 2014.** The Middlebury 2014 dataset [34] consists of 23 high-resolution (nearly 4K) indoor images with dense ground truth disparity maps.

**NuScenes.** The NuScenes dataset [4] is an autonomous driving dataset that provides multi-directional camera images along with depth information obtained from LiDAR and radar sensors. Among outdoor datasets, the Cityscapes dataset [6] provides higher resolution but suffers from blurred edges since its disparity maps are generated us-

Metric	PatchFusion [23]	PatchRefiner [24]	PRO (Ours)
CE↓	0.364	0.347	<b>0.049</b>

Table 2. Comparison of PatchFusion [23], PatchRefiner [24], and PRO (Ours) in terms of Consistency Error (CE).

ing a stereo matching algorithm instead of LiDAR sensors. Therefore, we use the NuScenes dataset for testing.

**DIS-5K.** The DIS-5K dataset [28] is a high-resolution image segmentation dataset that provides accurately annotated masks, making it suitable for evaluating edge accuracy.

**Evaluation Metrics.** We use the commonly adopted metrics for DE, Absolute Relative Error (AbsRel) and  $\delta_1$ . Additionally, to evaluate the edge quality, we adopt D<sup>3</sup>R metric on the Middlebury 2014 dataset following [26], and measure Boundary Recall (BR) metric on the DIS-5K dataset following [3]. We utilize Consistency Error (CE) [23] to evaluate the consistency of the DE results across patches.

### 4.2. Implementation Details

We adopt the pretrained DepthAnythingV2 (DA2) [47] as the baseline model ( $\Psi$  in Fig. 4). PRO, PatchFusion [23], and PatchRefiner [24] are (re)trained based on DA2, while BoostingDepth [26] uses DA2 without additional training. The input resolution to  $\Psi$  is fixed at  $518 \times 518$ . When we use the GPCT strategy, adjacent patches are cropped with an overlap of 224 pixels, ensuring an overlap of  $h \times 224$  for horizontally adjacent patches and  $224 \times w$  for vertically adjacent patches, where  $h$  and  $w$  denote the height and width of the patch, respectively. In our BFM, the dilation kernel size is empirically set to (10, 20). Training is conducted on 7,592 UnrealStereo4K samples for 8 epochs with a batch size of 64 (with gradient accumulation), taking 10 hours on a single RTX 4090 GPU. Additional details are described in *supplementary material*.

### 4.3. Performance Comparison

At test time, each input image is divided into a  $4 \times 4$  grid of patches. The depth of each patch is refined individually, and the refined patches are then reassembled to generate the final depth map.

**Zero-shot performance on high-resolution datasets.** We evaluate four patch-wise DE models: three are depth refinement models including our PRO model, BoostingDepth

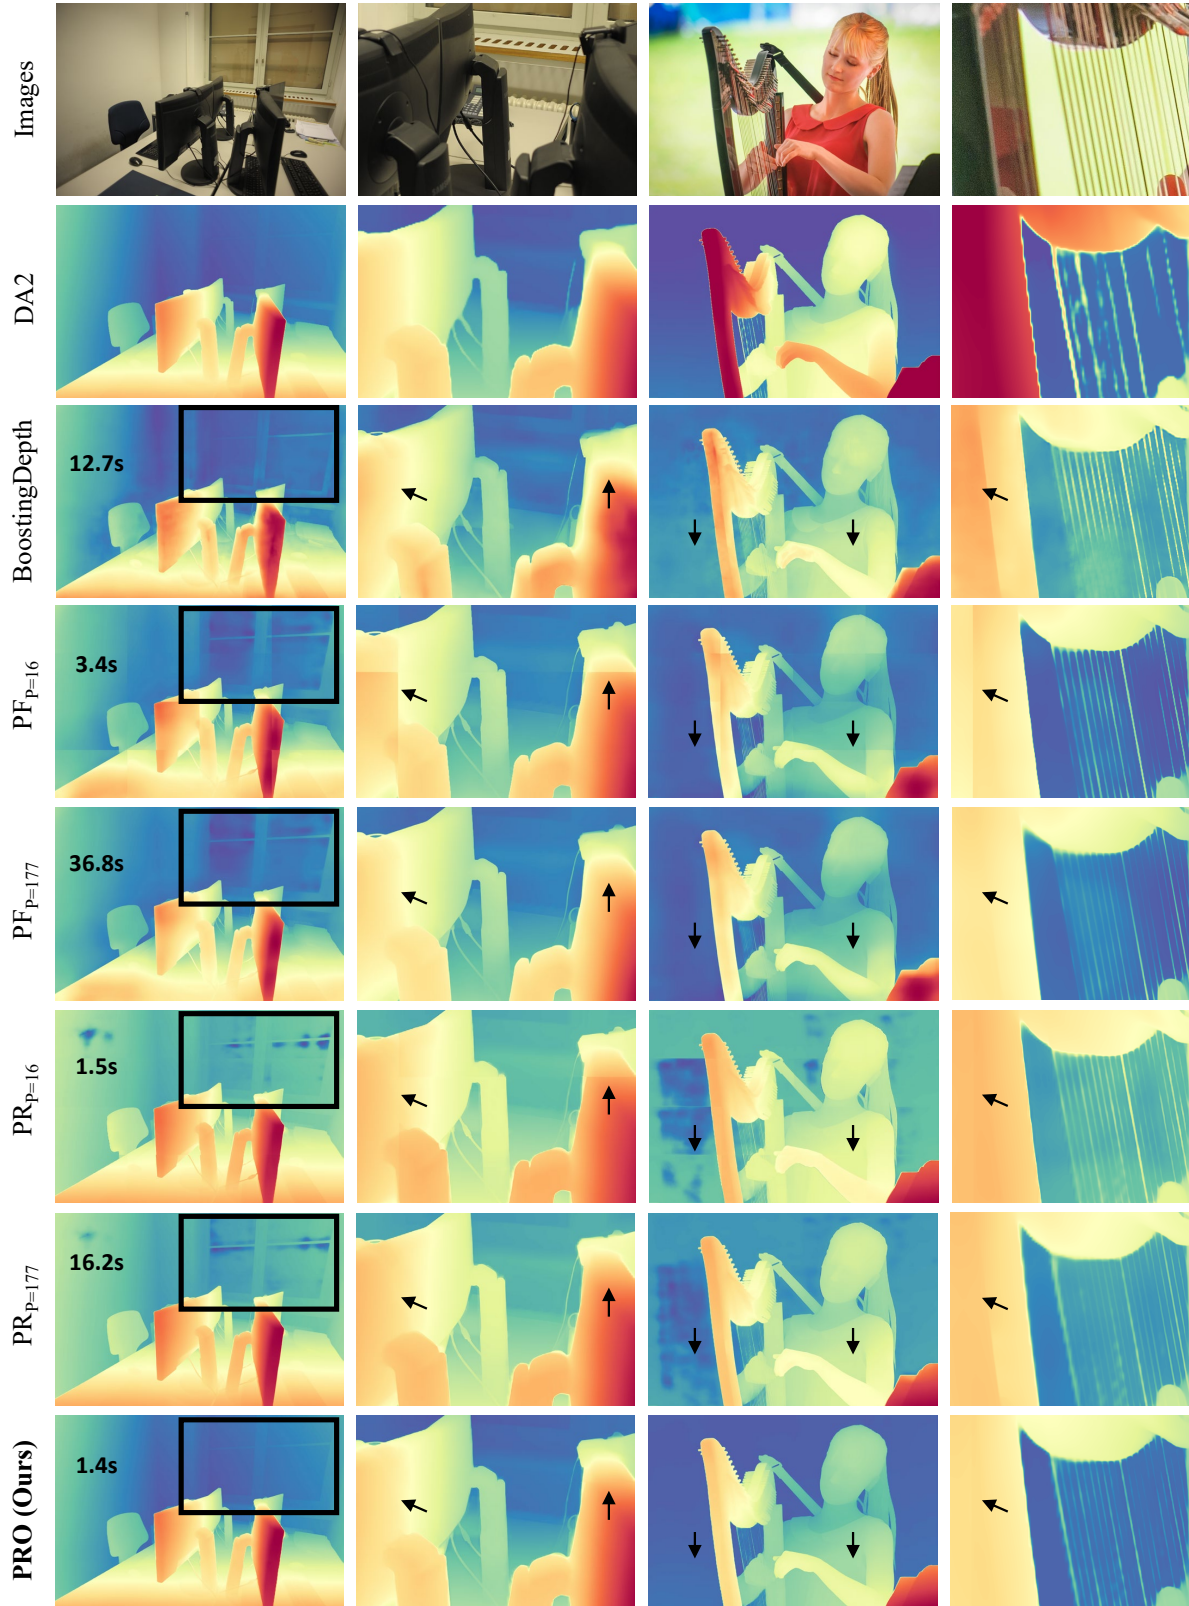


Figure 5. **Qualitative comparisons for patch-wise DE methods on ETH3D [35] and DIS-5K [28].** We compare PRO (Ours) with BoostingDepth [26], PatchFusion (PF) [23], and PatchRefiner (PR) [24]. The time displayed in the leftmost depth column represents the inference time. Black rectangle area represents the transparent object region. Black arrows indicate patch boundaries. Zoom in for details.

Model	GPCT	BFM	Booster		ETH3D		Middle14			NuScenes		CE↓
			AbsRel↓	$\delta_1$ ↑	AbsRel↓	$\delta_1$ ↑	AbsRel↓	$\delta_1$ ↑	D <sup>3</sup> R ↓	AbsRel↓	$\delta_1$ ↑	
(a)			0.0385	0.987	0.0428	<b>0.985</b>	0.0292	0.995	0.0807	0.105	0.882	0.208
(b)		✓	<b>0.0303</b>	<b>0.994</b>	0.0426	<b>0.985</b>	0.0290	0.995	0.0837	0.105	0.882	0.117
(c)	✓		0.0313	0.993	0.0425	<b>0.985</b>	0.0288	<b>0.996</b>	<b>0.0792</b>	0.105	0.882	0.058
(d)	✓	✓	0.0304	<b>0.994</b>	<b>0.0422</b>	<b>0.985</b>	<b>0.0287</b>	<b>0.996</b>	0.0803	<b>0.104</b>	<b>0.883</b>	<b>0.049</b>

Table 3. Ablation on GPCT and BFM of our PRO for Booster [29], ETH3D [35], Middlebury 2014 [34], and NuScenes [4] datasets.

[26], PatchRefiner [24], and the other one is a direct DE model that is PatchFusion [23]. For PatchFusion and PatchRefiner, we conduct evaluations under two settings: (i) ‘P=16’ setting that uses the same patch number as our method, and (ii) ‘P=177’ setting, where additional patches are used for test-time ensembling. Note that P denotes the number of patches used to reassemble final depth maps during inference. As shown in Table 1, our PRO achieves SOTA performance across all depth metrics while maintaining the lowest inference time. Notably, our BFM prevents overfitting to synthetic dataset biases, leading to a 9.5% improvement in AbsRel on the Booster dataset that contains transparent objects. In terms of edge quality, our method performs worse than BoostingDepth and PatchFusion in Boundary Recall (BR). However, as shown in the depth metrics, these two methods prioritize enhancing edge sharpness rather than focusing on enhancing overall depth accuracy. In contrast, when compared to PatchRefiner that exhibits similar depth performance, our PRO achieves superior results in terms of the edge metric (BR). To further demonstrate the generalizability of our method, we include additional experiments using alternative base models such as MiDaS v3.1 [2] and DepthAnythingV1 [46], as detailed in the *supplementary material*.

In terms of consistency, our PRO achieves the best result, with an 85.9% improvement, as shown in Table 2. It demonstrates the effectiveness of our GPCT strategy in maintaining consistency between individually processed patches.

**Qualitative Comparisons.** Fig. 5 shows qualitative comparisons for four patch-wise DE models. As shown in Fig. 5, PatchFusion and PatchRefiner trained on the synthetic dataset without masking exhibit artifacts in the window regions. Additionally, two patch-based methods [23, 24] produce noticeable boundary artifacts when test ensembling is not applied. In contrast, despite refining each patch only once, our PRO demonstrates minimal depth discontinuity while achieving the fastest inference time.

#### 4.4. Ablation Studies

We analyze the effectiveness of the core components of our PRO through ablation studies: Grouped Patch Consistency Training (GPCT) and Bias Free Masking (BFM).

**Ablation on GPCT and BFM.** Table 3 presents the ablation results on GPCT and BFM. When BFM is applied to the baseline (a), the improvements are relatively minor for most datasets. However, for the Booster dataset [29] that contains transparent objects, we observe a significant im-

Overlap	CE↓	ETH3D	Middle14
		AbsRel↓	AbsRel↓
28	0.108	0.0423	0.0288
56	0.113	0.0424	0.0289
112	0.065	0.0425	0.0288
224	<b>0.049</b>	<b>0.0422</b>	<b>0.0287</b>
448	0.060	0.0423	<b>0.0287</b>

Table 4. Ablations on overlap sizes in Consistency Error (CE) for ETH3D [35] and Middlebury 2014 [34] datasets.

provement of 21.3% in AbsRel metric, demonstrating that our BFM method effectively identifies unreliable regions. Even when only BFM is applied without GPCT, a reduction in consistency error (CE) suggests that supervision limited to reliable regions helps mitigate unnecessary refinement, thereby reducing confusion in the depth refinement process. Furthermore, applying GPCT to the baseline yields a 50.4% improvement in consistency error, demonstrating its effectiveness in maintaining consistency between independently processed patches. Although model (c) achieves the best performance in the D<sup>3</sup>R metric, it introduces artifacts in transparent regions. Consequently, model (d), which incorporates both GPCT and BFM, achieves the best overall performance across most metrics.

**Ablation on Overlap Sizes in GPCT.** Table 4 shows the effect of different overlap sizes in GPCT. Each overlap value denotes the size (in pixels) of the square overlapping region shared between adjacent patches. As shown, increasing the overlap generally reduces the CE. However, when the overlap reaches 448 pixels, CE begins to increase. We attribute this to the fact that highly overlapped patches capture nearly identical scene contents, resulting in fewer inconsistencies between patches. So, the depth consistency loss provides less meaningful supervision, reducing its effectiveness.

## 5. Conclusion

In this paper, we propose the Patch Refine Once (PRO) model for depth refinement on high-resolution images. To address the depth discontinuity problem, we introduce the Grouped Patch Consistency Training (GPCT) strategy, ensuring consistency between independently processed patches. In addition, we propose Bias Free Masking (BFM) to prevent the depth refinement model from overfitting to dataset-specific biases. Through these strategies, our PRO achieves superior zero-shot performance while maintaining a low inference time (12× faster than PatchRefiner with 177 patches), outperforming recent state-of-the-art methods across diverse datasets.

**Acknowledgements.** This work was supported by IITP grant funded by the Korea government(MSIT) (No.RS2022-00144444, Deep Learning Based Visual Representational Learning and Rendering of Static and Dynamic Scenes).

## References

- [1] 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. 3
- [2] Reiner Birkel, Diana Wofk, and Matthias Müller. Midas v3. 1—a model zoo for robust monocular relative depth estimation. *arXiv preprint arXiv:2307.14460*, 2023. 2, 3, 8, 1
- [3] 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. *arXiv preprint arXiv:2410.02073*, 2024. 6
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 3, 6, 8, 2
- [5] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8628–8638, 2021. 3
- [6] 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 *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [7] Yaqiao Dai, Renjiao Yi, Chenyang Zhu, Hongjun He, and Kai Xu. Multi-resolution monocular depth map fusion by self-supervised gradient-based composition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 488–496, 2023. 2, 3
- [8] Ingrid Daubechies. The wavelet transform, time-frequency localization and signal analysis. *IEEE transactions on information theory*, 36(5):961–1005, 1990. 4
- [9] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10786–10796, 2021. 2, 3
- [10] David Eigen, Christian Puhres, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014. 2, 3
- [11] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In *European Conference on Computer Vision*, pages 241–258. Springer, 2024. 2, 3
- [12] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11):1231–1237, 2013. 3
- [13] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 270–279, 2017. 2
- [14] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3828–3838, 2019. 2
- [15] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Rantotas, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2485–2494, 2020. 3
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 4
- [17] Yiwen Hua, Puneet Kohli, Pritish Uplavikar, Anand Ravi, Saravana Gunaseelan, Jason Orozco, and Edward Li. Holopix50k: A large-scale in-the-wild stereo image dataset. *arXiv preprint arXiv:2003.11172*, 2020. 2
- [18] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2821–2830, 2018. 2
- [19] 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. 2, 3
- [20] Tobias Koch, Lukas Liebel, Friedrich Fraundorfer, and Marco Korner. Evaluation of cnn-based single-image depth estimation methods. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 2, 3
- [21] Jiaqi Li, Yiran Wang, Jinghong Zheng, Zihao Huang, Ke Xian, Zhiguo Cao, and Jianming Zhang. Self-distilled depth refinement with noisy poisson fusion. *Advances in Neural Information Processing Systems*, 37:69999–70025, 2025. 2, 3
- [22] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018. 3, 5
- [23] Zhenyu Li, Shariq Farooq Bhat, and Peter Wonka. Patch-fusion: An end-to-end tile-based framework for high-resolution monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*

- Pattern Recognition*, pages 10016–10025, 2024. 1, 2, 3, 4, 6, 7, 8
- [24] Zhenyu Li, Shariq Farooq Bhat, and Peter Wonka. Patchrefiner: Leveraging synthetic data for real-domain high-resolution monocular metric depth estimation. In *European Conference on Computer Vision*, pages 250–267. Springer, 2024. 1, 2, 3, 4, 6, 7, 8
- [25] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. Binsformer: Revisiting adaptive bins for monocular depth estimation. *IEEE Transactions on Image Processing*, 2024. 3
- [26] S Mahdi H Miangoleh, Sebastian Dille, Long Mai, Sylvain Paris, and Yagiz Aksoy. Boosting monocular depth estimation models to high-resolution via content-adaptive multi-resolution merging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9685–9694, 2021. 2, 3, 6, 7, 8
- [27] Jaeho Moon, Juan Luis Gonzalez Bello, Byeongjun Kwon, and Munchurl Kim. From-ground-to-objects: Coarse-to-fine self-supervised monocular depth estimation of dynamic objects with ground contact prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10519–10529, 2024. 2
- [28] Xuebin Qin, Hang Dai, Xiaobin Hu, Deng-Ping Fan, Ling Shao, and Luc Van Gool. Highly accurate dichotomous image segmentation. In *European Conference on Computer Vision*, pages 38–56. Springer, 2022. 6, 7, 1, 2, 4, 5
- [29] Pierluigi Zama Ramirez, Alex Costanzino, Fabio Tosi, Matteo Poggi, Samuele Salti, Stefano Mattoccia, and Luigi Di Stefano. Booster: a benchmark for depth from images of specular and transparent surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(1):85–102, 2023. 3, 6, 8, 1, 2, 4, 5
- [30] 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. 2, 3
- [31] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021. 2, 3
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [33] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE transactions on pattern analysis and machine intelligence*, 31(5):824–840, 2008. 2, 3
- [34] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nešić, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *Pattern Recognition: 36th German Conference, GCPR 2014, Münster, Germany, September 2-5, 2014, Proceedings 36*, pages 31–42. Springer, 2014. 2, 3, 6, 8, 1
- [35] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3260–3269, 2017. 3, 6, 7, 8, 1, 2, 4, 5
- [36] Shuwei Shao, Zhongcai Pei, Weihai Chen, Xingming Wu, and Zhengguo Li. Ndddepth: Normal-distance assisted monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7931–7940, 2023. 2
- [37] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d 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. 3
- [38] Fabio Tosi, Yiyi Liao, Carolin Schmitt, and Andreas Geiger. Smd-nets: Stereo mixture density networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8942–8952, 2021. 2, 3, 4, 6
- [39] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web stereo video supervision for depth prediction from dynamic scenes. In *2019 International Conference on 3D Vision (3DV)*, pages 348–357. IEEE, 2019. 2
- [40] 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 *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2021. 2
- [41] 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
- [42] Yiran Wang, Min Shi, Jiaqi Li, Zihao Huang, Zhiguo Cao, Jianming Zhang, Ke Xian, and Guosheng Lin. Neural video depth stabilizer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9466–9476, 2023. 2
- [43] Jamie Watson, Oisín Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1164–1174, 2021. 2
- [44] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, and Zhenbo Luo. Monocular relative depth perception with web stereo data supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 311–320, 2018. 3
- [45] Chenxi Xie, Changqun Xia, Mingcan Ma, Zhirui Zhao, Xiaowu Chen, and Jia Li. Pyramid grafting network for one-stage high resolution saliency detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11717–11726, 2022. 2, 3, 6

- [46] 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. [2](#), [3](#), [8](#), [1](#)
- [47] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *Advances in Neural Information Processing Systems*, 37:21875–21911, 2025. [2](#), [3](#), [6](#), [1](#)
- [48] Wei Yin, Xinlong Wang, Chunhua Shen, Yifan Liu, Zhi Tian, Songcen Xu, Changming Sun, and Dou Renyin. Diversedepth: Affine-invariant depth prediction using diverse data. *arXiv preprint arXiv:2002.00569*, 2020. [3](#)
- [49] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. Neural window fully-connected crfs for monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3916–3925, 2022. [3](#)