

## Layered Depth Refinement with Mask Guidance

Soo Ye Kim<sup>1</sup> Jianming Zhang<sup>2</sup> Simon Niklaus<sup>2</sup> Yifei Fan<sup>2</sup>  
Simon Chen<sup>2</sup> Zhe Lin<sup>2</sup> Munchurl Kim<sup>1</sup>  
<sup>1</sup>KAIST, Republic of Korea <sup>2</sup>Adobe Inc., USA

### Abstract

Depth maps are used in a wide range of applications from 3D rendering to 2D image effects such as Bokeh. However, those predicted by single image depth estimation (SIDE) models often fail to capture isolated holes in objects and/or have inaccurate boundary regions. Meanwhile, high-quality masks are much easier to obtain, using commercial auto-masking tools or off-the-shelf methods of segmentation and matting or even by manual editing. Hence, in this paper, we formulate a novel problem of mask-guided depth refinement that utilizes a generic mask to refine the depth prediction of SIDE models. Our framework performs layered refinement and inpainting/outpainting, decomposing the depth map into two separate layers signified by the mask and the inverse mask. As datasets with both depth and mask annotations are scarce, we propose a self-supervised learning scheme that uses arbitrary masks and RGB-D datasets. We empirically show that our method is robust to different types of masks and initial depth predictions, accurately refining depth values in inner and outer mask boundary regions. We further analyze our model with an ablation study and demonstrate results on real applications. More information can be found on our project page.<sup>1</sup>

### 1. Introduction

Recent progress in deep learning has enabled the prediction of fairly reliable depth maps from single RGB images [20, 31, 32, 47]. However, despite the specialized network architectures [11, 29, 31] and training strategies [32, 46] in single image depth estimation (SIDE) models, the estimated depth maps are still inadequate in the following aspects: (i) depth boundaries tend to be blurry and inaccurate; (ii) thin structures such as poles and wires are often missing; and (iii) depth values in narrow or isolated background regions (e.g., between body parts in humans) are often imprecise, as shown in the initial depth estimation in Figure 1. Addressing these issues within a single SIDE model can be

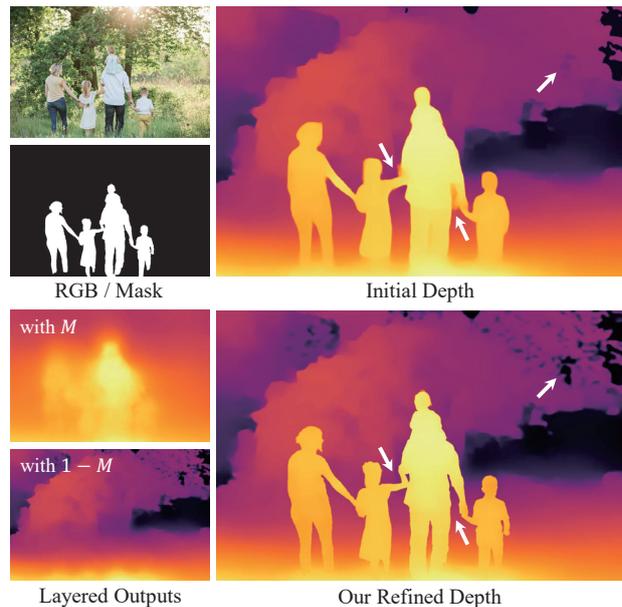


Figure 1. Our layered depth refinement result on an initial prediction by DPT [31]. Aided by a high-quality mask  $M$  generated with an auto-masking tool [33], our method is able to accurately refine mask boundaries and correct depth values in isolated hole regions between body parts. Regions in  $M$  and  $1 - M$  are refined and inpainted/outpainted separately with our layered approach.

very challenging due to limited model capacity and the lack of high-quality RGB-D datasets.

Therefore, we take a novel approach of utilizing an additional cue of a *high-quality mask* to refine depth maps predicted by SIDE methods. The provided mask can be hard (binary) or soft (e.g., from matting) and can be of objects or other parts of the image such as the sky. As high-quality auto-masking tools are very accessible nowadays, such masks can be easily obtained with commercial tools (e.g., removebg [33] or Photoshop) or off-the-shelf segmentation models [14, 30, 52, 57]. Segmentation masks can also be annotated by humans [7, 41, 49], and accurate datasets are easier to obtain than RGB-D data, which facilitates the training of auto-masking models.

<sup>1</sup><https://sooyekim.github.io/MaskDepth/>

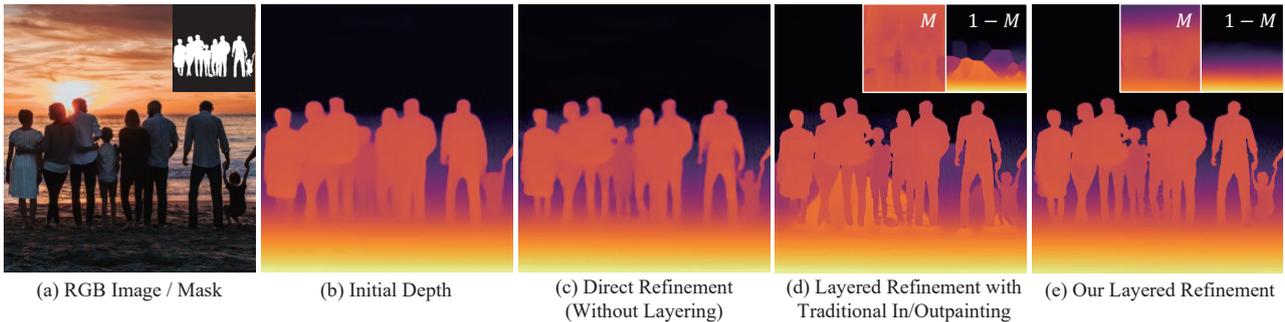


Figure 2. Refined depth maps with the guidance of a high-quality mask. (b) The initial depth prediction [31] has blurry boundaries and misses isolated hole regions between human body parts. (c) Direct refinement by training on a paired dataset [34] improves the initial depth but still has blurry boundaries. Layered refinement results in sharp edges due to the final compositing step using the mask, although (d) naive in/outpainting [36] generates artifacts in the background. (e) Our method successfully corrects the inaccurate depth values while in/outpainting each region with the guidance of the mask. Intermediate layered outputs are shown on the top right for the layered models.

However, even with such accurate masks, how to effectively train the depth refinement model remains an open issue. As shown in Figure 2(c), directly adding the mask as an input channel to the refinement model still results in blurrier boundaries than the given mask. Therefore, we propose a *layered refinement* strategy: The mask ( $M$ ) and inverse mask ( $1 - M$ ) regions are processed separately to interpolate or extrapolate the depth values beyond the mask boundary, leading to two layers of depth maps. As shown in Figure 2(e), the refined output is the composite of the two layers using the mask  $M$ , which fully preserves the boundary details of the mask, as well as filling in the correct depth values for the isolated background regions.

A naïve baseline for layered depth refinement would be using an off-the-shelf inpainting method to generate the depth map layers for  $M$  and  $1 - M$ . Unfortunately, as shown in Figure 2(d), generic inpainting may not work well for filling in large holes in a depth map. Moreover, deriving an appropriate region for hole-filling on an imperfect initial depth prediction based on the mask is a non-trivial problem. The hole-filling region often needs to be expanded to cover uncertain regions along the mask boundary, as otherwise, the erroneous depth values may propagate in the hole. However, too much expansion will make the hole-filling task much more challenging as it may overwrite the original depth structure in the scene (see the  $1 - M$  layer in Figure 2(d)).

To address the challenge, we propose a framework for degradation-aware layered depth completion and refinement, which learns to identify and correct inaccurate regions based on the context of the mask and the image. Our framework does not require additional input or heuristics to expand the hole-filling region. Furthermore, we devise a *self-supervised* learning scheme that uses RGB-D training data without paired mask annotations. We demonstrate that our method is robust under various conditions by empirically validating it on synthetic datasets and real images *in the wild*. We further

provide results on real-world downstream applications.

Our contributions are three-fold:

- We propose a novel mask-guided depth refinement framework that refines the depth estimations of SIDE models guided by a generic high-quality mask.
- We propose a novel *layered refinement* approach, generating sharp and accurate results in challenging areas without additional input or heuristics.
- We devise a self-supervised learning scheme that uses RGB-D training data *without* paired mask annotations.

## 2. Related Work

**Single Image Depth Estimation** Single image depth estimation (SIDE), also commonly termed monocular depth estimation, aims to predict a depth map from an RGB image. A common approach is to train a deep neural network on RGB-D datasets to learn the non-linear mapping from RGB to depth [20, 31, 32, 47]. As for the model architecture, convolutional neural networks (CNNs) are a popular choice [32, 47], with a transformer-based model [31] being recently proposed to overcome the limited receptive field size of CNNs. Transformer models [10, 37] leverage self-attention [39], expanding the receptive field to the entire image at every level. We also base our model architecture on transformers to benefit from the enlarged receptive field.

For training SIDE models, datasets are often augmented with synthetic datasets [4, 27, 43, 44, 50] and relative depths computed from stereo images [20, 40, 46]. Numerous supervision schemes [1, 5, 12, 13, 24, 26, 45, 53, 55, 56] and loss functions [17, 19, 20, 47] have been proposed to optimize the model training for SIDE. Several methods [26, 42, 56] attempt to exploit the relation between image segmentation and SIDE, with Zhu *et al.* [56] proposing regularizing depth boundaries with segmentation map boundaries in the loss function to enforce sharper edges in the resulting depth maps. However, even with sophisticated framework designs, cap-

turing highly accurate depth boundaries remains a challenge due to the ill-posed nature of the problem and the lack of pixel-perfect ground truth depth data.

**Depth Inpainting** Inpainting depth maps is often necessary in novel view synthesis for 3D photography to naturally fill in disoccluded regions [16, 27, 35]. Such methods apply joint RGB and depth inpainting in the background region near object edges. Another line of research is depth completion, which aims to fill in unknown depth values from sparsely known annotations. Imran *et al.* [15] proposed a layered approach, extrapolating foreground and background regions separately from LiDAR data. In our depth refinement method, both the mask and inverse mask regions are inpainted/outpainted *while* correcting inaccurate depth values and merged afterward to obtain accurate boundaries.

**Depth Refinement** In an inspirational work [25], Miangoleh *et al.* proposed boosting high-frequency details in SIDE results by merging multiple depth predictions at various resolutions, exploiting the limited receptive field size of CNNs. However, their merging algorithm tends to generate inconsistent depth values in foreground objects, and its refinement degrades with recent transformer architectures as it is based on a fundamental assumption related to CNNs. Furthermore, capturing very thin boundaries and generating accurate depth values in hole regions are still challenging.

In this paper, we explore a novel direction of using generic masks as guidance for depth refinement. Unlike previous methods that upscale or enhance details in the entire depth map, we focus on delicately refining along the boundary and hole regions of the mask. Handling such regions is often important in downstream applications such as Bokeh effect synthesis. Our method is generic and can refine depth maps generated by any SIDE model regardless of the model architecture, as long as the provided mask contains better boundaries than the initial depth map. Note that our method operates in the inverse depth space as many prior works [25, 31, 32], although we continue using the term depth.

### 3. Proposed Method

We propose a layered depth refinement framework for enhancing the initial depth prediction of SIDE models using the guidance of a quasi-accurate mask and an RGB image.

#### 3.1. Data Generation

**Random composition** With an RGB-D dataset consisting of an RGB image  $I$  and its depth map  $D$ , a general depth refinement model can be optimized in a self-supervised way by applying random perturbations  $\mathcal{P}$  on  $D$ , which inversely simulate initial depth predictions. A neural network  $\mathcal{R}$  can then be trained to predict the refined depth map  $\hat{D} = \mathcal{R}(\mathcal{P}(D), I)$  with an appropriate loss function  $\mathcal{L}(\hat{D}, D)$ .

However, collecting a dataset for training a *mask-guided*

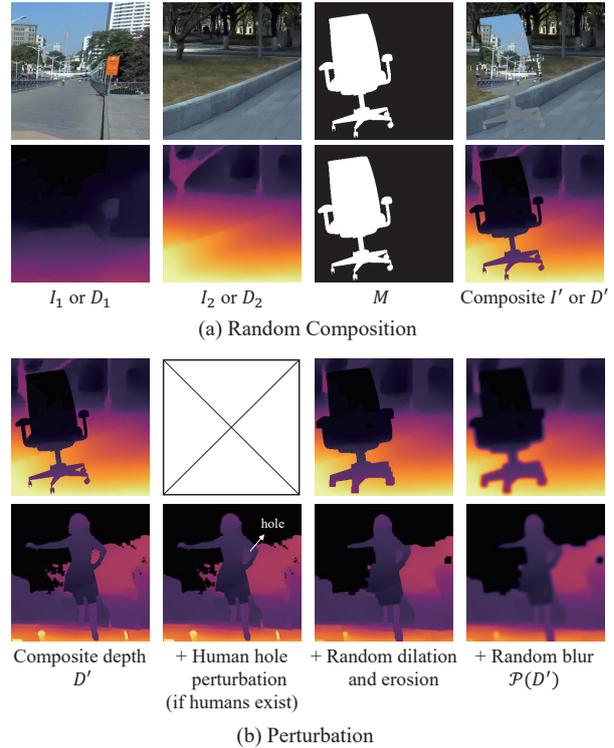


Figure 3. Data generation scheme. RGB-D patches are randomly composited using an arbitrary binary mask. Perturbations are applied to simulate depth estimates, resulting in isolated regions being covered up and thin structures being lost.

depth refinement model is challenging as datasets containing masks along with the RGB-D information are scarce. Hence, we devise a data generation scheme that does *not* require paired depth and mask annotations. Specifically, a composite depth map  $D'$  is randomly synthesized from two arbitrary depth maps  $D_1$  and  $D_2$  using an arbitrary binary mask  $M$  with  $m_{ij} \in \{0, 1\}$ , by  $D' = M \cdot D_1 + (1 - M) \cdot D_2$ . Likewise, the corresponding composite RGB image  $I'$  is computed by  $I' = M \cdot I_1 + (1 - M) \cdot I_2$ , where  $I_1$  and  $I_2$  are the RGB images corresponding to  $D_1$  and  $D_2$ , respectively. Examples of  $D'$  and  $I'$  are shown in Figure 3(a). Applying perturbations to  $D'$  leads to  $\mathcal{P}(D')$ , and the mask-guided refinement model  $\mathcal{R}_m$  can then be trained with  $\mathcal{L}(\hat{D}', D')$ , where  $\hat{D}' = \mathcal{R}_m(\mathcal{P}(D'), I', M)$ .

In this way, we can obtain a synthesized depth map  $D'$  and an RGB image  $I'$  that are aligned to  $M$  from any RGB-D dataset and arbitrary masks. Diverse types of masks can be mixed and used, including object and stuff masks from segmentation datasets [21, 54]. Furthermore, we can effortlessly acquire the ground truths for inpainting/outpainting ( $D_1$  and  $D_2$ ), which are essential for our layered refinement approach, explained in more detail in the next section.

**Perturbations** As shown in Figure 3(b), we apply three

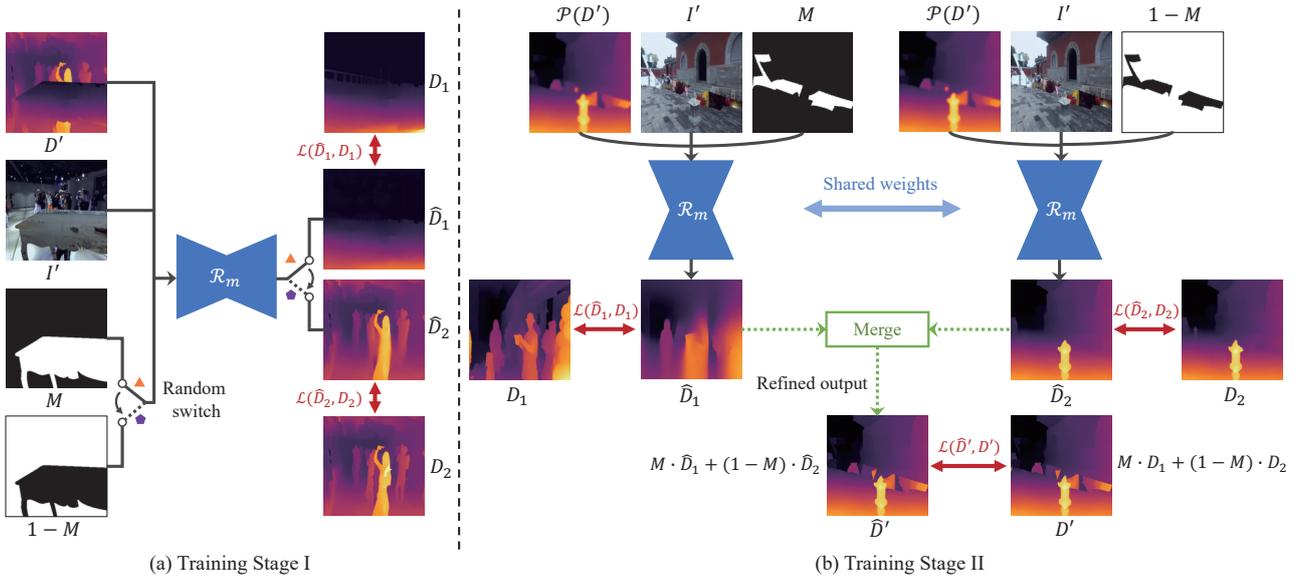


Figure 4. An overview of the proposed two-stage training strategy. In the first stage, the network is trained to complete regions with 0 based on regions with 1 in the given mask. The mask is randomly flipped and the corresponding depth ( $D_1$  or  $D_2$ ) is given as the ground truth. In stage II, we run the network twice to obtain  $\hat{D}_1$  and  $\hat{D}_2$  and merge them based on the mask to produce the refined output  $\hat{D}'$ . The network learns to remove perturbations while inpainting/outpainting the depth. During inference, the refined output is obtained following stage II.

types of perturbations on  $D'$  to simulate typical inaccuracies in SIDE model predictions. First, random dilation and erosion are applied in a random order so that the perturbed depth map lacks thin structures, and its depth boundaries are not always aligned with the RGB image or the mask. In Figure 3(b), it can be observed that thin structures (hand of the person) are lost, and isolated regions are covered up (between the arm and the main frame of the chair) after random dilation and erosion. Second, we apply random amounts of Gaussian blur on the depth map as estimated depth maps tend to have blurry boundaries. Lastly, we design a *human hole perturbation scheme* that detects isolated regions and assigns a random value between the mean depth values surrounding the hole and inside the original hole, simulating the often-missing isolated regions inside human bodies in estimated depth maps. More details of the perturbation scheme are provided in the supplementary material.

### 3.2. Training Strategy

**Two-stage training for layered refinement** Although depth refinement with an accurate mask may appear straightforward after data pairs are obtained with the proposed data generation scheme, directly predicting the refined depth map from concatenated RGB-D and mask inputs leads to sub-optimal results, as shown in Figure 2. To explicitly benefit from the accurate mask, we propose a *layered refinement* approach that refines regions specified by  $M$  and  $1 - M$  separately and merges two individual results based on  $M$ . In this way, the model can focus on correcting the depth values

in each region, and mask boundaries can be fully preserved after the merging stage.

We train our model in two stages shown in Figure 4. In the first stage, the model  $\mathcal{R}_m$  is trained for image completion by randomly providing  $M$  or  $1 - M$  and optimizing either  $\mathcal{L}(\mathcal{R}_m(D', I', M), D_1)$  or  $\mathcal{L}(\mathcal{R}_m(D', I', 1 - M), D_2)$ . Note that a single model is trained for both inpainting and outpainting the depth input to always complete regions with 0 based on regions with 1 signified by the given mask  $M$  or  $1 - M$ . Then in the second stage, we add perturbations  $\mathcal{P}$  and run the network twice with  $M$  and  $1 - M$  to obtain two outputs  $\hat{D}_1$  and  $\hat{D}_2$ , given by

$$\hat{D}_1 = \mathcal{R}_m(\mathcal{P}(D'), I', M) \quad \text{and} \quad (1)$$

$$\hat{D}_2 = \mathcal{R}_m(\mathcal{P}(D'), I', 1 - M). \quad (2)$$

Reasonable  $\hat{D}_1$  and  $\hat{D}_2$  are generated from the beginning of the second stage as the model has been pretrained for inpainting/outpainting in the first stage. Finally,  $\hat{D}_1$  and  $\hat{D}_2$  are merged to yield the refined output  $\hat{D}'$  as follows:

$$\hat{D}' = M \cdot \hat{D}_1 + (1 - M) \cdot \hat{D}_2. \quad (3)$$

Our model is optimized with three losses at this stage:  $\mathcal{L}(\hat{D}_1, D_1)$ ,  $\mathcal{L}(\hat{D}_2, D_2)$ , and  $\mathcal{L}(\hat{D}', D')$ . As a result, the network learns to remove perturbations while generating completed depth maps under a *unified* framework. Although we only utilize composite depth maps as input during training, the randomness in composition (*random* depth maps

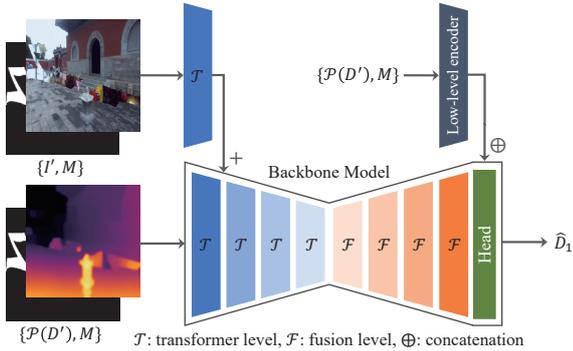


Figure 5. Our network architecture with DPT [31] as the backbone model. We add a low-level encoder and a branch for the RGB input.

composed with a *random* mask) and random perturbations lead to a robust model that generalizes well to real depth estimations and diverse masks.

**Loss function** The loss  $\mathcal{L}$  is comprised of three different loss terms summed with unit scale: L1 loss, L2 loss, and a multi-scale gradient loss with four scale levels [20]. The gradient loss is adopted to enforce sharp depth boundaries.

### 3.3. Model Architecture

We base our model architecture on the dense prediction transformer (DPT) [31] with four transformer encoder levels [10]  $l \in \{1, 2, 3, 4\}$  and four fusion decoder levels. At each encoder level, overlapping patches are extracted and embedded to dimensions  $d_l \in \{64, 128, 320, 512\}$  and fed into  $t_l \in \{3, 4, 18, 3\}$  transformer layers each with self-attention, LayerNorm [3] and MLP layers. The spatial resolution is decreased by a scale factor of  $s_l \in \{4, 2, 2, 2\}$  at each level. On the decoder side, features are fused with residual convolutional units at each fusion level, followed by a monocular depth estimation head at the end as in [31].

As shown in Figure 5, we insert an additional encoder branch with a single transformer level to the original backbone so that  $D'$  (or  $\mathcal{P}(D')$ ) and  $M$  (or  $1 - M$ ) are concatenated and fed into the main branch, and  $I'$  concatenated with  $M$  (or  $1 - M$ ) are fed into the additional branch. The outputs are simply summed after the first transformer level. Additionally, a lightweight low-level encoder is introduced to encode the low-level features of the input depth map. These features are concatenated with the features from the main decoder branch and entered into the head, ensuring that the network does not forget the initial depth values.

## 4. Experiments

### 4.1. Implementation Details

We train our model for 500K iterations for the first stage and another 500K iterations for the second stage following

the training strategy described in Sec. 3.2. We used a training patch size of  $320 \times 320$  and a batch size of 32. The model is optimized with AdamW [22] at an initial learning rate of  $10^{-4}$ , which is decreased by 1/10 at 60% and 80% of the total number of iterations. Our model is implemented using PyTorch and trained on 4 NVIDIA V100 GPUs.

For data augmentation, we apply random horizontal flipping and resizing to the input depth maps and RGB images. RGB images are further augmented with random contrast, saturation, brightness, JPEG compression, and grayscale conversions to make our model more robust to various types of inputs. Our model is trained on diverse indoor and outdoor natural RGB-D images, with depth maps scaled to  $[0, 10]$  as in [51] and RGB images normalized using ImageNet [9] mean and standard deviation. Furthermore, to benefit from the proposed self-supervised learning scheme that supports diversifying the types of masks, we sample 50% of masks from diverse object masks, 20% from sky masks and 30% from human masks, where humans with holes are selected 50% of the time (15% of all masks) during training.

### 4.2. Evaluation Datasets

For a quantitative evaluation, datasets with both depth and mask annotations are needed to exclude potential errors caused by inaccurate masking. Furthermore, the ground truth depth should be accurate for reliable evaluations on fine boundaries and object holes. Thus, we use Hypersim (CC-BY SA 3.0 License) [34] and TartanAir (3-Clause BSD License) [44], which are recently released synthetic datasets that contain dense and accurate depth values and also have instance segmentation maps. We select the first frame in each camera trajectory per scene for Hypersim and the 100-th frame for each trajectory in *Easy* difficulty per environment for TartanAir as the test set, which results in 456 images and 206 images in total for Hypersim and TartanAir, respectively. Other datasets such as Cityscapes [8] are not appropriate as the ground truth depth is noisy, often inaccurate around edges and misses thin structures. Additionally, we qualitatively evaluate our refinement method on various freely licensed images from the web [28, 38] with an auto-masking tool [33].

**Zero-shot cross-dataset transfer** We follow the experiment protocol in [32] for evaluation. None of the compared methods or our method have seen the RGB-D images in Hypersim [34] or TartanAir [44] during training. Predictions are scaled and shifted using  $l_2$  minimization to match the ground truth depth.

**Inference using segmentation maps** To use segmentation maps in a mask-guided framework, we take the following steps: (i) compute a binary mask  $M_i$  for each instance  $i$  with more than 1% of the total number of pixels in the instance segmentation map, (ii) run the model  $N$  times with  $M_i$ , and (iii) merge the refined outputs  $\hat{D}_i$  per each pixel by  $\hat{D} = \operatorname{argmax}_{\hat{D}_i} (|D' - \hat{D}_i|)$ , where  $D'$  is initial depth.

Method	Hypersim [34]						TartanAir [44]					
	R <sup>3</sup> ↑	MBE↓	ε <sub>acc</sub> ↓	ε <sub>comp</sub> ↓	WHDR↓	RMSE↓	R <sup>3</sup> ↑	MBE↓	ε <sub>acc</sub> ↓	ε <sub>comp</sub> ↓	WHDR↓	RMSE↓
MiDaS v2.1 [32]	-	0.0973	2.521	7.074	0.1496	0.0966	-	0.0596	3.483	6.913	<b>0.1207</b>	0.0533
+ Direct-composite	3.771	0.0941	1.915	6.233	0.1490	0.0961	5.897	0.0594	3.183	6.363	0.1209	0.0534
+ Direct-paired	-	-	-	-	-	-	3.507	0.0575	3.153	6.304	0.1196	<b>0.0525</b>
+ Layered-propagation	1.097	0.1044	1.942	6.284	0.1629	0.1028	3.642	0.0608	3.128	6.358	0.1255	0.0550
+ Layered-ours	2.332	0.1000	<b>1.871</b>	6.396	0.1560	0.0999	6.939	0.0580	3.243	6.437	0.1230	0.0539
+ Ours (proposed)	<b>5.209</b>	<b>0.0906</b>	1.888	<b>5.931</b>	<b>0.1481</b>	<b>0.0958</b>	<b>16.569</b>	<b>0.0579</b>	<b>2.851</b>	<b>6.272</b>	<b>0.1207</b>	0.0538
DPT-Large [31]	-	0.0936	2.071	6.190	0.1347	0.0911	-	0.0496	2.574	5.677	0.1091	0.0414
+ Direct-composite	2.574	0.0891	1.599	5.411	0.1339	0.0903	4.773	0.0486	2.462	5.480	0.1086	0.0411
+ Direct-paired	-	-	-	-	-	-	2.413	0.0485	2.519	5.394	0.1105	0.0412
+ Layered-propagation	1.188	0.1007	1.792	5.636	0.1502	0.0986	2.347	0.0524	2.579	5.527	0.1162	0.0442
+ Layered-ours	1.996	0.0954	1.606	5.605	0.1433	0.0953	5.626	0.0484	2.447	5.342	0.1116	0.0423
+ Ours (proposed)	<b>4.455</b>	<b>0.0840</b>	<b>1.491</b>	<b>5.087</b>	<b>0.1333</b>	<b>0.0896</b>	<b>8.767</b>	<b>0.0474</b>	<b>2.282</b>	<b>5.245</b>	<b>0.1078</b>	<b>0.0408</b>

Table 1. Quantitative results on Hypersim [34] and TartanAir [44] comparing mask-guided depth refinement models. Best values in **bold**.

### 4.3. Evaluation Metrics

We evaluate the overall error of the output depth maps with the RMSE and the Weighted Human Disagreement Rate (WHDR) [6] measured on 10K randomly sampled point pairs. To evaluate the boundary quality, we report the depth boundary error [18] on accuracy ( $\epsilon_{acc}$ ) and completeness ( $\epsilon_{comp}$ ). In addition, we propose two metrics, mask boundary error (MBE) and relative refinement ratio ( $R^3$ ). All metrics are measured in the inverse depth space.

MBE computes the average RMSE on mask boundary pixels over the  $N$  instances. Mask boundary  $M_i^b$  is obtained by subtracting the eroded  $M_i$  from  $M_i$  and dilating it with a  $5 \times 5$  kernel. The MBE is then given by

$$\text{MBE} = \frac{1}{N} \sum_{i=1}^N \sqrt{\frac{1}{N_i^b} \sum (M_i^b \cdot D - M_i^b \cdot \hat{D})^2}, \quad (4)$$

where  $N_i^b$  is the number of boundary pixels for each instance  $i$ . With  $\epsilon_{acc}$ ,  $\epsilon_{comp}$  and MBE, we can comprehensively measure the boundary accuracy of the refined depth map:  $\epsilon_{acc}$  and  $\epsilon_{comp}$  focusing on *depth* boundaries and MBE on the *mask* boundaries of depth maps. Furthermore, we define  $R^3$  (relative refinement ratio) as the ratio of the number of pixels improved by more than a threshold  $t$  to the number of pixels worsened by more than  $t$ , in terms of absolute error. We set  $t = 0.05$  and compute  $R^3$  of refined results over initial results by base models [31, 32].  $R^3$  is a meaningful indicator for assessing the refinement performance.

### 4.4. Compared Methods

To evaluate the refinement performance, we apply our method to the initial depth predictions of two SIDE models: CNN-based MiDaS v2.1 [32] and SOTA transformer-based DPT-Large [31]. Since there are no existing methods that perform mask-guided depth refinement, we set up the following

baselines using masks for comparison:

- *Direct-composite* produces the refined output without layering and is trained on the same dataset as ours (with composite images and the mask).
- *Direct-paired* also refines without layering but is trained on paired RGB-D and masks in Hypersim [34]. Hence, we only evaluate on TartanAir [44] for this method.
- Layered models (*Layered-propagation* and *Layered-ours*) either apply a propagation-based image completion algorithm [36] or use our model from stage I training, once with the dilated mask for inpainting and the second time with the eroded mask for outpainting. The inpainted/outpainted results are then merged with the mask, similar to our proposed approach.

The network architecture used for *Direct-composite* and *Direct-paired* is the same as our encoder-decoder-style transformer model in Figure 5. For the layered models, we set the dilation and erosion kernel to  $5 \times 5$  for evaluation with segmentation maps. For images *in the wild*, we manually tweak the kernel sizes for each image to obtain the best results.

Additionally, we compare to bilateral median filtering (BMF) with parameters from [35] (previously used for refining depth maps in [23, 35]) and Miangoleh *et al.*'s recent depth refinement method [25]. These approaches do not use masks as guidance. For all compared methods, we use the officially released code and weights.

### 4.5. Analysis

In Table 1, we provide the quantitative results on mask-guided refinement methods. Our method improves both MiDaS v2.1 [32] and DPT-Large [31] on all edge-related metrics ( $\epsilon_{acc}$ ,  $\epsilon_{comp}$  and MBE) and results in high  $R^3$  values of at most 16.569. WHDR and RMSE values are not very discriminative between mask-guided refinement methods

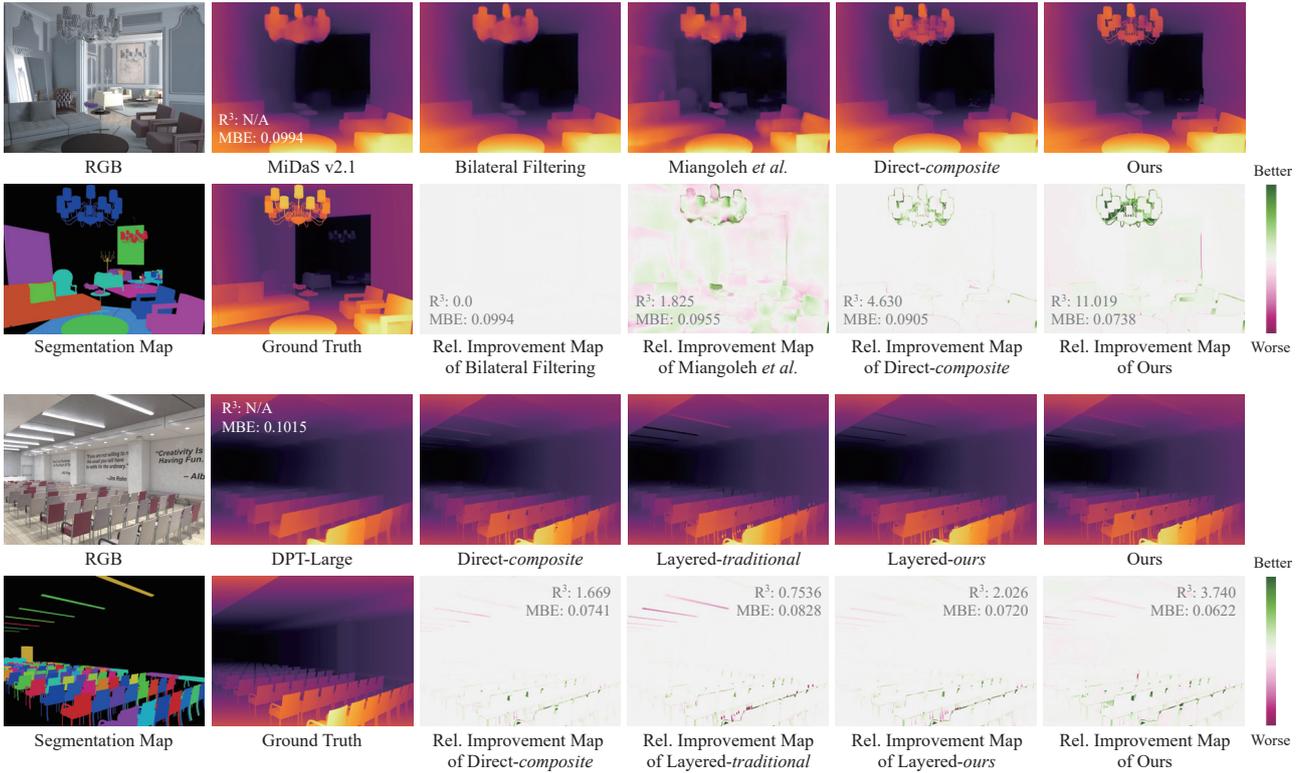


Figure 6. Qualitative results on Hypersim [34]. The relative improvement maps visualize where the refinement method improved and worsened the initial depth estimation by [32] or [31]. Our method focuses on the edges and hole regions, accurately refining fine structures.

as they measure the average error over all pixels, whereas mask-guided refinement methods aim at refining along mask boundaries and leave most internal regions as is. Our method outperforms all baselines in  $R^3$  and MBE, demonstrating the power of our layered refinement approach.

In Table 2, we compare to automatic depth refinement methods without mask-guidance. Conventional image filtering fails to enhance the edge-related metrics. Miangoleh *et al.*'s method [25] is at times better on the global edge metrics ( $\epsilon_{acc}$  and  $\epsilon_{comp}$ ) as it enhances all edges in the depth map. However, as it also carries the risk of distorting the original values,  $R^3$  values tend to be lower compared to ours, which mostly refines along mask boundaries and leaves other regions intact. Furthermore, as [25] heavily relies on the base model's behavior, its generalization capability is limited for other architecture types such as a transformer [31]. Our method works well regardless of the base model architecture and generalizes well to both datasets, leading to the best metric values when coupled with [31].

In Figure 6, we show the qualitative results on Hypersim [34]. We also visualize the relative improvement maps showing where the absolute error decreased compared to the base model MiDaS v2.1 [32] or DPT [31]. Our method focuses on refining edges and hole regions and leaves most other

	Hypersim [34]				TartanAir [44]			
	$R^3 \uparrow$	MBE $\downarrow$	$\epsilon_{acc} \downarrow$	$\epsilon_{comp} \downarrow$	$R^3 \uparrow$	MBE $\downarrow$	$\epsilon_{acc} \downarrow$	$\epsilon_{comp} \downarrow$
[32]	-	0.0973	2.521	7.074	-	0.0596	3.483	6.913
+ BMF	0.7784	0.0974	2.574	7.089	1.032	0.0597	3.489	6.947
+ [25]	4.671	0.0923	<b>1.551</b>	<b>5.837</b>	4.721	0.0602	3.605	7.287
+ Ours	<b>5.209</b>	<b>0.0906</b>	1.888	5.931	<b>16.569</b>	<b>0.0579</b>	<b>2.851</b>	<b>6.272</b>
[31]	-	0.0936	2.071	6.190	-	0.0496	2.574	5.677
+ BMF	0.9444	0.0937	2.094	6.203	0.6875	0.0497	2.667	5.836
+ [25]	1.843	0.0905	1.681	5.633	4.013	0.0496	2.414	5.569
+ Ours	<b>4.455</b>	<b>0.0840</b>	<b>1.491</b>	<b>5.087</b>	<b>8.767</b>	<b>0.0474</b>	<b>2.282</b>	<b>5.245</b>

BMF: Bilateral Median Filtering

Table 2. Comparison to automatic refinement methods. Our method refines mask boundaries and leaves other regions intact whereas [25] refines all regions at the risk of distorting original values.

regions untouched, whereas Miangoleh *et al.*'s method [25] often worsens homogeneous regions. Compared to other baselines, our layered refinement approach within a unified framework helps to correct low-level details effectively.

**Images in the wild** We further evaluate our model on real images *in the wild* to assess its generalization ability and robustness. Comparisons to baselines are shown in Figure 2 and more results are shown in Figures 1, 7, and 8. Our method is able to generate sharp depth maps consistent with

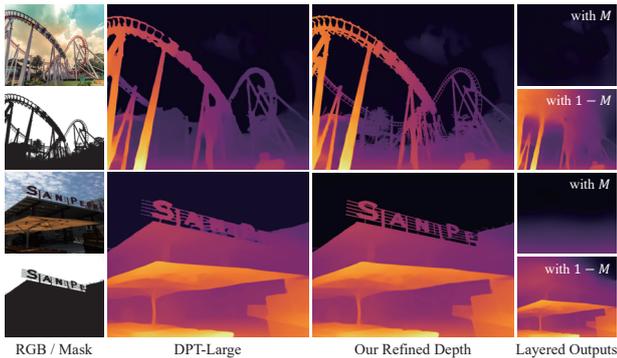


Figure 7. Refined results on real images with various masks.

Stage I	Stage II	HP	$R^3 \uparrow$	MBE $\downarrow$	$\epsilon_{acc} \downarrow$	$\epsilon_{comp} \downarrow$
	DPT-Large [31]		-	0.0936	2.071	6.190
✓			1.996	0.0954	1.606	5.605
	✓		2.016	0.0890	1.915	5.320
	✓	✓	2.613	0.0861	1.670	5.161
✓	✓		<b>5.384</b>	0.0846	<b>1.438</b>	5.100
✓	✓	✓	4.455	<b>0.0840</b>	1.491	<b>5.087</b>

HP: Hole Perturbation

Table 3. Ablation study on Hypersim [34]. Best values in bold.

the mask for various real images. All portrait images are free-licensed images from unsplash [38] and pixabay [28], and masks are generated with removebg [33]. Sky images are licensed by Adobe Stock [2], and their masks are annotated using a commercial photo editing tool.

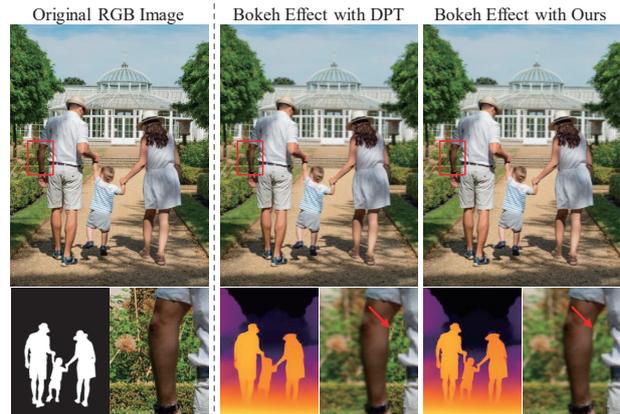
**Ablation study** We provide an ablation study of our model in Table 3 by removing different components in our framework. Stage I helps start from better-initialized parameters, and Stage II is necessary to train our model for layered refinement under a unified framework. Ablating either of them results in performance degradation. Although the quantitative results with or without hole perturbations are similar, hole perturbations are crucial in improving holes in humans.

**Results on downstream applications** More accurate depth maps can improve the outcomes of downstream applications. In Figure 8(a), edges and holes are improved with our refined depth map in point cloud representations of a novel view. In Figure 8(b), we apply Bokeh effect [48] using initial and refined depth maps. Inaccurate depth values in the initial prediction result in an unnatural sharp background region. With our refined depth map, it is corrected and blurry.

**Analysis on mask quality** We provide a visual comparison using different masks coupled with the same image and a numerical analysis with degraded masks in the supplementary material. We show that our method can improve the depth quality as long as the given mask contains more accurate details than the original depth map.



(a) Point Cloud Representation



(b) Bokeh Effect

Figure 8. Point cloud and Bokeh effect [48] using initial depth by DPT [31] and refined depth by Ours. Better viewed with zoom-in.

## 5. Conclusion

Although depth maps are widely used in many practical applications, obtaining sharp and accurate depths from a single RGB image is highly challenging. In this paper, we presented the novel problem of mask-guided depth refinement and proposed a layered refinement approach that can be trained in a self-supervised fashion. Our method can significantly enhance initial depth maps quantitatively and qualitatively. We extensively validated our method by comparing it to mask-guided depth refinement baselines and existing automatic refinement methods. Furthermore, we verified that our method works well on real images with various masks and improves the results of downstream applications. We believe that our method can be potentially extended to other types of dense predictions such as normals and optical flow. More results are provided in the supplementary material.

**Limitations** Since our method relies on a high-quality mask for refinement, its refinement performance is bounded by the mask quality. Although many auto-masking tools are available, capturing extremely fine-grained details may require manual work. Furthermore, as our method refines along mask boundaries, initially wrong depth values inside objects are likely to be left unaltered.

## References

- [1] Amir Atapour Abarghouei and Toby P. Breckon. Real-Time Monocular Depth Estimation Using Synthetic Data With Domain Adaptation via Image Style Transfer. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [2] AdobeStock. Website with a collection of licensed images. <https://stock.adobe.com/>. [Online; accessed 16-November-2021]. 8
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *arXiv/1607.06450*, 2016. 5
- [4] Daniel J. Butler, Jonas Wulff, Garrett B. Stanley, and Michael J. Black. A Naturalistic Open Source Movie for Optical Flow Evaluation. In *European Conference on Computer Vision*, 2012. 2
- [5] Tian Chen, Shijie An, Yuan Zhang, Chongyang Ma, Huayan Wang, Xiaoyan Guo, and Wen Zheng. Improving Monocular Depth Estimation by Leveraging Structural Awareness and Complementary Datasets. In *European Conference on Computer Vision*, 2020. 2
- [6] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-Image Depth Perception in the Wild. In *Advances in Neural Information Processing Systems*, 2016. 6
- [7] Ming-Ming Cheng, Niloy J. Mitra, Xiaolei Huang, Philip H. S. Torr, and Shi-Min Hu. Global Contrast Based Salient Region Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):569–582, 2015. 1
- [8] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 5
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 5
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, 2021. 2, 5
- [11] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep Ordinal Regression Network for Monocular Depth Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1
- [12] Ravi Garg, B. G. Vijay Kumar, Gustavo Carneiro, and Ian D. Reid. Unsupervised CNN for Single View Depth Estimation: Geometry to the Rescue. In *European Conference on Computer Vision*, 2016. 2
- [13] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised Monocular Depth Estimation With Left-Right Consistency. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [14] Qibin Hou, Ming-Ming Cheng, Xiaowei Hu, Ali Borji, Zhuowen Tu, and Philip H. S. Torr. Deeply Supervised Salient Object Detection with Short Connections. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(4):815–828, 2019. 1
- [15] Saif Imran, Xiaoming Liu, and Daniel Morris. Depth Completion with Twin-Surface Extrapolation at Occlusion Boundaries. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 3
- [16] Varun Jampani, Huiwen Chang, Kyle Sargent, Abhishek Kar, Richard Tucker, Michael Krainin, Dominik Kaeser, William T. Freeman, David Salesin, Brian Curless, and Ce Liu. SLIDE: Single Image 3D Photography with Soft Layering and Depth-aware Inpainting. In *IEEE International Conference on Computer Vision*, 2021. 3
- [17] Jianbo Jiao, Ying Cao, Yibing Song, and Rynson Lau. Look Deeper into Depth: Monocular Depth Estimation with Semantic Booster and Attention-Driven Loss. In *European Conference on Computer Vision*, 2018. 2
- [18] Tobias Koch, Lukas Liebel, Friedrich Fraundorfer, and Marco Körner. Evaluation of CNN-based Single-Image Depth Estimation Methods. In *European Conference on Computer Vision Workshops*, 2018. 6
- [19] Jae-Han Lee and Chang-Su Kim. Multi-Loss Rebalancing Algorithm for Monocular Depth Estimation. In *European Conference on Computer Vision*, 2020. 2
- [20] Zhengqi Li and Noah Snavely. MegaDepth: Learning Single-View Depth Prediction From Internet Photos. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2, 5
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft COCO: Common Objects in Context. *arXiv/1405.0312*, 2014. 3
- [22] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, 2019. 5
- [23] Ziyang Ma, Kaiming He, Yichen Wei, Jian Sun, and Enhua Wu. Constant Time Weighted Median Filtering for Stereo Matching and Beyond. In *IEEE International Conference on Computer Vision*, 2013. 6
- [24] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised Learning of Depth and Ego-Motion From Monocular Video Using 3D Geometric Constraints. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [25] S. Mahdi H. Miangoleh, Sebastian Dille, Long Mai, Sylvain Paris, and Yağız Aksoy. Boosting Monocular Depth Estimation Models to High-Resolution via Content-Adaptive Multi-Resolution Merging. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 3, 6, 7
- [26] Arsalan Mousavian, Hamed Pirsiavash, and Jana Kosecka. Joint Semantic Segmentation and Depth Estimation With Deep Convolutional Networks. In *International Conference on 3D Vision*, 2016. 2
- [27] Simon Niklaus, Long Mai, Jimei Yang, and Feng Liu. 3D Ken Burns Effect from a Single Image. *ACM Transactions on Graphics*, 38(6):184:1–184:15, 2019. 2, 3

- [28] pixabay. Website with free images that can be used for commercial and non-commercial purposes. <https://pixabay.com/>. [Online; accessed 16-November-2021]. 5, 8
- [29] Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, and Jiaya Jia. GeoNet: Geometric Neural Network for Joint Depth and Surface Normal Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1
- [30] Yu Qiao, Yuhao Liu, Xin Yang, Dongsheng Zhou, Mingliang Xu, Qiang Zhang, and Xiaopeng Wei. Attention-Guided Hierarchical Structure Aggregation for Image Matting. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1
- [31] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision Transformers for Dense Prediction. In *IEEE International Conference on Computer Vision*, 2021. 1, 2, 3, 5, 6, 7, 8
- [32] 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*, 2020. 1, 2, 3, 5, 6, 7
- [33] remove.bg. Software for removing background from images. <https://www.remove.bg/upload>. [Online; accessed 16-November-2021]. 1, 5, 8
- [34] 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 *IEEE International Conference on Computer Vision*, 2021. 2, 5, 6, 7, 8
- [35] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3D Photography Using Context-Aware Layered Depth Inpainting. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 3, 6
- [36] Alexandru Telea. An Image Inpainting Technique Based on the Fast Marching Method. *Journal of Graphics Tools*, 9(1):23–34, 2004. 2, 6
- [37] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. MLP-Mixer: An all-MLP Architecture for Vision. *arXiv/2105.01601*, 2021. 2
- [38] unsplash. Website with free images that can be used for commercial and non-commercial purposes. <https://unsplash.com/>. [Online; accessed 16-November-2021]. 5, 8
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Advances in Neural Information Processing Systems*, 2017. 2
- [40] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web Stereo Video Supervision for Depth Prediction from Dynamic Scenes. In *International Conference on 3D Vision*, 2019. 2
- [41] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to Detect Salient Objects with Image-Level Supervision. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1
- [42] Lijun Wang, Jianming Zhang, Oliver Wang, Zhe Lin, and Huchuan Lu. SDC-Depth: Semantic Divide-and-Conquer Network for Monocular Depth Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [43] 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. *arXiv/1912.09678*, 2019. 2
- [44] 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 *IEEE International Conference on Intelligent Robots and Systems*, 2020. 2, 5, 6, 7
- [45] Alex Wong, Byung-Woo Hong, and Stefano Soatto. Bilateral Cyclic Constraint and Adaptive Regularization for Unsupervised Monocular Depth Prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [46] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2
- [47] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-Guided Ranking Loss for Single Image Depth Prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1, 2
- [48] Lei Xiao, Anton Kaplanyan, Alexander Fix, Matthew Chapman, and Douglas Lanman. DeepFocus: Learned Image Synthesis for Computational Displays. *ACM Transactions on Graphics*, 37(6), 2018. 8
- [49] Ning Xu, Brian L. Price, Scott Cohen, and Thomas S. Huang. Deep Image Matting. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1
- [50] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. BlendedMVS: A Large-scale Dataset for Generalized Multi-view Stereo Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [51] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to Recover 3D Scene Shape from a Single Image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 5
- [52] Haichao Yu, Ning Xu, Zilong Huang, Yuqian Zhou, and Humphrey Shi. High-resolution deep image matting. In *AAAI Conference on Artificial Intelligence*, 2021. 1
- [53] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. T2Net: Synthetic-To-Realistic Translation for Solving Single-Image Depth Estimation Tasks. In *European Conference on Computer Vision*, 2018. 2
- [54] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene Parsing through ADE20K Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 3

- [55] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised Learning of Depth and Ego-Motion From Video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [56] Shengjie Zhu, Garrick Brazil, and Xiaoming Liu. The Edge of Depth: Explicit Constraints between Segmentation and Depth. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [57] Yunzhi Zhuge, Yu Zeng, and Huchuan Lu. Deep Embedding Features for Salient Object Detection. In *AAAI Conference on Artificial Intelligence*, 2019. 1