

PanopticDepth: A Unified Framework for Depth-aware Panoptic Segmentation

Naiyu Gao^{1,2}, Fei He^{1,2}, Jian Jia^{1,2}, Yanhu Shan⁴, Haoyang Zhang⁴, Xin Zhao^{1,2*}, Kaiqi Huang^{1,2,3}

¹ CRISE, Institute of Automation, Chinese Academy of Sciences

² School of Artificial Intelligence, University of Chinese Academy of Sciences

³ CAS Center for Excellence in Brain Science and Intelligence Technology, ⁴ Horizon Robotics, Inc.

{gaonaiyu2017, hefei2018, jiajian2018}@ia.ac.cn, {xzhaoh, kaiqi.huang}@nlpr.ia.ac.cn

{yanhu.shan, haoyang.zhang}@horizon.ai

Abstract

This paper presents a unified framework for depth-aware panoptic segmentation (DPS), which aims to reconstruct 3D scene with instance-level semantics from one single image. Prior works address this problem by simply adding a dense depth regression head to panoptic segmentation (PS) networks, resulting in two independent task branches. This neglects the mutually-beneficial relations between these two tasks, thus failing to exploit handy instance-level semantic cues to boost depth accuracy while also producing sub-optimal depth maps. To overcome these limitations, we propose a unified framework for the DPS task by applying a dynamic convolution technique to both the PS and depth prediction tasks. Specifically, instead of predicting depth for all pixels at a time, we generate instance-specific kernels to predict depth and segmentation masks for each instance. Moreover, leveraging the instance-wise depth estimation scheme, we add additional instance-level depth cues to assist with supervising the depth learning via a new depth loss. Extensive experiments on Cityscapes-DPS and SemKITTI-DPS show the effectiveness and promise of our method. We hope our unified solution to DPS can lead a new paradigm in this area. Code is available at <https://github.com/NaiyuGao/PanopticDepth>.

1. Introduction

Depth-aware panoptic segmentation (DPS) is a new challenging task in scene understanding, attempting to build 3D scene with instance-level semantic understanding from a single image. Its goal is to assign each pixel a depth value, a semantic class label and an instance ID. Thus, solving this problem involves monocular depth estimation [1] and panoptic segmentation [2]. Naturally, a straightforward solution to DPS is to add a dense depth regression head to

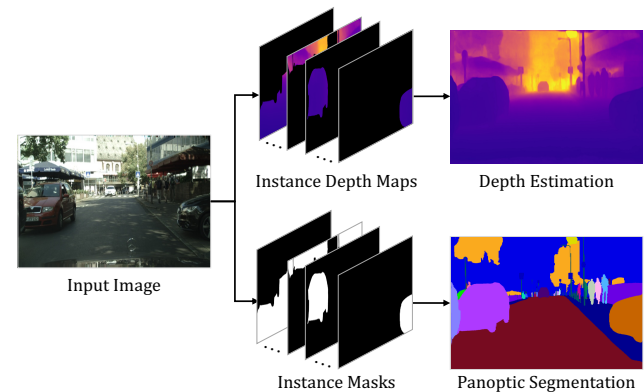


Figure 1. Illustration of our unified solution to depth-aware panoptic segmentation which requires assigning each pixel in a single image a depth value, a semantic class label and an instance ID. Instead of predicting pixel-wise depth, we predict instance-wise depth by instance-specific convolution kernels, which shares the same manner of instance mask generation.

panoptic segmentation (PS) networks [3–5], producing a depth value for each labelled pixel.

This method is intuitive yet sub-optimal. Since it tackles the two tasks with two independent branches, it does not explore the mutually-beneficial relations between them, especially failing to exploit handy instance-level semantic cues to boost depth accuracy. We observe that pixels of adjacent instances generally have discontinuous depth. For examples, two vehicles in a line are likely to have different depth. Therefore, it is hard to predict accurate depth for both vehicles using the same pixel-wise depth regressor. On the other hand, as indicated by previous works [6–9], considering that these pixels are of different vehicles, it benefits the depth estimation if separate regressors are used respectively.

Following the train of thought above, we propose PanopticDepth in this paper, a unified model that predicts the

*Corresponding author

mask and depth values in the same instance-wise manner (Figure 1). In contrast to predicting depth values for all pixels at a time, we manage to estimate depth for each thing/stuff instance, which also shares the way of generating instance masks. To this end, we employ the technique of dynamic convolution [10–13] in producing both instance masks and depth, which unifies the pipelines of mask generation and depth estimation.

Specifically, we first generate instance-specific mask and depth kernels for each instance concurrently, which is similar to the method used in [12]. Then, we apply the mask kernels to the mask embeddings and the depth kernels to the depth embeddings, producing the mask and depth map for each instance, respectively. Finally, we merge individual instance masks into a panoptic segmentation map, following a similar process presented in [14]. According to the panoptic segmentation results, we then aggregate each instance’s depth into a whole depth map. As a result, we get both the panoptic segmentation and depth map for one image. Figure 2 shows the pipeline.

Our method unifies the depth estimation and panoptic segmentation approaches via an instance-specific convolution kernel technique, which also in turn improves the performance on both tasks (Table 3 and Table 5). Thanks to the dynamic convolution kernel technique [12], the learned instance depth regressor aggregates not only global context, but local information, such as instance shapes, scales and positions, into instance depth prediction. Such information turns out essential to obtain accurate depth values especially those at instance boundaries (Figure 3).

Furthermore, in order to ease the depth estimation, inspired by Batch Normalization [15], we propose to represent each instance depth map as a triplet, *i.e.* a normalized depth map, a depth range and a depth shift. In general, depth values of different instances may vary greatly, like a long vehicle with a length of 70m *v.s.* a small car with a length of 4.5m. This large variation in scale may cause difficulties in learning shared depth embedding. To tackle this problem, we propose to represent the depth map with the aforementioned triplet and normalize the values of original instance depth map to $[0, 1]$. This improves the learning effectiveness (Table 3 and Table 5). At the same time, in addition to the traditional pixel-level depth supervision, we add instance-level depth statistics based on the new depth map representation, *e.g.* the depth shift, to reinforce the depth supervision. We also propose a corresponding depth loss to accommodate this new supervision (Sec:3.3.2), which helps improve the depth prediction.

Through extensive experiments on Cityscapes-DPS [3] and SemKITTI-DPS [3], we demonstrate the effectiveness of our unified solution to the depth-aware panoptic segmentation. We hope our unified framework can lead a new paradigm in this challenging task.

2. Related Work

2.1. Panoptic Segmentation

Panoptic segmentation [2] (PS) is a recently proposed vision task, which requires generating a coherent scene segmentation for an image. It unifies the tasks of semantic segmentation and instance segmentation. Early approaches [14, 16–22] to PS mainly rely on individual semantic segmentation models [23–25] to get pixel-level segmentation for stuff classes and two-stage instance segmentation models [26] to get instance-level segmentation for things, and then fuse them [2, 27] to get the final panoptic segmentation. Instead of adopting two-stage instance segmentation models, recent works [28–35] attempt to solve panoptic segmentation in a one-stage manner. Among them, DeeperLab [30] proposes to solve panoptic segmentation by predicting key-point and multi-range offset heat-maps, followed by a grouping process [36]. Similarly, the series of Panoptic-DeepLab [37–39] predict instance centers as well as offsets from each pixel to its corresponding center. More recently, inspired by the instance-specific kernel method used in [10, 11, 40, 41], some works [12, 13, 42–44] start to apply this technique in PS, achieving unified modelling of both stuff and thing segmentation.

In this work, we extend the PS task by predicting additional depth for each instance, to build a 3D semantic scene. Similar to [12, 13, 42–44], our proposed unified framework for this problem is also inspired by the instance-specific kernel technique, especially the work [12].

2.2. Monocular Depth Estimation

Monocular depth estimation aims to predict depth from a single image. Before deep learning era, there have already been many attempts [45–48]. Recently, a variety of better-performing methods based on deep networks have been proposed. Laina *et al.* [49] propose a fully convolutional architecture with up-projection blocks for high-dimensional depth regression, while Li *et al.* [50] design a two-streamed network to predict fine-scaled depth. Besides supervised learning, other learning techniques have also been explored to improve depth estimation generalization, including self-supervised learning [51–55], transfer learning with synthetic images [56, 57], and learning relative depth orders from web data for depth perception [58–63].

Most aforementioned works predict depth from pixel perspective, in contrast, our method predicts depth for each instance, which incorporates instance-level cues into depth estimation.

Enhance depth with semantic segmentation. There are works [6–8, 64–66] that attempt to improve monocular depth estimation with semantic segmentation. Liu *et al.* [47] propose to boost depth estimation with semantic

labels by Markov random fields. Pad-Net [7] leverages intermediate depth and segmentation predictions to refine the final output. Jiao *et al.* [9] refine depth predictions with attention-driven loss or task-level interactions. SigNet [8] integrates semantic segmentation with instance edges to model depth estimation in an unsupervised manner.

These works show the mutual benefit of simultaneously conducting depth estimation and semantic segmentation, but we argue that we can achieve better results if we perform instance-wise depth estimation with panoptic segmentation, since we now have instance-level cues to facilitate depth estimation.

Enhance depth with panoptic segmentation. More recently, some works [3, 5, 67, 68] have started to explore the joint learning of monocular depth estimation and panoptic segmentation, but most of them utilize panoptic segmentation to enhance depth just by sharing backbone features [3] or adding constraints [5, 68]. By contrast, we unify the panoptic segmentation and instance-wise depth estimation in one framework through the instance-specific convolution technique.

Among them, SDC-Depth [67] also explores the idea of instance-wise depth estimation, but it is different from our work in many aspects. First of all, our target task is different from its. The goal of SDC-Depth is just to estimate depth where instance masks is used only to provide clues for improving the depth map. In contrast, our DPS task requires producing both high-quality masks and depth maps, which is more challenging. Besides, we adopt a different instance depth normalization scheme, which allows us to supervise learning depth estimation with both pixel- and instance-level cues at the meantime. Finally, our approach is built upon a dynamic-convolution-based panoptic segmentation model [12], which generates more precise segment boundaries, rather than the two-stage approach employed in SDC-Depth.

3. PanopticDepth

We propose PanopticDepth, a unified model for depth-aware panoptic segmentation, which predicts the mask and depth values in the same instance-wise manner. Apart from backbone and the feature pyramid network [69], it consists of three main sub-networks, including a kernel producer for generating instance classification, instance-specific mask and depth convolution kernels, a panoptic segmentation model for generating instance masks, and an instance-wise depth map generator for estimating instance depths. The architecture of our networks is shown in Figure 2. We elaborate each module in the following sections.

3.1. Kernel Producer

In our system, we first generate instance classification, mask convolution kernels and depth estimation kernels by

a kernel producer sub-network (upper part of Figure 2). Our kernel producer is built on the state-of-the-art panoptic segmentation model, PanopticFCN [12], which adopts a dynamic convolution technique for PS but requires less training time and GPU memory compared to other recent methods [13, 42]. We briefly describe the process of kernel generation here. For more details please refer to [12].

Given multiple stage FPN features X , our kernel producer outputs N instance classifications, N mask kernels, and N depth kernels. It has two stages, that is, kernel generator and kernel fusion.

At the kernel generator stage, taking as input a single stage feature X_i from the i -th stage in FPN, the generator produces a kernel weight map G_i , and two position maps for things L_i^{th} and stuff L_i^{st} respectively, where each thing is represented by an object center and stuff by a region.

Given the position maps and kernel weight map from each FPN stage, at the kernel fusion stage, we merge repetitive kernel weights from multiple FPN stages. This is done through a proposed adaptive kernel fusion (AKF) operation, which is an improved fusion mechanism over the original average clustering used in PanopticFCN. Our adaptive fusion mechanism improves 1.2% PQ without introducing extra parameters. Specifically, to generate a mask kernel K_k^m for the k -th instance, we compute it through:

$$K_k^m = \frac{\sum(G * R_k)}{\sum R_k}, \quad (1)$$

where $*$ denotes element-wise multiplication, and R represents positions with peak scores in L^{th} for things and positions with the highest score for corresponding categories in L^{st} for stuff, respectively. Kernels with large cosine similarities are also fused, as done in [12].

For depth kernel generation, the similar process is performed. In this way, given FPN features, our kernel producer generates instance classification $C \in R^{N \times c}$, mask kernels $K^m \in R^{N \times e^m}$, and depth kernels $K^d \in R^{N \times e^d}$, where c is the number of categories, e^m and e^d are dimensions of mask and depth kernels, respectively.

3.2. Panoptic Segmentation

We adopt an instance-specific kernel method to perform panoptic segmentation [12]. As shown in the bottom part of Figure 2, the masks M of thing and stuff instances are derived by convolving a shared high-resolution mask embedding map $E^m \in R^{e^m \times H/4 \times W/4}$ with the mask kernels K^m , followed the Sigmoid activation:

$$M = \text{Sigmoid}(K^m \otimes E^m). \quad (2)$$

Redundant instance masks are firstly discarded as done in [14]. After that, all remaining instance masks are merged

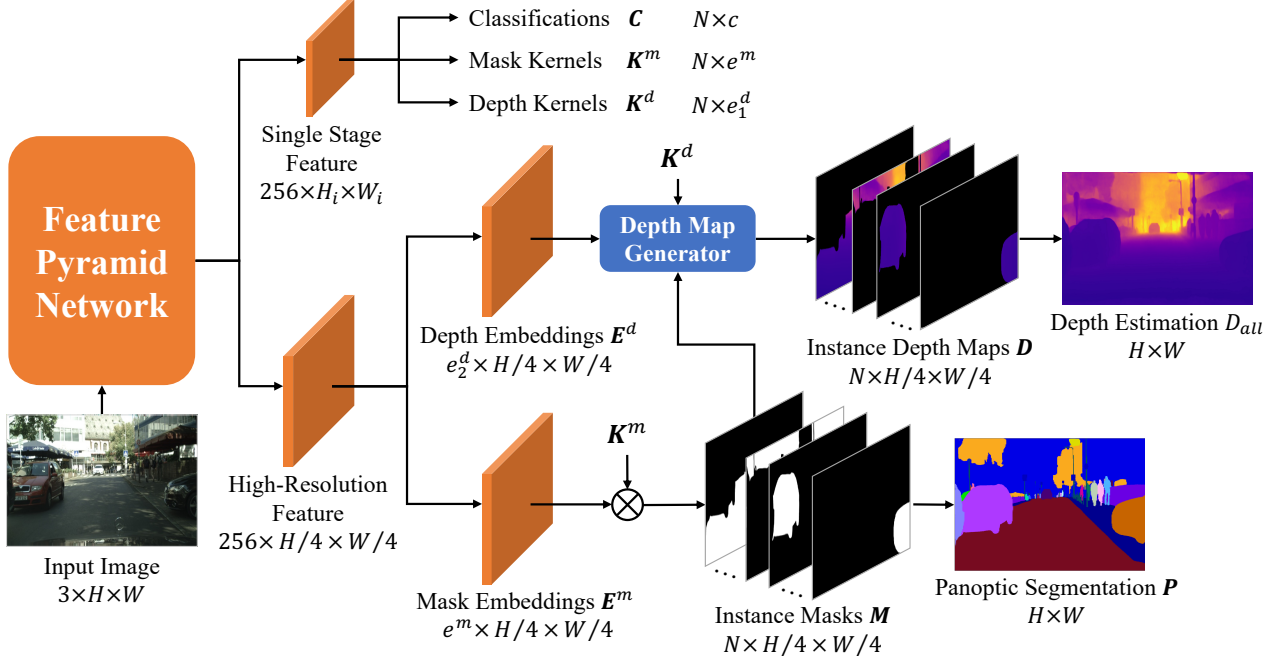


Figure 2. The framework of our proposed PanopticDepth. H and W are the height and width of the input image respectively. N denotes the number of thing and stuff instances. c is the number of categories. e^m , e_1^d , and e_2^d are embedding dimensions of mask kernels, depth kernels and the depth embedding map, respectively. \otimes represents convolution.

with argmax to yield the non-overlapped panoptic segmentation result, so that each pixel is assigned to a thing or stuff segment, with no pixel labeled as ‘VOID’.

Moreover, we propose an additional training process to bridge the performance gap between training and testing. Concretely, we find distant instances are often fused by the learned model, which is caused by the widely employed image cropping strategy. Some previous works [3, 37] tackle this problem by directly training models with original resolution images, which is effective but increases the GPU memory footprint dramatically. In contrast, we propose a more efficient training strategy, *i.e.*, fine-tuning the learned model at the full image scale, but with a small batch size. More detailed training process can be found in Subsection 4.3.

3.3. Instance-wise Depth Estimation

We predict depth for each instance through the same instance-specific kernel technique used in panoptic segmentation, which unifies the pipelines of depth estimation and panoptic segmentation. As shown in the middle part of Figure 2, we first run the depth kernels on the depth embedding to generate instance depth maps, and then merge these individual maps in accordance with the panoptic segmentation results to yield the final whole depth map. In this section, we first introduce the depth map generator and then a new depth loss.

3.3.1 Depth Map Generator

Given instance-specific depth kernels $K^d \in R^{N \times e_1^d}$ and the shared depth embedding $E^d \in R^{e_2^d \times H/4 \times W/4}$ (we set $e_1^d = e_2^d$), similar to the instance mask generation process, we first generate the normalized instance depth map D' through convolution and Sigmoid activation, and then unnormalize it to the depth map D through Eq 4 or Eq 5:

$$D' = \text{Sigmoid}(K^d \otimes E^d), \quad (3)$$

$$\mathcal{T}_1(D|D', d^r, d^s) = d_{\max} \times (d^r \times D' + d^s), \quad (4)$$

$$\mathcal{T}_2(D|D', d^r, d^s) = d_{\max} \times [d^r \times (D' - 0.5) + d^s], \quad (5)$$

where d_{\max} controls the depth scale and is set to 88 to be consistent with the depth range of Cityscapes-DPS and SemKITTI-DPS.

The reason why we normalize the depth map is that different instances have widely varied depth ranges, making it hard to learn effective shared depth embedding. To ease the learning of depth estimation and inspired by Batch Normalization [15], we predict a normalized depth map D' instead, which is obtained by normalizing the instance depth map D with two predicted instance-level depth variables, depth range $d^r \in R^{N \times 1}$ and depth shift $d^s \in R^{N \times 1}$. They describe the depth bias and variance of each instance, respectively, and can be derived from high-level features in parallel with depth kernels by simply setting $e_1^d = e_2^d + 2$. Note

that $0 \leq d^s \leq 1$ and $0 \leq d^r \leq 1$ after Sigmoid activation.

In this way, the normalized depth map D' only encodes relative depth values within each instance, and thus can be more easily learned. Besides, we develop two normalization schemes, *i.e.* Eq. 4 and Eq. 5, and find that the latter one works better.

After getting all instance depth maps, we aggregate them into a whole image depth map, according to the non-overlapped panoptic segmentation masks M . This generates precise depth values at instance boundaries.

3.3.2 Depth Loss

Following [3], we develop the depth loss function based on the combination of scale-invariant logarithmic error [70] and relative squared error [71], both of which are popular depth metrics and can be directly optimized end-to-end. Specifically,

$$\begin{aligned} \tilde{L}_{dep}(\mathbf{d}, \hat{\mathbf{d}}) &= \frac{1}{n} \sum_j (\log d_j - \log \hat{d}_j)^2 \\ &\quad - \frac{1}{n^2} \left(\sum_j \log d_j - \log \hat{d}_j \right)^2 \\ &\quad + \left[\frac{1}{n} \sum_j \left(\frac{d_j - \hat{d}_j}{\hat{d}_j} \right)^2 \right]^{0.5}, \end{aligned} \quad (6)$$

where \mathbf{d} and $\hat{\mathbf{d}}$ denote the predicted and the ground-truth depth, respectively.

Thanks to the instance-wise depth estimation means, we can now learn the depth prediction under both the traditional pixel-level supervision and extra instance-level supervision, which empirically improves the depth accuracy. To enable the dual supervision, our final depth loss L_{dep} includes two loss terms. One is the pixel-level depth loss L_{dep}^P and the other instance-level depth loss L_{dep}^I :

$$L_{dep} = L_{dep}^P + \lambda_{dep}^I L_{dep}^I, \quad (7)$$

where λ_{dep}^I controls the relative weighting and is set to 1 by default. The pixel-level depth loss L_{dep}^P calculates the depth error between depth predictions and ground truth inside each instance:

$$L_{dep}^P = \tilde{L}_{dep}(D_{all}, \hat{D}_{all}). \quad (8)$$

The instance-level depth loss L_{dep}^I computes the depth error between instance depth shift d^s and corresponding ground truth:

$$L_{dep}^I = \tilde{L}_{dep}(d^s, \hat{d}^s). \quad (9)$$

The minimum and average depth values within each instance mask are employed as ground truths of depth shift d^s in Eq. 4 and Eq. 5, respectively.

4. Experiments

4.1. Datasets

We evaluate our method on Cityscapes [72], Cityscapes-DPS [3], and SemKITTI-DPS [3] benchmarks. Cityscapes is a challenging dataset for image segmentation. In this dataset, 5,000 images of a high resolution $1,024 \times 2,048$ are annotated at the high-quality pixel level, and are divided into 2,975, 500, and 1,425 images for training, validation, and testing, respectively. Cityscapes panoptic segmentation benchmark evaluates 8 thing and 11 stuff classes. Recently, Qiao *et al.* [3] proposed a depth-aware panoptic segmentation dataset Cityscapes-DPS by supplementing Cityscapes with depth annotations, which are computed from the disparity maps via stereo images. 2,400 and 300 images are annotated for training and validation, respectively. SemKITTI-DPS [3] is converted from cloud points with semantic annotations in SemKITTI [73], consisting of 19,130 and 4,071 images for training and validation, respectively.

4.2. Metrics

Metric of PS. The results for panoptic segmentation are evaluated with the standard Panoptic Quality (PQ) metric, introduced by Kirillov *et al.* [2]. The formulation of PQ is:

$$PQ = \frac{\sum_{p,g \in TP} \text{IoU}(p,g)}{|TP| + \frac{1}{2}|FP| + \frac{1}{2}|FN|}, \quad (10)$$

where p and g denote the predicted and ground truth segment, respectively. TP , FN , and FP represent matched pairs of segments ($\text{IoU}(p,g) > 0.5$), unmatched ground truth segments, and unmatched predicted segments, respectively. PQ of both thing and stuff classes are reported.

Metric of DPS. The evaluation metric for depth-aware panoptic segmentation is DPQ [3], which quantifies the performance for segmentation and depth estimation simultaneously. Specifically, given prediction P , ground-truth \hat{P} , and depth threshold λ , DPQ^λ is computed as:

$$\text{DPQ}^\lambda(P, \hat{P}) = \text{PQ}(P^\lambda, \hat{P}). \quad (11)$$

$P^\lambda = P$ for pixels that have absolute relative depth errors under λ to filter out pixels that have large absolute relative depth errors. DPQ is calculated by averaging $\text{DPQ}^\lambda(P, \hat{P})$ over $\lambda = \{0.1, 0.25, 0.5\}$.

4.3. Implementation

Our models are implemented with PyTorch [78] and the Detectron2 [79] toolbox. Unless specified, ResNet-50 [80] with FPN [69] is employed as backbone. The dimension of mask embeddings E^m is set to 256 following PanopticFCN [12]. The dimension of depth embeddings E^d is set to 16.

Method	Backbone	Extra Data	PQ	PQ Th	PQ St	PQ [test]	PQ Th [test]	PQ St [test]
UPNet [18]	R-50	-	59.3	54.6	62.7	-	-	-
Seamless	R-50	-	59.8	54.6	63.6	-	-	-
UPNet [18]	R-50	COCO	60.5	57.0	63.0	-	-	-
SSAP [†] [32]	R-101	-	61.1	55.0	-	58.9	48.4	66.5
Unifying [22]	R-50	-	61.4	54.7	66.3	61.0	52.7	67.1
Panoptic-DeepLab [37]	Xcp-71 [74]	-	63.0	-	-	-	-	-
Panoptic-DeepLab [†]	Xcp-71	-	-	-	-	62.3	52.1	69.7
Baseline, original [12]	R-50	-	61.4	54.8	66.6	-	-	-
Baseline, our impl.	R-50	-	62.4	56.0	67.1	-	-	-
+ AKF	R-50	-	63.6	57.2	68.1	-	-	-
+ FSF	R-50	-	64.1	58.8	68.1	62.0	55.0	67.1

Table 1. Panoptic segmentation results on Cityscapes validation and test sets. ‘AKF’: adaptive kernel fusion. ‘FSF’: full-scale fine-tuning. †: test-time augmentation. Results are reported as percentages.

Method	Backbone	Extra Data	$\lambda=0.5$	$\lambda=0.25$	$\lambda=0.1$	DPQ
ViP-DeepLab [†]	WR-41 [39]	MV, CSV	68.7 / 61.4 / 74.0	66.5 / 60.4 / 71.0	50.5 / 45.8 / 53.9	61.9 / 55.9 / 66.3
Ours	R-50	-	65.6 / 59.2 / 70.2	62.3 / 57.0 / 66.1	43.2 / 40.7 / 45.1	57.0 / 52.3 / 60.5
Ours	Swin-T [75]	-	66.5 / 61.0 / 70.5	64.1 / 59.9 / 67.2	48.6 / 44.8 / 51.3	59.7 / 55.2 / 63.0
Ours	Swin-S [75]	-	67.4 / 62.5 / 71.0	65.0 / 61.4 / 67.7	48.8 / 44.2 / 52.3	60.4 / 56.0 / 63.6

Table 2. Depth-aware panoptic segmentation results on Cityscapes-DPS. ‘MV’: Mapillary Vistas [76]. ‘CSV’: Cityscapes videos with pseudo labels [77]. †: test-time augmentation. Each cell contains DPQ / DPQTh / DPQSt scores. Results are reported as percentages.

The training details of PS and DPS model are introduced below, where the trained PS model is used to initialize the DPS model.

Training details of PS model. The training of panoptic segmentation model consists of two steps, where the first step takes a large mini-batch of small cropped images, while the second one has a small mini-batch of large full-scale samples. Specifically, we first train PanopticFCN model with Adam [81] for 130k iterations using synchronized batch normalization [15]. The learning rate is initialized as 10^{-4} and the *poly* schedule with power 0.9 is adopted. The filters in the first two stages of the backbone are fixed in Detectron2 by default, but we find updating these parameters can boost around 0.5% PQ. The images are resized with random factors in [0.5, 2.0], and then cropped into 512×1024 . Each mini-batch contains 32 samples. Color augmentation [82] and horizontal flipping are employed during training. At the second step, we fine-tune the PS model with images scaled by [1.0, 1.5] and then cropped into 1024×2048 for 10k iterations. The batch size is 8. BN layers and the first two stage layers of the backbone are fixed in this and following training steps.

Training details of DPS model. We train the entire model on Cityscapes-DPS for 10k iterations. Images are resized by [0.8, 1.2] and depth annotations are scaled accordingly. We put 8 samples in one mini-batch, which are

center-cropped into 1024×2048 . Color augmentation and horizontal flipping are also employed. On SemKITTI-DPS, we use the pre-trained model from Cityscapes, but the training samples are cropped into 384×1280 at both pre-training and fine-tuning steps. Other settings remain the same.

The total training loss is:

$$L = \lambda_{pos}L_{pos} + \lambda_{seg}L_{seg} + \lambda_{dep}L_{dep}. \quad (12)$$

L_{dep} is defined in Eq. 7. L_{pos} and L_{seg} are losses for classification and segmentation separately as described in PanopticFCN [12]. λ_{pos} , λ_{seg} , and λ_{dep} are set to 1, 4, and 5, respectively. Training steps for PS and DPS model tasks 46 and 5 hours, respectively, on 8 Nvidia 2080Ti GPUs.

The predictions of our method are obtained by a single model and a single inference. No test-time augmentations such as horizontal flipping or multi-scale testing are adopted.

4.4. Main Results

Panoptic segmentation. The original PanopticFCN [12] achieves 61.4% PQ on the Cityscapes validation set, whereas our implementation achieves 62.4% PQ without introducing extra parameters. With the adaptive kernel fusion mechanism and the full-scale fine-tuning step, our full model achieves 64.1% PQ. Because the test set annotations are withheld, we submit the test set results to the online

Variants	IDE	IDN	L_{dep}^I	$\lambda=0.5$	$\lambda=0.25$	$\lambda=0.1$	DPQ
A				64.0 / 56.8 / 69.0	60.0 / 52.8 / 64.8	40.8 / 35.9 / 44.9	54.9 / 48.5 / 59.6
B	✓			63.8 / 56.4 / 69.2	60.0 / 52.6 / 65.4	41.0 / 34.1 / 46.0	54.9 / 47.7 / 60.2
C	✓	\mathcal{T}_1		64.9 / 57.1 / 70.2	61.0 / 53.2 / 66.2	42.4 / 36.5 / 45.5	56.1 / 48.9 / 60.6
D	✓	\mathcal{T}_2		65.0 / 57.3 / 70.1	60.9 / 53.2 / 66.2	42.5 / 36.7 / 45.3	56.1 / 49.1 / 60.5
E	✓	\mathcal{T}_1	✓	60.0 / 47.2 / 69.2	49.5 / 34.8 / 60.1	29.1 / 21.3 / 34.8	46.2 / 34.4 / 54.7
F	✓	\mathcal{T}_2	✓	65.6 / 59.2 / 70.2	62.3 / 57.0 / 66.1	43.2 / 40.7 / 45.1	57.0 / 52.3 / 60.5

Table 3. Ablation studies on Cityscapes-DPS. ‘IDE’: instance-wise depth estimation. ‘IDN’: instance depth normalization. Each cell contains DPQ / DPQTh / DPQSt scores. Results are reported as percentages.

Methods	RMSE [val]↓	RMSE [test]↓
MGNet [‡] [62]	8.30	-
Laina <i>et al.</i> [34]	-	7.27
Pad-Net [81]	-	7.12
Zhang <i>et al.</i> [84]	-	7.10
SDC-Depth [‡] [72]	-	6.92
Ours [‡]	6.91	6.69

Table 4. Monocular depth estimation results on Cityscapes. [‡]: methods utilize panoptic segmentation annotations.

Variants	IDE	IDN	L_{dep}^I	DPQ	DPQ Th	DPQ St
A				45.3	41.4	47.5
B	✓			45.8	42.9	47.5
C	✓	\mathcal{T}_1		46.7	45.6	47.6
D	✓	\mathcal{T}_2		46.7	45.5	47.6
E	✓	\mathcal{T}_1	✓	36.1	35.1	40.6
F	✓	\mathcal{T}_2	✓	46.9	46.0	47.6

Table 5. Ablation studies on SemKITTI-DPS. ‘IDE’: instance-wise depth estimation. ‘IDN’: instance depth normalization. Results are reported as percentages.

evaluation server. Results are shown in Table 1. Our method achieves the highest 55.0% PQTh.

Depth-aware panoptic segmentation. In Table 2, we compare our method with ViP-DeepLab, which is the only published work for depth-aware panoptic segmentation. It achieves a higher performance, however, please note that extra engineering tricks including pre-training with larger datasets [76], semi-supervised learning [77], AutoAug [83], and test-time augmentation are employed in ViP-DeepLab, which can also be utilized in our method to further improve the performance.

Monocular depth estimation. To further demonstrate the superiority of our method over SOTA methods, we report the results of the monocular depth estimation on Cityscapes. As shown in Table 4, the proposed approach shows a clear advantage even compared to methods that boost depth esti-

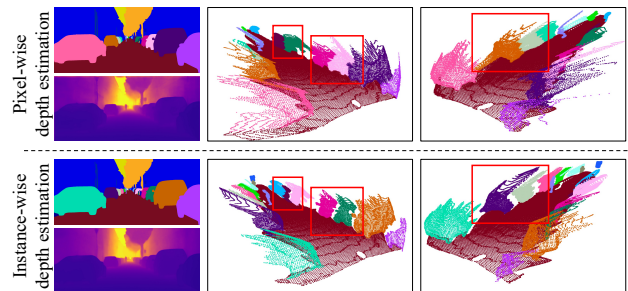


Figure 3. Pixel-wise depth estimation outputs smooth values at the boundary of two instances, whereas instance-wise depth estimation can generate more reasonable discontinuous depth values.

mation with panoptic segmentation annotations.

Visualization results. The panoptic segmentation, monocular depth estimation, and corresponding 3D point cloud results can be visualized in Figure 4.

4.5. Ablation Experiments

Influence of adaptive kernel fusion. Unlike PanopticFCN that generates stuff kernels by global average pooling the masked kernel map, our approach adaptively fuses stuff kernels (Eq. 1), avoiding the interference of varying stuff region sizes. As shown in Table 1, employing adaptive kernel fusion (AKF) mechanism improves PQTh and PQSt by 1.2% and 1.0%, respectively.

Full-scale fine-tuning. Fine-tuning the learned PS model with full-scale images (FSF) promotes the ability to distinguish distant instances. As shown in Table 1, FSF further boosts PQTh by 1.6% to 64.1% PQ.

Instance-wise depth estimation. The comparison of pixel- and instance-wise depth estimation is quantified with variant-A and -B in Table 3. Variant-A performs pixel-wise depth regression by setting the dimension e_2^d of the depth embedding map E^d to 1, achieving 54.9% DPQ. Variant-B conducts instance-wise depth estimation with dynamic kernels, which achieves similar performance to variant-A. However, as shown in Figure 3, pixel-wise depth estimation tends to output smooth values at object boundaries. By




Figure 4. Example predictions of our model. The panoptic segmentation and monocular depth estimation results are predicted from the input image with a unified model. Best viewed in color and zoom.

contrast, our instance-wise depth prediction method avoids this and can generate more reasonable discontinuous depth values at the boundary of two instances.

Instance depth normalization. In Table 3, variant-C and -D further normalizes the instance depth maps with depth shift and depth range respectively, as described in Eq. 4 and Eq. 5. These two variants achieve similar performances but are about 1.2% DPQ higher than variant-A and -B. This result demonstrates the effectiveness of leveraging instance depth statistics for depth estimation.

Instance-level depth loss. As shown in Table 3, variant-E and -F employs the instance-level depth loss upon variant-C and -D, in which the depth shift d^s is supervised with the minimum or mean depth value within each instance mask. Although variant-E drops the performance, variant-F significantly boosts DPQTh by 3.2%. This is because the ground truth depth is noisy, and minimum depth values are sensitive to noise while mean depth values not.

Results on SemKITTI-DPS. Table 5 shows our proposed instance-wise depth estimation (variant-F) achieves better performance again on the dataset of SemKITTI-DPS, especially for thing instances (46.0% v.s. 41.4% DPQTh), which is consistent with the results on Cityscapes-DPS (Table 3).

5. Conclusion

In this paper, a unified framework for depth-aware panoptic segmentation is proposed by predicting the mask and depth values for each thing/stuff in the same instance-wise manner. High-level object information is introduced into the depth estimation by adopting the dynamic kernel technique. Moreover, each instance depth map is normalized with a depth shift and a depth scale to ease the learning of shared depth embedding. Furthermore, a new depth loss is proposed to supervise the depth learning with instance-level depth cues. Experiments on Cityscapes-DPS and SemKITTI-DPS benchmarks demonstrate the effectiveness of the proposed approach. We hope our unified solution to depth-aware panoptic segmentation can lead a new paradigm in this area.

Acknowledgment

This work is supported in part by the National Natural Science Foundation of China (Grant No.61721004 and Grant No.61876181), the Projects of Chinese Academy of Science (Grant No. QYZDB-SSW-JSC006), the Strategic Priority Research Program of Chinese Academy of Sciences (Grant No. XDA27000000), and the Youth Innovation Promotion Association CAS.

References

- [1] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018.
- [2] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollar. Panoptic segmentation. In *CVPR*, 2019.
- [3] Siyuan Qiao, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Vip-deeplab: Learning visual perception with depth-aware video panoptic segmentation. In *CVPR*, 2021.
- [4] Sainan Liu, Vincent Nguyen, Yuan Gao, Subarna Tripathi, and Zhuowen Tu. Towards panoptic 3d parsing for single image in the wild. *arXiv preprint arXiv:2111.03039*, 2021.
- [5] Markus Schön, Michael Buchholz, and Klaus Dietmayer. Mgnet: Monocular geometric scene understanding for autonomous driving. In *ICCV*, 2021.
- [6] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Zequn Jie, Xiang Li, and Jian Yang. Joint task-recursive learning for semantic segmentation and depth estimation. In *ECCV*, 2018.
- [7] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *CVPR*, 2018.
- [8] Yue Meng, Yongxi Lu, Aman Raj, Samuel Sunarjo, Rui Guo, Tara Javidi, Gaurav Bansal, and Dinesh Bharadia. Signet: Semantic instance aided unsupervised 3d geometry perception. In *CVPR*, 2019.
- [9] Jianbo Jiao, Ying Cao, Yibing Song, and Rynson Lau. Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss. In *ECCV*, 2018.
- [10] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *ECCV*, 2020.
- [11] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. 2020.
- [12] Yanwei Li, Hengshuang Zhao, Xiaojuan Qi, Liwei Wang, Zeming Li, Jian Sun, and Jiaya Jia. Fully convolutional networks for panoptic segmentation. In *CVPR*, 2021.
- [13] Bowen Cheng, Alexander G Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *arXiv preprint arXiv:2107.06278*, 2021.
- [14] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollar. Panoptic feature pyramid networks. In *CVPR*, 2019.
- [15] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [16] Huanyu Liu, Chao Peng, Changqian Yu, Jingbo Wang, Xu Liu, Gang Yu, and Wei Jiang. An end-to-end network for panoptic segmentation. In *CVPR*, 2019.
- [17] Yanwei Li, Xinze Chen, Zheng Zhu, Lingxi Xie, Guan Huang, Dalong Du, and Xingang Wang. Attention-guided unified network for panoptic segmentation. In *CVPR*, 2019.
- [18] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. UPSNet: A unified panoptic segmentation network. In *CVPR*, 2019.
- [19] Lorenzo Porzi, Samuel Rota Buló, Aleksander Colovic, and Peter Kotschieder. Seamless scene segmentation. In *CVPR*, 2019.
- [20] Rohit Mohan and Abhinav Valada. Efficientps: Efficient panoptic segmentation. *IJCV*, 2021.
- [21] Jie Li, Allan Raventos, Arjun Bhargava, Takaaki Tagawa, and Adrien Gaidon. Learning to fuse things and stuff. *arXiv preprint arXiv:1812.01192*, 2018.
- [22] Qizhu Li, Xiaojuan Qi, and Philip H.S. Torr. Unifying training and inference for panoptic segmentation. In *CVPR*, 2020.
- [23] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017.
- [24] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv:1706.05587*, 2017.
- [25] Liangchih Chen, George Papandreou, Iasonas Kokkinos, Kevin P Murphy, and Alan L Yuille. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4), 2018.
- [26] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- [27] Justin Lazarow, Kwonjoon Lee, Kunyu Shi, and Zhuowen Tu. Learning instance occlusion for panoptic segmentation. In *CVPR*, 2020.
- [28] Rui Hou, Jie Li, Arjun Bhargava, Allan Raventos, Vitor Guizilini, Chao Fang, Jerome Lynch, and Adrien Gaidon. Real-time panoptic segmentation from dense detections. In *CVPR*, 2020.
- [29] Weixiang Hong, Qingpei Guo, Wei Zhang, Jingdong Chen, and Wei Chu. Lpsnet: A lightweight solution for fast panoptic segmentation. In *CVPR*, 2021.
- [30] Tien-Ju Yang, Maxwell D. Collins, Yukun Zhu, Jyh-Jing Hwang, Ting Liu, Xiao Zhang, Vivienne Sze, George Papandreou, and Liang-Chieh Chen. DeeperLab: Single-shot image parser. *arXiv:1902.05093*, 2019.
- [31] Yiding Liu, Siyu Yang, Bin Li, Wengang Zhou, Jizheng Xu, Houqiang Li, and Yan Lu. Affinity derivation and graph merge for instance segmentation. In *ECCV*, 2018.
- [32] Naiyu Gao, Yanhu Shan, Yupei Wang, Xin Zhao, Yanan Yu, Ming Yang, and Kaiqi Huang. Ssap: Single-shot instance segmentation with affinity pyramid. In *ICCV*, 2019.
- [33] Naiyu Gao, Yanhu Shan, Yupei Wang, Xin Zhao, and Kaiqi Huang. Ssap: Single-shot instance segmentation with affinity pyramid. *IEEE TCSVT*, 2021.
- [34] Konstantin Sofiiuk, Olga Barinova, and Anton Konushin. AdaptIS: Adaptive instance selection network. In *ICCV*, 2019.

- [35] Davy Neven, Bert De Brabandere, Marc Proesmans, and Luc Van Gool. Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In *CVPR*, 2019.
- [36] George Papandreou, Tyler Zhu, Liang-Chieh Chen, Spyros Gidaris, Jonathan Tompson, and Kevin Murphy. Person-Lab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In *ECCV*, 2018.
- [37] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-DeepLab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *CVPR*, 2020.
- [38] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *ECCV*, 2020.
- [39] Liang-Chieh Chen, Huiyu Wang, and Siyuan Qiao. Scaling wide residual networks for panoptic segmentation. *arXiv preprint arXiv:2011.11675*, 2020.
- [40] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT: Real-time instance segmentation. In *ICCV*, 2019.
- [41] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. Solo: Segmenting objects by locations. In *ECCV*, 2020.
- [42] Wenwei Zhang, Jiangmiao Pang, Kai Chen, and Chen Change Loy. K-net: Towards unified image segmentation. *arXiv preprint arXiv:2106.14855*, 2021.
- [43] Naiyu Gao, Yanhu Shan, Xin Zhao, and Kaiqi Huang. Learning category-and instance-aware pixel embedding for fast panoptic segmentation. *IEEE Transactions on Image Processing*, 2021.
- [44] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *CVPR*, 2021.
- [45] Derek Hoiem, Alexei A Efros, and Martial Hebert. Geometric context from a single image. In *ICCV*, 2005.
- [46] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE TPAMI*, 2008.
- [47] Beyang Liu, Stephen Gould, and Daphne Koller. Single image depth estimation from predicted semantic labels. In *CVPR*, 2010.
- [48] Miaomiao Liu, Mathieu Salzmann, and Xuming He. Discrete-continuous depth estimation from a single image. In *CVPR*, 2014.
- [49] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *2016 Fourth international conference on 3D vision (3DV)*, 2016.
- [50] Jun Li, Reinhard Klein, and Angela Yao. A two-streamed network for estimating fine-scaled depth maps from single rgb images. In *ICCV*, 2017.
- [51] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *ECCV*, 2016.
- [52] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *ICCV*, 2019.
- [53] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017.
- [54] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *CVPR*, 2018.
- [55] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.
- [56] Amir Atapour-Abarghouei and Toby P Breckon. Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer. In *CVPR*, 2018.
- [57] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. T2net: Synthetic-to-realistic translation for solving single-image depth estimation tasks. In *ECCV*, 2018.
- [58] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. 2016.
- [59] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *CVPR*, 2019.
- [60] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *CVPR*, 2018.
- [61] 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)*, 2019.
- [62] Lijun Wang, Xiaohui Shen, Jianming Zhang, Oliver Wang, Zhe Lin, Chih-Yao Hsieh, Sarah Kong, and Huchuan Lu. Deeplens: shallow depth of field from a single image. *ACM Transactions on Graphics (TOG)*, 2018.
- [63] 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 *CVPR*, 2018.
- [64] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015.
- [65] Arsalan Mousavian, Hamed Pirsiavash, and Jana Košecká. Joint semantic segmentation and depth estimation with deep convolutional networks. In *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, 2016.
- [66] Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan L Yuille. Towards unified depth and semantic prediction from a single image. In *CVPR*, 2015.
- [67] Lijun Wang, Jianming Zhang, Oliver Wang, Zhe Lin, and Huchuan Lu. Sdc-depth: Semantic divide-and-conquer network for monocular depth estimation. In *CVPR*, 2020.

- [68] Faraz Saeedan and Stefan Roth. Boosting monocular depth with panoptic segmentation maps. In *WACV*, 2021.
- [69] Tsungyi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [70] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *NeurIPS*, 2014.
- [71] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- [72] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [73] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *ICCV*, 2019.
- [74] Francois Chollet. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, 2017.
- [75] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021.
- [76] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulo, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017.
- [77] Liang-Chieh Chen, Raphael Gontijo Lopes, Bowen Cheng, Maxwell D Collins, Ekin D Cubuk, Barret Zoph, Hartwig Adam, and Jonathon Shlens. Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In *ECCV*, 2020.
- [78] Benoit Steiner, Zachary Devito, Soumith Chintala, Sam Gross, Adam Paszke, Francisco Massa, Adam Lerer, Gregory Chanan, Zeming Lin, Edward Yang, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- [79] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [80] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [81] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [82] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E Reed, Chengyang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016.
- [83] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, 2019.