

# MGNiceNet: Unified Monocular Geometric Scene Understanding

Markus Schön<sup>✉</sup>, Michael Buchholz<sup>✉</sup>, and Klaus Dietmayer<sup>✉</sup>

Institute of Measurement, Control, and Microtechnology, Ulm University, Germany  
{markus.schoen,michael.buchholz,klaus.dietmayer}@uni-ulm.de

**Abstract.** Monocular geometric scene understanding combines panoptic segmentation and self-supervised depth estimation, focusing on real-time application in autonomous vehicles. We introduce MGNiceNet, a unified approach that uses a linked kernel formulation for panoptic segmentation and self-supervised depth estimation. MGNiceNet is based on the state-of-the-art real-time panoptic segmentation method RT-K-Net and extends the architecture to cover both panoptic segmentation and self-supervised monocular depth estimation. To this end, we introduce a tightly coupled self-supervised depth estimation predictor that explicitly uses information from the panoptic path for depth prediction. Furthermore, we introduce a panoptic-guided motion masking method to improve depth estimation without relying on video panoptic segmentation annotations. We evaluate our method on two popular autonomous driving datasets, Cityscapes and KITTI. Our model shows state-of-the-art results compared to other real-time methods and closes the gap to computationally more demanding methods. Source code and trained models are available at <https://github.com/markusschoen/MGNiceNet>.

**Keywords:** Monocular Geometric Scene Understanding · Panoptic Segmentation · Self-supervised Depth Estimation · Multi-task Learning

## 1 Introduction

Panoptic segmentation [19] was introduced to unify the tasks of semantic segmentation and instance segmentation by segmenting both amorphous stuff regions and countable thing objects. Stuff regions encompass semantic categories such as road or vegetation, while thing objects encompass instance categories such as pedestrian or vehicle. This information is highly relevant for downstream tasks such as trajectory planning to navigate through the vehicle’s environment safely. The major drawback of panoptic segmentation for its application in autonomous driving systems is its limitation to the 2D image plane. Depth information is essential to lift 2D information into our 3D world. Therefore, it comes naturally to combine panoptic segmentation with monocular depth estimation [6], which aims to recover the distance to the camera for each pixel in the image. Both tasks share similarities, *e.g.*, depth discontinuities at object borders, which can be exploited. Combined methods [7, 13, 33, 41, 49] often show improved results compared to single-task methods. However, most of these methods share

two main limitations for wide application in autonomous driving systems. First, depth estimation is trained in a supervised way using either stereo disparities or projected lidar points as ground truth. Supervised methods have limited generalization ability to different camera setups, and data acquisition for large-scale depth datasets is challenging. It is time-consuming, and sensors need to be well-calibrated. Second, most methods focus on high accuracy rather than a fast inference speed. Methods need a reasonable inference speed to process incoming images quickly and keep the latency in the autonomous driving system low. Otherwise, the system cannot react to sudden changes in the vehicle’s environment, *e.g.*, pedestrians crossing the street.

In our previous work [33], we introduced monocular geometric scene understanding, *i.e.*, a combination of panoptic segmentation and self-supervised monocular depth estimation [55], to tackle both of these limitations. The self-supervised approach to depth estimation alleviates the problem of acquiring per-pixel ground truth depth using an image-to-image synthesis formulation, which requires only captured video sequences during training. Thus, self-supervised methods can easily be adapted to new cameras by retraining the model with videos captured from the new camera. Furthermore, we introduced MGNet, a network architecture which focuses on low latency for real-time applications. Despite showing reasonable performance, MGNet has one main limitation: The task combination is only handled implicitly using a shared encoder and separate task-specific decoders. The implicit formulation limits the model from exploiting the inherent task relations between panoptic segmentation and self-supervised depth estimation, ultimately limiting model performance.

Motivated by this limitation, we propose MGNiceNet, a unified monocular geometric scene understanding method that explicitly exploits the relations between the two tasks. MGNiceNet is based on our state-of-the-art real-time panoptic segmentation method RT-K-Net [34], which we extend to additionally perform the task of self-supervised depth estimation. Specifically, we introduce a set of self-supervised depth kernels and employ depth kernel updates similar to PolyphonicFormer [49] to produce instance-wise depth predictions. Inspired by the recent success of classification-regression approaches to depth estimation [1, 42], MGNiceNet also predicts depth using a discrete bin-wise depth representation. However, compared to [1, 42], we propose to use a bin-wise representation at the panoptic mask level, allowing for a better alignment between the panoptic segmentation and depth prediction. Assumptions for self-supervised depth learning are a moving camera in a static environment and no occlusions between frames [55]. Thus, dynamic objects break this assumption, resulting in high photometric errors and incorrect depth predictions. Several methods successfully incorporated motion masking techniques into self-supervised training [20, 29]. However, these methods either require a static frame selection [20] or video panoptic annotations [29]. Thus, we propose a novel panoptic-guided motion masking method that uses panoptic predictions instead of video panoptic annotations. In short, our main contributions can be summarized as follows:

- We introduce MGNiceNet, a unified method for monocular geometric scene understanding that extends our state-of-the-art real-time panoptic segmentation method RT-K-Net [34] with self-supervised depth estimation using explicit kernel linking between panoptic and depth kernels.
- We propose a bin-wise depth representation at the panoptic mask level using a depth predictor module with low computational overhead.
- We introduce a novel panoptic-guided motion masking technique based on warped panoptic segmentation masks with tracked instances, which do not need static frame selection or video panoptic annotations.
- We evaluate our method on two standard autonomous driving benchmarks, Cityscapes [5] and KITTI [8], and show the effectiveness of our contributions in extensive experiments.

## 2 Related Work

### 2.1 Panoptic Segmentation

Panoptic segmentation [19] was introduced to unify semantic segmentation and instance segmentation. Early works [2, 3, 18, 25, 30] often treated panoptic segmentation as a multi-task problem and used a shared encoder with separate decoders for semantic segmentation and instance segmentation. Recently, unified panoptic segmentation approaches [4, 16, 39, 47, 48, 51, 52] emerged using a mask classification approach to tackle both semantic segmentation and instance segmentation in a unified way. K-Net [52] was one of the first unified approaches. It uses learnable kernels to produce mask proposals for both stuff and thing classes, which can be used to predict panoptic segmentation maps directly. Current state-of-the-art methods such as OneFormer [16] use transformer queries instead of learnable kernels, a multi-scale transformer decoder, and a task-conditioned text encoder to produce high-quality panoptic segmentation maps. While such methods achieve high accuracy, they often rely on heavy architectures, making them unsuitable for real-time applications.

Only a few panoptic segmentation methods focus on real-time application [2, 15, 28, 34]. While most methods [2, 15, 28] rely on specialized formulations for semantic segmentation and instance segmentation, our previously introduced unified approach RT-K-Net [34] achieves state-of-the-art performance in real-time panoptic segmentation. Thus, MGNiceNet is based on RT-K-Net, which we extend to the task of self-supervised depth estimation.

### 2.2 Self-supervised Monocular Depth Estimation

SfmLearner [55] was the first method to use monocular video sequences to train a depth estimation model in a self-supervised manner. Godard *et al.* [10] propose to upsample the multi-scale depth maps before loss calculation and use the minimum photometric error to tackle occlusions. Furthermore, they propose an auto-masking strategy to avoid holes of infinite depth in low-texture regions and

regions with dynamic objects. Their contributions advanced the field considerably and were integrated by many recent methods [11, 42, 44, 46, 54], and we also build our framework on their foundation. Wang *et al.* [42] recently introduced a Self Query Layer to learn geometric object features from pixel features. The features are then decoded using a bin-wise discrete depth representation at the geometric object level. In contrast to [42], we propose to predict depth bins at the panoptic mask level, which better aligns with the panoptic prediction without the need for the computationally heavy Self Query Layer.

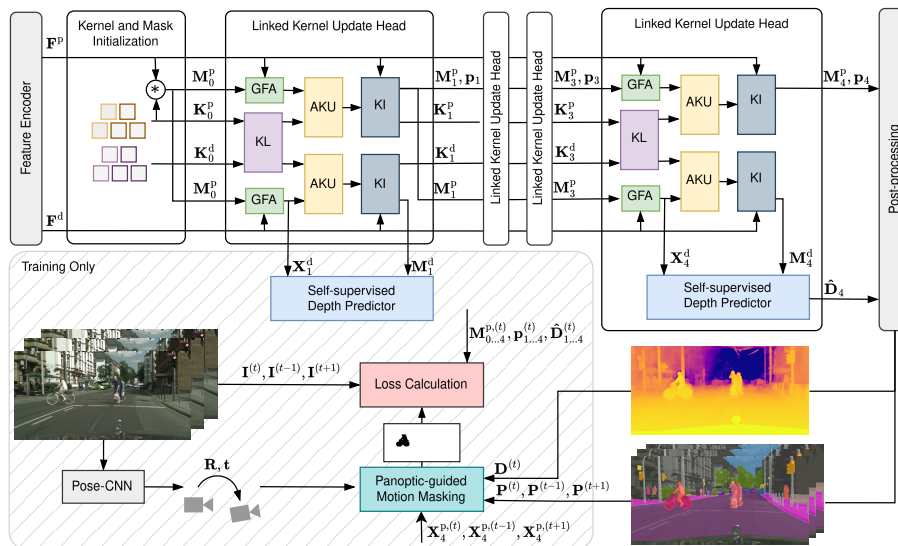
### 2.3 Combined Approaches

Some works focus on incorporating semantic [12, 20] or panoptic [32] annotations to boost the performance of self-supervised depth maps. For example, Klingner *et al.* [20] propose a multi-task model that performs semantic segmentation and self-supervised depth estimation. They use a motion masking technique based on the predicted semantic segmentation map to exclude dynamic objects from the photometric loss calculation. While they show the benefits of their masking technique, their approach is limited since object information is missing in the semantic segmentation prediction. Other works combine depth estimation with image [7, 13, 41, 45, 53] or video panoptic segmentation [17, 29, 31, 49]. Yuan *et al.* [49] introduce PolyphonicFormer, which uses query linking to boost depth estimation performance. Our method leverages kernel linking and combines it with a depth predictor for the self-supervised training objective. Petrovai and Nedeveschi [29] propose a combination of video panoptic segmentation and self-supervised depth estimation. They also propose a panoptic-guided motion masking technique, but their method relies on video panoptic annotations, which have a high annotation cost. In contrast, our motion masking technique only requires video sequence images since we calculate motion masks based on panoptic predictions. Our previous approach MGNet [33] combines panoptic segmentation with self-supervised depth estimation focusing on real-time application. However, MGNet uses separate task decoders for semantic segmentation, instance segmentation, and self-supervised depth estimation. For MGNiceNet, we propose to use a unified decoder approach using linked kernel update heads.

## 3 Method

### 3.1 Problem Formulation

Given a batch of  $B$  input color images  $\mathbf{I}^{(t)} \in \mathbb{R}^{B \times 3 \times H \times W}$  at time step  $t$  and spatial dimensions  $H$  and  $W$  the task of monocular geometric scene understanding is to infer a panoptic prediction  $\mathbf{P} \in \mathbb{N}^{B \times H \times W}$  and a depth prediction  $\mathbf{D} \in \mathbb{R}^{B \times H \times W}$ . The panoptic prediction assigns a unique class ID of a predefined number of classes  $N_c$  to each pixel in the image and a unique instance ID to each pixel corresponding to thing classes. The depth prediction assigns a metric depth value, *i.e.*, the metric distance to the camera, to each pixel in the image. Panoptic segmentation can be formulated as a mask classification problem [4]. The



**Fig. 1: Overview of our MGNiceNet architecture.** Images are fed into the Feature Encoder (Sec. 3.3) to produce high-resolution feature maps  $\mathbf{F}^p$  and  $\mathbf{F}^d$ . Next, the panoptic kernels  $\mathbf{K}_0^p$ , self-supervised depth kernels  $\mathbf{K}_0^d$ , and the panoptic masks  $\mathbf{M}_0^p$  are initialized (Sec. 3.4). Kernels and masks are updated iteratively in the four linked kernel update heads (Sec. 3.5). Each head contains Kernel Linking (KL), Group Feature Assembling (GFA), Adaptive Kernel Update (AKU), and Kernel Interaction (KI) stages and a self-supervised depth predictor (Sec. 3.6), which converts depth mask predictions  $\mathbf{M}_i^d$  and depth group features  $\mathbf{X}_i^d$  into an inverse depth prediction  $\hat{\mathbf{D}}_i^d$ . A post-processing stage (Sec. 3.7) converts mask predictions  $\mathbf{M}_4^p$ , class probability logits  $\mathbf{p}_4$ , and inverse depth predictions  $\hat{\mathbf{D}}_4^d$  into the final panoptic prediction  $\mathbf{P}$  and depth prediction  $\mathbf{D}$ . To improve optimization (Sec. 3.8), we introduce a panoptic-guided motion masking method (Sec. 3.9), which calculates a motion mask to mask out dynamic objects from the photometric loss.

idea is to predict a fixed number of  $N$  mask predictions  $\mathbf{M}^p \in \mathbb{R}^{B \times N \times H \times W}$  with corresponding class probability predictions  $\mathbf{p} \in \mathbb{R}^{B \times N \times N_c}$ . For stuff classes, a single mask per class is predicted, while each thing class can have multiple mask predictions depending on the number of objects present in the input image. During training, probability-mask pairs are assigned to ground truth using bipartite matching [4]. Depth estimation can be formulated in a self-supervised way as an image-to-image synthesis problem using monocular video sequences [55]. During training, based on the input reference frame  $\mathbf{I}^{(t)}$  at time step  $t$ , a depth prediction  $\mathbf{D}^{(t)}$  is used in conjunction with relative camera pose predictions to project adjacent video sequence frames, *i.e.*, context frames  $\mathbf{I}_c \in \{\mathbf{I}^{(t-1)}, \mathbf{I}^{(t+1)}\}$  into the reference frame. Optimization uses a photometric image reconstruction loss to maximize the similarity between the warped context frames and the reference frame.

### 3.2 Method Overview

An overview of our method is given in Fig. 1. MGNiceNet tackles monocular geometric scene understanding in a unified framework. The network architecture is built upon RT-K-Net [34], which is extended to self-supervised depth estimation. The following sections detail the proposed extension of the network architecture, the optimization process, and the proposed panoptic-guided motion masking.

### 3.3 Feature Encoder

Following RT-K-Net, RTFormer [40] is used as a feature encoder. RTFormer uses the input image batch to generate a single high-resolution feature map  $\mathbf{F}^p \in \mathbb{R}^{B \times C \times H' \times W'}$  with feature channel size  $C$  and  $H' = \frac{H}{8}, W' = \frac{W}{8}$ . For MGNiceNet, a second feature map  $\mathbf{F}^d \in \mathbb{R}^{B \times C \times H' \times W'}$  with the same dimensions is generated by adding a parallel convolution stage to the last encoder layer.

### 3.4 Kernel and Mask Initialization

RT-K-Net uses randomly initialized panoptic kernels  $\mathbf{K}_0^p \in \mathbb{R}^{B \times N \times C}$  to produce initial panoptic mask predictions by performing convolution  $\mathbf{M}_0^p = \mathbf{K}_0^p * \mathbf{F}^p$  using the initialized panoptic kernels and the panoptic feature map from the feature encoder. For MGNiceNet, self-supervised depth kernels  $\mathbf{K}_0^d \in \mathbb{R}^{B \times N \times C}$  are initialized similar to panoptic kernels. However, no depth mask predictions  $\mathbf{M}_0^d$  are initialized since panoptic masks are used to refine the self-supervised depth kernels in the linked kernel update heads instead.

### 3.5 Linked Kernel Update Heads

The initial kernels and mask predictions are not discriminative enough to perform panoptic segmentation [52]. Hence, RT-K-Net performs four iterative kernel updates to increase the representation ability of each kernel. The same holds for self-supervised depth kernels. Therefore, the kernel update heads of RT-K-Net are extended to update the self-supervised depth kernels. Each linked kernel update head  $i$  performs the following stages to update kernels and mask predictions from the last linked kernel update head  $i - 1$ :

**Kernel Linking (KL)** is performed similarly to query linking introduced in [49] to update self-supervised depth kernels by reusing the information present in panoptic kernels. The linking is unidirectional, *i.e.*, information is only passed from panoptic to self-supervised depth kernels using  $\tilde{\mathbf{K}}_i^d = \mathbf{K}_{i-1}^d + \mathbf{K}_{i-1}^p$ .

**Group Feature Assembling (GFA)** is used to reduce noise and speed up the learning process [52]. The panoptic mask predictions  $\mathbf{M}_{i-1}^p$  are first converted to binary mask predictions  $\mathbf{S}_{i-1}^p = \sigma(\mathbf{M}_{i-1}^p) > 0.5$ . Then, panoptic group features  $\mathbf{X}_i^p \in \mathbb{R}^{B \times N \times C}$  are formed using element-wise multiplication of the panoptic feature map  $\mathbf{F}^p$  and the normalized binary mask predictions

$$\mathbf{X}_i^p = \sum_u \sum_v \frac{\mathbf{S}_{i-1}^p(u, v)}{\sum_{H'W'} \mathbf{S}_{i-1}^p(u, v)} \mathbf{F}^p(u, v). \quad (1)$$

We use RT-K-Net’s normalization strategy to enable mixed-precision training and inference, which has been proven to reduce inference time [34]. Depth group features  $\mathbf{X}_i^d \in \mathbb{R}^{B \times N \times C}$  are generated similarly to Eq. (1) using the depth feature map  $\mathbf{F}^d$  and the normalized binarized mask predictions  $\mathbf{S}_{i-1}^p$ . This creates another link from the panoptic path to the depth path.

**Adaptive Kernel Update (AKU)** is used to re-weight the influence of panoptic kernels and panoptic group features [52] using learned gates  $G_{\mathbf{F}}, G_{\mathbf{K}}$  and linear transformations  $\psi_1, \psi_2$ . AKU is computationally inexpensive since it is independent of the image resolution. Thus, similar to [49], we use two parallel AKU modules to update the panoptic and self-supervised depth kernels

$$\tilde{\mathbf{K}}_i^p = G_{\mathbf{F}}^p \otimes \psi_1^p(\mathbf{X}_i^p) + G_{\mathbf{K}}^p \otimes \psi_2^p(\mathbf{K}_{i-1}^p) \quad (2)$$

$$\tilde{\mathbf{K}}_i^d = G_{\mathbf{F}}^d \otimes \psi_1^d(\mathbf{X}_i^d) + G_{\mathbf{K}}^d \otimes \psi_2^d(\tilde{\mathbf{K}}_i^d). \quad (3)$$

**Kernel Interaction (KI)** is performed to model the relationship between kernels, *i.e.*, between different pixel groups. Similar to AKU, this operation is inexpensive and thus extended by adding a second parallel KI module to update self-supervised depth kernels

$$\mathbf{K}_i^p = \text{FFN}_{\mathbf{K}}^p(\text{MHSA}^p(\tilde{\mathbf{K}}_i^p)), \quad \mathbf{K}_i^d = \text{FFN}_{\mathbf{K}}^d(\text{MHSA}^d(\tilde{\mathbf{K}}_i^d)) \quad (4)$$

using Multi-Head Self-Attention (MHSA) [38] followed by a Feed-Forward Neural Network (FFN). The updated kernels  $\mathbf{K}_i^p$  and  $\mathbf{K}_i^d$  are fed through separate FFNs

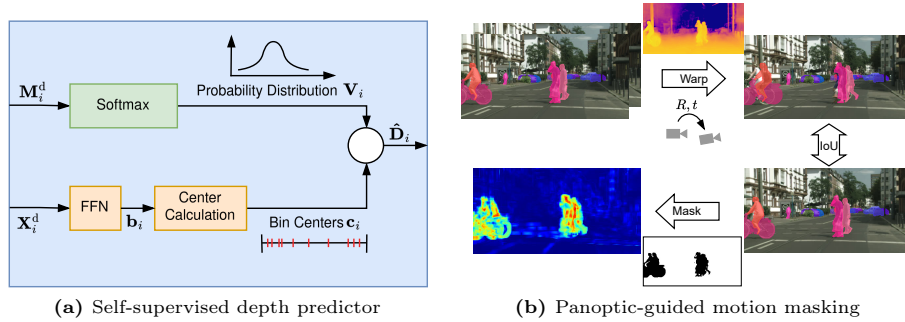
$$\mathbf{M}_i^p = \text{FFN}_{\mathbf{M}}^p(\mathbf{K}_i^p) * \mathbf{F}^p, \quad \mathbf{M}_i^d = \text{FFN}_{\mathbf{M}}^d(\mathbf{K}_i^d) * \mathbf{F}^d, \quad \mathbf{p}_i = \text{FFN}_{\mathbf{p}}(\mathbf{K}_i^p) \quad (5)$$

to generate refined panoptic mask predictions  $\mathbf{M}_i^p$ , self-supervised depth mask predictions  $\mathbf{M}_i^d$ , and a class probability prediction  $\mathbf{p}_i$ .

### 3.6 Self-supervised Depth Predictor

We introduce a self-supervised depth predictor module that efficiently generates an inverse depth prediction  $\hat{\mathbf{D}}_i$  based on depth mask predictions  $\mathbf{M}_i^d$  and depth group features  $\mathbf{X}_i^d$ . Inspired by recent approaches [1, 42], we reformulate depth estimation as a classification-regression problem and predict  $N$  depth bins  $\mathbf{b}_i \in \mathbb{R}^{B \times N}$  using a combination of depth planes and bin centers. Unlike recent methods, we directly use the encoded panoptic information present in our model to predict depth bins at the panoptic mask level. Figure 2a gives an overview of our self-supervised depth predictor module. The depth group features  $\mathbf{X}_i^d$  are calculated based on the depth feature map  $\mathbf{F}^d$  and the normalized binary panoptic mask predictions  $\mathbf{S}_{i-1}^p$  in the GFA stage of the linked kernel update heads. Thus, they encode object features at a panoptic mask level, which we leverage to predict depth bins at the panoptic mask level. To this end, we first reshape the depth group features to shape  $B \times NC$  and then feed them through a simple FFN

$$\mathbf{b}_i = \text{FFN}_{\mathbf{b}}^d(\mathbf{X}_i^d). \quad (6)$$



**Fig. 2:** Visualization of our self-supervised depth predictor module in (a) and our panoptic-guided motion masking in (b).

Following [1], we convert the predicted bins into bin centers  $\mathbf{c}_i \in \mathbb{R}^{B \times N}$

$$\mathbf{c}_i = d_{\min} + (d_{\max} - d_{\min}) \left( \frac{\mathbf{b}_i}{2} + \sum_{j=1}^{i-1} \mathbf{b}_j \right) \quad (7)$$

using the minimum depth  $d_{\min}$  and maximum depth  $d_{\max}$ . The depth mask predictions  $\mathbf{M}_i^d$  are converted into a probability distribution  $\mathbf{V}_i \in \mathbb{R}^{B \times N \times H' \times W'}$  over the predicted depth bins by performing a softmax operation over the mask predictions

$$\mathbf{V}_i = \text{softmax}(\mathbf{M}_i^d). \quad (8)$$

Finally, we use the linear combination of depth centers and softmax scores to calculate the inverse depth prediction  $\hat{\mathbf{D}}_i$ . Using the bin-wise formulation at the panoptic mask level has several advantages. First, it uses the panoptic information present in our model efficiently without relying on additional modules such as the Self-Query Layer used in [42], which is computationally expensive when processing high-resolution images ( $\mathcal{O}(H' \times W')$ ). Second, the resulting depth maps are aligned with the panoptic prediction and thus encompass fine-grained details.

### 3.7 Post-processing

We use the post-processing module introduced in RT-K-Net [34] to generate panoptic segmentation predictions based on the updated panoptic masks  $\mathbf{M}_4^p$  and class probability prediction  $\mathbf{p}_4$ . For depth estimation, we use the post-processing introduced in MGNet [33], where we use the panoptic road prediction as ground mask to calculate the depth scale factor based on the relation between the estimated median camera height for ground pixels and the calibrated camera height.



### 3.8 Optimization

For panoptic segmentation, we follow RT-K-Net and use bipartite matching to match mask predictions with ground-truth segments. The matched segments are used together with auxiliary outputs to calculate the panoptic loss

$$\mathcal{L}^P = \lambda_{\text{mask}}\mathcal{L}^{\text{mask}} + \lambda_{\text{dice}}\mathcal{L}^{\text{dice}} + \lambda_{\text{cls}}\mathcal{L}^{\text{cls}} + \lambda_{\text{rank}}\mathcal{L}^{\text{rank}} + \lambda_{\text{seg}}\mathcal{L}^{\text{seg}} + \lambda_{\text{inst}}\mathcal{L}^{\text{inst}} \quad (9)$$

with  $\mathcal{L}^{\text{mask}}$  being the binary cross-entropy loss,  $\mathcal{L}^{\text{dice}}$  being the dice loss [36],  $\mathcal{L}^{\text{cls}}$  being the focal loss for classification [22],  $\mathcal{L}^{\text{rank}}$  being the mask-id cross-entropy loss,  $\mathcal{L}^{\text{seg}}$  being the cross-entropy loss, and  $\mathcal{L}^{\text{inst}}$  being the instance discrimination loss [47]. The weight factors  $\lambda_{\text{mask}}, \lambda_{\text{dice}}, \lambda_{\text{cls}}, \lambda_{\text{rank}}, \lambda_{\text{seg}}$  and  $\lambda_{\text{inst}}$  are hyperparameters and used to balance the different loss terms. We follow common practice [10, 33] and use a combination of photometric loss and smoothness loss for self-supervised depth estimation

$$\mathcal{L}^d = \lambda_{\text{phot}}\mathcal{L}^{\text{phot}} + \lambda_{\text{smooth}}\mathcal{L}^{\text{smooth}}. \quad (10)$$

A self-supervised depth predictor module is added to each linked kernel update head to predict four inverse depth predictions  $\hat{\mathbf{D}}_{1..4}$ , which are used for loss calculation. We use a gradient-based smoothness loss [9]

$$\mathcal{L}^{\text{smooth}} = \sum_{i=1}^4 |\delta_x \hat{\mathbf{D}}_i^*| \exp(-|\delta_x \mathbf{I}^{(t)}|) + |\delta_y \hat{\mathbf{D}}_i^*| \exp(-|\delta_y \mathbf{I}^{(t)}|) \quad (11)$$

with predicted mean-normalized inverse depth  $\hat{\mathbf{D}}_i^*$  and reference image  $\mathbf{I}^{(t)}$ . For photometric loss calculation, we use a separate Pose-CNN [42] to predict the 6 DOF relative camera pose between adjacent frames. For each depth prediction  $\mathbf{D}_i = 1/\hat{\mathbf{D}}_i$ , the predicted poses are used to warp the context frames  $\mathbf{I}_c$  into the reference frame  $\mathbf{I}^{(t)}$ . For each warped image  $\hat{\mathbf{I}}_i^{(t)} \in \{\hat{\mathbf{I}}_i^{(t-1) \rightarrow (t)}, \hat{\mathbf{I}}_i^{(t+1) \rightarrow (t)}\}$ , the photometric loss is calculated according to

$$\mathcal{L}^{\text{phot}}(\mathbf{I}^{(t)}, \hat{\mathbf{I}}_i^{(t)}) = \alpha \frac{1 - \text{SSIM}(\mathbf{I}^{(t)}, \hat{\mathbf{I}}_i^{(t)})}{2} + (1 - \alpha) |\mathbf{I}^{(t)} - \hat{\mathbf{I}}_i^{(t)}| \quad (12)$$

using the Structured Similarity Index (SSIM) [43]. Following common practice [10], we set  $\alpha = 0.85$ . Furthermore, we use the minimum reprojection error between warped and unwarped frames to mask out static pixels and the auto masking strategy introduced in [10]. Finally, we calculate a panoptic-guided motion mask as described in the following section to mask out pixels that potentially belong to dynamic objects

$$\mathcal{L}^{\text{phot}} = \sum_{i=1}^4 \mathcal{L}^{\text{phot}}(\mathbf{I}^{(t)}, \hat{\mathbf{I}}_i^{(t)}) \odot \mathbf{M}_{\text{dyn}}. \quad (13)$$

The final loss is calculated as a combination of panoptic segmentation and depth loss terms

$$\mathcal{L} = \mathcal{L}^P + \lambda_{\text{depth}}\mathcal{L}^d \quad (14)$$

with weight factor  $\lambda_{\text{depth}}$  to balance both loss terms.

### 3.9 Panoptic-guided Motion Mask Calculation

To tackle the problem of dynamic objects, which break the static scene assumption in the self-supervised training objective, we introduce a panoptic-guided motion mask to mask out potential dynamic objects from the photometric loss calculation. Figure 2b gives an overview of the motion mask calculation. First, forward passes through the model are performed for the context images  $\mathbf{I}_c$  to receive panoptic predictions for both context images  $\mathbf{P}_c \in \{\mathbf{P}^{(t-1)}, \mathbf{P}^{(t+1)}\}$ . We employ a tracking module [35] to get consistent instance IDs for the three panoptic predictions  $\hat{\mathbf{P}}^{(t-1)}$ ,  $\hat{\mathbf{P}}^{(t)}$  and  $\hat{\mathbf{P}}^{(t+1)}$ . The advantage of this module is two-fold: First, it combines appearance-based matching and position-based matching for instance tracking. Second, the module does not require annotations with consistent instance IDs, which are costly to acquire for real-world applications. The tracked panoptic segmentation predictions are then warped into the reference frame using the predicted camera poses and depth prediction  $\mathbf{D}$ .

Finally, similar to [29], the intersection over unions (IoUs) between the warped and reference panoptic predictions are calculated for all segments corresponding to potential dynamic objects such as cars or pedestrians. Segments with a IoU below a threshold  $T$  are considered part of a dynamic object and thus added to the motion mask  $\mathbf{M}_{\text{dyn}}$ . In contrast to [29], we achieve the best results with a fixed threshold of  $T = 0.5$ . Furthermore, since we use predictions instead of ground truth segmentation maps, initial predictions contain much noise. Thus, we turn off the motion mask calculation for the first  $k = 1000$  training iterations until the model predicts reasonable object instances.

## 4 Experiments

### 4.1 Datasets

We provide experimental results on two popular autonomous driving benchmarks, Cityscapes [5] and KITTI [8]. The Cityscapes dataset provides video sequences recorded in 50 cities, primarily in Germany. It contains 5 000 images annotated with fine-grained panoptic labels for 30 classes, of which 19 are used during evaluation. The dataset is split into 2 975 images for training, 500 for validation, and 1 525 for testing. Depth annotations are given as pre-computed disparity maps from SGM [14] for the 5 000 fine annotated images, which are only used for depth evaluation. For KITTI, we use the KITTI Eigen split introduced in [6] with pre-processing from Zhou *et al.* [50] to remove static frames. This leads to an image split of 39 810 images for training and 4 424 for validation. Since the KITTI Eigen split does not provide panoptic annotations, we generate panoptic pseudo-labels using our best-performing model on Cityscapes and train our method using only the generated pseudo-labels. We evaluate depth estimation performance on KITTI on the 652 improved ground truth depth maps introduced in [37].

## 4.2 Implementation Details

All models are trained on 4 NVIDIA A6000 GPUs; inference times are measured on a single NVIDIA Titan RTX GPU without TensorRT [27] optimization. The AdamW optimizer [24] is used with the *poly* learning rate scheduler [23], and a base learning rate of 0.0001. We use a total batch size of 8, weight decay of 0.05, and train for 60k iterations. For data augmentation, we use random image scaling between 0.5 and 2.1, random cropping with the instance crop method introduced by RT-K-Net [34], random image color and brightness adjustment, and random left-right flip. For Cityscapes, we use a crop size of  $1024 \times 2048$  pixel; for KITTI, we resize images to a fixed size of  $384 \times 1280$  pixel. For panoptic loss weights, we follow RT-K-Net settings and set  $\lambda_{mask} = 1.0$ ,  $\lambda_{dice} = 4.0$ ,  $\lambda_{rank} = 0.1$ ,  $\lambda_{cls} = 2.0$ ,  $\lambda_{seg} = 1.0$ , and  $\lambda_{inst} = 1.0$ . For depth loss weights, we set  $\lambda_{phot} = 1.0$ ,  $\lambda_{smooth} = 0.001$ , and  $\lambda_{depth} = 25.0$  to balance panoptic and depth loss terms. We investigate the effect of different values for  $\lambda_{depth}$  in Sec. 4.4. We use  $N = 100$  and  $C = 256$ . Following recent combined methods [49, 51], we initialize our model with weights pre-trained on the Mapillary Vistas dataset [26] for panoptic segmentation unless stated otherwise. For Cityscapes, Panoptic Quality (PQ) [19] is the primary metric for panoptic segmentation, which can also be calculated for stuff (PQ<sub>st</sub>) and thing (PQ<sub>th</sub>) classes separately, while RMSE is the primary metric for self-supervised depth estimation. For KITTI, we report standard depth metrics as used in [6, 11, 42, 44, 46, 54]. The reader is referred to the respective publication [6, 19] for details on the respective metrics.

## 4.3 Experimental Results

**Cityscapes.** Table 1 shows results on the Cityscapes dataset compared to state-of-the-art methods. Table 1a shows the comparison with real-time panoptic segmentation methods while Table 1b shows the comparison with depth estimation methods. For a fair comparison, we report both results with and without Mapillary Vistas pre-training and inference times with and without automatic-mixed precision; denoted FP16 and FP32, respectively. For panoptic segmentation, MGNiceNet outperforms all other real-time capable methods regarding panoptic quality. Compared to our previous state-of-the-art approach RT-K-Net [34], PQ is improved by 1.5% without using additional data and 3.8 % when using Mapillary Vistas pre-training. Regarding runtime, MGNiceNet has a higher runtime of 42 ms than RT-K-Net’s 32 ms. Still, the computational overhead is marginal, given that MGNiceNet performs the additional task of self-supervised depth estimation. Compared to MGNet [33], MGNiceNet has a slightly higher runtime, but the performance gain in terms of accuracy is large enough to justify the higher runtime. A comparison of MGNiceNet to non-real-time panoptic segmentation methods, such as [49], is carried out in the paper’s supplementary. For depth estimation, MGNiceNet outperforms all other methods regarding the RMSE metric when using Mapillary Vistas pre-training. Without the pre-training, MGNiceNet still outperforms MGNet by 1 m RMSE and closes the gap to supervised methods [7, 41, 45, 53], which are not focused on real-time application, significantly.

**Table 1: Main results on the Cityscapes dataset.** A comparison to real-time panoptic segmentation methods on the Cityscapes validation set is shown in (a), while a comparison to depth estimation methods on the Cityscapes test set is shown in (b). Methods marked with † use additional data, *e.g.*, Mapillary Vistas [26] pre-training.

(a) Real-time panoptic segmentation						(b) Depth estimation		
Method	PQ †	PQ <sub>th</sub> †	PQ <sub>st</sub> †	GPU	Runtime [ms] ↓	Method	RMSE ↓	GT Depth
					FP16   FP32			
MGNet [33]†	55.7	45.3	63.1	Titan RTX	<u>36</u>   -	MGNet [33]†	8.3	
Panoptic-DeepLab [2]	58.4	-	-	V100	-   <b>63</b>	Pad-Net [45]	7.12	✓
Hou <i>et al.</i> [15]	58.8	52.1	63.7	V100	-   <u>99</u>	Zhang <i>et al.</i> [53]	7.10	✓
RT-K-Net [34]	60.2	51.5	66.5	Titan RTX	<b>32</b>   -	SDC-Depth [41]	6.92	✓
<b>Ours</b>	<u>61.7</u>	<u>54.6</u>	<u>66.8</u>	Titan RTX	42   <b>63</b>	PanopticDepth [7]†	<u>6.69</u>	✓
<b>Ours</b> †	<b>64.0</b>	<b>56.3</b>	<b>69.5</b>	Titan RTX	42   <b>63</b>	<b>Ours</b>	7.3	
						<b>Ours</b> †	<b>6.63</b>	

**Table 2: Main results on the KITTI Eigen split.** Depth estimation is evaluated for a distance up to 80 m using the improved ground truth maps from [37]. Runtimes marked with \* are calculated based on official code releases.

Method	DS	Resolution	Abs Rel ↓	RMSE ↓	$\delta < 1.25$ †	$\delta < 1.25^2$ †	$\delta < 1.25^3$ †	GPU	Runtime [ms] ↓	
									FP16   FP32	
Monodepth2 [10]	K	640 × 192	0.090	3.942	0.914	0.983	0.995	-	-	-
MGNet [33]	CS+K	1280 × 384	0.095	3.761	0.902	0.979	0.992	Titan RTX	<b>13*</b>	<u>34*</u>
SynDistNet [21]	K	640 × 192	0.076	3.406	0.931	0.988	0.996	-	-	-
PackNet-SfM [11]	CS+K	1280 × 384	0.071	3.153	0.944	0.990	0.997	V100	-	60
MonoViT [54]	K	1280 × 384	<u>0.067</u>	<u>3.108</u>	<u>0.950</u>	<u>0.992</u>	<b>0.998</b>	Titan RTX	41*	56*
SQLdepth [42]	K	1024 × 320	<b>0.058</b>	<b>2.925</b>	<b>0.962</b>	<b>0.993</b>	<b>0.998</b>	Titan RTX	99*	107*
<b>Ours</b>	CS+K	1280 × 384	0.069	3.143	0.949	<u>0.992</u>	<b>0.998</b>	Titan RTX	<u>21</u>	<b>28</b>

**KITTI Eigen.** Table 2 shows results on the KITTI Eigen split compared to state-of-the-art methods in self-supervised depth estimation. MGNiceNet outperforms most methods [10, 21, 33] by a large margin and even outperforms PackNet-SfM [11], which uses a much heavier architecture. Only the state-of-the-art approaches MonoViT [54] and SQLdepth [42] slightly outperform MGNiceNet in most metrics. We reason that this is primarily because our method is designed with a focus on low latency, while the aforementioned approaches rely on computationally expensive layers such as the Self Query Layer used in [42]. This becomes apparent when comparing inference times. Since most methods do not provide inference times in the paper, we use the author’s official codebases, if available, to get comparable inference times. MGNiceNet has an inference time of 21 ms when using automatic mixed-precision, which is around 75% faster than SQLdepth and 50 % faster than MonoViT. Thus, MGNiceNet provides a much better accuracy-speed trade-off than [42, 54]. Furthermore, MGNiceNet tackles the additional task of panoptic segmentation. Compared to MGNet, we see similar results as on Cityscapes, with MGNet slightly outperforming MGNiceNet in terms of runtime. We show qualitative results on Cityscapes and KITTI and direct qualitative comparisons in the paper’s supplementary.

Table 3: **Ablation studies for MGNiceNet.** All experiments are conducted on the Cityscapes validation dataset.

(a) Multi-task learning vs single-task performance						(b) Loss balancing		
Method	Depth	Panoptic	PQ $\uparrow$	RMSE $\downarrow$	Runtime [ms] $\downarrow$	$\lambda_{depth}$	PQ $\uparrow$	RMSE $\downarrow$
Depth	$\checkmark$	-	-	7.9	27	5.0	64.0	7.7
Panoptic	-	$\checkmark$	64.0	-	32	15.0	63.9	7.5
Ours	$\checkmark$	$\checkmark$	64.0	7.1	42	25.0	64.0	7.1
						30.0	63.3	7.1


(c) Motion masking				(d) Depth predictor choice			
Method	Motion masking	PQ $\uparrow$	RMSE $\downarrow$	Method	PQ $\uparrow$	RMSE $\downarrow$	Runtime [ms] $\downarrow$
Ours	-	63.6	7.8	Instance-Wise	63.3	7.3	45
	$\checkmark$	64.0	7.1	Ours	64.0	7.1	42

#### 4.4 Ablation Studies

**Multi-Task learning vs Single-Task Performance.** We compare our multi-task approach to single-task baselines by training MGNiceNet without depth estimation or panoptic segmentation, respectively. As seen in Tab. 3a, our multi-task approach has the same performance in terms of panoptic quality but improves depth performance compared to the single-task baselines. Furthermore, our method only adds marginal runtime while handling both tasks. We argue that this improved depth performance results from our contributions, namely our unified MGNiceNet architecture, the self-supervised depth predictor, and our panoptic-guided motion masking strategy.

**Loss Balancing.** We compare models trained with different depth loss weights  $\lambda_{depth} \in \{5.0, 15.0, 25.0, 30.0\}$  in Tab. 3b. As expected, depth estimation performance increases for higher values of  $\lambda_{depth}$  while panoptic segmentation performance decreases. Therefore, we choose  $\lambda_{depth} = 25.0$ , which achieves the best performance in both tasks.

**Motion Masking.** We investigate the effects of our panoptic-guided motion masking as described in Sec. 3.9 by training a model without and with our proposed motion masking approach. Tab. 3c shows that adding our panoptic-guided motion masking significantly improves depth estimation performance while panoptic quality remains nearly the same. This can also be seen in Fig. 3, which shows depth predictions of a model trained with and without our motion masking method. We want to especially highlight that the model trained without our panoptic-guided motion masking still uses MonoDepth2’s [10] auto-masking strategy, which fails for datasets with a high number of dynamic objects such as Cityscapes.



**Fig. 3:** Effect of our panoptic-guided motion masking. While the model predicts holes of infinite depth on pixels corresponding to the car driving with a similar velocity as the ego vehicle, the effect is reduced significantly when using our panoptic-guided motion masking method during training.

**Depth Predictor Choice.** We compare our self-supervised depth predictor as described in Sec. 3.6 with an instance-wise depth predictor. The instance-wise depth predictor uses depth mask predictions  $\mathbf{M}_4^d$  and binarized mask predictions  $\mathbf{S}_4^p$  to generate a depth prediction by gathering all pixels in  $\mathbf{M}_4^d$  where  $\mathbf{S}_4^p = 1$ . This is a similar strategy to PolyphonicFormer [49], which trains instance-wise depth predictions in a supervised way. Table 3d shows that our self-supervised depth predictor outperforms the instance-wised predictor in both metrics. Interestingly, the instance-wise predictor also negatively affects panoptic segmentation performance. We argue this is due to our method’s lack of instance-wise supervision with ground truth depth labels.

**Limitations and Future Work.** Our method is optimized to process high-resolution images at a fast inference speed. As described in Sec. 4.3, recent state-of-the-art methods still slightly outperform our approach. In future work, we aim to improve performance further, especially when dealing with lower-resolution input images, such as images from the KITTI dataset.

## 5 Conclusion

This work introduces MGNiceNet, a unified monocular geometric scene understanding approach. We extend our real-time panoptic segmentation method RT-K-Net to tackle the additional task of self-supervised depth estimation using linked kernel update heads and a lightweight depth predictor. Our self-supervised depth predictor uses a bin-wise depth representation at the panoptic mask level to predict high-quality depth maps that align well with panoptic predictions. We tackle the dynamic object problem in self-supervised depth estimation using a novel panoptic-guided motion masking method. We show our method’s effectiveness and limitations in extensive experiments on the Cityscapes and KITTI datasets. Our method provides an excellent speed-accuracy trade-off and closes the gap to more computationally demanding methods.

**Acknowledgments.** This research is accomplished within the project “AUTOftech.agil” (FKZ 01IS22088W). We acknowledge the financial support for the project by the Federal Ministry of Education and Research of Germany (BMBF).

## References

1. Bhat, S.F., Alhashim, I., Wonka, P.: Adabins: Depth estimation using adaptive bins. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 4009–4018 (2021)
2. Chen, L.C., Wang, H., Qiao, S.: Scaling wide residual networks for panoptic segmentation. *arXiv preprint arXiv:2011.11675* (2020)
3. Cheng, B., Collins, M.D., Zhu, Y., Liu, T., Huang, T.S., Adam, H., Chen, L.C.: Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 12472–12482 (2020)
4. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems (NeurIPS)* pp. 17864–17875 (2021)
5. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 3213–3223 (2016)
6. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. *Advances in Neural Information Processing Systems (NIPS)* **27** (2014)
7. Gao, N., He, F., Jia, J., Shan, Y., Zhang, H., Zhao, X., Huang, K.: Panopticdepth: A unified framework for depth-aware panoptic segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 1632–1642 (2022)
8. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)* **32**, 1231–1237 (2013)
9. Godard, C., Aodha, O.M., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 6602–6611 (2017)
10. Godard, C., Aodha, O.M., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* pp. 3827–3837 (2019)
11. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3d packing for self-supervised monocular depth estimation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 2482–2491 (2020)
12. Guizilini, V., Hou, R., Li, J., Ambrus, R., Gaidon, A.: Semantically-guided representation learning for self-supervised monocular depth. *International Conference on Learning Representations (ICLR)* (2020)
13. He, J., Wang, Y., Wang, L., Lu, H., Luo, B., He, J.Y., Lan, J.P., Geng, Y., Xie, X.: Towards deeply unified depth-aware panoptic segmentation with bi-directional guidance learning. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* pp. 4111–4121 (2023)
14. Hirschmuller, H.: Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**, 328–341 (2008)
15. Hou, R., Li, J., Bhargava, A., Raventos, A., Guizilini, V., Fang, C., Lynch, J., Gaidon, A.: Real-time panoptic segmentation from dense detections. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 8520–8529 (2020)

16. Jain, J., Li, J., Chiu, M.T., Hassani, A., Orlov, N., Shi, H.: Oneformer: One transformer to rule universal image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2989–2998 (2023)
17. Ji-Yeon, K., Hyun-Bin, O., Kim, D., Oh, T.H.: Mindvps: Minimal model for depth-aware video panoptic segmentation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2023)
18. Kirillov, A., Girshick, R., He, K., Dollar, P.: Panoptic feature pyramid networks. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 6392–6401 (2019)
19. Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 9404–9413 (2019)
20. Klingner, M., Termöhlen, J.A., Mikolajczyk, J., Fingscheidt, T.: Self-Supervised Monocular Depth Estimation: Solving the Dynamic Object Problem by Semantic Guidance, vol. 20, pp. 582–600. Springer (2020)
21. Kumar, V.R., Klingner, M., Yogamani, S., Milz, S., Fingscheidt, T., Mader, P.: Syndistnet: Self-supervised monocular fisheye camera distance estimation synergized with semantic segmentation for autonomous driving. Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) pp. 61–71 (2021)
22. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. Proceedings of the IEEE International Conference on Computer Vision (ICCV) pp. 2980–2988 (2017)
23. Liu, W., Rabinovich, A., Berg, A.C.: Parsenet: Looking wider to see better. International Conference on Learning Representations (ICLR) Workshops (2016)
24. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. International Conference on Learning Representations (ICLR) (2019)
25. Mohan, R., Valada, A.: Efficienttps: Efficient panoptic segmentation. International Journal of Computer Vision (IJCV) **129**, 1551–1579 (2021)
26. Neuhold, G., Ollmann, T., Bulo, S.R., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. Proceedings of the IEEE International Conference on Computer Vision (ICCV) pp. 5000–5009 (2017)
27. NVIDIA: Tensorrt library, <https://developer.nvidia.com/tensorrt>
28. Petrovai, A., Nedeveschi, S.: Real-time panoptic segmentation with prototype masks for automated driving. IEEE Intelligent Vehicles Symposium (IV) pp. 1400–1406 (2020)
29. Petrovai, A., Nedeveschi, S.: Monodvps: A self-supervised monocular depth estimation approach to depth-aware video panoptic segmentation. Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) pp. 3077–3086 (2023)
30. Porzi, L., Bulo, S.R., Colovic, A., Kotschieder, P.: Seamless scene segmentation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 8269–8278 (2019)
31. Qiao, S., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: Vip-deeplab: Learning visual perception with depth-aware video panoptic segmentation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 3997–4008 (2021)
32. Saedan, F., Roth, S.: Boosting monocular depth with panoptic segmentation maps. Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) pp. 3853–3862 (2021)



33. Schön, M., Buchholz, M., Dietmayer, K.: Mgnet: Monocular geometric scene understanding for autonomous driving. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* pp. 15784–15795 (2021)
34. Schön, M., Buchholz, M., Dietmayer, K.: Rt-k-net: Revisiting k-net for real-time panoptic segmentation. *IEEE Intelligent Vehicles Symposium (IV)* pp. 1–7 (2023)
35. Stolle, K., Dubbelman, G.: Unified perception: Efficient video panoptic segmentation with minimal annotation costs. *arXiv preprint arXiv:2303.01991* (2023)
36. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3.* pp. 240–248. Springer (2017)
37. Uhrig, J., Schneider, N., Schneider, L., Franke, U., Brox, T., Geiger, A.: Sparsity invariant cnns. *International Conference on 3D Vision (3DV)* pp. 11–20 (2017)
38. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)* **30** (2017)
39. Wang, H., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: Max-deeplab: End-to-end panoptic segmentation with mask transformers. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 5463–5474 (2021)
40. Wang, J., Gou, C., Wu, Q., Feng, H., Han, J., Ding, E., Wang, J.: Rtformer: Efficient design for real-time semantic segmentation with transformer. *Advances in Neural Information Processing Systems (NeurIPS)* pp. 7423–7436 (2022)
41. Wang, L., Zhang, J., Wang, O., Lin, Z., Lu, H.: Sdc-depth: Semantic divide-and-conquer network for monocular depth estimation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 541–550 (2020)
42. Wang, Y., Liang, Y., Xu, H., Jiao, S., Yu, H.: Sqldepth: Generalizable self-supervised fine-structured monocular depth estimation. *arXiv preprint arXiv:2309.00526* (2023)
43. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**, 600–612 (2004)
44. Watson, J., Aodha, O.M., Prisacariu, V., Brostow, G., Firman, M.: The temporal opportunist: Self-supervised multi-frame monocular depth. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 1164–1174 (2021)
45. Xu, D., Ouyang, W., Wang, X., Sebe, N.: Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 675–684 (2018)
46. Xue, F., Zhuo, G., Huang, Z., Fu, W., Wu, Z., Ang, M.H.: Toward hierarchical self-supervised monocular absolute depth estimation for autonomous driving applications. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* pp. 2330–2337 (2020)
47. Yu, Q., Wang, H., Kim, D., Qiao, S., Collins, M., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: Cmt-deeplab: Clustering mask transformers for panoptic segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 2560–2570 (2022)

48. Yu, Q., Wang, H., Qiao, S., Collins, M., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: k-means Mask Transformer, vol. 29, pp. 288–307. Springer (2022)
49. Yuan, H., Li, X., Yang, Y., Cheng, G., Zhang, J., Tong, Y., Zhang, L., Tao, D.: PolyphonicFormer: Unified Query Learning for Depth-aware Video Panoptic Segmentation, vol. 27, pp. 582–599. Springer (2022)
50. Zhan, H., Garg, R., Weerasekera, C.S., Li, K., Agarwal, H., Reid, I.: Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 340–349 (2018)
51. Zhang, J., Huang, J., Lu, S.: Unidaformer: Unified domain adaptive panoptic segmentation transformer via hierarchical mask calibration. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 11227–11237 (2023)
52. Zhang, W., Pang, J., Chen, K., Loy, C.C.: K-net: Towards unified image segmentation. Advances in Neural Information Processing Systems (NeurIPS) pp. 10326–10338 (2021)
53. Zhang, Z., Cui, Z., Xu, C., Jie, Z., Li, X., Yang, J.: Joint Task-Recursive Learning for Semantic Segmentation and Depth Estimation, vol. 10, pp. 238–255. Springer (2018)
54. Zhao, C., Zhang, Y., Poggi, M., Tosi, F., Guo, X., Zhu, Z., Huang, G., Tang, Y., Mattoccia, S.: Monovit: Self-supervised monocular depth estimation with a vision transformer. International Conference on 3D Vision (3DV) pp. 668–678 (2022)
55. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 1851–1858 (2017)