This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Video Class Agnostic Segmentation Benchmark for Autonomous Driving

Mennatullah Siam University of Alberta Edmonton, Canada mennatul@ualberta.ca Alex Kendall Wayve London, UK alex@wayve.ai Martin Jagersand University of Alberta Edmonton, Canada jag@cs.ualberta.ca

Abstract

Semantic segmentation approaches are typically trained on large-scale data with a closed finite set of known classes without considering unknown objects. In certain safety-critical robotics applications, especially autonomous driving, it is important to segment all objects, including those unknown at training time. We formalize the task of video class agnostic segmentation from monocular video sequences in autonomous driving to account for unknown objects. Video class agnostic segmentation can be formulated as an open-set or a motion segmentation problem. We discuss both formulations and provide datasets and benchmark different baseline approaches for both tracks. In the motionsegmentation track we benchmark real-time joint panoptic and motion instance segmentation, and evaluate the effect of ego-flow suppression. In the open-set segmentation track we evaluate baseline methods that combine appearance, and geometry to learn prototypes per semantic class. We then compare it to a model that uses an auxiliary contrastive loss to improve the discrimination between known and unknown objects. Datasets and models are publicly released at https://msiam.github.io/vca/.

1. Introduction

Semantic scene understanding is crucial in autonomous driving in both end-to-end and mediated perception approaches as described in [4]. Semantic segmentation which performs pixel-wise classification of the scene is mostly trained on large scale data with closed set of known classes [5]. A system trained on a limited set of classes would face difficulties in unexpected situations that could occur in different autonomous driving scenarios, such as the animal or parking scenario objects shown in Figure 1. These scenarios demonstrate potential objects outside the closed set of known classes that are mostly available in public datasets [5]. Some work was conducted on zeroshot/few-shot learning [2, 27] to enable deep neural networks to generalize to novel classes with limited data or



Figure 1: Video Class Agnostic Segmentation enables identification of objects outside the closed set of known classes. Predictions of known and unknown objects on Cityscapes-VPS [11], IDD [24], and our Synthetic Dataset.

side information, but does not address the main problem. For safety reasons, autonomous driving systems should be able to detect objects that are not known beforehand, which are called *"unknown unknown classes"* [7].

In this paper we formulate the task of video class agnostic segmentation to account for unknown objects and publicly release the necessary datasets and baselines for the different task formulation. Video class agnostic segmentation is defined as the task of segmenting objects without regards to their semantics, but combining appearance, motion and geometry from monocular video sequences. It is crucial first to identify how to formulate this problem. Since the main goal is to segment obstacles without regards to their semantics, we choose two main tracks for our benchmark: (1) The **motion segmentation track** that will identify moving objects regardless of their semantics. (2) The **open-set segmentation track** that will identify pixels belonging to objects outside the closed set of known classes without identifying its exact semantic class. The open-set segmentation formulation is harder than motion segmentation but has the clear advantage of identifying static objects. Examples of both tracks are shown in Figure 1.

Motion segmentation has been recently studied for autonomous driving in [21, 20, 16, 25]. Unlike previous literature, we focus on the joint panoptic and motion instance segmentation as a means to improve hollistic scene understanding. We provide motion segmentation labels on both Cityscapes-VPS [11] and KITTIMOTS [26] datasets. Our extended motion segmentation dataset increases the number of labelled sequences and object categories while providing manually labelled instance segmentation masks, unlike previous work that used weakly annotated masks [16]. On the other hand, the open-set segmentation task for autonomous driving has not been thoroughly studied in autonomous driving [30]. In our work we explore contrastive learning to improve the model's ability to segment known and unknown objects with any baseline model. We propose a contrastive learning method similar to [29], but their work focused solely on the classification task we extend it to the segmentation of unknown objects. To summarize, our main contributions are:

- We formalize the task of *Video Class Agnostic Segmentation* and benchmark different baseline approaches for both motion and open-set segmentation tracks.
- Our motion segmentation dataset offers more labelled sequences and object categories on both KITTI [26] and Cityscapes [11] while providing manually labelled dense masks.
- For the open-set segmentation track, we release a large-scale synthetic dataset for different autonomous driving scenarios. These scenarios can be included as part of the Carla Autonomous Driving Challenge [6] to drive forward research in both perception and policy learning.

2. Related Work

Datasets: Since our goal is performing *video class agnostic segmentation*, a related task is generic video object segmentation that is not focused on autonomous driving. The Densely Annotated VIdeo Segmentation (DAVIS) [18, 19] dataset provided a benchmark for both unsupervised and semi-supervised video object segmentation tracks. The Freiburg Berkley Motion Segmentation (FBMS) [1] and

SegTrack [23] provided a smaller scale video object segmentation dataset. However, all these generic datasets have much simpler scenes to segment than autonomous driving scenes. They mainly contain few visually salient objects in short video sequences that sometimes can be of low resolution. The camera motion is much simpler than what is experienced in autonomous driving which can significantly affect the methods trained for class agnostic segmentation based on motion.

More datasets in autonomous driving were recently released for the motion segmentation task [25, 21, 20]. However, they focused solely on vehicles which limits the generalization of the class agnostic segmentation to unknown moving objects. We resolve this issue in our provided dataset for motion segmentation. Recently, the Rare4D [30] dataset was collected for LIDAR data to segment different rare and unknown objects. However, the dataset was not publicly released and is focused on LIDAR data. We release the first class agnostic segmentation dataset focused on monocular video sequences from cameras, and provide further analysis on the relation between objects labelled as unknown between the training and testing phases.

Motion Segmentation: has been recently studied in the literature of autonomous driving with supervised learning methods [25, 21, 20, 16]. Vertens et. al.[25] proposed a joint model for semantic and motion segmentation with egoflow suppression, and released Cityscapes-KITTI-Motion dataset for only car category. Siam et. al.[21] concurrently released one of the earliest motion segmentation dataset KITTIMoSeg, that was extended by Rashed et al. [20] and with instance labels by Mohammed et. al. [16]. In our work we mainly focus on extending motion segmentation datasets to include more object categories to have eight object categories instead of only car and increase the number of sequences by $50\times$. These help to improve the generalization ability required to enable the segmentation of unknown moving objects. Our dataset also provides motion annotations on both KITTI and Cityscapes instead of focusing only on KITTI as [16]. We also improve the annotations by providing manually labelled segmentation masks, unlike weakly annotated ones [16]. Additionally, the corresponding tracking labels and panoptic segmentation labels from the original datasets are provided. A full comparison of these datasets is shown in Table 1.

Unknown Objects Segmentation: in autonomous driving has not been thoroughly studied in the literature with only two recent works [30, 17]. Wong et. al. [30] proposed a method inspired by [22] and trained on TOR4D, which is a large scale LIDAR dataset, and presented the Rare4D dataset to evaluate unknown objects. However, Rare4D dataset is not publicly available and shares no indication of the relation between unknown objects labelled between the training and testing datasets. Osep et al. [17] proposed a method that uses a video sequence of stereo images to predict 4D generic proposals for autonomous driving. But they evaluate only on 150 images for the autonomous driving dataset with unknown objects, which is insufficient to evaluate model scalability. They build their model on a twostage instance segmentation method, which is computationally inefficient, while most of the unknown objects belong to stuff classes such as cones, traffic warning, and barriers. These classes do not require the separation of instances as they will not be tracked. These limitations in the literature drove us to construct a benchmark for monocular video sequences for the open-set unknown objects track.

3. Motion Segmentation Track

3.1. Dataset

For motion segmentation we provide motion annotations to extend the original KITTIMOTS [26] and Cityscapes-VPS [11] datasets as shown in Figure 2. A trajectory annotation tool was built to annotate the trajectories for either moving or static and is further used to provide instance-wise motion masks. Table 1 shows the different statistics for our dataset in comparison to others. Our dataset has the main advantage of providing a larger variety of object categories and $50 \times$ increase in video sequences. These substantially improve the scalability of video class agnostic segmentation. Crucially, this pushes models to depend more on motion information and only use appearance information to detect objectness rather than learning to semantically detect cars only.

Although DAVIS provides a larger variety of object categories, the camera motion and scenes in these sequences are much simpler and easier to segment unlike our autonomous driving setting. We also provide instance masks unlike most previous literature which is crucial for motion instance segmentation. Concurrent to our work [16] provided instancewise motion masks, however we provide manually labelled segmentation masks unlike their weakly annotated ones. More importantly we provide annotations for both KITTI and Cityscapes, while their dataset is built solely on KITTI. This can affect the generalization ability of the class agnostic segmentation. Our dataset also benefits from the already provided panoptic labels for 3000 frames and tracking annotations for all frames which can aid in tracking moving objects. The statistics of moving and static instances for both datasets are shown in Table 2.

3.2. Baselines and Metrics

We propose a real-time two-stream multitask model that performs panoptic and motion instance segmentation as a baseline model. Our backbone network is a two-stream ResNet-50 [9] with a one-way feature pyramid network [13]. We extend the real-time instance segmenta-

tion model from SOLO [28] with a semantic segmentation head to segment the stuff classes and a class agnostic head that segments moving instances. The class agnostic head provides redundant signal for a safety-critical approach. Our panoptic segmentation model is initially trained on Cityscapes [5] and Cityscapes-VPS [11] datasets. Then the class agnostic head is trained on our motion segmentation datasets with fixed weights for the rest of the model. The input to the two-stream model is both appearance and motion as optical flow encoded in RGB. We compare two baselines the one with input optical flow estimated directly from Flownet-2 [10], while the other takes ego-flow suppressed one. The ego flow suppression uses the estimated depth and pose from [8], then computes the ego flow following the method from [25]. The ego-flow suppressed output is the result of subtracting the ego flow from the original estimated flow. The instance segmentation and class agnostic segmentation heads are trained with focal loss [14], and dice loss [15], while the segmentation head is trained with cross entropy.

We use the same evaluation metrics provided by [30] for measuring unknown instance segmentation. We report the class agnostic quality metric as shown in 1, which combines the segmentation and recognition quality. TP and FN are the true positives and false negatives respectively. Only instance masks that have intersection over union with the ground-truth above 0.5 are considered true positives.

$$SQ = \frac{\sum_{p,g \in TP} IoU(p,g)}{|TP|}$$
(1a)

$$RQ = \frac{|TP|}{|TP| + |FN|}$$
(1b)

$$CAQ = SQ.RQ$$
(1c)

4. Open-set Segmentation Track

In this section we discuss the synthetic dataset that we curated to conduct controlled experiments with a large-scale dataset to ensure the generalization ability of class agnostic segmentation of unknown objects.

4.1. Carla Scenarios and Dataset

In the open-set segmentation formulation we care about providing video sequences along with annotations for unknown objects in different autonomous driving scenarios. Thus, we build different scenarios within the Carla simulation environment [6]. The goal is to incorporate these scenarios as part of the Carla challenge for autonomous driving ¹ to benefit both perception and policy learning re-

¹See the Carla challenge here: https://leaderboard.carla.org/challenge/

Table 1: Comparison of different datasets for motion segmentation. M: stands for manually labelled annotations, W: stands for weakly labelled annotations. Seqs: sequences, Cats: categories. Our dataset provides instance-wise manually labelled annotations for moving objects and increases the sequences and categories unlike previous datasets.

Dataset	# Frames	# Seqs	# Cats	Instances	Panoptic	Tracks	Annotation
DAVIS [3]	6208	90	78	1	X	1	М
KITTI Motion [25]	455	-	1 (Car Only)	X	X	X	М
KITTI-MoSeg [21][20]	12919	~ 38	1 (Car Only)	X	X	X	W
InstKITTI-MoSeg[16]	12919	~ 38	5	1	X	X	W
Cityscapes Motion[25]	3475	-	1 (Car Only)	1	1	X	М
VCAS-Motion (ours)	11008	520	8	1	1	1	М



Figure 2: Our Extended Motion Annotations showing two consecutive frames on: (a) Top Row: Cityscapes datasets and (b) Bottom Row: KITTI. Red: moving instances, Blue: static instances.

Table 2: Dataset Statistics of moving and static instances.

	KITTI	Cityscapes
Moving	5788	41334
Static	6094	43981

searchers. Such scenarios serve as a way to evaluate the robustness of autonomous driving systems and the safetycritical design. Additionally, the available video dataset with dense labels (Cityscapes-VPS) has 3000 images and 300 for evaluation. Thus we build our aggregate data with approximately 70K frames.

First, we insert unknown objects and modify the basic virtual driving agent in Carla to avoid obstacles, in order to collect large-scale data with ground-truth depth and semantic segmentation labels. We then randomize object placement, weather condition, traffic (vehicles and pedestrian) and use different unknown objects and towns between training and testing as shown in Table 3. Figure 3 lists the three

main scenarios that are used to evaluate the open-set segmentation task, and Figure 5 shows the class statistics.

Table 3: Towns and Images collected in both training and testing splits for our open-set track in Carla.

	Training	Testing
	Town1	Town10
Towns	Town2	Town7
	Town3	
# Images	57972	1886

4.2. Analysis on the Unknown Objects Relations

We extend the carla environment to provide fine-grained labels for the specific set of unknown objects that we placed. The fine-grained labels are used to analyze the relation among unknown objects used during training and testing. This provides better understanding of task difficulty and model scalability. We propose to use the prototype fea-

Scenario	Unknown Objects									
	Cart with Bags	white white	II.							
Parking	Shopping Trolley	Sector NUS	-00-	÷						
	Garbage Bin		2K							
	Traffic Warning	. 👝 🦇 🍕	7							
Construction	Construction Cone									
Barrier	Traffic Pole	mili								
	Barrel	knimeller Andelick			_	_	_	 		/
	Traffic Cone	Summer Street	-	1					1	
Training	Traffic Barrier						- 🔼 🧍			
-	Static (others)									
	Dynamic (others)									

Figure 3: Different scenarios in Carla Simulation and objects considered as unknown in our synthetic data.



Figure 4: Dendrogram among unknown objects used during training and testing phases.

tures from masked average pooling following equation 2, where X_i is the i^{th} input image, F_i is the corresponding feature maps from ResNet-50 [9] with pretrained ImageNet weights θ , and M_i is its segmentation mask. P_c denotes the prototype for class c, that is computed on N images within the dataset. A pair-wise distance d(i, j) is measured among classes followed by agglomerative clustering.

A dendrogram among the classes is visualised to understand the relation between the unknown objects in the training and testing phases as shown in Figure 4. Classes that are visually similar, especially in texture, are closer such as street barrier and traffic cone. From the figure we can see less correlation among some of the different unknown objects used during testing and training especially in the parking scenario. Later, in the experiments we evaluate per scenario to confirm the generalization ability of the model.

$$F_i = f_\theta(X_i) \tag{2a}$$

$$P_c = \sum_{i=1}^{N} \sum_{x,y} \mathbb{1}[M_i^{x,y} = c] F_i^{x,y}$$
(2b)

$$d(i,j) = \|P_i - P_j\|_2$$
(2c)

4.3. Baselines and Metrics

The method we use for the open-set segmentation is similar to [30] but focused only on semantic segmentation without incorporating instances. We learn a representative prototype μ , σ per class. Then we use the distance to classify every pixel based on the matching prototype in 3a. We use a baseline that relies on both appearance and geometry, as we found depth to be an important signal to identify unknown objects. A two-stream backbone model f_{θ} with ResNet-50 backbone [9] and a feature pyramid network [13], which takes as input appearance and depth. This is followed by a semantic segmentation head f_{ϕ} with 4 convolutional modules with ReLU and group normalization, that learns prototypes μ , σ per class.

Let the input appearance and depth be denoted as x, the extracted features $h = f_{\theta}(x)$, and the embeddings from the segmentation head $m = f_{\phi}(h)$. The semantic segmentation head predicts the class probabilities following equation 3. Where the distance $d_{i,k}$ denotes the distance of pixel *i* features to the representative prototype of class k. A global learnable constant γ is used to estimate the unknown objects regardless of the objects' semantics similar to [30], then a softmax over C + 1 distances is used to estimate the probability of the pixel to belong to a certain class. For the contrastive learning baseline we learn a separate projection



Figure 5: Statistics for our Carla data for both known classes and unknown objects showing number of pixels per class.

head with auxiliary contrastive loss along with the segmentation loss as shown in equation 4.

$$d_{i,k} = \frac{-\|m_i - \mu_k\|^2}{2\sigma_k^2}$$
(3a)

$$d_{i,C+1} = \gamma \tag{3b}$$

$$\hat{y}_{i,k} = \frac{\exp(d_{i,k})}{\sum\limits_{i=1}^{C+1} \exp(d_{i,j})}$$
(3c)

$$l_{seg} = \frac{-1}{N} \sum_{i=1}^{N} \sum_{k=1}^{C+1} y_{i,k} \log \hat{y}_{i,k}$$
(3d)

$$L = L_{seg} + \lambda L_{cl} \tag{4}$$

As for the evaluation metric we report intersection over union in a class agnostic manner over the unknown objects in the test set. Where the predicted mask using the global learnable constant γ is used to label objects outside the closed set of known classes.

5. Experimental Results

5.1. Experimental Setup

Open-set Segmentation Setup: Throughout all experiments we use SGD with momentum optimizer with 0.005 learning rate and 0.9 momentum, and weight decay of 1×10^{-4} for 50 or 15 epochs for Cityscapes-VPS or Carla datasets respectively. A step learning rate scheduling which reduces the learning rate by 0.1 at epochs 30, 40 for Cityscapes-VPS, and 6, 8 for Carla is used. On

Cityscapes-VPS we resize images to 1024×512 then use random augmentations as random scales $\{0.8, 1.3\}$, random flipping and random cropping with 320×512 as crop sizes.

For datasets, we mainly use Cityscapes-VPS [11] as it has dense annotations for the video sequences and our collected dataset on Carla. In order to ensure no overfitting for the unknown objects segmentation occurs, we use different set of objects labelled as unknown during training and testing phases. In all experiments we include depth as another input modality to better segment unknown objects, in Carla we use the simulation's groundtruth and for Cityscapes-VPS we use the estimated depth from [8].

Motion Segmentation Setup: We use similar optimizer and learning rate to the open-set segmentation track, but train for 15 epochs on the combined KITTI-MOTS and Cityscapes-VPS motion datasets. Both datasets have their own training and testing splits, we train on the combined training splits then evaluate on the testing split from each. A step learning rate scheduling which reduces the learning rate with 0.1 at epochs 6, 8 is used. We report panoptic quality [12] and the class agnostic quality [30].

5.2. Motion Segmentation Track Results

Table 4 shows the results of our proposed multitask panoptic and class agnostic segmentation which we term as VCANet. The Pan-Base denotes our panoptic baseline without performing class agnostic segmentation. We initially report the class agnostic baseline trained on DAVIS [3] and compare it with the multi-task model VCANet to show that it maintains comparable segmentation quality. Our VCANet model with ego-flow suppression improves further the class agnostic quality (CAQ) metric on KITTIMOTS motion data. Our results act as a baseline

Table 4: Results of class agnostic and panoptic segmentation model. Tr-Te: Training and Testing motion segmentation datasets used. EFS: Ego Flow Suppression. D: DAVIS, K: KITTI our annotations, C: Cityscapes our annotations.

Model	Tr-Te	EFS	C	CA Metrics			Panoptic Metrics		
			SQ	RQ	CAQ	PQ_{All} PQ_{Th}		PQ_{St}	
Pan-Base	-	-	-	-	-	56.4	49.9	61.1	
CA-Base	D-D	X	76.6	60.6	46.4	-	-	-	
VCANet	D-D	X	72.2	53.2	38.6	56.4	49.9	61.1	
VCANet	KC-K	X	82.4	72.2	59.5	56.4	49.9	61.1	
VCANet	KC-K	1	82.4	75.4	62.1	56.4	49.9	61.1	
VCANet	KC-C	X	77.9	57.3	44.7	56.4	49.9	61.1	
VCANet	KC-C	1	78.0	56.2	43.8	56.4	49.9	61.1	



Figure 6: Predicted panoptic and CA segmentation on (a) KITTI. (b) Cityscapes. (c) IDD. Top: panoptic segmentation. Bottom: class agnostic segmentation.

for different approaches that tackle class agnostic instance segmentation. Figure 6 shows the qualitative analysis of VCANet on CityscapesVPS, KITTIMOTS and IDD. The model trained on our motion data is still able to segment moving objects in IDD sequences, which has an entirely different data distribution, and unknown objects such as animals that never appeared during training on Cityscapes-VPS known classes. Our multi-task model runs real-time at 5 fps on a 1080Ti GPU with 1024×512 image resolution.

5.3. Open-Set Segmentation Track Results

We conduct initial experiments on Cityscapes-VPS where four classes are withheld to be considered as unknown: Person, Rider, Motorcycle, Bicycle. During training, we label Person, Rider, and some of the ignored classes in Cityscapes as unknown objects for training the global unknown constant. Pixels belonging to Motorcycle and Bicycle are ignored in the cross entropy loss and do not contribute to the back-propagation. In the testing phase we evaluate on Motorcycle and Bicycle as the unknown objects and ignore the rest. Table 5 shows the results for two baseline models that follow the segmentation baseline we presented earlier. We found that using an auxiliary contrastive loss improves the results for the CA-IoU and mIoU metrics on Cityscapes-VPS. Figure 8 shows the results for segmenting both known classes and unknown objects, which demonstrates the ability to segment unknown bicycle and motorcycle classes.

In the same Table 5 we compare the baseline model with and without contrastive learning and train on our collected synthetic dataset. The CA-IoU is reported only on the unknown objects used during testing, which are labelled in Figure 4. We find that on the large-scale synthetic Carla dataset the baseline with contrastive learning does not improve, which could be explained because Carla dataset size is $25 \times$ of Cityscapes-VPS, and has more pixels labelled as unknown. We conduct experiments on a reduced version of the Carla dataset with only 2400 images as Cityscapes-VPS and less number of pixels labelled as unknown. This becomes more challenging to segment unknown objects and we see a clear gain from the auxiliary contrastive loss in low-data settings. This leaves an open question of how to devise ways to allow contrastive learning to improve performance even with abundant data. Figure 7 shows the CA-IoU over different scenarios and their qualitative results. It



Figure 7: CA-IoU reported per scenario. Predicted semantic and class agnostic segmentation on Carla Scenarios (a) Construction. (b) Barrier. (c) Parking. Top: semantic segmentation. Bottom: class agnostic segmentation



Figure 8: Predicted semantic and class agnostic segmentation on Cityscapes-VPS. Top: semantic segmentation. Bottom: class agnostic segmentation (Note: pedestrian - rider - bicycle - motorcycle are left out from training for known classes).

shows the model is able to segment some of the unknown objects that did not appear during training such as in the Parking Scenario. It is worth noting that some of the objects (Garbage Bin and Traffic Warning) have less correlation with unknown objects used during the training phase as shown in Figure 4. This confirms the model can scale to unknown objects not previously seen during training.

6. Conclusion

In this paper we formalized the video class agnostic segmentation task and provided necessary datasets and benchmarks for the two main tracks: (1) motion segmentation track, and (2) open-set segmentation track. In the motion segmentation track, an improved dataset for motion instance segmentation is provided and multiple motion segmentation baselines are benchmarked and publicly released. For the

Table 5: Open-Set Segmentation Results on Cityscapes-VPS and Carla. CL: contrastive learning. CA-IoU: class agnostic IoU on unknown objects.

	Dataset	mIoU	CA-IoU
No CL	Cityscapes-VPS	57.1	19.1
CL	Cityscapes-VPS	62.8	20.9
No CL	Carla	45.7	41.9
CL	Carla	44.2	37.2
No CL	Carla Reduced	38.5	6.5
CL	Carla Reduced	41.7	16.0

open-set segmentation track our synthetic dataset provides the means to assess model scalability to unknown objects and further study the relation among unknown objects used between training and testing.

References

- [1] Thomas Brox and Jitendra Malik. Object segmentation by long term analysis of point trajectories. In *European conference on computer vision*, pages 282–295. Springer, 2010.
- [2] Maxime Bucher, VU Tuan-Hung, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. In Advances in Neural Information Processing Systems, pages 468–479, 2019.
- [3] Sergi Caelles, Jordi Pont-Tuset, Federico Perazzi, Alberto Montes, Kevis-Kokitsi Maninis, and Luc Van Gool. The 2019 davis challenge on vos: Unsupervised multi-object segmentation. arXiv preprint arXiv:1905.00737, 2019.
- [4] Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE international conference on computer vision*, pages 2722– 2730, 2015.
- [5] 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. arXiv preprint arXiv:1604.01685, 2016.
- [6] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. arXiv preprint arXiv:1711.03938, 2017.
- [7] Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE trans*actions on pattern analysis and machine intelligence, 2020.
- [8] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3828–3838, 2019.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017.
- [11] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Video panoptic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9859–9868, 2020.
- [12] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceed*ings of the IEEE conference on computer vision and pattern recognition, pages 9404–9413, 2019.
- [13] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [14] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In Pro-

ceedings of the IEEE international conference on computer vision, pages 2980–2988, 2017.

- [15] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In 2016 fourth international conference on 3D vision (3DV), pages 565–571. IEEE, 2016.
- [16] Eslam Mohamed, Mahmoud Ewaisha, Mennatullah Siam, Hazem Rashed, Senthil Yogamani, and Ahmad El-Sallab. Instancemotseg: Real-time instance motion segmentation for autonomous driving. arXiv preprint arXiv:2008.07008, 2020.
- [17] Aljoša Ošep, Paul Voigtlaender, Mark Weber, Jonathon Luiten, and Bastian Leibe. 4d generic video object proposals. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 10031–10037. IEEE, 2020.
- [18] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Computer Vision and Pattern Recognition*, 2016.
- [19] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alexander Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. arXiv:1704.00675, 2017.
- [20] Hazem Rashed, Mohamed Ramzy, Victor Vaquero, Ahmad El Sallab, Ganesh Sistu, and Senthil Yogamani. Fusemodnet: Real-time camera and lidar based moving object detection for robust low-light autonomous driving. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- [21] Mennatullah Siam, Heba Mahgoub, Mohamed Zahran, Senthil Yogamani, Martin Jagersand, and Ahmad El-Sallab. Modnet: Moving object detection network with motion and appearance for autonomous driving. arXiv preprint arXiv:1709.04821, 2017.
- [22] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In Advances in neural information processing systems, pages 4077–4087, 2017.
- [23] David Tsai, Matthew Flagg, and James Rehg. Motion coherent tracking with multi-label mrf optimization. In *Proceedings of the British Machine Vision Conference*, pages 56.1– 56.11. BMVA Press, 2010. doi:10.5244/C.24.56.
- [24] Girish Varma, Anbumani Subramanian, Anoop Namboodiri, Manmohan Chandraker, and CV Jawahar. Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. In 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1743–1751. IEEE, 2019.
- [25] Johan Vertens, Abhinav Valada, and Wolfram Burgard. Smsnet: Semantic motion segmentation using deep convolutional neural networks. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, Vancouver, Canada, 2017.
- [26] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. Mots: Multi-object tracking and segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

- [27] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9197–9206, 2019.
- [28] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. Solo: Segmenting objects by locations. arXiv preprint arXiv:1912.04488, 2019.
- [29] Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, et al. Contrastive training for improved outof-distribution detection. arXiv preprint arXiv:2007.05566, 2020.
- [30] Kelvin Wong, Shenlong Wang, Mengye Ren, Ming Liang, and Raquel Urtasun. Identifying unknown instances for autonomous driving. In *Conference on Robot Learning*, pages 384–393. PMLR, 2020.