

Improving Online Performance Prediction for Semantic Segmentation

Marvin Klingner Andreas Bär Marcel Mross Tim Fingscheidt

{m.klingner, andreas.baer, m.mross, t.fingscheidt}@tu-bs.de

Technische Universität Braunschweig, Braunschweig, Germany

Abstract

In this work we address the task of observing the performance of a semantic segmentation deep neural network (DNN) during online operation, i.e., during inference, which is of high importance in safety-critical applications such as autonomous driving. Here, many high-level decisions rely on such DNNs, which are usually evaluated offline, while their performance in online operation remains unknown. To solve this problem, we propose an improved online performance prediction scheme, building on a recently proposed concept of predicting the primary semantic segmentation task's performance. This can be achieved by evaluating the auxiliary task of monocular depth estimation with a measurement supplied by a LiDAR sensor and a subsequent regression to the semantic segmentation performance. In particular, we propose (i) sequential training methods for both tasks in a multi-task training setup, (ii) to share the encoder as well as parts of the decoder between both task's networks for improved efficiency, and (iii) a temporal statistics aggregation method, which significantly reduces the performance prediction error at the cost of a small algorithmic latency. Evaluation on the KITTI dataset shows that all three aspects improve the performance prediction compared to previous approaches.

1. Introduction

The task of semantic segmentation is an essential component for camera-based environment perception applied in safety-critical applications such as autonomous driving, remote sensing, or virtual reality. Here, high-level decision systems require reliable information from the components of the environment perception. While deep neural networks (DNNs) nowadays achieve state-of-the-art performance on the semantic segmentation task [9, 57, 68], due to the lack of ground truth labels their performance during online deployment cannot be measured in a straightforward fashion. Moreover, the performance can drop significantly due to environment changes or domain shifts such as

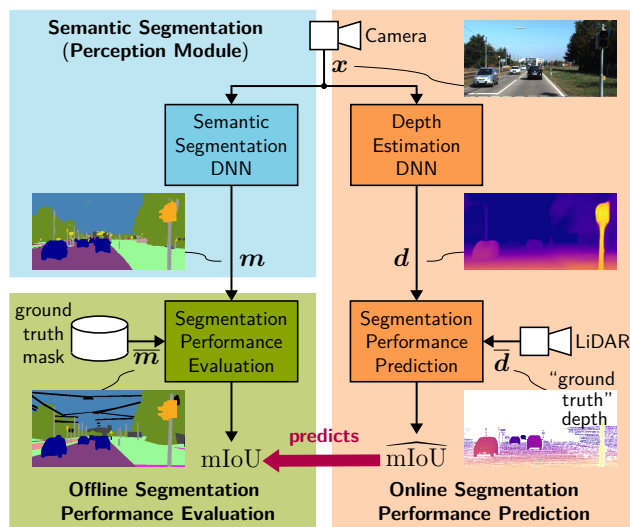


Figure 1. **Online performance prediction (mIoU)** concept for the semantic segmentation task **during online inference**. The typical offline evaluation requiring ground truth data is replaced by the auxiliary depth estimation task, which can be evaluated online using a LiDAR measurement.

lighting changes, camera noise, or even adversarial attacks [4, 22, 25, 39, 66]. Predicting the actual performance of semantic segmentation DNNs is therefore of highest interest.

Semantic segmentation DNNs are usually evaluated offline on well-prepared datasets [12, 43]. The observed offline performance (measured, e.g., by the mIoU a metric, cf. Fig. 1) is assumed to be maintained in online operation, as a direct evaluation is not possible without labels. Some initial works intend to mitigate the problem by temporal or network ensemble consistency checks [17, 51, 55], however, these approaches are not straightforward extendable to arbitrary tasks. Another approach uses an autoencoder to predict the performance based on the image reconstruction error [38]. This has, however, not been shown to be applicable on a few-shot basis, which would be necessary for online operation. Conclusively, there is still a desire for methods predicting DNN performance in online operation.

Addressing this task of online performance prediction, we build on the work of Klingner *et al.* [30], proposing a method, which relies on the auxiliary task of depth estimation also predicted by a DNN. As shown in Fig. 1, the semantic segmentation DNN (blue parts) is usually evaluated offline (green parts) by the mean intersection over union (mIoU) metric. In contrast, for online performance prediction (orange parts), the same input image is also given to a depth estimation DNN. The resulting depth estimate is evaluated afterwards. Finally, the depth performance can be connected to the segmentation performance by a simple polynomial regression, giving a prediction for the actual segmentation performance.

Our contribution with this work is threefold. Firstly, we introduce a *sequential* training protocol for both tasks in a multi-task network, improving the absolute performance as well as the performance prediction. Secondly, we analyze how many layers should be shared between both tasks to increase their correlation and thereby the performance prediction while still retaining a high performance of the semantic segmentation. Sharing more layers in the decoder also reduces the *computational complexity* induced by the performance prediction during inference, which can be of interest for computationally restricted settings. Thirdly, we investigate how a *temporal aggregation of statistics* can improve the prediction performance at the cost of a small decision latency but without affecting the absolute segmentation performance. Our analysis on the KITTI dataset [18] will show that all three methods significantly reduce the performance prediction error.

2. Related Work

In this section we give an overview on semantic segmentation techniques and how they profit from multi-task learning setups with depth estimation. Finally, we discuss other techniques for online performance prediction of DNNs.

Semantic Segmentation: The task of semantic segmentation has been facilitated by the increasing availability of public datasets [12, 43, 63] and the use of fully convolutional networks as outlined in the initial work of Long *et al.* [37]. Subsequent works adapted this concept and introduced improvements such as dilated convolutions [7, 62], atrous spatial pyramid pooling [9], state-of-the-art feature extractors [8, 45], label relaxation [68], or high resolution networks [57]. Other approaches explore unsupervised domain adaptation techniques [32, 27, 58, 61] to mitigate the domain shift problem for semantic segmentation and thereby facilitate the use of readily available synthetic datasets [50, 53]. Moreover, computational complexity can be reduced tremendously with only marginal loss in performance due to the use of efficient architectures. Examples are factorized convolutions [40, 52] or depthwise-separable convolutions [11], or a more efficient architectural design in

general [34, 45]. Another current research stream focuses on robustness of semantic segmentation DNNs to domain shifts or adversarial attacks [1, 2, 6, 24, 36]. In this work we do not approach semantic segmentation from a pure performance point of view but introduce a concept for an efficient online performance prediction.

Multi-Task Learning: The joint learning of depth estimation and semantic segmentation, either by shared network structures [13, 26] or by combined loss functions [26, 59, 64], has been shown to be beneficial for both tasks. Here, initial multi-task learning settings made use of supervised depth estimation. However, recent focus of depth estimation research shifted to self-supervised depth estimation techniques [65, 19, 20] due to the wide applicability of such methods not relying on ground truth depth labels. Conclusively, such techniques are also successfully combined with semantic segmentation [10, 23, 44]. On the one hand, the segmentation can be used to reduce the influence of dynamic objects on the depth estimation model assuming a static world [31, 49, 48] and sharpen depth edges at object boundaries [10, 67]. On the other hand, the depth estimation gives additional structural information to the semantic segmentation model [10, 44] and this multi-task setup can even be beneficial for unsupervised domain adaptation [31, 33, 56]. Finally, cross-task consistency checks [60, 41] and pixel-adaptive convolutions can be used to improve the performance of both tasks [23, 49].

While most works focus on performance gains of multi-task learning techniques, a recent work [29] also showed a positive effect on the robustness of a semantic segmentation DNN. Similarly, in this work we combine semantic segmentation with self-supervised depth estimation and show how this setting can be used to predict the performance of a semantic segmentation DNN in online operation.

Performance Prediction of DNNs: The usual way to evaluate the performance of a semantic segmentation DNN offline is to use a dataset containing ground truth labels [15, 12, 35], which is not possible in an online setting. To mitigate this problem one can model uncertainty in the output, *e.g.*, by Bayesian DNN approaches [17, 3, 51] which, however, involves the training and inference of an ensemble of networks and is therefore computationally expensive. Alternatively, temporal consistency of consecutive predictions can be checked [55], though the definition of consistency is highly task-specific. A recent successful approach of Löhdefink *et al.* [38] estimates the performance of a semantic segmentation model based on the image reconstruction error of an autoencoder DNN. Similarly, Klingner *et al.* [30] use the depth estimation task for this purpose. Both approaches, however, still have a high error in their segmentation performance predictions, which we significantly reduce in this work. In contrast to [38, 30] we therefore investigate, which training technique and multi-task network

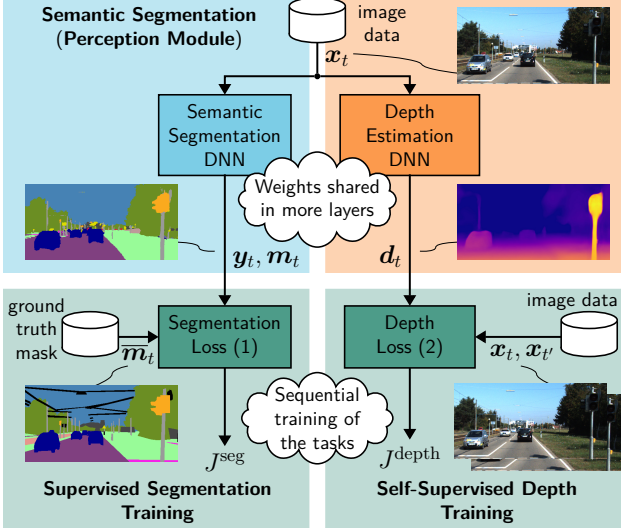


Figure 2. **Training setup** of the **online performance prediction** for the semantic segmentation task. Segmentation and depth tasks are initially trained by two separate loss functions. The final performance prediction can, however, be improved by sharing more network layers and—if a multi-task network is used—by training both tasks in a sequential fashion.

architecture is best suited for performance prediction. Also, we show how and under which conditions it is useful to employ a temporal aggregation of performance predictions.

3. Method Description

Here, we first describe our multi-task training procedure for semantic segmentation and depth estimation. Afterwards, we describe our online performance prediction, which can be applied after both tasks have been trained.

3.1. DNN Training

As shown in Fig. 2, we train both tasks simultaneously using task-specific loss functions. To find the optimal influence of each task, we do not weigh the loss functions, but the scaled gradients are flowing into the shared network parts (cf. the multi-task network in Fig. 3) as in [30, 31].

Semantic Segmentation: The semantic segmentation relies on a single input image $\mathbf{x}_t \in \mathbb{I}^{H \times W \times C}$, where H defines the height, W the width, and $C = 3$ the number of color channels. The range $\mathbb{I} = [0, 1]$ defines the normalized range of gray values. At inference, the neural network is used to predict a pixel-wise semantic segmentation map $\mathbf{m}_t \in \mathcal{S}^{H \times W}$ from the single image \mathbf{x}_t . Here, \mathcal{S} is the set of considered semantic classes. The DNN is trained by supervision from ground truth labels $\bar{\mathbf{m}}_t = (\bar{m}_{t,i}) \in \mathcal{S}^{H \times W}$, which are utilized in a one-hot encoded fashion as $\bar{\mathbf{y}}_t = (\bar{y}_{t,i,s}) \in \{0, 1\}^{H \times W \times |\mathcal{S}|}$, such that $\bar{m}_{t,i} = \arg\max_{s \in \mathcal{S}} \bar{y}_{t,i,s}$ at each pixel index

$i \in \mathcal{I} = \{1, \dots, H \cdot W\}$ in the image. During training, the network then predicts output probabilities $\mathbf{y}_t = (y_{t,i,s}) \in \mathbb{I}^{H \times W \times |\mathcal{S}|}$ with $|\mathcal{S}|$ semantic classes, which are optimized utilizing the cross-entropy loss

$$J^{\text{seg}} = -\frac{1}{H \cdot W} \sum_{i \in \mathcal{I}} \sum_{s \in \mathcal{S}} w_s \bar{y}_{t,i,s} \cdot \log(y_{t,i,s}), \quad (1)$$

where w_s are the class weights as defined in [46]. The final pixel-wise classes are then obtained by $m_{t,i} = \arg\max_{s \in \mathcal{S}} y_{t,i,s}$, yielding the estimated segmentation mask $\mathbf{m}_t = (m_{t,i}) \in \mathcal{S}^{H \times W}$.

Monocular Depth Estimation: The task of monocular depth estimation is to assign a depth value $d_{t,i} \in \mathbb{D}$ to each input image pixel $\mathbf{x}_{t,i} \in \mathbb{I}^C$. Thereby, a depth map $\mathbf{d}_t \in \mathbb{D}^{H \times W}$ is predicted from a monocular input image \mathbf{x}_t . Here, $\mathbb{D} = [d_{\min}, d_{\max}]$ defines the range of possible depth values, constrained by d_{\max} and d_{\min} , representing an upper bound and a lower bound, respectively.

To ensure that the depth estimation is applicable to a wide range of datasets, we build on recent works [65, 20], who train the depth estimation in a self-supervised fashion on sequential video images. To this end we consider consecutive pairs of images $(\mathbf{x}_t, \mathbf{x}_{t'})$ with $t' \in \mathcal{T}' = \{t-1, t+1\}$, which are given to a second network, predicting the six degrees of freedom of the relative pose between both images $\mathbf{T}_{t \rightarrow t'} \in SE(3)$. Here, $SE(3)$ represents the special Euclidean group of all possible rotations and translations [54] between the two images. The depth estimate \mathbf{d}_t predicted from \mathbf{x}_t together with the relative poses $\mathbf{T}_{t \rightarrow t'}$ can be used to project the image pixels $\mathbf{x}_{t'}$ to the pixel coordinates of the frame \mathbf{x}_t . The loss function is then based on the difference between the projected images $\mathbf{x}_{t' \rightarrow t}$ and the actual image \mathbf{x}_t . As previous approaches [20, 30], we use a mixture of absolute difference and structural similarity (SSIM) difference weighted by a factor $\alpha = 0.85$. The absolute difference $\|\cdot\|_1$ is computed over all color channels, while the SSIM difference $SSIM_i(\cdot) \in \mathbb{I}$ is additionally computed on 3×3 patches around the pixel index i of the images:

$$J_t^{\text{depth}} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \min_{t' \in \mathcal{T}'} \left(\frac{\alpha}{2} (1 - SSIM_i(\mathbf{x}_t, \mathbf{x}_{t' \rightarrow t})) + (1 - \alpha) \frac{1}{C} \|\mathbf{x}_{t,i} - \mathbf{x}_{t' \rightarrow t,i}\|_1 \right), \quad (2)$$

where we employ the minimum reprojection loss from [20] to handle occlusions between consecutive frames. Note, that the depth estimate \mathbf{d}_t is optimized implicitly as part of the projection model optimized through (2).

Additionally, we follow recent works in imposing a smoothness loss [19], allowing sharp depth boundaries only at pixel positions with strong color changes inside the input image. This additional loss is weighted by a factor of 10^{-3} as in [5, 19, 30]. Also, we apply the auto-masking technique from [20] to minimize the influence of

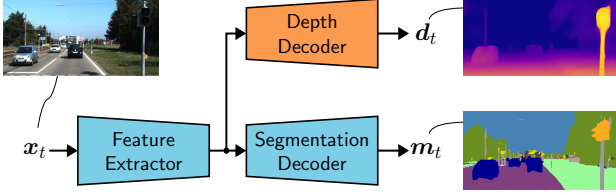


Figure 3. **Multi-task network architecture** used for the joint estimation of depth and segmentation outputs.

non-moving dynamic objects or frames. For more details on self-supervised depth estimation, the interested reader is referred to [20, 65], while [30, 31] contains more details on the multi-task training with semantic segmentation.

Improvement 1: Sequential Training Protocol: When training a multi-task network, several strategies can be applied: First, both decoders and the encoder (cf. Fig. 3) can be trained at once as in [30], which we dub MT. Alternatively, first one task’s decoder is trained and afterwards the second task’s decoder is trained, while the encoder weights are kept fixed during the second training stage. If the semantic segmentation network is first and afterwards the depth decoder, we dub this S,D. If we train the tasks in the opposite order, we dub this D,S. In this work we propose the sequential training protocol S,D,MT yielding an improved prediction performance as well as an improved absolute performance compared to the common MT training protocol.

Improvement 2: Sharing More Network Layers: As shown in Fig. 3, some layers of the semantic segmentation DNN and the depth estimation DNN can be shared to improve the correlation between both task’s outputs. To improve absolute performance, the common approach is to share the encoder and to use task-specific decoders [30]. However, for online performance prediction it can be beneficial to even share some decoder layers, if their architecture is identical in the first layers. Thereby, correlation and prediction performance can be improved even further.

3.2. Online Performance Prediction

To complete the performance prediction framework, the regression between the depth metric and the segmentation metric is calibrated in an offline regression setup [30]. In this work, we then propose to temporally aggregate the statistics from the predicted segmentation performance during online inference to improve prediction performance without compromising segmentation performance.

Network Input Perturbations: After the network has been trained, we simulate environment changes degrading network performance during inference by using various noise and adversarial attack types as shown in Fig. 4. These are added as perturbation $\mathbf{r}_\epsilon \in [-1, 1]^{H \times W \times C}$ onto the input image \mathbf{x} as

$$\mathbf{x}_\epsilon = \mathbf{x} + \mathbf{r}_\epsilon, \quad (3)$$

yielding the perturbed input image \mathbf{x}_ϵ . We then pass the perturbed image through the multi-task network estimating depth \mathbf{d}_ϵ and segmentation \mathbf{m}_ϵ , whose quality decreases dependent on the perturbation strength ϵ . We ensure $\mathbf{x}_\epsilon \in \mathbb{I}^{H \times W \times C}$ through clipping to the correct range. Moreover, the perturbation strength ϵ is defined by the signal-to-noise ratio as in [29] to make different perturbation types comparable. During evaluation, we apply several $\epsilon \in \mathcal{E}$, with the set of applied perturbation strengths \mathcal{E} .

Segmentation Evaluation: To implement a regression between the performance of the primary semantic segmentation task and the auxiliary depth estimation task, both tasks need to be evaluated. For semantic segmentation we employ the mean intersection-over-union (mIoU) metric [15], calculated from the number of true positives ($\text{TP}_{\epsilon,s}$), false negatives ($\text{FN}_{\epsilon,s}$), and false positives ($\text{FP}_{\epsilon,s}$) between prediction \mathbf{m}_ϵ and ground truth $\bar{\mathbf{m}}$ for each class s as

$$\text{mIoU}_\epsilon = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \frac{\text{TP}_{\epsilon,s}}{\text{TP}_{\epsilon,s} + \text{FP}_{\epsilon,s} + \text{FN}_{\epsilon,s}}. \quad (4)$$

The clean performance $\text{mIoU} = \text{mIoU}_0$ is given for a perturbation strength $\epsilon = 0$, while the regression is calculated from different $\epsilon \geq 0$. Note that $\text{TP}_{\epsilon,s}$, $\text{FN}_{\epsilon,s}$ and $\text{FP}_{\epsilon,s}$ are usually determined on an entire test set. As in this work performance prediction should, however, be applied on a single-image/few-shot basis, we calculate (4) on a single-image basis and report a test set average then.

Depth Evaluation: The depth estimation task is usually evaluated using seven different metrics (cf. [14, 65, 19, 5]). For comparability with previous works [30], we pick the accuracy metric ACC (in literature often referred to as “ $\delta < 1.25$ ”) for our performance prediction framework. It is calculated on a single-image basis by

$$\text{ACC}_\epsilon = \frac{1}{|\bar{\mathcal{I}}|} \sum_{i \in \bar{\mathcal{I}}} \left[\max \left(\frac{d_{\epsilon,i}}{\bar{d}_i}, \frac{\bar{d}_i}{d_{\epsilon,i}} \right) < 1.25 \right]_{\text{Iv}}, \quad (5)$$

where $d_{\epsilon,i}$ and \bar{d}_i are the pixel-wise elements of the predicted depth \mathbf{d}_ϵ and the ground truth depth $\bar{\mathbf{d}}$, respectively, $\bar{\mathcal{I}} \subset \mathcal{I}$ is the subset of all pixels for which the ground truth depth is available, and $|\bar{\mathcal{I}}|$ is the number of labeled pixels. The “ground truth” depth $\bar{\mathbf{d}}$ can be thought of as a LiDAR measurement, which is available during inference. Moreover, $[\cdot]_{\text{Iv}}$ represents the Iverson bracket being 1 or 0 if the condition inside the bracket is true or false, respectively. Also note that we determine the median of all image-wise scale factors on the validation set and apply this scale factor on the test set to make up for the global scale ambiguity in self-supervised monocular depth estimation.

Offline Regression Setup: The aim of the regression is to obtain a predicted performance $\widehat{\text{mIoU}}_\epsilon$ from the depth performance ACC_ϵ , which can be used to replace the actual

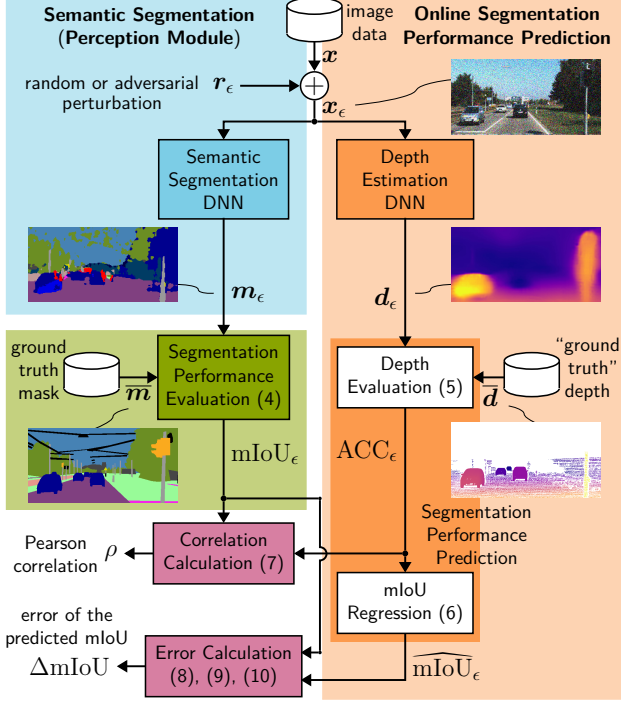


Figure 4. Offline **test setup** of the **online performance prediction** for the semantic segmentation task. The input image is artificially perturbed, leading to a performance degradation on both segmentation and depth estimation tasks. This can be used to obtain a regression between the segmentation performance and the depth performance. The depth estimation, the depth evaluation (with LiDAR “ground truth”), and the regression then make up the online performance prediction during inference, whose quality is reported by correlation and error metrics.

performance $mIoU_\epsilon$ during inference. Assuming a non-linear dependency between both metrics, we apply a polynomial regression of second order

$$\widehat{mIoU}_\epsilon = \sum_{k \in \mathcal{K}} \theta_k \cdot ACC_\epsilon^k, \quad (6)$$

where the parameters θ_k with $\mathcal{K} = \{0, 1, 2\}$ have to be optimized on the validation set. To this end, we consider the validation set $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n, \dots, \mathbf{x}_{|\mathcal{N}|}\}$ with $|\mathcal{N}|$ images from the set of images \mathcal{N} , where the actual performance $mIoU_{n,\epsilon}$ and the predicted performance $\widehat{mIoU}_{n,\epsilon}$ are calculated for each image \mathbf{x}_n . Accordingly, the regression parameters can be calibrated by minimizing the mean squared error $\sum_{n,\epsilon} (\widehat{mIoU}_{n,\epsilon} - mIoU_{n,\epsilon})^2$.

Online Inference Setup: During inference, the depth estimation, the depth evaluation (5) with LiDAR-based “ground truth” depth estimates \bar{d} , and the regression (6) can be used to predict the segmentation performance in online operation as shown in Fig. 4 (blue and orange parts only). In this work, we use a test set to simulate online operation and to evaluate our method (+ Fig. 4 green and red parts).

Thereby, we can also calculate the actual absolute performance by averaging the single-image results from (4) for $\epsilon = 0$ over the entire test set.

To evaluate the performance prediction, we use one correlation metric and two error metrics as defined in [30]. At first, the Pearson correlation metric is computed as

$$\rho = \frac{\sum_{n,\epsilon} (a_{n,\epsilon} - \mu_a)(b_{n,\epsilon} - \mu_b)}{\sqrt{\sum_{n,\epsilon} (a_{n,\epsilon} - \mu_a)^2} \sqrt{\sum_{n,\epsilon} (b_{n,\epsilon} - \mu_b)^2}}, \quad (7)$$

with abbreviations $a_{n,\epsilon} = mIoU_{n,\epsilon}$ and $b_{n,\epsilon} = ACC_{n,\epsilon}$, $\rho \in [-1, 1]$, and $\mu_a = \frac{1}{|\mathcal{N}| \cdot |\mathcal{E}|} \sum_{n,\epsilon} a_{n,\epsilon}$ (likewise for μ_b). A correlation coefficient of $\rho = 1$ indicates perfect positive correlation, while a correlation coefficient of $\rho = -1$ indicates perfect negative correlation. On the other hand, a correlation coefficient of $\rho = 0$ indicates no correlation at all between both entities.

We evaluate the performance prediction by using two error metrics between the predicted performance $\widehat{mIoU}_{n,\epsilon}$ and the actual performance $mIoU_{n,\epsilon}$ as shown in Fig. 4. To this end, we first calculate the image-wise differences

$$\Delta mIoU_{n,\epsilon} = \widehat{mIoU}_{n,\epsilon} - mIoU_{n,\epsilon}. \quad (8)$$

Afterwards, we calculate the mean absolute error

$$\Delta mIoU^M = \frac{1}{|\mathcal{N}| \cdot |\mathcal{E}|} \sum_{n \in \mathcal{N}} \sum_{\epsilon \in \mathcal{E}} |\Delta mIoU_{n,\epsilon}|, \quad (9)$$

as well as the root mean squared error

$$\Delta mIoU^R = \sqrt{\frac{1}{|\mathcal{N}| \cdot |\mathcal{E}|} \sum_{n \in \mathcal{N}} \sum_{\epsilon \in \mathcal{E}} (\Delta mIoU_{n,\epsilon})^2}. \quad (10)$$

The mean absolute error gives a very intuitive idea of the deviation between actual and predicted performance, while the root mean squared error is more sensitive to outliers.

Improvement 3: Aggregation of Statistics over Time:

While it is of course interesting to obtain performance predictions on a single-image basis, in this work we aim at improving performance by aggregation of statistics over a longer window of ΔN frames. To acquire these ΔN frames with indices $\nu \in \mathcal{N}_n$, every K -th frame is taken with frame n being in the center: $\mathcal{N}_n = \{n - \frac{\Delta N - 1}{2} \cdot K, \dots, n - K, n, n + K, n + K \cdot \frac{\Delta N - 1}{2}\}$. However, this comes at the cost of an algorithmic latency of $K \cdot \frac{\Delta N - 1}{2}$ frames, or $\bar{T} = \frac{\Delta N - 1}{2f} [s]$ with $f [\frac{1}{s}]$ being the frame rate of the deployed camera. This can be particularly tolerated in slowly changing environments such as, e.g., day/night shifts. In that case, one can use (predicted) quality averages over these time instants

$$mIoU_{n,\epsilon}^{(\Delta N)} = \frac{1}{\Delta N} \cdot \sum_{\nu \in \mathcal{N}_n} mIoU_{\nu,\epsilon} \quad (11)$$

$$\widehat{mIoU}_{n,\epsilon}^{(\Delta N)} = \frac{1}{\Delta N} \cdot \sum_{\nu \in \mathcal{N}_n} \widehat{mIoU}_{\nu,\epsilon} \quad (12)$$

for the actual and the predicted performance, respectively. Accordingly, the quantity $\widehat{\text{mIoU}}_{n,\epsilon}^{(\Delta N)}$ does represent the mean (not the single-image) predicted performance for ΔN frames $x_\nu, \nu \in \mathcal{N}_n$. As we intend to apply this method during offline testing on labelled image data (not necessarily videos), we randomly sample $\frac{\Delta N-1}{2}$ “preceding” and $\frac{\Delta N-1}{2}$ “successive” images indexed by \mathcal{N}_n from the whole test set. However, as these images are uncorrelated, we need to show which time $\frac{K_{\min}}{f}$ between sampled video frames $n, n+K_{\min}, \dots$ is necessary for this assumption to hold. The error of the prediction is then given by the difference

$$\Delta \text{mIoU}_{n,\epsilon} = \widehat{\text{mIoU}}_{n,\epsilon}^{(\Delta N)} - \text{mIoU}_{n,\epsilon}^{(\Delta N)}, \quad (13)$$

and subsequent application of (9) and (10). By this technique, one can improve the performance prediction at the cost of a certain algorithmic latency \tilde{T} during inference.

4. Experimental Setup

Describing our experimental setup, we outline our employed datasets, network architecture, training procedure, and evaluation metrics. All models are implemented in PyTorch [47] and have been trained and evaluated on an NVIDIA Tesla V100 graphics card.

Employed Datasets: During our experiments we have three stages: the training stage, the regression stage and the testing stage, which all require different datasets. For the training stage, we use the KITTI dataset [18] with the KITTI split $\mathcal{X}_{\text{train}}^{\text{KIT}}$ defined by [19] to train the depth estimation part of our method. Note, that due to the requirement of a preceding and a succeeding frame for our self-supervised training scheme we can use only 28,937 of the originally defined 29,000 images. Meanwhile, we use the 2,975 images of the training split $\mathcal{X}_{\text{train}}^{\text{CS}}$ from the Cityscapes dataset [12] to train the segmentation in a supervised fashion.

To evaluate our online performance prediction technique we need two data subsets with available labels for both depth and segmentation. One subset is used during the regression stage and the other one during the test stage. To achieve this, we split the 200 images of the KITTI Stereo 2015 training set [42] having the necessary labels for both tasks into two subsets. We use 50 images ($\mathcal{X}_{\text{val}}^{\text{KIT}}$) to determine the regression parameters and 150 images ($\mathcal{X}_{\text{test}}^{\text{KIT}}$) to test and evaluate our method.

Network Architecture: Our semantic segmentation DNN architecture is based on the encoder-decoder architecture with skip connections from [20]. We follow [31] in using the same architecture for the depth estimation DNN. The only difference is the output layer, where the depth decoder has a sigmoid output $\sigma_{n,i}$, where the final depth map is created by $\frac{1}{a\sigma_{n,i}+b}$ with a and b constraining the pixel-wise depth estimates in the range $[0.1, 100]$. The segmentation decoder on the other hand uses a softmax activation

function on the $|\mathcal{S}|$ feature maps such that the logits are converted to class probabilities. Using the same architecture for both tasks enables investigations on how many layers should be shared for online performance prediction.

Training Details: We train both tasks simultaneously using mixed batches with 6 images from both the depth and the segmentation dataset, i.e., 12 images in total per mini-batch. These are all passed through the multi-task network, while the losses are only calculated on the task-specific images. Additionally, for depth estimation, we use the full-resolution multi-scale loss from [20] and apply gradient scaling with factors of 0.9 and 0.1 for depth and segmentation, respectively, before the gradients reach the shared network parts. For optimization we apply the Adam optimizer [28] for 40 epochs with a learning rate of 10^{-4} , which is reduced to 10^{-5} after 30 epochs. For more training details the interested reader is referred to [30, 31].

Evaluation Details: During evaluation we apply four different perturbation types, namely Gaussian noise, salt and pepper noise [21], FGSM attacks [22], and PGD attacks [39]. As proposed by [30], we use a mix of them to calibrate the regression model and the mean over all four perturbation types to evaluate the performance prediction. We apply perturbation strengths of $\epsilon \in \mathcal{E} = \{0.25, 0.5, 1, 2, 4, 8, 12, 16, 20, 24, 28, 32\} \cdot \frac{1}{255}$ to cover the whole performance range of the segmentation model.

5. Experimental Evaluation

In this section, we will first evaluate and analyze the details of our regression used as part of the performance prediction framework. Afterwards, we outline how the performance prediction error can be reduced through a sequential training protocol, sharing of more layers, and (predicted) performance aggregation over a number of frames.

5.1. Regression Calibration

Analyzing our performance prediction method in Fig. 5, we show the dependency of the semantic segmentation metric $\text{mIoU}_{n,\epsilon}$ from the depth metric $\text{ACC}_{n,\epsilon}$, the value pairs being obtained on 20 images x_n perturbed with perturbations of strength $\epsilon \in \mathcal{E}$. For this experiment we use a model trained in a multi-task fashion as proposed by [30]. We can observe that the mix regression model (black curve), obtained on the validation set using results from all four perturbation types, is able to follow the trend of the data points to some extent. Surprisingly, we can also observe that the depth and segmentation metrics on clean images ($\epsilon = 0$, cf. green circles in Fig. 5) are not highly correlated as can be seen from their widely distributed positions in Fig. 5. Our conclusion is therefore that our method does indeed not exploit an intrinsic correlation on clean images between two task’s metrics. Our method rather relies on the correlation of the performance *degradation* in both tasks’ outputs. As

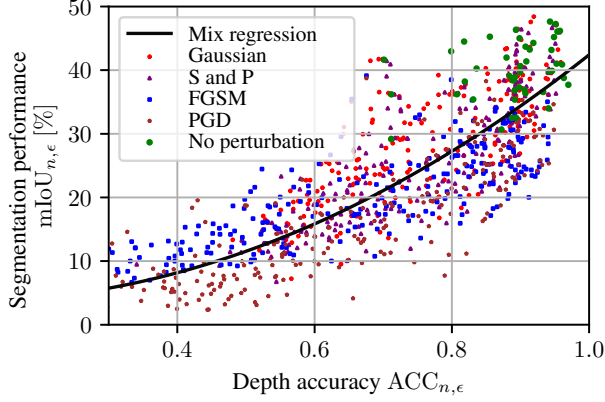


Figure 5. **Regression analysis:** We show a scatter plot of the image-wise calculated metrics $ACC_{n,\epsilon}$ (5) and $mIoU_{n,\epsilon}$ (4) on 20 images of the KITTI validation set \mathcal{X}_{val}^{KIT} . The results without any perturbation are marked as green circles. Also, we show the mix regression result (black curve), which is obtained using four different perturbation types and various perturbation strengths ϵ .

can be seen from the coherence in which the mix regression model captures the single-image results (cf. the regression in Fig. 5), this correlation is similar for different images and for different perturbation types and strengths.

5.2. Influence of Different Training Protocols

Now, we investigate which training protocol is best suited for our performance prediction method with results shown in Tab. 1. The advantage of only changing the training protocol is that no computational overhead is induced during inference and the network structure and complexity remain unchanged. When comparing the standard multi-task (MT) model to the S,D model, where first the segmentation network is trained and afterwards only the depth decoder, we observe a better prediction performance of the latter ($\Delta mIoU^M = 4.45\%$ compared to $\Delta mIoU^M = 6.01\%$) at the cost of a lower absolute performance ($mIoU = 34.97\%$ compared to $mIoU = 37.51\%$). By running S,D with subsequent MT in the S,D,MT training protocol, we can combine the strengths of both training schemes, i.e., we obtain a quite low performance prediction error $\Delta mIoU^M = 5.33\%$, while having an even higher absolute performance than the MT model with $mIoU = 39.43\%$. Please note that the best numbers of the depth performance ACC in Tabs. 1 and 2 are not written in boldface as we do not aim at optimizing this secondary task.

Still, the separate training of the depth decoder seems to be important, as the S,MT model has a similar absolute performance ($mIoU = 39.87\%$ compared to $mIoU = 39.43\%$) as the S,D,MT model, but a higher performance prediction error ($\Delta mIoU^M = 6.08\%$ compared to $\Delta mIoU^M = 5.33\%$). Also, training the depth network in a standalone fashion in the beginning (D,S model, D,MT model, and

Table 1. **Different training protocols:** Results on \mathcal{X}_{test}^{KIT} for multi-task training (MT), single-task depth training (D), and single-task segmentation training (S). If a second single-task training is applied (S after D or D after S), only the second decoder is trained. mIoU values in %; best results in **boldface**.

Training protocol	Performance		Prediction		
	ACC	mIoU	ρ	$\Delta mIoU^M$	$\Delta mIoU^R$
MT [30]	0.85	37.51	0.77	6.01	7.70
S,D [30]	0.86	34.97	0.86	4.45	6.16
S,MT	0.83	39.87	0.81	6.08	7.68
S,D,MT	0.86	39.43	0.84	5.33	7.01
D,S [30]	0.84	33.34	0.77	5.49	7.34
D,MT	0.82	39.81	0.78	6.14	8.00
D,S,MT	0.82	36.57	0.77	5.64	7.51

Table 2. **Sharing network parts:** Results on \mathcal{X}_{test}^{KIT} when different **network parts** are **shared** during training. We report results for completely separately trained networks, a shared encoder, as well as additionally shared decoder layers. All multi-task networks are trained according to the S,D,MT training protocol proposed in Tab. 1. mIoU values in %; best results in **boldface**.

Shared network parts	Performance		Prediction		
	ACC	mIoU	ρ	$\Delta mIoU^M$	$\Delta mIoU^R$
None [30]	0.84	34.97	0.67	7.55	10.10
+ Encoder	0.86	39.43	0.84	5.33	7.01
+ Dec. layer 1	0.86	40.79	0.82	5.87	7.71
+ Dec. layer 2	0.86	38.71	0.82	5.67	7.30
+ Dec. layer 3	0.86	35.80	0.83	4.96	6.48
+ Dec. layer 4	0.86	33.94	0.88	4.07	5.37

D,S,MT model in Tab. 1) offers no clear advantages over the MT model. Concluding, we show that *with the S,D,MT model we improve on the work of [30] (proposing the MT model) in terms of both performance prediction ($\Delta mIoU$ metrics) and absolute performance (mIoU) without computational overhead during inference.*

5.3. Influence of Sharing Network Parts

Another way to improve the performance prediction is to increase the correlation between both task’s outputs by sharing more network layers. The results in Tab. 2, where all models with shared network parts are trained according to the S,D,MT training protocol, clearly show that it is beneficial to share the encoder layers, as absolute performance improves from $mIoU = 34.97\%$ to $mIoU = 39.43\%$ and at the same time the prediction error reduces from $\Delta mIoU^M = 7.55\%$ to $\Delta mIoU^M = 5.33\%$. Sharing one additional decoder layer improves the absolute performance slightly but surprisingly not the performance prediction error. Only when sharing three or four decoder layers, we can reduce the performance prediction error significantly down to $\Delta mIoU^M = 4.07\%$, however, at the cost of a decreased absolute performance of $mIoU = 33.94\%$. This could be

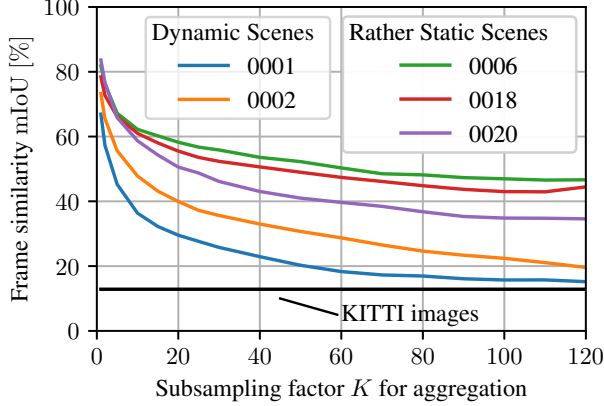


Figure 6. **Temporal correlation analysis:** Overview on the average mIoU between pairs of samples stemming from the KITTI dataset or from videos of the vKITTI dataset subsampled to every K -th frame. We observe that subsampling dynamic scenes by a factor of K in the range $K = 100 \dots 200$ yields similar mIoU for an image pair as any KITTI image pair.

because the decoders cannot learn optimal task-specific features in the reduced remaining number of task-specific layers. In conclusion, *sharing only the encoder already offers a good trade-off between absolute performance and performance prediction quality*, which is why we will stick to this model for our further analysis. If one aims at a significantly improved performance prediction or a more efficient performance prediction with less computational overhead, sharing more decoder layers is also a good option.

5.4. Influence of Aggregation of Statistics

Here, we will evaluate our simple aggregation method described by (11), (12), and (13), where the mean performance over ΔN frames is predicted. As labeled real-world video sequences are hard to obtain for semantic segmentation, we calculate the mIoU between any pair of images of the KITTI dataset as shown by the black line in Fig. 6. As these frames are rather uncorrelated, we use the video sequences of the vKITTI dataset [16] to check how large we have to choose K so that in a real world image sequence the same dissimilarity is obtained as if we used a random number of $\Delta N - 1$ extra images from our KITTI data for aggregation (11), (12). We can see that on the sequences 0006, 0018, and 0020 the approximation does not hold as these videos either show highway scenes with few variations or static scenes of a barely moving vehicle. On the other hand, for the more dynamic videos 0001 and 0002, the mIoU approaches the one present in the KITTI 2015 dataset (where segmentation ground truth is available) for a subsampling factor of $K = 100 \dots 200$. For this distance the evaluation on the KITTI 2015 dataset should be approximating a real video (for dynamically changing scenes).

The monitored time interval is given by $\Delta T = \frac{\Delta N \cdot K}{f}$,

Table 3. **Aggregation of statistics:** Results on $\mathcal{L}_{\text{test}}^{\text{KIT}}$ when the actual performance mIoU and the predicted performance $\widehat{\text{mIoU}}$ are averaged over several frames before the calculation of the difference in (8). We report results for different numbers of ΔN frames used for aggregation (first column). mIoU values in %.

ΔN	Decision latency \tilde{T} [s]	Performance		Prediction		
		ACC	mIoU	ρ	$\Delta \text{mIoU}^{\text{M}}$	$\Delta \text{mIoU}^{\text{R}}$
1	0	0.86	39.43	0.84	5.33	7.01
3	10	0.86	39.43	0.92	3.88	4.91
5	20	0.86	39.43	0.94	3.30	4.17
11	50	0.86	39.43	0.95	3.02	3.77
21	100	0.86	39.43	0.96	2.81	3.51
51	250	0.86	39.43	0.96	2.74	3.42
101	500	0.86	39.43	0.97	2.68	3.34

and the decision latency in Tab. 3 is $\tilde{T} = \frac{\Delta N - 1}{2} \cdot \frac{K}{f}$. The frame rate of the KITTI and vKITTI datasets is constant with $f = 10 \text{ s}^{-1}$ and we assume $K = 100$ frames distance. We can see that the mIoU prediction error decreases significantly already when considering only $\Delta N = 3$ frames, as the effect of outliers in both the predicted and the actual performance is reduced. and even more for a higher number of frames. A good trade-off is given by using $\Delta N = 11$ frames, where we get a significantly reduced prediction error of $\Delta \text{mIoU}^{\text{M}} = 3.02\%$ compared to $\Delta \text{mIoU}^{\text{M}} = 5.33\%$, while the absolute performance does obviously not change. This comes with a decision latency of $\tilde{T} = 50 \text{ s}$, which can be tolerated for slowly changing perturbations such as day/night changes, while for more rapid environment changes the single-image prediction method without aggregation should be preferred.

6. Conclusion

In this work we introduce a concept for online performance prediction of semantic segmentation DNNs based on the monocular estimation of depth, its subsequent evaluation using, *e.g.*, LiDAR output, and finally a regression to predict the segmentation performance. In particular, we show how the prediction error can be reduced by sequential training of both tasks, sharing encoder and decoder layers inside a multi-task network, and temporal aggregation of statistics. Thereby, we can reduce the online mIoU prediction error from 6.01% down to 5.33% (*i.e.*, 11% relative improvement), while even improving the absolute mIoU performance from 37.51% to 39.43%. At the cost of some decision latency, the mIoU error can be further reduced to 3.02% absolute or even below. Our online performance prediction concept is straightforward extendable to other perception tasks such as object detection or optical flow and provides valuable information for high-level decision systems inside safety critical applications such as autonomous driving or virtual reality.

References

- [1] Andreas Bär, Marvin Klingner, Serin Varghese, Fabian Hüber, Peter Schlicht, and Tim Fingscheidt. Robust Semantic Segmentation by Redundant Networks With a Layer-Specific Loss Contribution and Majority Vote. In *Proc. of CVPR - Workshops*, pages 1348–1358, Seattle, WA, USA, June 2020. 2
- [2] Andreas Bär, Jonas Löhdefink, Nikhil Kapoor, Serin John Varghese, Fabian Hüber, Peter Schlicht, and Tim Fingscheidt. The Vulnerability of Semantic Segmentation Networks to Adversarial Attacks in Autonomous Driving: Enhancing Extensive Environment Sensing. *IEEE Signal Processing Magazine*, 38(1):42–52, Jan. 2021. 2
- [3] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight Uncertainty in Neural Network. In *Proc. of ICML*, pages 1613–1622, Lille, France, July 2015. 2
- [4] Nicholas Carlini and David A. Wagner. Towards Evaluating the Robustness of Neural Networks. In *Proc. of SP*, pages 39–57, San Jose, CA, USA, May 2017. 1
- [5] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Depth Prediction Without the Sensors: Leveraging Structure for Unsupervised Learning from Monocular Videos. In *Proc. of AAAI*, pages 8001–8008, Honolulu, HI, USA, Jan. 2019. 3, 4
- [6] Hao-Yun Chen, Jhao-Hong Liang, Shih-Chieh Chang, Jia-Yu Pan, Yu-Ting Chen, Wei Wei, and Da-Cheng Juan. Improving Adversarial Robustness via Guided Complement Entropy. In *Proc. of ICCV*, pages 4881–4889, Seoul, Korea, Oct. 2019. 2
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic Image Segmentation With Deep Convolutional Nets and Fully Connected CRFs. In *Proc. of ICLR*, pages 1–14, San Diego, CA, USA, May 2015. 2
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic Image Segmentation With Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(4):834–848, Apr. 2018. 2
- [9] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-Decoder With Atrous Separable Convolution for Semantic Image Segmentation. In *Proc. of ECCV*, pages 801–818, Munich, Germany, Sept. 2018. 1, 2
- [10] Po-Yi Chen, Alexander H. Liu, Yen-Cheng Liu, and Yu-Chiang F. Wang. Towards Scene Understanding: Unsupervised Monocular Depth Estimation With Semantic-Aware Representation. In *Proc. of CVPR*, pages 2624–2632, Long Beach, CA, USA, June 2019. 2
- [11] François Chollet. Xception: Deep Learning With Depthwise Separable Convolutions. In *Proc. of CVPR*, pages 1063–6919, Honolulu, HI, USA, July 2017. 2
- [12] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *Proc. of CVPR*, pages 3213–3223, Las Vegas, NV, USA, June 2016. 1, 2, 6
- [13] David Eigen and Rob Fergus. Predicting Depth, Surface Normals and Semantic Labels With a Common Multi-Scale Convolutional Architecture. In *Proc. of ICCV*, pages 2650–2658, Santiago, Chile, Dec. 2015. 2
- [14] David Eigen, Christian Puhrsch, and Rob Fergus. Depth Map Prediction from a Single Image Using a Multi-Scale Deep Network. In *Proc. of NIPS*, pages 2366–2374, Montréal, QC, Canada, Dec. 2014. 4
- [15] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision (IJCV)*, 111(1):98–136, Jan. 2015. 2, 4
- [16] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual Worlds as Proxy for Multi-Object Tracking Analysis. In *Proc. of CVPR*, pages 4340–4349, Las Vegas, NV, USA, June 2016. 8
- [17] Y. Gal and Z. Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *Proc. of ICML*, pages 1050–1059, New York, NY, USA, June 2016. 1, 2
- [18] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision Meets Robotics: The KITTI Dataset. *International Journal of Robotics Research (IJRR)*, 32(11):1231–1237, Aug. 2013. 2, 6
- [19] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised Monocular Depth Estimation With Left-Right Consistency. In *Proc. of CVPR*, pages 270–279, Honolulu, HI, USA, July 2017. 2, 3, 4, 6
- [20] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging Into Self-Supervised Monocular Depth Estimation. In *Proc. of ICCV*, pages 3828–3838, Seoul, Korea, Oct. 2019. 2, 3, 4, 6
- [21] Rafael C. Gonzales and Richard E. Woods. *Digital Image Processing*. Prentice Hall, 2008. 6
- [22] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. In *Proc. of ICLR*, pages 1–10, San Diego, CA, USA, May 2015. 1, 6
- [23] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-Guided Representation Learning for Self-Supervised Monocular Depth. In *Proc. of ICLR*, pages 1–14, Addis Ababa, Ethiopia, Apr. 2020. 2
- [24] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using Self-Supervised Learning Can Improve Model Robustness and Uncertainty. In *Proc. of NeurIPS*, pages 15637–15648, Vancouver, BC, Canada, Dec. 2019. 2
- [25] Khaled F. Hussain, Mahmoud Affi, and Ghada Moussa. A Comprehensive Study of the Effect of Spatial Resolution and Color of Digital Images on Vehicle Classification. *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 20(3):1181–1190, 2019. 1
- [26] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. In *Proc. of CVPR*, pages 7482–7491, Salt Lake City, UT, USA, June 2018. 2

- [27] Myeongjin Kim and Hyeran Byun. Learning Texture Invariant Representation for Domain Adaptation of Semantic Segmentation. In *Proc. of CVPR*, pages 12975–12984, Seattle, WA, USA, June 2020. 2
- [28] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *Proc. of ICLR*, pages 1–15, San Diego, CA, USA, May 2015. 6
- [29] Marvin Klingner, Andreas Bär, and Tim Fingscheidt. Improved Noise and Attack Robustness for Semantic Segmentation by Using Multi-Task Training with Self-Supervised Depth Estimation. In *Proc. of CVPR - Workshops*, pages 1299–1309, Seattle, WA, USA, June 2020. 2, 4
- [30] Marvin Klingner and Tim Fingscheidt. Online Performance Prediction of Perception DNNs by Multi-Task Learning with Depth Estimation. *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, pages 1–14, 2021. 2, 3, 4, 5, 6, 7
- [31] Marvin Klingner, Jan-Aike Termöhlen, Jonas Mikolajczyk, and Tim Fingscheidt. Self-Supervised Monocular Depth Estimation: Solving the Dynamic Object Problem by Semantic Guidance. In *Proc. of ECCV*, pages 582–600, Glasgow, UK, Aug. 2020. 2, 3, 4, 6
- [32] Marvin Klingner, Jan-Aike Termöhlen, Jacob Ritterbach, and Tim Fingscheidt. Unsupervised BatchNorm Adaptation (UBNA): A Domain Adaptation Method for Semantic Segmentation Without Using Source Domain Representations. *arXiv*, (2011.08502), Nov. 2020. 2
- [33] Kuan-Hui Lee, German Ros, Jie Li, and Adrien Gaidon. SPIGAN: Privileged Adversarial Learning from Simulation. In *Proc. of ICLR*, pages 1–14, New Orleans, LA, USA, Apr. 2019. 2
- [34] Hanchao Li, Pengfei Xiong, Haoqiang Fan, and Jian Sun. DFANet: Deep Feature Aggregation for Real-Time Semantic Segmentation. In *Proc. of CVPR*, pages 9522–9531, Long Beach, CA, USA, June 2019. 2
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In *Proc. of ECCV*, pages 740–755, Zurich, Switzerland, Sept. 2014. 2
- [36] Zihao Liu, Qi Liu, Tao Liu, Nuo Xu, Xue Lin, Yanzhi Wang, and Wujie Wen. Feature Distillation: DNN-Oriented JPEG Compression Against Adversarial Examples. In *Proc. of CVPR*, pages 860–868, Long Beach, CA, USA, June 2019. 2
- [37] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proc. of CVPR*, pages 3431–3440, Boston, MA, USA, June 2015. 2
- [38] Jonas Löhdefink, Justin Fehrling, Marvin Klingner, Fabian Hüger, Peter Schlicht, Nico M. Schmidt, and Tim Fingscheidt. Self-Supervised Domain Mismatch Estimation for Autonomous Perception. In *Proc. of CVPR - Workshops*, pages 1359–1368, Seattle, WA, USA, June 2020. 1, 2
- [39] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks. In *Proc. of ICLR*, pages 1–28, Vancouver, BC, Canada, Apr. 2018. 1, 6
- [40] Sachin Mehta, Mohammad Rastegari, Anat Caspi, Linda Shapiro, and Hannaneh Hajishirzi. ESPNet: Efficient Spatial Pyramid of Dilated Convolutions for Semantic Segmentation. In *Proc. of ECCV*, pages 552–568, Munich, Germany, Sept. 2018. 2
- [41] 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 *Proc. of CVPR*, pages 9810–9820, Long Beach, CA, USA, June 2019. 2
- [42] Moritz Menze and Andreas Geiger. Object Scene Flow for Autonomous Vehicles. In *Proc. of CVPR*, pages 3061–3070, Boston, MA, USA, June 2015. 6
- [43] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kotschieder. The Mapillary Vistas Dataset for Semantic Understanding of Street Scenes. In *Proc. of ICCV*, pages 4990–4999, Venice, Italy, Oct. 2017. 1, 2
- [44] Jelena Novosel, Prashanth Viswanath, and Bruno Arsenali. Boosting Semantic Segmentation With Multi-Task Self-Supervised Learning for Autonomous Driving Applications. In *Proc. of NeurIPS - Workshops*, pages 1–11, Vancouver, BC, Canada, Dec. 2019. 2
- [45] Marin Oršić, Ivan Krešo, Petra Bevandić, and Siniša Šegvić. In Defense of Pre-Trained ImageNet Architectures for Real-Time Semantic Segmentation of Road-Driving Images. In *Proc. of CVPR*, pages 12607–12616, Long Beach, CA, USA, June 2019. 2
- [46] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. ENet: A Deep Neural Network Architecture for Real-Time Semantic Segmentation. *arXiv*, (1606.02147), June 2016. 3
- [47] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Proc. of NeurIPS*, pages 8024–8035, Vancouver, BC, Canada, Dec. 2019. 6
- [48] Varun Ravi Kumar, Marvin Klingner, Senthil Yogamani, Markus Bach, Stefan Milz, Tim Fingscheidt, and Patrick Mäder. SVDistNet: Self-Supervised Near-Field Distance Estimation on Surround View Fisheye Cameras. *arXiv*, (2104.04420), Apr. 2021. 2
- [49] Varun Ravi Kumar, Marvin Klingner, Senthil Yogamani, Stefan Milz, Tim Fingscheidt, and Patrick Mäder. SynDistNet: Self-Supervised Monocular Fisheye Camera Distance Estimation Synergized with Semantic Segmentation for Autonomous Driving. In *Proc. of WACV*, pages 61–71, Waikoloa, HI, USA, Jan. 2021. 2
- [50] Stephan Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for Data: Ground Truth from Computer Games. In *Proc. of ECCV*, pages 102–118, Amsterdam, Netherlands, Oct. 2016. 2
- [51] Hippolyt Ritter, Aleksandar Botev, and David Barber. A Scalable Laplace Approximation for Neural Networks. In *Proc. of ICLR*, pages 1–10, Vancouver, BC, Canada, Apr. 2018. 1, 2
- [52] Eduardo Romera, José M. Álvarez, Luis M. Bergasa, and Roberto Arroyo. ERFNet: Efficient Residual Factorized

- ConvNet for Real-Time Semantic Segmentation. *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 19(1):263–272, Jan. 2018. [2](#)
- [53] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes. In *Proc. of CVPR*, pages 3234–3243, Las Vegas, NV, USA, June 2016. [2](#)
- [54] Richard Szeliski. *Computer Vision: Algorithms and Applications*. Springer Science & Business Media, 2010. [3](#)
- [55] Serin Varghese, Yasin Bayzidi, Andreas Bär, Nikhil Kapoor, Sounak Lahiri, Jan David Schneider, Nico Schmidt, Fabian Hüger, Peter Schlicht, and Tim Fingscheidt. Unsupervised Temporal Consistency Metric for Video Segmentation in Highly-Automated Driving. In *Proc. of CVPR - Workshops*, pages 1369–1378, Seattle, WA, USA, June 2020. [1](#), [2](#)
- [56] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Perez. ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. In *Proc. of CVPR*, pages 2517–2526, Long Beach, CA, USA, June 2019. [2](#)
- [57] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep High-Resolution Representation Learning for Visual Recognition. In *Proc. of CVPR*, pages 1–17, Long Beach, CA, USA, June 2019. [1](#), [2](#)
- [58] Zhonghao Wang, Mo Yu, Yunchao Wei, Rogerio Feris, Jinjun Xiong, Wen-Mei Hwu, Thomas S. Huang, and Honghui Shi. Differential Treatment for Stuff and Things: A Simple Unsupervised Domain Adaptation Method for Semantic Segmentation. In *Proc. of CVPR*, pages 12635–12644, Seattle, WA, USA, June 2020. [2](#)
- [59] 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 *Proc. of CVPR*, pages 675–684, Salt Lake City, UT, USA, June 2018. [2](#)
- [60] Guorun Yang, Hengshuang Zhao, Jianping Shi, Zhidong Deng, and Jiaya Jia. SegStereo: Exploiting Semantic Information for Disparity Estimation. In *Proc. of ECCV*, pages 636–651, Munich, Germany, Sept. 2018. [2](#)
- [61] Jinyu Yang, Weizhi An, Sheng Wang, Xinliang Zhu, Chaochao Yan, and Junzhou Huang. Label-Driven Reconstruction for Domain Adaptation in Semantic Segmentation. In *Proc. of ECCV*, pages 480–498, Glasgow, UK, Aug. 2020. [2](#)
- [62] Fisher Yu and Vladlen Koltun. Multi-Scale Context Aggregation by Dilated Convolutions. In *Proc. of ICLR*, pages 1–13, San Juan, Puerto Rico, May 2016. [2](#)
- [63] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. BDD100K: A Diverse Driving Video Database With Scalable Annotation Tooling. *arXiv*, (1805.04687), Aug. 2018. [2](#)
- [64] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Yan Yan, Nicu Sebe, and Jian Yang. Pattern-Affinitive Propagation Across Depth, Surface Normal and Semantic Segmentation. In *Proc. of CVPR*, pages 4106–4115, Long Beach, CA, USA, June 2019. [2](#)
- [65] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised Learning of Depth and Ego-Motion from Video. In *Proc. of CVPR*, pages 1851–1860, Honolulu, HI, USA, July 2017. [2](#), [3](#), [4](#)
- [66] Wei Zhou, Julie Stephany Berrio, Stewart Worrall, and Eduardo Nebot. Automated Evaluation of Semantic Segmentation Robustness for Autonomous Driving. *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 21(5):1951–1963, 2019. [1](#)
- [67] Shengjie Zhu, Garrick Brazil, and Xiaoming Liu. The Edge of Depth: Explicit Constraints between Segmentation and Depth. In *Proc. of CVPR*, pages 13116–13125, Seattle, WA, USA, June 2020. [2](#)
- [68] Yi Zhu, Karan Sapra, Fitsum A. Reda, Kevin J. Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving Semantic Segmentation via Video Propagation and Label Relaxation. In *Proc. of CVPR*, pages 8856–8865, Long Beach, CA, USA, June 2019. [1](#), [2](#)