

SS-SFDA : Self-Supervised Source-Free Domain Adaptation for Road Segmentation in Hazardous Environments

Divya Kothandaraman, Rohan Chandra, Dinesh Manocha
University of Maryland, College Park

Tech Report, Code, and Video at <https://gamma.umd.edu/weatherSAFE/>

Abstract

We present a novel approach for unsupervised road segmentation in adverse weather conditions such as rain or fog. This includes a new algorithm for source-free domain adaptation (SFDA) using self-supervised learning. Moreover, our approach uses several techniques to address various challenges in SFDA and improve performance, including online generation of pseudo-labels and self-attention as well as use of curriculum learning, entropy minimization and model distillation. We have evaluated the performance on 6 datasets corresponding to real and synthetic adverse weather conditions. Our method outperforms all prior works on unsupervised road segmentation and SFDA by at least 10.26%, and improves the training time by 18 – 180×. Moreover, our self-supervised algorithm exhibits similar accuracy performance in terms of mIOU score as compared to prior supervised methods.

1. Introduction

Research in autonomous driving continues to advance in terms of improving the perception capabilities of self-driving cars for greater safety. This includes detecting both static and dynamic obstacles such as pedestrians [16], tracking and predicting the trajectories of other vehicles [4, 7, 6, 5], and scene segmentation [18, 62]. Immense progress along these lines has led to the deployment of level 2 and almost level 3 autonomous vehicles (AVs) in urban traffic environments [13]. However, these advances in perception technology have been primarily designed to work well in safe and clear weather conditions. Driving in adverse weather and lighting conditions such as snow, rain, or fog is challenging not just for autonomous vehicles but even for humans. These conditions result in a degradation in accuracy of perception techniques including road segmentation [43, 42]. Consequently, AVs are unable to distinguish drivable regions of the road from the non-driveable region (which may be affected by snow, rain, or fog), thereby increasing the likelihood of road accidents [36]. In this paper, we address the problem of road segmentation in adverse weather conditions.

The road segmentation [17, 48] problem corresponds to identifying the pixels in an RGB image or video that be-

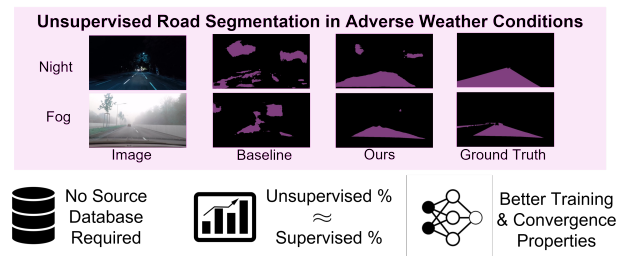


Figure 1: We highlight the results generated by SS-SFDA on night [42] and fog benchmarks [43], compared to the baseline source model pre-trained on clear weather CityScapes. The purple regions (right) denote the segmented road pixels. The overall accuracy of our self-supervised algorithm in terms of mIoU is (88 – 96%) of supervised methods.

long to the ‘road’ class. While general models designed for semantic segmentation in computer vision can be directly used for road segmentation, they suffer from the inability to capture semantic relationships between different objects due to the lack of unique labels for each class. The use of self-attention techniques [18, 60] can mitigate this issue by capturing long-range dependencies.

However, one major challenge in road segmentation in adverse weather is the lack of ground-truth annotations for road pixels. A common approach in deep learning for handling lack of training data is domain adaptation (DA) [22, 52]. However, DA-based methods assume access to source datasets (clear weather dataset in our context) at all times which can be prohibitive in terms of storage, memory, data corruption and privacy concerns. Recently, [26, 37] have proposed source-free domain adaptation (SFDA) in which deep neural networks (DNNs) do not require access to the source dataset during the adaptation stage; instead, DNNs are pre-trained on a source dataset (clear weather dataset) and the pre-trained model is directly used to adapt to the unlabeled target domain (adverse weather dataset).

Current methods for SFDA are used for image classification [28, 24, 29, 56] and may not work well for semantic segmentation due to the inherent differences between the classification and segmentation tasks. Moreover, many current SFDA methods use GANs to produce a “copy” of the original source domain distribution. In addition to being

computationally intensive due to the difficulty in training GANs, image generation for segmentation requires GANs to capture contextual information and semantic relationships between multiple objects and the background, which can be complicated in road scenes. As a result, prior SFDA techniques have not been used for road segmentation.

1.1. Main Contributions

We present a new approach for road segmentation in adverse weather conditions. Our approach is based on a novel algorithm for SFDA using self-supervised learning. We initialize our model with an auto-encoder baseline network using self-attention to generate a pre-trained model on the clear weather source dataset. Using self-attention improves the overall model by capturing long-range dependencies within the image (Section 3.1). Our novel contributions include:

1. We present a novel two-step self-supervised SFDA approach called SS-SFDA. In the first step, our method uses entropy minimization to enrich the noisy pseudo-labels generated by the pre-trained auto-encoder. In the second step, we use a novel self-training method that generates pseudo labels in an *online* manner, as opposed to iterative self-training use by prior methods (Section 3.2). We use curriculum learning to implement these two steps. This results in the following benefits compared to prior GAN-based approaches:
 - SS-SFDA directly exploits the pre-trained model and trains via curriculum learning to progressively bridge the domain gap between the pre-trained source domain and target domain and achieve faster training times.
 - Our online self-training scheme overcomes the saturation issues.
2. For heterogeneous adverse weather datasets, we propose a method that extends SS-SFDA by leveraging a few labeled images from the target domain to improve the accuracy using model distillation (Section 3.3).

We have evaluated our approach on 6 datasets corresponding to real and synthetic adverse weather conditions. Overall, our mIoU score is 88 – 96% of prior supervised methods. We also improve the training time over prior SFDA approaches by 18 – 180 \times . Finally, our improvement in terms of mIoU over the best SFDA approach is 10.26% on real adverse weather data.

2. Related Work

We discuss recent work related to road segmentation, domain adaptation and source-free domain adaptation, and self-supervised learning.

2.1. Road Segmentation

Research in deep learning for semantic segmentation [34, 57, 8, 9, 62, 18, 49] has paved the way for segmentation in urban traffic scenes like CityScapes [12]. These methods have been extended for supervised road segmentation [54, 64, 17, 48]. Our approach based on self-supervised learning is complimentary to these methods.

2.2. Domain Adaptation and Source Free Domain Adaptation

Traditional domain adaptation [23, 22, 45, 53, 50, 10] methods have achieved remarkable success in adapting models from one domain to another for clear weather conditions. However, these methods need access to the source data. Many domain specific solutions have been proposed for adverse weather conditions, including specific solutions for driving in rain, fog, etc. [40, 36, 43, 14, 44, 39] In contrast, we propose a generic method that neither relies on specific details from each domain, nor requires access to source data during the adaptation stage.

In source-image free domain adaptation (SFDA), a deep neural network (DNN) pre-trained on a source dataset is required to directly make predictions on the target domain dataset in an unsupervised manner. This approach has been primarily used for image classification tasks. In SFDA, generative approaches [28, 24, 29, 56, 31] are used to either emulate the source data by using the feature representations of the pre-trained model or create a negative source dataset during the pre-training stage. Non-generative approaches [26, 55, 33, 25] rely on computing adaptive class specific prototypes, and progressively learn on the target images. While generative approaches work well for classification tasks, they have not been effective for urban scene segmentation. This is because segmentation, being a pixel level task, requires the network to encode context, inter-class semantic relations, structure, and intricate boundaries [15, 59, 63]. In addition, training a generator for a complex segmentation task can lead to memory and computational overheads, and thereby adds to the difficulty of training GANs. While non-generative classification methods are easier to adapt for segmentation, the computation of prototypes results in similar issues due to the inherent differences between classification and segmentation tasks. Bateson et al. [1] explored SFDA in the context of medical segmentation and minimized entropy by the incorporation of class priors [1]. However, preservation of class priors does not extend well to urban road scenes because the number of road pixels in each image in the dataset can vary quite drastically. In contrast, we present a new method for SFDA using self-supervised learning that overcomes issues related to training and convergence that are freq generative processes. Our method is designed to be complimentary to GAN-based SFDA.

2.2.1 Self-Supervised Learning

Self-supervised learning has been used in semi-supervised learning [32, 3] and domain adaptation [46]. Most self-supervised learning methods are centered around the ideas of pseudo labeling [30, 11, 35, 65, 27], entropy minimization [52, 19], and curriculum learning [61]. Domain adaptation methods have access to source domain images. In contrast, we propose a novel self-training routine for SFDA and a completely unsupervised problem setting, where we have access to only a pre-trained model, and target domain images.

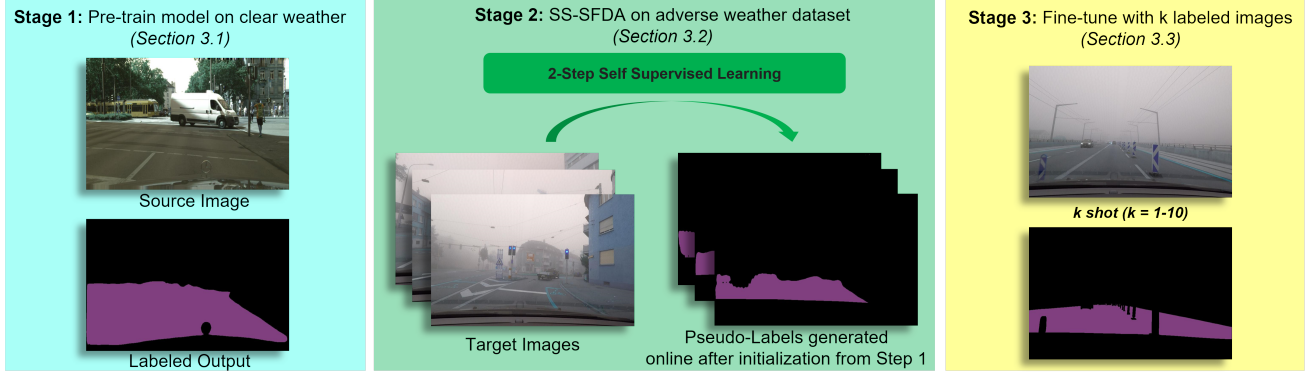


Figure 2: **Our Approach:** In stage 1, our model is pre-trained on a clear weather source dataset. In stage 2, our model is initialized with the pre-trained model from stage 1 and trained using our self-supervised algorithm, SS-SFDA, on the unlabeled adverse weather dataset. For heterogeneous weather datasets, we perform additional refinement steps based on model distillation (stage 3).

3. Our Approach

In this section, we present our approach for source-image free domain adaptive (SFDA) road segmentation based on self-supervised learning. Our approach consists of three main components:

1. **Pre-training using the self-attention auto-encoder:** During this stage, we train the self-attention auto-encoder architecture on a clear weather dataset. This generates a model that encapsulates knowledge about road pixels.
2. **SS-SFDA : A Self-Supervised learning algorithm for SFDA:** During this stage, we initialize the model using the pre-trained model from the previous step and the target domain images. We use a combination of curriculum learning and entropy minimization to bridge the domain gap between the pseudo-labels and the target domain images. We first sort the target domain images in the increasing order of entropy, and create mini-batches of the dataset. The next task is to execute the following steps on each mini-batch to progressively self-train the model:
 - Optimize the model with an entropy minimization constraint to bridge the domain gap.
 - Self-train the model by generating enriched pseudo-labels in an online manner.
3. **Few-Image Regularization:** For heterogeneous weather datasets, we use a very small number of labeled images (5 – 10) from the target domain to boost the performance of SS-SFDA via model distillation. (Section 3.3)

3.1. Pre-Training Baselines Using Self-Attention

The first step in SFDA is to pre-train a DNN on the source dataset for the task of road segmentation. In our case, the source dataset corresponds to traffic videos with clear weather conditions. While networks developed for semantic segmentation [49, 62, 8, 9] can be directly used for road segmentation, there is loss of context i.e. the model is unable to capture relationships between various semantic classes in the images like cars and roads, pedestrians and roads, sky and roads, etc. The loss of such context can lead

to local ambiguities in classifying pixels [15, 59, 63].

Self-attention benefits from its capability to capture long-range dependencies between various regions of the image. Thus, using self-attention in road segmentation can allow neural networks to alleviate the degradation in performance due to loss of context. We use a simple autoencoder self-attention architecture that can be combined with any “off-the-shelf” segmentation network. We begin by taking an input RGB image $\mathcal{I} \in \mathbb{R}^{w \times h \times 3}$, which is passed through an encoder E to generate feature maps F_{en} ($F_{en} = E(\mathcal{I})$). Next, we apply self-attention SA [60] on these feature maps to obtain *attention* maps F_{sa} (same dimensions as F_{en}). These feature maps encapsulate the semantic relationships between various parts of the image. These feature maps F_{sa} are used to learn the final predictions $\mathcal{P}_{out} \in \mathbb{R}^{w' \times h' \times 1}$, which corresponds to the probability that each pixel is classified as ‘road’. In the supervised setting where ground-truth labels $\mathcal{Y} \in \mathbb{Z}^{w' \times h' \times 1}$ are available, the network is optimized with a binary cross-entropy loss function,

$$\mathcal{L}_{CE} = - \sum_{h,w} \mathcal{Y} \log(P_{out}) + (1 - \mathcal{Y}) \log(1 - P_{out}). \quad (1)$$

We use this pre-trained model for SFDA on the adverse weather datasets [43, 42, 58, 21, 51].

3.2. SS-SFDA

In this section, we describe our two-step self-supervised learning algorithm for unsupervised source-image free domain adaptive road segmentation in adverse weather conditions. We use the pre-trained model from the previous step. The pseudo-labels generated by the model are noisy leading to a domain gap between the pseudo-labels generated by the pre-trained model (on the source dataset) and the target domain images. Thus, directly self-training using the pseudo labels can hamper the performance of the model. To counter this, we propose an entropy minimization step (Section 3.2.1) which encourages the network to generate more accurate pseudo labels.

In addition, to bridge the domain gap between the pre-

trained model and the target domain, we use curriculum learning [2, 20] in which the DNN is allowed to train on samples progressively in their increasing order of entropy of predictions. Given a probability map P denoting the probability that pixels are classified as road pixels, the entropy is computed as $-\sum P \times \log(P)$. This is because learning from samples with low entropy (low rain, for example) yields better pseudo labels on samples with higher entropy (high rain, for example) [14, 61, 61]. We create mini-batches of the dataset characterizing the difficulty of the images. For datasets which provide labels on the intensity (light rain vs heavy rain) of the weather condition, mini-batches can be created directly. For other datasets, we sort the images in increasing order of the entropy and then split them into m ($m \sim 4 - 5$, determined by hyperparameter tuning) equal mini-batches. The model is self-trained on the mini-batches in a sequential manner. For the first mini-batch, the model is initialized with the pre-trained model from Section 3.1. For subsequent mini-batches, our model is initialized with weights obtained by training the network on the previous mini-batch. For each mini-batch, the network is trained in two stages, as described below:

3.2.1 Step 1: Bridging the Domain Gap via Entropy Minimization

The pre-trained model from Section 3.1 has a low entropy (*i.e.* high prediction probability or better generalization) on images that are similar (for example, similar geography, light rain, light fog) to source domain images [42] and vice versa. Thus, initializing the network with these pre-trained weights, followed by training by entropy minimization [19, 41] allows the network to generate enriched pseudo-labels. The inputs to the network are images from the target domain. Let the predictions of the network be denoted by $\mathcal{P}_{out} \in \mathbb{R}^{w' \times h' \times 1}$, corresponding to the probability that each pixel is classified as ‘road’. The cost function for entropy minimization is given by,

$$L_{EM} = - \sum_{\forall \text{pixels}} P^{h,w} \log P^{h,w}, \quad (2)$$

where $P^{h,w}$ is the probability that a pixel belongs to a class ‘road’ at a given location, and $-P^{h,w} \log, P^{h,w}$ is the entropy.

3.2.2 Step 2: Online Self-Training Using Enriched Pseudo-Labels

The network trained in Step 1 generates enhanced pseudo labels with high probability, and is thus a better representative of the target domain than the pre-trained source model from Section 3.1. These enriched pseudo labels which can be used to self-train the model further to improve performance. A traditional method of self-training using pseudo-labels is iterative [38] in which the network is trained to convergence (validation loss less than a given threshold) over multiple iterations. In each iteration, pseudo labels from the trained model in the previous iteration are used to set up the

Dataset	Syn./Real	Weather
Rainy CityScapes [21]	Syn	Varying intensities (1mm - 200mm) of rain
Foggy CityScapes [21]	Syn	Varying intensities (750m - 30m) of fog
Foggy Zurich [43]	Real	Light and medium Fog
Dark Zurich [42]	Real	Twilight, Night
Raincover [51]	Real	Rain, night
BDD [58]	Real	Snow, Fog, Rain, Night

Table 1: **List of datasets:** The second column categorizes the datasets as synthetic or real and the third column describes the images contained in the dataset.

binary segmentation cost function. We show (Table 10) that iterative self training does not lead to any improvement in performance. This is because, the pre-training step, which is imperative for acquiring initial knowledge about road pixels since the problem is unsupervised in the target domain, causes the network to saturate quickly. Hence, we generate the pseudo labels in an online manner, *i.e.* the pseudo labels are generated from the network that is being trained. This allows the network to self-train from the improved pseudo labels as they are learnt.

The network in this stage is initialized with the weights obtained in Step 1. The inputs to the network are images from the target domain. Pseudo labels are generated in an online fashion from the network being trained as follows,

$$Y_{\text{pseudo}} = \begin{cases} 1 & \text{if } P^{h,w} \geq \tau, \\ 0 & \text{otherwise,} \end{cases}$$

where $P^{h,w}$ is the probability that a pixel belongs to the class ‘road’ at a given location and τ is a threshold. The network is optimized with these pseudo-labels using a binary cross entropy loss term, (similar to Equation 1).

3.3. Few-Image Fine-Tuning via Model Distillation

Some heterogeneous weather datasets like Raincover [51] and Berkeley Deep Drive (BDD) [58] contain a mixture of adversities within the same image (for instance night+rain in Raincover, see Table 1). Furthermore, these datasets are captured from different geographic conditions (*i.e.*, source and target datasets may be from different regions). To make our model robust against such factors, we use ground truth labels for a few images (order of 5 – 10 images) from the target dataset in a final refinement step described below. In a nutshell, given a model trained on the unlabeled target dataset using SS-SFDA, and $k \leq 10$ labeled images from the target domain images, our goal is to learn enhanced feature maps for the target domain in the presence of adversarial factors such as mixtures of adversities and different geographical regions.

We empirically observe that directly fine-tuning the SS-SFDA model on the k images is sub-optimal due to overfitting. To prevent overfitting, we propose a model distillation [31] regularizer. Let the weights of the SS-SFDA model be denoted by $\omega_{\text{SS-SFDA}}$, and the weights of the model being currently trained be denoted by ω_{fewIm} . The cost function for model distillation is given by,

$$L_{\text{model-distil}} = C(\omega_{\text{SS-SFDA}}, \omega_{\text{fewIm}}),$$

Model	Acc.(%) (w/o. SA)	Acc.(%) (w. SA)
DeepLabv2 [8]	89.59	90.54 (+0.95)
DeepLabv2 (E-1) [8]	87.50	90.78 (+3.28)
DeepLabv2 (E-2) [8]	88.13	88.62 (+0.49)
DRN-D-105 [57]	83.92	85.32 (+1.40)
DRN-D-38 [57]	90.69	91.44 (+0.75)

Table 2: **Effect of self-attention:** We show for various backbone architectures that self-attention improves the accuracy. DeepLabv2 (E-1) and DeepLabv2 (E-2) denote the removal of 1 and 2 layers from DeepLab respectively. We select DRN-D-38 with self-attention (**bolded**) as the baseline for all further experiments.

where C represents a distance function such as MSE distance or L1 distance. In our benchmarks, MSE distance works best.

The network is first initialized with weights of the SS-SFDA model. The model distillation term $L_{\text{model-distil}}$ with weight parameter $\lambda_{\text{model-distil}}$ is applied in conjunction with the binary cross entropy loss function (Equation 1) to constrain the probability predictions and ground-truth labels for k images. The $\lambda_{\text{model-distil}}$ term balances between extracting domain specific characteristics from the k images (such as mix of adverse weather, geographical features etc.) and prevents the weights of the model from diverging from the SS-SFDA weights (for better generalization). The overall equation follows as,

$$\mathcal{L}_{\text{overall}} = L_{CE} + \lambda_{\text{model-distil}} L_{\text{model-distil}} \quad (3)$$

4. Experiments and Results

We use the CityScapes dataset as the clear weather source domain. We conduct evaluation experiments on 6 datasets captured in adverse environmental conditions, described in Table 1. We evaluate our model using four metrics: mean Intersection over Union (mIoU), Recall (or accuracy), Precision, and F1 score. All our models are trained using one NVIDIA GeForce GPU, and we implement the model using the PyTorch framework. We will make all code publicly available. The hyperparameters generalize across our experiments on all datasets. For the segmentation model, we use the SGD optimizer with a learning rate of $2.5e - 4$, and momentum of 0.9 and weight decay of 0.0005. Dataset specific details are provided in the table below. Images are downsampled (by a factor of 2, where necessary) by bilinear sampling, and the corresponding ground-truth labels are downsampled by nearest neighbour down-sampling.

4.1. Analysing Pre-Trained Self-Attention-based AutoEncoder

In this section, we analyse the performance of our self-attention auto-encoder model (described in Section 3.1 and benchmark its performance on various datasets in the supervised setting. In Table 2 (I), we show that the usage of self-attention within various conventional semantic segmentation models to encode semantic relationships via capturing long-range dependencies improves performance over the corresponding baseline. We use the DRN-D-38 model

Intensity	mIoU	Recall	Prec.	F1
I. Synthetic Rain [21]				
1mm	93.23	96.13	96.86	96.49
5mm	94.06	96.95	96.92	96.93
17mm	93.92	96.87	96.85	96.86
25mm	93.07	96.4	96.42	96.41
50mm	92.45	95.72	96.43	96.08
75mm	92.05	96.25	95.47	95.86
100mm	91.57	96.66	94.56	95.60
200mm	90.57	95.69	94.42	95.05
II. Synthetic Fog [21]				
750m	95.53	97.46	97.97	97.71
375m	94.74	97.08	97.51	97.30
150m	92.72	95.71	96.74	96.22
75m	91.64	96.12	95.15	95.63
50m	90.59	94.97	95.16	95.06
40m	90.21	95.68	94.03	94.85
30m	89.00	94.33	94.02	94.18
III. Real datasets				
Raincouver [51]	71.88	80.36	87.19	83.64
BDD [58]	89.19	92.86	95.76	94.29
IV. SOTA Comparisons				
Dataset	Method	mIoU	Recall	F1
CityScapes	FCN [34]	89.90	95.70	94.68
	CA [18]	88.93	94.05	94.14
	BoT [47]	76.55	85.82	86.72
	DeepLabv3 [9]	90.87	95.35	95.21
	s-FCN-loc [54]	91.04	96.11	95.36
	Zohourian [64]	86.34	96.76	92.44
	RBA [48]	96.00	98.13	98.00
Ours	91.44	96.34	95.52	
100mm Rain	SNE-Seg [17]	90.80	96.80	95.80
	Ours	91.57	96.66	95.60
100mm Rain	SNE-Seg [17]	90.50	93.00	94.47
	Ours	91.64	96.12	95.63

Table 3: **Benchmarking and comparing the self-attention based pre-trained model:** We train our self-attention model from Section 3.1 on various datasets in a supervised manner. These supervised numbers help us conduct a relative study of the performance of SS-SFDA, which is unsupervised in the target domain. Additionally, we observe that our pre-trained model is comparable to the state-of-the-art (Experiment IV). Best results are in **bold** fonts, second best results are in **blue**.

[57] with self-attention in all further experiments. In Table 3 (I,II,III), we benchmark the model on various weather datasets in the supervised setting. We use these supervised numbers in the following subsections to perform a comparative study of our self-supervised model (See Figure 1 for a summary). In Table 3 (IV), we show that the self-attention autoencoder is comparable to the state-of-the-art on various datasets.

4.2. Results on Synthetic Datasets: Rain and Fog

We perform three evaluation experiments. Experiment A corresponds to testing the pre-trained CityScapes model on varying intensities of rain and fog. Experiment B corresponds to results obtained by training SS-SFDA without curriculum learning (*i.e.* by initializing with the CityScapes pre-trained model and training on higher intensities of rain and fog directly), and experiment C corresponds to results

Intensity	Experiment	mIoU	Recall	Prec.	F1
I. Synthetic Rain					
1 mm	A	94.73	97.16	97.42	97.29
5 mm	A	94.09	97.01	96.90	96.95
17 mm	A	93.95	96.83	96.93	96.88
25 mm	A	93.01	96.30	96.46	96.38
50mm	A	89.82	93.71	95.58	94.64
75mm	A	87.15	90.92	95.45	93.13
75mm	B	88.08	93.94	93.38	93.66
100mm	A	82.97	86.88	94.84	90.69
100mm	B	86.67	92.88	92.84	92.86
200mm	A	65.98	68.94	93.90	79.50
200mm	B	80.25	86.96	91.22	89.04
200mm	C	81.58	88.55	91.20	89.85
II. Synthetic Fog					
750m	A	94.90	96.68	98.09	97.38
375m	A	92.71	94.60	97.88	96.21
150m	A	85.55	87.57	97.38	92.21
150m	B	89.74	94.59	94.59	94.59
75m	A	70.81	72.51	96.79	82.9
75m	B	80.67	83.01	96.62	89.30
75m	C	87.48	92.99	93.65	93.32
50m	A	57.29	58.92	95.39	72.85
50m	B	71.64	73.69	96.24	83.47
50m	C	85.47	90.75	93.63	92.16
40m	A	46.27	48.45	91.12	63.27
40m	B	52.62	53.94	95.57	68.96
40m	C	83.65	89.19	93.09	91.10
30m	A	35.04	38.05	81.59	51.90
30m	B	22.24	22.92	88.29	36.39
30m	C	80.82	85.96	93.11	89.39

Table 4: **Results on synthetic rain and synthetic fog [21]:** We evaluate our model under three settings - A: (Baseline) Testing the clear weather CityScapes model, B: SS-SFDA without curriculum learning, C: SS-SFDA. We show that the CityScapes model generalizes well to light synthetic rain and fog, and its performance degrades as the intensity of rain and fog increases, which is restored by SS-SFDA. **Comparison to supervised results:** On synthetic rain, SS-SFDA achieves 90.07%-90.8% and 92.53%-97.6% of supervised mIoU and recall, respectively. On synthetic fog, SS-SFDA achieves 90.8%-96.78% and 91.12%-98.82% of supervised mIoU and recall respectively.

obtained by training SS-SFDA as proposed in Section 3.2 (*i.e.* with curriculum learning).

Results on Synthetic Rain: The results of SS-SFDA on Synthetic Rainy CityScapes [21] are shown in Table 4 (I). For low intensities of rain, the performance of the CityScapes model is more or less preserved (Table 3 (I)). For higher intensities of rain, there is a degradation in performance. We observe that the decrease in performance is highest for 200mm rain, at 27.15%. For 75mm, 100mm, and 200mm, Experiment B imparts an improvement of 1.03%, 4.45%, and 21.63% respectively over the corresponding baselines. Experiment C leads to a cumulative improvement of 23.64% over the corresponding baseline. Furthermore, we demonstrate that the mIoU and recall of our SS-SFDA, which is completely unsupervised, is 90.07% – 90.80% and 92.53% – 97.60% of the counterpart

supervised mIoU.

Results on Synthetic Fog: The results of our SS-SFDA on Synthetic Foggy CityScapes [21] are shown in Table 4II. Similar to synthetic rain, the performance of the CityScapes model on light fog is more or less preserved (Table 3 (II)). For higher intensities of fog, there is a degradation in performance. We notice that the degradation is very high for visibility distances less than 150m, and is the highest at 60.6% for 30m fog. Experiment B leads to an improvement of 4.89%, 13.9%, 25.04%, 13.7% on 150m, 75m, 50m, and 40m fog respectively. Direct application of the self-training algorithm on 30m fog (without curriculum learning) degrades performance since the generalization of CityScapes on 30m is very poor. Experiment C (Training SelfTr-Road Seg by curriculum learning *i.e.* progressively from 150m fog to 30m fog) cumulatively improves performance over the baselines (Experiment A) by 23.54%, 28.18%, 80.78% and 130.65% on 75m, 50m 40m, and 30m fog respectively. Furthermore, we demonstrate that the mIoU and recall of our SS-SFDA, which is completely unsupervised, is 90.8%-96.78% and 91.12%-98.82% of the supervised mIoU.

Benefits of Curriculum Learning: On a given dataset, the performance of our self-training algorithm depends heavily on the generalization capabilities of the pre-trained model used for initialization. Therefore, progressively initializing and training the model on increasing intensities of rain and fog will lead to the best accuracies (Table 4, Table 4). This implies that progressively training from low intensities to high intensities of rain and fog improves the quality of pseudo labels and prediction probabilities on high intensities of rain and fog. In simpler terms, for the 100mm rain dataset, a model trained on 75mm rain will work better than the CityScapes clear weather model. Similarly, for the 200mm rain dataset, a model trained on 100mm rain will work better than the 75mm rain model, which in turn will generalize better the CityScapes clear weather model. A similar intuition can be drawn for synthetic fog too. We validate this hypothesis in Table 5. The first and second columns correspond to the datasets that the model is trained and tested on respectively. We observe that curriculum learning progressively improves performance, which results in high quality pseudo labels with high confidence, a boon for self-training.

Generalization trends: We observe that our self-supervised SS-SFDA preserves the accuracy on the clear weather source dataset, CityScapes. The mIoU and accuracy of the fog model on clear weather CityScapes are at 96.9% and 99.00% of the supervised CityScapes model, respectively. The corresponding accuracy numbers for synthetic rain are 98.11% and 98.72%, respectively.

4.3. Results on Real datasets: Foggy Zurich and Dark Zurich

Analysis: Foggy Zurich The results are presented in Table 6. The pre-trained CityScapes model fails to generalize (Table 6 (I)) to the real fog dataset due to the domain gap. In accordance with the curriculum learning strategy, the self-training algorithm is first applied on images with light fog, and then on images with medium fog. We show results




Figure 3: **Qualitative results.** Our model generates results that closely resemble the ground-truth (GT) compared to the baseline CityScapes pre-trained model. Purple indicates the segmented road region. More results can be found in the supplementary material.

Training	Testing	mIoU	Recall	Prec.	F1
CS	100mm rain	82.97	86.88	94.84	90.69
75mm rain	100mm rain	86.31	92.63	92.68	92.65
CS	200mm rain	65.98	68.94	93.9	79.5
75 mm rain	200mm rain	78.97	85.86	90.77	88.25
100 mm	200mm rain	79.61	86.27	91.15	88.65
CS	75m fog	70.81	72.51	96.79	82.9
150m fog	75m fog	83.82	87.44	95.28	91.19
CS	50m fog	57.29	58.92	95.39	72.85
150m fog	50m fog	76.07	79.49	94.64	86.41
75m fog	50m fog	84.38	89.44	93.71	91.52
CS	40m fog	46.27	48.45	91.12	63.27
150m fog	40m fog	65.6	69.98	91.28	79.22
75m fog	40m fog	78.17	83.51	92.43	87.74
CS	30m fog	35.04	38.05	81.59	51.9
150m fog	30m fog	53.61	59.62	84.17	69.8
75m fog	30m fog	68.49	75.21	88.45	81.3

Table 5: **Curriculum learning improves performance by improving the accuracy of pseudo labels:** The first and second columns correspond to the datasets our model is trained and tested, respectively. Evaluating on 30m fog reveals that a model finetuned to 75m fog performs better than the counterpart 150m fog model, which in turn performs better than that the CityScapes model. Similar conclusions can be drawn for other intensities of rain/fog.

Experiment	mIoU	Recall	Prec.	F1
I. Baseline: CityScapes pre-trained model				
-	36.91	57.04	51.13	53.92
II. SS-SFDA on Light Fog; Init: CS model				
Step 1	59.1	68.27	81.48	74.29
Step 2	60.9	69.09	83.7	75.7
III. SS-SFDA on Medium Fog; Init: Light fog model;				
Step 1	73.34	76.22	95.1	84.62
Step 2	74.5	76.74	96.23	85.39

Table 6: **Results on the real fog dataset, Foggy Zurich.** We observe that SS-SFDA improves mIoU by 101.84% over the corresponding baseline. We notice that training the model on light fog followed by medium fog imparts the model with improvised self-training abilities, the mIoU improves by 22.3%. Additionally, the ablations on each minibatch indicate that the two step training procedure in SS-SFDA is indeed helpful.

for each of the two stages of SS-SFDA to demonstrate how the two-step training procedure gradually improves perfor-

Experiment	mIoU	Recall	Prec.	F1
I. Baseline: CityScapes pre-trained model				
-	54.9	63.89	79.6	70.88
II. SS-SFDA on Twilight Zurich; Init: CS model				
Step 1	69.11	74.75	90.15	81.74
Step 2	69.13	74.77	90.16	81.74
III. SS-SFDA on Night Zurich; Init: Twilight model				
Step 1	71.75	76.22	92.44	83.55
Step 2	72.18	76.51	92.72	83.84

Table 7: **Results on the night driving dataset, Dark Zurich.** We observe that SS-SFDA improves mIoU by 31.47. We observe that our two-stage training routine within the curriculum learning algorithm helps the network perform well.

Experiment	mIoU	Recall	Prec.	F1
I. SS-SFDA				
-	52.42	93.96	54.25	68.78
II. Results on FewIm-FT (Section 3.3)				
k=1	54.71	90.7	57.96	70.74
k=2	56.74	90.45	60.36	72.4
k=10	63.85	79.39	76.54	77.93
III. Effect of model distillation (MD); k=10				
Without MD	57.08	70.99	74.45	72.67
With MD	63.85	79.39	76.54	77.93

Table 8: **Results on Raincover:** SS-SFDA achieves 72.9% of supervised mIoU. We demonstrate the effectiveness of FewIm-FT (Section 3.3) over varying values of k in Experiment II. $k = 10$ results in an improvement of 21.8% over SS-SFDA, which is 88.82% of supervised mIoU. Additionally, we show the benefit of model distillation in Experiment III.

mance. Training with light fog improves the mIoU by 64.99% over the corresponding clear weather baseline. Initializing the model trained on light fog, and fine-tuning on medium fog using SS-SFDA further improves the performance by 22.3%, thus resulting in a cumulative improvement of 101.84%. We do not show comparisons with supervised methods due to a lack of labeled datasets.

Experiment	mIoU	Recall	Prec.	F1
CS pre-trained model	70.28	74.49	92.54	82.54
SS-SFDA	75.29	83.3	88.7	85.92
FewIm-FT, k=10	83.05	93.06	88.53	90.74

Table 9: **Results on BDD:** SS-SFDA improves baseline mIoU by 7.12%. Experiments on FewIm-FT (Section 3.3) demonstrate an improvement of 10.30% over SS-SFDA for k=10, which is 93.11% of supervised mIoU.

Experiment	Method	Source Data	mIoU	%(sup)
I. Synthetic Dataset: 100 mm rain				
Domain Adaptation	ADVENT [52]	✓	85.96	93.87%
Domain Adaptation	PrDA [26]	✓	86.12	94.04%
Self-training	CBST [65]		78.64	85.87%
SFDA	PrDA [26]		84.77	92.57%
SFDA	Ours		86.67	94.64%
II. Heterogeneous Real Dataset: Berkeley Deep Drive				
SFDA	PrDA [26]		75.32	82.43%
SFDA	Ours		83.05	93.11%

Table 10: **Comparisons against the state-of-the-art:** We adopt the state-of-the-art methods in self-training, domain adaptation, and SFDA classification to road segmentation, and present the results on Rain 100mm. We pick the best performing source-free model and conduct experiments on BDD. We observe an improvement of 10.26%.

Analysis-Dark Zurich: The results are presented in Table 7. Table 7 (I) shows that the CityScapes pre-trained model does not work well on Dark Zurich. SS-SFDA first trains on the twilight images, and then on the night images. Training with twilight images improves the mIoU by 24.91% over the corresponding clear weather baseline. Initializing the model with the model trained on twilight images, and fine-tuning on night images further improves performance by 4.41%, thus resulting in a cumulative improvement of 31.47%. We do not show comparisons to supervised accuracies due to the unavailability of training labels.

4.4. Heterogeneous Real Datasets: Raincouver and Berkeley Deep Drive

Raincouver [51] and Berkeley Deep Drive (BDD) [58] are complex datasets with images containing a mix of weather conditions in addition to scenes from different geographical regions. Raincouver consists of images captured under rain during the night. BDD consists of images in snow, fog, low light, glare, rain, etc. We show that SS-SFDA benefits from supervised finetuning with just 5 – 10 images using the procedure discussed in Section 3.3. The models converge in 40 iterations, which takes 2 minutes to train on one NVIDIA GeForce GPU with 11GB memory.

Analysis-Raincouver: The results are shown in Table 8. SS-SFDA results in an mIoU of 52.42, which is 72.9% of supervised mIoU. In Table 8 (II), we show the effectiveness of the fine-tuning step. As k (number of supervised images used by the algorithm) increases, the performance of the model improves. Using just 1 image ($k = 1$) results in an improvement of 4.3% (over SS-SFDA). $k = 10$ achieves 88.82% of supervised mIoU. In Table 8 (III), we demonstrate that model distillation improves mIoU by 11.86%. Hyperparameter tuning on the model distillation hyperparameter $\lambda_{\text{model-distil}}$ reveals that a value of 1.0 works best.

Analysis-Berkeley Deep Drive: The CityScapes model results in an mIoU of 70.28 (baseline). Training with SS-SFDA improves performance by 7.12% over the baseline, and is at 84.41% of supervised IoU. In Experiment III, we pick k random images from the dataset and finetune the network using the fine-tuning step. In concurrence with our intuition, we observe that the performance improves as the number of images increases. For $k = 10$, we demonstrate an mIoU improvement of 10.30% over SS-SFDA, which is 93.11% of supervised accuracy.

4.5. Training Time and Convergence

The pre-training step helps our model converges in 1/6 epoch (the accuracies are similar over multiple random runs), thus bringing the training time down to 15 minutes on an NVIDIA GeForce GPU with 11GB memory. Generative SFDA models [28] take 30 epochs to converge, while prior work on self-training [65] and SFDA [26] applied (where feasible) to road segmentation converge in 3 epochs. Thus, we improve training time over prior SFDA approaches by 18 – 180 \times .

4.6. Comparisons with Prior Work

We adapt (where feasible) the state-of-the-art method in the following categories: self-supervised learning [65], source-free DA[26], and DA [52] to road segmentation for comparisons. Other methods like feature alignment by class-wise prototype learning [33], generative methods [28, 24, 29, 56, 31], and models that optimize source domain class priors [20] do not scale well to SFDA road segmentation due to the problems highlighted in Section 2.2. We performed evaluations using each of the methods on synthetic rain (Table 10 (I)) and observe that our model outperforms all prior methods. We finally train the best performing SOTA model on BDD (Table 10 (II)) for comparisons on heterogeneous real datasets which carry the highest level of difficulty. On BDD our method outperforms prior methods by 10.26%.

5. Conclusion, Limitations, and Future Work

We propose a new method for road segmentation in adverse weather conditions using a novel self-supervised source-free domain adaptation approach. Through our evaluations on 6 real and synthetic datasets, we show that that our self-supervised model that has access to only a pre-trained clear weather model and unlabeled target images exhibits accuracy that is comparable to completely supervised models which have access to labels for all target domain images. In addition, we exhibit benefits in terms of faster training time and state-of-the-art performance.

There are a few limitations of our work. Currently, our approach is designed for binary segmentation and cannot perform multi-class segmentation. Extending the current method to multi-class segmentation would generalize the approach beyond road segmentation. Moreover, it would be interesting to investigate if SFDA via self-supervised learning could be extended to other computer vision problems such as object recognition, image classification, scene understanding, etc.

References

- [1] Mathilde Bateson, Hoel Kervadec, Jose Dolz, Hervé Lombaert, and Ismail Ben Ayed. Source-relaxed domain adaptation for image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 490–499. Springer, 2020. [2](#)
- [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009. [4](#)
- [3] Paola Cascante-Bonilla, Fuwen Tan, Yanjun Qi, and Vicente Ordonez. Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning. *arXiv preprint arXiv:2001.06001*, 2020. [2](#)
- [4] Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Taphic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8483–8492, 2019. [1](#)
- [5] Rohan Chandra, Uttaran Bhattacharya, Tanmay Randhavane, Aniket Bera, and Dinesh Manocha. Roadtrack: Realtime tracking of road agents in dense and heterogeneous environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1270–1277. IEEE, 2020. [1](#)
- [6] Rohan Chandra, Uttaran Bhattacharya, Christian Roncal, Aniket Bera, and Dinesh Manocha. Robusttp: End-to-end trajectory prediction for heterogeneous road-agents in dense traffic with noisy sensor inputs. In *ACM Computer Science in Cars Symposium*, pages 1–9, 2019. [1](#)
- [7] Rohan Chandra, Tianrui Guan, Srujan Panuganti, Trisha Mittal, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Forecasting trajectory and behavior of road-agents using spectral clustering in graph-lstms. *IEEE Robotics and Automation Letters*, 5(3):4882–4890, 2020. [1](#)
- [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*, 40(4):834–848, 2017. [2](#), [3](#), [5](#)
- [9] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. [2](#), [3](#), [5](#)
- [10] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1992–2001, 2017. [2](#)
- [11] Jaehoon Choi, Minki Jeong, Taekyung Kim, and Changick Kim. Pseudo-labeling curriculum for unsupervised domain adaptation. *arXiv preprint arXiv:1908.00262*, 2019. [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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. [2](#)
- [13] Michael A Cusumano. Self-driving vehicle technology: progress and promises. *Communications of the ACM*, 63(10):20–22, 2020. [1](#)
- [14] Dengxin Dai, Christos Sakaridis, Simon Hecker, and Luc Van Gool. Curriculum model adaptation with synthetic and real data for semantic foggy scene understanding. *International Journal of Computer Vision*, 128(5):1182–1204, 2020. [2](#), [4](#)
- [15] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. Context contrasted feature and gated multi-scale aggregation for scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2393–2402, 2018. [2](#), [3](#)
- [16] Piotr Dollar, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: An evaluation of the state of the art. *IEEE transactions on pattern analysis and machine intelligence*, 34(4):743–761, 2011. [1](#)
- [17] Rui Fan, Hengli Wang, Peide Cai, and Ming Liu. Sneroadseg: Incorporating surface normal information into semantic segmentation for accurate freespace detection. In *European Conference on Computer Vision*, pages 340–356. Springer, 2020. [1](#), [2](#), [5](#)
- [18] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3146–3154, 2019. [1](#), [2](#), [5](#)
- [19] Yves Grandvalet, Yoshua Bengio, et al. Semi-supervised learning by entropy minimization. In *CAP*, pages 281–296, 2005. [2](#), [4](#)
- [20] Guy Hacohen and Daphna Weinshall. On the power of curriculum learning in training deep networks. In *International Conference on Machine Learning*, pages 2535–2544. PMLR, 2019. [4](#), [8](#)
- [21] Shirsendu Sukanta Halder, Jean-François Lalonde, and Raoul de Charette. Physics-based rendering for improving robustness to rain. In *ICCV*, 2019. [3](#), [4](#), [5](#), [6](#)
- [22] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018. [1](#), [2](#)
- [23] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. [2](#)
- [24] Yunzhong Hou and Liang Zheng. Source free domain adaptation with image translation. *arXiv preprint arXiv:2008.07514*, 2020. [1](#), [2](#), [8](#)
- [25] Masato Ishii and Masashi Sugiyama. Source-free domain adaptation via distributional alignment by matching batch normalization statistics. *arXiv preprint arXiv:2101.10842*, 2021. [2](#)
- [26] Youngeun Kim, Sungeun Hong, Donghyeon Cho, Hyoungseob Park, and Priyadarshini Panda. Domain adaptation without source data. *arXiv preprint arXiv:2007.01524*, 2020. [1](#), [2](#), [8](#)
- [27] Divya Kothandaraman, Rohan Chandra, and Dinesh Manocha. Bomuda: Boundless multi-source domain adaptive segmentation in unconstrained environments. *arXiv preprint arXiv:2010.03523*, 2020. [2](#)
- [28] Jogendra Nath Kundu, Naveen Venkat, R Venkatesh Babu, et al. Universal source-free domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4544–4553, 2020. [1](#), [2](#), [8](#)
- [29] Vinod K Kurmi, Venkatesh K Subramanian, and Vinay P Namboodiri. Domain impression: A source data free domain adaptation method. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 615–625, 2021. [1](#), [2](#), [8](#)

- [30] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013. [2](#)
- [31] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9641–9650, 2020. [2](#), [4](#), [8](#)
- [32] Xinzhe Li, Qianru Sun, Yaoyao Liu, Shibao Zheng, Qin Zhou, Tat-Seng Chua, and Bernt Schiele. Learning to self-train for semi-supervised few-shot classification. *arXiv preprint arXiv:1906.00562*, 2019. [2](#)
- [33] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 6028–6039. PMLR, 2020. [2](#), [8](#)
- [34] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [2](#), [5](#)
- [35] Pietro Morerio, Riccardo Volpi, Ruggero Ragonese, and Vittorio Murino. Generative pseudo-label refinement for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3130–3139, 2020. [2](#)
- [36] Alexandra S Mueller and Lana M Trick. Driving in fog: The effects of driving experience and visibility on speed compensation and hazard avoidance. *Accident Analysis & Prevention*, 48:472–479, 2012. [1](#), [2](#)
- [37] Arun Reddy Nelakurthi, Ross Maciejewski, and Jingrui He. Source free domain adaptation using an off-the-shelf classifier. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 140–145. IEEE, 2018. [1](#)
- [38] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3764–3773, 2020. [4](#)
- [39] Fabio Pizzati, Raoul de Charette, Michela Zaccaria, and Pietro Cerri. Domain bridge for unpaired image-to-image translation and unsupervised domain adaptation. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 2990–2998, 2020. [2](#)
- [40] Horia Porav, Tom Bruls, and Paul Newman. I can see clearly now: Image restoration via de-raining. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7087–7093. IEEE, 2019. [2](#)
- [41] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8050–8058, 2019. [4](#)
- [42] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7374–7383, 2019. [1](#), [3](#), [4](#)
- [43] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 687–704, 2018. [1](#), [2](#), [3](#), [4](#)
- [44] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9):973–992, 2018. [2](#)
- [45] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Unsupervised domain adaptation for semantic segmentation with gans. *arXiv preprint arXiv:1711.06969*, 2:2, 2017. [2](#)
- [46] Inkyu Shin, Sanghyun Woo, Fei Pan, and In So Kweon. Two-phase pseudo label densification for self-training based domain adaptation. In *European Conference on Computer Vision*, pages 532–548. Springer, 2020. [2](#)
- [47] Aravind Srinivas, Tsung-Yi Lin, Niki Parmar, Jonathon Shlens, Pieter Abbeel, and Ashish Vaswani. Bottleneck transformers for visual recognition. *arXiv preprint arXiv:2101.11605*, 2021. [5](#)
- [48] Jee-Young Sun, Seung-Wook Kim, Sang-Won Lee, Ye-Won Kim, and Sung-Jea Ko. Reverse and boundary attention network for road segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. [1](#), [2](#), [5](#)
- [49] Towaki Takikawa, David Acuna, Varun Jampani, and Sanja Fidler. Gated-scnn: Gated shape cnns for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5229–5238, 2019. [2](#), [3](#)
- [50] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018. [2](#)
- [51] Frederick Tung, Jianhui Chen, Lili Meng, and James J Little. The raincover scene parsing benchmark for self-driving in adverse weather and at night. *IEEE Robotics and Automation Letters*, 2(4):2188–2193, 2017. [3](#), [4](#), [5](#), [8](#)
- [52] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2517–2526, 2019. [1](#), [2](#), [8](#)
- [53] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7364–7373, 2019. [2](#)
- [54] Qi Wang, Junyu Gao, and Yuan Yuan. Embedding structured contour and location prior in siamesed fully convolutional networks for road detection. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):230–241, 2017. [2](#), [5](#)
- [55] Shiqi Yang, Yaxing Wang, Joost van de Weijer, and Luis Herranz. Unsupervised domain adaptation without source data by casting a bait. *arXiv preprint arXiv:2010.12427*, 2020. [2](#)
- [56] Hao-Wei Yeh, Baoyao Yang, Pong C Yuen, and Tatsuya Harada. Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 474–483, 2021. [1](#), [2](#), [8](#)
- [57] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 472–480, 2017. [2](#), [5](#)
- [58] 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 preprint arXiv:1805.04687*, 2(5):6, 2018. [3](#), [4](#), [5](#), [8](#)
- [59] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Amrbrish Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7151–7160, 2018. [2](#), [3](#)
- [60] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International conference on machine learning*, pages 7354–7363. PMLR, 2019. [1](#), [3](#)
- [61] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2020–2030, 2017. [2](#), [4](#)
- [62] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [1](#), [2](#), [3](#)
- [63] Yizhou Zhou, Xiaoyan Sun, Zheng-Jun Zha, and Wenjun Zeng. Context-reinforced semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4046–4055, 2019. [2](#), [3](#)
- [64] Farnoush Zohourian, Borislav Antic, Jan Siegemund, Mirko Meuter, and Josef Pauli. Superpixel-based road segmentation for real-time systems using cnn. In *VISIGRAPP (5: VIS-APP)*, pages 257–265, 2018. [2](#), [5](#)
- [65] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018. [2](#), [8](#)