

FIFO: Learning Fog-invariant Features for Foggy Scene Segmentation

Sohyun Lee
GSAI, POSTECH

lshig96@postech.ac.kr

Taeyoung Son*
NALBI

taeyoung@nalbi.ai

Suha Kwak
GSAI, POSTECH

suha.kwak@postech.ac.kr

<http://cvlab.postech.ac.kr/research/FIFO/>

Abstract

Robust visual recognition under adverse weather conditions is of great importance in real-world applications. In this context, we propose a new method for learning semantic segmentation models robust against fog. Its key idea is to consider the fog condition of an image as its style and close the gap between images with different fog conditions in neural style spaces of a segmentation model. In particular, since the neural style of an image is in general affected by other factors as well as fog, we introduce a fog-pass filter module that learns to extract a fog-relevant factor from the style. Optimizing the fog-pass filter and the segmentation model alternately gradually closes the style gap between different fog conditions and allows to learn fog-invariant features in consequence. Our method substantially outperforms previous work on three real foggy image datasets. Moreover, it improves performance on both foggy and clear weather images, while existing methods often degrade performance on clear scenes.

1. Introduction

We have witnessed great advances in semantic segmentation for the last decade. However, most of existing models and datasets focus merely on improving accuracy under controlled environments, without considering image degradation caused by adverse weather conditions (*e.g.*, fog, rain, and snow), over- and under-exposure, motion blur, sensor noise, *etc.* The robustness of semantic segmentation models against these factors is of great importance in safety-critical applications and recently has gained increasing attention [3, 6, 8, 50–52, 56, 66].

Motivated by this, we study semantic segmentation of foggy scenes, whose goal and results are illustrated in Fig. 1. The task is challenging since fog often damages visibility of images seriously, leading to substantial performance degradation. Attaching a fog removal network to the front of

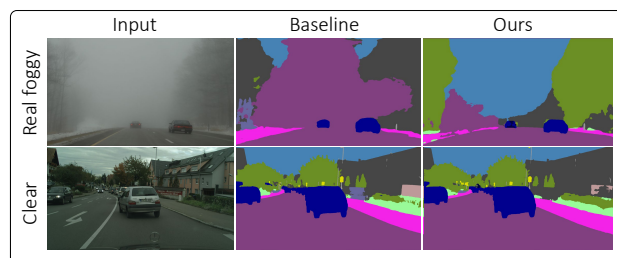


Figure 1. A summary of our results. Predictions of FIFO are accurate for both clear weather and real foggy images while the baseline, an ordinary segmentation model trained on clear weather images, fails to handle foggy images.

an existing model is not always useful for mitigating this issue [47, 52] as well as being heavy in computation and memory. The other reason for the difficulty is the absence of fully annotated data for the task. Collecting a large set of foggy scenes is not straightforward since they can be captured under only a specific condition, and it is hard to label them due to their limited visibility.

Existing methods [8, 51, 52] tackle these issues through synthetic foggy image datasets, which are obtained by applying realistic fog effects to fully annotated clear weather images and are used for supervised learning of semantic segmentation. Furthermore, they introduce curriculum learning approaches [8, 51] that gradually adapt a model from light synthetic fog to dense real fog using unlabeled real foggy images additionally. Although these methods have achieved impressive robustness, there remains room for further improvement in that their training strategies are limited to ordinary supervised learning. In addition, the curriculum adaptation demands external modules to control the fog density of real foggy images in training, and tends to make the final model biased to foggy scenes; it thus requires extra computation and additional hyper-parameters in training, and often degrades performance on clear images.

To resolve the above issues, we propose a new method that learns Fog-Invariant features for FOGgy scene segmentation, dubbed FIFO. Its overall pipeline is illustrated in Fig. 2. FIFO considers the fog condition of an image as its

*This work was done while Taeyoung Son was in POSTECH.

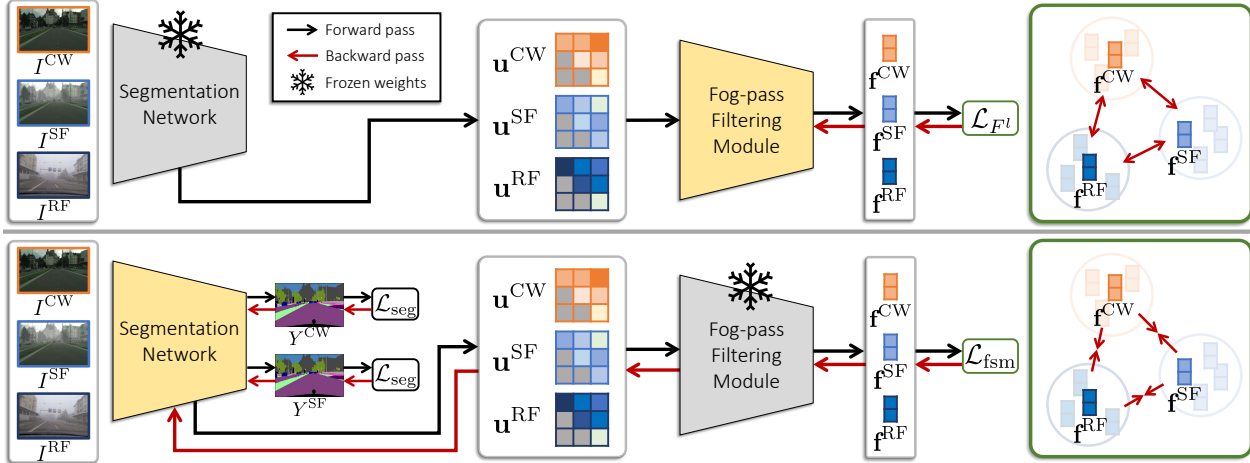


Figure 2. Overall pipeline of FIFO. For each iteration of training, the fog-pass filtering module and the segmentation network are updated alternately. (*top*) Given Gram matrices of feature maps of the segmentation network as input, the fog-pass filtering module learns to extract fog factors so that fog conditions of images are discriminated by their fog factors. (*bottom*) The segmentation network is trained by reducing the gap between fog factors of images with different fog conditions as well as by the segmentation loss. Note that the fog-pass filters are auxiliary modules used only in training.

style, ideally independent of its content, and aims to learn a segmentation model insensitive to fog style variation of input image. To this end, we first define three different domains of training images, *i.e.*, clear weather (CW), synthetic fog (SF), and real fog (RF), where images of the first two domains are labeled while those of the last one are not. FIFO then encourages the segmentation network to close the style discrepancy between different fog domains in feature spaces so that it learns fog-invariant features.

Then the success of FIFO depends heavily on the quality of the fog style representation. Unfortunately, existing style representation schemes [17, 60] are not desirable for our task since they are manually designed to capture the holistic style of an image that is affected also by factors other than fog (*e.g.*, when and where the image was taken) and even the content of the image [6]; the direct use of these neural styles thus introduce side-effects like content alteration and result in suboptimal solutions consequently.

To address this issue, we present *fog-pass filters*, learnable modules that take an ordinary neural style—the Gram matrix of a feature map [17] as input and extract only a fog-relevant information from the style precisely in the form of embedding vectors, called *fog factors*. In particular, they learn to draw fog factors of the same domain together and hold those of different domains apart so that they discriminate fog conditions of input images through their fog factors. The segmentation model is in turn encouraged to reduce the gap between fog factors of images from different domains during training. The alternating optimization of the fog-pass filter and the segmentation network gradually closes the fog style gap between different domains and eventually leads to fog-invariant features of the segmentation network. Note that the fog-pass filters are auxiliary

modules for training only, thus not required in testing.

FIFO has advantages over the previous work [8, 51, 52] in terms of both simplicity in training and efficacy in testing. Unlike the previous work, FIFO does not need to control fog density levels of synthetic and real foggy images in training, thus allowing end-to-end learning of a segmentation model with fewer hyper-parameters. More importantly, for the same reason, it demands no extra module for estimating and manipulating fog density of real foggy images in training. Regarding the effectiveness, FIFO clearly outperforms all existing records and improves performance on both foggy and clear weather domains, while existing methods often degrade performance on clear scenes.

2. Related Work

Semantic Foggy Scene Segmentation. Previous work [8, 51, 52] has developed fog simulators that are applied to clear images with full annotations to obtain labeled synthetic foggy images. Since supervised learning on the synthetic data limits performance due to the visual gap between synthetic and real foggy images, recent methods [8, 51] further employ curriculum learning to gradually adapt a model from light synthetic fog to dense real fog. However, the curriculum adaptation often degrades performance on clear weather images and demands extra modules to control density levels of real foggy images during training.

Image Dehazing. Fog damages visibility of image, and accordingly, degrades visual recognition performance substantially. Numerous dehazing algorithms have been proposed so far to restore latent clean image from foggy input [2, 5, 11, 12, 21, 36, 40, 67]. However, they are usually too heavy in computation to be attached to the front of recognition models. Further, recent studies [47, 52] suggest

that most dehazing models do not help improve recognition performance on foggy scenes. Hence, instead of having a separate dehazing module, our work learns a segmentation model whose features are invariant to fog condition.

Robustness. Robustness of recognition models against adverse conditions has been actively investigated due to its importance in real-world applications [13, 19, 23, 25, 27, 31, 49, 53, 66], and a variety of methods have been proposed to improve robustness [6, 35, 54–56, 64]. FIFO shares a similar idea with RobustNet [6], which regards adverse condition of input as its style and removes the effect of photometric transform of input from a neural style representation of a model so that the model becomes invariant to the transform. Compared to RobustNet, FIFO more explicitly quantifies the effect of adverse condition through a learnable module (*i.e.*, fog-pass filter), so is able to more precisely manipulate the adverse effects during training.

Style Transfer. Neural style transfer has been studied to comprehend the style of an image apart from its content [1, 10, 18, 28, 37, 57, 59]. In particular, a seminal work [17] studied the Gram matrix of a feature map as a neural style representation and showed that the style of an image can be transferred to another by approximating its Gram matrix; the efficacy of Gram matrix has been proven further in later studies [32, 41]. Also, Li *et al.* [38] proved that matching Gram matrices is useful for domain adaptation since it is equivalent to minimizing maximum mean discrepancy between domains. However, we found that Gram matrix is not appropriate as-is for quantifying the effect of fog in our task since it is affected other style factors or even content [6] as well as fog. We thus introduce the fog-pass filter to precisely capture only a fog-relevant factor from the Gram matrix of a feature map.

Unsupervised Domain Adaptation (UDA). Our work is also relevant to UDA since both adapt models to an unlabeled target domain. UDA methods for semantic segmentation can be categorized by the level at which adaptation is performed: Input-level [26, 33, 45, 48], feature-level [42, 58, 63], and output-level [43, 68, 69]. FIFO is related in particular to the feature-level adaptation that learns domain-invariant features. Most of existing methods in this category [42, 58, 63] train a discriminator together with a segmentation model so that the discriminator maximizes a discrepancy between source and target domains while the segmentation model learns to minimize the discrepancy. FIFO shares a similar idea with these methods, but as will be demonstrated, closing the gap between fog factors in FIFO is more effective than fooling a fog domain classifier in fog-invariant feature learning.

3. Configuration of Training Data

Training images for FIFO are categorized into three different domains according to their fog types: clear weather

(CW), synthetic fog (SF), and real fog (RF). For CW images, we adopt the Cityscapes dataset [7], which is fully annotated for supervised learning of semantic segmentation. Meanwhile, as SF images, we utilize the Foggy Cityscapes-DBF dataset [51], which is constructed by simulating realistic fog effects on images of the Cityscapes dataset, thus also fully annotated. Finally, RF images are taken from the Foggy Zurich dataset [51], which is a collection of unlabeled foggy scenes captured in the real world.

Note that the way FIFO uses the two foggy image datasets is different from that of existing methods [8, 51]. First, in the Foggy Cityscapes-DBF dataset, FIFO fixes the density level of synthetic fog by a single value (*i.e.*, the attenuation coefficient $\beta = 0.005$) and utilizes the entire dataset. On the other hand, the previous work adopts only a refined, high-quality subset of the dataset and varies the fog level during training for the curriculum learning. Second, FIFO utilizes the Foggy Zurich dataset as a whole, whereas the previous work divides it into multiple sets of different density levels using extra modules that estimate the fog density of the images. These differences allow the pipeline of FIFO to be more straightforward and concise.

4. Proposed Method

The training procedure for the fog-pass filters and the segmentation network is illustrated in Fig. 2 and their structures are presented in Fig. 3. The segmentation network is first pretrained on the Cityscapes dataset [7], and fog-pass filters initialized randomly are attached to different feature maps of the network. Thereafter, on the dataset of the three fog domains introduced in Sec. 3, the two parts of FIFO are trained alternately per mini-batch, except for the first 5K iterations where the fog-pass filters are solely trained to avoid cold-start. Note that the fog-pass filters are used only in training for fog-invariant feature learning of the segmentation network. In other words, FIFO imposes no additional inference-time complexity since it employs only the segmentation network for testing.

To construct a mini-batch, we randomly sample the same number of images from CW and RF domains, and choose SF counterparts of the sampled CW images. Given such a mini-batch, the fog-pass filters are learned to draw fog factors of the same fog domain together and hold those of different domains apart so that they discriminate input according to its fog domain. On the other hand, the segmentation network is optimized for closing the distance between fog factors of different domains as well as for minimizing the ordinary segmentation loss. This alternating optimization closes the fog style gap between different domains precisely, leading to fog-invariant features.

The remaining part of this section first gives details of training the fog-pass filtering modules and the segmentation network, then empirically verifies the key ideas of FIFO.

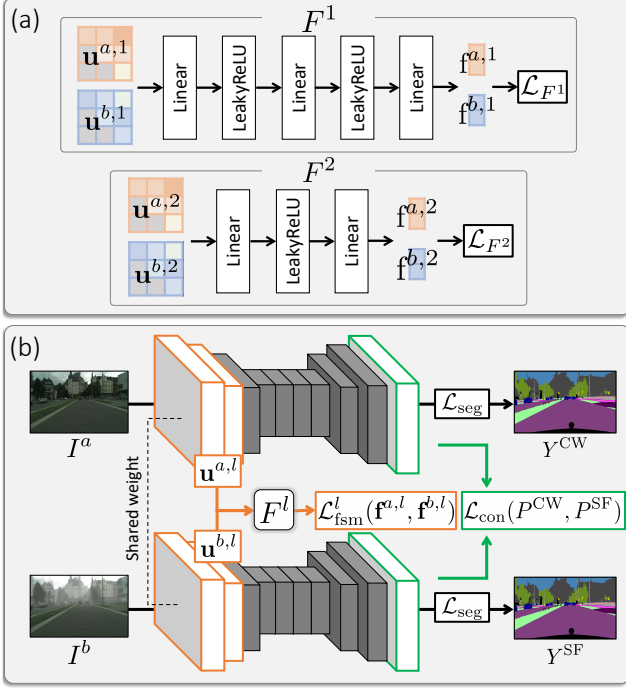


Figure 3. A schematic of FIFO. (a) The fog-pass filtering modules. Each of them takes as input the upper triangular part of a Gram matrix and returns a fog factor. The loss pulls or pushes a pair of fog factors according to the equivalence of their fog conditions. (b) Training of the segmentation network with the frozen fog-pass filters. Given a pair of images with different fog conditions as input, the network is trained by closing the gap between their fog factors and that between their segmentation predictions as well as by the ordinary segmentation loss.

4.1. Fog-pass Filtering Modules

Instead of a raw feature map of the segmentation network, a fog-pass filtering module takes a holistic style representation of the feature map as input in order to focus more on the style of image by filtering out most of its content information. In this context, the style representation can be considered as a hardwired layer [30] that encodes our prior knowledge. We in particular adopt the Gram matrix of the feature map [17] as the style representation as it provides richer style information than other methods, *e.g.*, channel-wise feature statistics [60]. The Gram matrix, denoted by $\mathbf{G} \in \mathbb{R}^{c \times c}$, captures correlations between c channels of its input feature map. The (i, j) element of \mathbf{G} indicates the correlation between i^{th} and j^{th} feature channels and is computed by $\mathbf{G}_{i,j} = \mathbf{a}_i^T \mathbf{a}_j$, where \mathbf{a}_i is the vector form of the i^{th} channel of the input feature map. Specifically, since the Gram matrix is symmetric, the vector form of only the upper triangular part of the matrix is used as input to the fog-pass filtering module.

Let I^a and I^b be a pair of images from the mini-batch and F^l denote the fog-pass filter attached to the l^{th} layer

of the segmentation network. Then the fog factors of the two images are computed by $\mathbf{f}^{a,l} = F^l(\mathbf{u}^{a,l})$ and $\mathbf{f}^{b,l} = F^l(\mathbf{u}^{b,l})$, respectively, where $\mathbf{u}^{a,l}$ and $\mathbf{u}^{b,l}$ denote the vectorized upper triangular parts of the Gram matrices computed from their l^{th} intermediate feature maps. The role of the fog-pass filter is to inform the segmentation network how I^a and I^b are different in terms of fog condition through $\mathbf{f}^{a,l}$ and $\mathbf{f}^{b,l}$. For this purpose, the fog-pass filter learns a space of fog factors where those of the same fog domain are grouped closely together and those of different domains are far from each other. Given the set of every image pair \mathcal{P} in the mini-batch, the loss function for F^l is designed as follows:

$$\mathcal{L}_{F^l} = \sum_{(a,b) \in \mathcal{P}} \left\{ (1 - \mathbb{I}(a,b)) \left[m - d(\mathbf{f}^{a,l}, \mathbf{f}^{b,l}) \right]_+^2 + \mathbb{I}(a,b) \left[d(\mathbf{f}^{a,l}, \mathbf{f}^{b,l}) - m \right]_+^2 \right\}, \quad (1)$$

where $d(\cdot)$ is the cosine distance, m is a margin, and $\mathbb{I}(a,b)$ denotes the indicator function that returns 1 if I^a and I^b are of the same fog domain and 0 otherwise.

4.2. Segmentation Network

The segmentation network is trained using three different objectives, which are designed for semantic segmentation, fog-invariant feature learning, and consistent prediction regardless of fog condition of input, respectively. We elaborate on each of the loss functions below.

Segmentation Loss. For learning semantic segmentation, we apply the pixel-wise cross-entropy loss to individual images. To be specific, the loss is given by

$$\mathcal{L}_{\text{seg}}(\mathbf{P}, \mathbf{Y}) = -\frac{1}{n} \sum_i \sum_j \mathbf{Y}_{i,j} \log \mathbf{P}_{i,j}, \quad (2)$$

where $\mathbf{P}_{i,j} \in \mathbb{R}$ and $\mathbf{Y}_{i,j} \in \{0, 1\}$ denote the predicted score and groundtruth label of class j at pixel i , respectively, while n is the number of pixels.

Fog Style Matching Loss. Given a pair of images from different fog domains, the segmentation network learns fog-invariant features that close the distance between their fog factors. To this end, the second loss matches the two fog factors given by the frozen fog-pass filters. Let $\mathbf{f}_i^{a,l}$ and $\mathbf{f}_i^{b,l}$ be the fog factors of the images computed by the fog-pass filter F^l . Then the loss is given by

$$\mathcal{L}_{\text{fsm}}^l(\mathbf{f}^{a,l}, \mathbf{f}^{b,l}) = \frac{1}{4d_l^2 n_l^2} \sum_{i=1}^{d_l} \left(\mathbf{f}_i^{a,l} - \mathbf{f}_i^{b,l} \right)^2, \quad (3)$$

where d_l and n_l denote the dimension of their fog factors and the spatial size of the l^{th} feature map, respectively.

Prediction Consistency Loss. A CW image and its SF counterpart have exactly the same semantic layout. By forcing predictions for these images being identical, we can

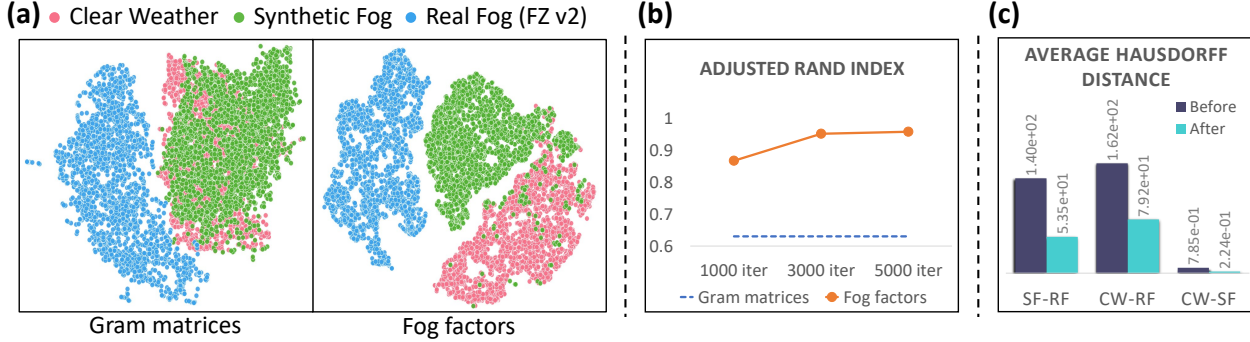


Figure 4. Empirical analysis on the impact of FIFO. (a) 2D visualization of distributions of Gram matrices and their fog factors. (b) Comparison between the quality of k -means clustering of the Gram matrices and that of the corresponding fog factors in adjusted Rand index. (c) The fog-style gap between different domains before and after training with FIFO, where the gap is measured by the average Hausdorff distance between two sets of fog factors.

align the CW and SF domains more aggressively in the learned representation. Hence, only for CW and SF images of the same origin, we encourage the model to predict the same segmentation map. Let $\mathbf{P}_i^{\text{CW}} \in \mathbb{R}^c$ and $\mathbf{P}_i^{\text{SF}} \in \mathbb{R}^c$ denote their class probability vectors predicted by the segmentation model for pixel i , where c is the number of classes. The third loss is designed to force the consistency between \mathbf{P}_i^{CW} and \mathbf{P}_i^{SF} for all pixels, and is given by

$$\mathcal{L}_{\text{con}}(\mathbf{P}^{\text{CW}}, \mathbf{P}^{\text{SF}}) = \sum_i \text{KLdiv}(\mathbf{P}_i^{\text{CW}}, \mathbf{P}_i^{\text{SF}}), \quad (4)$$

where $\text{KLdiv}(\cdot, \cdot)$ is the Kullback–Leibler divergence. This loss shares the same goal with the fog style matching loss in Eq. (3), but more strongly forces fog-invariance in the prediction level via a small number of CW–SF pairs. Also, it is complementary to the segmentation loss in Eq. (2) since the probability distributions in Eq. (4) provide information beyond the class labels used by the segmentation loss.

Training Strategy. Given a mini-batch, the same number of image pairs are sampled from each of the three different domain pairs, *i.e.*, CW–SF, CW–RF, and SF–RF. Note that images of each CW–SF pair are of the same semantic layout so that the prediction consistency loss is applied to them. For CW–SF pairs, the segmentation network is trained by minimizing

$$\begin{aligned} \mathcal{L}_S^{\text{CW-SF}} &= \mathcal{L}_{\text{seg}}(\mathbf{P}^{\text{CW}}, \mathbf{Y}^{\text{CW}}) + \mathcal{L}_{\text{seg}}(\mathbf{P}^{\text{SF}}, \mathbf{Y}^{\text{SF}}) \\ &+ \lambda_{\text{fsm}} \sum_l \mathcal{L}_{\text{fsm}}^l(\mathbf{f}^{\text{CW},l}, \mathbf{f}^{\text{SF},l}) + \lambda_{\text{con}} \mathcal{L}_{\text{con}}(\mathbf{P}^{\text{CW}}, \mathbf{P}^{\text{SF}}), \end{aligned} \quad (5)$$

where λ_{fsm} and λ_{con} are balancing hyper-parameters and $\mathbf{Y}^{\text{CW}} = \mathbf{Y}^{\text{SF}}$. On the other hand, for the other pairs of input domains including RF, the loss consists of the segmentation and fog style matching terms only, and is given by

$$\mathcal{L}_S^{\text{D-RF}} = \mathcal{L}_{\text{seg}}(\mathbf{P}^{\text{D}}, \mathbf{Y}^{\text{D}}) + \lambda_{\text{fsm}} \sum_l \mathcal{L}_{\text{fsm}}^l(\mathbf{f}^{\text{D},l}, \mathbf{f}^{\text{RF},l}), \quad (6)$$

where $\text{D} \in \{\text{CW}, \text{SF}\}$. Note that \mathcal{L}_{seg} is not applied to the

prediction for real foggy image \mathbf{P}^{RF} due to the absence of its segmentation label.

4.3. Empirical Verification

Impact of Fog-pass Filtering Modules. To justify the use of the fog-pass filters, we compare Gram matrices and their fog factors computed by a fog-pass filter in how well they disentangle the fog condition and the other aspects of an image. To this end, we examine distributions of Gram matrices and fog factors of CW, SF, and RF. To be specific, training images of the Cityscapes dataset [7], their synthetic foggy counterparts given by the fog simulator [8] with the attenuation coefficient $\beta = 0.005$, and those of the Foggy Zurich dataset [51] are adopted as CW, SF, and RF images, respectively; their Gram matrices are computed from ResBlock1 outputs of RefineNet-lw [46] pre-trained on the Cityscapes, and the corresponding fog factors are computed from the Gram matrices through the fog-pass filtering module. Fig. 4(a) presents t -SNE visualization [61] of the distributions, in which Gram matrices of CW and SF largely overlap each other while fog factors are well separated according to their fog domains. This result suggests that Gram matrices are affected substantially by image content while fog factors represent only fog-relevant information as desired. The same trend is observed in Fig. 4(b) that quantitatively evaluates the quality of k -means clusters [20] of the Gram matrices and fog factors via adjusted Rand index [29].

Fog-invariance Learned by FIFO. To investigate the impact of FIFO on fog-invariance learning, we first demonstrate that FIFO effectively reduces the gap between fog domains in the space of fog factors. To this end, each domain is represented as the set of fog factors of its images, and the gap between a pair of domains is measured by the average Hausdorff distance [9] between such sets of the domains before and after training with FIFO. Fig. 4(c) shows that FIFO closes the fog-style gap in all three domain pairs. The impact is also verified qualitatively through images reconstructed from intermediate features of the segmentation



Figure 5. Images reconstructed by the baseline, a variant of FIFO closing the gap between Gram matrices, and FIFO.

network trained by FIFO.¹ As a reconstruction model, we adopt RefineNet-lw with two additional upsampling layers; its decoder is first trained to reconstruct images while freezing its encoder pretrained on the Cityscapes dataset, then the encoder is replaced with that of the segmentation network trained by FIFO. Also, for comparisons, we reconstruct images using a baseline (*i.e.*, training only on the Cityscapes) and a variant of FIFO with no fog-pass filter (*i.e.*, training by directly reducing the gap between Gram matrices) in the same manner. Fig. 5 presents examples of the reconstructed images, which demonstrate that FIFO allows to sharpen images effectively, well emphasize object boundaries in particular, and make fewer artifacts.

5. Experiments

5.1. Implementation Details

Network Architecture. We adopt RefineNet-lw [46] with ResNet-101 [22] backbone as our segmentation network. Two fog-pass filtering modules are then respectively attached to the outputs of Conv1 and ResBlock1 layers of the segmentation network. As illustrated in Fig. 3(a), the two modules are implemented by multi-layer perceptrons with leaky ReLU activation functions [44].

Optimization and Hyper-parameters. The segmentation network is trained by SGD with a momentum of 0.9 and the initial learning rate of $6e-4$ for the encoder and $6e-3$ for the decoder; both learning rates are decreased by polynomial decay with a power of 0.5. The two fog-pass filtering modules are trained by Adamax [34] with initial learning rates of $5e-4$ (Conv1) and $1e-3$ (ResBlock1), respectively; the dimensionality of fog factors is set to 64. Each mini-batch is constructed by sampling 4 images from each fog domain, thus its size is 12. During training, images are resized, cropped to 600×600 , and flipped horizontally at random. Finally, the hyper-parameters λ_{fsm} , λ_{con} , and m are set to $5e-8$, $1e-4$, and 0.1, respectively.

¹We conduct this experiment only for qualitative analysis on the effect of FIFO. Note that FIFO directly learns fog-invariant features while bypassing explicit fog removal of input image, thus does not reconstruct images at all in both training and testing.

5.2. Datasets for Evaluation

FIFO is evaluated and compared with previous work on three real foggy datasets: Foggy Zurich (FZ) test v2 [51], Foggy Driving (FD) [52], and Foggy Driving Dense (FDD) [51]; FDD is a subset of FD. Images of these datasets depict various fog densities and are fully annotated. Also, they share the same class set with the Cityscapes dataset as described in [8]. We further apply FIFO and previous work to an unseen clear weather dataset, Cityscapes lindau 40 introduced in [8], to evaluate their performance on clear weather scenes as well.

5.3. Quantitative Analysis

Quantitative results of FIFO and previous arts are summarized in Table 1. As shown in the table, FIFO largely outperforms CMAda3+ [8], the current best performing model based on RefineNet backbone [39], on all the three foggy image datasets. These results indicate that our method of closing the fog style gap is superior to the curriculum adaptation on real images with various fog densities.

We also validate the performance of the models on the clear weather dataset. In this experiment, the accuracy of CMAda3+ drops substantially; we suspect this is a side effect of the curriculum adaptation, which may lead to overfitting to foggy scenes due to catastrophic forgetting. On the other hand, FIFO enhances performance on clear weather, probably because of the data augmentation effects of using the three domains altogether during training.

5.4. Qualitative Results

FIFO is qualitatively compared with two other methods. One is RefineNet-lw trained only on the Cityscapes dataset, which we call baseline. The other is a reduced version of FIFO using no fog-pass filtering module; this method learns the segmentation network while closing the gap between Gram matrices from different domains. Qualitative examples of their predictions are presented in Fig. 6. The baseline yields poor results in most cases. The reduced version of FIFO outperforms the baseline, especially for car, road, and vegetation classes, but FIFO using the fog-pass filtering modules clearly demonstrates the best segmentation results.

Method	Clear-weather Cityscapes [7]	Synthetic & Real fog SDBF [51] GoPro [51]	FZ test v2 [51] mIoU (%)	FDD [51] mIoU (%)	FD [52] mIoU (%)	Cityscapes lindau 40 [7] mIoU (%)
RefineNet [39]	✓		34.6	35.8	44.3	67.2
RefineNet-lw [46]	✓		28.5	35.9	43.6	63.8
AdSegNet [58]	✓	✓	25.0	15.8	29.7	-
AdvEnt [62]	✓	✓	39.7	41.7	46.9	61.7
FDA [65]	✓	✓	22.2	29.8	21.8	39.3
DANN [16]	✓	✓	43.1	41.4	46.0	60.1
DANN-Gram	✓	✓	43.4	43.3	47.3	67.1
CMAda2+ [8]	✓	✓	43.4	40.1	49.9	-
CMAda3+ [8]	✓	✓	46.8	43.0	49.8	59.6
FIFO	✓	✓	48.4	48.9	50.7	64.8

Table 1. Quantitative results in mean intersection over union (mIoU) on three real foggy datasets—Foggy Zurich (FZ) test v2, Foggy Driving Dense (FDD), Foggy Driving (FD), and a clear weather dataset—Cityscapes lindau 40.

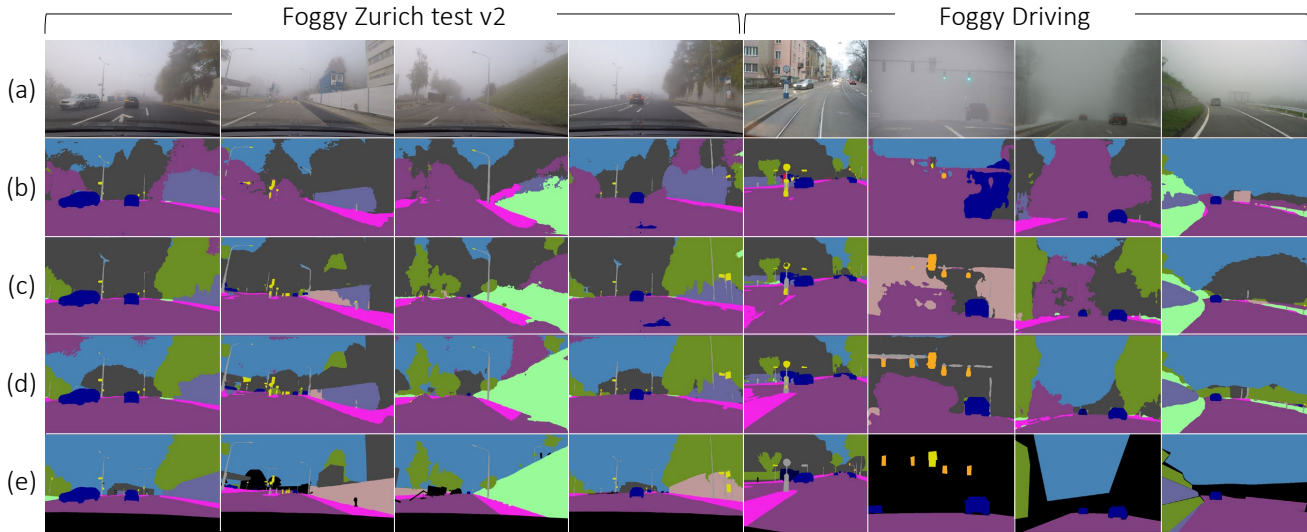


Figure 6. Qualitative results on the real foggy datasets. (a) Input images. (b) Baseline. (c) FIFO without the fog-pass filtering. (d) FIFO. (e) Groundtruth.

5.5. Comparison with UDA

The task of FIFO is ostensibly identical to that of unsupervised domain adaptation (UDA) since both of them adapt models to an unlabeled target domain. Hence, one may wonder how well UDA models work for semantic foggy scene segmentation under the setting of FIFO. In this context, FIFO is compared with multiple UDA methods based on various levels of adaptation: FDA [65] for input-level adaptation, AdSegNet [58] and AdvEnt [62] for output-level adaptation, and DANN [16] for feature-level adaptation. We also evaluate a variant of DANN whose domain classifier takes as input a Gram matrix, instead of a raw feature map, like FIFO; this variant is denoted by DANN-Gram. The UDA models are trained using the same datasets, *i.e.*, CW, SF, and RF. The input- and output-level adaptation models are pretrained on CW and further trained using SF while adapting to RF. DANN and DANN-Gram are trained using CW, SF, and RF at once like FIFO: Their discriminators are optimized to maximize the discrepancy

of fog domains like our fog-pass filters while their segmentation networks are learned to minimize the discrepancy.

The performance of these UDA methods is reported in Table 1. As shown in the table, the UDA models are all inferior to FIFO and CMAda3+, which suggests that, even though it is apparently similar to UDA for semantic segmentation, semantic foggy scene segmentation is of a different nature and has its own challenges. In the typical UDA setting, each source and target domain has its own style, by which it can be defined. However, the style of a foggy scene is not determined only by its fog condition, rather is the result of the chemical combination of fog and other style factors of the scene. The UDA methods that consider SF and RF as domains with unique styles are thus not well suited to the task. Moreover, fog substantially damages visibility, and enlarges intra-domain variations since the effect of fog varies significantly according to the 3D configuration of the scene [8]. Due to these differences and challenges, foggy scene segmentation demands dedicated solutions like FIFO.

CW-SF	CW-RF	SF-RF	FZ	FDD	FD	CW
✓			43.7	38.6	46.1	67.6
	✓		37.7	40.3	47.2	66.0
		✓	39.3	42.8	49.7	61.6
✓	✓		49.7	46.0	49.9	65.8
✓		✓	46.0	47.6	50.0	62.3
	✓	✓	47.4	38.2	47.0	64.3
✓	✓	✓	48.4	48.9	50.7	64.8

Table 2. Analysis on the impact of domain pairs. CW, SF and RF denote clear weather, synthetic fog, and real fog, respectively.

Method	FZ	FDD	FD	β	FZ	FDD	FD
Baseline	28.5	35.9	43.6	0	37.7	40.3	47.2
FIFO w/o \mathcal{L}_{fsm}	31.7	38.5	45.1	0.0025	45.5	40.4	45.0
FIFO w/o \mathcal{L}_{con}	41.6	45.4	48.9	0.01	42.4	45.4	50.0
FIFO w/ Gram	41.3	43.8	49.1	0.02	42.9	42.5	48.6
FIFO	48.4	48.9	50.7	0.005	48.4	48.9	50.7

Table 3. Analysis on the impact of the fog style matching loss, the prediction consistency loss, and the fog-pass filtering modules.

Table 4. Analysis on the impact of synthetic fog densities. $\beta = 0$ denotes a model trained with CW and RF.

5.6. Ablation Study

We conduct extensive experiments while varying domain pairs of FIFO to investigate their effects. Table 2 summarizes the results. We found that using more pairs in general boosts performance, which suggests that all the three pairs contribute to performance on real foggy images.

We also investigate contributions of the fog style matching loss \mathcal{L}_{fsm} , the prediction consistency loss \mathcal{L}_{con} , and the fog-pass filtering modules to the performance. Table 3 compares FIFO with its variants with and without \mathcal{L}_{fsm} , \mathcal{L}_{con} , and the fog-pass filters in terms of segmentation quality on real foggy images. Note that FIFO w/o \mathcal{L}_{fsm} is trained by the segmentation and prediction consistency losses only, dropping the fog style matching loss, while FIFO w/ Gram uses Gram matrices instead of fog factors when matching fog styles. The results in the tables suggest that all the losses and the fog-pass filtering contribute to the performance on all the three real foggy datasets, but the impact of the fog style matching is substantially larger than the others. Also, the gap between FIFO and FIFO w/ Gram demonstrates the superiority of fog factors over Gram matrices, which justifies the use of the fog-pass filters.

In addition, we demonstrate the effect of our chosen value of β , which is the attenuation coefficient used for generating synthetic fog [52]. Note that our method exploits a single value of β to resolve the issue of the curriculum adaptation [8, 51]. Table 4 summarizes the performance of variants of FIFO trained using different values of β ; the optimal value for β is 0.005.

Finally, we examine the impact of the layers to which

	FZ	FDD	FD	Method	RC
C1	45.3	41.0	48.3	†Baseline [24]	59.0
R1	45.1	39.1	47.0	†Cutmix [15]	61.9
C1+R1 (Ours)	48.4	48.9	50.7	†LP-BNN [14]	60.7
				†Superpixel-mix [13]	61.9
				Baseline [46]	57.6
				FIFO	67.6

Table 5. Analysis on the layers to which the fog style matching loss is applied. C1 and R1 indicate the output of the first convolution layer and that of the first residual block, respectively.

Table 6. Quantitative results on Rainy Cityscapes (RC).

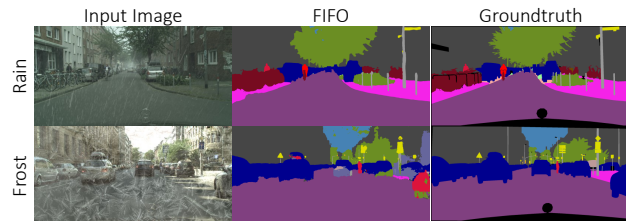


Figure 7. Qualitative results under rain and frost conditions.

the fog style matching loss \mathcal{L}_{fsm} is applied. Specifically, \mathcal{L}_{fsm} is applied to the output of the first convolutional layer (C1), the first residual block (R1), or both of them (C1+R1). As summarized in Table 5, the performance improves as more feature maps are affected by the loss. Overall, C1+R1, which is our final model, shows the best performance.

5.7. Generalization to Other Weather Conditions

We investigate the generalization ability of FIFO on other weather conditions, *rain* and *frost*. To this end, RefineNet-lw trained by FIFO in Sec. 4 is evaluated *as-is* on rainy [27] and frosty [23] versions of the Cityscapes dataset. On the rainy Cityscapes dataset, our model is compared with existing methods reported in [13] that aim at improving robustness using CW images only. As summarized in Table 6, our model largely outperforms the previous work based on a stronger segmentation network (*i.e.*, DeepLab v3+ [4], indicated by † in the table). Fig. 7 demonstrates that FIFO generalizes to the frosty images also. More results can be found in the supplementary material.

6. Conclusion

We have presented FIFO, a new approach to learning fog-invariant features for foggy scene segmentation. It precisely quantifies the fog style of an image through the fog-pass filtering modules and learns a segmentation network for closing the gap between images of different fog conditions in the fog style space. FIFO outperforms previous arts without sacrificing performance on clear weather images. Moreover, unlike the current best-performing method, it enables end-to-end learning and demands no extra module nor human intervention for training.

Acknowledgement: This work was supported by Samsung Research Funding & Incubation Center of Samsung Electronics under Project Number SRFC-IT1801-05.

References

- [1] Guillaume Berger and Roland Memisevic. Incorporating long-range consistency in cnn-based texture generation. In *Proc. International Conference on Learning Representations (ICLR)*, 2016. 3
- [2] Dana Berman, Tali Treibitz, and Shai Avidan. Non-local image dehazing. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [3] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 8
- [5] Wei-Ting Chen, Jian-Jiun Ding, and Sy-Yen Kuo. Pms-net: Robust haze removal based on patch map for single images. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [6] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3
- [7] 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. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3, 5, 7
- [8] 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 (IJCV)*, 2020. 1, 2, 3, 5, 6, 7, 8
- [9] M-P Dubuisson and Anil K Jain. A modified hausdorff distance for object matching. In *Proc. International Conference on Pattern Recognition (ICPR)*, 1994. 5
- [10] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. In *Proc. International Conference on Learning Representations (ICLR)*, 2016. 3
- [11] Raanan Fattal. Single image dehazing. *ACM transactions on graphics (TOG)*, 2008. 2
- [12] Raanan Fattal. Dehazing using color-lines. *ACM transactions on graphics (TOG)*, 2014. 2
- [13] Gianni Franchi, Nacim Belkhir, Mai Lan Ha, yufei hu, Andrei Bursuc, Volker Blanz, and Angela Yao. Robust semantic segmentation with superpixel-mix. In *Proc. British Machine Vision Conference (BMVC)*, 2021. 3, 8
- [14] Gianni Franchi, Andrei Bursuc, Emanuel Aldea, Séverine Dubuisson, and Isabelle Bloch. Encoding the latent posterior of bayesian neural networks for uncertainty quantification. *arXiv preprint arXiv:2012.02818*, 2020. 8
- [15] Geoff French, Samuli Laine, Timo Aila, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, varied perturbations. In *Proc. British Machine Vision Conference (BMVC)*, 2020. 8
- [16] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research (JMLR)*, 2016. 7
- [17] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 3, 4
- [18] Leon A Gatys, Alexander S Ecker, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. Controlling perceptual factors in neural style transfer. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [19] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proc. International Conference on Learning Representations (ICLR)*, 2015. 3
- [20] John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 1979. 5
- [21] Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2010. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [23] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *Proc. International Conference on Learning Representations (ICLR)*, 2019. 3, 8
- [24] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Proc. International Conference on Learning Representations (ICLR)*, 2017. 8
- [25] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *Proc. International Conference on Learning Representations (ICLR)*, 2018. 3
- [26] 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 *Proc. International Conference on Machine Learning (ICML)*, 2018. 3
- [27] Xiaowei Hu, Chi-Wing Fu, Lei Zhu, and Pheng-Ann Heng. Depth-attentional features for single-image rain removal. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3, 8
- [28] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017. 3

- [29] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of classification*, 1985. 5
- [30] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. In *Proc. International Conference on Machine Learning (ICML)*, 2010. 4
- [31] Jiongchao Jin, Arezou Fatemi, Wallace Lira, Fenggen Yu, Biao Leng, Rui Ma, Ali Mahdavi-Amiri, and Hao Zhang. Raidar: A rich annotated image dataset of rainy street scenes. *arXiv preprint arXiv:2104.04606*, 2021. 3
- [32] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Proc. European Conference on Computer Vision (ECCV)*, 2016. 3
- [33] Guoliang Kang, Yunchao Wei, Yi Yang, Yueting Zhuang, and Alexander G Hauptmann. Pixel-level cycle association: A new perspective for domain adaptive semantic segmentation. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. International Conference on Learning Representations (ICLR)*, 2015. 6
- [35] S. Lee, J. Kim, J. S. Yoon, S. Shin, O. Bailo, N. Kim, T. H. Lee, H. S. Hong, S. H. Han, and I. S. Kweon. Vpnet: Vanishing point guided network for lane and road marking detection and recognition. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017. 3
- [36] Boyi Li, Xiulian Peng, Zhangyang Wang, Jizheng Xu, and Dan Feng. Aod-net: All-in-one dehazing network. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [37] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2017. 3
- [38] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. *arXiv preprint arXiv:1701.01036*, 2017. 3
- [39] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6, 7
- [40] Xiaohong Liu, Yongrui Ma, Zhihao Shi, and Jun Chen. Grid-dehazenet: Attention-based multi-scale network for image dehazing. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [41] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [42] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [43] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [44] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. International Conference on Machine Learning (ICML)*, 2013. 6
- [45] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [46] Vladimir Nekrasov, Chunhua Shen, and Ian Reid. Lightweight refinenet for real-time semantic segmentation. In *Proc. British Machine Vision Conference (BMVC)*, 2018. 5, 6, 7, 8
- [47] Yanting Pei, Yaping Huang, Qi Zou, Yuhang Lu, and Song Wang. Does haze removal help cnn-based image classification? In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 1, 2
- [48] Fabio Pizzati, Raoul de Charette, Michela Zaccaria, and Pietro Cerri. Domain bridge for unpaired image-to-image translation and unsupervised domain adaptation. In *Proc. IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2020. 3
- [49] Horia Porav, Valentina-Nicoleta Musat, Tom Bruls, and Paul Newman. Rainy screens: Collecting rainy datasets, indoors. *arXiv preprint arXiv:2003.04742*, 2020. 3
- [50] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 1
- [51] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 1, 2, 3, 5, 6, 7, 8
- [52] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision (IJCV)*, 2018. 1, 2, 6, 7, 8
- [53] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 3
- [54] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [55] Baifeng Shi, Dinghui Zhang, Qi Dai, Zhanxing Zhu, Yadong Mu, and Jingdong Wang. Informative dropout for robust representation learning: A shape-bias perspective. In *Proc. International Conference on Machine Learning (ICML)*, 2020. 3
- [56] Taeyoung Son, Juwon Kang, Namyup Kim, Sunghyun Cho, and Suha Kwak. Urie: Universal image enhancement for visual recognition in the wild. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 1, 3

- [57] Chunjin Song, Zhijie Wu, Yang Zhou, Minglun Gong, and Hui Huang. Etnet: Error transition network for arbitrary style transfer. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2019. 3
- [58] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3, 7
- [59] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *Proc. International Conference on Machine Learning (ICML)*, 2016. 3
- [60] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. 2, 4
- [61] L.J.P van der Maaten and G.E. Hinton. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research (JMLR)*, 2008. 5
- [62] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 7
- [63] Haoran Wang, Tong Shen, Wei Zhang, Ling-Yu Duan, and Tao Mei. Classes matter: A fine-grained adversarial approach to cross-domain semantic segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 3
- [64] Z. Wang, S. Chang, Y. Yang, D. Liu, and T. S. Huang. Studying very low resolution recognition using deep networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [65] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7
- [66] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steiner, and Gustavo Fernandez Dominguez. Wilddash - creating hazard-aware benchmarks. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 1, 3
- [67] He Zhang and Vishal M Patel. Densely connected pyramid dehazing network. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [68] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Un-supervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 3
- [69] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 3