

WildNet: Learning Domain Generalized Semantic Segmentation from the Wild

Suhyeon Lee Hongje Seong Seongwon Lee Euntai Kim*
 School of Electrical and Electronic Engineering, Yonsei University, Seoul, Korea
 {hyeon93, hjseong, won4113, etkim}@yonsei.ac.kr

Abstract

We present a new domain generalized semantic segmentation network named *WildNet*, which learns domain-generalized features by leveraging a variety of contents and styles from the wild. In domain generalization, the low generalization ability for unseen target domains is clearly due to overfitting to the source domain. To address this problem, previous works have focused on generalizing the domain by removing or diversifying the styles of the source domain. These alleviated overfitting to the source-style but overlooked overfitting to the source-content. In this paper, we propose to diversify both the content and style of the source domain with the help of the wild. Our main idea is for networks to naturally learn domain-generalized semantic information from the wild. To this end, we diversify styles by augmenting source features to resemble wild styles and enable networks to adapt to a variety of styles. Furthermore, we encourage networks to learn class-discriminant features by providing semantic variations borrowed from the wild to source contents in the feature space. Finally, we regularize networks to capture consistent semantic information even when both the content and style of the source domain are extended to the wild. Extensive experiments on five different datasets validate the effectiveness of our *WildNet*, and we significantly outperform state-of-the-art methods. The source code and model are available online: <https://github.com/suhyeonlee/WildNet>.

1. Introduction

Domain generalized semantic segmentation aims to better predict pixel-level semantic labels on multiple unseen target domains while learning only on the source domain. Unfortunately, the domain shift between the source and target domains makes a segmentation model trained on the given source data behave stupidly on the unseen target data, as shown in Fig. 1b. In domain generalization (DG), the low generalization performance for unseen domains is obviously due to *overfitting to the source domain*. Since the

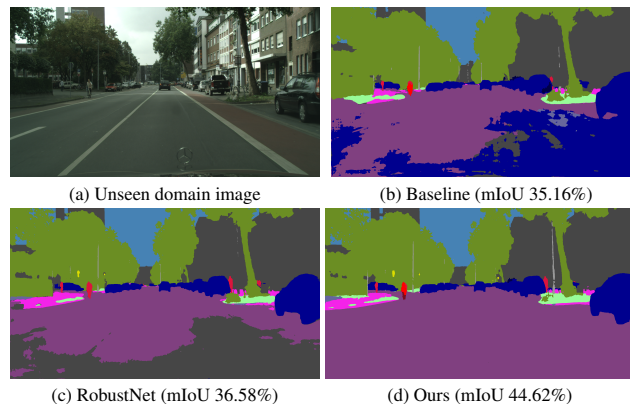


Figure 1. Semantic segmentation results on (a) an unseen domain image. The models are trained on GTAV [46] train set and validated on Cityscapes [10] validation set. (b) Baseline model overfits the source domain and performs poorly with mIoU 35.16% on the unseen target domain. (c) RobustNet [7], a state-of-the-art method, improved mIoU to 36.58% by whitening the style, but still has low generalization capability. (d) Our WildNet achieves superior generalization performance with mIoU 44.62% by learning various styles and contents from the wild. More qualitative results on other datasets are available in the supplementary material.

model cannot see any information about the target domains in the learning process and even unlabeled target images are not provided unlike domain adaptation (DA), it over-learns the statistical distribution of the given source data.

Recently, some studies [7, 29, 41, 42] have proposed learning the domain-generalized content feature by ‘removing’ domain-specific *style* information from the data to prevent overfitting to the source domain. Based on the correlation between the feature’s covariance matrix and style [13, 14], they assumed that only content features would remain if elements of features considered the domain-specific style were whitened [23, 30, 50, 53]. However, since the content and style are not orthogonal, whitening the style may cause a loss of semantic content, which is indispensable for semantic category prediction. As a result, they predict semantic categories from incomplete content features and have difficulty making accurate predictions, as shown in Fig. 1c.

In this paper, we propose a new domain generalized semantic segmentation network called *WildNet*, which learns

*Corresponding author.

the domain-generalized semantic feature by ‘extending’ *both content and style* to the *wild*. Although some previous works [22, 45, 62] utilized various styles from the wild, *e.g.*, ImageNet [11] for real styles and Painter by Numbers [38] for unreal styles, they overlooked that the high generalization ability comes from learning not only various styles but also various contents. In contrast to previous studies, our main idea is to naturally learn domain-generalized semantic information by leveraging a variety of contents and styles from the wild, without forcing whitening on domain-specific styles.

To extend both content and style to the wild, we present four effective learning methods. (i) Based on the relevance of style and feature statistics, *feature stylization* diversifies the style of the source feature by transferring the statistics of the wild feature to the source feature over several layers. (ii) To prevent overfitting to the source contents, we propose *content extension learning* to increase the intra-class content variability in the latent embedding space. Extending content from source to wild helps networks make generalized predictions on unseen contents. (iii) To prevent overfitting to the source style, we propose *style extension learning* to encourage networks to adapt to the various styles extended to the wild. (iv) Finally, *semantic consistency regularization* enables networks to capture consistent semantic information even when both the content and style of the source domain are extended to the wild. With the proposed learning methods, our WildNet learns domain-generalized semantic features by leveraging a variety of contents and styles from the wild. Extensive experiments over multiple domains show that our network achieves superior performance on domain generalization for semantic segmentation.

Our main contributions are as follows:

- We present a novel domain generalized semantic segmentation network named WildNet, which learns domain-generalized semantic features by leveraging a variety of contents and styles from the wild.
- We propose four learning techniques to train domain-generalized networks by extending both the content and style of the source domain to the wild. These enable our model to make reliable predictions on various unseen target domains without training on them.
- Our network achieves superior performance in extensive experiments on domain generalization for semantic segmentation constructed over multiple domains.

2. Related Work

2.1. Domain Adaptation and Generalization

Domain adaptation (DA) aims to increase the performance on the target domain by reducing the domain gap between the source and target domains. In semantic segmentation, DA is exploited to tackle the effort of annotating

pixel-level categories in an image. Most DA methods train networks using the ‘given’ target images via image translation [17, 20, 33, 60, 64], feature alignment [21, 44, 55–57], and self-training [16, 29, 32, 40, 63] strategies. However, it is hard to acquire target images from various environments during the learning process, and efforts to retrain networks are required whenever applying networks to a new target domain.

To overcome these limitations, domain generalization (DG) has recently attracted considerable attention. However, most DG studies have focused on image classification and there are only a few recent studies on semantic segmentation. In this study, we deal with DG for semantic segmentation. Unlike DA, DG does not have access to the target domains during the learning process. To make reliable predictions on various ‘unknown’ target domains, most existing studies focus on whitening [7], normalizing [41], and diversifying [22, 45, 62] styles to avoid overfitting to the style of the source domain. This paper focuses on extending both the content and style of the source domain to the wild [11], enabling networks to learn domain-generalized semantic features from diversified contents and styles.

2.2. Contrastive Learning

Contrastive learning [9, 39] is a strategy that minimizes the distance from a positive sample and maximizes the distance from a negative sample in the embedding space. Recently, He *et al.* [18] used a dynamic dictionary with a queue and Chen *et al.* [3] used two views of the same image as a positive pair to learn visual representations. To diversify a positive pair, a recent work [12] proposed to use the positive’s nearest neighbor in the latent space as a positive. After supervised contrastive learning [27] has been proposed, there are recent efforts to apply contrastive learning to fully- and semi-supervised semantic segmentation [1, 65, 66]. To obtain positive samples, these works perform image augmentation or store features using label information in a memory bank [58]. These enhance class discrimination in the seen source domain but do not guarantee improving class discrimination in various unseen domains. To adapt contrastive learning to DG for semantic segmentation, we propose a learning method using the wild-stylized feature and its closest wild content as positive samples.

2.3. Free ImageNet

Most studies regard ImageNet [11] as free and use it to pre-train networks. The ImageNet pre-trained model is commonly used in various fields such as object detection [34, 68], semantic segmentation [8, 54], panoptic segmentation [6, 35, 59], and video object segmentation [51], and is considered to be the same basis. The ImageNet pre-trained model is also used in most DA and DG for semantic segmentation methods, and ImageNet is used to borrow

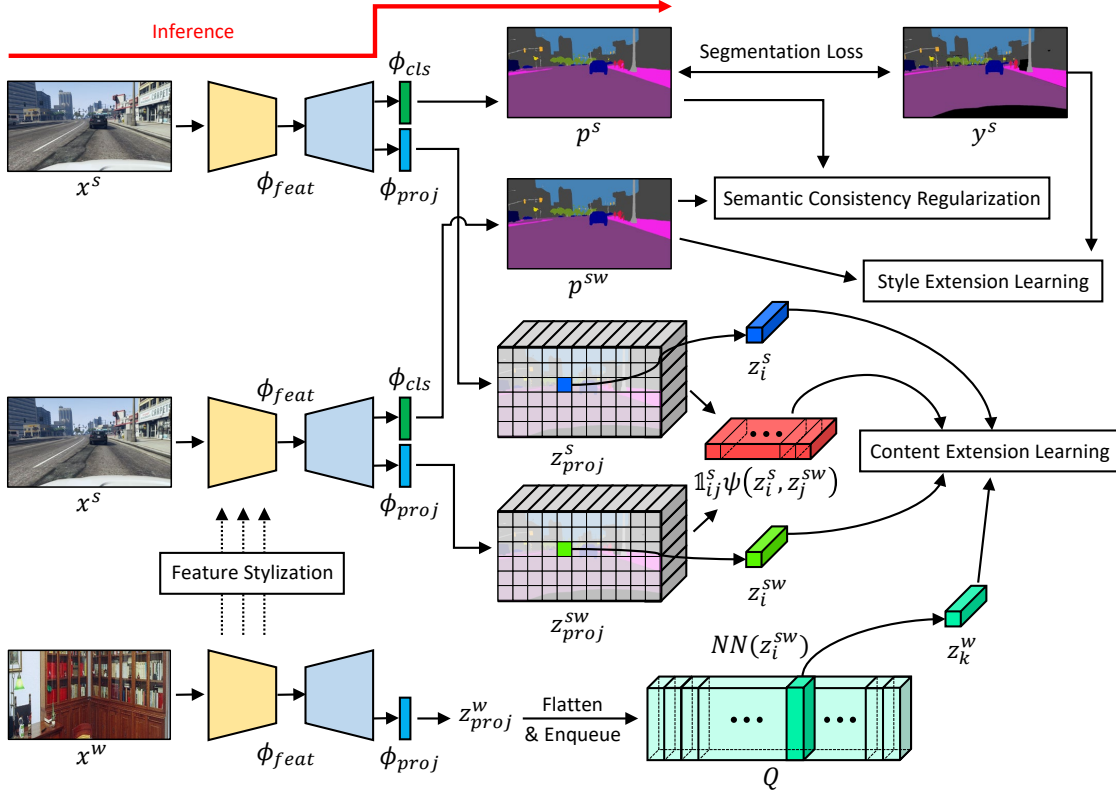


Figure 2. The overall learning process of WildNet. Our model is trained with four proposed methods: FS, CEL, SEL, and SCR. FS augments source features to resemble wild styles using the statistics of wild features, and the augmented features are used for CEL, SEL, and SCR. CEL performs contrastive learning using the augmented features and the wild features closest to them as positive samples and other class features as negative samples. SEL enables networks to learn task-specific information from features with diversified styles, and SCR regularizes networks to capture consistent semantic information from features with diversified contents and styles. With the proposed learning methods, our model learns domain-generalized semantic features by leveraging contents and styles from the wild.

various styles [22, 62]. In this paper, we focus on learning domain-generalized networks by leveraging a variety of contents and styles from ImageNet.

3. Proposed Method

In this section, we introduce four learning techniques consisting of Feature Stylization (FS), Content Extension Learning (CEL), Style Extension Learning (SEL), and Semantic Consistency Regularization (SCR) for learning domain-generalized features by extending both the content and style of the source domain to the wild. Our WildNet achieves superior generalization ability with them and the overall learning process is outlined in Fig. 2.

3.1. Problem Setup and Overview

Domain generalization (DG) aims to enhance the generalization capability on both the seen source domain \mathcal{S} and unseen target domains $\mathcal{T} = \{\mathcal{T}_1, \dots, \mathcal{T}_N\}$. Let ϕ be a semantic segmentation model that outputs pixel-wise category predictions p from image x . This model consists of a feature extractor ϕ_{feat} and classifier ϕ_{cls} . In DG, when we train the model, we have access to the source domain training dataset

$D^s = \{(x^s, y^s)\}$ while inaccessible to the target domains, where $x^s \in \mathbb{R}^{H \times W \times 3}$ is an image, $y^s \in \mathbb{R}^{H \times W \times K}$ is its pixel-wise label, and K is a number of semantic categories. The baseline model is trained with the segmentation loss

$$\mathcal{L}_{orig} = -\frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W \sum_{k=1}^K y_{hwk}^s \log(\phi(x^s)). \quad (1)$$

In this paper, we focus on extending both the content and style of the source domain to obtain high generalization performance on unknown target domains \mathcal{T} . We utilize the unlabeled wild dataset $D^w = \{x^w\}$, which has various contents and styles. At each training iteration, a random pair of source and wild images is provided as input, and the style and content of the source image are extended to the wild domain \mathcal{W} in the feature space. With the help of the wild, our network naturally learns domain-generalized semantic information from a variety of contents and styles. After the training, the model is evaluated on validation sets of both the seen source domain \mathcal{S} and unseen target domains \mathcal{T} .

3.2. Feature Stylization

As style is related to feature statistics [14, 15, 24, 26, 31] and the distributional shift due to style differences lies

mainly in shallow layers of networks [41], the styles of features can be diversified by adjusting the statistics of features from shallow layers. In this work, we diversify the styles of the source features with the help of wild styles by adding several AdaIN [24] layers to the feature extractor in the learning process. This enables us to augment source features to resemble wild styles without losing spatial information.

Let ϕ_l be the l -th layer of networks ϕ and let z_l be the feature output from ϕ_l when image x is input into ϕ . To allow networks to learn domain-generalized semantic information from various wild-style features, we swap the style of the source feature z_l^s from the source image x^s with the style of the wild feature z_l^w from the wild image x^w . In the l -th layer, we transfer the style of z_l^w to z_l^s and obtain the wild-stylized feature z_l^{sw} as

$$z_l^{sw} = \sigma(z_l^w) \frac{z_l^s - \mu(z_l^s)}{\sigma(z_l^s)} + \mu(z_l^w) \quad (2)$$

where $\mu(z_l)$ and $\sigma(z_l)$ are channel-wise mean and standard deviation of feature z_l , respectively. Because the distribution of z_l^s is re-normalized with channel-wise statistics of z_l^w , the style of z_l^s is swapped to the wild-style while maintaining the spatial information.

The wild-stylized feature z_l^{sw} is input into layer $l + 1$ and $z_{l+1}^{sw} = \phi_{l+1}(z_l^{sw})$ is output from the layer. z_{l+1}^{sw} can be swapped repeatedly in the style of z_{l+1}^w as

$$z_{l+1}^{sw} := \sigma(z_{l+1}^w) \frac{z_{l+1}^{sw} - \mu(z_{l+1}^{sw})}{\sigma(z_{l+1}^{sw})} + \mu(z_{l+1}^w). \quad (3)$$

By the above equation, feature z^{sw} is swapped in the style of z^w over multiple layers. As the layer deepens, semantic information should be captured more important than style, so FS applies only to some shallow layers in this work.

3.3. Content Extension Learning

In this subsection, we propose to extend the contents in the source domain to the wild. One of the reasons for overfitting to the source domain is that networks overlearn a limited amount of source content. We address this issue by increasing the intra-class content variability with content extension in the latent embedding space. To this end, we add a projection head ϕ_{proj} independently of the classification head ϕ_{cls} after the feature extractor ϕ_{feat} and extend the source contents to the wild in the embedding space.

When a source image x^s and wild image x^w enter the feature extractor ϕ_{feat} , it outputs the source feature z^s , wild-stylized source feature z^{sw} , and wild feature z^w . The projection head ϕ_{proj} receives features z^s , z^{sw} , and z^w and outputs pixel-level projected content features z_{proj}^s , z_{proj}^{sw} , and z_{proj}^w , respectively. All projected features are normalized by $z = z / \max(\|z\|_2, \epsilon)$ along the channel dimension. At each training iteration, z_{proj}^w is flattened and stored in the wild-content dictionary $Q \in \mathbb{R}^{C_q \times N_q}$ where C_q is the number of channels of projected features and N_q is the dictionary size. Our model uses the dynamic dictionary structure



Figure 3. Visualization of source-to-wild matching pixels on the cropped source and wild images. We extend the source content to the wild content closest to the wild-stylized source content corresponding to the source pixel. The source content is encouraged to come close to the wild content in the embedding space. This improves the generalization ability of our model to unseen contents.

in [18] as Q without a momentum update. We diversify contents by extending the source contents to the wild-stylized source contents and then to the wild domain using Q .

Here we focus on that the projected source feature z_i^s ¹ and projected wild-stylized feature z_i^{sw} corresponding to the i -th pixel of the source image x^s contain exactly the same semantic information, but content perturbation exists. In order to obtain reliable semantic information from unseen contents, networks should be able to cluster contents containing the same semantic information, distinguishing them from contents containing other semantic information. To achieve this objective, we adapt the contrastive learning strategy [39] to pixel-level instances in a supervised manner and define the source content extension loss for the i -th pixel as follows:

$$\mathcal{L}_{SCE}^i = -\log \frac{\psi(z_i^s, z_i^{sw})}{\psi(z_i^s, z_i^{sw}) + \sum_{j=1}^{N_z} \mathbb{1}_{ij} \psi(z_i^s, z_j^{sw})}, \quad (4)$$

$$\psi(z_i^s, z_j^{sw}) = \exp(z_i^s \cdot z_j^{sw} / \tau), \quad (5)$$

where $\mathbb{1}_{ij}^s$ is the negative pixel indicator that equals 1 if y_i^s and y_j^s are different and 0 if they are the same, N_z is the number of pixels and the temperature parameter τ is set to 0.07. We train the model only with reliable samples, ignoring ambiguous positive and negative samples by excluding pixels of unknown classes and pixels in other positions of the same class. Eq. (4) encourages z_i^s and z_i^{sw} to be close, while also encouraging z_i^s to move away from all negative class contents. Then, the pixel-wise loss can be applied to the entire source image by

$$\mathcal{L}_{SCE} = \frac{1}{N_z} \sum_{i=1}^{N_z} \mathcal{L}_{SCE}^i. \quad (6)$$

Eq. (6) encourages networks to make generalized predictions by reducing the distance in the embedding space between source contents and wild-perturbed source contents.

Next, we further extend the source contents to the wild by utilizing the wild-content dictionary Q . In the learning process, Q stores diverse pixel-level wild contents, which may not exist in the source domain. Thus, if we carefully select wild-content with semantic information that each pixel

¹The subscript *proj* is sometimes omitted for convenience.

needs to learn and then use it to train networks, networks become more robust to wild-content perturbations. Since there is no class information in the wild set D^w , Eq. (4) cannot be directly applied to this wild content extension. We address this issue from the perspective that similar semantic contents will be located close to each other in the embedding space. Inspired by [12], we take the wild content z_k^w closest to the wild-stylized source content z_i^{sw} from Q as

$$z_k^w = \arg \min_{q \in Q} \|z_i^{sw} - q\|_2 \quad (7)$$

and encourage the source content z_i^s to come close to it. Since z_i^{sw} and q are normalized early on, Eq. (7) can be calculated efficiently using a dot product and rewritten as

$$z_k^w = \arg \max_{q \in Q} (z_i^{sw} \cdot q). \quad (8)$$

Fig. 3 shows the wild content z_k^w matched to the source content z_i^s using the stylized source content z_i^{sw} . In this way, we provide various contents of the wild to the networks without category information. Now Eq. (4) can be adapted to the wild content extension as follows:

$$\mathcal{L}_{WCE}^i = -\log \frac{\psi(z_i^s, z_k^w)}{\psi(z_i^s, z_k^w) + \sum_{j=1}^{N_z} \mathbb{1}_{i \neq j}^s \psi(z_i^s, z_j^{sw})}. \quad (9)$$

In the wild content extension, we reuse the negative samples of the source content extension. Some negative extension approaches may give better performance and we left this for future work. Then we can apply the pixel-wise wild content extension loss to the entire source image by

$$\mathcal{L}_{WCE} = \frac{1}{N_z} \sum_{i=1}^{N_z} \mathcal{L}_{WCE}^i. \quad (10)$$

By combining the source content extension loss and wild content extension loss, the CEL loss is defined as

$$\mathcal{L}_{CEL} = \mathcal{L}_{SCE} + \mathcal{L}_{WCE}. \quad (11)$$

Our model learns to capture generalized semantic information from diverse contents by using the proposed CEL loss.

3.4. Style Extension Learning

Another reason for overfitting to the source domain is that networks overlearn a limited amount of the source style [36, 43, 67]. To address this issue, FS has diversified the styles of the source feature with the help of the wild. Interestingly, the style of the source features has changed while preserving spatial information, but networks fail to predict semantic categories from the wild-stylized feature as shown in Fig. 4f. In this subsection, we propose SEL for adapting networks to diversified styles. SEL aims to allow networks to naturally adapt to various styles by learning task-specific information from the wild-stylized feature.

When the wild-stylized source feature z^{sw} enters the classification head ϕ_{cls} , it outputs the pixel-wise softmax segmentation map $p^{sw} \in \mathbb{R}^{H \times W \times K}$. Because z^{sw} is the




Figure 4. Given the (a) source image and (d) label, FS stylizes the source feature with the wild feature from the (b) wild image. To visualize the wild-stylized feature, (c) we reconstructed an image from the wild-stylized feature using U-Net [48]. Contrary to (e) accurate predictions from the source feature, networks (f) fail to make correct predictions from the wild-stylized feature even though the spatial information of the source feature remains the same. To address this issue, we apply SEL loss to allow networks to learn task-specific information from the wild-stylized features.

wild-stylized source feature in which the style of the feature from x^s has been swapped with that of the feature from x^w , networks must predict the semantic label y^s of x^s from z^{sw} . For this objective, we train networks by minimizing the following SEL loss:

$$\mathcal{L}_{SEL} = -\frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W \sum_{k=1}^K y_{hwk}^s \log(p_{hwk}^{sw}). \quad (12)$$

Our model learns task-specific information from the wild-stylized features by applying the SEL loss. This enables our model to naturally learn domain-generalized semantic information from various styles.

3.5. Semantic Consistency Regularization

For high generalization capability on unseen domains, the classifier should capture consistent semantic information from features [5, 25], even if there are perturbations in both the style and content. However, as shown in Figs. 4e and 4f, the predicted result p^{sw} of the wild-stylized source feature z^{sw} differs from the predicted result p^s of the source feature z^s . Although SEL allows networks to learn task-specific information from z^{sw} , this does not guarantee that p^{sw} and p^s are identical. To address this issue, we propose SCR that regularizes networks to capture consistent semantic information even when both the content and style of the source domain are extended to the wild. SCR aims to train networks so that the predicted probability distributions $p^{sw} = \phi_{cls}(z^{sw})$ from the wild-stylized source features get closer to the $p^s = \phi_{cls}(z^s)$ from the source features. To this end, we adapt the Kullback-Leibler (KL) divergence loss as

$$\mathcal{L}_{SCR} = -\frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W \sum_{k=1}^K p_{hwk}^s \log \frac{p_{hwk}^s}{p_{hwk}^{sw}}. \quad (13)$$

With the SCR loss, our model learns consistent semantic information even with perturbations of style and content by the proposed wild extension methods.

4. Experiments

4.1. Datasets

Real semantic segmentation datasets. Cityscapes [10], BDD100K [61] and Mapillary [37] consist of 2975, 7000, and 18000 images for train set and 500, 1000, and 2000 for validation set. We consider 19 classes that are compatible with other datasets. In all of the tables, **C**, **B**, and **M** denote Cityscapes, BDD100K, and Mapillary, respectively.

Synthetic semantic segmentation datasets. GTAV [46] contains 24966 images rendered from the Grand Theft Auto V game engine. It has 12403, 6382, and 6181 images for train, validation, and test sets, respectively. SYNTHIA [49] contains 9400 images and we split it into 6580 and 2820 images for train and validation sets, following [7]. In all tables, **G** and **S** denote GTAV and SYNTHIA, respectively.

Wild dataset. ImageNet [11] is a large-scale image classification dataset used for network pre-training in most studies. In this paper, we use images without class labels as wild domain data. The generalization performance according to the number of images used for training our WildNet is presented in Table 6a.

4.2. Experimental Setup

We conduct extensive experiments over five different semantic segmentation datasets and report the mean intersection over union (mIoU) score on several domain generalization scenarios: $\text{GTAV} \rightarrow \{\text{Cityscapes}, \text{BDD100K}, \text{Mapillary}, \text{SYNTHIA}, \text{GTAV}\}$ and $\text{Cityscapes} \rightarrow \{\text{GTAV}, \text{BDD100K}, \text{Mapillary}, \text{SYNTHIA}, \text{Cityscapes}\}$. For fair comparisons with other DG methods, we re-implement IBN-Net [41] and RobustNet [7] on our baseline models and [†] denotes our re-implemented models. Our model is trained on one source domain train set (GTAV or Cityscapes) and validated on five domain validation sets (four unseen domains and one seen domain). To show the overall domain generalization performance, we additionally report the average value of mIoU on five domain validation sets (**Avg**). In all of the tables, the best results for each domain are marked in bold.

4.3. Implementation Details

We adapt ResNet-50, ResNet-101 [19], and VGG-16 [52] with DeepLabV3+ [2] as segmentation networks, and all backbones are pre-trained on ImageNet [11]. In the ResNet-based models, we use the SGD optimizer [47] with a momentum of 0.9 and weight decay of $5e-4$. The initial learning rate is set to $2.5e-3$ and is decreased using the polynomial policy with a power of 0.9. We train the models for 60K iterations with a batch size of 8. In the VGG-based models, we use the Adam optimizer [28] with a momentum of (0.9, 0.99). The initial learning rate is set to $1e-5$ and the batch size is set to 8. Following [7], we apply random scal-

Methods	C	B	M	S	G	Avg
Baseline [41]	22.20	-	-	-	61.00	-
IBN-Net [41]	29.60	-	-	-	64.20	-
Baseline [62]	32.45	26.73	25.66	-	-	-
DRPC [62]	37.42	32.14	34.12	-	-	-
Baseline [4]	23.29	-	-	-	-	-
ASG [4]	31.89	-	-	-	-	-
Baseline [7]	28.95	25.14	28.18	26.23	73.45	36.39
RobustNet [7]	36.58	35.20	40.33	28.30	72.10	42.50
Baseline [45]	31.70	-	-	-	-	-
GLTR [45]	38.60	-	-	-	-	-
Baseline	35.16	29.71	31.29	27.97	71.17	39.06
[†] IBN-Net [41]	36.52	34.18	38.74	30.41	70.78	42.12
[†] RobustNet [7]	38.78	35.64	40.38	28.97	70.16	42.78
WildNet (Ours)	44.62	38.42	46.09	31.34	71.20	46.33

Table 1. Comparison of mIoU(%) using ResNet-50 as backbone under the domain generalization setting $G \rightarrow \{C, B, M, S, G\}$.

Methods	C	B	M	S	G	Avg
Baseline [62]	33.56	27.76	28.33	-	-	-
DRPC [62]	42.53	38.72	38.05	-	-	-
Baseline [22]	33.40	27.30	27.90	-	-	-
FSDR [22]	44.80	41.20	43.40	-	-	-
Baseline [45]	34.00	28.10	28.60	-	-	-
GLTR [45]	43.70	39.60	39.10	-	-	-
Baseline	35.73	34.06	33.42	29.06	71.79	40.81
[†] IBN-Net [41]	37.68	36.64	36.75	30.84	70.39	42.46
[†] RobustNet [7]	37.26	38.66	38.09	30.17	70.53	42.94
WildNet (Ours)	45.79	41.73	47.08	32.51	71.91	47.81

Table 2. Comparison of mIoU(%) using ResNet-101 as backbone under the domain generalization setting $G \rightarrow \{C, B, M, S, G\}$.

ing within a range of [0.5, 2.0] and random cropping with a size of 768×768 . The output size of the projection head is 192×192 and we use uniformly sampled 64×64 size feature maps for CEL to prevent memory issues. For the diversity of the wild content dictionary, the wild feature maps are stored after uniform sampling with a size of 16×16 . The FS layer replaces first batch normalization and is added immediately after the addition operation of the first two residual blocks in ResNet, and it added right after the first ReLU after the first three maxpool layers in VGG. After training, all FS layers, projection head, and wild-content dictionary are removed, and our model can be applied to multiple unseen domains without further training on the target domains.

4.4. Comparison with DG methods

We compare our results with existing DG methods: IBN-Net [41], DRPC [62], ASG [4], FSDR [22], RobustNet [7], and GLTR [45]. Table 1 shows the generalization performance of the ResNet-50 model trained on GTAV. We evaluate models on five validation sets consisting of four unseen domains, including the Cityscapes, BDD100K, Map-

Methods	C	B	M	S	G	Avg
Baseline [62]	30.04	24.59	26.63	-	-	-
DRPC [62]	36.11	31.56	32.25	-	-	-
Baseline [4]	19.89	-	-	-	-	-
ASG [4]	31.47	-	-	-	-	-
Baseline [22]	-	-	-	-	-	-
FSDR [22]	38.30	34.40	37.60	-	-	-
Baseline [45]	31.40	-	-	-	-	-
GLTR [45]	37.20	-	-	-	-	-
Baseline	24.68	26.41	23.60	24.73	66.36	33.16
[†] IBN-Net [41]	30.25	30.09	31.87	26.22	65.47	36.78
[†] RobustNet [7]	30.13	29.22	33.96	26.16	64.73	36.84
WildNet (Ours)	39.18	34.49	40.75	27.25	64.57	41.25

Table 3. Comparison of mIoU(%) using VGG-16 as backbone under the domain generalization setting $G \rightarrow \{C, B, M, S, G\}$.

Methods	G	B	M	S	C	Avg
Baseline [41]	29.40	-	-	-	64.50	-
IBN-Net [41]	37.90	-	-	-	67.00	-
Baseline [7]	42.55	44.96	51.68	23.29	77.51	48.00
RobustNet [7]	45.00	50.73	58.64	26.20	76.41	51.40
Baseline	40.50	42.35	20.67	8.08	76.30	37.58
[†] IBN-Net [41]	45.28	46.61	56.78	26.41	75.47	50.11
[†] RobustNet [7]	45.28	48.21	56.97	26.59	74.91	50.39
WildNet (Ours)	47.01	50.94	58.79	27.95	75.59	52.06

Table 4. Comparison of mIoU(%) using ResNet-50 as backbone under the domain generalization setting $C \rightarrow \{G, B, M, S, C\}$.

\mathcal{L}_{orig}	\mathcal{L}_{CEL}	\mathcal{L}_{SEL}	\mathcal{L}_{SCR}	C	B	M	S	G	Avg
✓				35.16	29.71	31.29	27.97	71.17	39.06
✓	✓			41.25	35.95	40.06	31.26	68.75	43.46
✓	✓	✓		43.61	38.69	43.17	31.40	70.52	45.48
✓	✓	✓	✓	44.62	38.42	46.09	31.34	71.20	46.33

Table 5. Effect of the proposed losses on the domain generalization setting $G \rightarrow \{C, B, M, S, G\}$ using ResNet-50 as backbone in mIoU(%). Losses \mathcal{L}_{orig} , \mathcal{L}_{CEL} , \mathcal{L}_{SEL} , and \mathcal{L}_{SCR} are defined in Eq. (1), Eq. (11), Eq. (12), and Eq. (13), respectively.

illary, and SYNTHIA datasets, and one seen domain of GTAV. To demonstrate the high generalization ability over multiple domains, we also report the average value of the mIoU on the five domains. Our WildNet shows remarkably superior generalization capabilities, significantly outperforming other methods in all unseen target domains except the source domain. In particular, compared with the re-implemented results, we demonstrate that extending both the content and style is more effective in learning domain-generalized information than removing the domain-specific style. Given in Tables 2 and 3, we achieve superior generalization ability with ResNet-101 and VGG-16 models. Our model trained on Cityscapes also outperforms other DG methods as shown in Table 4. Extensive comparative experiments of different backbones on various domains demonstrate the superiority of our model.

4.5. Ablation Studies

In this subsection, extensive experiments with ResNet-50 model on the DG scenario from GTAV to Cityscapes, BDD100K, Mapillary, SYNTHIA, and GTAV are conducted to study the effectiveness of each component in the proposed method. Table 5 shows the effect of the proposed losses on domain generalization. The baseline model trained only with \mathcal{L}_{orig} overfits the source domain and has poor performance on unseen domains. Even with only \mathcal{L}_{CEL} applied, our model achieves an Avg of 43.46% with +4.40% improvement. This shows the importance of content diversification that is overlooked in many studies. Further, we make the wild-stylized features learn task-specific information with \mathcal{L}_{SEL} to achieve an Avg of 45.48%, and regularize the model to learn consistent semantic information with \mathcal{L}_{SCR} , finally achieving an Avg of 46.33%. Next, we conduct more ablations for important components.

Number of wild images. In Table 6a, the number of wild images used to train our model is considered. Even if only 10 wild images are used, the generalization performance is significantly enhanced by +5.54% compared with the baseline by preventing overfitting to the source domain. Moreover, the generalization performance of the model gradually improves as the number of wild images used increases. This shows that the extension of both content and style to the wild helps networks to learn domain-generalized semantic features.

Amount of FS. Table 6b shows the influence of the amount of FS on generalization performance. By replacing only the first batch normalization with FS, we can extend contents and styles to the wild based on diversified stylized features and improve generalization performance compared to baseline by +5.23%. Adding FS to some shallow layers boosts performance further. However, applying FS to deeper layers degrades performance slightly, as semantic content should be captured more important than style as the layer deepens. A suitable amount of FS, which does not disturb the semantic information, helps to train the generalized model by augmenting the source features to have various wild styles.

Size of wild content dictionary. Table 6c shows sensitivity to the size of the wild content dictionary. Extending the source content to the wild improves generalization performance, and even when extended to wild content within a mini-batch of size 2048 without a content dictionary, our model achieves higher generalization performance than without content extension. We take size of 393216.

FS with wild style. In Table 6d, we show the effect of FS using statistics of wild features on the generalization performance of the model. To apply FS without the help of the wild, the mean and standard deviation of the source features were multiplied by random values in the range [0.5, 1.5] and then used instead of the statistics of the wild features. The

Num.	C	B	M	S	G	Avg
Baseline	35.16	29.71	31.29	27.97	71.17	39.06
10	42.43	36.82	42.15	30.66	70.92	44.60
100	43.29	37.71	43.93	30.67	70.93	45.31
1000	43.70	38.27	43.56	30.80	70.94	45.45
10000	43.87	37.98	44.19	31.04	70.85	45.59
All	44.62	38.42	46.09	31.34	71.20	46.33

(a) Number of wild images used in the training process.

Residual Groups	C	B	M	S	G	Avg
Baseline	35.16	29.71	31.29	27.97	71.17	39.06
1	43.09	35.28	41.36	30.51	71.19	44.29
1-2	43.43	36.90	41.34	30.36	71.33	44.67
1-3	44.62	38.42	46.09	31.34	71.20	46.33
1-4	44.03	37.83	43.39	30.24	70.45	45.19
1-5	43.49	34.91	43.26	30.00	70.05	44.34

(b) FS applied to different amounts of residual groups.

Size	C	B	M	S	G	Avg
None	43.85	38.64	42.70	28.62	71.07	44.98
2048	43.64	38.51	43.94	29.10	71.14	45.27
49152	43.81	39.01	44.97	29.41	71.44	45.73
393216	44.62	38.42	46.09	31.34	71.20	46.33

(c) Size of the wild content dictionary.

Methods	C	B	M	S	G	Avg
Baseline	35.16	29.71	31.29	27.97	71.17	39.06
Random	42.67	34.84	38.71	30.36	71.14	43.54
Wild	44.62	38.42	46.09	31.34	71.20	46.33

(d) Effect of FS with the statistics of the wild features.

Methods	C	B	M	S	G	Avg
Random	43.83	38.75	45.10	30.65	71.10	45.89
Uniform	44.62	38.42	46.09	31.34	71.20	46.33

(e) Comparison of sampling methods for CEL.


Table 6. **Ablation Study.** For each setting, we report mIoU(%) using ResNet-50 as backbone in DG scenario: $G \rightarrow \{C, B, M, S, G\}$.

Figure 5. Visualization of extended wild contents. More visualizations are available in the supplementary material.

random FS improves performance compared to baseline by +4.48%, which shows the importance of diversifying styles. Furthermore, wild FS demonstrates that learning a natural style of the wild is much better with a gain of +7.27%.

Sampling methods. By using sampled feature maps, CEL stores various wild contents in the fixed-size dictionary and reduces memory consumption due to pixel-level contrastive loss calculations. Since two adjacent pixels have almost similar semantic information, uniform sampling makes learning more diverse contents than random sampling, leading to high generalization performance as shown in Table 6e.

5. Discussion

5.1. Qualitative Analysis

To analyze the wild content extension, we visualize the wild contents closest to the stylized source contents in Fig. 5. As can be seen in the figure, the source content is extended to wild content with semantic information similar to itself, e.g., the road under the car, the wheel of the bicycle, and the head of a man wearing a hat. Our model learns domain-generalized features by inducing source content closer to these wild content in the feature space. With

learning various wild contents, WildNet makes reliable predictions on unseen contents. Figs. 1 and 4 show segmentation results and visualization of wild-stylized features, and further analysis is provided in the supplementary material.

5.2. Limitations and Future Works

We have shown that the source content extends to wild content with semantic information similar to itself. However, similar semantic information in the two contents does not guarantee that the classes of the two contents are always the same, as observed for the rider of Fig. 5. Extending the rider with the hat to the person with the hat may bridge between the rider-class and person-class. Our future works will involve positive content selection using predicted class probabilities on wild images and negative content extension to further boost the class discrimination ability.

6. Conclusion

We presented WildNet which exploits unlabeled wild images for domain-generalized semantic segmentation. Our approach effectively extends style and content from source to wild, resulting in drastic performance improvement even we leverage 10 wild images. In contrast to previous studies that exploit generalization cues only from style, we additionally exploit the potential to generalize domain from content. We thoroughly ablated to demonstrate the efficacy of our WildNet and achieved superior segmentation performance under several domain generalization scenarios. We believe that our approach provides an opportunity to utilize huge amounts of unlabeled data for domain generalization.

Acknowledgement. This research was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF-2019R1A2C1007153).

References

- [1] Inigo Alonso, Alberto Sabater, David Ferstl, Luis Monteseano, and Ana C Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *ICCV*, 2021. 2
- [2] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, pages 801–818, 2018. 6
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607. PMLR, 2020. 2
- [4] Wuyang Chen, Zhiding Yu, Zhangyang Wang, and Animesh Anandkumar. Automated synthetic-to-real generalization. In *ICML*, pages 1746–1756. PMLR, 2020. 6, 7
- [5] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *CVPR*, pages 1791–1800, 2019. 5
- [6] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *CVPR*, pages 12475–12485, 2020. 2
- [7] 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 *CVPR*, pages 11580–11590, 2021. 1, 2, 6, 7
- [8] Sungha Choi, Joanne T Kim, and Jaegul Choo. Cars can't fly up in the sky: Improving urban-scene segmentation via height-driven attention networks. In *CVPR*, pages 9373–9383, 2020. 2
- [9] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, volume 1, pages 539–546. IEEE, 2005. 2
- [10] 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 *CVPR*, pages 3213–3223, 2016. 1, 6
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009. 2, 6
- [12] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *ICCV*, 2021. 2, 5
- [13] Leon Gatys, Alexander S Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. *NeurIPS*, 28:262–270, 2015. 1
- [14] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, pages 2414–2423, 2016. 1, 3
- [15] Leon A Gatys, Alexander S Ecker, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. Controlling perceptual factors in neural style transfer. In *CVPR*, pages 3985–3993, 2017. 3
- [16] Xiaoqing Guo, Chen Yang, Baopu Li, and Yixuan Yuan. Metacorection: Domain-aware meta loss correction for unsupervised domain adaptation in semantic segmentation. In *CVPR*, pages 3927–3936, 2021. 2
- [17] Jianzhong He, Xu Jia, Shuaijun Chen, and Jianzhuang Liu. Multi-source domain adaptation with collaborative learning for semantic segmentation. In *CVPR*, pages 11008–11017, 2021. 2
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 2, 4
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [20] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. 2
- [21] 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
- [22] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Fsd: Frequency space domain randomization for domain generalization. In *CVPR*, pages 6891–6902, 2021. 2, 3, 6, 7
- [23] Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. Decorrelated batch normalization. In *CVPR*, pages 791–800, 2018. 1
- [24] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, pages 1501–1510, 2017. 3, 4
- [25] Takashi Isobe, Xu Jia, Shuaijun Chen, Jianzhong He, Yongjie Shi, Jianzhuang Liu, Huchuan Lu, and Shengjin Wang. Multi-target domain adaptation with collaborative consistency learning. In *CVPR*, pages 8187–8196, 2021. 5
- [26] Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. Neural style transfer: A review. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3365–3385, 2019. 3
- [27] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, volume 33, 2020. 2
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [29] Suhyeon Lee, Junhyuk Hyun, Hongje Seong, and Euntae Kim. Unsupervised domain adaptation for semantic segmentation by content transfer. In *AAAI*, pages 8306–8315, 2021. 1, 2
- [30] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *NeurIPS*, pages 385–395, 2017. 1
- [31] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. In *IJCAI*, pages 2230–2236, 2017. 3

- [32] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *CVPR*, pages 6936–6945, 2019. 2
- [33] Haoyu Ma, Xiangru Lin, Zifeng Wu, and Yizhou Yu. Coarse-to-fine domain adaptive semantic segmentation with photometric alignment and category-center regularization. In *CVPR*, pages 4051–4060, 2021. 2
- [34] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, pages 116–131, 2018. 2
- [35] Rohit Mohan and Abhinav Valada. Efficientps: Efficient panoptic segmentation. *International Journal of Computer Vision*, 129(5):1551–1579, 2021. 2
- [36] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. In *CVPR*, pages 8690–8699, 2021. 5
- [37] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, pages 4990–4999, 2017. 6
- [38] Kiri Nichol. Painter by numbers, wikiart. <https://www.kaggle.com/c/painter-by-numbers>, 2016. 2
- [39] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 4
- [40] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *CVPR*, pages 3764–3773, 2020. 2
- [41] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, pages 464–479, 2018. 1, 2, 4, 6, 7
- [42] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaoou Tang, and Ping Luo. Switchable whitening for deep representation learning. In *ICCV*, pages 1863–1871, 2019. 1
- [43] Kwanyong Park, Sanghyun Woo, Inkyu Shin, and In So Kweon. Discover, hallucinate, and adapt: Open compound domain adaptation for semantic segmentation. In *NeurIPS*, volume 33, pages 10869–10880, 2020. 5
- [44] Sujoy Paul, Yi-Hsuan Tsai, Samuel Schulter, Amit K Roy-Chowdhury, and Manmohan Chandraker. Domain adaptive semantic segmentation using weak labels. In *ECCV*, pages 571–587. Springer, 2020. 2
- [45] Duo Peng, Yinjie Lei, Lingqiao Liu, Pingping Zhang, and Jun Liu. Global and local texture randomization for synthetic-to-real semantic segmentation. *IEEE Transactions on Image Processing*, 30:6594–6608, 2021. 2, 6, 7
- [46] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, pages 102–118. Springer, 2016. 1, 6
- [47] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951. 6
- [48] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015. 5
- [49] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, pages 3234–3243, 2016. 6
- [50] Subhankar Roy, Aliaksandr Siarohin, Enver Sangineto, Samuel Rota Buló, Nicu Sebe, and Elisa Ricci. Unsupervised domain adaptation using feature-whitening and consensus loss. In *CVPR*, pages 9471–9480, 2019. 1
- [51] Hongje Seong, Seoung Wug Oh, Joon-Young Lee, Seongwon Lee, Suhyeon Lee, and Euntai Kim. Hierarchical memory matching network for video object segmentation. In *ICCV*, pages 12889–12898, 2021. 2
- [52] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 6
- [53] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, pages 443–450. Springer, 2016. 1
- [54] Andrew Tao, Karan Sapra, and Bryan Catanzaro. Hierarchical multi-scale attention for semantic segmentation. *arXiv preprint arXiv:2005.10821*, 2020. 2
- [55] 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 *CVPR*, pages 7472–7481, 2018. 2
- [56] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *CVPR*, pages 2517–2526, 2019. 2
- [57] 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 *ECCV*, pages 642–659. Springer, 2020. 2
- [58] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, pages 3733–3742, 2018. 2
- [59] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. Upsnet: A unified panoptic segmentation network. In *CVPR*, pages 8818–8826, 2019. 2
- [60] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *CVPR*, pages 4085–4095, 2020. 2
- [61] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, pages 2636–2645, 2020. 6
- [62] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *ICCV*, pages 2100–2110, 2019. 2, 3, 6, 7
- [63] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *CVPR*, pages 12414–12424, 2021. 2

- [64] Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, and Tao Mei. Fully convolutional adaptation networks for semantic segmentation. In *CVPR*, pages 6810–6818, 2018. [2](#)
- [65] Xiangyun Zhao, Raviteja Vemulapalli, Philip Andrew Mansfield, Boqing Gong, Bradley Green, Lior Shapira, and Ying Wu. Contrastive learning for label efficient semantic segmentation. In *ICCV*, pages 10623–10633, 2021. [2](#)
- [66] Yuanyi Zhong, Bodi Yuan, Hong Wu, Zhiqiang Yuan, Jian Peng, and Yu-Xiong Wang. Pixel contrastive-consistent semi-supervised semantic segmentation. In *ICCV*, pages 7273–7282, 2021. [2](#)
- [67] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *ICLR*, 2021. [5](#)
- [68] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2020. [2](#)