

DiGA: *Distil to Generalize and then Adapt for Domain Adaptive Semantic Segmentation*

Fengyi Shen^{1,2,3}, Akhil Gurram², Ziyuan Liu², He Wang^{3*}, Alois Knoll^{1*}

¹Technical University of Munich, ²Huawei Munich Research Center, ³EPIC Lab, Peking University

¹fengyi.shen@tum.de, knoll@in.tum.de, ²{first.last}@huawei.com, ³hewang@pku.edu.cn

Abstract

Domain adaptive semantic segmentation methods commonly utilize stage-wise training, consisting of a warm-up and a self-training stage. However, this popular approach still faces several challenges in each stage: for warm-up, the widely adopted adversarial training often results in limited performance gain, due to blind feature alignment; for self-training, finding proper categorical thresholds is very tricky. To alleviate these issues, we first propose to replace the adversarial training in the warm-up stage by a novel symmetric knowledge distillation module that only accesses the source domain data and makes the model domain generalizable. Surprisingly, this domain generalizable warm-up model brings substantial performance improvement, which can be further amplified via our proposed cross-domain mixture data augmentation technique. Then, for the self-training stage, we propose a threshold-free dynamic pseudo-label selection mechanism to ease the aforementioned threshold problem and make the model better adapted to the target domain. Extensive experiments demonstrate that our framework achieves remarkable and consistent improvements compared to the prior arts on popular benchmarks. Codes and models are available at <https://github.com/fy-vision/DiGA>

1. Introduction

Semantic segmentation [7, 42, 45, 82] is an essential component in autonomous driving [20], image editing [54, 79], medical imaging [60], etc. However, for images in a specific domain, training deep neural networks [35, 37, 38, 62] for semantic segmentation often requires a vast amount of pixel-wisely annotated data, which is expensive and laborious. Therefore, **domain adaptive semantic segmentation**, *i.e.* learning semantic segmentation from a labelled source domain (either virtual data or an existing dataset) and then performing unsupervised domain adapta-

tion (UDA) [3, 15, 21, 46, 63] to the target domain, becomes an important research topic. Yet the remaining challenge is the severe model performance degradation caused by the visual domain gap between source and target domain data. In this work, we tackle this domain adaptive semantic segmentation problem, with the goal of aligning correct categorical information pixel-wisely from the labelled source domain onto the unlabelled target domain.

Currently, stage-wise training, composed of a warm-up and a self-training stage, has been widely adopted in domain adaptive semantic segmentation [1, 41, 50, 80, 81, 84] as it stabilizes the domain adaptive learning [81] and thus reduces the performance drop across domains. However, the gap is far from being closed and, in this work, we identify that there is still a large space to improve in both warm-up and self-training.

Regarding **warm-up**, recent works in this field [41, 50, 80, 81, 84] mostly adopt adversarial training [21, 67, 68] as their basic strategy, which usually contributes to limited adaptation improvements. Without knowing the target domain labels, this adversarial learning proposes to align the overall feature distributions across domains. Note that this alignment is class-unaware and fails to guarantee the features from the same semantic category are well aligned between the source and target domain, thus being sub-optimal as a warm-up strategy.

In contrast, we take an alternative perspective to improve the warm-up: simply enhancing the model’s domain generalizability without considering target data. To be specific, we introduce a pixel-wise symmetric knowledge distillation technique. The benefits are threefold: *i.* knowledge distillation is performed on the source domain where ground-truths are available, the learning thus becomes class-aware, which avoids the blind alignment as observed in adversarial training; *ii.* the soft labels created in the process of distillation can effectively avoid the model overfitting to domain-specific bias [71] and help to learn more generalized features across domains; *iii.* our symmetric proposal ensures the bidirectional consistency between the original source view and its augmented view, leading to more generaliz-

*corresponding author




Figure 1. **A systematic overview of DiGA framework.** For warm-up stage, instead of aligning features between the two domains from the beginning, we propose to first make the model generalizable to an unseen domain through our symmetric distillation scheme, which can be achieved even without access to the target domain data. Coupled with our target-aware CrDoMix data augmentation technique, a warm-up model of higher quality can be obtained. To make the model better adapted to the target domain, a threshold-free self-training stage is empowered by checking the consensus between feature-induced labels and probability-based labels that are initialized from the warm-up model. CrDoMix also contributes to the update of domain-generalized class centroids to strengthen the self-training stage.

able model performance in the warm-up stage. Our method achieves a significant improvement from 45.2 to 48.9 mIoU compared to adversarial training, meanwhile, it outperforms existing arts on domain generalized segmentation.

We further observe that making the data augmentation target-aware can help the model explore target domain characteristics and improve the adaptation. Hence, we propose cross-domain mixture (CrDoMix) data augmentation to better condition our warm-up model to the next stage.

In terms of **self-training** on the target domain, many works [1, 41, 50, 80, 81, 84] manage to optimize their self-training stage by finding proper thresholds for pseudo-labelling, which is onerous and not productive enough so far. In practice, however, self-training methods often get trapped into a performance bottleneck because the search for categorical thresholds is regarded as a trade-off between quantity and quality. Larger thresholds lead to insufficient learning, whereas smaller ones introduce noisy pseudo-labels in training.

To handle this, we propose bilateral-consensus pseudo-supervision, a threshold-free technique which selects pseudo-labels dynamically by checking the consensus between feature-induced and probability-based labels. The feature-induced labels come from pixel-to-centroid voting, focusing on local contexts of an input. The probability-based labels from the warm-up model are better at capturing global and regional contexts by the design nature of semantic segmentation architectures. Hence, by checking the consensus of these two types of labels generated by different mechanisms, the obtained pseudo-labels provide reliable and comprehensive estimation of target domain labels. Thus, an efficient self-training stage is enabled, leading to substantial performance gain against prior arts.

In this work, we present **DiGA**, a novel framework for domain adaptive semantic segmentation (see Fig.1). Our contributions can be summarized as follows:

- We introduce pixel-wise symmetric knowledge distillation solely on source domain, which results in a stronger warm-up model and turns out to be a better option than its adversarial counterpart.
- We introduce cross-domain mixture (CrDoMix), a novel data augmentation technique that brings further improvement to our warm-up model performance;
- We propose bilateral-consensus pseudo-supervision, empowering efficient self-training while abandoning categorical thresholds;
- Our method achieves remarkable and consistent performance gain over prior arts on popular benchmarks, e.g., GTA5- and Synthia-to-Cityscapes adaptation.

2. Related Work

Adversarial Domain Adaptation Adversarial learning [5, 21, 48, 55, 68] on UDA segmentation [67] is usually performed in GAN [2, 22, 32, 33, 49] fashion with a convolutional discriminator to force the output structure of the target segmentation maps to look like those from the source domain. However, adversarial training requires both source and target domain data as input and usually ends up with limited initial improvements due to blind feature alignment. In our work, in order to avoid this blind alignment, we propose pixel-wise symmetric knowledge distillation on source domain, which attains much better model generalization on target data even without observing them in training.

Knowledge Distillation Knowledge distillation technique [25, 44, 71–73] is first developed to perform model compression [4, 64] tasks. In conventional knowledge distillation pipeline, the student network learns, based on the same input, to mimic the soft output of a pretrained and fixed teacher. Therefore, the learning is unidirectional as the student has no feedback to the teacher. In our work, however, we point out that sufficient teacher-student interaction is useful in the UDA setting. We present a symmetric distillation approach to enhance the teacher-student collaboration and make the model domain generalizable.

Domain Generalization Domain generalization for semantic segmentation [12, 29, 36, 56, 83] assumes the accessibility only to one labelled source domain and aims to generalize the learned model to multiple target domains. SFDA [36] prepares different classification heads for different types of augmented source data, and the model generalization can be improved by assembling outputs of the heads at test time. SHADE [83] achieves domain generalization by utilizing a pretrained source domain teacher from the previous stage and places output consistency loss on the student network in the second stage, meanwhile, a style consistency loss is applied to the student outputs on different augmented styles.

Instead of using a fixed teacher in two-stage setting, we propose a simpler but more powerful end-to-end symmetric distillation strategy, where the teacher can get adaptively improved online by the student, gradually making the model domain generalizable.

Data Augmentation for UDA Segmentation Data augmentation [1, 14, 17, 18, 36, 52, 65, 66, 77, 78], which increases input diversity, is a promising strategy in UDA to explore out-of-source distribution and reduces the domain discrepancy at input level. Early dedications on GAN-based image domain transfer [10, 26, 30, 43, 76, 85] seek to train on target-like images with the same source labels to improve the model performance on the target domain. TIR [34] adds stylized source images carrying various texture changes to prevent the segmentation network from overfitting on one specific source texture. However, these approaches either introduce more training burdens (*i.e.*, each augmentation is an extra batch) or depend merely on augmented images without involving the equally informative source inputs. However, in our work, we leverage random masks taken from the source label map and simultaneously embed all meaningful augmentations in a single view without introducing multiple extra training batches.

Self-training for UDA Segmentation Self-training [11, 19, 53] is widely used in UDA as pseudo-labels [86, 87] can be generated for the target domain, such that it can be treated in the same fashion as source domain. In [31, 41, 86, 87], pseudo-label acquisition is based on class-wise thresholds determined by output uncertainty. However, improper thresholds often hamper the model to learn further. Seg-Uncertainty [84] presents a pseudo-label rectifying scheme according to the prediction variance, reducing the impact of thresholds, but obtains limited boost over the threshold-based arts. ProDA [80] seeks to rectify pseudo-labels softly online according to the prototypical context estimated by target domain features towards target class centroids. Nevertheless, all pixels in rectified pseudo-labels are treated equally during training, including the false labels, thus compromising the benefit of label rectification. In our work, we exploit to enhance the label quality while not sacrificing the quantity of target pseudo-labels in a threshold-free manner.

3. Method

This section describes our DiGA framework for domain adaptive semantic segmentation. For the warm up stage, we introduce the pixel-wise symmetric knowledge distillation technique in Section 3.1 and introduce our cross-domain mixture data augmentation technique in Section 3.2. Then, for the self-training stage, we describe our threshold-free pseudo-label selection mechanism in Section 3.3.

Our notations are shown as the following. Let $(\mathcal{X}_s, \mathcal{Y}_s)$ denote the source dataset and $x_s \in \mathcal{X}_s$ is a source RGB image with a semantic label map $y_s \in \mathcal{Y}_s$. And \mathcal{X}_t denotes

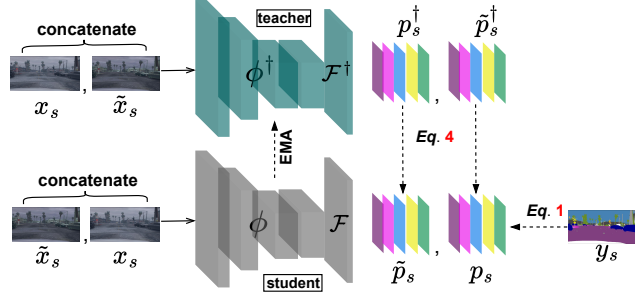


Figure 2. **Overview of the warm-up stage training.** The network takes x_s and \tilde{x}_s as input, then perform pixel-wise symmetric knowledge distillation solely on source domain.

the target domain dataset where $x_t \in \mathcal{X}_t$ stands for an unlabelled training image from the target domain. The goal is to train a segmentation network that can predict the correct per-pixel label for \mathcal{X}_t with the assist of $(\mathcal{X}_s, \mathcal{Y}_s)$.

3.1. Pixel-wise Symmetric Knowledge Distillation

In UDA segmentation, the purpose of a desired warm-up model is to provide strong support for the next training stage. We devise a new warm-up technique, pixel-wise symmetric knowledge distillation, to enhance the generalizability of the model towards the unseen target domain.

Supervised Loss As depicted in Fig. 2, our framework consists of a teacher network $\mathcal{F}^\dagger \circ \phi^\dagger$ and a student network $\mathcal{F} \circ \phi$, where $\{\phi^\dagger, \phi\}$ represent the feature encoders and $\{\mathcal{F}^\dagger, \mathcal{F}\}$ stand for the semantic classifiers. Now since ground-truths for source domain are available, we first train the student network $\mathcal{F} \circ \phi$ using $\{x_s, y_s\}$ and minimize the cross-entropy loss,

$$\mathcal{L}_s^{seg} = \sum_{h,w} \sum_c -y_s^{(c,h,w)} \log(p_s)^{(c,h,w)} \quad (1)$$

where h, w and c are height, width and number of semantic classes respectively.

Distillation Loss In parallel, we perform knowledge distillation solely on source domain data. Based on the x_s , we first create its augmented view \tilde{x}_s with basic operations (such as Gaussian blur, grayscale, color jitter, etc.) Then, we introduce a pixel-wise symmetric knowledge distillation scheme which interacts different views of the input image between the teacher and the student. By exchanging cross-view information alternately, we want the generalization of the model to be enhanced. Specifically, we pass x_s and \tilde{x}_s in a fused batch to the teacher $\mathcal{F}^\dagger \circ \phi^\dagger$ and the student $\mathcal{F} \circ \phi$ respectively to obtain the corresponding segmentation outputs (Fig. 2),

$$\{p_s^\dagger, \tilde{p}_s^\dagger\} = \sigma(\mathcal{F}^\dagger \circ \phi^\dagger(\{x_s, \tilde{x}_s\})) \quad (2)$$

$$\{\tilde{p}_s, p_s\} = \sigma(\mathcal{F} \circ \phi(\{\tilde{x}_s, x_s\})) \quad (3)$$



Figure 3. **An illustration of the unhelpful long-tail label from the source domain.** (a) is a patch cut from source ground-truth y_s , revealing the unhelpful label pixels in it; (b) shows teacher predicted label map $\sigma(\mathcal{F}^\dagger \circ \phi^\dagger(\tilde{x}_s))$ for the same patch from \tilde{x}_s ; (c) is the corresponding RGB image for this patch.

where $\sigma(\cdot)$ denotes the softmax operator, whose output stands for the segmentation map in probability space. Between the teacher and student outputs, we then introduce a pixel-wise knowledge distillation loss that is computed at all pixel locations on the segmentation map. In addition, the distillation loss is symmetric, meaning that \tilde{p}_s is supposed to get close to p_s^\dagger distribution while p_s close to \tilde{p}_s^\dagger ,

$$\mathcal{L}_s^{distil} = \overline{\mathcal{H}(p_s^\dagger, \tilde{p}_s)} + \alpha \overline{\mathcal{H}(\tilde{p}_s^\dagger, p_s)} \quad (4)$$

where $\mathcal{H}(a, b) = -a \log(b)$ is a cross-entropy loss, and the overline indicates that the loss is computed as the mean over all pixel locations. $\alpha \in (0, 1)$ is a scaling factor.

Following prior arts [6, 23, 66, 80], the gradient track of the teacher $\mathcal{F}^\dagger \circ \phi^\dagger$ is disabled, and its parameters $\Theta_{(\mathcal{F}^\dagger \circ \phi^\dagger)}$ are updated per iteration according to an exponential moving average (EMA) of $\mathcal{F} \circ \phi$ with a momentum of ξ ,

$$\Theta_{(\mathcal{F} \circ \phi)} \leftarrow \xi * \Theta_{(\mathcal{F}^\dagger \circ \phi^\dagger)} + (1 - \xi) * \Theta_{(\mathcal{F} \circ \phi)} \quad (5)$$

Why symmetric distillation works? *I.* Distilling the soft assignments p_s^\dagger to \tilde{p}_s : By softly forcing teacher-student prediction consistency between different views (*i.e.*, x_s and \tilde{x}_s), the network learns to produce stable outputs when input appearances are altered. This improves the model generalizability to a broader variants of source inputs and thereby, to data from an unknown domain distribution. *II.* Distilling the soft assignments \tilde{p}_s^\dagger to p_s : Not all source labels are helpful to learn a domain adaptable model [39]. As shown in Fig.3(a) and Fig.3(c), the virtual source label map is much too fine-grained by definition, such that it may provide labels that are hardly recognizable in the corresponding RGB image. Given this fact, training the network with those unhelpful long-tail source label pixels will introduce bias. However, the label map derived from the teacher output (Fig.3(b)) is smoother in those regions. Thus, introducing another symmetric distillation path on those regions of p_s can discourage the student to learn unhelpful source labels while maintaining the correct predictions under y_s 's supervision.

3.2. Cross-domain Mixture Data Augmentation

Other than basic data augmentation operations used in Sec. 3.1, the domain transferred images generated by neural




Figure 4. **A pictorial overview of our CrDoMix data augmentation technique.** A pretrained image domain translator is adopted.

networks provide a supplementary way of performing data augmentation as they provide useful information across domains. Therefore, we introduce cross-domain mixture (CrDoMix) data augmentation to create a target-aware novel view out of x_s (shown in Fig.4). Specifically, we take a **pre-trained and fixed** source-to-target image translator \mathcal{T}_{s2t} that is trained under CycleGAN [85] pipeline, but adding a semantic edge reconstruction loss when input comes from the source domain (detailed in Supplementary), and we pass x_s to \mathcal{T}_{s2t} to generate its target-like version $x_{s2t} = \mathcal{T}_{s2t}(x_s)$. Then, according to the source label map y_s , we randomly select half of its available classes c_{rs} and obtain a **binary mask** \mathcal{M} , with which we create CrDoMix augmentation x_{cdm} combining \tilde{x}_s and x_{s2t} ,

$$x_{cdm} = \tilde{x}_s \odot \mathcal{M} + x_{s2t} \odot (1 - \mathcal{M}) \quad (6)$$

where \odot is element-wise multiplication. The mixture of two images using CrDoMix simultaneously embeds diverse inter-domain effects on every single augmentation x_{cdm} without increasing the training batch size to cover all types of augmented images. Additionally, x_{cdm} does not break the geometric structure of the original input x_s , and the domain changes that happen around instance contours can randomly appear across the whole image. Inserting target-like appearances into data augmentation, CrDoMix helps to learn cross-domain knowledge and increases the target-awareness during warm-up training. Hence, we take x_{cdm} instead of just \tilde{x}_s for data augmentation in warm-up stage.

3.3. Threshold-free Self-training

After acquiring a warm-up model that demonstrates good generalizability on target data, we then consider further adapting it to the target domain assisted by pseudo-labels in a following **self-training (ST) stage**.

Label Preparation To prepare for the self-training stage, we initialize feature class centroids $\Lambda = \{\rho^{(k)}, k = 1, 2, \dots, c\}$ offline using encoder ϕ based on our warm-up model. According to the source ground-truth we get $\rho^{(k)}$,

$$\rho^{(k)} = \frac{\sum_{N_s} GAP(\phi(x_{cdm})^{(k)} \odot (y_s^{(k)} = 1))}{\sum_{N_s} \mathbb{1} \odot (y_s^{(k)} = 1)} \quad (7)$$

where k defines the specific semantic class, N_s is the number of images in \mathcal{X}_s , $GAP(\cdot)$ is the global average pooling

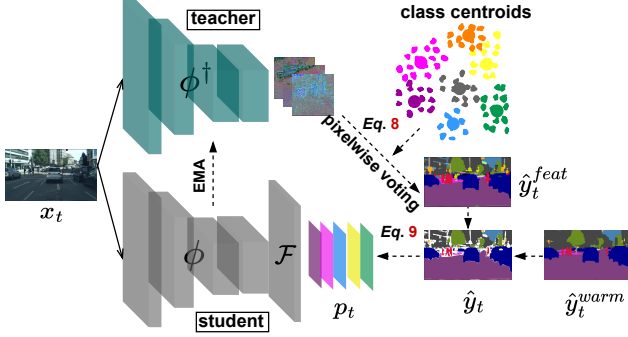


Figure 5. **Overview of the proposed bilateral-consensus pseudo-supervision.** Pseudo-labels for the target domain are determined dynamically according to consensus of probability-based label maps and the ones produced by centroid-guided voting in feature space.

operator, x_{cdm} is acquired using Eq.(6) to compute domain-robust centroids, and $\mathbb{1}$ is an indicator checking whether there exists class k in current image x_{cdm} .

In addition, a set of label maps $\hat{y}_t^{warm} \in \mathcal{Y}_t^{warm}$ for target domain can be obtained by processing the images from target dataset \mathcal{X}_t through the warm-up model. Then, we start to consider \mathcal{X}_t as input to our network for the ST stage. As shown in Fig.5, in each iteration, x_t is first encoded by ϕ^\dagger , and for each feature point at a specific pixel location we take feature centroids Λ as reference and vote this point to its nearest neighbor in Λ . Given that Λ are initialized based on CrDoMix inputs with source labels, we can confirm that they represent the correct categorical properties of semantic classes in feature space, such that pixel-to-centroid voting can lead to reasonable target label assignments. In this way, we acquire for target domain another type of label map induced from pixel-wise voting in feature space,

$$\hat{y}_t^{feat(jk)} = \mathcal{O}(\arg \min \|\phi^\dagger(x_t)^{(jk)} - \Lambda\|_2) \quad (8)$$

where \mathcal{O} stands for one-hot vectorization, $\|\cdot\|_2$ is L2 norm and $j \times k \in h \times w$.

Pseudo-Supervision Given the challenge of searching for categorical thresholds, can efficient self-training be achieved in a threshold-free manner? Here we give this question a positive answer. On top of the source domain distillation training described in Sec. 3.1, we introduce *bilateral-consensus pseudo-supervision* (BP) with a simple but effective dynamic pseudo-label generation procedure. Considering that \hat{y}_t^{feat} is acquired by pixel-wise voting in the feature space, it focuses more on local contexts but is likely to ignore the global semantics. For instance, \hat{y}_t^{feat} in Fig.5 spots the tiny objects such as traffic signs and lights, but it fails to segment building region as a whole part. On the contrary, \hat{y}_t^{warm} obtained from probability space is able to capture global structure of an image because the design of semantic classifiers [7, 42, 82] is usually able to cover larger

receptive fields rather than single pixels, such that neighboring region alignment can be better considered. Therefore, the correctness of pseudo-labels can be improved if we find pixel locations in \hat{y}_t^{feat} and \hat{y}_t^{warm} , which are defined from different mechanisms, sharing common label pixels. To this end, we select pseudo-labels \hat{y}_t dynamically by checking the consensus of label pixels between \hat{y}_t^{feat} and \hat{y}_t^{warm} , and compute the self-training loss on the target domain,

$$\hat{\mathcal{L}}_t^{seg} = \sum_{h,w} \sum_c -\hat{y}_t^{(c,h,w)} \log(p_t)^{(c,h,w)} \quad (9)$$

where $\hat{y}_t = \hat{y}_t^{feat} \cap \hat{y}_t^{warm}$. Hence, the self-training can be performed without any effort of finding class-wise thresholds. To compensate some wrong labels from warm-up models, we update \hat{y}_t^{warm} after 50 epochs with the current student output and repeat the training till it is finished. During training, the network parameters get constantly updated, thus the class centroids are dynamically altering from the perspective of the feature encoder. Therefore, in each iteration we update the feature class centroids based on EMA,

$$\rho^{(k)} \leftarrow \delta(\delta\rho^{(k)} + (1-\delta)\rho_s^{\prime(k)}) + (1-\delta)\rho_t^{\prime(k)} \quad (10)$$

where δ is the momentum and $\rho_s^{\prime(k)}$ is computed according to the latest x_{cdm} . Moreover, as \hat{y}_t becomes growingly more reliable, we also include $\rho_t^{\prime(k)}$ for obtaining domain-generalized centroids.

Full Objective So far, the full loss for training our DiGA framework can be summarized as following,

$$\mathcal{L}_{DiGA} = \lambda_s^{distil} \mathcal{L}_s^{distil} + \lambda^{seg} (\mathcal{L}_s^{seg} + \hat{\mathcal{L}}_t^{seg}) \quad (11)$$

4. Experiment and Discussion

4.1. Datasets and Implementation Details

For the source domain, we adopt the GTA5 dataset [58] consisting of 24,966 images with 1914×1052 resolution taken from the game engine, and SYNTHIA-RANDCITYSCAPES dataset [61] composed of 9,400 images of 1280×760 resolution with fine-grained segmentation labels. We adopt as target domain the 2975 urban scene training images in Cityscapes dataset [13] and its labelled validation set containing 500 images.

We implement DiGA on an NVIDIA Quadro RTX 8000 with 48 GB memory. For fair comparison we first use ImageNet [16] pretrained ResNet-101 [24] as backbone feature extractor and adopt Deeplab-V2 [7] for semantic segmentation. However, to test the architectural generalizability of our method, we also train on OCRNet [75] with HRNet-W48 [70] backbone, as well as on the transformer-based architectures DAFormer [27] and HRDA [28]. For the transformer-based approaches, in order to align with the




Figure 6. **Qualitative results of GTA5-to-Cityscapes adaptation on Cityscapes validation set.** Columns from left to right are: target domain inputs; ground-truth labels; segmentation predictions of BDL [41], ProDA [80], CPSL [40] and DiGA (ResNet).

Method	road	sdwk	bdng	wall	fence	pole	light	sign	veg	trm	sky	psn	rider	car	truck	bus	train	moto	bike	mIoU
BDL [41]	91.0	44.7	84.2	34.6	27.6	30.2	36.0	36.0	85.0	<u>43.6</u>	83.0	58.6	31.6	83.3	35.3	49.7	3.3	28.8	35.6	48.5
ProDA [‡] [80]	91.5	52.4	82.9	42.0	<u>35.7</u>	40.0	44.4	<u>43.3</u>	87.0	43.8	79.5	66.5	31.4	86.7	41.1	52.5	0.0	45.4	53.8	53.7
CPSL [‡] [40]	91.7	<u>52.9</u>	83.6	<u>43.0</u>	32.3	<u>43.7</u>	<u>51.3</u>	42.8	85.4	37.6	81.1	<u>69.5</u>	30.0	88.1	<u>44.1</u>	<u>59.9</u>	<u>24.9</u>	<u>47.2</u>	48.4	55.7
ProCA [31]	91.9	48.4	87.3	41.5	31.8	41.9	47.9	36.7	<u>86.5</u>	42.3	<u>84.7</u>	<u>68.4</u>	<u>43.1</u>	<u>88.1</u>	39.6	48.8	40.6	43.6	<u>56.9</u>	<u>56.3</u>
DiGA (Ours, ResNet)	95.6	67.4	89.8	51.6	38.1	52.0	59.0	51.5	86.4	34.5	87.7	75.6	48.8	92.5	66.5	63.8	19.7	49.6	61.6	62.7
DiGA (Ours, HRNet)	95.2	65.2	90.7	59.0	57.1	57.8	63.3	54.8	90.0	42.4	89.0	76.8	49.6	91.6	66.8	69.8	59.7	24.0	51.9	66.1
DAFormer [27]	95.7	70.2	89.4	53.5	48.1	49.6	55.8	59.4	89.9	47.9	92.5	72.2	44.7	92.3	74.5	78.2	65.1	55.9	61.8	68.3
DiGA (Ours + DAFormer)	95.7	70.4	89.8	54.8	47.8	51.3	57.8	63.9	90.3	48.8	91.8	73.1	46.6	92.6	78.5	81.3	74.8	57.3	63.2	70.0
HRDA [28]	96.4	74.4	91.0	61.6	51.5	57.1	63.9	69.3	91.3	48.4	94.2	79.0	52.9	93.9	84.1	85.7	75.9	63.9	67.5	73.8
DiGA (Ours + HRDA)	97.0	78.6	91.3	60.8	56.7	56.5	64.4	69.9	91.5	50.8	93.7	79.2	55.2	93.7	78.3	86.9	77.8	63.7	65.8	74.3

Table 1. **GTA5-to-Cityscapes adaptation results.** We compare our model performance with state-of-the-art methods. In all tables of Sec. 4.2, bold stands for **best** and underline for second-best. [‡] for fair comparison, we use their reported results after ST stage.

end-to-end training pipeline, we implement our \mathcal{L}_s^{distil} as an plug-and-play module, train distillation for 1 epoch and then adapt the model to target domain with self-training. During training, our input image size is 896×512 for both domains and batch size is 3 (containing both low and high resolution images). We augment the training data considering color jitter, color transfer [57], classmix [52, 66], grayscale and gaussian blur, etc. We use the SGD [59] optimizer with a default learning rate of 2.5×10^{-4} for ResNet but 1×10^{-3} for HRNet setting to train our network, and for transformer-based architectures we use AdamW [47] optimizer with a learning rate of 6×10^{-5} following [27, 28]. We use momentum $\xi = \delta = 0.999$ for EMA. During training, the losses are weighted by different hyperparameters, we set $\lambda^{seg} = 1$, $\alpha = 0.5$, and $\lambda_s^{distil} = 0.5$ for warm-up but 0.25 for ST stage as the model relies more on self-training in this stage. We report results based on multi-scale testing (MST) as mentioned in [31, 50, 69].

4.2. Evaluation on Benchmark Datasets

We compare our model with state-of-the-art approaches for domain adaptive semantic segmentation. As shown in Table 1, our DiGA framework shows leading performance among the state-of-the-art methods on GTA5-to-Cityscapes adaptation, achieving 62.7 mIoU after the ST stage. Note that our model still outperforms ProDA [80] and CPSL [40] even after they apply their proposed extra SimCLRv2 [8] distillation stage. Our model demonstrates superior results on many important classes (e.g., road, traffic light, traffic

sign, person, rider, car, truck, bus, bike etc.). A visual impression of segmentation examples generated by DiGA compared to other methods is shown in Fig.6. On Synthia-to-Cityscapes adaptation in Table 2, we also achieved state-of-the-art results (60.2 mIoU), outperforming other methods in segmenting sidewalk, traffic light as well as vehicles such as car and bicycle etc. In addition, the efficacy of our approach can be extended on more advanced architectures, reaching new milestones on both benchmarks, e.g., 66.1 (HRNet) and 74.3 (HRDA) mIoU for GTA5-to-Cityscapes adaptation, as well as 62.8 (HRNet) and 66.2 (HRDA) mIoU for Synthia-to-Cityscapes adaptation.

4.3. Ablation Study

Based on Deeplab-V2 with ResNet-101 setting, we provide ablative experiments in Table 3 to analyze the effect of each component for training our DiGA framework. Each experiment is evaluated on Cityscapes validation set to compute mIoU. Compared to the source-only approach, experiment in row (i) implies that distilling soft knowledge from p_s^\dagger to \tilde{p}_s on source domain according to the first half of Eq.(4) can lead to a considerable performance raise by 8.4%. On top of that, we observe that adding an additional symmetric path from \tilde{p}_s^\dagger to p_s (the second half of Eq.(4)) to our distillation brings a performance increase by 2.2%, reaching 48.9 mIoU (row (ii)). Combining CrDoMix, row (iii) gives the full model for training our warm-up stage, obtaining an increase from 48.9 to 51.1 mIoU, which already outperforms many existing self-training based meth-

Method	road	sdwk	blndg	wall*	fence*	pole*	light	sign	veg	sky	psn	rider	car	bus	mcycl	beycl	mIoU	mIoU*
BDL [41]	86.0	46.7	80.3	-	-	-	14.1	11.6	79.2	81.3	54.1	27.9	73.7	42.2	25.7	45.3	-	51.4
ProDA [‡] [80]	87.1	44.0	83.2	26.9	0.0	42.0	45.8	<u>34.2</u>	<u>86.7</u>	81.3	68.4	22.1	87.7	50.0	31.4	38.6	51.9	58.5
CPSL [‡] [40]	87.3	44.4	83.8	25.0	0.4	<u>42.9</u>	<u>47.5</u>	32.4	86.5	83.3	69.6	<u>29.1</u>	<u>89.4</u>	52.1	<u>42.6</u>	54.1	54.4	61.7
ProCA [31]	<u>90.5</u>	<u>52.1</u>	84.6	29.2	3.3	40.3	37.4	27.3	86.4	85.9	<u>69.8</u>	28.7	88.7	<u>53.7</u>	14.8	<u>54.8</u>	53.0	59.6
DiGA (Ours, ResNet)	89.1	53.4	86.1	<u>28.7</u>	<u>3.0</u>	49.6	50.6	34.9	88.2	<u>84.9</u>	71.3	40.9	91.6	75.1	50.3	65.8	60.2	67.9
DiGA (Ours, HRNet)	90.6	56.3	87.4	38.8	6.4	57.7	59.3	50.4	87.9	86.4	76.1	47.9	89.0	54.2	47.2	69.1	62.8	69.4
DAFormer [27]	84.5	40.7	88.4	41.5	6.5	50.0	55.0	54.6	86.0	89.8	73.2	48.2	87.2	53.2	53.9	61.7	60.9	67.4
DiGA (Ours + DAFormer)	85.2	41.4	88.2	42.6	7.5	52.1	57.5	47.7	87.8	90.8	75.0	50.8	87.8	58.0	58.5	63.0	62.1	68.6
HRDA [28]	85.2	47.7	88.8	49.5	4.8	57.2	65.7	60.9	85.3	92.9	79.4	52.8	89.0	64.7	63.9	64.9	65.8	72.4
DiGA (Ours + HRDA)	88.5	49.9	90.1	51.4	6.6	55.3	64.8	62.7	88.2	93.5	78.6	51.8	89.5	62.2	61.0	65.8	66.2	72.8

Table 2. **Synthia-to-Cityscapes adaptation results.** mIoU, mIoU* refer to 16-class and 13-class experimental settings, respectively. ‡ for fair comparison, we use their reported results after ST stage following [31].

Method	a	b	c	d	e	mIoU	Δ
Source-only						38.3	+0.0
Source-only	✓					38.9	+0.6
(i)	✓	✓				46.7	+8.4
(ii)	✓	✓	✓			48.9	+10.6
(iii)	✓	✓	✓	✓		51.1	+12.8
(iv)	✓	✓	✓	✓	✓	62.7	+24.4

Table 3. **DiGA components:** **a** \rightarrow MST, **b** \rightarrow $\overline{\mathcal{H}(p_s^\dagger, \tilde{p}_s)}$, **c** \rightarrow $\overline{\mathcal{H}(\tilde{p}_s^\dagger, p_s)}$, **d** \rightarrow CrDoMix, and **e** \rightarrow $\hat{\mathcal{L}}_t^{seg}$.

Strategy	Adv. [67]	Distil.	Adv. [67]+CrDoMix	Distil.+CrDoMix
mIoU	45.2	48.9	47.3	51.1

Table 4. **Warm-up model comparison** between adversarial training and our knowledge distillation w/ and w/o CrDoMix.

ods [11, 19, 53, 86]. Row (iv) indicates that, given a domain generalized warm-up model, our proposed bilateral-consensus pseudo-supervision strategy ensures high-quality self-training, boosting the mIoU from 51.1 to 62.7.

In Table 4, we compare warm-up models trained with different configurations. It indicates that our knowledge distillation technique provides a better warm-up solution than adversarial training [67], witnessing a substantial improvement no matter w/ or w/o our CrDoMix data augmentation.

4.4. Pseudo-labelling Comparison

We show that our proposed *bilateral-consensus pseudo-supervision* (BP) can be a more efficient strategy for pseudo-labelling in UDA segmentation. It assures a sufficient amount of reliable pixels in the resulting pseudo-labels but requires no effort to look for a proper threshold for each class. To verify this, we train ST stage models using different pseudo-labelling strategies, ablating on \hat{y}_t^{feat} , \hat{y}_t^{warm} and comparing with two influential strategies in this field, BDL [41] and ProDA [80]. For fair comparison, we start with our warm-up model in all experiments but only replace our BP with other strategies and evaluate the ST stage model performances on Cityscapes validation set. As shown in Table 5, training solely with

Strategy	(1) \hat{y}_t^{feat}	(2) \hat{y}_t^{warm}	(3)BDL	(4)ProDA	(5) $\hat{y}_t(ours)$
mIoU	52.1	53.8	56.2	59.5	62.7

Table 5. mIoU comparison of applying different pseudo-labelling techniques to train ST stage **based on our warm-up model.**

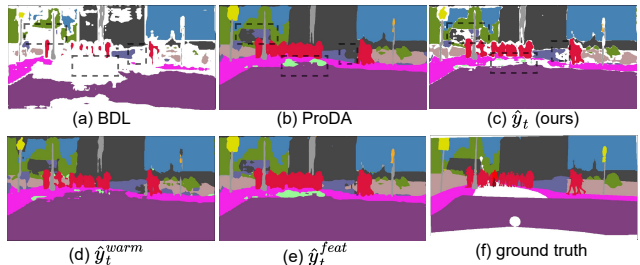


Figure 7. **Comparison of different pseudo-labelling techniques** given the same input image, and ground truth (f) is only adopted for comparison. Dashed black boxes reveal the major differences.

either \hat{y}_t^{feat} or \hat{y}_t^{warm} fails to show advantageous results (Exp.(1) & Exp.(2)), which indicates that it is required to check the consensus between \hat{y}_t^{feat} and \hat{y}_t^{warm} in order to achieve the best performance among all (Exp.(5)). We can also observe through Exp.(3) and Exp.(4) that replacing our BP strategy with BDL [41] and ProDA [80] strategies will lead to lower scores, obtaining 56.2 and 59.5 mIoU, respectively. Nevertheless, the scores are still higher than those reported in their original papers (48.5 and 53.7 mIoU), which also implies the superiority of our proposed warm-up strategy, confirming that a more advanced warm-up model contributes largely to the performance gain in ST stage.

A visual intuition can be obtained by comparing Fig.7 (a), (b) and (c) with the ground-truth (f). Given the same target domain input, BDL [41] manages to exclude predicted labels that are less confident than the class-wise thresholds, which misses a number of label pixels with correct predictions, leading to insufficient self-training. ProDA [80], on the other hand, keeps a full rectified label map by element-wisely re-weighting initial soft assignments using prototype-based soft assignments. Nevertheless, ProDA has




Figure 8. **Label uncertainty in ST stage.** From top to bottom: target input, prediction uncertainty (darker means lower confidence), pseudo-label, ground-truth (unused in training).

no label filtering mechanism, and the worst that can happen is that the model will be trained with wrong labels if false prototypical predictions override the correct initial ones. However, our BP strategy provides a better trade-off, showing superiority, both quantitatively and qualitatively, over the above methods in generating pseudo-labels.

4.5. Pseudo-labelling and Prediction Uncertainty

We reveal experimentally that our BP strategy is strongly correlated with pixel uncertainty in pseudo-label selection. During ST stage training, we visualize the model uncertainty on target inputs based on the output probabilities (See Fig. 8). For each mini-batch of data, the prediction uncertainty varies class-wisely and instance-wisely, which can hardly be described by some specific thresholds. However, we observe that BP strategy dynamically and efficiently identifies most pixel locations with relatively lower confidence ratio and successfully filters them out from the pseudo labels. Therefore, we confirm that, our BP strategy, despite of being threshold-free, can still obtain reliable pseudo-labels to improve self-training.

4.6. Extensive Experiments

Domain Generalization Even though we focus on solving the problem of UDA segmentation in this work, we also want to share that our warm-up strategy can offer a ‘free lunch’ for domain generalization on semantic segmentation and achieve superior results compared to SOTA methods. Decoupling CrDoMix (no x_{s2t}) from our data augmentation, meaning no information of target domain is touched, our warm-up distillation method works surprisingly well on domain generalization for semantic segmentation, *i.e.*, train on one labelled source domain and generalize to multiple target domains. Our model trained on GTA5 [58]

Method	Train on GTA5 (G)			
	→C	→B	→M	→S
ISW [12]	42.87	38.53	39.05	29.58
SFDA [71]	43.50	-	-	-
SAN-SAW [56]	45.33	41.18	40.77	31.84
SHADE [83]	46.66	43.66	45.50	-
Our Distillation	48.87	44.42	51.78	37.17

Table 6. **mIoU comparison with SOTA methods for domain generalization.** G, C, B, M and S denote GTA5, Cityscapes, BDD100k, Mapillary and Synthia, respectively. For fair comparison, all the listed methods are based on ResNet-101 backbone.

Method	Cityscapes			
	1/16 (186)	1/8 (372)	1/4 (744)	1/2 (1488)
CPS [9]	75.09	77.92	79.24	80.67
Ours	76.86	78.51	80.01	80.93

Table 7. **mIoU comparison of semi-supervised semantic segmentation** using HRNet backbone, based on which SOTA performance of CPS [9] is reported. Evaluation performed on Cityscapes validation set under different partition protocols.

dataset can be well generalized on the validation set of Cityscapes [13], BDD100k [74], Mapillary [51] and Synthia [61], outperforming existing SOTA methods of domain generalizable semantic segmentation [12, 56] by considerable margins (see Table 6). This **cannot** be accomplished by adversarial learning as it requires target data for training.

Semi-supervised Semantic Segmentation As shown in Table 7, following the same partition protocols on Cityscapes dataset, our stage-wise training pipeline also shows impressive performance on the task of semi-supervised semantic segmentation. In particular, the less labels available, the more advantageous DiGA is.

5. Conclusion

In this work, we propose DiGA framework for domain adaptive semantic segmentation. It first enhances the model generalization to the target domain by pixel-wise symmetric knowledge distillation performed on the source dataset. Supported by this strong warm-up model, our bilateral-consensus pseudo-supervision strategy reinforces the model adaptability during self-training without thresholds, demonstrating SOTA performances on popular benchmarks. Besides, through extensive experiments we also observe the efficacy of DiGA on domain generalized and semi-supervised semantic segmentation tasks. We believe DiGA can provide a universal solution for element-wise 2D (or 3D) classification problems under UDA or semi-supervised setting, into which we will investigate more in the future.

References

- [1] Nikita Araslanov and Stefan Roth. Self-supervised augmentation consistency for adapting semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15384–15394, 2021. [i](#), [ii](#), [iii](#)
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017. [ii](#)
- [3] Shai Ben-David, John Blitzer, Koby Crammer, Fernando Pereira, et al. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19:137, 2007. [i](#)
- [4] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541, 2006. [ii](#)
- [5] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 135–150, 2018. [ii](#)
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021. [iv](#)
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. [i](#), [v](#)
- [8] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020. [vi](#)
- [9] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2613–2622, 2021. [viii](#)
- [10] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1791–1800, 2019. [iii](#)
- [11] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6830–6840, 2019. [iii](#), [vii](#)
- [12] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11580–11590, 2021. [ii](#), [viii](#)
- [13] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. [v](#), [viii](#)
- [14] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020. [iii](#)
- [15] Hal Daumé III. Frustratingly easy domain adaptation. *arXiv preprint arXiv:0907.1815*, 2009. [i](#)
- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [v](#)
- [17] Terrance DeVries and Graham W Taylor. Dataset augmentation in feature space. *arXiv preprint arXiv:1702.05538*, 2017. [iii](#)
- [18] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. [iii](#)
- [19] Liang Du, Jingang Tan, Hongye Yang, Jianfeng Feng, Xiangyang Xue, Qibao Zheng, Xiaoqing Ye, and Xiaolin Zhang. Ssf-dan: Separated semantic feature based domain adaptation network for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 982–991, 2019. [iii](#), [vii](#)
- [20] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 22(3):1341–1360, 2020. [i](#)
- [21] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. [i](#), [ii](#)
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, 2014. [ii](#)
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [iv](#)
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [v](#)
- [25] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [ii](#)

- [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 *International conference on machine learning*, pages 1989–1998. PMLR, 2018. [iii](#)
- [27] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9924–9935, 2022. [v](#), [vi](#), [vii](#)
- [28] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXX*, pages 372–391. Springer, 2022. [v](#), [vi](#), [vii](#)
- [29] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Fsd: Frequency space domain randomization for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6891–6902, 2021. [ii](#)
- [30] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 172–189, 2018. [iii](#)
- [31] Zhengkai Jiang, Yuxi Li, Ceyuan Yang, Peng Gao, Yabiao Wang, Ying Tai, and Chengjie Wang. Prototypical contrast adaptation for domain adaptive semantic segmentation. *arXiv preprint arXiv:2207.06654*, 2022. [iii](#), [vi](#), [vii](#)
- [32] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard GAN. In *International Conference on Learning Representations*, 2019. [ii](#)
- [33] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. [ii](#)
- [34] Myeongjin Kim and Hyeran Byun. Learning texture invariant representation for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12975–12984, 2020. [iii](#)
- [35] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. [i](#)
- [36] Jogendra Nath Kundu, Akshay Kulkarni, Amit Singh, Varun Jampani, and R Venkatesh Babu. Generalize then adapt: Source-free domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7046–7056, 2021. [ii](#), [iii](#)
- [37] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015. [i](#)
- [38] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. [i](#)
- [39] Guangrui Li, Guoliang Kang, Wu Liu, Yunchao Wei, and Yi Yang. Content-consistent matching for domain adaptive semantic segmentation. In *European conference on computer vision*, pages 440–456. Springer, 2020. [iv](#)
- [40] Ruihuang Li, Shuai Li, Chenhang He, Yabin Zhang, Xu Jia, and Lei Zhang. Class-balanced pixel-level self-labeling for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11593–11603, 2022. [vi](#), [vii](#)
- [41] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6936–6945, 2019. [i](#), [ii](#), [iii](#), [vi](#), [vii](#)
- [42] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017. [i](#), [v](#)
- [43] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pages 700–708, Red Hook, NY, USA, 2017. [iii](#)
- [44] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2604–2613, 2019. [ii](#)
- [45] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [i](#)
- [46] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. *arXiv preprint arXiv:1602.04433*, 2016. [i](#)
- [47] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [vi](#)
- [48] Daniel Lowd and Christopher Meek. Adversarial learning. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 641–647, 2005. [ii](#)
- [49] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017. [ii](#)
- [50] Ke Mei, Chuang Zhu, Jiaqi Zou, and Shanghang Zhang. Instance adaptive self-training for unsupervised domain adaptation. *arXiv preprint arXiv:2008.12197*, 2020. [i](#), [ii](#), [vi](#)
- [51] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of the IEEE international conference on computer vision*, pages 4990–4999, 2017. [viii](#)
- [52] Viktor Olsson, Wilhelm Tranehed, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of*

- the *IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1369–1378, 2021. [iii](#), [vi](#)
- [53] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3764–3773, 2020. [iii](#), [vii](#)
- [54] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019. [i](#)
- [55] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *Thirty-second AAAI conference on artificial intelligence*, 2018. [ii](#)
- [56] Duo Peng, Yinjie Lei, Munawar Hayat, Yulan Guo, and Wen Li. Semantic-aware domain generalized segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2594–2605, 2022. [ii](#), [viii](#)
- [57] Erik Reinhard, Michael Adhikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer graphics and applications*, 21(5):34–41, 2001. [vi](#)
- [58] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016. [v](#), [viii](#)
- [59] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951. [vi](#)
- [60] 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. [i](#)
- [61] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3234–3243, 2016. [v](#), [viii](#)
- [62] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986. [i](#)
- [63] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3723–3732, 2018. [i](#)
- [64] Bharat Bhushan Sau and Vineeth N Balasubramanian. Deep model compression: Distilling knowledge from noisy teachers. *arXiv preprint arXiv:1610.09650*, 2016. [ii](#)
- [65] P.Y. Simard, D. Steinkraus, and J.C. Platt. Best practices for convolutional neural networks applied to visual document analysis. In *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.*, pages 958–963, 2003. [iii](#)
- [66] Wilhelm Trandheden, Viktor Olsson, Juliano Pinto, and Lennart Svensson. Dacs: Domain adaptation via cross-domain mixed sampling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1379–1389, 2021. [iii](#), [iv](#), [vi](#)
- [67] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7472–7481, 2018. [i](#), [ii](#), [vii](#)
- [68] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017. [i](#), [ii](#)
- [69] 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 *European conference on computer vision*, pages 642–659. Springer, 2020. [vi](#)
- [70] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020. [v](#)
- [71] Yufei Wang, Haoliang Li, Lap-pui Chau, and Alex C Kot. Embracing the dark knowledge: Domain generalization using regularized knowledge distillation. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 2595–2604, 2021. [i](#), [ii](#), [viii](#)
- [72] Yukang Wang, Wei Zhou, Tao Jiang, Xiang Bai, and Yongchao Xu. Intra-class feature variation distillation for semantic segmentation. In *European Conference on Computer Vision*, pages 346–362. Springer, 2020. [ii](#)
- [73] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4133–4141, 2017. [ii](#)
- [74] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. [viii](#)
- [75] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *European conference on computer vision*, pages 173–190. Springer, 2020. [v](#)
- [76] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2100–2110, 2019. [iii](#)
- [77] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regular-

- ization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019. [iii](#)
- [78] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. [iii](#)
- [79] Pan Zhang, Bo Zhang, Dong Chen, Lu Yuan, and Fang Wen. Cross-domain correspondence learning for exemplar-based image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5143–5153, 2020. [i](#)
- [80] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12414–12424, 2021. [i](#), [ii](#), [iii](#), [iv](#), [vi](#), [vii](#)
- [81] Qiming Zhang, Jing Zhang, Wei Liu, and Dacheng Tao. Category anchor-guided unsupervised domain adaptation for semantic segmentation. *arXiv preprint arXiv:1910.13049*, 2019. [i](#), [ii](#)
- [82] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [i](#), [v](#)
- [83] Yuyang Zhao, Zhun Zhong, Na Zhao, Nicu Sebe, and Gim Hee Lee. Style-hallucinated dual consistency learning for domain generalized semantic segmentation. *arXiv preprint arXiv:2204.02548*, 2022. [ii](#), [viii](#)
- [84] Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision*, 129(4):1106–1120, 2021. [i](#), [ii](#), [iii](#)
- [85] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [iii](#), [iv](#)
- [86] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018. [iii](#), [vii](#)
- [87] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5982–5991, 2019. [iii](#)