

Exploiting Diverse Characteristics and Adversarial Ambivalence for Domain Adaptive Segmentation

Bowen Cai^{1,2}, Huan Fu¹, Rongfei Jia¹, Binqiang Zhao¹, Hua Li² and Yinghui Xu¹

¹Alibaba Group

²Institute of Computing Technology, Chinese Academy of Sciences
 {kevin.cbw, fuhuan.fh, rongfei.jrf, binqiang.zhao}@alibaba-inc.com
 {lihua}@ict.ac.cn

Abstract

Adapting semantic segmentation models to new domains is an important but challenging problem. Recently enlightening progress has been made, but the performance of existing methods is unsatisfactory on real datasets where the new target domain comprises of heterogeneous sub-domains (*e.g.*, diverse weather characteristics). We point out that carefully reasoning about the multiple modalities in the target domain can improve the robustness of adaptation models. To this end, we propose a condition-guided adaptation framework that is empowered by a special attentive progressive adversarial training (APAT) mechanism and a novel self-training policy. The APAT strategy progressively performs condition-specific alignment and attentive global feature matching. The new self-training scheme exploits the adversarial ambivalences of easy and hard adaptation regions and the correlations among target sub-domains effectively. We evaluate our method (DCAA) on various adaptation scenarios where the target images vary in weather conditions. The comparisons against baselines and the state-of-the-art approaches demonstrate the superiority of DCAA over the competitors.

Introduction

Due to the expensive labeling costs, there is a growing interest in studying weakly-supervised or unsupervised approaches for semantic segmentation in recent years. One of the widely investigated techniques is synthetic-to-real unsupervised domain adaptation (UDA) (Hoffman et al. 2016; Tsai et al. 2018; Zhang, David, and Gong 2017; Zou et al. 2019), which expects to transfer knowledge from labeled synthetic datasets to unlabeled real domains. Recent solutions can be mainly divided into two streams, including feature alignment and self-training. The former utilizes adversarial training to align both the image styles and semantic feature distributions. As a promising alternative, self-training focuses on iteratively generating pseudo-labels with high confidence and fewer noises for target images, and re-train the networks using the produced labels. However, the performance of these models is still far from satisfactory on real datasets. One possible reason is that existing methods assume that images in the target domain are captured under similar conditions, while a new real domain can have

diverse characteristics. For example, the images taken in the same place with the same viewpoint may show significant differences because of weather changes.

The above analysis motivates us to study a more realistic UDA setting for semantic segmentation, *i.e.*, the images in a target domain has diverse characteristics. Taking “weather” as an example, a naive idea is to enlarge the synthetic dataset to cover as many weather conditions as possible and then directly adopt existing high-performing UDA methods. Unluckily, we find that the performance gains are not that remarkable. A possible reason is that the adversarial adaptation learning process may be more challenging when images in both domains show multiple modalities, resulting in semantic distortion translations and ambiguous feature alignment near weather condition boundaries. In fact, minimizing the discrepancy between synthetic and real domains is challenging (Saito et al. 2018), not to mention that each domain consists of several sub-domains (*e.g.*, Cityscapes*: {*Cloudy, Foggy, Rainy*}). Moreover, the learned pseudo-labels for target images contain significant noise when the atmosphere states are unconstrained.

In this paper, we argue that carefully reasoning about the diverse characteristics in the target domain may improve the effectiveness of adapted segmentation model. Towards this goal, we propose a condition guided adaptation framework (DCAA) empowered by a special attentive progressive adversarial training (APAT) mechanism and a novel self-training policy. We illustrate our observations and idea in Figure 1. For the APAT strategy, we first introduce a condition-guided style translator to generate realistic characteristic conditionable samples from synthetic images. Furthermore, we develop a condition attention module to exploit the complementarities of multi-modalities. Lastly, we study condition-specific adversarial training to align semantic feature distributions under each condition and further mitigate the global discrepancy.

For the novel self-training scheme, we investigate the correlations among target sub-domains to produce robust attentive pseudo-labels for target images. Importantly, we point out that the dense output from the discriminators implies adversarial ambivalence for hard adaptation regions. We thus consider the ambivalence to further improve the self-training process by revising the pseudo-labels and emerging the hard adaptation regions. We evaluate our DCAA on several chal-

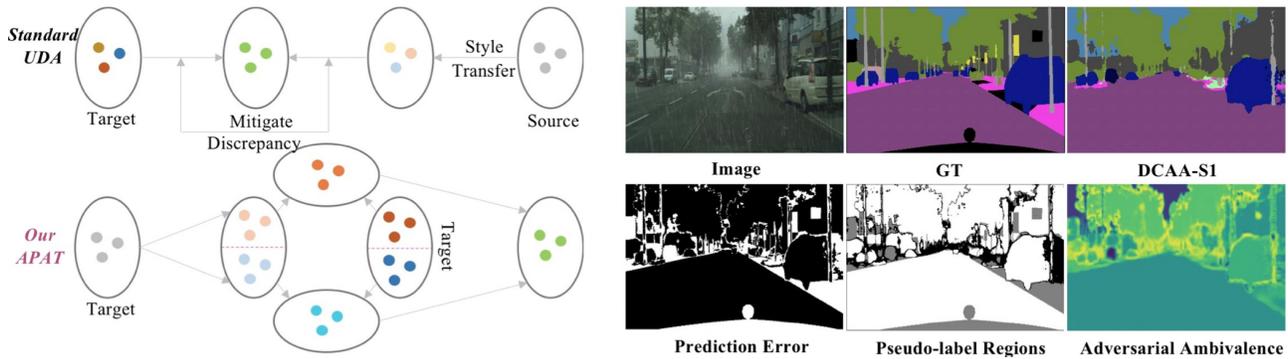


Figure 1: Left: A comparison between the typical domain adaptation process (Left Top) and our attentive progressive adversarial training (APAT) scheme (Left Bottom). Right: Adversarial Ambivalence and Pseudo-label. DCAA-S1: The first-stage prediction of DCAA. Prediction Error: The inconsistency between GT and DCAA-S1 (White). Pseudo-label Regions: “Black” represents the hard adaptation regions (without pseudo-labels); “Gray” and “White” denote regions with incorrect and correct pseudo-labels, respectively. Adversarial Ambivalence: The adversarial confidence obtained from the discriminator. “Yellow” means the regions are not well adapted. We find that the adversarial ambivalence shows consistency with the prediction error, thus can revise the pseudo-label related constraints. We refer to the “Introduction” section for more analyses of our motivations.

challenge adaptation scenarios, including GTA5→Cityscapes* (Richter et al. 2016; Cordts et al. 2016; Sakaridis et al. 2018; Hu et al. 2019), SYNTHIA→Cityscapes* (Ros et al. 2016), and GTA5→BDD100K (Yu et al. 2020a). Qualitative and quantitative comparisons against baselines and SOTA UDA methods demonstrate the superiority of our DCAA.

Related Work

Recent unsupervised domain adaptive segmentation solutions mainly focus on two points, *i.e.*, feature alignment (Du et al. 2019; Yang and Soatto 2020; Wang et al. 2020; Kim and Byun 2020; Xiao, Gu, and Zhang 2020; Li et al. 2020; Yang et al. 2019; Zhang and Wang 2020; Chen et al. 2019a; Wang, Cui, and Zhu 2018; Fu et al. 2019; Zhao et al. 2020; Gong et al. 2019a; Yu et al. 2020b) and self-training (Zhou et al. 2017; Zhang, David, and Gong 2017; Pan et al. 2020; Pei et al. 2018; Chen, Xue, and Cai 2019; Vu et al. 2019). Among the feature alignment approaches, FCN in the wild (Hoffman et al. 2016) was the first to introduce adversarial learning for global feature alignment and category statistic matching. AdaptSegNet (Tsai et al. 2018) modeled the structured spatial similarities on the semantic output space. Subsequently, many promising strategies has been devised to learn domain invariant features, such as conditional generator (Hong et al. 2018), task-specific decision boundaries (Saito et al. 2018), co-training (Luo et al. 2019), and patch distribution matching (Tsai et al. 2019). Other works show that style transfer (Zhu et al. 2017a) could largely mitigate the domain gap (Zhang et al. 2018; Sun et al. 2019). Thus, researchers exploited techniques such as semantic consistency (Hoffman et al. 2018), channel-wise feature alignment (Wu et al. 2018), symmetric adaptation consistency (Chen et al. 2019b), and controllable translation degrees (Gong et al. 2019b) to improve both the style translation and adaption processes. Self-training has recently gained considerable interests for UDA. For example, CBST (Zou et al. 2018) introduced an iterative class-balanced self-

placed learning policy to alternatively performed balanced pseudo-label generation and model retraining. The following works demonstrated the superiority of self-training by studying confidence regularization (Zou et al. 2019), bidirectional adversarial translation and self-supervised learning (Li, Yuan, and Vasconcelos 2019), and category-wise centroids (Zhang et al. 2019). These works always overlooked the diverse characteristics of natural images in a target domain, thus might produce sensitive models.

There are two adaptation settings that are closely related to our studied setting, *i.e.*, multi-target domain adaptation (MTDA) (Gholami et al. 2020) and compound domain adaptation (CDA) (Liu et al. 2020; Chen et al. 2019c). In MTDA, researchers expected to learn a general domain-agnostic model for largely unconstrained target domains. For example, the widely studied PACS dataset (Li et al. 2017) contains domains such as photo, cartoon, sketch, and art painting. In CDA, algorithms were developed to discover the unknown sub-domains in a target domain. In contrast, we study the setting where the sub-domains (from a single target domain) are strongly correlated in content, and the sub-domain labels are available (*e.g.*, BDD100K). We mainly take efforts into modeling the correlations (complementarities) of the multimodalities. Thus, our proposed approach is largely different from these MTDA and CDA methods. We clarify that assigning image-level (sub-domain) labels to a target domain is inexpensive in practice. We could even learn a condition classification network for the purpose.

Methodology

In the section, we take a special case that the target images are under two weather conditions to illustrate our method. Let $X^S = \{x_i^s\}_{i=1}^N$ and $Y^S = \{y_i^s\}_{i=1}^N$ be the source images and corresponding dense labels. The target images consist of two parts, *i.e.*, $X^T = \{X^{T1}, X^{T2}\}$, where T_1 and T_2 represent the special environmental conditions in this case. Our goal is to transfer knowledge from S to T so that a learned

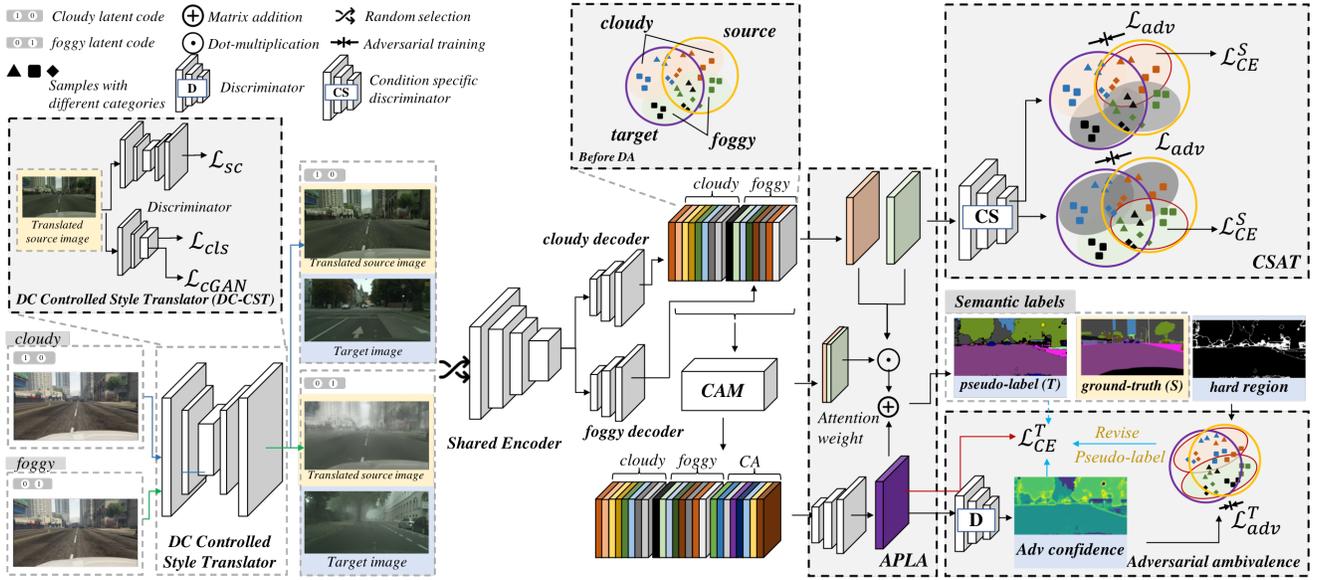


Figure 2: DCAA. We take GTA5→Cityscapes: $\{Cloudy, Foggy\}$ as an example to illustrate our DCAA. The framework is empowered by a special attentive progressive adversarial training scheme (CGST + CAM + CSAT) and a novel self-training policy (APLA + Adversarial Ambivalence).

segmentation model can produce promising parsing for images in the target domain. To this end, we develop a condition guided adaptation framework (DCAA), as shown in Figure 2. In the following, we will present our attentive progressive adversarial training (APAT) mechanism and self-training policy in detail.

Attentive Progressive Adversarial Training

Previous works (Luo et al. 2019; Hoffman et al. 2018; Tsai et al. 2018, 2019) preferred to take adversarial training to straightforwardly mitigate the source and target discrepancy. In contrast, we find that progressively reducing the domain shift by considering the weather conditions in the target domain would improve the adaptation process. Here, our APAT solution are constructed with three policies including condition guided style transfer, condition attention, and condition-specific adversarial training.

Condition Guided Style Transfer (CGST). We are here to learn a conditional generator $G_{ST}(X^S, c) : X^S \rightarrow X^T = \{X^{T_1}, X^{T_2}\}$ such that $G_{ST}(X^S)$ has the same distribution as X^{T_1} or X^{T_2} according to the specific latent variable $c \in \{0, 1\}$ (represents $\{T_1, T_2\}$). Thus, we develop a condition guided style transfer (CGST) following standard conditional GANs (Mirza and Osindero 2014) coupled with a semantic consistency constraint. We borrow concepts in StarGAN (Choi et al. 2018) and BicycleGAN (Zhu et al. 2017b) to improve the translation.

Specifically, we adopt a multi-level concatenation strategy (Zhu et al. 2017b) to inject c into G_{ST} as the translation guidance. We perform adversarial training to enforce G_{ST} producing realistic images with preferred styles. The adver-

sarial constraint takes the form as:

$$\mathcal{L}_{cGAN} = \mathbb{E}_{x^t \sim X^T} [\log(D_T(x^t))] + \mathbb{E}_{x^s \sim X^S} [\log(1 - D_T(G_{ST}(x^s, c)))] \quad (1)$$

where D_T is the discriminator with respect to X^T . The guidance from the injecting operation alone may not guarantee weather conditionable translations. Thus, we add an auxiliary condition classification loss (Choi et al. 2018) to further constraint the solution space:

$$\mathcal{L}_{cls} = \mathbb{E}_{x \sim X^T, c} [-\log(D_{cls}(c|x))] + \mathbb{E}_{x \sim X^T, c} [-\log(D_{cls}(c|G_{ST}(x)))], \quad (2)$$

where D_{cls} is the condition classifier which shares the same backbone with D^T .

We find that the obtained translator driven by constraints above would generate images with semantic distortion or inconsistency. For example, the trees in source images may become unreasonable building materials after the translation. We provide a remedy by incorporating a semantic consistency constraint which can be expressed as:

$$\mathcal{L}_{sc} = \mathbb{E}_{(x,y) \sim (X^S, Y^S)} [-\log(F_{seg}(y)|G_{ST}(x, c))], \quad (3)$$

where F_{seg} is a segmentation network. We use a light version of UperNet (Xiao et al. 2018) in the paper.

The full objective for our CGST can be expressed as:

$$\mathcal{L}_{CGST} = \mathcal{L}_{cGAN} + \mathcal{L}_{cls} + \lambda_{sc} \mathcal{L}_{sc}, \quad (4)$$

where the trade-off parameter λ_{sc} is set to 5.0 in our paper.

Condition Attention Module (CAM). After obtaining the stylized images $\hat{X}^S = \{\hat{X}^{S_1}, \hat{X}^{S_2}\} = G_{ST}(X^S, c)$, the trivial solution is to learn some multi-head networks to

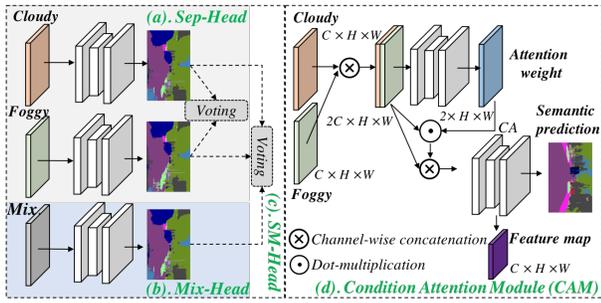


Figure 3: Several options in incorporating weather condition priors, including Sep-Head, Mix-Head, SM-Head, and our CAM. Dot lines refer that we use mean voting to capture the final prediction in the inference stage.

explicitly incorporate weather priors, as shown in Figure 3. We clarify that the effortless procedures may not sufficiently exploit the correlations of multi-modalities. We thus develop the condition attention module to address the issue.

In detail, we feed $x \in \{\hat{X}^S, X^T\}$ into a shared encoder (*Enc*) and two individual decoders (Dec^{C_1} and Dec^{C_2}) to capture:

$$\begin{aligned} f^{C_1}(x), p^{C_1}(x) &= Dec^{C_1}(Enc(x)), \\ f^{C_2}(x), p^{C_2}(x) &= Dec^{C_2}(Enc(x)), \end{aligned} \quad (5)$$

where $f^{C_1}(x), f^{C_2}(x) \in R^{D \times H \times W}$ are the semantic feature maps with dimension D , and $p^{C_1}(x), p^{C_2}(x) \in R^{L \times H \times W}$ are the probability predictions. In our paper, D is set to 256, and $L = 19$ is the number of semantic categories. Then, we learn attention weights $\mathcal{W} \in R^{2 \times H \times W}$ with a convolutional block (*Conv*) from $f^{C_1}(x)$ and $f^{C_2}(x)$ to capture the enhanced feature f . In further, we concatenate $f^{C_1}(x), f^{C_2}(x)$, and f together, and feed them to an additional parsing head (*CA*) to secure both $f^{CA}(x)$ and $p^{CA}(x)$. The operations can be expressed as:

$$\begin{aligned} \mathcal{W} &= Conv(f^{C_1}(x) \oplus f^{C_2}(x)), \\ f &= \mathcal{W}^{C_1} \otimes f^{C_1}(x) + \mathcal{W}^{C_2} \otimes f^{C_2}(x), \\ f^{CA}(x), p^{CA}(x) &= CA(f^{C_1}(x) \oplus f^{C_2}(x) \oplus f), \end{aligned} \quad (6)$$

where \oplus is the concatenation operation, \otimes is the element-wise product operation, and \mathcal{W}^{C_i} is the i th channel of \mathcal{W} .

Condition-Specific Adversarial Training (CSAT). To optimize the network, we introduce condition-specific adversarial training to encourage the decoders, *i.e.*, Dec^{C_1} , Dec^{C_2} , and *CA*, to focus on different characteristics. This allows alleviating the discrepancy between source and target feature distributions in a progressive manner. Specially, we employ three discriminators (D^{C_1} , D^{C_2} , and D^{CA}) for each decoder, respectively. D^{C_i} addresses the specific condition i , thus operates on \hat{X}^{S_i} and X^{T_i} to match the distributions $p^{C_i}(\hat{X}^{S_i})$ and $p^{C_i}(X^{T_i})$. And D^{CA} targets at transferring knowledge from \hat{X}^S to X^T , thus is a standard domain clas-

sifier. The encoder and decoders are optimized by:

$$\begin{aligned} \mathcal{L}_{adv}(x^t) &= - \sum_{h,w} \log(D^{CA}(p^{CA}(x_{hw}^t))) \\ &\quad - \sum_i \delta(x^t \in X^{T_i}) \sum_{h,w} \log(D^{C_i}(p^{C_i}(x_{hw}^t))), \end{aligned} \quad (7)$$

where $\delta(\cdot)$ is an indicator function which satisfies $\delta(True) = 1$ and $\delta(False) = 0$. Our discriminators follow the PatchGAN (Isola et al. 2017) fashion. Note that, instead of learning source-to-target adaptation, we perform target-to-source adaptation so that we can obtain the adversarial ambivalences for target images. We will introduce the adversarial ambivalence later.

Importantly, with the provided source-domain \hat{X}^S and Y^S pairs, we have a corresponding condition-specific segmentation loss:

$$\begin{aligned} \mathcal{L}_{CE}^S(\hat{x}^s) &= - \sum_{hw} \sum_{l=1}^L [y_{hw}^s \log(p_{hw}^{CA}(\hat{x}^s)) + \\ &\quad \sum_i \delta(\hat{x}^s \in \hat{X}^{S_i}) y_{hw}^s \log(p_{hw}^{C_i}(\hat{x}^s))]. \end{aligned} \quad (8)$$

Self-training

The proposed self-training policy aims to produce reliable pseudo-labels for target images and study adversarial ambivalence to revise the pseudo-labels and emerging the hard adaptation regions.

Attentive Pseudo-Label Assignment (APLA). Benefiting from the attentive progressive adversarial training scheme, we can obtain an adaptation model M_0 . As a result, we can secure pseudo-labels for target images where the correctness is guaranteed to a certain degree. In detail, given a target image x^t , M_0 will produce predictions $p^{C_1}(x^t)$, $p^{C_2}(x^t)$, and $p^{CA}(x^t)$, as well as the attention weights $\mathcal{W}(x^t)$. The attentive pseudo-labels y^t are generated as:

$$\begin{aligned} p(x^t) &= 0.5 * \sum_{i=1}^K \mathcal{W}^{C_i}(x^t) \otimes p^{C_i}(x^t) + 0.5 * p^{CA}(x^t) \\ y_{hw}^t(x^t) &= \begin{cases} l & l = \arg \max_l p_{hw}(l|x^t) \\ & \& p_{hw}(l|x^t) > \lambda_p \\ 0 & \text{otherwise} \end{cases}, \end{aligned} \quad (9)$$

where hw is the spatial location, l represents the category label, and λ_p denotes the pre-defined threshold.

Adversarial Ambivalence. With the pseudo-labels, we can directly transfer knowledge in a more explicit manner, *i.e.*, supervised learning:

$$\tilde{\mathcal{L}}_{CE}^T(x^t) = - \sum_{h,w} \sum_{l=1}^L y_{hw}^t(x^t) \log(p_{hw}^{CA}(x^t)). \quad (10)$$

However, the learned pseudo-labels are with many noises. Several previous works addressed the issue by adding regularization terms (Zou et al. 2019) or discovering hard adaptation images (Shin et al. 2020). For the latter, researchers

Method	Road	SW	Build	Wall	Fence	Pole	Light	Sign	Veg.	Terrace	Sky	Person	Rider	Car	Truck	Bus	Train	Moter	Bike	mIoU
GTA5 → Cityscapes*																				
Source Only	50.1	17.1	51.7	19.7	17.6	20.4	22.5	2.6	76.6	18.1	72.1	56.0	27.1	70.8	4.7	8.7	0.2	11.1	35.1	30.6
AdaptSegNet	80.4	16.4	71.9	14.3	17.2	18.0	30.1	20.3	75.3	15.0	74.8	53.0	26.0	76.0	12.8	32.4	2.7	24.1	29.5	36.3
CLAN	86.6	12.7	74.0	18.6	16.0	22.7	32.6	20.1	73.3	25.8	76.6	54.5	28.0	80.3	24.3	31.8	0.1	18.5	25.8	38.0
AdvEnt	88.2	33.1	78.3	26.0	11.1	31.3	34.7	26.3	74.5	22.0	63.4	55.3	29.9	82.6	25.0	25.0	0.01	20.0	31.7	39.9
CRST	90.6	42.8	78.5	21.0	12.1	35.1	35.8	25.9	69.7	14.6	78.2	61.3	30.2	81.7	16.0	34.8	5.04	28.7	47.5	42.6
CAG-UDA	87.1	37.9	82.7	26.3	29.9	35.1	41.2	39.1	82.3	27.8	76.3	61.6	32.1	83.6	17.8	34.3	18.8	19.2	30.1	45.4
AMEAN++	88.0	41.5	81.1	36.5	27.3	30.7	45.0	27.3	82.6	37.3	73.5	61.7	32.4	85.0	27.6	42.1	13.6	17.5	23.7	46.0
OCDA++	73.5	29.8	67.7	18.7	22.7	26.7	35.9	22.4	81.1	32.3	72.2	48.1	25.6	78.1	11.2	12.8	1.8	20.7	25.8	37.2
MDTA-ITA	80.1	23.6	76.0	17.6	24.2	25.4	34.5	14.2	77.7	18.3	73.5	49.7	12.3	79.7	16.9	22.9	5.1	6.8	10.8	35.2
DCAA	90.5	52.8	84.9	38.7	30.4	37.6	45.5	42.4	84.5	45.5	80.7	63.8	36.2	86.3	34.9	45.1	19.7	14.9	42.5	51.4
SYNTIA → Cityscapes*																				
Source Only	58.6	21.3	65.7	8.0	0.0	22.3	1.6	9.3	70.1	-	71.7	58.4	25.9	64.1	-	10.7	-	14.1	24.6	32.9
AdaptSegNet	84.7	44.1	74.1	8.7	0.2	17.5	3.8	4.3	71.3	-	69.5	47.8	16.4	62.2	-	23.5	-	10.5	31.5	35.6
CLAN	77.6	31.2	69.9	5.4	0.2	21.8	5.6	6.6	68.6	-	73.9	50.8	19.1	56.6	-	24.4	-	5.4	27.5	34.0
CRST	62.2	23.4	75.6	6.4	1.6	33.8	22.8	30.9	77.8	-	71.5	48.6	24.2	79.5	-	15.2	-	12.9	48.4	39.7
CAG-UDA	71.9	32.8	71.7	19.6	7.7	34.5	0.0	16.6	80.3	-	80.9	51.9	32.5	79.4	-	26.5	-	22.1	51.9	42.5
DCAA	67.7	30.7	74.7	22.9	14.1	33.2	38.5	30.4	82.4	-	81.6	60.1	31.2	80.8	-	40.2	-	14.5	54.7	47.4

Table 1: Performance on Cityscapes*. We reproduce other methods using their released codes following their best configurations. We clarify these reproduced scores under our setting are reliable considering the improvements over the baseline (Source Only). ++: We improve these CDA methods by providing the target sub-domain labels, and using our CGST.

compute the image-level confidence based on the pixel-level pseudo-label confidence, then employ pseudo-label supervision and adversarial learning for easy and hard samples, respectively. In contrast, we exploit the adversarial ambivalences by considering the dense discriminator outputs (PatchGAN) to further regulate the easy and hard adaptation regions. For each target image, we utilize pseudo-labels and the adversarial confidence to supervise the easy adaptation regions in a soft manner, and constraint the hard adaptation regions via adversarial loss. We take the regions where $y_{hw}^t(x^t) = 0$ as the hard adaptation regions.

Specially, we feed $p(x^t)$ into the discriminator D^{CA} to obtain an adaptation confidence map $d(x^t)$. Then, we rectify $\tilde{\mathcal{L}}_{CE}^T(x^t)$ by incorporating pixel-wise weights as:

$$\mathcal{L}_{CE}^T(x^t) = - \sum_{h,w} d_{hw}(x^t) \sum_{l=1}^L y_{hwl}^t(x^t) \log(p_{hwl}^{CA}(x^t)). \quad (11)$$

We filter out the easy adaptation regions to constraint the remained regions as:

$$\mathcal{L}_{adv}^T(x^t) = - \sum_{h,w} \delta(y_{hw}^t(x^t) = 0) \log(D^{CA}(p^{CA}(x_{hw}^t))). \quad (12)$$

In the self-training stage, \mathcal{L}_{CE}^S and \mathcal{L}_{CE}^T encourage the network produces domain invariant features, and \mathcal{L}_{adv}^T allows the network focusing on pushing hard adaptation regions to the feature space. We set λ_p to 0.6 to generate more pseudo-labels for target images benefiting from the adversarial ambivalence mechanism. Besides, we adopt M_0 (with fixed parameters) to produce pseudo-labels and adversarial ambivalences. We learn a new model M_1 that shares the same network architecture with M_0 . The parameters are ini-

tialized via the pre-trained M_0 .

Experiments

We refer to the supplemental materials for the implementation details and more analyses of our approach. Our codes will be made public available.

Adaptation Scenarios. We evaluate our proposed approach on three challenging adaptation scenarios, *i.e.*, GTA5→Cityscapes* and SYNTIA→Cityscapes*, and GTA5→BDD100K. Cityscapes* is a combined benchmark consisting of Cityscapes-cloudy (the commonly used Cityscapes dataset), Cityscapes-Foggy, and Cityscapes-Rainy. Specifically, Cityscapes-Cloudy provides 3,475 images with a resolution of 2048×1024 , which are officially split into 2,975 training images and 500 validation images. The images in Cityscapes-Foggy are captured by simulating different levels of foggy to images in Cityscapes-Cloudy. Cityscapes-Rainy selects 262 training images and 33 validation images in Cityscapes-Cloudy to attach 36 degrees of rain to generate 9,432\1,118 realistic rainy scenes. We randomly collect 12 rainy images for each reference image to construct our final Cityscapes-Rainy, *i.e.*, 3,144 images for training, and 396 images for evaluation. The enlarged target benchmark (Cityscapes*) is now has 9,094 (2975 + 2975 + 3144) images for training and 1,396 (500 + 500 + 396) images for evaluation. BDD100K is constructed for heterogeneous multitask learning where the images are with diverse weather conditions. We take BDD100K: {*Rainy, Snowy, Cloudy*} as our target domain following OCDA (Liu et al. 2020). There are 4,855\215, 5,307\242, and 4,535\346 training\test images for the three weather conditions, respectively. We also evaluate our trained model on the unseen BDD100K-Overcast domain (627 images).

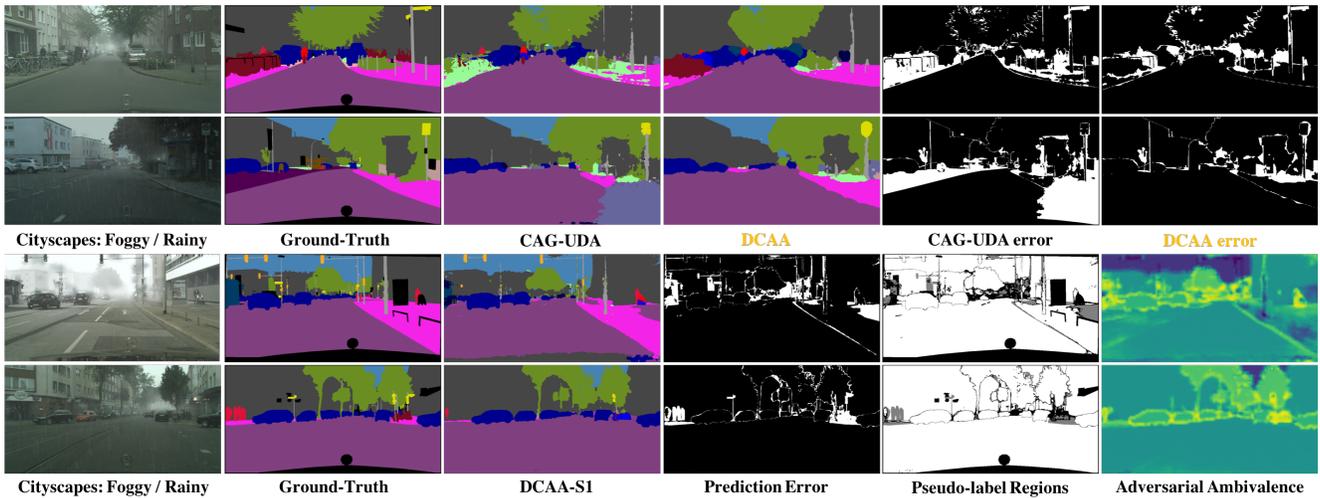


Figure 4: Qualitative Results. We refer to Figure 1 for the explanations of these maps. From the top part, our predictions contain more details compared with CAG-UDA. From the bottom part, the adversarial ambivalence is potent to indicate the incorrectly predicted regions and revise the pseudo-label related constraints.

The synthetic dataset GTA5 contains 24,966 images with a size of 1914×1052 , and has the consistent 19 semantic categories as Cityscapes. SYNTHIA is another synthetic benchmark that offers 9,400 images of size 1914×1052 . We use the 16 common category labels as Cityscapes for evaluation.

Compared Methods. We reproduce several widely investigated SOTA methods, including AdaptSegNet (Tsai et al. 2018), CLAN (Luo et al. 2019), CRST (Zou et al. 2019), ADVENT (Vu et al. 2019), and CAG-UDA (Zhang et al. 2019) using their shared codes and their best configurations for benchmark performance reporting. Directly evaluate their methods using the public models will result in lower scores since their models are trained only for Cityscapes-Cloudy. Moreover, we make the sub-domain labels available to improve the open domain adaptation approaches OCDA (Liu et al. 2020) and AMEAN (Chen et al. 2019c). We also compare our DCAA with a SOTA multi-target domain adaptation method MDTA-ITA (Gholami et al. 2020).

General Performance. The scores are reported in Table 2 and Table 1. Generally, DCAA ranks 1st on all the stud-

Method	Rainy	Snowy	Cloudy	Avg.	Overcast
Source Only	16.2	18.0	20.9	18.9	21.2
AdaptSegNet	20.2	21.2	23.8	22.1	25.1
CBST	21.3	20.6	23.9	22.2	24.7
IBN-Net	20.6	21.9	26.1	22.8	25.5
PyCDA	21.7	22.3	25.9	23.3	25.4
OCDA	22.0	22.9	27.0	24.5	27.9
OCDA++	23.1	22.8	29.2	25.7	29.8
MDTA-ITA	22.6	23.2	27.5	24.9	24.9
DCAA	31.0	34.6	37.4	34.7	38.3

Table 2: GTA5→BDD100K. During the training process, the target sub-domains consists of “Rainy”, “Snowy”, and “Cloudy”, while “Overcast” is unseen.

ied adaptation scenarios. For GTA5→BDD100K, DCAA achieves consistently promising improvements (about 9.0%) over all the target sub-domains, even for the unseen “Overcast” images. The results demonstrate that our approach can effectively exploit multi-modalities’ characteristics and is potent to produce high-performing adaptation models. For GTA5→Cityscapes*, DCAA outperforms the baseline (Source Only) by a large margin (51.4% vs. 30.6%), while yields a significant improvement (6.0%) over the SOTA method CAG-UDA. Besides, we find that OCDA and MDTA-ITA can not perform well on the segmentation adaptation scenarios. For example, OCDA only captures slight improvements over the AdaptSegNet baseline (1.1% ~ 2.4%) as reported in Table 1 and Table 2. A possible reason is that they focuses on disentangling the domain-style codes from the content codes. It may be easy to disentangle global style codes for image classification problems as studied in their paper. However, extracting content independent pixel-level style codes would be more challenging. For SYNTHIA→Cityscapes*, DCAA produces appealing semantic parsing for target images with a remarkable mean IoU (47.4%). Moreover, we present some qualitative results in Figure 4. Our DCAA can preserve more details in boundary areas and revise some incorrectly labeled regions predicted by other methods, especially for images with bad weather conditions. Furthermore, we find that the mIoU reported in previous publications (training and evaluating on Cityscapes-Cloudy) are consistently higher (about 5%) than the reproduced versions (training and evaluating on Cityscapes*). The large performance gap shows that trivially learning an adaptation model without explicitly exploiting the diverse weather conditions may not sufficiently discovering their correlations.

Method	Road	SW	Build	Wall	Fence	Pole	Light	Sign	Veg.	Terrace	Sky	Person	Rider	Car	Truck	Bus	Train	Motor	Bike	mIoU
Condition Attention Module (Stage1 Score)																				
Mix-Head	88.0	41.5	81.1	36.5	27.3	30.7	45.0	27.3	82.6	37.3	73.5	61.7	32.4	85.0	27.6	42.1	13.6	17.5	23.7	46.0
Sep-Head	83.6	41.1	76.9	31.8	29.4	30.5	44.2	30.8	83.1	28.6	77.4	60.7	32.4	84.9	25.0	40.5	11.0	17.9	30.2	45.3
SM-Head	85.1	41.3	79.0	33.7	30.2	31.2	45.3	30.3	82.7	30.1	75.8	61.2	31.9	84.8	25.1	39.1	9.6	19.8	29.5	45.6
CAM	89.9	46.4	82.5	37.7	29.3	34.1	45.7	25.4	83.1	33.5	76.8	62.3	32.9	85.3	28.3	39.9	2.6	23.3	37.4	47.2
Condition-Specific Adversarial Training (Building on CAM)																				
DAT-S1	84.0	38.0	80.1	28.2	29.8	28.9	44.9	31.7	82.7	28.9	77.4	61.2	30.7	82.2	22.2	30.4	17.3	22.8	39.8	45.3
CSAT-S1	89.9	46.4	82.5	37.7	29.3	34.1	45.7	25.4	83.1	33.5	76.8	62.3	32.9	85.3	28.3	39.9	2.6	23.3	37.4	47.2
DAT-S2	88.3	42.0	83.7	37.3	26.8	27.4	43.1	34.8	85.5	33.1	79.2	60.8	30.1	85.2	29.9	41.8	14.7	23.3	41.6	47.8
CSAT-S2	90.5	52.8	84.9	38.7	30.4	37.6	45.5	42.4	84.5	45.5	80.7	63.8	36.2	86.3	34.9	45.1	19.7	14.9	42.5	51.4
Attentive Pseudo-Label Assignment ($\lambda_p = 0.9, \mathcal{L}_{CE}^T$) (Baseline = DCAA-S1)																				
Baseline	89.9	46.4	82.5	37.7	29.3	34.1	45.7	25.4	83.1	33.5	76.8	62.3	32.9	85.3	28.3	39.9	2.6	23.3	37.4	47.2
MaxV	89.8	44.1	83.2	34.7	29.3	29.5	42.7	30.9	86.0	44.3	80.6	60.7	31.2	86.2	32.2	30.1	17.7	26.1	34.7	48.1
MeanV	89.6	48.4	83.6	38.8	30.1	27.1	38.9	28.4	86.1	43.0	79.8	62.4	29.4	85.8	37.0	32.7	16.7	32.7	34.9	48.7
APLA ⁺	90.5	47.0	83.7	35.0	17.0	31.2	43.6	36.6	85.9	44.3	77.7	63.8	34.3	88.2	30.6	38.7	8.6	34.9	54.9	49.8
Adversarial Ambivalence + APLA ($\lambda_p = 0.6$)																				
Baseline (\mathcal{L}_{CE}^T)	88.2	44.8	82.5	36.0	22.5	34.1	43.0	36.8	84.5	44.9	77.5	62.0	25.7	86.2	29.5	37.6	1.1	29.6	50.7	48.3
\mathcal{L}_{CE}^T	90.6	50.9	84.0	41.0	22.1	36.1	45.8	34.3	85.1	36.4	79.4	64.5	37.7	85.0	33.0	46.5	18.9	22.3	47.3	50.6
$\mathcal{L}_{CE}^T + \mathcal{L}_{adv}^T$	90.5	52.8	84.9	38.7	30.4	37.6	45.5	42.4	84.5	45.5	80.7	63.8	36.2	86.3	34.9	45.1	19.7	14.9	42.5	51.4

Table 3: Ablation Studies. Si: The i th stage results. APLA⁺: When studying APLA, we take λ_p as 0.9 to avoid introducing too many noisy pseudo-labels. MaxV and MeanV mean that we obtain the pseudo-labels via the simple max voting and mean voting operations.

Ablation Studies

We conduct ablation experiments based on the adaptation scenario GTA5→Cityscapes* to discuss our DCAA. We refer to the supplemental materials for other studies.

Condition Attention Module. We examine some other possible options that explicitly incorporate the weather priors, including Mix-Head, Sep-Head, and SM-Head. The differences are shown in Figure 3. From Table 3, either SM-head (45.6%) or Sep-Head (45.3%) even slightly decreases the performance of Mix-Head by 0.4% ~ 0.7%. This observation explicates that trivially incorporate weather priors may show negative influences. Instead, CAM (47.2%) effectively exploits the special characteristics and correlations of multi-modalities.

Condition-Specific Adversarial Training. We study the effectiveness of the standard adversarial training strategy and our progressive condition-specific adversarial training policy in minimizing the source and target discrepancy. In Table 3, DAT denotes that all the discriminators are standard domain classifiers to distinguish source features from target features. We find that the second-stage DAT-S2 (47.8%) only reaches a slight better mIoU over our first-stage CSAT-S1 (47.2%), and is marginally below CSAT-S2 (51.4%). This observation demonstrates that directly aligning source and target feature distributions is challenging when the target domain is heterogeneous.

Self-Training. We make comparisons with different baselines and variants to discuss our self-training policy. The scores and compared methods are presented in Table 3. We find that the produced pseudo-labels by MaxV and MeanV

can improve the baseline by 0.7% ~ 1.3%, while our APLA⁺ impressively yields a mIoU of 49.8% (+2.7%). This observation shows that trivial pseudo-label generation strategies may not sufficiently exploit the complementarities of multi-modal semantics. Furthermore, the revised \mathcal{L}_{CE}^T outperforms $\tilde{\mathcal{L}}_{CE}^T$ by a convincing margin (+2.3%) when the pseudo-labels are with many noises. At the same time, the proposed adversarial training for hard adaptation regions (\mathcal{L}_{adv}^T) can further boost the mIoU to 51.4%. The improvements demonstrate that the proposed adversarial ambivalence is potent to suppress the adverse impact of noisy pseudo-labels. It can also drive the network focusing on pushing hard adaptation regions to a shared feature space.

Conclusion

In the paper, we notice that images in a target domain may show diverse characteristics in typical domain adaptation studies. We clarify carefully reason about the multi-modalities may enhance the adaptation performance. Thus, we propose a novel adaptation approach DCAA that is empowered by an attentive progressive adversarial training mechanism and a self-training policy. The former migrate the source and target discrepancies in a progressive manner. The later first produces pseudo-labels for target images by investigating the correlations among target sub-domains, then exploiting the adversarial ambivalence to regularize the noisy pseudo-labels and constraint the hard adaptation regions. We conduct various comparisons and ablation studies to discuss our methods intensely. Experimental results on several challenging adaptation scenarios demonstrate the superiority of method in capturing high-performing models.

References

- Chen, C.; Chen, Z.; Jiang, B.; and Jin, X. 2019a. Joint domain alignment and discriminative feature learning for unsupervised deep domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3296–3303.
- Chen, M.; Xue, H.; and Cai, D. 2019. Domain adaptation for semantic segmentation with maximum squares loss. In *Proceedings of the IEEE International Conference on Computer Vision*, 2090–2099.
- Chen, Y.-C.; Lin, Y.-Y.; Yang, M.-H.; and Huang, J.-B. 2019b. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1791–1800.
- Chen, Z.; Zhuang, J.; Liang, X.; and Lin, L. 2019c. Blending-target domain adaptation by adversarial meta-adaptation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2248–2257.
- Choi, Y.; Choi, M.; Kim, M.; Ha, J.-W.; Kim, S.; and Choo, J. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8789–8797.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Du, L.; Tan, J.; Yang, H.; Feng, J.; Xue, X.; Zheng, Q.; Ye, X.; and Zhang, X. 2019. SSF-DAN: Separated Semantic Feature Based Domain Adaptation Network for Semantic Segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 982–991.
- Fu, H.; Gong, M.; Wang, C.; Batmanghelich, K.; Zhang, K.; and Tao, D. 2019. Geometry-consistent generative adversarial networks for one-sided unsupervised domain mapping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2427–2436.
- Gholami, B.; Sahu, P.; Rudovic, O.; Bousmalis, K.; and Pavlovic, V. 2020. Unsupervised multi-target domain adaptation: An information theoretic approach. *IEEE Transactions on Image Processing* 29: 3993–4002.
- Gong, M.; Xu, Y.; Li, C.; Zhang, K.; and Batmanghelich, K. 2019a. Twin auxiliary classifiers gan. In *Advances in neural information processing systems*, 1330–1339.
- Gong, R.; Li, W.; Chen, Y.; and Gool, L. V. 2019b. DLOW: Domain flow for adaptation and generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2477–2486.
- Hoffman, J.; Tzeng, E.; Park, T.; Zhu, J.-Y.; Isola, P.; Saenko, K.; Efros, A.; and Darrell, T. 2018. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, 1989–1998. PMLR.
- Hoffman, J.; Wang, D.; Yu, F.; and Darrell, T. 2016. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*.
- Hong, W.; Wang, Z.; Yang, M.; and Yuan, J. 2018. Conditional generative adversarial network for structured domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1335–1344.
- Hu, X.; Fu, C.-W.; Zhu, L.; and Heng, P.-A. 2019. Depth-attentional features for single-image rain removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8022–8031.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Kim, M.; and Byun, H. 2020. Learning Texture Invariant Representation for Domain Adaptation of Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12975–12984.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, 5542–5550.
- Li, G.; Kang, G.; Liu, W.; Wei, Y.; and Yang, Y. 2020. Content-consistent matching for domain adaptive semantic segmentation. In *European Conference on Computer Vision*, 440–456. Springer.
- Li, Y.; Yuan, L.; and Vasconcelos, N. 2019. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6936–6945.
- Liu, Z.; Miao, Z.; Pan, X.; Zhan, X.; Lin, D.; Yu, S. X.; and Gong, B. 2020. Open Compound Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12406–12415.
- Luo, Y.; Zheng, L.; Guan, T.; Yu, J.; and Yang, Y. 2019. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2507–2516.
- Mirza, M.; and Osindero, S. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Pan, F.; Shin, I.; Rameau, F.; Lee, S.; and Kweon, I. S. 2020. Unsupervised Intra-domain Adaptation for Semantic Segmentation through Self-Supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3764–3773.
- Pei, Z.; Cao, Z.; Long, M.; and Wang, J. 2018. Multi-adversarial domain adaptation. *arXiv preprint arXiv:1809.02176*.
- Richter, S. R.; Vineet, V.; Roth, S.; and Koltun, V. 2016. Playing for Data: Ground Truth from Computer Games. In Leibe, B.; Matas, J.; Sebe, N.; and Welling, M., eds., *European Conference on Computer Vision (ECCV)*, volume 9906 of LNCS, 102–118. Springer International Publishing.
- Ros, G.; Sellart, L.; Materzynska, J.; Vazquez, D.; and

- Lopez, A. M. 2016. The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3723–3732.
- Sakaridis, C.; Dai, D.; Hecker, S.; and Van Gool, L. 2018. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 687–704.
- Shin, I.; Woo, S.; Pan, F.; and Kweon, I. S. 2020. Two-phase Pseudo Label Densification for Self-training based Domain Adaptation. In *European Conference on Computer Vision*, 532–548. Springer.
- Sun, R.; Zhu, X.; Wu, C.; Huang, C.; Shi, J.; and Ma, L. 2019. Not all areas are equal: Transfer learning for semantic segmentation via hierarchical region selection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4360–4369.
- Tsai, Y.-H.; Hung, W.-C.; Schuler, S.; Sohn, K.; Yang, M.-H.; and Chandraker, M. 2018. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7472–7481.
- Tsai, Y.-H.; Sohn, K.; Schuler, S.; and Chandraker, M. 2019. Domain adaptation for structured output via discriminative patch representations. In *Proceedings of the IEEE International Conference on Computer Vision*, 1456–1465.
- Vu, T.-H.; Jain, H.; Bucher, M.; Cord, M.; and Pérez, P. 2019. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2517–2526.
- Wang, D.; Cui, P.; and Zhu, W. 2018. Deep Asymmetric Transfer Network for Unbalanced Domain Adaptation. In *AAAI*.
- Wang, Z.; Yu, M.; Wei, Y.; Feris, R.; Xiong, J.; Hwu, W.-m.; Huang, T. S.; and Shi, H. 2020. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12635–12644.
- Wu, Z.; Han, X.; Lin, Y.-L.; Gokhan Uzunbas, M.; Goldstein, T.; Nam Lim, S.; and Davis, L. S. 2018. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 518–534.
- Xiao, J.; Gu, S.; and Zhang, L. 2020. Multi-Domain Learning for Accurate and Few-Shot Color Constancy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3258–3267.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 418–434.
- Yang, J.; Xu, R.; Li, R.; Qi, X.; Shen, X.; Li, G.; and Lin, L. 2019. An Adversarial Perturbation Oriented Domain Adaptation Approach for Semantic Segmentation. *arXiv preprint arXiv:1912.08954*.
- Yang, Y.; and Soatto, S. 2020. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4085–4095.
- Yu, F.; Chen, H.; Wang, X.; Xian, W.; Chen, Y.; Liu, F.; Madhavan, V.; and Darrell, T. 2020a. BDD100K: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2636–2645.
- Yu, X.; Liu, T.; Gong, M.; Zhang, K.; Batmanghelich, K.; and Tao, D. 2020b. Label-noise robust domain adaptation. In *International Conference on Machine Learning*, 10913–10924. PMLR.
- Zhang, Q.; Zhang, J.; Liu, W.; and Tao, D. 2019. Category Anchor-Guided Unsupervised Domain Adaptation for Semantic Segmentation. In *Advances in Neural Information Processing Systems*, 433–443.
- Zhang, Y.; David, P.; and Gong, B. 2017. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, 2020–2030.
- Zhang, Y.; Qiu, Z.; Yao, T.; Liu, D.; and Mei, T. 2018. Fully convolutional adaptation networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6810–6818.
- Zhang, Y.; and Wang, Z. 2020. Joint Adversarial Learning for Domain Adaptation in Semantic Segmentation. In *AAAI*, 6877–6884.
- Zhao, S.; Gong, M.; Liu, T.; Fu, H.; and Tao, D. 2020. Domain Generalization via Entropy Regularization. volume 33.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene Parsing through ADE20K Dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017a. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.
- Zhu, J.-Y.; Zhang, R.; Pathak, D.; Darrell, T.; Efros, A. A.; Wang, O.; and Shechtman, E. 2017b. Toward multimodal image-to-image translation. In *Advances in neural information processing systems*, 465–476.
- Zou, Y.; Yu, Z.; Liu, X.; Kumar, B.; and Wang, J. 2019. Confidence regularized self-training. In *Proceedings of the IEEE International Conference on Computer Vision*, 5982–5991.
- Zou, Y.; Yu, Z.; Vijaya Kumar, B.; and Wang, J. 2018. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, 289–305.