

Focus on Your Target: A Dual Teacher-Student Framework for Domain-adaptive Semantic Segmentation

Xinyue Huo^{1,2} Lingxi Xie² Wengang Zhou¹ Houqiang Li¹ Qi Tian^{2*}

¹University of Science and Technology of China ²Huawei Inc.

xinyueh@mail.ustc.edu.cn 198808xc@gmail.com

{zhwg, lihq}@ustc.edu.cn tian.qil@huawei.com

Abstract

We study unsupervised domain adaptation (UDA) for semantic segmentation. Currently, a popular UDA framework lies in self-training which endows the model with two-fold abilities: (i) **learning** reliable semantics from the labeled images in the source domain, and (ii) **adapting** to the target domain via generating pseudo labels on the unlabeled images. We find that, by decreasing/increasing the proportion of training samples from the target domain, the ‘learning ability’ is strengthened/weakened while the ‘adapting ability’ goes in the opposite direction, implying a conflict between these two abilities, especially for a single model. To alleviate the issue, we propose a novel dual teacher-student (DTS) framework and equip it with a bidirectional learning strategy. By increasing the proportion of target-domain data, the second teacher-student model learns to ‘Focus on Your Target’ while the first model is not affected. DTS is easily plugged into existing self-training approaches. In a standard UDA scenario (training on synthetic, labeled data and real, unlabeled data), DTS shows consistent gains over the baselines and sets new state-of-the-art results of 76.5% and 75.1% mIoUs on $GTAv \rightarrow Cityscapes$ and $SYNTHIA \rightarrow Cityscapes$, respectively. The implementation is available at <https://github.com/xinyuehuo/DTS>.

1. Introduction

Deep neural networks [28] have been proven effective in learning from labeled, in-domain visual data, but their ability to adapt to unlabeled data or unknown domains is often not guaranteed. We delve into this topic by studying the unsupervised domain adaptation (UDA) problem for semantic segmentation, where densely (pixel-wise) labeled data are available for the source domain but only unlabeled images are provided for the target domain. The setting is useful

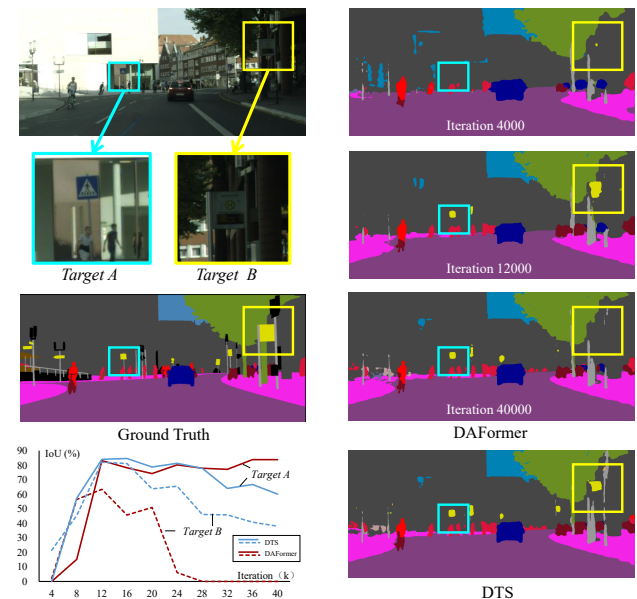


Figure 1. Segmentation results on a test image from Cityscapes (the model is trained on GTAv). There are two traffic signs with different styles, where *Target A* is similar to the training data in the source domain while *Target B* is not. The baseline (DAFormer) gradually improves the accuracy on *Target A* but achieves the best accuracy on *Target B* in the midst (at 12K iterations). Our method, DTS, achieves good accuracy for both targets and the trend is more stable throughout the training process (see the curves).

in real-world applications where new domains emerge frequently but annotating them is a major burden.

The state-of-the-art UDA approaches for semantic segmentation are mostly built upon a self-training framework. It first learns semantics from labeled data in the source domain and then generates pseudo labels in the target domain, so that the model adapts to the target domain by fitting the pseudo labels. Conceptually, the two steps of self-training endow the model with two-fold abilities, namely, **learning** knowledge (semantics) from the source domain and **adapt-**

*Corresponding Author.

ing the knowledge to the target domain.

However, these two abilities can conflict with each other. Figure 1 shows an example of two targets of the same class (traffic sign), where *Target A* is close to the typical training samples in the source domain, while *Target B* has a quite different appearance. Using the default setting of DAFormer [19] (a UDA baseline), the recognition accuracy of *Target A* gradually increases while that of *Target B* grows first but eventually drops to 0. Interestingly, if one increases the proportion of training samples from the target domain, the recognition quality of *Target B* will be better while the recognition accuracy of some unexpected classes (e.g., sidewalk) may drop significantly (more details and examples can be found in the Appendix). The conflict sets an obstacle to achieving higher UDA accuracy.

The above analysis prompts us to design a dual teacher-student (DTS) framework. The key is to train two individual models (with different network weights) so that the aforementioned conflict is alleviated. DTS can be built upon most existing self-training approaches in three steps: (1) making two copies of the original teacher-student model, (2) increasing the proportion of training data from the target domain in the second group, and (3) applying a bidirectional learning strategy so that two models supervise each other via pseudo labels. DTS allows the system to pursue a stronger **adapting** ability (with the second model) meanwhile the **learning** ability (guaranteed by the first model) is mostly unaffected.

We conduct semantic segmentation experiments with the standard UDA setting. Two synthetic datasets with dense labels, GTAv [38] and SYNTHIA [39], are used as the source domains, and a real dataset with only images, Cityscapes [6], is used as the target domain. We establish the DTS framework upon three recent self-training baselines, namely, DAFormer [19], HRDA [20], and MIC [21]. DTS improves the segmentation accuracy in every single trial, demonstrating its generalized ability across baselines and datasets. In particular, when integrated with MIC, the prior state-of-the-art, DTS reports 76.5% and 67.8% mIoUs on GTAv→Cityscapes and SYNTHIA→Cityscapes, respectively, setting new records for both benchmarks.

2. Related Work

Unsupervised domain adaptation (UDA), the task of transferring semantics from labeled source domains to unlabeled target domains, has been widely studied for many computer vision fields, e.g., classification [33, 58, 51], object detection [15, 24, 45], face recognition [25, 18, 43], semantic segmentation [17, 60], etc.

We briefly review two representative UDA approaches for semantic segmentation. **Adversarial learning** [47, 49] reduces the gap between the source and target domains by confusing a discriminator trained to judge which domain

the extracted visual features come from. It is closely related to style transfer methods [16, 55, 12, 56] that synthesize data to simulate the target domain. **Self-training** instead generates pseudo labels for the unlabeled target images and fits the model accordingly. There is a specific class of self-training methods [48, 10, 62] which further bridge the domain gap by creating mixed images that cover both domains [36, 1], inheriting benefits from style transferring. Many follow-up works [19, 20, 21], including this paper, were based on the self-training framework. Besides UDA, other data-efficient learning scenarios such as semi-supervised [22, 40] and few-shot [31, 35] learning also made use of self-training. Generating high-quality pseudo labels is crucial for self-training. Researchers developed various strategies, such as entropy minimization [4], class-balanced training [30], consistency regularization [21], auxiliary task learning [23, 53], etc., for this purpose.

A fundamental concept in self-training is **teacher-student optimization**. To boost the effect of teacher-student optimization, efforts were made in generating high-quality pseudo labels with stronger teacher models [46, 27, 29, 11] and/or preventing error transmission from the teacher model to the student model [26, 63, 22, 5]. There are also other improvements including using multiple student models [61], designing a better loss function [11], etc. Another closely related concept is co-training [37, 44], which facilitates cross-view consistency with a teacher-student loss. The idea of DTS is to alleviate the discrepancy between teacher and student models from different groups with an indirect consistency constraint.

3. Our Approach

We investigate the task of unsupervised domain adaptation (UDA) for semantic segmentation, where a labeled source domain \mathbb{S} and an unlabeled target domain \mathbb{T} are available. We denote the source dataset as $\mathcal{D}^{\mathbb{S}} = \{(\mathbf{x}_n^{\mathbb{S}}, \mathbf{y}_n^{\mathbb{S}})\}_{n=1}^{N_{\mathbb{S}}}$, where $\mathbf{x}_n^{\mathbb{S}}$ is an image and $\mathbf{y}_n^{\mathbb{S}}$ is the corresponding dense label map. The target dataset $\mathcal{D}^{\mathbb{T}}$ is another set of images without labels, denoted as $\{(\mathbf{x}_n^{\mathbb{T}})\}_{n=1}^{N_{\mathbb{T}}}$. The two domains share the same set of categories. The goal of UDA is to train a segmentation network $f(\mathbf{x}; \theta)$ that eventually works well in the target domain.

We first introduce a self-training baseline where a single teacher-student framework is established (Section 3.1). Then, we reveal a conflict when the model is facilitated to focus on the target domain (Section 3.2), which drives us to propose a dual teacher-student framework towards better UDA performance (Section 3.3).

3.1. Single Teacher-Student Baseline

We start with a standard self-training (i.e., teacher-student) baseline which trains a teacher model and a student model denoted as $f^{te}(\mathbf{x}; \theta^{te})$ and $f^{st}(\mathbf{x}; \theta^{st})$, respectively.

The teacher model is updated using the exponential moving average (EMA) strategy [46]:

$$\theta^{\text{te}} \leftarrow \theta^{\text{te}} \times \lambda + \theta^{\text{st}} \times (1 - \lambda), \quad (1)$$

where λ is a constant coefficient which is often close to 1. The student model is trained for segmentation in both the source and target domains. For the source images, semantic labels are available for supervised learning:

$$\mathcal{L}^{\mathbb{S}} = \mathcal{L}_{\text{CE}}(f^{\text{st}}(\mathbf{x}^{\mathbb{S}}), \mathbf{y}^{\mathbb{S}}), \quad (2)$$

where \mathcal{L}_{CE} denotes the pixel-wise cross-entropy loss. For the target images, the teacher model is used to generate a pseudo label $\tilde{\mathbf{y}}^{\mathbb{T}}$ for $\mathbf{x}^{\mathbb{T}}$, and the pixel-wise loss function is computed in a similar way:

$$\mathcal{L}^{\mathbb{T}} = \mathcal{L}_{\text{CE}}(f^{\text{st}}(\mathbf{x}^{\mathbb{T}}), \tilde{\mathbf{y}}^{\mathbb{T}}). \quad (3)$$

The overall loss is written as $\mathcal{L} = \mathcal{L}^{\mathbb{S}} + \mathcal{L}^{\mathbb{T}}$. In practice, a balancing mechanism is applied [48, 19]: all pixel-wise loss terms computed on the target domain is multiplied by a factor of γ , which is the proportion of the pseudo-label confidence surpassing a given threshold, τ .

Recently, a series of methods [48, 19, 20, 21] validated that introducing a mixed domain between the source and target domains is helpful for domain adaptation. A typical operation is named ClassMix [36] which samples mixed images by computing

$$\mathbf{x}_{i,j}^{\langle \mathbb{S}, \mathbb{T} \rangle} = \mathcal{M}(\mathbf{x}_i^{\mathbb{S}}, \mathbf{x}_j^{\mathbb{T}}) = \mathbf{x}_i^{\mathbb{S}} \odot \mathbf{M}_i^{\mathbb{S}} + \mathbf{x}_j^{\mathbb{T}} \odot (\mathbf{1} - \mathbf{M}_i^{\mathbb{S}}), \quad (4)$$

where $\langle \mathbb{S}, \mathbb{T} \rangle$ denotes a mixed domain, \odot is element-wise multiplication, and $\mathbf{M}^{\mathbb{S}}$ is a mask generated by randomly a few class-level masks from the $\mathbf{y}^{\mathbb{S}}$. Accordingly, the target domain in Eqn (3) and the overall loss is replaced with the mixed domain. Note that pseudo labels are only computed for the target portion of the mixed images – we refer the reader to [36] for technical details.

3.2. Conflict between Learning and Adapting

Although domain mixing improves the accuracy of UDA segmentation on several benchmarks, we observe that the trained model is easily biased towards the source domain and thus ignores the target domain to some extent. A typical example is shown in Figure 1.

We conjecture that it is the ‘over-focus’ on the source domain that prevents the model from achieving higher segmentation accuracy. To reveal this, we revisit the working mechanism of ClassMix. Compared to the original strategy that directly generates pseudo labels on the target images, ClassMix increases the proportion of data from the source domain (in a default setting of Eqn (4), an average of half target-domain pixels are replaced with source-domain pixels). This brings two differences. On the one hand, the

Label	$\langle \mathbb{S}, \mathbb{T} \rangle$	$\langle \mathbb{T}, \mathbb{T} \rangle$	<i>Supervised</i>
w/ pseudo labels	68.3	67.2	–
(pseudo acc.)	64.9	64.2	–
w/ <i>ground-truth</i>	74.4	76.9	77.8

Table 1. UDA segmentation mIoU (%) using DAFormer [19] on the GTAv→Cityscapes benchmark with different options. An *italic* number indicates a result that uses target-domain labels. The pseudo accuracy is computed on the training set of Cityscapes.

quality of pseudo labels is improved because part of them are inherited from the labels in the source domain; on the other hand, the focus on the target domain is weakened, increasing the risk of domain adaptation.

The above analysis prompts us to perform the ‘Focus on Your Target’ trail towards higher UDA accuracy. We replace the mixed domain of ClassMix, $\langle \mathbb{S}, \mathbb{T} \rangle$, with another domain, $\langle \mathbb{T}, \mathbb{T} \rangle$, in which the mixed samples are generated using two target images:

$$\mathbf{x}_{j_1, j_2}^{\langle \mathbb{T}, \mathbb{T} \rangle} = \mathcal{M}(\mathbf{x}_{j_1}^{\mathbb{T}}, \mathbf{x}_{j_2}^{\mathbb{T}}) = \mathbf{x}_{j_1}^{\mathbb{T}} \odot \mathbf{M}_{j_1}^{\mathbb{T}} + \mathbf{x}_{j_2}^{\mathbb{T}} \odot (\mathbf{1} - \mathbf{M}_{j_1}^{\mathbb{T}}), \quad (5)$$

and the pseudo label is adjusted accordingly.

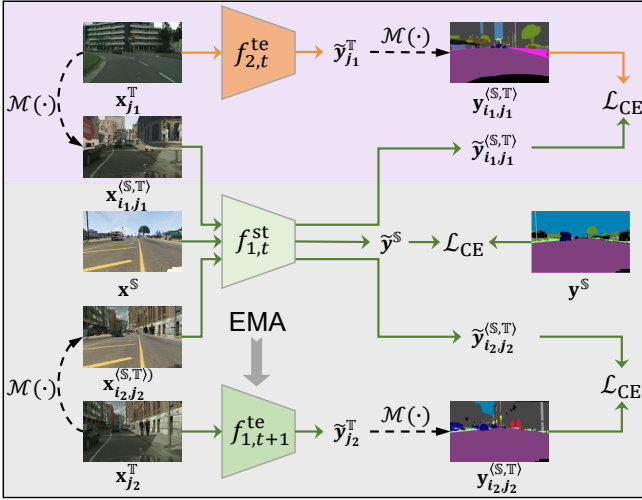
In Table 1, we run DAFormer [19] with either $\langle \mathbb{S}, \mathbb{T} \rangle$ or $\langle \mathbb{T}, \mathbb{T} \rangle$ for UDA segmentation on GTAv→Cityscapes. Compared to $\langle \mathbb{S}, \mathbb{T} \rangle$, $\langle \mathbb{T}, \mathbb{T} \rangle$ achieves an inferior mIoU of 67.2% (a 1.3% drop). But, interestingly, when ground-truth labels are used for the mixed images, the mIoU surpasses the original option by 2.5%, with merely a 0.9% mIoU lower than the supervised baseline. In other words, ‘Focus on Your Target’ improves the upper bound, but the models cannot achieve the upper bound because the pseudo labels deteriorate: it is a side effect of decreasing the proportion of training data from the source domain.

In summary, a strong algorithm for UDA shall be equipped with two critical abilities: (i) **learning** reliable semantics from the source domain, and (ii) **adapting** the learned semantics to the target domain. However, in the above self-training framework where a single teacher-student model is used, there seems a conflict between these two abilities, *e.g.*, increasing the proportion of target-domain data strengthens the ‘adapting ability’ but weakens the ‘learning ability’, setting an obstacle to better performance. In what follows, we offer a solution by training two teacher-student models simultaneously.

3.3. Dual Teacher-Student Framework

The proposed dual teacher-student (DTS) framework is illustrated in Figure 2. We denote the models used in the two teacher-student groups as $f_1^{\text{te}}(\mathbf{x}; \theta_1^{\text{te}})$, $f_1^{\text{st}}(\mathbf{x}; \theta_1^{\text{st}})$ and $f_2^{\text{te}}(\mathbf{x}; \theta_2^{\text{te}})$, $f_2^{\text{st}}(\mathbf{x}; \theta_2^{\text{st}})$, respectively, where we use a subscript (1 or 2) to indicate the group ID. As described above, we expect the first teacher-student model to stick to **learning** reliable semantics from the source domain, while the

Step 1 & Step 2



Step 3 & Step 4

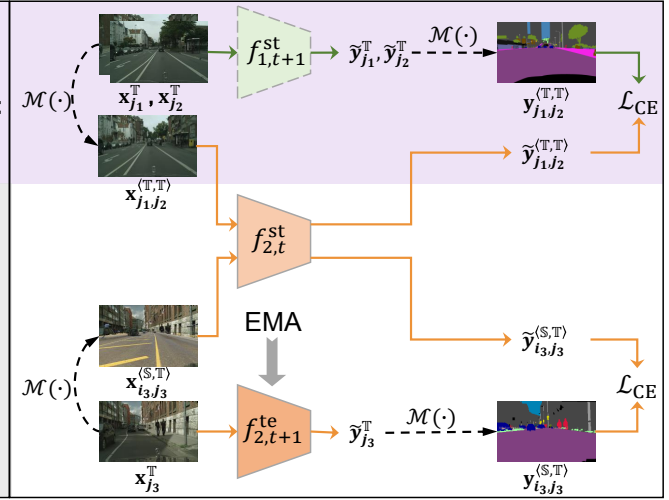


Figure 2. An illustration of the proposed DTS framework, where we use *Setting B* as an example. The left and right parts describe the optimization of the first and second group models, respectively. The bottom-left part (in gray) corresponds to a single teacher-student framework where ClassMix is used for domain mixing, and the top part (in purple) shows the bidirectional learning strategy, *i.e.*, models from the other group are used for pseudo label generation.

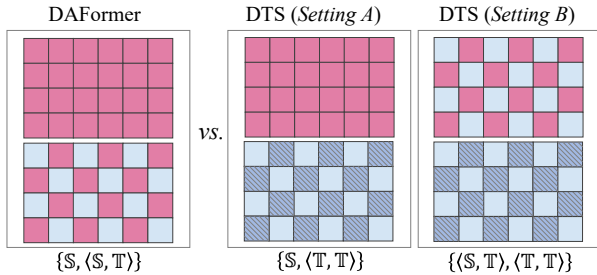


Figure 3. An illustration of different combinations of training data. We use magenta to represent source-domain patches and two types of blue to represent target-domain patches from different images. In practice, the images are not partitioned into regular grids, but into semantic regions according to the (real or pseudo) labels.

second teacher-student model learns to ‘Focus on Your Target’ to enhance the **adapting** ability.

We achieve the goal by using different data proportions for the two groups. An illustration is provided in Figure 3. The first group remains the same as in the baseline, *i.e.*, half of the training samples are simply source images and another half are ClassMix-ed images covering the source and target domains. Using the previously defined notations, we denote the combination as $\{S, \langle S, T \rangle\}$. To increase the proportion of the target-domain data, we first change the mixture of data from $\langle S, T \rangle$ to $\langle T, T \rangle$ (*i.e.*, same as that used in the ‘Focus on Your Target’ trial), and then replace the pure source domain with $\langle S, T \rangle$. That said, we create two settings for data combination, $\{S, \langle T, T \rangle\}$ and $\{\langle S, T \rangle, \langle T, T \rangle\}$. We name them *Setting A* and *Setting B* of the second group

throughout the remainder of this paper. Intuitively, the proportion of the target-domain data is increased from around 25% in the first group (as well as the baselines) to around 50%/75% using *Setting A/Setting B* in the second group.

There are two teacher models and two student models to be optimized. We perform iterative optimization and propose a **bidirectional learning** strategy so that each teacher-student model propagates information to the other by contributing to the generation of pseudo labels. The detailed training process is shown in Algorithm 1.

Note that, if $f_{2,t}^{te}(\cdot)$ is not used in Step 2 (*i.e.*, all pseudo labels are generated using $f_{1,t+1}^{st}(\cdot)$), the information is only propagated from the first teacher-student model to the second one, hence the optimization strategy degenerates from bidirectional to unidirectional. The advantage of the bidirectional strategy and other design choices will be diagnosed in the experimental section.

3.4. Technical Details

Choosing the proper option for data combination. We compute the averaged confidence of pseudo labels generated by the corresponding teacher models, $f_1^{st}(\cdot)$ and $f_2^{te}(\cdot)$. Following [48, 19], γ is defined as the proportion of pixels with the max class score surpassing a given threshold, τ . For each training image \mathbf{x} , we use $\gamma_1^{st}(\mathbf{x})$ and $\gamma_2^{te}(\mathbf{x})$ to denote the average γ values over all pixels when the teacher models are $f_1^{st}(\cdot)$ and $f_2^{te}(\cdot)$, respectively. Then, we estimate the chance that $f_2^{te}(\cdot)$ generates better pseudo labels than $f_1^{st}(\cdot)$:

$$\text{Prob} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}^\tau} [\mathbb{I}(\gamma_2^{te}(\mathbf{x}) > \gamma_1^{st}(\mathbf{x}))], \quad (6)$$

Algorithm 1: DTS: Dual Teacher-Student

Input : Source data $\{\mathbf{x}^{\mathcal{S}}, \mathbf{y}^{\mathcal{S}}\}$, target data $\{\mathbf{x}^{\mathcal{T}}\}$, the total iteration number T ;

- 1 **repeat**
- 2 Update $f_{1,t+1}^{\text{te}}(\cdot)$ using the EMA of $f_{1,t}^{\text{st}}(\cdot)$;
- 3 Get pseudo labels $\tilde{\mathbf{y}}^{\mathcal{T}}$ by $f_{1,t+1}^{\text{te}}(\cdot)$ and $f_{2,t}^{\text{te}}(\cdot)$;
- 4 Generate mixed data $\{\mathbf{x}^{\langle \mathcal{S}, \mathcal{T} \rangle}, \mathbf{y}^{\langle \mathcal{S}, \mathcal{T} \rangle}\}$, and train $f_{1,t+1}^{\text{st}}(\cdot)$ on $\{\mathcal{S}, \langle \mathcal{S}, \mathcal{T} \rangle\}$;
- 5 Update $f_{2,t+1}^{\text{te}}(\cdot)$ using the EMA of $f_{2,t}^{\text{st}}(\cdot)$;
- 6 Get pseudo labels $\tilde{\mathbf{y}}^{\mathcal{T}}$ by $f_{1,t+1}^{\text{st}}(\cdot)$ and $f_{2,t+1}^{\text{te}}(\cdot)$;
- 7 **if Setting A:** Generate data $\{\mathbf{x}^{\langle \mathcal{T}, \mathcal{T} \rangle}, \mathbf{y}^{\langle \mathcal{T}, \mathcal{T} \rangle}\}$, and Train $f_{2,t+1}^{\text{st}}(\cdot)$ on $\{\mathcal{S}, \langle \mathcal{T}, \mathcal{T} \rangle\}$;
- 8 **if Setting B:** Generate data $\{\mathbf{x}^{\langle \mathcal{S}, \mathcal{T} \rangle}, \mathbf{y}^{\langle \mathcal{S}, \mathcal{T} \rangle}, \mathbf{x}^{\langle \mathcal{T}, \mathcal{T} \rangle}, \mathbf{y}^{\langle \mathcal{T}, \mathcal{T} \rangle}\}$, and train $f_{2,t+1}^{\text{st}}(\cdot)$ on them;
- 9 $t \leftarrow t + 1$;
- 10 **until** $t = T$;

Return: $f_{2,T}^{\text{st}}(\cdot)$.

where $\mathbb{I}(\cdot)$ is the indicator function that outputs 1 if the statement is true and 0 otherwise. Note that the estimation is performed individually for two training procedures using *Setting A* and *Setting B*. Finally, we choose the setting with a higher chance that $f_2^{\text{te}}(\cdot)$ is more confident because $f_2^{\text{st}}(\cdot)$ is eventually used for inference. We shall see in Section 4.3 that this criterion always leads to the better choice.

Generating the mask. In the domain mixing operation, we compute the mask \mathbf{M} by choosing the class label of one image. If the $\langle \mathcal{S}, \mathcal{T} \rangle$ is used, the source image with ground-truth is chosen; if $\langle \mathcal{T}, \mathcal{T} \rangle$ is used, the image with a larger average value of γ is chosen.

Generating pseudo labels. The details are slightly different between *Setting A* and *Setting B*. When the batch size is k , the number of target images used for the second group in *Setting A* and *Setting B* are $2k$ and $3k$, respectively. For *Setting A*, k pseudo labels used are generated by $f_{1,t+1}^{\text{st}}(\cdot)$ and other k by $f_{2,t+1}^{\text{te}}(\cdot)$. For *Setting B*, the $2k$ pseudo labels used in $\langle \mathcal{T}, \mathcal{T} \rangle$ are generated by $f_{1,t+1}^{\text{st}}(\cdot)$ and the k labels used in $\langle \mathcal{S}, \mathcal{T} \rangle$ are generated by $f_{2,t+1}^{\text{te}}(\cdot)$.

4. Experiments

4.1. Datasets and Settings

Datasets. We train the models on a labeled source domain and an unlabeled target domain and test them on the target domain. We follow the convention to set the source domain to be a synthesized dataset (either GTA_v [38] which contains 24,966 densely labeled training images at a resolution of 1914×1052 or SYNTHIA [39] which contains 9,400 densely labeled training images at a resolution of

1280×760), and the target dataset to be a real dataset (Cityscapes [6] where only the finely-annotated subset is used, containing 2,975 training images and 500 validation images at a resolution of 2048×1024). Following the convention to report the mean IoU (mIoU) over 19 classes when GTA_v is used as the source domain, and 13 or 16 classes when SYNTHIA is the source domain.

Training details. Our framework is applied to different self-training frameworks. Specifically, we follow the DAFormer series [19, 20, 21] where the encoder is a MiT-B5 model pre-trained on ImageNet-1K [7] and the decoder involves a context-aware feature fusion module. We train the framework on a single NVIDIA Tesla-V100 GPU for 40K iterations with a batch size of 2 for both groups. We use an AdamW optimizer [34] with learning rates of 6×10^{-5} for the encoder and 6×10^{-4} for the decoder, a weight decay of 0.01, and linear warm-up of the learning rate for the first 1.5K iterations. The input image size is always 512×512 . During the training phase, we apply random rescaling and cropping to the source and target data, as well as color jittering and Gaussian blur as strong augmentations for the mixed data. The EMA coefficient, λ , is set to be 0.999. The pseudo-label confidence threshold, τ , is 0.968 as in [48].

4.2. Comparison to the State-of-the-art Methods

We compare DTS to state-of-the-art UDA approaches on the GTA_v→Cityscapes and SYNTHIA→Cityscapes benchmarks. Results are summarized in Table 2 and Table 3, respectively. We establish DTS upon three recently published works, namely, DAFormer [19], HRDA [20], and MIC [21].

DTS reports consistent gains in terms of segmentation accuracy beyond all three baselines. On GTA_v→Cityscapes and SYNTHIA→Cityscapes (13 classes), the mIoU gains over DAFormer, HRDA, MIC are 2.9%, 1.4%, 0.6% and 4.0%, 1.0%, 1.1% respectively. The gain generally becomes smaller when the baseline is stronger, which is as expected. Specifically, when added to MIC, the previous best approach, DTS achieves new records on these benchmarks: a 76.5% mIoU on GTA_v→Cityscapes and 67.8% and 75.1% mIoUs on SYNTHIA→Cityscapes with 16 and 13 classes, respectively.

Comparison to FST. We compare DTS to another plugin approach, FST [8]. On DAFormer, the gains brought by FST on the two benchmarks (SYNTHIA using 13 classes) are 1.0% and 1.3%, while the numbers for DTS are 2.9% and 4.0%, respectively. That said, besides the contribution of FST that allows us to ‘learn from future’, the higher accuracy also owes to the design of the DTS framework that alleviates the conflict of **learning** and **adapting**.

Class-level analysis. We investigate which classes are best improved by DTS. On GTA_v→Cityscapes, we detect two classes, **sign** and **truck**, on which the IoU gain on the three baselines is always higher than the mIoU gain. In the

Method	Road	Sidewalk	Building	Wall	Fence	Pole	Light	Sign	Veg	Terrain	Sky	Person	Rider	Car	Truck	Bus	Train	Motor	Bike	mIoU
AdaptSegNet [49]	86.5	36.0	79.9	23.4	23.3	23.9	35.2	14.8	83.4	33.3	75.6	58.5	27.6	73.7	32.5	35.4	3.9	30.1	28.1	42.4
PatchAligne [50]	92.3	51.9	82.1	29.2	25.1	24.5	33.8	33.0	82.4	32.8	82.2	58.6	27.2	84.3	33.4	46.3	2.2	29.5	32.3	46.5
FDA [56]	92.5	53.3	82.4	26.5	27.6	36.4	40.6	38.9	82.3	39.8	78.0	62.6	34.4	84.9	34.1	53.1	16.9	27.7	46.4	50.5
DACS [48]	89.9	39.7	87.9	39.7	39.5	38.5	46.4	52.8	88.0	44.0	88.8	67.2	35.8	84.5	45.7	50.2	0.0	27.3	34.0	52.1
ProDA [59]	87.8	56.0	79.7	46.3	44.8	45.6	53.5	53.5	88.6	45.2	82.1	70.7	39.2	88.8	45.5	50.4	1.0	48.9	56.4	57.5
DAP [23]	94.5	63.1	89.1	29.8	47.5	50.4	56.7	58.7	89.5	50.2	87.0	73.6	38.6	91.3	50.2	52.9	0.0	50.2	63.5	59.8
CPSL [30]	92.3	59.9	84.9	45.7	29.7	52.8	61.5	59.5	87.9	41.6	85.0	73.0	35.5	90.4	48.7	73.9	26.3	53.8	53.9	60.8
DAFormer [19]	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
DAFormer+FST [8]	95.3	67.7	89.3	55.5	47.1	50.1	57.2	58.6	89.9	51.0	92.9	72.7	46.3	92.5	78.0	81.6	74.4	57.7	62.6	69.3
DAFormer+CAMix [62]	96.0	73.1	89.5	53.9	50.8	51.7	58.7	64.9	90.0	51.2	92.2	71.8	44.0	92.8	78.7	82.3	70.9	54.1	64.3	70.0
DAFormer+DTS (ours)	96.8	76.0	90.0	56.5	50.1	51.9	57.4	65.1	90.5	50.5	92.4	73.5	46.9	93.0	80.4	85.2	74.0	58.7	63.5	71.2
HRDA [20]	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
HRDA+DTS (ours)	96.7	76.3	91.3	60.7	55.9	59.9	67.3	72.5	91.8	50.1	94.4	80.7	56.8	94.4	86.0	86.0	73.9	65.7	68.4	75.2
MIC [21]	97.4	80.1	91.7	61.2	56.9	59.7	66.0	71.3	91.7	51.4	94.3	79.8	56.1	94.6	85.4	90.3	80.4	64.5	68.5	75.9
MIC+DTS (ours)	97.0	80.4	91.8	60.6	58.7	61.7	67.9	73.2	92.0	45.4	94.3	81.3	58.7	95.0	87.9	90.7	82.2	65.7	69.0	76.5

Table 2. UDA segmentation accuracy (mIoU, %) on the GTAv→Cityscapes benchmark. The highest number in each column is marked with **bold**. The results in the first group are obtained on a ResNet101 backbone and the others are obtained on a MiT-B5 backbone.

Method	Road	Sidewalk	Building	Wall*	Fence*	Pole*	Light	Sign	Veg	Sky	Person	Rider	Car	Bus	Motor	Bike	mIoU	mIoU*
PatchAlign [50]	82.4	38.0	78.6	8.7	0.6	26.0	3.9	11.1	75.5	84.6	53.5	21.6	71.4	32.6	19.3	31.7	40.0	46.5
AdaptSegNet [49]	84.3	42.7	77.5	-	-	-	4.7	7.0	77.9	82.5	54.3	21.0	72.3	32.2	18.9	32.3	-	46.7
FDA [56]	79.3	35.0	73.2	-	-	-	19.9	24.0	61.7	82.6	61.4	31.1	83.9	40.8	38.4	51.1	-	52.5
DACS [48]	80.6	25.1	81.9	21.5	2.9	37.2	22.7	24.0	83.7	90.8	67.6	38.3	82.9	38.9	28.5	47.6	-	54.8
ProDA [59]	87.8	45.7	84.6	37.1	0.6	44.0	54.6	37.0	88.1	84.4	74.2	24.3	88.2	51.1	40.5	45.6	-	62.0
DAP [23]	84.2	46.5	82.5	35.1	0.2	46.7	53.6	45.7	89.3	87.5	75.7	34.6	91.7	73.5	49.4	60.5	-	64.3
CPSL [30]	87.2	43.9	85.5	33.6	0.3	47.7	57.4	37.2	87.8	88.5	79.0	32.0	90.6	49.4	50.8	59.8	-	65.3
DAFormer [19]	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
DAFormer+FST [8]	88.3	46.1	88.0	41.7	7.3	50.1	53.6	52.5	87.4	91.5	73.9	48.1	85.3	58.6	55.9	63.4	61.9	68.7
DAFormer+CAMix [62]	87.4	47.5	88.8	-	-	-	55.2	55.4	87.0	91.7	72.0	49.3	86.9	57.0	57.5	63.6	-	69.2
DAFormer+DTS (ours)	88.5	51.4	88.0	37.2	4.8	51.4	59.4	59.0	88.1	92.7	71.3	51.3	89.2	67.7	58.4	62.8	63.8	71.4
HRDA [20]	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
HRDA+DTS (ours)	88.5	53.1	88.7	44.0	5.8	59.2	68.3	60.7	87.7	93.4	79.5	46.4	91.3	64.2	66.8	65.9	66.5	73.4
MIC [21]	86.6	50.5	89.3	47.9	7.8	59.4	66.7	63.4	87.1	94.6	81.0	58.9	90.1	61.9	67.1	64.3	67.3	74.0
MIC+DTS (ours)	89.1	54.9	89.0	39.1	8.7	61.6	67.4	64.3	88.8	94.0	82.2	60.7	89.6	62.6	68.5	64.9	67.8	75.1

Table 3. UDA segmentation accuracy (mIoU, %) on the SYNTHIA→Cityscapes benchmark. The highest number in each column is marked with **bold**. The results in the first group are obtained on a ResNet101 backbone and the others are obtained on a MiT-B5 backbone. The mIoU and mIoU* columns denote the averaged IoU over 16 and 13 classes, where the latter does not contain the classes marked with *.

Appendix material, we show that these classes suffer larger variations between the source-domain and target-domain appearances, and hence the learned features are more difficult to transfer. They are largely benefited by the ‘Focus on Your Target’ ability of DTS.

Visualization. We visualize the segmentation results of DAFormer and DAFormer+DTS in Figure 4. For the real images, the advantage of DTS is significant on the classes that are easily confused with others (e.g., sidewalk which is visually similar to road) or suffer large domain shifts (e.g., sign or truck). For the mixed images, we find that the baseline model is easily confused by the unusual stitching and/or discontinuous textures because it has never seen such pat-

terns in the source domain – in such cases, the model of the second group is generally better and sometimes it can propagate the advantages to the model of the first group, revealing the advantage of bidirectional learning.

4.3. Diagnostic Studies

Ablative studies. We first diagnose the proposed approach by ablating the contribution of three components, namely, applying the dual teacher-student (DTS) framework, increasing the proportion of data from the target domain (i.e., ‘Focus on Your Target’), and adopting the bidirectional learning strategy. Results are summarized in Table 4. Note that some combinations are not valid, e.g.,

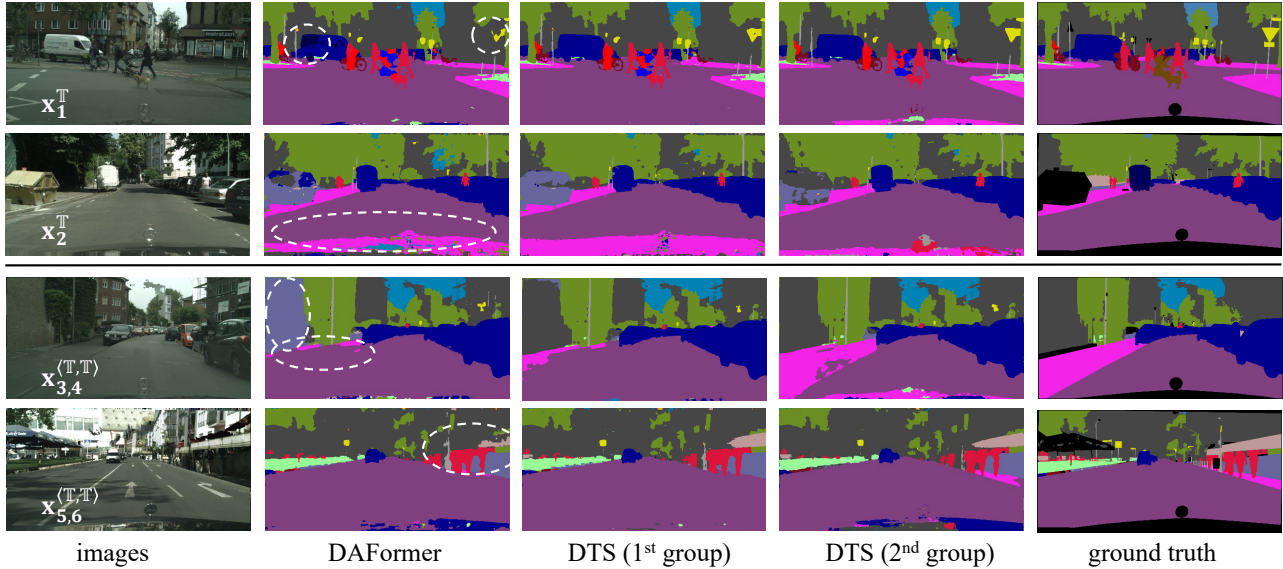


Figure 4. Visualization of the segmentation results of DAFormer [19] and DTS. The top two rows display two real images and the bottom two rows display two synthesized images from $\langle T, T \rangle$. In each group, the first example is from GTAv→Cityscapes and the second one is from SYNTHIA→Cityscapes. The white, dashed circles highlight the regions that are significantly improved by DTS.

Setting	Choice and Accuracy				
Dual Teacher-Student		✓	✓	✓	✓
'Focus on Your Target'	✓	✓			✓
Bidirectional Learning				✓	✓
GTA _v →Cityscapes	68.3	67.2	69.7	68.6	71.2
SYNTHIA→Cityscapes	60.9	58.4	63.3	61.2	63.7

Table 4. Ablative studies of DTS with respect to the UDA segmentation accuracy (mIoU, %) on both benchmarks. All these experiments are based on DAFormer [19]. For SYNTHIA→Cityscapes, the mIoUs are computed over 16 classes. Note that DTS alone (without the latter two options) is equivalent to the baseline (the leftmost column) and is simply ignored in the table.

solely applying DTS (without adjusting the proportion of training data nor applying bidirectional learning) does not make any difference from the baseline; also, bidirectional learning must be built upon DTS.

From Columns 1–3, one can learn the importance of the DTS framework. When DTS is not established, simply increasing the proportion of target data (Col 2) incurs an accuracy drop on both GTA_v→Cityscapes and SYNTHIA→Cityscapes. This validates our motivation, *i.e.*, there is a conflict between the **learning** and **adapting** abilities. However, after the dual teacher-student framework is established (Col 3, where the first teacher-student group focuses on **learning** and the second group focuses on **adapting**), the conflict is weakened and the accuracy gain becomes clear (1.6% on GTA_v→Cityscapes and 2.4% on SYNTHIA→Cityscapes).

	Baseline	DAFormer	HRDA	MIC
GTA _v	Setting A	70.3	74.7	76.5
	Setting B	71.2	75.2	76.2
SYNTHIA	Setting A	63.8	66.5	67.8
	Setting B	61.2	66.0	67.2

Table 5. The results of DTS with different data mixing options. For SYNTHIA→Cityscapes, the mIoUs are computed over 16 classes.

From Columns 3–5, one can observe the contribution of 'Focus on Your Target' (*i.e.*, increasing the proportion of target data) and bidirectional learning on top of the DTS framework. The best practice is obtained when all the modules are used (Col 5), aligning with the results reported in Tables 2 and 3. Interestingly, when bidirectional learning is adopted alone (Col 4), the accuracy gain is much smaller than that of using 'Focus on Your Target' alone, because the two groups are using the same proportion of data mixing, and thus the bidirectional learning carries limited complementary information.

Choose the option for data combination. We validate that the data combination option can be chosen by comparing the pseudo-label confidence, as described in Section 3.4. We plot the Prob curves throughout training procedure of Setting A and Setting B in Figure 5. On GTA_v→Cityscapes, the Prob value of Setting B is eventually higher than that of Setting A, and DTS with Setting B outperforms DTS with Setting A by a 0.9% mIoU (see Table 5). However, on SYNTHIA→Cityscapes, the Prob value of Setting A is always higher than that of Setting

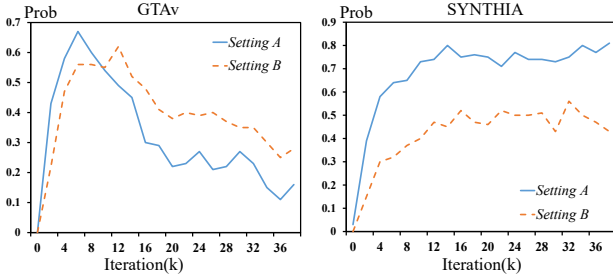


Figure 5. The curves of Prob computed on the GTAv→Cityscapes and SYNTHIA→Cityscapes benchmarks. DAFormer is used as the baseline. These curves suggest that, on DAFormer, *Setting B* is better for GTAv→Cityscapes while *Setting A* is better for SYNTHIA→Cityscapes.

Options	GTAvg	SYNTHIA
{⟨T, T⟩}	69.1	61.1
{⟨S, T⟩}	69.3	59.4
{S, ⟨T, T⟩}	70.3	63.8
{⟨S, T⟩, ⟨T, T⟩}	71.2	61.2

Table 6. The results of DTS with more data mixing options. We show the mIoU of the 16 classes in SYNTHIA and make this comparison based on the DAFormer baseline.

B, and DTS with *Setting A* outperforms DTS with *Setting B* by a 2.4% mIoU. Table 5 also shows the segmentation accuracy on the other two baselines. All choices of advantageous results align with the aforementioned criterion. From an intuitive view, *Setting B* has a smaller portion of source data, so it works better in the GTAv→Cityscapes scenario because GTAv has a smaller domain gap to Cityscapes than SYNTHIA, which can be validated by vanilla transfer experiments (GTAvg→Cityscapes: 46.2%, SYNTHIA→Cityscapes: 41.3%).

Except for *Setting A* and *Setting B*, there are two other options to increase the focus on the target domain as well, namely {⟨T, T⟩} and {⟨S, T⟩}. The two options do not include any mixing data and have a disadvantage compared to the previous two combinations. {⟨T, T⟩} means that the second teacher-student group achieves the ability of **learning** only from the pseudo labels provided by $f_1^{st}(x; \theta_1^{st})$. As a result, noise can easily persist without accurate supervision from the source data. {⟨S, T⟩} has a similar data proportion as {S, ⟨T, T⟩} but cannot solve the problem of inconsistency between the training and testing data. Using ⟨T, T⟩ as the training domain is an effective way to bridge the gap between the training and testing phases. Therefore, the two combinations containing two different types of data used in DTS are more suitable for balancing the **learning** and **adapting** abilities.

Models for pseudo label generation. We study the impact of different models for generating pseudo labels. Re-

Group #1		Group #2		mIoU
$f_{2,t}^{st}(\cdot)$	$f_{2,t}^{te}(\cdot)$	$f_{1,t+1}^{st}(\cdot)$	$f_{1,t+1}^{te}(\cdot)$	
	✓		✓	68.5 [†]
✓			✓	69.0
	✓		✓	69.5
✓		✓		70.2
	✓	✓		71.3

Table 7. UDA segmentation accuracy (mIoU, %) using different models to generate pseudo labels on GTAv→Cityscapes. [†]: In the first row, each group only uses the teacher model from the other group; in other cases, each group also uses its own teacher models (not presented in the table).

Method	GTAvg→City.		SYNTHIA→City.	
	DACS	DAFormer	DACS	DAFormer
w/o DTS	53.9	56.0	54.2	53.5
w DTS	55.2	59.9	56.6	56.9

Table 8. UDA segmentation accuracy (mIoU, %) on two CNN-based backbones and two benchmarks.

sults are summarized in Table 7. When each teacher-student group only uses the teacher model from another group (the first row), there is almost no accuracy gain (the baseline is 68.3%), indicating the necessity of the supervision signals from its own teacher.

Then, we discuss the situations where each group uses its teacher and one model from the other group. For the second teacher-student group, using $f_{1,t+1}^{st}(\cdot)$ as the teacher is superior to using $f_{1,t+1}^{te}(\cdot)$, because the former model is optimized one more iteration and thus produces better semantics (*i.e.*, ‘learning from future’). For the first teacher-student group, only models from the current iteration of the second group can be used, and $f_{2,t}^{te}(\cdot)$ always works better than $f_{2,t}^{st}(\cdot)$, implying the advantage of EMA.

DTS applied to other backbones. To show that DTS adapts to different network backbones (*e.g.*, CNNs), we follow the previous approaches [56, 48, 23] to inherit a ResNet101 [14] pre-trained on ImageNet and MSCOCO and train it using DeepLabV2 [2]. The other settings remain unchanged. Results are listed in Table 8. Still, DTS brings consistent accuracy gains, improving DACS and DAFormer by 1.3% and 3.9% on GTAv→Cityscapes, and by 2.4% and 3.4% on SYNTHIA→Cityscapes. These experiments validate that DTS is a generalized UDA framework that does not rely on specific network architectures.

Comparison of training time and compute. DTS trains two teacher-student groups, so the computational cost is about 2× of the baseline (*e.g.*, on a single V100 GPU, DAFormer/DTS needs 13.5/30 hours). The inference cost remains unchanged. We further compare DTS and DAFormer under the same training cost. DTS with half (20K) iterations achieves favorable accuracy (68.6% and

	GTA _v 2BDD	GTA _v 2IDD	CS2BDD	CS2IDD
Baseline	46.3	54.7	53.7	51.6
DTS	53.6	59.6	57.5	56.1
	CS2FZ	CS2DZ		
Baseline	39.7	34.0	42.1*	43.6*
DTS	40.4	37.6	45.2	45.6

Table 9. UDA segmentation accuracy (mIoU, %) on more driving datasets. ‘BDD’, ‘IDD’, ‘FZ’, ‘DZ’, and ‘CS’ for Berkeley Driving Dataset, India Driving Dataset, Foggy Zurich, Dark Zurich, and Cityscapes, respectively. Baselines are DAFormer except for the last two columns in CS2DZ, which are HRDA and MIC. *: are numbers inherited from [20, 21].

61.8% on GTA_v→Cityscapes and SYNTHIA→Cityscapes) over DAFormer with full (40K) iterations (68.3% and 60.9%), meanwhile, DTS with full (40K) iterations reports 71.2% and 63.8%, surpassing DAFormer with doubled (80K) iterations (69.3% and 61.2%).

Performance on more driving datasets. We conduct experiments from synthesized and real datasets to more driving datasets, *i.e.*, Berkeley Driving Dataset (BDD) [57], India Driving Dataset (IDD) [52], Foggy Zurich (FZ) [41], and Dark Zurich (DZ) [42]. The comparable results on validation sets of target domain datasets are summarized in Table 9. DTS reports consistent improvement in these transfer scenarios. On the test set of CS2DZ, DTS is also superior to DAFormer (54.2% vs 53.8%) and HRDA (58.9% vs 55.9%).

4.4. Semi-supervised Semantic Segmentation

Dataset. As a generalized self-training framework, DTS can also be applied to semi-supervised segmentation. We perform semi-supervised segmentation on the widely used PASCAL VOC dataset [9], which comprises 20 object classes and a background class. The training, validation, and testing sets contain 1,464, 1,449, and 1,456 images, respectively. To align with prior works, we use the augmented set, which contains 10,582 images, as the complete training set. During the training, we apply random flipping and cropping to augment the image data and resize the image to 513×513 before feeding it to the models.

Implementation. We perform all experiments using the DeepLabV3+ [3] network with a ResNet-101 backbone pre-trained on ImageNet. An SGD optimizer, with a learning rate of 0.0001 for the encoder and 0.001 for the decoder, and a weight decay of 0.0001, is used. The batch size of both labeled and unlabeled data is 16 for each group teacher-student. We train the models for 40 epochs and report the average of three runs for all the results. Since there is no domain gap between the labeled and unlabeled parts, only a dual teacher-student framework and bidirectional learning are applied in the semi-supervised experiments.

Methods	1/16 (662)	1/8 (1323)	1/4 (2646)
Mean Teacher [46]	71.3	75.4	76.5
FST [8]	73.9	76.1	78.1
USRN [13]	72.3	-	-
PS-MT [32]	75.5	78.2	78.7
U ² PL [54]	77.2	79.0	79.3
DTS (ours)	74.8	77.1	77.3
U ² PL + DTS (ours)	78.9	80.1	80.4

Table 10. Semi-supervised segmentation accuracy (mIoU, %) on PASCAL VOC.

Comparison. The segmentation results are summarized in Table 10. Using 1/16, 1/8, 1/4 labeled data, DTS produces consistent accuracy gains of 3.5%, 1.7%, 0.8% over the Mean Teacher baseline [46], and 1.7%, 1.1%, 1.1% over the U²PL baseline [54]. This validates that semi-supervised learning also benefits from bidirectional learning which effectively improves the ability to learn from the pseudo labels generated from other groups.

5. Conclusions

In this paper, we study the UDA segmentation problem from the perspective of endowing the model with two-fold abilities, *i.e.*, **learning** from the reliable semantics, and **adapting** to the target domain. We show that these two abilities can conflict with each other in a single teacher-student model, and hence we propose a dual teacher-student (DTS) framework, where the two models stick to the ‘learning ability’ and ‘adapting ability’, respectively. Specifically, the second model learns to ‘focus on your target’ to improve its accuracy in the target domain. We further design a bidirectional learning strategy to facilitate the two models to communicate via generating pseudo labels. We validate the effectiveness of DTS and set new records on two standard UDA segmentation benchmarks.

As the accuracy on the two benchmarks fast approaches the supervised baseline (*e.g.*, a 77.8% mIoU on Cityscapes), we look forward to building larger and higher-quality synthesized datasets to break through the baseline. This is of value to real-world applications, as synthesized data can be generated at low costs. We expect that the proposed methodology can be applied to other scenarios such as domain generalization, with the second group focusing on a generic, unbiased feature space.

Acknowledge This work was supported by NSFC under Contract U20A20183 and 62021001. It was also supported by the GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC, and the Supercomputing Center of the USTC.

References

- [1] Lin Chen, Zhixiang Wei, Xin Jin, Huaian Chen, Miao Zheng, Kai Chen, and Yi Jin. Deliberated domain bridging for domain adaptive semantic segmentation. *Proceedings of the Advances in Neural Information Processing Systems*, 35, 2022.
- [2] 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.
- [3] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [4] Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [5] 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*, 2021.
- [6] 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/CVF Conference on Computer Vision and Pattern Recognition*, 2016.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 2009.
- [8] Ye Du, Yujun Shen, Haochen Wang, Jingjing Fei, Wei Li, Liwei Wu, Rui Zhao, Zehua Fu, and Qingjie Liu. Learning from future: A novel self-training framework for semantic segmentation. *arXiv preprint arXiv:2209.06993*, 2022.
- [9] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111:98–136, 2015.
- [10] Li Gao, Jing Zhang, Lefei Zhang, and Dacheng Tao. Dsp: Dual soft-paste for unsupervised domain adaptive semantic segmentation. *arXiv preprint arXiv:2107.09600*, 2021.
- [11] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. *arXiv preprint arXiv:2001.01526*, 2020.
- [12] Rui Gong, Wen Li, Yuhua Chen, and Luc Van Gool. Dlow: Domain flow for adaptation and generalization. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 2019.
- [13] Dayan Guan, Jiaying Huang, Aoran Xiao, and Shijian Lu. Unbiased subclass regularization for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.
- [15] Judy Hoffman, Sergio Guadarrama, Eric S Tzeng, Ronghang Hu, Jeff Donahue, Ross Girshick, Trevor Darrell, and Kate Saenko. Lsda: Large scale detection through adaptation. *Proceedings of the Advances in Neural Information Processing Systems*, 27, 2014.
- [16] 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 *Proceedings of the International Conference on Machine Learning*, 2018.
- [17] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016.
- [18] Sungeun Hong, Woobin Im, Jongbin Ryu, and Hyun S Yang. Ssp-dan: Deep domain adaptation network for face recognition with single sample per person. In *Proceedings of the IEEE International Conference on Image Processing*, pages 825–829. IEEE, 2017.
- [19] 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*, 2022.
- [20] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. *arXiv preprint arXiv:2204.13132*, 2022.
- [21] Lukas Hoyer, Dengxin Dai, Haoran Wang, and Luc Van Gool. Mic: Masked image consistency for context-enhanced domain adaptation. *arXiv preprint arXiv:2212.01322*, 2022.
- [22] Xinyue Huo, Lingxi Xie, Jianzhong He, Zijie Yang, Wengang Zhou, Houqiang Li, and Qi Tian. Atso: Asynchronous teacher-student optimization for semi-supervised image segmentation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 2021.
- [23] Xinyue Huo, Lingxi Xie, Hengtong Hu, Wengang Zhou, Houqiang Li, and Qi Tian. Domain-agnostic prior for transfer semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [24] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [25] Meina Kan, Shiguang Shan, and Xilin Chen. Bi-shifting auto-encoder for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2015.
- [26] Zhanghan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *Proceedings of the European Conference on Computer Vision*, 2020.

- [27] Zhanghan Ke, Daoye Wang, Qiong Yan, Jimmy Ren, and Rynson WH Lau. Dual student: Breaking the limits of the teacher in semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [28] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [29] Kang Li, Shujun Wang, Lequan Yu, and Pheng-Ann Heng. Dual-teacher: Integrating intra-domain and inter-domain teachers for annotation-efficient cardiac segmentation. In *Proceedings of the Medical Image Computing and Computer Assisted Intervention, Proceedings, Part I 23*, pages 418–427. Springer, 2020.
- [30] 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*, 2022.
- [31] Xinzhe Li, Qianru Sun, Yaoyao Liu, Qin Zhou, Shibao Zheng, Tat-Seng Chua, and Bernt Schiele. Learning to self-train for semi-supervised few-shot classification. *Proceedings of the Advances in Neural Information Processing Systems*, 32, 2019.
- [32] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [33] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *Proceedings of the International Conference on Machine Learning*, 2017.
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [35] Subhabrata Mukherjee and Ahmed Awadallah. Uncertainty-aware self-training for few-shot text classification. *Proceedings of the Advances in Neural Information Processing Systems*, 33, 2020.
- [36] Viktor Olsson, Wilhelm Tranheden, 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*, 2021.
- [37] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan Yuille. Deep co-training for semi-supervised image recognition. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [38] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, 2016.
- [39] 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/CVF Conference on Computer Vision and Pattern Recognition*, 2016.
- [40] Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. Semi-supervised self-training of object detection models. 2005.
- [41] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [42] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Map-guided curriculum domain adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3139–3153, 2020.
- [43] Kihyuk Sohn, Sifei Liu, Guangyu Zhong, Xiang Yu, Ming-Hsuan Yang, and Manmohan Chandraker. Unsupervised domain adaptation for face recognition in unlabeled videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2017.
- [44] Ricardo Sousa and Joao Gama. Comparison between co-training and self-training for single-target regression in data streams using amrules. In *IOTSTREAMING@PKDD/ECML*, 2017.
- [45] Yuxing Tang, Josiah Wang, Boyang Gao, Emmanuel Delandrea, Robert Gaizauskas, and Liming Chen. Large scale semi-supervised object detection using visual and semantic knowledge transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.
- [46] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*, 2017.
- [47] Marco Toldo, Umberto Michieli, Gianluca Agresti, and Pietro Zanuttigh. Unsupervised domain adaptation for mobile semantic segmentation based on cycle consistency and feature alignment. *Image and Vision Computing*, 95:103889, 2020.
- [48] Wilhelm Tranheden, 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*, 2021.
- [49] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [50] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [51] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2015.
- [52] Girish Varma, Anbumani Subramanian, Anoop Namboodiri, Manmohan Chandraker, and CV Jawahar. Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2019.

- [53] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [54] Yuchao Wang, Haochen Wang, Yujun Shen, Jingjing Fei, Wei Li, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xinyi Le. Semi-supervised semantic segmentation using unreliable pseudo-labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [55] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S Davis. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [56] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [57] 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*, 2020.
- [58] Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. Central moment discrepancy (cmd) for domain-invariant representation learning. *arXiv preprint arXiv:1702.08811*, 2017.
- [59] 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*, 2021.
- [60] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2017.
- [61] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 2018.
- [62] Qianyu Zhou, Zhengyang Feng, Qiqi Gu, Jiangmiao Pang, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. Context-aware mixup for domain adaptive semantic segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*, 2022.
- [63] Yuliang Zou, Zizhao Zhang, Han Zhang, Chun-Liang Li, Xiao Bian, Jia-Bin Huang, and Tomas Pfister. Pseudoseg: Designing pseudo labels for semantic segmentation. *arXiv preprint arXiv:2010.09713*, 2020.