

# Rethinking Ensemble-Distillation for Semantic Segmentation Based Unsupervised Domain Adaption

Chen-Hao Chao, Bo-Wun Cheng, and Chun-Yi Lee

Elsa Lab, Department of Computer Science, National Tsing Hua University, Hsinchu, Taiwan

{lance\_chao, bobcheng15, cylee}@gapp.nthu.edu.tw

## Abstract

Recent researches on unsupervised domain adaptation (UDA) have demonstrated that end-to-end ensemble learning frameworks serve as a compelling option for UDA tasks. Nevertheless, these end-to-end ensemble learning methods often lack flexibility as any modification to the ensemble requires retraining of their frameworks. To address this problem, we propose a flexible ensemble-distillation framework for performing semantic segmentation based UDA, allowing any arbitrary composition of the members in the ensemble while still maintaining its superior performance. To achieve such flexibility, our framework is designed to be robust against the output inconsistency and the performance variation of the members within the ensemble. To examine the effectiveness and the robustness of our method, we perform an extensive set of experiments on both GTA5→Cityscapes and SYNTHIA→Cityscapes benchmarks to quantitatively inspect the improvements achievable by our method. We further provide detailed analyses to validate that our design choices are practical and beneficial. The experimental evidence validates that the proposed method indeed offer superior performance, robustness and flexibility in semantic segmentation based UDA tasks against contemporary baseline methods.

## 1 Introduction

In the past few years, semantic segmentation has been attracting the attention of computer vision researchers. Many supervised semantic segmentation methods have been proposed and achieved remarkable performance [1–13]. Typically, those supervised semantic segmentation methods require abundant labeled training data, which are usually expensive to annotate and are commonly unavailable in most real-world scenarios. To resolve this problem, semantic segmentation based unsupervised domain adaptation (UDA) methods [14–36] have been introduced to bridge different domains. These semantic segmentation based UDA models learn to generalize to a target domain by training with the annotated data from a source domain and the unlabeled

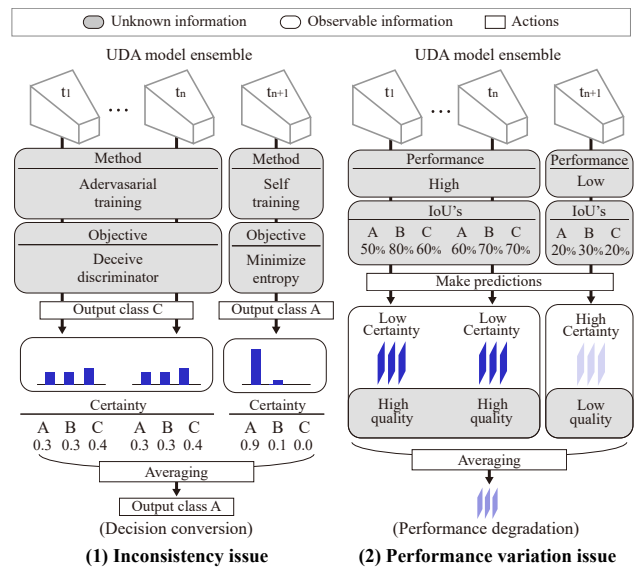


Figure 1: An illustrative example of (1) the inconsistency issue and (2) the performance variation issue mentioned in Section 1. For (1), the models  $t_1, \dots, t_n$  are trained with ADA methods, which produce relatively low-certainty outputs. In contrast, the model  $t_{n+1}$  is trained with self-training method and outputs high-certainty predictions. After averaging, the ensemble predicts class A as its final prediction instead of the majority consensus, i.e., class C. For (2), the models  $t_1, \dots, t_n$  are assumed to be high-performing members, while  $t_{n+1}$  is an under-performing one. After averaging, the high-certainty predictions from  $t_{n+1}$  could dominate the output of the ensemble, and cause the overall performance to degrade.

data from a target domain. Among these works, the authors in [14–25, 31, 37–39] resorted to adversarial domain adaptation (ADA) methods, through which the domain discrepancy is minimized by using their adversarial training schemes. Another branch of works has opted for self-training frameworks [26–30], which aim to improve the stability of their models during deployment by minimizing the entropy of the models’ predictions in a target domain. These ADA and self-training methods have demonstrated how a single model is able to learn to generalize to an annotation-less target domain. However, they only learn from a single distribution, leaving space for further improvements. Recently, in light of

the potential benefits of combining multiple UDA models, a number of works [40, 41] have attempted to borrow the concepts from ensemble learning. These works demonstrated how a group of UDA models can be trained simultaneously in an end-to-end fashion to learn different distributions of semantic information, and meanwhile transferring the knowledge to a compact student model. Despite their successes in bridging the domain gaps with multiple learners, these ensemble learning methods often lack flexibility as any modification to the teacher ensemble requires complete retraining of the whole framework.

To address such a problem, the concept of ensemble-distillation [42–50] can be leveraged since its focus is on designing an effective distillation process instead of a costly end-to-end ensemble learning framework. Typically, these ensemble distillation frameworks view the members in an ensemble as probabilistic models, and transfer the knowledge using expected certainty outputs. Nonetheless, the robustness of these methods is not guaranteed as they do not carefully take into account the followings: (1) the inconsistency in the scale of the output certainty values among the members in an ensemble (abbreviated as the ‘**inconsistency issue**’ hereafter), and (2) the performance variations across the members in an ensemble (abbreviated as the ‘**performance variation issue**’ hereafter). An example of these two issues is illustrated in Fig. 1. For the former, since each teacher model in the ensemble can be trained independently using different methods (e.g., ADA, self-training, data augmentation, or compound usage of them), the scale of the output certainty values may not be consistent across the ensemble. This may result in a situation that few members’ decisions with high certainty values dominate the entire ensemble’s output. As a result, the outputs from the members in an ensemble should be treated in an equal manner, as the inconsistency in their certainty values may come from their different training objectives instead of the real data distribution in the target domain. For the latter, since the performance (either per-class or average accuracy) of each teacher model in the ensemble may vary substantially, few underperforming members in the ensemble may cause the quality of the combined prediction to degrade significantly. This problem is especially severe under the context of UDA, since the ground truth labels in the target domain are unavailable, and the performance of the ensemble in the target domain is actually unknown. The above observations suggest that an effective mechanism is necessary to deal with these two issues and prevent them from influencing the quality of the combined predictions.

Being aware of these problems, we introduce a novel ensemble-distillation framework to avoid the aforementioned pitfalls. First, to tackle the certainty inconsistency issue, we introduce an output unification method in the framework to reduce the impact of the inconsistent scales of

the certainty outputs. Next, we embrace a new category of fusion function in our framework, named channel-wise fusion, to resolve the performance variation issue. Moreover, we design a method to determine the fusion policy of the proposed channel-wise fusion function to further enhance its effectiveness. To validate our designs, we evaluate the proposed framework with two commonly-adopted metrics, GTA5 [51]→Cityscapes [52] and SYNTHIA [53]→Cityscapes, to demonstrate the effectiveness and robustness of our framework against a number of baselines. The contributions are summarized as follows:

- We introduce a flexible UDA ensemble-distillation framework which is robust against the inconsistency in the scale of the output certainty values and the performance variations among the members in an ensemble.
- We propose a new category of fusion function, called channel-wise fusion, along with a fusion policy selection strategy as well as a conflict resolving mechanism to enhance its effectiveness.
- We evaluate our framework under various configurations, and demonstrate that it is able to outperform the baselines in terms of its robustness and effectiveness.

## 2 Related Works

**Unsupervised Domain Adaptation:** A number of methods have been proposed to bridge the discrepancy between different domains. One branch of these works adopted ADA frameworks to learn representations of their target domains [14–25, 31, 37–39]. These approaches typically employ a generator and a discriminator trained against each other to minimize the domain gap, and have shown significant improvements over those trained directly in the source domains. Another line of works has turned their attention to self-training and data augmentation measures to tackle UDA problems. For those works utilizing self-training, the concentration was mainly on preventing overfitting by using regularization [28, 29] or class-balancing [27] when minimizing the uncertainty in their target domains. The authors in [30] extended the concept of self-training and proposed a data augmentation technique. Their proposed method fine-tunes a model with mixed labels generated by combining ground truth annotations from a source domain and pseudo labels from a target domain. Recent researchers employed ensemble learning frameworks to resolve UDA problems [40, 41]. The authors in [41] proposed an end-to-end ensemble framework to solve UDA classification problems. The authors in [40] extended the idea of ensemble learning and proposed a joint learning ensemble framework to solve person re-identification UDA problems. These works showed how the ensemble learning frameworks can be integrated into UDA.

**Pseudo Labeling:** Pseudo labeling is a self-training method originally proposed to improve the performance of clas-

sification networks [54], and is usually accomplished by minimizing the entropy of a model’s predictions on unseen data. Pseudo labeling enables a better decision boundary to be achieved as the certainty of a model’s prediction increases [54]. This concept has been further extended to the field of semantic segmentation, and has gained success by incorporating the information of unlabeled data. Since self-training via pseudo labeling and UDA share many similar characteristics in terms of their problem formulations, it has recently been used to solve UDA problems [28–30].

**Ensemble-Distillation Method:** Ensemble-distillation is an extension of knowledge distillation. The authors in [43, 46] studied how the knowledge of a teacher model ensemble can be transferred to a student by training it with the soft predictions of the ensemble. They adopted averaging operation for combining the predictions and used KL-divergence as the loss function to transfer the knowledge. The authors in [50] aimed at resolving the diversity collapse issue in the ensemble-distillation problem. They argued that the averaging operation harms the diversity of the models in an ensemble and proposed to use a prior network [55] to estimate the distributions of their output uncertainties. These works have demonstrated their effectiveness under supervised training settings. However, the existing ensemble-distillation methods are not designed to handle unsupervised tasks, and are susceptible to the issues introduced in Section 1.

### 3 Preliminary

**Problem Definition:** For semantic segmentation based UDA problems, a model has access to the image-label pairs,  $x_{src}, y_{src}$ , from a source domain dataset  $\mathcal{D}_{src}$ , but only the images  $x_{tgt}$  from a target domain dataset  $\mathcal{D}_{tgt}$ . The training objective is to train the model such that its predictions can best estimate the ground truth labels  $y_{tgt}$  in the target domain. In other words, the mean intersection-over-union (mIoU) between the predictions of the model and  $y_{tgt}$  should be maximized. In the problem formulation concerned by this paper, a pretrained model ensemble  $\mathcal{T}$  is given, where each member in  $\mathcal{T}$  is separately trained using any arbitrary semantic segmentation based UDA method. The goal is to develop an ensemble-distillation strategy that can effectively integrate the knowledge from  $\mathcal{T}$  and distill it into a single student model, in a way that the mIoU of the student’s predictions for the instances in  $\mathcal{D}_{tgt}$  is maximized.

**Previous Ensemble-Distillation Method:** In this section, we explain how the concepts of the previous ensemble-distillation works [42–49] can be borrowed to perform semantic segmentation based UDA ensemble-distillation tasks. Typically, these works view  $\mathcal{T}$  as a set of probabilistic models, and complete the ensemble-distillation process through minimizing the negative log-likelihood loss  $\mathcal{L}_{KL}$  between the expected outputs from the ensemble and the student




Figure 2: A comparison between (a) the proposed ensemble-distillation framework and (b) the baseline framework.

model, as depicted in Fig. 2 (b). The distillation process of these methods under the settings of semantic segmentation based UDA can be formulated as:

$$\mathcal{L}_{KL} = - \sum_{p \in \mathcal{I}} \sum_{c \in \mathcal{C}} \tilde{s}^{(p,c)} \log(r^{(p,c)}), \quad (1)$$

where  $\mathcal{I}$  is a set of pixels in an image,  $\mathcal{C}$  is a given set of semantic classes.  $r^{(p,c)} \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{C}|}$  is the student’s certainty output, and  $\tilde{s}^{(p,c)} \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{C}|}$  represents the expected probabilistic prediction for class  $c \in \mathcal{C}$  at pixel  $p \in \mathcal{I}$ . A common way to capture  $\tilde{s}$  is through averaging, expressed as follows:

$$\tilde{s}^{(p,c)} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \hat{s}^{(p,c,t)}, \quad (2)$$

where  $\hat{s}^{(p,c,t)} \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{C}| \times |\mathcal{T}|}$  is the probabilistic prediction from  $t \in \mathcal{T}$  for class  $c \in \mathcal{C}$  at pixel  $p \in \mathcal{I}$  on target instances. As a result, the knowledge of the teacher ensemble can be transferred through minimizing  $\mathcal{L}_{KL}$  in a target domain.

**Pitfalls:** As discussed in Section 1, directly adopting previous methods to solve semantic segmentation based UDA problems is problematic because of the inconsistency issue and the performance variation issue. For the former, the scale of  $\hat{s}$  may vary across the models in  $\mathcal{T}$  under our problem formulation. This suggests that a direct operation, such as the averaging operation in Eq. (2), is inappropriate as a few  $\hat{s}$  with high certainty values may dominate the ensemble’s output decision. For the performance variation issue, the per-class or the average performance of each member in  $\mathcal{T}$  can vary substantially under our problem formulation. Since the fusion function formulated in Eq. (2) fuses the predictions of *all* teacher models, the under-performing members in  $\mathcal{T}$  can influence the quality of the fused results. Therefore, the adoption of such a fusion function is inappropriate as it may be sensitive to the performance variations within  $\mathcal{T}$ .

### 4 Methodology

To address the aforementioned problems, we introduce a new ensemble-distillation framework, and illustrate it in Fig. 2 (a). The main difference between the proposed method and the previous ones lies in two aspects: *Output Unification* and *Fusion Function*. In Sections 4.1 and 4.2, we




Figure 3: A illustration of the pixel-wise and channel-wise fusions.

walk through the designs of the output unification operation as well as the fusion function, and discuss how they may contribute to resolving the aforementioned issues. Finally, in Section 4.3, we formulate and summarize the proposed ensemble-distillation framework.

## 4.1 Output Unification

To resolve the inconsistency issue, we argue that the soft predictions  $\hat{s}$  in Eq. (2) should be unified first in the target domain, as illustrated in Fig. 2. This additional unification operation ensures that the raw output certainty values from the models in  $\mathcal{T}$  do not directly influence the subsequent fusion results. To accomplish this, we unify the soft predictions by converting them to pseudo labels, so as to make them all bear the same scale, i.e., representing the final decisions of the models in  $\mathcal{T}$ . The unification operation is formulated as:

$$\hat{y}^{(p,c,t)} = \begin{cases} 1, & \text{if } c = \arg \max_{c \in \mathcal{C}} \{\hat{s}^{(p,c,t)}\} \\ 0, & \text{otherwise} \end{cases}, \quad (3)$$

where  $\hat{y}^{(p,c,t)}$  is the unified output prediction from  $t \in \mathcal{T}$  for class  $c \in \mathcal{C}$  at pixel  $p \in \mathcal{I}$  in the target domain. This operation ensures that the subsequent fusion function can operate on items with a consistent scale, and thus eliminates the impact of the inconsistency in the original certainty outputs.

## 4.2 Fusion Function

We next move on to focus on investigating a fusion function that can take advantage of the unified predictions to achieve robustness against the performance variations issue. We compare two categories of fusion functions: *pixel-wise fusion* and *channel-wise fusion*. The former is a direct conversion from Eq. (2) and is used as our baseline method. The latter is the proposed method and is adopted to address the performance variation issue. Both pixel-wise fusion and channel-wise fusion are mapping functions  $f: \mathcal{I} \rightarrow \mathcal{C}_0$  that assign a class label  $c \in \mathcal{C}_0$  for the pixel  $p \in \mathcal{I}$  in the fusion output based on  $\hat{y}^{(p,c,t)}$ , where  $\mathcal{C}_0 := \mathcal{C} \cup \{c_0\}$  is a set that includes all  $c \in \mathcal{C}$  as well as the unlabeled symbol  $c_0$ .

### 4.2.1 Pixel-Wise Fusion

Pixel-wise fusion ( $f^{Pixel}$ ) adopts a statistical view on  $\mathcal{T}$ , and is designed to capture the average behavior of the ensemble. As depicted in Fig. 3, pixel-wise fusion treats each pixel in a semantic segmentation map as the basic unit of the fusion operation. Specifically, the fused result of each




Figure 4: An illustrative example of the three scenarios in Eq. (5).




Figure 5: An illustrative example of  $\pi$  used in  $f^{Channel}$ .

pixel is determined by taking majority voting among the predictions from  $\mathcal{T}$ , and is implemented as the following:

$$f^{Pixel}(p) = \arg \max_{c \in \mathcal{C}} \sum_{t \in \mathcal{T}} \hat{y}^{(p,c,t)}, \quad (4)$$

where  $\hat{y}$  is the unified output generated according to Eq. (3).

### 4.2.2 Channel-Wise Fusion

Based on a different perspective, channel-wise fusion ( $f^{Channel}$ ) treats each class channel as the basis for fusion, as depicted in Fig. 3. Instead of fusing the outputs of all  $t \in \mathcal{T}$ , channel-wise fusion relies on a fusion policy  $\pi: \mathcal{C} \rightarrow \mathcal{T}$ , which is a mapping function for recombining the unified outputs from different teacher models. More specifically, for each class  $c \in \mathcal{C}$ , the fusion policy  $\pi$  selects that class channel from the unified output  $\hat{y}^{(p,c,t)}$  of a teacher model  $t \in \mathcal{T}$ , as illustrated in the example shown in Fig. 5. Given such a  $\pi$ , the channel-wise fusion function is formulated as:

$$f^{Channel}(p) = \begin{cases} \text{(i) } \varepsilon, & \text{if } p \in A_o^\pi \\ \text{(ii) } c, & \text{if } p \in A_c^\pi \setminus A_o^\pi, \\ \text{(iii) } c_0, & \text{otherwise} \end{cases} \quad (5)$$

where (i) is the condition that  $p$  is labeled by multiple teachers, (ii) is the condition that  $p$  labeled as a certain class by a single teacher, and (iii) is the condition that  $p$  is unlabeled, as illustrated in Fig. 4. In Eq. (5),  $\varepsilon$  denotes a class label to be assigned in scenario (i),  $c$  is a class in  $\mathcal{C}$  in scenario (ii), and  $c_0$  denotes the unlabeled symbol in scenario (iii).  $A_c^\pi := \{p \mid p \in \mathcal{I}, \pi(c) = t, \hat{y}^{(p,c,t)} = 1\}$  is a set of pixels comprising  $\hat{y}^{(p,c,t)}$  for a given class  $c$  generated by a teacher  $t$  selected according to the given fusion policy  $\pi$ . Since different  $A_c^\pi$  may be produced by different  $t \in \mathcal{T}$  for different  $c \in \mathcal{C}$ , the pixels they cover are not necessarily mutually exclusive. Therefore,  $A_o^\pi$  is defined to represent the set of pixels labeled by multiple teachers, expressed as follows:

$$A_o^\pi = \bigcup_{\substack{c_1 \neq c_2, \\ c_1, c_2 \in \mathcal{C}}} (A_{c_1}^\pi \cap A_{c_2}^\pi). \quad (6)$$




Figure 6: An illustration of how the quality of a teacher model’s pseudo labels can affect the output certainty values of the student model. If the student model is trained with high-quality pseudo labels (i.e., pseudo labels with high IoU’s w.r.t.  $y_{tgt}$ ), it can learn a mapping from input features to the segmentation mask effectively, and generates high-certainty predictions. In contrast, if the student model is trained with low-quality pseudo labels (i.e., pseudo labels with low IoU’s w.r.t.  $y_{tgt}$ ) that are mismatched with the input features, the student’s output certainty values are likely to degrade.

The mechanism that assigns the value of  $\varepsilon$  for all  $p \in A_o^\pi$  is referred as the conflict-resolving mechanism. In this work, we employ a spatially-aware conflict-resolving mechanism that assigns a class label for each pixel in  $A_o^\pi$  using majority voting on a kernel. The size of the kernel is denoted as  $\kappa$ , and the set of pixels covered by the kernel centered at  $p$  is referred to as  $B_p^\kappa$ . The mechanism is formulated as follows:

$$\varepsilon = \arg \max_{c \in \mathcal{C}_p^\pi} |B_p^\kappa \cap A_c^\pi|. \quad (7)$$

where  $\mathcal{C}_p^\pi := \{c \mid c \in \mathcal{C}; p \in A_c^\pi\}$  represents a set of class(es) assigned to a pixel  $p$  under  $A_c^\pi$ .

### 4.2.3 Theoretical Properties of Channel-Wise Fusion

Let  $\Phi^{(c,t)}$  and  $\tilde{\Phi}^{(c,\pi(c))}$  be the per-class IoU’s w.r.t.  $y_{tgt}$  for the pseudo labels generated by  $t \in \mathcal{T}$  and the fused pseudo labels generated by  $f^{Channel}$ , respectively. Channel-wise fusion conforms to the following properties:

**Proposition 1.** Consider an arbitrary fusion policy  $\pi$ . Given a constant  $\alpha \in (0, 1)$  and classes  $c_1, \dots, c_n \in \mathcal{C}$ . If  $\Phi^{(c_i,t)} \geq \alpha, \forall i \in \{1, \dots, n\}, \forall t \in \mathcal{T}$  and  $|A_o^\pi| = 0$ , we have:

$$mIoU = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \tilde{\Phi}^{(c,\pi(c))} \geq \frac{n\alpha}{|\mathcal{C}|}. \quad (8)$$

**Proposition 2.** Consider an optimal fusion policy  $\pi^*(c) = \arg \max_{t \in \mathcal{T}} \{\Phi^{(c,t)}\}$ . Assume  $|A_o^{\pi^*}| = 0$ , we have:

$$mIoU = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \tilde{\Phi}^{(c,\pi^*(c))} \geq \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \Phi^{(c,t)}, \forall t \in \mathcal{T}. \quad (9)$$

For the detailed elaborations and proofs with regard to these properties, please refer to the supplementary materials. **Proposition 1** states the condition when the mIoU lower bound can be ensured. On the other hand, **Proposition 2** describes how the effectiveness of channel-wise fusion can be maximized. In the next section, we discuss how an effective  $\pi$  can be determined.




Figure 7: The cosine similarity between the per-class IoU of each  $t \in \mathcal{T}$  and the per-class certainty values of the student evaluated on the training set of Cityscapes, i.e., the cosine similarity between  $\Phi^{(c,t)}$  in Proposition 2 and  $\rho^{(c,t)}$  in Eq. (10). The experimental results reveal that the two variables are positively correlated to each other, as the cosine similarity values for all  $c \in \mathcal{C}$  are greater than 0 and close to 1. For the detailed settings, please refer to Section 5.

### 4.2.4 Certainty-Aware Policy Selection Strategy

Since the fusion policy  $\pi$  determines which teacher model is allowed to involve in  $f^{Channel}$ , choosing an appropriate  $\pi$  is therefore crucial to the robustness of our framework. In order to achieve this objective, a suitable measure is necessary for evaluating the quality of each teacher’s predictions in the target domain without using any target domain ground truth. The experimental clue illustrated and explained in Fig. 6 and 7 offers an empirical manner for the above purpose. From Fig. 6, it is observed that the quality of the unified outputs  $\hat{y}^{(p,c,t)}$ , i.e., the pseudo labels, from  $t \in \mathcal{T}$  are positively correlated to the output certainty values of distilled student model. This correlation suggests that any low-quality  $\hat{y}^{(p,c,t)}$ , which might be generated by some under-performing teacher models, can confuse the student model and causes its output certainty values to degrade. This correlation between the quality of the pseudo labels from  $t \in \mathcal{T}$  and the student’s output certainty values thus shed light on the development of the measure for approximating a teacher model’s performance without using target domain ground truth. In practice, an offline fusion policy selection strategy is adopted. Our framework first performs knowledge distillation on all  $t \in \mathcal{T}$  and transfers their knowledge to  $|\mathcal{T}|$  identical student models using the unified outputs. Then, their output certainty values are measured to obtain the approximated performance of their corresponding teacher models. Finally, for each  $c \in \mathcal{C}$ , a  $t \in \mathcal{T}$  that maximizes the student model’s output certainty values is selected. The fusion policy  $\pi$  is written as follows:

$$\pi(c) = \arg \max_{t \in \mathcal{T}} \{\rho^{(c,t)}\}, \quad (10)$$

where  $\rho^{(c,t)} \in \mathbb{R}^{|\mathcal{C}| \times |\mathcal{T}|}$  refers to the average certainty outputs of class  $c \in \mathcal{C}$  from the student model trained with the unified outputs generated by teacher  $t \in \mathcal{T}$ .

## 4.3 The Proposed Framework

Based on the formulations of output unification and fusion function, the loss function  $\mathcal{L}_{CE}$  for performing the ensemble-distillation in our framework is defined as follows:

$$\mathcal{L}_{CE} = - \sum_{p \in \mathcal{I}} \sum_{c \in \mathcal{C}} \tilde{y}^{(p,c)} \log(r^{(p,c)}), \quad (11)$$

where  $\tilde{y}^{(p,c)} \in \{0, 1\}^{|\mathcal{I}| \times |C|}$  is the fused results, defined as:

$$\tilde{y}^{(p,c)} = \begin{cases} 1, & \text{if } c = f^{Channel}(p) \\ 0, & \text{otherwise} \end{cases}. \quad (12)$$

The pseudo code of the ensemble-distillation method in our framework is summarized in the supplementary materials.

## 5 Experimental Setup

**Baselines and Evaluation Methods:** In this work, we evaluate and compare the experimental results in terms of the effectiveness and the robustness. To examine the effectiveness of the proposed framework, we compare our method against two ensemble-distillation schemes and a number of UDA baselines. The ensemble-distillation schemes include EnD [43] and its recent revision EnD<sup>2</sup> [50]. The semantic segmentation based UDA baselines cover APODA [20], PatchAlign [21], AdvEnt [22], FDA-MBT [36], PIT [35], CBST [27], MRKLD [28], R-MRNet [29], and DACS [30]. To examine the robustness of our framework, we select the members for  $\mathcal{T}$  based on two criteria: (1) each member in  $\mathcal{T}$  is trained with different UDA methods; and (2) there exists large per-class and average performance variations among the members in  $\mathcal{T}$ . According to these criteria, we select DACS [30] (data augmentation), R-MRNet [29] (adversarial training), MRKLD [28] (self-training), and CBST [27] (self-training) to form  $\mathcal{T}$  in our experiments. We evaluate the proposed framework and the baselines on two commonly adopted benchmarks: GTA5 [51]→Cityscapes [52] and SYNTHIA [53]→Cityscapes. For the former, the models have access to 24,966 image-label pairs from the training set of GTA5, and 2,975 images from the training set of Cityscapes. We evaluate the student model’s per-class IoU’s of the 19 semantic classes as well as the its mIoU’s on the validation set of Cityscapes. For the latter, the models have access to 9,400 image-label pairs from the training set of SYNTHIA, and 2,975 images from the training set of Cityscapes. In a similar fashion, we evaluate the student model’s per-class IoU’s of 13 and 16 semantic classes as well as its mIoU’s on the validation set of Cityscapes.

**Implementation Details:** For the student model, we adopt Deeplabv3+ [12] architecture with DRN-D-54 [1] as our backbone, which is trained using SGD with a learning rate initialized to  $2.5 \times 10^{-4}$  and decreased with a factor of 0.9. The weight decay is set to  $5 \times 10^{-3}$ , the momentum is set to 0.9, and the batch size is set to 10 for 100K iterations with early stopping. The value of  $\kappa$  in  $f^{Channel}$  is set to 13. During the process of certainty-aware policy selection strategy, 500 images from the training set of Cityscapes are used for measuring  $\rho^{(c,t)}$ , while the other 2475 images are used for training the student. The student model is pre-trained in the source domain, and fine-tuned with  $(x_{src}, y_{src})$  and  $(x_{tgt}, \tilde{y})$  during the distillation process.

## 6 Experimental Results

In this section, we present a number of experiments to validate the design of our framework. First, we compare our framework against a number of baselines and demonstrate its superior performance. Next, according to the experimental results, we show that the output unification operation can provide robustness against the inconsistency issue. Then, we present another experiment, in which under-performing members are added to  $\mathcal{T}$  to create performance variations, to demonstrate the robustness of our framework with  $f^{Channel}$  against the performance variation issue. In addition, we perform experiments to validate the effectiveness of the fusion policy selection strategy and analyze how the value of  $\kappa$  in the conflict resolving mechanism can impact the performance. Finally, we explore the flexibility of our framework by adding additional teacher models in an iterative manner, and show that our framework is able to evolve with time.

### 6.1 Quantitative Results on the Benchmarks

Table 1 first demonstrates the quantitative results of the proposed framework against a number of baselines mentioned in Section 5 on the two benchmarks: GTA5→Cityscapes and SYNTHIA→Cityscapes. It is observed that the student models trained under our proposed framework with  $f^{Channel}$  (i.e., ‘Ours (Channel)’) is able to outperform the previous ensemble-distillation baselines, i.e., EnD [43] and EnD<sup>2</sup> [50], by a margin of 6.64% mIoU and 6.02% mIoU on GTA5→Cityscapes, and 6.41% mIoU and 4.57% mIoU on SYNTHIA→Cityscapes, respectively. In addition, it is also observed that the student model trained with the proposed framework with  $f^{Channel}$  (i.e., ‘Ours (Channel)’) is able to outperform that with the baseline  $f^{Pixel}$  (i.e., ‘Ours (Pixel)’) by a margin of 4.72% mIoU on GTA5→Cityscapes, and 4.52% mIoU on SYNTHIA→Cityscapes.

### 6.2 Robustness of the Proposed Framework

In this section, we validate the robustness of the proposed framework against the two issues discussed in Section 1. First, to verify that the proposed output unification operation can provide robustness against the inconsistency issue, we leverage the insights from an experiment conducted on  $\mathcal{T}$ , with the members of  $\mathcal{T}$  bearing substantial output certainty scale variations, as shown in Fig. 8. The results from Table 1 reveal that the performance of the student models trained with ‘Ours (Pixel)’, which is basically the ensemble-distillation baseline EnD [43] equipped with the proposed output unification method, is able to outperform EnD by a noticeable margin on both benchmarks. This implies that the adoption of the output unification method is able to provide robustness against the inconsistency issue. Nevertheless, as discussed in Section 4.2, ‘Ours (Pixel)’ may still be vulnerable to the performance variations of the members in  $\mathcal{T}$ .

GTA5 → Cityscapes																				
Method	Road	SideW	Build	Wall	Fence	Pole	Light	Sign	Veg	Terrain	Sky	Person	Rider	Car	Truck	Bus	Train	Motor	Bike	mIoU
APODA [20]	85.6	32.8	79.0	29.5	25.5	26.8	34.6	19.9	83.7	40.6	77.9	59.2	28.3	84.6	34.6	49.2	8.0	32.6	39.6	45.9
PatchAlign [21]	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
AdvEnt [22]	89.4	33.1	81.0	26.6	26.8	27.2	33.5	24.7	83.9	36.7	78.8	58.7	30.5	84.8	38.5	44.5	1.7	31.6	32.4	45.5
FDA-MBT [36]	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
PIT [35]	87.5	43.4	78.8	31.2	30.2	36.3	39.9	42.0	79.2	37.1	79.3	65.4	37.5	83.2	46.0	45.6	25.7	23.5	49.9	50.6
CBST [27]	91.8	53.5	80.5	32.7	21.0	34.0	28.9	20.4	83.9	34.2	80.9	53.1	24.0	82.7	30.3	35.9	16.0	25.9	42.8	45.9
MRKLD [28]	91.0	55.4	80.0	33.7	21.4	37.3	32.9	24.5	85.0	34.1	80.8	57.7	24.6	84.1	27.8	30.1	<b>26.9</b>	26.0	42.3	47.1
R-MRNet [29]	90.4	31.2	85.1	36.9	25.6	37.5	48.8	48.5	85.3	34.8	81.1	64.4	36.8	86.3	34.9	52.2	1.7	29.0	46.6	50.3
DACS [30]	89.90	39.66	87.87	30.71	39.52	38.52	46.43	52.79	87.98	43.96	<b>88.76</b>	67.20	35.78	84.45	45.73	50.19	0.00	27.25	33.96	52.14
Source Only	57.40	21.43	56.80	8.93	22.14	32.38	34.62	24.90	78.98	15.92	63.71	55.55	13.83	58.11	21.99	29.78	2.36	28.41	33.98	34.80
EnD [43]	92.17	53.12	84.85	24.77	29.76	40.38	40.98	49.35	86.21	42.85	79.74	62.79	35.98	85.72	42.10	44.45	0.26	28.27	51.80	51.34
EnD <sup>2</sup> [50]	92.39	53.84	85.34	24.51	30.53	40.28	42.40	50.28	86.19	43.39	80.55	63.26	36.75	86.15	43.95	43.91	0.20	30.17	53.22	51.96
Ours (Pixel)	92.29	57.34	84.09	36.75	29.17	41.37	48.96	42.26	86.91	39.95	82.81	66.29	37.42	86.94	35.21	48.82	1.48	40.78	53.02	53.26
Ours (Channel)	<b>94.43</b>	<b>60.90</b>	<b>88.07</b>	<b>39.46</b>	<b>41.80</b>	<b>43.24</b>	<b>49.08</b>	<b>56.00</b>	<b>88.01</b>	<b>45.83</b>	87.79	<b>67.58</b>	<b>38.05</b>	<b>90.08</b>	<b>57.64</b>	<b>51.90</b>	0.00	<b>46.57</b>	<b>53.28</b>	<b>57.98</b>

SYNTIA → Cityscapes																					
Method	Road	SideW	Build	Wall <sup>*</sup>	Fence <sup>*</sup>	Pole <sup>*</sup>	Light	Sign	Veg	Terrain	Sky	Person	Rider	Car	Truck	Bus	Train	Motor	Bike	mIoU	mIoU <sup>*</sup>
APODA [20]	86.4	41.3	79.3	-	-	22.6	17.3	80.3	-	81.6	56.9	21.0	84.1	-	49.1	-	24.6	45.7	-	53.1	
PatchAlign [21]	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
AdvEnt [22]	85.6	42.2	79.7	8.7	0.4	25.9	5.4	8.1	80.4	-	84.1	57.9	23.8	73.3	-	36.4	-	14.2	33.0	41.2	48.0
FDA-MBT [36]	79.3	35.0	73.2	-	-	-	19.9	24.0	61.7	-	82.6	61.4	31.1	83.9	-	40.8	-	<b>38.4</b>	51.1	-	52.5
PIT [35]	83.1	27.6	81.5	8.9	0.3	21.8	26.4	33.8	76.4	-	78.8	64.2	27.6	79.6	-	31.2	-	31.0	31.3	44.0	51.8
CBST [27]	68.0	29.9	76.3	10.8	1.4	33.9	22.8	29.5	77.6	-	78.3	60.6	28.3	81.6	-	23.5	-	18.8	39.8	42.6	48.9
MRKLD [28]	67.7	32.2	73.9	10.7	1.6	<b>37.4</b>	22.2	31.2	80.8	-	80.5	60.8	29.1	82.8	-	25.0	-	19.4	45.3	43.8	50.1
R-MRNet [29]	87.6	41.9	83.1	14.7	1.7	36.2	31.3	19.9	81.6	-	80.6	63.0	21.8	86.2	-	40.7	-	23.6	53.1	47.9	54.9
DACS [30]	80.56	25.12	81.90	21.46	2.85	37.20	22.67	23.99	<b>83.69</b>	-	<b>90.77</b>	<b>67.61</b>	<b>38.33</b>	82.92	-	38.90	-	28.49	47.58	48.34	54.81
Source Only	25.40	15.55	59.70	18.07	0.66	26.35	19.36	30.22	72.50	-	74.28	48.11	13.67	74.62	-	36.94	-	13.92	36.45	35.36	40.06
EnD [43]	92.17	53.12	84.85	24.77	29.76	40.38	40.98	49.35	86.21	42.85	79.74	62.79	35.98	85.72	42.10	44.45	0.26	28.27	51.80	51.34	
EnD <sup>2</sup> [50]	92.39	53.84	85.34	24.51	30.53	40.28	42.40	50.28	86.19	43.39	80.55	63.26	36.75	86.15	43.95	43.91	0.20	30.17	53.22	51.96	
Ours (Pixel)	86.98	44.18	80.95	19.38	1.52	30.47	25.64	30.39	79.92	-	78.84	56.55	27.14	84.47	-	44.52	-	26.03	55.08	48.25	55.43
Ours (Channel)	<b>88.65</b>	<b>46.69</b>	<b>83.79</b>	<b>22.66</b>	<b>4.14</b>	35.01	<b>35.93</b>	<b>36.19</b>	82.80	-	81.35	61.61	32.13	<b>87.93</b>	-	<b>52.79</b>	-	31.95	<b>57.65</b>	<b>52.58</b>	<b>59.95</b>

Table 1: The quantitative results evaluated on the GTA5→Cityscapes and SYNTIA→Cityscapes benchmarks. The numbers presented in the middle and the last two columns correspond to per-class IoUs, mIoU, and mIoU\*, respectively. mIoU\* represents the mean IoU over all the semantic classes excluding those with superscript \*, and is adopted by a few baseline methods in their original papers. The models used in our semantic segmentation based UDA model ensemble  $\mathcal{T}$  are highlighted in blue. The setting ‘Source Only’ indicates that the student model is trained only with the source domain ground truth annotations. The evaluation results of EnD [43] and EnD<sup>2</sup> [50] are obtained from our self-implemented models, while those of the remaining baselines are directly obtained from their original papers.




Figure 8: The distribution of the pixel output certainty values from the models in  $\mathcal{T}$ , which are trained on GTA5→Cityscapes, with their output certainty values normalized to the range [0, 1] using softmax operation and evaluated on the training set of Cityscapes.




Figure 9: The performance comparison of our framework with  $f^{Channel}$  (‘Ours (Channel)’), our framework with  $f^{Pixel}$  (‘Ours (Pixel)’), and the baseline EnD and EnD<sup>2</sup> methods, with under-performing members added to  $\mathcal{T}$ . In this experiment, the members in  $\mathcal{T}$  are evaluated on the GTA5→Cityscapes benchmark.

To inspect the robustness of the proposed framework with channel-wise fusion (‘Ours (Channel)’), against the performance variation issue, we next look into an experi-

ment that introduces performance variations by adding under-performing members into  $\mathcal{T}$ . Specifically, we use the models trained using only source domain instances, i.e., ‘Source Only’ in Table 1, as the under-performing members. The analysis is presented in Fig. 9. It is observed that the performance of the students trained with ‘Ours (Pixel)’, EnD, and EnD<sup>2</sup> all degrade when the number of under-performing members increases. In contrast, the students trained using ‘Ours (Channel)’ is able to maintain their performance despite the inclusion of the under-performing members in  $\mathcal{T}$ . These results demonstrate the significance of preventing clueless incorporation of information from all  $t \in \mathcal{T}$ , and highlight the effectiveness of  $f^{Channel}$  in providing robustness against the unfavorable performance variation issue.

### 6.3 Analysis on Channel-Wise Fusion

**Fusion Policy:** To examine the effectiveness of  $\pi$  selected according to the strategy described in Section 4.2.4, we design two additional fusion policies,  $\pi^{rnd}$  and  $\pi^{tgt}$ , for comparison purposes.  $\pi^{rnd}$  designates a  $t \in \mathcal{T}$  for each  $c \in \mathcal{C}$  randomly, while  $\pi^{tgt}$  carries out this designation greedily according to the oracle performance of the models in  $\mathcal{T}$  in the target domain (i.e.,  $\pi^*$  described in Proposition 2). As shown in Table 3, the mIoU of the fused pseudo labels generated based on the  $\pi$  selected according to the proposed strategy is 7.09% higher than that with  $\pi^{rnd}$ . It is also observed that the mIoU of the fused pseudo labels generated

$\kappa$	1	5	7	13	21	27
mIoU Gains	+0.00	+0.15	+0.37	+0.73	+0.31	+0.09

Table 2: The mIoU gains of  $f^{Channel}$  with  $\pi^{rnd}$  for different  $\kappa$ .




Figure 10: The visualized results of the pseudo labels generated by the proposed channel-wise fusion with different choices of  $\kappa$ .

Policy	$\pi^{rnd}$ (Random)	$\pi$ (Ours)	$\pi^{tgt}$ (Oracle)
mIoU	49.22	56.31	56.48
Overlapped Area	$ A_0^{\pi^{rnd}} $	$ A_0^{\pi} $	$ A_0^{\pi^{tgt}} $
Ratio	4.6%	3.7%	2.6%
mIoU Gains	+0.73	+1.13	+1.21

Table 3: The mIoU’s of the fused pseudo labels generated by channel-wise fusion  $f^{Channel}$  under different fusion policies w.r.t. the ground truth annotations of the training set of Cityscapes.  $\pi^{rnd}$ ,  $\pi$ , and  $\pi^{tgt}$  refer to the fusion policies mentioned in Section 6.3. ‘Ratio’ refers to the average proportion of the overlapped area in an image (i.e.,  $\frac{|A_0^{\pi}|}{|Z|}$ ). The ‘mIoU gains’ represents the mIoU gains from the adoption of the proposed conflict-resolving mechanism.

based on this  $\pi$  is very close to the mIoU obtained from  $\pi^{tgt}$ . These pieces of evidence validate that the proposed policy selection strategy is able to generate a fusion policy  $\pi$  that is very close to the optimal policy  $\pi^*$  without leveraging any sort of unavailable target domain ground truth annotations.

**Conflict-Resolving Mechanism:** To investigate the effectiveness of the conflict-resolving mechanism described in Section 4.2 and Eq. (7), we perform parameter analysis on  $f^{Channel}$  with different choices of  $\kappa$ . In this analysis, the fusion policy adopted is  $\pi^{rnd}$  to ensure that results are independent of the design of  $\pi$ . Fig. 10 shows the visualized results of the fused pseudo labels with different  $\kappa$ . Table 2 reports the mIoU gained for various  $\kappa$  with respect to the condition  $\kappa = 1$ . It is observed that pixel predictions in the overlapped area can be better determined if a moderate collection of predictions from the neighboring pixels are taken into account. However, if  $\kappa$  becomes too large, an excessive amount of unrelated semantic information is included in the conflict resolving mechanism, and degrades the quality of the fused pseudo labels. Table 3 shows that different policies, e.g.,  $\pi$  and  $\pi^{tgt}$ , can also benefit from the conflict-resolving mechanism, thus justifying it from a different perspective.

## 6.4 Flexibility of the Proposed Framework

In comparison to end-to-end UDA ensemble learning methods, our framework is more flexible as the proposed framework can operate on any arbitrary compositions of  $\mathcal{T}$ . Such flexibility enables our framework to evolve with time as (1) a model trained with any newly developed UDA method can be integrated into our framework, and (2) the student models can be added back to  $\mathcal{T}$  and further improve




Figure 11: The performance of the student models trained with the proposed framework using  $f^{Channel}$  with different compositions of  $\mathcal{T}$ . The experiment is performed on GTA5→Cityscapes. For settings ‘S4’, ‘S5’ and ‘S6’, the student models trained under the previous settings, i.e., ‘S3’, ‘S4’, and ‘S5’, are added back to  $\mathcal{T}$ .

the overall performance. To demonstrate these two merits of our framework, we evaluate our framework with different composition of  $\mathcal{T}$ , as shown in Fig. 11. It is observed that the student models trained under the settings ‘S1’, ‘S2’, and ‘S3’, which simulate the addition of the members trained with newly developed UDA methods, can outperform their corresponding teacher models in  $\mathcal{T}$  (whose mIoU’s are reported in Table 1). This implies that our framework offers the potential to evolve with time. In addition, ‘S4’, ‘S5’, and ‘S6’ in Fig. 11 showcase that the performance of our framework can be further improved if the students are added back to  $\mathcal{T}$ , indicating that our framework can be applied in an iterative manner to produce better results. These results highlight the flexibility of our framework as any UDA methods can be incorporated and potentially enhance its performance.

## 7 Conclusion

In this paper, we presented a flexible ensemble-distillation framework to address the common pitfalls, i.e., the lack of robustness, of previous methods. We incorporated an output unification operation into the proposed framework to ensure that the fused outputs of the ensemble are free of the influence from the certainty inconsistency among the models in the ensemble. In addition, to tackle the performance variation issue, we proposed a channel-wise fusion function that is robust against this issue. As our framework is able to integrate different types of UDA methods while maintaining its robustness, it therefore pioneers a new direction for future semantic segmentation based UDA researches.

**Acknowledgement:** The authors acknowledge the support from the Ministry of Science and Technology (MOST) in Taiwan under grant nos. MOST 110-2636-E-007-010 and MOST 110-2634-F-007-019, as well as the support from MediaTek Inc., Taiwan. The authors thank the donation of the GPUs from NVIDIA Corporation and NVIDIA AI Technology Center, and the National Center for High-performance Computing of National Applied Research Laboratories in Taiwan for providing computational and storage resources.



## References

- [1] F. Yu, V. Koltun, and T. Funkhouser. Dilated residual networks. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 472–480, Jul. 2017. 1, 6
- [2] G. Lin, A. Milan, C. Shen, and I. Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 1925–1934, Jul. 2017.
- [3] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *Proc. Int. Conf. Learning Representations (ICLR)*, May 2016.
- [4] V. Badrinarayanan, A. Kendall, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, 39(12):2481–2495, Jan. 2017.
- [5] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, Jun. 2015.
- [6] Y. Yuan, X. Chen, and J. Wang. Object-contextual representations for semantic segmentation. In *Proc. European Conf. Computer Vision (ECCV)*, pages 173–190, Oct. 2020.
- [7] Z. Wu, C. Shen, and A. Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019.
- [8] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, Jun. 2018.
- [9] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected CRFs. *arXiv:1412.7062*, Jul. 2014.
- [10] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, 40(4):834–848, Apr. 2017.
- [11] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam. Re-thinking atrous convolution for semantic image segmentation. *arXiv:1706.05587*, Dec. 2017.
- [12] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proc. European Conf. Computer Vision (ECCV)*, pages 801–818, Sep. 2018. 6
- [13] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 2881–2890, Jul. 2017. 1
- [14] J. Hoffman, D. Wang, F. Yu, and T. Darrell. FCNs in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv:1612.02649*, Dec. 2016. 1, 2
- [15] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *Proc. Int. Conf. Machine Learning (ICML)*, pages 1989–1998, Jul. 2018.
- [16] Y. Luo, P. Liu, T. Guan, J. Yu, and Y. Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pages 6778–6787, Oct. 2019.
- [17] R. Gong, W. Li, Y. Chen, and L. V. Gool. Dlow: Domain flow for adaptation and generalization. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 2477–2486, Jun. 2019.
- [18] Z. Wu, X. Han, Y.-L. Lin, M. Gokhan Uzunbas, T. Goldstein, S. Nam Lim, and L. S. Davis. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *Proc. European Conf. Computer Vision (ECCV)*, pages 518–534, Oct. 2018.
- [19] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7472–7481, 2018.
- [20] J. Yang, R. Xu, R. Li, X. Qi, X. Shen, G. Li, and L. Lin. An adversarial perturbation oriented domain adaptation approach for semantic segmentation. In *Proc. the Thirty-Fourth AAAI Conf. Artificial Intelligence (AAAI)*, pages 12613–12620, Feb. 2020. 6, 7
- [21] Y.-H. Tsai, K. Sohn, S. Schulter, and M. Chandraker. Domain adaptation for structured output via discriminative patch representations. In *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pages 1456–1465, Oct. 2019. 6, 7
- [22] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 2517–2526, Jun. 2019. 6, 7
- [23] Y. Zhang, Z. Qiu, T. Yao, D. Liu, and T. Mei. Fully convolutional adaptation networks for semantic segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 6810–6818, Jun. 2018.
- [24] Y.-H. Chen, W.-Y. Chen, Y.-T. Chen, B.-C. Tsai, Y.-C. Frank Wang, and M. Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pages 1992–2001, Oct. 2017.
- [25] Z. Zheng and Y. Yang. Unsupervised scene adaptation with memory regularization in vivo. *Proc. Int. Joint Conf. Artificial Intelligence (IJCAI)*, Jul. 2020. 1, 2

- [26] J. Choi, T. Kim, and C. Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pages 6830–6840, Oct. 2019. [1](#)
- [27] Y. Zou, Z. Yu, B. Vijaya Kumar, and J. Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proc. European Conf. Computer Vision (ECCV)*, pages 289–305, Sep. 2018. [2](#), [6](#), [7](#)
- [28] Y. Zou, Z. Yu, X. Liu, B. Kumar, and J. Wang. Confidence regularized self-training. In *Proc. of the IEEE Int. Conf. on Computer Vision (ICCV)*, pages 5982–5991, Oct. 2019. [2](#), [3](#), [6](#), [7](#)
- [29] Z. Zheng and Y. Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *Int. Journal of Computer Vision (IJCV)*, 2020. doi: 10.1007/s11263-020-01395-y. [2](#), [6](#), [7](#)
- [30] W. Tranheden, V. Olsson, J. Pinto, and L. Svensson. DACS: Domain adaptation via cross-domain mixed sampling. In *Proc. IEEE Winter Conf. on Applications of Computer Vision (WACV)*, pages 1379–1389, Jan. 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [31] Y. Luo, L. Zheng, T. Guan, J. Yu, and Y. Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2507–2516, 2019. [1](#), [2](#)
- [32] K.-H. Lee, G. Ros, J. Li, and A. Gaidon. Spigan: Privileged adversarial learning from simulation. *arXiv:1810.03756*, 2018.
- [33] Y. Chen, W. Li, X. Chen, and L. V. Gool. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1841–1850, 2019.
- [34] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, pages 7364–7373, 2019.
- [35] F. Lv, T. Liang, X. Chen, and G. Lin. Cross-domain semantic segmentation via domain-invariant interactive relation transfer. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4334–4343, 2020. [6](#), [7](#)
- [36] Y. Yang and S. Soatto. FDA: Fourier domain adaptation for semantic segmentation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4085–4095, 2020. [1](#), [6](#), [7](#)
- [37] Y.-C. Chen, Y.-Y. Lin, M.-H. Yang, and J.-B. Huang. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 1791–1800, Jun. 2019. [1](#), [2](#)
- [38] Y. Li, L. Yuan, and N. Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 6936–6945, Jun. 2019.
- [39] L. Du, J. Tan, H. Yang, J. Feng, X. Xue, Q. Zheng, X. Ye, and X. Zhang. Ssf-dan: Separated semantic feature based domain adaptation network for semantic segmentation. In *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pages 982–991, Oct. 2019. [1](#), [2](#)
- [40] L. T. Nguyen-Meidine, A. Belal, M. Kiran, J. Dolz, L.-A. Blais-Morin, and E. Granger. Unsupervised multi-target domain adaptation through knowledge distillation. In *Proc. IEEE Winter Conf. on Applications of Computer Vision (WACV)*, pages 1339–1347, 2021. [2](#)
- [41] G. Kang, L. Jiang, Y. Yang, and A. G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proc. the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4893–4902, 2019. [2](#)
- [42] C. Buciluă, R. Caruana, and A. Niculescu-Mizil. Model compression. In *Proc. 12th ACM SIGKDD Int. Conf. Knowledge Discovery and Data mining (KDD-06)*, pages 535–541, Aug. 2006. [2](#), [3](#)
- [43] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *arXiv:1503.02531*, Mar. 2015. [3](#), [6](#), [7](#)
- [44] J. H. Cho and B. Hariharan. On the efficacy of knowledge distillation. In *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pages 4794–4802, Oct. 2019.
- [45] T. Furlanello, Z. C. Lipton, M. Tschannen, L. Itti, and A. Anandkumar. Born again neural networks. *Proc. Int. Conf. Machine Learning (ICML)*, Jul. 2018.
- [46] A. K. Balan, V. Rathod, K. P. Murphy, and M. Welling. Bayesian dark knowledge. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 3438–3446, Dec. 2015. [3](#)
- [47] Le Thanh Nguyen-Meidine, Éric Granger, M. Kiran, J. Dolz, and Louis-Antoine Blais-Morin. Joint progressive knowledge distillation and unsupervised domain adaptation. *Int. Joint Conf. on Neural Networks (IJCNN)*, pages 1–8, 2020.
- [48] M. Orbes-Arteainst, J. Cardoso, L. Sørensen, C. Igel, S. Ourselin, M. Modat, M. Nielsen, and A. Pai. Knowledge distillation for semi-supervised domain adaptation. In *OR 2.0 Context-Aware Operating Theaters and Machine Learning in Clinical Neuroimaging*, pages 68–76, 2019.
- [49] Y. Liu, K. Chen, C. Liu, Z. Qin, Z. Luo, and J. Wang. Structured knowledge distillation for semantic segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 2604–2613, Jun. 2019. [3](#)
- [50] A. Malinin, B. Mlodozieniec, and M. Gales. Ensemble distribution distillation. In *Proc. Int. Conf. Learning Representations (ICLR)*, 2020. [2](#), [3](#), [6](#), [7](#)

- [51] S. R. Richter, V. Vineet, S. Roth, and V. Koltun. Playing for data: Ground truth from computer games. In *Proc. European Conf. Computer Vision (ECCV)*, pages 102–118, Oct. 2016. [2, 6](#)
- [52] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016. [2, 6](#)
- [53] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez. The SYNTHIA Dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 3234–3243, Jun. 2016. [2, 6](#)
- [54] D.-H. Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, Int. Conf. on Machine Learning (ICML)*, volume 3, Jun. 2013. [3](#)
- [55] A. Malinin and M. Gales. Predictive uncertainty estimation via prior networks. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pages 7047–7058, 2018. [3](#)