

# Domain Adaptive Semantic Segmentation Using Weak Labels

Sujoy Paul<sup>1</sup>, Yi-Hsuan Tsai<sup>2</sup>, Samuel Schulter<sup>2</sup>,  
Amit K. Roy-Chowdhury<sup>1</sup>, Manmohan Chandraker<sup>2,3</sup>

<sup>1</sup>UC Riverside      <sup>2</sup>NEC Labs America      <sup>3</sup>UC San Diego

**Abstract.** Learning semantic segmentation models requires a huge amount of pixel-wise labeling. However, labeled data may only be available abundantly in a domain different from the desired target domain, which only has minimal or no annotations. In this work, we propose a novel framework for domain adaptation in semantic segmentation with image-level weak labels in the target domain. The weak labels may be obtained based on a model prediction for unsupervised domain adaptation (UDA), or from a human annotator in a new weakly-supervised domain adaptation (WDA) paradigm for semantic segmentation. Using weak labels is both practical and useful, since (i) collecting image-level target annotations is comparably cheap in WDA and incurs no cost in UDA, and (ii) it opens the opportunity for category-wise domain alignment. Our framework uses weak labels to enable the interplay between feature alignment and pseudo-labeling, improving both in the process of domain adaptation. Specifically, we develop a weak-label classification module to enforce the network to attend to certain categories, and then use such training signals to guide the proposed category-wise alignment method. In experiments, we show considerable improvements with respect to the existing state-of-the-arts in UDA and present a new benchmark in the WDA setting. Project page is at <http://www.nec-labs.com/~mas/WeakSegDA>.

## 1 Introduction

Unsupervised domain adaptation (UDA) methods for semantic segmentation have been developed to tackle the issue of domain gap. Existing methods aim to adapt a model learned on the source domain with pixel-wise ground truth annotations, e.g., from a simulator which requires the least annotation efforts, to the target domain that does not have any form of annotations. These UDA methods in the literature for semantic segmentation are developed mainly using two mechanisms: pseudo label self-training and distribution alignment between the source and target domains. For the first mechanism, pixel-wise pseudo labels are generated via strategies such as confidence scores [29,23] or self-paced learning [60], but such pseudo-labels are specific to the target domain, and do not consider alignment between domains. For the second mechanism, numerous spaces could be considered to operate the alignment procedure, such as pixel [21,35], feature [22,57], output [49,9], and patch [50] spaces. However, alignment performed by

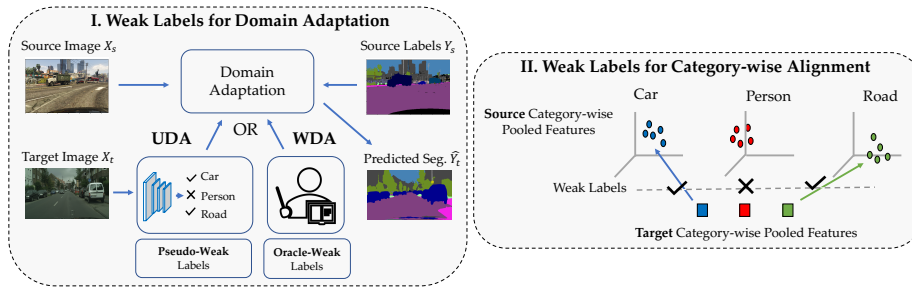


Fig. 1: Our work introduces two key ideas to adapt semantic segmentation models across domains. I: Using image-level weak annotations for domain adaptation, either estimated, i.e., pseudo-weak labels (Unsupervised Domain Adaptation, UDA) or acquired from a human oracle (Weakly-supervised Domain Adaptation (WDA)). II: We utilize weak labels to improve the category-wise feature alignment between the source and target domains.  $\checkmark/\times$  depicts weak labels, i.e., the categories present/absent in an image.

these methods are agnostic to the category, which may be problematic as the domain gap may vary across categories.

To alleviate the issue of lacking annotations in the target domain, we propose a concept of utilizing *weak labels* on the domain adaptation task for semantic segmentation, in the form of image- or point-level annotations in the target domain. Such weak labels can be used for category-wise alignment between the source and target domain, and also to enforce constraints on the categories present in an image. It is important to note that our weak labels could be estimated from the model prediction in the UDA setting, or provided by the human oracle in the weakly-supervised domain adaptation (WDA) paradigm (see left of Fig. 1). We are the first to introduce the WDA setting for semantic segmentation with image-level weak-labels, which is practically useful as collecting such annotations is much easier than pixel-wise annotations on the target domain. Benefiting from the concept of weak labels introduced in this paper, we aim to utilize such weak labels to act as an enabler for the interplay between the alignment and pseudo labeling procedures, as they are much less noisy compared to pixel-wise pseudo labels. Specifically, we use weak labels to perform both 1) image-level classification to identify the presence/absence of categories in an image as a regularization, and 2) category-wise domain alignment using such categorical labels. For the image-level classification task, weak labels help our model obtain a better pixel-wise attention map per category. Then, we utilize the category-wise attention maps as the guidance to further pool category-wise features for proposed domain alignment procedure (right of Fig. 1). Note that, although weak labels have been used in domain adaptation for object detection [24], our motivation is different from theirs. More specifically, [24] uses the weak labels to choose pseudo labels for self-training, while we formulate a general framework to learn from weak labels with different forms, i.e., UDA and WDA (image-level or point

supervision), as well as to improve feature alignment across domains using weak labels. We conduct experiments on the road scene segmentation problem from GTA5 [39]/SYNTHIA [40] to Cityscapes [11]. We perform extensive experiments to verify the usefulness of each component in the proposed framework, and show that our approach performs favorably against state-of-the-art algorithms for UDA. In addition, we show that our proposed method can be used for WDA and present its experimental results as a new benchmark. For the WDA setting, we also show that our method can incorporate various types of weak labels, such as image-level or point supervision. The **main contributions** of our work are: 1) we propose a concept of using weak labels to help domain adaptation for semantic segmentation; 2) we utilize weak labels to improve category-wise alignment for better feature space adaptation; and 3) we demonstrate that our method is applicable to both UDA and WDA settings.

## 2 Related Work

In this section, we discuss the literature of unsupervised domain adaptation (UDA) for image classification and semantic segmentation. In addition, we also discuss weakly-supervised methods for semantic segmentation.

**UDA for Image Classification.** The UDA task for image classification has been developed via aligning distributions across source and target domains. To this end, hand-crafted features [15,18] and deep features [16,51] have been considered to minimize the domain discrepancy and learn domain-invariant features. To further enhance the alignment procedure, maximum mean discrepancy [32] and adversarial learning [17,52] based approaches have been proposed. Recently, several algorithms focus on improving deep models [33,41,28,13], combining distance metric learning [44,45], utilizing pixel-level adaptation [3,47], or incorporating active learning [46].

**UDA for Semantic Segmentation.** Existing UDA methods in literature for semantic segmentation can be categorized primarily into two groups: domain alignment and pseudo-label self-training. For domain alignment, numerous algorithms focus on aligning distributions in the pixel [4,10,21,35,56,58], feature [7,22,57], and output [49,9] spaces. For pseudo-label re-training, current methods [43,60,30] aim to generate pixel-wise pseudo labels on the target images, which is utilized to finetune the segmentation model trained on the source domain.

To achieve better performance, recent works [14,29,50,54] attempt to combine the above two mechanisms. AdvEnt [54] adopts adversarial alignment and self-training in the entropy space, while BDL [29] combines output space and pixel-level adaptation with pseudo-label self-training in an iterative updating scheme. Moreover, Tsai et al. [50] propose a patch-level alignment method and show that their approach is complementary to existing modules such as output space adaptation and pseudo-label self-training. Similarly, Du et al. [14] integrate category-wise adversarial alignment with pixel-wise pseudo-labels, which may be noisy, leading to incorrect alignment. In addition, [14] needs to progressively

change a ratio for selecting pseudo-labels, and the final performance is sensitive to this chosen parameter.

Compared to the above-mentioned approaches, we propose to exploit weak labels by learning an image classification task, while improving domain alignment through category-wise attention maps. Furthermore, we show that our approach can be utilized even in the case where oracle-weak labels are available on the target domain, in which case the performance will be further improved.

**Weakly-supervised Semantic Segmentation.** In this paper, since we are specifically interested in how weak labels can help domain adaptation, we also discuss the literature for weakly-supervised semantic segmentation, which has been tackled through different types of weak labels, such as image-level [1,5,27,37,38], video-level [8,48,59], bounding box [36,12,25], scribble [31,53], and point [2] supervisions. Under this setting, these methods train the model using ground truth weak labels and perform testing in the same domain, which does not require domain adaptation. In contrast, we use a source domain with pixel-wise ground truth labels, but in the target domain, we consider pseudo-weak labels (UDA) or oracle-weak labels (WDA). As a result, we note that performance of weakly-supervised semantic segmentation methods which do not utilize any source domain, is usually much lower than the domain adaptation setting adopted in this paper, e.g., the mean IoU on Cityscapes is only 24.9% as shown in [42].

### 3 Domain Adaptation with Weak Labels

In this section, we first introduce the problem and then describe details of the proposed framework - the image-level classification module and category-wise alignment method using weak labels. Finally, we present our method of obtaining the weak labels for the UDA and WDA settings.

#### 3.1 Problem Definition

In the source domain, we have images and pixel-wise labels denoted as  $\mathcal{I}_s = \{X_s^i, Y_s^i\}_{i=1}^{N_s}$ . Whereas, our target dataset contains images and only image-level labels as  $\mathcal{I}_t = \{X_t^i, y_t^i\}_{i=1}^{N_t}$ . Note that  $X_s, X_t \in \mathbb{R}^{H \times W \times 3}$ ,  $Y_s \in \mathbb{B}^{H \times W \times C}$  with pixel-wise one-hot vectors,  $y_t \in \mathbb{B}^C$  is a multi-hot vector representing the categories present in the image and  $C$  is the number of categories, same for both the source and target datasets. Such image-level labels  $y_t$  are often termed as weak labels. We can either estimate them, in which case we call them pseudo-weak labels (Unsupervised Domain Adptation, UDA) or acquire them from a human oracle that is called oracle-weak labels (Weakly-supervised Domain Adaptation, WDA). We will further discuss details of acquiring weak labels in Section 3.6. Given such data, the problem is to adapt a segmentation model  $\mathbf{G}$  learned on the source dataset  $\mathcal{I}_s$  to the target dataset  $\mathcal{I}_t$ .

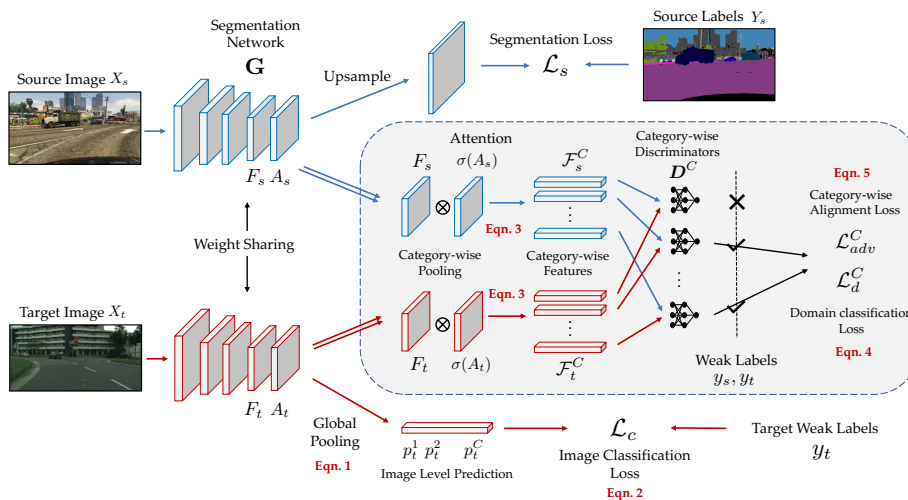


Fig. 2: The proposed architecture consists of the segmentation network  $G$  and the weak label module. We compute the pixel-wise segmentation loss  $\mathcal{L}_s$  for the source images and image classification loss  $\mathcal{L}_c$  using the weak labels  $y_t$  for the target images. Note that the weak labels can be estimated as pseudo-weak labels or provided by a human oracle. We then use the output prediction  $A$ , convert it to an attention map  $\sigma(A)$  and pool category-wise features  $\mathcal{F}^C$ . Next, these features are aligned between source and target domains using the category-wise alignment loss  $\mathcal{L}_{adv}^C$  guided by the category-wise discriminators  $D^C$  learned via the domain classification loss  $\mathcal{L}_d^C$ .

### 3.2 Algorithm Overview

Fig. 2 presents an overview of our proposed method. We first pass both the source and target images through the segmentation network  $G$  and obtain their features  $F_s, F_t \in \mathbb{R}^{H' \times W' \times 2048}$ , segmentation predictions  $A_s, A_t \in \mathbb{R}^{H' \times W' \times C}$ , and the upsampled pixel-wise predictions  $O_s, O_t \in \mathbb{R}^{H \times W \times C}$ . Note that  $H' (< H), W' (< W)$  are the downsampled spatial dimensions of the image after passing through the segmentation network. As a baseline, we use the source pixel-wise annotations to learn  $G$ , while aligning the output space distribution  $O_s$  and  $O_t$ , following [49].

In addition to having pixel-wise labels on the source data, we also have image-level weak labels on the target data. As discussed before, such weak labels can be either estimated (UDA) or acquired from an oracle (WDA). We then utilize these weak labels to update the segmentation network  $G$  in two different ways. First, we introduce a module which learns to predict the categories that are present in a target image. Second, we formulate a mechanism to align the features of each individual category between source and target domains. To this end, we use category-specific domain discriminators  $D^c$  guided by the weak labels to determine which categories should be aligned. In the following sections, we present these two modules in more detail.

### 3.3 Weak Labels for Category Classification

In order to predict whether a category is absent/present in a particular image, we define an image classification task using the weak labels, such that the segmentation network  $\mathbf{G}$  can discover those categories. Specifically, we use the weak labels  $y_t$  and learn to predict the categories present/absent in the target images. We first feed the target images  $X_t$  through  $\mathbf{G}$  to obtain the predictions  $A_t$  and then apply a global pooling layer to obtain a single vector of predictions for each category:

$$p_t^c = \sigma_s \left[ \frac{1}{k} \log \frac{1}{H'W'} \sum_{h',w'} \exp k A_t^{(h',w',c)} \right], \quad (1)$$

where  $\sigma_s$  is the sigmoid function such that  $p_t$  represents the probability that a particular category appears in an image. Note that (1) is a smooth approximation of the  $\max$  function. The higher the value of  $k$ , the better it approximates to  $\max$ . We set  $k = 1$  as we do not want the network to focus only on the maximum value of the prediction, which may be noisy, but also on other predictions that may have high values. Using  $p_t$  and the weak labels  $y_t$ , we can compute the category-wise binary cross-entropy loss:

$$\mathcal{L}_c(X_t; \mathbf{G}) = \sum_{c=1}^C -y_t^c \log(p_t^c) - (1 - y_t^c) \log(1 - p_t^c). \quad (2)$$

This is shown at the bottom stream of Fig. 2. This loss function  $\mathcal{L}_c$  helps to identify the categories which are absent/present in a particular image and enforces the segmentation network  $\mathbf{G}$  to pay attention to those objects/stuff that are partially identified when the source model is used directly on the target images.

### 3.4 Weak Labels for Feature Alignment

The classification loss using weak labels introduced in (2) regularizes the network focusing on certain categories. However, distribution alignment across the source and target domains is not considered yet. As discussed in the previous section, methods in literature either align feature space [22] or output space [49] across domains. However, such alignment is agnostic to the category, so it may align features of categories that are not present in certain images. Moreover, features belonging to different categories may have different domain gaps. Thereby, performing category-wise alignment could be beneficial but has not been widely studied in UDA for semantic segmentation. Although an existing work [14] attempts to align category-wise features, it utilizes pixel-wise pseudo labels, which may be noisy, and performs alignment in a high-dimensional feature space, which is not only difficult to optimize but also requires more computations (more discussions are provided in the experimental section).

To alleviate all the above issues, we use image-level weak labels to perform category-wise alignment in the feature space. Specifically, we obtain the category-wise features for each image via an attention map, i.e., segmentation prediction,

guided by our classification module using weak labels, and then align these features between the source and target domains. We next discuss the category-wise feature pooling mechanism followed by the adversarial alignment technique.

**Category-wise Feature Pooling.** Given the last layer features  $F$  and the segmentation prediction  $A$ , we obtain the category-wise features by using the prediction as an attention over the features. Specifically, we obtain the category-wise feature  $\mathcal{F}^c$  as a 2048-dimensional vector for the  $c^{th}$  category as follows:

$$\mathcal{F}^c = \sum_{h', w'} \sigma(A)^{(h', w', c)} F^{(h', w')}, \quad (3)$$

where  $\sigma(A)$  is a tensor of dimension  $H' \times W' \times C$ , with each channel along the category dimension representing the category-wise attention obtained by the softmax operation  $\sigma$  over the spatial dimensions. As a result,  $\sigma(A)^{(h', w', c)}$  is a scalar and  $F^{(h', w')}$  is a 2048-dimensional vector, while  $\mathcal{F}^c$  is the summed feature of  $F^{(h', w')}$  weighted by  $\sigma(A)^{(h', w', c)}$  over the spatial map  $H' \times W'$ . Note that we drop the subscripts  $s, t$  for source and target, as we employ the same operation to obtain the category-wise features for both domains. We next present the mechanism to align these features across domains. Note that we will use  $\mathcal{F}^c$  to denote the pooled feature for the  $c^{th}$  category and  $\mathcal{F}^C$  to denote the set of pooled features for all the categories. Category-wise feature pooling is shown in the middle of Fig. 2.

**Category-wise Feature Alignment.** To learn the segmentation network  $\mathbf{G}$  such that the source and target category-wise features are aligned, we use an adversarial loss while using category-specific discriminators  $\mathbf{D}^C = \{\mathbf{D}^c\}_{c=1}^C$ . The reason of using category-specific discriminators is to ensure that the feature distribution for each category could be aligned independently, which avoids the noisy distribution modeling from a mixture of categories. In practice, we train  $C$  distinct category-specific discriminators to distinguish between category-wise features drawn from the source and target images. The loss function to train the discriminators  $\mathbf{D}^C$  is as follows:

$$\mathcal{L}_d^C(\mathcal{F}_s^C, \mathcal{F}_t^C; \mathbf{D}^C) = \sum_{c=1}^C -y_s^c \log \mathbf{D}^c(\mathcal{F}_s^c) - y_t^c \log (1 - \mathbf{D}^c(\mathcal{F}_t^c)). \quad (4)$$

Note that, while training the discriminators, we only compute the loss for those categories which are present in the particular image via the weak labels  $y_s, y_t \in \mathbb{B}^C$  that indicate whether a category occurs in an image or not. Then, the adversarial loss for the target images to train the segmentation network  $\mathbf{G}$  can be expressed as follows:

$$\mathcal{L}_{adv}^C(\mathcal{F}_t^C; \mathbf{G}, \mathbf{D}^C) = \sum_{c=1}^C -y_t^c \log \mathbf{D}^c(\mathcal{F}_t^c). \quad (5)$$

Similarly, we use the target weak labels  $y_t$  to align only those categories present in the target image. By minimizing  $\mathcal{L}_{adv}^C$ , the segmentation network tries to fool the

discriminator by maximizing the probability of the target category-wise feature being considered as drawn from the source distribution. These loss functions in (4) and (5) are obtained in the right of the middle box in Fig. 2.

### 3.5 Network Optimization

**Discriminator Training.** We learn a set of  $C$  distinct discriminators for each category  $c$ . We use the source and target images to train the discriminators, which learn to distinguish between the category-wise features drawn from either the source or the target domain. The optimization problem to train the discriminator can be expressed as:  $\min_{\mathcal{D}^C} \mathcal{L}_d^C(\mathcal{F}_s^C, \mathcal{F}_t^C)$ . Note that each discriminator is trained only with features pooled specific to that particular category. Therefore, given an image, we only update those discriminators corresponding to those categories which are present in the image and ignore the rest.

**Segmentation Network Training.** We train the segmentation network with the pixel-wise cross-entropy loss  $\mathcal{L}_s$  on the source images, image classification loss  $\mathcal{L}_c$  and adversarial loss  $\mathcal{L}_{adv}^C$  on the target images. We combine these loss functions to learn  $\mathbf{G}$  as follows :

$$\min_{\mathbf{G}} \mathcal{L}_s(X_s) + \lambda_c \mathcal{L}_c(X_t) + \lambda_d \mathcal{L}_{adv}^C(\mathcal{F}_t^C). \quad (6)$$

We follow the standard GAN training procedure [19] to alternatively update  $\mathbf{G}$  and  $\mathcal{D}^C$ . Note that, computing  $\mathcal{L}_{adv}^C$  involves the category-wise discriminators  $\mathcal{D}^C$ . Therefore, we fix  $\mathcal{D}^C$  and backpropagate gradients only for the segmentation network  $\mathbf{G}$ .

### 3.6 Acquiring Weak Labels

In the above sections, we have proposed a mechanism to utilize image-level weak labels of the target images and adapt the segmentation model between source and target domains. In this section, we explain two methods to obtain such image-level weak labels.

**Pseudo-Weak Labels (UDA).** One way of obtaining weak labels is to directly estimate them using the data we have, i.e., source images/labels and target images, which is the unsupervised domain adaptation (UDA) setting. In this work, we utilize the baseline model [49] to adapt a model learned from the source to the target domain, and then obtain the weak labels of the target images as follows:

$$y_t^c = \begin{cases} 1, & \text{if } p_t^c > T, \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where  $p_t^c$  is the probability for category  $c$  as computed in (1) and  $T$  is a threshold, which we set to 0.2 in all the experiments unless specified otherwise. In practice, we compute the weak labels online during training and avoid any additional



inference step. Specifically, we forward a target image, obtain the weak labels using (7), and then compute the loss functions in (6). As the weak labels obtained in this manner do not require human supervision, adaptation using such labels is unsupervised.

**Oracle-Weak Labels (WDA).** In this form, we obtain the weak labels by querying a human oracle to provide a list of the categories that occur in the target image. As we use supervision from an oracle on the target images, we refer to this as weakly-supervised domain adaptation (WDA). It is worth mentioning that the WDA setting could be practically useful, as collecting such human annotated weak labels is much easier than pixel-wise annotations. Also, there has not been any prior research involving this setting for domain adaptation.

To show that our method can use different forms of oracle-weak labels, we further introduce the point supervision as in [2], which only increases effort by a small amount compared to the image-level supervision. In this scenario, we randomly obtain one pixel coordinate of each category that belongs in the image, i.e., the set of tuples  $\{(h^c, w^c, c) | \forall y_t^c = 1\}$ . For an image, we compute the loss as follows:  $\mathcal{L}_{point} = -\sum_{\forall y_t^c=1} y_t^c \log(O_t^{(h^c, w^c, c)})$ , where  $O_t \in \mathbb{R}^{H \times W \times C}$  is the output prediction of target after pixel-wise softmax.

## 4 Experimental Results

In this section, we perform an evaluation of our domain adaptation framework for semantic segmentation. We present the results for using both pseudo-weak labels, i.e., unsupervised domain adaptation (UDA) and human oracle-weak labels, i.e., weakly-supervised domain adaptation (WDA) and compare it with existing state-of-the-art methods. We also perform ablation studies to analyse the benefit of using pseudo/oracle-weak labels via our proposed weak-label classification module and category-wise alignment.

**Datasets and Metric.** We evaluate our domain adaptation method under the Sim-to-Real case with two different source-target scenarios. First, we adapt from GTA5 [39] to the Cityscapes dataset [11]. Second, we use SYNTHIA [40] as the source and Cityscapes as the target, which has a larger domain gap than the former case. For all experiments, we use the Intersection-over-Union (IoU) ratio as the metric. For SYNTHIA  $\rightarrow$  Cityscapes, following the literature [54], we report the performance averaged over 16 categories (listed in Table 2) and 13 categories (removing wall, fence and pole), which we denote as mIoU\*.

**Network Architectures.** For the segmentation network  $\mathcal{G}$ , to have a fair comparison with works in literature, we use the DeepLab-v2 framework [6] with the ResNet-101 [20] architecture. We extract features  $F_s, F_t$  before the Atrous Spatial Pyramid Pooling (ASPP) layer. For the category-wise discriminators  $D^C = \{D^c\}_{c=1}^C$ , we use  $C$  separate networks, where each consists of three fully-connected layers, having number of nodes  $\{2048, 2048, 1\}$  with ReLU activation.

**Training Details.** We implement our framework using PyTorch on a single Titan X GPU with 12G memory for all our experiments. We use the SGD method to optimize the segmentation network and the Adam optimizer [26] to train the discriminators. We set the initial learning rates to be  $2.5 \times 10^{-4}$  and  $1 \times 10^{-4}$  for the segmentation network and discriminators, with polynomial decay of power 0.9 [6]. As a common practice in weakly-supervised semantic segmentation [1], we use Dropout of 0.1 and 0.3 for oracle-weak labels and pseudo-weak labels respectively, on the spatial predictions before computing the loss  $\mathcal{L}_c$ . We choose  $\lambda_c$  to be 0.2 for oracle-weak labels and use a smaller  $\lambda_c = 0.01$  for pseudo-weak labels to account for its inaccurate prediction. For the weight on the category-wise adversarial loss  $\mathcal{L}_{adv}^C$ , we set  $\lambda_{adv} = 0.001$ . For experiments using pseudo weak labels, to avoid noisy pseudo weak label prediction in the early training stage, we first train the segmentation baseline network using [49] for 60K iterations. Then, we include the proposed weak-label classification and alignment procedure, and train the entire framework.

#### 4.1 Comparison with State-of-the-art Methods

**Unsupervised Domain Adaptation (UDA).** We compare our method with existing state-of-the-art UDA methods in Table 1 for GTA5→Cityscapes and in Table 2 for SYNTHIA→Cityscapes. Recent methods [4,21,29,50] show that adapting images from source to target on the pixel level and then adding those translated source images in training enhances the performance. We follow this practice in the final model via adding these adapted images to the source dataset, as their pixel-wise annotations do not change after adaptation. Thus adaptation using weak labels aligns the features not only between the original source and target images, but also between the translated source images and the target images. We show that our method is also complementary to pixel-level adaptation.

**Discussions.** In terms of applied techniques, e.g, pseudo-label re-training and domain alignment, the closest comparisons to our method are DISE [4], BDL [29], and Patch Space alignment [50]. We show that our method performs favorably against these approaches on both benchmarks. This can be attributed to our introduced concept of using weak labels, in which our UDA model explores pseudo-weak image-level labels, instead of using pixel-level pseudo-labels [50,29] that may be noisy and degrade the performance. In addition, these methods do not perform domain alignment guided by such pseudo labels, whereas we use weak labels to enable our category-wise alignment procedure.

The only prior work that adopts category-wise feature alignment is SSF-DAN [14]. However, our method is different from theirs in three aspects: 1) We introduce the weak-label classification module to take advantage of image-level weak labels that enables an efficient feature alignment process and the novel WDA setting; 2) Our unified framework can be applied for both UDA and WDA settings with various types of supervisions; 3) Due to the introduced weak-label module, our category-wise feature alignment is operated in the pooled feature space in (3) guided by an attention map, rather than in a much higher-dimensional spatial

Table 1: Results of adapting GTA5 to Cityscapes. The top group is for UDA, while the bottom group presents our method’s performance using the oracle-weak labels for WDA that use either image-level or point supervision.

		GTA5 → Cityscapes																			
Method	road	sidewalk	building	wall	fence	pole	light	sign	veg	terrain	sky	person	rider	car	truck	bus	train	mbike	bike	mIoU	
No Adapt.	75.8	16.8	77.2	12.5	21.0	25.5	30.1	20.1	81.3	24.6	70.3	53.8	26.4	49.9	17.2	25.9	6.5	25.3	36.0	36.6	
Road [9]	76.3	36.1	69.6	28.6	22.4	28.6	29.3	14.8	82.3	35.3	72.9	54.4	17.8	78.9	27.7	30.3	4.0	24.9	12.6	39.4	
AdaptOutput [49]	86.5	25.9	79.8	22.1	20.0	23.6	33.1	21.8	81.8	25.9	75.9	57.3	26.2	76.3	29.8	32.1	7.2	29.5	32.5	41.4	
AdvEnt [54]	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	<b>38.5</b>	44.5	1.7	31.6	32.4	45.5	
CLAN [34]	87.0	27.1	79.6	27.3	23.3	28.3	35.5	24.2	83.6	27.4	74.2	58.6	28.0	76.2	33.1	36.7	6.7	<b>31.9</b>	31.4	43.2	
SWD [28]	<b>92.0</b>	46.4	82.4	24.8	24.0	<b>35.1</b>	33.4	34.2	83.6	30.4	80.9	56.9	21.9	82.0	24.4	28.7	6.1	25.0	33.6	44.5	
SSF-DAN [14]	90.3	38.9	81.7	24.8	22.9	30.5	37.0	21.2	<b>84.8</b>	38.8	76.9	58.8	30.7	<b>85.7</b>	30.6	38.1	5.9	28.3	36.9	45.4	
DISE [4]	91.5	47.5	82.5	31.3	25.6	33.0	33.7	25.8	82.7	28.8	82.7	<b>62.4</b>	<b>30.8</b>	85.2	27.7	34.5	6.4	25.2	24.4	45.4	
BDL [29]	91.4	47.9	<b>84.2</b>	<b>32.4</b>	26.0	31.8	37.3	33.0	83.3	<b>39.2</b>	79.2	57.7	25.6	81.3	36.3	39.7	2.6	31.3	33.5	47.2	
AdaptPatch [50]	92.3	<b>51.9</b>	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	<b>46.3</b>	2.2	29.5	32.3	46.5	
Ours (UDA)	91.6	47.4	84.0	30.4	<b>28.3</b>	31.4	<b>37.4</b>	<b>35.4</b>	83.9	38.3	<b>83.9</b>	61.2	28.2	83.7	28.8	41.3	<b>8.8</b>	24.7	<b>46.4</b>	<b>48.2</b>	
Ours (WDA: Image)	89.5	54.1	83.2	31.7	34.2	37.1	43.2	39.1	85.1	39.6	85.9	61.3	34.1	82.3	42.3	51.9	34.4	33.1	45.4	53.0	
Ours (WDA: Point)	94.0	62.7	86.3	36.5	32.8	38.4	44.9	51.0	86.1	43.4	87.7	66.4	36.5	87.9	44.1	58.8	23.2	35.6	55.9	56.4	

space as in [14] that uses pixel-wise pseudo-labels. This essentially improves the training efficiency compared to [14], which requires a GPU with 16 GB memory as their discriminator needs much more computation time ( $> 20\times$ ) and GPU memory ( $> 8\times$ ) compared to our combined output space and category-wise discriminators. Also, the discriminators in [14] require 130 GFLOPS, whereas our discriminators require a total of only 0.5 GFLOPS.

## 4.2 Weakly-supervised Domain Adaptation (WDA)

**Image-level Supervision.** We present the results of our method when using oracle-weak labels (obtained from the ground truth of the training set) in the last rows of Table 1 and 2. To the best of our knowledge, we are the first to work on WDA, i.e., using human oracle-weak labels on domain adaptation for semantic segmentation, and there are no other methods to compare against in the literature. From the results, it is interesting to note that the major boost in performance using WDA compared to UDA occurs for categories such as truck, bus, train, and motorbike for both cases using GTA5 and SYNTHIA as the source domain. One reason is that those categories are most underrepresented in both the source and the target datasets. Thus, they are not predicted in most of the target images, but using the oracle-weak labels helps to identify them better.

**Point Supervision.** We introduce another interesting setting of point supervision as in [2], which adds only a slight increase of annotation time compared to the image-level supervision. We follow [2] and randomly sample one pixel per category in each target image as the supervision. Note that, all the details and

Table 2: Results of adapting SYNTHIA to Cityscapes. The top group is for UDA, while the bottom group presents the WDA setting using oracle-weak labels. mIoU and mIoU\* are averaged over 16 and 13 categories.

SYNTHIA $\rightarrow$ Cityscapes																		
Method	road	sidewalk	building	wall	fence	pole	light	sign	veg	sky	person	rider	car	bus	mbike	bike	mIoU	mIoU*
No Adapt.	55.6	23.8	74.6	9.2	0.2	24.4	6.1	12.1	74.8	79.0	55.3	19.1	39.6	23.3	13.7	25.0	33.5	38.6
AdaptOutput [49]	79.2	37.2	78.8	10.5	0.3	25.1	9.9	10.5	78.2	80.5	53.5	19.6	67.0	29.5	21.6	31.3	39.5	45.9
AdvEnt [54]	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
CLAN [34]	81.3	37.0	80.1	-	-	-	16.1	13.7	78.2	81.5	53.4	21.2	73.0	32.9	22.6	30.7	-	47.8
SWD [28]	82.4	33.2	<b>82.5</b>	-	-	-	<b>22.6</b>	<b>19.7</b>	<b>83.7</b>	78.8	44.0	17.9	75.4	30.2	14.4	39.9	-	48.1
DADA [55]	89.2	44.8	81.4	6.8	0.3	26.2	8.6	11.1	81.8	84.0	54.7	19.3	79.7	<b>40.7</b>	14.0	38.8	42.6	49.8
SSF-DAN [14]	84.6	41.7	80.8	-	-	-	11.5	14.7	80.8	<b>85.3</b>	57.5	21.6	82.0	36.0	19.3	34.5	-	50.0
DISE [4]	91.7	53.5	77.1	2.5	0.2	<b>27.1</b>	6.2	7.6	78.4	81.2	55.8	19.2	<b>82.3</b>	30.3	17.1	34.3	41.5	48.8
AdaptPatch [50]	82.4	38.0	78.6	8.7	<b>0.6</b>	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
Ours (UDA)	<b>92.0</b>	<b>53.5</b>	80.9	<b>11.4</b>	0.4	21.8	3.8	6.0	81.6	84.4	<b>60.8</b>	<b>24.4</b>	80.5	39.0	<b>26.0</b>	<b>41.7</b>	<b>44.3</b>	<b>51.9</b>
Ours (WDA: Image)	92.3	51.9	81.9	21.1	1.1	26.6	22.0	24.8	81.7	87.0	63.1	33.3	83.6	50.7	33.5	54.7	50.6	58.5
Ours (WDA: Point)	94.9	63.2	85.0	27.3	24.2	34.9	37.3	50.8	84.4	88.2	60.6	36.3	86.4	43.2	36.5	61.3	57.2	63.7

Table 3: Ablation of the proposed loss functions for GTA5 $\rightarrow$ Cityscapes.

GTA5 $\rightarrow$ Cityscapes				
	Supervision	$\mathcal{L}_c$	$\mathcal{L}_{adv}^C$	mIoU
UDA	No Adapt.			36.6
	Baseline [49]			41.4
	Pseudo-Weak	✓		44.2
		✓	✓	<b>45.6</b>
WDA	Oracle-Weak	✓		50.8
		✓	✓	<b>52.1</b>

Table 4: Ablation of the proposed loss functions for SYNTHIA $\rightarrow$ Cityscapes.

SYNTHIA $\rightarrow$ Cityscapes					
	Supervision	$\mathcal{L}_c$	$\mathcal{L}_{adv}^C$	mIoU	mIoU*
UDA	No Adapt.			33.5	38.6
	Baseline [49]			39.5	45.9
	Pseudo-Weak	✓		41.7	49.0
		✓	✓	<b>42.7</b>	<b>49.9</b>
WDA	Oracle-Weak	✓		47.8	56.0
		✓	✓	<b>49.2</b>	<b>57.2</b>

the modules are the same during training in this setting. In Table 1 and 2, the results show that using point supervision improves performance (3.4 – 6.6%) on both benchmarks compared to the image-level supervision. This shows that our method is a general framework that can be applied to the conventional UDA setting as well as the WDA setting using either image-level or point supervision, while all the settings achieve consistent performance gains.

Fig. 3b shows a comparison of annotation time v.s. performance for various levels of supervision. With low annotation cost in WDA cases, our model bridges the gap in performance between UDA and full supervision ones (more results are shown in the supplementary material). Note that, other forms of weak labels such as object count and density can also be effective.

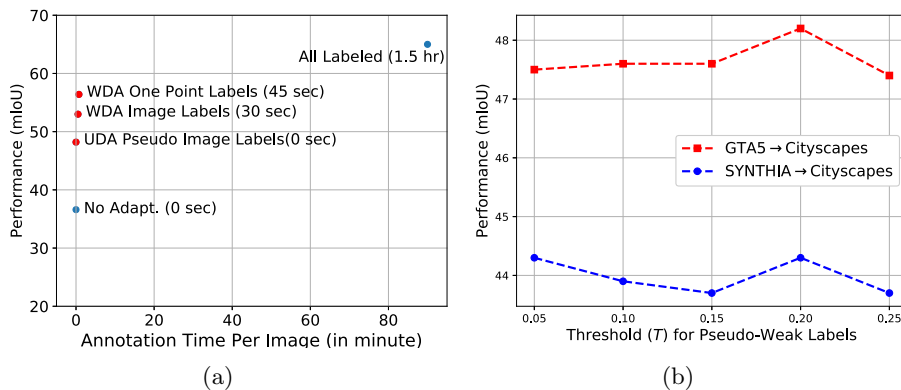


Fig. 3: (a) Performance comparison on GTA5→Cityscapes with different levels of supervision on target images: no target labels (“No Adapt.” and “UDA”), weak image labels (30 seconds), one point labels (45 seconds), and fully-supervised setting with all pixels labeled (“All Labeled”) that takes 1.5 hours per image according to [11]. (b) Performance of our method on GTA5→Cityscapes with variations in the threshold, i.e.,  $T$  in (7), for obtaining the pseudo-weak labels.

### 4.3 Ablation Study

**Effect of Weak Labels.** We show results for using both pseudo-weak labels as well as human oracle-weak labels. Table 3 and 4 present the results for different combinations of the modules used in our framework without pixel-level adaptation [21]. It is interesting to note that on GTA5→Cityscapes, even when using pseudo-weak labels, our method obtains a 4.2% boost in performance (41.4 → 45.6), as well as a 3 – 4% boost for SYNTHIA→Cityscapes. In addition, as expected, using oracle-weak labels performs better than pseudo-weak labels by 6.5% on GTA5→Cityscapes and 6.5 – 7.3% on SYNTHIA→Cityscapes. It is also interesting to note that using the category-wise alignment consistently improves the performance for all the cases, i.e., different types of weak labels and for different datasets.

**Effect of Pseudo-Weak Label Threshold.** We use a threshold  $T$  in (7) to convert the image-level prediction probability to a multi-hot vector denoting the pseudo-weak labels that indicates absence/presence of the categories. Note that the threshold is on a probability between 0 and 1. We then study the effect of  $T$  by varying it and plot the performance in Fig. 3a on GTA5→Cityscapes. The figure shows that our model generally works well with  $T$  in a range of 0.05 to 0.25. However, when we make  $T$  larger than 0.3, the performance starts to drop significantly, as in this case, the recall of the pseudo-weak labels would be very low compared with the oracle-weak labels (i.e., ground truths), which makes the segmentation network fail to predict most categories.

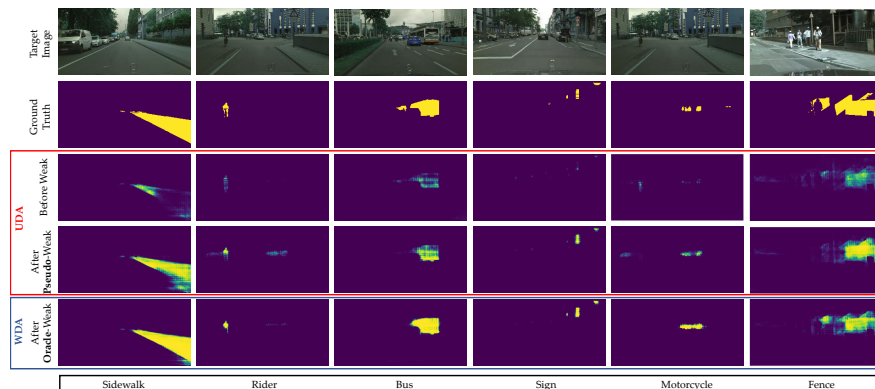


Fig. 4: Visualizations of category-wise segmentation prediction probability before and after using the pseudo-weak labels on GTA5→Cityscapes. Before adaptation, the network only highlights the areas partially with low probability, while using the pseudo-weak labels helps the adapted model obtain much better segments, and is closer to the model using oracle-weak labels.

**Output Space Visualization.** We present some visualizations of the segmentation prediction probability for each category in Fig. 4. Before using any weak labels (third row), the probabilities may be low, even though there is a category present in that image. However, based on these initial predictions, our model can estimate the categories and then enforce their presence/absence explicitly in the proposed classification loss and alignment loss. The fourth row in Fig. 4 shows that such pseudo-weak labels help the network discover object/stuff regions towards better segmentation. For example, the fourth and fifth column shows that, although the original prediction probabilities are quite low, results using pseudo-weak labels are estimated correctly. Moreover, the last row shows that the predictions can be further improved when we have oracle-weak labels.

## 5 Conclusions

In this paper, we use weak labels to improve domain adaptation for semantic segmentation in both the UDA and WDA settings, with the latter being a novel setting. Specifically, we design an image-level classification module using weak labels, enforcing the network to pay attention to categories that are present in the image. With such a guidance from weak labels, we further utilize a category-wise alignment method to improve adversarial alignment in the feature space. Based on these two mechanisms, our formulation generalizes both to pseudo-weak and oracle-weak labels. We conduct extensive ablation studies to validate our approach against state-of-the-art UDA approaches.

**Acknowledgment.** This work was a part of Sujoy Paul’s internship at NEC Labs America. This work was also partially funded by NSF grant 1724341.

## References

1. Ahn, J., Kwak, S.: Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In: CVPR (2018) 4, 10
2. Bearman, A., Russakovsky, O., Ferrari, V., Fei-Fei, L.: What's the point: Semantic segmentation with point supervision. In: ECCV (2016) 4, 9, 11
3. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: CVPR (2017) 3
4. Chang, W.L., Wang, H.P., Peng, W.H., Chiu, W.C.: All about structure: Adapting structural information across domains for boosting semantic segmentation. In: CVPR (2019) 3, 10, 11, 12
5. Chang, Y.T., Wang, Q., Hung, W.C., Piramuthu, R., Tsai, Y.H., Yang, M.H.: Weakly-supervised semantic segmentation via sub-category exploration. In: CVPR (2020) 4
6. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. CoRR abs/1606.00915 (2016) 9, 10
7. Chen, Y.H., Chen, W.Y., Chen, Y.T., Tsai, B.C., Wang, Y.C.F., Sun, M.: No more discrimination: Cross city adaptation of road scene segmenters. In: ICCV (2017) 3
8. Chen, Y.W., Tsai, Y.H., Lin, Y.Y., Yang, M.H.: Vostr: Video object segmentation via transferable representations. International Journal of Computer Vision (02 2020) 4
9. Chen, Y., Li, W., Gool, L.V.: Road: Reality oriented adaptation for semantic segmentation of urban scenes. In: CVPR (2018) 1, 3, 11
10. Choi, J., Kim, T., Kim, C.: Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In: ICCV (2019) 3
11. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016) 3, 9, 13
12. Dai, J., He, K., Sun, J.: Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In: ICCV (2015) 4
13. Dai, S., Sohn, K., Tsai, Y.H., Carin, L., Chandraker, M.: Adaptation across extreme variations using unlabeled domain bridges. arXiv preprint arXiv:1906.02238 (2019) 3
14. Du, L., Tan, J., Yang, H., Feng, J., Xue, X., Zheng, Q., Ye, X., Zhang, X.: Ssf-dan: Separated semantic feature based domain adaptation network for semantic segmentation. In: ICCV (2019) 3, 6, 10, 11, 12
15. Fernando, B., Habrard, A., Sebban, M., Tuytelaars, T.: Unsupervised visual domain adaptation using subspace alignment. In: ICCV (2013) 3
16. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: ICML (2015) 3
17. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. In: JMLR (2016) 3
18. Gong, B., Shi, Y., Sha, F., Grauman, K.: Geodesic flow kernel for unsupervised domain adaptation. In: CVPR (2012) 3
19. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014) 8

20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) [9](#)
21. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A.A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: ICML (2018) [1](#), [3](#), [10](#), [13](#)
22. Hoffman, J., Wang, D., Yu, F., Darrell, T.: Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. CoRR [abs/1612.02649](#) (2016) [1](#), [3](#), [6](#)
23. Hung, W.C., Tsai, Y.H., Liou, Y.T., Lin, Y.Y., Yang, M.H.: Adversarial learning for semi-supervised semantic segmentation. In: BMVC (2018) [1](#)
24. Inoue, N., Furuta, R., Yamasaki, T., Aizawa, K.: Cross-domain weakly-supervised object detection through progressive domain adaptation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5001–5009 (2018) [2](#)
25. Khoreva, A., Benenson, R., Hosang, J., Hein, M., Schiele, B.: Simple does it: Weakly supervised instance and semantic segmentation. In: CVPR (2017) [4](#)
26. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015) [10](#)
27. Kolesnikov, A., Lampert, C.H.: Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In: ECCV (2016) [4](#)
28. Lee, C.Y., Batra, T., Baig, M.H., Ulbricht, D.: Sliced wasserstein discrepancy for unsupervised domain adaptation. In: CVPR (2019) [3](#), [11](#), [12](#)
29. Li, Y., Yuan, L., Vasconcelos, N.: Bidirectional learning for domain adaptation of semantic segmentation. In: CVPR (2019) [1](#), [3](#), [10](#), [11](#)
30. Lian, Q., Lv, F., Duan, L., Gong, B.: Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: A non-adversarial approach. In: ICCV (2019) [3](#)
31. Lin, D., Dai, J., Jia, J., He, K., Sun, J.: Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In: CVPR (2016) [4](#)
32. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: ICML (2015) [3](#)
33. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. In: NIPS (2016) [3](#)
34. Luo, Y., Zheng, L., Guan, T., Yu, J., Yang, Y.: Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In: CVPR (2019) [11](#), [12](#)
35. Murez, Z., Kolouri, S., Kriegman, D., Ramamoorthi, R., Kim, K.: Image to image translation for domain adaptation. In: CVPR (2018) [1](#), [3](#)
36. Papandreou, G., Chen, L.C., Murphy, K., Yuille, A.L.: Weakly-and semi-supervised learning of a dcnn for semantic image segmentation. In: ICCV (2015) [4](#)
37. Pathak, D., Krahenbuhl, P., Darrell, T.: Constrained convolutional neural networks for weakly supervised segmentation. In: ICCV (2015) [4](#)
38. Pinheiro, P.O., Collobert, R.: From image-level to pixel-level labeling with convolutional networks. In: CVPR (2015) [4](#)
39. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV (2016) [3](#), [9](#)
40. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The SYNTHIA Dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: CVPR (2016) [3](#), [9](#)
41. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: CVPR (2018) [3](#)



42. Saleh, F.S., Aliakbarian, M.S., Salzmann, M., Petersson, L., Alvarez, J.M.: Bringing background into the foreground: Making all classes equal in weakly-supervised video semantic segmentation. In: ICCV (2017) [4](#)
43. Saleh, F.S., Aliakbarian, M.S., Salzmann, M., Petersson, L., Alvarez, J.M.: Effective use of synthetic data for urban scene semantic segmentation. In: ECCV (2018) [3](#)
44. Sohn, K., Liu, S., Zhong, G., Yu, X., Yang, M.H., Chandraker, M.: Unsupervised domain adaptation for face recognition in unlabeled videos. In: ICCV (2017) [3](#)
45. Sohn, K., Shang, W., Yu, X., Chandraker, M.: Unsupervised domain adaptation for distance metric learning. In: ICLR (2019) [3](#)
46. Su, J.C., Tsai, Y.H., Sohn, K., Liu, B., Maji, S., Chandraker, M.: Active adversarial domain adaptation. In: WACV (2020) [3](#)
47. Tran, L., Sohn, K., Yu, X., Liu, X., Chandraker, M.: Gotta adapt 'em all: Joint pixel and feature-level domain adaptation for recognition in the wild. In: CVPR (2019) [3](#)
48. Tsai, Y.H., Zhong, G., Yang, M.H.: Semantic co-segmentation in videos. In: ECCV (2016) [4](#)
49. Tsai, Y.H., Hung, W.C., Schuler, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: CVPR (2018) [1](#), [3](#), [5](#), [6](#), [8](#), [10](#), [11](#), [12](#)
50. Tsai, Y.H., Sohn, K., Schuler, S., Chandraker, M.: Domain adaptation for structured output via discriminative patch representations. In: ICCV (2019) [1](#), [3](#), [10](#), [11](#), [12](#)
51. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: ICCV (2015) [3](#)
52. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: CVPR (2017) [3](#)
53. Vernaza, P., Chandraker, M.: Learning random-walk label propagation for weakly-supervised semantic segmentation. In: CVPR (2017) [4](#)
54. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: CVPR (2019) [3](#), [9](#), [11](#), [12](#)
55. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Dada: Depth-aware domain adaptation in semantic segmentation. In: ICCV (2019) [12](#)
56. Wu, Z., Han, X., Lin, Y.L., Uzunbas, M.G., Goldstein, T., Lim, S.N., Davis, L.S.: Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In: ECCV (2018) [3](#)
57. Zhang, Y., David, P., Gong, B.: Curriculum domain adaptation for semantic segmentation of urban scenes. In: ICCV (2017) [1](#), [3](#)
58. Zhang, Y., Qiu, Z., Yao, T., Liu, D., Mei, T.: Fully convolutional adaptation networks for semantic segmentation. In: CVPR (2018) [3](#)
59. Zhong, G., Tsai, Y.H., Yang, M.H.: Weakly-supervised video scene co-parsing. In: ACCV (2016) [4](#)
60. Zou, Y., Yu, Z., Kumar, B.V.K.V., Wang, J.: Domain adaptation for semantic segmentation via class-balanced self-training. In: ECCV (2018) [1](#), [3](#)