

# Bi-directional Contrastive Learning for Domain Adaptive Semantic Segmentation

Geon Lee, Chanho Eom, Wonkyung Lee, Hyekang Park, and Bumsub Ham\*

<https://cvlab.yonsei.ac.kr/projects/DASS>

Yonsei University

**Abstract.** We present a novel unsupervised domain adaptation method for semantic segmentation that generalizes a model trained with source images and corresponding ground-truth labels to a target domain. A key to domain adaptive semantic segmentation is to learn domain-invariant and discriminative features without target ground-truth labels. To this end, we propose a bi-directional pixel-prototype contrastive learning framework that minimizes intra-class variations of features for the same object class, while maximizing inter-class variations for different ones, regardless of domains. Specifically, our framework aligns pixel-level features and a prototype of the same object class in target and source images (i.e., positive pairs), respectively, sets them apart for different classes (i.e., negative pairs), and performs the alignment and separation processes toward the other direction with pixel-level features in the source image and a prototype in the target image. The cross-domain matching encourages domain-invariant feature representations, while the bidirectional pixel-prototype correspondences aggregate features for the same object class, providing discriminative features. To establish training pairs for contrastive learning, we propose to generate dynamic pseudo labels of target images using a non-parametric label transfer, that is, pixel-prototype correspondences across different domains. We also present a calibration method compensating class-wise domain biases of prototypes gradually during training. Experimental results on standard benchmarks including GTA5  $\rightarrow$  Cityscapes and SYNTHIA  $\rightarrow$  Cityscapes demonstrate the effectiveness of our framework.

**Keywords:** Bi-directional contrastive learning, domain adaptive semantic segmentation, dynamic pseudo label

## 1 Introduction

Semantic segmentation is to assign a semantic label to each pixel in an image. In the past decade, supervised methods based on convolutional neural networks (CNNs) [1, 26, 56, 33, 17, 41] have achieved remarkable improvements in semantic segmentation. Training networks for the dense prediction task generally requires lots of pixel-level labels. Annotating pixel-level labels of high-resolution

---

\* Corresponding author.

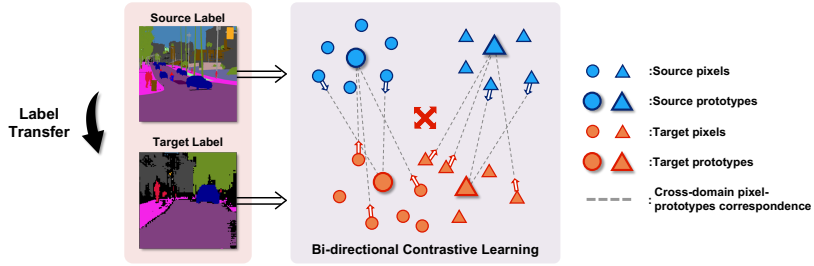


Fig. 1: An illustration of our framework for UDASS. We generate pseudo labels of target images using a nonparametric label transfer. We then perform bi-directional pixel-prototype contrastive learning. This encourages pixel-level features in a target image and a prototype of the same object class in a source domain to pull each other, while setting them apart for different ones. We also perform the alignment and separation in a reverse direction, with pixel-level features of source images and prototypes of a target domain.

images is, however, significantly labor-intensive and time-consuming. For example, annotating the labels for an image of size  $2048 \times 1024$  in Cityscapes [6] takes about 90 minutes. One alternative is to leverage synthetic datasets, *e.g.*, GTA5 [32] and SYNTHIA [34], that contain realistic images and corresponding pixel-level labels. The annotation cost is much cheaper than the manual labeling, but CNNs trained with synthetic datasets do not work well on real images, due to the domain discrepancy between synthetic and real images.

To reduce the domain discrepancy, several methods [43, 15, 23, 4, 16, 44] have exploited an unsupervised domain adaptation approach. It transfers knowledge learned from a source domain (*e.g.*, a synthetic dataset) to a target one (*e.g.*, a real dataset), with labels for the source domain alone. Many unsupervised domain adaptation methods leverage an adversarial training scheme [12] that aligns distributions of source and target domains by fooling a domain classifier [15, 30, 5, 23, 49, 16, 3, 36, 4, 55, 10, 27, 44, 28, 42, 48]. However, they typically focus on reducing the domain discrepancy globally, and fail to keep pixel-level semantics [50]. For example, regions corresponding to a car class in a source image might align with those for a bus class in a target image. Self-training methods [58, 59, 51, 50, 22] enable a class-aware alignment. They generate pseudo labels for target images iteratively in a parametric approach, typically using CNNs trained with a source dataset, and then retrain a segmentation model on both source and target samples with the pseudo labels. This aligns cross-domain features in a class-level, improving performance of the model on target images progressively. The pseudo labels obtained using a parametric approach have the following drawbacks: First, they are very sparse, since low confident predictions are discarded to obtain reliable labels. Second, estimating pseudo labels is also computationally demanding, making them not to be updated frequently during training. These problems cause the segmentation model to overfit to pseudo la-

bels, resulting in a large bias and a variance of predictions. In the following, we will call the labels estimated using a parametric approach as *static pseudo labels*.

We present a novel contrastive learning framework using cross-domain pixel-prototype correspondences for unsupervised domain adaptive semantic segmentation (UDASS). It aligns pixel-level features of each object class in target images, obtained by pseudo labels, with prototypes of corresponding class in a source domain, computed by ground truth, while setting them apart for different classes (Fig. 1). The alignment and separation process is also performed in a reverse direction, with pixel-level features of source images and prototypes of a target domain. The cross-domain matching encourages domain-invariant feature representations, and the bidirectional pixel-prototype correspondences provide compact and discriminative representations. We also present a nonparametric approach to generating *dynamic pseudo labels* using pixel-prototype correspondences. Specifically, we calibrate prototypes of individual object classes in a source domain, while considering the domain discrepancy in target images, and establish correspondences for each prototype with individual pixel-level features in target images. We then transfer ground-truth labels of prototypes to corresponding pixels in target images. In contrast to the parametric approach in current self-training methods, our nonparametric approach provides denser pseudo labels, and generates the labels dynamically, whenever source images are changed during training. This helps to obtain more accurate pseudo labels, and prevents the overfitting problem. Experimental results on standard benchmarks including GTA5-to-Cityscapes [32, 6] and SYNTHIA-to-Cityscapes [34, 6] demonstrate that our contrastive learning framework provides domain-invariant and discriminative features for UDASS. The main contributions can be summarized as follows:

- We introduce a novel contrastive learning framework using bi-directional pixel-prototype correspondences to learn domain-invariant and discriminative feature representations for UDASS.
- We propose a nonparametric approach to generating dynamic pseudo labels. We also present a calibration method to reduce domain biases for pixel-prototype correspondences between target and source domains.
- We set a new state of the art on standard benchmarks for UDASS, and demonstrate the effectiveness of our contrast learning framework.

## 2 Related work

**UDASS.** UDASS leverages knowledge learned from a label-rich source domain to predict semantic labels of a scene in a target domain, where ground-truth annotations are not available. Synthetic images (*e.g.*, GTA5 [32] and SYNTHIA [34]) are widely used as source samples, as pixel-level labels can be generated automatically using computer graphics engines. The key factor for UDASS is hence to learn domain-invariant features to reduce the discrepancy between source and target domains. To this end, many UDASS methods adopt an adversarial learning framework [12] to fool a domain discriminator. They can gen-

erally be categorized into image-level and feature-level alignment methods. Motivated by image translation techniques [18, 57], image-level alignment methods [15, 5, 23, 30, 49] transfer the styles (*e.g.*, texture and illumination) of target images to the source, so that segmentation models can accommodate both domains. Feature-level alignment methods [16, 3, 36, 4, 55, 10, 27, 44, 28, 42, 48] align the feature distributions of source and target images explicitly. These adversarial approaches, however, align source and target distributions globally. Namely, they perform a class-agnostic alignment, and ignore positional information of a scene. This suggests that the adversarial approaches fail to transfer pixel-level semantics, related to the structural information of a scene, from source to target domains.

UDASS methods based on self-training [58, 59, 22] have recently been introduced. The self-training approach first segments target images using a model trained on a source dataset, and obtains pseudo labels if the confidence of semantic labels predicted by the model exceeds a pre-defined threshold. It then retraining the model iteratively with both ground-truth and pseudo labels of source and target datasets, respectively. The representative work of [58] proposes to use different thresholds for individual object categories to consider a class imbalance problem. In [59], soft pseudo labels have been introduced, together with a confidence regularization technique that helps transfer discriminative feature representations from source to target domains. The self-training approaches [58, 59, 22, 52] are, however, likely to overfit to pseudo labels. The reasons are as follows: (1) Pseudo labels are fixed for a few epochs during training, due to computational overheads, which accumulates error from incorrect pseudo labels; (2) Pseudo labels are very sparse, as high confident predictions are chosen only as the labels. Our method alleviates these limitations by generating denser pseudo labels dynamically in a nonparametric way using pixel-prototype correspondences. Most similar to ours is PLCA [20] using pixel-wise matches. It adopts a contrastive learning scheme to reduce the distances between source and target features directly at a pixel-level. The pixel-level domain alignment, however, does not consider contextual information, and fails to obtain compact representations between corresponding object categories in source and target domains. Our method instead uses bidirectional pixel-prototype correspondences for contrastive learning, which encourages intra-class compactness and inter-class separability across domains.

**Prototypical learning.** The seminal work of [40] introduces prototypical networks that extract prototype representations for individual object categories. The prototypical features have proven useful in the limited-data regime for the task of, *e.g.*, few-/zero-shot classification. PL [9] extends the idea of prototypical learning for few-shot semantic segmentation in such a way that class prototypes obtained from a support set are matched to pixel-level features in a query image. PANet [46] presents a bidirectional framework exploiting correspondences between prototypical features for a support set and pixel-level ones for query images, and vice versa, for few-shot semantic segmentation. Similar to these methods, we exploit prototypical features for semantic segmentation. Differently, we leverage them within a framework of contrastive learning for UDASS. We use

pixel-prototype correspondences to obtain domain-invariant and discriminative feature representations. We also leverage the correspondences to obtain dynamic pseudo labels, which alleviates the limitations of current self-training methods using static pseudo labels.

**Contrastive learning.** Contrastive learning [2, 13] is a de facto approach to learning generic feature representations in a self-supervised way. The basic idea is to encourage positive pairs with the same label to be close, while negative ones with different labels to be distant. In order to set positive and negative pairs without ground-truth labels, contrastive learning augments a single input image, *e.g.*, using random cropping and color jittering. It then considers the original image and the augmented one as a positive pair, while setting the pairs composed of the original and other images as negative ones. Similar to ours, CANet [19] adopts contrastive learning for unsupervised domain adaptive classification. It computes the domain discrepancies using image-level features, and then performs a class-wise alignment using target labels obtained by a clustering method. Differently, our method leverages contrastive learning using correspondences between pixels and prototypes across domains. Optimizing bidirectional correspondences jointly in our method also enables aggregating features for the same object category, regardless of domains.

**Nonparametric label transfer.** Label transfer has been widely used in object localization [29], scene segmentation [35, 25, 31, 39], automatic image annotation [45], and image translation [38]. Label transfer methods first search visually similar images or patches in large datasets for given queries, and then transfer labels of retrieved samples to the queries. Similar to our approach, the work of [8] adopts a nonparametric label transfer method for scene parsing under different domains (*e.g.*, weather or illumination). Specifically, it extracts features from query images with pre-trained networks, finds the best matching images using SIFT flow [24], and transfers labels of the images to the queries via a probabilistic MRF model, suggesting that this approach requires source images and ground-truth labels at both training and test time. Our method, on the other hand, uses source images and corresponding ground-truth labels only at training time. Namely, we leverage non-parametric label transfer to train a parametric segmentation model.

## 3 Approach

### 3.1 Overview

We introduce a cross-domain contrastive learning framework for UDASS using pixel-prototype correspondences (Fig. 2). It first extracts feature maps from source and target images, respectively, using a siamese network. We obtain prototypes of source and target domains using ground-truth labels of source images and pseudo labels of target ones, respectively. Our method then establishes correspondences between the prototypes and pixel-level features across domains, and leverages them to learn domain-invariant and discriminative representations via

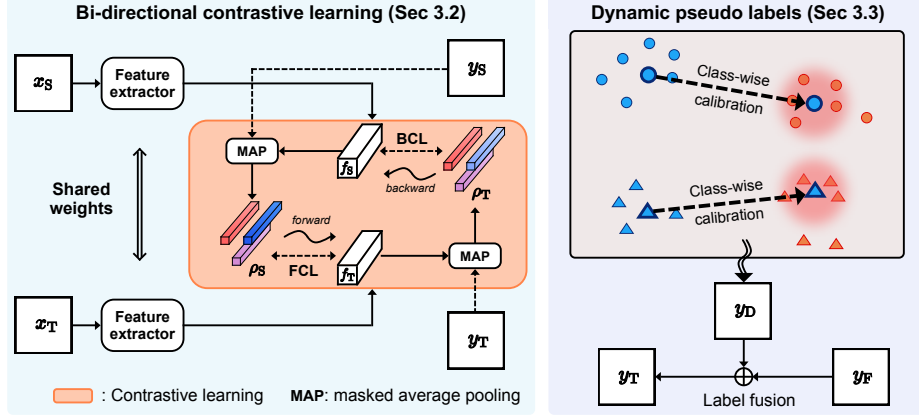


Fig. 2: An overview of our framework. (Left) Bi-directional contrastive learning: We first extract feature maps,  $f_S$  and  $f_T$ , from source and target images,  $x_S$  and  $x_T$ , respectively. We then obtain prototypes in a source domain,  $\rho_S$  using ground-truth labels of source images  $y_S$ . Prototypes in a target domain  $\rho_T$  are similarly computed but with dynamic pseudo labels of target images  $y_T$ . Bidirectional contrastive terms, FCL and BCL, exploit pixel-prototype correspondences across domains to learn domain-invariant and discriminative features for UDASS. (Right) Hybrid pseudo labels: We generate dynamic pseudo labels  $y_D$  using pixel-prototype correspondences across domains, while calibrating the prototypes to reduce domain discrepancies. We then combine them with static ones  $y_F$  using a parametric approach to obtain hybrid pseudo labels  $y_T$ .

contrastive learning. To this end, we introduce a bi-directional contrastive loss that consists of a forward contrastive term (FCL) and a backward contrastive term (BCL). FCL matches individual pixel-level features of a target image with prototypes of a source domain, and enforces pixel-prototype pairs with the same class labels to be aligned closely than other ones. BCL performs the alignment process in a reverse direction, with pixel-level features of a source image and prototypes of a target domain, encouraging our model to provide discriminative and compact features. In order to establish training pairs for computing the bi-directional contrastive loss, we require pseudo labels of target images. To this end, we use dynamic pseudo labels obtained by a nonparametric label transfer, addressing the drawbacks of static pseudo labels. Specifically, given a pair of source-target images, we establish correspondences between prototypes of a source domain and pixel-level features of a target image, while calibrating the prototypes progressively during training to compensate domain discrepancies. We then set the pseudo labels of pixel-level features to the class labels of the corresponding prototypes in a source domain. Unlike static pseudo labels estimated by a parametric approach [58, 22], our approach can generate novel pseudo labels of target images dynamically, whenever a pair of source-target images are changed, during training. We estimate hybrid pseudo labels by com-

binning dynamic and static labels, and use them for the bi-directional contrastive learning.

### 3.2 Bi-directional contrastive learning

Given a pair of source and target images, our goal is to aggregate pixel-level features for the same object class, regardless of domains, to learn domain-invariant and discriminative feature representations. To this end, we formulate UDASS as bi-directional pixel-prototype contrastive learning. Let us denote by  $\mathcal{C}$  the set of object classes. We obtain prototypes of source and target domains for the class  $c \in \mathcal{C}$ ,  $\rho_S(c)$  and  $\rho_T(c)$ , using masked average pooling (MAP) as follows:

$$\rho_S(c) = \frac{\sum_p f_S(p) y_S(p, c)}{\sum_p y_S(p, c)}, \rho_T(c) = \frac{\sum_p f_T(p) y_T(p, c)}{\sum_p y_T(p, c)}, \quad (1)$$

where we denote by  $f_S(p)$  and  $f_T(p)$  pixel-level features of source and target images, respectively, at position  $p$ .  $y_S(p, c)$  and  $y_T(p, c)$  are one-hot labels, *i.e.*, 1 if the class label at position  $p$  correspond to  $c$  and 0 otherwise. Note that we use ground-truth labels of source images  $y_S$  and hybrid pseudo labels of target ones  $y_T$  to set the labels,  $y_S$  and  $y_T$ , respectively. Using the prototypes of source and target domains, we perform cross-domain contrastive learning in a bi-directional way. We leverage a bi-directional constative loss that consists of FCL and BCL. FCL exploits prototypes of a source domain and pixel-level features of a target image. To be specific, given pixel-level features of a target image, we select the prototypes of a source domain having the same class labels as the features, and set them as positive pairs, while other prototypes are used to set negative ones. FCL maximizes the similarities between positive pairs as follows:

$$\mathcal{L}_{FC} = - \sum_c \sum_p y_T(p, c) \log \frac{\exp(s(f_T(p), \rho_S(c))/\tau)}{\sum_c \exp(s(f_T(p), \rho_S(c))/\tau)}, \quad (2)$$

where  $\tau$  is a temperature parameter, and  $s(\cdot, \cdot)$  computes cosine similarity. Similarly, BCL exploits prototypes of a target domain and pixel-level features of a source image. It encourages positive pairs sharing the same labels to pull each other, while making others set apart as follows:

$$\mathcal{L}_{BC} = - \sum_c \sum_p y_S(p, c) \log \frac{\exp(s(f_S(p), \rho_T(c))/\tau)}{\sum_c \exp(s(f_S(p), \rho_T(c))/\tau)}. \quad (3)$$

In summary, using the bidirectional contrastive loss, pixel-level features for the same object class are embedded closely, regardless of domains, while those for different classes are distinguished from each other. That is, by jointly optimizing FCL and BCL, we can minimize intra-class variations and maximize inter-class variations of pixel-level features progressively during training. In contrast to current UDASS methods [54, 58, 59, 51, 50, 22] that do not consider such variations for domain adaptation, our approach provides more discriminative and compact features. This in turn allows to perform more accurate class-wise alignments across domains, and enables our model to generalize better on a target domain.



### 3.3 Dynamic pseudo labels

Current self-training methods [54, 58, 59, 51, 50, 22] employ a parametric model trained with ground-truth labels of source images to obtain static pseudo labels of target images. Specifically, using the parametric segmentation model, confidence scores for individual object classes are computed for each pixel-level feature from entire target images. The pixel-level features with high confidence scores are chosen, and corresponding object classes are used as pseudo labels. Although exploiting static pseudo labels of target images enables performing class-aware UDASS, they have the following drawbacks: First, computing the pixel-level confidence scores for all target images to obtain the pseudo labels is computationally demanding. Current self-training methods perform this process for a few iterations (*e.g.*, 10000) during training, and update the pseudo labels of target images very occasionally. The error from incorrect pseudo labels might hence be accumulated. Second, current self-training methods choose highly confident pixel-level features only for static pseudo labels, and thus they are very sparse. These problems cause a model to overfit to the static pseudo labels, and induce suboptimal class-wise alignments between domains. To overcome the limitations, we introduce a novel approach to generating dynamic pseudo labels. It leverages a nonparametric label transfer technique using pixel-prototype correspondences between source and target images. That is, we estimate pseudo labels using pairs of source and target images. This suggests that our approach generates pseudo labels of target images dynamically, whenever source images are changed during training. In other words, the pseudo labels for the same target image could be different, depending on which source images are used to establish pixel-prototype correspondences w.r.t the target one (Fig. 3).

Concretely, given a pair of source and target images, we establish correspondences between prototypes of a source image and pixel-level features of a target one. To obtain reliable correspondences, we alleviate domain biases between source and target domains. We could estimate the degree of domain biases by calculating average class-wise features for each domain using all source and target images, followed by computing differences between the average features, which however requires lots of computational overheads. We instead leverage prototypes of source and target images. We first update prototypes of source and target domains progressively during training using an exponential moving average with a momentum parameter of  $\lambda$  as follows:

$$\mu_S(c) \leftarrow \lambda \mu_S(c) + (1 - \lambda) \rho_S(c), \quad (4)$$

$$\mu_T(c) \leftarrow \lambda \mu_T(c) + (1 - \lambda) \rho_T(c), \quad (5)$$

where we denote by  $\mu_S(c)$  and  $\mu_T(c)$  updated prototypes of source and target domains, respectively, for the class  $c$ . We then estimate class-wise domain biases:

$$\xi(c) = \mu_T(c) - \mu_S(c), \quad (6)$$

and obtain calibrated prototypes for each object class in a source domain as follows:

$$\rho_{S \rightarrow T}(c) = \rho_S(c) + \xi(c). \quad (7)$$






Fig. 3: Visual comparison of static, dynamic, and hybrid pseudo labels for a target image. In contrast to the static label (the first column), estimated using a parametric segmentation model, dynamic labels are obtained by a nonparametric label transfer between source and target images. This suggests that they are denser and cheap to update. We can also obtain different dynamic labels (the second column), according to source images (the fourth column). We combine both labels to get hybrid pseudo labels (the third column), and use them to augment the number of positive and negative pairs for contrastive learning.

Using the calibrated prototypes, we can establish more correct correspondences across domains. We consider the correspondences are correct, if similarity scores between the pixel-prototype matches are larger than a pre-defined threshold, and set dynamic pseudo labels of target images to corresponding object categories of the prototypes, as follows:

$$y_D(p, c) = \begin{cases} 1, & \text{if } s(f_T(p)), \rho_{S \rightarrow T}(c)) > \mathcal{T} \text{ and } c = c' \\ 0, & \text{otherwise} \end{cases}, \quad (8)$$

where  $y_D(p, c)$  is a dynamic pseudo label for the class  $c$  at position  $p$ ,  $\mathcal{T}$  is a pre-defined threshold, and

$$c' = \underset{c}{\operatorname{argmax}}(s(f_T(p)), \rho_{S \rightarrow T}(c)). \quad (9)$$

**Hybrid pseudo labels.** We can obtain diverse pseudo labels even for the same target image every iteration, and the dynamic labels are much denser than static ones. Static pseudo labels, on the other hand, are sparse but reliable. In order to take advantage of both, we combine them and obtain hybrid pseudo labels  $y_T$  as follows:

$$y_T(p, c) = \begin{cases} y_D(p, c), & \text{if } y_D(p, c) = 1 \\ y_F(p, c), & \text{if } y_D(p, c') = 0 \text{ for } c' \in \mathcal{C}, \text{ and } y_F(p, c) = 1 \\ 0, & \text{otherwise} \end{cases}, \quad (10)$$

where  $y_F(p, c)$  is a static pseudo label for the class  $c$  at position  $p$ .

### 3.4 Training loss

Following the previous works [58, 59, 51, 50, 22], we exploit segmentation and entropy terms using ground-truth and pseudo labels of source and target images, respectively. The former encourages our model to provide accurate pixel-wise predictions, and the latter minimizes the entropy of the predictions. We define a loss for training a baseline model as follows:

$$\mathcal{L}_{\text{base}} = \lambda_{\text{seg}}^{\text{S}} \mathcal{L}_{\text{seg}}^{\text{S}} + \lambda_{\text{seg}}^{\text{T}} \mathcal{L}_{\text{seg}}^{\text{T}} + \lambda_{\text{ent}}^{\text{S}} \mathcal{L}_{\text{ent}}^{\text{S}} + \lambda_{\text{ent}}^{\text{T}} \mathcal{L}_{\text{ent}}^{\text{T}}, \quad (11)$$

where  $\mathcal{L}_{\text{seg}}^{\text{S}}$  and  $\mathcal{L}_{\text{seg}}^{\text{T}}$  are segmentation losses for source and target domains, respectively.  $\mathcal{L}_{\text{ent}}^{\text{S}}$  and  $\mathcal{L}_{\text{ent}}^{\text{T}}$  are entropy terms for source and target domains, respectively.  $\lambda_{\text{seg}}^{\text{S}}$ ,  $\lambda_{\text{seg}}^{\text{T}}$ ,  $\lambda_{\text{ent}}^{\text{S}}$ , and  $\lambda_{\text{ent}}^{\text{T}}$  are balance parameters for each term. For the baseline, we obtain static pseudo labels using the method of [58]. As our final model, we additionally use a bi-directional contrastive loss to learn domain-invariant and discriminative representations as follows:

$$\mathcal{L} = \mathcal{L}_{\text{base}} + \lambda_{\text{FC}} \mathcal{L}_{\text{FC}} + \lambda_{\text{BC}} \mathcal{L}_{\text{BC}}, \quad (12)$$

where  $\lambda_{\text{FC}}$  and  $\lambda_{\text{BC}}$  are weighting factors for forward and backward contrastive terms, respectively.

## 4 Experiments

### 4.1 Implementation details

**Dataset and evaluation metric.** We evaluate our framework on two standard benchmarks (GTA5 [32]  $\rightarrow$  Cityscapes [6], and SYNTHIA [34]  $\rightarrow$  Cityscapes [6]). GTA5 and SYNTHIA provide 24,996 and 9400 images, respectively. Cityscapes consists of 2975, 500, and 1525 images for training, validation, and testing, respectively. Following the standard protocol in [15, 43, 58, 59], we report the mean intersection over Union (mIoU) on 19 classes for GTA5  $\rightarrow$  Cityscapes and 13 (or 16) classes for Synthia  $\rightarrow$  Cityscapes.

**Training.** We adopt the DeepLab-V2 [1] architecture with ResNet-101 [14] as a backbone network pretrained for ImageNet classification [7]. We first train DeepLab-V2 with a source dataset, and use it as an initial segmentation model for UDASS. We train the model for 100k iterations with a batch size of 4, using stochastic gradient descent (SGD) [21] of a momentum of 0.9 and weight decay of  $5 \times 10^{-4}$ . We use a poly learning rate scheduling with an initial learning rate of  $7.5 \times 10^{-5}$ . We update static pseudo labels  $y_{\text{F}}$  every 10k iterations. We resize a shorter side of images to 850, and crop them into a patch of size  $730 \times 730$ . For data augmentation, we use horizontal flipping and random scaling with the factor of [0.8, 1.2]. We use a weighted sampling strategy to select source images containing objects that rarely appear in a source domain, mitigating low co-occurrence rates for the rare object categories. Following [52, 11], we additionally apply a self-distillation technique to our final model. Detailed descriptions for the weighted sampling and hyperparameter settings are available in the supplement.

GTA5 $\rightarrow$ Cityscapes																							
Split	Methods	Type	road	side.	build.	wall	fence	pole	light	sign	veg.	terrian	sky	person	rider	car	truck	bus	train	motor	bike	mIoU	
Validation	Source-only	-	45.4	16.5	66.4	14.4	21.6	25.1	36.3	17.2	80.1	16.3	69.1	61.4	24.9	68.6	28.4	4.7	4.4	40.8	27.5	35.2	
	AdaptSeg [43]	AT	86.5	36.0	79.9	23.4	23.3	23.9	35.2	14.8	83.4	33.3	75.6	58.5	27.6	73.7	32.5	35.4	3.9	30.1	28.1	42.4	
	CBST [58]	ST	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	
	CRST [59]	ST	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	26.9	26.0	42.3	47.1	
	PLCA [20]	-	84.0	30.4	82.4	35.3	24.8	32.2	36.8	24.5	85.5	37.2	78.6	66.9	32.8	85.5	40.4	48.0	8.8	29.8	41.8	47.7	
	CAG-UDA [53]	ST	90.4	51.6	83.8	34.2	27.8	38.4	25.3	48.4	85.4	38.2	78.1	58.6	34.6	84.7	21.9	42.7	41.1	29.3	37.2	50.2	
	FDA [51]	ST	92.5	53.5	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	
	TPLD [37]	ST	94.2	60.5	82.8	36.6	16.6	39.3	29.0	25.5	85.6	44.9	84.4	60.6	27.4	84.1	37.0	47.0	31.2	36.1	46.4	51.2	
	CorDA [47]	ST	<b>94.7</b>	<b>63.1</b>	87.6	30.7	40.6	40.2	47.8	51.6	87.6	<b>47.0</b>	89.7	66.7	35.9	90.2	<b>48.9</b>	57.5	0.0	39.8	56.0	56.6	
	ProDA [52]	ST	87.1	55.1	78.1	<b>45.6</b>	<b>43.8</b>	<b>44.6</b>	52.5	<b>53.4</b>	89.1	44.7	82.1	70.1	<b>39.1</b>	88.4	43.8	<b>59.1</b>	1.0	<b>48.7</b>	54.4	56.5	
	Ours	ST	93.5	60.2	<b>88.1</b>	31.1	37.0	41.9	<b>54.7</b>	37.8	<b>89.9</b>	45.5	<b>89.9</b>	<b>72.7</b>	38.2	<b>90.7</b>	34.3	53.2	4.4	47.2	<b>58.5</b>	<b>57.1</b>	
Test	AdaptSeg [43]	AT	88.5	40.4	81.0	26.3	20.6	25.6	36.0	12.9	84.8	45.5	87.2	63.7	35.8	76.4	27.7	28.0	2.9	33.0	26.1	44.3	
	CBST [58]	ST	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	
	CRST [59]	ST	93.5	57.6	84.6	<b>39.3</b>	24.1	25.2	35.0	17.3	85.0	40.6	86.5	58.7	28.7	85.8	<b>49.0</b>	56.4	5.4	31.9	43.2	49.9	
	FDA-MBT [51]	ST	93.4	55.8	83.6	25.4	23.1	33.2	39.0	36.9	84.0	47.2	88.8	66.3	40.6	87.4	26.9	49.6	12.8	35.2	42.8	51.2	
	CorDA [47]	ST	<b>94.2</b>	<b>62.9</b>	88.1	30.2	41.2	40.1	49.1	49.9	89.1	49.1	90.1	69.1	28.9	86.2	46.2	59.5	1.2	35.2	52.3	57.5	
	ProDA [52]	ST	88.1	57.1	81.2	46.1	<b>45.2</b>	<b>41.5</b>	<b>55.1</b>	<b>56.2</b>	86.1	45.1	78.1	73.2	40.1	88.8	48.7	<b>60.1</b>	1.1	50.3	53.1	57.6	
	Ours	ST	93.8	59.7	<b>90.1</b>	38.0	33.4	39.9	45.3	30.5	<b>92.2</b>	<b>58.2</b>	<b>94.8</b>	<b>81.9</b>	<b>47.9</b>	<b>93.2</b>	40.1	53.1	13.1	<b>51.2</b>	<b>58.2</b>	<b>58.5</b>	

Table 1: Quantitative comparison with state-of-the-art methods on GTA5  $\rightarrow$  Cityscapes in terms of mIoU. AT: methods based on adversarial training; ST: methods based on self-training. †: a method using a different network architecture.

SYNTHIA $\rightarrow$ Cityscapes																			
Methods	Type	road	side.	build.	wall*	fence*	pole*	light	sign	veg.	sky	person	rider	car	bus	motor	bike	mIoU	mIoU*
Source-only	AT	53.4	23.4	73.0	5.5	0.0	25.7	6.6	7.0	77.9	55.3	52.9	21.0	60.9	6.6	21.8	33.7	32.5	37.6
AdaptSeg [43]	AT	84.3	42.7	77.5	-	-	-	4.7	7.0	77.9	82.5	54.3	21.0	72.3	32.2	18.9	32.3	-	46.7
CBST [58]	ST	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	38.9	42.6
CRST [59]	ST	67.7	32.2	73.9	10.7	1.6	37.4	22.2	31.2	80.8	80.5	60.8	29.1	82.8	25.0	19.4	45.3	43.8	50.1
CAG-UDA [53]	ST	84.7	40.8	81.7	7.8	0.0	35.1	13.3	22.7	84.5	77.6	64.2	27.8	80.9	19.7	22.7	48.3	44.5	51.5
FDA [51]	ST	79.3	35.0	73.2	-	-	-	19.9	24.0	61.7	82.6	61.4	31.1	83.9	40.8	38.4	51.1	-	52.5
PLCA [20]	-	82.6	29.0	81.0	11.2	0.2	33.6	24.9	18.3	82.8	82.3	62.1	26.5	85.6	48.9	26.8	52.2	46.8	54.0
TPLD [37]	ST	80.9	44.3	82.2	19.9	0.3	40.6	20.5	30.1	77.2	80.9	60.6	25.5	84.8	41.1	24.7	43.7	47.3	53.5
CorDA [47]	ST	<b>93.3</b>	<b>61.6</b>	85.3	19.6	5.1	37.8	36.6	<b>42.8</b>	84.9	90.4	69.7	41.8	85.6	38.4	32.6	53.9	55.0	62.8
ProDA [52]	ST	87.3	45.1	84.2	<b>36.5</b>	0.0	<b>43.3</b>	<b>54.7</b>	36.0	88.3	83.1	71.5	24.4	88.4	<b>50.1</b>	40.1	45.6	55.1	61.3
Ours	ST	83.8	42.2	<b>85.3</b>	16.4	<b>5.7</b>	43.1	48.3	30.2	<b>89.3</b>	<b>92.1</b>	<b>68.2</b>	<b>43.1</b>	<b>89.7</b>	47.2	<b>42.2</b>	<b>54.2</b>	<b>55.6</b>	<b>62.9</b>

Table 2: Quantitative comparison with state-of-the-art methods on SYNTHIA  $\rightarrow$  Cityscapes results in terms of mIoU. We report the results for 13 classes (mIoU\*) and 16 classes (mIoU). AT: methods based on adversarial training; ST: methods based on self-training.

## 4.2 Results

**Quantitative results.** We compare our method with the state-of-the-art methods on GTA5  $\rightarrow$  Cityscapes and SYNTHIA  $\rightarrow$  Cityscapes in Tables 1 and 2, respectively. Note that all methods in the tables are based on the DeepLab-V2 [1] architecture with ResNet-101, except for CAG-UDA [53]. For a fair comparison, we report the results of ProDA [52] using the same network architecture as other methods, reproduced using an official source code. CBST [58] uses a self-training-based method to perform a class-aware alignment. This method is similar to our baseline, but it uses a limited number of pseudo labels, being outperformed by our approach on both benchmarks. PLCA [20] uses a pixel-wise association method to align source and target domains in a pixel-level. This method, however, fails to obtain compact feature representations, and it is hence outperformed by our approach on both benchmarks. CorDA [47] uses depth maps of source and target domains to transfer the knowledge of a source domain to




Fig. 4: Qualitative comparisons on GTA5  $\rightarrow$  Cityscapes. Our model gives better results than the baseline. (Best viewed in color)




Fig. 5: t-SNE visualization of a baseline (a) and our model (b). (Best viewed in color)

Table 3: Quantitative results for variants of our model. We report mIoU scores for 19 and 16 classes on GTA5  $\rightarrow$  Cityscapes and SYNTHIA  $\rightarrow$  Cityscapes, respectively.

$\mathcal{L}_{\text{base}}$	$\mathcal{L}_{FC}$	$\mathcal{L}_{BC}$	$+y_D$ (w/o cal.)	$+y_D$ (w/ cal.)	Source dataset	
					GTA5	SYNTHIA
✓					49.5	45.1
✓	✓				51.2	48.8
✓	✓	✓			53.5	51.3
✓	✓	✓	✓		55.3	53.5
✓	✓	✓		✓	57.1	55.6

a target one. Our method outperforms CorDA [47] on both benchmarks even without using the depth information, indicating that our contrastive learning framework effectively transfers the knowledge across domains using pseudo labels. ProDA [52] focuses on removing false-positives of pseudo labels [58] and uses sparse labels. Different from ProDA [52], we are interested in generating additional labels based upon the ones obtained by the approach of [58]. That is, our method focuses on obtaining more true-positives and generating denser labels using pixel-prototype correspondences. Other than ProDA [52], we additionally use the bi-directional contrastive loss to minimize intra-class variations and maximize inter-class variations of pixel-level features. We achieve mIoU gains of 0.6% and 1.6% for GTA5  $\rightarrow$  Cityscapes and SYNTHIA  $\rightarrow$  Cityscapes, respectively, compared to ProDA [52]. The results imply that our method effectively learns domain-invariant and discriminative representations with denser pseudo labels, improving the mIoU performance of semantic segmentation. Additional comparisons with ProDA [52] are available in the supplement. We also report mIoU scores for the test split of Cityscapes, obtained from an official evaluation server, which has been ignored by most previous works. We use official source codes provided by the authors to obtain the results of state-of-the-art methods. We achieve non-trivial mIoU gains over CorDA [47] and ProDA [52] for the test split of Cityscapes, demonstrating that ours can generalize better than them. Considering the performance gains of recent UDASS methods, the results are significant. For example, FDA [51] achieves a mIoU gain of 0.3% over

CAG\_UDA [53], and TPLD [37] gets the gain of 0.7% over FDA [51]. CorDA [47] and ProDA [52] provide large mIoU gains compared to other methods, but the improvements mainly come from exploiting additional depth maps and applying post-processing method, respectively.

**Qualitative results.** We show in Fig. 4 segmentation results on the GTA5  $\rightarrow$  Cityscapes task. Compared to the baseline model, our model provides more accurate segmentation results (*e.g.*, the bus in the first row, and the road and the rider in the second row). We show in Fig. 5 the t-SNE plot of feature representations of our model and the baseline. We visualize features of source and target images for each method by red and blue circles, respectively. The results show that our method successfully aligns the features for the same object category and separates them for different ones. That is, it minimizes intra-class variations, and maximizes inter-class variations, regardless of domains.

**Ablation study.** We present in Table 3 an ablation analysis for each component of our framework on GTA5  $\rightarrow$  Cityscapes and SYNTHIA  $\rightarrow$  Cityscapes. We show mIoU scores for variants of our model on the validation split of Cityscapes. As a baseline in the first row, we use static pseudo labels, obtained by the method of [58], to perform a class-aware alignment between source and target domains. We can see from the second row that FCL gives better mIoU scores, demonstrating the effectiveness of our approach to aligning prototypes and pixel-level features across domains. From the first and third rows, we can clearly see that jointly optimizing two contrastive terms is effective to UDASS. The fourth row demonstrates that leveraging additional dynamic pseudo labels provides better results than exploiting the static ones alone in terms of the mIoU score, even without the calibration (w/o cal.). We can observe from the fifth row that the calibration (w/ cal.) reduces domain discrepancies, and further improves the performance significantly.

**Comparison of pseudo labels.** We measure the densities of various pseudo labels and corresponding label accuracies, and report the results in Table 4. We can see that the densities of dynamic pseudo labels are slightly higher than that of a static one, even without calibrating domain biases, while maintaining the label accuracies. Using pixel-prototype correspondences between target and source domains leads to obtaining denser labels than [58]. The calibration process largely densifies dynamic pseudo labels. We can establish more correct correspondences between source and target domains by using the calibration process. The approach of [58] neglects the biases between source and target domains. Different from [58], ours compensate for the class-wise domain biases and generate more accurate and denser labels than [58]. Hybrid pseudo labels that combine static and dynamic ones provide the best result in terms of the label density and accuracy. When obtaining hybrid pseudo labels, we can reduce the number of incorrect static labels [58] by comparing them with dynamic ones. We show in Fig. 6 examples of dynamic pseudo labels obtained with and without the class-wise calibration. The results show that calibrating class-wise domain biases for source prototypes leads to establishing more correct pixel-prototype correspondences, providing denser and more accurate pseudo labels.



(a) w/o cal. (b) w/ cal. (c) GT labels.

Fig. 6: Visualization of dynamic pseudo labels. (a-b) Pseudo labels obtained without and with calibrating prototypes of a source domain; (c) Target labels.

Table 4: Quantitative results for various pseudo labels of a target domain. We report the densities of static, dynamic, and hybrid pseudo labels and corresponding label accuracies.

Pseudo labels	Density(%)	Accuracy(%)
Static [58]	20.1	98.5
Dyn. (w/o cal.)	22.2	98.6
Dyn. (w/ cal.)	34.3	98.6
Hybrid	42.3	98.8

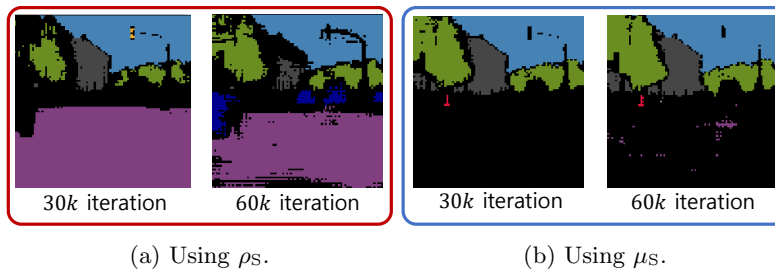
(a) Using  $\rho_S$ .(b) Using  $\mu_S$ .

Fig. 7: Pseudo labels at 30k and 60k iterations using  $\rho_S$  (a) and  $\mu_S$  (b), respectively.

In Fig. 7, we compare generated pseudo labels using instance-wise prototypes  $\rho_S$  and momentum-based ones  $\mu_S$ . We can see that using instance-wise prototypes  $\rho_S$  provides more diverse pseudo labels. They are more various than the other ones  $\mu_S$  slowly moving with momentum, and lead our model to establish diverse pixel-prototype correspondences.

## 5 Conclusion

We have introduced a novel contrastive learning framework for UDASS. Our key idea is to use cross-domain pixel-prototype correspondences to learn domain-invariant and discriminative representations. We have introduced a bi-directional contrastive loss to align the features for the same object category and separate them for different ones. We have also introduced an approach to generating pseudo labels dynamically in a nonparametric way using pixel-prototype correspondences, while compensating class-wise domain biases between source and target domains. Experimental results show the effectiveness of our framework, setting a new state of the art on standard benchmarks.

**Acknowledgements.** This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.RS-2022-00143524, Development of Fundamental Technology and Integrated Solution for Next-Generation Automatic Artificial Intelligence System, and No.2022-0-00124, Development of Artificial Intelligence Technology for Self-Improving Competency-Aware Learning Capabilities), and the Yonsei Signature Research Cluster Program of 2022 (2022-22-0002).

## References

1. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected CRFs. ICLR (2015) [1](#), [10](#), [11](#)
2. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020) [5](#)
3. Chen, Y.H., Chen, W.Y., Chen, Y.T., Tsai, B.C., Frank Wang, Y.C., Sun, M.: No more discrimination: Cross city adaptation of road scene segmenters. In: ICCV (2017) [2](#), [4](#)
4. Chen, Y., Li, W., Van Gool, L.: ROAD: Reality oriented adaptation for semantic segmentation of urban scenes. In: CVPR (2018) [2](#), [4](#)
5. Chen, Y.C., Lin, Y.Y., Yang, M.H., Huang, J.B.: CrDoCo: Pixel-level domain transfer with cross-domain consistency. In: CVPR (2019) [2](#), [4](#)
6. Cordts, M., Omran, M., Ramos, S., Scharwächter, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset. In: CVPR Workshop (2015) [2](#), [3](#), [10](#)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: CVPR (2009) [10](#)
8. Di, S., Zhang, H., Li, C.G., Mei, X., Prokhorov, D., Ling, H.: Cross-domain traffic scene understanding: A dense correspondence-based transfer learning approach. IEEE Trans. ITS **19**(3) (2017) [5](#)
9. Dong, N., Xing, E.P.: Few-shot semantic segmentation with prototype learning. In: BMVC (2018) [4](#)
10. 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) [2](#), [4](#)
11. Fang, Z., Wang, J., Wang, L., Zhang, L., Yang, Y., Liu, Z.: SEED: Self-supervised distillation for visual representation. ICLR (2021) [10](#)
12. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. JMLR (2016) [2](#), [3](#)
13. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020) [5](#)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) [10](#)
15. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: CyCADA: Cycle-consistent adversarial domain adaptation. In: ICML (2018) [2](#), [4](#), [10](#)
16. Hoffman, J., Wang, D., Yu, F., Darrell, T.: FCNs in the wild: Pixel-level adversarial and constraint-based adaptation. arXiv preprint arXiv:1612.02649 [2](#), [4](#)
17. Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: CCNet: Criss-cross attention for semantic segmentation. In: ICCV (2019) [1](#)
18. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017) [4](#)
19. Kang, G., Jiang, L., Yang, Y., Hauptmann, A.G.: Contrastive adaptation network for unsupervised domain adaptation. In: CVPR (2019) [5](#)
20. Kang, G., Wei, Y., Yang, Y., Zhuang, Y., Hauptmann, A.: Pixel-level cycle association: A new perspective for domain adaptive semantic segmentation. In: NeurIPS (2020) [4](#), [11](#)



21. Kiefer, J., Wolfowitz, J.: Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics* (1952) 10
22. Li, G., Kang, G., Liu, W., Wei, Y., Yang, Y.: Content-consistent matching for domain adaptive semantic segmentation. In: *ECCV* (2020) 2, 4, 6, 7, 8, 10
23. Li, Y., Yuan, L., Vasconcelos, N.: Bidirectional learning for domain adaptation of semantic segmentation. In: *CVPR* (2019) 2, 4
24. Liu, C., Yuen, J., Torralba, A.: SIFT flow: Dense correspondence across scenes and its applications. *IEEE Trans. PAMI* **33**(5) (2010) 5
25. Liu, C., Yuen, J., Torralba, A.: Nonparametric scene parsing via label transfer. *IEEE Trans. PAMI* **33**(12) (2011) 5
26. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *CVPR* (2015) 1
27. Luo, Y., Liu, P., Guan, T., Yu, J., Yang, Y.: Significance-aware information bottleneck for domain adaptive semantic segmentation. In: *ICCV* (2019) 2, 4
28. 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) 2, 4
29. Malisiewicz, T., Gupta, A., Efros, A.A.: Ensemble of Exemplar-SVMs for object detection and beyond. In: *ICCV* (2011) 5
30. Murez, Z., Kolouri, S., Kriegman, D., Ramamoorthi, R., Kim, K.: Image to image translation for domain adaptation. In: *CVPR* (2018) 2, 4
31. Najafi, M., Namin, S.T., Salzmann, M., Petersson, L.: Sample and filter: nonparametric scene parsing via efficient filtering. In: *CVPR* (2016) 5
32. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: *ECCV* (2016) 2, 3, 10
33. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *MICCAI* (2015) 1
34. 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) 2, 3, 10
35. Russell, B.C., Efros, A., Sivic, J., Freeman, W.T., Zisserman, A.: Segmenting scenes by matching image composites (2009) 5
36. Sankaranarayanan, S., Balaji, Y., Jain, A., Lim, S.N., Chellappa, R.: Unsupervised domain adaptation for semantic segmentation with gans. *arXiv preprint arXiv:1711.06969* 2, 4
37. Shin, I., Woo, S., Pan, F., Kweon, I.S.: Two-phase pseudo label densification for self-training based domain adaptation. In: *ECCV* (2020) 11, 13
38. Shrivastava, A., Malisiewicz, T., Gupta, A., Efros, A.A.: Data-driven visual similarity for cross-domain image matching. In: *SIGGRAPH Asia* (2011) 5
39. Singh, G., Kosecka, J.: Nonparametric scene parsing with adaptive feature relevance and semantic context. In: *CVPR* (2013) 5
40. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. *NeurIPS* (2017) 4
41. Song, L., Li, Y., Li, Z., Yu, G., Sun, H., Sun, J., Zheng, N.: Learnable tree filter for structure-preserving feature transform. *NeurIPS* (2019) 1
42. Sun, R., Zhu, X., Wu, C., Huang, C., Shi, J., Ma, L.: Not all areas are equal: Transfer learning for semantic segmentation via hierarchical region selection. In: *CVPR* (2019) 2, 4
43. Tsai, Y.H., Hung, W.C., Schulter, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: *CVPR* (2018) 2, 10, 11

44. Tsai, Y.H., Sohn, K., Schuler, S., Chandraker, M.: Domain adaptation for structured output via discriminative patch representations. In: ICCV (2019) [2](#), [4](#)
45. Uricchio, T., Ballan, L., Seidenari, L., Del Bimbo, A.: Automatic image annotation via label transfer in the semantic space. PR **71** (2017) [5](#)
46. Wang, K., Liew, J.H., Zou, Y., Zhou, D., Feng, J.: PANet: Few-shot image semantic segmentation with prototype alignment. In: ICCV (2019) [4](#)
47. Wang, Q., Dai, D., Hoyer, L., Van Gool, L., Fink, O.: Domain adaptive semantic segmentation with self-supervised depth estimation. In: ICCV (2021) [11](#), [12](#), [13](#)
48. Wang, Z., Yu, M., Wei, Y., Feris, R., Xiong, J., Hwu, W.m., Huang, T.S., Shi, H.: Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In: CVPR (2020) [2](#), [4](#)
49. Yang, J., An, W., Wang, S., Zhu, X., Yan, C., Huang, J.: Label-driven reconstruction for domain adaptation in semantic segmentation. In: ECCV (2020) [2](#), [4](#)
50. Yang, Y., Lao, D., Sundaramoorthi, G., Soatto, S.: Phase consistent ecological domain adaptation. In: CVPR (2020) [2](#), [7](#), [8](#), [10](#)
51. Yang, Y., Soatto, S.: FDA: Fourier domain adaptation for semantic segmentation. In: CVPR (2020) [2](#), [7](#), [8](#), [10](#), [11](#), [12](#), [13](#)
52. Zhang, P., Zhang, B., Zhang, T., Chen, D., Wang, Y., Wen, F.: Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In: CVPR (2021) [4](#), [10](#), [11](#), [12](#), [13](#)
53. Zhang, Q., Zhang, J., Liu, W., Tao, D.: Category anchor-guided unsupervised domain adaptation for semantic segmentation. NeurIPS (2019) [11](#), [13](#)
54. Zhang, Y., David, P., Gong, B.: Curriculum domain adaptation for semantic segmentation of urban scenes. In: ICCV (2017) [7](#), [8](#)
55. Zhang, Y., Qiu, Z., Yao, T., Liu, D., Mei, T.: Fully convolutional adaptation networks for semantic segmentation. In: CVPR (2018) [2](#), [4](#)
56. Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., Torr, P.H.: Conditional random fields as recurrent neural networks. In: ICCV (2015) [1](#)
57. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017) [4](#)
58. Zou, Y., Yu, Z., Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: ECCV (2018) [2](#), [4](#), [6](#), [7](#), [8](#), [10](#), [11](#), [12](#), [13](#), [14](#)
59. Zou, Y., Yu, Z., Liu, X., Kumar, B., Wang, J.: Confidence regularized self-training. In: ICCV (2019) [2](#), [4](#), [7](#), [8](#), [10](#), [11](#)