

# Unsupervised Domain Adaptation in Semantic Segmentation via Orthogonal and Clustered Embeddings

Marco Toldo, Umberto Michieli, Pietro Zanuttigh  
Department of Information Engineering, University of Padova  
{toldomarco, umberto.michieli, zanuttigh}@dei.unipd.it

## Abstract

Deep learning frameworks allowed for a remarkable advancement in semantic segmentation, but the data hungry nature of convolutional networks has rapidly raised the demand for adaptation techniques able to transfer learned knowledge from label-abundant domains to unlabeled ones. In this paper we propose an effective Unsupervised Domain Adaptation (UDA) strategy, based on a feature clustering method that captures the different semantic modes of the feature distribution and groups features of the same class into tight and well-separated clusters. Furthermore, we introduce two novel learning objectives to enhance the discriminative clustering performance: an orthogonality loss forces spaced out individual representations to be orthogonal, while a sparsity loss reduces class-wise the number of active feature channels. The joint effect of these modules is to regularize the structure of the feature space. Extensive evaluations in the synthetic-to-real scenario show that we achieve state-of-the-art performance.

## 1. Introduction

Semantic segmentation is one of the most challenging prediction tasks towards complete scene understanding and has achieved substantial improvements thanks to deep learning architectures. State-of-the-art approaches typically rely on an auto-encoder structure, where an encoder extracts meaningful compact representations of the scene and a decoder processes them to obtain a dense segmentation map. Starting from the well-known FCN architecture [29], many models have been proposed, such as PSPNet [73], DRN [70] and DeepLab [6, 5, 4]. The main drawback of such architectures is their high complexity: indeed, their success is strictly related to the availability of massive amounts of labeled data. For this reason, many datasets have been created (e.g., Cityscapes [13] or Mapillary [35] for urban scene




Figure 1: The proposed domain adaptation scheme is driven by 3 main components, i.e., feature clustering, orthogonality and sparsity. These push features in the previous step (in light gray) to new locations (colored) where features of the same class are clustered, while features of distinct classes are pushed away. To further improve performances, features of distinct classes are forced to be orthogonal and sparse.

understanding), but the pixel-wise annotation procedure is highly expensive and time consuming. To avoid this demanding process, UDA has come into play in order to exploit knowledge extracted from (related) data sources where labels are easily accessible to the problem at hand.

Three main levels to which adaptation may occur can be identified [55]: namely, at the input, features or output stages. A popular solution has become to bridge the domain gap at an intermediate or output representation level by means of adversarial techniques [22, 21, 33]. The major drawback of these kind of approaches is that they usually perform a semantically unaware alignment, as they neglect the underlying class-conditional data distribution. Additionally, they typically require a long training time to converge and the process may be unstable.

Differently from these techniques, our approach is sim-

Our work was in part supported by the Italian Minister for Education (MIUR) under the “Departments of Excellence” initiative (Law 232/2016).

ple and does not require complex adversarial learning schemes: it entails only a slight increase of computation time with respect to the sole supervised learning. In this work we focus on feature-level UDA. The main idea is depicted in Fig. 1: we devise a domain adaptation technique specifically targeted to guide the latent space organization and driven by 3 main components. The first is a feature clustering method developed to group together features of the same class, while pushing apart features belonging to different classes. This constraint, which works on both domains, is similar in spirit to the recent progresses in contrastive learning for classification problems [8]; however, it has been developed aiming at a simpler computation, as the number of features per image is significantly larger than in the classification task. The second is a novel orthogonality requirement for the feature space, aiming at reducing the cross-talk between features belonging to different classes. Finally, a sparsity constraint is added to reduce the number of active channels for each feature vector: our aim is to enforce the capability of deep learning architectures to learn a compact representation of the scene. The combined effect of these modules allows to regularize the structure of the latent space in order to encompass the source and target domains in a shared representation.

Summarizing, our main contributions are: (1) we extend feature clustering (similarly to contrastive learning) to semantic segmentation; (2) we introduce orthogonality and sparsity objectives to force a regular structure of the embedding space; (3) we achieve state-of-the-art results on feature-level adaptation on two widely used benchmarks.

## 2. Related Work

**Unsupervised Domain Adaptation.** The majority of the UDA techniques focus on a distribution matching to bridge the statistical gap affecting data representations from different domains. While early approaches directly quantify the domain shift in terms of first or second order statistics [60, 30, 31, 51, 52], extremely successful has been the introduction of adversarial learning as an adaptation technique [19]: a domain discriminator implicitly captures the domain discrepancy, which is then minimized in an adversarial fashion. Adversarial adaptation has been first introduced for image classification, both at pixel and feature levels, in order to produce a target-like artificial supervision [28, 3] or to learn a domain invariant latent space [19, 59]. A variety of techniques have been proposed to tackle the UDA task, ranging from completely non-adversarial approaches [67, 16, 1, 68, 12, 43] to enhanced adversarial strategies [14, 63].

**UDA for Semantic Segmentation.** The aforementioned adaptation methods have been developed in the context of image classification. Nonetheless, semantic segmentation introduces further challenges, as dense classification de-

mands for structured predictions that must capture high-level global contextual information, while simultaneously reaching pixel-level precision. Thus, domain adaptation in semantic segmentation calls for more sophisticated techniques to face the additional complexities. In turn, the adaptation is overly rewarding, being human effort in pixel-level labeling extremely expensive and time consuming.

Multiple methods have been recently proposed to tackle this scenario. Some works seek for statistical matching in the input (image) space to achieve cross-domain uniformity of visual appearance. This has been attained through generative adversarial techniques [21, 34, 46, 10, 56, 39] and with style transfer solutions [18, 66, 11]. Alternatively, domain distribution alignment has been enforced over semantic representations, both on intermediate or output network levels. In this direction, adversarial learning has been widely adopted [22, 45, 36, 24, 57, 58, 17, 37], as well as entropy minimization [7, 61], self-training [2, 33, 50, 75, 74, 69] or curriculum style approaches [72, 26].

**Clustering in UDA.** A few approaches have been proposed recently to address UDA in image classification by resorting to a discriminative clustering algorithm, whose goal is to disclose target class-conditional modes in the feature space. A group of works [23, 27, 62, 54] embed variations of the K-means algorithm in the overall adaptation framework, where clusters of geometrically close latent representations are identified, revealing the semantic modes from the unlabeled target domain. Being developed to address image classification, these clustering strategies may lose efficacy when dealing with semantic segmentation. Moreover, they deeply rely on a geometric measure of feature similarity to assign target pseudo-labels, which may not be feasible in a very high-dimensional space. For this reason, this type of clustering technique is often combined with a learnable projection to discover a more easily tractable lower dimensional latent space [27, 62, 54].

To overcome the lack of target semantic supervision, other approaches [67, 16] resort to pseudo-labels directly from network predictions to discover target class-conditional structures in the feature space. Those structures are then exploited to perform a within-domain feature clusterization [16] and cross-domain feature alignment by centroid matching [67, 16]. Starting from analogous premises, we extend a similar form of inter and intra class adaptation to the semantic segmentation scenario, by introducing additional modules that help to address the inherent increased complexity. Avoiding the need for target pseudo-labels, [43, 44] propose a self-supervised clustering technique to discover target modes without any form of supervision. However, their approach is not easily scalable to semantic segmentation, as it requires to store feature embeddings of past samples, which is rather impractical when each data instance is associated with thousands of latent representations.

Quite recently, Tang et al. [53] argue that a direct class-wise alignment over source and target features could harm the discriminative structure of target data. Thus, they perform intrinsic feature alignment by a joint, yet distinct, model training with both source and target data. Nonetheless, in a more complex semantic segmentation scenario, an implicit adaptation could be not enough to bridge the domain gap that affects the effectiveness of target predictions.

**Orthogonality and Sparsity.** Deep neural networks are trained to learn a compact representation of the scene. However, no constraint is posed on the orthogonality among feature vectors belonging to different classes or on the sparsity of their activations [48, 64]. The orthogonality between feature vectors has been recently proposed for UDA in [38, 65]. In these works, a penalty term is introduced to force the prototypes (i.e., the class centers) to be orthogonal. Differently from these approaches, we apply a feature-wise entropy minimization constraint, imposing that each feature is orthogonal with respect to all but one centroid.

To minimize the interference among features we drive them to be channel-wise sparse. To the best of our knowledge, there are no prior works employing channel-wise sparsification in deep learning models. However, some prior techniques exist for domain adaptation on linear models exploiting sparse codes on a shared dictionary between the domains [47, 71]. Additionally, in [40] an expectation maximization approach is proposed to compute sparse code vectors which minimize the energy of the model. Although the approach is applied to linear encoders and decoders for simplicity, it could be extended to non-linear models.

### 3. Method

In this section, we provide an in depth description of the core modules of the proposed method. Our approach leverages a clustering objective applied over the individual feature representations, with novel orthogonality and sparsity constraints. Specifically, inter and intra class alignments are enforced by grouping together features of the same semantic class, while simultaneously pushing away those of different categories. By enforcing the clustering objective on both source and target representations, we drive the model towards feature-level domain alignment. We further regularize the distribution of latent representations by the joint application of an orthogonality and a sparsity losses. The orthogonality module has a two-fold objective: first, it forces feature vectors of kindred semantic connotations to activate the same channels, while turning off the remaining ones; second, it constrains feature vectors of dissimilar semantic connotations to activate different channels, i.e., with no overlap, to reduce cross interference. The sparsity objective further encourages a lower volume of active feature channels from latent representations, i.e., it concentrates the energy of the features on few dimensions.

A graphical outline of the approach with all its components is shown in Fig. 2: the training objective is given by the combination of the standard supervised loss with the proposed adaptation modules, i.e., it is computed as:

$$\mathcal{L}'_{tot} = \mathcal{L}_{ce} + \lambda_{cl} \cdot \mathcal{L}_{cl} + \lambda_{or} \cdot \mathcal{L}_{or} + \lambda_{sp} \cdot \mathcal{L}_{sp} \quad (1)$$

where  $\mathcal{L}_{ce}$  is the standard supervised cross entropy loss. The other components will be detailed in the following sections: the main clustering objective ( $\mathcal{L}_{cl}$ ) is introduced in Section 3.1. The orthogonality constraint ( $\mathcal{L}_{or}$ ) is discussed in Section 3.2 and finally the sparsity constraint ( $\mathcal{L}_{sp}$ ) is detailed in Section 3.3. The  $\lambda$  parameters balance the multiple losses and are experimentally chosen using a validation set. In addition, we further integrate the proposed adaptation method with an off-the-shelf entropy-minimization like objective ( $\mathcal{L}_{em}$ ), to provide an extra regularizing action over the segmentation feature space and ultimately achieve an improved performance in some evaluation scenarios. In particular, we adopt the simple, yet effective, maximum squares objective of [7], in its *image-wise class-balanced* version. Hence, we can define the ultimate training objective comprising the entropy module as:

$$\mathcal{L}_{tot} = \mathcal{L}'_{tot} + \lambda_{em} \cdot \mathcal{L}_{em} \quad (2)$$

#### 3.1. Discriminative Clustering

In the considered UDA setting, we are provided with plenty of samples  $\mathbf{X}_n^s \in \mathbb{R}^{H \times W \times 3}$  from a source dataset, in conjunction with their semantic maps  $\mathbf{Y}_n^s \in \mathbb{R}^{H \times W}$ . Those semantic maps contain at each spatial location a ground truth index belonging to the set of possible classes  $\mathcal{C}$ , which denotes the semantic category of the associated pixel. Currently, we have at our disposal target training samples  $\mathbf{X}_n^t \in \mathbb{R}^{H \times W \times 3}$  with no label maps (we allow only the availability of a small amount of target labels for validation and testing purposes). Despite sharing similar high-level semantic content, the source and training samples are distributed differently, preventing a source-based model to achieve a satisfying prediction accuracy on target data without adaptation. We denote as  $S = F \circ C$  the segmentation network composed of an encoder and a decoder modules, namely the feature extractor  $F$  and the classifier  $C$ . Notice that the proposed method is agnostic to the employed deep learning model, except for the assumption of an auto-encoder structure and of positive feature values as provided by ReLU activations that are typically placed at the encoder output (as almost all the current state-of-the-art approaches for semantic segmentation).

To bridge the domain gap between the source and target datasets we operate at the feature level. The discrepancy of input statistics across domains is reflected into a shift of feature distribution in the latent space spanned by the feature extractor. This ultimately may cause the source-trained

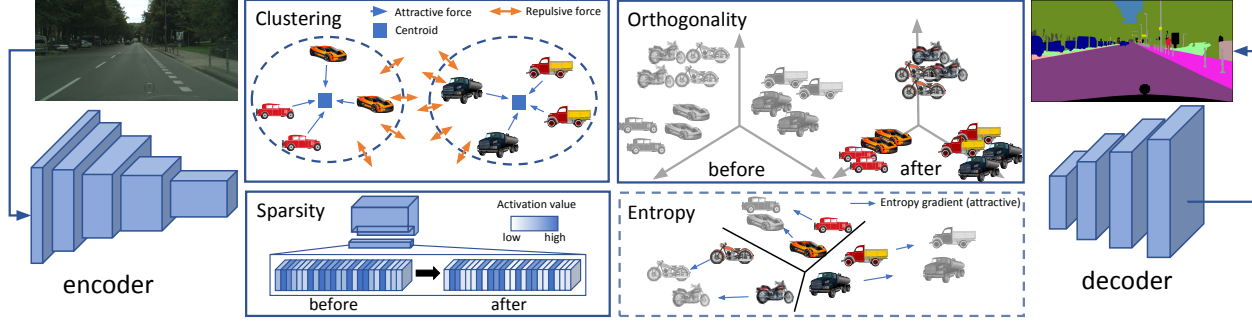


Figure 2: Overview of the proposed approach. Features after supervised training on the source domain are represented in light gray, while features of the current step are colored. A set of techniques is employed to better shape the latent feature space spanned by the encoder. Features are clustered and the clusters are forced to be disjoint. At the same time, features belonging to different classes are forced to be orthogonal with respect to each other. Additionally, features are forced to be sparse and an entropy minimization loss could also be added to guide target samples far from the decision boundaries.

classifier to draw decision boundaries crossing high density regions of the target latent space [61], since it is inherently unaware of the target semantic modes extracted from unlabeled target data. Thus, the classification performance over the target domain is strongly degraded when compared to the upper bound of the source prediction accuracy.

We cope with this performance degradation by resorting to a clustering module, that serves as constraint towards a class-conditional feature alignment between domains. Given a batch of source ( $\mathbf{X}_n^s$ ) and target ( $\mathbf{X}_n^t$ ) training images (for ease of notation we pick a single image per domain), we first extract the feature tensors  $\mathbf{F}_n^s = F(\mathbf{X}_n^s)$  and  $\mathbf{F}_n^t = F(\mathbf{X}_n^t)$ , along with the computed output segmentation maps  $\mathbf{S}_n^s = S(\mathbf{X}_n^s)$  and  $\mathbf{S}_n^t = S(\mathbf{X}_n^t)$ . The clustering loss is then computed as:

$$\mathcal{L}_{cl} = \frac{1}{|\mathbf{F}_n^{s,t}|} \sum_{\substack{\mathbf{f}_i \in \mathbf{F}_n^{s,t} \\ \hat{y}_i \in \mathbf{S}_n^{s,t}}} d(\mathbf{f}_i, \mathbf{c}_{\hat{y}_i}) - \frac{1}{|\mathcal{C}|(|\mathcal{C}|-1)} \sum_{j \in \mathcal{C}} \sum_{\substack{k \in \mathcal{C} \\ k \neq j}} d(\mathbf{c}_j, \mathbf{c}_k) \quad (3)$$

where  $\mathbf{f}_i$  is an individual feature vector corresponding to a single spatial location from either source or target domain and  $\hat{y}_i$  is the corresponding predicted class (to compute  $\hat{y}_i$  the segmentation map  $\mathbf{S}_n^{s,t}$  is downsampled to match the feature tensor spatial dimensions). The function  $d(\cdot)$  represents a generic distance measure, that we set to the  $L1$  norm (we also tried the  $L2$  norm but it yielded lower results). Finally,  $\mathbf{c}_j$  denotes the centroid of semantic class  $j \in \mathcal{C}$  computed according to the standard formula:

$$\mathbf{c}_j = \frac{\sum_{\mathbf{f}_i} \sum_{\hat{y}_i} \delta_{j, \hat{y}_i} \mathbf{f}_i}{\sum_{\hat{y}_i} \delta_{j, \hat{y}_i}}, \quad j \in \mathcal{C} \quad (4)$$

where  $\delta_{j, \hat{y}_i}$  is equal to 1 if  $\hat{y}_i = j$ , and to 0 otherwise.

The clustering objective is composed of two terms, the first measures how close features are from their respective

centroids and the second how spaced out clusters corresponding to different semantic classes are. Hence, the effect provided by the loss minimization is twofold: firstly, feature vectors from the same class but different domains are tightened around class feature centroids; secondly, features from separate classes are subject to a repulsive force applied to feature centroids, moving them apart.

### 3.2. Orthogonality

As opposed to previous works on clustering-based adaptation methods for image classification, in semantic segmentation additional complexity is brought by the dense structured classification. To this end, we first introduce an orthogonality constraint in the form of a training objective. More precisely, feature vectors from either domains, but of different semantic classes according to the network predictions, are forced to be orthogonal, meaning that their scalar product should be small. On the contrary, features sharing semantic classification should carry high similarity, i.e., large scalar product. Yet, feature tensors associated to training samples enclose thousands of feature vectors to cover the entire spatial extent of the scene and to reach pixel-level classification. Thus, since measuring pair-wise similarities requires a significant computational effort, we calculate the scalar product between each feature vector and every class centroid  $\mathbf{c}_j$  (centroids are computed using Eq. 4). Inspired by [43, 44], we devise the orthogonality objective as an entropy minimization loss that forces each feature to be orthogonal with respect to all the centroids but one:

$$\mathcal{L}_{or} = - \sum_{\mathbf{f}_i \in F(\mathbf{X}_n^{s,t})} \sum_{j \in \mathcal{C}} p_j(\mathbf{f}_i) \log p_j(\mathbf{f}_i) \quad (5)$$

where  $\{p_j(\mathbf{f}_i)\}$  denotes a probability distribution derived as:

$$p_j(\mathbf{f}_i) = \frac{e^{\langle \mathbf{f}_i, \mathbf{c}_j \rangle}}{\sum_{k \in \mathcal{C}} e^{\langle \mathbf{f}_i, \mathbf{c}_k \rangle}}, \quad j \in \mathcal{C} \quad (6)$$



The loss minimization forces a peaked distribution of the probabilities  $\{p_j(\mathbf{f}_i)\}$ , promoting the orthogonality property as described above, since each feature vector is compelled to carry a high similarity score with a single class centroid. The overall effect of the orthogonality objective is to promote a regularized feature distribution, which should ultimately boost the clustering efficacy in performing domain feature alignment.

### 3.3. Sparsity

To strengthen the regularizing effect brought by the orthogonality constraint, we introduce a further training objective to better shape class-wise feature structures inside the latent space. In particular, we propose a sparsity loss, with the intent of decreasing the number of active feature channels of latent vectors. The objective is defined as follows:

$$\mathcal{L}_{sp} = - \sum_{i \in \mathcal{C}} \|\tilde{\mathbf{c}}_i - \boldsymbol{\rho}\|_2^2 \quad (7)$$

where  $\tilde{\mathbf{c}}_i$  stands for the normalized centroid  $\mathbf{c}_i$  in  $[0, 1]^D$  and  $D$  denotes the number of feature maps in the encoder output. We also empirically set  $\boldsymbol{\rho} = [0.5]^D$ . It can be noted that the sparsifying action is delivered on class centroids, thus applying an indirect, yet homogeneous, influence over all feature vectors from the same semantic category. The result is a semantically-consistent suppression of weak activations, while rather active ones are jointly raised.

While the orthogonality objective aims at promoting different sets of activations on feature vectors from separate semantic classes, the sparsity loss seeks to narrow those sets to a limited amount of units. Again, the goal is to ease the clustering loss task in creating tight and well distanced aggregations of features of similar semantic connotation from either source and target domains, by providing an improved regularity to the class-conditional semantic structures inside the feature space.

## 4. Experimental Setup

**Datasets.** Supervised training is performed on the synthetic datasets GTA5 [41] and SYNTHIA [42]. We employ Cityscapes [13] as target domain. The GTA5 [41] contains 24966 synthetic images taken from a car perspective in US-style virtual cities, with a high level of quality and realism. On the other hand, the SYNTHIA [42] offers 9400 images from virtual European-style towns with large scene variability under various light and weather conditions, but little visual quality. For evaluation on real datasets, respectively 19 and 16 compatible classes are considered when adapting from GTA5 and SYNTHIA. The Cityscapes [13] is instead provided with the limited amount of 2975 images acquired in 50 European cities. The original training set (without labels) is used for unsupervised adaptation, while the 500 images in the original validation set are used as a test set.

**Model Architecture.** The modules introduced in this work are agnostic to the underlying network architecture and can be extended to other scenarios. For fair comparison with previous works [57, 7, 61] we employ the DeepLab-V2, a fully convolutional segmentation network with ResNet-101 [20] or VGG-16 [49] as backbones. Further details on the segmentation network architecture can be found in [57, 61], as we follow the same implementation adopted in those works. We initialize the two encoder networks with ImageNet [15] pretrained weights. In addition, prior to the actual adaptation phase, we supervisedly train the segmentation network on source data.

**Training Details.** The model is trained with the starting learning rate set to  $2.5 \times 10^{-4}$  and decreased with a polynomial decay rule of power 0.9. We employ weight decay regularization of  $5 \times 10^{-4}$ . Following [7], we also randomly apply mirroring and gaussian blurring for data augmentation during the training stage. To accommodate for GPU memory limitations, we resize images from the GTA5 dataset up to a resolution of  $1280 \times 720$  px, as done by [57]. SYNTHIA images are instead kept to the original size of  $1280 \times 780$  px. As for the target Cityscapes dataset, training unlabeled images are resized to  $1024 \times 512$  px, whereas the results of the testing stage are reported at the original image resolution ( $2048 \times 1024$  px). We use a batch size of 1 and the hyper-parameters are tuned by resorting to a small subset of labeled target data that we set aside from the original target training set and reserve only for parameter selection. As evaluation metric, we employ the mean Intersection over Union (mIoU). The entire model is developed using PyTorch and trained with a single GPU. The code is available at [https://lttm.dei.unipd.it/paper\\_data/UDAclustering/](https://lttm.dei.unipd.it/paper_data/UDAclustering/).

## 5. Results

We evaluate the performance of our approach on two widely used synthetic-to-real adaptation scenarios, namely the GTA5  $\rightarrow$  Cityscapes and SYNTHIA  $\rightarrow$  Cityscapes benchmarks. Table 1 reports the numerical results of the experimental evaluation. We compare our model to several state-of-the-art methods, which, similarly to our approach, resort to a direct or indirect form of feature-level regularization and distribution alignment to achieve domain adaptation. With *source only* we indicate the naïve fine-tuning approach, in which no form of target adaptation assists the standard source supervision.

### 5.1. GTA5 $\rightarrow$ Cityscapes

For the GTA5  $\rightarrow$  Cityscapes and ResNet-101 configuration, our approach shows state-of-the-art performances in feature-level UDA for semantic segmentation, achieving 45.3% of mIoU, which is further boosted up to 45.9% by the entropy minimization objective. By looking at Ta-

B	Method	Road	Sidewalk	Building	Wall*	Fence*	Pole*	T. Light	T. Sign	Vegetation	Terrain	Sky	Person	Rider	Car	Truck	Bus	Train	Motorbike	Bicycle	mIoU (all)	mIoU* (13-cl)
GTA5 → Cityscapes	Source Only	26.5	13.3	45.1	6.0	15.2	16.5	21.3	8.5	78.0	8.3	59.7	45.0	10.5	69.1	22.8	17.9	0.0	16.4	2.7	25.4	-
	VGG16																					
	FCNs ITW [22]	70.4	<b>32.4</b>	62.1	14.9	5.4	10.9	14.2	2.7	79.2	21.3	64.6	44.1	4.2	70.4	8.0	7.3	0.0	3.5	0.0	27.1	-
	CyCADA (feat) [21]	85.6	30.7	74.7	14.4	13.0	17.6	13.7	5.8	74.6	15.8	69.9	38.2	3.5	72.3	16.0	5.0	0.1	3.6	0.0	29.2	-
	CBST [75]	66.7	26.8	73.7	14.8	9.5	<b>28.3</b>	25.9	10.1	75.5	15.7	51.6	<b>47.2</b>	6.2	71.9	3.7	2.2	<b>5.4</b>	<b>18.9</b>	<b>32.4</b>	30.9	-
	MinEnt [61]	85.1	18.9	76.3	<b>32.4</b>	<b>19.7</b>	19.9	21.0	8.9	76.3	<b>26.2</b>	63.1	42.8	5.9	<b>80.8</b>	<b>20.2</b>	9.8	0.0	14.8	0.6	32.8	-
	MaxSquare IW <sup>(r)</sup> [7]	81.4	20.0	75.4	19.4	19.1	16.1	24.4	7.9	78.8	22.9	65.9	45.0	12.3	74.6	16.1	10.3	0.2	11.3	1.0	31.7	-
	Ours ( $\mathcal{L}'_{tot}$ )	83.6	16.6	79.0	19.8	18.7	21.5	27.3	<b>15.9</b>	80.2	14.3	<b>72.6</b>	47.0	17.5	76.8	16.6	<b>13.9</b>	0.1	16.0	3.4	33.7	-
	Ours ( $\mathcal{L}_{tot}$ )	<b>86.0</b>	13.5	<b>79.4</b>	20.4	18.5	21.5	<b>27.6</b>	15.2	<b>80.8</b>	21.9	<b>72.6</b>	46.3	<b>18.1</b>	80.0	16.9	13.1	1.0	14.6	2.0	<b>34.2</b>	-
	ResNet101																					
	Source Only	81.8	16.3	74.4	18.6	12.7	23.5	29.3	18.1	73.5	21.4	77.6	55.6	25.6	74.1	28.6	10.2	3.0	25.8	32.7	37.0	-
	AdaptSegNet (feat) [57]	83.7	27.6	75.5	20.3	19.9	27.4	28.3	27.4	79.0	28.4	70.1	55.1	20.2	72.9	22.5	35.7	<b>8.3</b>	20.6	23.0	39.3	-
	MinEnt [61]	84.4	18.7	80.6	23.8	23.2	28.4	36.9	23.4	83.2	25.2	<b>79.4</b>	59.0	<b>29.9</b>	78.5	33.7	29.6	1.7	29.9	33.6	42.3	-
	SAPNet [25]	88.4	38.7	79.5	<b>29.4</b>	<b>24.7</b>	27.3	32.6	20.4	82.2	32.9	73.3	55.5	26.9	82.4	31.8	41.8	2.4	26.5	24.1	43.2	-
	MaxSquare IW [7]	89.3	<b>40.5</b>	81.2	29.0	20.4	25.6	34.4	19.0	83.6	34.4	76.5	59.2	27.4	83.8	<b>38.4</b>	43.6	7.1	<b>32.2</b>	32.5	45.2	-
	Ours ( $\mathcal{L}'_{tot}$ )	88.7	32.2	81.8	24.1	22.1	<b>30.8</b>	<b>37.6</b>	<b>32.8</b>	83.4	36.3	76.0	60.0	27.0	81.0	34.2	43.0	8.0	23.4	38.1	45.3	-
	Ours ( $\mathcal{L}_{tot}$ )	<b>89.4</b>	30.7	<b>82.1</b>	23.0	22.0	29.2	<b>37.6</b>	31.7	<b>83.9</b>	<b>37.9</b>	78.3	<b>60.7</b>	27.4	<b>84.6</b>	37.6	<b>44.7</b>	7.3	26.0	<b>38.9</b>	<b>45.9</b>	-
SYNTHIA → Cityscapes	Source Only	7.8	13.7	66.6	2.2	0.0	23.9	4.8	13.3	71.2	-	76.5	49.2	12.1	67.1	-	24.5	-	9.8	9.2	28.3	32.8
	VGG16																					
	FCNs ITW [22]	11.5	19.6	30.8	4.4	0.0	20.3	0.1	11.7	42.3	-	68.7	51.2	3.8	54.0	-	3.2	-	0.2	0.6	20.2	22.9
	Cross-City [9]	62.7	25.6	<b>78.3</b>	-	-	-	1.2	5.4	81.3	-	81.0	37.4	6.4	63.5	-	16.1	-	1.2	4.6	-	35.7
	CBST [75]	69.6	28.7	69.5	<b>12.1</b>	<b>0.1</b>	<b>25.4</b>	11.9	13.6	<b>82.0</b>	-	<b>81.9</b>	49.1	14.5	66.0	-	6.6	-	3.7	<b>32.4</b>	35.4	36.1
	MinEnt [61]	37.8	18.2	65.8	2.0	0.0	15.5	0.0	0.0	76.0	-	73.9	45.7	11.3	66.6	-	13.3	-	1.5	<b>13.1</b>	27.5	32.5
	MaxSquare IW <sup>(r)</sup> [7]	9.1	12.7	72.5	1.0	0.0	22.3	7.0	8.4	80.0	-	77.9	49.4	10.0	71.8	-	23.8	-	6.0	13.5	29.1	34.0
	Ours ( $\mathcal{L}'_{tot}$ )	<b>78.5</b>	29.9	77.7	1.2	<b>0.1</b>	24.1	11.9	<b>15.0</b>	78.7	-	78.5	51.0	15.4	73.7	-	<b>24.7</b>	-	<b>10.1</b>	23.5	<b>37.1</b>	<b>43.7</b>
	Ours ( $\mathcal{L}_{tot}$ )	<b>78.3</b>	<b>30.1</b>	78.0	1.7	<b>0.1</b>	24.1	<b>12.0</b>	14.6	79.7	-	79.1	<b>51.4</b>	<b>15.5</b>	<b>74.4</b>	-	23.7	-	9.1	22.7	<b>37.1</b>	<b>43.7</b>
	ResNet101																					
	Source Only	39.5	18.1	75.5	10.5	0.1	26.3	9.0	11.7	78.6	-	81.6	57.7	21.0	59.9	-	30.1	-	15.7	28.2	35.2	40.5
	AdaptSegNet (feat) [57]	62.4	21.9	76.3	-	-	-	11.7	11.4	75.3	-	80.9	53.7	18.5	59.7	-	13.7	-	<b>20.6</b>	24.0	-	40.8
	MinEnt [61]	73.5	29.2	77.1	7.7	0.2	27.0	7.1	11.4	76.7	-	82.1	57.2	<b>21.3</b>	69.4	-	29.2	-	12.9	27.9	38.1	44.2
	SAPNet [25]	81.7	33.5	75.9	-	-	-	7.0	6.3	74.8	-	78.9	52.1	<b>21.3</b>	75.7	-	<b>30.6</b>	-	10.8	28.0	-	44.3
	MaxSquare IW [7]	78.5	34.7	76.3	6.5	0.1	<b>30.4</b>	12.4	12.2	<b>82.2</b>	-	<b>84.3</b>	<b>59.9</b>	17.9	80.6	-	24.1	-	15.2	31.2	40.4	46.9
	Ours ( $\mathcal{L}'_{tot}$ )	64.4	25.5	77.3	<b>14.3</b>	<b>0.9</b>	29.6	<b>21.2</b>	<b>24.2</b>	76.6	-	79.7	53.7	15.5	79.7	-	11.0	-	11.0	<b>35.2</b>	38.7	44.2
	Ours ( $\mathcal{L}_{tot}$ )	<b>88.3</b>	<b>42.2</b>	<b>79.1</b>	7.1	0.2	24.4	16.8	16.5	80.0	-	<b>84.3</b>	56.2	15.0	<b>83.5</b>	-	27.2	-	6.3	30.7	<b>41.1</b>	<b>48.2</b>

Table 1: Numerical evaluation of the GTA5 and SYNTHIA to Cityscapes adaptation scenarios in terms of per-class and mean IoU. Evaluations are performed on the validation set of the Cityscapes dataset. In all the experiments the DeepLab-V2 segmentation network is employed, with VGG-16 (top) or ResNet-101 (bottom) backbones. The mIoU\* results in the last column refer to the 13-classes configuration, i.e., classes marked with \* are ignored. MaxSquares IW <sup>(r)</sup> denotes our re-implementation, as original results are provided only for the ResNet-101 backbone.

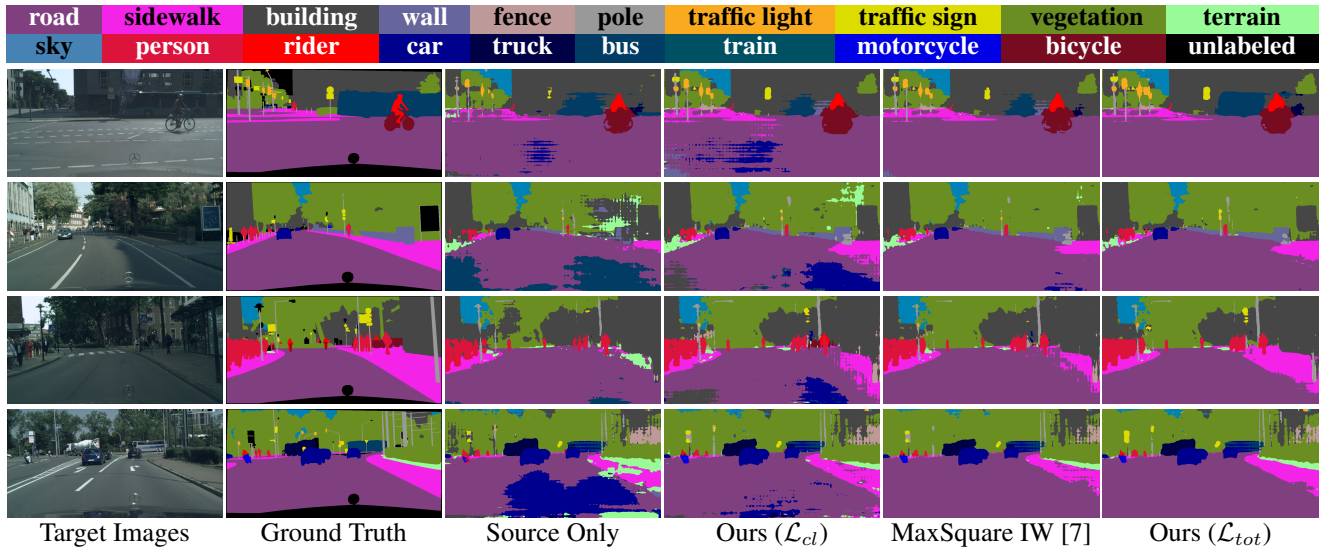


Figure 3: Semantic segmentation of some sample scenes from the Cityscapes validation dataset when adaptation is performed from the GTA5 source dataset and the DeepLab-V2 with ResNet-101 backbone is employed (*best viewed in colors*).

ble 1, we observe about 9% increase over the *source only* baseline, with the improvement well distributed over all the classes. A similar behavior can be noted when switching to the less performing VGG-16 backbone: we achieve 33.7% mean IoU with  $\mathcal{L}'_{tot}$  (i.e., without the entropy minimization objective) and 34.2% with  $\mathcal{L}_{tot}$  (i.e., with all components enabled) starting from the 25.4% of the baseline scenario without adaptation.

When compared with other approaches, our method performs better than standard feature-level adversarial techniques [22, 21, 57, 25]. For example, with the VGG-16 backbone there is a gain of 4.5% w.r.t. [21]. This proves that a more effective class-conditional alignment has been ultimately achieved in the latent space by our approach. Due to the similar regularizing effect over feature distribution and comparable ease of implementation, we also compare our framework with some entropy minimization and self-training techniques [61, 7, 75], further showing the effectiveness of our adaptation strategy even if the gap here is a bit more limited. With both backbones, our novel feature level modules ( $\mathcal{L}'_{tot}$ ) perform better than *MaxSquare IW* [7]. Moreover, adding the entropy minimization objective from [7] to  $\mathcal{L}'_{tot}$  provides a slight but consistent improvement. It is worth noting that our method does not rely on additional trainable modules (e.g., adversarial discriminators [22, 21, 57, 25]) and the whole adaptation process is end-to-end, not requiring multiple separate steps to be re-iterated (e.g., pseudo-labeling in self-training [75]). Moreover, being focused solely on feature level adaptation, it could be easily integrated with other adaptation techniques working at different network levels, such as the input (e.g., generative approaches) or output (e.g., self-training), as shown by the addition of the output-level entropy-minimization loss.

Fig. 3 displays some qualitative results on the Cityscapes validation set of the adaptation process when the ResNet-101 backbone is used. We observe that the introduction of the clustering module is beneficial to the target segmentation accuracy w.r.t. the *source only* case. Some small prediction inaccuracies remain, that are corrected with the introduction of the orthogonality, sparsity and entropy modules in the complete framework. By looking at the last two columns, we also notice that our entire framework shows an improvement over the individual entropy-minimization like objective from [7], which is reflected in a better detection accuracy both on frequent (e.g. *road*, *vegetation*) and less frequent (e.g. *traffic sign*, *bus*) classes.

## 5.2. SYNTHIA → Cityscapes

To further prove the efficacy of our method, we evaluate it on the more challenging SYNTHIA → Cityscapes benchmark, where a larger domain gap exists. Once more, our approach proves to be successful in performing domain alignment with both ResNet-101 and VGG-16 backbones, reach-

$\mathcal{L}_{cl}$	$\mathcal{L}_{or}$	$\mathcal{L}_{sp}$	$\mathcal{L}_{em}$	mIoU
				37.0
✓				42.3
	✓			43.2
		✓		43.7
			✓	44.8
✓	✓	✓		45.3
✓	✓	✓	✓	45.9

Table 2: Ablation results on the contribution of each adaptation module in the GTA5 to Cityscapes scenario and with ResNet-101 as backbone.

ing state-of-the-art results for feature-level UDA in both configurations (see Table 1). When ResNet-101 is used, the mIoU\* on the 13 classes setting is pushed up to 48.2% from the original 40.5% of *source only*, while the VGG-16 scenario witnesses an even more improved performance gain of almost 11% over the no adaptation baseline till a final value of 43.7%. Differently from the GTA5 → Cityscapes case, here the contribution of the entropy-minimization module varies for the two backbones. The induced benefit is absent with VGG-16, since the clustering, orthogonality and sparsity jointly enforced already carry the whole adaptation effort. Besides, even the  $\mathcal{L}_{em}$  objective alone (i.e., *MaxSquares IW*<sup>(r)</sup> [7]) displays quite limited gain over the no adaptation baseline. On the contrary, the regularizing effect of the entropy objective is strongly valuable in case the ResNet-101 backbone is used. Yet, the combination of all modules together actually provides a noticeable boost over both the entropy and feature-level modules separately applied. As for the GTA5 scenario, our model shows better performance than feature-level adversarial adaptation [22, 9, 57, 25] and output-level approaches [61, 75] comparable in computational ease. Qualitative results of adaptation from SYNTHIA are in the Supplementary Material.

## 5.3. Ablation Study

To verify the robustness of the framework, we perform an extensive ablation study on the adaptation from GTA5 to Cityscapes with ResNet-101 as backbone. First, we examine the contribution of each loss to the final mIoU; then, we investigate the effect of each novel loss component. Further considerations are reported in the Supplementary Material.

The contribution of each loss to the adaptation module is shown in Table 2. Every loss component largely improves the final mIoU results from 37.0% of the *source only* scenario up to a maximum of 44.8%. Combining the 3 novel modules of this work, we achieve a mIoU of 45.3%, which is higher than all the losses alone, but lower than our complete framework with all the losses enabled (45.9%).

To investigate the effect of the clustering module ( $\mathcal{L}_{cl}$ )




Figure 4: T-SNE computed over features of a single image of the Cityscapes validation set when adapting from GTA5 (best viewed in colors).

we show a t-SNE [32] plot of features extracted from a sample image of the Cityscapes validation set. In particular, we provide a comparison of the discovered low-dimensional feature distribution for two distinct training settings, i.e. *source only* and adaptation with  $\mathcal{L}_{cl}$  only. The results are reported in Fig. 4, where sparse points in the *source only* plot turn out more tightly clustered and spaced apart in the  $\mathcal{L}_{cl}$  approach (e.g., look at the *person* class).

We analyze the orthogonality constraint ( $\mathcal{L}_{or}$ ) via a similarity score defined as an average class-wise cosine similarity measure (Fig. 5). The cosine distance is first computed for every pair of feature vectors from a single target image. Then, the average values are taken over all features from the same class to get a score for each pair of semantic classes. The final values are computed by averaging over all images from the Cityscapes validation set. The score computation is performed for two different configurations, i.e.  $\mathcal{L}_{cl} + \mathcal{L}_{or} + \mathcal{L}_{sp}$  and  $\mathcal{L}_{cl} + \mathcal{L}_{sp}$ , to highlight the effect induced by the orthogonality module. Here we report the intra-class similarity scores, whereas the full matrix with also the inter-class values is in the Supplementary Material. The results in Fig. 5 show that  $\mathcal{L}_{or}$  causes the similarity score to significantly increase within almost all the classes.

To understand the efficacy of the sparsity loss ( $\mathcal{L}_{sp}$ ) we compute the sparsity scores as the fraction of activations in normalized feature vectors close to 0 or 1 (Fig. 6). Closeness is quantified as being distant from 0 or 1 less than a threshold, which we set to  $10^{-4}$ . As for the similarity scores, the sparsity measures for a single target image are obtained through averaging over all feature vectors from the same class. The final results correspond to the mean values over the entire Cityscapes validation set. We compute the scores for two different configurations, i.e.  $\mathcal{L}_{cl} + \mathcal{L}_{or} + \mathcal{L}_{sp}$  and  $\mathcal{L}_{cl} + \mathcal{L}_{or}$ , so that we can inspect the effect that the sparsity module is providing on feature distribution. From Fig. 6 we can appreciate that  $\mathcal{L}_{sp}$  effectively achieves higher values of sparseness for all the classes.




Figure 5: Similarity scores computed over images of the Cityscapes validation set when adapting from GTA5.




Figure 6: Sparsity scores computed over images of the Cityscapes validation set when adapting from GTA5.

## 6. Conclusions

In this paper we propose a novel feature oriented UDA framework for semantic segmentation. Our approach comprises 3 main objectives. First, features of same class and separate domains are clustered together, whilst features of different classes are spaced apart. Second, an orthogonality requirement over the latent space discourages the overlapping of active channels among feature vectors of different classes. Third, a sparsity constraint further reduces feature-wise the number of the active channels. All combined, these modules allow to reach a regularized disposition of latent embeddings, while providing a semantically consistent domain alignment over the feature space. We extensively evaluated our framework in the synthetic-to-real scenario, achieving state-of-the-art results in feature level UDA. For future work, we intend to explore new techniques aiming at the refinement of the pseudo-labeling based classification of feature vectors, and to integrate our feature space adaptation with other approaches targeting different network levels. Finally, we would like to investigate the regularizing effect of the proposed techniques when applied to the single-domain standard semantic segmentation.

## References

- [1] Yogesh Balaji, Rama Chellappa, and Soheil Feizi. Normalized wasserstein for mixture distributions with applications in adversarial learning and domain adaptation. In *Proceedings of International Conference on Computer Vision*, pages 6500–6508, 2019.
- [2] Matteo Basettoni, Umberto Michieli, Gianluca Agresti, and Pietro Zanuttigh. Unsupervised Domain Adaptation for Semantic Segmentation of Urban Scenes. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1211–1220, 2019.
- [3] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3722–3731, 2017.
- [4] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of European Conference on Computer Vision*, pages 833–851, 2018.
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:834–848, 2018.
- [7] Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *Proceedings of International Conference on Computer Vision*, pages 2090–2099, 2019.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the International Conference on Machine Learning*, pages 10709–10719, 2020.
- [9] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *Proceedings of International Conference on Computer Vision*, pages 2011–2020, 2017.
- [10] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1791–1800, 2019.
- [11] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *Proceedings of International Conference on Computer Vision*, pages 6830–6840, 2019.
- [12] Safa Cicek and Stefano Soatto. Unsupervised domain adaptation via regularized conditional alignment. In *Proceedings of International Conference on Computer Vision*, pages 1416–1425, 2019.
- [13] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [14] Shuhao Cui, Shuhui Wang, Junbao Zhuo, Chi Su, Qingming Huang, and Tian Qi. Gradually vanishing bridge for adversarial domain adaptation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 12452–12461, 2020.
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [16] Zhijie Deng, Yucen Luo, and Jun Zhu. Cluster alignment with a teacher for unsupervised domain adaptation. In *Proceedings of International Conference on Computer Vision*, pages 9944–9953, 2019.
- [17] Liang Du, Jingang Tan, Hongye Yang, Jianfeng Feng, Xiangyang Xue, Qibao Zheng, Xiaoqing Ye, and Xiaolin Zhang. SSF-DAN: separated semantic feature based domain adaptation network for semantic segmentation. In *Proceedings of International Conference on Computer Vision*, pages 982–991, 2019.
- [18] Aysegul Dundar, Ming-Yu Liu, Ting-Chun Wang, John Zedlewski, and Jan Kautz. Domain stylization: A strong, simple baseline for synthetic to real image domain adaptation. *arXiv preprint arXiv:1807.09384*, 2018.
- [19] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the International Conference on Machine Learning*, pages 1180–1189, 2015.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [21] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *Proceedings of the International Conference on Machine Learning*, pages 1994–2003, 2018.
- [22] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. FCNs in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016.
- [23] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 4893–4902, 2019.
- [24] Seungmin Lee, Dongwan Kim, Namil Kim, and Seong-Gyun Jeong. Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In *Proceedings of International Conference on Computer Vision*, pages 91–100, 2019.
- [25] Congcong Li, Dawei Du, Libo Zhang, Longyin Wen, Tiejian Luo, Yanjun Wu, and Pengfei Zhu. Spatial attention pyramid

- network for unsupervised domain adaptation. In *Proceedings of European Conference on Computer Vision*, 2020.
- [26] Qing Lian, Fengmao Lv, Lixin Duan, and Boqing Gong. Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: A non-adversarial approach. In *Proceedings of International Conference on Computer Vision*, pages 6758–6767, 2019.
- [27] Jian Liang, Ran He, Zhenan Sun, and Tieniu Tan. Distant supervised centroid shift: A simple and efficient approach to visual domain adaptation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2975–2984, 2019.
- [28] Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 469–477, 2016.
- [29] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [30] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *Proceedings of the International Conference on Machine Learning*, pages 97–105, 2015.
- [31] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2016.
- [32] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9:2579–2605, 2008.
- [33] Umberto Michieli, Matteo Bassetton, Gianluca Agresti, and Pietro Zanuttigh. Adversarial learning and self-teaching techniques for domain adaptation in semantic segmentation. *IEEE Transaction on Intelligent Vehicles*, 5:508–518, 2020.
- [34] Zak Murez, Soheil Kolouri, David J. Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 4500–4509, 2018.
- [35] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The Mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of International Conference on Computer Vision*, pages 4990–4999, 2017.
- [36] Sungrae Park, Jun-Keon Park, Su-Jin Shin, and Il-Chul Moon. Adversarial dropout for supervised and semi-supervised learning. In *Proceedings of AAAI Conference on Artificial Intelligence*, pages 3917–3924, 2018.
- [37] Sujoy Paul, Yi-Hsuan Tsai, Samuel Schuster, Amit K. Roy-Chowdhury, and Manmohan Chandraker. Domain adaptive semantic segmentation using weak labels. In *Proceedings of European Conference on Computer Vision*, 2020.
- [38] Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 8004–8013, 2018.
- [39] Fabio Pizzati, Raoul de Charette, Michela Zaccaria, and Pietro Cerri. Domain bridge for unpaired image-to-image translation and unsupervised domain adaptation. In *Proceedings of Winter Conference on Applications of Computer Vision*, pages 2990–2998, 2020.
- [40] Marc’Aurelio Ranzato, Christopher Poultney, Sumit Chopra, and Yann L Cun. Efficient learning of sparse representations with an energy-based model. In *Advances in Neural Information Processing Systems*, pages 1137–1144, 2007.
- [41] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *Proceedings of European Conference on Computer Vision*, pages 102–118, 2016.
- [42] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3234–3243, 2016.
- [43] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of International Conference on Computer Vision*, pages 8049–8057, 2019.
- [44] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, and Kate Saenko. Universal domain adaptation through self supervision. *arXiv preprint arXiv:2002.07953*, 2020.
- [45] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. In *International Conference on Learning Representations*, 2018.
- [46] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3752–3761, 2018.
- [47] Sumit Shekhar, Vishal M Patel, Hien V Nguyen, and Rama Chellappa. Generalized domain-adaptive dictionaries. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 361–368, 2013.
- [48] Weiwei Shi, Yihong Gong, De Cheng, Xiaoyu Tao, and Nanning Zheng. Entropy and orthogonality based deep discriminative feature learning for object recognition. *Pattern Recognition*, 81:71–80, 2018.
- [49] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- [50] Teo Spadotto, Marco Toldo, Umberto Michieli, and Pietro Zanuttigh. Unsupervised domain adaptation with multiple domain discriminators and adaptive self-training. In *Proceedings of International Conference on Pattern Recognition*, 2020.
- [51] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of AAAI Conference on Artificial Intelligence*, pages 2058–2065, 2016.
- [52] Baochen Sun and Kate Saenko. Deep CORAL: correlation alignment for deep domain adaptation. In *Proceedings of European Conference on Computer Vision Workshops*, pages 443–450, 2016.



- [53] Hui Tang, Ke Chen, and Kui Jia. Unsupervised domain adaptation via structurally regularized deep clustering. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 8722–8732, 2020.
- [54] Lei Tian, Yongqiang Tang, Liangchen Hu, Zhida Ren, and Wensheng Zhang. Domain adaptation by class centroid matching and local manifold self-learning. *arXiv preprint arXiv:2003.09391*, 2020.
- [55] Marco Toldo, Andrea Maracani, Umberto Michieli, and Pietro Zanuttigh. Unsupervised domain adaptation in semantic segmentation: a review. *Technologies*, 8(2), 2020.
- [56] Marco Toldo, Umberto Michieli, Gianluca Agresti, and Pietro Zanuttigh. Unsupervised domain adaptation for mobile semantic segmentation based on cycle consistency and feature alignment. *Image and Vision Computing*, 95, 2020.
- [57] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018.
- [58] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schuster, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *Proceedings of International Conference on Computer Vision*, pages 1456–1465, 2019.
- [59] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [60] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- [61] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019.
- [62] Qian Wang and Toby P. Breckon. Unsupervised domain adaptation via structured prediction based selective pseudo-labeling. In *Proceedings of AAAI Conference on Artificial Intelligence*, pages 6243–6250, 2020.
- [63] Sinan Wang, Xinyang Chen, Yunbo Wang, Mingsheng Long, and Jianmin Wang. Progressive adversarial networks for fine-grained domain adaptation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 9210–9219, 2020.
- [64] Wei Wang, Dan Yang, Feiyu Chen, Yunsheng Pang, Sheng Huang, and Yongxin Ge. Clustering with orthogonal autoencoder. *IEEE Access*, 7:62421–62432, 2019.
- [65] Si Wu, Jian Zhong, Wenming Cao, Rui Li, Zhiwen Yu, and Hau-San Wong. Improving domain-specific classification by collaborative learning with adaptation networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5450–5457, 2019.
- [66] Zuxuan Wu, Xin Wang, Joseph E. Gonzalez, Tom Goldstein, and Larry S. Davis. ACE: adapting to changing environments for semantic segmentation. In *Proceedings of International Conference on Computer Vision*, pages 2121–2130, 2019.
- [67] Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *Proceedings of the International Conference on Machine Learning*, pages 5419–5428, 2018.
- [68] Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *Proceedings of International Conference on Computer Vision*, pages 1426–1435, 2019.
- [69] Yanchao Yang and Stefano Soatto. FDA: fourier domain adaptation for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 4084–4094, 2020.
- [70] Fisher Yu, Vladlen Koltun, and Thomas A. Funkhouser. Dilated residual networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 636–644, 2017.
- [71] Heng Zhang, Vishal M Patel, Sumit Shekhar, and Rama Chellappa. Domain adaptive sparse representation-based classification. In *IEEE international conference and workshops on automatic face and gesture recognition (FG)*, pages 1–8, 2015.
- [72] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of International Conference on Computer Vision*, pages 2020–2030, 2017.
- [73] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2881–2890, 2017.
- [74] Yang Zou, Zhiding Yu, Xiaofeng Liu, B.V.K. Vijaya Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of International Conference on Computer Vision*, pages 5982–5991, 2019.
- [75] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of European Conference on Computer Vision*, pages 289–305, 2018.