

Latent Space Regularization for Unsupervised Domain Adaptation in Semantic Segmentation

Francesco Barbato, Marco Toldo, Umberto Michieli, Pietro Zanuttigh

Department of Information Engineering, University of Padova

{barbatofra, toldomarco, umberto.michieli, zanuttigh}@dei.unipd.it

Abstract

Deep convolutional neural networks for semantic segmentation achieve outstanding accuracy, however they also have a couple of major drawbacks: first, they do not generalize well to distributions slightly different from the one of the training data; second, they require a huge amount of labeled data for their optimization. In this paper, we introduce feature-level space-shaping regularization strategies to reduce the domain discrepancy in semantic segmentation. In particular, for this purpose we jointly enforce a clustering objective, a perpendicularity constraint and a norm alignment goal on the feature vectors corresponding to source and target samples. Additionally, we propose a novel measure able to capture the relative efficacy of an adaptation strategy compared to supervised training. We verify the effectiveness of such methods in the autonomous driving setting achieving state-of-the-art results in multiple synthetic-to-real road scenes benchmarks.

1. Introduction

Semantic segmentation is a key tool to address challenging scene understanding problems, like those connected to autonomous driving. Reliable solutions have started to appear with the rise of deep learning: most state-of-the-art approaches exploit Convolutional Neural Networks (CNNs) with an encoder-decoder architecture, starting from the pioneering FCN work [21] up to recent highly performing schemes like PSPNet [53] and DeepLab [4–6]. These architectures typically exploit CNNs originally designed for image classification as feature extractors, removing their tailing fully-connected layers and adding a decoding network to generate full-size segmentation maps. The outstanding results of these models come at the cost of an expensive training process requiring massive amounts of labeled data.

In spite of this, their generalization properties are not always satisfactory. When applied on domains similar but not identical to the training one, the models suffer from significant accuracy degradation: the so-called domain shift issue.

To alleviate this problem, many Unsupervised Domain Adaptation (UDA) solutions have been proposed, exploiting unlabeled samples from the target domain to aid the generalization aptitude of the network. The adaptation can be performed at different stages of the network, *i.e.*, at the input, feature or output level [40]. Deep networks typically solve complex tasks by building some compact latent representations of the inputs, which are representative of the classifier output. These internal representations are extremely meaningful for the subsequent decision process [1, 13]; nevertheless current UDA approaches for semantic segmentation hardly operate at this level due to the high dimensionality of the latent space. We propose a novel strategy working at the less-explored feature level: our aim is to reduce the performance discrepancy by employing latent-space shaping objectives between source and target domains.

First of all, a clustering-based objective forces the feature vectors of each class to be closer to the corresponding prototype centroids. While based on source supervision for prototype estimation, its action is delivered to both source and target representations to achieve class-conditional domain alignment. Then, an additional component enforces the perpendicularity of class prototypes, thus assembling features into well-distanced class clusters and, at the same time, promoting disjoint activation sets between semantic categories. Finally, we account for the fact that, as noticed in [52], feature vectors computed from target domain samples tend to have smaller norms than source domain ones. This latter claim is due to domain-specific features, which the network relies on to solve the source-supervised classification. Yet, those features may be missing in the target domain and, therefore, may lead to a weakened response of neuron activations in target latent representations. To address this issue, we introduce a regularization objective that promotes uniform vector norms across source and target representations, while jointly inducing progressively

Our work was in part supported by the Italian Ministry for Education (MIUR) under the “Departments of Excellence” initiative (Law 232/2016) and by the SID project “Semantic Segmentation in the Wild”.

increased norm values. Furthermore, the inter-class norm alignment has shown to remove distribution biases towards the most frequent classes, whose higher classification confidence is typically accompanied by bigger feature norms.

Since the proposed techniques require to set a strong relationship between predicted segmentation maps and feature representations, we additionally develop a novel strategy to propagate semantic information from the labels to the lower resolution feature space.

2. Related Works

Unsupervised Domain Adaptation is a challenging setting in domain adaptation where only unlabeled samples from the target domain are used. The goal is to limit the performance degradation due to the distribution discrepancy between source and target data (*i.e.*, the *domain shift* [38]), with no supervision in the target domain.

Early techniques focused on whole-image classification [3, 28, 36, 47], while, recently, the domain adaptation field has witnessed a rapid increase in interest, resulting in a multitude of different approaches for various tasks. From a general point of view, we can identify three major categories of works [40], according to the network location in which they act, namely: input-level, feature-level and output-level.

Usually, adaptation at the input level aims to reduce the domain shift by acting directly on input images and trying to match source and target visual appearance (*i.e.*, low level feature distribution), typically using generative adversarial schemes [8, 15, 16, 26, 30, 41]. At the output level, self-training concepts [54, 55] have been explored, where target network predictions in the form of pseudo-labels guide the learning process in a self-supervised manner. Alternatively, some works introduce entropy minimization techniques [7, 48], which force the network to be more confident in the segmentation of target samples, thus mimicking the behavior shown in the source domain. Other approaches [2, 24, 37] further exploit an adversarial discriminator to reduce the perceived discrepancy between segmentation maps produced by source and target domains.

A different line of works operates at the feature level, *e.g.*, by enforcing the extraction of more discriminative features. For example, some works resort to dropout regularization [18, 27, 34] (either channel-wise or point-wise) to push decision boundaries away from high density regions, while others opt for domain adversarial feature alignment [12, 35, 44, 45] either by acting on them directly or by training the network on reconstructed images.

Latent Space Regularization is a family of techniques that can be used to reduce the domain shift and have been applied in many semantic segmentation tasks such as UDA [17, 39], continual learning [25] and few-shot learning [11, 49]. In general, strategies belonging to this class make use of additional constraints imposed on the feature

vectors, effectively reducing the extent of space each of them can occupy. Such reduction has shown to promote more overlap between the source and target distributions, thus reducing the domain shift [42]. In UDA, these techniques are generally applied in class-conditional manner, hence relying on the exclusive supervision of source samples. However, it is reasonable to assume that their effect is reflected also on target samples. These techniques have been exploited in various ways. Different kinds of automatic feature clustering by embedding variations of the K-Means algorithm in the training procedure have been proposed [17, 20, 39, 50]. In [42] the authors further refine this idea by proposing an explicit clustering objective between feature vectors and the appropriate class prototypes. Another work [9] proposes feature-level orthogonality as an alternative for the standard cross-entropy optimization objective in image classification, trying to reduce the number of redundant features extracted by a CNN. Approaches closer to our strategy are [29], where orthogonal class prototypes are used as a medium through which classification is performed in an unsupervised domain adaptation setting, and [51], where an orthogonality constraint over the prototypes is exploited.

3. Problem Setup

In this section we detail our setup. Formally, we denote the input image space as $\mathcal{X} \subset \mathbb{R}^{H \times W \times 3}$ and the associated output label space as $\mathcal{Y} \subset \mathcal{C}^{H \times W}$, where H and W represent the spatial dimensions and \mathcal{C} the set of classes. Given a first training set $\mathcal{T}^s = \{(\mathbf{X}_n^s, \mathbf{Y}_n^s)\}_{n=1}^{N_s}$, where labeled samples $(\mathbf{X}_n^s, \mathbf{Y}_n^s) \in \mathcal{X}^s \times \mathcal{Y}^s$ originate from a supervised source domain, together with a second set of unlabeled input samples $\mathcal{T}^t = \{\mathbf{X}_n^t\}_{n=1}^{N_t}$, from a target domain ($\mathbf{X}_n^t \in \mathcal{X}^t$), our goal is to transfer knowledge on the segmentation task learned on the source domain to the unsupervised target domain (*i.e.*, without any label on the target set). Superscripts s and t specify the domain: source and target, respectively.

We assume that the segmentation network $S = D \circ E$ is based on an encoder-decoder architecture (as most recent approaches for semantic segmentation), *i.e.*, made by the concatenation of two logical blocks: the encoder network E , consisting of the feature extractor, and a decoder network D , which is the actual classifier producing the segmentation map. Moreover, we call the features extracted from a generic input image \mathbf{X} as $E(\mathbf{X}) = \mathbf{F} \in \mathbb{R}_{0+}^{H' \times W' \times K}$, where K denotes the number of channels and $H' \times W'$ denotes the low-dimensional latent spatial resolution. Given the structure of encoder-decoder convolutional segmentation networks, we can assume that each class is mapped to a reference representation in the latent space, that should be as invariant as possible to the domain shift. The techniques that will be introduced in Section 4 try to enforce this by comparing the extracted features with some *prototypes* for

the various classes. In the rest of this section we show how to associate feature vectors to semantic classes and how to compute the prototypes.

Histogram-Aware Downsampling. Since the spatial information of an image is mostly preserved while its content travels through an encoder-decoder network, we can infer a strict relationship between any feature vector and the semantic labeling of the corresponding image region.

Therefore, the first step of the extraction process is to identify a way to propagate the labeling information to latent representations (decimation), preserving the semantic content of the image region (window) associated to each feature vector. Otherwise, the generation of erroneous associations would significantly impair the estimation objective. For this task, we design a non-linear pooling function: instead of computing a simple subsampling (*e.g.*, nearest neighbor), we compute a frequency histogram over the labels of all the pixels in each window. Such histograms are then used to select appropriate classification labels for the downsampled windows, producing feature-level label maps $\{\mathbf{I}_n^{s,t}\}_{n=1}^{N_{s,t}}$. Specifically, the choice is made by selecting the label corresponding to the frequency peak in each window, if such peak is distinctive enough, *i.e.*, if any other peak is smaller than T_h times the biggest one (in a similar fashion to the orientation assignment step in the SIFT feature extractor [22]). Empirically, we set $T_h = 0.5$. A key feature of this technique is its ability to introduce void-class samples when a considered window cannot be assigned to a unique class, *i.e.*, it contains mixed classification labels. This procedure can be naturally extended to pseudo-labels (*i.e.*, network-generated segmentation maps) via a confidence measure over the maps that preserves only reliable predictions. In our case, such measure is computed efficiently by average pooling over the map of output probability peaks and used to mask the raw low-resolution pseudo-labels, *i.e.*, we select only confident labeling with average probability value greater than $T_p = 0.5$, empirically.

Prototype Extraction. Once computed, the feature-level label maps $\{\mathbf{I}_n^{s,t}\}_{n=1}^{N_{s,t}}$ can be used to extract the set \mathcal{F}_c of feature vectors belonging to a generic class $c \in \mathcal{C}$ in a training batch \mathcal{B} :

$$\mathcal{F}_c^{s,t} = \{\mathbf{F}_n^{s,t}[h, w] \in \mathbb{R}_{0+}^K \mid \mathbf{I}_n^{s,t}[h, w] = c, \forall n \in \mathcal{B}\}, \quad (1)$$

where $[h, w]$ denote all possible spatial locations over a feature map, *i.e.*, $0 \leq h < H'$ and $0 \leq w < W'$. Exploiting this definition, we can identify the set of all feature vectors in batch \mathcal{B} as the union $\mathcal{F}^{s,t} = (\bigcup_c \mathcal{F}_c^{s,t}) \cup \mathcal{F}_v^{s,t}$ where $\mathcal{F}_v^{s,t}$ are the sets of void-class samples. The class-wise sets are then used to estimate the per-batch class prototypes on labeled source data by simply computing their centroids:

$$\mathbf{p}_c[i] = \frac{1}{|\mathcal{F}_c^{s,t}|} \sum_{\mathbf{f} \in \mathcal{F}_c^{s,t}} \mathbf{f}[i] \quad \forall i, \quad 1 \leq i \leq K. \quad (2)$$

Finally, to reduce estimation noise and obtain more stable and reliable prototypes, we apply exponential smoothing:

$$\hat{\mathbf{p}}_c = \eta \hat{\mathbf{p}}'_c + (1 - \eta) \mathbf{p}_c. \quad (3)$$

Where $\hat{\mathbf{p}}_c$ and $\hat{\mathbf{p}}'_c$ are the estimates of class c prototype respectively at current and previous optimization steps. We initialized $\hat{\mathbf{p}}_c = \mathbf{0}$ and empirically set $\eta = 0.8$. This strategy allows us to keep track of classes that are not present in the current batch of source samples (in this case we set $\eta = 1$ to propagate the previous estimate), aiding significantly in the unsupervised target tasks.

4. Method

In this section, we provide a detailed description of our approach, which is based on the idea of aiding the standard cross-entropy loss with additional components to enforce the regularization of the latent space. While the source supervised cross-entropy provides task discriminativeness to the model, the additional objectives jointly imposed on source and target representations drive towards feature-level domain invariance, ultimately reducing domain bias. In particular, we add three feature-space shaping constraints to the standard source-supervised cross-entropy loss \mathcal{L}_{CE}^s , whose combined effect can be mathematically expressed by:

$$\mathcal{L} = \mathcal{L}_{CE}^s + \lambda_C \cdot \mathcal{L}_C^{s,t} + \lambda_P \cdot \mathcal{L}_P^s + \lambda_N \cdot \mathcal{L}_N^{s,t}. \quad (4)$$

Here, \mathcal{L}_C is the clustering loss on latent representations (Sec. 4.1), \mathcal{L}_P is the perpendicularity of class prototypes loss (Sec. 4.2) and \mathcal{L}_N is the norm alignment and enhancement loss (Sec. 4.3). For ease of notation, Eq. 4 reports each loss term once (the s, t superscript here indicates the sum of source and target loss instances). For an improved performance and to show that our approach can be applied on top of existing methods, we also extended our objective with the entropy minimization strategy proposed in [7], leading to $\mathcal{L}^+ = \mathcal{L} + \lambda_{EM} \cdot \mathcal{L}_{EM}$.

An overview of the complete approach is reported in Fig. 1.

4.1. Clustering of Latent Representations

The domain shift between source and target data is reflected into a discrepancy in distribution of latent representations from separate domains. Moreover, as the lack of target supervision inherently leads to a bias towards the source domain, it is very likely for the classifier to trace decision boundaries tight around source embeddings, regardless of the disposition of unlabeled target instances.

Thus, the misalignment of class-conditional feature statistic inevitably leads the model towards incorrect classification over target representations, in turn degrading the segmentation accuracy on the target domain. To cope with this issue, we start by introducing a clustering objective over

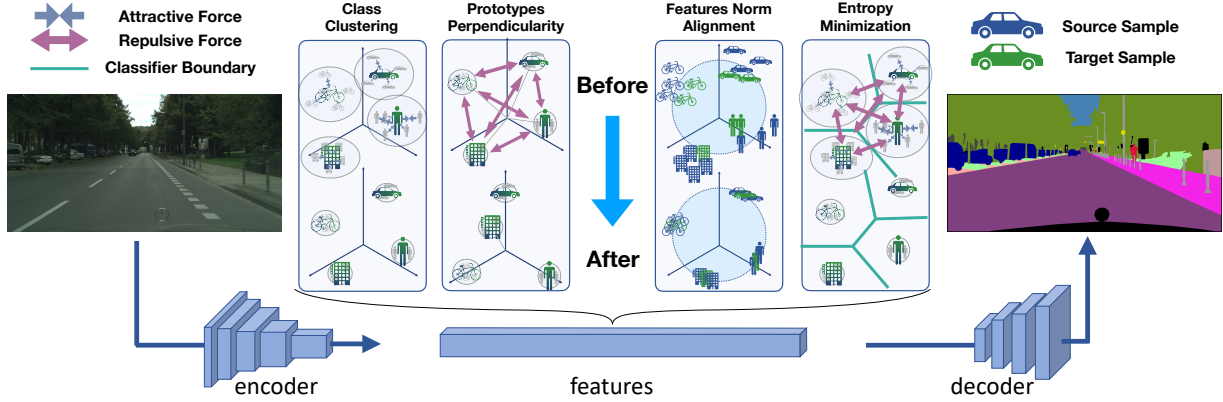


Figure 1: Visual summary of our strategy and of the effect of its application on the feature space vectors. The three proposed space shaping constraint are from left to right: Class Clustering (4.1), Prototypes Perpendicularity (4.2), Norm Alignment and Enhancement (4.3). Furthermore we apply also entropy minimization [7].

the latent space, in order to achieve class-conditional alignment of feature distribution. By exploiting the prototypes introduced in Sec. 3 and forcing source and target feature vectors to tightly assemble around them, we regularize the structure of the latent space, adapting representations into a domain independent class-wise distribution.

Mathematically, we define the clustering objective as:

$$\mathcal{L}_C^{s,t} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \frac{1}{|\mathcal{F}_c^{s,t}|} \sum_{\mathbf{f} \in \mathcal{F}_c^{s,t}} \|\hat{\mathbf{p}}_c - \mathbf{f}\|^2, \quad (5)$$

where $\|\cdot\|$ denotes the L2 norm. This loss has multiple purposes: first, to better cluster representations in the latent space in a supervised manner, thus reducing the probability of erroneous classification. Second, to perform semi-supervised clustering on target samples exploiting network predictions as pseudo-labels (notice that we used only the most confident labeling as detailed in Sec. 3). Finally, to improve prototype estimates, since forcing tighter clusters will result in more stable batch-wise centroids, which will be closer to the moving-averaged prototypes.

4.2. Perpendicularity of Latent Representations

We further enhance the space shaping action induced by the clustering objective by introducing a prototype perpendicularity loss. The idea is to improve the segmentation accuracy by better separating the tight and domain-invariant clusters on both domains. By doing so, we allow the classifiers to increase the margin between decision boundaries and feature clusters, and, consequently, we reduce the likelihood of those boundaries to cross target high-density regions of the feature space (*i.e.*, regions populated by many target samples). Unlike previous works [42], class clusters of latent embeddings are forced to adopt a regular disposition, whilst being spaced out. In fact, we directly encourage

a class-wise orthogonality property, by pushing prototypes to be perpendicular. In this way, not only we increase the distance among class clusters, but we jointly regularize the latent space and encourage channel-wise disjoint activations between different semantic categories.

To quantify perpendicularity in the loss value, we exploit the inner product in the euclidean space and its relationship with the angle θ between two vectors \mathbf{j} and \mathbf{k} , *i.e.*, $\mathbf{j} \cdot \mathbf{k} = \|\mathbf{j}\| \|\mathbf{k}\| \cos \theta$. Minimizing their normalized product is equivalent to maximizing the angle between them, since feature vectors have non-negative values. To capture this, we enforce cross-perpendicularity between any couple of prototypes:

$$\mathcal{L}_P^s = \frac{1}{|\mathcal{C}|(|\mathcal{C}| - 1)} \sum_{c_i, c_j \in \mathcal{C}, i \neq j} \frac{\mathbf{p}_{c_i}}{\|\mathbf{p}_{c_i}\|} \cdot \frac{\mathbf{p}_{c_j}}{\|\mathbf{p}_{c_j}\|}. \quad (6)$$

Eq. 6 computes the sum of the cosines over the set of all couples of non-void classes. We use the per-batch computation of prototypes on source samples \mathbf{p}_c (notice the missing hat on the prototype symbols, see Eq. 2), guaranteeing a stronger gradient flow through the network. In addition, thanks to the tight geometric relation between prototype estimates and feature vectors enforced by $\mathcal{L}_C^{s,t}$, the effect induced by the orthogonality constraint on the prototypes is propagated to the vectors associated to them.

The net result is the application of the shaping action to all feature vectors of each class, thus promoting perpendicularity between all individual components of distinct clusters. The loss seeks to increase the angular distance between latent representations of separate classes, which is achieved when distinct sets of active feature channels are associated to distinct semantic categories. A similar orthogonality constraint has been proposed in [42], but here we directly enforce the perpendicularity property between clusters, whereas in [42] each feature vector is considered inde-

pendently without accounting for its semantic labeling, as the constraint is imposed in an unsupervised fashion.

4.3. Latent Norm Alignment Constraint

The last constraint we propose acts on the norm of source and target feature vectors. In particular, we promote the extraction of latent representations with uniform norm values across domains. Our objective is twofold. First, we aim at increasing the classification confidence during target prediction, similarly to what achieved by adaptation strategies based on entropy minimization over the output space [48]. In particular, recent studies in image classification [52] highlight how the norm of target feature vectors tends towards smaller values than source ones, generally leading to reduced prediction confidence and potentially erroneous classifications. Second, we assist the perpendicularity loss by reducing the number of domain-specific feature channels exploited to perform classification. We argue, in fact, that by forcing the network to produce consistent feature norms, we reduce the number of channel activations switched on for only one of the two domains, as they would cause norm discrepancies. Formally, we define two separate objectives for source and target domains:

$$\mathcal{L}_N^s = \frac{1}{|\mathcal{F}^s|} \sum_{\mathbf{f} \in \mathcal{F}^s} |(\bar{f}_s + \Delta_f) - \|\mathbf{f}\||, \quad (7)$$

$$\mathcal{L}_N^t = \frac{1}{|\mathcal{F}^t|} \sum_{\mathbf{f} \in \mathcal{F}^t} \max(0, (\bar{f}_s + \Delta_f) - \|\mathbf{f}\|), \quad (8)$$

where \bar{f}_s is the mean of the feature vector norms computed from source samples in the previous optimization step and Δ_f dictates the enhancement step (we experimentally tuned it, *e.g.*, $\Delta_f = 0.002$). Feature vectors are pushed towards the same global average norm value, regardless of their labeling. This removes any bias generated by heterogeneous pixel-class distribution in semantic labels, which, for example, would cause the most frequent classes to show larger norm than the average. The source-specific constraint (Eq. 7) forces both the inter-class alignment and enhancement step, *i.e.*, it ensures that norms are progressively increased throughout the training process, towards a common value for all the classes. On the other hand, the target objective (Eq. 8) focuses on domain-alignment, by enforcing the target norms to be similar to the source ones. Furthermore, since target features have typically smaller norms, we do not penalize target norms exceeding the reference value.

5. Implementation Details

Datasets and Setup. We test our model (LSR, Latent Space Regularization) on synthetic-to-real UDA on road-view semantic segmentation. As (synthetic) source domains we employ the *GTAV* [32] and the *SYNTHIA* [33] datasets.

The former is comprised of 24,966 densely labeled, high-resolution (1914×1052 px) images, taken from a video sequence produced by the *GTAV* game engine, while the latter provides 9,500 densely labeled samples with resolution 1280×760 px, produced using the homonym software. As target domain, we choose the *Cityscapes* [10] dataset, which contains 5,000 densely labeled high-resolution (2048×1024 px) images, acquired in European cities.

We train the architecture in a closed-set [40] setup, *i.e.*, source and target class sets coincide. Therefore, we use the 19 and 16 common classes for *GTAV* and *SYNTHIA*, respectively. *GTAV* and *Cityscapes* images are resized during training to 1280×720 px and 1024×512 px, respectively, while *SYNTHIA* samples are kept at the original resolution.

Baseline Model. We used the common [7, 42–44, 48] DeepLabV2 network [4–6], with ResNet101 [14] as the backbone (with $K = 2048$ channels at the last level of the encoder) and stride 8. We pre-train the model on source-only samples with a batch size of 10 for more stable training, using patches of 512×512 px and data augmentation to remove visual biases introduced by the running mean components of batch-normalization layers when full images are employed. As an example, a dark patch on the bottom half of the image will often be interpreted as *road*, while a light patch on the top half will often be interpreted as *sky*, which is not always true (see the random camera angles in the *SYNTHIA* dataset) and preserving such behavior may be detrimental for some applications. The final goal is to reduce color and texture-based biases that could be introduced during training on a single source dataset.

Training Procedure. We optimize the network using SGD with momentum of rate 0.9 and weight decay regularization of 5×10^{-4} . The learning rate was scheduled according to a polynomial decay of power 0.9 starting from 2.5×10^{-4} over 250k steps, following [7]. A subset of the original training set was exploited as validation set for the hyper-parameters search in our loss terms. To reduce overfitting we employ various dataset augmentation strategies: random left-right flip; white point rebalancing $\propto \mathcal{U}([-75, 75])$; color jittering $\propto \mathcal{U}([-25, 25])$ (both applied independently over color channels) and random Gaussian blur [7, 55]. We used one NVIDIA Titan RTX GPU, with batch size of 2 (1 source and 1 target samples), training the network for 27,450 steps (*i.e.*, 10 epochs of the *Cityscapes* [10] dataset) and employing early stopping based on the validation set.

6. Results

In this section, we report the quantitative and qualitative results produced by the proposed approach (LSR), comparing it with several feature-level approaches [19, 42, 44] and with some works, like entropy minimization strategies [7, 48], whose effect on feature distribution is found to

Setup	Configuration	Road	Sidewalk	Building	Wall ¹	Fence ¹	Pole ¹	Traffic Light	Traffic Sign	Vegetation	Terrain	Sky	Person	Rider	Car	Truck	Bus	Train	Motorbike	Bicycle	mIoU	mIoU ¹	mASR	mASR ¹
From GTAV	Target Only	96.5	73.8	88.4	42.2	43.7	40.7	46.1	58.6	88.5	54.9	91.9	68.7	46.2	90.7	68.8	69.9	48.8	47.6	64.5	64.8	-	100	100
	Baseline [42]	71.4	15.3	74.0	21.1	14.4	22.8	33.9	18.6	80.7	20.9	68.5	56.6	27.1	67.4	32.8	5.6	7.7	28.4	33.8	36.9	-	54.0	-
	ASN (feat) [44]	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	8.3	20.6	23.0	39.0	-	56.9	-
	MinEnt [48]	84.4	18.7	80.6	23.8	23.2	28.4	36.9	23.4	83.2	25.2	79.4	59.0	29.9	78.5	33.7	29.6	1.7	29.9	33.6	42.3	-	61.9	-
	SAPNet [19]	88.4	38.7	79.5	<u>29.4</u>	24.7	27.3	32.6	20.4	82.2	32.9	73.3	55.5	26.9	<u>82.4</u>	31.8	41.8	2.4	26.5	24.1	43.2	-	63.1	-
	MaxSquareIW [7]	<u>87.7</u>	25.2	82.9	30.9	<u>24.0</u>	29.0	35.4	24.2	<u>84.2</u>	<u>38.2</u>	79.2	59.0	<u>27.7</u>	<u>79.5</u>	34.6	44.2	<u>7.5</u>	31.1	40.3	45.5	-	62.2	-
	UDA OCE [42]	89.4	30.7	82.1	23.0	<u>22.0</u>	29.2	37.6	31.7	83.9	37.9	78.3	60.7	27.4	84.6	<u>37.6</u>	<u>44.7</u>	7.3	26.0	38.9	45.9	-	<u>67.3</u>	-
	LSR (Ours)	87.7	<u>32.6</u>	<u>82.6</u>	29.1	23.0	28.5	36.1	<u>28.5</u>	84.8	41.8	80.1	<u>59.4</u>	23.8	76.5	38.4	45.8	7.1	28.5	<u>40.1</u>	46.0	-	67.7	-
From SYNTHIA	Baseline [42]	17.7	15.0	74.3	10.1	0.1	25.5	6.3	10.2	75.5	-	77.9	57.1	19.2	31.2	-	31.2	-	10.0	20.1	30.1	34.3	41.7	44.6
	ASN (feat) [44]	62.4	21.9	76.3	-	-	-	11.7	11.4	75.3	-	80.9	53.7	18.5	59.7	-	13.7	-	20.6	24.0	-	40.8	-	52.5
	MinEnt [48]	73.5	29.2	77.1	7.7	<u>0.2</u>	27.0	7.1	11.4	76.7	-	82.1	57.2	21.3	69.4	-	29.2	-	12.9	27.9	38.1	44.2	51.1	56.3
	SAPNet [19]	81.7	33.5	75.9	-	-	-	7.0	6.3	74.8	-	78.9	52.1	<u>21.3</u>	75.7	-	30.6	-	10.8	28.0	-	44.3	-	56.0
	MaxSquareIW [7]	78.9	33.5	75.3	15.0	0.3	<u>27.5</u>	<u>13.1</u>	16.7	73.8	-	77.7	50.4	<u>19.9</u>	66.7	-	36.1	-	13.7	<u>32.1</u>	39.4	45.2	53.8	58.3
	UDA OCE [42]	88.3	42.2	79.1	7.1	<u>0.2</u>	24.4	16.8	16.5	80.0	-	84.3	56.2	15.0	83.5	-	27.2	-	6.3	30.7	<u>41.1</u>	48.2	<u>54.3</u>	60.9
	LSR (Ours)	81.0	<u>36.9</u>	79.5	<u>13.4</u>	<u>0.2</u>	28.7	9.0	16.1	<u>79.1</u>	-	81.7	57.9	21.6	<u>77.2</u>	-	<u>35.3</u>	-	<u>14.2</u>	35.4	41.7	<u>48.1</u>	56.5	61.6

Table 1: Comparison of adaptation strategies in terms of IoU, mIoU and mASR % (Sec. 6). Best in **bold**, runner-up underlined. mIoU¹ and mASR¹ restrict to 13 classes, ignoring the ones with same superscript. We use same setup and codebase of [7, 42].

be similar to ours. An ablation study and a discussion of the effects brought by the proposed loss terms is also included.

Notice that our method is trained end-to-end, thus we can seamlessly add other adaptation techniques, *e.g.*, entropy minimization or adversarial input or output level approaches. To prove such compatibility, we introduce an entropy-minimization loss [7] in our framework.

6.1. Measuring the Adaptation Performances

For a better evaluation, we introduce a novel measure, called mASR (*mean Adapted-to-Supervised Ratio*), to capture the relative performance between an adapted architecture and its target supervised counterpart. We start by computing the ratio between the IoU score of the adapted network for each class $c \in \mathcal{C}$ (IoU_{adapt}^c) and the IoU of supervised training on target data (IoU_{sup}^c), that is a reasonable upper bound estimate. Notice that higher means better, with a value of 1 meaning that the adapted architecture has the same performances of supervised training. Finally we average the scores over all the classes in the dataset:

$$\text{mASR} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \text{ASR}^c, \quad \text{ASR}^c \stackrel{\text{def}}{=} \frac{\text{IoU}_{adapt}^c}{\text{IoU}_{sup}^c}. \quad (9)$$

Differently from other recent measures [23, 31], the mASR measure captures the relative performance between an adapted architecture and its supervised counterpart, allowing an evaluation of adaptation schemes more independent from the overall performances of the network on the various datasets. The per-class ASR values allows to quickly discover the most challenging classes for the adaptation task. Notice how, if compared with the standard mIoU, a key difference is that the contribution of each class is inversely proportional to the capacity of the segmentation model to learn

it in the supervised reference scenario, thus emphasizing the most challenging semantic categories.

6.2. Adaptation from GTAV to Cityscapes

In the *GTA*V \rightarrow *Cityscapes* setup our approach achieves a score of 46.0% mIoU, with a gain of 9.1% compared to the baseline. In addition, LSR outperforms all competitors, with only the very recent works of [42] and [7] able to get close to our result, while there is a quite relevant gap compared to other methods. Such improvement is quite stable across most single-class IoU scores, and is particularly evident in difficult classes, such as *terrain*, where our strategy shows the highest percentage gains. As an index of robustness and performance balance, we use the standard deviation of the per-class IoUs. LSR reduces it by 0.8 compared to the latest state-of-the-art [42] (from 24.8 to 24.0). Furthermore, LSR surpasses the same strategy by 0.4% in terms of mASR, reaching 67.7%, meaning that our approach shows improved accuracy over more challenging classes, thanks to the enhanced latent space regularization.

Some qualitative results are reported in the top half of Fig. 2 and on the paper webpage http://lstm.dei.unipd.it/paper_data/LSR. From visual inspection, we can verify the increased precision of *t. sign*, *t. light*, *pole* and *person* borders in both images. Furthermore, our approach correctly classifies the *wall* in the first image (labeled as *building* by competitors) and the *building* in the second image (confused for *fence* by competitors).

6.3. Adaptation from SYNTHIA to Cityscapes

In the *SYNTHIA* \rightarrow *Cityscapes* setup LSR surpasses all the competitors in the 16-classes configuration, reaching 41.7% of mIoU, with a gap compared to the best competing approach larger than in the previous setting. Once more,

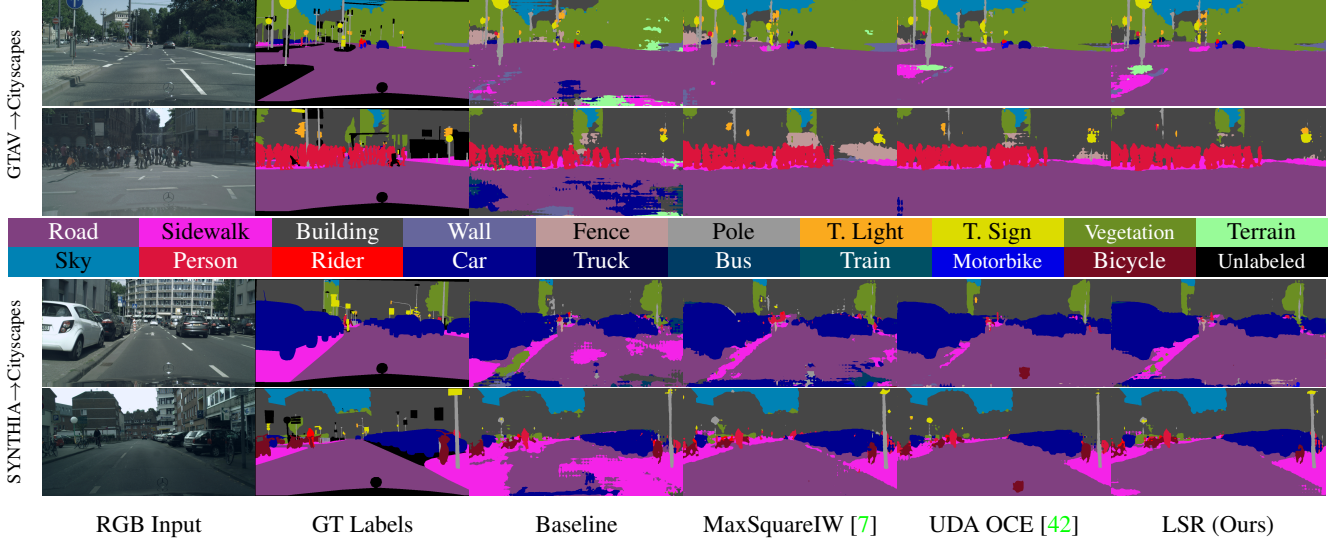


Figure 2: Qualitative results on sample scenes taken from the *Cityscapes* validation split.




Figure 3: (a) Ablation analyses. (b) Image of Fig. 1 downsampled nearest (top) or frequency-aware (bottom).

our method reduces the standard deviation of the IoU distribution (27.7 compared to 29.7 of [42]). At the same time, LSR shows the highest mASR, surpassing the second best approach (on 16 classes) by 2.2%. On the 13-classes setup our strategy outperforms all competitors in terms of mASR with only a marginal loss in mIoU score, confirming our previous claim of better performance balance across classes and reduced gap with respect to supervised learning.

Similarly to the *GTAV* case, the performance gain is visible also from the qualitative results in the bottom half of Fig. 2. The segmentation maps show an overall improvement in the shape of *sidewalk* and *pole*. Our method can correctly detect the *car* and *person* behind the *pole* in the last image, which are missed or wrongly classified by competing strategies (e.g., the *car* is confused as *person* in [7]), and can accurately predict the *t. sign*, missed by some strategies. Furthermore, we remark that LSR can correct the region around the car logo (bottom-part of each figure), which is often confused with *bicycle*, *car* or *bus* by competitors.





Figure 4: Class distribution of segmentation maps downsampled either via histogram-aware or via nearest neighbor.

6.4. Ablation Studies

In this section, we evaluate the impact of each constraint on the final accuracy. Quantitative results are reported in Fig. 3a, where we evaluate our strategy by removing each constraint independently and evaluating the impact on the final accuracy. In particular, we show how the absence of each of our losses reduces the final performance by a minimum of 0.8% mIoU and an average of 1% mIoU. Each module brings a significant improvement in terms of accuracy and all the components are needed for the best results. Concerning the novel downsampling scheme (Sec. 3), the goal of the proposed frequency-aware setup is to label only feature locations with a clear class assignment. This aims to reduce cross-talk between neighboring features of different classes, thus improving class discriminativeness at the latent space. We can observe this phenomenon in Fig. 3b, where the label map downsampled via our frequency-aware scheme (bottom) marks some features close to the edges of objects as *unlabeled*. This is confirmed by the class distribution of the downsampled segmentation maps (i.e., to match the spatial resolution of the feature level), reported in Fig. 4 for both the histogram-aware scheme (ours) or the standard



(a) Baseline. (b) LSR.

Figure 5: Features distribution before and after adaptation.

nearest neighbors one. In particular, the histogram-aware scheme generally seldom preserves small classes, promoting *unlabeled* classification when discrimination between classes is uncertain.

6.5. Analysis of the Latent Space Regularization

Norm Alignment. We analyze the effect of the norm alignment constraint in Fig. 5, where we show a plot of some feature vectors after projecting them to a 2D space for better visualization. This and the subsequent plots were produced using a balanced subset of feature vectors (350 vectors per class) extracted from the *Cityscapes* validation set for a fair comparative analysis across the classes.

In Fig. 5 the norm of each vector is represented in the radial axis in log scale, the inner angle between any point and its prototype’s direction is represented in the angular axis. The original direction in the high dimensional space of each centroid is ignored, as a meaningful representation would be very difficult to achieve in 2D. Instead, we assign to each centroid a reference angle (as shown by the colored lines) and plot the associated feature vectors centered on it. The plots also reports a confidence interval for the global average norm, to highlight how the proposed norm alignment constraint (Sec. 4.3) effectively promotes uniform norm values: the dark gray line represents the median of the distribution and the shaded gray area represents the 95% confidence interval. The effect of our space shaping strategy is clearly visible, in that the norms align very tightly around the global mean value (smaller shaded region around the unique gray line). The distribution of the points around each prototype is shrunk, thanks to the clustering objective, while the centroids themselves are pushed away from each other, thanks to the perpendicularity constraint (not visible from this plot, but appreciable in Fig. 6 as we discuss next).

Perpendicularity. To analyze the effect of the perpendicularity constraint, Fig. 6 shows the distribution of the average inner angle between a prototype’s direction and the direction of each other prototype. Ideally, we aim at producing as perpendicular prototypes as possible, in order to reduce the overlap of different semantic classes over feature channels (*i.e.*, cross-talk). The red dashed line at 90 degrees shows the target value for perpendicularity, which is also the





Figure 6: Average inter-prototype angle, comparison between Source Only and LSR.



(a) Baseline. (b) LSR.

Figure 7: t-SNE plots comparing feature vectors distributions before and after adaptation. Points are color coded according to the legend of Fig. 2.

upper bound, as our feature vectors have all non-negative coordinates. The figure shows that our strategy leads to an average increase of more than 5 degrees.

Clustering. Finally, we analyze our clustering objective by means of a t-SNE [46] embedding produced on the normalized features (to remove the norm information, enhancing the angular one) and we report it in Fig. 7. Our strategy increases significantly the cluster separation in the high dimensional space and the spacing between clusters belonging to different classes. As a side effect, this also reduces the probability of confusing visually similar classes (*e.g.*, the *truck* class with the *bus* and *train* ones).

7. Conclusions

In this work, we proposed a new set of latent-space regularization techniques to address the domain shift in an unsupervised fashion. We achieve domain invariance by means of multiple feature space shaping constraints: namely, class clustering, class perpendicularity and norm alignment. Our constraints can be flawlessly applied on top of existing frameworks as they are separate modules trained end-to-end. We achieved state-of-the-art results in feature-level adaptation on two commonly used benchmarks paving the way to employment of a new family of feature-level techniques to enhance discrimination ability of deep neural networks. Future research will concern the design of novel feature-level techniques, the analysis of the proposed adaptation strategies on source accuracy and the evaluation of their generalization ability to different tasks.

References

- [1] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [2] Matteo Basettoni, Umberto Michieli, Gianluca Agresti, and Pietro Zanuttigh. Unsupervised Domain Adaptation for Semantic Segmentation of Urban Scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1211–1220, 2019.
- [3] Karsten M Borgwardt, Arthur Gretton, Malte J Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57, 2006.
- [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 the 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 the International Conference on Computer Vision*, pages 2090–2099, 2019.
- [8] 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 the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1791–1800, 2019.
- [9] Hongjun Choi, Anirudh Som, and Pavan Turaga. Role of orthogonality constraints in improving properties of deep networks for image classification. *arXiv preprint arXiv:2009.10762*, 2020.
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [11] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *Proceedings of the British Machine Vision Conference*, volume 3, 2018.
- [12] 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 the International Conference on Computer Vision*, pages 982–991, 2019.
- [13] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 580–587, 2014.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [15] 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.
- [16] 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.
- [17] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G. Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4893–4902, 2019.
- [18] Seungmin Lee, Dongwan Kim, Namil Kim, and Seong-Gyun Jeong. Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In *Proceedings of the International Conference on Computer Vision*, pages 91–100, 2019.
- [19] 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 the European Conference on Computer Vision*, 2020.
- [20] 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 the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2975–2984, 2019.
- [21] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [22] D. G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1150–1157 vol.2, 1999.
- [23] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. Hota: A higher order metric for evaluating multi-object tracking. *International Journal of Computer Vision*, 129(2):548–578, 2021.
- [24] Umberto Michieli, Matteo Basettoni, 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.
- [25] Umberto Michieli and Pietro Zanuttigh. Continual semantic segmentation via repulsion-attraction of sparse and disentangled latent representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

- [26] Zak Murez, Soheil Kolouri, David J. Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4500–4509, 2018.
- [27] Sungrae Park, Jun-Keon Park, Su-Jin Shin, and Il-Chul Moon. Adversarial dropout for supervised and semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3917–3924, 2018.
- [28] Deepak Pathak, Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional multi-class multiple instance learning. *arXiv preprint arXiv:1412.7144*, 2014.
- [29] Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8004–8013, 2018.
- [30] 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 the Winter Conference on Applications of Computer Vision*, pages 2990–2998, 2020.
- [31] Hamid Rezaatoughi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 658–666, 2019.
- [32] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *Proceedings of the European Conference on Computer Vision*, pages 102–118, 2016.
- [33] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3234–3243, 2016.
- [34] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. In *Proceedings of the International Conference on Learning Representations*, 2018.
- [35] 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 the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3752–3761, 2018.
- [36] Jamie Shotton, Matthew Johnson, and Roberto Cipolla. Semantic textron forests for image categorization and segmentation. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2008.
- [37] Teo Spadotto, Marco Toldo, Umberto Michieli, and Pietro Zanuttigh. Unsupervised domain adaptation with multiple domain discriminators and adaptive self-training. In *Proceedings of the International Conference on Pattern Recognition*, 2020.
- [38] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2058–2065, 2016.
- [39] 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.
- [40] Marco Toldo, Andrea Maracani, Umberto Michieli, and Pietro Zanuttigh. Unsupervised domain adaptation in semantic segmentation: a review. *Technologies*, 8(2), 2020.
- [41] 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.
- [42] Marco Toldo, Umberto Michieli, and Pietro Zanuttigh. Unsupervised domain adaptation in semantic segmentation via orthogonal and clustered embeddings. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1358–1368, 2021.
- [43] Wilhelm Tranheden, Viktor Olsson, Juliano Pinto, and Lennart Svensson. Dacs: Domain adaptation via cross-domain mixed sampling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1379–1389, 2021.
- [44] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018.
- [45] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *Proceedings of the International Conference on Computer Vision*, pages 1456–1465, 2019.
- [46] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [47] Alexander Vezhnevets and Joachim M Buhmann. Towards weakly supervised semantic segmentation by means of multiple instance and multitask learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3249–3256. IEEE, 2010.
- [48] 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 the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019.
- [49] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the International Conference on Computer Vision*, pages 9197–9206, 2019.
- [50] Qian Wang and Toby P. Breckon. Unsupervised domain adaptation via structured prediction based selective pseudo-labeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6243–6250, 2020.
- [51] 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 *Proceed-*

ings of the AAAI Conference on Artificial Intelligence, pages 5450–5457, 2019.

- [52] 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 the International Conference on Computer Vision*, pages 1426–1435, 2019.
- [53] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2881–2890, 2017.
- [54] Yang Zou, Zhiding Yu, Xiaofeng Liu, B.V.K. Vijaya Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the International Conference on Computer Vision*, pages 5982–5991, 2019.
- [55] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision*, pages 289–305, 2018.