

Adapting Segmentation Networks to New Domains by Disentangling Latent Representations

Francesco Barbato, *Student Member, IEEE*, Umberto Michieli, *Graduate Student Member, IEEE*,
Marco Toldo and Pietro Zanuttigh, *Member, IEEE*.

Abstract—Deep learning models achieve outstanding accuracy in semantic segmentation, however they require a huge amount of labeled data for their optimization. Hence, domain adaptation approaches have come into play to transfer knowledge acquired on a label-abundant source domain to a related label-scarce target domain. However, such models do not generalize well to data with statistical properties not perfectly matching the ones of the training samples. In this work, we design and carefully analyze multiple latent space-shaping regularization strategies that work in conjunction to reduce the domain discrepancy in semantic segmentation. In particular, we devise a feature clustering strategy to increase domain alignment, a feature perpendicularity constraint to space apart feature belonging to different semantic classes, including those not present in the current batch, and a feature norm alignment strategy to separate active and inactive channels. Additionally, we propose a novel performance metric to capture the relative efficacy of an adaptation strategy compared to supervised training. We verify the effectiveness of our framework in synthetic-to-real and real-to-real adaptation scenarios, outperforming previous state-of-the-art methods on multiple road scenes benchmarks and using different backbones.

Index Terms—Domain Adaptation, Semantic Segmentation, Latent Space, Clustering.

I. INTRODUCTION

Nowadays, semantic segmentation is commonly tackled with deep convolutional neural networks (DCNNs): recent approaches such as FCN [1], U-Net [2], PSPNet [3] and DeepLab [4]–[6] have achieved outstanding results in image understanding tasks, provided that a sufficiently large number of labeled examples are available from the current input domain distribution. On the other side, the annotation of thousands of input images is highly expensive, time-consuming, error-prone and, possibly, worthless, since the test data can show a domain shift with respect to the labeled samples available at training time. Hence, recently, a strong requirement emerged: namely, of being able to train DCNNs with a combination of labeled source samples (*e.g.*, synthetic) and unlabeled target samples (*e.g.*, real-world), with the aim of getting high performances on data following the target distribution. The need for large quantities of labeled target data is superseded by data coming from a source domain where samples are abundantly available and annotations are much faster and cheaper to generate.

F. Barbato, U. Michieli, M. Toldo and P. Zanuttigh are with University of Padova. Corresponding e-mail: umberto.michieli@dei.unipd.it

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

Unfortunately, DCNNs are prone to failure when they are shown an input domain distribution other than the training one (*domain shift* phenomenon). To alleviate this problem, many Unsupervised Domain Adaptation (UDA) solutions have been proposed adapting the networks at different stages, *i.e.*, at the input, feature or output level [7].

Deep learning models for semantic segmentation are mostly based on encoder-decoder architectures, *i.e.*, they build some compact latent representations of the inputs, which are highly correlated with the classifier output. As such, they are used in the subsequent decision process [8], [9] that reconstructs the full resolution segmentation map. Nevertheless, current UDA approaches for semantic segmentation hardly operate on the internal (latent) feature space due to its high dimensionality. In this paper, we propose a novel set of strategies working on the feature level building on top of our previous conference work [10]. By employing latent-space shaping objective, our aims are to promote class-aware features extraction and features invariance between source and target domains.

First of all, a clustering objective groups feature vectors of each class tightly around their prototypical representation. Second, a perpendicularity constraint over the class prototypes promotes disjoint filter activation sets across different semantic categories. Finally, a regularization-based norm alignment objective promotes uniform vector norms across source and target representations, while jointly inducing progressively increased norm values. This, together with the perpendicularity constraint, is able to reduce the entropy associated with the feature vector channel activations.

We remark that the proposed techniques require the generation of accurate class prototypes and the imposition of a strong relationship between predicted segmentation maps and feature representations. Hence, we additionally develop a novel strategy to propagate semantic information from the labels to the lower-resolution feature space (annotations downsampling).

This paper moves from our previous work [10], which already achieved state-of-the-art results on feature-level UDA in semantic segmentation. Compared to the conference version, this journal extension introduces several novel contributions.

First of all the computation of prototypes and feature vector extraction were refined. The first now considers the prototype trajectory evolution for a better estimation (Sec. III-B), while the second exploits target information to reduce the domain shift (Sec. III-C), additionally a class-weighting scheme is used in the source supervision (Sec. III-A).

Then, each of the three proposed space-shaping constraints was improved and additional ablation studies were performed

both for the approach (Sec. VIII) and for the evaluation metric (Sec. VI). In particular, the clustering objective was modified to be more resilient to outliers (Sec. IV-A); the perpendicularity constraint now accounts for classes not present in the current batch (Sec. IV-B); the norm augmentation now ignores low-activated channels (Sec. IV-C).

Finally, extensive experiments were conducted on many different scenarios. The results are reported on 4 backbones and 6 setups (2 synthetic-to-real and 4 real-to-real). Additional results using the unlabeled Cityscapes coarse set [11] are reported, showing significant performance gains (see Table I).

II. RELATED WORKS

Unsupervised Domain Adaptation consists in transferring knowledge extracted from a label-rich source domain to a completely unlabeled target domain. The ultimate objective is to address the performance degradation caused by domain shift, which negatively affects the generalization capabilities of deep neural networks. The problem was initially studied for the classification task, but recently many works dealt with the unsupervised adaptation problem in relation to semantic segmentation. Although several methods have been proposed to tackle the adaptation task, they all share an underlying search for a form of domain distribution alignment over some representation space. Some methods pursue distribution matching inside the input image space via style transfer or image generation techniques, others aims at bridging the statistical gap between source and target representations produced by the task model, whether manipulating some output representations, or operating inside a latent feature space [7].

Input-space adaptation has been commonly addressed resorting to image-to-image translation [12]–[17]. By transferring visual attributes across source and target samples, domain invariance is achieved in terms of visual appearance. Source supervision can thus be safely exploited in the shared image space, retaining consistent accuracy on source and target data.

As concerns feature and output-space adaptation, adversarial learning has been largely employed to bridge the statistical domain gap [18]–[24]. With the help of a domain discriminator, the task network is forced to provide statistically indistinguishable source and target representations, typically drawn from a latent feature space [18]–[20] or in the form of probability maps at the output of the segmentation pipeline [20]–[24]. More recently, some works focusing on feature-level regularization have been proposed [10], [25]. In [25] a class-conditional domain alignment is achieved by means of a discriminative clustering module, paired with orthogonality constraints to enhance class separability. The approach of [10] relies on conditional clustering adaptation, enhanced by a perpendicularity objective over class prototypical representations and a novel norm alignment loss to improve class separability at the latent space. As an alternative form of feature-level adaptation, dropout regularization has been explored [26]–[28]; decision boundaries are pushed away from target high density regions in the latent space without direct supervision, where proximity to decision boundaries of latent representations is shown to be related to the prediction variability of the classifier.

Output-space adaptation has been further pursued resorting to self-training [29], [30], where target network predictions in the form of pseudo-labels guide the learning process in a self-supervised manner. Self-supervision has been proposed in a curriculum learning fashion as well [31], [32]. First, simple tasks that are less sensitive to domain shift are solved, by inferring some useful properties related to the target domain. Then, the extracted information is exploited to address more complex learning tasks (*e.g.*, semantic segmentation). Alternatively, some works introduce entropy minimization techniques [33], [34], which force more confident network predictions over target data, thus encouraging the behavior shown in the supervised source domain.

Latent Space Regularization has been shown to ease the semantic segmentation tasks in different settings, such as UDA [35], [36], continual learning [37] and few-shot learning [38], [39]. The idea is to embed additional constraints on feature representations during the training process, enforcing a regular semantic structure on latent spaces of the deep neural classifier. In UDA, where target semantic supervision is missing, regularization can be applied in class-conditional manner by relying on the exclusive supervision of source samples, while indirectly propagating its effect to target representations as well. Such improved regularity has, in fact, shown to promote generalization properties, leading to statistical alignment between the source and target distributions when regularization is jointly applied over both domains [10], [25].

A multitude of feature clustering techniques based on the K-Means algorithm have been proposed [35], [36], [40], [41] to address the adaptation task. Those works are mainly focused on image classification and resort to a projection to a more easily-tractable lower-dimensional latent space where to perform pseudo labeling of the original target representations extracted by the task model [36], [40], [41]. In [10], [25] the idea is further refined and applied to semantic segmentation by proposing an explicit clustering objective paired with orthogonality constraints to force feature vectors to cluster around the respective class prototypes. Feature-level orthogonality has been also explored in [42] to limit the redundancy of the information encoded in feature representations. Approaches closer to our strategy are [43], [44], where UDA is promoted via an orthogonality objective over class prototypes. Nonetheless, [42]–[44] all limit their focus to the image classification task.

III. PROBLEM SETUP

In this section we overview our setup, detailing the mathematical notation used throughout the paper. We start by denoting 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. Furthermore, we assume to have a training set $\mathcal{T} = \mathcal{T}^s \cup \mathcal{T}^t$, where $\mathcal{T}^s = \{(\mathbf{X}_n^s, \mathbf{Y}_n^s)\}_{n=1}^{N_s}$ contains labeled samples $(\mathbf{X}_n^s, \mathbf{Y}_n^s) \in \mathcal{X}^s \times \mathcal{Y}^s$ originated from a supervised source domain, while a second set of unlabeled input samples $\mathcal{T}^t = \{\mathbf{X}_n^t\}_{n=1}^{N_t}$ are drawn from a target domain $(\mathbf{X}_n^t \in \mathcal{X}^t)$. We transfer the knowledge of semantic segmentation 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.

As done by most recent approaches for semantic segmentation, we assume a task model $S = D \circ E$ based on an encoder-decoder architecture *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. We denote 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 refers to the number of channels and $H' \times W'$ to the low-dimensional latent space 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 Sec. IV enforces this goal by comparing the extracted features with some *prototypes* of the various classes. In the remainder of this section we show how to associate feature vectors to semantic classes and how to compute the prototypes.

A. Weighted 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 (*i.e.*, the vector of features associated to a single spatial location within the feature tensor) 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 weighted frequency histogram over the labels of all the pixels in the window corresponding to a low-resolution feature location. The weights are inversely proportional to the class-frequency in the source training dataset. Such histograms are then used to select appropriate classification labels for the downsampled windows, producing source feature-level label maps $\{\mathbf{I}_n^s\}_{n=1}^{N_s}$. For what concerns the computation of the target counterparts ($\{\mathbf{I}_n^t\}_{n=1}^{N_t}$) see Sec. III-C and we remark that each $\mathbf{I}_n^{s,t} \in \mathcal{C}^{H' \times W'}$. Specifically, the choice is made by selecting the label corresponding to the frequency peak in each window, only 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 [45]). 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.

B. 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 $\mathcal{F}_c^{s,t}$ 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:

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

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), allowing for a more robust prototype estimation.

C. Feature pseudo-labeling. While the histogram strategy can be seamlessly extended to be used with pseudo-labels (*i.e.*, network estimates for the unlabeled target samples, as was our strategy in the previous work [10]), this approach can introduce instability in the training procedure. To avoid such issue, we devise a novel way of extracting the target feature-level label maps $\{\mathbf{I}_n^t\}_{n=1}^{N_t}$.

Our strategy exploits the euclidean distance in the latent space, computing a clustering of the feature vectors around their prototype (see Fig. 1). More in detail, we compute an initial classification exploiting the prototypes computed over the source labeled data, which, due to the domain shift, will not be adequately representative of the target distribution:

$$\begin{aligned} \tilde{\mathcal{F}}_c^t &= \{\mathbf{F}_n^t[h, w] \text{ if} \\ &\quad \sigma_c(-\|\mathbf{F}_n^t[h, w] - \hat{\mathbf{p}}_c\|) > 0.5 \forall n \in \mathcal{B}\} \\ \mathbf{p}_c^t[i] &= \frac{1}{|\tilde{\mathcal{F}}_c^t|} \sum_{\mathbf{f} \in \tilde{\mathcal{F}}_c^t} \mathbf{f}[i] \quad \forall i, 1 \leq i \leq K. \end{aligned} \quad (4)$$

Where $\sigma_c(\cdot)$ is the softmax function computed over the classes. Then, we refine the classification keeping only those vectors that have a high classification confidence according to a probability distribution attained through a softmax function:

$$\mathbf{I}_n^t[h, w] = \begin{cases} c & \sigma_c(-\|\mathbf{F}_n^t[h, w] - \mathbf{p}_c^t\|) > 0.5 \\ \text{void} & \text{otherwise} \end{cases} \quad (5)$$

IV. METHOD

In this section, we provide a detailed description of our approach, highlighting the key differences with respect to our previous work. Our investigation moves from the fact that the discriminative effect acquired by the model with the source supervised cross-entropy objective may not be propagated to the target domain due to the distribution shift. To tackle such problem, in [10] we proposed to use additional space-shaping

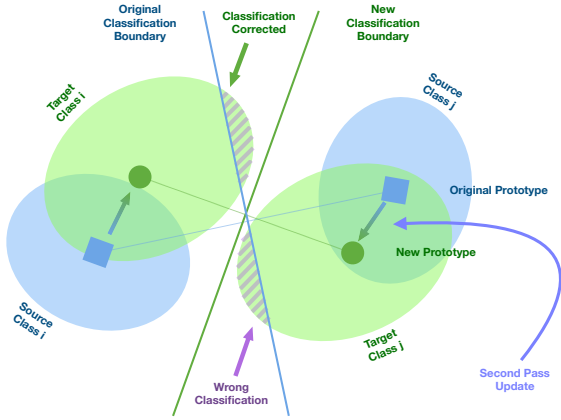


Fig. 1: Visual representation of our two-pass feature vector classification strategy. The initial source-based classification (in blue) can lead to erroneously classified target samples (purple shaded areas). This problem is tackled by computing target prototypes as the centroids of the partitioned vectors (notice the shift compared to the original source prototype), these prototypes are used as new classification centers (green boundary), producing a correct segmentation.

objectives to increase the network generalization capability, therefore improving robustness to distribution shifts from the original source training data. In particular, we added three feature-space shaping constraints to the standard source-supervision (\mathcal{L}_{CE}^s), whose combined effect can be mathematically expressed by:

$$\mathcal{L} = \mathcal{L}_{CE}^s + \lambda_C \cdot \mathcal{L}_C^{s,t} + \lambda_P \cdot \mathcal{L}_P^s + \lambda_N \cdot \mathcal{L}_N^{s,t}. \quad (6)$$

Here, \mathcal{L}_C represents the clustering objective acting on the feature vectors (Sec. IV-A), \mathcal{L}_P the perpendicularity constraint applied to class prototypes (Sec. IV-B) and \mathcal{L}_N the norm alignment goal (Sec. IV-C). For ease of notation, Eq. 6 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 [33], leading to $\mathcal{L}^+ = \mathcal{L} + \lambda_{EM} \cdot \mathcal{L}_{EM}$. By doing so, we also show that our space-shaping objectives provide a different and complementary effect on the feature vectors when compared to the entropy minimization constraint. An overview of the complete approach is reported in Fig. 2.

A. Clustering of Latent Representations

Due to the domain shift between source and target domains, feature vectors originating from the two distributions are misaligned. This inevitably causes some incorrect classifications of target representations, in turn degrading the segmentation accuracy in the target domain. We introduce our first loss, a clustering objective over the latent space, to mitigate this problem, seeking for class-conditional alignment of feature distribution. We do so by exploiting the prototypical representations introduced in Sec. III and forcing source and target

feature vectors to tightly assemble around them, regularizing the structure of the latent space and adapting representations into a common class-wise distribution.

Differently from the previous work, we define the clustering objective as the L1 distance between feature vectors and their associated class-prototype. This results in a more stable training evolution and lower error rate in clustering, thanks to the outlier-rejecting properties of the L1 norm. In particular, due to the quadratic nature of MSE, outliers with distances greater than 1 have a strong push towards the clusters even when they should not. On the other hand, the L1 loss is stronger than MSE for close samples, which are more representative of each class, and is significantly gentler than L2 for distant outliers. The loss can be expressed mathematically as:

$$\mathcal{L}_C^{s,t} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \frac{1}{|\mathcal{F}_c^{s,t}|} \sum_{\mathbf{f} \in \mathcal{F}_c^{s,t}} \frac{1}{K} \sum_{k=1}^K |\hat{\mathbf{p}}_c[k] - \mathbf{f}[k]|, \quad (7)$$

This loss has multiple purposes: first, to better cluster representations in the latent space in a supervised manner, thus reducing the probability of erroneous network classification. Second, to perform self-supervised clustering on target samples exploiting our two-pass pseudo-labeling strategy (see Sec. III-C). 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.

B. Perpendicularity of Latent Representations

To further enhance the space-shaping action of the clustering objective, we introduce 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). We directly encourage a class-wise orthogonality property, not only increasing the distance among class clusters, but also encouraging channel-wise disjoint activations between different semantic categories.

To account for perpendicularity in the loss, 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:

$$\mathcal{L}_P^s = \frac{1}{|\mathcal{C}|(|\mathcal{C}| - 1)} \sum_{c_i, c_j \in \mathcal{C}, i \neq j} \frac{\hat{\mathbf{p}}_{c_i}}{\|\hat{\mathbf{p}}_{c_i}\|} \cdot \frac{\hat{\mathbf{p}}_{c_j}}{\|\hat{\mathbf{p}}_{c_j}\|}. \quad (8)$$

Eq. (8) computes the sum of the cosines over the set of all couples of non-void classes. 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

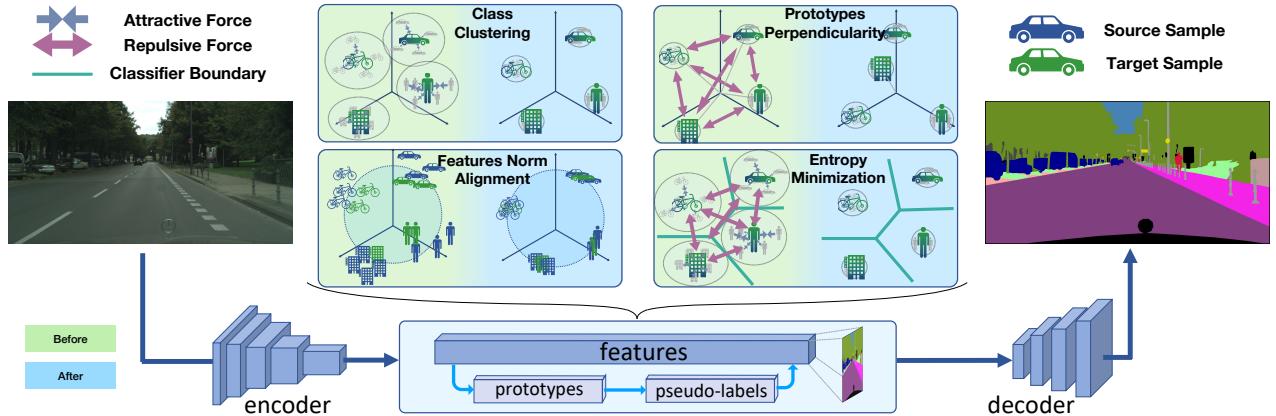


Fig. 2: Visual summary of our strategy. Features are associated to semantic classes and prototypes are computed from them (III). The three proposed space shaping constraint are: Class Clustering (IV-A), Prototypes Perpendicularity (IV-B), Norm Alignment and Enhancement (IV-C). Furthermore we apply also entropy minimization [33].

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.

In contrast to our previous paper [10], we compute the loss on the exponentially smoothed version of the prototypes (*i.e.*, from Eq. 3). This guarantees that the space will be more evenly occupied by the classes, since all directions are considered in the computation of the loss, instead of considering only the ones in the current batch.

C. 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 [34]. 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. Moreover, to reduce the possible decrease in norm value during the alignment process, we introduce a regularization term that promotes norm increase. Differently from [10], here the norm objective is encoded as a relative difference with a regularization term inversely proportional to the norm value. This allows to obtain a value-independent loss where norm values higher than the target are less discouraged. Moreover, we introduce a *norm filtering* strategy to reduce the negative effects a careless increase in norm could imply. In particular, we suppress low channel activations, stopping the gradient flow through them and preventing the norm alignment procedure to increase their value, in contrast to what source supervision indicates.

Formally, we define the loss term as:

$$\mathcal{L}_N^{s,t} = \frac{1}{|\mathcal{F}_*^{s,t}|} \sum_{\mathbf{f} \in \mathcal{F}_*^{s,t}} \frac{|(\bar{f}_s + \Delta_f) - \|\mathbf{f}\||}{\bar{f}_s}, \quad (9)$$

where \bar{f}_s is the mean of the feature vector norms computed from source samples in the previous optimization step, Δ_f dictates the regularization strength (experimentally tuned to 0.1) and $\mathcal{F}_*^{s,t}$ is a thresholded version of $\mathcal{F}^{s,t}$ where we set to 0 the low-activated channels of each feature vector, stopping the gradient propagation:

$$\mathcal{F}_*^{s,t} = \{\phi(\mathbf{f}) \mid \forall \mathbf{f} \in \mathcal{F}^{s,t}\},$$

$$\phi(\mathbf{f})_i = \begin{cases} \mathbf{f}_i & \mathbf{f}_i \geq \frac{1}{K} \sum_{j=1}^K \mathbf{f}_j, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

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 constraint of Eq. 9 forces the inter-class alignment step, *i.e.*, it ensures that norms are progressively aligned throughout the training process towards a common value for all the classes, while guaranteeing the value does not decrease significantly. In other words, the goal value is computed on the source samples and it is the same for both datasets. An additional benefit of rescaling the loss by the norm target is that the loss gradients will be limited in magnitude and, therefore, more stable.

V. IMPLEMENTATION DETAILS

Datasets and Setup. We test our model (LSR⁺, Latent Space Regularization) on synthetic-to-real and real-to-real UDA on road scenes semantic segmentation. As (synthetic) source domains we employ the *GTAV* [46] and the *SYNTIA* [47] 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* [11] dataset, which contains 5,000 densely labeled high-resolution (2048×1024 px) images and an additional set of 20,000 coarsely labeled samples, acquired in European cities. When considering only unlabeled samples the two versions are equivalent and can be merged (obtaining a dataset we refer as *CS-full*) increasing the adaptation process (as we show in Table I). In the real-to-real setup we used the Cross-City benchmark, where the *Cityscapes* dataset takes the role of source domain, while the Cross-City dataset [48] takes the role of target. Such dataset is comprised of 12,800 ($4 \times 3,200$) high resolution (2048×1024 px) images taken in four major cities (*Rome, Rio, Tokyo, Taipei*).

We train the architecture in a closed-set [7] setup, *i.e.*, source and target class sets coincide. Therefore, we use the 19, 16, 13 common classes for *GTAV*, *SYNTHIA* and *Cross-City*, respectively. *GTAV*, *Cityscapes* and *Cross-City* images are resized during training to 1280×720 px, 1280×640 px and 1280×640 px, respectively, while *SYNTHIA* samples are kept at the original resolution.

Baseline Model. We used the common [20], [25], [33], [34], [49] DeepLabV2 network [4]–[6], with ResNet101 [50] as the backbone (with $K = 2048$ channels at the last level of the encoder) and stride 8. We pre-train the model following the same procedure as our previous work [10], and employing the same data augmentation techniques used during adaptation.

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 follows a polynomial decay of power 0.9 starting from 2.5×10^{-4} over $250k$ steps, following [33]. 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 re-balancing $\propto \mathcal{U}([-75, 75])$; color jittering $\propto \mathcal{U}([-25, 25])$ (both applied independently over color channels) and random Gaussian blur [29], [33]. We used one NVIDIA Titan RTX GPU, with batch size of 2 (1 source and 1 target samples), training the network for 24,750 steps (*i.e.*, 10 epochs of the *Cityscapes*-fine train split) and employing early stopping based on the validation set.

The code developed for this work is publicly available at the following link: <https://github.com/LTTM/LSR>.

VI. MEAN ADAPTED-TO-SUPERVISED RATIO METRIC

In this section we introduce a novel measure, called mASR (*mean Adapted-to-Supervised Ratio*), in order to better evaluate the domain adaptation task than allowed by the usual mIoU.

The idea behind the new metric sparks from realizing that the mIoU is missing a key component to evaluate an adaptation method: *i.e.*, it does not account for the starting accuracy for the different classes in supervised training. In particular, the objective of domain adaptation is to transfer the knowledge learned on a source dataset to a target one, trying to get as close as possible to the results attainable




Fig. 3: mASR score as a function of the injected noise intensity.

through supervised learning on the target domain. We design mASR to capture the relative performance between an adapted architecture and its target supervised counterpart, which we identify as a reasonable upper bound. Therefore mASR focuses less on the absolute-term performance and more on the relative accuracy obtained by an adapted architecture when compared to traditional supervised training.

We compare the per-class IoU score of the adapted network for each $c \in \mathcal{C}$ (IoU_{adapt}^c) with the results of supervised training on target data (IoU_{sup}^c) and we compute mASR by:

$$\text{mASR} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \text{ASR}^c, \quad \text{ASR}^c \stackrel{\text{def}}{=} \frac{\text{IoU}_{adapt}^c}{\text{IoU}_{sup}^c} \cdot 100. \quad (11)$$

In mASR, 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 and producing a more class-agnostic adaptation score. In this metric, higher means better and when the adapted network has the same performance as supervised training the score is 100%.

As an example, the mASR scores reported in the last two columns of Table I allow to identify at a glance the algorithms that more faithfully match the target performance.

To validate the new metric, we used as reference the supervised training on the *Cityscapes* dataset and compared it with the training on corrupted versions of the same dataset using the introduced mASR metric to evaluate the relative performances and so, indirectly, the domain shift introduced by the perturbations. In Fig. 3 we identified 5 types of perturbations which are likely to be encountered by an agent moving outdoor (*i.e.*, Gaussian noise, motion blur, snow, fog, brightness) and we set 5 levels of noise intensity as defined by [51]. As expected, the higher is the noise intensity and the lower is the adaptation score computed by mASR. Furthermore, we can also have a hint of the most detrimental types of noise for adapting source knowledge: namely, Gaussian noise, snow, motion blur. This can help us identify which set of samples we should consider more in order to obtain a reliable model capable of handling these situations. On the other hand, brightness and fog influence less the final results.

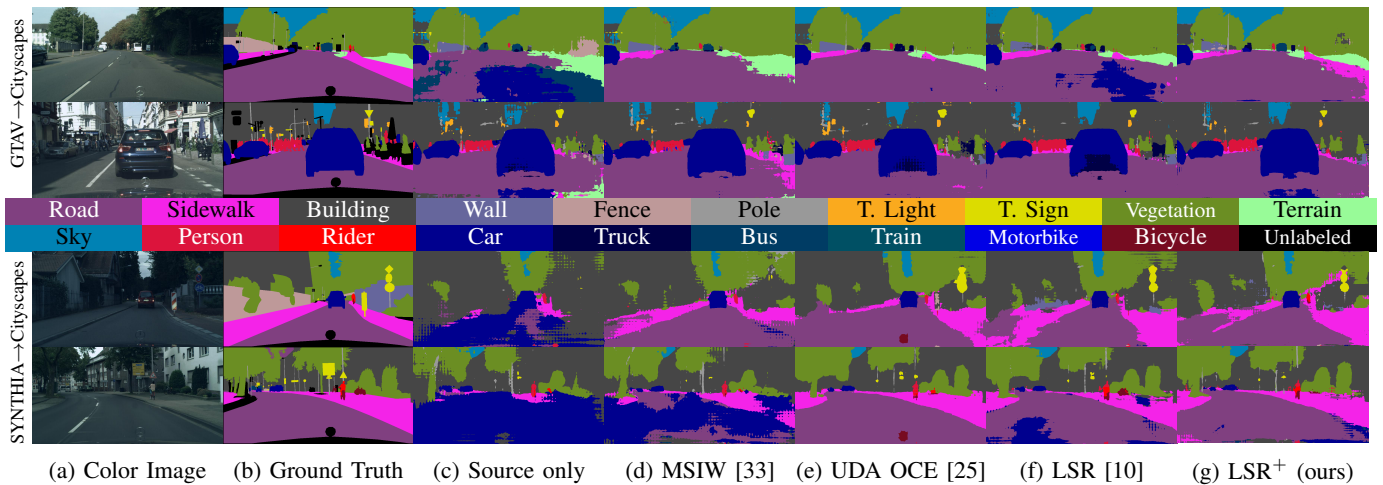


Fig. 4: Qualitative results on sample scenes taken from the *Cityscapes* validation split.

of [25] and [33] able to get close to our result, while there is a quite relevant gap compared to all the other methods. Such improvement is quite stable across most per-class IoU scores, and is particularly evident in challenging classes, such as *terrain* and *t. light* where our strategy shows very high percentage gains, and on *train* where we significantly outperform the competitors by doubling the score of the second-best strategy.

Some qualitative results are reported in the top half of Fig. 4. From visual inspection, we can verify the increased precision of edges in the *t. sign*, *t. light*, *pole* and *person* classes in both images. Furthermore, our approach is the only one to correctly classify the *bus* in the right of the first image as such (confused as *truck* by the other strategies). Importantly, we can also see the effects of our two-pass labeling (see Sec III-C) on the left of the top image (where part of the *fence* is correctly classified by our strategy, while being missed by all competitors) and of the second image (where LSR⁺ significantly reduces the confusion between *sky* and the white building).

In the SYNTHIA to Cityscapes setup, LSR⁺ surpasses its conference version (LSR) by about 1% of mIoU in the 16-classes setup and by 0.6% in the 13-classes one, achieving a final score of 42.6% and 48.7%, respectively. It also outperforms all the other competitors, with a slight margin of 1% on average w.r.t [25] and a larger one (more than 3%) w.r.t. all the other approaches.

Qualitative results are reported in the bottom half of Fig. 4, where the overall increase in segmentation accuracy for many classes such as *car*, *road* and *sidewalk* is evident. In the first image (third row of Fig. 4) we can see how LSR⁺ is the only strategy to correctly classify both *rider* and *bike*, whereas other strategies even miss the *t. sign* in the foreground. Similarly, in the second image we note improvements on the prediction on such classes and, fundamentally, of the *road* in foreground (confused for *car* and *bicycle* by the competitors).

B. Adaptation from Cityscapes to Cross-City

In this subsection we discuss the contents of Table II, where the performance on the *Cross-City* real-to-real benchmark is

reported. This benchmark is comprised of 4 cities: Rome, Rio, Tokyo and Taipei. When evaluated on those setups, our strategy reaches an mIoU score of 56.2%, 52.3%, 50.0% and 50.0% surpassing the source only model by 5.2%, 3.4%, 2.2% and 3.7%, respectively. Importantly, our approach achieves consistent results across the setups (LSR⁺ is the top scorer in 3 out of 4 setups and second in the remaining one) surpassing the average best competitor score by 0.5% mIoU (52.1% versus 51.6%). We remark that the best competitor changes depending on the setup, being [25], [25], [20] and [20] for Rome, Rio, Tokyo and Taipei, respectively, underlining the unstable performances of many approaches usually associated with this benchmark.

Looking at the per-class IoU scores, we can see how our strategy significantly outperforms the competitors in *t. light* and *rider* in the Cityscapes→Rome setup (increase of 6% of IoU), in *person* and *rider* in the Cityscapes→Rio setup (increase of 4% of IoU) and in *motorbike* in the Cityscapes→Taipei setup (increase of 9.3% of IoU). Qualitative results on this benchmark are shown in the Supplementary Material.

C. Results with Different Backbones

Table III shows the performance of our strategy on GTAV→Cityscapes using multiple encoder-decoder backbones in order to evaluate the generalization of the approach to different networks. Here we can see how LSR⁺ outperforms the source-only models (*i.e.*, without adaptation) by 13.3%, 12.7% and 7.8% using ResNet50, VGG-16 and VGG-13, respectively. Even more importantly, we can see how the performance improvement is consistent across all backbones, in opposition to what happens to competing strategies. Finally, we remark the stability of the mASR score of our strategy, hovering around a mean of 57.0% with a very tight standard deviation of 1.4% (the other strategies have means 48.5% and 42.2%, and standard deviations of 2.45% and 31.3%, respectively). Per-class IoUs are reported in the Supplementary Material.

TABLE III: Additional quantitative results with multiple backbones, GTAV→Cityscapes setup. (r) indicates that the strategy was re-trained, starting from the official code.

Backbone	Configuration	mIoU	mASR
ResNet 50	Target only	65.2	100
	Source only	27.6	39.1
	MaxSquareIW [33] (r)	36.8	52.0
	UDA OCE [25] (r)	36.6	51.7
	LSR ⁺ (ours)	40.9	58.6
VGG 16	Target only	59.6	100
	Source only	25.5	42.4
	MaxSquareIW [33] (r)	31.7	46.9
	UDA OCE [25] (r)	34.2	51.5
	LSR ⁺ (ours)	37.2	57.2
VGG 13	Target only	59.5	100
	Source only	28.5	42.6
	MaxSquareIW [33] (r)	31.6	46.7
	UDA OCE [25] (r)	16.8	23.3
	LSR ⁺ (ours)	36.3	55.5

TABLE IV: Ablation Studies, mIoU and mASR scores comparison when removing any of the losses.

\mathcal{L}_C	\mathcal{L}_P	\mathcal{L}_N	\mathcal{L}_{EM}	mIoU	mASR
				42.8	64.4
	✓	✓	✓	44.9	66.3
✓		✓	✓	45.3	66.7
✓	✓		✓	46.0	68.3
✓	✓	✓		44.5	66.1
✓	✓	✓	✓	46.9	69.5

VIII. ABLATION STUDIES

In this section, we evaluate the impact of each component of the approach on the final accuracy.

A. Analysis of the Latent Space Regularization

For visualization purposes, the plots of this section are computed on a balanced subset of feature vectors (250 vectors per class) extracted from the Cityscapes validation set for a fair comparative analysis across the classes.

Two-pass prototypes and clustering. To investigate the semantic feature representation learning produced by our approach we computed a shared t-SNE [53] embedding on the prototypes sampled during the training procedure and of the target features produced by the final model. We remind the reader that, in order to more effectively shift target features closer to the source ones, we resort to a two-stage label assignment procedure which recovers target awareness (by averaging target-extracted features) from prototypes computed on the source domain (by centroid computation) as reported in Sec. III-C. In the left plot of Fig. 5 we report the learned prototype trajectory embeddings, and on the right the respective feature vectors. Here we can appreciate how prototypes get further apart while training goes on and how features extracted from the target domain lie in a neighborhood of the prototype, which we recall is computed exclusively via source-supervision. This underlines the effectiveness of our clustering

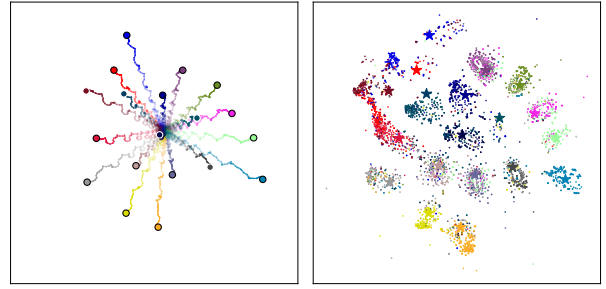


Fig. 5: t-SNE embedding of the target feature vectors: trajectories of prototypes sampled over 200 training steps (on the left), features produced by the final model embedded according to the shared t-SNE projection (right).




Fig. 6: t-SNE embeddings of the normalized feature vectors.

strategy, which is able to shift the target feature distribution closer to the source one.

Finally, to further analyze our clustering objective we produce additional t-SNE embeddings starting from the normalized features (to remove the norm information, focusing on the angular one), which is reported in Fig. 6. Our strategy increases significantly the cluster separation in the high dimensional space and the spacing between clusters belonging to different classes, promoting features disentanglement. 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).

Finally, PCA embeddings are reported in the Supplementary Material to evaluate the effect of latent-spacing techniques when projected to a lower dimension via a linear function.

Weighted histogram-aware downsampling. In this work, we extended the scheme proposed in [10] by adding class weights inversely proportional to the class-frequency in the training dataset (see Sec. III). 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. 7, where the label map downsampled via our frequency-aware schemes (middle and right) marks some features close to the edges of objects as *unlabeled*, keeping only faithful features. As expected, class-weighting (right plot of Fig. 7) promotes rarer classes at the feature level compared to the version without it [10] (middle plot of Fig. 7): for instance, compare the traffic sign (in yellow). This is confirmed




Fig. 7: Sample image downsampled nearest (left), frequency-aware [10] (middle) or weighted frequency-aware (LSR⁺).




Fig. 8: Class frequency of the downsampled feature-level segmentation maps.

by the class distribution of the downsampled segmentation maps (*i.e.*, to match the spatial resolution of the latent space), reported in Fig. 8 for our weighted histogram-aware scheme, the previous un-weighted histogram-aware scheme of the conference version [10] and the standard nearest neighbor. In particular, the schemes based on histogram-awareness generally seldom preserve small object classes, promoting *unlabeled* classification when discrimination between classes is uncertain. Our weighted histogram-aware scheme improves uniformity across rarer or smaller semantic categories, which were over-penalized by the previous approach [10], where all classes were treated equally, regardless of their occurrence.

Perpendicularity. To analyze the effect of the perpendicularity constraint, Fig. 9 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 upper bound for the angle, as our feature vectors have all non-negative coordinates. From the figure, it emerges clearly that LSR-based approaches increase the inter-prototypical angle and that LSR⁺ makes prototypes even more orthogonal with an improvement of more than 2 degrees on average.

Norm Alignment. We analyze the effect of the norm alignment constraint in Fig. 10, where we show the mean channel entropy for each class. We observe that the entropy corresponding to feature vectors produced by LSR⁺ is significantly reduced, meaning that features are characterized by more relevant peaks and fewer poorly-activated channels.

IX. CONCLUSIONS

In this paper, we introduced a set of latent-space regularization techniques to address the domain shift in an unsupervised fashion. We achieved domain invariance by means of multiple latent space-shaping constraints (namely, class clustering, class




Fig. 9: Average inter-prototype angle.




Fig. 10: Average channel entropy.

perpendicularity and norm alignment), to space apart features belonging to different classes while clustering together features of the same class in a consistent way on both the source and target domain. To support their computation, we introduced a novel target pseudo-labeling scheme and a weighted label decimation strategy. Results have been evaluated using both the standard mIoU and a novel metric (mASR), which captures the relative performance between an adapted model and its target supervised counterpart. We surpassed state-of-the-art results in feature-level adaptation on two commonly used synthetic-to-real benchmarks and one real-to-real setup, paving the way to employment of a new family of feature-level techniques to enhance the discrimination ability of deep neural networks.

Future research will focus on designing novel feature-level techniques and on evaluating their generalization ability to different tasks. The adaptation from multiple source domains to multiple target ones will also be considered.

REFERENCES

- [1] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3431–3440.
- [2] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [3] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2881–2890.
- [4] L. Chen, G. Papandreou, F. Schroff, and H. Adam, “Rethinking atrous convolution for semantic image segmentation,” *arXiv preprint arXiv:1706.05587*, 2017.
- [5] L. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Proceedings of the European Conference on Computer Vision*, 2018, pp. 833–851.
- [6] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, pp. 834–848, 2018.

- [7] M. Toldo, A. Maracani, U. Michieli, and P. Zanuttigh, "Unsupervised domain adaptation in semantic segmentation: a review," *Technologies*, vol. 8, no. 2, 2020.
- [8] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [9] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 580–587.
- [10] F. Barbato, M. Toldo, U. Michieli, and P. Zanuttigh, "Latent space regularization for unsupervised domain adaptation in semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2021.
- [11] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The Cityscapes dataset for semantic urban scene understanding," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 3213–3223.
- [12] Y.-C. Chen, Y.-Y. Lin, M.-H. Yang, and J.-B. Huang, "Crdoco: Pixel-level domain transfer with cross-domain consistency," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1791–1800.
- [13] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "Cycada: Cycle-consistent adversarial domain adaptation," in *Proceedings of the International Conference on Machine Learning*, 2018, pp. 1994–2003.
- [14] J. Hoffman, D. Wang, F. Yu, and T. Darrell, "FCNs in the wild: Pixel-level adversarial and constraint-based adaptation," *arXiv preprint arXiv:1612.02649*, 2016.
- [15] Z. Murez, S. Kolouri, D. J. Kriegman, R. Ramamoorthi, and K. Kim, "Image to image translation for domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4500–4509.
- [16] M. Toldo, U. Michieli, G. Agresti, and P. Zanuttigh, "Unsupervised domain adaptation for mobile semantic segmentation based on cycle consistency and feature alignment," *Image and Vision Computing*, vol. 95, 2020.
- [17] F. Pizzati, R. d. Charette, M. Zaccaria, and P. Cerri, "Domain bridge for unpaired image-to-image translation and unsupervised domain adaptation," in *Proceedings of the Winter Conference on Applications of Computer Vision*, 2020, pp. 2990–2998.
- [18] L. Du, J. Tan, H. Yang, J. Feng, X. Xue, Q. Zheng, X. Ye, and X. Zhang, "SSF-DAN: separated semantic feature based domain adaptation network for semantic segmentation," in *Proceedings of the International Conference on Computer Vision*, 2019, pp. 982–991.
- [19] S. Sankaranarayanan, Y. Balaji, A. Jain, S. Nam Lim, and R. Chellappa, "Learning from synthetic data: Addressing domain shift for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3752–3761.
- [20] Y.-H. Tsai, W.-C. Hung, S. Schuster, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7472–7481.
- [21] Y.-H. Tsai, K. Sohn, S. Schuster, and M. Chandraker, "Domain adaptation for structured output via discriminative patch representations," in *Proceedings of the International Conference on Computer Vision*, 2019, pp. 1456–1465.
- [22] M. Biassetton, U. Michieli, G. Agresti, and P. Zanuttigh, "Unsupervised Domain Adaptation for Semantic Segmentation of Urban Scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 1211–1220.
- [23] U. Michieli, M. Biassetton, G. Agresti, and P. Zanuttigh, "Adversarial learning and self-teaching techniques for domain adaptation in semantic segmentation," *IEEE Transaction on Intelligent Vehicles*, vol. 5, pp. 508–518, 2020.
- [24] T. Spadotto, M. Toldo, U. Michieli, and P. Zanuttigh, "Unsupervised domain adaptation with multiple domain discriminators and adaptive self-training," in *Proceedings of the International Conference on Pattern Recognition*, 2020.
- [25] M. Toldo, U. Michieli, and P. 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*, 2021, pp. 1358–1368.
- [26] S. Lee, D. Kim, N. Kim, and S.-G. Jeong, "Drop to adapt: Learning discriminative features for unsupervised domain adaptation," in *Proceedings of the International Conference on Computer Vision*, 2019, pp. 91–100.
- [27] S. Park, J. Park, S. Shin, and I. Moon, "Adversarial dropout for supervised and semi-supervised learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018, pp. 3917–3924.
- [28] K. Saito, Y. Ushiku, T. Harada, and K. Saenko, "Adversarial dropout regularization," in *Proceedings of the International Conference on Learning Representations*, 2018.
- [29] Y. Zou, Z. Yu, B. Vijaya Kumar, and J. Wang, "Unsupervised domain adaptation for semantic segmentation via class-balanced self-training," in *Proceedings of the European Conference on Computer Vision*, 2018, pp. 289–305.
- [30] Y. Zou, Z. Yu, X. Liu, B. V. Kumar, and J. Wang, "Confidence regularized self-training," in *Proceedings of the International Conference on Computer Vision*, 2019, pp. 5982–5991.
- [31] Y. Zhang, P. David, and B. Gong, "Curriculum domain adaptation for semantic segmentation of urban scenes," in *Proceedings of the International Conference on Computer Vision*, 2017, pp. 2020–2030.
- [32] Y. Zhang, P. David, H. Foroosh, and B. Gong, "A curriculum domain adaptation approach to the semantic segmentation of urban scenes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [33] M. Chen, H. Xue, and D. Cai, "Domain adaptation for semantic segmentation with maximum squares loss," in *Proceedings of the International Conference on Computer Vision*, 2019, pp. 2090–2099.
- [34] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez, "Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2517–2526.
- [35] G. Kang, L. Jiang, Y. Yang, and A. G. Hauptmann, "Contrastive adaptation network for unsupervised domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4893–4902.
- [36] L. Tian, Y. Tang, L. Hu, Z. Ren, and W. Zhang, "Domain adaptation by class centroid matching and local manifold self-learning," *arXiv preprint arXiv:2003.09391*, 2020.
- [37] U. Michieli and P. 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.
- [38] N. Dong and E. P. Xing, "Few-shot semantic segmentation with prototype learning," in *Proceedings of the British Machine Vision Conference*, vol. 3, 2018.
- [39] K. Wang, J. H. Liew, Y. Zou, D. Zhou, and J. Feng, "Panet: Few-shot image semantic segmentation with prototype alignment," in *Proceedings of the International Conference on Computer Vision*, 2019, pp. 9197–9206.
- [40] J. Liang, R. He, Z. Sun, and T. 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*, 2019, pp. 2975–2984.
- [41] Q. Wang and T. P. Breckon, "Unsupervised domain adaptation via structured prediction based selective pseudo-labeling," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 6243–6250.
- [42] H. Choi, A. Som, and P. Turaga, "Role of orthogonality constraints in improving properties of deep networks for image classification," *arXiv preprint arXiv:2009.10762*, 2020.
- [43] P. O. Pinheiro, "Unsupervised domain adaptation with similarity learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8004–8013.
- [44] S. Wu, J. Zhong, W. Cao, R. Li, Z. Yu, and H.-S. Wong, "Improving domain-specific classification by collaborative learning with adaptation networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, pp. 5450–5457.
- [45] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, 1999, pp. 1150–1157 vol.2.
- [46] S. R. Richter, V. Vineet, S. Roth, and V. Koltun, "Playing for data: Ground truth from computer games," in *Proceedings of the European Conference on Computer Vision*, 2016, pp. 102–118.
- [47] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez, "The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 3234–3243.
- [48] Y.-H. Chen, W.-Y. Chen, Y.-T. Chen, B.-C. Tsai, Y.-C. Frank Wang, and M. Sun, "No more discrimination: Cross city adaptation of road scene segmenters," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1992–2001.

- [49] W. Tranheden, V. Olsson, J. Pinto, and L. Svensson, "Dacs: Domain adaptation via cross-domain mixed sampling," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1379–1389.
- [50] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [51] D. Hendrycks and T. G. Dietterich, "Benchmarking neural network robustness to common corruptions and surface variations," 2019.
- [52] C. Li, D. Du, L. Zhang, L. Wen, T. Luo, Y. Wu, and P. Zhu, "Spatial attention pyramid network for unsupervised domain adaptation," in *Proceedings of the European Conference on Computer Vision*, 2020.
- [53] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of machine learning research*, vol. 9, no. 11, 2008.



Francesco Barbato received the M.Sc. degree in Telecommunication Engineering from the University of Padova in 2020. He is currently a graduate student with research grant at the same University. His research focuses on domain adaptation and continual learning applied to computer vision tasks, particularly to semantic segmentation for autonomous vehicles.



Umberto Michieli received the M.Sc. degree in Telecommunication Engineering from the University of Padova in 2018. He is currently a final-year Ph.D. student at the same University. In 2018, he spent 6 months as a Visiting Researcher with the Technische Universität Dresden. In 2020 he interned as Research Engineer for 8 months at Samsung Research UK. His research focuses on transfer learning techniques for semantic segmentation, in particular on domain adaptation and on incremental learning.



Marco Toldo Marco Toldo received the M.Sc. degree in ICT for Internet and Multimedia in 2019 at the University of Padova. At present, he is doing his Ph.D. at the Department of Information Engineering of the same university. He is also doing an internship as Research Engineer at Samsung Research UK. His research involves domain adaptation and continual learning applied to computer vision.



Pietro Zanuttigh received a Master degree in Computer Engineering at the University of Padova in 2003 where he also got the Ph.D. degree in 2007. Currently he is an associate professor at the Department of Information Engineering. He works in the computer vision field, with a special focus on domain adaptation and incremental learning in semantic segmentation, 3D acquisition with ToF sensors, depth data processing, sensor fusion and hand gesture recognition.