



# Road scenes segmentation across different domains by disentangling latent representations

Francesco Barbato<sup>1</sup> · Umberto Michieli<sup>1</sup> · Marco Toldo<sup>1</sup> · Pietro Zanuttigh<sup>1</sup>

Accepted: 22 February 2023  
© The Author(s) 2023

## Abstract

Deep learning models obtain impressive accuracy in road scene understanding; however, they need a large number of labeled samples for their training. Additionally, such models do not generalize well to environments where the statistical properties of data do not perfectly match those of training scenes, and this can be a significant problem for intelligent vehicles. Hence, domain adaptation approaches have been introduced to transfer knowledge acquired on a label-abundant source domain to a related label-scarce target domain. In this work, we design and carefully analyze multiple latent space-shaping regularization strategies that work together to reduce the domain shift. More in detail, we devise a feature clustering strategy to increase domain alignment, a feature perpendicularity constraint to space apart features 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. In addition, we propose a novel evaluation metric to capture the relative performance of an adapted model with respect to supervised training. We validate our framework in driving scenarios, considering both synthetic-to-real and real-to-real adaptation, outperforming previous feature-level state-of-the-art methods on multiple road scenes benchmarks.

**Keywords** Semantic segmentation · Domain adaptation · Latent space shaping · Representation learning

## 1 Introduction

One of the key components of a self-driving vehicle is the capability to understand the surrounding environment from sensory input data. Semantic segmentation enables accurate scene understanding, by assigning all pixels of the input images to a semantic category corresponding to key elements to be detected, such as the road, other vehicles or traffic lights and signs [1]. Nowadays, such a task is commonly tackled

with deep convolutional neural networks (DCNNs), which have achieved outstanding results in image understanding tasks, provided that a sufficiently large number of labeled examples are available from the target input domain distribution. On the other side, the annotation of thousands of images of road scenes is highly expensive, time-consuming, error-prone and, possibly, worthless, since the test data can show a domain shift with respect to the training labeled samples. Therefore, recently, a new requirement has emerged for DCNN-based scene understanding systems embedded in autonomous driving vehicles, namely allowing training with a combination of labeled source samples (*e.g.*, synthetic from ad hoc simulators or driving video games) and unlabeled target samples (*e.g.*, real-world acquisitions from cameras mounted on cars), with the aim of getting high performance on data following the target distribution. With this strategy, the need for large quantities of labeled real-world data can be avoided by using samples from a source domain where they are abundantly available and annotations are faster and cheaper to generate.

Unfortunately, DCNNs are prone to failure when they are shown an input domain distribution other than the training one (*domain shift* phenomenon). In order to deal with this

---

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*.

✉ Francesco Barbato  
francesco.barbato@dei.unipd.it

Umberto Michieli  
umberto.michieli@dei.unipd.it

Marco Toldo  
marco.toldo.3@phd.unipd.it

Pietro Zanuttigh  
zanuttigh@dei.unipd.it

<sup>1</sup> Department of Information Engineering, University of Padova, Via Gradenigo 6/b, 35131 Padova, Italy

problem, various unsupervised domain adaptation (UDA) techniques have been developed to adapt networks at different stages (the most common are the input, feature and output levels) [2].

Deep learning models for semantic segmentation are mostly based on encoder–decoder architectures, *i.e.*, they build some concise latent representations of the inputs, which are highly correlated with the classifier output. As such, they are used in the subsequent classification process [3,4] that reconstructs the full-resolution segmentation map. Nevertheless, only some UDA techniques for semantic segmentation work in the feature space because of its large dimensionality.

In this paper, we propose an approach focusing on feature space adaptation, introducing a new set of strategies working at the latent space level, built on top of our previous conference work [5]. Employing a shaping objective at that level we are able to promote class-aware feature extraction and feature invariance between source and target domains. More in detail, our target is to push the feature vectors on the boundary of an  $N$ -dimensional sphere, assigning to each class a direction orthogonal to all others. Such a shaping objective can be implemented using three major components:

Firstly, a clustering constraint groups feature vectors of each class tightly around their prototypical representation.

Secondly, a perpendicularity objective over the class prototypes promotes disjoint filter activation sets across different semantic categories.

Finally, a regularization-based norm alignment objective enforces consistent vector norms in the source and target domains, while jointly forcing progressively increasing norm values. This, in combination with the perpendicularity constraint, is able to reduce the entropy associated with feature vector channel activations.

Importantly, the proposed techniques require the generation of accurate class prototypes and the imposition of a strong correlation between feature representations and predicted segmentation maps. Hence, we also propose a novel strategy to map semantic information from the labeling maps to the low-resolution feature space (annotations downsampling).

## 1.1 Limitations in existing work

The major limitations of our previous approach [5] can be found in the inherent instability of pseudo-labels in the UDA setting. Since the employed architecture is trained on a different domain, the predictions and features it produces when analyzing target samples may vary significantly between batches. This leads to a loss in performance, especially when considering class-conditional clustering and feature perpendicularity. A further problem can be found in the norm alignment constraint, which enforced the same norm value for all features, regardless of the semantic content and of

the specific channels' activation profile. This strategy may introduce noise, since channels that the network may want to disable are still forced to produce a nonzero result, especially in feature vectors that correspond to pixel windows with multiple class labels (usually found along object borders) or for which there is no clear network response.

## 1.2 Contribution

To tackle the first problem, we focused on the computation of prototypes and feature vector extraction. The first now considers the prototype trajectory evolution for a better estimation (Sect. 3.2), while the second exploits target information to reduce the domain shift (Sect. 3.3); additionally, a class-weighting scheme is used in the source supervision (Sect. 3.1). In particular, the clustering objective was modified to be more resilient to outliers (Sect. 4.1); the perpendicularity constraint now accounts for classes not present in the current batch (Sect. 4.2); and the norm alignment now ignores low-activated channels (Sect. 4.3).

To analyze the improvements to the proposed space-shaping constraints, we report ablation studies on our approach, LSR<sup>+</sup> (Sect. 8), and for the proposed evaluation metric, mASR (Sect. 6).

Extensive experiments have been conducted on many road scenarios, expanding the set of experiments reported in [5]. The results are evaluated on 4 backbones in 6 different setups. These include not only 2 synthetic-to-real ones, commonly used in related works, but also 4 real-to-real settings addressing the critical issue of generalizing autonomous driving systems across different cities and types of roads in different regions of the world. Additional results using the unlabeled *Cityscapes* coarse set [6] are reported, showing significant performance gains when more unlabeled data are used (see Table 1).

Overall, our approach shows significant improvement over the conference version [5], gaining, on average, 1% mIoU on all settings. Importantly, the performance gains are well spread across the classes, as is confirmed by the very high mASR score. Such achievement confirms that our space-shaping approach is able to bridge the gap between source and target domains very effectively, even when small and rare classes are considered.

## 2 Related works

*Semantic segmentation of road scenes* is a very active research field. Many semantic segmentation architectures based on the widely used encoder–decoder architecture [7–12] or more recently on vision transformers [13–15] have been applied to road scenes for two interconnected motivations: First, there is a large interest into the target application

of self-driving vehicles, and second, there is the availability of large real world (Cityscapes [6], Mapillary [16], IDD [17], Cross-City [18], CamVid [19]) and synthetic (GTA5 [20], SYNTHIA [21], SELMA [22] and SHIFT[23]) datasets that can be employed to train deep learning architectures. These two motivations also provided a strong push for road-scenes-specific architectures [24–29], which have recently started to be proposed by the research community.

**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 decline 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 and in other similarly complex tasks [30]. 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 [31] or image generation techniques, while others aim 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 [2].

Input-space adaptation has been commonly addressed resorting to image-to-image translation [32–39]. 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 the source and target data.

As concerns feature and output-space adaptation, adversarial learning has been largely employed to bridge the statistical domain gap [40–46]. 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 [40–42] or in the form of probability maps at the output of the segmentation pipeline [42–46]. More recently, some works focusing on feature-level regularization have been proposed [5,47]. In [47], 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 [5] 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 [48–50]: decision boundaries are pushed away from target high-density regions in the latent space without direct supervision.

Output-space adaptation has been further pursued by resorting to self-training [51,52], where the learning process is guided (in a self-supervised manner) by pseudo-labels extracted from target network predictions. Self-supervision has been proposed in a curriculum learning fashion as well [53–55]. 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 [56–58], 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 [59,60], continual learning [61], federated learning [62] and few-shot learning [63,64]. 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 a 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 [5,47].

A multitude of feature clustering techniques based on the K-means algorithm has been proposed [59,60,65,66] 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 [60,65,66]. In [5,47], 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 [67] to limit the redundancy of the information encoded in feature representations. Approaches closer to our strategy are [68,69], where UDA is promoted via an orthogonality objective over class prototypes. Nonetheless, [67–69] all limit their focus to the image classification task.

### 3 Problem setting

In this section, we overview our setup, detailing the mathematical notation used throughout the paper. We start by identifying the input space as  $\mathcal{X} \subset \mathbb{R}^{H \times W \times 3}$  and the corresponding label space as  $\mathcal{Y} \subset \mathcal{C}^{H \times W}$ , where  $H$  and  $W$

represent the image resolution and  $\mathcal{C}$  the class set. 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 source domain, while an additional set of input samples  $\mathcal{T}^t = \{\mathbf{X}_n^t\}_{n=1}^{N_t}$  is drawn from an unlabeled target domain ( $\mathbf{X}_n^t \in \mathcal{X}^t$ ). We adapt the knowledge of semantic segmentation learned on the source domain to the unsupervised target domain. Superscript  $s$  identifies the source domain, while  $t$  the target.

As done by most recent approaches for semantic segmentation, we assume a model  $S = D \circ E$  based on an encoder–decoder architecture, that is, made by the consecutive application of an encoder network  $E$  (referred to as backbone, which acts as feature extractor) and a decoder network  $D$ , which actually performs the classification and produces the segmentation map. We denote the features extracted from an 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 (feature-level) spatial resolution. Thanks to the topology of encoder–decoder DCNNs for semantic segmentation, classes are encoded into ideal latent representations, invariant with respect to the domain shift. The strategies presented in Sect. 4 enforce this goal by comparing the extracted features belonging to each class with the respective prototypical representations. In the following paragraphs, we present the techniques used to compute the prototypes and associate feature vectors with semantic classes.

### 3.1 Weighted histogram-aware downsampling

Given that the spatial arrangement of the pixels of an image is maintained while it is processed by an encoder–decoder dilated residual network [70] (even if at reduced resolution in the internal feature representation), we can assume a tight relationship between any feature vector (*i.e.*, the vector of features associated with a single spatial location within the feature tensor) and the semantic labeling of the corresponding image region.

Hence, the extraction process begins with the identification of a way to propagate the labeling information to latent representations (decimation), without losing the semantic content of the window (image region) corresponding to each feature vector. Issues in the mapping typical of naïve approaches would strongly affect the whole following procedure. Our solution is a nonlinear pooling function, which instead of computing a simple subsampling (*e.g.*, nearest neighbor) extracts 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. Then, these metrics are used to select the most

appropriate class for each image region, producing source feature-level label maps  $\{\mathbf{I}_n^s\}_{n=1}^{N_s}$ . The computation of the target counterparts ( $\{\mathbf{I}_n^t\}_{n=1}^{N_t}$ ) is discussed in Sect. 3.3 and we remark that each  $\mathbf{I}_n^{s,t}$  belongs to the set  $\mathcal{C}^{H' \times W'}$ . In particular, we choose the label with the highest frequency peak in the windows, only if such a peak is relevant enough, *i.e.*, if all other peaks are smaller than  $T_h$  times it. (A similar approach is found in the orientation assignment step of the SIFT feature extractor [71].) Empirically, we set  $T_h = 0.5$ . Finally, we remark on a useful side effect of this technique: whenever a window cannot be uniquely assigned to a class (that is, it contains multiple labels), the procedure automatically assigns it to the *void* class.

### 3.2 Prototype extraction

The feature-level label maps  $\{\mathbf{I}_n^{s,t}\}_{n=1}^{N_{s,t}}$  allow to identify the set  $\mathcal{F}_c^{s,t}$  of feature vectors belonging to class  $c \in \mathcal{C}$  in training batch  $\mathcal{B}$ :

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

where the couple  $[h, w]$  denotes the spatial location ( $0 \leq h < H'$  and  $0 \leq w < W'$ ). The definition is further expanded into the set of all feature vectors in batch  $\mathcal{B}$  by taking their union with the set  $\mathcal{F}_v^{s,t}$  of samples belonging to class *void*:  $\mathcal{F}^{s,t} = (\bigcup_c \mathcal{F}_c^{s,t}) \cup \mathcal{F}_v^{s,t}$ . From these sets, we can extract the batch-wise prototypes of each class (note that we use feature vectors exclusively from the source):

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


Moreover, with the goal of obtaining more stable and reliable prototypes, and reducing estimation noise, we consider the exponentially smoothed vectors:

$$\hat{\mathbf{p}}_{c,new} = (1 - \eta)\mathbf{p}_c + \eta\hat{\mathbf{p}}_{c,old} \quad (3)$$

The parameters are initialized with  $\hat{\mathbf{p}}_c = \mathbf{0}$  and  $\eta = 0.8$  (empirically). In our notation,  $\hat{\mathbf{p}}_{c,old}$  represents the smoothed estimate up to the previous optimization step,  $\mathbf{p}_c$  the prototypes estimated on the current batch and  $\hat{\mathbf{p}}_{c,new}$  the new smoothed estimate. This way, by setting  $\eta > 0$ , we can propagate the previous estimates to the current batch, allowing to consider classes absent from  $\mathbf{Y}_n^s$  in the loss computation.

### 3.3 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



**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

work [5]), this approach can introduce instability in the training procedure. To avoid such an 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:

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

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

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

## 4 Methodology

The proposed approach is detailed in this section, 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 a problem,

in [5] we proposed to use additional space-shaping 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 expressed as:

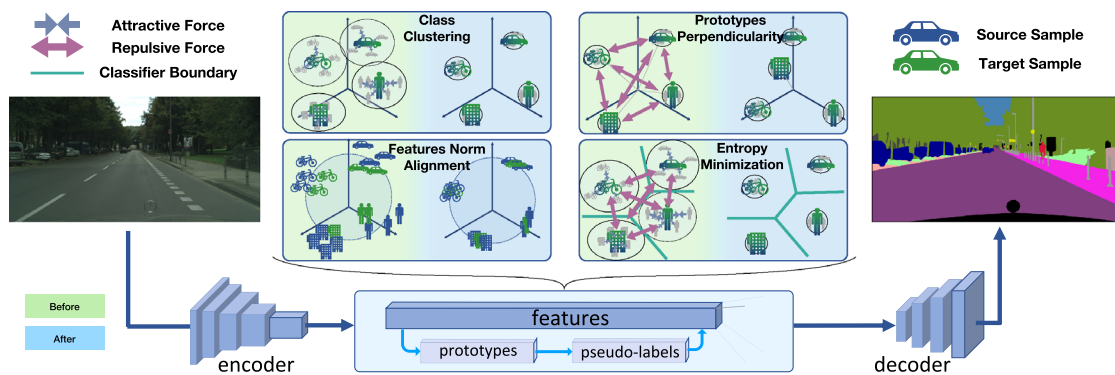
$$\mathcal{L} = \mathcal{L}_{CE}^s + \lambda_C \cdot \mathcal{L}_C^{s,t} + \lambda_P \cdot \mathcal{L}_P^s + \lambda_N \cdot \mathcal{L}_N^{s,t} \quad (6)$$

Here,  $\mathcal{L}_C$  represents the clustering objective acting on the feature vectors (Sect. 4.1),  $\mathcal{L}_P$  the perpendicularity constraint applied to class prototypes (Sect. 4.2) and  $\mathcal{L}_N$  the norm alignment goal (Sect. 4.3). To simplify the notation, Eq. (6) contains each loss component with  $s, t$  superscripts to indicate the sum of the loss on source and target samples. To further improve the performance and to show how the proposed techniques can be used on top of existing strategies, we also add to the optimization target the entropy minimization loss introduced by Chen *et al.* [56], thus obtaining:  $\mathcal{L}^+ = \mathcal{L} + \lambda_{EM} \cdot \mathcal{L}_{EM}$ . 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 proposed strategy is presented in Fig. 2.

### 4.1 Clustering of latent representations

Due to the distribution discrepancy between source and target domains, feature vectors originating from them will be 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 class-conditional alignment of feature distribution. We do so by exploiting the prototypical representations discussed in Sect. 3 and forcing the feature vectors from source and target representations to tightly cluster around them: representations are adapted into a common class-wise distribution and the discrimination capacity of the latent space is increased.

Differently from the previous work, we define the clustering objective as the L1 distance between feature vectors and their associated class prototypes. 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 the L2 loss, outliers with distances greater than 1 have a strong push toward the clusters. On the other hand, the L1 loss is stronger than L2 for close samples, which are more representative of each class, and is significantly gentler than L2 for distant outliers. The



**Fig. 2** Visual summary of our strategy. Features are associated with semantic classes and prototypes are computed from them (Sect. 3). Class clustering (Sect. 4.1), prototypes perpendicularity (Sect. 4.2) and

norm alignment and enhancement (Sect. 4.3) are the three proposed space-shaping constraints. Additionally, we apply on top an entropy minimization objective [56]

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 targets: the first is the increased clustering of the latent representations thanks to label supervision, which reduces the tendency to erroneous predictions. The second one is to perform self-supervised clustering on target samples using our two-pass pseudo-labeling strategy (see Sect. 3.3). Finally, it leads to better prototype estimates, due to the fact that forcing tighter clusters will lead to more stable batch-wise centroids, which in turn will get closer to the moving-averaged prototypes.

### 4.2 Perpendicularity of latent representations

A prototype perpendicularity loss is further proposed to aid the latent space regularization brought by the clustering objective. Our goal is to induce compact and domain-aligned feature clusters, in order to boost the accuracy of network segmentation maps. As a direct consequence, the margin between classification boundaries and feature clusters is expanded, thus decreasing the probability that target high-density regions are traversed by such boundaries. We directly encourage a class-wise orthogonality property, not only increasing the distance among class clusters but also reducing class cross talk by discouraging shared channel activations in distinct categories.

In the loss, we encode the perpendicularity score exploiting the definition of euclidean space inner product:  $\mathbf{j} \cdot \mathbf{k} = \|\mathbf{j}\| \|\mathbf{k}\| \cos \theta$ , where  $\theta$  is the angle between the two vectors  $\mathbf{j}$  and  $\mathbf{k}$ . To maximize  $\theta$  we just need to minimize the vectors' normalized product (recall that  $\mathbf{j}, \mathbf{k} \in \mathbb{R}_{0+}^K$ ).

Therefore, the cross-perpendicularity between prototypes is encoded as:

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

Equation (8) computes the sum of the cosines over the set of all couples of non-void classes. The influence of the orthogonality objective indirectly reaches all feature vectors, as prototypical representations and single feature instances share a strong geometric bound promoted by  $\mathcal{L}_C^{s,t}$ . What we ultimately achieve is thus to enforce a perpendicularity constraint among instances of different clusters, with a homogeneous action over all latent representations from the same semantic class. In other words, the angular gap among distinct semantic categories in the feature space is enlarged, by inducing disjoint patterns of activated feature channels between distinct classes.

In contrast to our previous paper [5], 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.

### 4.3 Latent norm alignment constraint

This loss term is computed by exploiting source and target feature vector norms. More in detail, we enforce norm consistency between the latent representations extracted from the two domains. This has two objectives: firstly, we aim at improved classification confidence on target predictions, as done by adaptation techniques using entropy minimization in the output space [57]. Secondly, we assist the perpendicu-

larity constraint by reducing the number of domain-specific feature channels used by the network for classification. Thirdly, we reduce the number of channels enabled only on one of the domains, which would lead to norm misalignment. Moreover, to reduce the possible decrease in norm value during the alignment process, we introduce a regularization term that promotes norm increase. Differently from [5], here the norm objective is encoded as a relative difference with a regularization term inversely proportional to the norm value. This allows obtaining 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 average source vector norm (extracted in the previous optimization step),  $\Delta_f$  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:

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

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

This objective is applied in a completely unsupervised manner, and the vector norms are forced to align to the same value regardless of their class. In this way, we remove the bias generated by heterogeneous pixel-class distribution in semantic labels, that, for example, lead features of the most frequent classes to have larger norms than average. The constraint of Eq. (9) forces the inter-class alignment step, *i.e.*, it promotes gradual alignment of the norms toward a target common to all categories, while discouraging the value of such a target to decrease too rapidly. 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.

## 5 Implementation details

### 5.1 Training data

We evaluated our approach (LSR<sup>+</sup>, latent space regularization) on road scenes segmentation in various synthetic-to-real and real-to-real unsupervised adaptation tasks. As source domains we used the synthetic datasets *GTAV* [20] and *SYNTHIA* [21]. The first contains 24,966 labeled images at a resolution of 1914 × 1052 px, produced with the rendering engine of the *GTAV* video game, while the second contains 9,500 labeled images at a resolution of 1280 × 760 px, rendered with a custom software. As target domain, we selected the real-world dataset *Cityscapes* [6]. It contains 5,000 labeled images at a resolution of 2048 × 1024 px 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 to as *CS-full*) improving the adaptation process (as we show in Table 1). In the real-to-real setup, we used the *Cross-City* benchmark, where the *Cityscapes* dataset takes the role of the source domain, while the *Cross-City* dataset [18] takes the role of target. Such dataset is comprised of 12,800 high resolution (2048 × 1024 px) images taken in four major cities (*Rome*, *Rio*, *Tokyo* and *Taipei*).

We trained the model in a closed-set [2] setting, *i.e.*, with the same source and target class sets. More in detail, we used the 19, 16 and 13 classes in common for *GTAV*, *SYNTHIA* and *Cross-City*, respectively. *GTAV*, *Cityscapes* and *Cross-City* images have been rescaled for training to 1280 × 720 px, 1280 × 640 px and 1280 × 640 px, respectively, while the resolution of *SYNTHIA* images has not been changed.

### 5.2 Segmentation network

We employed the well-known [42,47,56,57,72] DeepLabV2 network [10–12], with ResNet101 [73] as the backbone (using 2048 channels in the last stage of the encoder) and a stride of 8. We pre-train the network as in [5], employing the same data augmentation techniques used during adaptation.

### 5.3 Network training

We optimize the model with SGD (using a momentum of 0.9 and a weight decay regularization of  $5 \times 10^{-4}$ ). The

**Table 1** Comparison of adaptation strategies in terms of IoU, mIoU and mASR (Sect. 6). Best in **bold**, runner-up underlined. mIoU<sup>1</sup> and mASR<sup>1</sup> restricted to 13 classes, ignoring the classes with the same superscript

UDA setup	Configuration	Road	Sidewalk	Building	Wall <sup>1</sup>	Fence <sup>1</sup>	Pole <sup>1</sup>	Traffic light	Traffic sign	Vegetation	Terrain			
From GTAV	Target only	96.5	73.8	88.4	42.2	43.7	40.7	46.1	58.6	88.5	54.9			
	Baseline [47]	71.4	15.3	74.0	21.1	14.4	22.8	33.9	18.6	80.7	20.9			
	ASN (feat) [42]	83.7	27.6	75.5	20.3	19.9	27.4	28.3	27.4	79.0	28.4			
	MinEnt [57]	84.4	18.7	80.6	23.8	23.2	28.4	36.9	23.4	83.2	25.2			
	SAPNet [75]	88.4	<b>38.7</b>	79.5	<u>29.4</u>	<b>24.7</b>	27.3	32.6	20.4	82.2	32.9			
	MaxSquareIW [56]	87.7	25.2	<b>82.9</b>	<b>30.9</b>	24.0	29.0	35.4	24.2	84.2	38.2			
	UDA OCE [47]	<b>89.4</b>	30.7	82.1	23.0	22.0	<u>29.2</u>	<u>37.6</u>	<b>31.7</b>	83.9	37.9			
	LSR [5]	87.7	<u>32.6</u>	<u>82.6</u>	29.1	23.0	28.5	36.1	<u>28.5</u>	<b>84.8</b>	<u>41.8</u>			
	LSR <sup>+</sup> (ours)	<u>88.9</u>	26.6	82.0	21.0	<u>24.4</u>	<b>30.1</b>	<b>41.1</b>	27.0	<u>84.7</u>	<b>42.7</b>			
LSR <sup>+</sup> on <i>CS-full</i>	89.3	28.7	82.1	25.2	27.5	31.9	40.3	33.2	84.7	38.7				
From SYNTHIA	Baseline [47]	17.7	15.0	74.3	10.1	0.1	25.5	6.3	10.2	75.5	–			
	ASN (feat) [42]	62.4	21.9	76.3	–	–	–	11.7	11.4	75.3	–			
	MinEnt [57]	73.5	29.2	77.1	7.7	<u>0.2</u>	27.0	7.1	11.4	76.7	–			
	SAPNet [75]	81.7	33.5	75.9	–	–	–	7.0	6.3	74.8	–			
	MaxSquareIW [56]	78.9	33.5	75.3	<u>15.0</u>	<b>0.3</b>	27.5	<u>13.1</u>	<b>16.7</b>	73.8	–			
	UDA OCE [47]	<b>88.3</b>	<b>42.2</b>	79.1	7.1	<u>0.2</u>	24.4	<b>16.8</b>	<u>16.5</u>	<u>80.0</u>	–			
	LSR [5]	81.0	36.9	<b>79.5</b>	13.4	<u>0.2</u>	<u>28.7</u>	9.0	16.1	79.1	–			
	LSR <sup>+</sup> (ours)	<u>82.6</u>	<u>38.4</u>	<b>80.6</b>	<b>15.5</b>	<b>0.3</b>	<b>31.8</b>	6.7	16.3	<b>81.7</b>	–			
	LSR <sup>+</sup> on <i>CS-full</i>	89.4	47.9	79.4	13.9	0.4	29.5	10.0	16.5	79.5	–			
Backbone setup	Configuration	Sky	Person	Rider	Car	Truck	Bus	Train	Motorbike	Bicycle	mIoU	mIoU <sup>1</sup>	mASR	mASR <sup>1</sup>
From GTAV	Target only	91.9	68.7	46.2	90.7	68.8	69.9	48.8	47.6	64.5	64.8	–	100	100
	Baseline [47]	68.5	56.6	27.1	67.4	32.8	5.6	7.7	28.4	33.8	36.9	–	54.0	–
	ASN (feat) [42]	70.1	55.1	20.2	72.9	22.5	35.7	<u>8.3</u>	20.6	23.0	39.0	–	56.9	–
	MinEnt [57]	<u>79.4</u>	59.0	<b>29.9</b>	78.5	33.7	29.6	1.7	29.9	33.6	42.3	–	61.9	–
	SAPNet [75]	73.3	55.5	26.9	82.4	31.8	41.8	2.4	26.5	24.1	43.2	–	63.1	–
	MaxSquareIW [56]	79.2	59.0	<u>27.7</u>	79.5	34.6	44.2	7.5	<u>31.1</u>	<u>40.3</u>	45.5	–	67.2	–
	UDA OCE [47]	78.3	<u>60.7</u>	27.4	<b>84.6</b>	<u>37.6</u>	<u>44.7</u>	7.3	26.0	38.9	45.9	–	67.3	–
	LSR [5]	<b>80.1</b>	59.4	23.8	76.5	<b>38.4</b>	<b>45.8</b>	7.1	28.5	40.1	<u>46.0</u>	–	<u>67.7</u>	–
	LSR <sup>+</sup> (ours)	<b>80.1</b>	<b>63.0</b>	26.4	<u>83.1</u>	30.4	44.3	<b>16.8</b>	<b>35.8</b>	<b>42.4</b>	<b>46.9</b>	–	<b>69.5</b>	–
LSR <sup>+</sup> on <i>CS-full</i>	81.2	63.2	27.2	85.2	34.7	43.9	9.8	37.2	47.7	48.0	–	71.3	–	
From SYNTHIA	Baseline [47]	77.9	57.1	19.2	31.2	–	31.2	–	10.0	20.1	30.1	34.3	41.7	44.6
	ASN (feat) [42]	80.9	53.7	18.5	59.7	–	13.7	–	<b>20.6</b>	24.0	–	40.8	–	52.5
	MinEnt [57]	82.1	57.2	<u>21.3</u>	69.4	–	29.2	–	12.9	27.9	38.1	44.2	51.1	56.3
	SAPNet [75]	78.9	52.1	<u>21.3</u>	75.7	–	30.6	–	10.8	28.0	–	44.3	–	56.0
	MaxSquareIW [56]	77.7	50.4	19.9	66.7	–	<b>36.1</b>	–	13.7	32.1	39.4	45.2	53.8	58.3
	UDA OCE [47]	<b>84.3</b>	56.2	15.0	<b>83.5</b>	–	27.2	–	6.3	30.7	41.1	<u>48.2</u>	54.3	60.9
	LSR [5]	81.7	<u>57.9</u>	<b>21.6</b>	77.2	–	<u>35.3</u>	–	14.2	<u>35.4</u>	<u>41.7</u>	48.1	<u>56.5</u>	<u>61.6</u>
	LSR <sup>+</sup> (ours)	<u>82.5</u>	<b>58.4</b>	20.2	<u>81.3</u>	–	32.7	–	<u>15.3</u>	<b>36.7</b>	<b>42.6</b>	<b>48.7</b>	<b>57.7</b>	<b>62.1</b>
	LSR <sup>+</sup> on <i>CS-full</i>	83.3	57.7	17.0	84.3	–	37.7	–	21.5	28.6	43.5	50.2	58.8	64.2



learning rate starts from  $2.5 \times 10^{-4}$  and uses a polynomial decay of power 0.9 over 250k steps, as in [56]. We used for validation a subset of the original training set to tune the hyper-parameters of our loss components. To tackle overfitting, we used some data augmentation strategies: random left–right flipping; white point re-balancing  $\propto \mathcal{U}([-75, 75])$ ; and color jittering  $\propto \mathcal{U}([-25, 25])$  (the last two applied independently in the R, G and B channels) and random Gaussian blur [51,56]. We perform training on an NVIDIA Titan RTX, using a batch size of 2 (1 source and 1 target samples) for 24, 750 steps (*i.e.*, 10 epochs of the *Cityscapes* train set). We also exploited the validation set for early stopping.

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

## 6 Adapted-to-supervised ratio metric

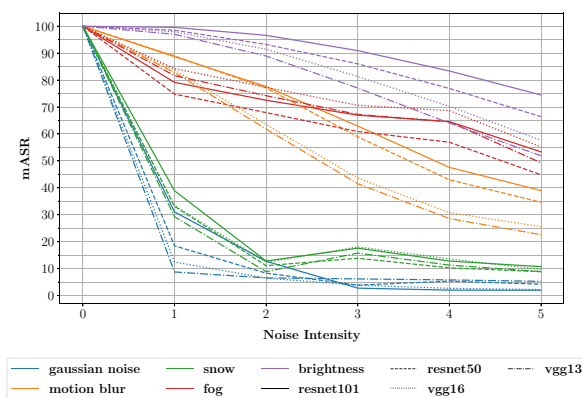
In this section, we introduce a novel measure, called mASR (*mean Adapted-to-Supervised Ratio*), in order to better evaluate the domain adaptation task with respect to the usual mIoU.

The idea behind the new metric sparks from realizing that the mIoU is missing a key component to evaluating an adaptation method: *i.e.*, it does not account for the starting accuracy of 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 through supervised learning on the target domain. We design mASR to capture the relative accuracy of an adapted architecture with respect to 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  $c \in \mathcal{C}$  ( $\text{IoU}_{adapt}^c$ ) with the results of supervised training on target data ( $\text{IoU}_{sup}^c$ ), 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 corresponding to each semantic category is inversely proportional to the accuracy of the segmentation model on it in the supervised scenario, thus giving more relevance to the most challenging classes and producing a more class-agnostic adaptation score. Furthermore, notice how the most challenging classes in driving scenarios are typically associated with small objects like traffic lights or pedestrians and bicycles, which are very critical for autonomous navigation. In this metric, higher means bet-



**Fig. 3** mASR score as a function of the injected noise intensity. Best viewed in color

ter, 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 1 allow identifying at a glance the algorithms that more faithfully match the target performance.

To validate the new metric, we used as a 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 that are likely to be encountered by an agent moving outdoors (*i.e.*, Gaussian noise, motion blur, snow, fog, brightness) and we set 5 levels of noise intensity as defined by [74]. As expected, the higher the noise intensity, the lower 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 and 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 scores.

## 7 Experimental evaluation

The qualitative and quantitative results achieved by the proposed approach (LSR<sup>+</sup>) in various driving contexts will be presented in this section, where it will be compared with several other feature-level approaches (*i.e.*, [42,47,75]), with some entropy minimization strategies (*i.e.*, [56,57]) that have a similar effect on feature distribution, and finally with the conference version of our work [5]. The key feature of these approaches is the training efficiency; indeed, the addition of such constraints does not increase the training computational complexity, differently from expensive adversarial approaches or modified architectures.

Our end-to-end method allows straightforward integration with other strategies, *e.g.*, adversarial approaches at input or output level, or entropy minimization. In order to verify such compatibility, we introduce an additional entropy minimization loss [56] in our setup. We start by considering two widely used synthetic-to-real benchmarks and a standard ResNet101 as backbone architecture obtaining the results shown in Table 1. Then, a real-to-real benchmark [18] has also been used (see Table 2). To further verify the robustness of our setup, in Table 3 we report some results using different backbones (*i.e.*, ResNet50, VGG16 and VGG13).

## 7.1 Adaptation from synthetic data to Cityscapes

When adapting source knowledge from the *GTA5* dataset to the *Cityscapes* one, our approach (LSR<sup>+</sup>) achieves a mIoU of 46.9%, with a gain of 10% compared to the baseline and of 0.9% compared to the conference version (LSR) [5], thanks to the refined space-shaping objectives. In particular, the classes which enjoyed the greatest improvements are those targeted by our approach, *i.e.*, the less common or the ones corresponding to small objects (*pole*, *t. light*, *train*, etc.). Our new pipeline improves recognition of these types of classes since we actively re-weight the labels during downsampling and employ outlier-resistant strategies, differently from the previous approach. In general, our approach outperforms all compared strategies: the only techniques able to get close to its performance are the works by [47] and [56], while the other competitors see a significant score drop. The performance improvement is quite stable across per-class IoUs and is particularly noticeable in challenging classes, like *terrain* and *t. light* where our strategy shows very high accuracy 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 that our method increases the precision of edges on *t. sign*, *t. light*, *pole* and *person* classes in both images. Furthermore, our approach is the only one to correctly classify the *bus* on the right of the first image, which is confused as *truck* by the other strategies. Importantly, we can also see the effects of our two-pass labeling (see Sect. 3.3) 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<sup>+</sup> significantly reduces the confusion between *sky* and the white building, and is able to correctly recognize the *car*, which was confused for *truck* by the old approach).

In the *SYNTHIA* to *Cityscapes* setup, LSR<sup>+</sup> surpasses its conference version (LSR) by about 1% of mIoU in the 16-class setup and by 0.6% in the 13-class one, achieving a final score of 42.6 and 48.7%, respectively. Again, the most noticeable improvement can be found in the difficult classes, such

as *wall*, *pole* or *bicycle*. Such gains over the older approach confirm once again the effectiveness of our new formulation. Similarly to the *GTA5-to-Cityscapes* experiments, also in this setting LSR<sup>+</sup> outperforms all the other competitors, with a slight margin of 1% on average with respect to [47] and a larger one (more than 3%) compared to 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<sup>+</sup> 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 in the prediction of such classes and, fundamentally, of the *road* in the foreground (confused for *car* and *bicycle* by the competitors).

## 7.2 Adaptation from Cityscapes to Cross-City

Besides using synthetic data, another key requirement is the capability of adapting networks trained on road scenes coming from certain geographical areas to other regions. However, the great variability of road scenes across the world limits the wide application of locally trained models on a global scale. To investigate the capability of our approach to cope with this problem, we evaluate the performance on the *Cross-City* real-to-real benchmark in Table 2. This benchmark is comprised of 4 cities with a completely different type of urban setting: 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<sup>+</sup> 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 [42] for Rome and Taipei and [47] for Rio and Tokyo, underlining the unstable performances of many approaches usually associated with this benchmark.

This adaptation setup is particularly suited for evaluating approaches since it presents data with varying degrees of domain shift. The fact that our approach is able to achieve good performance in all settings highlights how the new space-shaping objectives are able to work effectively even in the presence of a significant distribution shift, thanks to the updated feature vector estimation. 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

**Table 2** Quantitative results on the *Cross-City* real-to-real benchmark. (r) indicates that the strategy was retrained, starting from the official code. Best in **bold**, runner-up underlined

Target city	Configuration	Road	Sidewalk	Building	Traffic light	Traffic sign	Vegetation		
Rome	Source only [56]	85.0	34.7	86.4	17.5	39.0	84.9		
	<i>Cross-City</i> [18]	79.5	29.3	84.5	0.0	22.2	80.6		
	ASN (feat) [42]	83.9	34.2	<b>88.3</b>	18.8	40.2	<b>86.2</b>		
	MaxSquareIW [56] (r)	<b>86.2</b>	<b>37.8</b>	86.4	22.3	39.5	85.4		
	UDA OCE [47] (r)	<u>85.6</u>	<u>35.0</u>	87.9	<u>23.1</u>	<u>42.0</u>	<u>85.9</u>		
	LSR <sup>+</sup> (ours)	83.4	34.5	<u>88.1</u>	<b>29.0</b>	<b>44.5</b>	85.5		
Rio	Source only [56]	74.2	42.2	84.0	12.1	20.4	78.3		
	<i>Cross-City</i> [18]	74.2	43.9	79.0	2.4	7.5	77.8		
	ASN (feat) [42]	76.2	44.7	<u>84.6</u>	9.3	<b>25.5</b>	<b>81.8</b>		
	MaxSquareIW [56] (r)	<b>79.5</b>	<u>50.7</u>	84.5	<b>14.9</b>	17.7	80.8		
	UDA OCE [47] (r)	<u>78.9</u>	48.5	<b>85.3</b>	<u>14.2</u>	<u>24.4</u>	<u>81.3</u>		
	LSR <sup>+</sup> (ours)	<b>79.5</b>	<b>52.2</b>	83.7	10.2	23.1	79.3		
Tokyo	Source only [56]	81.4	28.4	<u>78.1</u>	14.5	19.6	81.4		
	<i>Cross-City</i> [18]	83.4	<b>35.4</b>	72.8	12.3	12.7	77.4		
	ASN (feat) [42]	81.5	26.0	77.8	<b>17.8</b>	<b>26.8</b>	<u>82.7</u>		
	MaxSquareIW [56] (r)	84.1	32.9	76.7	11.3	23.8	82.3		
	UDA OCE [47] (r)	<b>85.0</b>	33.3	77.9	8.5	<u>25.5</u>	82.5		
	LSR <sup>+</sup> (ours)	<u>84.2</u>	<u>34.6</u>	<b>78.2</b>	<u>16.8</u>	22.6	<b>83.3</b>		
Taipei	Source only [56]	<b>82.6</b>	<b>33.0</b>	<u>86.3</u>	16.0	<b>16.5</b>	<u>78.3</u>		
	<i>Cross-City</i> [18]	78.6	28.6	80.0	13.1	7.6	68.2		
	ASN (feat) [42]	81.7	29.5	85.2	<b>26.4</b>	<u>15.6</u>	76.7		
	MaxSquareIW [56] (r)	80.9	31.3	83.3	12.9	13.4	75.4		
	UDA OCE [47] (r)	81.4	30.1	84.3	16.7	13.4	75.4		
	LSR <sup>+</sup> (ours)	<u>81.8</u>	<u>32.9</u>	<b>86.8</b>	<u>19.1</u>	14.2	<b>79.3</b>		
Target city	Configuration	Sky	Person	Rider	Car	Bus	Motorbike	Bicycle	mIoU
Rome	Source only [56]	85.4	43.8	15.5	81.8	46.3	38.4	4.8	51.0
	<i>Cross-City</i> [18]	82.8	29.5	13.0	71.7	37.5	25.9	1.0	42.9
	ASN (feat) [42]	<u>93.1</u>	47.8	21.7	80.9	47.8	48.3	8.6	53.8
	MaxSquareIW [56] (r)	84.0	<u>49.5</u>	21.2	82.7	<b>55.3</b>	<u>48.5</u>	<b>9.5</b>	54.5
	UDA OCE [47] (r)	89.2	49.3	<u>24.3</u>	<u>82.8</u>	<u>48.8</u>	<u>48.5</u>	<u>9.0</u>	<u>54.7</u>
	LSR <sup>+</sup> (ours)	<b>93.9</b>	<b>51.9</b>	<b>31.3</b>	<b>83.2</b>	44.7	<b>51.5</b>	8.8	<b>56.2</b>
Rio	Source only [56]	<b>87.9</b>	50.1	25.6	<b>76.6</b>	<u>40.0</u>	27.6	17.0	48.9
	<i>Cross-City</i> [18]	69.5	39.3	10.3	67.9	<b>41.2</b>	27.9	10.9	42.5
	ASN (feat) [42]	<u>87.3</u>	55.3	32.7	74.3	28.9	<b>43.0</b>	27.6	51.6
	MaxSquareIW [56] (r)	85.7	54.5	29.6	<u>75.1</u>	37.0	40.6	24.5	51.9
	UDA OCE [47] (r)	87.0	<u>55.9</u>	<u>36.2</u>	74.3	29.7	<u>41.8</u>	27.9	<b>52.7</b>
	LSR <sup>+</sup> (ours)	82.3	<b>59.8</b>	<b>40.0</b>	75.0	23.0	<b>43.0</b>	<b>29.0</b>	<u>52.3</u>
Tokyo	Source only [56]	86.5	51.9	22.0	70.4	<u>18.2</u>	<u>22.3</u>	46.4	47.8
	<i>Cross-City</i> [18]	64.3	42.7	21.5	64.1	<b>20.8</b>	8.9	40.3	42.8
	ASN (feat) [42]	<b>90.9</b>	<u>55.8</u>	<b>38.0</b>	72.1	4.2	<b>24.5</b>	50.8	<u>49.9</u>
	MaxSquareIW [56] (r)	87.4	55.3	30.0	72.0	8.6	18.9	47.1	48.5
	UDA OCE [47] (r)	<u>89.4</u>	<b>56.1</b>	29.2	<b>72.4</b>	2.1	12.3	41.9	47.4
	LSR <sup>+</sup> (ours)	89.3	55.0	<u>33.2</u>	72.0	8.6	20.5	<b>52.2</b>	<b>50.0</b>

**Table 2** continued

Target city	Configuration	Sky	Person	Rider	Car	Bus	Motorbike	Bicycle	mIoU
Taipei	Source only [56]	83.3	26.5	8.4	70.7	36.1	47.9	15.7	46.3
	Cross-City [18]	82.1	16.8	9.4	60.4	34.0	26.5	9.9	39.6
	ASN (feat) [42]	91.7	31.0	<b>12.5</b>	<u>71.5</u>	41.1	47.3	27.7	<u>49.1</u>
	MaxSquareIW [56] (r)	89.5	31.8	3.9	69.0	<b>44.3</b>	<u>49.4</u>	<b>33.3</b>	47.6
	UDA OCE [47] (r)	<b>91.9</b>	<u>32.5</u>	4.6	71.0	<u>41.4</u>	48.0	<b>33.3</b>	48.0
	LSR <sup>+</sup> (ours)	<u>91.8</u>	<b>35.1</b>	<u>11.6</u>	<b>72.8</b>	33.8	<b>58.7</b>	<u>31.6</u>	<b>50.0</b>

(increase of 9.3% of IoU). Qualitative results on this benchmark are presented in Fig. 5. From a visual inspection of the images, we can see an overall increase in the discrimination of the object borders, particularly for classes such as *car*, *road*, *building*, *vegetation* and *person*. In Rome, we see how LSR<sup>+</sup> is the only strategy that correctly identifies the rider behind the cars in the second image. In Rio, our architecture significantly reduces the amount of confusion regarding the building on the left of the second image. Again, in Tokyo, we note how LSR<sup>+</sup> is the only technique able to recognize the traffic sign on the right of the first image. Finally, in Taipei, we see how our approach is the only to correctly identify the person and motorcycle in the second image.

### 7.3 Results with different backbones

Table 3 shows the performance of our strategy on *GTAV* → *Cityscapes* using multiple encoder–decoder backbones in order to evaluate the generalization properties of the approach to different network architectures. Here we can see how LSR<sup>+</sup> outperforms the source-only models (*i.e.*, without adaptation) by 13.3, 12.7 and 7.8% using ResNet50, VGG16 and VGG13, 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 on 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 mean values of the other strategies are 48.5 and 42.2%, and the standard deviations are 2.5 and 31.3%, respectively.) We can see how the consistent performance of our approach is preserved even in the class-wise IoU scores, particularly in the *train* class, where we significantly outperform the competitors, gaining an average of 7% of IoU across the four backbones with respect to the second-best strategy.

## 8 Ablation studies

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

are reported in Tables 4 and 5. In the first, we evaluate our strategy by removing each constraint independently and evaluating the impact on the score, while also comparing the results with [5], highlighting the improvement over the older version of the approach. In the second, we report the performance attained when sweeping across values of each loss learning rate while keeping the optimum value for the other two losses. Several conclusions can be drawn from the tables: firstly, the removal of any of our losses incurs in a noticeable drop in performance ranging from 0.9% mIoU to around 2%, meaning that the three objectives should always be employed together, since only when working together they can effectively shape the latent space into the desired arrangement; secondly, the novel implementation offers significantly reduced space-shaping overlap with the minimum entropy constraint, leading to a higher score when  $\mathcal{L}_{EM}$  is included in the pipeline; thirdly, all three objectives offer an improvement over the baseline score regardless of the chosen weighting factor (the worst scores are 43.9, 44.6 and 45.7% for  $\mathcal{L}_C$ ,  $\mathcal{L}_P$  and  $\mathcal{L}_N$ , respectively, which correspond to improvements of 1.1, 1.8 and 2.9%).

### 8.1 Analysis of the latent space regularization

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

#### 8.1.1 Two-pass prototypes and clustering

To investigate the semantic feature representation learning produced by our approach, we computed a shared t-SNE [76] embedding of 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 Sect. 3.3. In the left plot of Fig. 6, we

**Table 3** Quantitative results with different backbones,  $GTAV \rightarrow Cityscapes$  setup. (r) indicates that the strategy was retained, starting from the official code

Backbone	Configuration	Road	Sidewalk	Building	Wall <sup>1</sup>	Fence <sup>1</sup>	Pole <sup>1</sup>	Traffic Light	Traffic Sign	Vegetation	Terrain
ResNet101	Target only	96.5	73.8	88.4	42.2	43.7	40.7	46.1	58.6	88.5	54.9
	Source only [47]	71.4	15.3	74.0	21.1	14.4	22.8	33.9	18.6	80.7	20.9
	MaxSquareIW [56] (r)	87.7	25.2	82.9	30.9	24.0	29.0	35.4	24.2	84.2	38.2
	UDA OCE [47] (r)	89.4	30.7	82.1	23.0	22.0	29.2	37.6	31.7	83.9	37.9
	LSR <sup>+</sup> (ours)	88.9	26.6	82.0	21.0	24.4	30.1	41.1	27.0	84.7	42.7
ResNet50	Target Only	97.0	76.6	88.6	48.0	43.5	45.0	47.3	61.4	89.2	52.8
	Source only	57.4	8.9	74.7	13.8	21.7	23.0	24.2	11.6	73.5	10.3
	MaxSquareIW [56] (r)	83.6	17.3	79.4	20.8	17.9	24.4	27.2	15.6	80.1	31.5
	UDA OCE [47] (r)	83.9	14.4	80.0	22.8	20.8	24.7	26.4	16.5	80.3	31.7
	LSR <sup>+</sup> (ours)	88.6	31.1	79.5	23.9	24.1	26.8	30.8	15.9	84.5	33.0
VGG16	Target only	96.5	73.7	86.8	39.3	41.2	35.2	40.5	51.4	87.4	49.4
	Source only	26.5	13.3	45.1	6.0	15.2	16.5	21.3	8.5	78.0	8.3
	MaxSquareIW [56] (r)	81.4	20.0	75.4	19.4	19.1	16.1	24.4	7.9	78.8	22.9
	UDA OCE [47] (r)	86.0	13.5	79.4	20.4	18.5	21.5	27.6	15.2	80.8	21.9
	LSR <sup>+</sup> (ours)	88.9	32.2	79.3	25.9	23.6	26.9	30.6	10.7	83.3	33
VGG13	Target only	96.5	74.3	86.5	34.2	41.1	35.9	39.5	51.7	87.4	52.8
	Source only	62.3	15.1	67.8	12.1	29.8	16.6	19.1	6.5	75.8	12.8
	MaxSquareIW [56] (r)	82.3	3.8	78.8	16.2	31.4	12.1	18.5	4.7	79.9	28.6
	UDA OCE [47] (r)	0.7	0.3	68.0	0.1	11.9	6.0	18.4	4.7	76.1	10.9
	LSR <sup>+</sup> (ours)	86.4	33.1	79.8	26.2	17.6	23.6	23.5	7.6	80.4	27.9

Table 3 continued

Backbone	Configuration	Sky	Person	Rider	Car	Truck	Bus	Train	Motorbike	Bicycle	mIoU	mASR
ResNet101	Target only	91.9	68.7	46.2	90.7	68.8	69.9	48.8	47.6	64.5	64.8	100
	Source only [47]	68.5	56.6	27.1	67.4	32.8	5.6	7.7	28.4	33.8	36.9	54.0
	MaxSquareIW [56]	79.2	59.0	27.7	79.5	34.6	44.2	7.5	31.1	40.3	45.5	62.2
	UDA OCE [47]	78.3	60.7	27.4	84.6	37.6	44.7	7.3	26.0	38.9	45.9	67.3
	LSR <sup>+</sup> (ours)	80.1	63.0	26.4	83.1	30.4	44.3	16.8	35.8	42.4	<b>46.9</b>	<b>69.5</b>
ResNet50	Target Only	92.0	69.0	46.9	90.8	58.2	69.6	49.2	48.9	65.0	65.2	100
	Source only	66.8	48.6	13.1	31.6	16.6	8.5	0.0	15.8	5.0	27.6	39.1
	MaxSquareIW [56]	79.9	58.6	21.3	67.5	25.1	12.8	1.8	19.7	14.3	36.8	52.0
	UDA OCE [47]	80.1	57.9	21.6	72	26.7	14.5	3.6	13.2	3.6	36.6	51.7
	LSR <sup>+</sup> (ours)	75.7	55.0	17.6	80.0	23.4	34.6	20.3	17.5	15.4	<b>40.9</b>	<b>58.6</b>
VGG16	Target only	89.1	63.8	40.5	88.8	46.2	63.5	37.2	41.4	60.6	59.6	100
	Source only	59.7	45.0	10.5	69.1	22.8	17.9	0.0	16.4	2.7	25.4	38.7
	MaxSquareIW [56]	65.9	45.0	12.3	74.6	16.1	10.3	0.2	11.3	1.0	31.7	46.9
	UDA OCE [47]	72.6	46.3	18.1	80.0	16.9	13.1	1.0	14.6	2.0	34.2	51.5
	LSR <sup>+</sup> (ours)	73.1	47.6	20.5	81.3	18.2	10.7	0.1	14.7	5.9	<b>37.2</b>	<b>57.2</b>
VGG13	Target only	89.2	63.1	39.1	88.3	44.2	61.9	43.6	40.1	61.5	59.5	100
	Source only	75.5	48.5	4.6	60.8	16.2	3.6	0.0	12.8	1.8	28.5	42.4
	MaxSquareIW [56]	74.9	42.6	2.8	79.4	23.2	10.7	0.0	10.9	0.3	31.6	46.7
	UDA OCE [47]	75.2	17.5	2.2	22.0	0.4	0.6	0.0	3.9	0.0	16.8	23.3
	LSR <sup>+</sup> (ours)	70.8	50.5	12.0	82.0	24.9	26.3	2.8	12.0	3.0	<b>36.3</b>	<b>55.5</b>

Best values are in bold

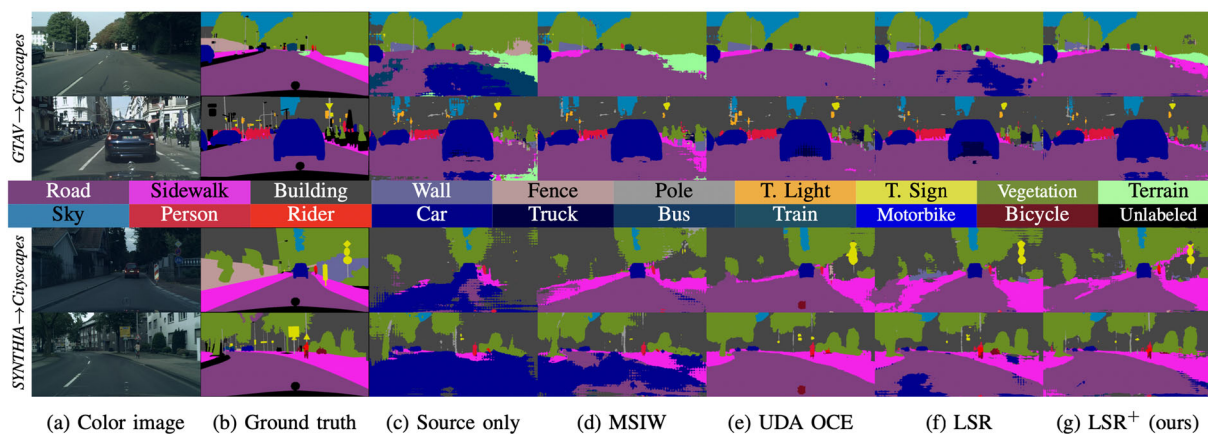


Fig. 4 Qualitative results sampled from the *Cityscapes* validation split. Best viewed in color

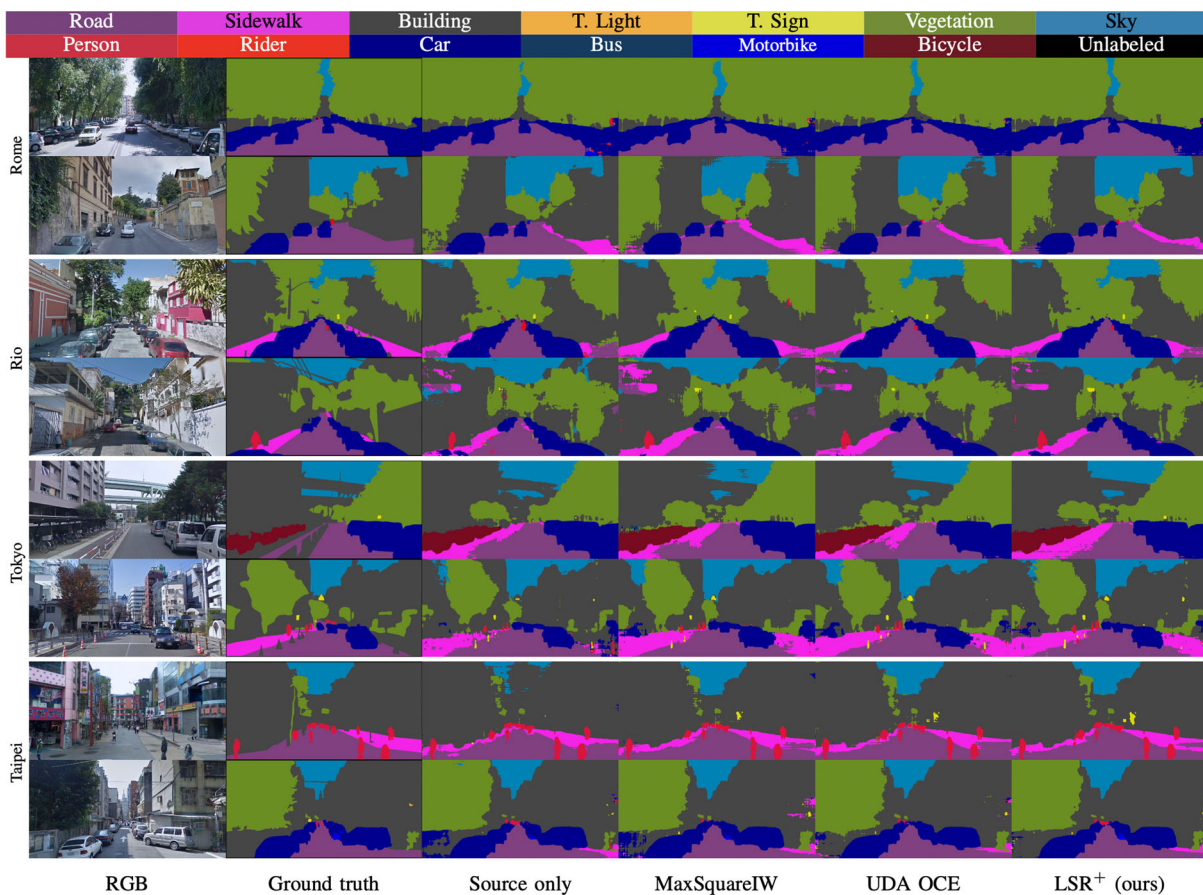


Fig. 5 Qualitative results on the Cross-City benchmark. Best viewed in color

report the learned prototype trajectory embeddings and on the right the respective feature vectors. Here we can appreciate how prototypes get farther 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 effec-

tiveness of our clustering strategy, which is able to shift the target feature distribution closer to the source one.

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. 8. Our strategy significantly improves the cluster separation in the embedded


**Table 4** Ablation studies, mIoU and mASR scores comparison when removing any of the losses. Implementations of losses from [5] are compared with the new ones in this work.

$\mathcal{L}_C$		$\mathcal{L}_P$		$\mathcal{L}_N$		$\mathcal{L}_{EM}$	mIoU	mASR
[5]	Ours	[5]	Ours	[5]	Ours			
							42.8	64.4
		✓		✓		✓	44.8	–
			✓		✓	✓	44.9	66.3
✓				✓		✓	44.9	–
	✓				✓	✓	45.3	66.7
✓		✓				✓	45.2	–
	✓		✓			✓	46.0	68.3
✓		✓		✓			44.2	–
	✓		✓		✓		44.5	66.1
✓		✓		✓		✓	46.0	67.7
	✓		✓		✓	✓	<b>46.9</b>	<b>69.5</b>

Best values are in bold

**Table 5** Ablation on the weighting factors of the losses, we fix the best value for two weighting coefficients and sweep values for the third one.  $\lambda_{EM}$  is fixed to the best value (i.e., 0.1, according to [56])

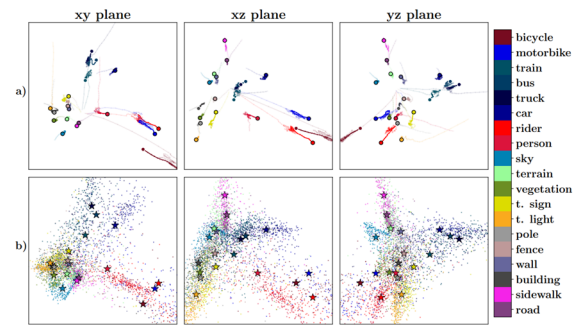
Loss	Off	$\div 10$	$\div 2$	Best	$\times 2$	$\times 10$
$\lambda_C$	0	0.01	0.05	<b>0.1</b>	0.2	1
	44.9	45.3	46.0	<b>46.9</b>	45.5	43.9
$\lambda_P$	0	0.02	0.04	<b>0.2</b>	0.4	2
	45.3	45.5	45.3	<b>46.9</b>	45.6	44.6
$\lambda_N$	0	0.0025	0.0125	<b>0.025</b>	0.05	0.25
	46.0	45.7	45.8	<b>46.9</b>	46.4	46.3




**Fig. 6** t-SNE embedding of the target feature vectors. On the left: trajectories of prototypes sampled over 200 training steps. On the right: features produced by the final model embedded according to the shared t-SNE projection. Best viewed in color

space and increases the spacing between clusters belonging to different classes, promoting features’ disentanglement. This cross-talk reduction is also reflected in the decreased probability of confusing visually similar classes (e.g., the *truck* class with the *bus* and *train* ones).

Finally, PCA embeddings are reported in Fig. 7 to evaluate the effect of latent-spacing techniques when projected to a



**Fig. 7** Prototypes trajectories and target feature vectors projected via PCA. Projection is three-dimensional; here, we report the three *xy*, *xz* and *yz* planes. Best viewed in color



(a) Baseline (b) LSR (c) LSR<sup>+</sup>

**Fig. 8** t-SNE embeddings of the normalized feature vectors. Best viewed in color



(a) Nearest (b) Histogram (c) Weighted Histogram


**Fig. 9** Sample image downsampled nearest (left), frequency-aware [5] (middle) or weighted frequency-aware (LSR<sup>+</sup>). Best viewed in color

lower dimension via a linear function, which confirms the findings highlighted in the t-SNE embeddings.


### 8.1.2 Weighted histogram-aware downsampling

In this work, we extended the scheme proposed in [5] by adding class weights inversely proportional to the class frequency in the training dataset (see Sect. 3). Our goal is to provide labeling only to spatial locations in feature maps where a clear class association can be performed, by relying on a frequency-aware scheme. By doing so, we seek the disentanglement of activations belonging to different classes, even when their feature vectors are neighbors in a given label map. This effect can be noted in Fig. 9, where our downsampling algorithms enhanced with frequency awareness are able to identify some feature locations close to class edges as *unlabeled* in the downsampled label map (middle and right), keeping only faithful features. As expected, class-weighting (right plot of Fig. 9) promotes rarer classes at the feature level






**Fig. 10** Class frequency of the downsampled feature-level segmentation maps



**Fig. 11** Mean inter-prototype angle



**Fig. 12** Average channel entropy

compared to the version without it [5] (middle plot of Fig. 9): for instance, compare the traffic sign (in yellow).

Further evidence of this can be found in the class distribution of segmentation maps (computed after their downsampling to the latent space spatial resolution), which we reported in Fig. 10 for our weighted histogram-aware scheme, the previous unweighted histogram-aware scheme of the conference version [5] 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 overpenalized by the previous approach [5], where all classes were treated equally, regardless of their occurrence.

### 8.1.3 Perpendicularity

is analyzed in Fig. 11 where we display the average angular distance between each prototype and all the remaining ones. Our goal is to achieve prototype perpendicularity, such that we minimize the overlap (*i.e.*, cross-talk) among dis-

tinct semantic categories over feature activations. By the red dashed line, we highlight the upper bound to the angular distance, which is set to 90 degrees since we assume feature vectors to have nonnegative entries. From the figure, it emerges clearly that LSR-based approaches increase the inter-prototypical angle and that LSR<sup>+</sup> makes prototypes even more orthogonal with an improvement of more than 2 degrees on average.

### 8.1.4 Norm alignment

is analyzed in Fig. 12, where we show the mean channel entropy for each class. We observe that the entropy corresponding to feature vectors produced by LSR<sup>+</sup> is significantly reduced, meaning that features are characterized by more relevant peaks and fewer poorly activated channels.

## 9 Conclusions

In this work, we tackled domain adaptation of road scene segmentation models by introducing a set of latent space regularization strategies for unsupervised domain adaptation. We improved domain invariance using different latent space-shaping constraints (*i.e.*, class clustering, class 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 outperformed state-of-the-art methods in feature-level adaptation on two widely used synthetic-to-real road scene benchmarks and in real-to-real setups, paving the way to a new set of feature-level adaptation strategies capable to improve the discrimination ability of road scene understanding approaches.

Future work will focus on designing novel feature-level techniques and on evaluating their capability of generalizing to various tasks in driving scenarios. The adaptation from multiple source domains to multiple target ones will also be considered together with the application also to multimodal data (e.g., LIDARs or depth cameras) mounted on cars.

**Funding** Open access funding provided by Università degli Studi di Padova within the CRUI-CARE Agreement. This work has also been partially funded by the University of Padova SID2020 project “Semantic Segmentation in the Wild” and by the Italian Ministry for Education (MIUR) under the “Departments of Excellence” initiative (Law 232/2016).

**Data Availability** All datasets used in the work are publicly available and can be accessed as described in each of the referenced presentation papers.

## Declarations

**Conflict of interest** The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Rizzoli, G., Barbato, F., Zanuttigh, P.: Multimodal semantic segmentation in autonomous driving: a review of current approaches and future perspectives. *Technologies* **10**(4), 90 (2022). <https://doi.org/10.3390/technologies1004090>
- Toldo, M., Maracani, A., Michieli, U., Zanuttigh, P.: Unsupervised domain adaptation in semantic segmentation: a review. *Technologies* **8**(2), 35 (2020)
- Bengio, Y., Courville, A., Vincent, P.: Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1798–1828 (2013)
- Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 580–587 (2014)
- Barbato, F., Toldo, M., Michieli, U., Zanuttigh, P.: Latent space regularization for unsupervised domain adaptation in semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (2021)
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The Cityscapes dataset for semantic urban scene understanding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3213–3223 (2016)
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3431–3440 (2015)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-assisted Intervention*, Springer, pp. 234–241 (2015)
- Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2881–2890 (2017)
- Chen, L., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587* (2017)
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFS. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**, 834–848 (2018)
- Chen, L., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: *Proceedings of the European Conference on Computer Vision*, pp. 833–851 (2018)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030* (2021)
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *Adv. Neural. Inf. Process. Syst.* **34**, 12077–12090 (2021)
- Neuhold, G., Ollmann, T., Rota Bulo, S., Kotschieder, P.: The Mapillary vistas dataset for semantic understanding of street scenes. In: *Proceedings of the International Conference on Computer Vision*, pp. 4990–4999 (2017)
- Varma, G., Subramanian, A., Nambodiri, A., Chandraker, M., Jawahar, C.: IDD: A dataset for exploring problems of autonomous navigation in unconstrained environments. In: *Proceedings of the Winter Conference on Applications of Computer Vision*, IEEE, pp. 1743–1751 (2019)
- Chen, Y., Chen, W., Chen, Y., Tsai, B., Wang, Y.F., Sun, M.: No more discrimination: cross city adaptation of road scene segmenters. In: *Proceedings of the International Conference on Computer Vision*, pp. 2011–2020 (2017)
- Brostow, G.J., Fauqueur, J., Cipolla, R.: Semantic object classes in video: a high-definition ground truth database. *Pattern Recogn. Lett.* **30**(2), 88–97 (2009)
- Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: ground truth from computer games. In: *Proceedings of the European Conference on Computer Vision*, pp. 102–118 (2016)
- Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3234–3243 (2016)
- Testolina, P., Barbato, F., Michieli, U., Giordani, M., Zanuttigh, P., Zorzi, M.: Selma: Semantic large-scale multimodal acquisitions in variable weather, daytime and viewpoints. *arXiv preprint arXiv:2204.09788* (2022)
- Sun, T., Segu, M., Postels, J., Wang, Y., Van Gool, L., Schiele, B., Tombari, F., Yu, F.: SHIFT: a synthetic driving dataset for continuous multi-task domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21371–21382 (2022)
- Lin, Z., Sun, W., Tang, B., Li, J., Yao, X., Li, Y.: Semantic segmentation network with multi-path structure, attention reweighting and multi-scale encoding. *Vis. Comput.* **39**(2), 1–12 (2022)
- Wang, K., Yang, J., Yuan, S., Li, M.: A lightweight network with attention decoder for real-time semantic segmentation. *Vis. Comput.* **38**(7), 2329–2339 (2022)
- Jiqing, C., Depeng, W., Teng, L., Tian, L., Huabin, W.: All-weather road drivable area segmentation method based on cycleGAN. *Vis. Comput.* **24**, 1–17 (2022)
- Barbato, F., Rizzoli, G., Zanuttigh, P.: Depthformer: Multimodal positional encodings and cross-input attention for transformer-based segmentation networks. *arXiv preprint arXiv:2211.04188* (2022)

28. Wang, H., Chen, Y., Cai, Y., Chen, L., Li, Y., Sotelo, M.A., Li, Z.: SFNet-N: An improved SFNet algorithm for semantic segmentation of low-light autonomous driving road scenes. *IEEE Trans. Intell. Trans. Syst.* **23**(11), 21405–17 (2022)
29. Lei, Y., Emaru, T., Ravankar, A.A., Kobayashi, Y., Wang, S.: Semantic image segmentation on snow driving scenarios. In: 2020 IEEE International Conference on Mechatronics and Automation, IEEE, pp. 1094–1100 (2020)
30. Ding, Y., Duan, Z., Li, S.: Source-free unsupervised multi-source domain adaptation via proxy task for person re-identification. *Vis. Comput.* **38**(6), 1871–1882 (2022)
31. Toldo, M., Michieli, U., Zanuttigh, P.: Learning with style: continual semantic segmentation across tasks and domains. *arXiv preprint arXiv:2210.07016* (2022)
32. Chen, Y.-C., Lin, Y.-Y., Yang, M.-H., Huang, J.-B.: Crdoco: Pixel-level domain transfer with cross-domain consistency. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1791–1800 (2019)
33. Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: Proceedings of the International Conference on Machine Learning, pp. 1994–2003 (2018)
34. Hoffman, J., Wang, D., Yu, F., Darrell, T.: FCNs in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649* (2016)
35. Murez, Z., Kolouri, S., Kriegman, D.J., Ramamoorthi, R., Kim, K.: Image to image translation for domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4500–4509 (2018)
36. Toldo, M., Michieli, U., Agresti, G., Zanuttigh, P.: Unsupervised domain adaptation for mobile semantic segmentation based on cycle consistency and feature alignment. *Image Vis. Comput.* **95**, 103889 (2020)
37. Pizzati, F., Charette, R.d., Zaccaria, M., Cerri, P.: Domain bridge for unpaired image-to-image translation and unsupervised domain adaptation. In: Proceedings of the Winter Conference on Applications of Computer Vision, pp. 2990–2998 (2020)
38. Zhou, Q., Feng, Z., Gu, Q., Pang, J., Cheng, G., Lu, X., Shi, J., Ma, L.: Context-aware mixup for domain adaptive semantic segmentation. *arXiv preprint arXiv:2108.03557* (2021)
39. Kunert, C., Schwandt, T., Nadar, C.R., Broll, W.: Neural network adaption for depth sensor replication. *Vis. Comput.* **23**, 1–11 (2022)
40. Du, L., Tan, J., Yang, H., Feng, J., Xue, X., Zheng, Q., Ye, X., Zhang, X.: SSF-DAN: separated semantic feature based domain adaptation network for semantic segmentation. In: Proceedings of the International Conference on Computer Vision, pp. 982–991 (2019)
41. Sankaranarayanan, S., Balaji, Y., Jain, A., Nam Lim, S., Chellappa, R.: Learning from synthetic data: addressing domain shift for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3752–3761 (2018)
42. Tsai, Y.-H., Hung, W.-C., Schuster, S., Sohn, K., Yang, M.-H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7472–7481 (2018)
43. Tsai, Y.-H., Sohn, K., Schuster, S., Chandraker, M.: Domain adaptation for structured output via discriminative patch representations. In: Proceedings of the International Conference on Computer Vision, pp. 1456–1465 (2019)
44. Biassetton, M., Michieli, U., Agresti, G., Zanuttigh, P.: Unsupervised domain adaptation for semantic segmentation of Urban scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 1211–1220 (2019)
45. Michieli, U., Biassetton, M., Agresti, G., Zanuttigh, P.: Adversarial learning and self-teaching techniques for domain adaptation in semantic segmentation. *IEEE Trans. Intell. Vehicles* **5**, 508–518 (2020)
46. Spadotto, T., Toldo, M., Michieli, U., Zanuttigh, P.: Unsupervised domain adaptation with multiple domain discriminators and adaptive self-training. In: Proceedings of the International Conference on Pattern Recognition (2020)
47. Toldo, M., Michieli, U., Zanuttigh, P.: Unsupervised domain adaptation in semantic segmentation via orthogonal and clustered embeddings. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 1358–1368 (2021)
48. Lee, S., Kim, D., Kim, N., Jeong, S.-G.: Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In: Proceedings of the International Conference on Computer Vision, pp. 91–100 (2019)
49. Park, S., Park, J., Shin, S., Moon, I.: Adversarial dropout for supervised and semi-supervised learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 3917–3924 (2018)
50. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Adversarial dropout regularization. In: Proceedings of the International Conference on Learning Representations (2018)
51. Zou, Y., Yu, Z., Vijaya Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: Proceedings of the European Conference on Computer Vision, pp. 289–305 (2018)
52. Zou, Y., Yu, Z., Liu, X., Kumar, B.V.K.V., Wang, J.: Confidence regularized self-training. In: Proceedings of the International Conference on Computer Vision, pp. 5982–5991 (2019)
53. Zhang, Y., David, P., Gong, B.: Curriculum domain adaptation for semantic segmentation of urban scenes. In: Proceedings of the International Conference on Computer Vision, pp. 2020–2030 (2017)
54. Zhang, Y., David, P., Foroosh, H., Gong, B.: A curriculum domain adaptation approach to the semantic segmentation of urban scenes. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**(8), 1823–41 (2020)
55. Khindkar, V., Arora, C., Balasubramanian, V.N., Subramanian, A., Saluja, R., Jawahar, C.V.: To miss-attend is to misalign! residual self-attentive feature alignment for adapting object detectors. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 3632–3642 (2022)
56. Chen, M., Xue, H., Cai, D.: Domain adaptation for semantic segmentation with maximum squares loss. In: Proceedings of the International Conference on Computer Vision, pp. 2090–2099 (2019)
57. Vu, T.-H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2517–2526 (2019)
58. Truong, T.-D., Duong, C.N., Le, N., Phung, S.L., Rainwater, C., Luu, K.: Bimal: Bijective maximum likelihood approach to domain adaptation in semantic scene segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8548–8557 (2021)
59. Kang, G., Jiang, L., Yang, Y., Hauptmann, A.G.: Contrastive adaptation network for unsupervised domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4893–4902 (2019)
60. Tian, L., Tang, Y., Hu, L., Ren, Z., Zhang, W.: Domain adaptation by class centroid matching and local manifold self-learning. *arXiv preprint arXiv:2003.09391* (2020)
61. Michieli, U., Zanuttigh, P.: 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)
62. Michieli, U., Toldo, M., Ozay, M.: Federated learning via attentive margin of semantic feature representations. *IEEE Internet Things*

- J. **10**(2), 1517–1535 (2023). <https://doi.org/10.1109/JIOT.2022.3209865>
63. Dong, N., Xing, E.P.: Few-shot semantic segmentation with prototype learning. In: Proceedings of the British Machine Vision Conference, vol. 3 (2018)
  64. Wang, K., Liew, J.H., Zou, Y., Zhou, D., Feng, J.: Panet: Few-shot image semantic segmentation with prototype alignment. In: Proceedings of the International Conference on Computer Vision, pp. 9197–9206 (2019)
  65. Liang, J., He, R., Sun, Z., Tan, T.: 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, pp. 2975–2984 (2019)
  66. Wang, Q., Breckon, T.P.: Unsupervised domain adaptation via structured prediction based selective pseudo-labeling. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 6243–6250 (2020)
  67. Choi, H., Som, A., Turaga, P.: Role of orthogonality constraints in improving properties of deep networks for image classification. arXiv preprint [arXiv:2009.10762](https://arxiv.org/abs/2009.10762) (2020)
  68. Pinheiro, P.O.: Unsupervised domain adaptation with similarity learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8004–8013 (2018)
  69. Wu, S., Zhong, J., Cao, W., Li, R., Yu, Z., Wong, H.-S.: Improving domain-specific classification by collaborative learning with adaptation networks. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 5450–5457 (2019)
  70. Yu, F., Koltun, V., Funkhouser, T.A.: Dilated residual networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 636–644 (2017)
  71. Lowe, D.G.: Object recognition from local scale-invariant features. In: Proceedings of the Seventh IEEE International Conference on Computer Vision **2**, 1150–11572 (1999). <https://doi.org/10.1109/ICCV.1999.jspta790410>
  72. Tranheden, W., Olsson, V., Pinto, J., Svensson, L.: Dacs: Domain adaptation via cross-domain mixed sampling. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 1379–1389 (2021)
  73. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
  74. Hendrycks, D., Dietterich, T.G.: Benchmarking neural network robustness to common corruptions and surface variations (2019)
  75. Li, C., Du, D., Zhang, L., Wen, L., Luo, T., Wu, Y., Zhu, P.: Spatial attention pyramid network for unsupervised domain adaptation. In: Proceedings of the European Conference on Computer Vision (2020)
  76. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**(11), 2579–2605 (2008)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Francesco Barbato** received the M.Sc. degree in Telecommunication Engineering from the University of Padova in 2020 with honors. He is currently a Ph.D. student at the same University. His research focuses on domain adaptation, multimodal learning 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 senior engineer at Samsung Research UK. 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** received the M.Sc. degree in ICT for Internet and Multimedia in 2019 at the University of Padova. In 2023 he has completed his Ph.D. at the Department of Information Engineering of the same university. During the PhD he attended an internship as Research Engineer at Samsung Research UK. He is currently working as a Research Engineer at Microtec Venice. 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.