

Language-Guided Instance-Aware Domain-Adaptive Panoptic Segmentation

Elham Amin Mansour¹ Ozan Unal^{1,2} Suman Saha^{1,3} Benjamin Bejar³ Luc Van Gool^{1,4,5}
¹ETH Zurich, ²Huawei Technologies, ³PSI, ⁴KU Leuven, ⁵INSAIT

Abstract

The increasing relevance of panoptic segmentation is tied to the advancements in autonomous driving and AR/VR applications. However, the deployment of such models has been limited due to the expensive nature of dense data annotation, giving rise to unsupervised domain adaptation (UDA). A key challenge in panoptic UDA is reducing the domain gap between a labeled source and an unlabeled target domain while harmonizing the subtasks of semantic and instance segmentation to limit catastrophic interference. While considerable progress has been achieved, existing approaches mainly focus on the adaptation of semantic segmentation. In this work, we focus on incorporating instance-level adaptation via a novel instance-aware cross-domain mixing strategy IMix. IMix significantly enhances the panoptic quality by improving instance segmentation performance. Specifically, we propose inserting high-confidence predicted instances from the target domain onto source images, retaining the exhaustiveness of the resulting pseudo-labels while reducing the injected confirmation bias. Nevertheless, such an enhancement comes at the cost of degraded semantic performance, attributed to catastrophic forgetting. To mitigate this issue, we regularize our semantic branch by employing CLIP-based domain alignment (CDA), exploiting the domain-robustness of natural language prompts. Finally, we present an end-to-end model incorporating these two mechanisms called LIDAPS, achieving state-of-the-art results on all popular panoptic UDA benchmarks. <https://github.com/elhamAm/LIDAPS>

1. Introduction

Panoptic segmentation [32] unifies semantic and instance segmentation by not only assigning a class label to each pixel but also segmenting each object into its own instance. The common approach when tackling panoptic segmentation is to deconstruct it into two subtasks and later fuse the resulting dense predictions [28, 52]. The challenge in such an approach lies in the contradictory nature of the individual task objectives [28]. While semantic segmentation seeks to map the embeddings of semantically similar object instances into a class-specific representation, in-

stance segmentation aims to learn discriminative features to separate instances from one another, resulting in conflicting gradients from two different objectives. Despite the apparent challenges, the rich semantic information with instance-level discrimination is crucial for downstream applications such as autonomous driving or AR/VR. Yet, the complexity and cost of acquiring such panoptic annotations heavily hinder the real-world deployability of such models. Furthermore, given the variance in data distribution between different domains caused by geographical changes, object selection, weather conditions, or sensor setups, models trained on previously acquired annotated data often perform poorly in new domains. This phenomenon, known as the “domain gap”, remains a further limiting factor. To this end, recent works have focused on incorporating data-efficiency into panoptic segmentation through the task of unsupervised domain adaptation (UDA) [28, 52, 79]. In contrast to the aforementioned supervised setting, in panoptic UDA a model is trained on labeled source domain images and unlabeled target domain images with supervision only available on the source domain. This allows (i) available labeled data to be used to tackle further domains (real-to-real adaptation) or (ii) to reduce annotation requirements altogether (synthetic-to-real adaptation).

However, under a panoptic UDA setting, balancing both tasks and limiting the effects that arise from the contradictory objectives becomes more challenging due to the lack of a supervisory signal on the target domain. In Tab. 1, we provide an overview of SOTA panoptic UDA methods [28, 52, 79] based on different criteria. Apart from CVRN [28] that avoids the problem by completely decoupling the two tasks and training individual networks, i.e., fully rely on task-specific representations (TR), previously proposed methods that utilize more memory-efficient unified network architecture (e.g. exploiting both shared (SR) and task-specific representations) have tackled panoptic UDA by only adapting the semantic segmentation branch to improve panoptic quality. Specifically, EDAPS [52] utilizes ClassMix [48] to generate semantically cross-domain mixed inputs that align the target domain to the source, and, UniDAformer [79] hierarchically calibrates the semantic masks across generated regions, superpixels, and pixels. Such SOTA methods for panoptic UDA are thus able



Figure 1. While previous SOTA methods for panoptic UDA such as EDAPS [52] achieve good semantic segmentation performance, they struggle to predict correct object boundaries and thus instance segmentation masks.

Table 1. Comparison of LIDAPS with SOTA on different aspects such as self-training (ST) type; ST feature space: semantic (Sem) vs. instance (Inst); shared (SR) vs. task-specific (TR) representations; sampling strategies: ClassMix [48] vs. proposed IMix (Sec. 3.3); and proposed CLIP-based domain alignment (CDA).

Method	ST	SemST	InstST	SR	TR	ClassMix	IMix	CDA
CVRN [28]	Offline	✓	✓		✓			
UniDAformer [79]	Online	✓	✓	✓				
EDAPS [52]	Online	✓		✓	✓	✓		
LIDAPS (Ours)	Online	✓	✓	✓	✓	✓	✓	✓

to learn good semantic segmentation masks in the target domain, however, are prone to predict inaccurate instance segmentation masks due to the conflicting objectives. This problem is more prominent when multiple overlapping or occluded object instances are present in a scene. An example is shown in Fig. 1 where it can be seen that while EDAPS correctly predicts the semantic segmentation masks for the “car” (top-left) and “person” (bottom-right) classes, it fails to identify individual instance boundaries resulting in the merging of objects. This limitation is expected, given the lack of adaptation for instance segmentation. In the current literature, the adaptation on an instance-level for panoptic UDA is heavily underexplored [28, 79], with no work in instance-level cross-domain mixing.

In this work, we propose a novel instance-aware mixing strategy IMix (Sec. 3.3), to improve the recognition quality of a panoptic UDA model directly. With IMix, we leverage the panoptic predictions of a model to generate a cross-domain input image consisting of high-confidence instances from the target domain pasted onto a source image and finetune itself through self-supervision. By employing target-to-source mixing, we retain the exhaustiveness of the generated panoptic pseudo-label, i.e. each object within the scene always has an associated instance label. This allows us to reduce the *confirmation bias* while directly learning target instance segmentation on a simpler source background.

While IMix enhances panoptic quality via improved instance segmentation performance, the enhancement is limited due to a drop in semantic segmentation performance (Table 3). The model finetuned with IMix becomes subject to catastrophic interference, yielding the ability to map se-

mantically similar objects into a joint embedding in favor of increased instance separability [21]. To remedy this, we propose employing CLIP-based domain alignment (CDA) to act as a regularizer on the semantic branch (Sec. 3.4). In essence, CDA continually aligns both the target and source domains with a pre-trained frozen CLIP [50] model. Specifically, we leverage the rich feature space of CLIP to construct class-wise embeddings from a set of static text prompts. We then compute their inner product with the semantic decoder features to generate per-pixel-text similarity maps following DenseCLIP [51] that are then supervised via ground truth or pseudo-target labels.

Finally, we combine our two proposed modules with a unified transformer backbone and individual task decoders to construct LIDAPS, a language-guided instance-aware domain-adapted panoptic segmentation model. Our proposed LIDAPS, while improving instance segmentation, is also able to enhance the semantic quality through CDA. For example, LIDAPS predicts correct semantic and instance segmentation masks for the motor-bike (top-right) and the rider (bottom-left), while EDAPS fails to do so.

In summary, our contributions are as follows:

1. We introduce IMix, a novel target-to-source instance-aware cross-domain mixing strategy that generates exhaustively labeled source images with target instances for improved recognition quality (i.e. reduced false positives and negatives).
2. We reduce the catastrophic forgetting that arises when training with IMix by introducing CLIP-based domain alignment (CDA) as a regularizer.
3. We combine both proposed modules to form LIDAPS, a language-guided instance-aware domain-adapted panoptic segmentation model that achieves SOTA results across multiple benchmarks.

While we propose an end-to-end model with LIDAPS, our individual contributions remain orthogonal to the development of better panoptic UDA frameworks and are model-agnostic. Furthermore, both contributions can be detached during inference and thus do not induce any memory or computational constraints during inference.

2. Related works

Panoptic segmentation [4, 9, 26, 32, 39, 64, 67, 68] is becoming increasingly important with the rise of autonomous driving and AR/VR.

Unsupervised Domain Adaptation (UDA) panoptic segmentation trains on source domain labeled data and target domain unlabeled data. To achieve good performance on the unlabeled target images, these methods incorporate UDA techniques. A common one is self-training on pseudo-labels generated from target images by a teacher network [5, 6, 15, 28, 30, 31, 72, 79] or the model itself [33, 85]. Some works also refine and stabilize pseudo-labels [15, 73, 78, 80]. Other approaches include adversarial strategies [14, 19, 20, 43, 60, 62, 63], contrastive learning [6, 10, 37], regularizers [7, 29, 57, 58, 81], multiple resolution [22, 23] and domain adaptive architecture design [44, 79]. In UDA works, only a few [28, 52, 79] address panoptic segmentation. However, unlike existing work, we explore instance-aware cross-domain mixing to adapt the instance branch while simplifying the learning of difficult target objects by pasting them onto easy-to-segment source backgrounds. Furthermore, we are the first to exploit the domain-robustness of language-vision models to further align the source and target domains for panoptic UDA.

Language in Segmentation Several segmentation studies incorporate language to enhance their performance. This trend originated with DenseCLIP [51], an extension of CLIP [50] designed for dense downstream applications. While we also leverage dense per-pixel text similarity maps similar to DenseCLIP, as opposed to applying alignment on supervised images, we utilize the maps to align both the source and target domains via ground truth and generated pseudo-labels with domain-invariant CLIP text embeddings. Importantly, unlike DenseCLIP which applies this knowledge distillation to the encoder features, we apply deep in the semantic decoder to prevent losing class-agnostic features in the shared encoder that are key for the task of instance segmentation. Open-vocabulary segmentation works also largely integrate language into their architecture [16, 40, 55, 66, 68, 70, 71, 75]. These works do not perform unsupervised domain alignment. In some previous works [27, 69, 74, 76], mask annotations are unavailable, and large vision-language models are solely relied on for knowledge distillation. In contrast, we leverage the direct supervision available from a source domain. Some domain generalization segmentation works [13, 61] also incorporate language to align their source embeddings with large language-vision embeddings to generalize to the target domain. However, these works [13, 61] do not address instance segmentation which is specifically challenging given that CLIP mainly consists of semantic knowledge. Moreover, while some works [34, 56, 77] investigate the incorporation of CLIP in UDA, only a few explore its effects in UDA seg-

mentation. For instance, Chapman *et. al* [5] uses CLIP for UDA instance segmentation on an image level. In contrast, our work utilizes CLIP in a panoptic setting and calculates the text similarity on a pixel level.

Augmented Data for Domain Adaptation A key strategy in UDA segmentation involves training on augmented images. A common approach is the stylization and augmentation of images [1, 24, 28, 36, 37, 46] or the features of source images [12, 41, 61]. Another approach is to leverage diffusion models and GANs to translate the style of source images or to synthesize training images [10, 11, 35, 38, 47, 49, 54, 65, 82]. Zhao *et. al* [82] generates images of instances using generative models and then crops the instances using pre-trained segmentation models in order to do cross-domain mixing. This is not applied in a UDA setting and furthermore uses pre-trained segmentation models while we use our own EMA teacher network to predict the pseudo-masks of the target images for cropping. An alternative mainstream tactic is cross domain mix sampling (CDMS) [3, 59, 84]. ClassMix [48], a CDMS technique, pastes pixels from half of the source image semantic classes onto a target image [2, 8, 21, 29, 59, 83]. Cardace *et. al* [3], another CDMS technique, pastes semantic masks from target to source. However, instance-aware mixing for the domain invariance enhancement of the instance decoder, specifically in panoptic segmentation models, remains largely unexplored. Lu *et. al* [42] explores instance mixing from source-to-target for UDA in action detection but neglects to refine the pseudo-masks. In contrast, we employ confidence-based thresholding to refine the pseudo-instance-masks which we find is key to reduce the confirmation bias. Furthermore, we apply the mixing in the opposite direction which yields a considerable performance gain by avoiding further bias injected due to an incomplete set of pseudo-labels arising from false negative predictions.

3. Method

3.1. Preliminary

Panoptic segmentation is commonly tackled by breaking it down into its subtasks: semantic and instance segmentation. A panoptic segmentation model is thereby trained on a panoptic segmentation loss \mathcal{L}_{pan} given by the sum of a semantic and an instance loss:

$$\mathcal{L}_{\text{pan}} = \mathcal{L}_{\text{sem}} + \mathcal{L}_{\text{inst}} \quad (1)$$

In this work, for semantic segmentation, we use pixel-wise categorical cross-entropy loss, while for instance segmentation, we follow a top-down approach, and compute RPN and RoIAlign box regression and classification losses following MaskRCNN [18]. **Panoptic UDA** is the task of transferring knowledge from a learned source domain to a target domain. In this setup, a machine learning model ϕ

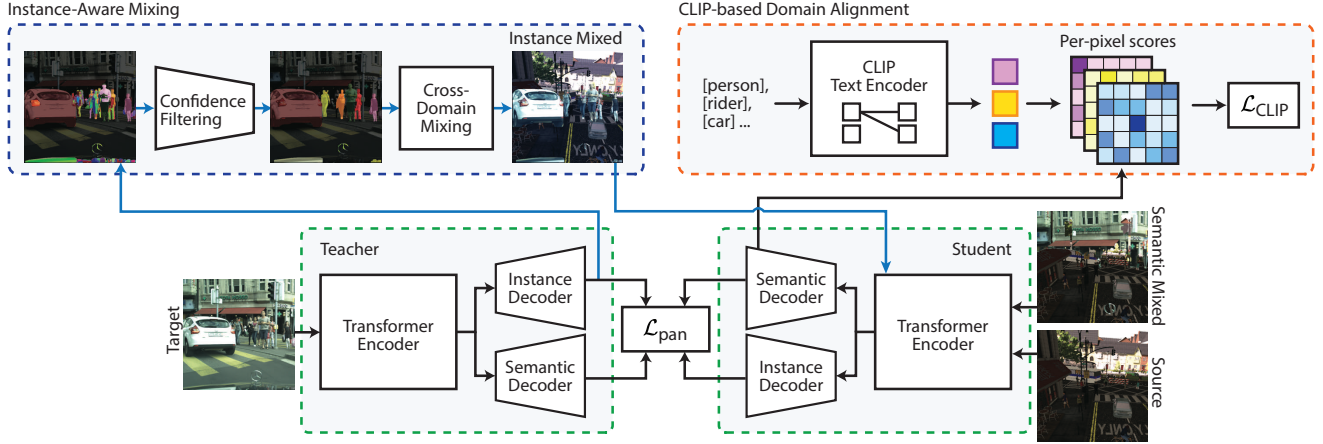


Figure 2. Illustration of the LIDAPS pipeline. (Green) The baseline panoptic UDA model is built on a mean-teacher framework and consists of a common transformer encoder and individual task decoders. The student model is supervised directly from source domain labels as well as semantically mixed inputs whose labels are generated by the teacher model. (Blue) We apply IMix to further adapt the instance segmentation branch of LIDAPS, mixing high-confidence predicted target instances with source images. Blue paths are only active when self-training with IMix is enabled. (Orange) We regularize the semantic branch via CLIP-based domain alignment that utilizes similarity maps to reduce catastrophic forgetting.

is trained on both source $\mathcal{D}^s = \{x_i^s, y_i^s\}_{i=1}^{N^s}$ and target domain images $\mathcal{D}^t = \{x_i^t\}_{i=1}^{N^t}$, with direct human annotated supervision only available on the source domain via semantic $y_{\text{sem}}^s \in \mathbb{R}^{H \times W \times C}$ and instance labels $y_{\text{inst}}^s \in \mathbb{R}^{H \times W \times B}$. Here C denotes the number of semantic classes and B denotes the number of ground truth instances given images of size $H \times W$. The naïve approach to tackling panoptic UDA is to treat the problem similar to standard supervised training, and thus only train with the supervision from the source labels via the source loss \mathcal{L}^s . However, the variance in data distribution between the source and target images, i.e. the “domain gap”, severely limits the transferability of learned knowledge across domains. Such a naïve approach consisting of only a source loss thus remains inadequate for achieving a good performance on the target domain.

Self-training is a common technique used to reduce the domain gap between source and target by leveraging a model’s own predictions to extend the supervision to the target domain [21, 25, 45, 52, 59, 80, 81, 86]. In this work, we adopt a self-training approach that entails both the supervised loss on the source domain $\mathcal{L}_{\text{pan}}^s$, and a self-supervised loss $\mathcal{L}_{\text{pan}}^{ss}$, resulting in the final training objective:

$$\arg \min_{\phi} \mathcal{L}_{\text{pan}}^s + \mathcal{L}_{\text{pan}}^{ss} \quad (2)$$

3.2. Establishing a Baseline for Panoptic UDA

In a self-training framework, a model learns from its own predictions. This however can result in *confirmation bias* as the model trains on incorrect pseudo-labels, therefore commonly, predictions are refined prior to application [15, 73]. The mean-teacher framework [58] proposes a simple but effective way to generate stabilized on-the-fly pseudo-labels

by leveraging the fact that the stochastic averaging of a model’s weights yields a more accurate model than using the final training weights directly. A mean-teacher framework is therefore built with two models, namely the student that is trained (e.g. via gradient decent), and the teacher θ whose weights are updated based on the exponential moving average (EMA) of successive student weights:

$$\phi_{t+1} \leftarrow \alpha \phi_t + (1 - \alpha) \theta_t^1 \quad (3)$$

for time step t and α that denotes the smoothing coefficient. While the mean-teacher extends the supervision to the target domain, the supervisory signal remains highly noisy and may still destabilize the training process. A common solution applied in UDA setups is to employ cross-domain mixing to generate images that contain both noisy target domain information and clean source domain ground truth annotations [2, 8, 21, 29, 59, 83]. Specifically, the teacher network θ predicts the pseudo-labels for the target image that forms the augmented input for the student via a cut-and-paste operation on the source image. Formally, given a binary mask of semantic labels to be cut $\mathbf{M}_{\text{sem}} \in \{0, 1\}^{H \times W}$ and target semantic pseudo-labels generated by the teacher model y_{sem}^t , the semantic cross-domain mixed sampling (DACS) can be defined as:

$$\begin{aligned} \tilde{x} &= \mathbf{M}_{\text{sem}} \odot x^s + (1 - \mathbf{M}_{\text{sem}}) \odot x^t \\ \tilde{y}_{\text{sem}} &= \mathbf{M}_{\text{sem}} \odot y_{\text{sem}}^s + (1 - \mathbf{M}_{\text{sem}}) \odot y_{\text{sem}}^t \end{aligned} \quad (4)$$

with \odot denoting a dot product and $\tilde{\cdot}$ indicating the mixed domain. Such DACS operations leveraging ClassMix [48]

¹We use the notation for the model and its weights interchangeably for readability.

coupled with self-training have shown significant performance gains when tackling semantic UDA [21], with the core idea stemming from consistency regularization [48, 53, 57, 58] which states that predictions for unlabelled data should be invariant to perturbations or augmentation. Firstly, the supervised semantic loss consists of :

$$\mathcal{L}_{\text{sem}}^s(\hat{y}_{\text{sem}}^s, y_{\text{sem}}^s) = - \sum_{i,j,c} (y_{\text{sem}}^s \log(\hat{y}_{\text{sem}}^s))_{i,j,c} \quad (5)$$

Formally, the self-supervised loss for the semantically adapted self-training baseline, built on a weighted cross-entropy, is given by:

$$\mathcal{L}_{\text{pan}}^{ss} = \mathcal{L}_{\text{sem}}^{ss}(\hat{y}_{\text{sem}}, \tilde{y}_{\text{sem}}) \quad (6)$$

with $\hat{\cdot}$ denoting the prediction of the model and

$$\mathcal{L}_{\text{sem}}^{ss}(\hat{y}_{\text{sem}}, \tilde{y}_{\text{sem}}) = \begin{cases} \mathcal{L}_{\text{sem}}^s(\hat{y}_{\text{sem}}, y_{\text{sem}}^s), & \text{if } \mathbf{M}_{\text{sem}}^{(h,w,c)} = 1, \\ - \sum k_{(h,w)}^t (y_{\text{sem}}^t \log(\hat{y}_{\text{sem}}))_{(h,w,c)}, & \text{otherwise} \end{cases} \quad (7)$$

with k^t the model’s estimated confidence for the pseudo-label [59]. Specifically, we apply a cross-entropy loss on pixels coming from the source image ($\mathbf{M}_{\text{sem}}^{(h,w,c)} = 1$), and apply a weighted loss on the pixels coming from the target image, supervised via the teacher generated pseudo-label y_{sem}^t . We illustrate this baseline in Fig. 2 - green.

Given that semantic segmentation forms one-half of panoptic segmentation, such a baseline approach that adapts the semantic maps between the source and target domains via DACS can significantly reduce the domain gap for panoptic segmentation. However, such an approach forgoes a crucial element of panoptic UDA altogether, adapting the instance segmentation task between two domains. In fact, DACS does not generate augmented images containing sufficient instance-specific information to adapt the instance branch. In the following section, we tackle the adaptation of instance segmentation between a source and target domain to improve panoptic segmentation performance.

3.3. Instance-Aware Mixing (IMix)

We propose a novel mixing strategy called IMix, to reduce the domain gap when tackling instance segmentation. The goal of IMix is to apply cross-domain mixed sampling while not only retaining instance-level information but also simplifying the recognition of target objects by presenting them within source environments. A sample image utilizing IMix is compared to DACS in Fig. 3.

However, mixing source and target domain information on an instance level raises a crucial challenge, stemming from how the two tasks are supervised. Unlike semantic segmentation where losses are applied on a pixel level, instance segmentation is typically supervised by the injective

function that maps the set of ground truth objects to the set of predicted instances. Therefore an instance segmentation model remains prone to confirmation bias if ground truth label exhaustiveness is not guaranteed, i.e. the model will learn to incorrectly identify objects as background if every visible object within the scene does not have an associated instance mask (see supp. for false negative example).

A mixing operation must account for such a challenge. We thus construct IMix such that it is handled from target-to-source, avoiding the incompleteness of instance labels that may emerge from false negative predictions. Formally:

$$\begin{aligned} \tilde{x} &= \mathbf{M}_{\text{inst}} \odot x^t + (1 - \mathbf{M}_{\text{inst}}) \odot x^s \\ \tilde{y}_{\text{inst}} &= \mathbf{M}_{\text{inst}} \odot y_{\text{inst}}^t + (1 - \mathbf{M}_{\text{inst}}) \odot y_{\text{inst}}^s \end{aligned} \quad (8)$$

with y_{inst}^t and y_{inst}^s denoting the target pseudo-label and source ground truth label respectively, and $\mathbf{M}_{\text{inst}} \in \{0, 1\}^{H \times W}$ the sum of binary instance masks based on the teacher’s prediction.

Specifically, we cut the instances from the teacher model’s output and paste them onto a source image, constructing the mixed pseudo-label by merging the ground truth instance labels with the teacher’s predictions. Utilizing the source image as a background ensures all visible objects have corresponding mask annotations. Furthermore, since a model learns a source domain much more efficiently thanks to direct supervision, with IMix, we simplify the recognition task of target instances by presenting them on easy-to-separate source domain environments.

However, while retaining exhaustiveness limits confirmation bias caused by false negative predictions, a self-supervised model is still prone to such effects due to false positives. In other words, if incorrect instance masks are pasted on the mixed image, the model will learn to affirm its preexisting biases, causing an increased number of false positive predictions. To minimize such cases, we propose a simple but effective confidence filtering step. We predict a confidence score alongside the instance masks of each object [18]. We apply filtering based on the predicted confidence values to redefine the joint mixing mask as:

$$\mathbf{M}_{\text{inst}} = \sum_{i \in I} \mathbb{1}[h_i^t > \tau] y_{\text{inst},i}^t \quad (9)$$

with I denoting the set of predicted instances, $y_{\text{inst},i}^t$ the predicted i ’th instance mask, h_i^t the corresponding confidence score and τ the threshold hyperparameter. Thus, our self-supervised panoptic loss from Eq. 6 can be updated as:

$$\mathcal{L}_{\text{pan}}^{ss} = \mathcal{L}_{\text{sem}}^{ss}(\hat{y}_{\text{sem}}, \tilde{y}_{\text{sem}}) + \mathcal{L}_{\text{inst}}^{ss}(\hat{y}_{\text{inst}}, \tilde{y}_{\text{inst}}) \quad (10)$$

As commonly seen in multitask frameworks, the increase in supervisory signals from one task may cause catastrophic forgetting for another, i.e. the weights in the network that




Figure 3. Our Proposed IMix cuts and pastes pseudo-instances with a confidence score above a certain threshold from source to target while in a reversed DACS, half of the semantic classes are pasted from source to target without preserving instance-level information.

are important for one task may be changed to meet the objectives of another [34]. We observe similar behavior in our training when fine-tuning LIDAPS on IMix (please refer to Sec. 4). Specifically, the performance gains of our model for panoptic segmentation are hindered by the drop in semantic quality. In the following section, we address this problem by introducing a language-based regularization for semantic segmentation.

3.4. CLIP-based Domain Alignment (CDA)

A simple but effective solution to reducing catastrophic forgetting when multitask learning is to leverage the embedding space of a pre-trained model as an anchor, i.e. the intermediate features as continual auxiliary targets, which is also commonly employed in unsupervised domain adaptation frameworks to limit overfitting onto the source domain [21, 52]. In this work we exploit both use cases for weight anchoring by relying on CLIP [50] embeddings to regularize the semantic branch of our network. CLIP is trained on a very large-scale image-text pair dataset, providing a diversified, robust world model. We argue that by semantically aligning each domain to the CLIP embedding space, we can implicitly enforce domain invariance. In other words, we train our model such that the features of a source or target image both aim to generate high similarities to a joint CLIP embedding. An illustration of CLIP-based domain alignment (CDA) can be seen in Fig. 2 - orange.

However, to be able to exploit CLIP features to avoid the divergence of semantic features of source and target images, regularization needs to be applied deep within the network. This of course imposes limitations on the expressibility of the features or effectiveness of the regularization when directly using CLIP embeddings as targets. To this end, we construct a pixel-level representation from natural language prompts following DenseCLIP [51] and only supervise the similarity to the semantic decoder features. Specifically, our CLIP-based domain alignment strategy follows two steps as illustrated in Fig. 4. We first generate class-wise mean CLIP features by mean pooling over the CLIP embeddings generated from P text prompts for C semantic classes following set precedent [51], with P denoting the number of text prompts per class (Fig. 4a-c). Each row in the resulting matrix represents a CLIP embedding that encodes meaningful

semantic information about a particular class. These embeddings act as anchors within our alignment module, with each generated semantic feature (Fig. 4d) aiming to achieve high similarity with a semantically corresponding vector. Finally, we compute the per-pixel text similarity maps σ^{sim} through the inner product of the decoder features and mean CLIP features (Fig. 4e).

Formally, the CLIP-based domain alignment loss can be stated as follows:

$$\mathcal{L}_{\text{CLIP}} = -\frac{1}{HWC} \sum_{h=1}^H \sum_{w=1}^W \sum_{c=1}^C \mathbb{1}[y_{(h,w)} = c] \log \left(\hat{y}_{(h,w,c)}^{\text{sim}} \right), \quad (11)$$

with $\mathbb{1}[\cdot]$ denoting the indicator function and \hat{y}^{sim} denoting the embedding-text similarity probability given by:

$$\hat{y}_{(h,w,c)}^{\text{sim}} = \frac{\exp \left(\sigma_{(h,w,c)}^{\text{sim}} \right)}{\sum_{c'=1}^C \exp \left(\sigma_{(h,w,c')}^{\text{sim}} \right)}. \quad (12)$$

In our proposed LIDAPS model, the CLIP loss is incorporated in our \mathcal{L}_{sem} loss with a weight loss of 1.0.

4. Experiments

We follow standard evaluation protocols and datasets for panoptic UDA following set precedent [28, 52, 79]. All reported values are the averaged scores over three runs with three different seeds (1, 2, 3). For further details including the implementation details including architecture and hyper-parameter details please see the supplementary.

4.1. Results

We compare our proposed LIDAPS with other state-of-the-art (SOTA) UDA panoptic segmentation methods on four different benchmarks including SYNTHIA \rightarrow Cityscapes (S \rightarrow C), SYNTHIA \rightarrow Mapillary Vistas (S \rightarrow M), Cityscapes \rightarrow Mapillary Vistas (C \rightarrow M) and Cityscapes \rightarrow Foggy Cityscapes (C \rightarrow F). As seen in Tab. 2, our model consistently outperforms existing works across the board, exceeding the performance of previous SOTA by up to +3.6 mPQ. Furthermore in Fig. 5, we provide qualitative results demonstrating the capabilities of LIDAPS. Compared to EDAPS [52], LIDAPS can better separate semantically similar neighboring instances by leveraging

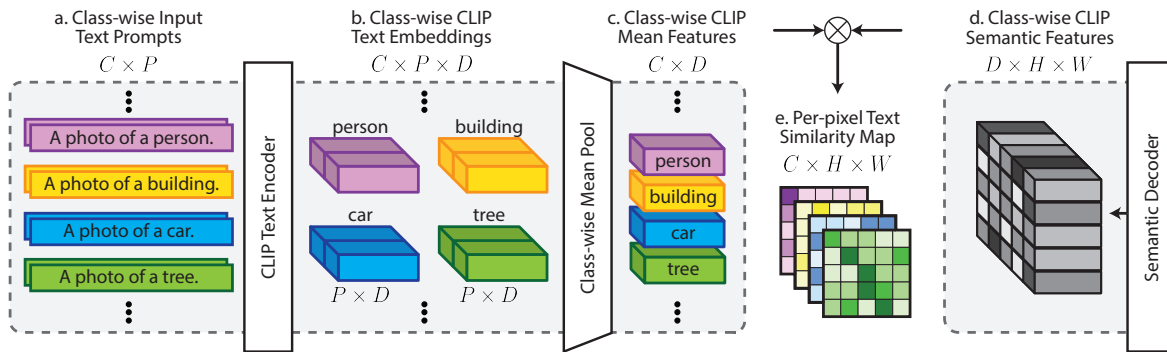


Figure 4. Pipeline used to compute the pixel-text similarity map for CLIP-based domain alignment. We generate class-wise CLIP mean features from a series of fixed text prompts (a-c). C : #semantic classes, D : dimension of text encodings, H , W : width and height of image, P : #prompts (e.g. A painting of a person). The similarity maps are given by the inner product with the semantic decoder features.




Figure 5. Qualitative results from SYNTHIA \rightarrow Cityscape comparing EDAPS [52] to our proposed LIDAPS.

instance-aware adaptation via IMix and retain its semantic quality via DCA.

4.2. Ablation Studies

In this section, we conduct ablation studies on SYNTHIA \rightarrow Cityscapes to demonstrate the effectiveness of our proposed components. The corresponding tables with the standard deviations alongside a hyper-parameter study for the confidence score threshold is found in the supplement.

Effects of Network Components: In Tab. 3, we isolate the effects of the different modules of our proposed pipeline. Starting from the baseline EDAPS*, we introduce our cross-domain instance mixing (IMix) which improves the panoptic segmentation performance (+1.3% mPQ) through significant instance segmentation gain (+5.4% mAP). Due to their contradictory goals, the improvement in instance segmentation comes in lieu of semantic performance (-0.3% mIoU) that becomes subject to catastrophic interference.

To remedy this we propose CLIP-based domain alignment (CDA). First, we separately introduce CDA to understand its isolated effects. As observed, the module aids the panoptic segmentation performance with a +1.9% mPQ improvement, which solely stems from the gains in semantic segmentation (+1.6% mIoU increase as opposed to the relatively unchanged mAP). Next, we showcase the combined effects of the two components which improve both the semantic and instance segmentation performance of our baseline, allowing LIDAPS to achieve 44.8% mPQ. As seen, the final model demonstrates significant gains in instance segmentation (+8.5 mAP) thanks to IMix, while retaining its semantic segmentation gains from the CLIP-based domain alignment (+1.6% mIoU). CDA alone increases the final mPQ more than IMix alone because semantic seg. includes 16 classes (both “thing” and “stuff” classes), while instance seg. considers only 6 thing (countable) classes where “stuff” classes (e.g. sky) occupy a significant number of pixels.

Cross-Domain Mixing: We investigate the impact of the mixing direction for IMix. Specifically, we compare the effects of source-to-target mixing, in which we cut ground truth instance masks from a source image and paste them to a target image, to our proposed target-to-source mixing, where we rely on the filtered predicted instances from the target domain to augment a source image. As seen in Tab. 4 (i), cross-mixing from source-to-target substantially degrades the panoptic performance, specifically the instance segmentation quality. Target-to-target pasting is not included because intra-domain mixing does not apply to UDA and is used in semi-supervised setting.

We argue that this performance degradation stems from the incompleteness of the target labels. When training directly on a target image, the pseudo-labels generated from the teacher model are not guaranteed to be exhaustive, i.e. false negative predictions will yield objects on the image with no

Table 2. Class-wise comparison to SOTA on four different benchmarks for UDA panoptic segmentation. Reported results are averaged over three runs with three different seeds.

Method	road	sidewalk	building	wall	fence	pole	light	sign	veg	sky	person	trailer	car	bus	m.bike	bike	mSQ	mRQ	mPQ
SYNTHIA → Cityscapes																			
FDA [72]	79.0	22.0	61.8	1.1	0.0	5.6	5.5	9.5	51.6	70.7	23.4	16.3	34.1	31.0	5.2	8.8	65.0	35.5	26.6
CRST [85]	75.4	19.0	70.8	1.4	0.0	7.3	0.0	5.2	74.1	69.2	23.7	19.9	33.4	26.6	2.4	4.8	60.3	35.6	27.1
AdvEnt [63]	87.1	32.4	69.7	1.1	0.0	3.8	0.7	2.3	71.7	72.0	28.2	17.7	31.0	21.1	6.3	4.9	65.6	36.3	28.1
CVRN [28]	86.6	33.8	74.6	3.4	0.0	10.0	5.7	13.5	80.3	76.3	26.0	18.0	34.1	37.4	7.3	6.2	66.6	40.9	32.1
UniDAformer [79]	73.7	26.5	71.9	1.0	0.0	7.6	9.9	12.4	81.4	77.4	27.4	23.1	47.0	40.9	12.6	15.4	64.7	42.2	33.0
EDAPS [52]	77.5	36.9	80.1	17.2	1.8	29.2	33.5	40.9	82.6	80.4	43.5	33.8	45.6	35.6	18.0	2.8	72.7	53.6	41.2
LIDAPS(ours)	80.8	48.8	80.8	17.6	2.5	29.9	34.6	42.9	82.8	82.9	44.4	40.5	51.7	39.2	27.4	10.7	74.4	57.6	44.8
Cityscapes → Foggy Cityscapes																			
DAF [7]	94.0	54.5	57.7	6.7	10.0	7.0	6.6	25.5	44.6	59.1	26.7	16.7	42.2	36.6	4.5	16.9	70.6	41.7	31.8
FDA [72]	93.8	53.1	62.2	8.2	13.4	7.3	7.6	28.9	50.8	49.7	25.0	22.6	42.9	36.3	10.3	15.2	71.4	43.5	33.0
AdvEnt [63]	93.8	52.7	56.3	5.7	13.5	10.0	10.9	27.7	40.7	57.9	27.8	29.4	44.7	28.6	11.6	20.8	72.3	43.7	33.3
CRST [85]	91.8	49.7	66.1	6.4	14.5	5.2	8.6	21.5	56.3	50.7	30.5	30.7	46.3	34.2	11.7	22.1	72.2	44.9	34.1
SVMIn [17]	93.4	53.4	62.2	12.3	15.5	7.0	8.5	18.0	54.3	57.1	31.2	29.6	45.2	35.6	11.5	22.7	72.4	45.5	34.8
CVRN [28]	93.6	52.3	65.3	7.5	15.9	5.2	7.4	22.3	57.8	48.7	32.9	30.9	49.6	38.9	18.0	25.2	72.7	46.7	35.7
UniDAformer [79]	93.9	53.1	63.9	8.7	14.0	3.8	10.0	26.0	53.5	49.6	38.0	35.4	57.5	44.2	28.9	29.8	72.9	49.5	37.6
EDAPS [52]	91.0	68.5	80.9	24.1	29.0	50.1	47.2	67.0	85.3	71.8	50.9	51.2	64.7	47.7	36.9	41.5	79.2	70.5	56.7
LIDAPS(ours)	92.3	70.0	83.2	23.8	31.9	56.4	47.7	68.8	86.6	82.6	72.5	53.2	53.6	68.0	56.6	45.9	80.2	73.2	59.6
SYNTHIA → Mapillary Vistas																			
FDA [72]	44.1	7.1	26.6	1.3	0.0	3.2	0.2	5.5	45.2	61.3	30.1	13.9	39.4	12.1	8.5	7.0	63.8	26.1	19.1
CRST [85]	36.0	6.4	29.1	0.2	0.0	2.8	0.5	4.6	47.7	68.9	28.3	13.0	42.4	13.6	5.1	2.0	63.9	25.2	18.8
AdvEnt [63]	27.7	6.1	28.1	0.3	0.0	3.4	1.6	5.2	48.1	66.5	28.4	13.4	40.5	14.6	5.2	3.3	63.6	24.7	18.3
CVRN [28]	33.4	7.4	32.9	1.6	0.0	4.3	0.4	6.5	50.8	76.8	30.6	15.2	44.8	18.8	7.9	9.5	65.3	28.1	21.3
EDAPS [52]	77.5	25.3	59.9	14.9	0.0	27.5	33.1	37.1	72.6	92.2	32.9	16.4	47.5	31.4	13.9	3.7	71.7	46.1	36.6
LIDAPS(ours)	76.5	25.2	64.2	14.0	0.2	29.1	35.6	35.3	72.1	94.4	33.8	18.3	50.3	33.9	19.3	5.9	73.9	47.7	38.0
Cityscapes → Mapillary Vistas																			
CRST [85]	77.0	22.6	40.2	7.8	10.5	5.5	11.3	21.8	56.5	77.6	29.4	18.4	56.0	27.7	11.9	18.4	72.4	39.9	30.8
FDA [72]	74.3	23.4	42.3	9.6	11.2	6.4	15.4	23.5	60.4	78.5	33.9	19.9	52.9	8.4	17.5	16.0	72.3	40.3	30.9
AdvEnt [63]	76.2	20.5	42.6	6.8	9.4	4.6	12.7	24.1	59.9	83.1	34.1	22.9	54.1	16.0	13.5	18.6	72.7	40.3	31.2
CVRN [28]	77.3	21.0	47.8	10.5	13.4	7.5	14.1	25.1	62.1	86.4	37.7	20.4	55.0	21.7	14.3	21.4	73.8	42.8	33.5
EDAPS [52]	58.8	43.4	57.1	25.6	29.1	34.3	35.5	41.2	77.8	59.1	35.0	23.8	56.7	36.0	24.3	25.5	75.9	53.4	41.2
LIDAPS(ours)	49.1	44.3	70.1	26.5	29.9	37.4	37.2	43.2	80.0	46.1	35.9	25.0	57.1	41.6	29.6	28.4	76.6	54.9	42.6

Table 3. Ablation study on proposed modules. Starting from a baseline EDAPS*, we introduce our instance-aware cross-domain mixing (IMix) and CLIP-based domain alignment (CDA).

EDAPS*	IMix	CDA	mSQ	mRQ	mPQ	mIoU	mAP
✓			72.3	53.3	41.0	58.0	34.1
✓	✓		73.0	54.7	42.3	57.7	39.5
✓		✓	73.9	55.0	42.9	59.6	34.4
✓	✓	✓	74.4	57.6	44.8	59.6	42.6

Table 4. Ablation study on mixing strategy for panoptic segmentation comparing (i) the mixing direction when applying IMix, (ii) the effects of ClassMix [48] when applied from target-to-source. S stands for source and T stands for target. The baseline is EDAPS*+CDA.

Method	Copy	Paste	mSQ	mRQ	mPQ	mIoU	mAP
Baseline	-	-	73.9	55.0	42.9	59.6	34.4
(i) + IMix	S	T	62.0	37.6	29.3	56.2	1.9
	T	S	74.4	57.6	44.8	59.6	42.6
(ii) + ClassMix	T	S	73.5	53.9	42.1	58.6	34.8

associated instance label (see the supplement for an example). Within a self-supervised setting, this results in the introduction of confirmation bias that severely degrades performance. On the contrary, we see significant benefits in favor of target-to-source mixing that fully avoids the incompleteness of labels, resulting in +1.8% mPQ improvement which is mainly attributed to the gains in instance segmen-

tation performance (+8.2% mAP). Furthermore in Tab. 4 (i) and (ii), we isolate the effects of the mixing task by fixing the mixing direction. Specifically, we compare an inverted ClassMix [48] that cuts and pastes semantic masks from target-to-source, to our proposed IMix strategy that works on an instance level. As seen, the inverted ClassMix slightly underperforms compared to the baseline model and significantly underperforms compared to IMix (-2.7% mPQ).

5. Conclusion

In this work, we tackle the task of UDA for panoptic segmentation. To this end, we introduce a framework LIDAPS that reduces the domain gap between target and source images by leveraging instance-aware cross-domain mixing. Specifically, we propose a novel mixing strategy IMix, that cuts and pastes confidence-filtered instance predictions from the target to the source domain, and ensures the exhaustiveness of the resulting pseudo-labels while reducing the confirmation bias. To limit the effects of catastrophic forgetting on the semantic branch, we propose a CLIP-based domain alignment mechanism that employs CLIP embeddings as anchors for both the source and target domain. Combined, our model consistently outperforms existing SOTA on popular UDA panoptic benchmarks.

Limitations: We discuss the limitations in the supplement.

References

- [1] Araslanov, N., Roth, S.: Self-supervised augmentation consistency for adapting semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15384–15394 (2021) [3](#)
- [2] Arnaudo, E., Tavera, A., Masone, C., Dominici, F., Caputo, B.: Hierarchical instance mixing across domains in aerial segmentation. *IEEE Access* **11**, 13324–13333 (2023) [3, 4](#)
- [3] Cardace, A., Ramirez, P.Z., Salti, S., Di Stefano, L.: Shallow features guide unsupervised domain adaptation for semantic segmentation at class boundaries. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1160–1170 (2022) [3](#)
- [4] Chang, C.Y., Chang, S.E., Hsiao, P.Y., Fu, L.C.: Epsnet: Efficient panoptic segmentation network with cross-layer attention fusion. In: Proceedings of the Asian Conference on Computer Vision (ACCV) (November 2020) [3](#)
- [5] Chapman, N.H., Dayoub, F., Browne, W., Lehnert, C.: Predicting class distribution shift for reliable domain adaptive object detection. *arXiv preprint arXiv:2302.06039* (2023) [3](#)
- [6] Chen, M., Zheng, Z., Yang, Y., Chua, T.S.: Pipa: Pixel- and patch-wise self-supervised learning for domain adaptive semantic segmentation. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 1905–1914 (2023) [3](#)
- [7] Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3339–3348 (2018) [3, 8](#)
- [8] Chen, Z., Ding, Z., Gregory, J.M., Liu, L.: Ida: Informed domain adaptive semantic segmentation. *arXiv preprint arXiv:2303.02741* (2023) [3, 4](#)
- [9] Cheng, B., Collins, M.D., Zhu, Y., Liu, T., Huang, T.S., Adam, H., Chen, L.C.: Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12475–12485 (2020) [3](#)
- [10] Cheng, X., Zhu, C., Yuan, L., Zhao, S.: Cross-modal domain adaptive instance segmentation in sar images via instance-aware adaptation. In: Chinese Conference on Image and Graphics Technologies. pp. 413–424. Springer (2023) [3](#)
- [11] Choi, J., Kim, T., Kim, C.: Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6830–6840 (2019) [3](#)
- [12] Fahes, M., Vu, T.H., Bursuc, A., Pérez, P., de Charette, R.: Poda: Prompt-driven zero-shot domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18623–18633 (2023) [3](#)
- [13] Gong, R., Danelljan, M., Sun, H., Mangas, J.D., Van Gool, L.: Prompting diffusion representations for cross-domain semantic segmentation. *arXiv preprint arXiv:2307.02138* (2023) [3](#)
- [14] Gong, R., Li, W., Chen, Y., Dai, D., Van Gool, L.: Dlow: Domain flow and applications. *International Journal of Computer Vision* **129**(10), 2865–2888 (2021) [3](#)
- [15] Gong, R., Wang, Q., Danelljan, M., Dai, D., Van Gool, L.: Continuous pseudo-label rectified domain adaptive semantic segmentation with implicit neural representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7225–7235 (2023) [3, 4](#)
- [16] Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921* (2021) [3](#)
- [17] Guan, D., Huang, J., Lu, S., Xiao, A.: Scale variance minimization for unsupervised domain adaptation in image segmentation. *Pattern Recognition* **112**, 107764 (2021) [8](#)
- [18] He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017) [3, 5](#)
- [19] Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: International conference on machine learning. pp. 1989–1998. Pmlr (2018) [3](#)
- [20] 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) [3](#)
- [21] Hoyer, L., Dai, D., Van Gool, L.: Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9924–9935 (2022) [2, 3, 4, 5, 6](#)
- [22] Hoyer, L., Dai, D., Van Gool, L.: Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. In: European Conference on Computer Vision. pp. 372–391. Springer (2022) [3](#)
- [23] Hoyer, L., Dai, D., Van Gool, L.: Domain adaptive and generalizable network architectures and training strategies for semantic image segmentation. *arXiv preprint arXiv:2304.13615* (2023) [3](#)
- [24] Hoyer, L., Dai, D., Wang, H., Van Gool, L.: Mic: Masked image consistency for context-enhanced domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11721–11732 (2023) [3](#)
- [25] Hoyer, L., Dai, D., Wang, Q., Chen, Y., Van Gool, L.: Improving semi-supervised and domain-adaptive semantic segmentation with self-supervised depth estimation. *International Journal of Computer Vision* (2021) [4](#)
- [26] Hu, J., Huang, L., Ren, T., Zhang, S., Ji, R., Cao, L.: You only segment once: Towards real-time panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 17819–17829 (June 2023) [3](#)
- [27] Hu, X., Zhang, K., Xia, L., Chen, A., Luo, J., Sun, Y., Wang, K., Qiao, N., Zeng, X., Sun, M., et al.: Reclip: Refine contrastive language image pre-training with source free domain

- adaptation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2994–3003 (2024) [3](#)
- [28] Huang, J., Guan, D., Xiao, A., Lu, S.: Cross-view regularization for domain adaptive panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10133–10144 (2021) [1](#), [2](#), [3](#), [6](#), [8](#)
- [29] Huo, X., Xie, L., Hu, H., Zhou, W., Li, H., Tian, Q.: Domain-agnostic prior for transfer semantic segmentation. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 7075–7085 (2022) [3](#), [4](#)
- [30] Huo, X., Xie, L., Zhou, W., Li, H., Tian, Q.: Focus on your target: A dual teacher-student framework for domain-adaptive semantic segmentation. arXiv preprint arXiv:2303.09083 (2023) [3](#)
- [31] Kennerley, M., Wang, J.G., Veeravalli, B., Tan, R.T.: 2pnet: Two-phase consistency training for day-to-night unsupervised domain adaptive object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11484–11493 (2023) [3](#)
- [32] Kirillov, A., He, K., Girshick, R., Rother, C., Dollar, P.: Panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019) [1](#), [3](#)
- [33] Kothandaraman, D., Chandra, R., Manocha, D.: Ss-sfda: Self-supervised source-free domain adaptation for road segmentation in hazardous environments. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3049–3059 (2021) [3](#)
- [34] Lai, Z., Vesdapunt, N., Zhou, N., Wu, J., Huynh, C.P., Li, X., Fu, K.K., Chuah, C.N.: Padclip: Pseudo-labeling with adaptive debiasing in clip for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16155–16165 (October 2023) [3](#), [6](#)
- [35] Lee, K.H., Ros, G., Li, J., Gaidon, A.: Spigan: Privileged adversarial learning from simulation. arXiv preprint arXiv:1810.03756 (2018) [3](#)
- [36] Lee, S., Seong, H., Lee, S., Kim, E.: Wildnet: Learning domain generalized semantic segmentation from the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9936–9946 (2022) [3](#)
- [37] Li, T., Roy, S., Zhou, H., Lu, H., Lathuilière, S.: Contrast, stylize and adapt: Unsupervised contrastive learning framework for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4868–4878 (2023) [3](#)
- [38] Li, Y., Yuan, L., Vasconcelos, N.: Bidirectional learning for domain adaptation of semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6936–6945 (2019) [3](#)
- [39] Li, Z., Wang, W., Xie, E., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P., Lu, T.: Panoptic segformer: Delving deeper into panoptic segmentation with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1280–1289 (June 2022) [3](#)
- [40] Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7061–7070 (2023) [3](#)
- [41] Liu, Y., Zhang, S., Li, Y., Yang, J.: Learning to adapt via latent domains for adaptive semantic segmentation. Advances in Neural Information Processing Systems **34**, 1167–1178 (2021) [3](#)
- [42] Lu, Y., Singh, G., Saha, S., Van Gool, L.: Exploiting instance-based mixed sampling via auxiliary source domain supervision for domain-adaptive action detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 4145–4156 (2023) [3](#)
- [43] Luo, Y., Zheng, L., Guan, T., Yu, J., Yang, Y.: Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2507–2516 (2019) [3](#)
- [44] Mascaro, R., Teixeira, L., Chli, M.: Domain-adaptive semantic segmentation with memory-efficient cross-domain transformers. In: The 34th British Machine Vision Conference (BMVC) (2023) [3](#)
- [45] Mei, K., Zhu, C., Zou, J., Zhang, S.: Instance adaptive self-training for unsupervised domain adaptation. In: Proceedings of the European Conference on Computer Vision. pp. 415–430. Springer (2020) [4](#)
- [46] Melas-Kyriazi, L., Manrai, A.K.: Pixmatch: Unsupervised domain adaptation via pixelwise consistency training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12435–12445 (2021) [3](#)
- [47] Niemeijer, J., Schwonberg, M., Termöhlen, J.A., Schmidt, N.M., Fingscheidt, T.: Generalization by adaptation: Diffusion-based domain extension for domain-generalized semantic segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2830–2840 (January 2024) [3](#)
- [48] Olsson, V., Tranheden, W., Pinto, J., Svensson, L.: Classmix: Segmentation-based data augmentation for semi-supervised learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1369–1378 (2021) [1](#), [2](#), [3](#), [4](#), [5](#), [8](#)
- [49] Peng, D., Hu, P., Ke, Q., Liu, J.: Diffusion-based image translation with label guidance for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 808–820 (2023) [3](#)
- [50] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) [2](#), [3](#), [6](#)

- [51] Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: Denseclip: Language-guided dense prediction with context-aware prompting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18082–18091 (2022) [2](#), [3](#), [6](#)
- [52] Saha, S., Hoyer, L., Obukhov, A., Dai, D., Van Gool, L.: Edaps: Enhanced domain-adaptive panoptic segmentation. arXiv preprint arXiv:2304.14291 (2023) [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [53] Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in Neural Information Processing Systems* **29** (2016) [5](#)
- [54] Shen, F., Zhou, L., Kucukaytekin, K., Liu, Z., Wang, H., Knoll, A.: Controluda: Controllable diffusion-assisted unsupervised domain adaptation for cross-weather semantic segmentation. arXiv preprint arXiv:2402.06446 (2024) [3](#)
- [55] Shin, G., Albanie, S., Xie, W.: Zero-shot unsupervised transfer instance segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4847–4857 (2023) [3](#)
- [56] Singha, M., Pal, H., Jha, A., Banerjee, B.: Ad-clip: Adapting domains in prompt space using clip. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4355–4364 (2023) [3](#)
- [57] Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems* **33**, 596–608 (2020) [3](#), [5](#)
- [58] Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* **30** (2017) [3](#), [4](#), [5](#)
- [59] 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) [3](#), [4](#), [5](#)
- [60] Tsai, Y.H., Hung, W.C., Schuler, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7472–7481 (2018) [3](#)
- [61] Vidit, V., Engilberge, M., Salzmann, M.: Clip the gap: A single domain generalization approach for object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3219–3229 (2023) [3](#)
- [62] 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/CVF conference on computer vision and pattern recognition. pp. 2517–2526 (2019) [3](#)
- [63] 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/CVF conference on computer vision and pattern recognition. pp. 2517–2526 (2019) [3](#), [8](#)
- [64] Wang, H., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: Max-deeplab: End-to-end panoptic segmentation with mask transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5463–5474 (June 2021) [3](#)
- [65] Wang, Z., Yu, M., Wei, Y., Feris, R., Xiong, J., Hwu, W.m., Huang, T.S., Shi, H.: Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12635–12644 (2020) [3](#)
- [66] Xian, P., Po, L.M., Zhao, Y., Yu, W.Y., Cheung, K.W.: Clip driven few-shot panoptic segmentation. *IEEE Access* (2023) [3](#)
- [67] Xiong, Y., Liao, R., Zhao, H., Hu, R., Bai, M., Yumer, E., Urtasun, R.: Upsnet: A unified panoptic segmentation network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019) [3](#)
- [68] Xu, J., Liu, S., Vahdat, A., Byeon, W., Wang, X., De Mello, S.: Open-vocabulary panoptic segmentation with text-to-image diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2955–2966 (2023) [3](#)
- [69] Xu, J., Hou, J., Zhang, Y., Feng, R., Wang, Y., Qiao, Y., Xie, W.: Learning open-vocabulary semantic segmentation models from natural language supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2935–2944 (2023) [3](#)
- [70] Xu, M., Zhang, Z., Wei, F., Hu, H., Bai, X.: Side adapter network for open-vocabulary semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2945–2954 (2023) [3](#)
- [71] Xu, M., Zhang, Z., Wei, F., Lin, Y., Cao, Y., Hu, H., Bai, X.: A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In: European Conference on Computer Vision. pp. 736–753. Springer (2022) [3](#)
- [72] Yang, Y., Soatto, S.: Fda: Fourier domain adaptation for semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4085–4095 (2020) [3](#), [8](#)
- [73] Yao, D., Li, B.: Dual-level interaction for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4527–4536 (2023) [3](#), [4](#)
- [74] Yi, M., Cui, Q., Wu, H., Yang, C., Yoshie, O., Lu, H.: A simple framework for text-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7071–7080 (2023) [3](#)
- [75] Yu, Q., He, J., Deng, X., Shen, X., Chen, L.C.: Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. arXiv preprint arXiv:2308.02487 (2023) [3](#)

- [76] Yun, S., Park, S.H., Seo, P.H., Shin, J.: Ifseg: Image-free semantic segmentation via vision-language model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2967–2977 (2023) [3](#)
- [77] Zara, G., Roy, S., Rota, P., Ricci, E.: Autolabel: Clip-based framework for open-set video domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11504–11513 (2023) [3](#)
- [78] Zhang, J., Huang, J., Jiang, X., Lu, S.: Black-box unsupervised domain adaptation with bi-directional atkinson-shiffrin memory. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11771–11782 (2023) [3](#)
- [79] Zhang, J., Huang, J., Zhang, X., Lu, S.: Unidaformer: Unified domain adaptive panoptic segmentation transformer via hierarchical mask calibration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11227–11237 (2023) [1](#), [2](#), [3](#), [6](#), [8](#)
- [80] Zhang, P., Zhang, B., Zhang, T., Chen, D., Wang, Y., Wen, F.: Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12414–12424 (2021) [3](#), [4](#)
- [81] Zhang, Q., Zhang, J., Liu, W., Tao, D.: Category anchor-guided unsupervised domain adaptation for semantic segmentation. *Advances in neural information processing systems* **32** (2019) [3](#), [4](#)
- [82] Zhao, H., Sheng, D., Bao, J., Chen, D., Chen, D., Wen, F., Yuan, L., Liu, C., Zhou, W., Chu, Q., et al.: X-paste: revisiting scalable copy-paste for instance segmentation using clip and stablediffusion. In: International Conference on Machine Learning. pp. 42098–42109. PMLR (2023) [3](#)
- [83] Zhao, X., Mithun, N.C., Rajvanshi, A., Chiu, H.P., Samarasekera, S.: Unsupervised domain adaptation for semantic segmentation with pseudo label self-refinement. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2399–2409 (January 2024) [3](#), [4](#)
- [84] Zhou, Q., Feng, Z., Gu, Q., Pang, J., Cheng, G., Lu, X., Shi, J., Ma, L.: Context-aware mixup for domain adaptive semantic segmentation. *IEEE Transactions on Circuits and Systems for Video Technology* **33**(2), 804–817 (2022) [3](#)
- [85] Zou, Y., Yu, Z., Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: Proceedings of the European conference on computer vision (ECCV). pp. 289–305 (2018) [3](#), [8](#)
- [86] Zou, Y., Yu, Z., 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) [4](#)