

UniDAformer: Unified Domain Adaptive Panoptic Segmentation Transformer via Hierarchical Mask Calibration

Jingyi Zhang^{1,*} Jiaxing Huang^{1,*} Xiaoqin Zhang² Shijian Lu^{1,†}
¹ S-lab, Nanyang Technological University ² Wenzhou University

Abstract

Domain adaptive panoptic segmentation aims to mitigate data annotation challenge by leveraging off-the-shelf annotated data in one or multiple related source domains. However, existing studies employ two separate networks for instance segmentation and semantic segmentation which lead to excessive network parameters as well as complicated and computationally intensive training and inference processes. We design UniDAformer, a unified domain adaptive panoptic segmentation transformer that is simple but can achieve domain adaptive instance segmentation and semantic segmentation simultaneously within a single network. UniDAformer introduces Hierarchical Mask Calibration (HMC) that rectifies inaccurate predictions at the level of regions, superpixels and pixels via online self-training on the fly. It has three unique features: 1) it enables unified domain adaptive panoptic adaptation; 2) it mitigates false predictions and improves domain adaptive panoptic segmentation effectively; 3) it is end-to-end trainable with a much simpler training and inference pipeline. Extensive experiments over multiple public benchmarks show that UniDAformer achieves superior domain adaptive panoptic segmentation as compared with the state-of-the-art.

1. Introduction

Panoptic segmentation [30] performs instance segmentation for *things* and semantic segmentation for *stuff*, which assigns each image pixel with a semantic category and a unique identity simultaneously. With the advance of deep neural networks [5, 17–19, 31, 43], panoptic segmentation [4, 7–9, 29, 30, 33, 35, 54, 55] has achieved very impressive performance under the supervision of plenty of densely-annotated training data. However, collecting densely-annotated panoptic data is prohibitively laborious and time-consuming [10, 11, 37] which has become one major constraint along this line of research. One alternative is

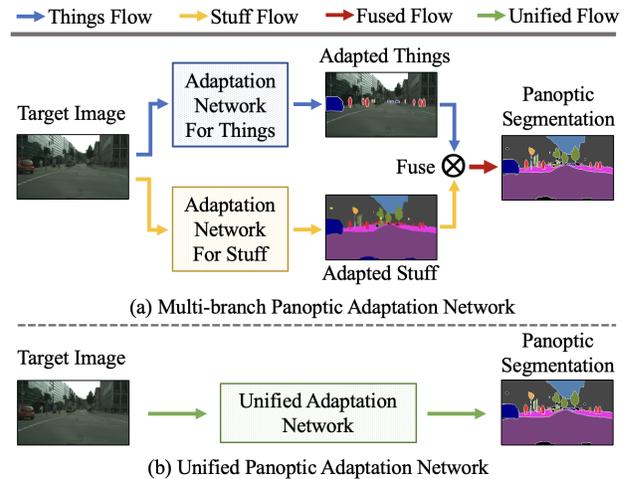


Figure 1. Existing domain adaptive panoptic segmentation [20] adapts things and stuff separately with two isolated networks (for instance segmentation and semantic segmentation) and fuses their outputs to produce the final panoptic segmentation as in (a), leading to excessive network parameters as well as complicated and computationally intensive training and inference. Differently, UniDAformer employs a single unified network to jointly adapt things and stuff as in (b), which involves much less parameters and simplifies the training and inference pipeline greatly.

to leverage off-the-shelf labeled data from one or multiple *source domains*. Nevertheless, the source-trained models often experience clear performance drop while applied to various *target domains* that usually have different data distributions as compared with the *source domains* [20].

Domain adaptive panoptic segmentation can mitigate the inter-domain discrepancy by aligning one or multiple labeled source domains and an unlabeled target domain [20]. To the best of our knowledge, CVRN [20] is the only work that tackles domain adaptive panoptic segmentation challenges by exploiting the distinct natures of instance segmentation and semantic segmentation. Specifically, CVRN introduces cross-view regularization to guide the two segmentation tasks to complement and regularize each other and achieves very impressive performance. However, CVRN re-

*Equal contribution, {jingyi.zhang, jiaxing.huang}@ntu.edu.sg.

†Corresponding author, shijian.lu@ntu.edu.sg.

	Multi-branch Architecture			Unified Architecture								
	PSN [30]			Panoptic FCN [35]			MaskFormer [9]			DETR [4]		
	mSQ	mRQ	mPQ	mSQ	mRQ	mPQ	mSQ	mRQ	mPQ	mSQ	mRQ	mPQ
Supervised Setup	75.5	60.2	47.7	79.7	73.1	59.6	79.1	62.6	51.1	79.1	64.1	51.9
Adaptation Setup	59.0	27.8	20.1	47.5	19.7	15.8	56.6	19.2	16.2	56.4	21.8	18.3
Performance Drop	-16.5	-32.4	-27.6	-32.2	-53.4	-43.8	-22.5	-43.4	-34.9	-22.7	-42.3	-33.6

Table 1. Panoptic segmentation with traditional multi-branch architecture [30] and recent unified architectures [4, 9, 35]: The *Supervised Setup* trains with the Cityscapes [10] and tests on the same dataset. The *UDA Setup* trains with the SYNTHIA [45] and tests on Cityscapes. It can be seen that the *performance drops* between the two learning setups come more from mRQ than from mSQ consistently across different architectures. In addition, such a phenomenon is more severe for unified architectures. This demonstrates a clear false prediction issue in unified domain adaptive panoptic segmentation as mRQ is computed with false positives and false negatives.

lies on a multi-branch architecture that adopts a two-phase pipeline with two separate networks as illustrated in Fig. 1 (a). This sophisticated design directly doubles network parameters, slows down the training, and hinders it from being end-to-end trainable. It is desirable to have a unified panoptic adaptation network that can effectively handle the two segmentation tasks with a single network.

We design a unified domain adaptive panoptic segmentation transformer (UniDAformer) as shown in Fig. 1 (b). Our design is based on the observation that one major issue in unified panoptic adaptation comes from a severe false prediction problem. As shown in Table 1, most recent unified panoptic segmentation architectures [4, 9, 35] outperform traditional multi-branch architectures [30] by large margins under the supervised setup. However, the situation inverts completely under unsupervised domain adaptation setup. Such contradictory results are more severe for the recognition quality in mRQ. This shows that the panoptic quality drop of unified architecture mainly comes from False Positives (FP) and False Negatives (FN) as mRQ is computed from all predictions (True Positives, False Negatives and False Positives) while the segmentation quality in mSQ is computed with True Positives (TP) only.

In the proposed UniDAformer, we mitigate the false prediction issue by introducing Hierarchical Mask Calibration (HMC) that calibrates inaccurate predictions at the level of regions, superpixels, and pixels. With the corrected masks, UniDAformer re-trains the network via an online self-training process on the fly. Specifically, HMC treats both things and stuff predictions as masks uniformly and corrects each predicted pseudo mask hierarchically in a coarse-to-fine manner, *i.e.*, from region level that calibrates the overall category of each mask to superpixel and pixel levels that calibrate the superpixel and pixels around the boundary of each mask (which are more susceptible to prediction errors). UniDAformer has three unique features. *First*, it achieves unified panoptic adaptation by treating things and stuff as masks and adapting them uniformly. *Second*, it mitigates the false prediction issue effectively

by calibrating the predicted pseudo masks iteratively and progressively. *Third*, it is end-to-end trainable with much less parameters and simpler training and inference pipeline. Besides, HMC introduces little computation overhead and could be used as a plug-in.

The contributions of this work can be summarized in three aspects. *First*, we propose UniDAformer that enables concurrent domain adaptive instance segmentation and semantic segmentation within a single network. It is the first end-to-end unified domain adaptive panoptic segmentation transformer to the best our knowledge. *Second*, we design Hierarchical Mask Calibration with online self-training, which allows to calibrate the predicted pseudo masks on the fly during self-training. *Third*, extensive experiments over multiple public benchmarks show that the proposed UniDAformer achieves superior segmentation accuracy and efficiency as compared with the state-of-the-art.

2. Related Work

Panoptic Segmentation is a challenging task that assigns each image pixel with a semantic category and a unique identity. The pioneer work [30] employs two networks for instance segmentation and semantic segmentation separately, and then combines the outputs of the two segmentation networks to acquire panoptic segmentation. The later studies [4, 7–9, 29, 33, 35, 54, 55] simplify the complex pipeline by unifying the segmentation of things and stuff within single network. For example, DETR [4] predicts boxes around both things and stuff classes, and makes a final panoptic prediction by adding an FPN-style segmentation head. Panoptic segmentation has achieved very impressive accuracy but requires a large amount of densely-annotated training data that are often laborious and time-consuming to collect. Domain adaptive panoptic segmentation (DAPS), which leverages off-the-shelf annotated data for mitigating the data annotation constraint, is instead largely neglected.

Unsupervised Domain Adaptation (UDA) aims to exploit labeled source-domain data to learn a well-performing

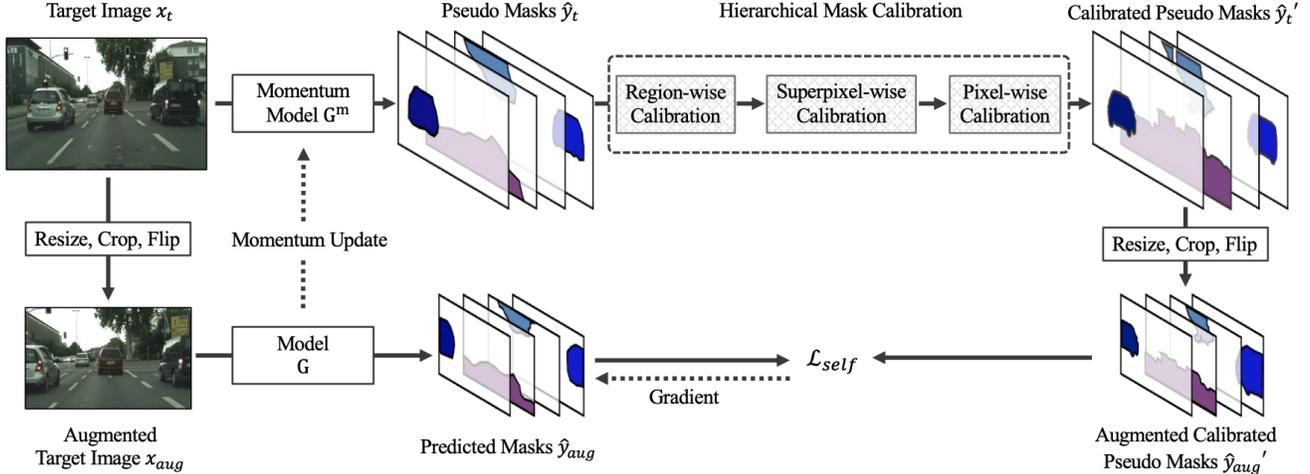


Figure 2. Overview of proposed unified domain adaptive panoptic segmentation transformer (UniDAformer): it involves two flows, *i.e.*, a pseudo mask generation flow that calibrates pseudo masks with momentum model G^m , and an unsupervised training flow that optimizes model G with the calibrated pseudo masks. For pseudo mask calibration, we feed a given unlabeled target image x_t into the momentum model G^m to calibrate pseudo masks \hat{y}_t with HMC via a coarse-to-fine manner (*i.e.*, from region level to superpixel and pixel levels). For network optimization, we conduct simple augmentations (*i.e.*, resize, crop and flip) for x_t and its calibrated pseudo masks \hat{y}_t' , and then optimize model G with self-training loss \mathcal{L}_{self} .

model on unlabeled target-domain data. In recent years, it has been studied extensively for various computer vision tasks, including image classification [12, 14, 38, 39, 42, 42, 47, 49, 51, 65], instance segmentation/detection [3, 6, 16, 21, 25, 32, 46, 48, 56, 59, 63] and semantic segmentation [22–24, 27, 34, 41, 52, 57, 60–62, 65]. On the other hand, domain adaptive panoptic segmentation is largely neglected despite its great values in various visual tasks and practical applications. To the best of our knowledge, CVRN [20] is the only work, which exploits the distinct natures of instance segmentation and semantic segmentation and introduces cross-view regularization to guide the two tasks to complement and regularize each other for panoptic adaptation. However, CVRN achieves panoptic adaptation by using two separate adaptation networks for things and stuff respectively, which directly doubles network parameters, slows down the network, and hinders it from being end-to-end trainable. In contrast, our proposed UniDAformer greatly simplifies training and inference pipeline by unifying the adaptation of things and stuff in a single panoptic adaptation network.

Self-training is a mainstream unsupervised domain adaptation technique that retrains networks with pseudo-labeled target-domain data. Most existing self-training methods [22, 25, 26, 28, 36, 53, 58, 60, 64, 65] involve an iterative retraining process for effective learning from pseudo-labeled data. In each training iteration, an offline pseudo label generation process is involved which predicts and selects pseudo labels according to their confidence. For example, [64] proposes class-balanced self-training (CBST)

that globally selects the same proportion of predictions as pseudo labels for each category for overcoming class-imbalance issues. To sidestep the cumbersome multi-round and offline training process, several studies [2, 40] explore ‘online’ self-training for semantic segmentation by directly enforcing pixel-wise consistency of predictions from different data augmentations. Differently, the proposed UniDAformer focuses on the false prediction issue in unified domain adaptive panoptic segmentation. It achieves effective ‘online’ self-training with a Hierarchical Mask Calibration technique which allows pseudo label calibration and correction on the fly.

3. Method

3.1. Task Definition

This work focuses on domain adaptive panoptic segmentation. The training data involves a labeled source domain $\mathcal{D}_s = \{(x_s^i, y_s^i)\}_{i=1}^{N_s}$ (y_s^i is the panoptic annotation of sample x_s^i) and an unlabeled target domain $\mathcal{D}_t = \{x_t^i\}_{i=1}^{N_t}$. The goal is to learn a model G from \mathcal{D}_s and \mathcal{D}_t that well performs in \mathcal{D}_t . The baseline model is trained with the source domain data \mathcal{D}_s only:

$$\mathcal{L}_{sup} = l(G(x_s), y_s), \quad (1)$$

where $l(\cdot)$ denotes the panoptic segmentation loss that consists of a matching cost and a Hungarian loss [4].

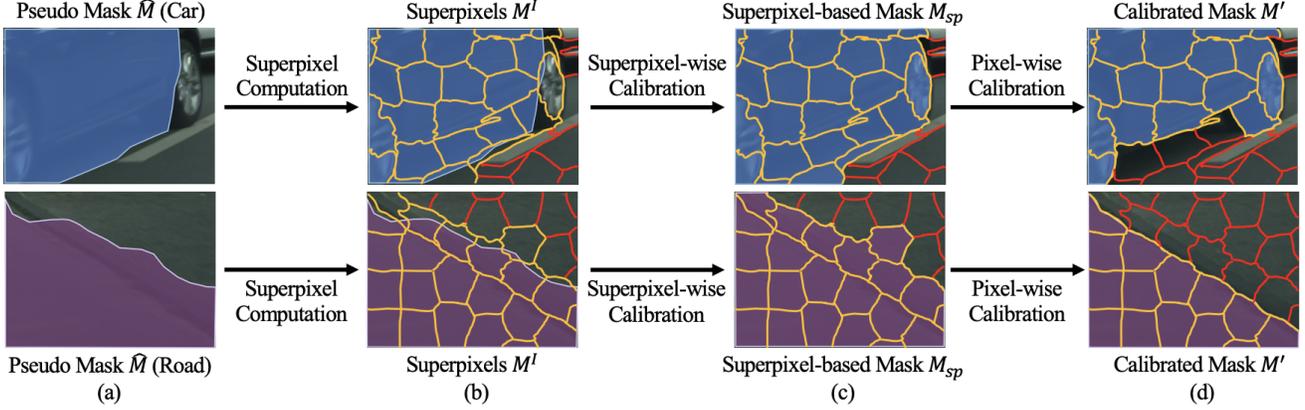


Figure 3. Overview of Hierarchical Mask Calibration: it consists of three sub-modules, *i.e.*, Region-wise Calibration Module, Superpixel-wise Calibration Module and Pixel-wise Calibration Module. For simplicity, we skip visualizing Region-wise Calibration and directly present pseudo mask after region-wise calibration in (a). In Superpixel-wise Calibration, we first compute superpixels M^I and select the superpixels (marked with yellow lines as in (b)) that overlap with the pseudo mask, based on which the pseudo mask is expanded into superpixel-based mask M_{sp} as in (c). In Pixel-wise Calibration, we discard the superpixels that are inconsistent with the calibrated overall category through a pixel-wise voting mechanism to form the final calibrated mask M' as in (d).

3.2. UniDAformer Overview

This subsection presents the overall framework of proposed UniDAformer, which consists of a supervised training process over the labeled source domain and an unsupervised training process over the unlabeled target domain. For the *supervised training*, the source samples (x_s, y_s) are fed to a panoptic segmentation model G that is optimized via the supervised loss \mathcal{L}_{sup} as defined in Eq. 1.

The *unsupervised training* involves two flows as illustrated in Fig. 2. The first flow calibrates pseudo masks with the momentum model G^m (the moving averaged of G , *i.e.*, $\theta_{G^m} \leftarrow \gamma \theta_{G^m} + (1 - \gamma)\theta_G$, and γ is a momentum coefficient). In pseudo mask calibration, we first feed target image x_t into the momentum model G^m and calibrate the pseudo masks \hat{y}_t . The pseudo masks \hat{y}_t are then forwarded to the Hierarchical Mask Calibration (HMC) module that produces corrected pseudo masks \hat{y}'_t via coarse-to-fine calibration. The second flow optimizes G with the calibrated pseudo masks. Specifically, we first apply simple data augmentations (*i.e.*, resize, crop and flip) to x_t and \hat{y}'_t to obtain x_{aug} and \hat{y}'_{aug} . The network model G is then optimized with the augmented data and the self-training loss \mathcal{L}_{self} as defined in Eq. 9.

3.3. Hierarchical Mask Calibration

One key component in the proposed UniDAformer is HMC that calibrates the predicted pseudo masks and enables effective pseudo-label retraining on the fly. HMC treats both things and stuff predictions as masks uniformly and corrects each predicted mask hierarchically in a coarse-to-fine manner. The correction involves three consecutive

stages of calibration including a Region-wise Calibration, a Superpixel-wise Calibration and a Pixel-wise Calibration as illustrated in Fig. 2. First, Region-wise Calibration corrects the overall category of each mask by adaptively re-weighting its category-wise probabilities. Leveraging the feature that superpixels adhere well to the boundaries of things and stuff, Superpixel-wise Calibration then adjusts the shape of each mask by considering the boundary of the computed superpixels. Finally, Pixel-wise Calibration introduces pixel-level categorization information and further refines the boundary of each mask with a simple pixel-wise voting mechanism.

As the proposed calibration technique works for all predicted pseudo masks (*i.e.*, things and stuff) uniformly, we take one pseudo mask \hat{y}_t^k from $\hat{y}_t = \{\hat{y}_t^k\}_{k=1}^K$ as an example for illustration. Each pseudo mask \hat{y}_t^k includes a predicted category $\hat{c}_t^k = \arg \max_c p^c$ ($p^c \in \{p^c\}_{c=1}^C$ is the probability of belonging to the c -th category) and a predicted binary mask \hat{M}_t^k of size $H \times W$.

Region-wise Calibration corrects the predicted category \hat{c}_t^k by re-weighting its category-wise probability p^c as following:

$$c_t^{k'} = \arg \max_c (w^{(c,k)} \otimes p^c), \quad (2)$$

where \otimes denotes the element-wise multiplication and $w^{(c,k)} \in \{w^{(c,k)}\}_{c=1}^C$ is the calibration weight of the corresponding c -th category probability for each pseudo mask.

The calibration weight w is calculated according to the distance between \hat{y}_t^k and the centroids δ over feature space (*i.e.*, feature center calculated and updated in Eqs. 5 and 6). Given the feature $f_t^k \in \mathbb{R}^{E \times H \times W}$ (E , H and W denote the feature's channel, height and width respectively) as gen-

erated by the momentum backbone, we pool the feature within the mask \hat{M} into a region-wise vector $v_r^k \in \mathbb{R}^E$ (subscript r denotes it is a region-wise vector) as follows:

$$v_r^k = \text{GAP}(\hat{M}_t^k \otimes f_t^k), \quad (3)$$

where $\text{GAP}(\cdot)$ denotes the global average pooling operation.

Generally, if the region-wise vector v_r^k is far from the c -th centroid δ^c , the pseudo mask \hat{y}_t^k should be assigned with a lower probability of belonging to the c -th category, and vice versa. Therefore, the calibration weight in Eq. 2 is defined as follows:

$$w^{(c,k)} = \text{Softmax}(-\|v_r^k - \delta^c\|_1), \quad (4)$$

where the distance is measured using L1 distance and softmax operation is performed along the category dimension.

Here we demonstrate how we compute and update the mask centroids along the training process. The mask centroids are first initialized by all target predictions from the baseline model. For each category, the mask centroid δ^c is defined as follows:

$$\delta^c = \frac{\sum_{x_t \in \mathcal{D}_t} \sum_{k \in K} v_m^r \cdot \mathbb{1}(\hat{c}^k = c)}{\sum_{x_t \in \mathcal{D}_t} \sum_{k \in K} \mathbb{1}(\hat{c}^k = c)}, \quad (5)$$

where $\mathbb{1}$ is an indicator function that returns ‘1’ if the vector v_m^k belongs to c -th category, and ‘0’ otherwise.

Along training process, we update the mask centroids with the current batch of data:

$$\delta^c \leftarrow \gamma' \delta^c + (1 - \gamma') \delta_*^c, \quad (6)$$

where δ_*^c is the mask centroid calculated with the current data and model, and γ' is a update coefficient for smooth centroid update.

Supervoxel-wise Calibration: Following region-wise calibration, we first correct the shape of the pseudo mask by exploiting supervoxels that adhere well to the boundaries of things and stuff [1]. To this end, we first compute a supervoxel map M^I which includes total I supervoxels $M^{(i)}$ for target image x_t . Then, we select the supervoxels that overlap with the original mask \hat{M}_t^k to form an adjusted binary mask M_{sp}^k as follows:

$$M_{sp}^k = \bigcup_{i \in I} M^{(i)} \cdot \mathbb{1}(A^{(i)} > 0), \quad (7)$$

where $\mathbb{1}$ is an indicator function and we denote the overlapping area between i -th supervoxel and the mask \hat{M}_t^k as $A^{(i)}$.

The supervoxel-based mask M_{sp}^k adjusts the original mask \hat{M} with the computed supervoxels which adheres better to the edge of things or stuff, as illustrated in Figs. 3 (b) and (c).

Self-train.	Region	Supervoxel	Pixel	mSQ	mRQ	mPQ
				56.4	21.8	18.3
✓				59.5	29.9	22.6
✓	✓			61.2	36.9	28.7
✓		✓		63.0	32.4	26.2
✓			✓	62.6	31.8	24.2
✓	✓	✓		63.4	39.9	30.9
✓	✓		✓	63.2	38.9	30.1
✓		✓	✓	64.3	32.7	26.9
✓	✓	✓	✓	64.7	42.2	33.0

Table 2. Ablation study of the proposed Hierarchical Mask Calibration technique over task SYNTHIA \rightarrow Cityscapes, where ‘Region’, ‘Supervoxel’ and ‘Pixel’ stand for region-wise calibration, supervoxel-wise calibration and pixel-wise calibration, respectively.

Pixel-wise Calibration: Following supervoxel-wise calibration, we further introduce pixel-level categorization information and refine the boundary of the supervoxel-based mask M_{sp}^k in a more precise manner.

We design a simple pixel-wise voting mechanism to achieve the fine-grained calibration: the supervoxels are discarded from the M_{sp}^k if a majority of pixel-wise feature vectors within it are inconsistent with the overall category of pseudo mask \hat{y}_t^k as illustrated in Fig. 3 (d). Let $v_p^{(k,j)}$ denote a pixel-wise feature vector within supervoxel $M^{(j)} \in M_{sp}$ (subscript p denotes it is a pixel-wise vector), and we determine that it is inconsistent with the pseudo mask if it does not satisfy the following condition:

$$\arg \min_c (\|v_p^{(k,j)} - \delta^c\|) = c_t^{k'}, \quad (8)$$

where $c_t^{k'}$ is the corrected category of the pseudo mask \hat{y}_t^k . Such pixel-wise voting effectively suppresses the side effect of outlier pixels by enforcing that pixels within the same supervoxel should share the same semantic category [1].

The final calibrated pseudo masks \hat{y}_t' for target image x_t can be acquired by repeating the above-mentioned processes for all the calibrated pseudo mask (*i.e.*, $\hat{y}_t' = \{\hat{y}_t^{k'}\}_{k=1}^K$, where $\hat{y}_t^{k'} = \{c_t^{k'}, M_t^{k'}\}$).

3.4. Network Optimization

With the calibrated pseudo masks \hat{y}_t' , the self-training loss \mathcal{L}_{self} can be formulated as follows:

$$\mathcal{L}_{self} = l(G(x_{aug}), \hat{y}'_{aug}), \quad (9)$$

where $l(\cdot)$ denotes the panoptic segmentation loss that consists of a matching cost and a Hungarian loss [4]. \hat{y}'_{aug} and x_{aug} are the simple augmentations (*i.e.*, resize, crop and flip) of \hat{y}_t' and x_t , respectively.

The overall training objective is defined by minimizing the supervised loss \mathcal{L}_{sup} and the unsupervised loss \mathcal{L}_{self} :

$$\arg \min_G \mathcal{L}_{sup} + \mathcal{L}_{self}. \quad (10)$$

SYNTHIA → Cityscapes Panoptic Segmentation																			
Methods	road	side.	build.	wall	fence	pole	light	sign	vege.	sky	pers.	rider	car	bus	mot.	bike	mSQ	mRQ	mPQ
Baseline [4]	33.9	7.3	45.6	0.0	0.0	2.9	5.2	7.6	65.0	57.1	19.2	5.4	22.3	14.9	1.3	4.8	56.4	21.8	18.3
DAF [6]	34.6	7.5	53.8	0.0	0.0	2.1	3.1	2.6	72.8	67.5	21.3	9.4	28.2	22.7	4.4	4.4	59.0	28.3	20.9
FDA [57]	29.2	5.8	63.6	0.1	0.0	4.9	4.0	4.8	73.9	62.5	24.5	11.4	32.1	27.1	5.2	8.0	59.1	30.3	22.3
CRST [65]	46.8	11.8	56.8	0.9	0.0	4.0	2.6	3.5	70.9	64.3	20.7	11.8	32.9	32.2	6.2	7.9	62.5	31.9	23.3
SVMIn [15]	48.0	11.8	56.9	0.5	0.0	3.9	3.8	4.4	72.5	68.9	26.4	15.0	35.3	25.8	6.3	7.6	63.3	32.6	24.2
AdvEnt [53]	55.9	14.4	64.0	0.0	0.0	4.6	3.3	2.7	75.5	72.3	24.9	9.7	33.8	26.7	5.2	7.1	60.2	33.0	25.0
CVRN [20]	66.2	19.4	72.5	2.1	0.0	3.8	6.5	4.4	79.7	75.1	26.5	11.5	36.6	34.1	7.1	8.2	61.4	35.9	27.9
UniDAformer	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

Table 3. Experiments with unified panoptic segmentation architecture [4] over task SYNTHIA → Cityscapes. PQ is computed for each category. Mean SQ (mSQ), mean RQ (mRQ), mean PQ (mPQ) are computed over all categories.

Cityscapes → Foggy Cityscapes Panoptic Segmentation																			
Methods	road	side.	build.	wall	fence	pole	light	sign	vege.	sky	pers.	rider	car	bus	mot.	bike	mSQ	mRQ	mPQ
Baseline [4]	92.5	48.9	60.6	6.0	10.7	5.3	9.9	23.6	49.7	55.6	22.3	15.4	38.5	23.7	1.6	2.8	70.0	38.6	29.2
DAF [6]	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 [57]	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 [53]	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 [65]	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 [15]	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 [20]	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	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

Table 4. Experiments with unified panoptic segmentation architecture [4] over task Cityscapes → Foggy cityscapes. PQ is computed for each category. Mean SQ (mSQ), mean RQ (mRQ), mean PQ (mPQ) are computed over all categories.

4. Experiment

This section presents experiments including datasets, evaluation metric, ablation studies, comparisons with the state-of-the-art and discussions. Due to the space limit, the implementation details are provided in the appendix.

4.1. Datasets

We evaluate UniDAformer over three widely used domain adaptation tasks with four datasets:

- 1) SYNTHIA [45] → Cityscapes [10] which aims for domain adaptation from synthetic images to real-world images. The training set in SYNTHIA are adopted as source domain and the training set in Cityscapes are considered as target domain. The evaluation is performed on the validation set of Cityscapes.
- 2) Cityscapes [10] → Foggy cityscapes [50] which aims for domain adaptation across different weather conditions, where Cityscapes is used as source domain and Foggy Cityscapes is considered as target domain. The adaptation performance is evaluated over the validation set of Foggy Cityscapes.
- 3) VIPER [44] → Cityscapes [10] which aims for domain adaptation from synthetic images to real-world images. We adopt the training set of VIPER as source domain and the training set in Cityscapes as target domain. The evaluation is performed on the validation set of Cityscapes.

In evaluations, we adopt three panoptic segmentation metrics [30] including segmentation quality (SQ), recognition quality (RQ) and panoptic quality (PQ) as in [20, 29, 30, 35]. For each category, PQ can be computed as the multiplication of the corresponding SQ term and RQ term as follows:

$$PQ = \underbrace{\frac{\sum_{(p,g) \in TP} \text{IoU}(p,g)}{|TP|}}_{\text{segmentation quality (SQ)}} \times \underbrace{\frac{|TP|}{|TP| + \frac{1}{2}|FP| + \frac{1}{2}|FN|}}_{\text{recognition quality (RQ)}}, \quad (11)$$

where g is the ground truth segment and p is the matched prediction. TP, FP and FN denote true positives, false positives and false negatives, respectively. IoU is the insertion over union metric [13] which is widely used in semantic segmentation evaluations. With the above definitions, RQ captures the proportion of TP in all predictions, SQ captures the segmentation quality within TP while PQ integrates PQ and SQ and captures the overall panoptic segmentation quality.

4.2. Ablation Studies

The core of UniDAformer is Hierarchical Mask Calibration that consists of a Region-wise Calibration, a Superpixel-wise Calibration and a Pixel-wise Calibration. We first study the three calibration modules to examine how they contribute to the overall domain adaptive panoptic seg-

VIPER → Cityscapes Panoptic Segmentation																
Methods	road	side.	build.	fence	light	sign	vege.	sky	pers.	car	bus	mot.	bike	mSQ	mRQ	mPQ
Baseline [4]	25.2	5.7	35.1	0.0	5.9	3.9	75.3	68.7	21.2	39.7	21.4	11.4	0.0	59.7	32.0	24.1
DAF [6]	56.6	7.3	41.0	0.0	3.5	2.7	76.4	70.2	19.0	34.3	14.2	6.3	0.0	61.1	33.3	25.5
FDA [57]	50.0	7.6	59.4	0.0	6.2	6.1	73.3	65.9	19.4	38.2	15.5	8.1	0.0	61.0	35.2	26.9
AdvEnt [53]	52.6	10.8	51.0	0.0	2.0	4.8	73.9	70.1	15.9	38.2	19.9	12.4	0.0	61.2	35.4	27.0
CRST [65]	68.7	9.1	54.4	0.0	2.4	2.7	76.3	69.9	21.2	34.0	21.9	7.7	0.0	61.0	36.5	28.3
SVMin [15]	87.6	14.2	70.7	0.0	4.1	6.3	74.4	70.0	16.9	32.5	2.4	11.0	1.2	61.3	37.5	29.9
CVRN [20]	75.1	18.8	59.9	0.0	9.1	6.5	76.8	71.1	22.3	37.0	15.5	8.6	3.8	66.4	40.2	31.1
UniDAformer	87.1	22.1	71.1	0.0	8.2	8.6	78.3	71.8	25.4	46.8	13.7	12.8	2.8	68.9	43.0	34.5

Table 5. Experiments with unified panoptic segmentation architecture [4] over task VIPER → Cityscapes. PQ is computed for each category. Mean SQ (mSQ), mean RQ (mRQ), mean PQ (mPQ) are computed over all categories.

SYNTIA → Cityscapes Panoptic Segmentation																			
Methods	road	side.	build.	wall	fence	pole	light	sign	vege.	sky	pers.	rider	car	bus	mot.	bike	mSQ	mRQ	mPQ
PSN [30]	32.3	5.1	58.5	0.9	0.0	0.9	0.0	4.6	61.7	61.3	27.6	9.5	32.8	22.6	1.0	2.7	59.0	27.8	20.1
FDA [57]	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 [65]	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 [53]	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 [20]	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	87.7	34.0	73.2	1.3	0.0	8.1	9.9	6.7	78.2	74.0	37.6	25.3	40.7	37.4	15.0	18.8	66.9	44.3	34.2

Table 6. Experiments with multi-branch panoptic segmentation architecture [30] over task SYNTIA → Cityscapes. Mean SQ (mSQ), mean RQ (mRQ), mean PQ (mPQ) are computed over all categories.

mentation.

Table 2 shows experimental results over task SYNTIA → Cityscapes. It can be seen that the baseline in the 1st Row (trained with the labeled source data only) does not perform well due to domain shifts. Including self-training over unlabeled target data in the 2nd Row improves the baseline from 18.3 to 22.6 in mPQ. On top of the self-training, including any of the three calibration modules improves the segmentation consistently as shown in Rows 3-5. Specifically, region-wise calibration improves mRQ more (15.1 above the baseline) than the other two calibration modules (10.6 and 10.0), showing that region-wise calibration suppresses false predictions effectively by calibrating the overall category of each mask. On the other hand, superpixel-wise and pixel-wise calibrations improve mSQ more than region-wise calibration (6.6 and 6.2 vs 4.8), showing that superpixel-wise and pixel-wise calibrations focus on refining the boundary of each mask.

The three calibration modules correct pseudo masks from different levels which complement each in domain adaptive panoptic segmentation. We can observe that combining any two modules further improves mSQ, mRQ and mPQ consistently as shown in Rows 6-8, and combining all three achieves the best mSQ, mRQ and mPQ. Such experimental results are well aligned with the motivation and design of the proposed hierarchical mask calibration.

4.3. Comparisons with the State-of-the-art

Due to the lack of prior studies on unified domain adaptive panoptic segmentation, we conduct two sets of experiments to benchmark UniDAformer with the state-of-the-art.

In the first set of experiments, we benchmark UniDAformer over the unified panoptic segmentation architecture (*i.e.*, DETR [4]) by reproducing the state-of-the-art [20] with DETR. Specifically, we re-implement the cross-style regularization (one of two cross-view designs) in DETR to reproduce CVRN (cross-task regularization relies on multi-branch architecture and cannot work in the unified architecture). Following [20], we also reproduce several domain adaptive methods by directly implementing their adaptation module in DETR. We perform comparisons over three domain adaptive panoptic segmentation tasks as shown in Tables 3-5. It can be seen that UniDAformer improves the baseline [4] by large margins (8.3, 20.4 and 14.7 in mSQ, mRQ and mPQ) and it also outperforms the state-of-the-art clearly for SYNTIA → Cityscapes. In particular, UniDAformer improves more in mRQ as compared with the state-of-the-art, indicating that it corrects more false predictions effectively. Similar experimental results are observed on the other two tasks as shown in Tables 4 and 5.

In the second set of experiments, we benchmark UniDAformer over the multi-branch panoptic segmentation architecture (*i.e.*, PSN [30]). Since HMC introduces little extra computation overhead and can be incorporated as a plug-in, we directly apply HMC (with the online self-

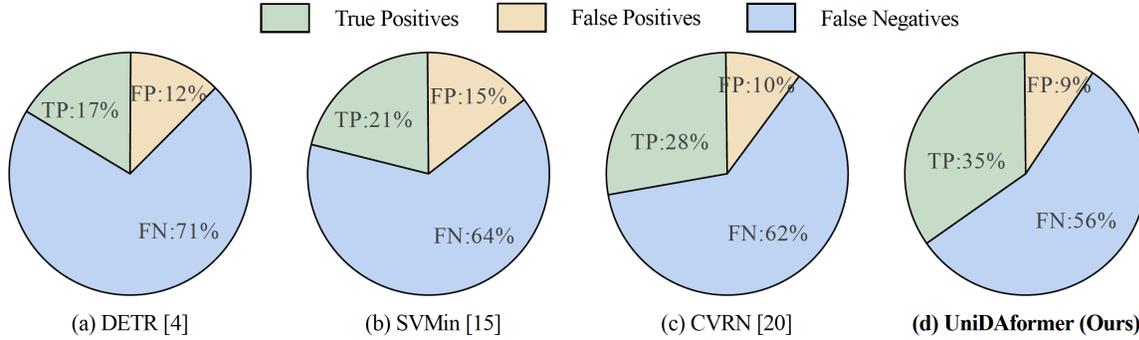


Figure 4. Prediction quality analysis over task SYNTHIA \rightarrow Cityscapes.

Order	R \rightarrow S \rightarrow P			P \rightarrow S \rightarrow R			S \rightarrow P \rightarrow R		
	mSQ	mRQ	mPQ	mSQ	mRQ	mPQ	mSQ	mRQ	mPQ
Results	64.7	42.2	33.0	62.0	38.9	29.9	62.2	39.9	30.1

Table 7. The calibration order affects domain adaptation performance. The experiments are conducted over task SYNTHIA \rightarrow Cityscapes. R, S and P denote region-wise calibration, superpixel-wise calibration and pixel-wise calibration, respectively.

Methods	Architecture	Parameter	Training Speed	Inference Speed
CVRN [20]	Multi-branch	185.58 M	0.27 fps	0.36 fps
UniDAformer	Unified	77.68 M	1.53 fps	5.23 fps

Table 8. Efficiency comparison with multi-branch panoptic adaptation network CVRN [20] in terms of parameter number, training speed and inference speed.

training loss) on the multi-branch architecture for benchmarking. Table 6 shows experimental results on SYNTHIA \rightarrow Cityscapes. We can see that UniDAformer outperforms CVRN in mSQ, mRQ and mPQ consistently. In addition, it similarly improves mRQ by large margins, which further verifies the motivation and design of the proposed HMC.

4.4. Discussions

Prediction Quality Analysis. UniDAformer suppresses false predictions effectively via HMC. We examine it over task SYNTHIA \rightarrow Cityscapes with DETR [4]. As discussed in Section 4.1, the predictions in panoptic segmentation consists of three parts including TP, FP and FN. We compute the proportion of each part over all predictions and Fig. 4 shows experimental results. We can observe that UniDAformer produces clearly more TP and less FN and FP as compared with both baseline [4] and the state-of-the-art [15, 20]. This demonstrates the superiority of UniDAformer in suppressing false predictions in domain adaptive panoptic segmentation.

The Calibration Order Matters. The proposed HMC calibrates predicted pseudo masks in a coarse-to-fine manner

(*i.e.*, from region level to superpixel and pixel levels). We study how calibration order affects panoptic segmentation by testing two reversed calibration orders as shown in Table 7. It can be seen that reversing calibration order leads to clear performance drops, indicating the benefits of the coarse-to-fine calibration in our design.

Efficiency Comparison with CVRN [20]. Beyond segmentation accuracy, we also benchmark UniDAformer with multi-branch panoptic adaptation network CVRN [20] in parameter number, training speed and inference speed. As Table 8 shows, UniDAformer has clearly less parameters and its training and inference time is much shorter than CVRN as well, demonstrating its great simplicity and efficiency.

5. Conclusion

This paper presents UniDAformer, a unified domain adaptive panoptic segmentation transformer. UniDAformer introduces a Hierarchical Mask Calibration (HMC) technique to calibrate the predicted pseudo masks on the fly during re-training. UniDAformer has three unique features: 1) it achieves unified panoptic adaptation by treating things and stuff as masks and adapting them uniformly; 2) it mitigates the severe false prediction issue effectively by calibrating the predicted pseudo masks iteratively and progressively; 3) it is end-to-end trainable with much less parameters and simpler training and inference pipeline. Besides, the proposed HMC introduces little extra computation overhead and could be used as a plug-in. Extensive experiments over multiple public benchmarks show that UniDAformer achieves superior segmentation accuracy and efficiency as compared with the state-of-the-art. Moving forwards, we plan to continue to investigate simple yet effective techniques for unified domain adaptive panoptic segmentation.

Acknowledgement. This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2274–2282, 2012. **5**
- [2] Nikita Araslanov and Stefan Roth. Self-supervised augmentation consistency for adapting semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15384–15394, 2021. **3**
- [3] Qi Cai, Yingwei Pan, Chong-Wah Ngo, Xinmei Tian, Lingyu Duan, and Ting Yao. Exploring object relation in mean teacher for cross-domain detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11457–11466, 2019. **3**
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. **1, 2, 3, 5, 6, 7, 8**
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. **1**
- [6] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018. **3, 6, 7**
- [7] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. 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*, pages 12475–12485, 2020. **1, 2**
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. *arXiv preprint arXiv:2112.01527*, 2021. **1, 2**
- [9] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. **1, 2**
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. **1, 2, 6**
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. **1**
- [12] Zhekai Du, Jingjing Li, Hongzu Su, Lei Zhu, and Ke Lu. Cross-domain gradient discrepancy minimization for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3937–3946, 2021. **3**
- [13] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. **6**
- [14] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. **3**
- [15] Dayan Guan, Jiaying Huang, Shijian Lu, and Aoran Xiao. Scale variance minimization for unsupervised domain adaptation in image segmentation. *Pattern Recognition*, 112:107764, 2021. **6, 7, 8**
- [16] Dayan Guan, Jiaying Huang, Aoran Xiao, Shijian Lu, and Yanpeng Cao. Uncertainty-aware unsupervised domain adaptation in object detection. *IEEE Transactions on Multimedia*, 2021. **3**
- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. **1**
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **1**
- [19] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. **1**
- [20] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Cross-view regularization for domain adaptive panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10133–10144, 2021. **1, 3, 6, 7, 8**
- [21] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Fskr: Frequency space domain randomization for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6891–6902, 2021. **3**
- [22] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Model adaptation: Historical contrastive learning for unsupervised domain adaptation without source data. *Advances in Neural Information Processing Systems*, 34:3635–3649, 2021. **3**
- [23] Jiaying Huang, Dayan Guan, Aoran Xiao, Shijian Lu, and Ling Shao. Category contrast for unsupervised domain adaptation in visual tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1203–1214, 2022. **3**
- [24] Jiaying Huang, Shijian Lu, Dayan Guan, and Xiaobing Zhang. Contextual-relation consistent domain adaptation for semantic segmentation. In *European Conference on Computer Vision*, pages 705–722. Springer, 2020. **3**
- [25] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5001–5009, 2018. **3**

- [26] Javed Iqbal and Mohsen Ali. Mlsl: Multi-level self-supervised learning for domain adaptation with spatially independent and semantically consistent labeling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1864–1873, 2020. 3
- [27] Myeongjin Kim and Hyeran Byun. Learning texture invariant representation for domain adaptation of semantic segmentation. *arXiv preprint arXiv:2003.00867*, 2020. 3
- [28] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6092–6101, 2019. 3
- [29] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6399–6408, 2019. 1, 2, 6
- [30] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. 1, 2, 6, 7
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 1
- [32] Congcong Li, Dawei Du, Libo Zhang, Longyin Wen, Tiejian Luo, Yanjun Wu, and Pengfei Zhu. Spatial attention pyramid network for unsupervised domain adaptation. In *European Conference on Computer Vision*, pages 481–497. Springer, 2020. 3
- [33] Yanwei Li, Xinze Chen, Zheng Zhu, Lingxi Xie, Guan Huang, Dalong Du, and Xingang Wang. Attention-guided unified network for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7026–7035, 2019. 1, 2
- [34] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6936–6945, 2019. 3
- [35] Yanwei Li, Hengshuang Zhao, Xiaojuan Qi, Liwei Wang, Zeming Li, Jian Sun, and Jiaya Jia. Fully convolutional networks for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 214–223, 2021. 1, 2, 6
- [36] Qing Lian, Fengmao Lv, Lixin Duan, and Boqing Gong. Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: A non-adversarial approach. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6758–6767, 2019. 3
- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James HaysPietro, PeronaDeva Ramanan, Piotr DollárC, and Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755, 2014. 1
- [38] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *International conference on machine learning*, pages 2208–2217. PMLR, 2017. 3
- [39] Zhihe Lu, Yongxin Yang, Xiatian Zhu, Cong Liu, Yi-Zhe Song, and Tao Xiang. Stochastic classifiers for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9111–9120, 2020. 3
- [40] Luke Melas-Kyriazi and Arjun K Manrai. Pixmatch: Unsupervised domain adaptation via pixelwise consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12435–12445, 2021. 3
- [41] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. *arXiv preprint arXiv:2004.07703*, 2020. 3
- [42] Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8004–8013, 2018. 3
- [43] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015. 1
- [44] Stephan R Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2213–2222, 2017. 6
- [45] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3234–3243, 2016. 2, 6
- [46] Suman Saha, Anton Obukhov, Danda Pani Paudel, Menelaos Kanakis, Yuhua Chen, Stamatios Georgoulis, and Luc Van Gool. Learning to relate depth and semantics for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8197–8207, 2021. 3
- [47] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8050–8058, 2019. 3
- [48] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6956–6965, 2019. 3
- [49] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2018. 3
- [50] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9):973–992, 2018. 6
- [51] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains

- using generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8503–8512, 2018. 3
- [52] M Naseer Subhani and Mohsen Ali. Learning from scale-invariant examples for domain adaptation in semantic segmentation. *arXiv preprint arXiv:2007.14449*, 2020. 3
- [53] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019. 3, 6, 7
- [54] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5463–5474, 2021. 1, 2
- [55] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. Upsnet: A unified panoptic segmentation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8818–8826, 2019. 1, 2
- [56] Chang-Dong Xu, Xing-Ran Zhao, Xin Jin, and Xiu-Shen Wei. Exploring categorical regularization for domain adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11724–11733, 2020. 3
- [57] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020. 3, 6, 7
- [58] Fuxun Yu, Di Wang, Yinpeng Chen, Nikolaos Karianakakis, Pei Yu, Dimitrios Lymberopoulos, and Xiang Chen. Unsupervised domain adaptation for object detection via cross-domain semi-supervised learning. *arXiv preprint arXiv:1911.07158*, 2019. 3
- [59] Jingyi Zhang, Jiaying Huang, Zichen Tian, and Shijian Lu. Spectral unsupervised domain adaptation for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9829–9840, 2022. 3
- [60] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. 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*, pages 12414–12424, 2021. 3
- [61] Qiming Zhang, Jing Zhang, Wei Liu, and Dacheng Tao. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In *Advances in Neural Information Processing Systems*, pages 433–443, 2019. 3
- [62] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2020–2030, 2017. 3
- [63] Yangtao Zheng, Di Huang, Songtao Liu, and Yunhong Wang. Cross-domain object detection through coarse-to-fine feature adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13766–13775, 2020. 3
- [64] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018. 3
- [65] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5982–5991, 2019. 3, 6, 7