

Detect, Augment, Compose, and Adapt: Four Steps for Unsupervised Domain Adaptation in Object Detection

Mohamed L. Mekhalfi
mmekhalfi@fbk.eu

Davide Boscaini
dboscaini@fbk.eu

Fabio Poiesi
poiesi@fbk.eu

Technologies of Vision Lab
Fondazione Bruno Kessler, Trento, Italy

Abstract

Unsupervised domain adaptation (UDA) plays a crucial role in object detection when adapting a source-trained detector to a target domain without annotated data. In this paper, we propose a novel and effective four-step UDA approach that leverages self-supervision and trains source and target data concurrently. We harness self-supervised learning to mitigate the lack of ground truth in the target domain. Our method consists of the following steps: (1) identify the region with the highest-confidence set of detections in each target image, which serve as our pseudo-labels; (2) crop the identified region and generate a collection of its augmented versions; (3) combine these latter into a composite image; (4) adapt the network to the target domain using the composed image. Through extensive experiments under cross-camera, cross-weather, and synthetic-to-real scenarios, our approach achieves state-of-the-art performance, improving upon the nearest competitor by more than 2% in terms of mean Average Precision (mAP). The code is available at <https://github.com/MohamedTEV/DACA>.

1 Introduction

Domain Adaptation in object detection aims to adapt an object detection model, trained on a data distribution (source), to perform well on a different data distribution (target domain). The challenge of this process arises not only from domain shift but also from factors such as the high variability in object statistics within the target domain. For example, target objects with a different appearance from that of source objects are more difficult to adapt. Another challenge is the data bias across domains, such as the disparity in the quantity of source objects versus the abundance of objects in the target domain. When the target domain possesses a few ground-truth annotations, these can be used as initial seeds to enhance the model's generalization capabilities on the target, which has been shown to be effective in past research works [11, 82]. However, in real applications the target domain often lacks annotations, and this leads to Unsupervised Domain Adaptation (UDA) techniques [42].




Figure 1: Comparison between our method (DACA) and ConfMix [20]. Green bounding boxes are ground-truth annotations, red bounding boxes are pseudo-detections. S_{rowcol} are source regions, T_{rowcol} are target regions. T_{12} in this example corresponds to the most confident target region, while the remaining transparent regions of the target image are discarded. $f_1(T_{12})$, $f_2(T_{12})$, $f_3(T_{12})$, $f_4(T_{12})$ are augmented versions of T_{12} . Unlike DACA, ConfMix requires source images to build the mixed image.

UDA can be tackled through style transfer techniques to mitigate domain shift [23, 25]. Style transfer involves modifying the appearance of a source image to resemble the aesthetic of a target image, while retaining the semantic content. Then, a detector is trained on this intermediate data domain and deployed on the target domain. Despite its simplicity, this strategy can encounter limitations when the target domain’s style is inconsistent. In particular, the source may successfully emulate the style of a subset of the target while remaining misaligned with other underlying subsets. An alternative approach that can address this issue is feature alignment, which tackles the problem at the feature level rather than at the image appearance level [52, 41, 48, 50]. In addition to feature alignment, other UDA methods can adapt a detector’s model by using detections as pseudo-labels that are generated from target images [20, 42]. However, a too inclusive pseudo-label selection (i.e. lenient selection criterion) may accumulate false positives, causing unreliable self-training. Conversely, implementing strict measures for pseudo-label selection may lead to few true positives, making the detector to remain overly reliant on the source domain. In either scenario, the pseudo-labels could mislead the detector or cause overfitting. Recent research has shown that data augmentation can yield plausible improvements in classification and segmentation tasks [9, 36, 43, 44]. Approaches based on data-augmentation tackle the distribution shift problem at input level, without requiring modifications to the detector’s architecture. Our approach is a novel technique that belongs to the family of augmentation-based approaches.

In this paper, we propose a novel approach for UDA in object detection. Our approach is based on four simple processing steps: *detect*, *augment*, *compose*, and *adapt*. We refer to our approach as DACA. DACA is inspired by the idea that adaptation can be made effective by generating high-quality pseudo-labels from the target and using them to supervise augmented versions of the target itself to align the model to the target data distribution. Specifically, we exploit the region of target images with the most confident object detections to gather pseudo-labels for self-training. We create a composite image by applying a collection of random augmentations to the original target image region. This mechanism allows the detector to self-train with confident pseudo-labels applied to augmented images under domain shift. We evaluate DACA on the popular benchmarks Cityscapes [2], FoggyCityscapes [24], Sim10K [6], and KITTI [1]. DACA achieves state-of-the-art performance, improving upon the nearest competitor by more than 2% in terms of mean Average Precision (mAP). Fig. 1

illustrates the idea of DACA. To summarize, our contributions are:

- Our approach is the first alternative to mix up approaches that does not mix images from different domains, but instead generates difficult and informative composite images only from the unsupervised target images.
- We devise a novel approach for UDA based on self-supervision. DACA generates the composite image based on augmented versions of the target image region with the most confident detections, making the adaptation more effective.

2 Related Work

2.1 UDA for object detection

Domain adaptation has been an active area of research to compensate for the distribution mismatch between a source and target data in classification and segmentation endeavours [5, 11, 13, 19, 32, 35, 46, 52]. It has assumed several adaptation forms such as multi-source adaptation [27], few-shot adaptation [20], and source-free adaptation [39]. For instance, Inoue *et al.* [8] tackle the problem of weakly-supervised domain adaptation where instance-level annotations are provided at the source and image-level annotations are envisioned at the target. Wang *et al.* [33] leverage synthetic data to carry out the adaptation.

As per UDA in object detection, Zhu *et al.* [50] align features of object-prone regions across disjoint domains. Li *et al.* [18] present a student-teacher approach, where the student model ensures cross-domain consistent features via adversarial training. SC-UDA [42] translates source images to an intermediate domain by CycleGAN to close the style gap with the target, while content gap is reduced by fusing multiple pseudo-detections produced with embedded stochastic inference. Soviany *et al.* [45] adopt curriculum learning, which aims to progressively adapt a source-trained model to a target, starting with easy images. Image difficulty was determined based on the number of objects and their size. CycleGAN was also trained to stylize source images to minimize style shift.

In this context, several UDA for object detection works seem to focus on feature alignment across domains. Recent approaches mix up source and target images to conduct the adaptation procedure. To the best of our knowledge, ConfMix [24] is the first approach in this direction. ConfMix copy-pastes confident regions along with its detections from the target image to the source image, and supervises the model with a self-supervised loss on the mixed image, scaled by a confidence-induced factor, as well as a supervised loss the source image. Although ConfMix is effective, we argue that this is a rather preliminary approach for UDA and there is significant room for improvement. For example, the source image is used redundantly during training: the network processes it nearly twice at each iteration, once for the “source pass”, and once for the “target pass” in its mix up form (see Fig. 1). We believe that this may hinder the network to adapt its statistics to those of the target domain. Moreover, ConfMix does not leverage data augmentation on the target domain to make adaptation more effective. DACA trains a detector with confident target image regions only. We believe that the source data is unnecessary and potentially deleterious during the target pass.

2.2 Mixing under domain shift

The benefit of data augmentation is evident in prior art [49, 53]. Cutout [8] is a regularization strategy that models object occlusion. It simply masks out random square regions of training

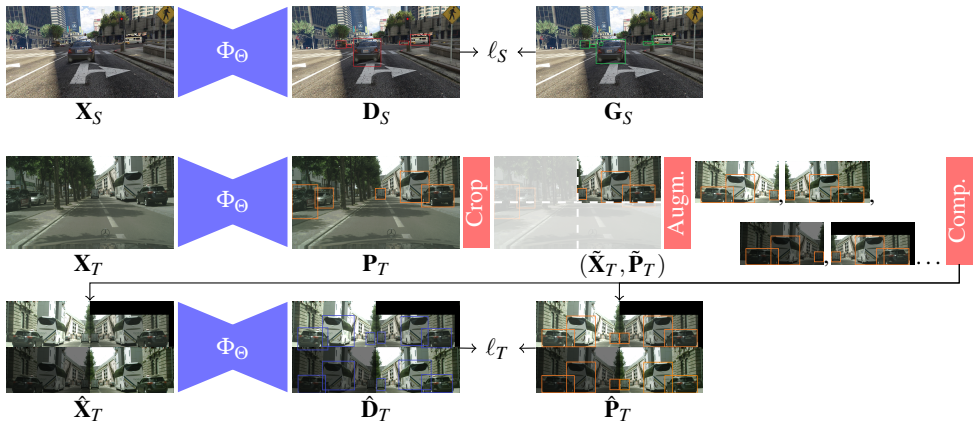


Figure 2: Overview of our proposed approach (DACA). The detector maintains its knowledge on the source by supervision via ℓ_S , and jointly adapts towards the target via ℓ_T using the augmented target image and its pseudo-detections. **Keys:** Comp.: composition. Augm.: augmentation. Bounding boxes colors: red: source detections, green: source ground-truth, orange: target pseudo-detections, blue: composite target detections.

images and has shown to improve the overall performance of deep models. Mixup [24] produces element-wise convex combinations of pairs of image examples and their labels. These techniques have been common to address regularization issues in vision tasks so far [1]. Regarding augmentation-based UDA for object detection, the literature is still developing [5, 4]. For instance, CutMix [43] was devised primarily for classification purposes, and then tailored to object detection [20]. It replaces a random portion of an image with a randomly sampled portion from another image. However, this may override important source information (e.g. objects and their bounding boxes, as well as background context) for the adaptation procedure. Moreover, the region that is randomly cropped from the target may not carry detections to be used as pseudo-labels. ConfMix [20] partially addresses the above problem by selecting a confident region from the target instead of a random one (see Fig. 1). See description of ConfMix and its limitations in the previous section.

DACA aims to settle the shortcomings of mix up approaches at once. The target image combines random augmentations of a confident target region only. We empirically show that our approach achieves state-of-the-art results.

3 Our approach

3.1 Overview

Given a detector Φ_Θ pre-trained on source data, DACA performs adaptation by self-training Φ_Θ on challenging composite images obtained by combining different augmentations of the same crop extracted from target data. We adapt the detector to the target distribution by imposing consistency between the detections predicted on the composite image and the detections of the most confident image region transformed accordingly to the augmentations.

3.2 Detect, augment, compose, and adapt

We define a pair of annotated source data as $(\mathbf{X}_S, \mathbf{G}_S)$, where $\mathbf{X}_S \in \mathbb{R}^{H \times W \times 3}$ is an RGB image sampled from the source data distribution S and $\mathbf{G}_S = \{(\mathbf{B}_i, c_i) : i = 1, \dots, N_S\}$ is the list of its N_S ground-truth bounding boxes, where $\mathbf{B}_i \in \mathbb{R}^4$ contains the coordinates of two opposite corners of the i th bounding box, and $c_i \in [1, \dots, C]$ is a label representing the category of the object contained inside it. We denote an RGB image sampled from the target data distribution T as $\mathbf{X}_T \in \mathbb{R}^{H \times W \times 3}$. Without loss of generality, we assume \mathbf{X}_S and \mathbf{X}_T to have the same size. A neural network model designed to perform object detection can be defined as a parametric function $\Phi_\Theta : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^{M \times 4} \times [1, \dots, C]^M$ with learnable parameters Θ that takes as input an RGB image \mathbf{X} and produces as output a list of M detections of the objects of interest contained in it, i.e. $\Phi_\Theta(\mathbf{X}) = \mathbf{D}$.

3.2.1 Finding reliable pseudo-detections

Given a target image \mathbf{X}_T , we first feed it to the detector Φ_Θ (pre-trained on source data), to obtain the detections $\mathbf{P}_T \in \mathbb{R}^{M_T \times 4} \times [1, \dots, C]^{M_T}$. Because we want to use \mathbf{P}_T as pseudo-labels to self-train our model, it is convenient to find a way to retain good detections and filter out the ones with low confidence which are more likely to be false positives. To this end, we first divide the input image \mathbf{X}_T into S_{row} rows and S_{col} columns, creating a grid of $S_{\text{row}} \times S_{\text{col}}$ cells. We then assign to each cell $\mathbf{X}_T^{ij} \in \mathbb{R}^{H/S_{\text{row}} \times W/S_{\text{col}} \times 3}$, $i = 1, \dots, S_{\text{row}}$, $j = 1, \dots, S_{\text{col}}$, a confidence value obtained by averaging the confidence of the bounding boxes whose center belongs to it. The most confident cell $\tilde{\mathbf{X}}_T$ is cropped from \mathbf{X}_T and the portions of its bounding box detections that fall outside the selected region are trimmed accordingly to obtain $\tilde{\mathbf{P}}_T$ (Fig. 2, Crop).

3.2.2 Composing augmented versions

Since the most reliable detections $\tilde{\mathbf{P}}_T$ occupy only a small portion of \mathbf{X}_T , we perform the following operations: (i) we fill the remaining cells with useful information to avoid wasting computation, (ii) we sample such information from the target domain to balance source and target data during adaptation, and (iii) we fill remaining cells with augmented versions of $\tilde{\mathbf{X}}_T$ and preserve pseudo-labels reliability because we transform $\tilde{\mathbf{P}}_T$ according to the augmentation function chosen, as opposed to estimating them over potentially highly distorted versions of $\tilde{\mathbf{X}}_T$. To this end, we first perturb the pair $(\tilde{\mathbf{X}}_T, \tilde{\mathbf{P}}_T)$ by different data augmentations $f_1, f_2, \dots, f_{S_{\text{row}} S_{\text{col}}}$, randomly selected from a list of predefined ones (Fig. 2, Augm.). Throughout our experiments the list of predefined data augmentations includes horizontal flipping, cropping, blurring, color jittering, downscaling, and perturbing brightness and contrast (Tab. 1), but it can be expanded or tailored according to the downstream task. A composite image $\hat{\mathbf{X}}_T = g(\tilde{\mathbf{X}}_T) \in \mathbb{R}^{H \times W \times 3}$ is then created by composing the augmented versions of $\tilde{\mathbf{X}}_T$ in a $S_{\text{row}} \times S_{\text{col}}$ grid as

$$g: \mathbb{R}^{H/S_{\text{row}} \times W/S_{\text{col}} \times 3} \rightarrow \mathbb{R}^{H \times W \times 3}, \quad g(\tilde{\mathbf{X}}_T) = \begin{bmatrix} f_1(\tilde{\mathbf{X}}_T) & \dots & f_{S_{\text{col}}}(\tilde{\mathbf{X}}_T) \\ \vdots & \ddots & \vdots \\ f_{(S_{\text{row}}-1)S_{\text{col}}+1}(\tilde{\mathbf{X}}_T) & \dots & f_{S_{\text{row}}S_{\text{col}}}(\tilde{\mathbf{X}}_T) \end{bmatrix},$$

obtaining an image of the same size of \mathbf{X}_T (Fig. 2, Comp.). Composite pseudo-detections are obtained in the same way, i.e. $\hat{\mathbf{P}}_T = g(\tilde{\mathbf{P}}_T)$ (Fig. 2, Comp.).

3.2.3 Loss function

The core idea behind DACA is to use $\hat{\mathbf{P}}_T$ as pseudo-detections to self-train Φ_Θ . To do so, we pass $\hat{\mathbf{X}}_T$ in input to Φ_Θ to obtain the detections $\hat{\mathbf{D}}_T$ and promote consistency between $\hat{\mathbf{D}}_T$ and $\hat{\mathbf{P}}_T$ by minimizing the loss ℓ_T . Performing adaptation by minimizing ℓ_T alone may incur in two problems: (i) catastrophic forgetting [14], i.e. the knowledge learned by the detector on S during the pre-training phase can be quickly forgotten, and (ii) potential false positive pseudo-detections on target data can undermine the performance. To cope with both issues, during the adaptation phase we minimize the combined loss $\ell = \ell_S + \ell_T$, where ℓ_S promotes consistency between source detections $\mathbf{D}_S = \Phi_\Theta(\mathbf{X}_S)$ and its ground-truth annotations \mathbf{G}_S (Fig. 2, first row).

4 Experiments

4.1 Datasets

We evaluate and compare our method under three benchmark adaptation scenarios, which include four datasets: (i) Cityscapes (C) [14]: it contains urban images acquired from different cities. Annotations are provided for the eight object classes Person, Rider, Car, Truck, Bus, Train, Motorcycle, and Bicycle; (ii) FoggyCityscapes (F) [24]: it is a variation of Cityscapes obtained by applying a synthetic fog filter with three different intensity levels; (iii) Sim10K (S) [9]: it is a synthetic dataset produced by the GTA-V game engine; (iv) KITTI (K) [9]: it is a traffic scenes dataset. We follow the same experimental convention as in previous works [15, 22, 30, 42, 50]. In particular, the first experimental scenario envisions a synthetic-to-real adaptation task on the Car category where S is Sim10K and T is Cityscapes (S→C). The second experimental scenario is cross-camera adaptation on the Car category where S is KITTI and T is Cityscapes (K→C). The last experimental scenario is weather-adaptation on Cityscapes’ eight categories (C→F).

4.2 Implementation details

We compare our method with previous works [9, 14, 15, 16, 17, 22, 26, 29, 30, 38, 42, 47, 50] that used Faster R-CNN [8] and FCOS [28] detectors for their evaluation. It is worth-mentioning that, for fair comparison in UDA, it is also important to compare based on the same detector, backbone and code reproduction details (e.g. initial seeds). We directly compare with ConfMix [20], as it is one of the most recent UDA methods for object detection. In particular, we use the one-stage detector YOLOv5s as in [20], using COCO pre-trained weights as initialization and with a confidence threshold of 0.25 to select pseudo-labels. We adopt a batch size of 2 including one source image and one target image, with an image size of 600×600. We train our detector with source data for 20 epochs, and then train with source and target data for 50 more epochs by using DACA. This number of epochs is the same as in [20]. We use Average Precision (AP) and Mean Average Precision (mAP) as evaluation metrics for all the experiments with the Intersection over Union threshold set to 0.5.

4.3 Quantitative results

We report the baseline scores (Source only) in which the model is trained on the source data and applied on the test set of the target, which represents our lower bound of the performance

Transformation	Acronym	Operation	Hyper-parameters
HorizontalFlip	HF	Randomly flips the image horizontally	probability(p)=0.5
BBoxSafeRandomCrop	RC	Randomly crops the image without compromising the bounding boxes	$p=0.2$
Blur	B	Randomly blurs out the image	$p=0.5$
ColorJitter	CJ	Randomly applies color jitter	$p=0.5$, brightness=0.2, contrast=0.2, saturation=0.2, hue=0.2
Downscale	D	Randomly downscales the image	$p=0.5$, minimum scale=0.5, maximum scale=0.99
RandomBrightnessContrast	BC	Randomly alters brightness/contrast of the image	$p=0.5$, brightness limit=0.1, contrast limit=0.1

Table 1: List of image augmentation operations to build our composite input image.

Method	Detector	Backbone	S→C	K→C	C→F	Avg.
CTRP [14]	Faster R-CNN	VGG-16	44.5	43.6	50.1	46.1
CDN [15]	Faster R-CNN	VGG-16	49.3	44.9	50.9	48.4
FL-UDA [16]	Faster R-CNN	VGG-16	43.1	44.6	44.4	44.0
SC-UDA [17]	Faster R-CNN	VGG-16	52.4	46.4	56.0	51.6
SAPN [18]	Faster R-CNN	VGG-16	44.9	43.4	59.8	49.4
MeGA-CDA [19]	Faster R-CNN	VGG-16	44.8	43.0	52.4	46.7
MGA [50]	Faster R-CNN	VGG-16	54.6	48.5	60.6	54.6
IRGG [80]	Faster R-CNN	ResNet-50	43.2	45.7	51.9	46.9
GIPA [68]	Faster R-CNN	ResNet-50	47.6	47.9	54.1	49.9
SSOD [42]	Faster R-CNN	ResNet-50	49.3	47.6	57.2	51.4
EPM [0]	FCOS	ResNet-101	51.2	45.0	57.1	51.1
SCAN [13]	FCOS	VGG-16	52.6	45.8	57.3	51.9
SIGMA [14]	FCOS	VGG-16	53.7	45.8	63.7	54.4
Baseline (Source only)	YOLOv5	CSP-Darknet53	50.4	42.9	54.9	49.4
Oracle (Target only)	YOLOv5	CSP-Darknet53	69.5	69.5	67.9	69.0
ConfMix [44]	YOLOv5	CSP-Darknet53	56.2	51.6	63.0	56.9
DACA (Ours)	YOLOv5	CSP-Darknet53	60.6	54.2	63.0	59.3

Table 2: Comparison of Car detection performance (AP) on the S→C, K→C, and C→F adaptation benchmarks. Keys: C: Cityscapes. F: FoggyCityscapes. S: Sim10K. K: KITTI. Avg.: average across the three adaptation scenarios.

range. We also report the Oracle scores, where the model is trained on the target using its ground truth and tested on the test set of the target, which represents our upper bound of the performance range. We start the experiments with a 2×2 target grid division layout in which the confident target region is augmented four times based on a list of typical image augmentations as listed in Tab. 1.

Tab. 2 shows the quantitative results obtained on the S→C, K→C, and C→F benchmarks for the Car category. DACA outperforms all the previous works. We observe gains of 4.4% and 2.6% in the S→C and K→C cases, respectively, compared to the closest competitor method (ConfMix), achieving new state-of-the-art on both scenarios. In particular, across the three scenarios, our method improves over ConfMix by 2.4% on average.

Tab. 3 shows the quantitative results obtained on the C→F benchmark. DACA improves by 0.3% over ConfMix, while both underperform some of other methods that use different detectors and different backbones (e.g. MGA and SIGMA [14, 50]). Although this adaptation case is particularly challenging due to the high occlusion of objects and high domain shift owed to the presence of fog (confirmed also by the low mAP achieved by the baseline), the gap between DACA and the oracle score is reasonably narrow (3.3%).

We display qualitative examples on the class Car for the K→C and S→S scenarios in Fig. 3 to highlight the fact that DACA produces less false positives compared to ConfMix.

More results can be found in the Supplementary Material.

Method	Detector	Backbone	Person	Rider	Car	Truck	Bus	Train	Motorcycle	Bicycle	mAP
IRGG [60]	Faster R-CNN	ResNet-50	37.4	45.2	51.9	24.4	39.6	25.2	31.5	41.6	37.1
GIPA [55]	Faster R-CNN	ResNet-50	32.9	46.7	54.1	24.7	45.7	41.1	32.4	38.7	39.5
SSOD [62]	Faster R-CNN	ResNet-50	38.8	45.9	57.2	29.9	50.2	51.9	31.9	40.9	43.3
CTRP [16]	Faster R-CNN	VGG-16	32.7	44.4	50.1	21.7	45.6	25.4	30.1	36.8	35.9
CDN [42]	Faster R-CNN	VGG-16	35.8	45.7	50.9	30.1	42.5	29.8	30.8	36.5	36.6
FL-UDA [14]	Faster R-CNN	VGG-16	30.4	51.9	44.4	34.1	25.7	30.3	37.2	41.8	37.0
SC-UDA [14]	Faster R-CNN	VGG-16	38.5	43.7	56.0	27.1	43.8	29.7	31.2	39.5	38.7
SAPN [14]	Faster R-CNN	VGG-16	40.8	46.7	59.8	24.3	46.8	37.5	30.4	40.7	40.9
MeGA-CDA [44]	Faster R-CNN	VGG-16	37.7	49.0	52.4	25.4	49.2	46.9	34.5	39.0	41.8
MGA [61]	Faster R-CNN	VGG-16	43.9	49.6	60.6	29.6	50.7	39.0	38.3	42.8	44.3
EPM [8]	FCOS	ResNet-101	41.5	43.6	57.1	29.4	44.9	39.7	29.0	36.1	40.2
SCAN [13]	FCOS	VGG-16	41.7	43.9	57.3	28.7	48.6	48.7	31.0	37.3	42.1
SIGMA [15]	FCOS	ResNet-50	44.0	43.9	60.3	31.6	50.4	51.5	31.7	40.6	44.2
Baseline (Source only)	YOLOv5	CSP-Darknet53	39.2	38.0	54.9	12.4	33.1	6.2	19.9	33.6	29.7
Oracle (Target only)	YOLOv5	CSP-Darknet53	45.6	43.0	67.9	30.2	48.0	39.4	30.3	37.5	42.7
ConfMix [20]	YOLOv5	CSP-Darknet53	44.0	43.3	63.0	30.1	43.0	29.6	25.5	34.4	39.1
DACA (Ours)	YOLOv5	CSP-Darknet53	41.9	40.8	63.0	29.4	42.2	37.2	27.8	33.0	39.4

Table 3: Object detection performance (APs) on the C→F scenario. The last column reports the mean average precision (mAP) across all the object categories. DACA improves slightly over its direct competitor ConfMix [20] while both score the highest AP on the Car class.




Figure 3: Qualitative instances of DACA as compared to the ground truth and ConfMix [20].

4.4 Ablation study

4.4.1 Effect of different augmentation types

Tab. 4 (left) shows an ablation study on the six augmentation types listed in Tab. 1 and their combinations. We can observe that horizontal flipping (HF) performs better than the other augmentations on average. Downscale (D) and blurring (B) are the next best performing augmentations, they improve the model’s robustness to scale and occlusion. Combining all the augmentations together performs the best in all benchmarks. Interestingly, turning off all the augmentations and using raw images (the None case in Tab. 4 (left)) lowers the mAP significantly, which underlines the benefit of using data augmentations for UDA.

4.4.2 Effect of grid layout

Tab. 4 (top right) shows an ablation on the grid layout used to find reliable detections for our best setting (the All case in Tab. 4 (left)). We can see that a 3×3 grid scores the lowest mAP,

Transf.	C→F	K→C	S→C	Avg.
None	33.5	52.8	57.4	47.9
HF	38.0	52.9	59.7	50.2
RC	34.3	52.2	59.4	48.7
B	35.9	53.2	58.4	49.2
CJ	34.6	52.5	58.2	48.5
D	35.3	54.1	59.8	49.8
BC	33.9	52.6	57.5	48.0
HF+B	38.9	53.9	59.3	50.7
HF+D	35.4	53.3	56.7	48.5
D+B	36.8	53.5	59.9	50.0
HF+D+B	37.8	54.0	57.1	49.6
All	39.4	54.2	60.6	51.4

Grid layout	C→F	K→C	S→C	Avg.
3×3	37.8	51.2	57.7	48.9
2×3	38.5	51.7	58.7	49.6
3×2	38.6	53.6	59.9	50.7
2×2	39.4	54.2	60.6	51.4

Num. regions	C→F	K→C	S→C	Avg.
1	35.4	51.5	56.6	47.8
2	38.3	52.2	58.5	49.7
3	39.1	53.1	60.2	50.8
4	39.4	54.2	60.6	51.4

Table 4: Ablation results in terms of mAP for the three adaptation benchmarks. Left: Effect of augmentation types. Top right: Effect of grid layout. Bottom right: Effect of the number of augmented regions (in a 2×2 grid) to build the composite image. C: Cityscapes, F: FoggyCityscapes, S: Sim10K, and K: KITTI. Avg. Average across the three adaptation scenarios




Figure 4: Qualitative instances depicting the effect of the number of augmented regions used to build the composite image.

we believe this happens because it produces a small confident region and consequently few objects and small spatial context. On the other hand, a 3×2 grid is slightly better than a 2×3 one, suggesting that the horizontal context is more informative than the vertical one due to the presence of cars and roads in the target datasets. Finally, a 2×2 grid performs best, being a good trade-off of the previous layouts.

4.4.3 Does the number of augmented regions matter?

Tab. 4 (bottom right) shows an ablation on the number of augmented regions for our best setting (the All case in Tab. 4 (left)). We can observe that lowering the number of augmented regions that are used to build the composite image incurs a drastic performance decay. Fig. 4 shows some qualitative results obtained in the S→C adaptation scenario when using only one augmented region versus when four regions are used in a 2×2 grid layout. Using only one region compromises the robustness of the detector and yields more false positive detections.

5 Conclusions

We presented a four-step approach for unsupervised domain adaptation in object detection. We exploit regions of target images that have high-confidence object detections to collect pseudo-labels for self-training. We generate a new image through the composition of different

augmented versions of the selected image region and use its pseudo-labels to supervise the training of this composite image, which allows us to adapt the detector towards the target domain in a self-supervised fashion. The adaptation of the detectors is carried out simultaneously with the training of source images to prevent knowledge drift of the detector. We evaluated our approach on several benchmarks for UDA and compared it against alternative approaches. Results show that our approach can outperform the other approaches.

DACA augments a confident region four times, which entails that potential false positive detections will multiply four times too, thus accumulating spurious information in the self-supervision. To remedy this, one can anticipate DACA with an image style translation step (□) to mitigate the severity of cross-domain discrepancy. Another potential direction is to consider confident bounding boxes instead of confident regions to improve DACA's composition step by reducing the number of false positives negatively affecting the self-training process.

Acknowledgements

We are very grateful to the support by European Union's Horizon Europe research and innovation programme under grant agreement No. 101092043, project AGILEHAND (Smart Grading, Handling and Packaging Solutions for Soft and Deformable Products in Agile and Reconfigurable Lines).

References

- [1] M. Chen, M. Lin, Z. Lin, Y. Zhang, F. Chao, and R. Ji. SMMix: Self-Motivated Image Mixing for Vision Transformers. *arXiv preprint arXiv:2212.12977*, 2022.
- [2] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The Cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [3] T. DeVries and G. W. Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [4] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *CVPR*, 2012.
- [5] M. Ghifary, W. B. Kleijn, M. Zhang, D. Balduzzi, and W. Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *ECCV*, 2016.
- [6] R. Girshick. Fast R-CNN. In *ICCV*, 2015.
- [7] C. Hsu, Y. Tsai, Y. Lin, and M. Yang. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *ECCV*, 2020.
- [8] N. Inoue, R. Furuta, T. Yamasaki, and K. Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *ICCV*, 2018.
- [9] M. Johnson-Roberson, C. Barto, R. Mehta, S. N. Sridhar, K. Rosaen, and R. Vasudevan. Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? In *ICRA*, 2017.

- [10] B. Kang, Z. Liu, X. Wang, F. Yu, J. Feng, and T. Darrell. Few-shot object detection via feature reweighting. In *ICCV*, 2019.
- [11] G. Kang, L. Jiang, Y. Yang, and A. G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *CVPR*, 2019.
- [12] R. Kemker, M. McClure, A. Abitino, T. Hayes, and C. Kanan. Measuring catastrophic forgetting in neural networks. In *AAAI*, 2018.
- [13] W.M. Kouw and M. Loog. A review of domain adaptation without target labels. *TPAMI*, 43(3), 2019.
- [14] C. Li, D. Du, L. Zhang, L. Wen, T. Luo, Y. Wu, and P. Zhu. Spatial attention pyramid network for unsupervised domain adaptation. In *ECCV*, 2020.
- [15] W. Li, X. Liu, X. Yao, and Y. Yuan. Scan: Cross domain object detection with semantic conditioned adaptation. In *AAAI*, 2022.
- [16] W. Li, X. Liu, and Y. Yuan. Sigma: Semantic-complete graph matching for domain adaptive object detection. In *CVPR*, 2022.
- [17] X. Li, W. Chen, D. Xie, S. Yang, P. Yuan, S. Pu, and Y. Zhuang. A free lunch for unsupervised domain adaptive object detection without source data. In *AAAI*, 2021.
- [18] Y. Li, X. Dai, C. Ma, Y. Liu, K. Chen, B. Wu, Z. He, K. Kitani, and P. Vajda. Cross-domain adaptive teacher for object detection. In *CVPR*, 2022.
- [19] D. Mahapatra, S. Korevaar, B. Bozorgtabar, and R. Tennakoon. Unsupervised domain adaptation using feature disentanglement and gcns for medical image classification. In *ECCV*, 2023.
- [20] G. Mattolin, L. Zanella, E. Ricci, and Y. Wang. ConfMix: Unsupervised Domain Adaptation for Object Detection via Confidence-based Mixing. In *WACV*, 2023.
- [21] S. Motiian, Q. Jones, S. Iranmanesh, and G. Doretto. Few-shot adversarial domain adaptation. In *NIPS*, 2017.
- [22] F. Rezaeianaran, R. Shetty, R. Aljundi, D.O. Reino, S. Zhang, and B. Schiele. Seeking similarities over differences: Similarity-based domain alignment for adaptive object detection. In *ICCV*, 2021.
- [23] A.L. Rodriguez and K. Mikolajczyk. Domain adaptation for object detection via style consistency. *arXiv preprint arXiv:1911.10033*, 2019.
- [24] C. Sakaridis, D. Dai, and L. Van Gool. Semantic foggy scene understanding with synthetic data. *IJCV*, 2018.
- [25] P. Soviany, R. Ionescu, P. Rota, and N. Sebe. Curriculum self-paced learning for cross-domain object detection. *CVIU*, 2021.
- [26] P. Su, K. Wang, X. Zeng, S. Tang, D. Chen, D. Qiu, and X. Wang. Adapting object detectors with conditional domain normalization. In *ECCV*, 2020.

- [27] S. Sun, H. Shi, and Y. Wu. A survey of multi-source domain adaptation. *Inf. Fusion*, 2015.
- [28] Z. Tian, C. Shen, H. Chen, and T. He. FCOS: Fully convolutional one-stage object detection. In *ICCV*, 2019.
- [29] V. S. Vibashan, V. Gupta, P. Oza, V.A. Sindagi, and V.M. Patel. Mega-cda: Memory guided attention for category-aware unsupervised domain adaptive object detection. In *CVPR*, 2021.
- [30] V.S. Vibashan, P. Oza, and V.M. Patel. Instance Relation Graph Guided Source-Free Domain Adaptive Object Detection. In *CVPR*, 2023.
- [31] T. Vu, B. Sun, B. Yuan, A. Ngai, Y. Li, and J. Frahm. LossMix: Simplify and Generalize Mixup for Object Detection and Beyond. *arXiv preprint arXiv:2303.10343*, 2023.
- [32] T. Wang, X. Zhang, L. Yuan, and J. Feng. Few-shot adaptive faster R-CNN. In *CVPR*, 2019.
- [33] Y. Wang, R. Guerrero, and V. Pavlovic. D2f2wod: Learning object proposals for weakly-supervised object detection via progressive domain adaptation. In *WACV*, 2023.
- [34] G. Wei, C. Lan, W. Zeng, and Z. Chen. Metaalign: Coordinating domain alignment and classification for unsupervised domain adaptation. In *CVPR*, 2021.
- [35] X. Wu, Z. Wu, H. Guo, L. Ju, and S. Wang. DanNet: A one-stage domain adaptation network for unsupervised nighttime semantic segmentation. In *CVPR*, 2021.
- [36] Y. Wu, D. Inkpen, and A. El-Roby. Dual mixup regularized learning for adversarial domain adaptation. In *ECCV*, 2020.
- [37] L. Xiong, M. Ye, D. Zhang, Y. Gan, and Y. Liu. Source data-free domain adaptation for a faster r-cnn. *Pattern Recognit.*, 2022.
- [38] M. Xu, H. Wang, B. Ni, Q. Tian, and W. Zhang. Cross-domain detection via graph-induced prototype alignment. In *CVPR*, 2020.
- [39] S. Yang, Y. Wang, Joost V. D. W., Luis Herranz, and Shangling Jui. Generalized source-free domain adaptation. In *ICCV*, 2021.
- [40] Y. Yang and S. Soatto. FDA: Fourier Domain Adaptation for Semantic Segmentation. In *CVPR*, 2020.
- [41] S. Yao, Q. Kang, M. Zhou, M. J. Rawa, and A. Albeshri. Discriminative manifold distribution alignment for domain adaptation. *IEEE Trans. Syst. Man Cybern.: Syst.*, 2022.
- [42] F. Yu, D. Wang, Y. Chen, N. Karianakis, T. Shen, P. Yu, D. Lymberopoulos, S. Lu, W. Shi, and X. Chen. Sc-uda: Style and content gaps aware unsupervised domain adaptation for object detection. In *WACV*, 2022.
- [43] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019.

- [44] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018.
- [45] W. Zhang, J. Wang, Y. Wang, and F. Wang. ParaUDA: Invariant feature learning with auxiliary synthetic samples for unsupervised domain adaptation. *IEEE Trans. Intell. Transp. Syst.*, 2022.
- [46] Y. Zhang, P. David, H. Foroosh, and B. Gong. A curriculum domain adaptation approach to the semantic segmentation of urban scenes. *TPAMI*, 2019.
- [47] G. Zhao, G. Li, R. Xu, and L. Lin. Collaborative training between region proposal localization and classification for domain adaptive object detection. In *ECCV*, 2020.
- [48] L. Zhao and L. Wang. Task-specific inconsistency alignment for domain adaptive object detection. In *CVPR*, 2022.
- [49] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. Random erasing data augmentation. In *AAAI*, 2020.
- [50] W. Zhou, D. Du, L. Zhang, T. Luo, and Y. Wu. Multi-granularity alignment domain adaptation for object detection. In *CVPR*, 2022.
- [51] X. Zhu, J. Pang, C. Yang, J. Shi, and D. Lin. Adapting object detectors via selective cross-domain alignment. In *CVPR*, 2019.
- [52] Y. Zhu, F. Zhuang, J. Wang, J. Chen, Z. Shi, W. Wu, and Q. He. Multi-representation adaptation network for cross-domain image classification. *Neural Netw.*, 2019.
- [53] B. Zoph, E. D. Cubuk, G. Ghiasi, T. Lin, J. Shlens, and Q. V. Le. Learning data augmentation strategies for object detection. In *ECCV*, 2020.