

ConfMix: Unsupervised Domain Adaptation for Object Detection via Confidence-based Mixing

Giulio Mattolin¹, Luca Zanella², Elisa Ricci^{1,2}, Yiming Wang²

¹University of Trento, Trento, Italy ²Fondazione Bruno Kessler, Trento, Italy

lzanella@fbk.eu

Abstract

Unsupervised Domain Adaptation (UDA) for object detection aims to adapt a model trained on a source domain to detect instances from a new target domain for which annotations are not available. Different from traditional approaches, we propose ConfMix, the first method that introduces a sample mixing strategy based on region-level detection confidence for adaptive object detector learning. We mix the local region of the target sample that corresponds to the most confident pseudo detections with a source image, and apply an additional consistency loss term to gradually adapt towards the target data distribution. In order to robustly define a confidence score for a region, we exploit the confidence score per pseudo detection that accounts for both the detector-dependent confidence and the bounding box uncertainty. Moreover, we propose a novel pseudo labelling scheme that progressively filters the pseudo target detections using the confidence metric that varies from a loose to strict manner along the training. We perform extensive experiments with three datasets, achieving state-of-the-art performance in two of them and approaching the supervised target model performance in the other. Code is available at <https://github.com/giuliomattolin/ConfMix>.

1. Introduction

Object detection is a fundamental task in computer vision which involves the classification and localisation, *e.g.* by bounding boxes, of objects of interest belonging to certain predefined categories. Due to its importance in many applications such as autonomous driving, video surveillance and robotic perception, object detection has received significant attention, leading to the development of several different models [15, 34, 41, 35]. However, as detectors mostly

This work has been supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 957337, and the European Commission Internal Security Fund for Police under grant agreement No. ISFP-2020-AG-PROTECT-101034216-PROTECTOR.

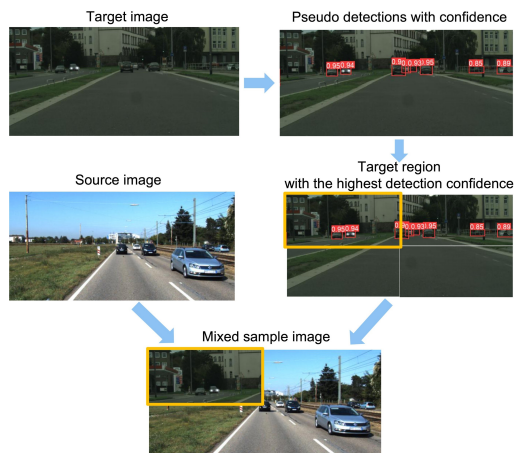


Figure 1. ConfMix is based on a novel sample mixing strategy which combines the source image and the target region (orange box) with the highest pseudo detection confidence.

rely on deep learning, it is a well known fact that they suffer from severe performance degradation when being tested on images that are *visually different* from the ones encountered during training, due to the *domain shift* [4].

To address this problem, recent research efforts have been put on devising Unsupervised Domain Adaptation (UDA) techniques for building deep models that can adapt from an annotated source dataset to a target one without tedious manual annotations [40, 11, 1, 16, 9]. The vast majority of UDA methods for detection resort on adversarial training and on exploiting the Gradient Reversal Layer (GRL) [11] to perform adaptation both at image-level and instance-level [4, 37, 55, 39, 43]. Other approaches mostly focus on robustly producing pseudo detections in order to effectively finetune the model on the target data [49, 52, 44]. In general, while over the last few years several solutions have been proposed in the literature for adapting two-stage object detectors, we argue that devising UDA approaches which can also be applied to one-stage detectors would be desirable. Indeed, the latter methods are more appropriate in applications such as autonomous driving that necessitate of real-time processing and high computational efficiency.

Concurrently, recent works in computer vision have shown the benefit of adopting sophisticated data augmentation techniques by synthesising mixed samples with target and source images in order to improve generalisation ability of deep architectures [51, 50, 17]. These methods have been considered in the context of UDA for classification [45, 30] and semantic segmentation [10, 32, 20, 5], demonstrating some empirical advantage. However, extending these approaches to UDA for detection is far from trivial.

Inspired by these previous works, in this paper we propose **ConfMix**, the first mixing-based UDA approach for object detection based on the regional confidence of pseudo detections. The main idea behind **ConfMix** is illustrated in Figure 1. Specifically, we propose to artificially generate samples by combining the region of target images where the model is most confident with source images. We also introduce during training an associated consistency loss to enforce coherent predictions among generated images. Our intuition is that, by combining source and target images and forming new mixed samples, we are training our model on novel, synthetically generated sample images with reliable pseudo detections and with visual appearance close to the samples of target domain, thus improving the generalisation capabilities of the detector. Moreover, the quality of pseudo detections plays an essential role during adaptation, and is tightly related to the confidence metric. By exploiting a stricter confidence metric, e.g. enriching the detector-dependent confidence with bounding box uncertainty [6], one can obtain more reliable pseudo detections, however with a reduced number. To mitigate this, we propose to progressively restrict the confidence metric for pseudo labelling. With a less strict confidence metric at the initial adaptation phase, we allow more pseudo detections in order to learn the representation of the target domain, while with a gradually stricter confidence metric, we aim to improve the detection accuracy with more trustworthy pseudo detections. We conduct extensive experiments on different datasets (Cityscapes [7] → FoggyCityscapes [38], Sim10K [19] → Cityscapes and KITTI [13] → Cityscapes) and we show that our approach outperforms existing algorithms in most setups.

We summarise our main contributions as below:

- We introduce the first sample-mixing UDA method for object detection. Our approach, named **ConfMix**, mixes samples from source and target domains based on the regional confidence of target pseudo detections.
- We propose a novel Progressive Pseudo-labelling scheme by gradually restricting the confidence metric along the adaptive learning, which allows for a smooth transition when learning target representation, thus improving detection accuracy.
- **ConfMix** scores the new state-of-the-art adaptation performance, achieving +1.7% on Sim10k →

Cityscapes, and +3.7% on KITTI → Cityscapes in terms of mean Average Precision (mAP).

2. Related work

Object Detection. Current object detection models can be grouped into two main categories: one-stage and two-stage approaches. One-stage object detectors, such as YOLO [34] and FCOS [41], adopt a unified framework to obtain final results directly from the feature maps generated by a CNN backbone. These frameworks are very computationally efficient and are able to achieve near real-time speed during inference. On the other hand, two-stage object detectors, such as RCNN [15], generate predictions by first extracting region proposals and then, leveraging this information, produce classification labels and bounding box coordinates. Such models are widely adopted for their high performance but, although research has been conducted to improve detection speed [14, 35, 8], they are considerably slower compared to one-stage detectors.

Unsupervised Domain Adaptation. Given a labelled source domain and an unlabelled target domain, UDA aims to use the available data to produce a model that is able to generalise and perform well on the target domain. A conventional approach is to reduce the domain gap by directly minimising the distance between feature distributions using discrepancy loss functions [29, 40]. On the other hand, adversarial-based methods [11, 12, 42], employ a domain discriminator and a feature extractor that learns to produce domain-invariant feature representations by fooling the discriminator. Many works demonstrated the benefit of using pseudo labels to maximally leverage information from the target domain [28, 23, 24], eventually considering a gradual scheme for incorporating them [48]. Other works have focused on adopting sample mixing techniques, such as mixup [51] or CutMix [50], to improve generalisation. For instance, in [45, 47] domain-level mixup regularisation is applied to ensure domain invariance in the learned feature representations, while in [3, 33] the model’s attention is used to re-assign the confidence of saliency-guided samples and labels. Similar ideas are implemented in previous works considering the segmentation task [10, 32, 20, 5, 31]. However, to the best of our knowledge, no previous works have been proposed to exploit mixing techniques for UDA in the context of object detection.

UDA for Object Detection. In the context of object detection, UDA was recently introduced by [4], which proposed image- and instance-level alignment using two GRLs [11] on Faster R-CNN. Subsequently, several methods started to address this problem mainly using two-stage detectors. Focusing on image-level, [37] showed that strong-local alignment and weak-global alignment of the features extracted from the backbone improve adaptation, while [55], focusing on instance-level, exploited RPN proposals to perform

region-level alignment. To adapt the source-biased decision boundary to the target data, [2] combined adversarial training with image-to-image translation by generating interpolated samples using Cycle-GAN [54]. Other recent works have proposed applying self-training with pseudo detections to perform the adaptation. To address the risk of performance degradation caused by overfitting noisy pseudo detections, [49] introduces an uncertainty-based fusion of pseudo detections sets generated via stochastic inference, [27] proposes self-entropy descent (SED) as a metric to search for an appropriate confidence threshold for reliable pseudo detections, while [44] uses a student-teacher framework and gradually updates the source-trained model.

Few works have addressed UDA for one-stage detectors, e.g. FCOS [25, 26, 18] or SSD [21]. In particular, adopting a self-training procedure reduces the negative effects of inaccurate pseudo detections by performing hard negative pseudo detections mining followed by a weak negative mining strategy, where instance-level scores are computed for each detection considering all neighbouring boxes [21]. In addition, adversarial learning is employed using GRL [11] and a discriminator with the aim of extracting discriminative background features and reducing the domain shift. However, our approach is radically different, as it does not require additional architectural components to the network, but proposes a mixing-based data augmentation strategy to promote regularisation of the model.

3. Method

The proposed **ConfMix**, as illustrated in Figure 2, synthesises an image $x_M \in R^{W \times H \times C}$ by mixing a source image $x_S \in R^{W \times H \times C}$ and the local region of a target image $x_T \in R^{W \times H \times C}$ with the most reliable pseudo detections. We first predict a set of N_T pseudo detections $\tilde{\mathbf{y}}_T = \{\tilde{y}_T^i | i \in [1, N_T]\}$ on the target image and compute the confidence per pseudo detection using the detector network $F(\Theta)$ that is parameterised with Θ and is originally trained only on the source data. We opt to follow a Gaussian modelling of the bounding box predictions, instead of the deterministic one, in order to improve the reliability of the detector confidence with the uncertainty of the bounding box prediction. Next, we divide the target image x_T into regions of equal size and select the region with the highest average confidence of pseudo detections to mix with the source sample x_S , forming the mixed sample x_M .

We pass x_T , x_S , and x_M to the detector $F(\Theta)$ and obtain their corresponding detections $\tilde{\mathbf{y}}_T$, $\tilde{\mathbf{y}}_S$ and $\tilde{\mathbf{y}}_M$, respectively. The detector then learns to adapt to the target domain by imposing a consistency loss L_{cons} which promotes the similarity between $\tilde{\mathbf{y}}_M$ and the combined detections $\tilde{\mathbf{y}}_{S,T}$ by merging the source $\tilde{\mathbf{y}}_S$ and target $\tilde{\mathbf{y}}_T$ detections according to how the two sample images are mixed. The supervision of source ground-truth detections \mathbf{y}_S is achieved with the

detector-related loss L_{det} in order to maintain the detector capability during adaptation.

In the following sections, we describe our proposed **ConfMix** in details, where we first introduce the estimation of the Gaussian-based detection confidence in Sec. 3.1, followed by the confidence-based region mixing strategy for synthesising training samples in Sec. 3.2 and the progressive pseudo labelling in Sec. 3.3. Finally, we present the training objectives with losses in Sec. 3.4.

3.1. Gaussian-based detection confidence

Conventional object detectors, such as YOLO [34], Faster R-CNN [35] and FCOS [41], compute and assign to each detection a confidence score $C_{det} \in [0, 1]$ that is often detector-dependent and is used to filter out unreliable predictions via non-maximum suppression. However, such confidence score does not account for the reliability of the predicted bounding box $\mathbf{b} = [b_x, b_y, b_h, b_w]$, where $[b_x, b_y]$ are the position of bounding box on the image and b_h and b_w represent the height and width, respectively. As suggested in [6], by taking into consideration both the detector-dependent confidence and the confidence that is derived from the uncertainty of bounding box prediction, one can improve the reliability of pseudo detections and reduce the number of false positives.

In order to compute the bounding box uncertainty, \mathbf{b} requires a Gaussian-based modelling. Specifically, for each element in \mathbf{b} , the detector model predicts both a mean μ and a variance Σ , where the variance represents the localisation uncertainty. Thus, we can express the Gaussian-based bounding box $\hat{\mathbf{b}}$ as:

$$\hat{\mathbf{b}} = [\mu_{bx}, \mu_{by}, \mu_{bh}, \mu_{bw}, \Sigma_{bx}, \Sigma_{by}, \Sigma_{bh}, \Sigma_{bw}], \quad (1)$$

where both the means $\hat{\mathbf{b}}_\mu = [\mu_{bx}, \mu_{by}, \mu_{bh}, \mu_{bw}]$ and the variances $\hat{\mathbf{b}}_\Sigma = [\Sigma_{bx}, \Sigma_{by}, \Sigma_{bh}, \Sigma_{bw}]$ are predicted by the detector with an updated regression loss (see details in Sec. 3.4). Note that a sigmoid function $\sigma(\cdot)$ is applied to the predicted variance value to ensure its range is between 0 and 1.

As a larger variance value implies a higher uncertainty, the confidence of a bounding box is computed as:

$$C_{bbx} = 1 - \text{mean}(\hat{\mathbf{b}}_\Sigma), \quad (2)$$

where $\text{mean}(\cdot)$ computes the average variance of $\hat{\mathbf{b}}_\Sigma$.

The combined confidence can thus be computed as:

$$C_{comb} = C_{det} \cdot C_{bbx}. \quad (3)$$

3.2. Confidence-based Region Mixing

With the estimated confidence for each pseudo detection on the target image, we design a novel mixing strategy to synthesise new training samples with highly reliable pseudo

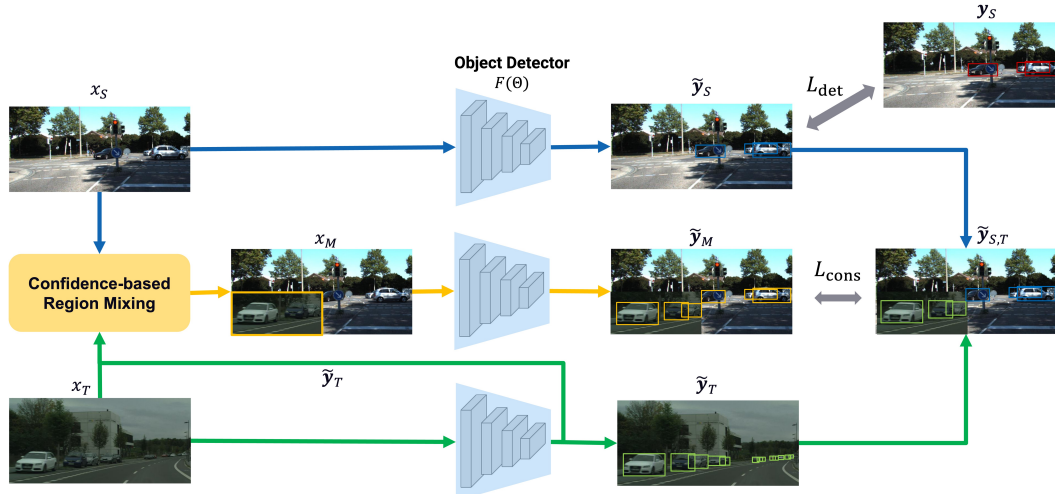


Figure 2. Overview of the proposed **ConfMix** method. We pass the source sample x_S and target sample x_T to the detector model $F(\Theta)$, obtaining the prediction \tilde{y}_S and \tilde{y}_T , respectively. We select the target region with the highest regional confidence to form the mixed sample x_M with the source image, which is then fed to the detector model $F(\Theta)$, producing predictions \tilde{y}_M . We train the model with the supervised detection loss using the source annotations \mathbf{y}_S and the self-supervised consistency loss by comparing \tilde{y}_M with the combined source and target predictions $\tilde{y}_{S,T}$.

detections. Instead of extracting only pseudo detections or randomly selecting part of the target image [50] to mix, we propose a novel region-level mixing strategy whose synthesised samples contain both the foreground and background features from the two domains, contributing to a more effective adaptation towards the target domain.

Specifically, we randomly sample a source image x_S and a target image x_T . The target x_T is then passed to the object detector $F(\Theta)$, producing the predictions \tilde{y}_T . The target image x_T is then evenly divided into 4 regions as shown in Figure 1. Each region is considered to contain a prediction \tilde{y}_T^i if its centre coordinate resides within the region. The region confidence is computed as the average of the confidences of all the pseudo detections that lie within the region. We select the region with the highest region confidence to mix with the sampled source image x_S and generate the synthesised image x_M :

$$x_M = M_T \odot x_S + (1 - M_T) \odot x_T, \quad (4)$$

where $M_T \in R^{W \times H}$ is the mask matrix indicating which pixels of the target image should be masked.

3.3. Progressive Pseudo Labelling

The correctness of the pseudo detections are tightly related to the confidence metric which is used for filtering the detections. At the early stage of the adaptation, the confidence tends to be less reliable and in general of a lower value due to the large domain gap. Thus, the non-maximum suppression ends up in filtering out most of the pseudo detections if a strict confidence metric, such as C_{comb} , is applied. Therefore, we propose to perform a gradual transition

of the confidence metric from a loose to strict manner, to first learn an initial representation of the target domain by allowing more pseudo detections and then gradually shift towards a stricter confidence metric to improve detection accuracy with more reliable pseudo detections.

To this end, we start with the loose confidence metric C_{det} for filtering the pseudo detections. As iterations continue, we progressively assign more importance to C_{comb} with a shifting weight δ :

$$C = (1 - \delta) \cdot C_{det} + \delta \cdot C_{comb}. \quad (5)$$

The shifting weight δ varies based on the progress of the training, thus it is dependent on the iteration t , epoch e and the number of batches in one epoch N_b . We devise δ with a non-linear function to gradually increase from 0 to 1:

$$\delta = \frac{2}{1 + \exp(-\alpha \cdot r)} - 1, \quad (6)$$

$$r = \frac{t}{N_b \cdot e}, \quad (7)$$

where r is the ratio of the current iteration to the total number of iterations, with its scale modulated by α .

The pseudo detection with a confidence value that is higher than a predefined threshold value C_{th} , i.e. $C > C_{th}$, is considered a valid detection, and it will be accounted during the confidence-based region mixing and the training for detector adaptation.

3.4. Adaptive detector training

To facilitate the adaptive learning of the detector $F(\Theta)$, we rely on two main losses: a self-supervised consistency

loss term L_{cons} on the mixed samples and a supervised detector loss term L_{det} on the labelled source samples. L_{det} aims to maintain the task-specific knowledge during adaptation, while the consistency loss L_{cons} aims to adapt towards the target representation by penalising the difference between $\tilde{\mathbf{y}}_M$ and the combined detections $\tilde{\mathbf{y}}_{S,T}$ by merging the source $\tilde{\mathbf{y}}_S$ and target $\tilde{\mathbf{y}}_T$ detections based on how the mixed samples are formed.

Specifically, let $\tilde{\mathbf{y}}_T^R$ be the set of target pseudo detections residing within the selected target region, while $\tilde{\mathbf{y}}_S^{R-}$ be the set of source pseudo detections residing outside the selected target region. The combined detections $\tilde{\mathbf{y}}_{S,T}$ is the union of the two sets, i.e. $\tilde{\mathbf{y}}_{S,T} = \{\tilde{\mathbf{y}}_T^R, \tilde{\mathbf{y}}_S^{R-}\}$. It can happen that the bounding box dimension of $\tilde{y}_T^i \in \tilde{\mathbf{y}}_T^R$ (or $\tilde{y}_S^i \in \tilde{\mathbf{y}}_S^{R-}$) can exceed the selected target (or source) region, leading to inaccurate pseudo detections. We therefore clip such bounding boxes by their corresponding region boundary.

We define L_{cons} and L_{det} as $L_{cons} = \mathcal{L}(\tilde{\mathbf{y}}_M, \tilde{\mathbf{y}}_{S,T})$ and $L_{det} = \mathcal{L}(\tilde{\mathbf{y}}_S, \mathbf{y}_S)$, where both the supervised detection loss L_{det} and the self-supervised consistency loss L_{cons} share the same loss function $\mathcal{L}(\cdot)$. While L_{det} aims to penalise the difference between the predicted detections $\tilde{\mathbf{y}}_S$ and the ground-truth detections \mathbf{y}_S on the source samples, L_{cons} aims to penalise the difference between the predicted detections $\tilde{\mathbf{y}}_M$ and the pseudo detections $\tilde{\mathbf{y}}_{S,T}$ on the mixed samples. Note that $\mathcal{L}(\cdot)$ is dependent on the employed object detector. In the case of the one-stage YOLOv5, $\mathcal{L}(\cdot)$ is a combination of three terms: L_{box} is a Complete-IOU (CIoU) loss for regressing the bounding box coordinates, L_{obj} is the Binary Cross Entropy (BCE) loss for the objectness score and L_{cl} is a BCE loss for the classification score.

In particular, as our predicted bounding box follows a Gaussian modelling, the regression loss per sample image is updated as follows:

$$L_{box} = \frac{1}{N} \sum_{i=1}^N (1 - \text{mean}(\mathcal{N}(y^i | \hat{\mathbf{b}}_\mu^i, \hat{\mathbf{b}}_\Sigma^i))), \quad (8)$$

where $\mathcal{N}(\cdot)$ is the probability density function of a normal distribution for calculating the conditional probability of obtaining the ground-truth $y^i \in \mathbf{y}_S$ for L_{det} , or pseudo detection $y^i \in \tilde{\mathbf{y}}_{S,T}$ for L_{cons} , given the respective means $\hat{\mathbf{b}}_\mu^i$ and variances $\hat{\mathbf{b}}_\Sigma^i$ predicted by the object detector. N represents the total number of $\{y^i\}$.

Finally, the total loss is expressed as a weighted sum of L_{det} and L_{cons} :

$$L_{total} = L_{det} + \gamma L_{cons}, \quad (9)$$

where γ is a hyperparameter to balance the supervised and self-supervised terms. The consistency loss L_{cons} can have a greater importance when the pseudo detections are more reliable, and vice versa. We therefore define γ as the ratio

of the number of pseudo detections on x_M with confidence greater than C_{th}^γ and the total number of pseudo detections after non-maximum suppression, to reflect the reliability of the pseudo detections:

$$\gamma = \frac{\left| \left\{ \tilde{y}_{S,T}^i \in \tilde{\mathbf{y}}_{S,T}^i : C^i \geq C_{th}^\gamma \right\} \right|}{|\tilde{\mathbf{y}}_{S,T}|}, \quad (10)$$

where $|\cdot|$ is the cardinality of a set.

4. Experiments

We evaluate our proposed method **ConfMix** against state-of-the-art methods on three common benchmark adaptation scenarios, together with extensive ablation studies to prove the effectiveness of our design choices.

Datasets. We evaluate our method on the four datasets:

- **Cityscapes [7]** is a collection of urban street scenes for semantic understanding. Images were collected in 50 cities over several months, during the day, and in good weather conditions. Single instance annotations are available for the following 8 categories: person, rider, car, truck, bus, train, motorcycle and bicycle.
- **FoggyCityscapes [38]** is an extension of Cityscapes in which images are augmented by applying a fog filter. FoggyCityscapes includes the same images and 8 categories as Cityscapes.
- **Sim10K [19]** is a synthetic dataset consisting of 10,000 images derived from the video game Grand Theft Auto V, including only the car category.
- **KITTI [13]** is a dataset of several hours of traffic video recorded by high-resolution colour and greyscale cameras, containing 7481 training images with annotations provided for 8 categories: car, van, truck, pedestrian, person sitting, cyclist, tram and misc.

Following [18], we experiment on the benchmark Cityscapes \rightarrow FoggyCityscapes regarding weather adaptation, Sim10K \rightarrow Cityscapes regarding synthetic-to-real adaptation, and KITTI \rightarrow Cityscapes regarding cross-camera adaptation. In the latter synthetic-to-real and cross-camera adaptations, we only consider the car category, while for Cityscapes \rightarrow FoggyCityscapes, we consider the complete 8 categories.

Evaluation Metrics. We evaluate our proposed method on the target domain in terms of Average Precision (AP), which is computed by combining precision and recall for each object category separately. We obtain the mean AP (mAP) by averaging the AP across all object categories.

Implementation details. We based our experiments on the YOLOv5s architecture for its lightness among the YOLOv5 series, using PyTorch and the default settings. We set the batch size to 2, with each batch containing a source image and a target image of the size of 600×600 . In all our experiments, we pre-train the model on the source domain for

20 epochs with the COCO-pretrained weights as initialisation, and perform adaptive learning for 50 epochs. In the non-maximum suppression stage, we set the IoU threshold to 0.5, and the confidence threshold C_{th} to 0.25 for producing pseudo detections. For the computation of γ , we set the confidence threshold C_{th}^γ to 0.5. Please refer details regarding the hyper-parameters in the Supplementary Material.

4.1. Comparisons

We compare **ConfMix** against recent state-of-the-art UDA approaches for adaptive object detection on three benchmarks. In particular, we compare with adversarial feature learning methods, such as MGA [53], MeGA-CDA [43], SSOD [36], EPM [18], CDN [39], SAPN [22]; pseudo-label based self-training techniques such as SC-UDA [49], IRL [44], FL-UDA [27], CTRP [52]; and graph reasoning works such as SCAN [25], SIGMA [26], GIPA [46]. We also include ‘‘Source only’’, the detector model that is trained only with labelled source data, serving as the performance lower-bound, and ‘‘Oracle’’, the detector model that is trained with labelled target data, serving as the performance upper-bound.

Result Discussion. Table 1 reports the results of **ConfMix** and all compared methods in the synthetic-to-real scenario Sim10k \rightarrow Cityscapes and the cross-camera scenario KITTI \rightarrow Cityscapes. On both benchmarks, our **ConfMix** scores the new state-of-the-art adaptation performance, achieving +1.7% on Sim10k \rightarrow Cityscapes, and +3.7% on KITTI \rightarrow Cityscapes in terms of mAP.

Table 2 reports the per-class detection performance of **ConfMix** and all compared methods in the weather adaptation scenario Cityscapes \rightarrow FoggyCityscapes. The gap between **ConfMix** and our upper-bound ‘‘Oracle’’ is rather narrow, i.e. -1.5%. Limited by the low ‘‘Oracle’’ performance on this benchmark, which is lower than MGA [53], it is quite unlikely our method can outperform existing approaches regarding mAP. Nevertheless, our method achieves a +1% AP gain in the person class compared to SIGMA [26] and a +2% AP gain in the car class compared to MGA [53], which are the approaches obtaining the second best AP in these classes, respectively.

Real-time analysis. We perform all experiments on a single NVIDIA Tesla V100. At adaptive training time, each epoch took \sim 13 minutes with a batch size of 2, while training YOLOv5 per epoch the time is \sim 6 minutes on Sim10K \rightarrow Cityscapes. On the Cityscape dataset, the detection speed of **ConfMix** is 76 frames per second, almost equal to the 79 frames per second of YOLOv5.

4.2. Ablation study

We verify the impact of the main design choices of our method with an ablation study on Sim10K \rightarrow Cityscapes.

Does the confidence-based region mixing help? We anal-

Method	Detector	Backbone	Sim10K \rightarrow	KITTI \rightarrow
			Cityscapes mAP	Cityscapes mAP
Source only	YOLOv5	CSP-Darknet53	49.5	39.9
SC-UDA [49]	Faster R-CNN	VGG-16	52.4	46.4
MeGA-CDA [43]	Faster R-CNN	VGG-16	44.8	43.0
FL-UDA [27]	Faster R-CNN	VGG-16	43.1	44.6
CDN [39]	Faster R-CNN	VGG-16	49.3	44.9
SAPN [22]	Faster R-CNN	VGG-16	44.9	43.4
CTRP [52]	Faster R-CNN	VGG-16	44.5	43.6
IRGG [44]	Faster R-CNN	ResNet-50	43.2	45.7
SSOD [36]	Faster R-CNN	ResNet-50	49.3	47.6
GIPA [46]	Faster R-CNN	ResNet-50	47.6	47.9
MGA [53]	FCOS	VGG-16	54.6	48.5
SIGMA [26]	FCOS	VGG-16	53.7	45.8
SCAN [25]	FCOS	VGG-16	52.6	45.8
EPM [18]	FCOS	ResNet-101	51.2	45.0
ConfMix (Ours)	YOLOv5	CSP-Darknet53	56.3	52.2
Oracle	YOLOv5	CSP-Darknet53	70.3	70.3

Table 1. Quantitative results (mAP) for Sim10K/KITTI \rightarrow Cityscapes benchmark.

yse a variety of different mixing strategies and their impact in terms of object detection performance after adaptation. Specifically, we run our method using the C_{comb} confidence without the proposed progressive pseudo labelling. We vary the mixing strategy with 5 different options (as shown in Fig. 4): **CutMix** [50] randomly cuts a target region and mixes it with the source image; **ConfMix** (Vertical Mix) vertically cuts both source and target images in the middle and mixes the most confident target region; **ConfMix** (Horizontal Mix) horizontally cuts both source and target images in the middle and mixes the most confident target region; **ConfMix** (2-region Mix) selects the two most confident regions of the target image for mixing; and **ConfMix** selects only the most confident region of the target image for mixing. We further examine the 4-division scheme of **ConfMix** by varying the number of divided regions into 6 (2×3) and 9 (3×3), respectively.

As can be seen in Table 3, our **ConfMix** achieves a mAP gain of +5.6% compared to CutMix, meaning that considering the most confident target region to mix is more beneficial than randomly cutting a target region and mixing it with the source image for adaptation. **ConfMix** shows also a superior performance than **ConfMix** (Vertical Mix), **ConfMix** (Horizontal Mix), and **ConfMix** (2-region Mix). This means that mixing more target regions with the source image can impact the adaptation performance negatively, which might be due to the inclusion of a larger amount of less confident target pseudo detections. Interestingly, we notice that the cutting direction, i.e. vertical or horizontal, impacts the adaptation performance, where a vertical mixing demonstrates a better adaptation performance than the horizontal mixing; we believe that this phenomenon is scenario-dependent. In particular, for datasets concerning autonomous driving scenarios, a vertical mix always includes the road area of both target and source samples, thus being more likely to include objects, while a horizontal mix might only include further scenes where it is less likely to

Method	Detector	Backbone	person	rider	car	truck	bus	train	motorcycle	bicycle	mAP
Source only	YOLOv5	CSP-Darknet53	34.8	37.6	48.7	14.3	30.1	8.8	14.6	28.1	27.1
SC-UDA [49]	Faster R-CNN	VGG-16	38.5	43.7	56.0	27.1	43.8	29.7	31.2	39.5	38.7
MeGA-CDA [43]	Faster R-CNN	VGG-16	37.7	49.0	52.4	25.4	49.2	46.9	34.5	39.0	41.8
FL-UDA [27]	Faster R-CNN	VGG-16	30.4	51.9	44.4	34.1	25.7	30.3	37.2	41.8	37.0
CDN [39]	Faster R-CNN	VGG-16	35.8	45.7	50.9	30.1	42.5	29.8	30.8	36.5	36.6
SAPN [22]	Faster R-CNN	VGG-16	40.8	46.7	59.8	24.3	46.8	37.5	30.4	40.7	40.9
CTRP [52]	Faster R-CNN	VGG-16	32.7	44.4	50.1	21.7	45.6	25.4	30.1	36.8	35.9
MGA [53]	Faster R-CNN	VGG-16	43.9	49.6	60.6	29.6	50.7	39.0	38.3	42.8	44.3
IRGG [44]	Faster R-CNN	ResNet-50	37.4	45.2	51.9	24.4	39.6	25.2	31.5	41.6	37.1
SSOD [36]	Faster R-CNN	ResNet-50	38.8	45.9	57.2	29.9	50.2	51.9	31.9	40.9	43.3
GIPA [46]	Faster R-CNN	ResNet-50	32.9	46.7	54.1	24.7	45.7	41.1	32.4	38.7	39.5
SCAN [25]	FCOS	VGG-16	41.7	43.9	57.3	28.7	48.6	48.7	31.0	37.3	42.1
SIGMA [26]	FCOS	ResNet-50	44.0	43.9	60.3	31.6	50.4	51.5	31.7	40.6	44.2
EPM [18]	FCOS	ResNet-101	41.5	43.6	57.1	29.4	44.9	39.7	29.0	36.1	40.2
ConfMix (Ours)	YOLOv5	CSP-Darknet53	45.0	43.4	62.6	27.3	45.8	40.0	28.6	33.5	40.8
Oracle	YOLOv5	CSP-Darknet53	42.3	43.9	65.9	33.6	45.0	37.5	29.3	36.7	42.3

Table 2. Quantitative results (mAP) for Cityscapes \rightarrow Foggy Cityscapes benchmark.



Figure 3. Qualitative results on Sim10K \rightarrow Cityscapes scenario of **ConfMix** models trained with different confidence settings. We visualise true positives in blue bbx, false negatives in red bbx and false positives in orange bbx. By introducing a gradual transition between confidence metrics, we achieve less false positive detections compared to training only using C_{det} and less false negative detections compared to training only using C_{comb} .

include objects. Finally, our 4-region division scheme leads to the best adaptation performance, with a mAP improvement of +0.6% and +1.2% compared to 6 and 9 divisions, respectively. This is supported by the intuition that smaller regions produce mixed samples containing a larger portion of the source domain, and increase the probability of occluded target objects, thus limiting the adaptive learning of their complete representation.

Method	mAP
CutMix [50]	49.1
ConfMix (Vertical Mix)	53.6
ConfMix (Horizontal Mix)	39.6
ConfMix (2-region Mix)	41.1
ConfMix (6-division)	54.1
ConfMix (9-division)	53.5
ConfMix	54.7

Table 3. Target detection accuracy with various mixing strategies.

Does the progressive pseudo labelling help? We investigate a set of variants of pseudo labelling in order to verify the proposed progressive pseudo labelling strategy. Specifically, we ablate the usage of only C_{det} or C_{comb} for thresholding the pseudo detections. We also investigate different directions for the weight adjustment, i.e. $C_{det} \rightarrow C_{comb}$ and $C_{comb} \rightarrow C_{det}$, as well as different shifting weights with r (in Equation 7) representing a linear decay and δ (in Equation 6) representing a non-linear decay. Our proposed strategy ($C_{det} \rightarrow C_{comb}$ (δ)) gradually shifts from C_{det} to the stricter C_{comb} using the proposed shifting weight δ .

As shown in Table 4, using only C_{comb} demonstrates to be more advantageous than using only C_{det} for adaptation, and this is mainly due to more reliable pseudo detections. Moreover, we achieve the best result by gradually exploiting from C_{det} to C_{comb} with the non-linear weight δ . A less restrictive confidence metric C_{det} at the early adaptation allows more target pseudo detections and this can help




Figure 4. Illustration of different mixing strategies. **CutMix** [50] randomly cuts a target region and mixes it with the source image. **ConfMix** (Vertical Mix) vertically cuts the source and target image in the middle and mixes the most confident target region. **ConfMix** (Horizontal Mix) horizontally cuts the source and target image in the middle and mixes the most confident target region. **ConfMix** (2-region Mix) selects the two most confident regions of the target image for mixing. Finally, **ConfMix** (6-division), **ConfMix** (9-division), and **ConfMix** divides the target image into 6, 9 and 4 regions, respectively, and selects only the most confident target region for mixing.

with the target representation learning, while by gradually shifting to the usage of C_{comb} , we improve the reliability of pseudo detection, thus benefiting the detector accuracy.

Figure 3 shows qualitative detection results when using only C_{det} , only C_{comb} , and our proposed strategy $C_{det} \rightarrow C_{comb}(\delta)$. As can be observed, the model is more likely to predict false positives when using only C_{det} , whereas the model using only C_{comb} generates more false negatives. Our proposed gradual transition strategy can better combine the two confidence metrics for adaptive training, achieving the best adaptation performance on the target domain.

Confidence	C_{det}	C_{comb}	$C_{comb} \rightarrow C_{det}(t)$	$C_{det} \rightarrow C_{comb}(t)$	$C_{comb} \rightarrow C_{det}(\delta)$	$C_{det} \rightarrow C_{comb}(\delta)$
mAP	52.7	54.7	55.0	54.9	54.3	56.3

Table 4. Target detection accuracy with various confidence-based pseudo labelling.

Does the weight of consistency loss matter? As pseudo detections are inevitably noisy, we are motivated to weight the consistency loss appropriately in order to avoid the introduction of pseudo detection errors. We therefore investigate how the weight of consistency loss impacts the adaptation performance by using a set of constant weights, in comparison to our dynamic weight γ . Regarding the constant weight, we vary it from 0.2 to 1 in the step of 0.2. As shown in Table 5, using γ as the consistency weight leads to the best mAP performance, with an improvement of +0.7% compared to the best constant weight of 0.6. Therefore, the use of a dynamic weight whose value varies depending on the pseudo detections confidence, can improve the performance of the model by stabilising the training and mitigating the problem of over-fitting unreliable pseudo detections.

Does the number of pseudo detections before mixing matter? We analyse the number of pseudo detections be-

Weight	0.2	0.4	0.6	0.8	1	γ
mAP	53.8	54.7	55.6	54.9	47.7	56.3

Table 5. Target detection accuracy with different weights on the consistency loss.

fore mixing by retaining with only 25%, 50%, 75% and 100% (i.e. our setting) of the most reliable pseudo detections in \tilde{y}_T , with the confidence threshold C_{th} fixed for filtering the detections. As shown in Table 6, decreasing the number of pseudo detections consistently leads to a worse result than using all of them, i.e. our setting. Thanks to our progressive pseudo labelling scheme, most of the false positives could be filtered out already, thus detections with a confidence greater than C_{th} are generally useful for the model to learn the target features.

Pseudo detections (%)	25%	50%	75%	100%
mAP	45.1	48.9	51.7	56.3

Table 6. Target detection accuracy at varying number of pseudo detections.

Does sample mixing performs better than simple self-training? We compare **ConfMix** with a simple baseline that consists in applying naive finetuning with pseudo detections. With the model trained on the source domain, we generate pseudo detections on the target dataset using the same confidence threshold C_{th} used in **ConfMix** to filter the boxes on top of non-maximum suppression. We then expand the training dataset containing both source and target samples for further training. With such self-training, we achieve a mAP of 30.5 (Cityscapes→Foggy Cityscapes), 55.4 (Sim10K→Cityscapes) and 46.4 (KITTI→ Cityscapes), which are much inferior than **ConfMix** as reported in Table 1 and 2, confirming the effectiveness of our proposal.

5. Conclusion

We proposed **ConfMix**, a novel confidence-based mixing method for adapting object detectors trained on a source domain to a target domain in an unsupervised manner. We introduced a region-level strategy for sample data augmentation by mixing the region of the target image with the most confident pseudo detections with the source image, and achieved adaptation with a consistency loss on the pseudo detections. We also introduced the progressive pseudo labelling scheme by gradually restricting the confidence metric in order to facilitate the smooth transition from learning the target representation to improving detection accuracy. We compared our approach with state-of-the-art methods, demonstrating its superior performance on two benchmarks. As future work, we will apply our method in other practical scenarios, other than autonomous driving, and improve its compatibility with different object detection frameworks.

A. Supplementary Material

In this Supplementary Material, we provide additional experiments to demonstrate the effect of pretrained detector backbone, i.e. with/without COCO-pretrained weights, on the adaptation performance using our proposed **ConfMix**. Moreover, we also justify key hyperparameter choices in the implementation, including the confidence threshold C_{th} to filter the detections on top of non-maximum suppression, and α that is used to scale the impact of the current iteration to the total number of iterations for the calculation of the shifting weight δ . Consistently with the main manuscript, we conduct the ablation study on the Sim10K \rightarrow Cityscapes setup.

Does backbone initialisation negate the effect of domain adaptation? As there are works using ImageNet-pretrained networks [27, 22, 44, 46, 25, 52, 53, 26, 18] and works that do not specify whether they are pretrained or not [49, 43, 39, 36], we are motivated to examine how the initialisation of our backbone affects the proposed adaptation strategy. Therefore, we experiment with the initialisation of the backbone with random weights. Compared to the setting described in the main manuscript, we only vary the confidence threshold C_{th} used to filter the detections on top of the non-maximum suppression from 0.25 to 0.3 for the random weight setting, in order to account for less reliable predictions at the initial training phase. As shown in Table 7, with random weights initialisation, Source only, **ConfMix** and Oracle achieve lower performance than their corresponding ones with the COCO pretrained weights. Moreover, we notice that **ConfMix** with random weights obtains a mAP gain of +12.3% and +28.4% compared to its Source-only counterpart, on Sim10K \rightarrow Cityscapes and KITTI \rightarrow Cityscapes, respectively. While with COCO-pretrained weights, **ConfMix** achieves a mAP gain of +6.8% and +12.3% compared to its Source-only counterpart, on Sim10K \rightarrow Cityscapes and KITTI \rightarrow Cityscapes, respectively. This shows that the adaptation of **ConfMix** is more effective when the backbone is not pretrained, although its general detection performance on the target domain is bounded by the Oracle’s performance.

How does C_{th} for filtering out detections affect adaptation? We experiment with varying confidence threshold C_{th} , i.e. 0.1, 0.25 (our setting), 0.5, and 0.7, to filter the detections on top of non-maximum suppression. The weight used to balance the consistency loss $\gamma = 1$ is fixed throughout this experiment. As shown in Table 8, using $C_{th} = 0.5$ leads to the best adaptation performance, with a mAP improvement of +10.6% and +10.9% compared to $C_{th} = 0.1$ and $C_{th} = 0.7$, respectively. While the use of low values for C_{th} leads the model to keep erroneous pseudo detections, the use of high values filters out those pseudo detections that are useful for the model to learn the target features.

However, the use of C_{th} is closely related to C_{th}^γ , which

Method	Detector	Backbone	Pretrained	Sim10K \rightarrow	KITTI \rightarrow
				Cityscapes	Cityscapes
Source only	YOLOv5	CSP-Darknet53	No	33.9	21.7
ConfMix (Ours)	YOLOv5	CSP-Darknet53	No	46.2	50.1
Oracle	YOLOv5	CSP-Darknet53	No	64.1	64.1
Source only	YOLOv5	CSP-Darknet53	COCO	49.5	39.9
ConfMix (Ours)	YOLOv5	CSP-Darknet53	COCO	56.3	52.2
Oracle	YOLOv5	CSP-Darknet53	COCO	70.3	70.3

Table 7. Quantitative results (mAP) for Sim10K/KITTI \rightarrow Cityscapes benchmark.

C_{th}	0.1	0.25	0.5	0.7
mAP	42.1	47.7	52.7	41.8

Table 8. Target detection accuracy with different confidence thresholds C_{th} .

C_{th}	0.25	0.25	0.25	0.25	0.25	0.25	0.5	0.5	0.5	0.5
C_{th}^γ	0.4	0.5	0.6	0.7	0.8	0.9	0.6	0.7	0.8	0.9
mAP	55.1	56.3	55.3	56.0	55.6	53.3	55.0	54.4	53.8	51.5

Table 9. Target detection accuracy with different confidence thresholds C_{th}^γ under $C_{th} = 0.25$ and $C_{th} = 0.5$.

α	1	3	5	10
mAP	55.5	55.7	56.3	55.4

Table 10. Target detection accuracy with different confidence transition magnitudes.

is the confidence threshold used to calculate the reliability of pseudo detections γ , defined as the ratio of valid detections after non-maximum suppression with confidence greater than C_{th}^γ . We therefore vary C_{th}^γ from 0.4 to 0.9 under $C_{th} = 0.25$, and from 0.6 to 0.9 under $C_{th} = 0.5$ (note that the value of C_{th}^γ should be larger than C_{th} to function meaningfully). From the results reported in Table 9, we can see that the best performance is given by the combination of $C_{th} = 0.25$ and $C_{th}^\gamma = 0.5$, which is our experimental setting reported in the main manuscript.

How does α for the confidence transition affect adaptation? We analyse the effect of the hyperparameter α used to calculate the shifting weight δ for the smooth transition from C_{det} to C_{comb} during training. As illustrated in Figure 5, a lower value of α gives a greater importance to the less strict confidence C_{det} , while a higher value of α gives a greater importance to the stricter confidence C_{comb} . As shown in Table 10, using $\alpha = 5$ leads to the best mAP performance, with an improvement of +0.8% compared to $\alpha = 1$, +0.6% compared to $\alpha = 3$, and +0.9% compared to $\alpha = 10$.

References

- [1] Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 627–636, 2019.




Figure 5. Visualisation of the evolution of δ throughout the training iterations with different values of α . Note that $\delta = r$ is the linear function already ablated in the main manuscript.

- [2] Chaoqi Chen, Zebiao Zheng, Xinghao Ding, Yue Huang, and Qi Dou. Harmonizing transferability and discriminability for adapting object detectors. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8869–8878, 2020.
- [3] Jie-Neng Chen, Shuyang Sun, Ju He, Philip HS Torr, Alan Yuille, and Song Bai. Transmix: Attend to mix for vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12135–12144, 2022.
- [4] 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 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3339–3348, 2018.
- [5] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jiabin Huang. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1791–1800, 2019.
- [6] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee. Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pages 502–511, 2019.
- [7] 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 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [8] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. *Advances in neural information processing systems*, 29, 2016.
- [9] Zhongying Deng, Kaiyang Zhou, Yongxin Yang, and Tao Xiang. Domain attention consistency for multi-source domain adaptation. *arXiv preprint arXiv:2111.03911*, 2021.
- [10] Geoff French, Samuli Laine, Timo Aila, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, varied perturbations. *arXiv preprint arXiv:1906.01916*, 2019.
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of International Conference on Machine Learning*, pages 1180–1189. PMLR, 2015.
- [12] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- [13] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012.
- [14] Ross Girshick. Fast r-cnn. In *Proceedings of IEEE International Conference on Computer Vision*, pages 1440–1448, 2015.
- [15] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 580–587, 2014.
- [16] Abel Gonzalez-Garcia, Joost Van De Weijer, and Yoshua Bengio. Image-to-image translation for cross-domain disentanglement. *Advances in neural information processing systems*, 31, 2018.
- [17] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019.
- [18] Cheng-Chun Hsu, Yi-Hsuan Tsai, Yen-Yu Lin, and Ming-Hsuan Yang. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *Proceedings of European Conference on Computer Vision*, pages 733–748. Springer, 2020.
- [19] Matthew Johnson-Roberson, Charles Barto, Rounak Mehta, Sharath Nittur Sridhar, Karl Rosaen, and Ram Vasudevan. Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? *arXiv preprint arXiv:1610.01983*, 2016.
- [20] Jongmok Kim, Jooyoung Jang, and Hyunwoo Park. Structured consistency loss for semi-supervised semantic segmentation. *arXiv preprint arXiv:2001.04647*, 2020.
- [21] 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 IEEE/CVF International Conference on Computer Vision*, pages 6092–6101, 2019.
- [22] Congcong Li, Dawei Du, Libo Zhang, Longyin Wen, Tiejian Luo, Yanjun Wu, and Pengfei Zhu. Spatial attention pyra-

- mid network for unsupervised domain adaptation. In *Proceedings of European Conference on Computer Vision*, pages 481–497. Springer, 2020.
- [23] Jichang Li, Guanbin Li, Yemin Shi, and Yizhou Yu. Cross-domain adaptive clustering for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [24] Kai Li, Chang Liu, Handong Zhao, Yulun Zhang, and Yun Fu. Ecacl: A holistic framework for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [25] Wuyang Li, Xinyu Liu, Xiwen Yao, and Yixuan Yuan. Scan: Cross domain object detection with semantic conditioned adaptation. In *AAAI*, volume 6, page 7, 2022.
- [26] Wuyang Li, Xinyu Liu, and Yixuan Yuan. Sigma: Semantic-complete graph matching for domain adaptive object detection. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5291–5300, 2022.
- [27] Xianfeng Li, Weijie Chen, Di Xie, Shicai Yang, Peng Yuan, Shiliang Pu, and Yueting Zhuang. A free lunch for unsupervised domain adaptive object detection without source data. In *Proceedings of Conference on Artificial Intelligence*, volume 35, pages 8474–8481, 2021.
- [28] Jian Liang, Dapeng Hu, and Jiashi Feng. Domain adaptation with auxiliary target domain-oriented classifier. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16632–16642, 2021.
- [29] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *Proceedings of International Conference on Machine Learning*, pages 97–105. PMLR, 2015.
- [30] Xudong Mao, Yun Ma, Zhenguo Yang, Yangbin Chen, and Qing Li. Virtual mixup training for unsupervised domain adaptation. *arXiv preprint arXiv:1905.04215*, 2019.
- [31] Luke Melas-Kyriazi and Arjun K Manrai. Pixmatch: Unsupervised domain adaptation via pixelwise consistency training. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12435–12445, 2021.
- [32] Viktor Olsson, Wilhelm Tranehed, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1369–1378, 2021.
- [33] Joonhyung Park, June Yong Yang, Jinwoo Shin, Sung Ju Hwang, and Eunho Yang. Saliency grafting: Innocuous attribution-guided mixup with calibrated label mixing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 7957–7965, 2022.
- [34] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [36] Farzaneh Rezaeianaran, Rakshith Shetty, Rahaf Aljundi, Daniel Olmeda Reino, Shanshan Zhang, and Bernt Schiele. Seeking similarities over differences: Similarity-based domain alignment for adaptive object detection. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pages 9204–9213, 2021.
- [37] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6956–6965, 2019.
- [38] 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.
- [39] Peng Su, Kun Wang, Xingyu Zeng, Shixiang Tang, Dapeng Chen, Di Qiu, and Xiaogang Wang. Adapting object detectors with conditional domain normalization. In *Proceedings of European Conference on Computer Vision*, pages 403–419. Springer, 2020.
- [40] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *Proceedings of European Conference on Computer Vision*, pages 443–450. Springer, 2016.
- [41] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pages 9627–9636, 2019.
- [42] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [43] Vibashan Vs, Vikram Gupta, Poojan Oza, Vishwanath A Sindagi, and Vishal M Patel. Mega-cda: Memory guided attention for category-aware unsupervised domain adaptive object detection. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4516–4526, 2021.
- [44] Vibashan VS, Poojan Oza, and Vishal M Patel. Instance relation graph guided source-free domain adaptive object detection. *arXiv preprint arXiv:2203.15793*, 2022.
- [45] Yuan Wu, Diana Inkpen, and Ahmed El-Roby. Dual mixup regularized learning for adversarial domain adaptation. In *Proceedings of European Conference on Computer Vision*, pages 540–555. Springer, 2020.
- [46] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12355–12364, 2020.
- [47] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6502–6509, 2020.
- [48] Jeongbeen Yoon, Dahyun Kang, and Minsu Cho. Semi-supervised domain adaptation via sample-to-sample self-distillation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1978–1987, 2022.

- [49] Fuxun Yu, Di Wang, Yinpeng Chen, Nikolaos Karianakis, Tong Shen, Pei Yu, Dimitrios Lymberopoulos, Sidi Lu, Weisong Shi, and Xiang Chen. Sc-uda: Style and content gaps aware unsupervised domain adaptation for object detection. In *Proceedings of IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1061–1070, 2022.
- [50] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019.
- [51] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [52] Ganlong Zhao, Guanbin Li, Ruijia Xu, and Liang Lin. Collaborative training between region proposal localization and classification for domain adaptive object detection. In *Proceedings of European Conference on Computer Vision*, pages 86–102. Springer, 2020.
- [53] Wenzhang Zhou, Dawei Du, Libo Zhang, Tiejian Luo, and Yanjun Wu. Multi-granularity alignment domain adaptation for object detection. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9581–9590, 2022.
- [54] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of IEEE International Conference on Computer Vision*, pages 2223–2232, 2017.
- [55] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 687–696, 2019.