

# Domain Adaptation for Object Detection via Style Consistency

Adrian Lopez Rodriguez  
adrian.lopez-rodriguez15@imperial.ac.uk  
Krystian Mikolajczyk  
k.mikolajczyk@imperial.ac.uk

MatchLab  
Imperial College London  
London, UK

---

## Abstract

We propose a domain adaptation approach for object detection. We introduce a two-step method: the first step makes the detector robust to low-level differences and the second step adapts the classifiers to changes in the high-level features. For the first step, we use a style transfer method for pixel-adaptation of source images to the target domain. We find that enforcing low distance in the high-level features of the object detector between the style transferred images and the source images improves the performance in the target domain. For the second step, we propose a robust pseudo labelling approach to reduce the noise in both positive and negative sampling. Experimental evaluation is performed using the detector SSD300 on PASCAL VOC extended with the dataset proposed in [18], where the target domain images are of different styles. Our approach significantly improves the state-of-the-art performance in this benchmark.

## 1 Introduction

Object detection has been extensively studied in recent years, resulting in the development of highly accurate models such as Faster R-CNN [6] or SSD [25]. Typically, object detectors rely on a generic backbone (e.g. VGG [69] or ResNet [14]) pretrained in ImageNet [5] for classification tasks, and then finetuned on datasets with similar images for both training and evaluation, such as MS COCO [23] and PASCAL VOC [7]. However, such models do not behave well when trained and tested in different domains as shown in [2, 13, 67].

Several approaches have been developed for unsupervised domain adaptation, aiming to train a model using a fully-annotated source dataset  $\mathcal{S}$  and a non-annotated target dataset  $\mathcal{T}$  to perform well in  $\mathcal{T}$ . The proposed methods range from high-level feature adaptation via domain confusion [8, 42] to image adaptation using a generative model [16, 43].

Image adaptation has been shown to be beneficial for closing the domain gap between the source and target domain, especially when the two domains are largely dissimilar in low-level statistics. The generative models usually employed for the domain adaptation task, such as CycleGAN [48], only provide one translated image per source image and cannot translate the source instance to a specific style in the target domain. Some datasets may contain several styles as shown in Figure 1, e.g. changes in texture, shapes, patterns, illumination or colors. In those cases, having the capability of generating several translated source images to any of the target styles may help to close the domain gap. Style transfer methods are capable of



Figure 1: Overview of the unsupervised domain adaptation object detection task. We have a labelled source set  $\mathcal{S}$  and an unlabelled target set  $\mathcal{T}$  that shares classes with  $\mathcal{S}$ . We aim to use both sets to train a detector with good accuracy in  $\mathcal{T}$ .

adapting the input source image in an instance-level manner, making them suitable for the domain adaptation task. Apart from low-level changes, the object appearance can also vary at a high level, *e.g.* shape changes. For good performance on the target domain, we also need to adapt the high-level features once the model is robust to the low-level changes, which can be done pseudo labelling the target domain instances using the detector trained in the source domain.

We propose an approach for domain adaptation of object detectors that relies on two main contributions. First, we use a style transfer method to adapt the source images to the target domain. We force feature consistency between the non-transferred image and the different styles, aiming to make the detector robust to low-level variations. Secondly, we propose robust pseudo labelling to provide labelled training examples from the unannotated target domain, using positive and negative example sampling with low label noise. The pseudo labels are used along with the style transfer method to train on style transferred target images.

The remainder of the paper is organized as follows. We first review related works in the area of object detection, style transfer and domain adaptation in Section 2. Section 3 presents the style transfer method and the object detector used. We then introduce our approach in Section 4. Experimental results and discussions are in Section 5.

## 2 Related work

We present the relevant literature on general object detection task, style transfer and domain adaptation.

**Object detectors** based on deep learning can be categorized in one-stage or two-stage methods. Two well-known two-stage detectors are R-CNN [10] and Fast R-CNN [11], both of which used a low-level vision algorithm for the first stage. Faster R-CNN [12] introduced a Region Proposal Network (RPN), making the model end-to-end trainable. Examples of one-stage detectors are YOLO [13] and SSD [14], which do not have a ROI proposal network, leading to a more efficient inference.

Some works tackled the problem of weakly-supervised object detection [15, 16]. A related problem was investigated in [17, 18, 19], where a set of classes with annotated bounding-boxes are used to detect another set of classes with only image-level labels available.

**Style Transfer** method was proposed in [20], which minimized a content and a style loss. Further extensions, such as [6, 21, 22] used a feed-forward method instead. Contrary to the

previous approaches, in [17] a method was proposed to generalize to unseen styles. Recent developments focused on adapting the style of a real image (e.g. illumination or weather) to another real image [2, 4].

**Domain adaptation** has been actively researched for image classification problems. The approaches developed in [8, 41] sought to confuse a domain classifier to align the features. A generalized framework using an adversarial method and untied weight sharing was introduced in [42]. An adversarial loss was also used in [35] to force alignment of  $\mathcal{T}$  and  $\mathcal{S}$  samples. In [38] Virtual Adversarial Training was used along with a teacher network to refine the decision boundaries. In [34] a co-training approach was applied with three classifiers, trained with source and agreed pseudo labelled target examples. The consistency between predictions for transformations of the input data was used in [33]. Joint pixel and feature adaptation was attempted in [16] using a similar approach to CycleGAN [48].

Domain adaptation for object detection is more challenging and only recently gained some attention. Some works used synthetically generated data to train an object detector for real images [13, 28]. A recent approach [29] focused on merging different fully-annotated datasets in different domains and with different object classes. In [2], instance-level and image-level adaptation was done for Faster R-CNN. In [18] a two-step training method was proposed, where CycleGAN was applied for pixel-level adaptation of PASCAL VOC images to other domains, then used for training a detector on domain adapted images. A new dataset for object detection in domain adaptation was also introduced in [18] with significantly different domain styles (comic, watercolor and clipart), sharing the object categories with PASCAL VOC. A recent approach in [43] developed a framework to adapt GANs for the pixel-level translation in an object detection setting. In [36], the focal loss [24] was used to adapt the high-level features that were more similar between the source and target domains. A style transfer approach that is related to our method was used in [46] for semantic segmentation, however no style consistency on the feature level was enforced.

### 3 Preliminaries

In this section we present the style transfer method and the object detector used in our experiments. We use for most of our experiments PASCAL VOC [7] as  $\mathcal{S}$  and the dataset proposed in [18] as  $\mathcal{T}$ , which contains three domains: Clipart1k, Watercolor2k and Comic2k. In each of these domains, there is a large variation of individual styles in the different instances. In [18], image adaptation from VOC to each of these three domains was performed using a CycleGAN. We will refer to these adapted images as Domain Transferred Images (*DTI*). A more detailed overview of the datasets is given in the supplementary material.

**Style Transfer via Adaptive Instance Normalization.** We adopt the method from [17] to transfer the style of an image  $x_{\mathcal{T}}$  to the image  $x_{\mathcal{S}}$  via a feed-forward network resulting in  $x_{\mathcal{S} \rightarrow \mathcal{T}}$ . A feature map is generated using a VGG19 pretrained in ImageNet up to the layer `relu4_1` for target  $f_v(x_{\mathcal{T}})$ , and source image  $f_v(x_{\mathcal{S}})$ . Next,  $f_v(x_{\mathcal{S}})$  undergoes an affine transformation to match the mean and variance of  $f_v(x_{\mathcal{T}})$  resulting in  $f'_v(x_{\mathcal{S}})$ , which is passed to a decoder that generates  $x_{\mathcal{S} \rightarrow \mathcal{T}}$ . The style transfer loss is the sum of a content loss  $\mathcal{L}_{Cont}$  and a style loss  $\mathcal{L}_{Style}$ :

$$\mathcal{L}_{ST} = \mathcal{L}_{Cont} + \mathcal{L}_{Style} \quad (1)$$

The content loss is defined as  $\mathcal{L}_{Cont} = \|f'_v(x_{\mathcal{S}}) - f_v(x_{\mathcal{S} \rightarrow \mathcal{T}})\|^2$  where  $f_v(x_{\mathcal{S} \rightarrow \mathcal{T}})$  is the VGG19 feature map for  $x_{\mathcal{S} \rightarrow \mathcal{T}}$ . The style loss  $\mathcal{L}_{Style}$  is defined as the Euclidean distance between the

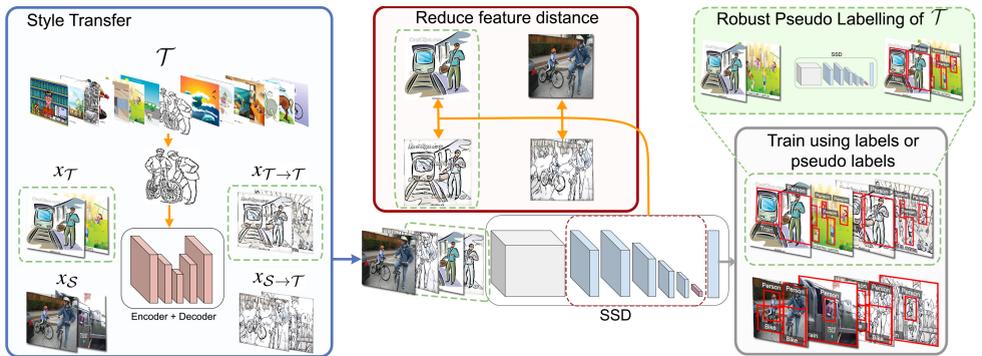


Figure 2: Overview of the proposed approach for domain adaptation of object detectors. We first make the detector robust to low-level variations by using a style transfer method for pixel-level adaptation of the source images, and then forcing feature consistency in the object detector between the style transferred image and the original image. Next, we pseudo label  $\mathcal{T}$ , and retrain following the same method with both images from  $\mathcal{S}$  and  $\mathcal{T}$ .

means and the standard deviations of the feature maps of  $x_{\mathcal{T}}$  and  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  in layers `relu1_1`, `relu2_1`, `relu3_1` and `relu4_1`.

**Single Shot Detector (SSD).** We build our approach on SSD [45] but other methods such as Faster R-CNN [61] or YOLO [60] can be used as our approach addresses the training methodology rather than a specific detector. SSD is a single stage approach with a set of anchor boxes from different feature maps. The loss used during training with image  $x$  is:

$$\mathcal{L}_{SSD}(x) = \frac{1}{N} \sum_b (\mathcal{L}_{conf}(b, s_c) + \alpha \mathcal{L}_{loc}(b, g)) \quad (2)$$

where confidence loss  $\mathcal{L}_{conf}$  is the cross-entropy loss of score  $s_c$  for ground-truth class label  $c$  and matched box  $b$  (Equation 3) averaged over all  $N$  positive matched boxes in image  $x$ .  $\mathcal{L}_{loc}$  is a smooth  $l_1$  loss between the predicted geometric parameters of box  $b$  and the ground-truth box  $g$  for positive examples. Hyperparameters in our experiments include  $\alpha = 1$  and ratio 3:1 of negative to positive object boxes in a given image. More details can be found in [45].  $\mathcal{L}_{conf}$  for a single box  $b$  with predicted score  $s_c$  for correct class label  $c$  is defined as:

$$\mathcal{L}_{conf}(b, s_c) = -\log(s_c) \quad (3)$$

where  $s_c$  is the normalized score applying the softmax function.

## 4 Method

In this section we present our approach, consisting of two steps. In the first step, we train the object detector via style transfer and feature consistency to improve robustness to low-level variations. To adapt it to high-level changes, we provide the detector with pseudo labelled training examples from  $\mathcal{T}$  and finetune again. Figure 2 shows an overview of the method.

## 4.1 Style Transfer with Feature Consistency

Recent studies [10] have revealed that Deep Learning approaches are heavily reliant on low-level features for classification, making low-level adaptation crucial for domain adaptation. Thus, we use the style transfer approach presented in Section 3 for low-level adaptation. During training, we sample at random a style image  $x_{\mathcal{T}}$  and use it to transfer each of the input images  $x_{\mathcal{S}}$  into the target style, resulting in  $x_{\mathcal{S} \rightarrow \mathcal{T}}$ . Image  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  shares the ground-truth labels with  $x_{\mathcal{S}}$ , which we use to compute loss  $\mathcal{L}_{SSD}$  for both  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  and  $x_{\mathcal{S}}$ . Low-level adaptation was done in [18] with CycleGAN for obtaining one transferred image per  $x_{\mathcal{S}}$ , but with our method we are capable of generating several instances  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  out of the same  $x_{\mathcal{S}}$ .

**Enforcing feature consistency.** Our motivation to train the detector with  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  instances is to force the network to recognize the objects under various low-level transformations. However, no constraint is introduced on the feature level. Compared to  $x_{\mathcal{S}}$ , the generated  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  contain low-level variations, *e.g.* color, texture or edges, but the location and shape of the objects is mostly preserved. Therefore, we assume that high-level (*i.e.* semantic) features in the detector should remain invariant to the style change of the image. Following that assumption, we introduce a constraint to reduce the distance between high-level features of  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  and  $x_{\mathcal{S}}$  aiming to increase the robustness of high-level features to low-level changes. Hence, the feature consistency loss is

$$\mathcal{L}_{Cons} = \sum_{k=1}^K \|f_k(x_{\mathcal{S}}) - f_k(x_{\mathcal{S} \rightarrow \mathcal{T}})\|^2 \quad (4)$$

where  $f_k(x)$  refers to one the last  $K = 6$  feature maps of SSD300, where the object detections are generated. Figure 2 shows how  $\mathcal{L}_{Cons}$  is applied in the SSD300 detector.

## 4.2 Robust Pseudo Labelling

The approach presented in Section 4.1 makes the high-level features extracted by the detector robust to low-level changes, *i.e.* the model extracts similar high-level features for both  $x_{\mathcal{S}}$  and  $x_{\mathcal{S} \rightarrow \mathcal{T}}$ . However, the style transfer method does not incorporate any high-level semantic information in the image transfer process, meaning that there will still be a domain gap between the high-level features extracted by the detector from the images coming from  $\mathcal{S}$  and the images coming from  $\mathcal{T}$ . We aim to adapt the classifiers to those high-level differences via a pseudo labelling approach. A possible domain adaptation approach is to automatically label the target data using the source trained model  $h_{\mathcal{S}}(x_{\mathcal{T}})$  and then obtain an improved model  $h_{\mathcal{S}+\mathcal{T}}$  by training with such pseudo labelled examples (refer to [21] for an in-depth explanation of why pseudo labels work). As past works [9, 19], we use the confidence of  $h_{\mathcal{S}}(x_{\mathcal{T}})$  to label  $\mathcal{T}$ . However, in an object detection problem, we also need to be careful when selecting negative examples.

**Positive examples sampling** has been done in [18] by taking top-1 object per class in a target image in a weakly supervised setting, discarding other positive instances. We argue that the actual value of the label score  $s_c = h(b)$  for box  $b$  and class  $c$  output by the detector can be used to filter out the less confident detections. We therefore propose to accept a pseudo labelled example if the probability of  $c$  being correct is  $P(c|h(b)) \geq \varepsilon_+$ . The parameter  $\varepsilon_+$  defines the precision of the sampling process, which we expect to remain high in different domains for a given high  $h(b)$ , while the recall may drop significantly. We test this assumption by evaluating a detector trained on source VOC and tested in *DTI* VOC domain,

which present strong low-level feature changes. We report the results in the supplementary material. The precision curve is high for  $\mathcal{S}$  and  $DTI$  images *e.g.* the detector confidence of 0.8 corresponds to  $P(c|h(b) \geq 0.8) \geq 0.8$  in both domains, while recall drops from 0.6 to 0.16. Experimentally, we directly set the threshold  $h(b) \geq \hat{\varepsilon}_c$  that is then used to include the predictions as pseudo labelled examples. The pseudo labelling process is only run once.

**Negative examples sampling** is equally important. Hard-negative mining during training of SSD in [25] is performed by choosing for a given instance the top-k boxes with the highest confidence  $h(b_k)$  that were detected in the background regions, which gives good results when all ground-truth annotations are available. However, in the unsupervised setting with positive target examples pseudo labelled by  $P(c|h(b)) \geq \varepsilon_+$  at high precision point, many false negative examples will have  $P(c|h(b))$  just below  $\varepsilon_+$  therefore including the rejected top detections as negatives may result in positive examples within the negative set. Hence, we set another threshold  $h(b) \geq \hat{\varepsilon}_-$  for selecting the negative samples. We maintain the ratio of 3:1 negative to positive number of examples as in [25] and negative sampling is performed for each instance during training.

**Combining Pseudo Labels and Style Transfer.** An image from the target domain,  $x_{\mathcal{T}}$ , can be transformed to have the style of another target image  $x_{\mathcal{T}}$  using the style transfer model, generating a new instance  $x_{\mathcal{T} \rightarrow \mathcal{T}}$ , with the same but unknown labels as  $x_{\mathcal{T}}$ . The pseudo labels from  $x_{\mathcal{T}}$  can be used with  $x_{\mathcal{T} \rightarrow \mathcal{T}}$  for training the detector. Thus, the model detects high-level concepts in the target domain under different low-level variations, generated from individual image styles present in the target domain. This approach is related to the consistency constraints applied in [82], however, instead of perturbations such as color jitter or adding noise, we apply a transformation in the target example to adapt it to the style of another target example. Our loss for Robust Pseudo Labelling  $\mathcal{L}_{RPL}$  has the same expression as  $\mathcal{L}_{SSD}$  loss (see Equation 2) but applied to  $x_{\mathcal{T}}$  and  $x_{\mathcal{T} \rightarrow \mathcal{T}}$  instances where the pseudo labels are obtained with the strategy presented above.

### 4.3 Combined Domain Adaptation Loss

Our final loss combines all the previously discussed terms:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{ST} + \lambda_2 \mathcal{L}_{Cons} + \lambda_3 \mathcal{L}_{\mathcal{S}} + \lambda_4 \mathcal{L}_{RPL} \quad (5)$$

In summary, our approach applies the style transfer loss  $\mathcal{L}_{ST}$  (Equation 1) enforcing consistency of VGG19 feature maps between  $x_{\mathcal{S}}$  and  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  for the content loss. The consistency loss  $\mathcal{L}_{Cons}$  (Equation 4) performs similar operation on high-level SSD feature maps also between  $x_{\mathcal{S}}$  and  $x_{\mathcal{S} \rightarrow \mathcal{T}}$ .  $\mathcal{L}_{\mathcal{S}}$  refers to detector loss  $\mathcal{L}_{SSD}$  (Equation 2) for both  $x_{\mathcal{S}}$  and  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  images. Finally,  $\mathcal{L}_{RPL}$  is also  $\mathcal{L}_{SSD}$  but for  $x_{\mathcal{T}}$  and  $x_{\mathcal{T} \rightarrow \mathcal{T}}$  images based on our robust pseudo labelling approach (Section 4.2).  $\mathcal{L}_{ST}$  and  $\mathcal{L}_{Cons}$  are also computed for  $x_{\mathcal{T}}$  and  $x_{\mathcal{T} \rightarrow \mathcal{T}}$  in the second step of training, *i.e.* after pseudo labelling.  $\lambda_1, \lambda_2, \lambda_3$  and  $\lambda_4$  are hyperparameters.

## 5 Experiments

**Architecture.** We use the ChainerCV [27] implementation of SSD300, which employs a VGG16 as feature extractor and is pretrained on the *trainval* subset of VOC2007 and VOC2012. The input images are resized to 300x300 and during inference we discard regions with a confidence lower than 0.01. An intersection over union (IoU) threshold of 0.45

		Pascal VOC→Clipart1k																				
Model		aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	prsn	plnt	sheep	sofa	train	tv	mAP
<i>Fully supervised</i>																						
$\mathcal{T}$ [18]		50.5	60.3	40.1	55.9	34.8	79.7	61.9	13.5	56.2	76.1	57.7	36.8	63.5	92.3	76.2	49.8	40.2	28.1	60.3	74.4	55.4
<i>Weakly supervised</i>																						
DTI+PL [18]		35.7	61.9	26.2	45.9	29.9	74.0	48.7	2.8	53.0	72.7	50.2	19.3	40.9	83.3	62.4	42.4	22.8	38.5	49.3	59.5	46.0
<i>Unsupervised</i>																						
$\mathcal{S}$ [18]		19.8	49.5	20.1	23.0	11.3	38.6	34.2	2.5	39.1	21.6	27.3	10.8	32.5	54.1	45.3	31.2	19.0	19.5	19.1	17.9	26.8
ADDA [18]		20.1	50.2	20.5	23.6	11.4	40.5	34.9	2.3	39.7	22.3	27.1	10.4	31.7	53.6	46.6	32.1	18.0	21.1	23.6	18.3	27.4
DTI [18]		23.3	<b>60.1</b>	24.9	41.5	26.4	53.0	44.0	4.1	45.3	51.5	39.5	11.6	40.4	62.2	61.1	37.1	20.9	39.6	38.4	36.0	38.0
DA-Faster* [0]		15.0	34.6	12.4	11.9	19.8	21.1	23.2	3.1	22.1	26.3	10.6	10.0	19.6	39.4	34.6	29.3	1.0	17.1	19.7	24.8	19.8
Strong-Weak* [18]		26.2	48.5	<b>32.6</b>	<b>33.7</b>	<b>38.5</b>	54.3	37.1	<b>18.6</b>	34.8	<b>58.3</b>	17.0	12.5	33.8	65.5	61.6	<b>52.0</b>	9.3	24.9	<b>54.1</b>	49.1	38.1
<i>Ours</i>																						
ST		32.0	56.9	24.9	36.0	29.3	51.2	46.7	4.0	47.8	55.7	38.9	16.1	43.7	79.2	64.2	39.2	<b>21.8</b>	43.3	38.6	47.9	40.9
ST+C		35.0	57.3	24.7	41.9	28.0	56.8	49.1	9.9	49.3	55.6	44.0	<b>16.5</b>	42.3	83.1	65.0	42.8	17.7	43.9	42.0	52.6	42.9
ST+C+PL <sup>†</sup> [18]		34.8	51.9	23.2	43.8	29.8	58.6	48.2	6.9	46.1	56.2	40.6	20.0	37.6	84.2	60.3	40.2	19.9	37.8	34.6	48.8	41.2
ST+C+PL [18]		32.4	50.2	22.9	43.2	28.6	60.2	46.3	5.6	46.7	57.1	40.9	21.4	35.7	87.9	63.3	41.4	19.4	37.8	35.2	54.3	41.5
ST+C+RPL <sup>†</sup>		35.5	58.8	24.9	41.9	25.3	59.1	48.4	10.2	49.4	55.7	41.5	15.9	42.4	<b>84.3</b>	66.1	42.2	18.7	43.8	44.0	52.6	43.0
ST+C+RPL		<b>36.9</b>	55.1	26.4	<b>42.7</b>	23.6	<b>64.4</b>	<b>52.1</b>	10.1	<b>50.9</b>	57.2	<b>48.2</b>	16.2	<b>45.9</b>	83.7	<b>69.5</b>	41.5	21.6	<b>46.1</b>	48.3	<b>55.7</b>	<b>44.8</b>

Table 1: Object detection results in Average Precision (%) in the domain Clipart1k. We report results for different combinations of the loss proposed in Section 4, and we use  $\mathcal{S}$  for source domain data,  $\mathcal{T}$  for target,  $DTI$  for domain transferred images,  $ST$  for pixel level adaptation using style transfer,  $C$  for consistency loss,  $PL$  for the pseudo labelling method from [18] and  $RPL$  for robust pseudo labelling (<sup>†</sup> means no style transfer applied to  $\mathcal{T}$  examples). Methods marked with \* use a Faster R-CNN with images of 600 pixels on the shorter side. Our approach brings significant improvements over baseline and other state-of-the-art methods from [0, 18, 36, 42].

is used for Non-Maximum Suppression as in [25]. For the Style Transfer model we follow the implementation of [17], where the encoder is a VGG19 pretrained on ImageNet and frozen during training. The decoder mirrors the encoder and is pretrained using WikiArt [26] as style images and MS COCO [23] as content images.

**Training details.** We finetune the detector using SGD with a learning rate of  $10^{-5}$  and a momentum of 0.9 as in [18]. The training is done in two steps. In the first step, we set  $\lambda_1 = 1.0$ ,  $\lambda_3 = 1.0$  and  $\lambda_4 = 0$ , and  $\lambda_2 = 1.0$  or  $\lambda_2 = 0.0$  depending if feature consistency is enforced. Before the second step, we pseudo label the target set with the updated object detector, and train again setting  $\lambda_4 = 1.0$ . We only pseudo label  $\mathcal{T}$  once, using a threshold of  $\hat{s}_c = 0.7$  (see Section 4.2). For negative sampling, we set  $\hat{s}_- = 0.9$ . With the higher  $\hat{s}_-$  compared to  $\hat{s}_c$  we aim to avoid including many false negatives as, contrary to positive sampling, negative sampling is performed during training by selecting the hardest negative boxes satisfying  $h(b) \geq \hat{s}_-$ . However, the method is stable to lower values of  $\hat{s}_-$  as discussed later in this section. In the first step, the batch only contains  $\mathcal{S}$  instances, and in the second step half of the examples come from  $\mathcal{S}$  and half from  $\mathcal{T}$ . We use examples coming from  $\mathcal{S}$  also in the second step to mitigate the effects of the incorrect pseudo-labels present in  $\mathcal{T}$ . The batch size is 6 and we finetune the detector for 10000 iterations in the first step and for 5000 iterations in the second step. Due to the small number of test images, and few instances for some classes, the obtained mean Average Precision (mAP) has a high variance. Thus, we perform mAP evaluation in the test set every 100 iterations and we report the average mAP of the last 10 evaluations.

**Comparison.** Our main experiments are performed in the cross-domain dataset [18] mentioned in Section 3. We report the per-class average precision and mAP using the evaluation

		Pascal VOC→Watercolor2k						Pascal VOC→Comic2k						Sim10k→Cityscapes		Cityscapes→Foggy Cityscapes																				
Model		bike	bird	car	cat	dog	prsn	mAP	bike	bird	car	cat	dog	prsn	mAP	Model		car	bike	bus	car	mbike	prsn	rider	train	truck	mAP									
<i>Fully sup.</i>																																				
T [18]		76.0	60.0	52.7	41.0	43.8	77.3	58.4	55.9	26.8	40.4	42.3	43.0	70.1	46.4	Fully sup.		T	60.1	35.4	53.3	56.9	31.2	28.6	35.4	43.6	34.0	39.8								
<i>Weakly sup.</i>																																				
DTI+PL [18]		76.5	54.9	46.0	37.4	38.5	72.3	54.3	55.2	18.5	38.2	22.9	34.1	54.5	37.2	Unsup.		S	43.7	29.1	28.5	38.8	23.8	22.1	25.7	18.8	17.6	25.6								
Unsup.																S*		[18]	34.6	26.5	22.3	34.3	15.3	24.1	33.1	3.0	4.1	20.3								
DA-Faster*		[18]	75.2	40.6	48.0	31.5	20.6	60.0	46.0															DA-Faster*		[18]	38.9	31.0	25.0	40.5	22.1	35.3	20.2	20.0	27.1	27.6
Str-weak*		[18]	82.3	55.9	46.5	32.7	35.5	66.7	53.3	29.6	20.1	26.3	20.3	22.3	51.6	28.4	Str-weak*		[18]	42.3	<b>35.3</b>	<b>36.2</b>	<b>43.5</b>	<b>30.0</b>	<b>29.9</b>	<b>42.3</b>	<b>32.6</b>	<b>24.5</b>	<b>34.3</b>							
<i>Ours</i>																																				
ST		78.2	55.7	47.0	41.2	33.0	67.0	53.7	49.5	16.4	37.6	20.8	30.9	57.1	35.4	ST		43.7	29.7	32.6	40.8	26.1	24.1	30.3	20.9	23.4	28.5									
ST+C		81.4	54.3	47.5	40.5	35.7	68.3	54.6	51.4	17.3	39.9	21.4	<b>31.9</b>	56.1	36.3	ST+C		44.0	30.1	32.2	39.6	25.6	24.1	29.8	23.0	23.0	28.4									
ST+RPL		81.4	56.1	48.2	41.0	40.4	72.0	56.5	55.0	18.4	40.9	22.7	31.5	59.5	38.0	ST+RPL		44.1	29.6	32.3	42.1	26.9	24.4	30.2	24.3	22.4	29.0									
ST+C+RPL <sup>†</sup>		81.3	56.1	48.2	41.0	40.4	72.0	56.5	55.0	18.4	40.9	22.7	31.5	59.5	38.0	ST+C+RPL <sup>†</sup>		42.1	28.4	31.3	39.0	23.4	22.0	24.7	21.1	22.0	26.5									
ST+C+RPL		79.9	<b>56.5</b>	<b>48.6</b>	<b>42.1</b>	<b>42.9</b>	<b>73.7</b>	<b>57.3</b>	<b>55.9</b>	19.7	<b>42.3</b>	<b>23.6</b>	31.5	<b>63.4</b>	<b>39.4</b>	ST+C+RPL		<b>44.2</b>	30.3	35.4	41.5	26.9	24.2	29.2	26.7	23.1	29.7									

(a) Results in Watercolor2k and Comic2k

(b) Results in driving datasets

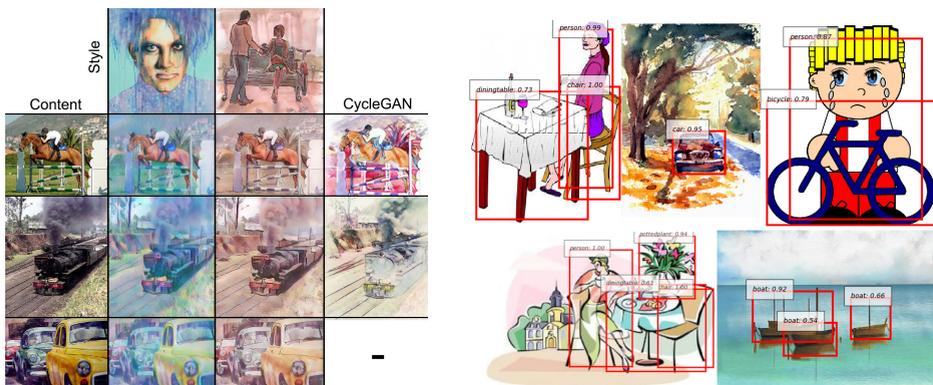
Table 2: Average Precision (%) in Watercolor2k and Comic2k, and in driving datasets.

code in ChainerCV as [18]. We include the results from [18] for the weakly supervised method, referred to as *DTI+PL*, and the ablation study of using only the *DTI*. For Clipart1k and Watercolor2k, we include the results from [36] for their method, *Strong-Weak*, and for [18], *DA-Faster*. Both [18] and [36] use Faster R-CNN with images of 600 pixels on the shorter side. For Comic2k we run the official code of [36] with global and local adaptation and context regularization. We also include the results for *ADDA* [12] in Clipart1k reported in [18].

## 5.1 Domain Adaptation for Object Detection Results

Table 1 shows the results obtained in Clipart1k and Table 2a in Watercolor2k and Comic2k. Compared to *DTI*, the performance increases in all three datasets by using the style transfer method, especially in Comic2k with an increase of 5.6%. The improved performance shows that generating multiple  $x_{S \rightarrow \mathcal{T}}$  from a single  $x_S$  is beneficial in datasets with various styles. Applying the consistency loss proves to be favourable in the three datasets, notably in Clipart1k we achieve an increase of 2.0%. We explored the impact of sampling different styles by restricting the style to a randomly chosen single image from  $\mathcal{T}$  in Clipart1k (with consistency loss) achieving 36.3%, that is 6.3% lower compared to using all of the styles. In the three datasets, using  $x_{\mathcal{T} \rightarrow \mathcal{T}}$  instances along with their pseudo labels leads to an improved performance compared to using only  $x_{\mathcal{T}}$  instances, referred to as *ST+C+RPL*<sup>†</sup>. Even though the pseudo labelling method proposed in [18] was developed for the case where image-level labels were available, we also tested its performance for the unsupervised case, *i.e.* taking the top-1 in confidence for each of the classes and not using any threshold for the negative sampling. We show in Table 1 that using the pseudo labelling method from [18] for our second step results in lower performance compared to not using any pseudo labels. The drop in performance is due to three reasons. First, by only pseudo labelling the top-1 prediction per class, a high number of predictions with high-confidence is discarded. Second, the method from [18] does not impose any constraint in the negative sampling process, resulting in selecting a high amount of false negatives during training. Third, when no image-level labels are given, the method from [18] assumes that every class is present in the image, which increases the number of false positives.

Using all of our proposed modifications we obtain a model with a higher mAP performance than the state-of-the-art methods on the three datasets: by 6.7% in Clipart1k, 4.0% in Watercolor2k and 9.6% in Comic2k. We also improve upon the weakly supervised method



(a) Examples of style transferred images

(b) Examples of detections in the test set of  $\mathcal{T}$ 

Figure 3: Qualitative results. Left: Examples of style transferred images along with their corresponding CycleGAN generated image in  $[\mathcal{X}]$ . Right: Examples of object detections in the target domain.

in  $[\mathcal{X}]$  for both Watercolor2k and Comic2k. In particular, in Watercolor2k the performance is only 1.1% lower compared to the fully supervised approach.

**Qualitative results.** In Figure 3 we show some  $x_{\mathcal{S} \rightarrow \mathcal{T}}$  samples, along with the corresponding image generated by CycleGAN. The last row shows an example of  $x_{\mathcal{T} \rightarrow \mathcal{T}}$ . Figure 3 also shows examples of detections in  $\mathcal{T}$ . The model detects objects with similar shapes but different low-level statistics than the  $\mathcal{S}$  domain, *e.g.* the car in the upper row, objects with shape changes, *e.g.* the person in the bottom-left image, or objects with a simplified shape, *e.g.* the bicycle in the upper-right image.

**Hyperparameter sensitivity.** Figure 4 shows the effect of varying  $\hat{s}_c$  and  $\hat{s}_-$  thresholds for pseudo labelling (see Section 4.2) in Clipart1k. The performance is stable for  $\hat{s}_c \geq 0.2$  and low values of  $\hat{s}_-$  give slightly lower performance. Additionally, for the standard hard-negative mining case, *i.e.*  $\hat{s}_- = 0$ , we obtained 44.0%, that is 0.8% lower compared to  $\hat{s}_- = 0.9$ , showing the importance of carefully selecting the negative samples. Figure 4 also shows the effect of varying  $\lambda_i$  in Clipart1k (see Section 4.3). For the first step of training, we vary the value of  $\lambda_1$  and  $\lambda_2$ , which are the weights for  $\mathcal{L}_{ST}$  and  $\mathcal{L}_{Cons}$ . We find that the performance increases when the style loss is given more weight. As for  $\mathcal{L}_{Cons}$ ,  $\lambda_2 = 1.0$  achieves the best result and  $\lambda_2 = 10.0$  is the only value that drops the performance compared to  $\lambda_2 = 0.0$  (see Table 1). For the second step, we vary the weight of  $\mathcal{L}_S$  and  $\mathcal{L}_{RPL}$  in the loss finding that the optimal setting is weighting them equally with  $\lambda_3 = 1.0$  and  $\lambda_4 = 1.0$ .

**Driving datasets.** We follow the same protocol in  $[\mathcal{B}2]$  to test our method in two experiments: adaptation from Sim10k  $[\mathcal{B}1]$  (simulated data) to Cityscapes  $[\mathcal{B}]$  (real data) and from Cityscapes to Foggy Cityscapes  $[\mathcal{B}7]$  (synthetic fog added to Cityscapes). A detailed overview of the datasets is included in the supplementary material. We employed the SSD512 model, which uses images of 512x512 pixels, due to small objects being present in the images. The batch size is reduced to 2 due to memory constraints, and for Cityscapes  $\rightarrow$  Foggy Cityscapes we doubled the number of iterations (20000 and 10000 iterations for first and second step). Table 2b shows the results. We also report the results given in  $[\mathcal{B}6]$  when using only source data for the Faster R-CNN method due to the difference with SSD512. The improvement is limited for Sim10k  $\rightarrow$  Cityscapes, whereas in Cityscapes  $\rightarrow$

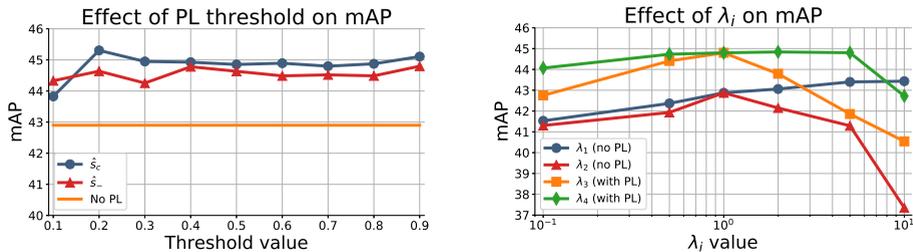


Figure 4: Left: Effect of the threshold used for pseudo labelling in Clipart1k. Right: Effect of varying  $\lambda_i$  values in Clipart1k.

Foggy Cityscapes we see an improvement of 4.1% over the source domain trained model. Style transfer adaptation is greatly beneficial for Cityscapes  $\rightarrow$  Foggy Cityscapes as Foggy Cityscapes only differs from Cityscapes in the low-level statistics, *i.e.* added fog. Consistency only helps when using  $\mathcal{T}$  instances, where it increases the performance by 0.7% compared to not applying it. In these datasets, the improvements are limited for several reasons. First, the target datasets do not have a large variation in style in the instances, which also limits the benefits of forcing feature consistency among the different styles. Second, the low-level feature mismatch is lower compared to Pascal VOC and the cross-domain dataset. Lastly, the style transfer method is not suited to generate realistic images and the output images contain distortions (examples of style-transferred images included in the supplementary material). Using methods adapted for realistic images such as [22, 47] may improve the performance. In both experiments, the best result is obtained when all of the proposed blocks are used. Notably, using pseudo labels without  $x_{\mathcal{T} \rightarrow \mathcal{T}}$  instances makes the performance drop in both experiments, showing that forcing consistency in the detections between  $x_{\mathcal{T}}$  and  $x_{\mathcal{T} \rightarrow \mathcal{T}}$  instances is also beneficial in this setting.

## 6 Conclusion

We presented an approach for domain adaptation of an object detector based on low-level adaptation via style transfer and high-level adaptation via robust pseudo labelling. We demonstrated that each of these contributions brings an improvement with the highest performance obtained when all are combined. Our approach gives the highest mean Average Precision in the cross-domain dataset, where we proved the benefits of generating several translated source images to any of the target styles. We also showed that using style transfer with pseudo labelled instances is also key to improve the results. The method was also tested in driving datasets using the same parameters, showing an improvement over the source trained model.

**Acknowledgment.** This research was supported by FACER2VM EPSRC EP/N007743/1.

## References

- [1] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2846–2854, 2016.

- [2] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3339–3348, 2018.
- [3] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Weakly supervised object localization with multi-fold multiple instance learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 39(1):189–203, 2017.
- [4] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3213–3223, 2016.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [6] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *International Conference on Learning Representations (ICLR)*, 2, 2017.
- [7] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.
- [8] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning (ICML)*, pages 1180–1189, 2015.
- [9] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2414–2423, 2016.
- [10] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations (ICLR)*, 2019. URL <https://openreview.net/forum?id=Bygh9j09KX>.
- [11] Ross Girshick. Fast r-cnn. In *IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [12] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 580–587, 2014.
- [13] Hironori Hattori, Vishnu Naresh Boddeti, Kris M Kitani, and Takeo Kanade. Learning scene-specific pedestrian detectors without real data. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3819–3827, 2015.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.

- [15] Judy Hoffman, Sergio Guadarrama, Eric S Tzeng, Ronghang Hu, Jeff Donahue, Ross Girshick, Trevor Darrell, and Kate Saenko. Lsda: Large scale detection through adaptation. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3536–3544, 2014.
- [16] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning (ICML)*, 2018.
- [17] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1501–1510, 2017.
- [18] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5001–5009, 2018.
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 694–711. Springer, 2016.
- [20] M. 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? In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8, 2017.
- [21] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*, volume 3, page 2, 2013.
- [22] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *European Conference on Computer Vision (ECCV)*, pages 453–468, 2018.
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, pages 740–755. Springer, 2014.
- [24] Tsung-Yi Lin, Priyank Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2018.
- [25] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European Conference on Computer Vision (ECCV)*, pages 21–37. Springer, 2016.
- [26] Kiri Nichol. Painter by numbers, 2016. URL <https://www.kaggle.com/c/painter-by-numbers>.

- [27] Yusuke Niitani, Toru Ogawa, Shunta Saito, and Masaki Saito. Chainercv: a library for deep learning in computer vision. In *2017 ACM on Multimedia Conference*, pages 1217–1220. ACM, 2017.
- [28] Xingchao Peng, Baochen Sun, Karim Ali, and Kate Saenko. Learning deep object detectors from 3d models. In *IEEE International Conference on Computer Vision (CVPR)*, pages 1278–1286, 2015.
- [29] Alexandre Rame, Emilien Garreau, Hedi Ben-Younes, and Charles Ollion. Omnia faster r-cnn: Detection in the wild through dataset merging and soft distillation. *arXiv preprint arXiv:1812.02611*, 2018.
- [30] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, 2016.
- [31] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 91–99, 2015.
- [32] Subhankar Roy, Aliaksandr Siarohin, Enver Sangineto, Samuel Rota Buló, Nicu Sebe, and Elisa Ricci. Unsupervised domain adaptation using feature-whitening and consensus loss. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [33] Subhankar Roy, Aliaksandr Siarohin, Enver Sangineto, Samuel Rota Buló, Nicu Sebe, and Elisa Ricci. Unsupervised domain adaptation using feature-whitening and consensus loss. *arXiv preprint arXiv:1903.03215*, 2019.
- [34] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *International Conference on Machine Learning (ICML)*, pages 2988–2997, 2017.
- [35] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [36] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, page To appear, 2019.
- [37] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision (IJCV)*, 126(9):973–992, Sep 2018. URL <https://doi.org/10.1007/s11263-018-1072-8>.
- [38] Rui Shu, Hung Bui, Hirokazu Narui, and Stefano Ermon. A DIRT-t approach to unsupervised domain adaptation. In *International Conference on Learning Representations (ICLR)*, 2018. URL <https://openreview.net/forum?id=H1q-TM-AW>.
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

- [40] Yuxing Tang, Josiah Wang, Boyang Gao, Emmanuel Dellandréa, Robert Gaizauskas, and Liming Chen. Large scale semi-supervised object detection using visual and semantic knowledge transfer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2119–2128, 2016.
- [41] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *IEEE International Conference on Computer Vision (ICCV)*, pages 4068–4076, 2015.
- [42] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2962–2971, 2017.
- [43] Eric Tzeng, Kaylee Burns, Kate Saenko, and Trevor Darrell. Splat: Semantic pixel-level adaptation transforms for detection. *arXiv preprint arXiv:1812.00929*, 2018.
- [44] Jasper Uijlings, Stefan Popov, and Vittorio Ferrari. Revisiting knowledge transfer for training object class detectors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1101–1110, 2018.
- [45] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *International Conference on Machine Learning (ICML)*, volume 1, page 4, 2016.
- [46] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S Davis. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *European Conference on Computer Vision (ECCV)*, pages 518–534, 2018.
- [47] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. *arXiv preprint arXiv:1903.09760*, 2019.
- [48] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [49] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *European Conference on Computer Vision (ECCV)*, pages 289–305, 2018.