

# Zero-Annotation Object Detection with Web Knowledge Transfer

Qingyi Tao<sup>1,2</sup>(✉), Hao Yang<sup>3\*</sup>, and Jianfei Cai<sup>1</sup>

<sup>1</sup> Nanyang Technological University, Singapore  
qtao002@e.ntu.edu.sg, asjfcai@ntu.edu.sg

<sup>2</sup> NVIDIA AI Technology Center

<sup>3</sup> Amazon Rekognition lance1ot365@gmail.com

**Abstract.** Object detection is one of the major problems in computer vision, and has been extensively studied. Most of the existing detection works rely on labor-intensive supervision, such as ground truth bounding boxes of objects or at least image-level annotations. On the contrary, we propose an object detection method that does not require any form of human annotation on target tasks, by exploiting freely available web images. In order to facilitate effective knowledge transfer from web images, we introduce a multi-instance multi-label domain adaption learning framework with two key innovations. First of all, we propose an instance-level adversarial domain adaptation network with attention on foreground objects to transfer the object appearances from web domain to target domain. Second, to preserve the class-specific semantic structure of transferred object features, we propose a simultaneous transfer mechanism to transfer the supervision across domains through pseudo strong label generation. With our end-to-end framework that simultaneously learns a weakly supervised detector and transfers knowledge across domains, we achieved significant improvements over baseline methods on the benchmark datasets.

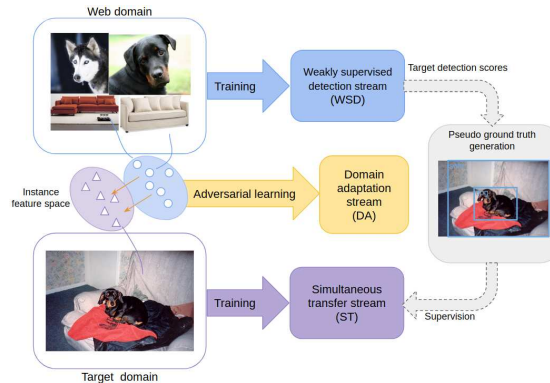
**Keywords:** Object detection, domain adaptation, web knowledge transfer.

## 1 Introduction

In recent years, with the advances of deep convolutional neural networks (DCNN), object detection tasks have attracted significant attention and have achieved great improvements in performance and efficiency. State-of-the-art works such as Faster R-CNN [25], SSD [21], FPN [20] achieve high accuracy but require labour-intensive bounding box annotations for training. To alleviate the large labour cost for annotating ground truth bounding boxes, weakly supervised object detection methods that only rely on image-level human annotations have also been extensively studied [2–4, 8, 14, 15, 27, 16]. However, for large-scale multi-object detection problem, even annotating just image-level labels could deem to be too

---

\* This work was done when Hao Yang was at NTU, Singapore



**Fig. 1.** Overall idea of object detection without human annotations. First of all, we mine freely available web images through automatic retrieval with respect to a given set of object categories. Our framework then facilitates knowledge transfer from these web images to the target task using a multi-stream network with three major components: 1) a weakly supervised detection stream (WSD) to train the detection model from web images; 2) an instance-level domain adaptation (DA) stream to minimize the feature discrepancy across domains at instance-level feature space; 3) a simultaneous transfer (ST) stream that learns to discriminate unsupervised target examples by transferring supervision from web detection model. These three streams are trained simultaneously to effectively transfer the learning of web images to the target task.

expensive. This motivates us to develop an object detection method with no human annotations involved. Our basic idea is to transfer knowledge from free web resources to the target tasks.

With the similar motivation, zero-shot learning (ZSL) problem has been proposed for unsupervised learning. Many works [17, 18, 10, 24, 11, 1] have been proposed to utilize side information such as attributes, Wikipedia or WordNet to jointly encode semantic space and image feature space for solving zero-shot recognition problems. However, although textual side information could help zero-shot object recognition with exploiting the intrinsic semantic relations between categories, it is hard to learn a class-specific object detector that can accurately differentiate objects from the background as well as different objects with just semantic descriptions. In contrast, our direction is to exploit freely available web images as a much stronger side information to solve the object detection problem without human annotations, considering that there are huge amount of image resources from the web and plenty of works studying the automatic collection of these web imagery resources [19, 7, 34, 28].

One baseline approach for learning detectors with web images is to simply use the web images and their image “labels” (essentially the pre-defined labels used as search phrases to retrieve the images) to train a web object detector using some weakly supervised detection (WSD) methods and apply them on target images. This naive learning scheme is referred as webly supervised learning in

previous works [9, 6]. However, directly applying the web models to the target data produces poor results. The major reason is that it ignores the domain discrepancies between web images and target images. As shown in Fig. 1, web images from image search engines are mostly studio-shot images, which are simple, clear and unblocked. In contrast, the target images (e.g. Pascal VOC images) usually contain multiple objects of different classes that are often occluded with cluttered scenes. Hence it is necessary to properly transfer the models learned from web images to the target images.

To address this domain discrepancy problem, we need to adapt the source (i.e. web domain) and target domain object appearances, for which unsupervised domain adaptation is the common way [29, 12, 22, 30, 5]. Although many unsupervised domain adaptation methods have been proposed, they all focus on image-level domain adaptation for image classification problems. What we consider here is the domain adaptation at instance level (i.e. object proposal level), which is non-trivial to solve. Inspired by the recent adversarial domain adaptation works [29, 12, 30], we propose an instance-level adversarial domain adaptation network to reduce the domain discrepancies particularly at instance level. Our adversarial domain adaptation network includes a domain discriminator that differentiates object features from web domain and target domain, and a feature generator that projects source and target objects to the same manifold in the feature space so that the discriminator can no longer tell their differences.

In addition, we introduce an innovative component in our domain adaptation network: attention on foreground objects. As weakly supervised detection is essentially a multi-instance multi-label learning problem, each image actually is a bag of instances, where each instance corresponds to a bounding box proposal. Equally treating all proposals in each image when training adversarial domain adaptation network will lead to sub-par results, as we care more about proposals containing objects than proposals that are largely background. Therefore, we introduce an attention mechanism to emphasize the transfer of object proposals and suppress the transfer of background proposals.

However, the introduced instance-level domain adaptation network brings in a side effect, i.e. the feature generator is likely to ignore the semantic structure of different object classes, since there is no class-specific constraint. As a result, it not only brings features from different domains together to the same manifold, but also mixes up the sub-manifolds from different classes. For example, the “cow” from web domain will be confused with the “sheep” from target domain through the domain adaptation. To address this issue, we further introduce simultaneous learning towards class-specific pseudo labels to preserve the semantic structure during the domain adaptation. This component compensates the side effect of the domain adaptation component so that the domain shift will be guided in a class-specific manner. In this way, our overall architecture including the web object detector, the domain adaptation component and the simultaneous transfer component significantly boosts up the object detection results on unsupervised target data.

We would like to highlight that the rationale of studying this problem lies in that such detector can be trained without any human labour and therefore the whole process could be fully automated. Different from fully supervised and weakly supervised object detection, our object detector allows the training of the detection models to be highly scalable in term of categories. For example, in the Pascal VOC dataset, if we want to add the object class “keyboard”, which exists in some of the images but is not annotated, we need to re-annotate all the images in the training data by providing respective labels at bounding box level (for supervised detector) or image level (for weakly supervised detector). Another example is that if we want to further break down the “bird” class into multiple classes such as “parrot”, “goose”, “hawk” and etc, we also need to revise the annotations for all images containing “bird” objects. In contrast, our solution can automatically search the web and progressively transfer the web knowledge to learn the detector without any human intervention or any modification in the target domain dataset. The training of such detector can be a completely self-taught process. Hence, we think this problem is highly meaningful and worth to be studied.

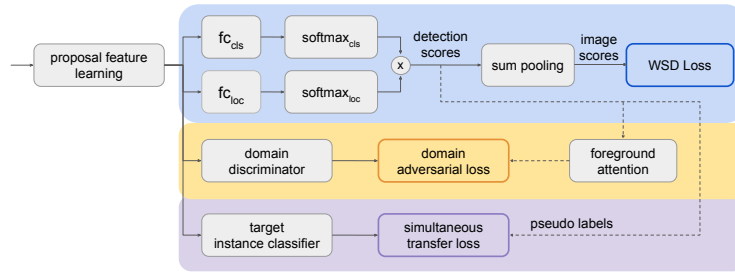
Overall, the main contributions of our work can be summarized as follows:

- We propose a new problem of knowledge transfer in object detection for *un-supervised* data, which enables learning an object detector from free web images and alleviates any forms of human annotations for target domain. By studying this problem, the learning of object detectors can be fully automated and highly scalable with categories.
- We propose an *instance-level* domain adaptation method to transfer web knowledge to unsupervised target dataset. The proposed domain adaptation framework includes: 1) an instance-level adversarial domain adaptation network with attention on foreground objects; 2) a simultaneous transfer stream to preserve the semantic structure of classes by transferring the pseudo labels obtained from the web domain detector to the target domain detector.
- Our method significantly reduces the gap between unsupervised object detection (i.e. train a detector using only web images and then directly apply it on target images) and the upper bound (i.e. train a detector using image-level labels of target data) by 3.6% in detection mAP.

## 2 Related Works

Our work is related to a few computer vision and machine learning areas. We will review these related topics in this section.

**Weakly supervised object detection:** Recent works on weakly supervised object detection aim to reduce the intensive human labour cost by using only image-level labels instead of bounding box annotations [2, 3, 8]. They are more cost-effective than the fully supervised object detection methods since image-level labels are easier to obtain compared with the bounding box annotations. These works formulate the weakly supervised object detection task as a multi-instance learning (MIL) problem in which the model will be learned alternatively



**Fig. 2.** The proposed network branches into three streams after the proposal feature learning layers. The first stream (in blue) is the weakly supervised detection (WSD) network which is further divided into recognition and localization streams. The middle stream (in yellow) is the instance-level domain adaptation (DA) stream that optimizes an adversarial loss to enforce domain invariant feature representation. The last stream (in purple) is the simultaneous transfer (ST) stream to preserve semantic structure of target data with pseudo labels.

to recognize the object categories and to find the object locations of each category. The recent work [4] is the first one introducing an end-to-end network with two separate branches for object recognition and localization respectively. Later, [15] introduced context information to the weakly supervised detection network in the localization stream. [27] proposed online classifier refinement to refine the instance classification based on image-level labels.

Our work is related to these works as the web data will be trained in a weakly supervised way with their weak labels. In this paper, we use WSDDN in [4] as the base model for our work.

**Learning from web data:** Web data is a free source of training samples that can be collected automatically for various tasks [6, 9, 35, 26]. Previous works [6, 9, 35] study the web data collection approaches and further evaluate their data collection methods by training those web data for different tasks. They focus on reducing the effects of noises from web images and thereby construct robust and clean web datasets. While learning for the target tasks, these prior works simply treat the web dataset as the substitute of the training dataset in the target task without considering the domain shifts between web data and target data, which is similar as our baseline approach. Apart from that, web data are often used as complementary data to improve the training of target dataset. In [33], web images are used to produce pseudo masks for pre-training the semantic segmentation network. In [32], an object interaction dataset with web images is created to facilitate the semantic segmentation task as additional data. In their approaches, the image-level labels (in [33]) or pixel-level ground truth masks (in [32]) of target images are required and web images are utilized as additional knowledge to improve the segmentation model performance. In our work, we attempt to solve the detection problem using only the web images without any forms of annotations from target dataset.

**Domain adaptation:** Our work is also closely related to the domain adaptation works [29, 12, 30, 5, 36]. [12] introduced the domain adversarial training of neural networks. The domain adaptation is achieved by introducing a domain classifier to classify features to their corresponding domains and applying a gradient reversal layer between the feature extractor and the domain classifier. With this reversal layer, when the domain classifier learns to distinguish the features from different domains, the feature extractor learns in the reverse way to make the feature distributions as indistinguishable as possible. Hence, this domain adversarial training can result in a domain-invariant feature representation. [29] also uses a similar method for domain transfer in image classification task. In [29], a domain classification loss and a domain confusion loss influence the training in an adversarial manner. They also added a soft label layer while learning the source examples in order to transfer correlations between classes to the target examples. Later, [30] proposed to untie the weight sharing between two domains. These previous works have validated the effectiveness of the adversarial domain adaptation methods in the image classification problem. In our work, we follow the principles of the end-to-end adversarial methods but for our zero-annotation detection task with the domain transfer of proposal-level features to reduce the domain mismatch between web data and target data.

### 3 Problem Definition and Notations

In this section, we formally define our problem of zero-annotation object detection with web knowledge transfer. Essentially, we define this problem as an unsupervised multi-instance multi-label domain adaptation problem. Specifically, we consider two domains, the web domain  $D^w$  representing web images and target domain  $D^t$  representing target tasks (e.g. Pascal VOC and MS COCO). The source data  $\{X_j^w, y_j\}_{j=1}^{n_w}$  is sampled from  $D^w$ , where  $X_j^w$  is the  $j$ -th image,  $y_j \in \mathbb{R}^C$  sampled from label space  $Y$  is the corresponding  $C$  dimensional binary label vector and  $n_w$  is the number of source images. For object detection problems, it is natural to decompose each image to a bag of instances, i.e., object proposals, through dense sampling or objectness techniques. Thus,  $X_j^w$  can be represented as  $X_j^w = \{x_i^{w,j}\}_{i=1}^{m_j^w}$ , where  $x_i^{w,j}$  is the  $i$ -th proposal in  $X_j^w$  and  $m_j^w$  is the number of total proposals of  $X_j^w$ . Similarly, the target data sampled from  $D^t$  can be denoted as  $\{X_j^t\}_{j=1}^{n_t}$ , and  $X_j^t = \{x_i^{t,j}\}_{i=1}^{m_j^t}$ . Note that since we do not have annotations for target data, effective knowledge transfer from the web domain is necessary.

Traditional domain adaptation methods usually optimize an objective function  $f : X^w, X^t \rightarrow Y$ , which jointly learns a classifier for source/web domain and transfers the knowledge to target domain at *image-level*. However, for object detection, we need to go deeper to *instance-level*. In particular, we need to learn  $f : x^w, x^t \rightarrow Y$ . Therefore, we will need a backbone structure to learn from image-level labels and propagate knowledge to instances, and an effective way to transfer knowledge from the web domain to the target domain at instance-level.

## 4 Methodology

Fig. 2 shows the diagram of the proposed framework for zero-annotation object detection with web knowledge transfer. The entire framework branches into three streams after feature representation, including WSD, DA, and ST. In the following, we describe each stream one by one.

### 4.1 Weakly supervised detection trained on web images

Our weakly supervised detection backbone is based on the basic WSDDN [4] (blue region in Fig. 2). Note that other end-to-end WSD methods can be easily applied as well. Specifically, for WSDDN, the proposal features  $x_i \in \mathbb{R}^d$  are obtained through an ROI pooling layer on the feature map of the image, followed by two fully connected layers, similar to Fast-RCNN [13]. Then we represent each image  $X$  as the concatenation of its proposal features, i.e.,  $X = \text{concat}(x_i), \forall i \in [1, m]$ , thus  $X \in \mathbb{R}^{m \times d}$ , where  $m$  denotes the number of proposals in the image. Note that here we abuse the notation  $x_i$  to represent both proposal and its corresponding feature, and  $X$  to represent both image and its corresponding concatenated feature matrix.

Following the proposal feature learning, the WSD network breaks into two branches of fully connected ( $fc$ ) layers to produce two score matrices  $S^{cls}$  and  $S^{loc} \in \mathbb{R}^{m \times C}$ , where  $C$  is the number of object classes. Then  $S^{cls}$  and  $S^{loc}$  are passed to two *softmax* layers with different axes, i.e.  $S^{cls}$  is normalized in the class dimension to produce the class probability of each proposal and  $S^{loc}$  is normalized in the proposal dimension to find the most responsive proposal for each class among all candidate proposals. For proposal  $i$  and class  $c$ , we respectively denote the outputs of these two *softmax* layers as  $p_{i,c}^{cls}$  and  $p_{i,c}^{loc}$ , which are defined as

$$p_{i,c}^{cls} = \frac{e^{s_{i,c}^{cls}}}{\sum_{k=1}^C e^{s_{i,k}^{cls}}}, \quad p_{i,c}^{loc} = \frac{e^{s_{i,c}^{loc}}}{\sum_{k=1}^m e^{s_{k,c}^{loc}}} \quad (1)$$

Then the detection probability  $p_{i,c}$  of each proposal can be computed by element-wise products of the normalized probabilities from the two branches:


$$p_{i,c} = p_{i,c}^{cls} \cdot p_{i,c}^{loc}. \quad (2)$$

The image classification probability  $p_c$  is calculated by summing up the detection probabilities of all proposals:

$$p_c = \sum_{i=1}^m p_{i,c}. \quad (3)$$

Finally, the multi-class cross entropy loss is adopted as the loss function of WSD, which is defined as

$$L_{WSD} = - \sum_{c=1}^C [y(c) \log(p_c) + (1 - y(c)) \log(1 - p_c)] \quad (4)$$



**Fig. 3.** Instance-level domain adaptation stream with foreground attention. We visualize the attended image regions produced by the foreground attention mechanism. The examples show that the foreground object regions are well attended and the background regions are suppressed during domain adversarial learning.

where  $y(c) \in \{0, 1\}$  is the web image label for class  $c$ .

Note that since we do not have any label in the target domain, this WSD loss is only optimized by training with web images.

## 4.2 Instance-level adversarial domain adaptation

The purpose of this instance-level domain adaptation (DA) stream (yellow region in Fig. 2) is to close the feature discrepancies between the two domains. Fig. 3 gives the detailed structure of this DA stream. In particular, it includes two players with adversarial goals: a discriminator trained to differentiate the domains where input features come from, and a feature learner shared with the WSD stream trained to align features from both domains so as to confuse the discriminator.

In particular, the proposed discriminator consists of a fully connected layer  $fc_d$  that classifies the input proposal features  $x_i$  in  $i$ -th row of  $X$  to their domains  $y_i^t \in \{0, 1\}$ . Here we define  $y_i^t = 0$  for  $x_i$  from the web domain  $D^w$  and  $y_i^t = 1$  for  $x_i$  from target domain  $D^t$ . Through a *softmax* operation, we can compute the domain probability as  $p_i^t$ , i.e.  $prob(y_i^t = 1)$ . The adversarial loss can then be written as

$$\begin{aligned}
 \min_{\phi_f^w} \max_{\phi_{fc_d}} & \mathbb{E}_{x \sim D^t} [\log(p^t)] + \mathbb{E}_{x \sim D^w} [\log(1 - p^t)], \\
 \mathbb{E}_{x \sim D^t} [\log(p^t)] &= \sum_i \mathbb{1}[y_i^t = 1] \log(p_i^t), \\
 \mathbb{E}_{x \sim D^w} [\log(1 - p^t)] &= \sum_i \mathbb{1}[y_i^t = 0] \log(1 - p_i^t),
 \end{aligned} \tag{5}$$

where  $\phi_f^w$  denotes the parameters of the feature learner,  $\phi_{fc_d}$  denotes the parameters of the discriminator  $fc_d$ , and  $\mathbb{1}[\cdot]$  is the indication function.

The optimization of the minimax domain adversarial loss in (5) is achieved by training alternatively between the following two steps. First, we update  $\phi_{fc_d}$  to distinguish proposal features from  $D^w$  and  $D^t$  to seek for maximizing the loss.






Fig. 4. Simultaneous transfer stream with pseudo ground truth generation.

Then we fix  $\phi_{f_{c_d}}$  and learn the feature representation  $\phi_f^w$  to minimize the loss so as to confuse the discriminator. In practice, we only shift the web domain  $D^w$  towards the target domain and  $\phi_f^w$  is updated by training only web images.

Moreover, unlike the existing domain adaptation works for image classification [29, 12], which focus on aligning image-level features, here we need to align instance-level features instead, especially for important instances that are more likely to contain objects. Specifically, while adapting the instance-level features, we care more about the foreground features than those background features in order to learn common object appearances. Therefore, we introduce an attention mechanism to focus on the adaptation of foreground features and suppress the effects for background features. As shown in Fig. 3, our foreground attention model uses the detection scores from the WSD stream and computes the foreground probability  $p_i^f$  for proposal  $i$  by summing up  $p_{i,c}$  in (2) over all the classes (i.e.  $\sum_{c=1}^C p_{i,c}$ ) followed by a *softmax* operation for the normalization over all the proposals. This is to find out the most responsive proposals regardless which object classes they belong to, and the responsive proposals with high  $p^f$  scores are highly likely to be foreground. Finally, we use the foreground probability as the attention weight, and modify the minimax adversarial loss as

$$\min_{\phi_f^w} \max_{\phi_{f_{c_d}}} \mathbb{E}_{x \sim D^t} [p^f \cdot \log(p^t)] + \mathbb{E}_{x \sim D^w} [p^f \cdot \log(1 - p^t)]. \quad (6)$$

### 4.3 Simultaneous transfer by pseudo supervision

Ideally, the domain adaptation stream should produce domain invariant features and improve the detection results while applying on the target dataset. However, it is observed that it fails to perform domain shift with class-specific directions. Specifically, it could encourage the features to be indistinguishable across not only domains but also classes. This ill effect of DA stream eventually makes features to be non-discriminative. Therefore, to preserve the semantic structure across different categories, we introduce the simultaneous transfer (ST) stream (purple part in Fig. 2) and use the pseudo labels generated from the WSD network as the supervision to preserve or even enhance the discriminative power of the learned features. The network details are shown in Fig 4.

To generate the pseudo ground truth for each target image, we use the detection scores  $p_{i,c}$  in (2) from the WSD stream. We select to highest scoring

proposal for each object class  $c$ , denoted as  $i_c = \operatorname{argmax}_i p_{i,c}$ . We set a threshold  $t$  to determine the presence of a class in an image. If  $p_{i_c,c} \geq t$ , the corresponding proposal  $i_c$  is selected as the pseudo ground truth bounding box. Given the pseudo ground truth boxes, we then sample the boxes with large overlaps with the pseudo ground truth boxes as positive examples and randomly sample a few background examples from the remaining bounding boxes.

Finally, we use the *softmaxloss* as the ST loss function:

$$L_{ST} = - \sum_{i \in P} \sum_{c=0}^C \mathbb{1}[y_i^{ST} = c] \log(p_{i,c}^{ST}), \quad (7)$$

where  $y^{ST} \in \{0, 1, 2, \dots, C\}$  are the class labels (0 is the class label of background),  $P$  is the set of the selected proposals, and  $p_{i,c}^{ST}$  is the class probability output from the fully connected layer  $fc_t$  followed by a *softmax* operation.

Conditional adversarial loss is also a common way in GAN to enable class-specific domain adaptation. However, here the conditions are instance-level pseudo labels, which are noisy labels. It will be more stable to detach the class conditional learning with the domain adversarial learning.

## 5 Experiments

In this section, we conduct various experiments to evaluate the effectiveness of our proposed zero-annotation object detection with web knowledge transfer.

### 5.1 Datasets and experiment setup

We evaluate our method on two object detection benchmark datasets: Pascal VOC 2007 and 2012. These two datasets contain images of 20 object classes. The web images we used are from the STC dataset [33], whose images can be freely obtained from Internet without human labour. Similar as most supervised detection works, mean average precision (mAP) is used as the evaluation metric. Following the common standard, the IoU threshold is set to be 0.5 between ground truths and correctly predicted boxes.

**Implementation details.** Our method is built upon two pre-trained networks on imagenet: VGG\_M and VGG\_16. We use selective search [31] to generate proposals for source and target images. In the WSD stream, we follow the details in the basic model of WSD as described in Section 4.1. The ROI features from the web domain are passed to the WSD stream to optimized the WSD loss whereas the ROI features from the target domain are only forwarded up to the detection score layer to generate foreground attention weights for the DA stream and pseudo labels for the ST stream. The DA stream takes the inputs from both source and target domains. It alternates between training the discriminator and the feature generator each time after training 5000 images. Lastly, the ST stream takes the inputs from the target domain and uses the detection scores generated from the WSD stream to generate pseudo ground truths as described in Section 4.3.

## 5.2 Baseline and upper bound

**Table 1.** Baseline(wt.web data) and upper-bound(wt.VOC labels) on VOC 2007.

Method	mAP
WSD(wt.web data)-VGG_M	21.5
WSD(wt.web data)-VGG16	21.8
WSD(wt.VOC labels)-VGG_M	30.2
WSD(wt.VOC labels)-VGG16	29.3

The baseline of our method is the basic WSD network [4] trained using only web images with web image labels. As shown in Table 1, due to the domain mismatch, the results are only 21.5 for VGG\_M and 21.8 for VGG16.

The upper bound of our method is to train the basic WSD network with VOC image-level labels similar as [4]. Our obtained upper bound result for VGG\_M is quite close to that reported in [4] with selective search proposals, while our result of 29.3 for VGG16 is higher than that of 24.3 reported in [4] with selective search proposals. Also, we have the same finding as [4] that VGG 16 performs slightly worse than VGG\_M. This could be because the image level labels might not give sufficient supervision for a very deep network for the MIL problem.

Overall, there are significant gaps between the results without VOC labels and those with VOC labels. We aim to reduce the gap between the unsupervised and weakly supervised detection by transferring the knowledge of web domain to target domain with our proposed method.

## 5.3 Detailed results and analysis

Table 2 shows the detailed detection results of different combinations of the three streams developed in our method on VOC2007 test set. All of these methods are evaluated against the baseline, ‘WSD(Baseline)’, that uses web images to train the WSD network alone. Before training the DA and ST streams, we train the WSD for one epoch first. This will give a more stable initialization to get the foreground attention weights for DA and pseudo labels for ST.

From Table 2, we can see that adding DA alone, ‘WSD+DA’, results in a slight drop in mAP. As discussed in Section 4.2, DA could result in an unexpected feature confusion among object classes with similar appearance, such as vehicle classes and animal classes. Only for classes that are different from all the other classes, such as “tv monitor”, DA shows its contribution to the detection results.

It can be seen that by further adding the ST stream, ‘WSD+DA+ST’, the detection results improve significantly, by 2.7% for VGG\_M and 3.2% for VGG16, compared with the baselines. In addition, inspired by the idea of [27], we also evaluate the performance of adding multiple pseudo label transfer streams one by one. Specifically, the pseudo labels generated by the first ST stream are used

**Table 2.** Average precision results (%) of different component combinations on VOC2007 test set.

	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
WSD(Baseline)-VGG_M	31.1	27.1	<b>18.6</b>	10.0	9.1	29.9	37.7	21.5	2.7	15.8	21.5	<b>27.8</b>	30.0	35.7	<b>10.8</b>	9.9	17.6	28.9	23.1	21.1	21.5
WSD+DA-VGG_M	30.3	24.1	15.6	<b>13.8</b>	9.1	32.7	<b>39.0</b>	21.4	2.9	19.0	26.4	25.5	24.7	32.9	4.3	8.2	15.6	28.7	24.5	25.1	21.2
WSD+DA+ST-VGG_M	33.3	<b>31.5</b>	16.9	<b>13.8</b>	<b>10.8</b>	39.5	36.2	<b>30.8</b>	8	19.9	33.4	18.4	26.4	37.8	8.3	<b>13.1</b>	15.5	32.1	25.0	33.8	24.2
WSD+DA+2ST-VGG_M	34.3	31.3	18.5	9.4	10.6	39.6	37.7	17.9	<b>10.2</b>	16.7	<b>34.7</b>	19.8	<b>31.8</b>	40.7	7.4	12.5	18.6	<b>33.0</b>	26.8	34.6	24.3
WSD+DA+3ST-VGG_M	<b>35.6</b>	31.3	18.2	7.7	9.1	<b>40.4</b>	38.4	23.8	9.7	<b>20.1</b>	33.4	22.5	30.9	<b>41.4</b>	9.8	10.8	<b>18.7</b>	28.7	<b>27.1</b>	<b>34.7</b>	<b>24.6</b>
WSD(Baseline)-VGG16	<b>45.8</b>	28.2	11.1	8.5	2.5	42.8	<b>41.5</b>	25.9	4.2	15.9	13.0	16.9	<b>28.0</b>	40.8	3.6	5.5	11.0	38.5	28.4	23.2	21.8
WSD+DA-VGG16	33.8	22.4	13.1	13.4	9.1	38.1	36.5	25.8	9.2	20.1	12.6	<b>19.8</b>	19.9	34.4	4.4	10.8	13.8	30	26.8	25.1	21.0
WSD+DA+ST-VGG16	43.7	30.8	15.7	10.6	<b>13.4</b>	41.3	39.5	23.9	<b>12.8</b>	<b>20.7</b>	27.9	13.9	23.4	39.7	<b>10.3</b>	12.7	<b>21.3</b>	<b>39.6</b>	28.1	30.7	25.0
WSD+DA+2ST-VGG16	44.7	<b>31.0</b>	12.1	15.7	11.8	38.8	40.6	29.1	12.0	17.9	<b>32.2</b>	9.1	24.1	42.8	7.6	<b>13.7</b>	17.0	33.4	30.6	<b>33.5</b>	24.9
WSD+DA+3ST-VGG16	40.6	30.1	<b>17.8</b>	<b>15.9</b>	6.4	<b>42.9</b>	40.5	<b>31.5</b>	11.4	20.3	27.4	15.7	24.1	<b>43.8</b>	8.9	12.2	17.7	37.3	<b>32.1</b>	31.0	<b>25.4</b>

**Table 3.** Average precision results (%) on VOC2012 test set.

	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
WSD(Baseline)-VGG_M	39.7	25.4	12.6	5.8	2.3	32.3	<b>25.0</b>	20.7	1.6	17.9	9.6	<b>29.0</b>	24.3	42.4	3.8	4.6	10.6	16.6	22.5	11.4	17.9
WSD+DA+3ST-VGG_M	<b>44.3</b>	<b>29.8</b>	<b>15.6</b>	<b>6.6</b>	<b>6.0</b>	<b>34.4</b>	24.2	<b>25.1</b>	<b>5.7</b>	<b>20.3</b>	<b>22.3</b>	24.9	<b>29.1</b>	<b>45.2</b>	<b>7.8</b>	<b>9.4</b>	<b>12.4</b>	<b>21.4</b>	<b>22.6</b>	<b>26.0</b>	<b>21.7</b>
WSD(Baseline)-VGG16	47.9	29.2	14.8	<b>7.9</b>	3.5	<b>39.6</b>	<b>27.3</b>	24.6	2.3	15.9	4.9	18.3	25.5	<b>47.5</b>	3.8	4.3	9.4	22.2	19.3	16.0	19.2
WSD+DA+3ST-VGG16	<b>48.8</b>	<b>32.8</b>	<b>16.6</b>	6.3	<b>7.7</b>	39.0	26.2	<b>32.6</b>	<b>7.8</b>	<b>18.3</b>	<b>12.4</b>	<b>22.1</b>	<b>29.7</b>	45.9	<b>9.6</b>	<b>9.0</b>	<b>14.5</b>	<b>24.0</b>	<b>26.8</b>	<b>28.1</b>	<b>22.9</b>


as the supervision of the second ST stream, whose generated pseudo labels are then used as the supervision of the third ST stream. The results of appending multiple ST streams, ‘WSD+DA+2ST’ and ‘WSD+DA+3ST’, are also shown in Table 2. We can see that adding one additional ST stream generally leads to slight improvements. Overall, by adding the ST streams, our method brings up the results for most categories, especially difficult classes such as “chairs” and “dining tables”. These classes are usually in cluttered scenes and the single WSD learned from clean web images can hardly capture the objects from the environment.

The overall performance gains from the best combinations are 3.1% for VGG\_M and 3.6% for VGG16. These results show that our proposed method improves the baseline weakly supervised detection model significantly by introducing the DA and ST streams. In VGG16, it brings up the unsupervised results to 25.4% without any labels from the target dataset, much closer to the weakly supervised result of 29.3% that requires image-level labels from the target dataset.

In addition to the mAP results for detection, we also measured the correct localization (CorLoc) result on the VOC 2007 trainval set (see Table. 5 ) and compare it with the best reported CorLoc results of the WSD works [3, 4, 27]. Note that all of these WSD methods use image labels of VOC trainval set during the training and CorLoc is measured on these training images. In our method, we do not include any VOC training labels and we can still achieve a good localization model, 44.3% images are with correctly localized objects, which is even better than [3].

## 5.4 Ablation experiments

In the following sections, we analyze the effectiveness of each component, including the domain adaptation stream and the simultaneous transfer streams.



**Fig. 5.** Visualization of features in 2D space by t-SNE [23]. We randomly sample some object proposals from target and web domains and extract fc7 features (VGG\_M) using different methods. Then we use PCA and t-SNE to reduce the dimension to 2. We plot the scatter diagrams for all mammal animal classes. Left: WSD (baseline). Middle: WSD+DA. Right: WSD+DA+3ST.

**Analysis of the DA stream.** To further verify the effects of the DA stream, we visualizing the feature distributions of ‘WSD+DA’ in 2D space by t-SNE [23] in Fig. 5. Although this visualization of high dimensional features in 2D space may not be accurate, we can still have some ideas that the DA stream does help shift the features closer to the same region across domains.

We further examine the results by removing the DA stream from the overall structure. As shown in Table 4, the WSD with the ST streams only cannot achieve as high detection mAP as our overall network with the DA stream, which demonstrates the contribution of DA to the overall network. In Table 4, we also evaluate the effectiveness of the foreground attention mechanism (FA) for the DA stream. It can be seen that the result of DA without FA, ‘WSD+DA(w/o.FA)+3ST’, is even worse than of no DA, ‘WSD+3ST’, which suggests that treating all proposals equally during DA does not help.

**Table 4.** Comparing the results (mAP in %) on VOC 2007 test set with different settings of the DA stream.

Method	mAP
WSD+3ST-VGG_M	23.5
WSD+DA(w/o.FA)+3ST-VGG_M	23.3
WSD+DA+3ST-VGG_M	<b>24.6</b>

**Table 5.** CorLoc results on VOC 2007 compared with WSD methods.

Method	CorLoc
Bilen et al [3]	43.7
Bilen et al [4]	56.1
Tang et al [27]	60.6
WSD+DA+3ST-VGG16	44.3

**Analysis of the ST stream.** We also visualized the features of WSD+DA+3ST in 2D space in Fig. 5. It can be seen that by adding both DA and ST, we are able to move the cross-domain features closer while making the classes in target domain more separable.

We would like to point out that the incremental gains of our method with multiple ST streams are not as much as [27] that also use multiple refinement streams. This is due to the following reason. In [27], the positive samples are selected by image-level labels of target dataset and their purpose is to refine the instance classifier for multiple times. However, our method does not use the image labels of VOC dataset and our purpose is to prevent the unexpected distribution shift among similar classes. In other words, the gain of pseudo label transfer in our scenario is mainly from the effects of preserving the semantic structure among classes rather than refining the instance classifiers again and again.

One insight of the ST stream is that our framework trains the WSD model from web domain and selects pseudo ground truth samples of target domain based on the current WSD model at the same time. In other words, the ST stream is trained simultaneously with the WSD stream. In this way, it shares the feature learning between the WSD stream for web image training and the ST stream for target dataset training. An alternative way of transferring the pseudo labels is to train on the two datasets in an isolated way. In particular, we can first pre-train the WSD using web images, then use this pre-trained WSD model to generate the pseudo ground truths for the target dataset and finally use these pseudo ground truths to train a detector for target dataset. We conduct the experiment using such isolated method and obtain an mAP of 22.5%. This implies that the simultaneous weights sharing is important for the learning transfer across domains.

### 5.5 More results

We also evaluate our method on VOC 2012 dataset and the results are shown in Table 3. The baseline result shows that the detection model trained using only web images gives poor results for VOC 2012 test images. By adding our DA stream and ST streams, the results are largely improved for most classes. Overall, we achieve significant increases of 3.8% and 3.7% in mAP with VGG\_M and VGG16 respectively for VOC 2012 dataset.

## 6 Conclusion

In conclusion, we introduced an annotation-free object detection method by learning from web image resources. Particularly, to solve the domain mismatch problem between the web domain objects and the target domain objects, we proposed an instance-level domain adaptation stream with foreground attention, together with a simultaneous transfer stream that simultaneously learns target data from pseudo labels. Through these novel components, we achieved significant improvements in detection results and successfully reduced the performance gap between the baseline detectors trained with and without human annotations.

**Acknowledgements.** This project is partially supported by MoE Tier-2 Grant (MOE2016-T2-2-065).

## References

1. Al-Halah, Z., Stiefelhagen, R.: Automatic discovery, association estimation and learning of semantic attributes for a thousand categories. arXiv preprint arXiv:1704.03607 (2017)
2. Bilen, H., Pedersoli, M., Tuytelaars, T.: Weakly supervised object detection with posterior regularization. In: Proceedings BMVC 2014. pp. 1–12 (2014)
3. Bilen, H., Pedersoli, M., Tuytelaars, T.: Weakly supervised object detection with convex clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1081–1089 (2015)
4. Bilen, H., Vedaldi, A.: Weakly supervised deep detection networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2846–2854 (2016)
5. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). vol. 1, p. 7 (2017)
6. Chen, X., Gupta, A.: Webly supervised learning of convolutional networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1431–1439 (2015)
7. Chen, X., Shrivastava, A., Gupta, A.: Neil: Extracting visual knowledge from web data. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1409–1416 (2013)
8. Cinbis, R.G., Verbeek, J., Schmid, C.: Weakly supervised object localization with multi-fold multiple instance learning. *IEEE transactions on pattern analysis and machine intelligence* **39**(1), 189–203 (2017)
9. Divvala, S.K., Farhadi, A., Guestrin, C.: Learning everything about anything: Webly-supervised visual concept learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3270–3277 (2014)
10. Ferrari, V., Zisserman, A.: Learning visual attributes. In: Advances in Neural Information Processing Systems. pp. 433–440 (2008)
11. Fu, Z., Xiang, T., Kodirov, E., Gong, S.: Zero-shot object recognition by semantic manifold distance. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2635–2644 (2015)
12. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International Conference on Machine Learning. pp. 1180–1189 (2015)
13. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 1440–1448 (2015)
14. Jie, Z., Wei, Y., Jin, X., Feng, J., Liu, W.: Deep self-taught learning for weakly supervised object localization. arXiv preprint arXiv:1704.05188 (2017)
15. Kantorov, V., Oquab, M., Cho, M., Laptev, I.: Contextlocnet: Context-aware deep network models for weakly supervised localization. In: European Conference on Computer Vision. pp. 350–365. Springer (2016)
16. Kumar Singh, K., Xiao, F., Jae Lee, Y.: Track and transfer: Watching videos to simulate strong human supervision for weakly-supervised object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3548–3556 (2016)
17. Lampert, C.H., Nickisch, H., Harmeling, S.: Learning to detect unseen object classes by between-class attribute transfer. In: Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on. pp. 951–958. IEEE (2009)

18. Lampert, C.H., Nickisch, H., Harmeling, S.: Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(3), 453–465 (2014)
19. Li, L.J., Fei-Fei, L.: Optimol: automatic online picture collection via incremental model learning. *International journal of computer vision* **88**(2), 147–168 (2010)
20. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection
21. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: *European conference on computer vision*. pp. 21–37. Springer (2016)
22. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Deep transfer learning with joint adaptation networks. *arXiv preprint arXiv:1605.06636* (2016)
23. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* **9**(Nov), 2579–2605 (2008)
24. Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. pp. 1532–1543 (2014)
25. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: *Advances in neural information processing systems*. pp. 91–99 (2015)
26. Sultani, W., Shah, M.: What if we do not have multiple videos of the same action? video action localization using web images. In: *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*. pp. 1077–1085. IEEE (2016)
27. Tang, P., Wang, X., Bai, X., Liu, W.: Multiple instance detection network with online instance classifier refinement. In: *CVPR* (2017)
28. Tao, Q., Yang, H., Cai, J.: Exploiting web images for weakly supervised object detection. *arXiv preprint arXiv:1707.08721* (2017)
29. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4068–4076 (2015)
30. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: *Computer Vision and Pattern Recognition (CVPR)*. vol. 1, p. 4 (2017)
31. Uijlings, J.R., Van De Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. *International journal of computer vision* **104**(2), 154–171 (2013)
32. Wang, G., Luo, P., Lin, L., Wang, X.: Learning object interactions and descriptions for semantic image segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5859–5867 (2017)
33. Wei, Y., Liang, X., Chen, Y., Shen, X., Cheng, M.M., Feng, J., Zhao, Y., Yan, S.: Stc: A simple to complex framework for weakly-supervised semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence* **39**(11), 2314–2320 (2017)
34. Xia, Y., Cao, X., Wen, F., Sun, J.: Well begun is half done: Generating high-quality seeds for automatic image dataset construction from web. In: *European Conference on Computer Vision*. pp. 387–400. Springer (2014)
35. Xu, Z., Huang, S., Zhang, Y., Tao, D.: Augmenting strong supervision using web data for fine-grained categorization. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2524–2532 (2015)
36. Yang, X., Zhang, H., Cai, J.: Shuffle-then-assemble: Learning object-agnostic visual relationship features. In: *ECCV* (2018)