

Cross-Domain Weakly-Supervised Object Detection through Progressive Domain Adaptation

Naoto Inoue Ryosuke Furuta Toshihiko Yamasaki Kiyoharu Aizawa
The University of Tokyo, Japan

{inoue, furuta, yamasaki, aizawa}@hal.t.u-tokyo.ac.jp

Abstract

Can we detect common objects in a variety of image domains without instance-level annotations? In this paper, we present a framework for a novel task, cross-domain weakly supervised object detection, which addresses this question. For this paper, we have access to images with instance-level annotations in a source domain (e.g., natural image) and images with image-level annotations in a target domain (e.g., watercolor). In addition, the classes to be detected in the target domain are all or a subset of those in the source domain. Starting from a fully supervised object detector, which is pre-trained on the source domain, we propose a two-step progressive domain adaptation technique by fine-tuning the detector on two types of artificially and automatically generated samples. We test our methods on our newly collected datasets¹ containing three image domains, and achieve an improvement of approximately 5 to 20 percentage points in terms of mean average precision (mAP) compared to the best-performing baselines.

1. Introduction

Object detection is a task to localize instances of particular object classes in an image. It is a fundamental task and has advanced rapidly due to the development of convolutional neural networks (CNNs). Best-performing detectors [9, 28, 22, 27, 19, 20] are fully supervised detectors (FSDs). They are highly data-hungry and typically learned from many images with instance-level annotations. An instance-level annotation is composed of a label (i.e., the object class of an instance) and a bounding box (i.e., the location of the instance).

While object detection in a natural image domain has achieved outstanding performance, less attention has been paid to the detection in other domains such as watercolor. This is because it is often difficult and unrealistic to construct

¹Datasets and codes are available at https://naoto0804.github.io/cross_domain_detection

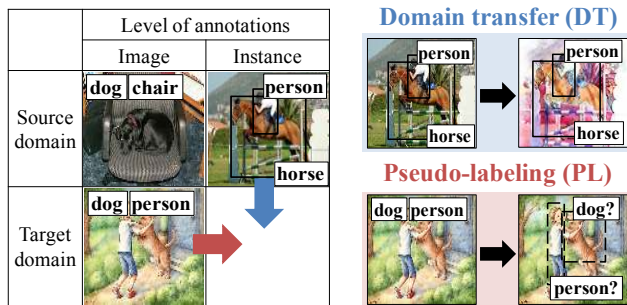


Figure 1: Left: the situation of the cross-domain weakly supervised object detection; Right: Our methods to generate instance-level annotated samples in the target domain.

a large dataset with instance-level annotations in many image domains. There are many obstacles such as lack of image sources, copyright issues, and the cost of annotation.

We tackle a novel task, cross-domain weakly supervised object detection. The task is described as follows: (i) *instance*-level annotations are available in a *source* domain; (ii) only *image*-level annotations are available in a *target* domain; (iii) the classes to be detected in the target domain are all or a subset of those in the source domain. The objective of the task is to detect objects as accurately as possible in the target domain under these conditions by using sufficient instance-level annotations in the source domain and a small number of image-level annotations in the target domain. This assumption is reasonable as it is easier to collect image-level annotations than instance-level annotations from existing datasets or an image search engine.

We will describe a framework to solve the proposed task. Starting from an FSD trained on images with instance-level annotations in the source domain, we fine-tune the FSD in the target domain, as this is the most straightforward and promising approach. However, there are no *instance*-level annotations available in the target domain. Instead, as shown in Fig. 1, we present two methods to generate images with *instance*-level annotations artificially and automatically, and fine-tune the FSD on them. The first method, *domain trans-*

fer (DT), is used to generate images that look like those in the target domain from images in the source domain having *instance*-level annotations. This generation is achieved by image-to-image translation methods from unpaired examples such as CycleGAN [40]. The second method, *pseudo-labeling* (PL), is used to generate pseudo *instance*-level annotations. Given images with *image*-level annotations in the target domain and the FSD which is fine-tuned on the artificially generated samples by DT, these annotations and predictions of the FSD are combined. We achieve a two-step progressive domain adaptation by sequentially fine-tuning the FSD on the artificially generated samples. Our framework is general to cross-domain weakly supervised object detection across any image domain and is relatively scalable to many classes and instances.

Since there is no dataset for the target domain that is suitable to evaluate the proposed task, we construct new datasets with instance-level annotations, which we call *Clipart1k*, *Watercolor2k*, and *Comic2k*. Each dataset comprises 1,000, 2,000, and 2,000 images of clipart, watercolor, and comic, respectively. The validity of our methods is demonstrated using these datasets. We show that the proposed two-step adaptation achieves an improvement of approximately 5 to 20 percentage points as compared to the best-performing baselines' mAP across all datasets. We believe that this paper itself can be a strong baseline for cross-domain weakly supervised object detection.

Our main contributions are as follows:

- We propose a framework for a novel task, cross-domain weakly supervised object detection. We achieve a two-step progressive domain adaptation by sequentially fine-tuning the FSD on the artificially generated samples by the proposed domain transfer and pseudo-labeling.
- We construct novel, fully instance-level annotated datasets with multiple instances of various object classes across three domains that are far from natural images.
- Our experimental results show that our framework outperforms the best-performing baselines by approximately 5 to 20 percentage points in terms of mAP.

2. Related Work

2.1. Fully Supervised Detection

Standard methods in fully supervised object detection, such as R-CNN [10], Fast R-CNN [9], and Faster R-CNN [28], are based on a two-stage approach: generating region proposals and then, classifying them. Recently, single-stage object detectors such as SSD [22], YOLOv2 [27], and RetinaNet [20] have also emerged. All of these detectors require large datasets with instance-level annotations such as PASCAL VOC [6], Microsoft Common Objects in Context (MSCOCO) [21], and OpenImages [17].

Dataset construction for a new image domain becomes harder as the number of images and classes increases. [32] reported that it took 35 seconds for a worker to annotate one bounding box. Recently, [26] reduced it to 7 seconds through extreme clicking, while it still takes much time and effort to obtain large-scale datasets. On the contrary, our framework does not require instance-level annotations for the new target domain at all.

2.2. Weakly Supervised Detection

One possible approach addressing the lack of large-scale instance-level annotations for object detection is to use a weakly supervised detector (WSD). In weakly supervised object detection, only pairs of an image and an image-level annotation (i.e., labels of objects in each image) are provided for training. Many existing methods are built upon region-of-interest (RoI) extraction methods such as selective search [36]. Feature extraction for each region, region selection, and classification of the selected region are performed through multiple instance learning (MIL) [30, 11, 31, 1, 18] or two-stream CNN [2, 15, 33]. However, WSDs are poor at accurately localizing the object boundary. Our framework uses image-level annotations in the target domain by pseudo-labeling the image.

2.3. Cross-domain Object Detection

Using an object detector that is neither trained nor fine-tuned for the target domain causes a significant drop in performance as shown in [38]. Therefore, adapting the detector with the help of some information on the target domain is essential. [13, 5] are the some of the best works closely related to this paper. Our methods and [13] are similar as they propose to learn from a combination of instance- and image-level annotations. However, we address the adaptation of the detector from one domain to another, whereas [13] addresses the classifier-to-detector adaptation for weakly labeled object classes within one domain. This paper and [5] are similar as they tackle the adaptation of the detector from one domain to another. However, only image-level annotations are available in the source domain in [5]. This is the first work to propose the cross-domain weakly supervised object detection.

For evaluating the cross-domain object detection method, the existing datasets for detecting common objects in various domains seem to have limitations. People-Art [37] is used only for single-class detection in an artwork. Photo-Art [39] assumes only one instance per image, which is unrealistic. Besides, we introduce the fully instance-level annotated datasets for object detection which comprises multiple common classes to be detected various visual domains.

2.4. Unsupervised Domain Adaptation

Unsupervised domain adaptation (UDA) in an image is a task used for learning domain invariant models, where

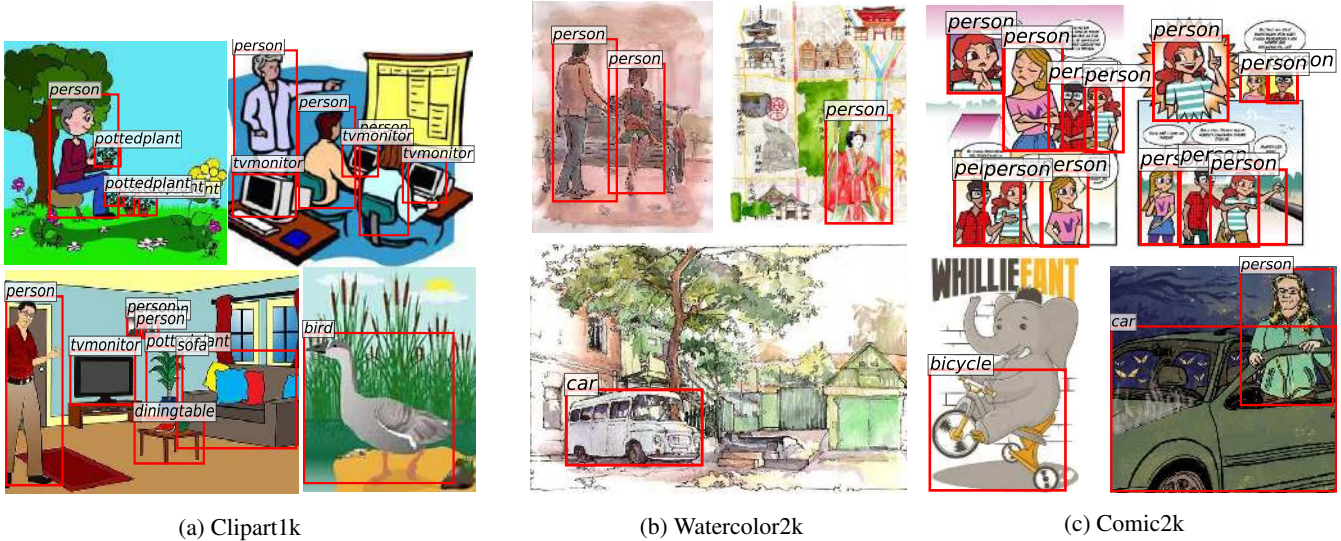


Figure 2: Examples of datasets that we collected across three domains; The images usually contain not only the target objects but also various other objects and complex backgrounds.

pairs of an image and annotation are available in the source domain while only images are available in the target domain. Previous works for UDA in image classification is mostly distribution-matching-based, in which features extracted from two domains are made to closely resemble each other using the maximum mean discrepancy (MMD) [12] or a domain classifier network [23, 8, 24, 35]. Although current distribution-matching-based methods are applicable, it is primarily challenging to fully align the distribution for tasks that require structured outputs, such as object detection. This is because it is essential to keep the spatial information in the feature map. Our framework employs image-to-image translation and fine-tuning to avoid this problem.

3. Dataset

Our objective is to detect objects in a target domain by adapting an FSD that is originally trained on a source domain. The classes to be detected in the target domain are all or a subset of the classes defined in the dataset which is in the source domain. In this paper, PASCAL VOC [6], which contains twenty classes, was used for the source domain, natural image. As no suitable dataset for the target domain of our task was available, we constructed three original datasets, Clipart1k, Watercolor2k, and Comic2k using Amazon Mechanical Turk.

Examples of the images are shown in Fig. 2. These images usually contain multiple objects per image, and some instances are small or partially occluded by the other objects. The statistics of the three datasets are shown in Table 1. We collected a total of 5,000 images and 12,869 instance-level annotations. We believe these datasets are good benchmarks

Table 1: List of the datasets that we constructed for the target domains in this paper.

Dataset	#classes	#images	#instances
Clipart1k	20	1,000	3,165
Watercolor2k	6	2,000	3,315
Comic2k	6	2,000	6,389

not only for domain adaptation but also for fully and weakly supervised, semi-supervised detection tasks. For the more detailed statistics, please refer to the supplementary material. In the following subsection, we will briefly describe each dataset, and the data collection method.

3.1. Clipart1k

In Clipart1k, the target domain classes to be detected were the same as those in the source domain. All the images for a clipart domain were collected from one dataset (i.e., CMPlaces [4]) and two image search engines (i.e., Openclipart² and Pixabay³). When we collected the images from the search engines, we used queries of the 205 scene classes (e.g., pasture) used in CMPlaces to collect various objects and scenes with complex backgrounds.

3.2. Comic2k and Watercolor2k

In Comic2k and Watercolor2k, the classes to be detected in the target domain were the subset of those in the source domain. The images were collected from BAM! [38].

²<https://openclipart.org/>

³<https://pixabay.com/>

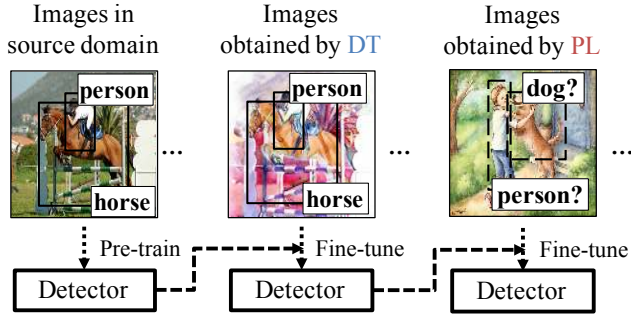


Figure 3: The workflow of our framework.

In BAM!, millions of images with slightly noisy (80%–90% in precision) image-level attributes regarding object classes, domain, and emotion are provided in a human-in-the-loop fashion. Specifically, the target classes are `bicycle`, `bird`, `cat`, `car`, `dog`, and `person`, which are representative of the intersection of the classes in VOC and those in BAM!. We chose the watercolor and comic domains as the other domains in BAM! are not suitable for object detection. For example, oil paint images are unsuitable as they usually depict a single person in the center of the image, making object detection a trivial task.

As collecting instance-level annotations for all images in BAM! is difficult, we annotated the images in the following way: (i) images that contained at least one of the six target classes were extracted. Note that we relied on the labels provided by BAM! and did not conduct any other filtering process. We obtained 17,814 and 52,790 images for watercolor and comic domains, respectively. (ii) as many as 2,000 images are randomly sampled and assigned instance-level annotations for each domain. The remaining 15,814 and 50,790 images, which we called `extra` dataset, are still useful as they possess many image-level annotations. Although they are noisy and incomplete, they provided room for further improvement with respect to detector performance as shown in Sec. 5.3.

4. Proposed Method

We propose a framework to adapt an FSD that is pre-trained on a source domain. The adaptation is achieved through fine-tuning the FSD on artificially generated samples with instance-level annotations in a target domain. We propose two methods to generate the samples as shown in Fig. 1: (i) domain transfer (DT), transferring images with instance-level annotations from the source domain to the target domain, and (ii) pseudo-labeling (PL), pseudo-labeling the images with image-level annotations in the target domain. The samples generated by these two methods display different properties. Although the samples generated by (i) are not high-quality images with respect to their similarities to

target-domain images, bounding boxes are correctly annotated. On the contrary, although the samples generated by (ii) do not have accurate bounding boxes, image quality is guaranteed as they are completely target-domain images.

We progressively fine-tune an FSD using these examples as shown in Fig. 3. First, we pre-train it while using instance-level annotations in the source domain. Second, we fine-tune it while using the images obtained by DT. Lastly, we fine-tune it while using the images obtained by PL. We would like to emphasize that the sequential execution of the two fine-tuning steps is critical as the performance of PL highly depends on the used FSD.

4.1. Domain Transfer (DT)

The differences between the source and target domains tackled in this paper mainly lie in their low-level features, such as color and texture. We generate images that look like those in the target domain to capture such differences and then, make the FSD robust to such differences by fine-tuning the FSD on the generated images.

To achieve this goal, an unpaired image-to-image translation method called CycleGAN [40] is employed. With CycleGAN, the goal is to learn the mapping functions between two image domains \mathcal{X} and \mathcal{Y} with unpaired examples. In practice, a mapping $G : \mathcal{X} \rightarrow \mathcal{Y}$ and an inverse mapping $F : \mathcal{Y} \rightarrow \mathcal{X}$ are jointly learned using CNN. We train CycleGAN to learn the mapping functions between the source domain, \mathcal{X}_s , and the target domain, \mathcal{X}_t . Once the mapping functions are trained, we convert images in the source domain that are used in the pre-training and obtain domain-transferred images that accompany instance-level annotations. Using these images, the FSD is fine-tuned.

4.2. Pseudo-Labeling (PL)

In the target domain, if we use an FSD that is trained only on the source domain for object detection, then the FSD mainly fails because of confusion with other classes and backgrounds rather than inaccurate localization. We will later show this trend in Fig. 4. Fine-tuning FSD on images obtained by PL dramatically reduces such confusion. PL is simple and applicable in any FSD as it does not access intermediate layers of an FSD.

Formally, the objective of PL is to obtain a pseudo instance-level annotation \mathbf{G} for each image \mathbf{x} from the target domain \mathcal{X}_t . Let $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ denote an RGB image, where H and W are the image’s height and width, respectively. \mathcal{C} indicates a set of object classes. \mathbf{z} indicates an image-level annotation: the set of classes in \mathbf{x} . Further, \mathbf{G} comprises $g = (\mathbf{b}, c)$, where $\mathbf{b} \in \mathbb{R}^4$ is a bounding box, and $c \in \mathcal{C}$. First, we obtain FSD outputs \mathbf{D} . \mathbf{D} comprises each detection $d = (p, \mathbf{b}, c)$, where $c \in \mathcal{C}$ and $p \in \mathbb{R}$ indicates the probability of \mathbf{b} belonging to c . Second, for each class $c \in \mathcal{C}$, we take the top-1 confident detection $d = (p, \mathbf{b}, c) \in \mathbf{D}$

and add (b, c) to G . We fine-tune the FSD using pairs of x and G . Note that no layers of the FSD are replaced to preserve the original network’s detection ability. The FSD that is trained on images obtained by DT was subsequently fine-tuned on these images.

5. Experiments

In Sec. 5.1, we explain the details of the implementations, the compared methods, and the evaluation metrics. In Sec. 5.2, we test our methods using Clipart1k and conduct error analysis and ablation studies on the FSDs. In Sec. 5.3, we confirm that our framework is generalized for a variety of domains using Watercolor2k and Comic2k. In Sec. 5.4, we show actual detection results and the generated domain-transferred images for further discussion.

5.1. Implementation and Evaluation Metrics

Our methods were implemented using Chainer [34]. We evaluated our methods using average precision (AP) and its mean, i.e., mAP.

Dataset Arrangement for Training and Evaluation VOC2007-trainval and VOC2012-trainval [6] were used as images in the source domain (i.e., natural image) in all the experiments. For the target-domain images, the ones with instance-level annotations were used when discussing the performance gap between our methods and the ideal case quantitatively. The target-domain images were split into training set and test set by a ratio of 1:1. For the training set, the bounding box information was ignored to meet the proposed situation. For the test set, the labels and the bounding boxes of the annotations were used to evaluate the performance of the compared methods and our methods.

Comparison We compared our methods against the following methods:

- **Baseline:** SSD300 [22] was used as our baseline FSD. We used the implementation provided by ChainerCV [25], and we obtained an SSD300, which was pre-trained on VOC2007-trainval and VOC2012-trainval, and skipped the pre-training. We followed the original paper on hyper-parameters of SSD300 unless specified. The input images were resized to 300×300 in SSD300. The IoU threshold for NMS (0.45) and the confidence threshold for discarding low confidence detections (0.01) were employed.
- **Ideal case:** In this case, we had access to the instance-level annotations for the training set of the target domain dataset. We simply fine-tuned the baseline FSD using these annotations. This experiment was to confirm the weak upper-bound performance of our methods.

- **Weakly supervised detection (WSD):** ContextLocNet (CLNet) [15] and WSDDN [2] were chosen as the compared WSDs. Each WSD was trained on the images with image-level annotations from the training set of the dataset in the target domain.

- **Unsupervised domain adaptation (UDA):** We tested one of the state-of-the-art UDA methods, ADDA [35]. We aligned the distributions at the `relu4_3` layer in SSD300. We trained the model with a batch size of 32 and a learning rate of 1.0×10^{-6} for 1,000 iterations using Adam [16].

- **Ensemble:** In image classification, unweighted averaging, which uses the unweighted average of the output score / probability of all the base models as the output, is a reasonable approach to boost performance. We accumulated all the detections produced by multiple detectors and applied non-maximum suppression (NMS)⁴ to them. The parameters of NMS were the same as those used in SSD300.

Details of Training We trained CycleGAN with a learning rate of 1.0×10^{-5} for the first ten epochs and a linear decay rate to zero over the next ten epochs. We followed the original paper on the other hyper-parameters of CycleGAN. When fine-tuning SSD300, we employed a learning rate of 1.0×10^{-5} , which is the same as the final learning rate for the original SSD300 training. Fine-tuning, using the images obtained by DT and PL, was conducted for one epoch and 10,000 iterations, respectively.

5.2. Quantitative Results on Clipart1k

Table 2 shows the comparison of AP for each class and mAP among our methods against the baseline FSD and the comparable methods. We observe that SSD300 performs better than the WSDs in terms of mAP, although SSD300 is not trained to adapt to the target domain. On the contrary, WSDs perform poorer due to insufficient data and their poor localization ability. The ensembling of WSDs with SSD300 have almost no effect as shown in the case of Ensemble. Conventional distribution-matching-based methods do not work well as shown in the case of ADDA.

The results of our methods based on SSD300 are shown in the bottom half of Table 2. To quantify the relative contribution of each step, the performances of our methods are examined using different configurations.

- **DT+PL:** the proposed two-step fine-tuning.
- **DT:** only fine-tuning using images obtained by DT.
- **PL:** only fine-tuning using images obtained by PL. Note that the baseline FSD is used for PL.

PL provides an improvement of 9.6 percentage points improvement from the baseline SSD300 in terms of mAP. Further, DT+PL achieves 46.0 % in terms of mAP. This result

⁴Further details about NMS can be found in [7].

Table 2: Comparison of all the methods in terms of AP [%] using SSD300 as the baseline FSD in Clipart1k. Ensemble denotes an ensemble of SSD300, CLNet, and WSDDN.

Method	AP for each class																			mAP	
	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train		tv
Baseline	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
<i>Compared</i>																					
WSDDN [2]	1.6	3.6	0.6	2.3	0.1	11.7	4.5	0.0	3.2	0.1	2.8	2.3	0.9	0.1	14.4	16.0	4.5	0.7	1.2	18.3	4.4
CLNet [15]	3.2	22.3	2.2	0.7	4.6	4.8	17.5	0.2	4.8	1.6	6.4	0.6	4.7	0.6	12.5	13.1	14.1	4.1	8.0	29.7	7.8
Ensemble	20.6	49.6	20.5	23.4	11.3	39.3	35.2	2.6	39.0	22.8	27.3	11.2	33.2	54.7	34.0	30.7	21.0	20.3	20.3	18.3	26.7
ADDA [35]	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
<i>Proposed</i>																					
PL w/o label	18.6	40.3	17.1	16.7	4.9	35.3	36.1	1.1	36.0	22.9	29.1	14.7	31.5	52.6	43.8	28.6	13.3	14.6	32.8	15.1	25.3
PL	24.2	59.8	22.0	26.6	25.0	54.7	51.3	3.9	47.4	44.5	40.3	14.3	33.6	55.1	50.8	41.1	23.2	26.3	40.5	43.2	36.4
DT	23.3	60.1	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
DT+PL w/o label	16.8	53.7	19.7	31.9	21.3	39.3	39.8	2.2	42.7	46.3	24.5	13.0	42.8	50.4	53.3	38.5	14.9	25.1	41.5	37.3	32.7
DT+PL	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
Ideal case	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

ensures that both of our methods work and are complementary. The mAP of the combination of DT+PL is 19.2 percentage points higher than that of the baseline SSD300 and is approximately 18 percentage points greater than the ensemble of the detectors. We emphasize that this performance is only 9.4 percentage points lower than `Ideal case`.

Ablation Study We considered a setting where we can obtain only images with no annotation in the target domain.

DT is applicable without any modification. DT provides an improvement of 11.2 percentage points improvement from the baseline SSD300 in terms of mAP in Table 2.

PL is not directly applicable as we do not have access to the image-label annotations. To address this problem, only one detection d_{best} , which has the highest probability p among all detections, can be pseudo-labeled. The results are shown as `PL w/o label` and `DT+PL w/o label` in Table 2. Fine-tuning the FSD on the images labeled by this method harms the performance as the result of pseudo-labeling contains a lot of inaccuracy. Therefore, image-level annotations in the target domain are essential for PL, which greatly improves the detection performance.

Generality across Detectors We investigated our framework on other FSDs such as Faster R-CNN [28] and YOLOv2 [27]. Please refer to the supplementary material about details of the hyper-parameters. The result further emphasizes the generality of our framework across all baseline FSDs, as shown in Table 3. We additionally found that the ensembling of SSD300 and Faster R-CNN yields 30.2 % in terms of mAP and that of all the three FSDs yields 31.0 % in terms of mAP, which is not so remarkable compared to

Table 3: Results of our methods on the different baseline FSDs in terms of mAP [%] in Clipart1k.

Method	SSD300	YOLOv2	Faster R-CNN
Baseline	26.8	25.5	26.2
DT	38.0	31.5	32.1
PL	36.4	34.0	29.8
DT+PL	46.0	39.9	34.9
Ideal case	55.4	51.2	50.0

the improvement by DT+PL. The performance gain is significant in SSD300 compared to YOLOv2 and Faster R-CNN. This result suggests the importance of data augmentation (e.g., the zoom-in and zoom-out features implemented in SSD300) during the process of training FSDs with pseudo-labeled annotations, which are often noisy and incomplete.

Performance Analysis Focusing on Errors The tool from [14] was used to understand the type of the detection error that is reduced by our methods. The classes within the brackets were regarded as the same category: {all vehicles}, {all animals including person}, {chair, dining table, sofa}(furniture), {aeroplane, bird}(air objects). Considering the class, the category, and the IoU between the predicted bounding box and the ground truth bounding box, the detections were classified into five groups as listed below:

- Correct (Cor): correct class and $IoU > .5$
- Localization (Loc): correct class, misaligned bounding

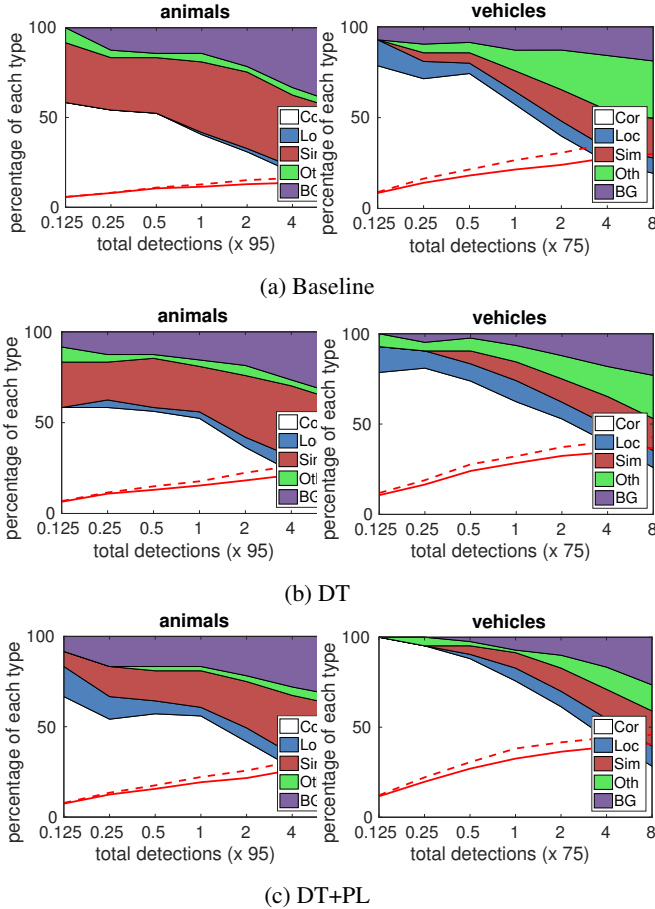


Figure 4: Visualization of performance for various methods on animals and vehicles in the test set of Clipart1k using SSD300 as the baseline FSD. The solid red line and dashed red line reflect the change of recall with strong criteria (0.5 jaccard overlap) and weak criteria (0.1 jaccard overlap) as the number of detections increases, respectively.

box ($.1 < IoU < .5$)

- Similar (Sim): wrong class, correct category, $IoU > .1$
- Other (Oth): wrong class, wrong category, $IoU > .1$
- Background (BG): $IoU < .1$ for any object

Fig. 4 shows the example of the error analysis in the Clipart1k test set. Comparing the baseline and DT, we observe that fine-tuning the FSD on images obtained by DT improves the detection performance, especially in less-confident detections. Comparing DT and DT+PL, we observe that the confusion, which emerged with the other classes (Sim and Oth), especially in more confident detections, is greatly reduced by PL, which uses the image-level annotations in the target domain to remove such confusions with the FSD.

Table 4: Comparison in terms of AP [%] using SSD300 as the baseline FSD in Watercolor2k.

Method	AP for each class						
	bike	bird	car	cat	dog	person	mAP
Baseline	79.8	49.5	38.1	35.1	30.4	65.1	49.6
<i>Compared</i>							
WSDDN [2]	1.5	26.0	14.6	0.4	0.5	33.3	12.7
CLNet [15]	4.5	27.9	19.6	14.3	6.4	31.4	17.4
Ensemble	79.8	49.6	38.1	35.2	30.4	58.7	48.6
ADDA [35]	79.9	49.5	39.5	35.3	29.4	65.1	49.8
<i>Proposed</i>							
PL	76.3	54.9	46.6	37.5	36.9	71.7	54.0
DT	82.8	47.0	40.2	34.6	35.3	62.5	50.4
DT+PL	76.5	54.9	46.0	37.4	38.5	72.3	54.3
PL (+extra)	84.8	57.7	48.0	44.9	46.6	72.6	59.1
DT+PL (+extra)	86.3	57.3	48.5	43.0	46.5	73.2	59.1
Ideal case	76.0	60.0	52.7	41.0	43.8	77.3	58.4

Table 5: Comparison in terms of AP [%] using SSD300 as the baseline FSD in Comic2k.

Method	AP for each class						
	bike	bird	car	cat	dog	person	mAP
Baseline	43.9	10.0	19.4	12.9	20.3	42.6	24.9
<i>Compared</i>							
WSDDN [2]	1.5	0.1	11.9	6.9	1.4	12.1	5.6
CLNet [15]	0.0	0.0	2.0	4.7	1.2	14.9	3.8
Ensemble	44.0	10.0	19.4	14.5	20.7	42.9	25.3
ADDA [35]	39.5	9.8	17.2	12.7	20.4	43.3	23.8
<i>Proposed</i>							
PL	52.9	13.7	35.3	16.2	28.9	50.8	32.9
DT	43.6	13.6	30.2	16.0	26.9	48.3	29.8
DT+PL	55.2	18.5	38.2	22.9	34.1	54.5	37.2
PL (+extra)	53.4	19.0	35.0	30.0	30.5	53.7	36.9
DT+PL (+extra)	56.6	24.0	40.7	35.8	39.0	57.3	42.2
Ideal case	55.9	26.8	40.4	42.3	43.0	70.1	46.4

5.3. Quantitative Results on Watercolor2k and Comic2k

The comparison among our methods against the baseline FSD and the comparable methods is shown in Table 4 and Table 5. In Watercolor2k, the learning rate was set to 10^{-6} as the fine-tuning overfitted in 10^{-5} even in Ideal case. Both our methods work in the two domains.

+extra in both tables indicates the use of extra

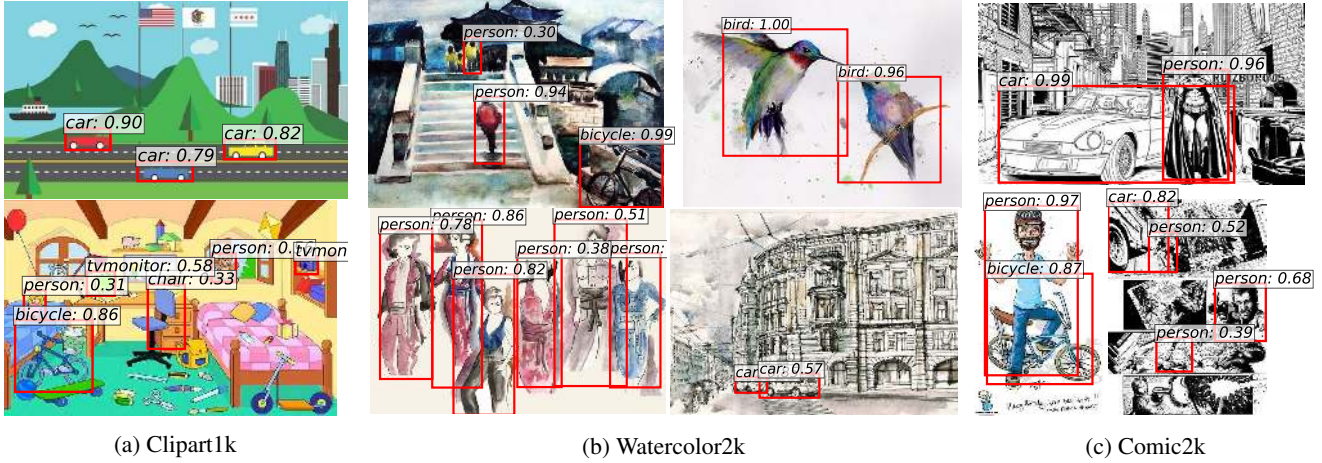


Figure 5: Example outputs for our DT+PA in the test set of each dataset. We only show windows whose scores are over 0.25 to maintain visibility.




Figure 6: Example images generated by DT.

BAM! images with raw noisy image-level labels of the target classes as described in Sec. 3.2. These images were pseudo-labeled and used for fine-tuning the FSD. With a substantial number of images, the training of `+extra` methods underwent 30000 iterations. The methods using extra noisy labels in BAM! significantly improved the detection performance and sometimes proved to be better than Ideal case trained on 1,000 clean instance-level annotations. Without any manual annotation, our framework can use large-scale images with noisy labels.

5.4. Qualitative Results

Fig. 6 shows the example images generated by DT. There was no mode collapse in the training of CycleGAN. Visibly, the perfect mapping is not accomplished in this experiment as the representation gap between a natural image domain and the other domains used in this paper is too wide as compared with the gap between synthetic and real images

tackled in recent studies, such as [29, 3]. CycleGAN seems to transfer color and texture while keeping most of the edges and semantics of the input image. The result of fine-tuning the FSD on these domain-transferred images in Table 2, Table 4, and Table 5 confirms the validity of our domain transfer method. Moreover, our methods are valid for various depiction styles as shown in Fig. 5. For more results, please refer to the supplementary material.

6. Discussion

In PL, only the top-1 bounding box for each class is employed. The other instances, if any, can be considered as negative samples. This issue is our future work. Moreover, if we could extract features with the same size corresponding to each detection, using the standard MIL paradigm in WSD such as [11, 18], we would improve the localization accuracy in pseudo-labeling, which is also our future work.

7. Conclusion

We proposed the novel task, cross-domain weakly supervised object detection, and the novel framework performing the two-step progressive domain adaptation to address this task. To evaluate our methods, we constructed original datasets comprising images with instance-level annotations in three visual domains. The results suggested that our methods were better than the other existing comparable methods and provided a simple but solid baseline.

Acknowledgments This work was partially supported by JST-CREST (JPMJCR1686) and Microsoft IJARC core13. N. Inoue is supported by GCL program of The Univ. of Tokyo by JSPS. R. Furuta is supported by the Grants-in-Aid for Scientific Research (16J07267) from JSPS.

References

- [1] H. Bilen, M. Pedersoli, and T. Tuytelaars. Weakly supervised object detection with convex clustering. In *CVPR*, 2015. 2
- [2] H. Bilen and A. Vedaldi. Weakly supervised deep detection networks. In *CVPR*, 2016. 2, 5, 6, 7
- [3] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, 2017. 8
- [4] L. Castrejon, Y. Aytar, C. Vondrick, H. Pirsiavash, and A. Torralba. Learning aligned cross-modal representations from weakly aligned data. In *CVPR*, 2016. 3
- [5] X. Chen and A. Gupta. Weakly supervised learning of convolutional networks. In *ICCV*, 2015. 2
- [6] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2), 2010. 2, 3, 5
- [7] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *TPAMI*, 32(9), 2010. 5
- [8] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(59), 2016. 3
- [9] R. Girshick. Fast R-CNN. In *ICCV*, 2015. 1, 2
- [10] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 2
- [11] R. Gokberk Cinbis, J. Verbeek, and C. Schmid. Multi-fold MIL training for weakly supervised object localization. In *CVPR*, 2014. 2, 8
- [12] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *JMLR*, 13(Mar), 2012. 3
- [13] J. Hoffman, D. Pathak, T. Darrell, and K. Saenko. Detector discovery in the wild: Joint multiple instance and representation learning. In *CVPR*, 2015. 2
- [14] D. Hoiem, Y. Chodpathumwan, and Q. Dai. Diagnosing error in object detectors. In *ECCV*, 2012. 6
- [15] V. Kantorov, M. Oquab, M. Cho, and I. Laptev. Context-LocNet: Context-aware deep network models for weakly supervised localization. In *ECCV*, 2016. 2, 5, 6, 7
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [17] I. Krasin, T. Duerig, N. Alldrin, V. Ferrari, S. Abu-El-Haija, A. Kuznetsova, H. Rom, J. Uijlings, S. Popov, A. Veit, S. Belongie, V. Gomes, A. Gupta, C. Sun, G. Chechik, D. Cai, Z. Feng, D. Narayanan, and K. Murphy. Openimages: A public dataset for large-scale multi-label and multi-class image classification. <https://github.com/openimages>, 2017. 2
- [18] D. Li, J.-B. Huang, Y. Li, S. Wang, and M.-H. Yang. Weakly supervised object localization with progressive domain adaptation. In *CVPR*, 2016. 2, 8
- [19] Y. Li, K. He, J. Sun, et al. R-FCN: Object detection via region-based fully convolutional networks. In *NIPS*, 2016. 1
- [20] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 1, 2
- [21] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2
- [22] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. SSD: Single shot multibox detector. In *ECCV*, 2016. 1, 2, 5
- [23] M. Long, Y. Cao, J. Wang, and M. I. Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 3
- [24] M. Long, H. Zhu, J. Wang, and M. I. Jordan. Unsupervised domain adaptation with residual transfer networks. In *NIPS*, 2016. 3
- [25] Y. Niitani, T. Ogawa, S. Saito, and M. Saito. ChainerCV: a library for deep learning in computer vision. In *ACM Multimedia*, 2017. 5
- [26] D. P. Papadopoulos, J. R. Uijlings, F. Keller, and V. Ferrari. Extreme clicking for efficient object annotation. In *ICCV*, 2017. 2
- [27] J. Redmon and A. Farhadi. YOLO9000: Better, Faster, Stronger. In *CVPR*, 2017. 1, 2, 6
- [28] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 1, 2, 6
- [29] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. In *CVPR*, 2017. 8
- [30] H. O. Song, R. B. Girshick, S. Jegelka, J. Mairal, Z. Harchaoui, T. Darrell, et al. On learning to localize objects with minimal supervision. In *ICML*, 2014. 2
- [31] H. O. Song, Y. J. Lee, S. Jegelka, and T. Darrell. Weakly-supervised discovery of visual pattern configurations. In *NIPS*, 2014. 2
- [32] H. Su, J. Deng, and L. Fei-Fei. Crowdsourcing annotations for visual object detection. In *AAAI workshop*, 2012. 2
- [33] P. Tang, X. Wang, X. Bai, and W. Liu. Multiple instance detection network with online instance classifier refinement. In *CVPR*, 2017. 2
- [34] S. Tokui, K. Oono, S. Hido, and J. Clayton. Chainer: a next-generation open source framework for deep learning. In *NIPS workshop*, 2015. 5
- [35] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. 3, 5, 6, 7
- [36] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders. Selective search for object recognition. *IJCV*, 104(2):154–171, 2013. 2
- [37] N. Westlake, H. Cai, and P. Hall. Detecting people in artwork with cnns. In *ECCV workshop*, 2016. 2
- [38] M. J. Wilber, C. Fang, H. Jin, A. Hertzmann, J. Collomosse, and S. Belongie. BAM! the behance artistic media dataset for recognition beyond photography. In *ICCV*, 2017. 2, 3
- [39] Q. Wu, H. Cai, and P. Hall. Learning graphs to model visual objects across different depictive styles. In *ECCV*, 2014. 2
- [40] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 2, 4