



Training Object Detectors from Few Weakly-Labeled and Many Unlabeled Images

Zhaohui Yang, Miaojing Shi, Yannis Avrithis, Chao Xu, Vittorio Ferrari

► To cite this version:

Zhaohui Yang, Miaojing Shi, Yannis Avrithis, Chao Xu, Vittorio Ferrari. Training Object Detectors from Few Weakly-Labeled and Many Unlabeled Images. 2019. hal-02393688

HAL Id: hal-02393688

<https://inria.hal.science/hal-02393688>

Preprint submitted on 4 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Training Object Detectors from Few Weakly-Labeled and Many Unlabeled Images

Zhaohui Yang¹, Miaoqing Shi², Yannis Avrithis², Chao Xu¹, Vittorio Ferrari³

¹ Key Laboratory of Machine Perception (Ministry of Education), Peking University.

² Univ Rennes, Inria, CNRS, IRISA. ³ Google Research.

zhaohuiyang@pku.edu.cn; miaoqing.shi@inria.fr; yannis@avrithis.net

xuchao@cis.pku.edu.cn; vittoferrari@google.com

Abstract

Weakly-supervised object detection attempts to limit the amount of supervision by dispensing the need for bounding boxes, but still assumes image-level labels on the entire training set are available. In this work, we study the problem of training an object detector from one or few clean images with image-level labels and a larger set of completely unlabeled images. This is an extreme case of semi-supervised learning where the labeled data are not enough to bootstrap the learning of a classifier or detector.

Our solution is to use a standard weakly-supervised pipeline to train a student model from image-level pseudo-labels generated on the unlabeled set by a teacher model, bootstrapped by region-level similarities to clean labeled images. By using the recent pipeline of PCL [47] and more unlabeled images, we achieve performance competitive or superior to many state of the art weakly-supervised detection solutions.

1. Introduction

The objective of visual object detection is to place a tight bounding box on every instance of an object class. With the advent of deep learning, recent methods [11, 35, 27, 34] have significantly boosted the detection performance. Most are fully supervised, using a large amount of data with carefully annotated bounding boxes. However, annotating bounding boxes is expensive. Humans can easily recognize a multitude of objects in images when only told once, without any spatial annotation. This ability is still a challenge for machine perception.

In terms of reducing the amount of supervision, the most common setting is *weakly-supervised object detection* (WSOD) [12, 4, 40, 48, 57]. In this setting, we are given a set of images known to contain instances of a certain class as specified by labels, but we do not know the object lo-

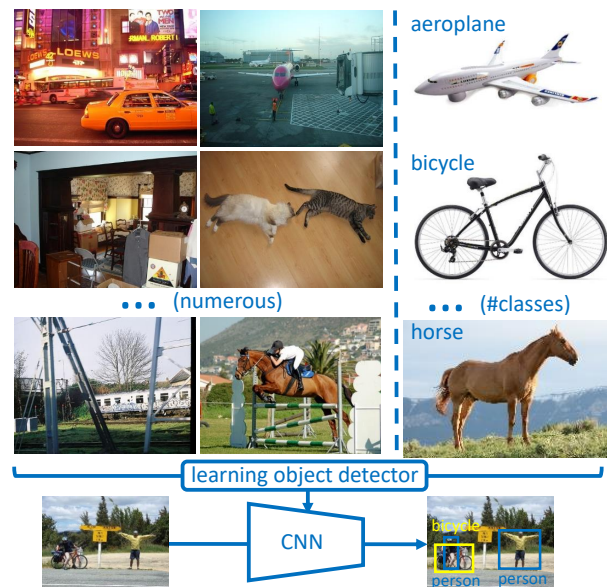


Figure 1. Given a set of completely unlabeled images (left) and very few or just one clean image per class with one image-level label per image and no other information (right), we learn an object detector. In practice, as clean images we just use the top-ranking results of web image search, given the class names.

cations in the form of bounding boxes or otherwise. Many works [3, 12, 40, 41] formulate weakly supervised object detection as *multiple instance learning* (MIL) [1], which has been extended to be learnable end-to-end [4, 48].

There are *mixed approaches* where a small number of images are annotated with bounding boxes and labels, and a large amount of images have only image-level labels [50, 18, 19, 57]. Many are related to transfer learning where some classes come with strong supervision (e.g. bounding boxes) while others with weak supervision (e.g. image-level labels). This is often referred as a *semi-supervised* setting [50, 57], but there is no consensus.

Semi-supervised learning [5] refers to using a small




Figure 2. Examples of web images. We show one image per class corresponding to the 20 object classes in PASCAL VOC detection task, *i.e.* aeroplane, bicycle, bird, . . . , train, tv monitor (from top left to bottom right).

amount of labeled data and a large amount of unlabeled data. It is traditionally studied for classification [56, 23, 33, 22], with one class label per image and no bounding boxes. In object detection, this would normally translate to a small number of images having labels and bounding boxes, and a large number of images having no annotation at all. This problem has been studied for the case where the fully annotated data (with bounding boxes) are enough to train a detector in the first place [6, 31], resulting in two-stage learning. But what if these data are very scarce?

In this work, we study what we consider to be the minimal supervision required for object detection: only one or few images per class with one image-level class label per image, and a large amount of images with no annotations at all (Fig. 1). This is an extreme case of semi-supervised learning. In practice, instead of randomly labeling few images per class, we use the top-ranking images from *web image search* with the class name as the query.

The motivation for using top-ranking images from web search is that they are clean, *i.e.* they mostly contain objects against a simple background and in a canonical pose and viewpoint, without clutter or occlusion (Fig. 2). This is true only for the top-ranking images: as the ranked list grows, severe noise appears [13]. We only provide one image-level label per image and we specify no bounding box or other information: these images are still *weakly labeled*.

Our solution, called *nano-supervised object detection* (NSOD), is conceptually simple. We begin by computing region-level class scores based on the similarity between the unlabeled images and the clean web images, which we then pool into image-level class probabilities. This yields image-level *pseudo-labels* on the entire unlabeled set, which we use to train a *teacher* model on a classification task. Then, by predicting new image-level pseudo-labels on the unlabeled set, we train a *student* model on a detection task using a weakly-supervised object detection pipeline.

Contributions. To our knowledge, we study for the first time the problem of training an object detector from few clean images with one image-level label per image and many images with no annotation at all. We introduce a new method for this problem that is simple, efficient (computational cost comparable to standard weakly supervised object detection), and modular (can build on any weakly-

supervised object detection pipeline). By using the recent pipeline of PCL [47] and more unlabeled images, we achieve performance competitive or superior to many state of the art weakly-supervised detection solutions. On PASCAL VOC 2007 test set for instance, we get mAP 42.0 using labels only for the few web images, compared to 43.5 of PCL using image-level labels on the entire training set.

2. Related Work

Weakly supervised object detection (WSOD). In this setting, all training images have image-level class labels. A major approach is to formulate it as multiple instance learning (MIL) [1], considering each training image as a “bag” and iteratively selecting high-scoring object proposals from each bag, treating them as ground truth to learn an object detector. MIL fails in large-scale datasets with poor initial solutions. Many works thus focus on providing robust initialization, using either hand crafted or CNN features [12, 46, 39].

Bilen and Vedaldi [4] propose *weakly-supervised deep detection network* (WSDDN), which pools region-level scores into image-level class probabilities and enables end-to-end learning from image-level labels. Building upon WSDDN, Zhang *et al.* [58] propose a *zigzag* strategy to discover reliable object instances and train them by *curriculum learning* [2]. Shen *et al.* [38] jointly learn weakly supervised detection and segmentation, using failures of one task to complement the other. Tang *et al.* [48] extend WSDDN to multiple instance detection network including *online instance classifier refinement* (OICR) and introduce a weakly-supervised region proposal network as a plugin [49]. Furthermore, in *proposal cluster learning* (PCL) [47], pre-clustering of object proposals followed by OICR accelerates learning and boosts performance. Our work builds on PCL, training it with image-level pseudo-labels.

Our problem is different from WSOD in that the few clean labeled images have no bounding boxes and the bulk of the training set is completely unlabeled.

Extra information/data. Besides improvements in the network architecture, there are attempts to incorporate additional cues into WSOD that are still weaker than bounding boxes, *e.g.* object size [40], count [10] or center points [29].

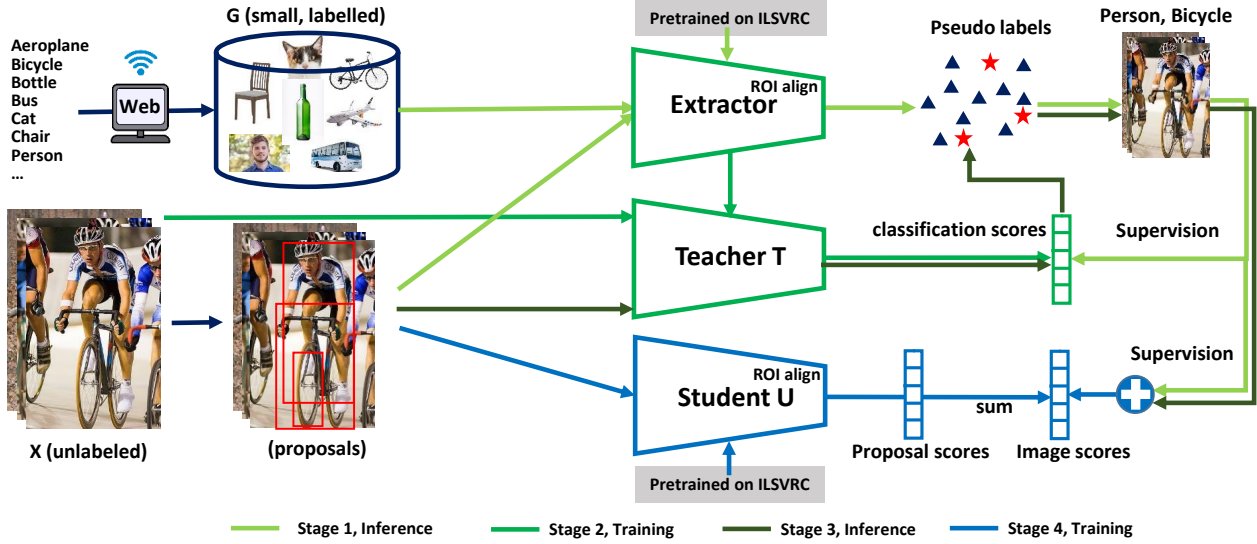


Figure 3. Overview of our *nano-supervised object detection* (NSOD) framework. We are given a large unlabeled benchmark set X and we collect the support set G of one or few clean images per class by web search, using class names as queries. Using the images in G and a feature extractor pre-trained on classification, we infer image-level class probabilities of images in X (stage 1). We then extract pseudo-labels on X and train a *teacher* network T on a C -way classification task (stage 2). T is used to classify each proposal of images in X , resulting in new image-level class probabilities (stage 3). We average these with the ones obtained in stage 1, based on G . Finally, we extract multi-class pseudo-labels on X and train a *student* network U on weakly-supervised detection by PCL [47] (stage 4).

It is also common to use extra data to *transfer knowledge* from a source domain and help localize objects in the target domain [36, 39, 53]. Large-scale weakly-supervised *web images* [13, 51] and *videos* [44, 26] with noisy labels are also common as extra data. In this work we only use a few weakly-labeled images. We do not even use the image-level labels from the standard WSOD training set.

Semi-supervised learning. There are several works that assume a few images are annotated with object bounding boxes and the rest still have image-level labels as in WSOD [50, 18, 19, 57]. These are often called *semi-supervised* [50, 57]. However, semi-supervised may also refer to the situation where some images are labeled (at image-level or with bounding boxes) and the rest have no annotation at all [6, 31]. This situation is consistent with the standard definition of *semi-supervised learning* [5].

Despite significant advances in deep semi-supervised learning [56, 23, 33, 17, 22, 52, 20], most such work focuses on classification tasks. This includes *pseudo-label* [23], where classifier predictions on unlabeled data are used as labels along with true labels on labeled data. The few exceptions focusing on detection [6, 31] still assume there are enough labeled images to learn an object detector in the first place, which is not the case in our work. Dong *et al.* [7] use few images with object bounding boxes and class labels along with many unlabeled images. However, this method relies on several models and iterative training, which is computationally expensive. In contrast, we develop a new efficient solution that allows us to use more

unlabeled images.

Few-shot learning is a related area where, after learning on a set of base classes with abundant examples, new tasks are given with only few examples of novel (unseen) classes [54, 45, 14]. In fact, few-shot learning has been applied to both object detection [21, 9] and semantic segmentation [32, 59, 28, 42]. However, we are also using a large training set of unlabeled images, while our target classes may not be unseen.

Knowledge distillation [16], initially introduced for transferring knowledge from an ensemble into a single model, has found several uses in problems ranging from incremental learning [25] and learning with noisy labels [24] to semi-supervised learning [31] and adversarial defenses [30]. In this work we transfer knowledge from one set to another and from one model to another, though we do not use *soft targets* like [16] but rather *hard targets* like [31].

3. Method

3.1. Overview

Problem. We are given a *benchmark* set of images X , with no annotations at all. Each image contains one or more objects belonging to one of C classes, along with background clutter. For each class, we search by its name through a web search engine to collect k images per class, where k is very small (typically less than 20). This results in a *support set* of images G and their associated image-level labels over C

classes. There is no bounding box or any other information in either X nor G . Using these data and a feature extractor ϕ pre-trained on classification, the problem is to learn a detector to recognize instances of the C classes and localize them with bounding boxes in new images.

Motivation. The problem relates to both weakly-supervised detection and semi-supervised classification. Similar to the former, we study the multiple instance learning problem but without image-level labels in the benchmark set. It also differs from the latter, at least in its common setting where thousands of examples are used [33, 22], in that G is too small to bootstrap the learning of a classifier or detector. In our setting, k can be as few as one example per class. For this reason, we propagate labels from G to X to initiate training.

Method overview. As shown in Fig. 3, we start by collecting G via web image search, (section 3.2); then infer image-level class probabilities on the benchmark set X (section 3.3). In particular, we extract object proposals [61] from images in X and compare region-level features extracted from a feature extractor ϕ against global features of images in G . We estimate class probabilities on X by propagating these similarities to the image level (stage 1 in Fig. 3). We infer pseudo-labels on X and train a *teacher* network T inherited from ϕ on a C -way classification task (stage 2, section 3.4). The fine-tuned T is used to classify regions in images of X , resulting in new image-level class probabilities (stage 3), which we average with the ones obtained in stage 1 on G . Finally, we extract multi-class pseudo-labels on X and train a *student* network U on a weakly-supervised detection task by PCL [47] (stage 4).

3.2. Web image collection (support set G)

An image in X may in general contain multiple objects and complex backgrounds, making it hard to mine localized object information from X alone. In contrast, we observe that the top-ranking images of web image search often contain a single instance taking up most of the image and with clean background (sometimes even fully uniform).

Given the object class names as text queries, we form the support set G by collecting the top- k results per class from Google Images. This process provides us with a reliable set of examples. We only focus on very few top-ranking images, which goes along with our setting of minimal supervision, as well as avoiding any noise in the search engine results (typical for lower-ranked images). Examples for $k = 1$ are shown in Fig. 2. Note that the returned images are not perfect, e.g. class *cow* comes with grass in the background, and class *person* is only depicting the face. This is however the best information we can get with the least effort/supervision.

3.3. Inferring class probabilities for benchmark X

Given the support set G and corresponding labels, the next step is to estimate image-level class probabilities for the benchmark set X . For each image \mathbf{x} in X , we use *edge boxes* [61] to extract a collection of R object proposals (regions). Ideally, we would like to have a label per region so we can train an object detector. Since the supervision in our case is very limited, it is not realistic to assign an accurate label per region based only on G . Inspired by the two-stream CNN architecture of WSDN [4], we rather infer class probabilities globally for image \mathbf{x} .

Similarity. We extract a feature vector $\phi(\mathbf{r})$ for each region \mathbf{r} of \mathbf{x} . We do the same for each image \mathbf{g} in G , extracting a feature vector $\phi(\mathbf{g})$. This is a global feature vector, as images in G are clean and object-centric. Let G_j be the support images labeled as class j , with $|G_j| = k$. Let also \mathbf{r}_i be the i -th region of \mathbf{x} . We define the $R \times C$ *similarity matrix* $S = \{s_{ij}\}$ with elements

$$s_{ij} := \frac{1}{k} \sum_{\mathbf{g} \in G_j} c(\phi(\mathbf{r}_i), \phi(\mathbf{g})), \quad (1)$$

where c denotes cosine similarity.

Voting. Following [4], we form $R \times C$ classification matrix $\sigma_{\text{cls}}(S)$ with each row being the softmax of the same row of S , implying competition over classes per region; similarly, we form $R \times C$ detection matrix $\sigma_{\text{det}}(S)$ with each column being the softmax of the same column of S , implying competition over regions per class:

$$\sigma_{\text{cls}}(S)_{ij} := \frac{e^{s_{ij}}}{\sum_{j=1}^C e^{s_{ij}}}, \quad \sigma_{\text{det}}(S)_{ij} := \frac{e^{s_{ij}}}{\sum_{i=1}^R e^{s_{ij}}}. \quad (2)$$

Here, the i -th row of $\sigma_{\text{cls}}(S)$ expresses a vector of class probabilities for proposal \mathbf{r}_i , while the j -th column of $\sigma_{\text{det}}(S)$ a vector of probabilities over regions (spatial distribution) of class j .

The final image-level class scores $\sigma(S)$ are obtained by element-wise product of $\sigma_{\text{cls}}(S)$ and $\sigma_{\text{det}}(S)$ followed by sum pooling over regions

$$\sigma(S)_j := \sum_{i=1}^R \sigma_{\text{cls}}(S)_{ij} \sigma_{\text{det}}(S)_{ij}. \quad (3)$$

Each score $\sigma(S)_j$ is in $[0, 1]$ and can be interpreted as the probability of object class j occurring in image \mathbf{x} .

Discussion. The above is a robust *voting strategy* which propagates proposal-level information to the image-level while suppressing noise. Formula (1) suggests that region \mathbf{r}_i will respond for class j if it is similar to any of the support images in G_j . While this response is noisy since it is only based on a few examples, it is only maintained if it is among the strongest over all classes and all regions in that image. Note that in [4], the softmax is applied on two separate data streams whereas we apply it to the *same matrix*.

3.4. Teacher and student training

Having class probability vectors (3) per image in X , a next step would be to convert them to multi-class pseudo-labels and train the student directly on a weakly-supervised detection task. Nevertheless, probabilities generated this way rely on the few support images in G for classification, while the object information in the benchmark set X is not exploited. Therefore we introduce a training strategy to *distill* knowledge from the support set G into the benchmark set X using the teacher, and then distill this knowledge from the teacher to the student.

Data distillation. We form the teacher T as the feature extractor network ϕ followed by a randomly initialized C -output fully-connected layer and softmax. Then we fine-tune T on a C -way classification task on X . The probabilities (3) are meant for multi-label classification (C independent binary classifiers), while here we are learning a single C -way classifier, *i.e.* for mutually exclusive labels. Given the class probability vector $\sigma(S)$ for each image \mathbf{x} in X , we take the most likely class $\arg \max_j \sigma(S)_j$ as a C -way *pseudo-label*. We fine-tune T on these pseudo-labels with a standard cross-entropy loss.

We have also tried several multi-label variants [55, 60], which are inferior to the simple C -way cross-entropy loss in our experiments. This may be attributed to the class sample distribution in X being unbalanced.

Knowledge distillation. The fine-tuned teacher T encodes object information of X into its network parameters. Directly using its image-level predictions on X would not be appropriate to train the student U for detection, because the latter would need multi-class labels. On the other hand, using it as feature extractor to repeat the process of Sec. 3.3 would not make much difference either, as it still produces class probabilities based on G . Instead, we use T to *directly classify object proposals* in X . Each proposal ideally contains one object, so it is particularly suitable to use T as it was designed: a C -way classifier.

Given an input image \mathbf{x} in X , we collect output class probabilities of T on each region \mathbf{r}_i of \mathbf{x} into a $R \times C$ matrix A with element a_{ij} being the probability of class j . From this matrix, it is possible to estimate new image-level class probabilities by $\sigma(A)$, similar to (3). Because it is based on T being trained on X as classifier, while $\sigma(S)$ (3) is based on G alone, we combine their strength by averaging both into a probability vector

$$\hat{\mathbf{q}} := \frac{1}{2} (\sigma(S) + \sigma(A)) \quad (4)$$

corresponding to \mathbf{x} .

An image-level *multi-class pseudo-label* $\hat{\mathbf{y}} \in \{0, 1\}^C$ is then obtained from $\hat{\mathbf{q}}$ by element-wise thresholding. An element $\hat{y}_j = 1$ specifies that an object of class j occurs in

image \mathbf{x} . In the absence of prior knowledge or validation data, we choose $\frac{1}{2}$ as threshold. Those image-level pseudo-labels are all that is needed to obtain an object detector if we use any weakly-supervised detection pipeline. In particular, we train the student model U on weakly-supervised detection on X using PCL [47].

4. Experiments

4.1. Experimental Setup

Datasets. We construct the support set G by collecting the top-ranking images per class from Google Image Search. We choose the standard object detection datasets PASCAL VOC 2007 and 2012 [8] for the benchmark sets. They have 20 object classes in total. Each contains a *trainval* set and a *test* set. For VOC 2007, the *trainval* set has 5011 images while *test* set 4952 images. For VOC 2012, the size of *trainval* and *test* sets are 11540 and 10991, respectively. Importantly, we do not use any labels, not even image-level labels. We use the *trainval* sets as X to train the object detector. We evaluate this detector on the *test* set.

Networks. We choose by default VGG16 [43] as our student U , which is consistent with most weakly-supervised detection methods [4, 48, 49, 38, 47]. As for teacher network T (feature extractor ϕ), since it is not used at inference time, we choose the more powerful ResNet-152 [15]. Both are pre-trained on the ILSVRC classification task [37].

Implementation details. We use $k = 1, 10, 20$ clean images per class in G , with $k = 20$ being the default. Following representative WSOD methods [4, 48, 47, 58], we adopt *edge boxes* [61] to extract on average 2000 proposals per image in X . For the default teacher model T , we first resize the input image to 256 pixels on the short side and then crop it to 224×224 . We set the batch size to 128 and the learning rate to 10^{-3} initially with cosine decay. For the default student model U , we feed the network with one image per batch. Training lasts for 50,000 iterations in total; learning rate starts at 10^{-5} and decays at 35,000 iterations by an order of magnitude.

Evaluation protocol. We evaluate the performance of our NSOD framework on both image classification and object detection. For image classification, we measure on the *trainval* set of X the *average prevision* (AP) and *mean AP* (mAP) for *multi-class predictions* [55, 60], as well as the accuracy of the top-1 class prediction per image. For object detection, we quantify localization performance on the *trainval* set by *CorLoc* [4, 40, 48, 58] and detection performance on the *test* set by mAP. At test time, the detector can localize multiple instances of the same class per image and this mAP metric is identical to what used to evaluate fully supervised object detectors.

METHOD	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	persn	plant	sheep	sofa	train	tv	mAP
NSOD _G	88.8	85.8	98.0	67.8	79.4	68.4	96.8	95.1	80.6	72.1	38.9	93.4	82.3	65.2	98.0	56.7	70.1	55.6	72.0	60.2	76.3
NSOD _X	86.4	96.9	97.1	71.4	98.5	67.1	89.9	95.1	80.0	66.8	36.5	92.9	74.2	62.9	96.9	53.1	59.9	58.8	70.1	78.9	76.7
NSOD	91.2	90.7	98.0	71.1	94.3	73.8	95.8	95.5	80.5	74.7	39.1	95.3	81.2	66.9	98.4	58.7	73.8	59.7	75.6	70.4	79.2
METHOD	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	persn	plant	sheep	sofa	train	tv	mAcc
NSOD _G	92.2	97.7	99.1	78.7	100.0	73.0	93.2	98.5	89.2	82.0	41.8	97.7	77.7	72.0	99.7	63.7	68.7	63.0	77.5	87.5	82.7
NSOD _X	93.1	93.4	98.2	79.2	100.0	78.9	96.3	96.7	84.0	83.5	45.1	95.7	84.2	72.9	98.5	73.3	77.2	66.1	83.3	87.7	84.3
NSOD	93.8	92.4	99.3	80.4	100.0	81.1	97.8	97.1	78.6	86.7	49.7	97.1	88.2	77.2	99.6	79.7	79.1	67.9	87.5	85.6	85.9

Table 1. Classification mAP for multi-class prediction (top) and classification mAcc for top-1 class prediction (bottom) on the *trainval* set of PASCAL VOC 2007. NSOD: our Nano-supervised object detection framework.

METHOD	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	persn	plant	sheep	sofa	train	tv	mAP
NSOD _G	57.2	52.7	36.0	14.1	11.0	50.6	46.9	35.8	5.7	47.1	16.1	52.8	34.3	54.4	14.8	11.4	29.0	48.8	43.4	13.9	33.9
NSOD _X	58.5	51.5	37.5	11.6	10.6	55.3	48.2	40.4	5.8	49.9	16.0	51.3	31.6	56.3	14.6	9.0	34.3	45.5	42.2	20.3	34.5
NSOD	57.9	59.7	43.2	10.5	13.1	62.7	58.6	43.9	10.6	51.1	25.7	49.8	39.3	60.6	14.9	10.9	33.5	45.2	42.5	27.8	38.0
NSOD (VOC07+12)	51.5	65.2	48.9	13.2	19.7	64.8	59.3	55.5	12.4	59.3	24.3	54.1	47.4	62.8	20.7	15.0	39.5	51.3	53.8	21.4	42.0
NSOD ($k = 1$)	53.0	58.0	24.4	13.3	11.3	41.3	43.8	43.6	2.3	50.3	6.1	32.4	19.0	50.5	15.0	8.7	35.7	41.7	42.8	6.2	30.0
NSOD ($k = 10$)	57.2	27.8	40.4	9.7	11.2	61.2	57.0	25.9	13.4	47.2	6.2	45.5	35.7	53.0	21.2	14.1	34.8	43.7	39.8	19.8	33.2
NSOD ($k = 20$)	57.9	59.7	43.2	10.5	13.1	62.7	58.6	43.9	10.6	51.1	25.7	49.8	39.3	60.6	14.9	10.9	33.5	45.2	42.5	27.8	38.0
NSMT ($z = 300$)	49.6	33.9	29.6	15.5	9.5	47.9	32.9	49.1	0.2	13.2	21.1	34.4	19.7	31.5	9.6	9.9	35.6	43.1	38.9	15.0	27.0
WSDN [4]	39.4	50.1	31.5	16.3	12.6	64.5	42.8	42.6	10.1	35.7	24.9	38.2	34.4	55.6	9.4	14.7	30.2	40.7	54.7	46.9	34.8
OICR [48]	58.0	62.4	31.1	19.4	13.0	65.1	62.2	28.4	24.8	44.7	30.6	25.3	37.8	65.5	15.7	24.1	41.7	46.9	64.3	62.6	41.2
WSRPN [49]	57.9	70.5	37.8	5.7	21.0	66.1	69.2	59.4	3.4	57.1	57.3	35.2	64.2	68.6	32.8	28.6	50.8	49.5	41.1	30.0	45.3
PCL [47]	54.4	69.0	39.3	19.2	15.7	62.9	64.4	30.0	25.1	52.5	44.4	19.6	39.3	67.7	17.8	22.9	46.6	57.5	58.6	63.0	43.5
WS-JDS [38]	52.0	64.5	45.5	26.7	27.9	60.5	47.8	59.7	13.0	50.4	46.4	56.3	49.6	60.7	25.4	28.2	50.0	51.4	66.5	29.7	45.6

Table 2. Detection mAP on the *test* set of PASCAL VOC 2007. NSOD: our Nano-supervised object detection framework. NSMT: Nano-supervised mean teacher. Unless otherwise stated, NSOD is using by default $k = 20$ support images per class. All compared methods [4, 48, 49, 47, 38] utilize the image-level labels in the benchmark set X while NSOD and NSMT do not.

4.2. PASCAL VOC 2007

Ablation study. We begin by conducting an ablation study on our labeling strategy, support and benchmark set size.

Labeling strategy (classification). Referring to (3) and (4), we first ablate combining $\sigma(S)$ and $\sigma(A)$ to generate image-level pseudo-labels. $\sigma(S)$ is computed based on G alone, while $\sigma(A)$ is computed based on the teacher model trained on X . We apply a hard threshold of $\frac{1}{2}$ on the predicted class probabilities of $\sigma(S)$ and $\sigma(A)$ to generate two sets of image-level pseudo-labels. We train two different student models separately on the two sets of pseudo-labels, which we denote as NSOD_G and NSOD_X respectively.

The classification accuracy of the two sets of pseudo-labels is first evaluated on the *trainval* set of VOC 2007 and shown in Table 1. It can be seen that NSOD_G and NSOD_X produce a similar classification mAP of 76.3 vs. 76.7, while the AP on individual classes differs. However, in terms of top-1 class accuracy, NSOD_X is clearly better than NSOD_G. This is reasonable as NSOD_X is fine-tuned as a C -way classifier, which takes the top-1 class predictions of $\sigma(S)$ as pseudo-labels. The two sets of pseudo-labels are clearly complementary by averaging $\sigma(S)$ and $\sigma(A)$ according to (4), denoted by NSOD. This improves both multi-class and top-1 class predictions, reaching the highest scores of 79.2 and 85.9, respectively.

Labeling strategy (detection). To further investigate the

complementary effect of NSOD_G and NSOD_X, we evaluate their detection result on the *test* set of VOC 2007 (Table 2). The mAP of NSOD_G (33.9) is slightly greater than that of NSOD_X (34.5). Their combination (our full method NSOD) further increases mAP by +3.5% to 38.0. The detection result on the *test* set is consistent with the classification result on the *trainval* set, which validates our idea of distilling knowledge from the support set to the benchmark set and from the teacher to the student model.

Support set size. We evaluate performance for different number k of images per class of the support set G in Table 2: mAP is 30.0 for $k = 1$, 33.2 for $k = 10$ and 38.0 for $k = 20$. Further increasing to $k = 30$ hardly improves, presumably bringing more noisy rather than diverse examples. How to deal with large-scale noisy web images/videos is an open problem [13, 51, 44, 26]. We keep G small to avoid bringing too many noisy images, while at the same time using the unlabeled benchmark set X for more diversity.

Benchmark set size. Since X is unlabeled and our method is computationally efficient, we can easily improve performance by simply using more unlabeled data. As shown in Table 2, NSOD on the union of VOC 2007 and VOC 2012 (VOC07+12) yields a mAP of 42%, which is an improvement by +4% over using VOC 2007 alone. Since neither set is labeled, this improvement comes at almost no cost.

Comparison to weakly-supervised methods. We compare to several representative weakly-supervised object de-

METHOD	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	persn	plant	sheep	sofa	train	tv	mAP
WSDDN [4]	65.1	58.8	58.5	33.1	39.8	68.3	60.2	59.6	34.8	64.5	30.5	43.0	56.8	82.4	25.5	41.6	61.5	55.9	65.9	63.7	53.5
OICR [48]	81.7	80.4	48.7	49.5	32.8	81.7	85.4	40.1	40.6	79.5	35.7	33.7	60.5	88.8	21.8	57.9	76.3	59.9	75.3	81.4	60.6
WSRPN [49]	77.5	81.2	55.3	19.7	44.3	80.2	86.6	69.5	10.1	87.7	68.4	52.1	84.4	91.6	57.4	63.4	77.3	58.1	57.0	53.8	63.8
PCL [47]	79.6	85.5	62.2	47.9	37.0	83.8	83.4	43.0	38.3	80.1	50.6	30.9	57.8	90.8	27.0	58.2	75.3	68.5	75.7	78.9	62.7
WS-JDS [38]	82.9	74.0	73.4	47.1	60.9	80.4	77.5	78.8	18.6	70.0	56.7	67.0	64.5	84.0	47.0	50.1	71.9	57.6	83.3	43.5	64.5
NSOD	80.0	73.3	66.1	34.0	29.0	72.6	76.5	56.4	17.7	74.7	47.5	61.4	60.5	86.4	31.9	36.6	60.8	59.1	57.4	49.1	56.6
NSOD (VOC07+12)	78.3	78.4	70.3	34.0	34.0	75.1	76.6	66.9	24.8	76.0	45.6	69.8	67.7	88.8	34.4	41.4	67.0	62.1	67.3	40.9	60.0

Table 3. Detection CorLoc on the *trainval* set of PASCAL VOC 2007. All compared methods [4, 48, 49, 47, 38] utilize the image-level labels in the benchmark set X , whereas our NSOD does not.

METHOD	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	persn	plant	sheep	sofa	train	tv	mAP
OICR [48]	67.7	61.2	41.5	25.6	22.2	54.6	49.7	25.4	19.9	47.0	18.1	26.0	38.9	67.7	2.0	22.6	41.1	34.3	37.9	55.3	37.9
ZLDN [58]	54.3	63.7	43.1	16.9	21.5	57.8	60.4	50.9	1.2	51.5	44.4	36.6	63.6	59.3	12.8	25.6	47.8	47.2	48.9	50.6	42.9
PCL [47]	58.2	66.0	41.8	24.8	27.2	55.7	55.2	28.5	16.6	51.0	17.5	28.6	49.7	70.5	7.1	25.7	47.5	36.6	44.1	59.2	40.6
NSOD	56.3	27.6	42.2	10.9	23.8	55.1	46.2	36.6	5.6	51.8	15.5	55.9	54.0	63.6	23.5	10.8	43.1	39.2	49.0	21.5	36.6
NSOD (VOC07+12)	57.3	50.7	49.2	11.3	21.2	56.8	46.4	55.0	6.6	52.7	12.8	61.8	45.8	64.7	18.9	10.5	34.9	41.0	48.1	19.9	38.3

Table 4. Detection mAP on *test* set of PASCAL VOC 2012. All compared methods [48, 47, 58] utilize the image-level labels in the benchmark set X while our NSOD does not.

METHOD	CorLoc
OICR [48]	62.1
ZLDN [58]	61.5
PCL [47]	63.2
NSOD	58.7
NSOD (VOC07+12)	59.9

Table 5. Detection CorLoc on the *trainval* set of PASCAL VOC 2012. All compared methods [48, 47, 58] utilize the image-level labels in the benchmark set while our NSOD does not.

tection methods [4, 48, 49, 47, 38] in Table 2. For a fair comparison, all these methods use the same VGG16 backbone as we do, without bells and whistles. Recall that NSOD requires no annotation on the benchmark set X , while weakly-supervised methods assume image-level labels for all images in X .

One directly competing method is PCL trained on ground truth image-level labels in X . Despite using no annotation on X , NSOD achieves an mAP that is only 5.5% below that of PCL (38.0 vs. 43.5). Most importantly, our method naturally improves simply by using more unlabeled data for X : on VOC07+12, NSOD is only 1.5% below PCL (42.0 vs. 43.5). This result is already competitive with state-of-the-art weakly-supervised methods, and outperforms WSDDN [4] and OICR [48] when trained on VOC 2007 with image-level labels. This is a strong result that confirms the value of our core contribution.

There are also some WSOD works employing large-scale web images/videos as extra data. For instance, [51] and [44] build their works on the WSDDN pipeline [4] and produce mAP 36.8 and 39.4 on VOC 2007, respectively. Unlike these works, our NSOD uses few web images, an unlabeled benchmark set, and an advanced WSOD pipeline.

Importantly, NSOD also delivers higher mAP.

In Table 3 we report CorLoc on the *trainval* set of VOC 2007, as is common for weakly-supervised detection methods [4, 40, 48, 58]. Notice that the CorLoc of NSOD on VOC07+12 is only 2.7% below PCL (60.0 vs. 62.7) and generally among the best-performing WSOD methods (e.g. OICR has 60.6). This is consistent with the mAP measurements on the *test* set. Finally, Fig. 4 gives some examples of detection results of NSOD on PASCAL VOC 2007.

Comparison to an adapted semi-supervised method [52]. We adapt the *mean teacher* [52] semi-supervised classification method to our setting. Using image-level class probabilities \hat{q} (4), we select the top- z scored images as positive for each class and the rest we treat as negative. With those pseudo-labels, we train PCL on VGG16, applying the consistency loss of [52] to image-level predictions on X . We call this approach *nano-supervised mean teacher* (NSMT).

We choose $z = 300$ as it works the best in practice. NSMT then yields an mAP of 27.0 as shown in Table 2. This result is lower than our NSOD by 11.0% (27.0 vs. 38.0). This suggests it is not straightforward to transfer a successful semi-supervised approach from the classification to the detection task.

We have also tried to directly infer object bounding boxes on the *test* set of X using naive approaches like k -NN classification on regions directly. Those fail, producing mAP lower than 10. We should emphasize the importance of propagating similarity scores from region-level to image-level as we do in NSOD.

4.3. PASCAL VOC 2012

In addition to VOC 2007, it is interesting to report results on VOC 2012 as well. Using the same support set G , we first train an object detector with our NSOD on VOC 2012 alone. The mAP is reported on the *test* set of VOC 2012




Figure 4. Detection results of our NSOD on PASCAL VOC 2007, using default settings ($k = 20$). Failure cases are illustrated in the last row with white boxes.

and compared to representative weakly-supervised object detection methods [48, 47, 58] in Table 4. Despite not using any VOC 2012 labels, NSOD is only 4.0% below PCL (36.6 vs. 40.6). Furthermore, we report the detection result by training a detector with NSOD on VOC 2012 and 2007 combined (always unlabeled). The mAP increases to 38.3, now outperforming OICR. The AP of our NSOD on some individual classes is particularly competitive.

Finally, CorLoc on the *trainval* set of VOC 2012 is reported in Table 5. Again, NSOD on unlabeled VOC07+12 is very competitive with weakly supervised methods trained on VOC 2012 with image-level labels.

5. Discussion

We have studied for the first time a particular setting for object detection that uses minimal supervision: a few clean images per class with only image-level labels, plus a large set of completely unlabeled images. This is an extreme case of semi-supervised learning for object detection, where the labeled images are too few to initiate detector training.

Our nano-supervised object detection framework (NSOD) basically begins with a combination of *few-shot* and *semi-supervised* classification. The former is using the clean images as *class prototypes* [45] to estimate class probabilities per region, which are propagated at image level using the *voting process* of WSDN [4]. The latter

is generating *pseudo-labels* on the unlabeled set from these probabilities to train a classifier, as is standard in semi-supervised classification [23].

By using the PCL pipeline [47] and extending the unlabeled set to both VOC 2007 and VOC 2012, our NSOD achieves detection mAP very close to PCL itself trained on VOC 2007 with image-level labels. Moreover, our result is already competitive or superior to many state of the art weakly-supervised detection solutions. It is natural to expect further improvement when applying our method to very large unlabeled collections. Moreover, since NSOD produces image-level pseudo-labels that can be used to train any weakly-supervised detection network, it is reasonable to expect further improvement when feeding those to more advanced WSOD methods such as [38].

References

- [1] S. Andrews, I. Tsochantaridis, and T. Hofmann. Support vector machines for multiple-instance learning. In *NIPS*, 2003. 1, 2
- [2] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *ICML*, 2009. 2
- [3] H. Bilen, M. Pedersoli, and T. Tuytelaars. Weakly supervised object detection with posterior regularization. In *BMVC*, 2014. 1
- [4] H. Bilen and A. Vedaldi. Weakly supervised deep detection networks. In *CVPR*, 2016. 1, 2, 4, 5, 6, 7, 8
- [5] O. Chapelle, B. Schölkopf, and A. Zien. *Semi-Supervised Learning*. MIT Press, 2006. 1, 3
- [6] M.-K. Choi, J. Park, J. Jung, H. Jung, J.-H. Lee, W. J. Won, W. Y. Jung, J. Kim, and S. Kwon. Co-occurrence matrix analysis-based semi-supervised training for object detection. *arXiv preprint arXiv:1802.06964*, 2018. 2, 3
- [7] X. Dong, L. Zheng, F. Ma, Y. Yang, and D. Meng. Few-example object detection with model communication. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1641–1654, 2018. 3
- [8] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 5
- [9] Q. Fan, W. Zhuo, and Y.-W. Tai. Few-shot object detection with attention-RPN and multi-relation detector. *arXiv preprint arXiv:1908.01998*, 2019. 3
- [10] M. Gao, A. Li, R. Yu, V. I. Morariu, and L. S. Davis. C-wsl: Count-guided weakly supervised localization. In *ECCV*, 2018. 2
- [11] R. Girshick. Fast r-cnn. In *ICCV*, 2015. 1
- [12] R. Gokberk Cinbis, J. Verbeek, and C. Schmid. Multi-fold mil training for weakly supervised object localization. In *CVPR*, 2014. 1, 2
- [13] S. Guo, W. Huang, H. Zhang, C. Zhuang, D. Dong, M. R. Scott, and D. Huang. Curriculumnet: Weakly supervised learning from large-scale web images. In *ECCV*, 2018. 2, 3, 6
- [14] B. Hariharan and R. B. Girshick. Low-shot visual recognition by shrinking and hallucinating features. *ICCV*, 2017. 3
- [15] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [16] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [17] E. Hoffer and N. Ailon. Semi-supervised deep learning by metric embedding. *arXiv preprint arXiv:1611.01449*, 2016. 3
- [18] J. Hoffman, S. Guadarrama, E. S. Tzeng, R. Hu, J. Donahue, R. Girshick, T. Darrell, and K. Saenko. Lsda: Large scale detection through adaptation. In *NIPS*, 2014. 1, 3
- [19] J. Hoffman, D. Pathak, T. Darrell, and K. Saenko. Detector discovery in the wild: Joint multiple instance and representation learning. In *CVPR*, 2015. 1, 3
- [20] A. Iscen, G. Tolias, Y. Avrithis, and O. Chum. Label propagation for deep semi-supervised learning. In *CVPR*, pages 5070–5079, 2019. 3
- [21] B. Kang, Z. Liu, X. Wang, F. Yu, J. Feng, and T. Darrell. Few-shot object detection via feature reweighting. In *ICCV*, October 2019. 3
- [22] S. Laine and T. Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2, 3, 4
- [23] D.-H. Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICMLW*, 2013. 2, 3, 8
- [24] Y. Li, J. Yang, Y. Song, L. Cao, J. Luo, and J. Li. Learning from noisy labels with distillation. *arXiv preprint arXiv:1703.02391*, 2017. 3
- [25] Z. Li and D. Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, Dec 2018. 3
- [26] X. Liang, S. Liu, Y. Wei, L. Liu, L. Lin, and S. Yan. Towards computational baby learning: A weakly-supervised approach for object detection. In *CVPR*, 2015. 3, 6
- [27] 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
- [28] K. Nguyen and S. Todorovic. Feature weighting and boosting for few-shot segmentation. In *ICCV*, October 2019. 3
- [29] D. P. Papadopoulos, J. R. Uijlings, F. Keller, and V. Ferrari. Training object class detectors with click supervision. In *CVPR*, pages 6374–6383, 2017. 2
- [30] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami. Distillation as a defense to adversarial perturbations against deep neural networks. *arXiv preprint arXiv:1511.04508*, 2015. 3
- [31] I. Radosavovic, P. Dollar, R. Girshick, G. Gkioxari, and K. He. Data distillation: Towards omni-supervised learning. In *CVPR*, June 2018. 2, 3
- [32] K. Rakelly, E. Shelhamer, T. Darrell, A. A. Efros, and S. Levine. Few-shot segmentation propagation with guided networks. *arXiv preprint arXiv:1806.07373*, 2018. 3
- [33] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko. Semi-supervised learning with ladder networks. In *NIPS*, 2015. 2, 3, 4
- [34] J. Redmon and A. Farhadi. Yolo9000: better, faster, stronger. In *CVPR*, 2017. 1
- [35] 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
- [36] M. Ročan and Y. Wang. Weakly supervised localization of novel objects using appearance transfer. In *CVPR*, pages 4315–4324, 2015. 3
- [37] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein. Imagenet large scale visual recognition challenge. *arXiv preprint arXiv:1409.0575*, 2014. 5
- [38] Y. Shen, R. Ji, Y. Wang, Y. Wu, and L. Cao. Cyclic guidance for weakly supervised joint detection and segmentation. In *CVPR*, 2019. 2, 5, 6, 7, 8
- [39] M. Shi, H. Caesar, and V. Ferrari. Weakly supervised object localization using things and stuff transfer. In *ICCV*, 2017. 2, 3

- [40] M. Shi and V. Ferrari. Weakly supervised object localization using size estimates. In *ECCV*, 2016. 1, 2, 5, 7
- [41] Z. Shi, P. Siva, and T. Xiang. Transfer learning by ranking for weakly supervised object annotation. *arXiv preprint arXiv:1705.00873*, 2017. 1
- [42] M. Siam, B. N. Oreshkin, and M. Jagersand. AMP: Adaptive masked proxies for few-shot segmentation. In *ICCV*, October 2019. 3
- [43] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [44] K. K. Singh and Y. J. Lee. You reap what you sow: Using videos to generate high precision object proposals for weakly-supervised object detection. In *CVPR*, 2019. 3, 6, 7
- [45] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017. 3, 8
- [46] H. O. Song, Y. J. Lee, S. Jegelka, and T. Darrell. Weakly-supervised discovery of visual pattern configurations. In *NIPS*, 2014. 2
- [47] P. Tang, X. Wang, S. Bai, W. Shen, X. Bai, W. Liu, and A. L. Yuille. Pcl: Proposal cluster learning for weakly supervised object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 1, 2, 3, 4, 5, 6, 7, 8
- [48] P. Tang, X. Wang, X. Bai, and W. Liu. Multiple instance detection network with online instance classifier refinement. In *CVPR*, 2017. 1, 2, 5, 6, 7, 8
- [49] P. Tang, X. Wang, A. Wang, Y. Yan, W. Liu, J. Huang, and A. Yuille. Weakly supervised region proposal network and object detection. In *ECCV*, 2018. 2, 5, 6, 7
- [50] Y. Tang, J. Wang, B. Gao, E. Dellandréa, R. Gaizauskas, and L. Chen. Large scale semi-supervised object detection using visual and semantic knowledge transfer. In *CVPR*, 2016. 1, 3
- [51] Q. Tao, H. Yang, and J. Cai. Exploiting web images for weakly supervised object detection. *IEEE Transactions on Multimedia*, 2018. 3, 6, 7
- [52] A. Tarvainen and H. Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, 2017. 3, 7
- [53] J. Uijlings, S. Popov, and V. Ferrari. Revisiting knowledge transfer for training object class detectors. In *CVPR*, pages 1101–1110, 2018. 3
- [54] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al. Matching networks for one shot learning. In *NIPS*, 2016. 3
- [55] Y. Wei, W. Xia, M. Lin, J. Huang, B. Ni, J. Dong, Y. Zhao, and S. Yan. Hcp: A flexible cnn framework for multi-label image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9):1901–1907, 2016. 5
- [56] J. Weston, F. Ratle, and R. Collobert. Deep learning via semi-supervised embedding. In *ICML*, pages 1168–1175, 2008. 2, 3
- [57] Z. Yan, J. Liang, W. Pan, J. Li, and C. Zhang. Weakly-and semi-supervised object detection with expectation-maximization algorithm. *arXiv preprint arXiv:1702.08740*, 2017. 1, 3
- [58] X. Zhang, J. Feng, H. Xiong, and Q. Tian. Zigzag learning for weakly supervised object detection. In *CVPR*, 2018. 2, 5, 7, 8
- [59] X. Zhang, Y. Wei, Y. Yang, and T. Huang. Sg-one: Similarity guidance network for one-shot semantic segmentation. *arXiv preprint arXiv:1810.09091*, 2018. 3
- [60] F. Zhu, H. Li, W. Ouyang, N. Yu, and X. Wang. Learning spatial regularization with image-level supervisions for multi-label image classification. In *CVPR*, 2017. 5
- [61] C. L. Zitnick and P. Dollár. Edge boxes: Locating object proposals from edges. In *ECCV*, 2014. 4, 5