

Improving Object Detection with *Selective Self-Supervised Self-Training*

Yandong Li^{1,2,*[0000-0003-2448-1294]}, Di Huang^{2[0000-0001-6021-5175]},
Danfeng Qin^{2[0000-0002-7756-2074]}, Liqiang Wang^{1[0000-0002-1265-4656]}, and
Boqing Gong^{2[0000-0003-3915-5977]}

¹University of Central Florida, ²Google Inc
lyndon.leeseu@outlook.com, lwang@cs.ucf.edu,
{dihuang,qind,bgong}@google.com

Abstract. We study how to leverage Web images to augment human-curated object detection datasets. Our approach is two-pronged. On the one hand, we retrieve Web images by image-to-image search, which incurs less domain shift from the curated data than other search methods. The Web images are diverse, supplying a wide variety of object poses, appearances, their interactions with the context, etc. On the other hand, we propose a novel learning method motivated by two parallel lines of work that explore unlabeled data for image classification: self-training and self-supervised learning. They fail to improve object detectors in their vanilla forms due to the domain gap between the Web images and curated datasets. To tackle this challenge, we propose a selective net to rectify the supervision signals in Web images. It not only identifies positive bounding boxes but also creates a safe zone for mining hard negative boxes. We report state-of-the-art results on detecting backpacks and chairs from everyday scenes, along with other challenging object classes.

1 Introduction

Object detection is a fundamental task in computer vision. It has achieved unprecedented performance for many objects, partially thanks to the recently developed deep neural detectors [48,32,47,53]. Some detectors have made their way into real-world applications, such as smart mobile phones and self-driving cars.

However, upon a careful investigation into the top three teams’ class-wise detection results on 80 common objects in context (COCO) [34], we find that they still fall short of detecting backpacks, handbags, and chairs, among other *functional* objects. As of the paper submission, they report an average precision [34] of less than 0.30 on detecting backpacks and less than 0.40 on chairs.

These man-made objects are defined by their functionalities more than visual appearances, leading to high intra-class variation. Minsky [38] wrote that “there’s little we can find in common to all chairs except for their intended use.”

* This work was done while Yandong Li was an intern at Google Inc.



Fig. 1. Chairs (top) and backpacks (bottom) are difficult to detect in daily scenes. (Images and groundtruth boxes (red) are from COCO [34], and the blue boxes are predicted by the best detectors in our experiments.)

Grabner *et al.* [20] quantitatively evaluated the challenges of detecting functional objects like chairs. Another potential reason that contributes to the low performance in detecting backpacks and chairs is that they are too common to draw photographers' attention. Consequently, they often sit out of the camera focus, appear small, and become occluded in context (cf. Figure 1).

How to improve the detection of backpacks, chairs, and other common, “less eye-catching” functional objects? We believe the answer resides on both the quality of training data and the inductive bias of advanced detectors. In this paper, we focus on the data aspect and mainly study the potential of *unlabeled* Web images for improving the detection results of backpacks and chairs, without heavily taxing human raters.

In other words, we study how to leverage the *unlabeled* Web images to augment human-curated object detection datasets. Web images are diverse and massive, supplying a wide variety of object poses, appearances, their interactions with the context, etc., which may lack in the curated object detection datasets. However, the Web images are *out of the distribution* of the curated datasets. The domain gap between them calls for a careful design of methods to effectively take advantage of the signals in the Web data.

Our approach is two-pronged. On the one hand, we retrieve a big pool of candidate Web images via Google Image (<https://images.google.com>) by using its image-to-image search. The query set consists of all training images in the original human-curated dataset. Compared with text-based search, which mainly returns iconic photos, the image-based search gives rise to more natural images with diverse scenes, schematically reducing the domain mismatch between the retrieved images and the original datasets.

On the other hand, we propose a novel learning method to utilize the Web images for object detection inspired by self-training [50,66] and self-supervised learning [42,11,21,64,6,18], both of which are popular in semi-supervised learning. Our problem is similar to semi-supervised learning, but there exists a domain gap between the Web images and the curated datasets. We find that the domain

gap fails both self-training and self-supervised learning in their vanilla forms because the out-of-domain Web images give rise to many inaccurate candidate boxes and uncalibrated box classification scores. To tackle these challenges, we propose a selective net to identify high-quality positive boxes and a safe zone for mining hard negative boxes [33,48]. It rectifies the supervision signals in Web images, enabling self-training and self-supervised learning to improve neural object detectors by leveraging the Web images.

The main contributions in the paper are as follows. First, we customize self-training for the object detection task by a selective net, which identifies positive bounding boxes and assigns some negative boxes to a safe zone to avoid messing up the hard negative mining in the training of object detectors. Second, we improve the consistency-based [28,62,9] semi-supervised object detection [26] by the selective net under our self-training framework. Third, to the best of our knowledge, this work is the first to explore *unlabeled, out-of-domain* Web images to augment curated object detection datasets. We report state-of-the-art results for detecting backpacks and chairs, along with other challenging objects.

2 Augmenting COCO detection with Web images

We augment the training set of COCO detection [34] by retrieving relevant Web images through Google Image (<https://images.google.com>). We focus on the backpack and chair classes in this paper. They represent non-rigid and rigid man-made objects, respectively, and the existing results of detecting them are still unsatisfactory (less than 0.40 *AP* on COCO as of March 5th, 2020).

COCO-backpack, COCO-chair, Web-backpack, and Web-chair. COCO is a widely used dataset for object detection, which contains 118k training images and 5k validation images [34]. Out of them, there are 8,714 backpacks in 5,528 training images. We name these images the COCO-backpack query set. Similarly, we have a COCO-chair query set that contains 12,774 images with 38,073 chair instances. Using the images in COCO-backpack and COCO-chair to query Google Image, we obtain 70,438 and 186,192 unlabeled Web images named Web-backpack and Web-chair, respectively. We have removed the Web images that are nearly duplicate with any image in the COCO training and validation sets. Figure 2 shows two query images and the retrieved Web images.

Labeling Web-backpack¹. To facilitate the evaluation of our approach and future research, we label a subset of the Web-backpack images. This subset contains 16,128 images and is selected as follows. We apply a pre-trained R101-FPN object detector [63,32] to all Web-backpack images and then keep the ones that contain at least one backpack box detected with the confidence score higher than 0.7. We ask three raters to label each survived images. One rater draws bounding boxes over all backpacks in an image. The other two examine the results sequentially. They modify the boxes if they find any problem with the



Fig. 2. Examples of the top three retrieved images from image-to-image search.



Fig. 3. Labeling errors in COCO-backpack. Two images in one group mean that they contain conflict annotations.

previous rater’s annotation. Please see the supplementary materials for more details of the annotation procedure, including a full annotation instruction we provided to the raters.

Relabeling COCO-backpack¹. Using the same annotation procedure above, we also relabel the backpacks in COCO training and validation sets. The main reason for the relabeling is to mitigate the annotation mismatch between Web-backpack and COCO-backpack caused by different annotation protocols. Another reason is that we observe inconsistent bounding boxes in the existing COCO detection dataset. As Figure 3 shows, some raters label the sling bags, while others do not, and some raters enclose the straps in the bounding boxes and others do not. We still ask three raters to label the COCO-backpack training images. For the validation set, we tighten the quality control and ask five raters to screen each image.

¹ We use [Google Cloud Data Labeling Service](#) for all the labeling work.

Table 1. Statistics of the Web images in this paper and their counterparts in COCO

Dataset	# Images	# Boxes by COCO	# Boxes by us	# Raters
COCO-backpack	5,528	8,714	7,170	3
COCO-backpack validation	5,000	371	436	5
Web-backpack	70,438	–	–	–
Web-backpack labeled	16,128	–	23,683	3
COCO-chair	12,774	38,073	–	–
COCO-chair validation	5,000	1,771	–	–
Web-chair	186,192	–	–	–

Statistics. Table 1 shows the statistics of the datasets used in this paper. We augment the COCO-backpack (chair) by Web-backpack (chair), whose size is about 15 times as the former. “Web-backpack labeled” is for evaluation only.

3 Selective self-supervised self-training

In this section, we describe our learning method named Selective Self-Supervised Self-training for Object Detection (S⁴OD). Without loss of generality, we consider detecting only one class of objects from an input image. Denote by $\mathcal{D} = \{(\mathbf{I}_1, \mathcal{T}_1), (\mathbf{I}_2, \mathcal{T}_2), \dots, (\mathbf{I}_n, \mathcal{T}_n)\}$ the labeled image set and $\mathcal{U} = \{\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2, \dots, \tilde{\mathbf{I}}_m\}$ the crawled Web images, where $t_i^j = \{x_i^j, y_i^j, w_i^j, h_i^j\} \in \mathcal{T}_i$ contains the top-left coordinate, width, and height of the j -th ground-truth bounding box in the i -th image \mathbf{I}_i . The labeled image set \mathcal{D} is significantly smaller than the set of Web images \mathcal{U} . Besides, there exists a domain shift between the two sets, although we have tried to mitigate the mismatch by using the image-to-image search. Finally, some Web images could contain zero objects of the class being considered. In the following, we first customize vanilla self-training for object detection, discuss its limitations and fixations by a selective net, and then arrive at the full S⁴OD algorithm. Figure 4 illustrates a diagram of our approach.

3.1 Self-training for object detection (SOD)

Given the labeled set \mathcal{D} and unlabeled set \mathcal{U} , it is natural to test how self-training performs, especially given that it has recently achieved remarkable results [66] on ImageNet [8]. Following the procedure in [66], we first train a teacher object detector $f(\mathbf{I}, \theta_t^*)$ from the labeled images, where θ_t^* stands for the network weights. We then produce pseudo boxes for each unlabeled Web image $\tilde{\mathbf{I}}_i \in \mathcal{U}$:

$$\tilde{\mathcal{T}}_i, \tilde{\mathcal{S}}_i \leftarrow f(\tilde{\mathbf{I}}_i, \theta_t^*), \quad i = 1, \dots, m \quad (1)$$

where each pseudo box $\tilde{t}_i^j \in \tilde{\mathcal{T}}_i$ is also associated with a confidence score $\tilde{s}_i^j \in \tilde{\mathcal{S}}_i$. We obtain the confidence score from the detector’s classification head. Finally, we train a student detector in a pre-training-and-fine-tuning manner. The idea is to pre-train the student detector using the Web images \mathcal{U} along with the pseudo bounding boxes, followed by fine-tuning it on the curated set \mathcal{D} .

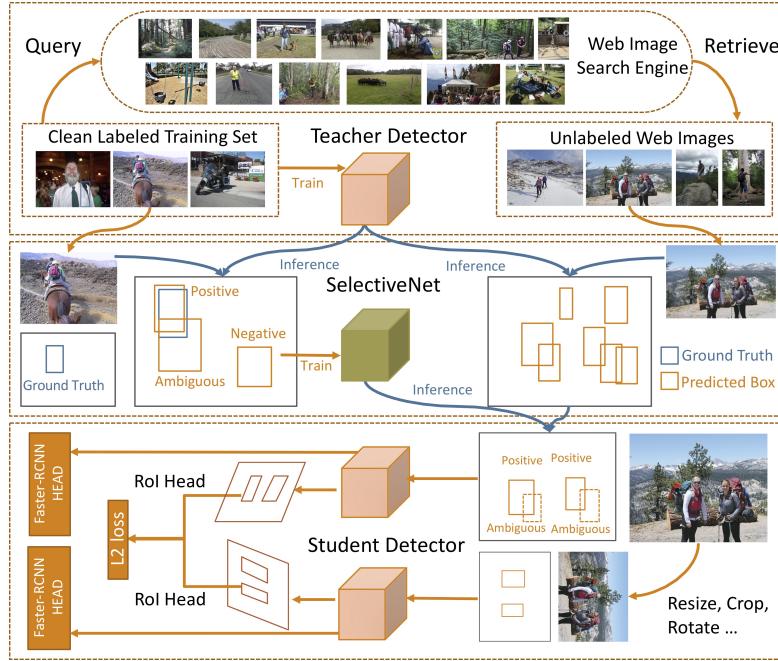


Fig. 4. Overview of the proposed approach. Top: using a small curated dataset to retrieve relevant Web images and to train a teacher detector. Middle: training a selective net to group pseudo boxes predicted by the teacher into positive, negative, and ambiguous groups. Bottom: learning a student detector from the Web-augmented training set with a self-supervised loss.

Modern object detectors generate hundreds of object candidates per image to ensure high recall even after non-maximum suppression [48], implying that many of the predicted pseudo boxes are incorrect. Traditional self-training used in image classification [66,50] disregards low-confidence labels when they train the student model. In the same spirit, we only keep the pseudo boxes whose confidence scores are higher than 0.7.

3.2 Selective self-training for object detection (S^2OD)

In SOD described above, a crucial step is the selection of pseudo boxes by thresholding the confidence scores. The effectiveness of SOD largely depends on the quality of the selected boxes, which, unfortunately, poorly correlates with the confidence score. Figure 5 shows some examples where the pseudo boxes tightly bound the backpacks, but the teacher detector assigns them very low confidence scores. As a result, those boxes would be removed before SOD trains the student detector, under-utilizing the Web images \mathcal{U} . What's worse is that the mistakenly removed boxes could be discovered as false hard negatives during training.



Fig. 5. Top: images in the COCO training set with groundtruth boxes. Bottom: boxes, their IoUs with the groundtruth, and confidence scores predicted by the teacher detector. The correlation between IoU and confidence is low.

To tackle the challenges, we propose a selective net g to calibrate the confidence scores of the pseudo boxes in the Web images \mathcal{U} . The main idea is to automatically group the boxes into three categories: positive, negative, and ambiguity. Denote by $g \circ \tilde{\mathcal{T}}_i$ the grouping results of the pseudo boxes in $\tilde{\mathcal{T}}_i$ for the i -th Web image \mathbf{I}_i . We pre-train the student detector $f(\mathbf{I}, \theta_s^*)$ by the following (and then fine-tune it on the curated training set \mathcal{D}),

$$\theta_s^* \leftarrow \arg \min_{\theta_s} \frac{1}{m} \sum_{i=1}^m \ell(g \circ \tilde{\mathcal{T}}_i, f(\mathbf{I}_i, \theta_s)), \quad (2)$$

where $\ell(\cdot, \cdot)$ is the conventional loss for training object detectors. For Faster-RCNN [48], the loss consists of regression, classification, objectiveness, etc. All the positive boxes predicted by our selective net g are used to activate those loss terms. In contrast, the ambiguous boxes, which could be correct but missed by the selective net, create a safe zone and do not contribute to any of the loss terms. This safe zone is especially useful when the learning algorithm has a hard negative mining scheme built in because it excludes potentially false “hard negatives” that fall in this ambiguity group.

Preparing training data for the selective net g . How do we learn the selective net g without knowing any groundtruth labels of the Web images? We seek answers by revisiting the labeled training set \mathcal{D} instead. After we learn the teacher detector, we apply it to the training images in \mathcal{D} and obtain a large pool of pseudo boxes. We assign each pseudo box \tilde{t}_i^j in the i -th image to one of the three groups by comparing it to the groundtruth boxes \mathcal{T}_i ,

$$g \circ \tilde{\mathcal{T}}_i^j = \begin{cases} \text{Negative,} & \max_{t \in \mathcal{T}_i} IoU(\tilde{t}_i^j, t) \leq \gamma_l, \\ \text{Positive,} & \max_{t \in \mathcal{T}_i} IoU(\tilde{t}_i^j, t) \geq \gamma_h, \\ \text{Ambiguity,} & \text{otherwise,} \end{cases} \quad (3)$$

where IoU is the intersection-over-union, a common evaluation metric in object detection and semantic segmentation [34,12], and $\gamma_h = 0.6$ and $\gamma_l = 0.05$ are

two thresholds for IoU as opposed for the confidence scores. Interestingly, we can choose γ_h by using the COCO evaluation protocol [34] as follows. Considering all the boxes in the positive group as the teacher detector’s final output, we can compute their mean average precision (mAP) over the labeled images \mathcal{D} . We choose the threshold of $\gamma_h = 0.6$ that maximizes the mAP.

Preparing features for the selective net g . We accumulate all potentially useful features to represent a pseudo box so that the selective net can have enough information to group the boxes. The feature vector for a box \tilde{t}_i^j is

$$\left(f_{RoI}(\tilde{t}_i^j, \mathbf{I}_i, \theta_t^*), \tilde{s}_i^j, \tilde{x}_i^j/W_i, \tilde{y}_i^j/H_i, \tilde{w}_i^j/W_i, \tilde{h}_i^j/H_i, W_i, H_i \right) \quad (4)$$

where $f_{RoI}(\tilde{t}_i^j, \mathbf{I}_i, \theta_t^*)$ is the RoI-pooled features [48] from the teacher detector, \tilde{s}_i^j is the confidence score, W_i and H_i are respectively the width and height of the i -th image, and the others are normalized box coordinate and size.

Training the selective net g . With the training data (eq. (3)) and features (eq. (4)) of the pseudo boxes, we learn the selective net by a three-way cross-entropy loss. We employ a straightforward architecture for the selective net. It comprises two towers. One is to process the normalized RoI-pooled features, and the other is to encode the remaining box features. They are both one-layer perceptrons with 512 and 128 hidden units, respectively. We then concatenate and feed their outputs into a three-way classifier.

One may concern that applying the teacher detector to the original training set \mathcal{D} may not give rise to informative training data for the selective net because the detector could have “overfitted” the training set. Somehow surprisingly, we find that it is extremely difficult to overfit a detector to the training set, probably due to inconsistent human annotations of the bounding boxes. At best, the detector plays the role of an “average rater” who still cannot reach 100% mAP on the training set whose bounding boxes are provided by different users.

3.3 Selective self-supervised self-training object detection (S^4OD)

Finally, we boost S^2OD by a self-supervised loss based on two considerations. One is that Xie et al. demonstrate that it is beneficial to enforce the student network to learn more knowledge than what the teacher provides (e.g., robustness to artificial noise) [66]. The other is that Jeong et al. show the effectiveness of adding a consistency regularization to semi-supervised object detection [26].

More concretely, we add the following loss to eq. (2) for each Web image $\tilde{\mathbf{I}}_i$,

$$\ell_i := \sum_{\tilde{t}_i^j \in \tilde{\mathcal{T}}_i(\text{Positive})} \left\| f_{RoI}(\tilde{t}_i^j, \tilde{\mathbf{I}}_i, \theta_s) - f_{RoI}(t_i^j, \mathbf{I}_i, \theta_s) \right\|_2 + \ell(g \circ \mathcal{T}_i, f(\mathbf{I}_i, \theta_s)), \quad (5)$$

which consists of a consistency term borrowed from [26] and the same detection loss as eq. (2) yet over a transformed Web image $\tilde{\mathbf{I}}_i$ — we explain the details

below. They are additional cues to the pseudo boxes provided by the teacher. By learning harder than the teacher from all the supervision using the extra Web data, we expect the student detector to outperform the teacher.

Given a noisy Web image $\tilde{\mathbf{I}} \in \mathcal{U}$, we use the selective net g to pick up positive boxes $\tilde{\mathcal{T}}(\text{Positive})$ and limit the consistency loss over them. This small change from [26], which is feasible due to the selective net, turns out vital to the final performance. We transform a Web image $\tilde{\mathbf{I}}$ to \mathbf{I} by randomly choosing an operation from $\{\text{rotation by } 90, 180, \text{ or } 270 \text{ degrees}\} \times \{\text{horizontal flip or not}\} \times \{\text{random crop}\}$. We use the crop ratio 0.9 and always avoid cutting through positive boxes. Accordingly, we can also obtain the transformed pseudo boxes $\mathbf{t}^j \in \mathcal{T}$. The bottom panel of Figure 4 exemplifies this transformation procedure.

4 Related Work

Weakly supervised object detection. To reduce the dependence on large-scale human-curated data of state-of-the-art object detectors [75,19,48,32,35,47,22], weakly supervised object detection (WSOD) [2,5,70,55,54,73] resort to image-level annotations to help localize objects. Besides, motion cues in videos are also explored by [71,45,52,27]. In addition, Tao *et al.* [57] incorporate Web images to learn a good feature representation for WSOD. Fine-grained segmentation [17,31] can also be used to guide WSOD. Although great progress has been made in WSOD, there is still quite a gap for them to catch up its supervised counterpart. The performance of fully supervised methods [63,48] are about 25 points better in terms of mean average precision compared to the weakly supervised ones [73].

Semi-supervised Classification. Semi-supervised learning [3,46,28,58,39,36,4,65] jointly explore labeled and unlabeled data in training machine learning models. The consistency of model predictions or latent representations is an effective tool in semi-supervised learning [36,4,65,39]. The consistency regularization [28,58,39] reach state-of-the-art performance [41]. Besides, pseudo-label based approaches are also popular in semi-supervised learning [24,51,29,1]. In this vein, Xie *et al.* [66] use self-training [50,67,49,72] and aggressively inject noise in the learning procedure. They report state-of-the-art results on ImageNet with 300M unlabeled images.

Semi-supervised Object Detection. In the literature, the semi-supervised object detection could refer to two categories of work: **weakly semi-supervised detectors** [68,56,16,69] and **complete semi-supervised detectors** [26,60]. The former uses fully annotated data with box annotations and weakly labeled data with only image-level annotations. The latter learns by using unlabeled data in combination with the box-level labeled data. There are only a few works in the complete semi-supervised detection. Wang *et al.* [60] present a self-supervised sample mining process in active learning to get reliable region proposals for enhancing the object detector. Jeong *et al.* [26] introduce a consistency loss over unlabeled images.

Unsupervised/Self-supervised Representation Learning. Unsupervised/self-supervised representation learning [10,21,14,13,15,74] from unlabeled data has attracted great attention recently. Some of them define a wide range of pretext tasks, such as recovering the input under some corruption [59,44], predicting rotation [18] or patch orderings [10,40] of an exemplar image, and tracking [61] or segmenting objects [43] in videos. Others utilize contrastive learning [42,11,21,64,6] by maximizing agreement between differently augmented views of the same data example. High-quality visual representations can help object detection [37,18], and self-supervised learning has been applied to replace the supervised ImageNet pretraining [43,25] for object detection. In addition, Lee *et al.* [30] propose a set of auxiliary tasks to make better use of scarce labels.

We study how to improve object detection with unlabeled, out-of-distribution Web images. While this problem is similar to semi-supervised object detection, neither self-training nor the self-supervised consistency loss, which are popular in semi-supervised learning, can cope with the Web images (cf. Table 2 where SOD and CSD can barely improve over the baselines (BD)). A major reason is that the out-of-distribution Web images give rise to many inaccurate candidate boxes and uncalibrated box classification scores. To tackle these challenges, we propose a selective net to identify high-quality (both positive and negative) boxes.

5 Experiments

We augment the COCO-backpack and COCO-chair training images with unlabeled Web images and run extensive experiments to test our approach on them. We note that the Web images are out of the distribution of COCO [34]. Some of them may contain no backpack or chair at all, and we rely on our selective net to identify the useful pseudo boxes in them produced by a teacher detector.

Besides, since S⁴OD is readily applicable to vanilla semi-supervised object detection, whose labeled and unlabeled sets follow the same distribution, we also test it following the experiment protocol of [26].

5.1 Augmenting COCO-backpack(chair) with Web-backpack(chair)

We compare S⁴OD to the following competing methods.

R101-FPN [32]: We use R101-FPN implemented in Detectron2 [63] as our object detector. Its *AP* on COCO-validation is 42.03, which is on par with the state of the arts. However, this detector still performs unsatisfactorily on the backpack and chair classes. As shown in Table 2, it only achieves 16.45 *AP* on detecting backpacks and 28.28 *AP* on chairs.

BD [32]: We finetune R101-FPN on COCO-backpack and COCO-chair, respectively, by changing the 80-class detector to a binary backpack or chair detector. We observe that the binary detector outperforms its 80-class counterpart by about 1% (0.9% *AP* for backpacks and 1.3% *AP* for chairs). We denote by BD the binary backpack (chair) detector and use it as the teacher detector for the remaining experiments.

Table 2. Comparison results for detecting backpacks and chairs.

Dataset	Method	Web	AP@[.5, .95]	AP@.5	AP@.75	AP _S	AP _M	AP _L
COCO-backpack	R101-FPN [32]	✗	16.45	33.63	15.26	18.29	18.23	17.61
	BD [32]	✗	17.34	34.71	15.13	19.03	19.46	21.31
	SOD [50]	✓	17.62	33.83	16.71	18.98	21.88	15.93
	CSD [26]	✓	17.87	32.97	16.55	20.42	22.05	20.45
	CSD-Selective (ours)	✓	18.32	33.35	17.05	20.84	21.08	21.11
	S ² OD (ours)	✓	18.28	35.35	18.46	19.23	21.76	23.94
	S ⁴ OD (ours)	✓	19.48	36.25	19.01	20.31	23.47	18.53
COCO-chair	R101-FPN [32]	✗	28.28	48.93	28.57	19.14	33.21	42.06
	BD [32]	✗	29.56	49.15	30.70	20.16	36.39	41.99
	SOD [50]	✓	30.01	49.05	31.53	20.32	37.10	43.12
	CSD [26]	✓	30.19	48.58	31.92	20.77	37.19	43.74
	CSD-Selective (ours)	✓	30.95	49.75	32.11	21.03	38.02	44.96
	S ² OD (ours)	✓	30.95	50.25	31.92	21.19	37.82	45.48
	S ⁴ OD (ours)	✓	31.74	51.15	33.29	20.57	38.66	46.93

CSD [26]: We include the recently published consistency-based semi-supervised object detection (CSD) in the experiments. We carefully re-implement it by using the same backbone detector (R101-FPN) as ours. CSD imposes the consistency loss over all candidate boxes.

CSD-Selective: We also report the results of applying our selective net to CSD.

We remove the loss terms in CSD over the boxes of negative and ambiguity groups predicted by the selective net. Table 2 shows that it increases the performance of the original CSD by 0.45 *AP* on detecting backpacks and 0.76 *AP* on chairs.

SOD [50]: We presented a vanilla self-training procedure for object detection (SOD) in Section 3.1. Unlike self-training in image classification [66], we find that SOD is sensitive to the threshold of confidence scores probably because the out-of-distribution Web images make the scores highly uncalibrated. We test all thresholds between 0.5 and 0.9 (with an interval of 0.2) and find that SOD can only beat BD at the threshold of 0.7 (reported in Table 2).

S²OD (ours): This is an improved method over SOD by employing the selective net (cf. Section 3.2). It is also an ablated version of our S⁴OD by removing the self-supervised loss.

Table 2 presents the comparison results on the validation sets of both COCO-backpack and COCO-chair. We can see that S²OD performs consistently better than its teacher detector (BD), the vanilla self-training (SOD), and the consistency-based self-supervised learning (CSD). Our full approach (S⁴OD) brings additional gains and gives rise to 19.48 *AP* on detecting COCO backpacks and 31.74 *AP* on detecting chairs — about 2% better than its original teacher (BD). Besides, the improvements of S²OD over SOD and CSD-selective over CSD both attribute to the selective net. Finally, Figure 6 shows some cases where the selective net correctly groups low(high)-confidence boxes into the positive (negative) group.



Fig. 6. The two panels on the left: Positive (negative) boxes chosen by the selective net yet with low (high) confidence scores predicted by the teacher detector. Rightmost: Ambiguous boxes according to the selective net.



Fig. 7. Example Web images for augmenting COCO-backpack.

In addition to the overall comparison results in Table 2, we next ablate our approach and examine some key components. We also report the “upper-bound” results for augmenting COCO-backpack with Web-backpack by using the bounding box labels we collected for a subset of Web-backpack.

Web-backpack *vs.* text-to-image search. The image-to-image search for Web images is the very first step of our approach, and it is superior over the text-to-image search in various aspects. As the leftmost panel in Figure 7 shows, most top-ranked Web images are iconic with salient objects sitting in clean backgrounds if we search using class names. For COCO [34], the images are collected from Flickr by complex object-object and object-scene queries. Using the same technique, we can retrieve more natural images. However, they are mostly recent and come from diverse sources, exhibiting a clear domain shift from the COCO dataset which is about six years old. In contrast, the image-to-image search well balances between the number of the retrieved Web images and their domain similarity to the query images (cf. the right panel in Figure 7).

Table 3 compares image-to-image search with text-to-image search by their effects on the final results. The S⁴OD-text row is the results obtained using 230k Web images crawled by text-to-image. While they are slightly better than BD’s results, they are significantly worse than what S⁴OD achieves with 70k Web images retrieved by image-to-image search (cf. Table 1).

Web-backpack: the size matters. We study how the number of the unlabeled Web images influences the proposed S⁴OD by training with 1/3 and 2/3 of the

Table 3. Results of S⁴OD with various Web data augmentations to COCO-backpack.

Method	AP@[.5, .95]	AP@.5	AP@.75	AP _S	AP _M	AP _L
BD [32]	17.34	34.71	15.13	19.03	19.46	21.31
S ⁴ OD-text	17.37	34.26	14.46	18.77	20.10	17.66
S ⁴ OD w/ 1/3 Web-backpack	18.40	34.12	17.91	19.25	21.20	21.56
S ⁴ OD w/ 2/3 Web-backpack	19.08	35.35	17.96	21.04	20.15	21.27
S ⁴ OD w/ full Web-backpack	19.48	36.25	19.01	20.31	23.47	18.53
S ⁴ OD-2 _{nd} Iteration	19.52	36.44	19.07	21.70	20.90	24.90

Table 4. Comparison results on the relabeled COCO-backpack.

Method	AP@[.5, .95]	AP@.5	AP@.75	AP _S	AP _M	AP _L
BD	18.75	36.03	16.98	11.90	21.62	31.58
Upper bound	22.63	40.89	21.15	15.97	26.53	32.11
S ⁴ OD	20.27	37.55	20.49	13.51	24.04	31.17

crawled Web-backpack. As shown in table 3, S⁴OD-1/3 can improve over BD by 1% AP, and S⁴OD-2/3 is better than BD by 1.7% AP. In contrast, S⁴OD with the full Web-backpack leads to about 2.1% AP improvement. Overall, we see that the more Web images, the larger boost in performance, implying that it is worth studying the data pipeline and expanding its coverage in future work.

Iterating S⁴OD. What happens if we use the detector trained by S⁴OD as the teacher and train another student detector using S⁴OD again? S⁴OD-2_{nd} Iteration in table 3 outperforms the single-iteration version only by a very small margin. It is probably because, in the second iteration, the student detector does not receive more supervision than what the teacher BD provides. We will explore some stochastic self-supervised losses in future work to enforce the student to learn more than what the teacher provides at every iteration.

An “upper bound” of S⁴OD. We further investigate the effectiveness of S⁴OD by comparing it to an “upper bound”. We run the experiments using the labeled subset of Web-backpack and the relabeled COCO-backpack to have consistent annotations across the two datasets. Recall that we have fixed some inconsistent bounding boxes in the original COCO-backpack in the relabeling process. As a result, if we train BD [48] on the relabeled COCO-backpack and evaluate on the relabeled validation set, the AP is 18.75 (cf. Table 4), in contrast to 17.34 AP (cf. Table 2) by BD trained and evaluated using the original COCO-backpack. Using this BD as the teacher, we train a student detector using S⁴OD. By the “upper bound”, we pool the labeled subset of Web-backpack and the relabeled COCO-backpack together and then train BD. The results in Table 4 indicate that S⁴OD is almost right in the middle of the lower-bound (BD) and the upper bound. The gap between S⁴OD and the upper bound is small. We will study better learning methods to close the gap and how the Web data volume could impact the performance in future work.

Table 5. Comparison esults on VOC2007 (* the numbers reported in [26])

Method	Labeled data	Unlabeled data	AP	Gain
Supervised [7,48]	VOC07	–	*73.9/74.1	–
Supervised [7,48]	VOC07&12	–	*79.4/80.3	–
CSD [26]	VOC07	VOC12	*74.7	*0.8
SOD	VOC07	VOC12	74.8	0.7
S ² OD	VOC07	VOC12	75.8	1.7
S ⁴ OD	VOC07	VOC12	76.4	2.3

5.2 Semi-supervised object detection on PASCAL VOC

It is straightforward to extend S⁴OD to multiple object detection tasks. Following the experiment protocol in [26], we further validate it in the setting of semi-supervised object detection, whose labeled and unlabeled sets are drawn from the same distribution. We use PASCAL VOC2007 trainval (5,011 images) as the labeled set and PASCAL VOC2012 trainval (11,540 images) as the unlabeled set. There are 20 classes of objects to detect. We test all detectors on the test set of PASCAL VOC2007 (4,952 images). In order to make our results comparable to what are reported in [26], we switch to Faster-RCNN [48] with ResNet-50 [23] as the base detector.

Table 5 shows the comparison results. We include both our results and those reported in [25] and mark the latter by *. Considering the object detector trained on VOC2007 as a baseline and the one trained on both datasets with labels as an upper bound, our S⁴OD is right in the middle, outperforming CSD by 1.5 AP. Although we propose the selective net mainly to handle noisy Web images, it is delightful that the resulting method also works well with the clean, in-domain VOC2012 images. It is probably because the consistency loss in CSD drives the detector toward high-entropy, inaccurate predictions (cf. more discussions in [66]). Our selective net avoids this caveat by supplying high-quality boxes to the consistency loss.

6 Conclusion

In this paper, we propose a novel approach to improving object detection with massive unlabeled Web images. We collect the Web images with image-to-image search, leading to smaller domain mismatch between the retrieved Web images and the curated dataset than text-to-image search does. Besides, we incorporate a principled selective net into self-training rather than threshold confidence scores as a simple heuristic for selecting bounding boxes. Moreover, we impose a self-supervised training loss over the high-quality boxes chosen by the selective net to make better use of the Web images. The improvement in detecting challenging objects is significant over the competing methods. Our approach works consistently well on not only the Web-augmented object detection but also the traditional semi-supervised object detection.

References

1. Arazo, E., Ortego, D., Albert, P., O'Connor, N.E., McGuinness, K.: Pseudo-labeling and confirmation bias in deep semi-supervised learning. arXiv preprint arXiv:1908.02983 (2019)
2. Arun, A., Jawahar, C., Kumar, M.P.: Dissimilarity coefficient based weakly supervised object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9432–9441 (2019)
3. Bachman, P., Alsharif, O., Precup, D.: Learning with pseudo-ensembles. In: Advances in neural information processing systems. pp. 3365–3373 (2014)
4. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: Advances in Neural Information Processing Systems. pp. 5050–5060 (2019)
5. 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)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. arXiv preprint arXiv:2002.05709 (2020)
7. Dai, J., Li, Y., He, K., Sun, J.: R-fcn: Object detection via region-based fully convolutional networks. In: Advances in neural information processing systems. pp. 379–387 (2016)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
9. Ding, Y., Wang, L., Fan, D., Gong, B.: A semi-supervised two-stage approach to learning from noisy labels. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1215–1224. IEEE (2018)
10. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1422–1430 (2015)
11. Donahue, J., Simonyan, K.: Large scale adversarial representation learning. In: Advances in Neural Information Processing Systems. pp. 10541–10551 (2019)
12. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International journal of computer vision **88**(2), 303–338 (2010)
13. Gan, C., Sun, C., Duan, L., Gong, B.: Webly-supervised video recognition by mutually voting for relevant web images and web video frames. In: ECCV. pp. 849–866 (2016)
14. Gan, C., Yao, T., Yang, K., Yang, Y., Mei, T.: You lead, we exceed: Labor-free video concept learning by jointly exploiting web videos and images. In: CVPR. pp. 923–932 (2016)
15. Gan, C., Zhao, H., Chen, P., Cox, D., Torralba, A.: Self-supervised moving vehicle tracking with stereo sound. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 7053–7062 (2019)
16. Gao, J., Wang, J., Dai, S., Li, L.J., Nevatia, R.: Note-rcnn: Noise tolerant ensemble rcnn for semi-supervised object detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9508–9517 (2019)
17. Gao, Y., Liu, B., Guo, N., Ye, X., Wan, F., You, H., Fan, D.: C-midn: Coupled multiple instance detection network with segmentation guidance for weakly supervised object detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9834–9843 (2019)

18. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728* (2018)
19. Girshick, R.: Fast r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1440–1448 (2015)
20. Grabner, H., Gall, J., Van Gool, L.: What makes a chair a chair? In: *CVPR 2011*. pp. 1529–1536. IEEE (2011)
21. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722* (2019)
22. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2961–2969 (2017)
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
24. Iscen, A., Tolias, G., Avrithis, Y., Chum, O.: Label propagation for deep semi-supervised learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5070–5079 (2019)
25. Jenni, S., Favaro, P.: Self-supervised feature learning by learning to spot artifacts. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2733–2742 (2018)
26. Jeong, J., Lee, S., Kim, J., Kwak, N.: Consistency-based semi-supervised learning for object detection. In: *Advances in Neural Information Processing Systems*. pp. 10758–10767 (2019)
27. 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)
28. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242* (2016)
29. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: *Workshop on challenges in representation learning, ICML*. vol. 3, p. 2 (2013)
30. Lee, W., Na, J., Kim, G.: Multi-task self-supervised object detection via recycling of bounding box annotations. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4984–4993 (2019)
31. Li, X., Kan, M., Shan, S., Chen, X.: Weakly supervised object detection with segmentation collaboration. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 9735–9744 (2019)
32. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2117–2125 (2017)
33. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2980–2988 (2017)
34. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer (2014)
35. 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)

36. Luo, Y., Zhu, J., Li, M., Ren, Y., Zhang, B.: Smooth neighbors on teacher graphs for semi-supervised learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8896–8905 (2018)
37. Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., van der Maaten, L.: Exploring the limits of weakly supervised pretraining. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 181–196 (2018)
38. Minsky, M.: Society of mind. Simon and Schuster (1988)
39. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* **41**(8), 1979–1993 (2018)
40. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European Conference on Computer Vision. pp. 69–84. Springer (2016)
41. Oliver, A., Odena, A., Raffel, C.A., Cubuk, E.D., Goodfellow, I.: Realistic evaluation of deep semi-supervised learning algorithms. In: Advances in Neural Information Processing Systems. pp. 3235–3246 (2018)
42. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018)
43. Pathak, D., Girshick, R., Dollár, P., Darrell, T., Hariharan, B.: Learning features by watching objects move. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2701–2710 (2017)
44. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2536–2544 (2016)
45. Prest, A., Leistner, C., Civera, J., Schmid, C., Ferrari, V.: Learning object class detectors from weakly annotated video. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 3282–3289. IEEE (2012)
46. Rasmus, A., Berglund, M., Honkala, M., Valpola, H., Raiko, T.: Semi-supervised learning with ladder networks. In: Advances in neural information processing systems. pp. 3546–3554 (2015)
47. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7263–7271 (2017)
48. 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)
49. Riloff, E., Wiebe, J.: Learning extraction patterns for subjective expressions. In: Proceedings of the 2003 conference on Empirical methods in natural language processing. pp. 105–112 (2003)
50. Scudder, H.: Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory* **11**(3), 363–371 (1965)
51. Shi, W., Gong, Y., Ding, C., MaXiaoyu Tao, Z., Zheng, N.: Transductive semi-supervised deep learning using min-max features. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 299–315 (2018)
52. Singh, K.K., Lee, Y.J.: You reap what you sow: Using videos to generate high precision object proposals for weakly-supervised object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9414–9422 (2019)
53. Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10781–10790 (2020)

54. Tang, P., Wang, X., Bai, X., Liu, W.: Multiple instance detection network with online instance classifier refinement. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2843–2851 (2017)
55. Tang, P., Wang, X., Wang, A., Yan, Y., Liu, W., Huang, J., Yuille, A.: Weakly supervised region proposal network and object detection. In: Proceedings of the European conference on computer vision (ECCV). pp. 352–368 (2018)
56. Tang, Y., Wang, J., Gao, B., Dellandréa, E., Gaizauskas, R., Chen, L.: Large scale semi-supervised object detection using visual and semantic knowledge transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2119–2128 (2016)
57. Tao, Q., Yang, H., Cai, J.: Zero-annotation object detection with web knowledge transfer. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 369–384 (2018)
58. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in neural information processing systems. pp. 1195–1204 (2017)
59. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th international conference on Machine learning. pp. 1096–1103 (2008)
60. Wang, K., Yan, X., Zhang, D., Zhang, L., Lin, L.: Towards human-machine cooperation: Self-supervised sample mining for object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1605–1613 (2018)
61. Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2794–2802 (2015)
62. Wei, X., Gong, B., Liu, Z., Lu, W., Wang, L.: Improving the improved training of wasserstein gans: A consistency term and its dual effect. arXiv preprint arXiv:1803.01541 (2018)
63. Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. <https://github.com/facebookresearch/detectron2> (2019)
64. Wu, Z., Xiong, Y., Yu, S., Lin, D.: Unsupervised feature learning via non-parametric instance-level discrimination. arXiv preprint arXiv:1805.01978 (2018)
65. Xie, Q., Dai, Z., Hovy, E., Luong, M.T., Le, Q.V.: Unsupervised data augmentation. arXiv preprint arXiv:1904.12848 (2019)
66. Xie, Q., Hovy, E., Luong, M.T., Le, Q.V.: Self-training with noisy student improves imagenet classification. arXiv preprint arXiv:1911.04252 (2019)
67. Yalniz, I.Z., Jégou, H., Chen, K., Paluri, M., Mahajan, D.: Billion-scale semi-supervised learning for image classification. arXiv preprint arXiv:1905.00546 (2019)
68. Yan, Z., Liang, J., Pan, W., Li, J., Zhang, C.: Weakly-and semi-supervised object detection with expectation-maximization algorithm. arXiv preprint arXiv:1702.08740 (2017)
69. Yang, H., Wu, H., Chen, H.: Detecting 11k classes: Large scale object detection without fine-grained bounding boxes. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9805–9813 (2019)
70. Yang, K., Li, D., Dou, Y.: Towards precise end-to-end weakly supervised object detection network. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 8372–8381 (2019)
71. Yang, Z., Mahajan, D., Ghadiyaram, D., Nevatia, R., Ramanathan, V.: Activity driven weakly supervised object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2917–2926 (2019)

72. Yarowsky, D.: Unsupervised word sense disambiguation rivaling supervised methods. In: 33rd annual meeting of the association for computational linguistics. pp. 189–196 (1995)
73. Zeng, Z., Liu, B., Fu, J., Chao, H., Zhang, L.: Wsod2: Learning bottom-up and top-down objectness distillation for weakly-supervised object detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 8292–8300 (2019)
74. Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., Torralba, A.: The sound of pixels. In: ECCV. pp. 570–586 (2018)
75. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. In: arXiv preprint arXiv:1904.07850 (2019)