

End-to-End Semi-Supervised Object Detection with Soft Teacher

Mengde Xu^{1†*} Zheng Zhang^{1,2*‡} Han Hu^{2‡} Jianfeng Wang² Lijuan Wang² Fangyun Wei²
Xiang Bai¹ Zicheng Liu²

¹Huazhong University of Science and Technology

{mdxu, xbai}@hust.edu.cn, macaroniz1990@gmail.com

²Microsoft

{zhez, hanhu, jianfw, lijuanw, fawe, zliu}@microsoft.com

Abstract

This paper presents an end-to-end semi-supervised object detection approach, in contrast to previous more complex multi-stage methods. The end-to-end training gradually improves pseudo label qualities during the curriculum, and the more and more accurate pseudo labels in turn benefit object detection training. We also propose two simple yet effective techniques within this framework: a soft teacher mechanism where the classification loss of each unlabeled bounding box is weighed by the classification score produced by the teacher network; a box jittering approach to select reliable pseudo boxes for the learning of box regression. On the COCO benchmark, the proposed approach outperforms previous methods by a large margin under various labeling ratios, i.e. 1%, 5% and 10%. Moreover, our approach proves to perform also well when the amount of labeled data is relatively large. For example, it can improve a 40.9 mAP baseline detector trained using the full COCO training set by +3.6 mAP, reaching 44.5 mAP, by leveraging the 123K unlabeled images of COCO. On the state-of-the-art Swin Transformer based object detector (58.9 mAP on test-dev), it can still significantly improve the detection accuracy by +1.5 mAP, reaching 60.4 mAP, and improve the instance segmentation accuracy by +1.2 mAP, reaching 52.4 mAP. Further incorporating with the Object365 pre-trained model, the detection accuracy reaches 61.3 mAP and the instance segmentation accuracy reaches 53.0 mAP, pushing the new state-of-the-art. The code and models will be made publicly available at <https://github.com/microsoft/SoftTeacher>.

1. Introduction

Data matters. In fact, large data such as ImageNet has largely triggered the boom of deep learning in computer vi-

*Equal contribution. †This work is done when Mengde Xu was intern in MSRA. ‡Contact person.

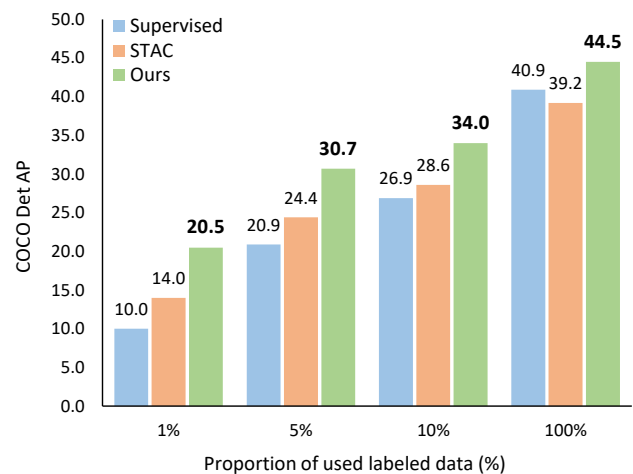


Figure 1. The proposed end-to-end pseudo-label based semi-supervised object detection method outperforms the STAC [27] by a large margin on MS-COCO benchmark.

sion. However, obtaining labels can be a bottleneck, due to the time-consuming and expensive annotation process. This has encouraged learning methods to leverage unlabeled data in training deep neural models, such as self-supervised learning and semi-supervised learning. This paper studies the problem of semi-supervised learning, in particular for object detection.

For semi-supervised object detection, we are concerned with the pseudo-label based approaches, which are the current state-of-the-art. These approaches [27, 36] conduct a multi-stage training schema, with the first stage training an initial detector using labeled data, followed by a pseudo-labeling process for unlabeled data and a re-training step based on the pseudo labeled unannotated data. These multi-stage approaches achieve reasonably good accuracy, however, the final performance is limited by the quality of pseudo labels generated by an initial and probably inaccurate detector trained using a small amount of labeled data.

detector	method	val2017		test-dev2017	
		mAP ^{det}	mAP ^{mask}	mAP ^{det}	mAP ^{mask}
HTC++(Swin-L) w/ single-scale	supervised	57.1	49.6	-	-
	ours	59.1	51.0	-	-
	ours*	60.1	51.9	-	-
HTC++(Swin-L) w/ multi-scale	supervised	58.2	50.5	58.9	51.2
	ours	59.9	51.9	60.4	52.4
	ours*	60.7	52.5	61.3	53.0

Table 1. On the state-of-the-art detector HTC++(Swin-L), our method surpasses the supervised learning on both val2017 and test-dev2017. * indicates that models are pre-trained with Object365 [24] dataset.

To address this issue, we present an end-to-end pseudo-label based semi-supervised object detection framework, which simultaneously performs pseudo-labeling for unlabeled images and trains a detector using these pseudo labels along with a few labeled ones at each iteration. Specifically, labeled and unlabeled images are randomly sampled with a preset ratio to form one data batch. Two models are applied on these images, with one conducting detection training and the other in charge of annotating pseudo labels for unlabeled images. The former is also referred to as a student, and the latter is a teacher, which is an exponential moving average (EMA) of the student model. This end-to-end approach avoids the complicated multi-stage training scheme. Moreover, it also enables a “flywheel effect” that the pseudo labeling and the detection training processes can mutually reinforce each other, so that both get better and better as the training goes on.

Another important benefit of this end-to-end framework is that it allows for greater leverage of the teacher model to guide the training of the student model, rather than just providing “some generated pseudo boxes with hard category labels” as in previous approaches [27, 36]. A *soft teacher* approach is proposed to implement this insight. In this approach, the teacher model is used to directly assess all the box candidates that are generated by the student model, rather than providing “pseudo boxes” to assign category labels and regression vectors to these student-generated box candidates. The direct assessment on these box candidates enables more extensive supervision information to be used in the student model training. Specifically, we first categorize the box candidates as foreground/background by their detection scores with a high foreground threshold to ensure a high precision of the positive pseudo labels, as in [27]. This high foreground threshold, however, results in many positive box candidates mistakenly assigned as background. To address this issue, we propose using a *reliability* measure to weight the loss of each “background” box candidate. We empirically find that a simple detection score produced by the teacher model can well serve as the reliability measure, and is used in our approach. We find that this approach measure performs significantly better than previous hard fore-

ground/background assignment methods (see Table 2 and Table 3), and we name it “*soft teacher*”.

Another approach instantiates this insight is to select reliable bounding boxes for the training of the student’s localization branch, by a *box jittering* approach. This approach first jitters a pseudo-foreground box candidate several times. Then these jittered boxes are regressed according the teacher model’s location branch, and the variance of these regressed boxes is used as a *reliability measure*. The box candidate with adequately high reliability will be used for the training of the student’s localization branch.

On MS-COCO object detection benchmark [16], our approach achieves 20.5 mAP, 30.7 mAP and 34.0 mAP on val2017 with 1%, 5% and 10% labeled data using the Faster R-CNN [22] framework with ResNet-50 [8] and FPN [14], surpassing previous best method STAC [27] by **+6.5**, **+6.4** and **+5.4** mAP, respectively.

In addition, we also perform evaluation on a more challenge setting where the labelled data has been adequately large to train a reasonably accurate object detector. Specifically, we adopt the complete COCO train2017 set as labeled data and the unlabeled2017 set as the unlabeled data. Under this setting, we improve the supervised baseline of a Faster R-CNN approach with ResNet-50 and ResNet-101 backbones by **+3.6** mAP and **+3.0** mAP, respectively.

Moreover, on a state-of-the-art Swin-Transformer [18] based detector which achieves 58.9 mAP for object detection and 51.2 mAP for instance segmentation on COCO test-dev2017, the proposed approach can still improve the accuracy by +1.5 mAP and +1.2 mAP, respectively, reaching **60.4** mAP and **52.4** mAP. Further incorporating with the Object365 [24] pre-trained model, the detection accuracy reaches **61.3** mAP and the instance segmentation accuracy reaches **53.0** mAP, which is the new state-of-the-art on this benchmark.

2. Related works

Semi-supervised learning in image classification Semi-supervised learning in image classification can be roughly categorized into two groups: consistency based and pseudo-label based. The consistency based methods [1, 23, 19, 11]

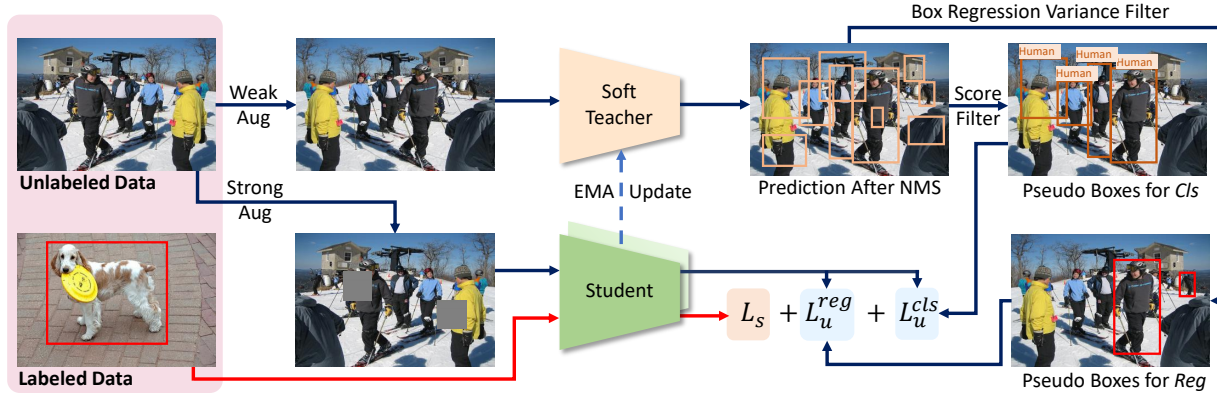


Figure 2. The overview of the end-to-end pseudo-labeling framework for semi-supervised object detection. Unlabeled images and labeled images form the training data batch. In each training iteration, a soft teacher is applied to perform pseudo-labeling on weak augmented unlabeled images on the fly. Two sets of pseudo boxes are produced: one is used for classification branch by filtering boxes according to the foreground score, and the other is used for box regression branch by filtering boxes according to box regression variance. The teacher model is updated by student model via exponential mean average (EMA) manner. The final loss is the sum of supervised detection loss L_s and unsupervised detection loss L_u .

leverage the unlabeled images to construct a regularization loss which encourages different perturbations of a same image to produce similar predictions. There are several ways to implement perturbations, including perturbing the model [1], augmenting the images [23] or adversarial training [19]. In [11], the training target is assembled by predicting different training steps. In [29], they develop [11] by ensembling the model itself instead of the model prediction, the so-called exponential mean average (EMA) of the student model. The pseudo-label approaches [33, 7, 12] (also named as self-training) annotate unlabeled images with pseudo labels by an initially trained classification model, and the detector is refined by these pseudo labeled images. Unlike our method which focuses on object detection, the pseudo-label does not have to solve the problem of assigning foreground/background labels and box regression when classifying images. Recently, some works [32, 3, 2, 26] explore the importance of data augmentation in semi-supervised learning, which inspire us to use the weak augmentation to generate pseudo-labels and the strong augmentation for the learning of detection models.

Semi-supervised learning in object detection Similar to the semi-supervised learning in image classification, semi-supervised object detection methods also have two categories: the consistency methods [10, 28] and pseudo-label methods [20, 36, 13, 27, 31]. Our method belongs to the pseudo-label category. In [20, 36], the predictions of different data augmentation are ensembled to form the pseudo labels for unlabeled images. In [13], a SelectiveNet is trained to select the pseudo-label. In [31], a box detected on an unlabeled image is pasted onto a labeled image, and

the localization consistency estimation is performed onto the pasted label image. As the image itself is modified, a very thorough detection process is required in [31]. In our method, only the lightweight detection head is processed. STAC [27] proposes to use a weak data augmentation for model training and a strong data augmentation is used for performing pseudo-label. However, like other pseudo-label methods [20, 36, 13, 27, 31], it also follows the multi-stage training scheme. In contrast, our method is an end-to-end pseudo-labeling framework, which avoids the complicated training process and also achieves better performance.

Object Detection Object detection focuses on designing efficient and accurate detection framework. There are two mainstreams: single-stage object detectors [17, 21, 30] and two-stage object detectors [6, 22, 14, 34, 35]. The main difference between the two types of methods is whether to use a cascade to filter a large number of object candidates (proposals). In theory, our method is compatible with both types of methods. However, to allow a fair comparison with previous works [28, 27] on semi-supervised object detection, we use Faster R-CNN [22] as our default detection framework to illustrate our method.

3. Methodology

Figure 2 illustrates an overview of our end-to-end training framework. There are two models, a student model and a teacher model. The student model is learned by both the detection losses on the labeled images and on the unlabeled images using pseudo boxes. The unlabeled images have two sets of pseudo boxes, which are used to drive the training of the classification branch and the regression branch, respec-

tively. The teacher model is an exponential moving average (EMA) of the student model. Within this end-to-end framework, there are two crucial designs: *soft teacher* and *box jittering*.

3.1. End-to-End Pseudo-Labeling Framework

We first introduce the end-to-end framework for pseudo-label based semi-supervised object detection. Our approach follows the teacher-student training scheme. In each training iteration, labeled images and unlabeled images are randomly sampled according to a data sampling ratio s_r to form a training data batch. The teacher model is performed to generate the pseudo boxes on unlabeled images and the student model is trained on both labeled images with the ground-truth and unlabeled images with the pseudo boxes as the ground-truth. Thus, the overall loss is defined as the weighted sum of supervised loss and unsupervised loss:

$$\mathcal{L} = \mathcal{L}_s + \alpha \mathcal{L}_u, \quad (1)$$

where \mathcal{L}_s and \mathcal{L}_u denote supervised loss of labeled images and unsupervised loss of unlabeled images respectively, α controls contribution of unsupervised loss. Both of them are normalized by the respective number of images in the training data batch:

$$\mathcal{L}_s = \frac{1}{N_l} \sum_{i=1}^{N_l} (\mathcal{L}_{\text{cls}}(I_l^i) + \mathcal{L}_{\text{reg}}(I_l^i)), \quad (2)$$

$$\mathcal{L}_u = \frac{1}{N_u} \sum_{i=1}^{N_u} (\mathcal{L}_{\text{cls}}(I_u^i) + \mathcal{L}_{\text{reg}}(I_u^i)), \quad (3)$$

where I_l^i indicates the i -th labeled image, I_u^i indicates the i -th unlabeled image, \mathcal{L}_{cls} is the classification loss, \mathcal{L}_{reg} is the box regression loss, N_l and N_u denote the number of labeled images and unlabeled images, respectively.

At the beginning of training, both the teacher model and student model are randomly initialized. As the training progresses, the teacher model is continuously updated by the student model, and we follow the common practices [29, 26] that the teacher model is updated by exponential moving average (EMA) strategy.

In contrast to taking a simple probability distribution as the pseudo-label in image classification, creating pseudo-label for object detection is more complicated since an image usually contains multiple objects and the annotation of objects consists of location and category. Given an unlabeled image, the teacher model is used to detect objects and thousands of box candidates are predicted. The non-maximum suppression (NMS) is then performed to eliminate redundancy. Although most redundant boxes are removed, there are still some non-foreground candidates left. Therefore, only candidates with the foreground score¹

¹The foreground score is defined as the maximum probability of all non-background categories.

higher than a threshold are retained as the pseudo boxes.

In order to generate high-quality pseudo boxes and to facilitate the training of the student model, we draw on FixMatch [26] which is the latest advancement in semi-supervised image classification task. Strong augmentation is applied for detection training of the student model and weak augmentation is used for pseudo-labeling of the teacher model.

In theory, our framework is applicable to mainstream object detectors, including single-stage object detectors [15, 17, 21, 30] and two-stage object detectors [22, 9, 5, 35, 34]. To allow a fair comparison with previous methods, we use Faster R-CNN [22] as our default detection framework to illustrate our method.

3.2. Soft Teacher

The performance of the detector depends on the quality of the pseudo-label. In practice, we find that using a higher threshold on foreground score to filter out most of the student-generated box candidates with low-confidence can achieve better results than using a lower threshold. As shown in Table. 8, the best performance is achieved when the threshold is set to 0.9. However, while the strict criteria (higher threshold) leads to higher foreground precision, the recall of the retained box candidates also falls off quickly. As shown in Figure. 3 (a), when the foreground threshold is set to 0.9, the recall is low, as 33%, while the precision reaches 89%. In this case, if we use IoU between student generated box candidates and teacher-generated pseudo boxes to assign foreground and background labels, as a general object detection framework does when real box annotations are provided, some foreground box candidates will be mistakenly assigned as negatives, which may hinder the training and harm the performance.

To alleviate this issue, we propose a *soft teacher* approach which leverages richer information from the teacher model, thanks to the flexibility of the end-to-end framework. Specifically, we assess the reliability of each student-generated box candidate to be a real background, which is then used to weigh its background classification loss. Given two box sets $\{b_i^{\text{fg}}\}$ and $\{b_i^{\text{bg}}\}$, with $\{b_i^{\text{fg}}\}$ denoting boxes assigned as foreground and $\{b_i^{\text{bg}}\}$ denoting the boxes assigned as background, the classification loss of an unlabeled image with the reliable weighting is defined as:

$$\mathcal{L}_u^{\text{cls}} = \frac{1}{N_b^{\text{fg}}} \sum_{i=1}^{N_b^{\text{fg}}} l_{\text{cls}}(b_i^{\text{fg}}, \mathcal{G}_{\text{cls}}) + \sum_{j=1}^{N_b^{\text{bg}}} w_j l_{\text{cls}}(b_j^{\text{bg}}, \mathcal{G}_{\text{cls}}), \quad (4)$$

$$w_j = \frac{r_j}{\sum_{k=1}^{N_b^{\text{bg}}} r_k}, \quad (5)$$

where \mathcal{G}_{cls} denotes the set of (teacher-generated) pseudo boxes used for classification, l_{cls} is the box classification

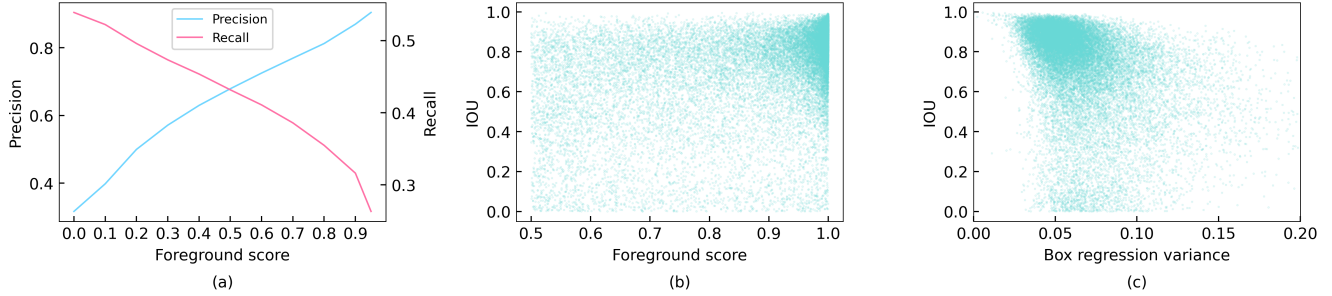


Figure 3. We randomly sampled 10k unlabeled training images from `train2017` to draw figures based on the model trained with 10% labeled images. (a) precision and recall of foreground under different foreground score thresholds. (b) the correlation between the IoU with ground-truth and box foreground score. (c) the correlation between the IoU with ground-truth and box regression variance. Each point in (b) and (c) represents a box candidate.

loss, r_j is the reliability score for j -th background box candidate, N_b^{fg} and N_b^{bg} are the number of box candidates of the box set $\{b_i^{\text{fg}}\}$ and $\{b_i^{\text{bg}}\}$, respectively.

Estimating the reliability score r is challenging. We find empirically that the background score produced by the teacher model with weak augmented image can well serve as a proxy indicator of r and is easily obtained in our end-to-end training framework. Specifically, given a student-generated box candidate, its background score can be obtained simply by using the teacher (BG-T) to process the box through its detection head. It is worth noting that this approach, unlike the widely used hard negative mining approaches, e.g., OHEM [25] or Focal Loss [15], is more like a “simple” negative mining. For comparison, we also examine several other indicators:

- *Background score of student model (BG-S)*: Another natural way to generate the background score is to use the prediction of student model directly.
- *Prediction difference (Pred-Diff)*: The prediction difference between the student model and teacher model is also a possible indicator. In our approach, we simply use the difference between the background scores of the two models to define the reliability score:

$$r = 1 - |p_S^{\text{bg}}(b) - p_T^{\text{bg}}(b)|, \quad (6)$$

where p_S^{bg} and p_T^{bg} are the predicted probability of the background class of the student and the teacher model, respectively.

- *Intersection-over-Union*: The IoU between ground-truths and box candidates is a commonly used criterion for foreground/background assignment. There are two different hypotheses about how to use IoU to measure whether a box candidate belongs to the background. In the first hypothesis, if the IoU between a box candidate and a ground-truth box is less than a threshold (e.g., 0.5), a larger IoU indicates the box candidate has greater probability of being background. This

can be viewed as an IoU-based hard negative mining which is adopted by Fast R-CNN [6] and Faster R-CNN [22] in the early implementation. In contrast, the other hypothesis suggests that box candidates with a smaller IoU with ground-truths are more likely to be backgrounds. In our experiments, we validate both hypotheses and name them as *IoU* and *Reverse-IoU*.

3.3. Box Jittering

As shown in Figure 3 (b), the localization accuracy and the foreground score of the box candidates do not show a strong positive correlation, which means that the boxes with high foreground score may not provide accurate localization information. This indicates that the selection of the teacher-generated pseudo boxes according to the foreground score is not suitable for box regression, and a better criterion is needed.

We introduce an intuitive approach to estimate the localization reliability of a candidate pseudo box by measuring the consistency of its regression prediction. Specifically, given a teacher-generated pseudo box candidate b_i , we sample a jittered box around b_i and feed the jittered box into the teacher model to obtain the refined box \hat{b}_i , which is formulated as follows:

$$\hat{b}_i = \text{refine}(\text{jitter}(b_i)). \quad (7)$$

The above procedure is repeated several times to collect a set of N_{jitter} refined jittered boxes $\{\hat{b}_{i,j}\}$, and we define the localization reliability as the box regression variance:

$$\bar{\sigma}_i = \frac{1}{4} \sum_{k=1}^4 \hat{\sigma}_k, \quad (8)$$

$$\hat{\sigma}_k = \frac{\sigma_k}{0.5(h(b_i) + w(b_i))}, \quad (9)$$

where σ_k is the standard derivation of the k -th coordinate of the refined jittered boxes set $\{\hat{b}_{i,j}\}$, $\hat{\sigma}_k$ is the normalized σ_k , $h(b_i)$ and $w(b_i)$ represent the height and width of box candidate b_i , respectively.

A smaller box regression variance indicates a higher localization reliability. However, computing the box regression variances of all pseudo box candidates is unbearable during training. Therefore, in practice, we only calculate the reliability for the boxes with a foreground score greater than 0.5. In this way, the number of boxes that need to be estimated is reduced from an average of hundreds to around 17 per image and thus the computation cost is almost negligible.

In Figure 3 (c), we illustrate the correlation between the localization accuracy and our box regression variance. Compared with the foreground score, the box regression variance can better measure the localization accuracy. This motivates us to select box candidates whose box regression variance is smaller than a threshold as pseudo-label to train the box regression branch on unlabeled images. Given the pseudo boxes \mathcal{G}_{reg} for training the box regression on unlabeled data, the regression loss is formulated as:

$$\mathcal{L}_u^{\text{reg}} = \frac{1}{N_b^{\text{fg}}} \sum_{i=1}^{N_b^{\text{fg}}} l_{\text{reg}}(b_i^{\text{fg}}, \mathcal{G}_{\text{reg}}), \quad (10)$$

where b_i^{fg} is i -th box assigned as foreground, N_b^{fg} is the total number of foreground box, l_{reg} is the box regression loss. Therefore, by substituting Equ. 4 and Equ. 10 into Equ. 3, the loss of unlabeled images is:

$$\mathcal{L}_u = \frac{1}{N_u} \sum_{i=1}^{N_u} (\mathcal{L}_u^{\text{cls}}(I_u^i, \mathcal{G}_{\text{cls}}^i) + \mathcal{L}_u^{\text{reg}}(I_u^i, \mathcal{G}_{\text{reg}}^i)). \quad (11)$$

Here we use the pseudo boxes \mathcal{G}_{cls} and \mathcal{G}_{reg} as the inputs of the loss to highlight the fact that the pseudo boxes used in classification and box regression are different in our approach.

4. Experiments

4.1. Dataset and Evaluation Protocol

We validate our method on the MS-COCO benchmark [16]. Two training datasets are provided, the `train2017` set contains 118k labeled images and the `unlabeled2017` set contains 123k unlabeled images. In addition, the `val2017` set with 5k images is also provided for validation. In previous methods [27, 28, 10], there are two settings for validating the performance:

Partially Labeled Data: STAC [27] first introduced this setting. 1%, 5% and 10% images of `train2017` set are sampled as the labeled training data, and the remaining unsampled images of `train2017` are used as the unlabeled data. For each protocol, STAC provides 5 different data folds and the final performance is the average of all 5 folds.

Fully Labeled Data: In this setting, the entire `train2017` is used as the labeled data and `unlabeled2017` is used as the additional unlabeled




Figure 4. The qualitative results of our method. (a), (c) are the results of the supervised baseline. (b), (d) are the results of our method.

data. This setting is more challenging. Its goal is to use the additional unlabeled data to improve a well-trained detector on large-scale labeled data.

We evaluate our method on both settings and follow the convention to report the performance on `val2017` with the standard mean average precision (mAP) as the evaluation metrics.

4.2. Implementation Details

We use the Faster R-CNN [22] equipped with FPN [14] (Feature Pyramid Network) as our default detection framework to evaluate the effectiveness of our method, and an ImageNet pre-trained ResNet-50 [8] is adopted as the backbone. Our implementation and hyper-parameters are based on MMDetection [4]. Since the amount of training data of **Partially Labeled Data** setting and **Full Labeled Data** setting has large differences, the training parameters under the two settings are slightly different, and we refer to more details in the supplementary material.

4.3. System Comparison

In this section, we compare our method with previous state-of-the-arts on MS-COCO. We first evaluate on the **Partially Labeled Data** setting and compare our method with STAC. For benchmarking, we compare the supervised baseline of our method with the results reported in STAC and find they perform similarly, the results are shown in Table 2. In this case, we further compare our method with STAC at the system level, and our method shows a significant performance improvement in different protocols. Specifically, our method outperforms the STAC by **6.5** points, **6.4** points and **5.4** points when there are 1%, 5%, and 10% labeled data, respectively. The qualitative results of our method compared with supervised baseline are shown in Figure 4.

Then we compare our method with other state-of-the-

Method	1%	5%	10%
Supervised baseline (Ours)	10.0 ± 0.26	20.92 ± 0.15	26.94 ± 0.111
Supervised baseline (STAC) [27]	9.83 ± 0.23	21.18 ± 0.20	26.18 ± 0.12
STAC [27]	13.97 ± 0.35	24.38 ± 0.12	28.64 ± 0.21
Ours	20.46±0.39	30.74±0.08	34.04±0.14

Table 2. System level comparison with STAC on `val2017` under the **Partially Labeled Data** setting. All the results are the average of all 5 folds. For benchmarking, we also compare the supervised benchmark performance between our method and STAC, and their performance is similar.

Method	Extra dataset	mAP
Proposal learning [28]	unlabeled2017	37.4 $\xrightarrow{+1.0}$ 38.4
STAC [27]	unlabeled2017	39.5 $\xrightarrow{-0.3}$ 39.2
Self-training [36]	ImageNet+OpenImages	41.1 $\xrightarrow{+0.8}$ 41.9
Ours	unlabeled2017	40.9 $\xrightarrow{+3.6}$ 44.5

Table 3. Comparison with other state-of-the-arts under the setting of using all data of `train2017` set. Particularly, Self-training uses ImageNet (1.2M images) and OpenImages (1.7M images) as additional unlabeled images, which is 20× larger than `unlabeled2017` (123k images).

art methods in **Fully Labeled Data** setting. Since the reported performance of supervised baseline varies in different works, we report the results of the comparison methods and their baseline at the same time. The results are shown in Table. 3.

We first compare with the Proposal Learning [28] and STAC [27] which also use `unlabeled2017` as additional unlabeled data. Because of the better hyper-parameters and more adequate training, our supervised baseline achieved better performance than other methods. Under the stronger baseline, our method still shows a greater performance gain (+3.6 points) than Proposal Learning (+1.0 points) and STAC (-0.3 points). Self-training [36] uses ImageNet (1.2M images) and OpenImages (1.7M images) as the additional unlabeled data, which is 20× larger than the `unlabeled2017` (123k images) that we use. With similar baseline performance, our method also shows better result with less unlabeled data.

In addition, we further evaluate our method on other stronger detectors, and the results evaluated on `val2017` set are shown in Table. 4. Our method consistently improves the performance of different detectors by a notable margin. Even in the state-of-the-art detector HTC++ with Swin-L backbone, we still show 1.8 improvement on detection AP and 1.4 improvement on mask AP. Moreover, we also report the results on `test-dev2017` set. As shown in Tabel. 1, our method improves the HTC++ with Swin-L backbone by 1.5 mAP on detection, which is the first work to surpass 60 mAP on COCO object detection benchmark.

4.4. Ablation Studies

In this section, we validate our key designs. If not specified, all the ablation experiments are performed on the single data fold provided by [27] with 10% labeled images

from `train2017` set.

Multi-Stage vs. End-to-End. We compare our end-to-end method with the multi-stage framework as shown in Table 5. By simply switching from the multi-stage framework to our end-to-end framework, performance is increased by 1.3 points. By updating the teacher model with the student model through the exponential moving average (EMA) strategy, our method further achieves 31.2 mAP.

Effects of Soft Teacher and Box Jittering. We ablate the effects of soft teacher and box jittering. The results are shown in Table. 6. Based on our end-to-end model equipped with EMA (E2E+EMA), integrating the soft teacher improves the performance by 2.4 points. Further applying the box jittering, the performance reaches 34.2 mAP, which is 3 points better than E2E+EMA.

Different Indicators in Soft Teacher. In Section. 3.2, several different indicators are explored for reliability estimation. Here, we evaluate the different indicators and the results are shown in Table. 7. The background score predicted by the teacher model achieves the best performance. Simply switching the model from teacher to student will make the performance worse. In addition, the improvement of IoU and Reverse-IoU is negligible compared with BG-T. These results prove the necessity of leveraging the teacher model.

Effects of other hyper-parameters. We study the effects of hyper-parameters used in our method. Table. 8 studies the effects of different foreground score thresholds. The best performance is achieved when the threshold is set to 0.9, and lower or higher thresholds will cause significant performance degradation. In Table. 9, we study the box regression variance threshold. The best performance is

detector	backbone	method	mAP ^{det}	mAP ^{mask}
Faster R-CNN	ResNet-50	supervised	40.9	-
		ours	44.5(+3.6)	-
Faster R-CNN	ResNet-101	supervised	43.8	-
		ours	46.8(+3.0)	-
HTC++	Swin-L	supervised	57.1	49.6
		ours	59.1(+2.0)	51.0(+1.4)
HTC++(multi-scale)	Swin-L	supervised	58.2	50.5
		ours	59.9(+1.7)	51.9(+1.4)

Table 4. Compared with various supervised trained detectors on val₂₀₁₇. The entire train₂₀₁₇ is used as the labeled images, and the unlabeled₂₀₁₇ is used as the additional unlabeled images.

Method	mAP	mAP@0.5	mAP@0.75
Supervised	27.1	44.6	28.6
Multi-Stage	28.7	47.0	30.9
E2E	30.0	47.4	32.4
E2E+EMA	31.2	48.8	34.0

Table 5. Multi-Stage vs. End-to-End. The end-to-end (E2E) method outperforms the multi-stage framework. Updating the teacher network through the exponential moving average (EMA) strategy further improves the performance.

Soft teacher	Box jittering	mAP	mAP@0.5	mAP@0.75
		31.2	48.8	34.0
✓		33.6	52.9	36.6
✓	✓	34.2	52.6	37.3

Table 6. We study the effects of *soft teacher* and *box jittering* techniques.

Indicator	mAP	mAP@0.5	mAP@0.75
w/o weight	31.2	48.8	34.0
IoU	31.7	51.4	34.2
Reverse-IoU	31.6	49.5	34.1
Pred-Diff	32.3	51.0	34.6
BG-S	25.9	44.4	27.0
BG-T	33.6	52.9	36.6

Table 7. Comparison of different indicators in soft teacher.

Threshold	mAP	mAP@0.5	mAP@0.75
0.70	29.9	48.6	32.1
0.80	33.2	52.8	35.9
0.90	33.6	52.9	36.6
0.95	32.1	50.6	34.7

Table 8. Ablation study on the effects of different foreground thresholds.

achieved when the threshold is set to 0.02. In Table. 10, we study the effects of different number of jittered boxes, and the performance is saturated when N_{jitter} is set to 10.

Threshold	mAP	mAP@0.5	mAP@0.75
0.04	33.8	52.3	36.7
0.03	34.0	52.5	36.9
0.02	34.2	52.6	37.3
0.01	32.9	52.2	35.8

Table 9. Ablation study on the effects of different thresholds for selecting pseudo boxes for box regression according to box regression variance.

N_{jitter}	mAP	mAP@0.5	mAP@0.75
5	34.0	52.3	37.0
10	34.2	52.6	37.3
15	34.2	52.5	37.4

Table 10. Ablation study on the effects of different number of jittered boxes used to estimate the box regression variance.

5. Conclusion

In this paper, we proposed an end-to-end training framework for semi-supervised object detection, which discards the complicated multi-stage schema adopted by previous approaches. Our method simultaneously improves the detector and pseudo labels by leveraging a student model for detection training, and a teacher model which is continuously updated by the student model through the exponential moving average strategy for online pseudo-labeling. Within the end-to-end training, we present two simple techniques named soft teacher and box jittering to facilitate the efficient leverage of the teacher model. The proposed framework outperforms the state-of-the-art methods by a large margin on MS-COCO benchmark in both partially labeled data and fully labeled data settings.

6. Acknowledgement

We would like to thank Yue Cao for his valuable suggestions and discussions; Yutong Lin and Yixuan Wei for help on Swin Transformer experiments.

This work was supported by National Key R&D Program of China (No. 2018YFB1004600)

References

- [1] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. *arXiv preprint arXiv:1412.4864*, 2014. 2, 3
- [2] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019. 3
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*, 2019. 3
- [4] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 6
- [5] Yihong Chen, Zheng Zhang, Yue Cao, Liwei Wang, Stephen Lin, and Han Hu. Reppoints v2: Verification meets regression for object detection. *NIPS*, 2020. 4
- [6] Ross Girshick. Fast r-cnn. In *ICCV*, 2015. 3, 5
- [7] Yves Grandvalet, Yoshua Bengio, et al. Semi-supervised learning by entropy minimization. In *CAP*, 2005. 3
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 6
- [9] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *CVPR*, 2018. 4
- [10] Jisoo Jeong, Seungeui Lee, Jeessoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. *NIPS*, 2019. 3, 6
- [11] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *ICLR*, 2016. 2, 3
- [12] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop of ICML*, 2013. 3
- [13] Yandong Li, Di Huang, Danfeng Qin, Liqiang Wang, and Boqing Gong. Improving object detection with selective self-supervised self-training. In *ECCV*, 2020. 3
- [14] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 2, 3, 6
- [15] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 4, 5
- [16] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 6
- [17] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 3, 4
- [18] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. 2
- [19] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *TPAMI*, 2018. 2, 3
- [20] Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omniscient supervised learning. In *CVPR*, 2018. 3
- [21] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 3, 4
- [22] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 2, 3, 4, 5, 6
- [23] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *arXiv preprint arXiv:1606.04586*, 2016. 2, 3
- [24] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *CVPR*, 2019. 2
- [25] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016. 5
- [26] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *NIPS*, 2020. 3, 4
- [27] Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020. 1, 2, 3, 6, 7
- [28] Peng Tang, Chetan Ramaiah, Yan Wang, Ran Xu, and Caiming Xiong. Proposal learning for semi-supervised object detection. In *WACV*, 2021. 3, 6, 7
- [29] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NIPS*, 2017. 3, 4
- [30] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, 2019. 3, 4
- [31] Keze Wang, Xiaopeng Yan, Dongyu Zhang, Lei Zhang, and Liang Lin. Towards human-machine cooperation: Self-supervised sample mining for object detection. In *CVPR*, 2018. 3
- [32] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. *NIPS*, 2020. 3
- [33] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, 2020. 3

- [34] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *ICCV*, 2019. [3](#), [4](#)
- [35] Ze Yang, Yinghao Xu, Han Xue, Zheng Zhang, Raquel Urtasun, Liwei Wang, Stephen Lin, and Han Hu. Dense reppoints: Representing visual objects with dense point sets. *ECCV*, 2019. [3](#), [4](#)
- [36] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin D Cubuk, and Quoc V Le. Rethinking pre-training and self-training. *NIPS*, 2020. [1](#), [2](#), [3](#), [7](#)