

Dense Teacher: Dense Pseudo-Labels for Semi-supervised Object Detection

Hongyu Zhou^{1,3*}, Zheng Ge^{1*}, Songtao Liu¹, Weixin Mao^{1,2},
Zeming Li¹, Haiyan Yu³, and Jian Sun¹

¹MEGVII Technology ²Waseda University

³Harbin Institute of Technology

Abstract. To date, the most powerful semi-supervised object detectors (SS-OD) are based on pseudo-boxes, which need a sequence of post-processing with fine-tuned hyper-parameters. In this work, we propose replacing the sparse pseudo-boxes with the dense prediction as a united and straightforward form of pseudo-label. Compared to the pseudo-boxes, our Dense Pseudo-Label (DPL) does not involve any post-processing method, thus retaining richer information. We also introduce a region selection technique to highlight the key information while suppressing the noise carried by dense labels. We name our proposed SS-OD algorithm that leverages the DPL as Dense Teacher. On COCO and VOC, Dense Teacher shows superior performance under various settings compared with the pseudo-box-based methods. Code is available at <https://github.com/Megvii-BaseDetection/DenseTeacher>

Keywords: Semi-supervised Object Detection, Dense Pseudo-Label

1 Introduction

Current high-performance object detection neural networks rely on a large amount of labeled data to ensure their generalization capability. However, labeling samples takes a high cost of human effort. Thus the industry and academia pay extensive attention to the use of relatively easy-to-obtain unlabeled data. An effective way to use these data is Semi-Supervised Learning (SSL), where at the training time, only part of the data is labeled while the rest are unlabeled. On image classification tasks, the dominated method of mining information from unlabeled data is “Consistency-based Pseudo-Labeling” [5,23,4,3]. Pseudo-Labeling [12] is a technique that utilizes trained models to generate labels for unlabeled data. Meanwhile, the Consistency-based regularization [1], from another perspective, forces a model to have similar output when giving a normal and a perturbed input with different data augmentations and perturbations like Dropout [25].

This pipeline has been successfully transferred to Semi-Supervised Object Detection (SS-OD) [24,18,30]. Specifically, the predicted boxes from a pre-trained “teacher” detector are used as the annotations of unlabeled images to train the

* Authors contributed equally to this work.

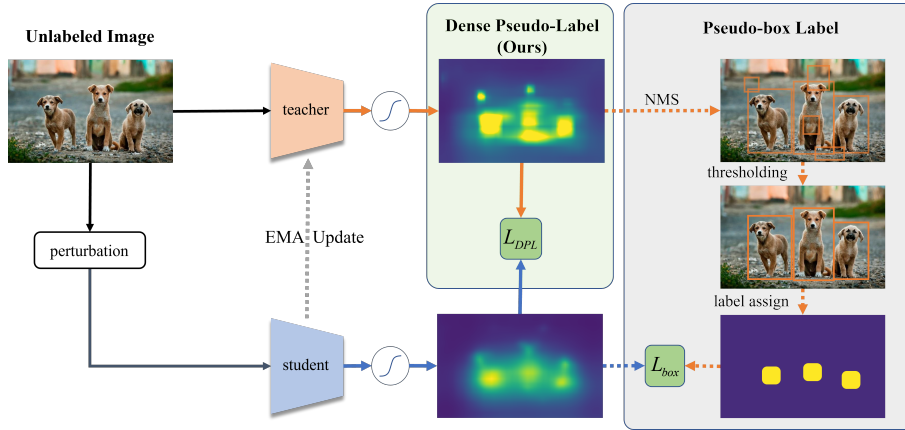


Fig. 1: The overview of our purposed pipeline for unlabeled data compared with traditional pseudo-box based pipeline. For each iteration, Dense Pseudo-Label (DPL) are generated by the teacher model on unlabeled images. The student model then calculates unsupervised loss on perturbed images and corresponding DPLs. By removing post-processing steps, DPL contains rich information from the teacher model. Note that the vanilla learning approach uses only labeled data (not plotted in the figure) and the total loss is the sum of both supervised and unsupervised loss.

“student” detector, where the same images are applied with different augmentations for the teacher and student model. This instinctive method has proven to be effective in SS-OD and has achieved state-of-the-art scores on benchmarks such as COCO [17] and Pascal VOC [6]. However, it is not reasonable to replicate all the empirics directly from the classification task. While the generated pseudo-label is a single and united class label for an image in classification, the object detectors predict a set of pseudo-boxes as the annotation of an image. As shown in Fig. 1, making direct supervision on the unlabeled image with these pseudo-boxes requires several additional steps, including Non-Maximum-Suppression (NMS), Thresholding, and Label Assignment. Such a lengthy label-generating procedure introduces many hyper-parameters, such as NMS threshold σ_{nms} and score threshold σ_t , substantially affecting the SS-OD performance.¹ This motivates us to explore a more simple and effective form of pseudo-labels for SS-OD.

In this work, we propose a new SS-OD pipeline named Dense Teacher with a *united* form of pseudo-label — Dense Pseudo-Label (DPL), which enables more efficient knowledge transfer between the teacher and student models. DPL is an integral label. Different from the existing box-like labels in a human-readable form, it is the original output from the network without ordinary post-processing. Our Dense Teacher, following existing Pseudo-Labeling paradigms, works in the following way: For each iteration, labeled and unlabeled images are randomly

¹ See also in Sec. 3.2 for a related discussion.

sampled to form a data batch. The “teacher” model, which is an Exponential Moving Average of the “student” model, generates DPL for unlabeled data. And the student model is trained on both ground truth on labeled images and DPL on unlabeled images. Since the DPL does not require any post-processing, the pipeline of Dense Teacher is extremely simple. The overall pipeline of Dense Teacher can be seen in Fig. 1.

Although DPL provides more information than pseudo-box labels, it contains high-level noise (*e.g.*, low-scoring predictions) as well. We show in Sec. 4.2 that learning to make those low-scoring predictions can distract the student model, resulting in poor detection performance. Therefore, we propose a region division method to suppress noise and highlight key regions where the student model should concentrate on. According to our experiments, the region division strategy can effectively utilize the rich information contained in hard negative regions to enhance training. As a result, our proposed Dense Teacher, together with the region division strategy, shows state-of-the-art performance on MS-COCO and Pascal VOC.

Our main contributions in this paper are:

- We conduct a thorough analysis on the drawbacks of pseudo-boxes in the SS-OD task.
- We propose a united form of pseudo-label named DPL to better fit the semi-supervised setting and the Dense Teacher framework to apply the DPL on one-stage detector.
- The proposed Dense Teacher achieves state-of-the-art performance on MS-COCO and Pascal VOC benchmarks under various settings. Gain analysis and ablation study are provided to verify the effectiveness of each part in Dense Teacher.

2 Related Works

2.1 Semi-Supervised Learning

Semi-Supervised Learning (SSL) means that a portion of the data is labeled at training time while the other is not. Currently, there are two main approaches to achieve this goal: pseudo-labeling and consistency regularization. Pseudo-label-based methods [12] first train a network on labeled data, then use the trained network as a teacher to make inferences on unlabeled data. This prediction result is then assigned to a specific class according to a threshold based on the predicted confidence and used as labeled data to train another student network. In [27], the teacher model is replaced by the Exponential Moving Average (EMA) of the student to conduct online pseudo-labeling. Consistency regularization based methods [1] construct a regularization loss to force predictions under a set of perturbations $\{T_i\}$ to be same. Perturbations can be implemented using augmentation [5, 21, 29], dropout [11]), or adversarial training [19]. This approach does not require annotation and can be used in combination with other methods; therefore, it is widely adopted in many SSL frameworks [2, 23, 4].

2.2 Object Detection

Object detectors can be divided into Anchor-based and Anchor-free paradigms. Anchor-based detectors predict the offsets and scales of target boxes from predefined anchor boxes. Although this approach has succeeded on many tasks, one needs to redefine new anchor boxes when applying such models to new data. In contrast to anchor-based detectors, predefined anchor boxes are not required for anchor-free detectors. These detectors directly predict the box size and location on the feature map. Take FCOS model as an example; this detector predicts the classification score, distances to four boundaries, and a quality score on each pixel of Feature Pyramid Network (FPN) [15]. A variety of subsequent improvements, such as the adaptive label assigning while training [33], boundary distribution modeling [14] were proposed to improve its performance. Considering the wide application, streamlined architecture, and excellent performance of FCOS, we will conduct our experiments under this framework.

2.3 Semi-Supervised Object Detection

The label type is the main difference between Semi-Supervised Object Detection (SS-OD) and SSL. Previous studies have transferred a great deal of experience from SSL works to the SS-OD domain. CSD [10] use a flipped image I' to introduce a consistency loss between $F(I)$ and $F(I')$, this regularization can be applied to unlabeled image. STAC [24] train a teacher detector on labeled images and generate pseudo-labels on unlabeled data using this static teacher. These pseudo-labels will then be selected and used for training like labeled data. Unbiased Teacher [18] use thresholding to filter pseudo-labels, Focal Loss [16] is also applied to address the pseudo-labeling bias issue. Adaptive Class-Rebalancing [32] artificially adds foreground targets to images to achieve inter-class balance. *Li, et al.* [13] propose dynamic thresholding and loss re-weighting for each category. Soft Teacher [30] proposed a score-weighted classification loss and box jittering approach to select and utilize regression loss of pseudo-boxes. While these methods successfully transferred paradigms from SSL to SS-OD, they ignored the unique characteristics of SS-OD. These pseudo-box-based strategies treat pseudo-boxes, or selected pseudo-boxes, as ordinary target boxes, and thus they invariably follow the detector’s label assign strategy.

3 Dense Teacher

In Sec. 3.1, we first introduce the existing Pseudo-Labeling SS-OD framework. Then we analyze the disadvantages of utilizing pseudo-boxes in Sec. 3.2. In the remaining part, we propose Dense Pseudo-Label to overcome the issues mentioned above and introduce our overall pipeline in detail, including the label generation, loss function, and the learning region selection strategy. Since our primary motivation is to show the superiority of dense pseudo-labels compared to pseudo-box labels, we naturally choose to verify our idea on dense detectors (*i.e.*, one-stage detectors).

3.1 Pseudo-Labeling Framework

Our Dense Teacher follows the existing pseudo-labeling framework [18,30] as shown in Fig. 1. Within each iteration:

1. Labeled and unlabeled images are randomly sampled to form a data batch.
2. The teacher model, an exponential moving average (EMA) of the student, takes the augmented unlabeled images to generate pseudo-labels.
3. The student model then takes the labeled data for vanilla training and calculates supervised loss \mathcal{L}_s , while the unlabeled data together with pseudo-labels are used to produce unsupervised loss \mathcal{L}_u .
4. Two losses are weighted and learned to update parameters of the student model. The student model updates the teacher model in an EMA manner.

Finally, the overall loss function is defined as:

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

where w_u is the unsupervised loss weight. Traditionally, the unsupervised loss \mathcal{L}_u is calculated with pseudo-boxes. However, in the following section, we point out that using processed boxes as pseudo-labels can be inefficient and sub-optimal.

3.2 Disadvantages of Pseudo-box Labels

In this part, we study the behavior of pseudo-box-based SS-OD algorithms on COCO [17], as well as CrowdHuman² [22] since the impact of the NMS threshold can be more clearly demonstrated in the crowd situation. We adopt Unbiased Teacher [18] as a representative algorithm to FCOS for these experiments.

Dilemma in Thresholding In SS-OD algorithms [18,30], the output of the teacher model is expected to play the role of ground-truth labels for unsupervised images. To this end, Thresholding is a key operation to screen out low-scoring boxes so that the quality of pseudo-box labels can be improved. However, our preliminary experiments show that the threshold σ_t introduced by this operation may substantially affect the entire training process. In Fig. 2(a), we present the training results of Unbiased Teacher under different σ_t . It shows that the detection performance fluctuates significantly on both datasets as the σ_t varies. Moreover, when σ_t is set to a high value (*e.g.*, 0.7 and 0.9), the training process even fails to converge. Such a phenomenon is possibly caused by a large number of false negatives in the teacher’s prediction, as shown in Fig. 2(c) and (d). When this is the case, Thresholding will eliminate many high-quality predictions and mislead the learning process of the student model. Conversely, when set σ_t to a low value such as 0.3, the performance shows apparent degradation due to the increasing number of false positives (see in Fig. 2(c) and (d) as well). As a result, one can not find a perfect threshold to ensure the quality of generated pseudo-boxes.

² CrowdHuman is a benchmark for detecting humans in a crowded situation, performance is measured by Log-average Miss Rate (mMR). The lower the better.

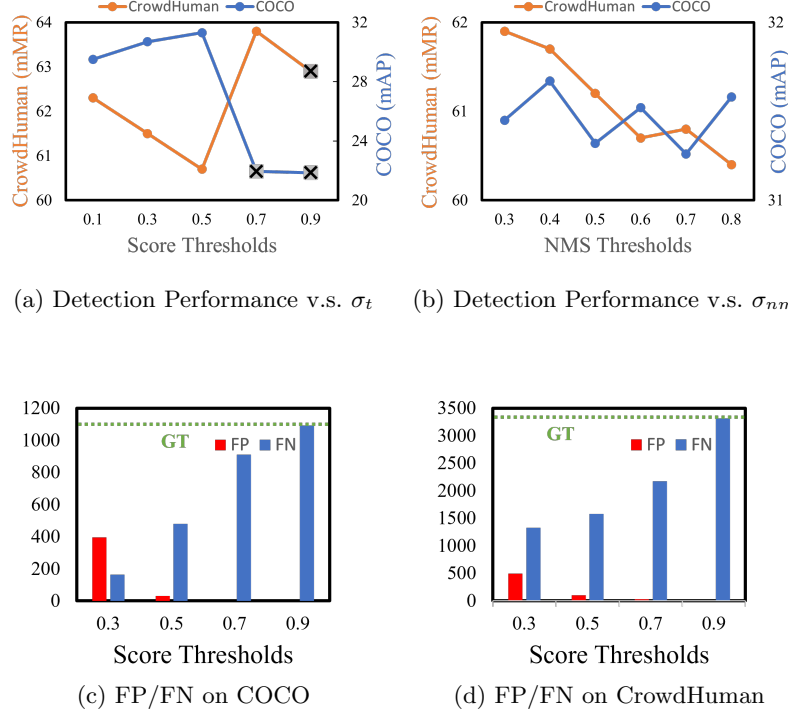


Fig. 2: Analysis of Pseudo-box based approaches. (a) and (b): Performances under different σ_t and σ_{NMS} . Note that the gray \boxtimes represents the training fails to converge. (c) and (d): False Positive and False Negative boxes on 128 images under different threshold on COCO and CrowdHuman, the green line denotes the ground truth box number

Dilemma in Non-Maximum Suppression (NMS) NMS is adopted on the detector’s original outputs for most object detection algorithms to remove redundant predictions. It is also indispensable to the teacher model in existing SS-OD frameworks, without which the resulting pseudo-labels will be a mess. NMS introduces a threshold σ_{nms} to control the degree of suppression. According to our experiments, we find that σ_{nms} also has a non-negligible effect on the SS-OD algorithms. Fig. 2(b) shows the relationship between σ_{nms} and performance of Unbiased Teacher. From this figure, we can tell that 1). different σ_{nms} may lead to fluctuations in the detection performance (especially on CrowdHuman). 2). the optimal σ_{nms} values for different datasets are different (*i.e.*, 0.7 on COCO and 0.8 on CrowdHuman), which will bring in extra workload for developers to tune the optimal σ_{nms} on their custom datasets. Moreover, previous works [8, 7] show that in a crowd scene like in the CrowdHuman dataset, there does not exist

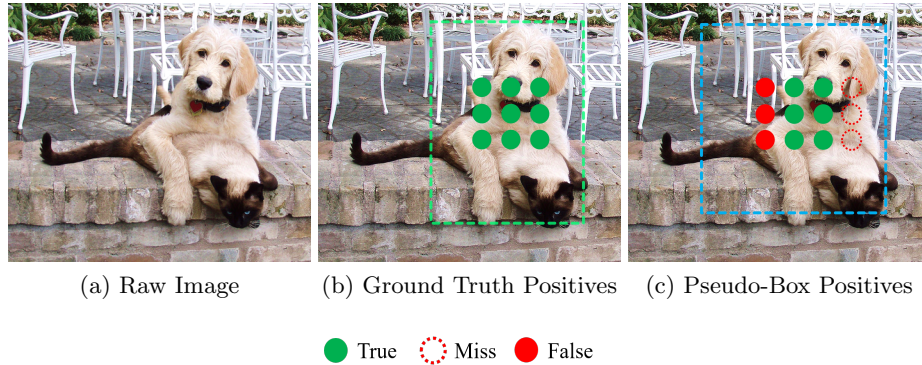


Fig. 3: Comparisons between (b) foreground pixels assigned by ground truth boxes and (c) foreground pixels assigned by pseudo-boxes

a perfect σ_{nms} that can keep all true positive predictions while suppressing all false positives. As a result, with NMS adopted, the unreliability of pseudo-box labels is further exacerbated.

Inconsistent Label Assignment As shown in Fig. 1, existing pseudo-label-based algorithms convert the sparse pseudo-boxes into a dense form by label assignment to form the final supervision. An anchor box (or point) will be assigned as either positive or negative during label assignment based on a particular pre-defined rule. Although this process is natural in the standard object detection task, we believe it is harmful to SS-OD tasks. The reason is quite simple: the pseudo-boxes may suffer from the inaccurate localization problem, making the label assigning results inconsistent with the potential ground-truth labels. In Fig. 3, we can find that although the predicted box matches the actual box under IoU threshold 0.5, a severe inconsistent assigning result appears due to the inaccurate pseudo-box. This inconsistency with the ground truth is likely to degrade the performance.

Due to the above three issues, we challenge the convention of using pseudo-box as the middle-ware of unsupervised learning and propose a new form of pseudo-label that is dense and free of post-processing.

3.3 Dense Pseudo-Label

To address the problems mentioned above, we propose Dense Pseudo-Label (DPL) that encompasses richer and undistorted supervising signals. Specifically, we adopt the post-sigmoid logits predicted by the trained model as our desired dense pseudo-label, as shown in the green box in Fig. 1. After bypassing those lengthy post-processing methods, one can naturally discover that our proposed

DPL reserves more detailed information from the teacher than its pseudo-box counterpart.

Since DPL represents information in continuous values (value between 0 and 1) and the standard Focal Loss [16] can only deal with discrete binary values (0 or 1), we adopt Quality Focal Loss [14] to conduct learning between dense pseudo-labels and the student’s predicting results. Let us denote $\vec{y}_i \triangleq \vec{p}_i^t$ as DPL (*i.e.*, teacher’s prediction) and denote \vec{p}_i^s as student’s prediction for i -th anchor³, we hope the prediction and the target to be similar for the same anchor. Therefore, we can write the classification loss on the i -th anchor for an unlabeled image as:

$$\mathcal{L}_i^{cls} = -|\vec{y}_i - \vec{p}_i^s|^\gamma * [\vec{y}_i \log(\vec{p}_i^s) + (1 - \vec{y}_i) \log(1 - \vec{p}_i^s)] \quad (2)$$

where γ is the suppression factor.

While DPL contains rich information, it also keeps many low-scoring predictions due to the absence of the thresholding operation. Since those low-scoring predictions usually involve the background regions, intuitively, the knowledge encompassed in them shall be less informative. In Sec. 4.4, we experimentally prove that learning to mimic the teacher’s response in those regions will hurt the SS-OD algorithm’s performance. Therefore, we propose to divide the whole input image into a learning region and a suppressing region(*e.t.*, negative region in positive-negative division) based on the teacher’s Feature Richness Score (FRS [34]). With the help of this richness score, we select the pixels with top $k\%$ scores as the learning region and the other regions will be suppressed to 0. As a result, our DPL is extended to:

$$S_i = \max_{c \in [1, C]} (p_{i,c}^t) \quad (3)$$

$$\vec{y}_i = \begin{cases} \vec{p}_i^t, & \text{if } S_i \text{ in top } k\%, \\ \vec{0}, & \text{otherwise.} \end{cases} \quad (4)$$

where $p_{i,c}^t$ denotes the score prediction of c -th class for i -th sample from the teacher, C denotes the total number of classes.

Besides, this design has other advantages:

1. By modifying the learning region, we can easily achieve Hard Negative Mining by selecting extra samples (see Fig. 4). In Sec. 4.4 we will analyze the gain from this part in detail.
2. Since the learning region is selected, unsupervised learning for regression branch can be easily achieved. We apply IoU Loss on this branch and analyze its gain in Sec. 4.2

³ “Anchor” stands for “anchor point” in anchor-free detectors and “anchor box” in anchor-based detectors.

4 Experiments

4.1 Datasets and Experiment Settings

Datasets. We present our experimental results on MS-COCO [17] and Pascal VOC [6] benchmarks. For MS-COCO, both labeled and unlabeled training datasets will be used. The `train2017` set contains 118k images with target bounding boxes and the `unlabeled2017` contains 123k unlabeled images. Validation is performed on the subset `val2017`. For Pascal VOC, training set uses `VOC07 train` and `VOC12 train` and validation set uses `VOC07 test`. The following three experimental settings are mainly studied:

- **COCO-Standard:** 1%, 2%, 5% and 10% of the `train2017` set are sampled as labeled data, respectively. The rest of images are viewed as unlabeled data while training. For fairness of comparison, we follow the same dataset division as in [18] which contains 5 different data folds. Mean score of all 5 folds are taken as the final performance.
- **COCO-Full:** `train2017` is used as labeled data while `unlabeled2017` is used as unlabeled data.
- **VOC Mixture:** `VOC07 train` is used as labeled data, while `VOC12 train` and `COCO20cls`⁴ are taken as unlabeled data.

Implementation Details. Without loss of generality, we take FCOS [28] as the representative anchor-free detector for experiments. ResNet-50 [9] pre-trained on ImageNet [20] is used as the backbone. We use batch-size 16 for both labeled and unlabeled images. The base learning rate and γ in QFL are set to 0.01 and 2 in all of our experiments. Loss weight w_u on unlabeled data is set to 4 on COCO-Standard and 2 on the other settings. Following previous works [18,24], we adopt “burnin” strategy to initialize the teacher model, same data augmentations as in [18] are applied.

4.2 Main Results

In this section, we progressively improve the Dense Teacher and analyze the performance gain from each part in detail. We adopt Unbiased Teacher as our baseline. Results on the COCO-Standard 10% setting are shown in Table. 1. We first replace Unbiased Teacher’s pseudo-boxes with our proposed Dense Pseudo-Labels without the region division strategy. It shows that this improves the mAP from 31.52% to 32.0%. Then, we apply our region division strategy on DPL and the mAP is further improved by 1.34%. Finally, we extend the unsupervised learning scheme to the regression branch as done by [30], and our final mAP comes to 35.11%. To the best of our knowledge, this is the new state-of-the-art under the COCO-Standard 10% setting. According to these results, we can

⁴ `COCO20cls` is the sampled COCO `train2017` set, only 20 classes same as in VOC are included.

Table 1: Performance under different model configurations on COCO-Standard 10%. * denotes our re-implemented result on FCOS. “Our Division” means the learning/suppression region division based on FRS score

Method	Learning Region	Cls	Reg	AP	AP50	AP75
Supervised	-	-	-	26.44	42.69	28.11
Unbiased Teacher* [18]	Predicted Positive	✓	×	31.52	48.80	33.57
Dense Teacher	All	✓	×	32.00	50.29	34.17
Dense Teacher	Our Division	✓	×	33.34	52.14	35.53
Unbiased Teacher* [18]	Predicted Positive	✓	✓	33.13	49.96	35.36
Dense Teacher	Our Division	✓	✓	35.11	53.35	37.79

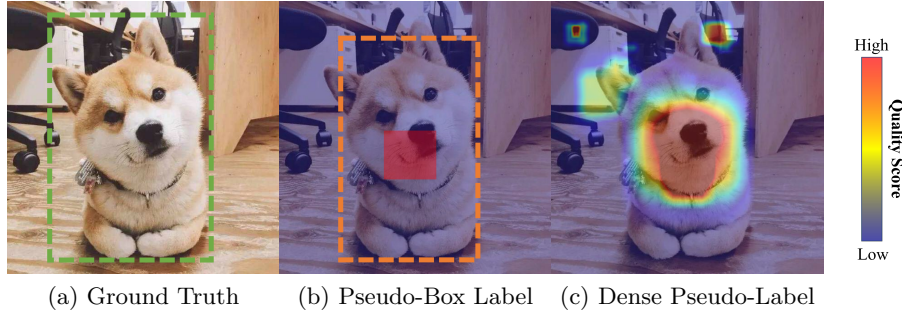


Fig. 4: Illustration of (c) our Dense Pseudo-Label compared with (b) pseudo-box label and its assigning result. Blue areas denote the assigned negative samples. In Dense Pseudo-Label, red means high quality scores, which denotes positive samples in pseudo-box label. It can be seen that Dense Pseudo-Label is able to leverage more hard negative regions compared to the Pseudo-Box based method

attribute the advantages of Dense Teacher over existing methods to two major improvements:

1). The new form of pseudo-label resolves the deterioration problem of the pseudo-box label as mentioned in Sec. 3.2. It is worth mentioning that by getting rid of the lengthy post-processing procedure, our Dense Teacher forms a much simpler SS-OD pipeline but still with better performance. However, the resulting improvement (from 31.52% to 32.00%) remains marginal without the advanced learning region division strategy.

2). Our region division strategy can efficiently utilize hard negative regions to enhance training. Specifically, we conduct label assignment on FCOS using ground truth annotation of COCO, finding that there are only about 0.4% of positive samples in the COCO train2017 set. By specifying $k = 1$, we take a fair amount of hard negative samples for unsupervised training. In Fig. 4, we can see that hard negatives samples distribute on meaningful background objects (chair

cushion, cabinet, and other parts of the dog) in the image. These responses from the teacher are valuable in improving the student’s performance.

4.3 Comparison with State-of-the-arts

Table 2: Experimental results on COCO-Standard. * means our re-implemented results on FCOS, ⊙ means large scale jittering is adopted when training

	COCO-Standard			
	1%	2%	5%	10%
Supervised	11.24 ± 0.18	15.04 ± 0.31	20.82 ± 0.13	26.44 ± 0.11
CSD [10]	10.51 ± 0.06	13.93 ± 0.12	18.63 ± 0.07	22.46 ± 0.08
STAC [24]	13.97 ± 0.35	18.25 ± 0.25	24.38 ± 0.12	28.64 ± 0.21
Instant Teaching [35]	18.05 ± 0.15	22.45 ± 0.30	26.75 ± 0.05	30.40 ± 0.05
ISMT [31]	18.88 ± 0.74	22.43 ± 0.56	26.37 ± 0.24	30.52 ± 0.52
Unbiased Teacher [18]	20.75 ± 0.12	24.30 ± 0.07	28.27 ± 0.11	31.50 ± 0.10
Humble Teacher [26]	16.96 ± 0.38	21.72 ± 0.24	27.70 ± 0.15	31.61 ± 0.28
<i>Li, et al.</i> [13]	19.02 ± 0.25	23.34 ± 0.18	28.40 ± 0.15	32.23 ± 0.14
Unbiased Teacher* [18]	18.31 ± 0.44	22.39 ± 0.26	27.73 ± 0.13	31.52 ± 0.15
Ours	19.64 ± 0.34	25.39 ± 0.13	30.83 ± 0.21	35.11 ± 0.12
Soft Teacher [⊙] [30]	20.46 ± 0.39	—	30.74 ± 0.08	34.04 ± 0.14
Ours [⊙]	22.38 ± 0.31	27.20 ± 0.20	33.01 ± 0.14	37.13 ± 0.12

COCO-Standard. We compare Dense Teacher with several existing methods under the COCO-Standard setting in Table. 2. When the labeled data varies from 2% to 10% our model consistently shows superior results. Whereas Under the 1% labeled setting, the performance of Dense Teacher is lower than Faster R-CNN based Unbiased Teacher. However, a more direct comparison between our method and Unbiased Teacher under the 1% setting on FCOS shows that Dense Teacher still leads by 1.2% mAP. Moreover, when applying large-scale jittering for augmentation following the implementation of Soft Teacher, our method obtains more significant improvements and becomes new the state-of-the-art.

VOC & COCO-Full. Results in Table. 3 and Table. 4 show that Dense Teacher lead the performance in both settings. On VOC dataset, Dense Teacher improves its supervised baseline by 8.2% and 10.0% on AP50 and mAP (*i.e.*, from AP50 to AP95). Under the COCO-Full setting, since the baseline reported in other works are not the same, we list the performance of each method in the form of “baseline→result”. Our approach obtains a boost of 3.1% from the 2017unlabeled set, which is much higher than CSD, STAC, and Unbiased Teacher. We finally apply the large-scale jittering trick and a longer training scheduler for a fair comparison with Soft Teacher, where Dense Teacher boosts mAP by 4.9%, reaching 46.12% mAP.

Table 3: Comparison with existing methods on Pascal VOC. Evaluations are performed on **VOC07 test**

Method	Labeled	Unlabeled	AP_{50}	$AP_{50:95}$
Supervised(Ours)	VOC07	None	71.69	45.87
CSD [10]			74.7	-
STAC [24]			77.45	44.64
ISMT [31]			77.23	46.23
Unbiased Teacher [18]	VOC07	VOC12	77.37	48.69
<i>Li, et al.</i> [13]			79.00	54.60
Instant Teaching [35]			79.20	50.00
Ours			79.89	55.87
CSD [10]			75.1	-
STAC [24]			79.08	46.01
ISMT [31]		VOC12	77.75	49.59
Unbiased Teacher[18]	VOC07	+	78.82	50.34
<i>Li, et al.</i> [13]		COCO20cls	79.60	56.10
Instant Teaching [35]			79.90	55.70
Ours			81.23	57.52

4.4 Ablation and Key Parameters

Effect of Hard Negative Samples In Dense Teacher, hard negative samples/anchors can be better utilized. We explore three different strategies to study the impact of these hard negative regions⁵, including “suppressing”, “ignoring” and “selecting”. Results are shown in Table. 5. We first suppress these samples to 0 and find a significant performance drop on both classification and regression branches compared to the original setting. Then, we ignore these samples when calculating loss. It turns out that this setting performs better than the “suppress” setting but still falls short when they are selected for training, indicating that learning to predict those hard negative samples can positively affect the model performance.

Regression Branch We have shown that unsupervised learning on the regression branch effectively improves model performance. However, since the output of deltas in the background region is not meaningful in FCOS, the quality of pseudo-labels in this branch is highly dependent on the design of the learning region. As can be seen in Table. 1, in our region division and pseudo-box based method, the model can obtain gains of 1.8% mAP and 1.6% mAP. When using ground truth positives as learning region (see Table. 5 “suppress” and “ignore”), the model gains about 2% mAP. Therefore, our region division strategy can produce sufficiently reliable regions for this task.

⁵ Since the “unlabeled images” under the COCO-Standard setting actually come with annotations, we can perform label assignments on images using these annotations. The difference between our division ($k = 1$) and the assigned foreground is defined as hard negatives.

Table 4: Experimental results on COCO-Full. Evaluations are done on COCO val2017. Note that 1x represents 90K training iterations, Nx represents Nx90K iterations. © means training with large scale jittering

Method	mAP
CSD [10] (3x)	40.20 $\xrightarrow{-1.38}$ 38.82
STAC [24] (6x)	39.48 $\xrightarrow{-0.27}$ 39.21
ISMT [31]	37.81 $\xrightarrow{+1.83}$ 39.64
Instant-Teaching [35]	37.63 $\xrightarrow{+2.57}$ 40.20
Unbiased Teacher [18] (3x)	40.20 $\xrightarrow{+1.10}$ 41.30
Humble Teacher [26] (3x)	37.63 $\xrightarrow{+4.74}$ 42.37
<i>Li, et al.</i> [13] (3x)	40.20 $\xrightarrow{+3.10}$ 43.30
Ours(3x)	41.22 $\xrightarrow{+2.66}$ 43.90
Ours(6x)	41.22 $\xrightarrow{+3.70}$ 44.94
Soft Teacher© [30] (8x)	40.90 $\xrightarrow{+3.70}$ 44.60
Ours(8x)©	41.24 $\xrightarrow{+4.88}$ 46.12

Effectiveness on other detectors and datasets Apart from comparison with state-of-the-arts, we also validate the effectiveness of our method on anchor-based detector and on CrowdedHuman [22] dataset. For anchor-based detector, we take RetinaNet as representation. When comparing with pseudo-box based Unbiased Teacher, our method stay ahead of the curve as show in 6, Dense Teacher achieved a 2.2% mAP improvement over Unbiased Teacher. On the CrowdedHuman dataset with FCOS detector, our method obtains a 2.1% mMR improvement as well.

Size of Learning Region We compare Dense Teacher’s performance under different selecting ratio k in Table. 7(a). As shown in the table, we obtain the best performance when selecting 1% of the samples (5 layers of FPN) for unsupervised learning. According to statistics, there are about 0.4% of positive samples in COCO under the label assigning rule of FCOS. Therefore, the optimal learning region not only contains ground truth positives, but also encompasses hard negatives that are valuable. This also suggests that although the model

Table 5: Impact of strategies dealing with hard negatives

HN samples	Cls	Reg	AP
Suppress	✓	×	31.56
Ignore	✓	×	32.64
Select	✓	×	33.34
Suppress	✓	✓	33.47
Ignore	✓	✓	34.72
Select	✓	✓	35.11

Table 6: Extensive comparison on different detectors and datasets

Method	Anchor	Dataset	AP/MR
UT [18]	✓	COCO	28.9
Ours	✓	COCO	31.1
UT [18]	×	CH[22]	62.8
Ours	×	CH[22]	60.7

Table 7: Ablation study on hyper-parameters introduced by our method. COCO 10% stands for COCO-Standard 10%

(a) Size of learning region					(b) Unsupervised weight factor w_u				
Setting	$k(\%)$	AP	AP50	AP75	Setting	w_u	AP	AP50	AP75
COCO 10%	0.1	25.76	41.79	27.34	COCO 10%	2	34.86	53.22	37.34
	0.5	34.11	51.94	36.80		4	35.11	53.35	37.79
	1	35.11	53.35	37.79		8	33.81	51.97	36.34
	3	34.47	53.13	36.87	COCO-	2	44.92	63.71	48.79
	5	33.85	53.00	36.03	Full	4	43.04	61.60	46.87

performance is affected by this hyperparameter, the statistical characteristics of the data set can help us determine the optimal value of this hyperparameter, mitigate model migration and deployment challenges.

Unsupervised Loss Weight The weight of unsupervised data also has an important impact on the training results. The experimental results in 7(b) turn out that: for a limited amount of supervised data like in the COCO-Standard 10% setting, a relatively large weight of 4 is favorable. Meanwhile, for the COCO-Full setting where much more labeled data are available, weight of 2 is enough. We attribute this phenomenon to the different degrees of over-fitting. When only a small amount of annotations are available, a relatively large weight of unsupervised parts could introduce stronger supervision. In contrast, when given a large amount of labeled data, a small unsupervised weight could address and better utilize supervision.

5 Conclusion

In this paper, we revisit the form of pseudo-labels in existing semi-supervised learning. By analyzing various flaws caused by the lengthy pseudo-box generation pipeline, we point out that pseudo-box is a sub-optimal choice for unlabeled data. To address this issue, we propose the Dense Teacher, a SS-OD framework which adopts dense predictions from the teacher model as pseudo-labels for unlabeled data. Our approach is simpler but stronger. We demonstrate its efficacy by comparing Dense Teacher with other pseudo-box based SS-OD algorithms on MS-COCO and Pascal VOC benchmarks. Results on both benchmarks show that our Dense Teacher achieves state-of-the-art performance.

References

1. Bachman, P., Alsharif, O., Precup, D.: Learning with pseudo-ensembles. *Advances in neural information processing systems* **27** (2014) 1, 3
2. Bachman, P., Alsharif, O., Precup, D.: Learning with pseudo-ensembles. *Advances in neural information processing systems* **27** (2014) 3
3. Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785* (2019) 1
4. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems* **32** (2019) 1, 3
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International conference on machine learning*. pp. 1597–1607. PMLR (2020) 1, 3
6. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html> 2, 9
7. Ge, Z., Hu, C., Huang, X., Qiu, B., Yoshie, O.: Dualbox: Generating bbox pair with strong correspondence via occlusion pattern clustering and proposal refinement. In: *2020 25th International Conference on Pattern Recognition (ICPR)*. pp. 2097–2102. IEEE (2021) 6
8. Ge, Z., Jie, Z., Huang, X., Xu, R., Yoshie, O.: Ps-rcnn: Detecting secondary human instances in a crowd via primary object suppression. In: *2020 IEEE international conference on multimedia and expo (ICME)*. pp. 1–6. IEEE (2020) 6
9. 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) 9
10. Jeong, J., Lee, S., Kim, J., Kwak, N.: Consistency-based semi-supervised learning for object detection. *Advances in neural information processing systems* **32** (2019) 4, 11, 12, 13
11. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242* (2016) 3
12. Lee, D.H., et al.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks 1, 3
13. Li, H., Wu, Z., Shrivastava, A., Davis, L.S.: Rethinking pseudo labels for semi-supervised object detection. *arXiv preprint arXiv:2106.00168* (2021) 4, 11, 12, 13
14. Li, X., Wang, W., Wu, L., Chen, S., Hu, X., Li, J., Tang, J., Yang, J.: Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *Advances in Neural Information Processing Systems* **33**, 21002–21012 (2020) 4, 8
15. 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) 4
16. 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) 4, 8

17. 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) 2, 5, 9
18. Liu, Y.C., Ma, C.Y., He, Z., Kuo, C.W., Chen, K., Zhang, P., Wu, B., Kira, Z., Vajda, P.: Unbiased teacher for semi-supervised object detection. arXiv preprint arXiv:2102.09480 (2021) 1, 4, 5, 9, 10, 11, 12, 13
19. 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) 3
20. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115(3), 211–252 (2015). <https://doi.org/10.1007/s11263-015-0816-y> 9
21. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems* 29 (2016) 3
22. Shao, S., Zhao, Z., Li, B., Xiao, T., Yu, G., Zhang, X., Sun, J.: Crowddhuman: A benchmark for detecting human in a crowd. arXiv preprint arXiv:1805.00123 (2018) 5, 13
23. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems* 33, 596–608 (2020) 1, 3
24. Sohn, K., Zhang, Z., Li, C.L., Zhang, H., Lee, C.Y., Pfister, T.: A simple semi-supervised learning framework for object detection. arXiv preprint arXiv:2005.04757 (2020) 1, 4, 9, 11, 12, 13
25. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15(1), 1929–1958 (2014) 1
26. Tang, Y., Chen, W., Luo, Y., Zhang, Y.: Humble teachers teach better students for semi-supervised object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3132–3141 (2021) 11, 13
27. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* 30 (2017) 3
28. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9627–9636 (2019) 9
29. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10687–10698 (2020) 3
30. Xu, M., Zhang, Z., Hu, H., Wang, J., Wang, L., Wei, F., Bai, X., Liu, Z.: End-to-end semi-supervised object detection with soft teacher. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3060–3069 (2021) 1, 4, 5, 9, 11, 13
31. Yang, Q., Wei, X., Wang, B., Hua, X.S., Zhang, L.: Interactive self-training with mean teachers for semi-supervised object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5941–5950 (2021) 11, 12, 13
32. Zhang, F., Pan, T., Wang, B.: Semi-supervised object detection with adaptive class-rebalancing self-training. arXiv preprint arXiv:2107.05031 (2021) 4

- 33. Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9759–9768 (2020) [4](#)
- 34. Zhixing, D., Zhang, R., Chang, M., Liu, S., Chen, T., Chen, Y., et al.: Distilling object detectors with feature richness. *Advances in Neural Information Processing Systems* **34** (2021) [8](#)
- 35. Zhou, Q., Yu, C., Wang, Z., Qian, Q., Li, H.: Instant-teaching: An end-to-end semi-supervised object detection framework. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4081–4090 (2021) [11](#), [12](#), [13](#)