

Semi-Supervised Object Detection with Object-wise Contrastive Learning and Regression Uncertainty

Honggyu Choi¹

honggyuchoi@kaist.ac.kr

Zhixiang Chen²

zhixiang.chen@sheffield.ac.uk

Xuepeng Shi³

x.shi19@imperial.ac.uk

Tae-Kyun Kim^{1,3}

kimtaekyun@kaist.ac.kr

¹ KAIST

² The University of Sheffield

³ Imperial College London

Abstract

Semi-supervised object detection (SSOD) aims to boost detection performance by leveraging extra unlabeled data. The teacher-student framework has been shown to be promising for SSOD, in which a teacher network generates pseudo-labels for unlabeled data to assist the training of a student network. Since the pseudo-labels are noisy, filtering the pseudo-labels is crucial to exploit the potential of such framework. Unlike existing suboptimal methods, we propose a two-step pseudo-label filtering for the classification and regression heads in a teacher-student framework. For the classification head, **OCL (Object-wise Contrastive Learning)** regularizes the object representation learning that utilizes unlabeled data to improve pseudo-label filtering by enhancing the discriminativeness of the classification score. This is designed to pull together objects in the same class and push away objects from different classes. For the regression head, we further propose **RUPL (Regression-Uncertainty-guided Pseudo-Labeling)** to learn the aleatoric uncertainty of object localization for label filtering. By jointly filtering the pseudo-labels for the classification and regression heads, the student network receives better guidance from the teacher network for object detection task. Experimental results on Pascal VOC and MS-COCO datasets demonstrate the superiority of our proposed method with competitive performance compared to existing methods.

1 Introduction

Semi-Supervised Object Detection (SSOD) leverages both labeled data and extra unlabeled data to learn object detectors. Previous works [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41] have proposed various methods to exploit the unlabeled data. Among them, one promising solution is to generate pseudo-labels for the unlabeled data by a pre-trained model and then use them as labeled data along with the original labeled data to train the target object detector. By applying different augmentations to an image, the consistency between




Figure 1: The correlation between the regression uncertainties and the ground-truth IoUs of bounding boxes. The ground-truth IoU represents the localization quality. Low uncertainty suggests accurate localization of the bounding box.

the labels of them can serve as extra knowledge to train the target model. The consistency can be exploited by using the generated pseudo-labels from weakly augmented images to regularize the model’s predictions of strongly augmented images [27]. A teacher-student framework [30] further learns the pseudo-labels by mutual learning teacher and student networks, where the teacher network generates pseudo-labels of unlabeled data to assist the training of the student network, and the student network transfers the updated knowledge to the weights of the teacher network.

To train the detector with the pseudo-labels, the quality of the generated pseudo-labels is critical to the detection performance. For the detection task, the pseudo-labels are defined for each region of interest and include the class labels and bounding boxes. Hence, the qualities of both the classification and the localization are important to filter out unreliable pseudo-labels. The classification score is widely adopted to select region proposals with confidence higher than a pre-defined threshold. The selected labeled regions are used to train both the classification and regression heads in the detector. There are some heuristic designs to measure the localization quality of the generated pseudo-labels, e.g., prediction consistency [36], interval classification uncertainty [48]. While these techniques together with the classification score are effective, we argue the existing classification scores and localization optimization is suboptimal, since the lack of labeled data makes classification score less discriminative, and the measurement of pseudo labels’ localization quality is less investigated. The classification scores mainly distill knowledge from labeled data to evaluate the quality of unlabeled data, which under exploits the unlabeled data. To tackle these challenges, we extend existing methods in self-supervised representation learning to SSOD task. Instead of using heuristic designs, we further adopt uncertainty to measure the localization quality, which turns out to be an effective indicator as shown in Fig. 1.

In this paper, we propose an effective method to generate pseudo-labels for semi-supervised object detection. Specifically, we present a two-step pseudo-label filtering for the classification and regression heads in a teacher-student framework. For the classification head, the classification score is used to filter out unreliable pseudo class labels. We aim to improve the classification filtering through enhancing the capacity of the classification branch by taking into consideration of the unlabeled data. We introduce an object-wise contrastive learning (OCL) loss for the feature extractor in the classification head. This contrastive loss is defined on object regions with semantic similar regions defined as positive pairs. This loss pull together object representations of the positive pairs and push away representations of negative pairs. For the regression head, we propose to use uncertainty as the indicator to filter out pseudo-labels with regression-uncertainty-guided pseudo-labeling (RUPL). We design an uncertainty head in parallel to the classification and regression heads to learn the aleatoric uncertainty for bounding boxes. We only select pseudo-labels with low uncertain-

ties to train our models. We combine OCL and RUPL in our framework to improve the quality of pseudo-labels for the detection task. To demonstrate effectiveness of our framework, we perform the experiments on standard object detection benchmark: PASCAL VOC [9] and MS-COCO [19]. We experimentally show that OCL and RUPL are complementary and have a synergetic effect. Our framework also achieves competitive performance compared to existing methods, without geometry or improved augmentations.

Our contributions are summarized as follows: (1) We propose an integrated framework that addresses filtering pseudo-labels for classification and regression in SSOD. (2) We propose OCL that improves the discriminativeness of the classification score to enhance the filtering the pseudo-labels for classification. (3) We propose RUPL that models bounding box localization quality via uncertainty and removes misplaced pseudo-labels for regression. (4) We experimentally show the synergetic effect between OCL and RUPL and demonstrate the effectiveness of the overall framework on standard object detection benchmarks.

2 Related Work

Semi-Supervised Object Detection. SSOD methods can be divided into two categories: pseudo-labeling [27, 54] methods and consistency regularization [11, 11, 28] methods. Pseudo-labeling methods regularize the model using predictions generated from the pre-trained model utilizing unlabeled data. Recent pseudo-labeling-based works [22, 29, 58, 41] adopt the teacher-student framework [30]. In this framework, a teacher’s predictions guide a student and the student weight parameters evolve the teacher, which leads to remarkable performance improvement. Since pseudo-labels may have noise that degrades the model performance, classification scores are used to eliminate unreliable ones. To improve the pseudo-label quality, IT [41] and ISTM [58] propose ensemble-based method, and ACRST [59] adopts a multi-label classifier at the image level to use high level information. Furthermore, to eliminate misaligned pseudo-labels, 3DIoUMatch [52] adopts IoU prediction [12] in semi-supervised 3D object detection, and ST [36] utilizes regression prediction consistency, and RPL [18] reformulates regression as classification. Concurrent work [23] adopts uncertainty estimation and guides the student if pseudo-labels have lower uncertainty than the student’s predictions. In contrast to existing works, we propose an integrated framework that improves pseudo-label quality by improving the discriminativeness of object-wise features and modeling regression uncertainty.

Contrastive Learning. Contrastive Learning (CL) decreases the distance between positive paired samples and increases the distance between negative ones. Representation learning [8, 4, 31] has succeeded through self-supervised CL that treats different views of the same image as a positive pair. By extension, SupCon [15] proposes supervised CL that makes a positive pair for images from the same class. Recent works [13, 57] apply CL that makes positive pairs if unlabeled images or points have the same predicted class. In SSOD, PL [28] adopts CL that treats overlapped region proposals as positive pairs and otherwise as negative pairs even if they have the same category, which hinders the representation learning. By contrast, in SSOD, we first introduce CL that leverages predictions to identify object regions from unlabeled images and makes positive pairs for objects from the same class to improve the discriminativeness.

Uncertainty Estimation. A seminal work [14] captures epistemic and aleatoric uncertainties in the deep learning frameworks of computer vision. Recent works [5, 9] employ aleatoric uncertainty in bounding box regression to identify well-localized bounding boxes.

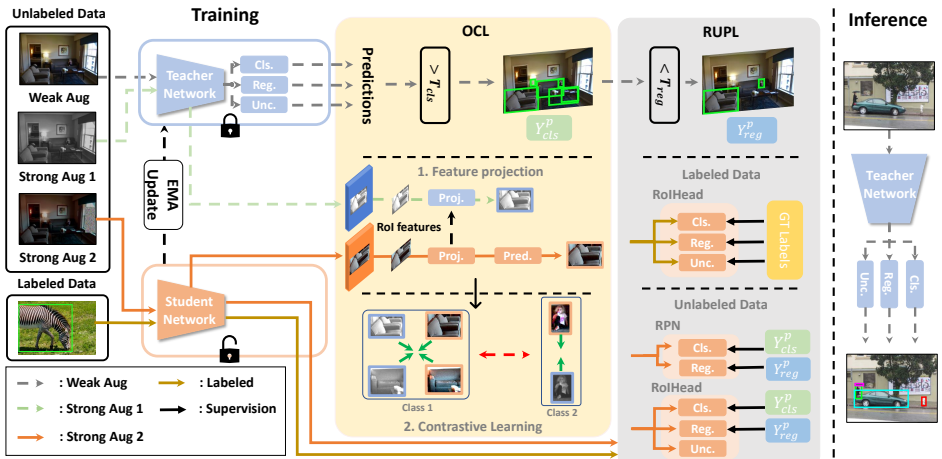


Figure 2: The overview of our proposed framework.

UPS [24] applies aleatoric uncertainty in semi-supervised image classification to remove noisy pseudo-labels. In contrast, our work adopts aleatoric uncertainty in SSOD to model the localization quality of pseudo-labels and select reliable pseudo-labels for regression.

3 Proposed Method

We first define the problem, then show the overall framework in Sec. 3.2. Finally we detail the proposed OCL in Sec. 3.3 and RUPL in Sec. 3.4.

3.1 Problem Definition

Compared with supervised object detection, semi-supervised object detection aims to utilize additional unlabeled images to improve the object detection accuracy. In the teacher-student framework, our task is to generate pseudo-labels from unlabeled data using the teacher network for to train the student network. Specifically, there are a labeled dataset $\{\mathbf{X}^l, \mathbf{Y}^l\} = \{(\mathbf{x}_i^l, \mathbf{y}_i^l)\}_{i=1}^{N_l}$ and an unlabeled dataset $\mathbf{X}^u = \{\mathbf{x}_i^u\}_{i=1}^{N_u}$, where N_l , N_u , \mathbf{x}^l , \mathbf{y}^l and \mathbf{x}^u denote the size of labeled dataset, the size of unlabeled dataset, labeled images, ground-truth labels, and unlabeled images, respectively. The ground-truth label \mathbf{y}^l consists of each object's class category c^l and bounding box b^l . The goal of pseudo-label filtering is to generate $\{\mathbf{X}_{cls}^p, \mathbf{Y}_{cls}^p\}$ and $\{\mathbf{X}_{reg}^p, \mathbf{Y}_{reg}^p\}$ from \mathbf{X}^u to train the classification and regression heads in the detection model, respectively.

3.2 Framework Overview

To utilize the unlabeled data in semi-supervised object detection, we propose object-wise contrastive learning (OCL) and regression-uncertainty-guided pseudo-labeling (RUPL) to select reliable pseudo-labels for the training of classification and regression, respectively. We apply our proposed modules to a simple baseline [22]. This baseline utilizes the teacher-student framework and pseudo-labeling, which are commonly adopted in SSOD frameworks.

Fig. 2 shows our proposed framework with OCL and RUPL.

Training stage. The teacher-student framework has a student network and a teacher network. At the pre-training stage, the student network is trained with a supervised loss (Eq. 1) using the labeled dataset $\{\mathbf{X}^l, \mathbf{Y}^l\}$.

$$\mathcal{L}_{sup} = \mathcal{L}_{rpn,cls} + \mathcal{L}_{rpn,reg} + \mathcal{L}_{roi,cls} + \mathcal{L}_{roi,reg} \quad (1)$$

The supervised loss consists of classification and regression losses for each RPN and ROI. Specifically, we follow [22] to use cross-entropy loss for $\mathcal{L}_{rpn,cls}$, smooth L1 loss [23] for $\mathcal{L}_{rpn,reg}$, focal loss [24] for $\mathcal{L}_{roi,cls}$. We adopt uncertainty-aware regression loss [25] for $\mathcal{L}_{roi,reg}$. At the mutual learning stage, the teacher network is initialized by the pre-trained student network. Both labeled and unlabeled images are used at this stage. The teacher network generates the pseudo-labels of unlabeled images to train the student network. The overall loss of the student network is

$$\mathcal{L} = \mathcal{L}_{sup} + \lambda_{unsup} \mathcal{L}_{unsup} + \lambda_{ocl} \mathcal{L}_{ocl}, \quad (2)$$

where \mathcal{L}_{unsup} and \mathcal{L}_{ocl} denote the unsupervised loss and object-wise contrastive loss, respectively. The unsupervised and object-wise contrastive losses are computed on unlabeled images. Our unsupervised loss consists of similar loss terms as the supervised loss. Here, $\mathcal{L}_{rpn,cls}$ and $\mathcal{L}_{roi,cls}$ are computed on the set of pseudo-labels for classification $\{\mathbf{X}_{cls}^p, \mathbf{Y}_{cls}^p\}$, $\mathcal{L}_{rpn,reg}$ and $\mathcal{L}_{roi,reg}$ are computed on the set of pseudo-labels for regression $\{\mathbf{X}_{reg}^p, \mathbf{Y}_{reg}^p\}$. And we use smoothL1 loss for $\mathcal{L}_{roi,reg}$. The teacher network is iteratively updated by the student network’s weight parameters via exponential moving average [26].

Inference Stage. The teacher network is used at the inference stage to detect objects for a given image. With the uncertainty branch, our model also predicts the aleatoric uncertainty of a bounding box, which can be interpreted as the confidence of the boundary location.

3.3 Object-wise Contrastive Learning

In the classification head, we propose the object-wise contrastive learning (OCL) to regularize the feature representation of ROI objects. We visualize the details of OCL and show an illustration in Fig. 3.

To use the contrastive loss, we construct positive and negative pairs of ROI objects within each batch during training. We first predict the class category c^p , classification score p^p , and bounding box b^p for each ROI object detected by the teacher network taking as input the weakly augmented unlabeled images $\mathcal{A}_0(\mathbf{x}^u)$. We denote the number of detected ROI objects in a batch as N_o . We then input two strong augmented images $\mathcal{A}_1(\mathbf{x}^u)$ and $\mathcal{A}_2(\mathbf{x}^u)$ to the teacher and student networks to get the ROI features \mathbf{z}^t and \mathbf{z}^s using predicted bounding boxes b^p , respectively. Then, \mathbf{z}^t and \mathbf{z}^s are projected to the low-dimensional feature space and normalized with L2-normalization to generate $\hat{\mathbf{z}}^t$ and $\hat{\mathbf{z}}^s$, respectively.

As the predicted class is noisy, we take the classification score into consideration to define ROI object pairs, as shown in Eq. 3. A pair is positive if the n -th object feature $\hat{\mathbf{z}}_n^t$ and the m -th object feature $\hat{\mathbf{z}}_m^t$ belong to the same ROI. We further measure the similarity if the two features are not from the same ROI but have the same class with high confidence.

$$w_{nm} = \begin{cases} 1 & \text{if } n = m, \\ p_n^p \cdot p_m^p & \text{if } n \neq m \text{ and } c_n^p = c_m^p \text{ and } p_n^p > T_{cont} \text{ and } p_m^p > T_{cont}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

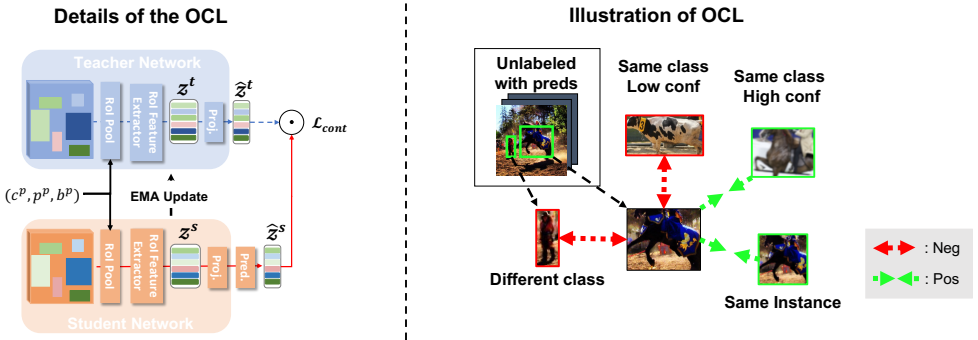


Figure 3: Left: Training details of OCL. Right: Intuitive illustration of the OCL.

T_{cont} is a classification score threshold for contrastive learning. Our contrastive loss is based on the supervised contrastive loss [15], and is formulated as

$$\mathcal{L}_{cont} = -\frac{1}{\mathcal{N}_o} \sum_{n=1}^{\mathcal{N}_o} \frac{1}{\sum_{m=1}^{\mathcal{N}_o} \mathbb{1}_{(w_{nm}>0)}} \sum_{m=1}^{\mathcal{N}_o} w_{nm} \cdot \log \frac{\exp(\hat{z}_n^s \cdot \hat{z}_m^t / \tau)}{\sum_{l=1}^{\mathcal{N}_o} \mathbb{1}_{(l \neq n)} \exp(\hat{z}_n^s \cdot \hat{z}_l^t / \tau)}, \quad (4)$$

where τ is the temperature parameter. $\mathbb{1}(\cdot)$ equals to 1 if the condition is true, otherwise zero. Note that, for the n -th ROI object, we normalize the loss by the number of positive pairs rather than the similarity. This aims to reduce the influence of mis-classified predictions. We adopt the symmetrized loss [15] to compute our object-wise contrastive loss as

$$\mathcal{L}_{ocl} = \frac{1}{2}(\mathcal{L}_{cont} + \mathcal{L}'_{cont}), \quad (5)$$

where \mathcal{L}'_{cont} is the same loss as \mathcal{L}_{cont} but with the augmentations for the teacher and student networks swapped. Through the OCL (Eq. 5), intra-class objects are encouraged to have similar feature representations, and inter-class objects are encouraged to have different feature representations. As a result, we can select pseudo-labels for classification task as

$$\{\mathbf{X}_{cls}^p, \mathbf{Y}_{cls}^p\} = \{(x_i^p, y_i^p) | p_{i,\cdot}^p > T_{cls}\}_{i=1}^{N_u}, \quad (6)$$

where y_i^p is the pseudo-labels for the i -th image and consists of the class category and bounding box for objects, $p_{i,\cdot}^p$ is the classification score for the corresponding object in the i -th image, and T_{cls} is a predefined threshold.

3.4 Regression-Uncertainty-guided Pseudo-Labeling

We propose regression-uncertainty-guided pseudo-labeling (RUPL) to utilize regression uncertainty to filter out unreliable bounding boxes. Following [9], we add an uncertainty head in parallel to classification and regression heads to predict regression uncertainty σ . We use the uncertainty-aware regression loss [14] as the the student network's bounding box regression loss ¹,

$$\mathcal{L}_{roi,reg} = \frac{\text{smooth}L_1(\hat{t}, t)}{\sigma^2} + \lambda_{unc} \log(\sigma^2), \quad (7)$$

¹This loss is only applied to labeled images.

where \hat{t} and t denote the predicted bounding box offset and ground truth bounding box offset. λ_{unc} is a hyperparameter to control the effect of the uncertainty term. During training, we apply this loss to each of the four boundaries of a bounding box. During inference, we use the average of these four uncertainties (denoted as $\bar{\sigma}$) as the regression uncertainty of the bounding box. We can then derive the selection of pseudo labels for regression as

$$\{\mathbf{X}_{reg}^p, \mathbf{Y}_{reg}^p\} = \{(\mathbf{x}_i^p, \mathbf{y}_i^p) | p_{i,c}^p > T_{cls} \text{ and } \bar{\sigma}_{i,c}^p < T_{reg}\}_{i=1}^{N_u}, \quad (8)$$

where \mathbf{y}_i^p is the pseudo-labels for the i -th image and consists of the class category and bounding box for objects, $p_{i,c}^p$ and $\bar{\sigma}_{i,c}^p$ are the classification score and regression uncertainty for the corresponding object in the i -th image, and T_{cls} T_{reg} are predefined thresholds. By learning the regression uncertainty, we provide an alternative way to capture the reliability of pseudo-labels without additional forwarding or reformulation.

4 Experiments

4.1 Experimental Settings

Following previous works [22], we evaluate our method on Pascal VOC [8] and MS-COCO [14]. We conduct experiments using different settings. (1) VOC [22]: VOC07-trainval and VOC12-trainval are used as labeled and unlabeled datasets, respectively. We also show results using COCO20cls [22] as an additional unlabeled dataset. VOC07-test is used to evaluate. (2) COCO-standard [22]: We randomly select 1%/5%/10% samples from COCO2017 train dataset as our labeled datasets, and the rest of them as unlabeled datasets. COCO2017 validation set is used to evaluate. (3) COCO-35k [42]: We use a subset of COCO2014 validation set as a labeled dataset and COCO2014 training set as an unlabeled dataset. COCO2014 minival is used to evaluate. (4) COCO-additional [22]: We use COCO2017 train dataset and COCO2017 unlabeled dataset as labeled and unlabeled datasets, respectively. COCO2017 validation set is used to evaluate.

4.2 Implementation details

Following [22, 27], our object detector is Faster-RCNN [23] with ResNet-50 backbone and FPN [20]. For VOC and VOC with COCO20cls settings, we train our model by 60k and 90k iterations, respectively, including 12k iterations for pre-training. For the coco-standard setting, we pre-train 5k/20k/40k iterations for 1%/5%/10% of COCO-standard and persist the training until 180k iterations. For the COCO-35k and COCO-additional settings, models are pre-trained/trained for 12k/180k and 90k/360k iterations, respectively. More training details, including augmentation strategies and model architecture, are in the supplementary.

4.3 Results

VOC. Experimental results on Pascal VOC [8] are shown in Tab. 1. We add our proposed OCL and RUPL to UBT [22] without additional augmentation strategies. When using VOC12-trainval as unlabeled data, our model outperforms UBT [22] by 3.35 mAP and 0.89 mAP on AP_{50:95} and AP₅₀, respectively. When using COCO20cls [22] as additional unlabeled data, our model has larger improvements and outperforms UBT [22] by 3.54 mAP and 1.25 mAP on AP_{50:95} and AP₅₀, respectively. This shows that our proposed method can

Method	Conference	Unlabeled	AP _{50:95}	AP ₅₀	Unlabeled	AP _{50:95}	AP ₅₀
Supervised [†] [14]	-	None	42.13	72.63	None	42.13	72.63
STAC [14]	-	VOC12-trainval	44.64	77.45	VOC12-trainval + COCO20cls	46.01	79.08
UBT [†] [14]	ICLR2021		48.69	77.37		50.34	78.82
IT [14]	CVPR2021		50.00	79.20		50.80	79.90
HT [14]	CVPR2021		53.04	80.94		54.41	81.29
CN [14]	NIPS2021		49.3	80.6		50.2	81.4
ST [14]	ICCV2021		-	80.32		-	-
ACRST [14]	AAAI2022		54.30	81.11		-	-
RPL [14]	AAAI2022		54.6	79.0		56.1	79.6
DDT [14]	AAAI2022		54.7	82.4		55.9	82.5
MUM [‡] [14]	CVPR2022		50.22	78.94		52.31	80.45
MA-GCP [14]	CVPR2022		-	81.72		-	-
SED [8]	CVPR2022		-	80.60		-	-
UBTv2 [‡] [14]	CVPR2022		56.87	81.29		58.08	82.04
Ours [†]	-		52.04	78.29		53.88	80.07
Ours [‡]	-		57.34	82.18		58.99	82.98

Table 1: Experimental results on Pascal VOC dataset [8]. Our method can outperform existing works. † and ‡ denote results that are evaluated by the COCOevaluator and VOCEvaluator in Detectron2 [35], respectively. Unmarked works are not implemented on Detectron2 [35], or are hard to check through the paper or released code. We give more information about the evaluator in our supplementary Sec. 3.

better benefit from unlabeled data. Moreover, our method outperforms other existing models [8, 16, 17, 18, 23, 27, 29, 33, 39, 40, 41] on both experiment settings, which shows the superiority of our method.

COCO-35k		COCO-standard				COCO-additional	
Method	AP _{50:95}	Method	1%	5%	10%	Method	AP _{50:95}
Supervised [14]	31.3	Supervised [14]	9.05	18.47	23.86	Supervised [14]	40.20
DD [14]	33.1	STAC [14]	13.97	24.38	28.64	STAC [14]	39.21
MP [14]	34.8	UBT [14]	20.75	28.27	31.50	PL [14]	38.40
MP + DD [14]	35.2	IT [14]	18.05	26.75	30.4	UBT [14]	41.30
UBT [‡]	36.36	RPL [14]	18.21	27.78	31.67	IT [14]	40.20
Ours	37.13	CN [14]	18.41	28.96	32.43	RPL [14]	43.30
		ST [†] [14]	20.46	30.74	34.04	CN [14]	43.20
		DDT [14]	18.62	29.24	32.80	ST [†] [14]	44.50
		MUM [14]	21.88	28.52	31.87	DDT [14]	41.90
		UBTv2 [†] [14]	25.40	31.85	35.08	MUM [14]	42.11
		Ours	21.63	30.66	33.53	UBTv2 [†] [14]	44.75
						Ours	41.89

Table 2: Experimental results on MS-COCO. We report the AP_{50:95} of different methods. † denotes using extra augmentation such as a scale-jittering augmentation strategy to improve the accuracy, while we do not use. It significantly improves the mAP by about 1.37AP (We refer the reader to Tab. 6 of the supplementary of UBTv2 [23]). ‡ denotes our implemented result based on the official repository.

MS-COCO. We also conduct experiments on MS-COCO [19] to verify the effectiveness of our method, shown in Tab. 2. Comparing with UBT [22], our model outperform it by 0.88/2.39/2.03 mAP, 0.77 mAP, and 0.59 mAP on COCO-standard, COCO-35k, and COCO-additional. respectively. Moreover, our method achieves comparable results on COCO-standard, COCO-35k, and COCO-additional compared to other state-of-the-arts. These results consistently support the effectiveness of our method.




Figure 4: mAP (AP_{50:95}) curves during training of the experiments in Tab 3.

Method	AP _{50:95}	AP ₅₀	AP ₇₅
w/o CL	18.05	34.45	17.09
Self-sup CL	18.24	34.89	17.22
OCL (Ours)	18.34	35.21	17.16

Table 4: Ablation studies of OCL. CL denotes contrastive learning. A pair in Self-sup CL is positive if and only if they are the same instance ($i = j$ in Eq. 3). We exclude RUPL to facilitate the comparison.

	OCL	RUPL	AP _{50:95}	AP ₅₀	AP ₇₅
(1)	✓	✓	19.42	34.65	19.36
(2)		✓	18.95	33.82	19.07
(3)	✓		18.34	35.21	17.16
(4)			18.05	34.45	17.09

Table 3: Ablation studies of different modules.

Method	AP _{50:95}	AP ₅₀	AP ₇₅
Box jittering [56]	18.15	34.84	17.37
Predicted IoU [52]	18.85	34.76	18.26
Aleatoric uncertainty [14] (Ours)	18.95	33.81	19.06

Table 5: Different methods capturing the localization quality of pseudo-labels. Aleatoric uncertainty [14] (Ours) is the best. We exclude OCL to facilitate comparison.

4.4 Ablation studies

We conduct ablation studies to investigate the effectiveness of our model using 1% COCO-standard. Because of the limitation of computing resources, all experiments in this section are conducted with batch size 12/12 (labeled/unlabeled) and training iteration 45k, as in [16].

Effectiveness of proposed modules. We remove each proposed module from our framework and report results in Tab. 3. Comparing (3) to (4), we can see introducing OCL can improve both AP₅₀ and AP₇₅. This result demonstrates that the model benefits from more accurate pseudo-labels w.r.t classification. Comparing (2) to (4), we can see introducing RUPL can improve AP₇₅ but decrease AP₅₀. We assume the reason is that RUPL makes the model focus more on regression than classification during training. Comparing (1) to (2-4), we can see the accuracy of (1) is significantly higher than that of (2-4), which shows applying both OCL and RUPL can lead to the best accuracy. We argue that the model can generate more precise pseudo-labels with more discriminative classification scores and the introduced localization uncertainties, which leads to the improvement of accuracy. We also show the mAP curves of different experiments during training in Fig. 4, which further supports the effectiveness of our proposed modules. We emphasize that the proposed OCL and RUPL can complement each other and synergistically improve the model performance.

Ablation studies of OCL. To verify the effectiveness of the class information of our OCL, we show the results of three models in Tab. 4, i.e., without, self-supervised, and, object-wise semi-supervised contrastive learning (our OCL). In Tab. 4, self-sup CL improves the model performance compared to w/o CL and OCL further improves the performance since it helps the model to learn more discriminative feature representation for objects from different classes. This supports the effectiveness of our OCL.

Different localization quality measurements. In Tab. 5, we compare different localization quality measurements. Specifically, we compare box jittering [56], predicted IoU [52], and aleatoric uncertainty [14] in our RUPL. We use grid-search to find the best thresholds of box

jittering [56] and predicted IoU [52] to filter pseudo-labels, and set them as 0.01 and 0.8, respectively. The model with aleatoric uncertainty [14] achieves the highest performance on the AP_{50:95}.

Different Regression Thresholds of RUPL. In this section, we show experimental results on how different thresholds of RUPL (T_{reg}) affect the detection accuracy. As shown in Tab. 6, the model achieves the highest performance when we set the threshold as 0.5. The detection accuracy decreases when we set a larger threshold or small threshold. With a larger threshold, the selected samples become more diverse but unreliable. While with a smaller threshold, the selected samples become more reliable but monotonous.

Threshold	AP _{50:95}	AP ₅₀	AP ₇₅
0.3	19.18	34.53	19.13
0.4	19.37	34.49	19.28
0.5	19.42	34.65	19.36
0.6	19.26	34.45	19.36
0.7	18.78	33.50	18.93

Table 6: Detection accuracy with different regression thresholds of RUPL.

5 Conclusion

In this paper, we propose a two-step pseudo-label filtering for SSOD. We deal with both the classification and regression heads in the detection model. For the classification head, we propose an object-wise contrastive learning loss to exploit the unlabeled data to enhance the discriminativeness of classification score for pseudo label filtering. For the regression head, we design an uncertainty branch to learn regression uncertainty to measure the localization quality for bounding box filtering. We experimentally show that the two components create a synergistic effect when integrated into the teacher-student framework. Our framework achieves remarkable performance gain against our baseline on both PASCAL VOC and MS-COCO without additional augmentation, and shows competitive results compared to other state-of-the-arts.

Acknowledgement

This work is in part sponsored by KAIA grant (22CTAP-C163793-02, MOLIT), NST grant (CRC 21011, MSIT), KOCCA grant (R2022020028, MCST) and the Samsung Display corporation.

References

- [1] Binbin Chen, Weijie Chen, Shicai Yang, Yunyi Xuan, Jie Song, Di Xie, Shiliang Pu, Mingli Song, and Yueting Zhuang. Label matching semi-supervised object detection. In *CVPR*, pages 14381–14390, June 2022.
- [2] Binghui Chen, Pengyu Li, Xiang Chen, Biao Wang, Lei Zhang, and Xian-Sheng Hua. Dense learning based semi-supervised object detection. In *CVPR*, pages 4815–4824, June 2022.

- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, volume 119, pages 1597–1607, 2020.
- [4] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, pages 9620–9629, 2021.
- [5] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee. Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving. In *ICCV*, pages 502–511, 2019.
- [6] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010.
- [7] Ross B. Girshick. Fast R-CNN. In *ICCV*, pages 1440–1448, 2015.
- [8] Qiushan Guo, Yao Mu, Jianyu Chen, Tianqi Wang, Yizhou Yu, and Ping Luo. Scale-equivalent distillation for semi-supervised object detection. In *CVPR*, pages 14522–14531, June 2022.
- [9] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *CVPR*, pages 2888–2897, 2019.
- [10] Jisoo Jeong, Seungeui Lee, Jeessoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. In *NeurIPS*, pages 10758–10767, 2019.
- [11] Jisoo Jeong, Vikas Verma, Minsung Hyun, Juho Kannala, and Nojun Kwak. Interpolation-based semi-supervised learning for object detection. In *CVPR*, pages 11602–11611, 2021.
- [12] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yuning Jiang. Acquisition of localization confidence for accurate object detection. In *ECCV*, volume 11218, pages 816–832, 2018.
- [13] Li Jiang, Shaoshuai Shi, Zhuotao Tian, Xin Lai, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *ICCV*, pages 6403–6412, 2021.
- [14] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, pages 5574–5584, 2017.
- [15] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020.
- [16] JongMok Kim, JooYoung Jang, Seunghyeon Seo, Jisoo Jeong, Jongkeun Na, and Nojun Kwak. Mum: Mix image tiles and unmix feature tiles for semi-supervised object detection. In *CVPR*, pages 14512–14521, June 2022.

- [17] Aoxue Li, Peng Yuan, and Zhenguo Li. Semi-supervised object detection via multi-instance alignment with global class prototypes. In *CVPR*, pages 9809–9818, June 2022.
- [18] Hengduo Li, Zuxuan Wu, Abhinav Shrivastava, and Larry S. Davis. Rethinking pseudo labels for semi-supervised object detection. In *AAAI*, pages 1314–1322, 2022.
- [19] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, pages 740–755, 2014.
- [20] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 936–944, 2017.
- [21] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE T-PAMI*, 42(2):318–327, 2020.
- [22] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. In *ICLR*, 2021.
- [23] Yen-Cheng Liu, Chih-Yao Ma, and Zsolt Kira. Unbiased teacher v2: Semi-supervised object detection for anchor-free and anchor-based detectors. In *CVPR*, pages 9819–9828, June 2022.
- [24] Peng Mi, Jianghang Lin, Yiyi Zhou, Yunhang Shen, Gen Luo, Xiaoshuai Sun, Liujuan Cao, Rongrong Fu, Qiang Xu, and Rongrong Ji. Active teacher for semi-supervised object detection. In *CVPR*, pages 14482–14491, June 2022.
- [25] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *NeurIPS*, pages 91–99, 2015.
- [26] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S. Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *ICLR*, 2021.
- [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. *CoRR*, abs/2005.04757, 2020.
- [28] Peng Tang, Chetan Ramaiah, Yan Wang, Ran Xu, and Caiming Xiong. Proposal learning for semi-supervised object detection. In *WACV*, pages 2290–2300, 2021.
- [29] Yihe Tang, Weifeng Chen, Yijun Luo, and Yuting Zhang. Humble teachers teach better students for semi-supervised object detection. In *CVPR*, pages 3132–3141, 2021.
- [30] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, pages 1195–1204, 2017.

- [31] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748, 2018.
- [32] He Wang, Yezhen Cong, Or Litany, Yue Gao, and Leonidas J. Guibas. 3dioumatch: Leveraging iou prediction for semi-supervised 3d object detection. In *CVPR*, pages 14615–14624, 2021.
- [33] Zhenyu Wang, Ya-Li Li, Ye Guo, and Shengjin Wang. Combating noise: Semi-supervised learning by region uncertainty quantification. In *NeurIPS*, pages 9534–9545, 2021.
- [34] Zhenyu Wang, Yali Li, Ye Guo, Lu Fang, and Shengjin Wang. Data-uncertainty guided multi-phase learning for semi-supervised object detection. In *CVPR*, pages 4568–4577, 2021.
- [35] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [36] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *ICCV*, pages 3040–3049, 2021.
- [37] Fan Yang, Kai Wu, Shuyi Zhang, Guannan Jiang, Yong Liu, Feng Zheng, Wei Zhang, Chengjie Wang, and Long Zeng. Class-aware contrastive semi-supervised learning. *CoRR*, abs/2203.02261, 2022.
- [38] Qize Yang, Xihan Wei, Biao Wang, Xian-Sheng Hua, and Lei Zhang. Interactive self-training with mean teachers for semi-supervised object detection. In *CVPR*, pages 5941–5950, 2021.
- [39] Fangyuan Zhang, Tianxiang Pan, and Bin Wang. Semi-supervised object detection with adaptive class-rebalancing self-training. In *AAAI*, pages 3252–3261, 2022.
- [40] Shida Zheng, Chenshu Chen, Xiaowei Cai, Tingqun Ye, and Wenming Tan. Dual decoupling training for semi-supervised object detection with noise-bypass head. In *AAAI*, pages 3526–3534, 2022.
- [41] Qiang Zhou, Chaohui Yu, Zhibin Wang, Qi Qian, and Hao Li. Instant-teaching: An end-to-end semi-supervised object detection framework. In *CVPR*, pages 4081–4090, 2021.