

Calibrated Teacher for Sparsely Annotated Object Detection

Haohan Wang^{1*}, Liang Liu^{2*†}, Boshen Zhang², Jiangning Zhang², Wuhao Zhang²,
Zhenye Gan², Yabiao Wang^{2†}, Chengjie Wang^{2,3†}, Haoqian Wang^{1†}

¹ Shenzhen International Graduate School, Tsinghua University

² Tencent YouTu Lab

³ Shanghai Jiao Tong University

wang-hh20@mails.tsinghua.edu.cn, {leoneliu, boshenzhang, vtzhang, wuhaozhang, wingzygan, caseywang, jasoncjwang}@tencent.com, wanghaoqian@tsinghua.edu.cn

Abstract

Fully supervised object detection requires training images in which all instances are annotated. This is actually impractical due to the high labor and time costs and the unavoidable missing annotations. As a result, the incomplete annotation in each image could provide misleading supervision and harm the training. Recent works on sparsely annotated object detection alleviate this problem by generating pseudo labels for the missing annotations. Such a mechanism is sensitive to the threshold of the pseudo label score. However, the effective threshold is different in different training stages and among different object detectors. Therefore, the current methods with fixed thresholds have sub-optimal performance, and are difficult to be applied to other detectors. In order to resolve this obstacle, we propose a *Calibrated Teacher*, of which the confidence estimation of the prediction is well calibrated to match its real precision. In this way, different detectors in different training stages would share a similar distribution of the output confidence, so that multiple detectors could share the same fixed threshold and achieve better performance. Furthermore, we present a simple but effective *Focal IoU Weight* (FIoU) for the classification loss. FIoU aims at reducing the loss weight of false negative samples caused by the missing annotation, and thus works as the complement of the teacher-student paradigm. Extensive experiments show that our methods set new state-of-the-art under all different sparse settings in COCO. *Code will be available at* <https://github.com/Whileherham/CalibratedTeacher>.

Introduction

Remarkable progress in fully supervised object detection has been witnessed in recent years. In order to train such an object detector, a large number of training images are required, and bounding boxes in each image are supposed to be exhaustively annotated. However, it could be practically infeasible due to the high labor and time costs with the volume of datasets increasing, and some of the annotations would be unavoidably missing due to human error.


Sparsely annotated object detection (SAOD) aims to handle such a task that each training image may miss some box

*These authors contributed equally.

†Corresponding authors.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(a) Before Calibration:



(b) After Calibration:




Figure 1: An example of confidence calibration. In order to filter out two *green pillars* which not belong to the detection classes, while keep the *dog* and two *fire hydrants* as pseudo labels, different score thresholds are needed. After calibration, different training stages and detectors are able to share the same score threshold. We show the confidence scores below the image with a color indicator. The dash or solid line indicates the score below or above the threshold.

annotations. It is highly relevant to a popular topic called semi-supervised object detection (SSOD), whose training data consists of a fully labeled part and an unlabeled part. It is worth mentioning that these two tasks are complementary, since even in the semi-supervised settings, the case of missing labels may still appear in the part of labeled data.

Considering the relevance of these two tasks, we start from one of the most widely studied methods called “pseudo labeling” in semi-supervised learning (Tarvainen and Valpola 2017). Pseudo labeling utilizes the predictions of a teacher network to produce supervision signals over unlabeled data for a student network. Therefore, it is natural to transfer this idea to the case of sparse annotation. In terms of pseudo labeling, the main difference between the semi-supervised task and sparsely annotated task is that the image in the former is either fully annotated or completely unla-

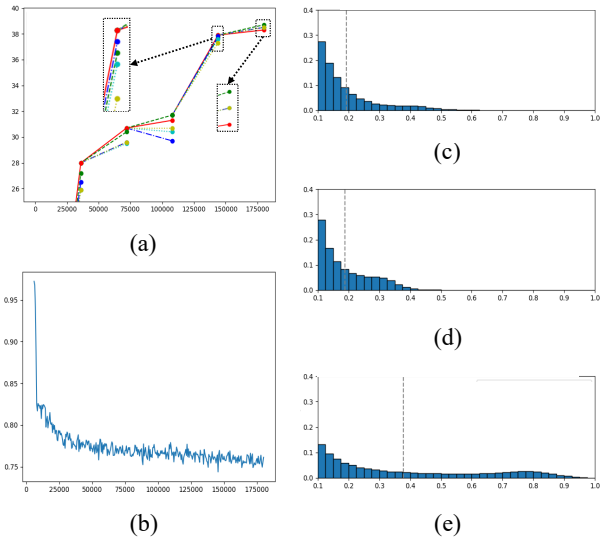


Figure 2: The meaning of confidence score varies in different training stages and among different detectors. (a) compares the effect of different score thresholds in different training stages, where the horizontal and vertical axes are iterations and mAP. The color red, green, blue, cyan and yellow represent threshold 0.5, 0.6, 0.7, 0.8 and the baseline, respectively. (b) shows the precision of pseudo labels during the training with a fixed score threshold, where the horizontal and vertical axes are iterations and precision. (c)~(e) demonstrates the score distribution of three models respectively, where the horizontal and vertical axes are confidence and sample proportion.

beled, while in the later, it is unknown how well the image is labeled. In order to guarantee the precision of pseudo labels, most of previous methods for the semi-supervised task obtain pseudo labels by taking a high score threshold to filter the detection boxes from teacher network. As for the sparsely annotated task, an additional IoU filter is indispensable to drop the predictions which have high overlap with the sparse annotations. The remain boxes along with the sparse annotations compose the final supervision signal.

Following this paradigm, some recent works have tried to generate pseudo labels for the missing annotation boxes to train the model. However, all of the previous methods (Wang et al. 2021; Rambhatla et al. 2022) adopt a fixed score threshold, which is usually set to a high value, to guarantee only a few accurate pseudo boxes could survive. Such a high score threshold drops lots of moderately confident predictions, which may contain more valuable information (e.g. tiny objects, with rare classes, partially occluded objects), leading to insufficient learning on the missing label. Therefore, a fixed score threshold is not a good choice.

Nevertheless, it is non-trivial to select a better strategy of score threshold, since we find the meaning of confidence score is changing all the time. Firstly we compare the most effective thresholds in different training stages. Concretely, we train RetinaNet under the teacher-student paradigm with

six different thresholds. During training we evaluate these models every 36k iterations, of which the best model is utilized to initialize all the models for the next stage. As shown in Fig. 2 (a) that the lower thresholds result in the better performance in the earlier stage, but the advantage is narrowing with the training going on. Finally the relatively high thresholds perform better at the end of the training, while the lowest threshold performs worse than not using pseudo labels. Besides, as shown in Fig. 2 (b), the precision of candidates decreases gradually, which reflects that a stricter threshold is required to ensure the quality of pseudo labels.

Furthermore, there are numerous factors impacting the meaning of confidence score predictions. Fig. 2 (c)~(e) show the score distribution of RetinaNet (Lin et al. 2017), FCOS (Tian et al. 2019), and Faster R-CNN (Ren et al. 2015) respectively. A wide variation is exhibited due to the different types of loss function (e.g. focal loss v.s. cross entropy loss), different score formulations (e.g. classification score multiply with centerness in FCOS), different model designs (e.g. single stage v.s. two stages). As shown in Fig. 1 (a), it turns into an engineering task to search for a proper threshold strategy to filter out the correct pseudo labels.

In this paper, we propose to give a unified and specific meaning for the confidence score to simplify the selection of score threshold in pseudo label generating. Inspired by the idea of confidence calibration (Guo et al. 2017), We propose to calibrate the predictions of the teacher model during training. After calibration, the confidence score is supposed to reflect the expectation of its precision, which is decoupled from the characters and capabilities of detectors. As shown in Fig. 1 (b), it becomes easier to set a unified threshold for different training stages and detectors after calibration.

Besides, in order to guarantee the precision of pseudo labels, it is impossible to cover all possible candidates. Hence there still exist some missing annotations not recalled. To resolve this problem, we propose a *Focal IoU Weight* (FIoU) for the classification loss. Supposing the missing annotations tend to have a low IoU with the existing annotations, FIoU down-weights the negative samples who have low IoU with the existing ground truth, with a mechanism similar to focal loss (Lin et al. 2017). Therefore, FIoU could be regarded as the complement of SAOD methods.

Our contributions are summarized as follows:

- We point out that the score threshold is the key factor in pseudo labeling methods, and analyze the obstacle of the existing methods using fixed score threshold in detail.
- We propose a general framework for sparsely annotated object detection called Calibrated Teacher, which transforms the confidence score predictions of the teacher network to fit the real precision, so that the model can adaptively generate proper pseudo labels during training.
- We propose Focal IoU Weight for classification loss to reduce the negative influence of missing annotations, which are not mined as pseudo labels.
- Extensive experiments show that our methods not only simplify the threshold tuning for different detectors, but also set new state-of-the-art in sparse settings of COCO.

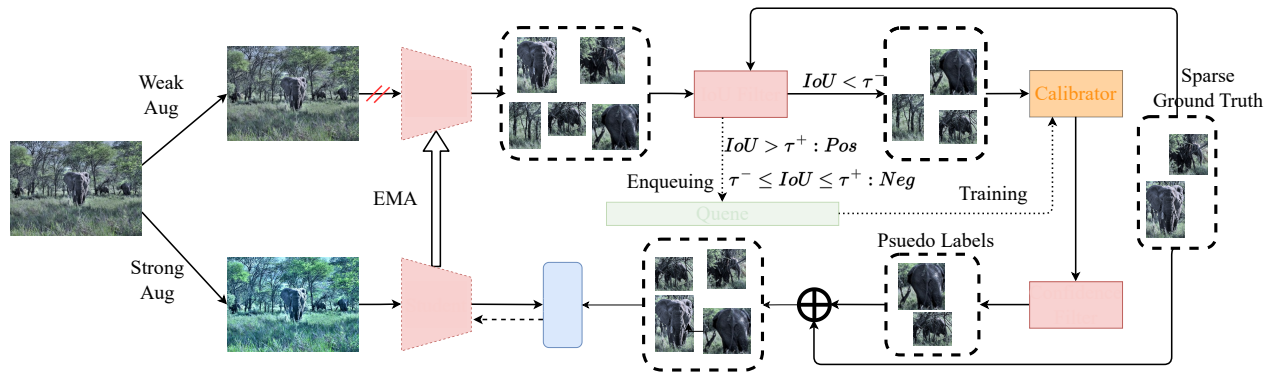


Figure 3: The overall pipeline of our framework. The solid and dotted arrows represent the data and gradient flow, respectively.

Related Work

Object Detection is one of the most fundamental problems in computer vision. Various architectures have been proposed such as one-stage methods (Redmon and Farhadi 2018; Liu et al. 2016) and two-stage methods (Ren et al. 2015). Besides, anchor-free methods draw attention gradually, which usually predict the bounding boxes based on the center points (Zhou, Wang, and Krähenbühl 2019), pseudo center points (Tian et al. 2019) or corner points (Law and Deng 2018). with the development of transformer, query-based methods (Carion et al. 2020; Zhu et al. 2020) are proposed to model the detection as a set prediction problem. Meanwhile, there exist numerous works on the design of loss function to focus on the class imbalance (Lin et al. 2017), the scale variation (Yu et al. 2016; Rezaatofghi et al. 2019) and uncertainty measurement (Li et al. 2020). However, numerous architectures and training strategies bring great difficulty to design a unified method for sparsely annotated object detection. In contrast, we propose a model-agnostic framework with calibrated confidence predictions.

Sparsely Annotated Object Detection aims at training detectors in a more practical setting, where every training image is likely to contain unannotated instances. Earlier works are usually concentrated on reducing the influence of false negative samples, which should have been assigned to the unannotated ground truth. (Wu et al. 2018; Niitani et al. 2019) down-weight the negative samples according to the IoU with existing ground truth, or the foreground probability predicted by a pre-trained detector. (Zhang et al. 2020) treat the hard negative samples as easy positive ones to avoid the large error signals. However, they could only reach the performance of models trained with partial annotations at best. Recently, (Yang, Liang, and Carin 2020) estimate the loss of anchors assigned to missing annotations by that of other positive anchors, but such rough estimation is far from the real loss and thus leads to the sub-optimal performance. (Wang et al. 2021; Rambhatla et al. 2022) adopt a two-stream detection framework to generate the pseudo labels for the missing annotation and improve the performance greatly. However, the performance is sensitive to the threshold of pseudo label selection, which brings difficulty for practical applications.

On the contrary, our method provides a general and adaptive strategy to generate pseudo labels.

Confidence Calibration for Object Detection focuses on making the confidence score of detectors consistent with the actual precision. Specifically, for the prediction with the confidence p , the expectation of its precision is also supposed to be p if the detector is well-calibrated. Although confidence calibration is rarely explored in object detection (Kuppers et al. 2020), the methodology could be referenced from image classification, such as temperature scaling (Guo et al. 2017), Platt scaling (Platt et al. 1999), histogram binning (Zadrozny and Elkan 2001), Bayesian binning (Naeini, Cooper, and Hauskrecht 2015), etc. However, all of these methods require an extra fully annotated validation set to optimize the calibrator after training. On the contrary, we optimize the calibrator with the sparsely annotated training set only, which releases the dependence on the extra clean data and makes it possible to update the calibrator during training. To our best knowledge, this is the first time to introduce the idea of confidence calibration into the sparsely annotated object detection.

Method

Overview

We propose a model agnostic framework *Calibrated Teacher* for sparsely annotated object detection. The overall pipeline is illustrated in Fig. 3. For clarity, we adopt RetinaNet (Lin et al. 2017) as the teacher and student detector if not otherwise specified. Notably, other detectors such as FCOS (Tian et al. 2019) and Faster R-CNN (Ren et al. 2015) could be replaced into our framework without any other modification.

Following the classical teacher-student paradigm (Tarvainen and Valpola 2017), given the input image I , the weakly augmented image and strongly augmented image are taken as the input of the teacher network and student network, respectively. The teacher network is updated by exponential moving average of student network, and the student network is updated in the fully supervised manner but with sparse ground truth and pseudo labels.

In order to mine the missing labels for the sparse ground truth annotations \mathbb{G} , the predictions of the teacher network

\mathbb{P}_t are split into two parts: one with low IoU with \mathbb{G} , which will serve as pseudo label candidates, and the other with high IoU, which will be pushed into a queue \mathbb{Q} with length L . After that, the confidence scores of candidates are transformed by the calibrator $\phi_c(\cdot|\theta_c)$, and will be further filtered by the score threshold τ_s to generate the final pseudo labels. Meanwhile, for every T iteration, the weights of the calibrator are trained with \mathbb{Q} . The final pseudo labels along with the sparse ground truth compose the target of the student network.

Dynamic Threshold: A Tedious Yet Effective Trick

As explained above, even for the same detector, the effective score threshold for pseudo label selection is varying during the training. A relatively low threshold could contribute to the faster convergence in the earlier stage, but harm the performance in the later stage. As a result, an intuitive solution is to adopt a dynamic threshold, which grows gradually with the training going on. In this work, the dynamic score threshold τ_s at epoch e is defined as follows:

$$\tau_s = \tau_0 + (1 - \tau_0) \cdot \frac{\ln(e - e^-)}{\ln(e^+ - e^-)} \quad (1)$$

where e^+ , e^- are the total epochs and the beginning epoch to adjust the threshold, respectively. τ_0 is the initial threshold, and we set a relatively low τ_0 for better performance.

We experimentally demonstrate the effectiveness of dynamic threshold using RetinaNet and FCOS, as shown in Tab. 2. However, this strategy still faces the challenge that different detectors do not share the hyper-parameters e^- and τ_0 , which requires extra effort to tune. Furthermore, it could be found in Tab. 2 that the performance of Faster R-CNN are especially sensitive to the threshold of pseudo label selection. Even a slight change of the threshold might lead to a performance drop over 10%. For such detectors, the pre-defined dynamic strategy has to be designed carefully, which extends the cycle of the model designing greatly.

Calibrated Teacher: Simplify Threshold Tuning

In order to address the shortcomings of the dynamic threshold, we first analyse the mechanism behind this method. An ideal threshold is closely related to the confidence distribution of the teacher network’s prediction, which varies in the different training stages and among different detectors. The dynamic strategy tries to fit the variation of this distribution and adjust the threshold accordingly. However, the variation of confidence distribution is hard to parameterized, which brings great difficulty of designing a pre-defined rules of adjusting threshold.

Hence, it is unwise to leverage a pre-defined rules to adapt to the confidence distribution changing continuously. Instead, we turn to focus on the adjustment of confidence distribution and leave the threshold unchanged. This is ideally equivalent to dynamic threshold. However, with the fixed threshold, all we need is to calibrate the confidence distribution into some specific form, which is more robust and easier than adapting the threshold to a varying and non-parametric distribution. Furthermore, once the confidence is calibrated, an effective threshold could also work in other cases, *i.e.*, in other training stages or using other detectors.

Then the problem turns into what the target confidence distribution looks like and how we transform the original distribution to that. For the first problem, considering that the precision is the most direct and crucial indicator of the quality of pseudo labels, we expect the confidence score to reflect the real precision. This is the same as the target of confidence calibration in object detection (Kuppers et al. 2020), which defines a calibrated confidence (unbiased confidence) as follows.

Supposing the prediction set is $\{\hat{c}_i, \hat{p}_i, \hat{b}_i\}$, where \hat{c}_i , \hat{p}_i and \hat{b}_i are the predicted category, the confidence and the bounding box, respectively. Accordingly, the ground truth assigned to each prediction is denoted as $\{c_i, b_i\}$. Then, the prediction is perfectly calibrated to be unbiased if

$$\mathbb{P}(m = 1|\hat{p} = p) = p. \quad (2)$$

where $m = 1$ denotes a correct match *i.e.*, $\hat{c}_i = c_i$ and $\text{IoU}(\hat{b}_i, b_i) > \tau$ (τ is the pre-defined threshold), while $m = 0$ indicates a mismatch.

Therefore, we naturally consider a regression model for confidence calibration. Specifically, we adopt a logistic calibration model $\phi_c(\cdot|\theta_c)$, and minimize the Negative Log Likelihood of $\phi_c(\hat{p}_i|\theta_c)$ to update the model.

Online Calibration

In the field of model calibration, the parameters of calibrators are optimized with the validation set at the end of training. However, there exist two extra challenges in our framework. The first challenge is that there is no validation set available. Although we could utilize training set instead, it is actually noisy due to the missing annotation. The calibrator optimized with such dataset would tend to output the low confidence, as it is trained to reduce the confidence of some correct prediction assigned with the missing ground truth. The second challenge is that the meaning of the score prediction is varying during training, and thus we have to calibrate the teacher network timely, which requires the calibrate model is updating during the training stage.

In order to tackle the first challenge, we divide the teacher’s prediction in the training set into two parts by an IoU filter. Bounding boxes having a high IoU with at least one of the existing ground truth are more likely to be assigned to the sparse annotation, while those having a low IoU with all the existing annotations possibly cover the missing annotation. Theoretically, the statistics of the two parts are similar, therefore we could optimize the calibrator with the sparse annotations and the corresponding predictions (the high IoU parts).

For the second challenge, we introduce a queue \mathbb{Q} with length L during training. At each iteration, for the teacher network’s predictions with high IoU, $\{\hat{p}_i, m_i\}$ are pushed into \mathbb{Q} , while those from the earliest image are popped, so that \mathbb{Q} only contains the training data in the most recent batches. Therefore, $\phi_c(\cdot|\theta_c)$ is updated with elements in \mathbb{Q} . For more details, please refer to Algorithm 1.

Focal IoU Weight

Although the calibrated teacher is able to generate pseudo labels properly, it is impossible to cover all the missing an-

Algorithm 1: Pseudo labels in Calibrated Teacher

Input: Image I , Sparse annotations $\mathbb{G} = \{c_j, b_j\}_{j=1}^M$,
Teacher model predictions at t step $\mathbb{P}_t = \{\hat{c}_i, \hat{p}_i, \hat{b}_i\}_{i=1}^N$,
Calibration queue \mathbb{Q} with length L , Calibration model ϕ_c
with weights θ_c , initialized as an identity mapping.
Parameter: Normal and stricter IoU thresholds τ^- , τ^+ ,
Score threshold τ_s , Calibration model update interval T .
Output: Pseudo labels \mathbb{L} , Calibration model weights θ_c , Up-
dated queue \mathbb{Q} ;

- 1: Let $\mathbb{L} = \mathbb{G}$.
- 2: **for** $i = 1, 2, 3, \dots, N$ **do**
- 3: $\text{iou} = \max_j (\text{IoU}(\hat{b}_i, b_j))$ s.t. $c_i = c_j$
- 4: $\text{score} = \phi_c(\hat{p}_i | \theta_c)$ // Calibrate model prediction
- 5: **if** $\text{iou} < \tau^-$ and $\text{score} > \tau_s$ **then**
- 6: $\mathbb{L} = \mathbb{L} \cup \{(\hat{c}_i, \hat{b}_i)\}$ // Add to pseudo labels
- 7: **end if**
- 8: **if** $\text{iou} > \tau^-$ **then**
- 9: **if** $\text{iou} > \tau^+$ **then**
- 10: $\mathbb{Q}.\text{enqueue}((\hat{p}_i, 1))$
- 11: **else**
- 12: $\mathbb{Q}.\text{enqueue}((\hat{p}_i, 0))$
- 13: **end if**
- 14: **end if**
- 15: **end for**
- 16: **if** $t \bmod T = 0$ **then**
- 17: // Update the weights of calibration model
- 18: $\theta_c = \arg \min_{\theta_c} (NLL(\phi_c(\mathbb{Q} | \theta_c)))$
- 19: **end if**
- 20: **return** $\mathbb{L}, \mathbb{Q}, \theta_c$.

notation. On the one hand, even the model trained with the complete annotation could fail to detect all instances in the training set. On the other hand, the high recall of missing annotation could also result in the low precision, since we have to set a low threshold and thus preserve many false positive predictions. Therefore, the missing but not recalled annotation would still make some positive samples regarded as negative samples. These false negative samples have little influence on the two-stage detectors like Faster R-CNN due to the sampling strategy. However, for the one-stage detector like RetinaNet, they are unlikely to be ignored. Furthermore, the false negative samples usually have a relatively high confidence of foreground, they would further harm the training due to the commonly used focal loss.

As a result, it is necessary to reduce the influence of false negative samples. Inspired by (Zhang et al. 2020), we adopt the IoU with the existing ground truth as the criterion for judging negative samples. It is safe to assign a high weight to those negative samples with high IoU with the existing ground truth labels, as they are less likely to match the missing labels, and vice versa. Analogous to the mechanism of focal loss, where high weights are assigned to negative samples with high foreground confidence, we propose a *Focal IoU Weight* (FIoU) for the classification loss. Considering the formulation of focal loss for the negative samples:

$$\text{FL}(x) = -\alpha_t x^\gamma \log(1 - x) \quad (3)$$

where x is the confidence of the foreground. α_t, γ are hyper-parameters. Similarly, for the given negative sample, the FIoU weight for the classification loss is defined as:

$$\text{FIoU}(\text{iou}) = w_0 + k \cdot \text{FL}(\text{iou}) \quad (4)$$

where w_0 and k are hyper-parameters, and the iou is the max IoU between the negative samples and the existing labels and pseudo labels.

Experimental Results

Datasets

Recent SAOD methods (Yang, Liang, and Carin 2020; Wang et al. 2021; Zhang et al. 2020; Rambhatla et al. 2022) are mainly evaluated on the challenging COCO-2017 dataset (Lin et al. 2014). However, these methods adopt different ways to reconstruct the original training set into the sparsely annotated one. For comprehensive comparison, we evaluate our method in almost all of the existing sparse settings:

Split-1: Following (Wang et al. 2021; Zhang et al. 2020; Rambhatla et al. 2022), for each category c in the training set, $p\%$ annotations are deleted randomly, where $p = \{30, 50, 70\}$.

Split-2: Following (Rambhatla et al. 2022), for each category c in the training set, images containing c are firstly selected. Then for each image, all annotations of c are deleted simultaneously with a probability $p\%$, where $p = \{30, 50, 70\}$. We make sure that each image contains at least one annotation.

Split-3: Following (Rambhatla et al. 2022), we deleted $p\%$ annotations randomly in a class-agnostic fashion, where $p = \{30, 50, 70\}$.

Easy / Hard / Extreme: These training sets come from (Wang et al. 2021). For each image in the Easy split, one annotation is deleted randomly; For each image in the Hard split, half of the annotations are deleted randomly; For each image in the Extreme split, only one annotation is preserved randomly.

Implementation Details

We take experiments on three common detection methods with ImageNet (Deng et al. 2009) pretrained ResNet101 (He et al. 2016) and ResNet50 for the comparison with state-of-the-art and ablation study, respectively. Our models are trained for $180k$ iterations with a total batch size 16. The learning rate is initialized as 0.01 and gradually decreases to 0.001 and 0.0001 at $120k$ and $160k$ iterations. Other hyper-parameters of the architecture and training schedule are consistent with the implementation in (Chen et al. 2019).

As for the confidence calibration, τ^+, τ^-, τ_s are set to 0.75, 0.6 and 0.7 for all detectors, respectively. However, most of the prediction of the teacher has an especially low confidence, *e.g.*, lower than 0.2. Prediction under such a low confidence could hardly cover a satisfying result. Therefore, we only consider the prediction whose original confidence is higher than 0.4 during training to reduce the cost of memory and computation. We adopt Plat scaling as the form of the calibrator for its simplicity and effectiveness. For the calibrator training, T is set to 500 and L is the number of predictions of 8000 images. For the FIoU, w_0 and k are set to

Method	Split-1			Split-2			Split-3		
	30%	50%	70%	30%	50%	70%	30%	50%	70%
Full Annotation									
Baseline	39.3	37.5	34.1	39.1	37.4	35.5	39.2	37.6	34.3
BRL (Zhang et al. 2020)	–	32.7	–	–	–	–	–	–	–
Co-mining (Wang et al. 2021)	36.4	32.8	24.9	36.7	33.0	24.8	36.8	32.5	25.0
Ours + RetinaNet	40.5	39.3	36.7	40.1	38.9	37.4	40.5	39.0	36.6
Unbiased Teacher [†] (Liu et al. 2021)	32.0	31.1	27.9	36.4	32.9	31.4	36.0	32.1	30.1
SAOD [†] (Rambhatla et al. 2022)	38.5	36.2	33.0	40.0	37.2	35.9	39.7	37.4	35.9
Ours + Faster R-CNN [†]	41.0	39.3	35.3	40.8	39.2	36.5	41.0	39.1	35.9

Table 1: Comparison with recent sparsely annotated object detection methods on three splits of COCO dataset. All the methods adopt ResNet 101 with FPN as backbone. [†] denotes using Faster R-CNN rather than RetinaNet as the detector.

Detector	Base	Fixed Threshold			DT	CT
		0.5	0.7	0.9		
RetinaNet	37.5	36.9	37.4	37.5	38.5	38.5
FCOS	37.9	38.0	38.2	38.0	38.7	39.0
Faster R-CNN	37.9	11.6	27.5	39.2	38.6	39.3

Table 2: Generalization of Calibrated Teacher on different detectors. Base is the abbreviations of baseline. DT and CT represent dynamic threshold and Calibrated Teacher, respectively.

0.5 and 1.5, while the α_t and γ stay consistent with focal loss.

The strong augmentation used in our framework contains random resize, contrast normalization, histogram equalization, random solarization, color balance, contrast, brightness, sharpness and posterization, while the weak augmentation includes nothing but random flipping.

Comparison with State-of-the-Art

In this subsection we compare our methods with state-of-the-art. Tab. 1 reports the results in Split-1,2,3. It could be seen that our framework with both RetinaNet and Faster R-CNN could set new state-of-the-art in all sparse settings. Specifically, when adopting RetinaNet as the detector, our methods surpass the counterparts by a large margin (12.8 mAP at most and 7.3 mAP on average). Meanwhile, our methods with Faster R-CNN have the superiority of 3.1 mAP at most and 1.6 mAP on average. It is encouraging to find that our framework improves the performance of baseline methods significantly, which is close to the models trained with the complete training set. This inspires us that in order to relieve the burden of annotating, a sparsely annotated dataset could replace the completely annotated one to some extent. Furthermore, different from the recent work (Rambhatla et al. 2022), which could only be applied on two-stage detectors due to the dependence on the architecture modification of RPN, our framework is model-agnostic, which is convenient and flexible for practical application. In conclusion, the strong performance demonstrates the effectiveness of our methods.

Calibrated Teacher	Focal IoU	Easy	Hard	Extreme
		37.1	35.3	26.7
✓		37.7	35.5	27.4
	✓	37.7	35.6	27.8
✓	✓	38.1	36.4	28.7

Table 3: Ablation study on Easy, Hard and Extreme sets

Generalization to Multiple Detectors

As explained before, one of the advantages of Calibrated Teacher over dynamic threshold is that different detectors are able to share the same threshold. In order to test this conjecture, we evaluate Calibrated Teacher and dynamic threshold with another two detectors FCOS and Faster R-CNN in Split-1 with 50% missing annotations. For dynamic threshold, we first train models with different fixed thresholds, then the threshold with best performance would serve as τ_0 for better performance. We set e^- to 18 for all three detectors to keep the same dynamic strategy. For Calibrated Teacher, we set the same τ_s as RetinaNet.

Tab. 2 summarizes the comparison. It could be found that the dynamic threshold does not always outperform the fixed threshold. This might result from the difficulty of designing a proper dynamic strategy, as different detectors may not share the same effective strategy due to the different confidence distribution. However, Calibrated Teacher is able to surpass both dynamic threshold and fixed threshold in all three detectors, even without adjusting hyper-parameters according to the property of them. Hence we could conclude that Calibrated Teacher is a general framework with better performance. It simplifies the design of hyper-parameters by providing an adaptive strategy to select proper pseudo labels in different training stages, and making it possible that different detectors could share the same hyper-parameters.

Ablation Study

Effectiveness of Each Module We further validate the effectiveness of each module in Easy, Hard and Extreme sets. As shown in Tab. 3, both Calibrated Teacher and Focal IoU outperform the baseline in all three settings when




Figure 4: Qualitative visualization for the results of sparsely annotated object detection. The 1st and 2nd rows visualize the pseudo labels generated during training, where the green boxes correspond to the available ground truth, and the red boxes indicate the mined pseudo labels. The 3rd row visualizes some failure while mining pseudo labels, where the green and red boxes are available ground truth and pseudo labels.

τ_s	1	0.55	0.65	0.7	0.75	0.85
mAP	37.5	37.6	38.3	38.5	38.1	38.0

Table 4: Sensitivity of τ_s for the score filter after calibration.

w_0	0.5	0.5	0.5	0	1
k	1.0	1.5	2.0	1.5	1.5
mAP	39.2	39.3	39.3	33.9	39.3

Table 5: Sensitivity of hyper-parameters in Focal IoU.

used alone, with the superiority of 0.7 mAP and 1.1 mAP at most, respectively. Furthermore, when combined these two modules, the performance gain is greater than the sum of the respective gains. Taking experiments in Hard set as an example, Calibrated Teacher and Focal IoU could only bring the improvement of 0.2 mAP and 0.3 mAP, but the combination could result in a performance gain of 1.1 mAP. This is because the two modules are complementary. Focal IoU down-weights the false negative anchors to avoid great misleading supervision, which is helpful to train a better model to generate better pseudo labels.

Impact of Hyper-parameters In this subsection we first compare the Calibrated Teacher with different τ_s in Split-1 with 50% missing annotations. It could be found in Tab. 4 that all five τ_s could outperform the baseline, which demonstrates the superiority over using a fixed threshold. Moreover, we set $\tau_s = 0.7$ due to its best performance.

Then we evaluate Focal IoU with different w_0 and k in Split-1 with 50% missing annotations. Tab. 5 reports the

L	4000	8000	16000
mAP/AP ₅₀	33.6/51.8	33.0/50.7	33.0/50.8

Table 6: Comparison in the queue size L .

T	250	500	1000
mAP/AP ₅₀	33.0/50.6	33.0/50.7	33.2/51.0

Table 7: Comparison in the training interval T .

results that Focal IoU is more robust to k than w_0 . When $w_0 = 0$, there exists a great drop of performance. This is because most negative samples have low IoU with ground truth bounding boxes, and of which most are true negative samples. These samples contribute a lot to the classification loss, but their weights tend to be zero when w_0 is very small. Therefore, we need a trade-off to balance down-weighting the false negative samples and saving the true negative samples. Hence we use $w_0 = 0.5$ and $k = 1.5$ in this work.

Next we analyze the effect of the queue size. Concretely, we train the models for 90k iterations in Hard set, and decrease the learning rate at 60k and 80k iterations. As shown in Tab. 6, when the queue size L is the number of predictions of 4000 images, the model performs best. This is because a smaller L means that the data stored in the queue comes from more recent iterations and thus has a more similar distribution to the current data. Therefore, if L is not too small to serve as a training set, a relatively small L could result in better performance.

Finally we experiment with different intervals at which




Figure 5: Comparison on the score distribution and ECE between original and calibrated teacher’s prediction during training. The first and second columns are the results in the last 500 iterations of RetinaNet and Faster R-CNN. For each sub-figure, the first row shows the score distribution, where the horizontal and vertical axes are confidence and sample proportion. The second row illustrates the precision of the pseudo label candidates in each score interval, where the horizontal and vertical axes are confidence and precision, while the area of red and orange regions represent the degree of the under-confidence and over-confidence, respectively.

the calibrator is trained in Hard set. Specifically the model is trained for $90k$ iterations and L is set 8000. Tab. 7 reports the results, and it turns out that our method is robust to the interval T .

Qualitative Results The first two rows of Fig. 4 show the qualitative results of the pseudo labels generated. Our methods are able to mine the small or occluded objects in dense scenes, which tend to miss in annotating. For other more obvious objects (like the zebras in the second row), our methods could provide pseudo labels with extremely high quality, releasing the burden of manual working.

Moreover we visualize the qualitative effect of confidence calibration. Concretely we compare the score distribution and Expected Calibration Error (ECE) (Naeini, Cooper, and Hauskrecht 2015) of the pseudo label candidates of RetinaNet and Faster R-CNN in Fig. 5. It could be found that after calibration the confidence score is closer to the precision in each interval, and the ECE of the prediction is lower. It indicates that our methods calibrate score distribution into a unified and meaningful form, so that one can design the score threshold more easily and effectively.

Failed Cases

In the 3rd row of Fig. 4, we further show some failed cases when generating the pseudo labels and analyse the possible reasons. In the 1st column we miss three hot dogs while recalling the others. Since the lowest calibrated confidence of the recalled hot dogs is close to the τ_d (0.71 v.s. 0.7), the

missing hot dogs are very likely to be refused by the confidence filter. In the 2nd column we miss one instance in a pile of donuts. This is possibly because the missing donuts are too close to be distinguished from its neighbors, making its corresponding prediction refused by the IoU filter. In the 3rd column the person behind the mesh wire is not recalled, which might result from the great variation in color. In the 4th column we fail to find the ball in the left mostly due to its extremely tiny volume. Finally the right person is excluded from the pseudo labels, which is possibly because only an arm is visible in the picture. Although our method fails to recall some instances, it is actually unavoidable when guarantee the precision of pseudo labels, which is why we need the Focal IoU Weight as a complement.

Conclusion and Limitations

In this work we experimentally explain the obstacle of recent methods of sparsely annotated object detection (SAOD). Hence we propose a Calibrated Teacher that adjusts the confidence of the teacher network’s prediction adaptively to provide reliable pseudo labels. We further present the Focal IoU Weight as a complement to down-weight the false negative samples. Extensive methods show that our methods not only outperform the state-of-the-art counterparts, but also simplify the threshold tuning greatly. One limitation is that we only adopt a simple strategy for calibration, it is unclear that whether other advanced calibrators could help the performance. We leave it for future work.

Acknowledgments

This work is partially supported by the NSFC fund (61831014), in part by the Shenzhen Science and Technology Project under Grant (CJGJZD20200617102601004, ZDYBH201900000002).

References

Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229.

Chen, K.; Wang, J.; Pang, J.; Cao, Y.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Xu, J.; et al. 2019. MMDetection: Open mmlab detection toolbox and benchmark. arXiv:1906.07155.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 248–255.

Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *International conference on machine learning*, 1321–1330.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

- Kuppers, F.; Kronenberger, J.; Shantia, A.; and Haselhoff, A. 2020. Multivariate confidence calibration for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 326–327.
- Law, H.; and Deng, J. 2018. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, 734–750.
- Li, X.; Wang, W.; Wu, L.; Chen, S.; Hu, X.; Li, J.; Tang, J.; and Yang, J. 2020. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *Advances in Neural Information Processing Systems*, 33: 21002–21012.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*, 21–37.
- Liu, Y.-C.; Ma, C.-Y.; He, Z.; Kuo, C.-W.; Chen, K.; Zhang, P.; Wu, B.; Kira, Z.; and Vajda, P. 2021. Unbiased teacher for semi-supervised object detection. arXiv:2102.09480.
- Naeini, M. P.; Cooper, G.; and Hauskrecht, M. 2015. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2901.
- Niitani, Y.; Akiba, T.; Kerola, T.; Ogawa, T.; Sano, S.; and Suzuki, S. 2019. Sampling techniques for large-scale object detection from sparsely annotated objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6510–6518.
- Platt, J.; et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3): 61–74.
- Rambhatla, S. S.; Suri, S.; Chellappa, R.; and Shrivastava, A. 2022. Sparsely Annotated Object Detection: A Region-based Semi-supervised Approach. arXiv:2201.04620.
- Redmon, J.; and Farhadi, A. 2018. Yolov3: An incremental improvement. arXiv:1804.02767.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Rezatofghi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 658–666.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30.
- Tian, Z.; Shen, C.; Chen, H.; and He, T. 2019. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9627–9636.
- Wang, T.; Yang, T.; Cao, J.; and Zhang, X. 2021. Co-mining: Self-Supervised Learning for Sparsely Annotated Object Detection. *Proceedings of the AAAI conference on artificial intelligence*.
- Wu, Z.; Bodla, N.; Singh, B.; Najibi, M.; Chellappa, R.; and Davis, L. S. 2018. Soft sampling for robust object detection. arXiv:1806.06986.
- Yang, Y.; Liang, K. J.; and Carin, L. 2020. Object detection as a positive-unlabeled problem. arXiv:2002.04672.
- Yu, J.; Jiang, Y.; Wang, Z.; Cao, Z.; and Huang, T. 2016. Unitbox: An advanced object detection network. In *Proceedings of the 24th ACM international conference on Multimedia*, 516–520.
- Zadrozny, B.; and Elkan, C. 2001. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *ICML*, volume 1, 609–616.
- Zhang, H.; Chen, F.; Shen, Z.; Hao, Q.; Zhu, C.; and Savvides, M. 2020. Solving missing-annotation object detection with background recalibration loss. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1888–1892.
- Zhou, X.; Wang, D.; and Krähenbühl, P. 2019. Objects as points. arXiv:1904.07850.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *International Conference on Learning Representations*.