





# Semi-supervised Object Detection via Virtual Category Learning

Changrui Chen<sup>1</sup>, Kurt Debattista<sup>1</sup>, and Jungong Han<sup>1,2</sup> 

<sup>1</sup> University of Warwick, WMG

{Changrui.Chen; K.Debattista}@warwick.ac.uk

<sup>2</sup> Aberystwyth University, Computer Science  
juh22@aber.ac.uk

**Abstract.** Due to the costliness of labelled data in real-world applications, semi-supervised object detectors, underpinned by pseudo labelling, are appealing. However, handling confusing samples is nontrivial: discarding valuable confusing samples would compromise the model generalisation while using them for training would exacerbate the confirmation bias issue caused by inevitable mislabelling. To solve this problem, this paper proposes to use confusing samples proactively without label correction. Specifically, a virtual category (VC) is assigned to each confusing sample such that they can safely contribute to the model optimisation even without a concrete label. It is attributed to specifying the embedding distance between the training sample and the virtual category as the lower bound of the inter-class distance. Moreover, we also modify the localisation loss to allow high-quality boundaries for location regression. Extensive experiments demonstrate that the proposed VC learning significantly surpasses the state-of-the-art, especially with small amounts of available labels.

**Keywords:** Semi-supervised Learning, Object Detection

## 1 Introduction

The Deep Learning community is suffering from the expensive labelling cost of large-scale datasets. Semi-supervised learning, which makes use of limited labelled data combined with large amounts of unlabelled data for training, has shown great potential to reduce the reliance on large amounts of labels [25,4,34,22]. In particular, this work focuses on object detection, a task that typically requires a significant effort to label, under the semi-supervised setting, especially with extremely small amounts of labelled data.

Pseudo labelling [11], one of the most advanced methods in semi-supervised learning, has been recently introduced to semi-supervised object detection (SSOD) [25,43,41,40]. Unlabelled data are automatically annotated by the detector itself [35] (or an exponential moving average version [25]) and then fed back to re-optimize the detector. Typically, due to the limited diversity of samples in an extremely small available set of labelled data, the detector is usually unable

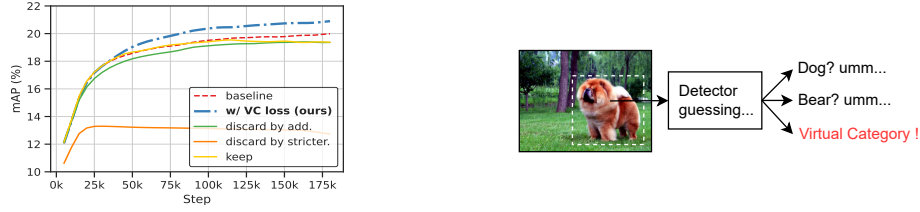


Fig. 1: Left: The mAP of a semi-supervised detector [25] with a preset confidence score filtering on 1% labelled MS COCO [44]. The mAP sees a decrease with all strategies for dealing with confusing samples (*e.g.*, the bear-like dog at the right) except for our VC learning. The additional filtering mechanism (add.) is the temporal stability verification proposed in this paper. Stricter filtering (stricter.) is simply raising the threshold in the score filtering from 0.7 to 0.8. Right: Illustration of the basic idea of our virtual category.

to make a confident decision on some unseen confusing samples when inferring their pseudo labels. Two strategies are usually adopted in SSOD to deal with the confusing samples: a) discarding all of them by a strict filtering mechanism [35], or b) retaining them with all potential labels [41]. However, none of these two options is optimal. Admittedly, the value of confusing hard samples is clear to see, since hard example mining [33] has successfully proven its effectiveness in fully-supervised learning. If all confusing samples are rejected by a strict filtering mechanism in SSOD, their positive contributions will be wasted, while the remaining well-fitted samples only marginally contribute to performance improvements. On the contrary, simply keeping them all does not work either due to the involvement of too many incorrect pseudo labels. Arbitrarily optimising semi-supervised detectors with these noisy labels will lead to the confirmation bias issue [2]. To demonstrate these points, we show in Fig. 1 the mean Average Precision (mAP) of a semi-supervised detector with different strategies on 1% labelled MS COCO [44]. We observe noticeable performance degradations when either choosing one stricter filtering mechanism (orange line) or adding an additional filtering (green line) to reject confusing samples. Likewise, simply keeping all pseudo labels of confusing samples (yellow line) also ends up with a decreased mAP since the unreliable pseudo labels aggravate the confirmation bias issue and even give rise to training collapse.

Therefore, efforts are dedicated to exploring how to correct the biased pseudo labels to utilise the confusing samples without any risks. Existing methods [38,22] initially investigate relatively easy tasks such as classification on CIFAR [18]. However, progress has not yet been made for a complex task, such as object detection with extremely small amounts of labelled data, due to its difficulty in correcting object pseudo labels. Then, a question arises: *what if we do not discard any confusing samples but consider their contributions, which may not necessarily need the concrete label information, during the model training?* This paper answers this question by proposing a novel Virtual Category (VC) learn-

ing, based on an observation that the confusing samples with doubtful pseudo labels can be profitable for training.

Specifically, we discover that building a potential category (PC) set consisting of the possible categories of a confusing sample  $x$ , compared to determining the exactly correct label for it, seems more feasible. Therefore, instead of selecting the correct one from the PC set, which is usually hard, we compromise by proposing a virtual category label to *take the place of all unreliable labels in the PC set*. As shown in the example of the right part of Fig. 1, the PC set of this object consists of two categories (*i.e.*, dog and bear). A new modified Cross Entropy loss function, namely VC loss, allows the detector to be optimised with the virtual category label. By doing so, ignoring the categories in the PC set will disable the corresponding output logits, thus avoiding any wrong gradients to mislead the model optimisation. Most importantly, the proposed virtual category specifies a reasonable lower bound for the inter-class distance. Hence, the decision boundary can consistently benefit from the confusing data without suffering from the confirmation bias issue. With regards to the PC set, we investigate multiple methods to build it, including temporal stability verification and cross-model verification. As can be seen in Fig. 1, the mAP of the model armed with the proposed VC loss (dot-dashed line) sees a significant increase due to the effective use of the confusing samples. Furthermore, we also enhance localisation training by decoupling the horizontal and vertical location qualities of a pseudo bounding box, which enables us to use the high-quality boundaries to optimise the bounding box regression head.

We apply our VC learning to a semi-supervised detector and evaluate it on MS COCO [44] and Pascal VOC [13] datasets with extremely low label ratios, including 0.5%, 1%, 2%, 5% and 10%. On MS COCO, our method achieves 19.46 mAP with only 586 labelled images; this outperforms some recently published semi-supervised detectors [35,43] with 1000+ labelled images. The contributions of this paper are summarised as follows:

- We propose VC learning to alleviate the confirmation bias issue caused by confusing samples in SSOD. The proposed method takes the initiative to use confusing samples with unreliable pseudo labels for the first time.
- A detailed explanation of the VC learning is provided to demonstrate the mathematical feasibility and intuitive effect. Our intriguing findings highlight the need to rethink the usage of confusing samples in SSOD tasks.
- We demonstrate how the proposed method surpasses state-of-the-arts by a significant margin on two benchmark datasets with different label ratios.

## 2 Related Works

**Object Detection** aims to distinguish foreground objects in images or videos and identify them. Object detectors so far can be divided into three types: 1) Two-stage detectors [31,7], represented by Faster RCNN [31]; 2) One-stage detectors [24,42], such as the YOLO series [30,6]; and 3) Point-based Detectors [12,37,21], such as Center Net [12]. The main difference between two-stage

and one-stage detectors lies in whether an additional module is used to generate candidate region proposals. Point-based detectors discard anchor boxes and instead use points and sizes to represent objects. In this paper, Faster RCNN, one of the most widely used detectors, serves as our baseline model.

**Semi-supervised Learning** is a training scheme that uses only a very small amount of labelled data and a large amount of unlabelled data to train a machine learning model. Recently proposed semi-supervised methods mainly focus on image classification [34,4]. Consistency regularisation is one of the most advanced methods. It requires models to produce consistent outputs when the inputs are perturbed. Image augmentations, such as flipping, Cutout [10], or Gaussian Blurring, are usually applied to perturb inputs [20,3]. Some solutions take advantage of adversarial learning and proposed learnable adversarial augmentations [27]. Inspired by the noisy label learning method [17], negative sampling [8] improved the performance of semi-supervised algorithms by randomly sampling a non-target label for an unlabelled data. Alternatively, Pseudo labelling [11], which is used in this paper, is another semi-supervised technique that has proven successful. It minimises the entropy of the predictions of unlabelled data by generating pseudo labels for them.

**Semi-supervised Object Detection (SSOD)** originates from semi-supervised classification, where only a small amount of bounding box labelled data and numerous unlabelled data are available for training a detector. Most of the recently proposed SSOD methods can be broadly categorised into two different types: 1) consistency-based and 2) self-supervised (pseudo-labelling) methods. Consistency-based semi-supervised object detectors, such as CSD-SSD [19], when dealing with unlabelled images, apply consistency regularisation on the predicted classification probability vectors and regression vectors of the input image and its mirror version. Alternatively, self-supervised detectors [25,43,41,35] were inspired by the pseudo-labelling classifier. Although the self-supervised detectors show good potential, they are struggling with the confirmation bias when training with confusing samples. This paper alleviates this issue via VC learning.

### 3 Methodology

In this section, the overall problem is first defined. The VC learning and its explanation are subsequently described.

#### 3.1 Problem Definition

In SSOD, two data subsets  $\mathcal{D}^l$  and  $\mathcal{D}^u$  are given for model optimisation, where  $\mathcal{D}^l = \{(x_n^l, b_n^l) | n=0\}^{N^l}$  is the subset with available ground truth bounding box  $b^l$ ,  $\mathcal{D}^u = \{x_n^u | n=0\}^{N^u}$  is the unlabelled subset.  $N^l$  and  $N^u$  are the number of labelled and unlabelled data. Usually,  $N^u \gg N^l$ .  $b = [x_1, y_1, x_2, y_2, cls]$  is the label – bounding box, where the first 4 numbers indicate the coordinates of the top-left and bottom-right points and  $cls$  is the index of the category label. This paper follows the self-supervised framework with a score filtering mechanism [25] to

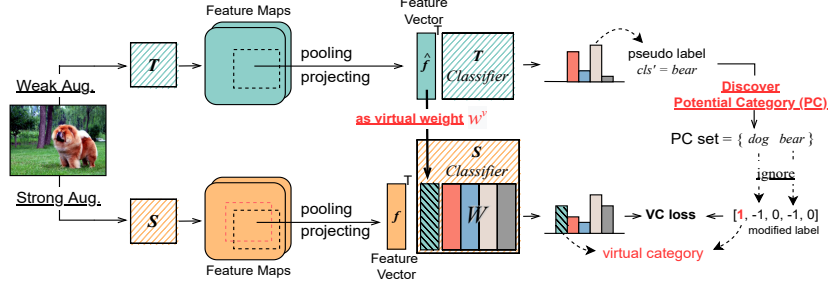


Fig. 2: The pipeline of the VC learning when dealing with a confusing sample. The regression branch is ignored for simplicity.  $T$  and  $S$  represent the teacher and student model. The black dashed line box on feature maps indicates the pseudo label position. When training the student classifier with a RPN-generated proposal box (pink dashed box in the figure) assigned to this pseudo label, the weight matrix  $W$  of the student classifier is extended by a virtual weight  $w^v$ , which is transformed from the teacher feature vector  $\hat{f}$  of the pseudo label area.

generate pseudo labels  $b^l$  of  $x^u$ . As shown in Fig. 2, two detectors  $T$  and  $S$  sharing the same architecture are introduced. The parameters of the teacher detector  $T$  are updated by the parameters of the student  $S$  with a momentum.

### 3.2 Virtual Category Learning

In Fig. 2, the teacher classifier produces the categorical probability of the candidate object (black dashed box in Fig. 2) according to the pooled and projected feature vector  $\hat{f}$  of it. Typically, the category with the highest probability, *e.g.*,  $cls' = bear$  here, will then be used as the pseudo label of the candidate proposal box assigned to this instance during the student training. However, incorrect pseudo labels may mislead the training.

In this paper, we propose VC learning which modifies the pseudo category label with an additional virtual category to allow the student model to be optimised safely by confusing samples. Once the initial pseudo label  $cls' = bear$  is obtained, a potential category (PC) discovery operation is performed to construct a PC set  $\{dog, bear\}$  for this training sample. We find that PC discovery is relatively feasible compared to designing a correction function:  $g(cls'|f) = GT$ , especially when the labelled subset  $\mathcal{D}^l$  is much smaller than  $\mathcal{D}^u$ . The discovery function will be introduced in the following Section 3.3.

If the PC set contains more than one category, it means that this sample is confusing to the detector. In Fig. 2, the bear-like dog is a confusing sample with  $PCset = \{dog, bear\}$ . To allow all the proposal boxes (*e.g.*, pink dashed box in Fig. 2) assigned to this bear-like dog consistently contribute to the optimisation of the student detector  $S$  rather than arbitrarily discarding them, the weight matrix  $W$  in the student classifier is extended by an additional weight vector  $w^v$  named *virtual weight*. The virtual weight  $w^v$  is obtained from the feature vector

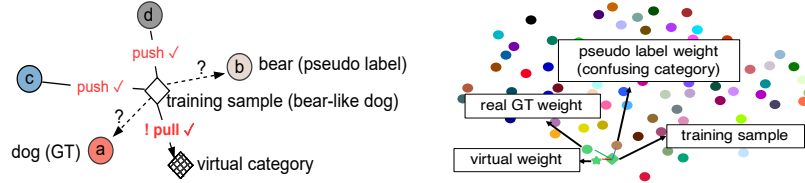


Fig. 3: left) Explanation of VC loss in the feature space. The circles are the embeddings of the category cluster centres. The hollow diamond is the embedding of the training sample. The cross-hatch diamond is the embedding of the virtual category. right) T-SNE visualisation of the virtual weight (star), learned classifier weights (circles), and a training sample (diamond).

$\hat{f}$ , which is the feature of this confusing sample in the teacher model. Thus, the virtual weights for different confusing samples are different. The normalisation and scaling are applied to  $\hat{f}$  to ensure the norm of the produced  $w^v$  and the norm of the weight vectors in  $W$  are in the same range. We investigate two scaling factors, *i.e.*, a constant value and an adaptive factor calculated by averaging the norm of all the weight vectors in  $W$ . With the extended weight matrix, the size of the student classifier output (*i.e.*, logits) is therefore increased by 1:

$$f^\top \cdot \overbrace{[w^v, w^{c_0}, \dots, w^{c_{N-1}}]}^{1+N} = \overbrace{[l^v, l^{c_0}, \dots, l^{c_{N-1}}]}^{1+N}, \quad (1)$$

where  $N$  is the number of the predefined categories,  $l^v$  and  $l^{c_i}$  are the logits of the virtual category and the class  $c_i$  respectively. To calculate the loss value of the extended logits, the one-hot label is modified by providing a positive label ‘1’ for the virtual category. Thanks to the virtual category taking on the responsibility of being the target category, the confusing categories in the PC set can be ignored (indicated by ‘-1’ in Fig. 2), thereby avoiding any possible misleading as we can hardly determine which one is the real ground truth.

**Explanation** How the proposed VC learning works can be interpreted from both the feature space aspect and the mathematical feasibility.

In the feature space, as shown in Fig. 3 left, let the circles indicate the centres of the predefined categories. Pulling the training sample (diamond) to the circles  $a$  or  $b$  is risky since we don’t know which one is the real ground truth. Given the virtual category, the decision boundary can still be optimised with VC learning, as the virtual category providing a safe optimising direction: pushing the training sample away from the circles  $c$  and  $d$  and pulling it closer to the diamond of the virtual category. Although one may suspect that our approach looks similar to the contrastive learning [15,9] in terms of the optimisation target, they differ in several aspects. Firstly, contrastive learning operates before the task-relevant layer (*i.e.*, the classifier). As a result, it only drives the backbone to extract better features but contributes nothing to the task-relevant layer. While our approach acts after the classifier so that the gradient of virtual category

can backpropagate to not only the backbone but also the weight vectors in the classifier. Secondly, the weight vectors of the other categories in the classifier naturally constitute negative samples such that there is no need to maintain a negative sample pool, which has been a worrying bottleneck for contrastive learning. Fig. 3 right presents the T-SNE [26] visualisation of the weight vectors  $w^{c_i}$  in the trained classifier, and the virtual weight  $w^v$  and the feature  $f$  of a sample. Our virtual category is clearly a better training target compared to the pseudo labels, which is in line with the feature space aspect explanation.

To explain our method from the mathematical aspect, we first define the VC loss starting from the cross entropy (CE) loss. We abbreviate the category  $c_i$  by its index  $i$  in the following. Assuming a batch size of 1, the CE loss is:

$$\mathcal{L}_{CE} = -\log\left(\frac{e^{f^\top \cdot w^{i=GT}}}{\sum_{i=0}^N e^{f^\top \cdot w^i}}\right) = \log\left(\sum_{i=0}^N e^{l^i - l^{GT}}\right), \quad (2)$$

where  $f \in \mathbb{R}^{channel \times 1}$  is the input feature vector of the last linear layer (*i.e.*, classifier),  $w^i \in \mathbb{R}^{channel \times 1}$  is the corresponding weight vector of the category  $c_i$  in the last linear layer,  $l^i = f^\top \cdot w^i$  is the logit of the category  $c_i$ ,  $N$  is the number of the predefined categories, and  $GT$  is the index of the ground truth. Here, we ignore the bias in the last linear layer for simplicity. With Eq. (1), the VC loss can be defined as follows according to Eq. (2):

$$\mathcal{L}_{VC} = \log\left(\sum_{i=0, i \notin PC}^{N+1} e^{l^i - l^v}\right), \quad (3)$$

where  $PC$  is the PC set.  $i = N$  indicates the virtual category and  $l^{i=N} = l^v$ .  $i \notin PC$  means the uncertain potential categories in the PC set are ignored.

The intuitive target of minimising CE loss is to get a large logit  $l^{GT}$  of the ground truth and small values for the rest of the categories  $l^{i \neq GT}$ . In a self-supervised scheme, the pseudo classification label  $cls'$  of a unlabelled data  $x_u$  is generated by the model itself or the EMA (Exponential Moving Average) model. As Eq. (2) can be a smooth approximation of the *max* function [32] (explanation without the approximation and more details are in the supplementary document), the CE loss with the pseudo label  $cls'$  can be expressed as:

$$\mathcal{L}_{CE} = \log\left(\sum_{i=0}^N e^{l^i - l^{cls'}}\right) \approx \max_{i \in \{0, \dots, N-1\}} (l^i - l^{cls'}), \quad (4)$$

minimising Eq. (4) is expected to satisfy:

$$(l^{i \neq cls'} - l^{cls'}) \leq (l^{i=cls'} - l^{cls'}) \triangleq 0, \quad i.e., \quad l^{i \neq cls'} \leq l^{cls'}. \quad (5)$$

Since  $cls'$  may not always be correct, when  $cls'$  is not equal to the real ground truth  $GT$ , satisfying Eq. (5) leads to  $l^{GT} \leq l^{cls'}$ , thereby aggravating the issue typically termed confirmation bias.

However, for VC loss, following the derivation of Eq. (4), we get:

$$\mathcal{L}_{VC} = \log\left(\sum_{i=0, i \notin PC}^{N+1} e^{l^i - l^v}\right) \approx \max_{i \in \{0, \dots, N\} \setminus PC} (l^i - l^v). \quad (6)$$

Similar to Eq. (5), minimising Eq. (6) is expecting:

$$(l^{i \neq v \wedge i \notin PC} - l^v) \leq (l^{i=v} - l^v) \triangleq 0, \text{ i.e., } l^{i \neq v \wedge i \notin PC} \leq l^v. \quad (7)$$

Comparing Eq. (7) with Eq. (5) reveals:

1.  $\mathcal{L}_{VC}$  firstly ignores the logits  $l^{i \in PC}$  of the confusing categories in the PC set when satisfying Eq. (7), thereby avoiding misleading the training.
2. Additionally, it provides an alternative upper bound  $l^v = f^\top \cdot w^v$  for all the rest logits  $l^{i \notin PC}$ . As before-mentioned,  $w^v$  is the normalised and scaled feature vector  $\hat{f}$  according to the length (norm of the vector) of the weight vectors  $w^i$ . The directions of  $w^v$  and all the  $w^i$  can represent the information of different categories [29]. The  $l^v = f^\top \cdot w^v$  should be larger since  $w^v$  are obtained by  $\hat{f}$  which is the feature in the teacher of the exactly same object of feature  $f$ . The shared information between  $f$  and  $\hat{f}$  should be the most. Thus,  $l^v$  can be a meaningful upper bound for all the rest logits  $l^{i \notin PC}$ .

### 3.3 Potential Category Set

Two methods are explored in this paper to discover the potential category.

#### a. Temporal stability

It is observed that the pseudo label set  $\mathcal{B}'$  of an image varies at different training iteration steps [41]. As shown in Fig. 4, by comparing  $\mathcal{B}'$  at different training steps, those changed pseudo labels reveal the potential categories. We select  $\mathcal{B}'$  of the current iteration step and  $\mathcal{B}'_{last}$  generated when the model viewed the current image the last time for the comparison. Specifically, for one pseudo label  $b' \in \mathcal{B}' \cup \mathcal{B}'_{last}$ , if there is any pseudo label  $\hat{b}' \in \mathcal{B}' \cup \mathcal{B}'_{last}$  that is close to  $b'$  and  $b'_{cls} = \hat{b}'_{cls}$ , no confusion will be recorded. Otherwise, if there is a nearby  $\hat{b}'$  but  $b'_{cls} \neq \hat{b}'_{cls}$ ,  $PCset = \{b'_{cls}, \hat{b}'_{cls}\}$ . No nearby pseudo label box means that the confusing category is ‘background’, thus  $PCset = \{b'_{cls}, bg\}$ . The IoU is adopted to determine whether two pseudo bounding boxes are close enough or not.

#### b. Cross-model verification

Inspired by the Co-training algorithm [5], comparing the predictions of two conditionally independent models for the same sample can also help to discover the potential category. Two detectors are initialised with different initial parameters. The orders of the training data for these two detectors are also different, ensuring that they will not collapse onto each other. For a pseudo bounding box  $b'_1$  from the detector-1, we use the detector-2 to re-determine the object in the region of  $b'_1$ . The decision of the detector-2 is represented by  $b'_2$ . Similar to the temporal stability verification, if  $b'_{cls_1} = b'_{cls_2}$ , the PC set contains only one category  $b'_{cls_1}$  (or  $b'_{cls_2}$ ). Otherwise,  $PCset = \{b'_{cls_1}, b'_{cls_2}\}$ .



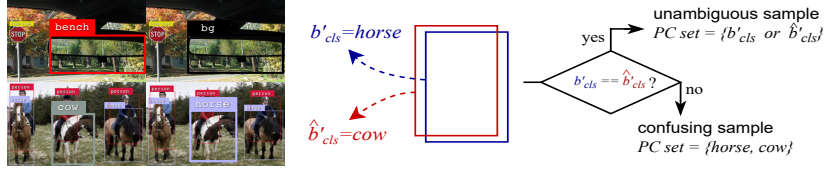


Fig. 4: Left: Examples of the confusing samples found by the temporal stability verification. Right: By comparing the current predictions (blue) with the last predictions (red) of the sample, the PC set should be  $\{horse, cow\}$ .

In summary, for the classification training, If the size of the PC set is not equal to 1, it means that this sample should be a confusing sample. The VC loss will take over the training of this confusing sample.

### 3.4 Localisation Loss

Since the classification confidence score is not qualified to indicate the location quality of pseudo labels, some of the previous works [25,35] disabled the localisation loss of unlabelled data. We find that the creating method for the PC set can also measure the location quality of pseudo labels. When we create the PC set for a pseudo box  $b$ , we evaluate its location shift with the nearby box  $\hat{b}$ . We propose to decouple the horizontal and vertical boundary quality instead of using the IoU as a comprehensive metric to filter out the whole bounding box with low IoU value. The reason is that the IoU value can be affected by one biased boundary, even if the remaining boundaries are good. The horizontal quality flag  $q_{hor}$  is calculated as:

$$q_{hor} = \begin{cases} 1, & \frac{(x_1 - \hat{x}_1)}{w} < t_{loc} \ \& \ \frac{(x_2 - \hat{x}_2)}{w} < t_{loc} \\ 0, & otherwise \end{cases}, \quad (8)$$

where  $x_1, x_2, \hat{x}_1, \hat{x}_2$  are the coordinates of the left and right boundary of the pseudo box  $b$  and the nearby box  $\hat{b}$ ,  $w$  is the width of  $b$ ,  $t_{loc}$  is the threshold for high-quality boundaries.  $q_{ver}$  is calculated in the same way.

The decoupling allows high-quality boundaries to contribute to the localisation training. For example, the regression of the left and right boundary can be trained when the horizontal boundary quality is satisfied, even the top and bottom boundaries are biased. The localisation loss of four Smooth-L1 [14] terms:

$$\mathcal{L}_{reg^*} = q_{hor}\mathcal{L}^x + q_{ver}\mathcal{L}^y + q_{hor}\mathcal{L}^w + q_{ver}\mathcal{L}^h. \quad (9)$$

In summary, the overall loss function is:

$$\mathcal{L} = \mathcal{L}_{CE}(\mathcal{D}_l) + \mathcal{L}_{reg}(\mathcal{D}_l) + \lambda(\mathcal{L}_{VC}(\mathcal{D}_u) + \mathcal{L}_{reg^*}(\mathcal{D}_u)), \quad (10)$$

where  $\lambda$  is the weight for the loss of the unlabelled data. In the implementation, we also use the VC loss function to handle the unambiguous samples by only ignoring the index of the virtual category.

## 4 Experiments

### 4.1 Datasets and Evaluation Protocol

To evaluate the proposed method, we assess it on two well-known object detection benchmark datasets – MS COCO [44] and Pascal VOC [13]. Following the mainstream evaluation setting, we use the subset index provided by Unbiased Teacher [25] to split the *train set* across five different labelled ratios: 0.5%, 1%, 2%, 5% and 10% (each ratio are with five random seeds to get the averaged mAP). We also report the performance on Pascal VOC with *VOC07-trainval* as the labelled subset and *VOC12-trainval* as the unlabelled subset. The performance is evaluated on *VOC07-test*. The evaluation metric of all the experiments reported in this paper is mAP calculated via the COCO evaluation kit [1].

### 4.2 Implementation

The proposed method is implemented in PyTorch framework [28]. Following the mainstream choice of the community, we adopted Faster-RCNN [31] with FPN [23] and ResNet-50 [16] as the object detector. The weight term in Focal Loss [24] is also integrated into our VC loss to address the class imbalance issue.

The training is operated on 8 GPUs with batch size 1/4 per GPU for labelled/unlabelled data. The training iteration is 180k including a warmup of 2k steps with only the labelled subset. The optimiser is SGD with a constant  $lr = 0.01$ . The default scaling factor for the virtual weight is a constant 3.5. The location quality threshold  $t_{loc}$  in  $\mathcal{L}_{reg^*}$  is 0.05. The default PC set creating method is the temporal stability verification. More details are introduced in the supplementary. The code can be found at: <https://github.com/GeoffreyChen777/VC>.

### 4.3 Performance

**MS COCO** We first evaluate our method on MS COCO with five label ratios. The 5 seeds averaged results are reported in Tab. 1. ‘Ours’ is our method with exactly the same settings of the solid baseline Unbiased Teacher for the sake of fairness. ‘Ours\*’ is obtained with some training tricks adopted by Soft Teacher. Results with † are obtained from the available official code. The significant improvements can be summarised as follows:

- 1) Compared with the supervised baseline, the mAP increases dramatically after training with the unlabelled data via our method.
- 2) Our method outperforms other state-of-the-art semi-supervised detectors on all the label ratios by a significant margin. The mAP of our method at a small label ratio is close to or even exceeds the mAP of some methods at a large ratio.

**Pascal VOC** We also evaluate our method with VOC07 as the labelled subset and VOC12 and COCO\* as the unlabelled subsets. We collect the images that contain objects in VOC predefined categories from MS COCO to build a subset *COCO\**. The results are presented in Tab. 2. Since the source codes of some methods are unavailable, the evaluation styles they used is unclear.

Table 1: The performance on MS COCO with different label ratios. Results with † are obtained from the available official code. Ours is the results of the exactly same settings of Unbiased Teacher for the sake of fairness. Ours\* is the results obtained by using some training tricks in Soft Teacher.

COCO label ratio	0.5%	1%	2%	5%	10%
Supervised	6.83	9.05	12.70	18.47	23.86
CSD [19]	7.41	10.51	13.93	18.63	22.46
STAC [35]	9.78	13.97	18.25	24.38	28.64
Instant Teaching [43]	-	18.05	22.45	26.75	30.40
Interactive [41]	-	18.88	22.43	26.37	30.53
Humble Teacher [36]	-	16.96	21.72	27.70	31.61
Combating Noise [39]	-	18.41	24.00	28.96	32.43
Soft Teacher [40]	15.04†	20.46	25.93†	30.74	34.04
Unbiased Teacher [25]	16.94	20.75	24.30	28.27	31.50
Ours	18.12	21.61	25.84	30.31	33.45
Ours*	<b>19.46</b>	<b>23.86</b>	<b>27.70</b>	<b>32.05</b>	<b>34.82</b>

Table 2: Experiment results on VOC.  $\mathcal{D}^l$  and  $\mathcal{D}^u$  are the labelled and unlabelled subset choices. COCO\* consists of the images from COCO that contains objects in VOC categories. Numbers in () are obtained with the VOC-style AP.

Method	$\frac{\mathcal{D}^l}{\mathcal{D}^u}$	VOC12	VOC07	VOC12 + COCO*
STAC [35]		44.64		46.01
Instant Teaching [43]		50.00		50.80
Interactive [41]		46.23		49.59
Humble Teacher [36]		53.04		54.41
Combating Noise [39]		49.30		50.20
Unbiased Teacher [25]		48.69		50.34
Ours		<b>50.40 (55.74)</b>		<b>51.44 (56.70)</b>

Usually, the results based on the VOC-style AP are higher. Thus, we evaluate our method with COCO-style mAP and VOC-style AP for the sake of fairness. Our method presents the best performance on these two unlabelled data subsets. Since VOC07 consists of more than 5K labelled images, and it is a relatively easy dataset, Tab. 2 indicates that our method can effectively further improve the performance, even if there is already sufficient labelled data. The similar conclusion can be drawn when we use the unlabelled COCO-additional and fully-labelled COCO-standard to validate our method (analysed in the supplementary).

#### 4.4 Analysis and Ablation Study

In this section, we choose 1% data of MS COCO as the labelled subset to analyse and validate our method in detail. All the experiments in this section are performed under the exactly same setting of the baseline model Unbiased Teacher except from the batch size. We adopt a smaller batch size to shorten the training time of each ablation study, therefore resulting in slightly decreased mAPs of all experiments compared to Tab. 1.

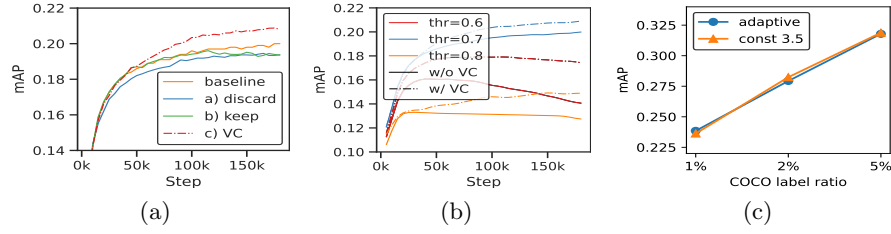


Fig. 5: a) Experiments of different strategies for dealing with confusing samples. As the pseudo labels of the valuable confusing samples are highly unreliable, it is not optimal to either discard or keep them. Our VC learning(dotdashed line) satisfies both demands, thereby resulting in a significant improvement. b) Experiments of different thresholds of the confidence score filtering w/ or w/o our VC learning. c) Comparison of the adaptive and the constant scaling factor.

**Analysis of Virtual Category** Here, we adopt the temporal stability verification to create the PC set for confusing samples. To analyse the effectiveness of the virtual category, we respectively report the mAP of the model under three policies: a) discarding all confusing samples (discard), b) retaining all potential labels for them (keep) and c) assigning our virtual category to replace the potential categories (VC). The baseline model is trained with vanilla pseudo labels (baseline) without the potential category discovery. As shown in Tab. 3, both discarding and retaining policies decrease the mAP. By analysing the mAP during the whole training presented in Fig. 5a, we noticed that rejecting the confusing samples (blue line) results in a low mAP at the very beginning of the training. The reason is that using this policy discards some confusing samples with correct pseudo labels that the model desperately needs. Then, as shown by the green line in Fig. 5a, training with all PCs gives a small performance boost at the early stage of training because more underfitted samples are introduced to the model, but ends up with a low mAP. We believe this is due to the confirmation bias issue caused by incorrect pseudo labels that gradually hurts the performance. Our approach effectively resolves this conflict by providing a virtual category for the confusing sample. The dotdashed line in Fig. 5a demonstrates that these confusing samples consistently benefit the model. The mAP sees a rise of 0.01 with our VC learning. The model with our VC learning exceeds the baseline early in the training and continues to lead until the end of training.

In addition, as can be seen from Fig. 5b, we evaluate our VC learning with different thresholds (indicated by three colours) of the confidence score filtering adopted by our baseline model Unbiased Teacher. The confusing samples always exist, no matter whether the filtering mechanism is strict or not. The model with VC learning (dot-dashed lines) outperforms the baseline (solid lines) on three thresholds. Notably, the slump in the mAP disappears when  $thr = 0.6$ , meaning that the confirmation bias has been effectively alleviated.

Table 3: Validation mAP with different strategies for dealing with the confusing samples.

Strategies	mAP
baseline	20.00
a) discard	19.37
b) keep	19.36
c) VC	<b>20.81</b>

Table 4: Ablation study on VC loss and modified localisation loss Reg\* Loss. The creating method for the potential category set in VC learning is temporal stability verification.

VC loss	Reg* Loss	mAP
✓		20.00
✓	✓	20.81
		<b>20.94</b>

Table 5: Ablation study: a) different image augmentation for virtual weights generation, b) different methods for creating the PC. We also report the performance with only co-training techniques. c) SC. is the scaling factor of the virtual weight. d)  $t_{loc}$  is the threshold of the location quality.

(a)		(b)		(c)		(d)	
Aug.	mAP	Methods	mAP	SC.	mAP	$t_{loc}$	mAP
none	20.80	baseline	20.00	2.5	20.63	0.03	20.92
flipping	<b>20.81</b>	Temporal	20.81	3.5	20.81	0.05	20.94
strong aug.	20.70	Cross	<b>20.96</b>	4.5	20.65	0.1	20.31
		co-training only	20.53				

The constant and adaptive scaling factors of the virtual weight are compared in Fig. 5c. The constant 3.5 achieves a mAP which is comparable to the adaptive factor. Interestingly, we discover that  $\frac{1}{ada.factor} \approx 3.5$ .

**Ablation Study** The overall ablation study is reported in Tab. 4. The model with VC learning and Reg\* Loss performs favourably against the baseline model.

*a. Creating methods of potential category set:* Section 3.3 explored two methods to create the PC set. We validate them and report the results in Tab. 5b. The cross-model verification achieves the best performance. The reason is that The cross-model verification is similar to the co-training technique which uses two independent models to provide pseudo labels for each other. it slightly alleviates the confirmation bias issue, thus resulting in additional improvement. As shown in Tab. 5b, the co-training can improve the mAP by 0.53 individually. In fairness to other methods without co-training, we use the temporal stability verification in all other experiments, although the cross-model verification performs better.

*b. Augmentation of the virtual weights:* We choose the feature vector from the teacher model to produce the virtual weight. It is natural to validate different augmentations for the input image of the teacher model to generate various virtual weights. We explore three different settings: no augmentation, horizontal flipping, and strong augmentation. The results are reported in Tab. 5a. No performance gap can be observed between no augmentation and only horizontal flip. Training with the virtual weight generated by the strong augmentation slightly

degrades the mAP. The possible reason is that the strong augmentation, especially the cutout, significantly perturbs the input image. Thus, in the feature space, the direction of the virtual weight is far away from the weight vector of the GT category.

*c. Hyperparameters:* The scaling factor ensures the norm of the virtual weight is in the similar range of the rest weight vectors. In addition to the adaptive version, we investigate the different constant scaling factor values here. The results (w/o  $\mathcal{L}_{reg^*}$ ) are reported in Tab. 5c.  $SC. = 3.5$  achieves the best mAP. The ablation study of the location quality threshold  $t_{loc}$  in the  $\mathcal{L}_{reg^*}$  are shown in Tab. 5d. A higher threshold will retain more unstable boundaries, thereby leading to worse performance.

## 5 Discussion

In this section, we discuss the limitations and potential of VC learning. The VC loss takes over the optimisation of the confusing samples. Thus, higher quality of the PC set yields more improvement. The proposed creating methods are straightforward and easy to implement. Although a significant increment in mAP was observed already, there is still room to further boost the mAP with a better PC set. The mAP would increase from 20.81 to 24.10 if we were to use the ground truth to create the PC set under the setting of 1% labelled COCO. Generating a better PC set can be worth exploring in the future. Moreover, by comparing the mAP gains of 10% and other small label ratios, the improvement of our method is slightly lower, but still remains first. This phenomenon is expected and reasonable. On the one hand, more labelled data means a better baseline detector. Thus, the room between the baseline and the fully-supervised upper bound is smaller. On the other hand, less unlabelled data and a better detector indicate the confusing samples are fewer. At the extreme, with 100% labelled data, our VC learning will be applied to no sample, thereby resulting in no improvement. Notably, this scenario is not the topic of this paper. We focus on the situation with very few labelled data.

## 6 Conclusion

This paper proposed VC learning, which exploits the confusing underfitted unlabelled data. We provide a virtual category label to a sample if its pseudo label is unreliable. It allows the model to be safely trained with confusing data for further improvement to achieve state-of-the-art performance. To the best of our knowledge, VC learning is the first method to positively utilise confusing samples in SSOD rather than discarding them. It can also serve as a stepping stone to future work for the community of other semi-supervised learning tasks.

**Acknowledgments** We thank China Scholarship Council for the funding. We also thank anonymous reviewers for their comments. Corresponding author: Jungong Han.

## References

1. Coco evaluation kit by facebook detectron2. [https://github.com/facebookresearch/detectron2/blob/main/detectron2/evaluation/coco\\_evaluation.py](https://github.com/facebookresearch/detectron2/blob/main/detectron2/evaluation/coco_evaluation.py)
2. Arazo, E., Ortego, D., Albert, P., Connor, N.E.O., McGuinness, K.: Pseudo-labeling and confirmation bias in deep semi-supervised learning. In: IJCNN (2020)
3. Bachman, P., Alsharif, O., Precup, D.: Learning with pseudo-ensembles. In: NeurIPS (2014)
4. Berthelot, D., Carlini, N., Goodfellow, I.J., Papernot, N., Oliver, A., Raffel, C.: Mixmatch - a holistic approach to semi-supervised learning. In: NeurIPS (2019)
5. Blum, A., Mitchell, T.M.: Combining labeled and unlabeled data with co-training. In: COLT (1998)
6. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4 - optimal speed and accuracy of object detection. arXiv (2020)
7. Cai, Z., Vasconcelos, N.: Cascade r-cnn - delving into high quality object detection. In: CVPR (2018)
8. Chen, J., Shah, V., Kyrillidis, A.: Negative sampling in semi-supervised learning. In: International Conference on Machine Learning (ICML). pp. 1704–1714 (2020)
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: ICML (2020)
10. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. arXiv (2017)
11. Dong-Hyun, L.: Pseudo-label : The simple and efficient semi-supervised learning method for deep neural networks. In: ICMLw (2013)
12. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet - keypoint triplets for object detection. In: ICCV (2019)
13. Everingham, M., Eslami, S.M.A., Gool, L.V., Williams, C.K.I., Winn, J.M., Zisserman, A.: The pascal visual object classes challenge - a retrospective. IJCVR (2015)
14. Girshick, R.B.: Fast r-cnn. In: IEEE International Conference on Computer Vision (ICCV) (2015)
15. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.B.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
17. Kim, Y., Yim, J., Yun, J., Kim, J.: Nlnl: Negative learning for noisy labels. In: IEEE International Conference on Computer Vision (ICCV). pp. 101–110 (2019)
18. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Citeseer (2009)
19. Kwak, J.J.S.L.J.K.N.: Consistency-based semi-supervised learning for object detection. In: NeurIPS (2019)
20. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: ICLR (2017)
21. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: ECCV. pp. 734–750 (2018)
22. Li, J., Xiong, C., Hoi, S.: Comatch - semi-supervised learning with contrastive graph regularization. ICCV (2020)
23. Lin, T.Y., Dollár, P., Girshick, R.B., He, K., Hariharan, B., Belongie, S.J.: Feature pyramid networks for object detection. In: CVPR (2017)

24. Lin, T.Y., Goyal, P., Girshick, R.B., He, K., Dollár, P.: Focal loss for dense object detection. TPAMI (2020)
25. 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. In: ICLR (2021)
26. der Maaten, V., Laurens, Hinton, G.: Visualizing data using t-sne. Journal of Machine Learning Research (2008)
27. Miyato, T., ichi Maeda, S., Koyama, M., Ishii, S.: Virtual adversarial training - a regularization method for supervised and semi-supervised learning. TPAMI (2019)
28. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E.Z., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch - an imperative style, high-performance deep learning library. In: NeurIPS (2019)
29. Qi, H., Brown, M., Lowe, D.G.: Low-shot learning with imprinted weights. In: CVPR (2018)
30. Redmon, J., Divvala, S.K., Girshick, R.B., Farhadi, A.: You only look once - unified, real-time object detection. In: CVPR (2016)
31. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster r-cnn - towards real-time object detection with region proposal networks. TPAMI (2017)
32. S, B., SP, B., L, V.: Convex optimization. Cambridge university press (2004), 1st ed., page 72
33. Shrivastava, A., Gupta, A., Girshick, R.B.: Training region-based object detectors with online hard example mining. In: CVPR (2016)
34. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch - simplifying semi-supervised learning with consistency and confidence. In: NeurIPS (2020)
35. Sohn, K., Zhang, Z., Li, C.L., Zhang, H., Lee, C.Y., Pfister, T.: A simple semi-supervised learning framework for object detection. arXiv (2020)
36. Tang, Y., Chen, W., Luo, Y., Zhang, Y.: Humble teachers teach better students for semi-supervised object detection. In: Computer Vision and Pattern Recognition (CVPR) (2021)
37. Tian, Z., Shen, C., Chen, H., He, T.: Fcos - fully convolutional one-stage object detection. In: ICCV (2019)
38. Wang, X., Hua, Y., Kodirov, E., Clifton, D.A., Robertson, N.M.: Proselflc - progressive self label correction for training robust deep neural networks. In: CVPR (2021)
39. Wang, Z., Li, Y., Guo, Y., Wang, S.: Combating noise: Semi-supervised learning by region uncertainty quantification. In: NeurIPS (2021)
40. 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. ICCV (2021)
41. Yang, Q., Wei, X., Wang, B., Hua, X.S., Zhang, L.: Interactive self-training with mean teachers for semi-supervised object detection. In: CVPR (2021)
42. Zagoruyko, N.C.F.M.G.S.N.U.A.K.S.: End-to-end object detection with transformers. In: ECCV (2020)
43. Zhou, Q., Yu, C., Wang, Z., Qian, Q., Li, H.: Instant-teaching - an end-to-end semi-supervised object detection framework. In: CVPR (2021)
44. Zitnick, T.Y.L.M.M.S.J.B.J.H.P.P.D.R.P.D.C.L.: Microsoft coco - common objects in context. In: ECCV (2014)