



This is a repository copy of *Virtual category learning: a semi-supervised learning method for dense prediction with extremely limited labels*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/209262/>

Version: Accepted Version

Article:

Chen, C., Han, J. and Debattista, K. (2024) Virtual category learning: a semi-supervised learning method for dense prediction with extremely limited labels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. ISSN 0162-8828

<https://doi.org/10.1109/TPAMI.2024.3367416>

© 2024 The Author(s). Except as otherwise noted, this author-accepted version of a journal article published in *IEEE Transactions on Pattern Analysis and Machine Intelligence* is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Virtual Category Learning: A Semi-Supervised Learning Method for Dense Prediction with Extremely Limited Labels

Changrui Chen, Jungong Han*, Kurt Debattista

Abstract—Due to the costliness of labelled data in real-world applications, semi-supervised learning, underpinned by pseudo labelling, is an appealing solution. However, handling confusing samples is nontrivial: discarding valuable confusing samples would compromise the model generalisation while using them for training would exacerbate the issue of confirmation bias caused by the resulting 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 in such a way that it can safely contribute to the model optimisation even without a concrete label. This provides an upper bound for inter-class information sharing capacity, which eventually leads to a better embedding space. Extensive experiments on two mainstream dense prediction tasks — semantic segmentation and object detection, demonstrate that the proposed VC learning significantly surpasses the state-of-the-art, especially when only very few labels are available. Our intriguing findings highlight the usage of VC learning in dense vision tasks.

Index Terms—Semi-supervised learning, Semantic Segmentation, Object Detection.

1 INTRODUCTION

DEEP Learning solutions are significantly disadvantaged by the expensive labelling cost of large-scale datasets, especially on dense prediction tasks. Annotating the data for tasks such as object detection and semantic segmentation takes substantially longer when compared to non-dense applications such as image classification. Though crowd-sourcing platforms facilitate data labelling for many common application scenarios, such as autonomous driving, data labelling in scientific applications usually requires expert labelling, which is not readily available. Semi-supervised learning, which makes use of limited labelled data in combination with large amounts of unlabelled data for training, has shown great potential to reduce the reliance on large amounts of data labelling [1], [2], [3], [4]. Diminishing the performance gap between fully- and semi-supervised methods enables the introduction of deep-learning models to more application topics. In existing semi-supervised learning frameworks, a challenging problem is: *how to best utilise the unlabelled data*.

Pseudo labelling (PL) [5] has recently emerged as a solution to the above problem [1], [3] and achieved state-of-the-art performance. Here, unlabelled data are automatically annotated by the model itself [6] (or via an exponential moving average version [1]) and then fed back to re-optimize the model. Despite their preliminary success, existing PL-based semi-supervised methods are not good at or even incapable of dealing with the extremely-scarce label setting.

- C. Chen (changrui.chen@warwick.ac.uk) and K. Debattista are with WMG, University of Warwick, UK.
- J. Han is with the University of Sheffield, UK.

Manuscript received xxx, xxx; revised xxx, xxx and xxx, xxx; accepted xxx, xxx. This research was supported by the funds of China Scholarship Council under Grant No. 202006330013. (Corresponding author: Jungong Han).

It motivates us to study the limited supervised case, which is imperative in real-life application scenarios.

Typically, due to the limited diversity of training samples in a very small available set of labelled data, the non-optimal decision boundary usually leads to an indecisive decision on some unseen confusing samples when inferring their pseudo labels. In PL, two strategies are usually adopted to deal with confusing samples: a) discarding all of them using a strict filtering mechanism [6], or b) retaining them with all potential labels [7]. However, neither of these two options is optimal, especially when the labelled training data are very limited. The value of confusing hard samples is clear to see since hard example mining [8] has successfully proven its effectiveness in fully-supervised learning. If all confusing samples are rejected by a strict filtering mechanism, 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 is ineffective due to the involvement of too many incorrect pseudo labels. Arbitrarily optimising semi-supervised detectors with these noisy labels results in confirmation bias issue [9]. To demonstrate these points, we show in Fig. 1a the mean Average Precision (mAP) of a semi-supervised object detector with different strategies on 1% labelled MS COCO [10]. Noticeable performance degradation can be observed when either choosing one stricter filtering mechanism (orange line) or adding an additional one (green line) to reject confusing samples. Likewise, simply keeping all confusing samples (yellow line) pseudo labels also ends up with a decreased mAP since the unreliable pseudo labels aggravate the confirmation bias issue and can result in training collapse. A similar phenomenon can be seen in semantic segmentation (Fig. 1b) as well.

In light of the above, efforts have been dedicated to exploring how to correct the biased pseudo labels to utilise




Fig. 1: a): The mAP of a semi-supervised detector [4] with a preset confidence score filtering on 1% labelled MS COCO [10]. The mAP sees a decrease with all strategies for dealing with confusing samples (e.g., the bear-like dog at the right) except for the VC learning. The additional filtering mechanism (add.) is the temporal stability verification proposed in this paper. Stricter filtering (stricter.) raises the threshold in the score filtering from 0.7 to 0.8. b): The mIoU of a semi-supervised semantic segmentor on 1/128 labelled Pascal VOC.

the confusing samples efficiently. Existing methods [4] initially investigate relatively straightforward tasks such as classification on CIFAR [11]. However, promising progress has not yet been made for a complex dense prediction task, such as object detection with extremely small amounts of labelled data. Then, a question arises: *what if we do not discard 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) learning based on the observation that there is an implicit and safe optimisation direction in PL models for confusing data.

Fig. 2 provides an example, in which the bear-like dog is a typical confusing sample due to its appearance. For a classification model, the arrow pointing towards the ‘dog’ is the best optimising direction, which results in the smallest testing error value. However, if the doubtful pseudo label is ‘bear’, the incorrect optimising direction would lead to worse performance. We discover that building a Potential Category (PC) set consisting of the possible categories of a confusing sample x , compared to determining the exact correct label, is much easier. The remaining task is to find a good optimising direction (labelled VC in Fig. 2) for the sample x without the guidance of the categories in the potential category set. Therefore, instead of selecting the correct one from the potential category set, which is usually challenging, we compromise by proposing a VC label to *take the place of all unreliable labels in the potential category set*. A new learning scheme, namely VC learning, allows the model to be optimised with the VC label. By ignoring the categories in the potential category, it will disable the gradient of the corresponding output logits, thus avoiding any wrong optimising direction that would mislead the model. Most importantly, the proposed VC specifies a reasonable upper bound for the inter-class sharing information capacity. Hence, the decision boundary can consistently benefit from the confusing data without suffering from the confirmation bias issue. With regards to the potential category set, we come up




Fig. 2: Illustration of the basic idea of the virtual category in the manifold of the optimisation space. The peak indicates a high testing error value.

with multiple methods to build it. As can be seen in Fig. 1, the performance of the model armed with the proposed VC learning (dot-dashed line) sees a significant increase due to the effective use of the confusing samples.

The proposed VC learning is applied to a semi-supervised detector and a semi-supervised segmentor, both with extremely limited labelled training data, to evaluate its effectiveness and generalisation capability on dense prediction tasks. In the object detection on MS COCO, VC learning achieves 19.46 mAP with only 586 labelled images; this even outperforms some recently published semi-supervised detectors [6], [12] with 1000+ labelled images. For semantic segmentation, we developed a powerful and straightforward pseudo-labelling framework. VC learning further boosts the baseline framework to achieve a mIoU of 55.37 on Pascal VOC with only 82 labelled images. It surpasses state-of-the-art methods by a large margin. The contributions of this article are summarised as follows:

- We take advantage of *confusing samples* with pseudo labels in a semi-supervised manner through VC learning. Our VC learning alleviates the confirmation bias issue caused by confusing samples. It works extremely well when only very limited labelled training data are available.
- We theoretically exhibit the feasibility of using VC learning for semi-supervised learning. The findings highlight the need to rethink the usage of confusing samples in semi-supervised tasks.
- Compared to our previous ECCV oral paper [13], we extend VC learning to semi-supervised semantic segmentation. More methods for the potential category set creation and an additional loss function form are introduced, *i.e.*, the mean squared error for VC learning. On top of it, we also propose a new module to generate the virtual weight in VC learning. The experiments show that VC learning can be well applied in semantic segmentation.
- We incorporate VC learning into a unified pseudo-labelling framework for semi-supervised learning, which can deal with multiple dense prediction tasks, including semantic segmentation and object detection. The proposed framework surpasses state-of-the-art methods by a significant margin on ALL tasks, which verifies the generalisation of VC learn-

ing.

2 RELATED WORKS

In this section, a literature review on the semi-supervised learning and the downstream tasks — segmentation and detection — is conducted.

2.1 Semi-supervised Learning

In the last few years, numerous deep backbones and modules training under a fully-supervised scheme have been proposed. VGGNet [14] adopted 3×3 convolution layer as the main model component. GoogLeNet [15] proposed the network-in-network structure for the first time, which allows scaling the width and depth of a convolution neural network (CNN). ResNet [16], proposed the residual block, which made it possible to optimise a very deep neural network. Recently, inspired by the attention module in neural language processing [17], the vision transformer [18] appealed to many researchers. Fully-supervised training is very close to the ability of human beings in many vision tasks [19]. However, in other research and application areas, it is usually difficult to build large-scale labelled datasets such as ImageNet [20] and MS COCO [10] to satisfy the training of fully supervised models. Semi-supervised learning, such as MeanTeacher [21] and FixMatch [3] etc., tackles this issue.

Semi-supervised learning is a training scheme that uses only a small amount of labelled data and a large amount of unlabelled data to train a model. It can be grouped into three main paradigms: a) generative models, b) graph-based methods, and c) pseudo-labelling models. Several unsupervised generative models were extended to solve the semi-supervised problems. For example, the stacked semi-supervised generative model proposed by Kingma *et al.* [22] appended a generative classifier to the latent representation produced by the encoder to enable variational autoencoder to tackle semi-supervised classification. The feature representation would benefit from the reconstruction proxy task of the auto-encoder with the unlabelled data. Generative Adversarial Network (GAN) has also been considered as semi-supervised learning methods [23]. By assigning a ‘fake’ class to all generated images, unlabelled images in the dataset can be labelled as ‘non-fake’ to train the classifier [23]. In addition to the generative model, some graph-based methods introduced data relationships into semi-supervised training [4], [24]. The intuitive motivation is that adjacent nodes in an embedding graph should have similar representations. Recently proposed semi-supervised methods mainly focus on the teacher-student framework via pseudo labelling [2], [3]. The overall idea of a teacher-student framework is to let the predictions of the teacher model be the pseudo labels for optimising the student model. It requires models to produce consistent outputs when the inputs are perturbed. Image augmentations, such as flipping, Cutout [25], or Gaussian Blurring, are usually applied to perturb input images. Some solutions take advantage of adversarial learning and proposed learnable adversarial augmentations [26]. The form of teacher model is in a variety of styles. An exponential moving averaged (EMA)

version [21] or even the student itself [3] was investigated to play the role of the teacher. Many different entities of consistency regularisation have been explored. For example, Jeong *et al.* [27] tried to minimise the discrepancy between the latent representations of perturbed inputs. Yang *et al.* [7] proposed to use the temporal ensembling predictions as the teacher predictions. FixMatch [3] adopted a weak and a strong augmentation to obtain the predictions of the teacher and student model, respectively. The teacher-student framework, which is used as the baseline in this paper, has proven to be successful in several downstream tasks. Despite their successes, the performances of such systems are far from satisfactory in the real scenario, where an extremely low label ratio, say below 1%, is provided.

2.2 Semantic Segmentation

Semantic segmentation can be seen as a dense classification task at the pixel level. Most recent segmentation models are inspired by FCN [28]. The performance of segmentation models is sensitive to the output resolution. Thus, some following works, such as U-Net [29], proposed an encoder-decoder framework to increase the output resolution without compromising efficiency. The receptive field is also crucial to segmentation accuracy. Atrous convolution used by the Deeplab series [30], [31], [32] decently enlarged the receptive field without the aid of stacking large convolutional kernels. The self-attention mechanism enables the segmentor to build long-range connections across the entire images, further exploiting contextual information [33]. Compared to the image-level labels, the cost required by such pixel-level dense labelling makes semi-supervised learning even more crucial. Consistency regularisation is also widely adopted in semi-supervised semantic segmentation. Ouali *et al.* [34] proposed to align the output of different decoders or models. The discriminator, which is usually used in adversarial learning, was considered to minimise the distribution distance between the predictions of unlabelled data and the ground truth of labelled data. The teacher-student model has been well introduced to self-supervised segmentation [35]. However, we discover that its training is very unstable. Thus, this article first investigates and solves this problem to build a strong and stable baseline model and then evaluates the proposed VC learning on it.

2.3 Object Detection

Object Detection, which finds significant applications in downstream tasks, aims to distinguish foreground objects in images or videos and identify them. Object detectors so far can generally be divided into three types: 1) two-stage detectors [36], [37], represented by Faster RCNN [36]; 2) one-stage detectors [38], [39], [40], [41], such as the YOLO series [38], [39]; and 3) point-based detectors [42], [43], [44], such as Center Net [42]. The main difference between two-stage and one-stage detectors lies in whether an additional module is used to generate candidate region proposals for classification and localisation. 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 detector to explore VC learning in semi-supervised object detection(SSOD).

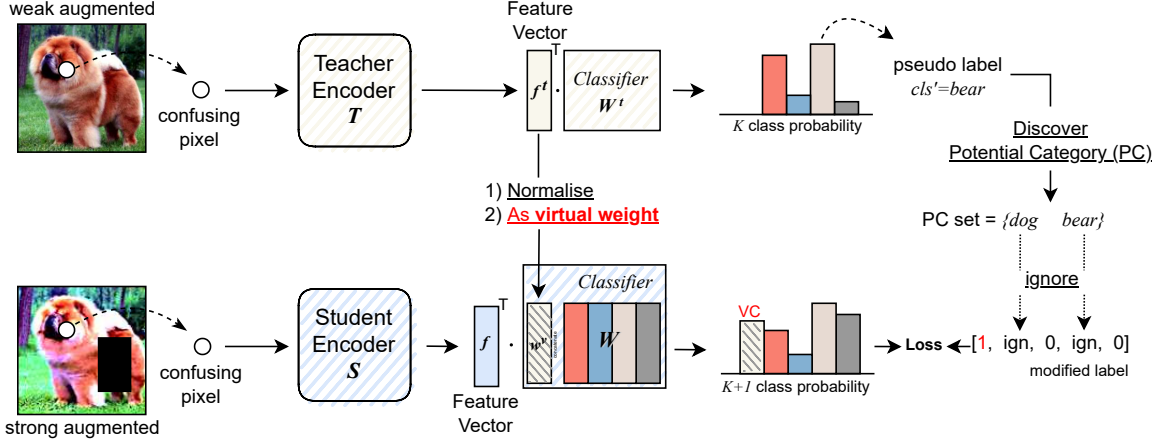


Fig. 3: The pipeline of the proposed VC learning when dealing with a confusing sample in semi-supervised one-pixel classification. T is the teacher model. S represents the student model. When training the student classifier with a confusing sample, the weight matrix W^s of the student classifier is extended by a virtual weight w^v , which is transformed from the corresponding teacher feature vector \hat{f} .

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 algorithms followed pseudo-labelling methods. For instance, CSD-SSD [27] applied consistency regularisation on the predicted classification probability vectors and regression vectors of the input image and its mirror version when dealing with unlabelled images. Several self-supervised detectors [1], [6], [7], [12], which provide teacher-produced pseudo-labels for student detectors, have emerged recently.

Although the teacher-student pseudo labelling technique shows good potential on both tasks, it is still struggling with the confirmation bias issue when training with confusing samples. This paper proactively utilises these confusing samples and alleviates the confirmation bias issue via the novel VC learning method.

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 the semi-supervised problem, two data subsets \mathcal{D}^l and \mathcal{D}^u are given for model optimisation, where $\mathcal{D}^l = \{(x_n^l, y_n^l) |_{n=0}^{N^l}\}$ is the subset with ground truth label y^l available, $\mathcal{D}^u = \{x_n^u |_{n=0}^{N^u}\}$ is the unlabelled subset. N^l and N^u are the numbers of labelled and unlabelled data, respectively. This paper mainly investigates the **limited-supervised learning** problem, *i.e.*, $N^u \gg N^l$, which can be considered as a sort of challenging subproblem of semi-supervised learning. For segmentation, the pixel-level class index $y \in \mathbb{N}_+^{H \times W}$ is the label. $y = [a_1, b_1, a_2, b_2, cls]$ is the label of the object detection task, where the first four 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 teacher-student framework to generate the pseudo label y' of x^u for re-optimising. As shown in Fig. 3, two encoders T and S , which share the same architecture, are introduced. The parameters of the teacher encoder T are updated by the parameters of the student S with a momentum parameter.

3.2 Virtual Category Learning

For ease of understanding, we abstract our VC learning framework for dense vision tasks as a general one-pixel classification task. It can be easily extended to a multi-pixel framework for semantic segmentation and an instance-level framework for object detection as shown in Fig. 4. In Fig. 3, two feature encoders T and S first embed the confusing pixel from the weakly and strongly augmented images into the feature space. The linear classifier parameterised by W^t in the teacher branch produces the categorical probability of the input data by performing a matrix multiplication of the feature vector $f^{t\top}$ and the weight matrix W^t . The bias parameter is ignored here for simplicity. Typically, the category with the highest probability, *e.g.*, $y' = bear$ here, will then be used as the pseudo label for the output of the student branch. However, an incorrect pseudo label 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 $y' = bear$ is obtained, a potential category discovery operation is performed to construct a set $\{dog, bear\}$ for this training sample. We find that the potential category discovery is relatively feasible compared to designing a mapping function $P(y = GT | y'; f)$ to correct wrong pseudo labels, especially when the labelled subset \mathcal{D}^l is much smaller than \mathcal{D}^u . The discovery method will be introduced in the following Section 3.5.

If $y' = bear$ is the only potential category we can find, this pixel is regarded as an unambiguous sample. If the potential category set contains more than one category, it

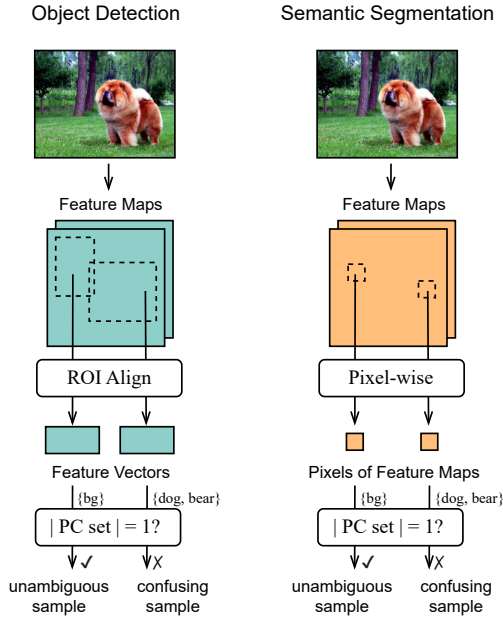


Fig. 4: Explanation of data sample preparation for different tasks. In object detection, we use ROI Align [45] to extract the feature vectors of each region of interest as the data samples. While VC learning is operated in a pixel-wise manner in semantic segmentation.

means that this sample is exactly a confusing sample to the model. In Fig. 3, a pixel (white circle) of the bear-like dog is an example confusing sample with $\{dog, bear\}$. To encourage the confusing pixel to consistently contribute to the optimisation of the student model S rather than arbitrarily discarding it, the weight matrix W in the student classifier is extended by a ‘personal’ weight vector w^v named *virtual weight*. The ingredient of the virtual weight w^v is the feature vector f^t , which is the feature of this confusing sample in the teacher model. Notably, the so-called ‘personal’ virtual weights for different confusing samples are various. With the extended weight matrix, the size of the student classifier output (*i.e.*, logits) is therefore increased by 1:

$$f^T \cdot [\overbrace{w^v, w^0, \dots, w^{K-1}}^{1+K}] = [\overbrace{l^v, l^0, \dots, l^{K-1}}^{1+K}], \quad (1)$$

where K is the number of the predefined categories, l^v and l^i are the logits of the virtual category and the predefined class i respectively.

Optimisation Objective To calculate the loss value of the extended logits, the pseudo label is modified by providing a positive label ‘1’ for the virtual category. The training target form is: $[1, 0, \dots, ign, \dots, ign, 0]$, where ‘ign’ means we ignore that class. Thanks to the virtual category taking on the responsibility of being the target category, the confusing labels in the potential category set can be ignored, thereby avoiding any potential misleading, as it is hard to determine which one is the real ground truth. Thus, the optimisation objective is to get a large logit value for the virtual category and small logit values for the rest of the categories except for those in the potential category set.




Fig. 5: Explanation of VC learning 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.

3.3 Explanation

This section describes how the proposed VC learning can be interpreted both from the aspect of the feature space and from the aspect of mathematical feasibility.

3.3.1 Feature Space

In the feature space, as shown in Fig. 5, 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 VC, the decision boundary can still be optimised with VC learning, as it provides 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 contrastive learning [46], [47] in terms of the optimisation objective, they differ in several aspects. First, contrastive learning operates before the task-relevant layer (*i.e.*, the classifier). As a result, it only drives the backbone encoder to extract better features but does not contribute anything to the task-relevant layer. While our approach acts after the classifier so that the gradient of the virtual category can backpropagate to not only the backbone but also the weight vectors in the classifier. Second, 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.

3.3.2 Mathematical Derivation

To explain our method from the mathematical perspective, we define the loss function of VC learning starting from cross entropy (CE) loss. Assuming a batch size of 1, the CE loss is:

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

where $f \in \mathbb{R}^{channel \times 1}$ is the input feature vector of the last linear layer (*i.e.*, the classifier), $w^i \in \mathbb{R}^{channel \times 1}$ is the corresponding weight vector of the category i in the last linear layer, $l^i = f^T \cdot w^i$ is the logit of the category i , K is the number of the predefined categories, and GT is the index of the ground truth.

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 teacher model infers the pseudo classification label y' of an unlabelled data sample x^u . As Eq. (2) can be a smooth approximation of the max function [48], the CE loss with the pseudo label y' can be expressed as:

$$\mathcal{L}_{CE} = \log\left(\sum_{i=0}^K e^{l^i - l^{y'}}\right) \approx \max_{i \in \{0, \dots, K-1\}} (l^i - l^{y'}). \quad (3)$$

Minimising Eq. (3) is expected to satisfy:

$$(l^{i \neq y'} - l^{y'}) \leq (l^{i=y'} - l^{y'}) \triangleq 0, \quad i.e., l^{i \neq y'} \leq l^{y'}. \quad (4)$$

Since y' may not always be correct, when y' is not equal to the real ground truth GT , satisfying Eq. (4) leads to $l^{GT} \leq l^{y'}$, thereby aggravating the issue typically termed confirmation bias. It can also be explained without the max function approximation. The $LogSumExp$ is monotonically increasing:

$$\frac{\partial}{\partial (l^j - l^{y'})} \log\left(\sum_{i=0}^K e^{l^i - l^{y'}}\right) = \frac{e^{l^j - l^{y'}}}{\sum_{i=0}^K e^{l^i - l^{y'}}} > 0. \quad (5)$$

Thus, when $y' \neq GT$, the logit $l^{y'}$ will be large and l^i of the other categories (including GT 's logit) will be small when $\log(\sum_{i=0}^K e^{l^i - l^{y'}})$ is converged.

With Eq. (1), the loss in VC learning can be defined as follows according to Eq. (2):

$$\mathcal{L}_{VC-CE} = \log\left(\sum_{i=0, i \notin PC}^{K+1} e^{l^i - l^v}\right), \quad (6)$$

where PC is the potential category set. $i = K$ is the index of the virtual category, i.e., $l^{i=K} = l^v$. $i \notin PC$ means the labels in the potential category set are ignored in the summation.

For VC learning, following the derivation of Eq. (3), we obtain:

$$\mathcal{L}_{VC-CE} = \log\left(\sum_{i=0, i \notin PC}^{K+1} e^{l^i - l^v}\right) \approx \max_{i \in \{0, \dots, K\} \setminus PC} (l^i - l^v). \quad (7)$$

Similar to Eq. (4), minimising Eq. (7) is expecting:

$$(l^{i \neq v \wedge i \notin PC} - l^v) \leq (l^{i=v} - l^v) \triangleq 0, \quad i.e., l^{i \neq v \wedge i \notin PC} \leq l^v. \quad (8)$$

Comparing Eq. (8) with Eq. (4) reveals:

- 1) VC loss first ignores the logits $l^{i \in PC}$ of the confusing labels in the uncertain potential set when satisfying the inequation in Eq. (8), thereby avoiding misleading the training.
- 2) Additionally, it provides an alternative upper bound $l^v = f^T \cdot w^v$ for all the rest of the logits $l^{i \notin PC}$. The information embedded by the classifier weight vector can be decoupled into two parts: the direction of the vector and the magnitude(norm) of the vector. Since the magnitude is controlled by a norm factor in this paper, VC learning tries to ensure that the cosine similarity of f and w^v is the maximum. The directions of the weight vectors in a linear classifier can represent the information of different categories [49]. Thus, $l^v = f^T \cdot w^v$ should be larger since w^v is obtained from f^t , which is the feature in




Fig. 6: The transformer layer for the virtual weight generation.

the teacher of the exact same data sample of feature f . f and f^t share a lot of information, thus leading to the largest cosine similarity. Consequently, l^v can be a meaningful upper bound for all the rest logits $l^{i \notin PC}$. The shared information between f and f^t is the upper bound of intra-class information sharing capacity.

The VC learning is applicable not only to the cross-entropy loss function. Mean squared error loss, which is a widely used loss function in several semi-supervised models [50], is also compatible as follows:

$$\mathcal{L}_{VC-MSE} = \sum_{i=0, i \notin PC}^{K+1} (\sigma(l^i) - t^i)^2, \quad (9)$$

where σ is the sigmoid function, t^i is the binary objective label. $t^i = 1$ if the sample is of the category i . Otherwise, $t^i = 0$. The objective label of VC is $t^K = 1$ if the sample is confusing.

3.4 Virtual Weight

As mentioned before, the weight vectors in a linear classifier can be decoupled into two parts: the direction and the magnitude. To make the virtual weight w^v , the most intuitive solution is to directly use the direction of the teacher feature vector f^t and scale it with a magnitude factor (the minimal norm of the pre-define classifier weights is adopted):

$$w^v = \frac{\hat{f}}{\|\hat{f}\|_2} * norm, \quad (10)$$

which is used in our former ECCV publication [13]. In this paper, we further explore a new option.

As shown in Fig. 6, a self-attention transformer layer is adopted to generate a learnable virtual weight. The input tokens consist of the weight vectors $w^{0 \sim K}$ in the student classifier and the feature vector f or f^t . We use the first output of the transformer layer as the virtual weight w^v . This module is trained by the available ground truth and the pseudo labels with high confidence scores. By doing so, the weight generator is expected to be aware of the inherent relation between $w^{0 \sim K}$ and the feature vector to produce the final virtual weight.

3.5 Potential Category Set

The potential category set consists of different predictions of one training sample. Any method that can give reasonably different predictions can be used to build the potential category set. For example, inspired by consistency regularisation, the predictions under different conditions reveal the potential categories of a confusing sample. The pseudo label inferring process can be represented by: $y' = g_\theta(x^u)$, where g is a neural network parameterised by θ , and x^u is an unlabelled input image. By adopting different g_θ , predictions under different conditions, which constitute the PC set, can be obtained. The size of the potential category set is used to determine whether a training sample is confusing to the model or not. If different predictions agree with each other, *i.e.*, the number of elements of the potential category set is one, the training sample is not considered a confusing sample. Otherwise, it is a confusing sample. In this paper, different policies are tailored and investigated for different tasks based on the specificity of each task.

3.5.1 Semantic segmentation

Top-2 probability. Most semi-supervised semantic segmentors usually adopt a high confidence score threshold, such as 0.95, to filter out confusing pixels. However, there are still many valuable low-confidence pixels. We use the categories of the top-2 probability to create the potential category set as the simplest policy for them in semantic segmentation.

Teacher-student mutual verification. Though the parameters of the teacher model are updated by the student’s parameters, comparing the segmentation result $y'^{(s)}$ inferred by the student with $y'^{(t)}$ produced by the teacher finds many confusing pixels in the pseudo labels of semantic segmentation.

3.5.2 Object detection

Unlike semantic segmentation, the location of pseudo labels plays a crucial role in the model training of object detection. Therefore, the confidence score-based methods, such as the top-2 probability policy, cannot be used in object detection since a pseudo bounding box with a low confidence score usually means its position is also unreliable. The prediction format for object detection consists of the coordinates and categories of bounding boxes. The principle of creating a PC set in object detection is comparing two individual prediction sets of one image. Two sets of predictions may have varying numbers of bounding boxes located at different positions. Consequently, we employ the Intersection over Union (IoU) metric to match bounding boxes between the two sets. When the IoU between two boxes exceeds a threshold (set at 0.5 in this paper), we compare their category predictions to determine if they are confusing samples. A bounding box in one set that doesn’t find a match in the other set is considered a confusing sample. For example, if the category of a bounding box in one set is ‘bear’, but finds no matching in the other set, the PC set is $\{bear, bg\}$.

We propose two easy-to-implement methods for object detection to get the two sets for the comparison:

Temporal stability. In object detection, the pseudo label of an image varies at different training iteration steps [7]. When pseudo-labels at different training steps are compared, those

Algorithm 1 The training process with VC learning.

```

"""
x: Tensor(1*3*H*W), unlabelled images
model_t: Teacher model
model_s: Student model
loss_fn: Classification loss function
"""

x_w = aug(x)
x_s = strongaug(x)

# Get logits of K classes and student features
logits, feats = model_s(x_s)

# Get pseudo labels and teacher features
logits_l, feats_l = model_t(x_w)
logits_2, feats_2 = model_s(x_w)
y_1, y_2 = argmax(logits_l), argmax(logits_2)

if y_1 == y_2:
    return loss_fn(softmax(logits), y_1)

# Get logit of virtual category
logit_vc = einsum(feats, vw_gen(feats_l), "nc,nc->n")
logits = cat((logits, logit_vc), dim=1)

# Masked Softmax, ignore potential categories.
probs, mask = masked_softmax(logits, ign_idx=[y_1, y_2])
loss = loss_fn(
    probs, num_classes,
    reduction="none") * mask
loss = loss.sum(dim=1).mean()

return loss

```

einsum: sums the product based on the Einstein summation convention;
masked_softmax: softmax with ignoring index argument;
cat: concatenation.

mismatched pseudo labels reveal the potential categories. We select the model of the current iteration step $g_{\theta^{cur}}$ and $g_{\theta^{last}}$ which is the checkpoint when the model viewed the current image the last time to produce two prediction sets $y'^{(1)}$ and $y'^{(2)}$ for the comparison. In the first training epoch, VC learning does not involve any data, as the model views all images for the first time.

Cross-model verification. Comparing the decisions of two conditionally independent models $g_{\theta^{(a)}}$ and $g_{\theta^{(b)}}$ for the same sample x can also be used to discover the potential categories. Two models are initialised with different initial parameters. The orders of the training data for these two models are also different, ensuring that they do not collapse on each other *i.e.*, $\theta^{(a)} \neq \theta^{(b)}$.

From the aspect of the design of PC set creation, the PC set creation method for object detection can also be used in semantic segmentation. The reason for using different methods for different tasks is rooted in the variances in the implementation aspects of the two tasks. For example, the cross-model verification in semantic segmentation requires an extremely high GPU memory to train two deep segmentation models simultaneously.

In summary, if the size of the potential category set is not equal to 1, it means that this sample should be considered as a confusing sample. The VC learning will take over the training of this confusing sample. Regular loss, such as the cross-entropy loss, is used for unambiguous samples.

4 EXPERIMENTS

In this section, experiments on two dense prediction tasks — semantic segmentation and object detection — are conducted to evaluate the proposed VC learning.

TABLE 1: Ablation study on different t_{low} . The values inside the parentheses indicate peak performance.

t_{low}	0.0	0.3	0.6
mIoU	49.02 (49.02)	49.25 (49.25)	48.94 (49.29)

4.1 Implementation Details

The teacher-student pseudo labelling method serves as the baseline in this section. Given an unlabelled image x^u , the softmax output $P(x^u)$ is inferred by the teacher model. In the semantic segmentation task, $P(x^u)$ is a pixel-level probability matrix. While in object detection, it is an instance-level probability distribution. A predefined confidence threshold t_{low} is used to filter out extremely noisy pseudo labels using $\max(P(x^u)) < t_{low}$. The prediction with a very high confidence score $\max(P(x^u)) > t$ will be fully trusted. We ablate this threshold and found that the final performance does not show significant differences (see Tab. 1). But with a threshold of 0.6 (we adopt in this paper), the convergence speed is faster. The remaining pseudo labels and all available ground truth labels participate in the semi-supervised training. The proposed method is implemented in the PyTorch framework [51]. The code can be found at the public repository¹.

The training process with VC learning is shown in Algorithm 1. Here we take a general classification model and the teacher-student mutual verification as an example for simplicity. If the pseudo labels produced by the teacher and the student agree with each other, VC learning returns the loss value without any additional process. Otherwise, the logit of the potential category will be ignored when the masked softmax function is performed. Finally, the target of the loss function is replaced by the virtual category (here it is K in Algorithm 1) to calculate the loss value.

4.2 Semantic Segmentation

4.2.1 Datasets and Evaluation Protocol

For the semantic segmentation task, we perform experiments on two well-known datasets, Pascal VOC and Cityscape. PascalVOC 2012 consists of 20 classes and one background class. The size of the training images is 1464. Hariharan *et al.* [52] augment PascalVOC 2012 with 9118 additional images. All the 10582 images are adopted in our experiments following the mainstream settings of the semi-supervised segmentation community. Cityscape contains 5k pixel-level labelled images of urban street scenes. The evaluation metric is the mean of Intersection over Union (mIoU).

4.2.2 Settings

The semantic segmentation model of all the following experiments is DeeplabV3+ [53] with a ResNet50 [16] backbone. We randomly divide the dataset with partitions 1/64, 1/96, and 1/128 for labelled/unlabelled data in Pascal VOC, respectively. 1/64, 1/72, and 1/80 are adopted for Cityscapes. The overall loss function is as follows:

$$\mathcal{L} = \mathcal{L}_i(x^l; y) + \beta \mathcal{L}_{VC}(x^u; y'). \quad (11)$$

Other details of the model set-up are shown in Tab. 2.

TABLE 2: The hyperparameters and augmentation settings of VC learning in semi-supervised semantic segmentation.

Hyper-parameter	Value
t	0.95
t_{low}	0.6 (Pascal VOC) / 0.8 (Cityscapes)
VC loss type	CE
EMA momentum	0.9996
optimiser	SGD
learning rate lr	0.001 (Pascal VOC) / 0.002 (Cityscapes)
weight decay	0.0001
momentum	0.9
iteration num	40K (Pascal VOC) / 80K (Cityscapes)
β	1
labelled data number	1/64, 1/96, 1/128 (Pascal VOC) 1/64, 1/72, 1/80 (Cityscapes)
batch size (labelled) bs^l	4
batch size (unlabelled) bs^u	4
Augmentation	Parameters
weak augmentations	
Random Flip	$p = 0.5$
Random Crop	512×512
Random Resize	Scales: [0.5, 0.75, 1, 1.5, 1.75, 2.0]
strong augmentations	
weak aug.	same as above
Random ColorJitter	
Random Grayscale	$p = 0.2$
Random Gaussian Blur	$p = 0.5$
Cutout	$p = 0.5$ s=(0.02, 0.4), r=(0.3, 3.3)




Fig. 7: The mIoU of the teacher-student baseline model with different BN momentums.

4.2.3 Baseline

The teacher-student architecture (mean teacher [21]) is widely adopted in semi-supervised learning. However, it does not work satisfactorily in semantic segmentation. At the beginning of training, the model performance sees a remarkable improvement, but then the mIoU drops off a cliff. We find that the batch normalisation layer is the key component for a strong baseline model. In the limited-supervised setting, it can greatly affect the quality of pseudo labels and aggravate the confirmation bias issue.

A batch normalisation (BN) layer consists of two statistical parameters — an estimated mean and an estimated standard deviation. However, different augmentations fit different BN parameters [57]. The statistical information of the data used for pseudo labelling and model training is inconsistent. The pseudo labelling phase usually adopts a very basic augmentation to process the input data. In contrast, some strong augmentations are used for the data processing of the model training stage. Thus, we manage

1. <https://github.com/GeoffreyChen777/VC>

TABLE 3: The performance of semantic segmentation on Pascal VOC with different label ratios.

Pascal VOC label ratio Fold	1/64				1/96				1/128			
	1	2	3	mean	1	2	3	mean	1	2	3	mean
CCT (CVPR20) [34]	39.12	42.82	50.65	44.20	37.81	38.79	33.19	36.60	23.89	23.89	28.22	25.33
GCT (ECCV20) [54]	41.30	41.34	50.10	44.25	41.07	35.67	40.80	39.18	28.91	35.46	33.85	32.74
CPS (CVPR21) [50]	45.34	44.15	47.83	45.78	40.06	33.78	37.99	37.28	29.71	32.54	31.22	31.16
USRN (CVPR22) [55]	61.70	50.44	53.22	55.12	46.02	44.51	46.40	45.64	39.59	37.43	N/A	38.41
ST++ (CVPR22) [56]	62.11	59.39	62.43	61.31	55.57	54.46	55.35	55.13	45.54	46.99	43.42	45.32
Our Baseline	62.91	60.45	57.61	60.32	56.01	57.84	55.86	56.57	46.08	53.40	52.83	50.77
VC	65.54	63.93	58.86	62.78	59.02	58.60	59.12	58.91	48.94	55.37	54.71	53.00

TABLE 4: The performance of semantic segmentation on Cityscape with different label ratios.

Cityscapes label ratio Fold	1/64				1/72				1/80			
	1	2	3	mean	1	2	3	mean	1	2	3	mean
CCT (CVPR20) [34]	52.59	54.24	51.27	52.70	51.44	48.68	52.50	50.87	51.10	54.01	54.26	53.12
GCT (ECCV20) [54]	52.14	51.52	49.20	50.95	47.43	46.62	52.23	48.76	47.46	51.06	51.31	49.94
CPS (CVPR21) [50]	46.80	48.05	48.02	47.62	42.10	46.69	48.16	45.65	46.36	40.63	44.89	43.96
USRN (CVPR22) [55]	N/A	50.54	N/A	50.54	N/A	N/A	N/A	N/A	N/A	N/A	49.92	49.92
ST++ (CVPR22) [56]	51.67	56.56	49.11	52.45	51.20	47.02	52.77	50.33	44.02	49.64	46.04	46.57
Our Baseline	51.72	56.63	49.84	52.73	53.38	48.96	51.20	51.18	48.35	59.59	53.81	53.92
VC	52.62	58.30	50.74	53.88	55.53	49.92	53.62	53.02	48.92	61.00	54.65	54.86

two groups of statistical parameters in the BN layers for each phase.

Moreover, the estimated statistics are updated with a momentum argument by $x_{new} = (1 - momentum) \times x_{old} + momentum \times x_{current}$ ². Most neural networks adopt a small momentum such as 0.1. As shown in Fig. 7, we find that it leads to a severe training collapse when the numbers of labelled data are very limited. Increasing the momentum alleviates this issue. The possible reason is that a small momentum exacerbates a statistics bias due to the repeating sampling of the labelled data. Each training batch in a teacher-student framework is composed of labelled and unlabelled images. Due to the limited size of the labelled data subset, oversampling is quickly encountered. A small momentum makes the statistics of the labelled subset dominant in the estimated mean and standard deviation.

Thus, we adopt a relatively large momentum for all BN layers. Doing so makes the training stable, leading to an effective straightforward baseline model. The details can be found in the public source code repository.

4.2.4 Performance

We report the performance of our semi-supervised baseline and the model armed with the proposed VC learning in Tabs. 3 and 4. Since different selections of labelled images yield very different results in the extremely scarce-label setting, we randomly select three data folds to make the experiments more convincing. Other methods’ results are reported in their papers or produced by their official codes. On the Cityscapes dataset, USRN cannot be optimised on some data folds due to the extremely unbalanced class distribution issue. Therefore, we use N/A in the table for those experimental results. To have a fair comparison, we use the final checkpoint to produce the results in Tabs. 3 and 4 rather than the checkpoint with the highest mIoU. The

strong teacher-student baseline model outperforms most of the state-of-the-art semi-supervised semantic segmentation algorithms. VC learning further boosts the baseline model to achieve a remarkable improvement. For example, on Pascal VOC fold 2 with a 1/128 label ratio, VC learning achieves a mIoU of 55.37, which is far higher than the recent ST++ (46.99). At most label ratios and data folds, VC learning surpasses others, indicating the superiority of VC learning.

In comparison to the VOC dataset, Cityscapes is a smaller (10k+ vs. 3k) and more challenging dataset. Adding some training data, such as increasing the label ratio from 1/80 to 1/64, does not significantly improve model performance. When we compare the three columns representing mean performance in Tab. 4, it’s evident that most methods achieve similar mIoU scores across different label ratios. There is no clear, consistent trend of performance decline as the label ratio decreases, as observed with methods like CCT, GCT, and ours. However, ST++ significantly underperforms at extremely low label ratios, and one possible reason is overfitting. ST++ introduced a scheme to compare multiple-step predictions to filter out images with unstable predictions. Instead of utilising all available training data to the maximum extent, they discard a significant portion of it and retrain the model with the remaining data. This means that ST++ substantially reduces the number of training images, especially at very low label ratios. As a result, the model is prone to overfitting due to the training data being limited, leading to worse performance. In contrast, our VC learning makes an effort to leverage not only the confident training samples but also the confusing ones, resulting in the best overall performance.

We also evaluate VC learning with different backbones and report the segmentation results in Tab. 5. It is clear that VC learning constantly improves the baseline model with both backbones. On 1/128 label ratio, the baseline model with ResNet101 backbone performs worse than the one with ResNet50, which is possibly caused by overfitting.

2. Here, the momentum is not the ema momentum to update the teacher parameters. It is the momentum in each BN layer for updating the estimated statistics.




Fig. 8: a) and b) Experiments of different strategies for dealing with confusing samples. t.s.m. means the teacher-student mutual verification policy for the potential category set creation. t.2 indicates the top-2 probability method. T means the virtual weight is produced by the weight generator formed by a transformer layer. c) Ratio of confusing samples. d) mIoU of pseudo labels w/ and w/o VC learning.

TABLE 5: Segmentation results of our baseline model and VC learning with different backbones on Pascal VOC.

Backbone	Lable Ratio	Setting	mIoU
ResNet50	1/64	baseline	62.91
		+ VC	65.54
	1/96	baseline	56.01
		+ VC	59.02
ResNet101	1/64	baseline	53.40
		+ VC	55.37
	1/96	baseline	65.73
		+ VC	66.11
ResNet101	1/96	baseline	56.78
		+ VC	60.03
ResNet101	1/128	baseline	50.79
		+ VC	55.23

4.2.5 Analysis and Ablation Study

We ablate VC learning using the Pascal VOC 1/128 label ratio fold1 setting. By doing so, the value of the confusing sample and the effect of VC learning can be observed.

A. Ratio of confusing samples.

The ratio of confusing samples is shown in Fig. 8c. At the very beginning of training, the ratio of confusing samples is nearly 20%. Fig. 8d shows that the mIoU of the pseudo-labels in the early stage is unsurprisingly low. Thus,

TABLE 6: Segmentation results of VC learning with CE loss and MSE loss on Pascal VOC.

Label Ratio	Loss	Fold			
		1	2	3	mean
1/64	CE	65.54	63.93	58.86	62.78
	MSE	65.14	63.19	58.28	62.20
	NEG	64.83	62.13	58.18	61.71
1/96	CE	59.02	58.60	59.12	58.91
	MSE	58.94	58.62	58.49	58.69
	NEG	58.15	58.22	57.54	57.97
1/128	CE	48.94	55.37	54.71	53.00
	MSE	49.45	54.45	54.34	52.75
	NEG	48.80	54.32	54.54	52.55

arbitrarily training the model with low-quality confusing samples makes it become a victim of confirmation bias. VC learning aims to alleviate this problem, thus resulting in better pseudo-label accuracy (w/ vc in Fig. 8d).

B. Policies for confusing samples.

In the previous section, we mentioned that neither discarding nor retaining is the optimal solution for confusing samples. In Fig. 8a, we plot the mIoU curves of these two solutions and our VC learning. The keeping strategy (line 1, also to be regarded as the baseline) performs significantly worse than the other two solutions, which indicates the performance issues that arise due to confirmation bias caused by wrong pseudo labels. By comparing line 3 with line 1, we can see that eliminating the influence of confirmation bias by discarding confusing samples can improve the model performance by far. However, it totally discards the potential contributions of confusing samples. Our VC learning (line 2) proactively utilises confusing samples, which further improves the mIoU to best the other runs. It supports our motivation for making use of confusing samples.

C. Creating methods of potential category set.

In Fig. 8a, The potential category set creation method is the teacher-student mutual verification. It is only applied to the pseudo labels of high confidence pixels (*i.e.*, > 0.95). We also propose a method termed top-2 probability to deal with the low confidence pixels (*i.e.*, < 0.95) and plot the mIoU in Fig. 8b as line 5. By comparing line 5 with line 1, it indicates that VC learning makes effective use of low-confidence data. We then combine the abovementioned two methods to gain an even better mIoU (line 4).

D. Virtual weight.

The previous ablation studies are conducted with the virtual weight proposed in our ECCV publication, *i.e.*, the normalised and scaled teacher feature vectors. This paper introduces a new version produced by a transformer layer. Line 6 in Fig. 8b is the mIoU of VC learning with the virtual weight generated by the transformer. We can see that it exceeds all others, indicating that a better virtual weight is worth exploring in semantic segmentation. Moreover, we ablate the feature vectors from the teacher and the student as the input token of the transformer layer. A similar performance is observed: 48.94 (*f*) v.s. 48.76 (*f*[†]).

To illustrate the individual contributions of direction and magnitude of the virtual weight, we employ cosine




Fig. 9: Visualisation of potential category sets. The row of Pseudo Label (S) is inferred by the student model. The predictions of the teacher model are indicated by Pseudo Label (T).




Fig. 10: Visualization results on the validation set of Pascal VOC.

similarity as the loss function, replacing the cross-entropy (CE) and mean squared error (MSE) losses used in the basic version of VC learning. The results are presented in Tab. 7.

TABLE 7: Cosine Sim. means using cosine similarity to replace the CE or MSE loss in VC learning.

Method	Baseline	Cosine Sim.	VC
mIoU	46.08	47.76	48.62

By maximizing the cosine similarity between the feature vector and the virtual weight while minimizing the cosine similarity between the feature vector and the classifier weight of negative categories, we disregard the magnitudes of both the virtual weight and the classifier weights. Cosine Sim. outperforms the baseline, highlighting the significance of the direction. In addition, it's important to note that VC demonstrates that magnitude is also crucial, as it achieves the best results.

E. Visualisation of potential category set in semantic segmentation.

To analyse what kind of areas VC learning will be mainly applied to, we visualise some demo pictures with their pseudo labels predicted by the student and teacher model in Fig. 9. Disputed pixels are usually located on the boundaries of objects or semantically similar objects. For example, the model makes an indecisive decision about the monitor's boundary in column 1 of Fig. 9. Some pixels belonging to the printer also introduce some false-positive predictions.

F. Qualitative results.

The qualitative visual results of our baseline model and the model with VC learning are in Fig. 10. The baseline model trained by limited labelled data produces many wrong predictions. On the contrary, the model with our VC learning processes these images very well.

G. Loss forms.

VC learning is compatible with not only the cross entropy (CE) loss function but also the mean squared error (MSE) loss function. The comparison results are reported in Tab. 6. The VC learning with CE loss function outperforms the version with MSE loss. They all stand out against other methods, indicating the superiority and compatibility of our approach.

H. Contribution of Virtual Category.

In VC learning, we propose VC as the training label while omitting the categories in the potential category set. With the MSE loss, it is feasible to train the model with only the negative categories outside the potential category set. The performance is reported in Tab. 6 of the label 'NEG'. As indicated in Tab. 6, employing MSE with VC produces superior results compared to using only NEG. This highlights the substantial contribution of VC.

4.3 Object Detection

4.3.1 Datasets and Evaluation Protocol

To evaluate the proposed method for object detection, we assess it on two well-known object detection benchmark datasets – MS COCO [10] and Pascal VOC [58]. Following the mainstream evaluation setting, we use the subset index provided by Unbiased Teacher [1] to split the *train set* across five different labelled ratios: 0.5%, 1%, 2%, 5% and 10% (each ratio uses five random seeds used to obtain an averaged mAP). We also report the performance on Pascal VOC with *VOC07-trainval* as the labelled subset and *VOC12-trainval* as the unlabelled subset. Performance is evaluated on *VOC07-test*. The evaluation metric for all the experiments reported in this subsection is mAP calculated by the COCO evaluation kit [59].

4.3.2 Settings

Following the mainstream choice of the community, we adopted Faster-RCNN [36] with FPN [65] and ResNet-50 [16] as the object detector. The training is conducted on 8 GPUs with batch size of 1/4 per GPU for labelled/unlabelled data. More details are introduced in Tab. 8.

Object detection consists broadly of two subtasks: classification and localisation. Since the classification confidence score is not qualified to indicate the location quality of

TABLE 8: The hyperparameters and augmentation settings of VC learning in semi-supervised object detection.

Hyper-parameter	Value
t	0.7
VC norm	const=3.5
VC loss type	CE with focal loss term
PC set discovery	T.S.
EMA momentum	0.9996
optimiser	SGD
learning rate lr	0.01
weight decay	0.02
momentum	0.9
α	4
β	4
iteration num	180k
labelled data number	0.5-10% (COCO) VOC07 as labelled data
batch size (labelled) bs^l	8
batch size (unlabelled) bs^u	32
Augmentation	Parameters
weak augmentations	
Random Flip	$p = 0.5$
Random Resize	Range: [400, 1200]
strong augmentations	
weak aug.	same as above
Random ColorJitter	
Random Grayscale	$p = 0.2$
Random Gaussian Blur	$p = 0.5$
Cutout	$p = 0.7$ s=(0.05, 0.2), r=(0.3, 3.3) $p = 0.5$ s=(0.02, 0.2), r=(0.1, 6.0) $p = 0.3$ s=(0.02, 0.2), r=(0.05, 8.0)

pseudo labels, some of the previous works [1], [6] disabled the localisation loss of unlabelled data. We find that the method of creating the potential category set can also measure the quality of the location of pseudo labels. When we create the potential category set for a pseudo box b , we evaluate its location shift with the nearby box \hat{b} . We utilise the Intersection over Union (IoU) metric to match the bounding boxes. When the IoU between two boxes b and \hat{b} of two different predictions is higher than a predefined threshold (which we have set at 0.5 in this paper) and is the max one, \hat{b} is considered the nearby box of 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:

$$q_{hor} = \begin{cases} 1, & \frac{(x_1 - \hat{x}_1)}{w} < t_{loc} \quad \& \quad \frac{(x_2 - \hat{x}_2)}{w} < t_{loc} \\ 0, & otherwise \end{cases}, \quad (12)$$

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 if the top and bottom boundaries

are biased. The localisation loss consists of four Smooth-L1 [66] loss terms:

$$\mathcal{L}^{reg^*} = q_{hor}\mathcal{L}^x + q_{ver}\mathcal{L}^y + q_{hor}\mathcal{L}^w + q_{ver}\mathcal{L}^h. \quad (13)$$

In summary, the overall loss function is as follows:

$$\mathcal{L} = \mathcal{L}_l^{cls}(x^l; y) + \mathcal{L}_l^{reg}(x^l; y) + \beta \mathcal{L}_{VC}^{cls}(x^u; y') + \beta \mathcal{L}^{reg^*}(x^u; y'). \quad (14)$$

The classification loss for unlabelled data is replaced by the proposed VC learning loss term. The default hyperparameters and augmentation settings are as shown in Tab. 8.

4.3.3 Performance

MS COCO We first evaluate our method on MS COCO with five label ratios. The results using 5 averaged random seeds are reported in Tab. 9. The results with † are obtained from the available official code. ‘Ours’ are obtained by the model with exactly identical settings of UnbiasedTeacher, which is our baseline model. Given that scale jittering, as used in SoftTeacher, has been shown to have a substantial positive impact, we have incorporated this technique into our approach and report the corresponding results as ‘Ours*’. Furthermore, we adopt a relatively smaller batch size for labelled data, as recommended by SoftTeacher, to accelerate the training process in ‘Ours*’. 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 using 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. 10. Since the source codes of some methods are unavailable, the evaluation styles they used are unclear. Usually, the results based on the VOC-style AP are higher. Thus, we evaluate our method with both 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. 10 indicates that our method can effectively further improve the performance, even if there is already sufficient labelled data.

4.3.4 Analysis and Ablation Study

In this subsection, we choose 1% data of MS COCO as the labelled subset in object detection 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 for 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. 9. The overall ablation study is reported in Tab. 12. The model with VC

TABLE 9: The performance on MS COCO with different label ratios. The results with † are obtained from the available official code. Ours are the results with the same settings of Unbiased Teacher for the sake of fairness. Ours* is the results obtained by using some training settings in Soft Teacher. Different styles of underlines highlight the fair comparisons.

COCO label ratio	0.5%	1%	2%	5%	10%
Supervised	6.83	9.05	12.70	18.47	23.86
CSD (NeurIPS19) [27]	7.41	10.51	13.93	18.63	22.46
STAC (arXiv) [6]	9.78	13.97	18.25	24.38	28.64
Instant Teaching (CVPR21) [12]	-	18.05	22.45	26.75	30.40
Interactive (CVPR21) [7]	-	18.88	22.43	26.37	30.53
Humble Teacher (CVPR21) [60]	-	16.96	21.72	27.70	31.61
Combating Noise (NeurIPS21) [61]	-	18.41	24.00	28.96	32.43
Unbiased Teacher (ICLR21) [1]	16.94	20.75	24.30	28.27	31.50
Soft Teacher (ICCV21) [62]	15.04†	20.46	25.93†	30.74	34.04
MUM (CVPR22) [63]	-	21.27	26.84	31.90	35.92
DTG (NeurIPS22) [64]	18.54	21.88	24.84	28.52	31.87
<u>Ours</u>	18.12	21.61	25.84	30.31	33.45
<u>Ours*</u>	19.46	23.86	27.70	32.05	34.82

TABLE 10: Results of the experiment in VOC. \mathcal{D}^l and \mathcal{D}^u are the labelled and unlabelled subset choices. COCO* consists of the images from COCO that contain objects in the VOC categories. Numbers in () are obtained with the VOC-style AP.

Method	\mathcal{D}^l \mathcal{D}^u	VOC07	
		VOC12	VOC12 + COCO*
STAC [6]		44.64	46.01
Instant Teaching [12]		50.00	50.80
Interactive [7]		46.23	49.59
Humble Teacher [60]		53.04	54.41
Combating Noise [61]		49.30	50.20
Unbiased Teacher [1]		48.69	50.34
MUM [63]		50.22	52.31
<u>Ours</u>		50.40 (55.74)	51.44 (56.70)

TABLE 11: Validation mAP with different strategies to deal with confusing samples.

Strategy	baseline	a) discard	b) keep	c) VC
mAP	20.00	19.37	19.36	20.81

learning and Reg* Loss performs favourably against the baseline model.

A. Ratio of confusing samples.

The ratio of confusing samples is shown in Fig. 11c. At the beginning of training, the ratio of confusing samples is increasing as there are many images the model encountered the first time. This ratio accounts for 20% throughout the entire training process. This illustrates that there are numerous uncertain samples present when the number of labels is extremely limited. Fig. 11d shows the mAP of the pseudo-labels in the early stage, which reveals training the model with low-quality confusing samples yields a worse pseudo label quality, thereby hindering the model performance.

B. Policies for confusing samples.

Here, we adopt the temporal stability verification to create the potential category 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

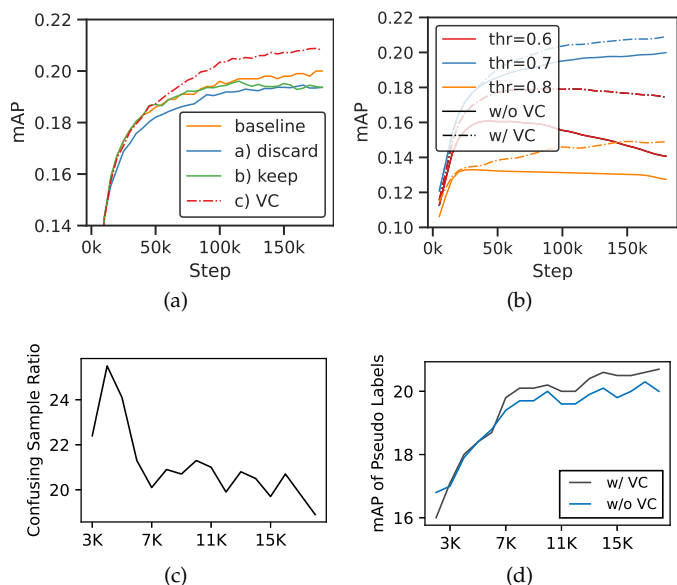


Fig. 11: a) Experiments of different strategies for dealing with confusing samples. As the pseudo labels of the valuable confusing samples are highly unreliable, it is not optimal to either discard or keep them. Our VC learning (dot-dashed line) satisfies both demands, thereby resulting in a significant improvement. b) Experiments of different thresholds of the confidence score filtering w/ or w/o our VC learning. c) Ratio of confusing samples. d) mAP of pseudo labels w/ and w/o VC learning.

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. 11, both discarding and retaining policies decrease the mAP. By analysing the mAP during the entire training presented in Fig. 11a, we noticed that rejecting 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 needs, as shown by the green line in Fig. 11a, training

TABLE 12: Ablation study on VC loss and modified localisation loss Reg* Loss. The method of creating the potential category set in VC learning is temporal stability verification.

VC loss	Reg* Loss	mAP
		20.00
✓		20.81
✓	✓	20.94

TABLE 13: Ablation study of different methods for creating the potential category. We also report the performance with only co-training techniques.

Method	mAP
baseline	20.00
Temporal	20.81
Cross	20.96
co-training w/o VC	20.53

with all potential categories gives a small performance boost at the early stage of training because more under-fitted 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 dot-dashed line in Fig. 11a demonstrates that these confusing samples consistently benefit the model. The mAP sees a rise of 0.81 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 the training.

In addition, as can be seen from Fig. 11b, we evaluate our VC learning with different thresholds (indicated by three colours) of the confidence score filtering adopted by our baseline model Unbiased Teacher. 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.

C. Creating methods of potential category set

Section 3.5.2 explored two methods to create the potential category set. We validate them and report the results in Tab. 13. 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. 13, the co-training can improve mAP by 0.53 individually. For the sake of fairness to other methods without co-training, we use the temporal stability verification in all other experiments, although the cross-model verification performs better.

D. Virtual weight

We choose the feature vector f^t 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

TABLE 14: Ablation study of different image augmentation for virtual weights generation.

Augmentation	none	flipping	strong aug.
mAP	20.80	20.81	20.70

TABLE 15: Ablation study on the contributions of omitting the loss of categories in potential category set and VC.

Method	Omi. (discard)	Omi. (detach)	VC
mIoU	19.37	19.70	20.81

settings: no augmentation, horizontal flipping, and strong augmentation. The results are reported in Tab. 14. No performance gap can be observed between no augmentation and only horizontal flipping. Training with the virtual weight generated by the strong augmentation slightly degrades the mAP. The possible reason might be 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.

In this paper, we introduce a transformer module to generate the virtual weight and its effect has been verified in the segmentation task. We also evaluate it in object detection on 1% labelled COCO. The experiments show that it achieves comparable performance to the basic version (20.74 v.s. 20.81), *i.e.*, the feature vector from the teacher branch. Two main reasons are as follows: 1) unlike semantic segmentation that each pixel must have a label, the position of pseudo labels in object detection is highly uncertain, especially for the background class. Therefore, it is difficult to select a proper position and size of the bounding box to extract local features to train the transformer. The imbalance of foreground and background will also lead to the collapse of its training. 2) the quality of the virtual weight is not the bottleneck of VC learning in object detection. We evaluate the upper bound by using the ground truth to select the real weight as the virtual weight. It indicates that the improvement room is only 0.48 (21.29 vs 20.81). In summary, it is a good choice to adopt the most straightforward solution for object detection, *i.e.*, using f^t , which not only avoids the abovementioned problems but also achieves good results.

E. Contribution of Virtual Category

Similar to the ablation study in segmentation (Sec. 4.2.5 H). We conduct experiments to present the contribution of VC and the discarding operation. As shown in Tab. 15, "Omi. (discard)" denotes that we remove all confusing training samples, while "Omi. (detach)" signifies that we omit the loss components associated with categories in the potential category set and disable the VC's contribution by detaching it from the PyTorch computational graphs. 'Omi. (discard)' yields the worst results. This strategy discards all contributions from confusing samples, even though some correct pseudo-labels are undoubtedly conducive to optimization. 'Omi. (detach)' performs better than 'Omi. (discard)'. This experiment confirms that the accuracy gain does not solely arise from ignoring the loss components of categories in the potential category set as it still falls short of the performance

TABLE 16: Ablation study of different hyperparameters. t_{loc} is the threshold for location quality.

t_{loc}	0.03	0.05	0.1
mAP	20.92	20.94	20.31



Fig. 12: Visualisation of potential category sets in object detection.

achieved by VC. ‘VC’ delivers the best results. Our VC learning strives to leverage all confusing samples to the fullest extent possible, thereby achieving superior performance.

F. Virtual Category learning in different stages

We attempted to incorporate VC learning into the first stage of FasterRCNN, specifically the Region Proposal Network (RPN). Since the RPN is primarily engaged in a binary classification task, distinguishing between foreground and background regions, the confusion arises when determining whether a given region of interest corresponds to an object within the dataset or not. We observed that many of these confusing regions of interest pertain to out-of-dataset objects. Importantly, including such training samples for RPN does not negatively impact the final detection results, as the second stage of detection can effectively handle them. On the contrary, including these out-of-dataset objects during RPN training proves beneficial to some extent in identifying potential objects, especially when there are limited available ground truth labels. The experimental results are detailed in Tab. 18.

G. Visualisation of potential category set.

We visualise some demo pictures with their confusing pseudo labels in Fig. 12. Unlike semantic segmentation, labels for object detection can appear anywhere in an image. Therefore, the teacher model inevitably produces false positive pseudo labels as shown in the first column of Fig. 12. In addition, similar semantics also make the model fail to give accurate predictions.

H. Hyperparameters of \mathcal{L}_{reg^*} Loss

The ablation study of the location quality threshold t_{loc} in the \mathcal{L}_{reg^*} are shown in Tab. 16. A higher threshold will

TABLE 17: The top-1 accuracy on miniImageNet (4k labels).

Method	Top-1 Acc.
MeanTeacher [21]	27.49
Label Propagation [67]	29.71
SimPLE [68]	49.39
Teacher-student Baseline	43.10
Ours	51.49

TABLE 18: Ablation study on VC learning at different stages.

Baseline	VC (stage 2)	VC (stage 1 and 2)
20.00	20.81	20.72

retain more unstable boundaries, leading to worse performance.

4.4 Others

In addition, we simply evaluate VC learning on miniImageNet [69] for a non-dense task — image classification. The top-1 accuracy of the proposed method is presented in Tab. 17. VC learning achieves 51.49% on miniImageNet, surpassing existing works by a large margin.

5 DISCUSSION AND CONCLUSION

In this section, we discuss the limitations of VC learning first. The VC learning takes over the optimisation of the confusing samples. In object detection, by comparing the mAP gains of 10% and other small label ratios, the improvement of our method is slightly lower but still rivals the 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, fewer unlabelled data and a better detector indicate that 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 more practical situations, where very limited labelled data are available. The experiments with very limited labels demonstrated that the effectiveness of VC learning is remarkable.

In conclusion, this paper proposed VC learning, which exploits the confusing under-fitted 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. It can serve as a stepping stone to future work for the community of semi-supervised learning, especially with very limited labels.

REFERENCES

- [1] Y.-C. Liu, C.-Y. Ma, Z. He, C.-W. Kuo, K. Chen, P. Zhang, B. Wu, Z. Kira, and P. Vajda, “Unbiased Teacher for Semi-Supervised Object Detection.” in *Int. Conf. Learn. Represent.*, 2021. **1, 2, 4, 11, 12, 13**
- [2] D. Berthelot, N. Carlini, I. J. Goodfellow, N. Papernot, A. Oliver, and C. Raffel, “Mixmatch: A Holistic Approach to Semi-Supervised Learning.” in *Adv. Neural Inform. Process. Syst.*, 2019, pp. 5050–5060. **1, 3**

- [3] K. Sohn, D. Berthelot, N. Carlini, Z. Zhang, H. Zhang, C. Raffel, E. D. Cubuk, A. Kurakin, and C.-L. Li, "Fixmatch: Simplifying Semi-Supervised Learning with Consistency and Confidence." in *Adv. Neural Inform. Process. Syst.*, 2020. **1, 3**
- [4] J. Li, C. Xiong, and S. C. H. Hoi, "Comatch: Semi-supervised Learning with Contrastive Graph Regularization." in *Int. Conf. Comput. Vis.*, 2021, pp. 9455–9464. **1, 2, 3**
- [5] D.-H. Lee, "Pseudo-Label : The Simple and Efficient Semi-Supervised Learning Method for Deep Neural Networks," in *Int. Conf. Mach. Learn. Workshops*, vol. 3, no. 2, 2013, p. 896. **1**
- [6] K. Sohn, Z. Zhang, C.-L. Li, H. Zhang, C.-Y. Lee, and T. Pfister, "A simple semi-supervised learning framework for object detection," *ArXiv*, vol. abs/2005.04757, 2020. **1, 2, 4, 12, 13**
- [7] Q. Yang, X. Wei, B. Wang, X.-S. Hua, and L. Zhang, "Interactive Self-Training With Mean Teachers for Semi-Supervised Object Detection." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 5941–5950. **1, 3, 4, 7, 13**
- [8] A. Shrivastava, A. Gupta, and R. B. Girshick, "Training Region-Based Object Detectors with Online Hard Example Mining." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 761–769. **1**
- [9] E. Arazo, D. Ortego, P. Albert, N. E. O. Connor, and K. McGuinness, "Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning." in *IEEE Int. Joint Conf. Neural Netw.*, 2020, pp. 1–8. **1**
- [10] T.-Y. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context." in *Eur. Conf. Comput. Vis.*, 2014, pp. 740–755. **1, 2, 3, 11**
- [11] A. Krizhevsky, G. Hinton, and others, "Learning multiple layers of features from tiny images," 2009. **2**
- [12] Q. Zhou, C. Yu, Z. Wang, Q. Qian, and H. Li, "Instant-Teaching: An End-to-End Semi-Supervised Object Detection Framework." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 4081–4090. **2, 4, 13**
- [13] C. Chen, K. Debattista, and J. Han, "Semi-supervised Object Detection via Virtual Category Learning," in *Eur. Conf. Comput. Vis.*, 2022, pp. 169–185. **2, 6**
- [14] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition." in *Int. Conf. Learn. Represent.*, 2015. **3**
- [15] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 1–9. **3**
- [16] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 770–778. **3, 8, 11**
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is All you Need." in *Adv. Neural Inform. Process. Syst.*, vol. 30, 2017, pp. 5998–6008. **3**
- [18] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale." in *Int. Conf. Learn. Represent.*, 2021. **3**
- [19] K. He, X. Zhang, S. Ren, and J. Sun, "Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification," in *Int. Conf. Comput. Vis.* IEEE, 2015, pp. 1026–1034. **3**
- [20] J. Deng, W. Dong, R. Socher, L.-J. Li, u. K. Li, and u. L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2009, pp. 248–255. **3**
- [21] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results." in *Int. Conf. Learn. Represent.*, 2017. **3, 8, 15**
- [22] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling, "Semi-supervised Learning with Deep Generative Models." in *Adv. Neural Inform. Process. Syst.*, 2014, pp. 3581–3589. **3**
- [23] A. Odena, "Semi-supervised learning with generative adversarial networks," *ArXiv*, vol. abs/1606.01583, 2016. **3**
- [24] Y. Luo, J. Zhu, M. Li, Y. Ren, and B. Zhang, "Smooth Neighbors on Teacher Graphs for Semi-Supervised Learning," in *IEEE Conf. Comput. Vis. Pattern Recog.* IEEE, 2018, pp. 8896–8905. **3**
- [25] T. DeVries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," *ArXiv*, vol. abs/1708.04552, 2017. **3**
- [26] T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii, "Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning." *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 8, pp. 1979–1993, 2019. **3**
- [27] J. Jeong, S. Lee, J. Kim, and N. Kwak, "Consistency-based Semi-supervised Learning for Object detection." in *Adv. Neural Inform. Process. Syst.*, 2019, pp. 10758–10767. **3, 4, 13**
- [28] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 3431–3440. **3**
- [29] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation." in *Med. Image. Comput. Assist. Interv.*, 2015, pp. 234–241. **3**
- [30] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs." in *Int. Conf. Learn. Represent.*, 2015. **3**
- [31] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," *ArXiv*, vol. abs/1706.05587, 2017. **3**
- [32] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deepplab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs." *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, 2018. **3**
- [33] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu, "Dual Attention Network for Scene Segmentation." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 3146–3154. **3**
- [34] J. Ouali, C. Hudelot, and M. Tami, "Semi-Supervised Semantic Segmentation With Cross-Consistency Training," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 12 671–12 681. **3, 9**
- [35] Y. Zou, Z. Zhang, H. Zhang, C.-L. Li, X. Bian, J.-B. Huang, and T. Pfister, "Pseudoseg: Designing Pseudo Labels for Semantic Segmentation." in *Int. Conf. Learn. Represent.*, 2021. **3**
- [36] S. Ren, K. He, R. B. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks." *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017. **3, 11**
- [37] Z. Cai and N. Vasconcelos, "Cascade R-CNN: Delving Into High Quality Object Detection." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 6154–6162. **3**
- [38] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 779–788. **3**
- [39] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "Yolov4: Optimal Speed and Accuracy of Object Detection," *ArXiv*, vol. abs/2004.10934, 2020. **3**
- [40] T.-Y. Lin, P. Goyal, R. B. Girshick, K. He, and P. Dollár, "Focal Loss for Dense Object Detection." in *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 2, 2020, pp. 318–327. **3**
- [41] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-End Object Detection with Transformers." in *Eur. Conf. Comput. Vis.*, 2020, pp. 213–229. **3**
- [42] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "Centernet: Keypoint Triplets for Object Detection." in *Int. Conf. Comput. Vis.*, 2019, pp. 6568–6577. **3**
- [43] Z. Tian, C. Shen, H. Chen, and T. He, "Fcos: Fully Convolutional One-Stage Object Detection." in *Int. Conf. Comput. Vis.*, 2019, pp. 9626–9635. **3**
- [44] H. Law and J. Deng, "Cornernet: Detecting objects as paired keypoints," in *Eur. Conf. Comput. Vis.*, 2018, pp. 734–750. **3**
- [45] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Int. Conf. Comput. Vis.*, 2017, pp. 2980–2988. **5**
- [46] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick, "Momentum Contrast for Unsupervised Visual Representation Learning." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 9726–9735. **5**
- [47] T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton, "A Simple Framework for Contrastive Learning of Visual Representations." in *Int. Conf. Mach. Learn.*, 2020, pp. 1597–1607. **5**
- [48] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004, vol. 51. **6**
- [49] H. Qi, M. Brown, and D. G. Lowe, "Low-Shot Learning With Imprinted Weights." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 5822–5830. **6**
- [50] X. Chen, Y. Yuan, G. Zeng, and J. Wang, "Semi-Supervised Semantic Segmentation With Cross Pseudo Supervision." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 2613–2622. **6, 9**
- [51] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison,

- A. Köpf, E. Z. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An Imperative Style, High-Performance Deep Learning Library." in *Adv. Neural Inform. Process. Syst.*, vol. 32, 2019, pp. 8024–8035. **8**
- [52] B. Hariharan, P. Arbelaez, L. D. Bourdev, S. Maji, and J. Malik, "Semantic contours from inverse detectors." in *Int. Conf. Comput. Vis.*, 2011, pp. 991–998. **8**
- [53] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation." in *Eur. Conf. Comput. Vis.*, 2018, pp. 833–851. **8**
- [54] Z. Ke, D. Qiu, K. Li, Q. Yan, and R. W. H. Lau, "Guided Collaborative Training for Pixel-Wise Semi-Supervised Learning." in *Eur. Conf. Comput. Vis.*, 2020, pp. 429–445. **9**
- [55] D. Guan, J. Huang, A. Xiao, and S. Lu, "Unbiased Subclass Regularization for Semi-Supervised Semantic Segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 9968–9978. **9**
- [56] L. Yang, W. Zhuo, L. Qi, Y. Shi, and Y. Gao, "St++: Make Self-training Work Better for Semi-supervised Semantic Segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 4268–4277. **9**
- [57] J. Yuan, Y. Liu, C. Shen, Z. Wang, and H. Li, "A Simple Baseline for Semi-supervised Semantic Segmentation with Strong Data Augmentation*." in *Int. Conf. Comput. Vis.*, 2021, pp. 8209–8218. **8**
- [58] M. Everingham, S. M. A. Eslami, L. V. Gool, C. K. I. Williams, J. M. Winn, and A. Zisserman, "The Pascal Visual Object Classes Challenge: A Retrospective." *Int. J. Comput. Vis.*, vol. 111, no. 1, pp. 98–136, 2015. **11**
- [59] "Coco evaluation kit by facebook detectron2," https://github.com/facebookresearch/detectron2/blob/main/detectron2/evaluation/coco_evaluation.py. **11**
- [60] Y. Tang, W. Chen, Y. Luo, and Y. Zhang, "Humble Teachers Teach Better Students for Semi-Supervised Object Detection." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 3132–3141. **13**
- [61] Z. Wang, Y.-L. Li, Y. Guo, and S. Wang, "Combating Noise: Semi-supervised Learning by Region Uncertainty Quantification." in *Adv. Neural Inform. Process. Syst.*, 2021, pp. 9534–9545. **13**
- [62] M. Xu, Z. Zhang, H. Hu, J. Wang, L. Wang, F. Wei, X. Bai, and Z. Liu, "End-to-End Semi-Supervised Object Detection with Soft Teacher." in *Int. Conf. Comput. Vis.*, 2021, pp. 3040–3049. **13**
- [63] J. Kim, J. Jang, S. Seo, J. Jeong, J. Na, and N. Kwak, "Mum : Mix Image Tiles and UnMix Feature Tiles for Semi-Supervised Object Detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 14512–14521. **13**
- [64] G. Li, X. Li, Y. Wang, S. Zhang, W. Yichao, and D. Liang, "Dtg-SSOD: Dense Teacher Guidance for Semi-Supervised Object Detection," in *Adv. Neural Inform. Process. Syst.*, 2022. **13**
- [65] T.-Y. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie, "Feature Pyramid Networks for Object Detection." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 936–944. **11**
- [66] R. B. Girshick, "Fast R-CNN." in *Int. Conf. Comput. Vis.*, 2015, pp. 1440–1448. **12**
- [67] A. Iscen, G. Toliás, Y. Avrithis, and O. Chum, "Label Propagation for Deep Semi-Supervised Learning," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 5070–5079. **15**
- [68] Z. Hu, Z. Yang, X. Hu, and R. Nevatia, "Simple: Similar Pseudo Label Exploitation for Semi-Supervised Classification." in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 15 099–15 108. **15**
- [69] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, "Matching Networks for One Shot Learning." in *Adv. Neural Inform. Process. Syst.*, 2016, pp. 3630–3638. **15**



Changrui Chen received the bachelor's and master's degrees in computer science and technology from the Ocean University of China in 2017 and 2020, respectively. He is currently pursuing the Ph.D. degree with the WMG, University of Warwick, U.K. His research interests are in computer vision and semi-supervised learning.



actions on Circuits and Systems for Video Technology, and Pattern Recognition.

Jungong Han is the Chair Professor in Computer Vision at the Department of Computer Science, University of Sheffield, U.K. He also holds an Honorary Professorship with the University of Warwick, U.K. His research interests include computer vision, artificial intelligence, and machine learning. He is a Fellow of the International Association of Pattern Recognition and serves as the Associate Editor for several prestigious journals, such as IEEE Transactions on Neural Networks and Learning Systems, IEEE Trans-



Kurt Debattista Kurt Debattista received a B.Sc. in mathematics and computer science, an M.Sc. in psychology, an M.Sc. degree in computer science, and a Ph.D. from the University of Bristol. He is currently a Professor with WMG, at the University of Warwick. His research interests are high-fidelity rendering, HDR imaging, machine learning, and applied perception.