

# POLITE TEACHER: SEMI-SUPERVISED INSTANCE SEGMENTATION WITH MUTUAL LEARNING AND PSEUDO-LABEL THRESHOLDING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We present Polite Teacher, a simple yet effective method for the task of semi-supervised instance segmentation. The proposed architecture relies on the Teacher-Student mutual learning framework. To filter out noisy pseudo-labels, we use confidence thresholding for bounding boxes and mask scoring for masks. The approach has been tested with CenterMask, a single-stage anchor-free detector. Tested on the COCO 2017 *val* dataset, our architecture significantly (approx. +8 pp. in mask AP) outperforms the baseline at different supervision regimes. To the best of our knowledge, this is one of the first works tackling the problem of semi-supervised instance segmentation and the first one devoted to an anchor-free detector. The code is available.

## 1 INTRODUCTION

The advent of deep learning transformed computer vision pipelines both in academia and industry. However, progress is often hindered, since deep learning models are expensive to train for several reasons. Leaving the hardware and computational expenses aside, the vast share of costs often comes from providing the right amount of samples to learn from. For a number of supervised problems in computer vision, it is relatively easy to obtain data. However, labelling them is often the real source of expenses. Semi-supervised learning methods are tailored to deal with the situation in which there are *enough* data samples, but access to the labels is severely limited.

Semantic segmentation (sometimes called dense classification) is a classical computer vision task of assigning each pixel a category. This enables clustering images into semantically coherent parts. Object detection is concerned with the location and identification of semantic objects on images. Instance segmentation combines these two – it is concerned with locating and identifying entities with pixel-wise accuracy. While an intense research activity can be observed in the areas of semi-supervised semantic segmentation (Tarvainen & Valpola, 2017; Chen et al., 2021) and object detection Liu et al. (2021; 2022b), very little attention has been devoted to semi-supervised instance segmentation methods in the computer vision community.

We propose Polite Teacher, a simple yet effective method for the task of semi-supervised instance segmentation. The architecture is built on the Teacher-Student framework. Its *politeness* in the name is an acronym for *pseudo-label thresholding*, which is concerned with filtering noisy pseudo-labels in detection and mask heads. Our contribution is two-fold:

- We present Polite Teacher – one of the first works devoted to semi-supervised instance segmentation and the first one devoted to modern anchor-free detectors. Our approach uses mask scoring (Huang et al., 2019) for the pseudo-mask thresholding.
- The presented method significantly (approx. +8 pp. in mask AP) improves baseline performance with different supervision regimes on COCO 2017 and sets the new baseline for further comparison, becoming de facto the new state-of-the-art for this dataset.

The paper is organised as follows. Section 2 contains a comprehensive survey of related research. In Section 3, we present the details of our method – Polite Teacher. The results of the evaluation are discussed in Section 4, along with the detailed analysis and ablation studies. The paper is concluded with a short summary in Section 5.

## 2 RELATED RESEARCH

This section presents a comprehensive survey of areas adjacent to our work. We present recent research in the area of instance segmentation. Then, we summarise the recent progress in semi-supervised learning. Finally, we combine these two areas and briefly discuss the body of knowledge for semi-supervised instance segmentation.

**Instance segmentation.** Instance segmentation is a computer vision problem concerned with pixel-wise delineating instances of semantic classes. As it can be perceived as a combination of object detection and semantic segmentation, the advances in instance segmentation are tightly coupled with the two aforementioned tasks (especially the former). In recent years, two kinds of object detectors were popular: single- and two-staged. A typical example of the latter category is Faster R-CNN (Ren et al., 2015). It consists of the backbone feature extractor (eg. ResNet) and two heads: the region proposals network (RPN) and the second one for final detections (RoI – the region of interests head). The proposal candidates are searched on a pre-defined set of anchors using RPN and they are later refined with the RoI head. He et al. (2017) introduced Mask R-CNN, which added the mask head to Faster R-CNN to solve segmentation tasks on predicted bounding boxes. Mask scoring (Huang et al., 2019) adds another head on top of that – it regresses the IoU (intersection over union) score of the predicted masks to improve model robustness. Single-stage detectors try to achieve the outcomes of the aforementioned architecture in a single pass. This often results in higher speed at the expense of precision. A notable examples of such detectors include the YOLO family (Redmon et al., 2016) or RetinaNet (Lin et al., 2017). More recently, Lee & Park (2020) proposed CenterMask, an anchor-free instance segmentation framework targeted at real-time applications. It is build on FCOS (Tian et al., 2020), which details are discussed in Section 3. The architecture of CenterMask2 introduces spatial attention-guided masks (SAG-Mask) along with backbone feature extractors tailored for instance segmentation. Recently, there is a surge of research on architectures utilising the concept of self-attention (also called transformers). Dosovitskiy et al. (2021) introduced the visual transformer (ViT), which successfully adapted self-attention to computer vision. This seminal work has sparked research interest in transformers in the vision community. For instance, DINO (Caron et al., 2021) adapts the Teacher-Student paradigm and self-supervised learning for various vision tasks, such as object detection. Recently, MaskDINO (Li et al., 2022) added mask prediction to DINO and topped several instance segmentation benchmarks. However, while displaying exceptional performance, solutions purely based on ViT suffer from quadratic computational complexity, which hinders their adoption.

**Semi-supervised learning.** Semi-supervised learning techniques can be framed as a middle ground between supervised and unsupervised learning since data with and without labels participate in the learning process. It is related to weakly supervised and self-supervised learning. Some approaches to the problem of semi-supervised learning such as  $\Gamma$  model (Rasmus et al., 2015),  $\Pi$  model or temporal ensembling (Laine & Aila, 2017) use the notion of self-ensembling. However, more modern ones are focused on the non-standard architecture during the training phase, often incorporating multiple subnetworks. For instance, Tarvainen & Valpola (2017) introduced Mean Teacher, which is a popular semi-supervised training framework. It overcomes the limitations of Temporal Ensembling and  $\Pi$  models. Instead of using the standard gradient-based approach, the teacher is updated using the exponential moving average (EMA). Unbiased Teacher (Liu et al., 2021) builds on top of the Mean Teacher framework – it does add focal loss and confidence thresholding of pseudo labels. Focal loss borrowed from the work of Lin et al. (2017) helps with the class imbalance, whereas confidence bounding box thresholding reduces the influence of noisy pseudo-labels. The recent Ubiased Teacher v2 (Liu et al., 2022b) extends it to anchor-free detectors and tackles the issue of the pseudo-labelling on bounding box regression. Besides the Teacher-Student paradigm, there are also other approaches. Cross pseudo supervision for semantic segmentation (Chen et al., 2021) is an example of a consistency regularisation method. Here, two networks are trained on the output of each other and are penalised for discrepancies in predictions. Contrastive methods constitute another approach to semi-supervised learning. For instance, regional contrast, abbreviated as ReCo (Liu et al., 2022a) belongs to this category. While using the Teacher-Student framework, this model introduces a dedicated loss function and utilises the semantic relationship between classes.

**Semi-supervised instance segmentation.** Contrary to object detection and semantic segmentation, instance segmentation in the semi-supervised setting received little attention among scholars so far. Concurrently to our work, Wang et al. (2022) presented Noisy Boundaries (NB). This framework also uses the Teacher-Student paradigm and introduces different bounding box thresholds per category, drawing from the work of Radosavovic et al. (2018). The NB architecture has also two special features: the noise-tolerant mask head and boundary preserving re-weighting. While the noise-tolerant head works with low-level resolution features to suppress the noise on mask boundaries, the boundary preserving map is focused on highlighting the boundary region for the segmentation part. At the time of writing this publication, the problem of semi-supervised instance segmentation with anchor-free detectors has never been tackled in the literature.

### 3 POLITE TEACHER

This section is devoted to the introduction of Polite Teacher. First, we formulate the problem we are solving – semi-supervised instance segmentation. Then, we introduce the architecture of our solution – used detectors, the Teacher-Student learning paradigm, and pseudo-label thresholding. The section is concluded with a detailed description of the used loss function.

**Problem formulation.** We consider the problem of *semi-supervised instance segmentation*. Instance segmentation is a computer vision task which combines object detection and semantic segmentation. Semi-supervised setting means that only part of the data available during the training phase is labelled. More formally, we consider training dataset  $\mathcal{D}$  consisting of a set of  $N_{\text{sup}}$  labelled ( $\mathcal{D}^{\text{sup}} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^{N_{\text{sup}}}$ ) and  $N_{\text{unsup}}$  unlabelled ( $\mathcal{D}^{\text{unsup}} = \{\mathbf{x}_i\}_{i=1}^{N_{\text{unsup}}}$ ) images. Here,  $\mathbf{x}_i$  and  $\mathbf{y}_i$  stand for images and their labels (instances categories along with their bounding boxes and masks) respectively. Typically,  $N_{\text{unsup}} \gg N_{\text{sup}}$ . In this work, we assume that  $\mathcal{D}^{\text{sup}}$  and  $\mathcal{D}^{\text{unsup}}$  come from the same distribution.

#### 3.1 ARCHITECTURE

The architecture of Polite Teacher depends on several components. The first one is the detector, which is used twice due to the Teacher-Student paradigm. We use CenterMask (Lee & Park, 2020). This is a single-stage anchor-free detector, which has a relatively simple architecture and therefore it is easy to tune. Two such networks are then framed in the Teacher-Student paradigm to handle both labelled and unlabelled data. Finally, two-fold pseudo-label thresholding takes place to remove noisy ones. The first one uses bounding box uncertainty, and the second one rejects masks with an estimated low IoU score.

**Detector.** To properly present CenterMask, FCOS should be discussed first. Tian et al. (2020) introduced Fully Convolutional One-Stage Object Detector (abbreviated as FCOS), an anchor-free object detector. In general, one-stage detectors due to the lack of the proposal generation phase have fewer hyper-parameters to tune and therefore there are easier to train. Being anchor-free means eliminating pre-defined anchors, which diminishes the computational burden related to calculating IoU scores. FCOS frames detection as a per-pixel prediction task, which resembles semantic segmentation. Three loss components are subject to optimisation: classification, regression, and centre-ness. While classification works similarly to other detectors, the regression targets are quite different. Instead of predicting bounding box corners (like in e.g. Faster R-CNN), the four regressed values are  $l$  (the distance from the centre of a bounding box to its left border),  $t$  (top),  $r$  (right),  $b$  (bottom). Finally, the centre-ness denotes the centre of a given bounding box. Ground-truth centre-ness for  $(l^*, t^*, r^*, b^*)$  is defined as follows:

$$\text{centreness}^* = \sqrt{\frac{\min(l^*, r^*)}{\max(l^*, r^*)} \times \frac{\min(t^*, b^*)}{\max(t^*, b^*)}}. \quad (1)$$

Intuitively, this approach promotes bounding boxes which are located at the centre of a given object.

Lee & Park (2020) introduced CenterMask, which extends FCOS for the task of instance segmentation. This is done similarly to how Faster R-CNN (Ren et al., 2015) is extended by Mask R-CNN (He et al., 2017). However, there are some differences. For instance, the RoI assignment function




Figure 1: Architecture of Polite Teacher (the unsupervised part).

is redefined due to the different levels of the feature pyramid (FPN) which are used. Instead of the mask head from Mask R-CNN, CenterMask utilises the spatial attention-guided mask (abbreviated as SAG-Mask). For  $\mathbf{x}$ , which here mean features extracted from RoI align, the attention-guided feature map is calculated as follows:

$$\mathbf{x}_{\text{sag}} = \sigma(\text{conv}_{3 \times 3}(\text{concat}(P_{\text{max}}, P_{\text{avg}}))) \odot \mathbf{x}, \quad (2)$$

where  $\sigma$  denotes sigmoid function,  $\text{conv}_{3 \times 3}$  is convolutional layer with  $3 \times 3$  filter,  $P_{\text{max}}$  and  $P_{\text{avg}}$  are the results of max and average pooling, and  $\text{concat}$  stands for the concatenation.

**Teacher-Student Learning.** We adopt a 2-step training procedure. In the first step, the model is trained using only labelled data ( $\mathcal{D}^{\text{sup}}$ ), which makes this part a standard supervised instance segmentation. Instead of using a fixed number of batches for this step – as *burn-in stage* in Unbiased Teacher (Liu et al., 2021) – we rather train it as long as it converges in terms of mask AP and take the best model  $\theta$  to ensure the highest results. Naturally, this step is expected to take longer with a higher number of supervised examples. In the second step, mutual Teacher-Student learning with pseudo-labels takes part. The best model from the first step is used and copied to be used as student and teacher models ( $\theta_s \leftarrow \theta, \theta_t \leftarrow \theta$ ). The model can be trained with the burn-in stage as well.

Teacher and student models receive the same input data – they are augmented differently, though. The teacher receives moderately augmented images (*weak* augmentations – we use random flipping), whereas the student consumes visibly perturbed images (*strong* augmentations – same as weak, plus colour jitter, random grayscale, gaussian blur, and random patch erasing). During the training, the predictions from the teacher model serve as pseudo-labels (bounding boxes with their classes and masks) for the student. The teacher is updated using the exponential moving average – see equations 6 and 7 in the next subsection. Figure 1 illustrates the unsupervised part of the process (the second step).

**Pseudo-Label Thresholding.** As the teacher is used to generate pseudo-labels  $\hat{y}$  in the semi-supervision regime, they can be noisy – especially with a high share of unsupervised data. Therefore, Polite Teacher uses two-step pseudo-label thresholding: one is concerned with bounding boxes, whereas the second one refines the masks. Similarly to STAC Sohn et al. (2020) and Unbiased Teacher (Liu et al., 2021), we introduce a bounding box confidence threshold –  $\tau_{\text{cls}}$ . Bounding boxes with a classification score smaller than  $\tau_{\text{cls}}$  are discarded and not used further in the training. The sigmoid output of the classification is treated here as confidence. Inspired by the work of Huang et al. (2019), we also use a mask-scoring mechanism. It regresses the IoU values of the generated masks and improves instance segmentation performance due to the prioritisation of more accurate masks. While not directly designed for the task of semi-supervised learning, the output of this block can be used for filtering noisy pseudo-masks in a straight-forward manner. That is, only masks satisfying  $\hat{y}_{\text{IoU}} > \tau_{\text{IoU}}$  are used in the unsupervised learning stage. The other ones are considered uncertain and receive zero gradients.

### 3.2 OPTIMISATION

The overall batch-wise loss  $\mathcal{L}$  for supervised  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{B_{\text{sup}}}$  and unsupervised  $\{(\mathbf{x}_j, \hat{\mathbf{y}}_j)\}_{j=1}^{B_{\text{unsup}}}$  examples in a batch is computed as follows:

$$\mathcal{L} = \sum_i^{B_{\text{sup}}} \mathcal{L}^{\text{sup}}(\mathbf{x}_i, \mathbf{y}_i) + \lambda \sum_j^{B_{\text{unsup}}} \mathcal{L}^{\text{unsup}}(\mathbf{x}_j, \hat{\mathbf{y}}_j), \quad (3)$$

where  $\mathcal{L}^{\text{sup}}$  is the loss of the supervised part, and  $\mathcal{L}^{\text{unsup}}$  is the loss of the unsupervised part. The unsupervised part is scaled by  $\lambda$ . The supervised component is calculated as follows:

$$\mathcal{L}^{\text{sup}}(\mathbf{x}, \mathbf{y}) = \mathcal{L}_{\text{cls}}^{\text{sup}}(\mathbf{x}, \mathbf{y}) + \mathcal{L}_{\text{centre}}^{\text{sup}}(\mathbf{x}, \mathbf{y}) + \mathcal{L}_{\text{box}}^{\text{sup}}(\mathbf{x}, \mathbf{y}) + \mathcal{L}_{\text{mask}}^{\text{sup}}(\mathbf{x}, \mathbf{y}) + \mathcal{L}_{\text{IoU}}^{\text{sup}}(\mathbf{x}, \mathbf{y}), \quad (4)$$

where  $\mathcal{L}_{\text{cls}}^{\text{sup}}$  is the supervised classification loss,  $\mathcal{L}_{\text{centre}}^{\text{sup}}$  is the supervised centre loss,  $\mathcal{L}_{\text{box}}^{\text{sup}}$  represents the supervised bounding box regression loss, and  $\mathcal{L}_{\text{mask}}^{\text{sup}}$  is the supervised segmentation mask loss, and  $\mathcal{L}_{\text{mask-IoU}}^{\text{sup}}$  is the supervised segmentation mask scoring loss. Regarding the pseudo-labelling loss, we use the following definition:

$$\mathcal{L}^{\text{unsup}}(\mathbf{x}, \hat{\mathbf{y}}) = \mathbb{1}_{\hat{\mathbf{y}}_{\text{cls}} > \tau_{\text{cls}}} \mathcal{L}_{\text{cls}}^{\text{unsup}}(\mathbf{x}, \hat{\mathbf{y}}) + \mathbb{1}_{\hat{\mathbf{y}}_{\text{IoU}} > \tau_{\text{IoU}}} \mathcal{L}_{\text{mask}}^{\text{unsup}}(\mathbf{x}, \hat{\mathbf{y}}) + \mathcal{L}_{\text{IoU}}^{\text{unsup}}(\mathbf{x}, \hat{\mathbf{y}}), \quad (5)$$

where  $\mathcal{L}_{\text{cls}}^{\text{unsup}}$  is the unsupervised classification loss,  $\mathcal{L}_{\text{mask}}^{\text{unsup}}$  is the unsupervised segmentation mask loss, and  $\mathcal{L}_{\text{IoU}}^{\text{unsup}}$  is the unsupervised segmentation mask scoring loss. Regarding the particular loss functions implementation, we follow FCOS and CenterMask:  $\mathcal{L}_{\text{cls}}^{\text{sup}}$  is focal loss (Lin et al., 2017),  $\mathcal{L}_{\text{box}}^{\text{sup}}$  is UnitBox IoU loss (Yu et al., 2016),  $\mathcal{L}_{\text{centre}}^{\text{sup}}$  is binary cross-entropy loss,  $\mathcal{L}_{\text{mask}}^{\text{sup}}$  is average binary cross-entropy loss (He et al., 2017), and  $\mathcal{L}_{\text{IoU}}^{\text{sup}}$  is  $L_2$  loss. The same losses are used for unsupervised components (where applicable).

The student is trained using a standard stochastic gradient descent, whereas the teacher can be perceived as an ensemble of the students:

$$\theta_s \leftarrow \theta_s - \gamma \frac{\partial (\mathcal{L}^{\text{sup}} + \lambda \mathcal{L}^{\text{unsup}})}{\partial \theta_s}, \quad (6)$$

$$\theta_t \leftarrow \alpha \theta_t + (1 - \alpha) \theta_s. \quad (7)$$

Here  $\theta_t$  and  $\theta_s$  represent the teacher and student model parameters respectively and  $\alpha$  is the EMA coefficient (a hyperparameter). Following Liu et al. (2021), the teacher trained in such a way is more robust to the sudden changes of decision boundaries caused by the minority classes in batches – especially in the presence of pseudo-labels. An important practical implication of this is the fact that there is no need to store gradients for the teacher model, which reduces GPU memory usage (compared to simply training two models).

## 4 EVALUATION

This section describes the evaluation of Polite Teacher. We start with discussing the training setup, implementation details, and dataset. Then, we present the result of our main experiment, which is followed by the detailed analysis and ablation studies of particular components of Polite Teacher.

### 4.1 SETUP

**Setup and implementation.** All the experiments were performed either on a2-highgpu-4g instances on Google Cloud Platform (4xA100, 40 GB RAM each) or various machines with 8 GPUs (up to 16 GB RAM each) on the proprietary cluster (each of which contained Titan V, RTX 2080 Ti, or Titan X GPUs). Polite Teacher was developed on CenterMask2 and Unbiased Teacher codebase – both built on the Detectron2 framework (Wu et al., 2019).

**Data and evaluation metrics.** We evaluate Polite Teacher on the MS-COCO 2017 dataset (Lin et al., 2014) using different supervision regimes (1%, 2%, 5% and 10% supervised). The supervised-unsupervised split is taken from the Unbiased Teacher Liu et al. (2021) – while it was originally meant to be used for evaluation of semi-supervised object detection, it can be used with our method as well. We report evaluation results on `val` subset, as the test one is not publicly available. The reported metric used in this study is mask mAP (mean average precision, simply called AP later on), which is calculated as an average of AP with IoU thresholds set from 0.5 to 0.95 (with 0.05 intervals). Bounding box AP is also reported for selected experiments.

Table 1: Results with ResNet50 backbone. Oracle results reported by Lee & Park (2020). The results for two-stage detectors are taken from the work of Wang et al. (2022). Notice that it uses a random split of the dataset – in particular, this is different from the one used by us. Therefore, these results cannot be directly compared, but we report them in order to trace the comparison with their baselines.

Architecture	% supervised			
	1	2	5	10
<i>Mask AP, single-stage detectors (oracle: 34.70%), COCO 2017 val, split from Liu et al. (2021)</i>				
CenterMask2 (Lee & Park, 2020)	10.07	13.46	18.04	22.08
Polite Teacher (ours)	18.33 <sup>(+8.26)</sup>	22.28 <sup>(+8.82)</sup>	26.46 <sup>(+8.42)</sup>	30.08 <sup>(+8.00)</sup>
<i>Mask AP, two-stage detectors (oracle: 34.50%), COCO 2017 val, split from Wang et al. (2022)</i>				
Mask R-CNN (He et al., 2017)	3.5	9.4	17.3	22.0
DD (Radosavovic et al., 2018)	3.8 (+0.30)	11.8 (+2.40)	20.4 (+3.10)	24.2 (+2.20)
NB (Wang et al., 2022)	7.7 (+4.20)	16.3 (+6.90)	24.9 (+7.60)	29.2 (+7.20)

**Hyperparameters.** As base hyperparameters, we used the ones set in CenterMask2. The EMA coefficient  $\alpha$  for Teacher learning is set to 0.9996. The models were trained for up to 270,000 batches with stochastic gradient descent. We used batch size 32 (16 supervised and 16 unsupervised samples) with a learning rate  $\gamma = 0.006$ , weight decay of 0.0001 and momentum 0.9. Similarly to the CenterMask2, the learning rate has been decreased by a factor of 10 on steps 210,000 and 250,000. However, such a long training was often not necessary, as models overfitted on much earlier stages. These experiments have been early stopped. Regarding the pseudo-label thresholding, we used  $\tau_{\text{cls}} = 0.6$  and  $\tau_{\text{IoU}} = 0.9$ . The unsupervised weight has been set to  $\lambda = 2$ . More details on the last three values are in Section 4.3. We use ResNet-50 backbone (He et al., 2016) for all the experiments.

## 4.2 RESULTS

Table 1 shows results for the main experiment conducted on MS-COCO 2017 validation dataset. Polite Teacher reached 18.33/22.28/26.46/30.08 mask AP on 1%/2%/5%/10% respectively, which stands for +8.26/+8.82/+8.42/+8.00 pp. change in this metric over the baseline CenterMask2 respectively. Figure 2 presents qualitative results from different models created in this experiment at different levels of supervision.

For the vast majority of our experiments, we thought that our method will be the first one devoted to semi-supervised instance segmentation. The recent Noisy Boundaries (NB) approach (Wang et al., 2022), a concurrent work to Polite Teacher, is also concerned with this problem and has been evaluated on a similar percentage of supervision on COCO 2017. However, these are different splits. We did not perform direct comparisons, as we were not aware of this work for the majority of our research – we report these results for scientific integrity, though. In general, Noisy Boundaries reported a smaller increase in the mask AP (especially with low supervision), although for fair comparison such claims should be made after running the models on *exactly* the same supervised/unsupervised data splits. It is also unclear how much of this difference can be attributed to the different detectors (a two-staged Mask R-CNN has been used). Following Wang et al. (2022), we also report the results for Data Distillation (DD) method Radosavovic et al. (2018), which was evaluated jointly with NB. It was developed for the task of *omni-supervised* (known also as *webly-supervised*) learning, a special case of semi-supervised learning in which unlabelled data from the Internet are considered during the training. At the heart of this approach lies the pipeline of different data transformations. The results are later ensembled to provide pseudo-labels.

Figure 2: Qualitative Polite Teacher results on COCO 2017 `val` with different supervision regimes.

Table 2: Influence of bounding box filtering threshold (5% supervision).

Metric	$\tau_{cls}$						
	0.5	0.55	0.6	0.65	0.7	0.8	0.9
AP (box)	25.77	26.91	27.46	25.67	23.71	21.16	18.73
AP (mask)	24.14	24.97	25.32	23.92	22.25	20.03	17.65

### 4.3 DETAILED ANALYSIS AND ABLATION STUDIES

In this section, a detailed analysis of the influence of hyperparameters and ablation studies is presented. Unless otherwise specified, all the configuration is the same as in Section 4.1. All experiments have been performed with the 5% supervision regime.

**Influence of bounding box filtering threshold.** In this experiment, we investigate the importance of bounding box filtering thresholds. To separate the influence of sole bounding box filtering, we did not include mask IoU in the optimisation – it is a subject of another experiment. That is,  $\mathcal{L}_{IoU}^{sup}(\mathbf{x}_i, \mathbf{y}_i)$  and  $\mathcal{L}_{IoU}^{unsup}(\mathbf{x}_j, \hat{\mathbf{y}}_j)$  has been not taken into account in equations 4 and 5 respectively. As it turns out, even this significantly improves mask AP over baselines. Table 2 and Figure 3 (left) present AP values for this experiment. The bounding box threshold value with the highest bounding box and mask AP was 0.6. Interestingly, this is a slightly smaller threshold than in the original Unbiased Teacher paper (0.7). The difference might stem from the different neural network architectures (Faster R-CNN vs CenterMask). Note that this experiment used suboptimal  $\lambda = 0.75$  from Equation 3 and hence the results are slightly worse compared to the following experiments. This is because the experiment to determine the correct unsupervised loss weight was yet to be carried out at this point.

**Influence of unsupervised loss weight.** We also examine the influence of the weight of unsupervised loss, which is denoted as  $\lambda$  in equation 3. Figure 3 (centre) and Table 3 presents detailed




Figure 3: Results for experiments controlled for different hyperparameters values in the 5% supervision regime (see Section 4.3, as well as Table 2, 3, and 4).

Table 3: Importance of unsupervised loss weight (5% supervision).

Metric	$\lambda$							
	0.5	1.0	1.5	2.0	2.5	3.0	4.0	8.0
AP (bbox)	26.29	27.50	27.30	28.45	27.63	27.41	27.08	21.88
AP (mask)	24.37	25.43	25.36	26.37	25.69	25.53	25.20	20.53

results of this study. Similarly to the previous experiment, we did not include  $\mathcal{L}_{IoU}^{\text{sup}}(\mathbf{x}_i, \mathbf{y}_i)$  and  $\mathcal{L}_{IoU}^{\text{unsup}}(\mathbf{x}_j, \hat{\mathbf{y}}_j)$  from equations 4 and 5 as the optimisation components. We used  $\tau_{\text{bbox}} = 0.6$ , which is the result of the previous experiment. The highest mask AP has been obtained at  $\lambda = 2.0$ . Interestingly, in Unbiased Teacher, which is a similar architecture, this parameter was set to  $\lambda = 4.0$ .

**Influence of mask scoring filtering threshold.** In this experiment, we investigate the importance of mask filtering threshold  $\tau_{IoU}$ . We use  $\tau_{cls} = 0.6$  and  $\lambda = 2.0$ , as these two values provided best results in the previous experiments. Figure 3 (right) and Table 4 present the detailed results of this experiment. Compared to the results without mask scoring, the best value ( $\tau_{IoU} = 0.9$ ) yielded insignificant differences in mask AP (+0.03 pp.) and bounding box AP (−0.03 pp.). Interestingly, the conducted experiment displayed a convex-like U-shaped relationship between  $\tau_{cls}$  and mask AP. Passing all pseudo-masks resulted in the highest bbox AP, whereas filtering most of them yielded the highest mask AP.

**Ablation on pseudo-bounding box thresholding.** For an ablation study, we compare the baseline CenterMask model to the Teacher-Student with bounding box thresholding. Essentially, such a model is very similar to Unbiased Teacher (Liu et al., 2021), which is proven to greatly improve results for semi-supervised object detection. While raw CenterMask achieved 18.04% on 5% supervision, Polite Teacher yielded 26.46% mask AP, which is a +8.42 pp. increase (see tables 1 and

Table 4: Influence of mask scoring filtering threshold.

Metric	$\tau_{IoU}$					
	0.0	0.5	0.6	0.7	0.8	0.9
AP (bbox)	28.60	27.13	27.34	27.31	28.16	28.42
AP (mask)	25.90	25.38	25.17	25.68	25.96	26.40



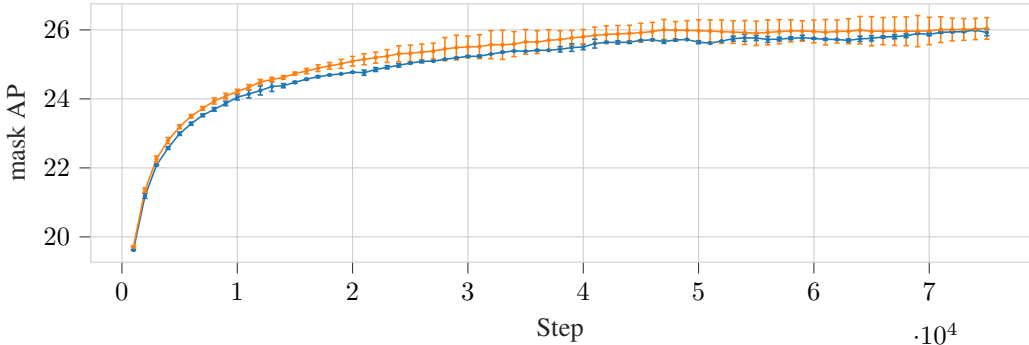


Figure 4: Mean results for Polite Teacher training with 5% supervised data, with (orange, 4 runs) and without (blue, 3 runs) mask scoring.

3). This suggests that much of the mask AP gain can be attributed to the Teacher-Student paradigm with bounding box thresholding.

**Ablation on pseudo-mask thresholding.** In this ablation, we compare the model with bounding box thresholding to the model with bounding box and mask thresholding (that is, Polite Teacher). Judging only by mask AP, the influence of the pseudo-mask filtering threshold on the final results can be easily neglected, as shown in Figure 3 (right). However, applying mask scoring resulted in visibly faster convergence. The model with mask scoring reached 26% mask AP in 40k steps, whereas the model without it needed 74K steps to reach the same value, which is almost two times longer. The highest mask AP values have been reached at 47k (26.39%) and 99k step (26.37%) respectively, which is also close to two times longer. The detailed figures for this run are in tables 3 and 4. In order to check the stability of this behaviour, we repeat these experiments (Figure 4).

**Variance examination.** Due to the computational limitations, we are not reporting results as a series of experiments with their means and standard deviations. However, to assess the variance of the proposed method we carried out a separate experiment, in which we ran Polite Teacher training with 5% supervised data several times – each time with a different seed value. Figure 4 presents mean evaluation results per each step (batch), along with the standard deviations. While the experiment has shown non-homogeneous variance for the variant with mask scoring, the maximum mask AP values are similar: 26.39, 26.39, 26.01, 25.85 with a mean of 26.16 ( $\sigma = 0.27$ ). The variant without mask scoring achieved 26.03, 25.99, 25.99 max mask API (mean 26.00,  $\sigma = 0.02$ ) – that is, an order of magnitude smaller variance, but at the expense of slower convergence and lower metrics value. Interestingly, for the model with the mask scoring head, without the last run, the mean would be 26.26 ( $\sigma = 0.22$ ). While it seems that the last run missed the local optima and much of the per-step variance can be attributed to it, we report all the obtained results for scientific integrity. The high variance might suggest that another hyperparameter should be introduced (e.g. mask scoring head weight) or learning rate should be readjusted.

## 5 SUMMARY

We presented Polite Teacher, a simple and effective architecture for semi-supervised instance segmentation. Tested with a CenterMask, a single-stage detector, our approach yielded approx. +8 pp. mask AP on different supervision regimes with COCO 2017, while it introduces only three hyperparameters to tune. A certain limitation of this study is the lack of validation on other datasets. Similarly, more single-stage detectors, as well as two-stage detectors can be taken into consideration. For more robust evaluation results, several runs of the experiments to explore variance on different values of supervision might be carried out. Therefore, future work should consider validating methods with more detectors, backbones and datasets. Providing a direct comparison with Noisy Boundaries (Wang et al., 2022) might be considered as well. A natural next step would consider taking pseudo-bounding boxes and pseudo-centre-ness regression into account.

## REPRODUCIBILITY STATEMENT

Along with the paper, we provide the source code for running our experiments. This allows recreating the results from our experiments (up to the randomness). The implementation is built on the Detectron2 framework (Wu et al., 2019). Our code is based on the freely available implementations of CenterMask<sup>1</sup> and Unbiased Teacher<sup>2</sup>, which were starting points for this research. After the publication, the code for Polite Teacher will be released on GitHub as well.

## ETHICS STATEMENT

To the best of our knowledge, this work does not possess any direct ethical issues or a negative social impact.

## REFERENCES

- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9650–9660, 2021.
- Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2613–2622, 2021.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.
- Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring r-cnn. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6409–6418, 2019.
- Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=BJ6oOfqge>.
- Youngwan Lee and Jongyoul Park. Centermask: Real-time anchor-free instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 13906–13915, 2020.
- Feng Li, Hao Zhang, Shilong Liu, Lei Zhang, Lionel M Ni, Heung-Yeung Shum, et al. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. *arXiv preprint arXiv:2206.02777*, 2022.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988, 2017.

---

<sup>1</sup><https://github.com/youngwanLEE/centermask2>

<sup>2</sup><https://github.com/facebookresearch/unbiased-teacher>

- Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew J Davison. Bootstrapping semantic segmentation with regional contrast. In *International Conference on Learning Representations*, 2022a.
- Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. In *International Conference on Learning Representations*, 2021. URL [https://openreview.net/forum?id=MJIve1zgR\\_](https://openreview.net/forum?id=MJIve1zgR_).
- Yen-Cheng Liu, Chih-Yao Ma, and Zsolt Kira. Unbiased teacher v2: Semi-supervised object detection for anchor-free and anchor-based detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9819–9828, 2022b.
- Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4119–4128, 2018.
- Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. *Advances in neural information processing systems*, 28, 2015.
- Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788, 2016.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020.
- Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017.
- Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: A simple and strong anchor-free object detector. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Zhenyu Wang, Yali Li, and Shengjin Wang. Noisy boundaries: Lemon or lemonade for semi-supervised instance segmentation? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16826–16835, 2022.
- Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- Jiahui Yu, Yuning Jiang, Zhangyang Wang, Zhimin Cao, and Thomas Huang. Unitbox: An advanced object detection network. In *Proceedings of the 24th ACM international conference on Multimedia*, pp. 516–520, 2016.