

BoxTeacher: Exploring High-Quality Pseudo Labels for Weakly Supervised Instance Segmentation

Tianheng Cheng^{1,*}, Xinggang Wang¹, Shaoyu Chen^{1,*}, Qian Zhang², Wenyu Liu^{1,†}

¹ School of EIC, Huazhong University of Science & Technology

² Horizon Robotics

<https://github.com/hustvl/BoxTeacher>

Abstract

Labeling objects with pixel-wise segmentation requires a huge amount of human labor compared to bounding boxes. Most existing methods for weakly supervised instance segmentation focus on designing heuristic losses with priors from bounding boxes. While, we find that box-supervised methods can produce some fine segmentation masks and we wonder whether the detectors could learn from these fine masks while ignoring low-quality masks. To answer this question, we present *BoxTeacher*, an efficient and end-to-end training framework for high-performance weakly supervised instance segmentation, which leverages a sophisticated teacher to generate high-quality masks as pseudo labels. Considering the massive noisy masks hurt the training, we present a mask-aware confidence score to estimate the quality of pseudo masks, and propose the noise-aware pixel loss and noise-reduced affinity loss to adaptively optimize the student with pseudo masks. Extensive experiments can demonstrate effectiveness of the proposed *BoxTeacher*. Without bells and whistles, *BoxTeacher* remarkably achieves 35.0 mask AP and 36.5 mask AP with *ResNet-50* and *ResNet-101* respectively on the challenging *COCO* dataset, which outperforms the previous state-of-the-art methods by a significant margin and bridges the gap between box-supervised and mask-supervised methods.

1. Introduction

Instance segmentation, aiming at recognizing and segmenting objects in images, is a fairly challenging task in computer vision. Fortunately, the rapid development of object detection methods [7, 40, 50] has greatly advanced the emergence of numbers of successful methods [5, 6, 23, 49, 54, 55] for effective and efficient instance segmenta-

* This work was done when Tianheng Cheng and Shaoyu Chen were interns at Horizon Robotics. † Wenyu Liu is the corresponding author: liuwuy@hust.edu.cn

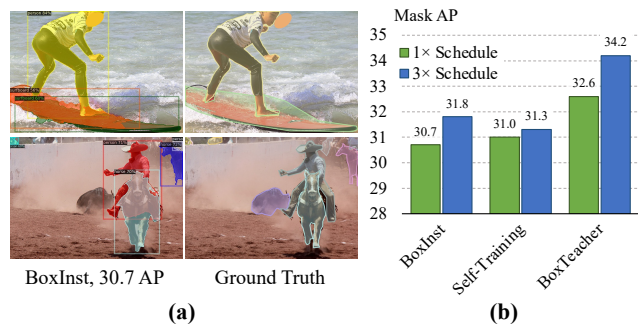


Figure 1. (a) **Segmentation Masks from BoxInst.** BoxInst (*ResNet-50* [24]) can produce some fine segmentation masks with weak supervisions from bounding boxes and images. (b) **Self-Training with Pseudo Masks on COCO val.** We explore the self-training to train a *CondInst* [49] with the pseudo labels generated by BoxInst. However, the improvements are limited

tion. With the fine-grained human annotations, recent instance segmentation methods can achieve impressive results on challenging the *COCO* dataset [34]. Nevertheless, labeling instance-level segmentation is much complicated and time-consuming, *e.g.*, labeling an object with polygon-based masks requires $10.3\times$ more time than that with a 4-point bounding box [11].

Recently, a few works [25, 31–33, 51, 53] explore weakly supervised instance segmentation with box annotations or low-level colors. These weakly supervised methods can effectively train instance segmentation methods [23, 49, 55] without pixel-wise or polygon-based annotations and obtain fine segmentation masks. As shown in Fig. 1(a), *BoxInst* [51] can output a few high-quality segmentation masks and segment well on the object boundary, *e.g.*, the person, even performs better than the ground-truth mask in details though other objects may be badly segmented. Naturally, we wonder if the generated masks of box-supervised methods, especially the high-quality masks, could be qualified as pseudo segmentation labels to further improve the performance of weakly supervised instance segmentation.

To answer this question, we first employ the naive

self-training to evaluate the performance of using box-supervised pseudo masks. Given the generated instance masks from BoxInst, we propose a simple yet effective *box-based pseudo mask assignment* to assign pseudo masks to ground-truth boxes. And then we train the CondInst [49] with the pseudo masks, which has the same architecture with BoxInst and consists of a detector [50] and a dynamic mask head. Fig. 1(b) shows that using self-training brings minor improvements and fails to unleash the power of high-quality pseudo masks, which can be attributed to two obstacles, *i.e.*, (1) the naive self-training fails to filter low-quality masks, and (2) the noisy pseudo masks hurt the training using fully-supervised pixel-wise loss. Besides, the multi-stage self-training is inefficient.

To address these problems, we present BoxTeacher, an end-to-end training framework, which takes advantage of high-quality pseudo masks produced by box supervision. BoxTeacher is composed of a sophisticated *Teacher* and a perturbed *Student*, in which the teacher generates high-quality pseudo instance masks along with the *mask-aware confidence scores* to estimate the quality of masks. Then the proposed *box-based pseudo mask assignment* will assign the pseudo masks to the ground-truth boxes. The student is normally optimized with the ground-truth boxes and pseudo masks through box-based loss and *noise-aware pseudo mask loss*, and then progressively updates the teacher via Exponential Moving Average (EMA). In contrast to the naive multi-stage self-training, BoxTeacher is more simple and efficient. The proposed mask-aware confidence score effectively reduces the impact of low-quality masks. More importantly, pseudo labeling can mutually improve the student and further enforce the teacher to generate higher-quality masks, hence pushing the limits of the box supervision. BoxTeacher can serve as a general training paradigm and is agnostic to the methods for instance segmentation.

To benchmark the proposed BoxTeacher, we adopt CondInst [49] as the basic segmentation method. On the challenging COCO dataset [34], BoxTeacher surprisingly achieves 35.0 and 36.5 mask AP based on ResNet-50 [24] and ResNet-101 respectively, which remarkably outperforms the counterparts. We provide extensive experiments on PASCAL VOC and Cityscapes to demonstrate its effectiveness and generalization ability. Furthermore, BoxTeacher with Swin Transformer [37] obtains 40.6 mask AP as a weakly approach for instance segmentation.

Overall, the contribution can be summarized as follows:

- We solve the box-supervised instance segmentation problem from a new perspective, *i.e.*, self-training with pseudo masks, and illustrate its effectiveness.
- We present BoxTeacher, a simple yet effective framework, which leverages pseudo masks with the mask-aware confidence score and noise-aware pseudo masks

loss. Besides, we propose a pseudo mask assignment to assign pseudo masks to ground-truth boxes.

- We improve the weakly supervised instance segmentation by large margins and bridge the gap between box-supervised and mask-supervised methods, *e.g.*, BoxTeacher achieves 36.5 mask AP on COCO compared to 39.1 AP obtained by CondInst.

2. Related Work

Instance Segmentation. Methods for instance segmentation can be roughly divided into two groups, *i.e.*, single-stage methods and two-stage methods. Single-stage methods [5, 49, 58, 62] tend to adopt single-stage object detectors [35, 50], to localize and recognize objects, and then generate segmentation masks through object embeddings or dynamic convolution [9]. Wang *et al.* present box-free SOLO [54] and SOLOv2 [55], which are independent of object detectors. SparseInst [13] and YOLACT [5], aiming for real-time inference, achieve great trade-off between speed and accuracy. Two-stage methods [14, 23, 27, 29] adopt bounding boxes from object detectors and RoIAlign [23] to extract the RoI (region-of-interest) features for object segmentation, *e.g.*, Mask R-CNN [23]. Several methods [14, 27, 29] based on Mask R-CNN are proposed to refine the segmentation masks for high-quality instance segmentation. Recently, many approaches [7, 10, 12, 17, 20, 63] based on transformers [18, 52] or the Hungarian algorithm [46] have made great progress in instance segmentation.

Weakly Supervised Instance Segmentation. Considering the huge cost of labeling instance segmentation, weakly supervised instance segmentation using image-level labels or bounding boxes gets lots of attention. Several methods [1, 2, 64, 66] exploit image-level labels to generate pseudo masks from activation maps. Khoreva *et al.* [28] propose to generate pseudo masks with GrabCut [42] from given bounding boxes. BoxCaseg [53] leverages a saliency model to generate pseudo object masks for training Mask R-CNN along with the multiple instance learning (MIL) loss. Recently, many box-supervised methods [25, 31, 33, 51] combines the MIL loss or pairwise relation loss from low-level features obtain impressing results with box annotations. In comparison with BoxInst [51], BoxTeacher inherits the box supervision [51] but concentrates more on the novel training paradigm and exploiting noisy pseudo masks for high-performance box-supervised instance segmentation with box annotations. Different from DiscoBox [31] based on mean teacher [48], BoxTeacher aims at a simple yet effective training framework with obtaining high-quality pseudo masks and learning from noisy masks.

Semi-supervised Learning. Pseudo labeling [3, 21, 41] and consistency regularization [4, 30, 43, 44, 59] have greatly

advanced the semi-supervised learning, which enables the training on large-scale unlabeled datasets. Recently, semi-supervised learning has been widely used in object detection [36, 45, 60] and semantic segmentation [8, 56, 61] and demonstrated its effectiveness. Motivated by high-quality masks from box supervision, we adopt the successful pseudo labeling and consistency regularization to develop a new training framework for weakly supervised instance segmentation. Compared to [22] which has similar motivation but aims for semi-supervised object detection with labeled images and extra point annotations, BoxTeacher addresses box-supervised instance segmentation with box-only annotations. Compared to [26, 47] which adopt multi-stage training and combine weakly supervised and semi-supervised learning, BoxTeacher is a one-stage framework without pre-trained labelers.

3. Naive Self-Training with Pseudo Masks

Revisiting Box-supervised Methods. Note that *box-only* annotations is sufficient to train an object detector, which can accurately localize and recognize objects. Box-supervised methods [31, 33, 51] based on object detectors mainly exploit two exquisite losses to supervise mask predictions, *i.e.*, the multiple instance learning (MIL) loss and the pairwise relation loss. Concretely, according to the bounding boxes, the MIL loss can determine the positive and negative bags of pixels of the predicted masks. Pairwise relation loss concentrates on the local relations of pixels from low-level colors or features, in which neighboring pixels have the similar color will be regarded as a positive pair and should output similar probabilities. The MIL loss and pairwise relation loss enables the box-supervised methods to produce the complete segmentation masks, and even some high-quality masks with fine details.

Naive Self-Training. Considering that the box-supervised methods can produce some high-quality masks without mask annotations, we adopt self-training to utilize the high-quality masks as pseudo labels to train an instance segmentation method with full supervision. Specifically, we adopt the successful BoxInst [51] to generate pseudo instance masks on the given dataset $\mathbb{X} = \{\mathcal{X}, \mathcal{B}^g\}$, which only contains the box annotations. For each input image \mathcal{X} , let $\{\mathcal{B}^p, \mathcal{C}^p, \mathcal{M}^p\}$ denote the predicted bounding boxes, confidence scores, and predicted instance masks, respectively. We propose a simple yet effective *Box-based Pseudo Mask Assignment* algorithm in Alg. 1 to assign the predicted instance masks to the box annotations via the confidence scores and intersection-over-union (IoU) between ground-truth boxes \mathcal{B}^g and predicted boxes \mathcal{B}^p . The hyper-parameters τ_{iou} and τ_c are set to 0.5 and 0.05, respectively. The assigned instance masks will be rectified by removing the parts beyond the bounding boxes. Then, we adopt the

dataset $\hat{\mathbb{X}} = \{\mathcal{X}, \mathcal{B}^g, \mathcal{M}^g\}$ with pseudo instance masks to train an approach, *e.g.*, CondInst [49].

Naive Self-Training is Limited. Fig. 1(b) and Tab. 7 provide the experimental results of using naive self-training pseudo masks. Compared to the pseudo labeler, using self-training brings minor improvements and even fails to surpass the pseudo labeler. We attribute the limited performance to two issues, *i.e.*, the naive self-training fails to exclude low-quality masks and the fully-supervised loss is sensitive to the noisy pseudo masks.

Algorithm 1: Box-based Pseudo Mask Assignment

Input: predicted boxes $\mathcal{B}^p \in \mathbb{R}^{N \times 4}$, predicted masks $\mathcal{M}^p \in \mathbb{R}^{N \times H \times W}$, confidence score $\mathcal{C}^p \in \mathbb{R}^N$, ground-truth boxes $\mathcal{B}^g \in \mathbb{R}^{K \times 4}$.

Parameter: IoU threshold τ_{iou} , confidence threshold τ_c .

Output: assigned pseudo masks $\mathcal{M}^g \in \mathbb{R}^{K \times H \times W}$.

- 1 Initialize output masks \mathcal{M}^g with empty (0), assignment index $A \in \mathbb{R}^K$ with -1 ;
- 2 Filter the predictions by the confidence threshold τ_c ;
- 3 Sort the confidence score \mathcal{C}^p in descending order with output indices $S \in \mathbb{N}^N$;
- 4 **foreach** prediction i in S **do**
- 5 Initialize $u \leftarrow -1, v \leftarrow -1$;
- 6 **for** $j = 1$ to K **do**
- 7 $iou_{ij} = \text{ComputeIoU}(\mathcal{B}_i^p, \mathcal{B}_j^g)$;
- 8 **if** $A_j > 0$ **then**
- 9 | continue;
- 10 **end**
- 11 **if** $iou_{ij} \geq \tau_{iou}$ **and** $iou_{ij} \geq u$ **then**
- 12 | $u \leftarrow iou_{ij}, v \leftarrow i$;
- 13 **end**
- 14 **if** $v > 0$ **then**
- 15 | Assign mask \mathcal{M}_i^p to mask \mathcal{M}_v^g ;
- 16 | $A_j \leftarrow i$;
- 17 **end**
- 18 **end**
- 19 **end**

4. BoxTeacher

In this section, we present BoxTeacher, an end-to-end training framework, which aims to unleash the power of pseudo masks. In contrast to multi-stage self-training, BoxTeacher, consisting of a teacher and a student, simultaneously facilitates the training of the student and pseudo labeling of the teacher. The mutual optimization is beneficial to both the teacher and the student, thus leading to higher performance for box-supervised instance segmentation.

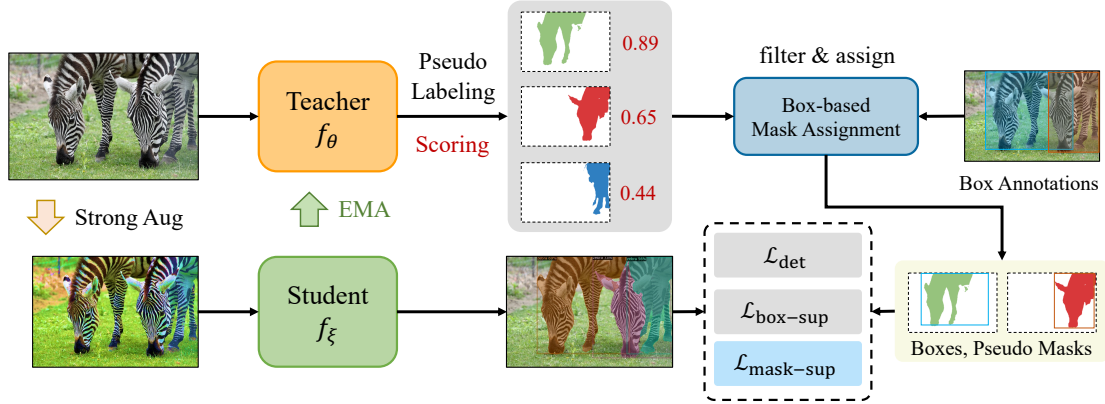


Figure 2. **The Architecture of BoxTeacher.** Images are firstly fed into the *Teacher* to obtain the pseudo masks and estimate the quality of masks. Then the box-based mask assignment filters and assigns pseudo masks to box annotations. The *Student* adopt the augmented images (*i.e.*, random scale or color jittering) and pseudo masks to update the parameters by gradient descent and then update the *Teacher* with exponential moving average (EMA).

4.1. Architecture

The overall architecture of BoxTeacher is depicted in Fig. 2. BoxTeacher is composed of a teacher and a student, which shares the same model. Given the input image, the teacher f_θ straightforwardly generates the predictions, including the bounding boxes, segmentation masks, and *mask-aware confidence scores*. Similarly, we apply the *box-based pseudo mask assignment* in Alg. 1 to assign the predicted masks to the ground-truth annotations. Inspired by consistency regularization [36, 44, 45, 59], we adopt strong augmentation for images (*e.g.*, color jittering) fed into the student f_ξ and the student is optimized under the box supervision and the mask supervision. To acquire high-quality pseudo masks, we adopt the exponential moving average to gradually update the teacher from student [48], *i.e.*, $f_\theta \leftarrow \alpha \cdot f_\theta + (1 - \alpha) \cdot f_\xi$ (α is empirically set 0.999).

Mask-aware Confidence Score. Considering that the generated pseudo masks are noisy and unreliable, which may hurt the performance, we define the mask-aware confidence score to estimate the quality of the pseudo masks. Inspired by [54], we denote $m_i^b \in \mathbb{R}^{H \times W}$ and $m_i \in \mathbb{R}^{H \times W}$ as the box-based binary masks and sigmoid probabilities of the i -th pseudo mask with the detection confidence c_i , the mask-aware confidence score s_i is defined as follows:

$$s_i = \sqrt{c_i \cdot \frac{\sum_{x,y}^{H,W} \mathbb{1}(m_{i,x,y} > \tau_m) \cdot m_{i,x,y} \cdot m_{i,x,y}^b}{\sum_{x,y}^{H,W} \mathbb{1}(m_{i,x,y} > \tau_m) \cdot m_{i,x,y}^b}}, \quad (1)$$

where $\mathbb{1}(\cdot)$ is the indicator function, τ_m is the threshold for binary masks and set to 0.5. The mask-aware score calculates the average probability score of the positive masks inside the ground-truth boxes, and the higher average probability means more confident pixels in the mask. In addition, we explore several kinds of quality scores and compare with the mask-aware score in experiments.

Training Loss. BoxTeacher can be end-to-end optimized with box annotations and the generated pseudo masks. The overall loss is defined as: $\mathcal{L} = \mathcal{L}_{\text{det}} + \mathcal{L}_{\text{box-sup}} + \mathcal{L}_{\text{mask-sup}}$, which consists of the standard detection loss \mathcal{L}_{det} , the box-supervised loss $\mathcal{L}_{\text{box-sup}}$, and the mask-supervised loss $\mathcal{L}_{\text{mask-sup}}$. We inherit the detection loss defined in FCOS [49], and we follow previous works [25, 31, 51] to adopt the max-projection loss and the color-based pairwise relation loss [51] for box-supervised mask loss $\mathcal{L}_{\text{box-sup}}$.

4.2. Noise-aware Pseudo Mask Loss

The goal of BoxTeacher is to take advantage of high-quality pseudo masks in a fully supervised manner while reduce the impact of the noisy or low-quality instance masks. To this end, we present the noise-aware pseudo mask loss in Eq. (2). Ideally, BoxTeacher can leverage the pseudo masks to calculate the fully-supervised pixel-wise segmentation loss, *e.g.*, dice loss [38]. Besides, we also propose a novel *noise-reduced mask affinity loss* $\mathcal{L}_{\text{affinity}}$ to enhance the pixel-wise segmentation with neighboring pixels. Further, we employ the proposed mask-aware confidence scores $\{s_i\}$ as weights for the pseudo mask loss, which adaptively scales the weights for pseudo masks of different qualities. The total pseudo mask loss is defined as follows:

$$\mathcal{L}_{\text{mask-sup}} = \frac{1}{N_p} \sum_{i=1}^{N_p} s_i \cdot (\lambda_p \mathcal{L}_{\text{pixel}}(m_i^p, m_i^g) + \lambda_a \mathcal{L}_{\text{affinity}}(m_i^p, m_i^g)), \quad (2)$$

where m_i^p and m_i^g denotes the i -th predicted masks and pseudo masks, N_p denotes the number of valid pseudo masks, λ_{pixel} and $\lambda_{\text{affinity}}$ are set to 0.5 and 0.1 respectively. To stabilize the training, we adopt a linear warmup strategy for pseudo mask loss at the beginning of the training, *i.e.*, the first 10k iterations.

Noise-reduced Mask Affinity Loss. considering pseudo masks contain much noise while neighboring pixels in local regions (*e.g.*, 3×3) tend to have similar semantics or labels, we exploit label affinities among pixels within local regions to alleviate label noise. Given the i -th pixel sigmoid probability g_i of the pseudo segmentation, we first calculate the refined pixel probability \tilde{g}_i with its neighboring pixels, which is defined as follows:

$$\tilde{g}_i = \frac{1}{2} \left(g_i + \frac{1}{|\mathcal{P}|} \sum_{j \in \mathcal{P}} g_j \right), \quad (3)$$

where \mathcal{P} denotes the set of neighboring pixels, *e.g.*, a 3×3 region. This refinement can reduce the outliers and enhance the pixels with local context. Then, We present a simple noise-reduced mask affinity loss and define the affinity μ_{ij} between i -th and j -th pixels as follows:

$$\mu_{ij} = \tilde{g}_i \cdot \tilde{g}_j + (1 - \tilde{g}_i) \cdot (1 - \tilde{g}_j), \quad (4)$$

where \tilde{g}_i and \tilde{g}_j are refined pixels which encode the local context. Then the *noise-reduced mask affinity loss* for i -th pixel is defined as follows:

$$\mathcal{L}_{\text{affinity}} = - \frac{\sum_{j \in \mathcal{P}} \mathbb{1}(\mu_{ij} > \tau_a) (\log(p_i \cdot p_j) + \log((1-p_i) \cdot (1-p_j)))}{\sum_{j \in \mathcal{P}} \mathbb{1}(\mu_{ij} > \tau_a)}, \quad (5)$$

where $j \in \mathcal{P}$ are the neighboring pixels of the i -th pixel and τ_a is set to 0.5 as default. p_i denotes the i -th pixel of the predicted mask.

5. Experiments

In this section, we mainly evaluate the proposed BoxTeacher on the COCO dataset [34], the Cityscapes dataset [15], and the PASCAL VOC 2012 dataset [19], and provide extensive ablations to analyze the proposed method.

Datasets. The COCO dataset contains 80 categories and 110k images for training, 5k images for validation, and 20k images for testing. The Cityscapes dataset, aiming for perception in driving scenes, consists of 5000 street-view high-resolution images, in which 2975, 500, and 1525 images are used for training, validation, and testing, respectively. Fore-ground objects in Cityscapes are categorized into 8 classes and fine-annotated with pixel-wise segmentation labels instead of polygons adopted in COCO, thus making the labeling process much costly. The PASCAL VOC 2012 dataset has 20 categories and consists of 10582 images for training and 1449 images for validation. For weakly supervised instance segmentation, we only keep the bounding boxes and ignore the segmentation masks during training.

Implementation Details. The proposed BoxTeacher is implemented based on PyTorch [39] and we mainly adopt

CondInst [49] as the meta method for instance segmentation. The backbone networks are initialized with the ImageNet-pretrained weights and the BatchNorm layers are frozen. All BoxTeacher models are trained over 8 GPUs.

Data Augmentation. For images input to the student, we adopt random horizontal flip and the multi-scale augmentation which randomly resizes images from 640 to 800 as the basic augmentation. In addition, we randomly apply color jittering, grayscale, and Gaussian blur for stronger augmentation. While the images fed into the teacher are fixed to 800×1333 without perturbation.

5.1. Instance Segmentation on COCO

Experimental Setup. Following the training recipes [49–51], BoxTeacher is trained with 16 images per batch. Unless specified, we adopt the standard $1 \times$ schedule (90k iterations) [23, 57] with the SGD and the initial learning rate 0.01. For comparisons with the state-of-art methods, we scale up the learning schedule to $3 \times$ (270k iterations).

Main Results. Tab. 1 shows the main results on COCO test-dev. In comparison with other state-of-the-art methods, we evaluate the proposed BoxTeacher with different backbone networks, *i.e.*, ResNet [24] and Swin Transformer [37], and under different training schedules, *i.e.*, $1 \times$ and $3 \times$. It’s clear that BoxTeacher with ResNet-50 achieves 32.9 mask AP, which outperforms other box-supervised methods [31, 51] even with longer schedules. Compared to recent BoxInst [51], BoxLevelSet [33] and DiscoBox [31], BoxTeacher significantly brings about 3.0 mask AP improvements based on ResNet-50 and ResNet-101 under the same setting. Remarkably, BoxTeacher also bridges the gap between mask-supervised methods and box-supervised methods, *e.g.*, the gap between BoxTeacher based on ResNet-101 and CondInst is reduced to 2.6 AP. With the stronger backbones, *e.g.*, Swin Transformer [37], BoxTeacher can surprisingly obtain 40.6 mask AP on COCO dataset, which is highly competitive as a weakly supervised method for instance segmentation.

5.2. Instance Segmentation on PASCAL VOC

Experimental Setup. BoxTeacher is trained for 15k iterations with 16 images per batch. Following previous works [31–33], we report COCO-style AP and AP under 4 thresholds, *i.e.*, $\{0.25, 0.5, 0.70, 0.75\}$.

Main Results. Tab. 2 shows the comparisons with the state-of-the-art methods on PASCAL VOC 2012. Compared to recent box-supervised methods [31, 33, 51], the proposed BoxTeacher achieves better results under different IoU thresholds, which remarkably outperforms BoxInst and DiscoBox by large margins. Notably, BoxTeacher obtains significant improvements under higher IoU thresholds.

Table 1. **Results on COCO Instance Segmentation.** Comparisons with state-of-the-art methods on COCO `test-dev`. With the same backbone or learning schedule, BoxTeacher surprisingly surpasses the counterparts by large margins (more than 2.0 mask AP). †: trained without strong augmentation. ‘R-50’ and ‘R-101’ denote ResNet-50 and ResNet-101, and ‘DCN’ denotes deformable convolution [16,65].

Method	Backbone	Schedule	AP	AP ₅₀	AP ₇₅	AP _s	AP _m	AP _l
<i>Mask-supervised methods.</i>								
Mask R-CNN [23]	R-50-FPN	1×	35.5	57.0	37.8	19.5	37.6	46.0
CondInst [49]	R-50-FPN	1×	35.9	57.0	38.2	19.0	38.6	46.7
CondInst [49]	R-50-FPN	3×	37.7	58.9	40.3	20.4	40.2	48.9
CondInst [49]	R-101-FPN	3×	39.1	60.9	42.0	21.5	41.7	50.9
SOLO [54]	R-101-FPN	6×	37.8	59.5	40.4	16.4	40.6	54.2
SOLOv2 [54]	R-101-FPN	6×	39.7	60.7	42.9	17.3	42.9	57.4
<i>Box-supervised methods.</i>								
BoxInst [51]	R-50-FPN	3×	32.1	55.1	32.4	15.6	34.3	43.5
DiscoBox [31]	R-50-FPN	3×	32.0	53.6	32.6	11.7	33.7	48.4
BoxTeacher [†]	R-50-FPN	1×	32.9	54.1	34.2	17.4	36.3	43.7
BoxTeacher	R-50-FPN	3×	35.0	56.8	36.7	19.0	38.5	45.9
BBTP [25]	R-101-FPN	1×	21.1	45.5	17.2	11.2	22.0	29.8
BBAM [32]	R-101-FPN	1×	25.7	50.0	23.3	-	-	-
BoxCaseg [53]	R-101-FPN	1×	30.9	54.3	30.8	12.1	32.8	46.3
BoxInst [51]	R-101-FPN	3×	33.2	56.5	33.6	16.2	35.3	45.1
BoxLevelSet [33]	R-101-FPN	3×	33.4	56.8	34.1	15.2	36.8	46.8
BoxLevelSet [33]	R-101-DCN-FPN	3×	35.4	59.1	36.7	16.8	38.5	51.3
DiscoBox [31]	R-101-DCN-FPN	3×	35.8	59.8	36.4	16.9	38.7	52.1
BoxTeacher	R-101-FPN	3×	36.5	59.1	38.4	20.1	40.2	47.9
BoxTeacher	R-101-DCN-FPN	3×	37.6	60.3	39.7	21.0	41.8	49.3
BoxTeacher	Swin-Base-FPN	3×	40.6	65.0	42.5	23.4	44.9	54.2

Table 2. **Results on PASCAL VOC.** Comparisons with the state-of-the-art methods on PASCAL VOC 2012 `val`. All methods adopt *box-only* annotations.

Method	Backbone	AP	AP ₂₅	AP ₅₀	AP ₇₀	AP ₇₅
SDI [28]	VGG-16	-	-	44.8	-	16.3
BoxInst [51]	R-50	34.3	-	59.1	-	34.2
DiscoBox [31]	R-50	-	71.4	59.8	41.7	35.5
BoxLevelSet [33]	R-50	36.3	76.3	64.2	43.9	35.9
BoxTeacher	R-50	38.6	77.6	66.4	46.1	38.7
BBTP [25]	R-101	-	75.0	58.9	30.4	21.6
Arun <i>et al.</i> [2]	R-101	-	73.1	57.7	33.5	31.2
BBAM [32]	R-101	-	76.8	63.7	39.5	31.8
BoxInst [51]	R-101	36.4	-	61.4	-	37.0
DiscoBox [31]	R-101	-	72.8	62.2	45.5	37.5
BoxLevelSet [33]	R-101	38.3	77.9	66.3	46.4	38.7
BoxTeacher	R-101	40.3	78.4	67.8	48.0	41.3

5.3. Instance Segmentation on Cityscapes

Experimental Setup. Following previous methods [23, 49], we train all models for 24k iterations with 8 images per batch. The initial learning rate is 0.005. Cityscapes contains high-resolution images (2048×1024), and we randomly resize images from 800 to 1024 for the student and keep the original size for the teacher during training. In addition, we

also adopt the COCO pre-trained models (1× schedule) to initialize the weights for higher performance.

Main Results. Tab. 3 shows the evaluation results on Cityscapes `val`. BoxTeacher outperforms previous box-supervised methods significantly, especially with the COCO pre-trained weights. Though performance gap between fully supervised methods and weakly supervised methods become larger than that in COCO, the human labour of labeling pixel-wise segmentation for a high-resolution Cityscapes image is much costly (90 minutes per image). And we hope future research can bridge the gap between box-supervised methods and mask-supervised methods for high-resolution images.

5.4. Ablation Experiments

Effects of Pseudo Mask Loss. In Tab. 4, we explore the different pseudo mask loss for BoxTeacher. Firstly, we apply the box-supervised loss proposed in [51] achieves 30.7 mask AP (gray row). As shown in Tab. 4, directly applying binary cross entropy (bce) loss with pseudo masks leads to severe performance degradation, which can be attributed to the foreground/background imbalance and noise in pseudo masks. Using dice loss to supervise the training with pseudo masks can bring significant improvements in comparison to the baseline. Adding mask affinity loss $\mathcal{L}_{\text{affinity}}$ provides 0.4



Figure 3. **Visualization Results.** We provide the visualization results of BoxTeacher with ResNet-101 on the COCO test-dev. The proposed BoxTeacher can produce the high-quality segmentation results, even in some complicated scenes.

Table 3. **Results on Cityscapes Instance Segmentation.** Comparisons with state-of-the-art methods for mask AP on Cityscapes val. †: our re-produced results on Cityscapes based on the public code. ‘fine’ denotes the Cityscapes train with fine annotations while ‘fine+COCO’ denotes using COCO pre-trained weights. For box-supervised methods, we remove the fine-grained mask annotations in Cityscapes.

Method	Data	AP	AP ₅₀
<i>Mask-supervised methods.</i>			
Mask R-CNN [23]	fine	31.5	-
CondInst [49]	fine	33.0	59.3
CondInst [49]	fine + COCO	37.8	63.4
<i>Box-supervised methods.</i>			
BoxInst [†] [51]	fine	19.0	41.8
BoxInst [†] [51]	fine + COCO	24.0	51.0
BoxLevelSet [†] [33]	fine	20.7	43.3
BoxLevelSet [†] [33]	fine + COCO	22.7	46.6
BoxTeacher	fine	21.7	47.5
BoxTeacher	fine + COCO	26.8	54.2

AP gain based on the dice loss. Moreover, we ablate the loss weights in pseudo mask loss in Tab. 5.

Effects of Mask-aware Confidence Score. Tab. 6 explores several different scores to estimate the quality of pseudo masks in an unsupervised manner, *i.e.*, (1) classification scores (cls), (2) matched IoU between predicted boxes and ground-truth boxes (iou), (3) mean entropy of the pixel probabilities of pseudo masks (mean-entropy: $s = 1 + \frac{1}{HW} \sum_{i,j} (p_{ij} \log p_{ij} + (1-p_{ij}) \log(1-p_{i,j}))$), (4) the proposed mask-aware score (mask-aware). As Tab. 6 shows, using the proposed mask-aware confidence score leads to better performance for BoxTeacher. Notably, measuring the quality of predicted masks is critical but challenging for leverage pseudo masks. Accurate quality estimation can effectively reduce the impact of noisy masks

Comparisons with Self-Training Paradigm. We adopt the box-supervised approach, *i.e.*, BoxInst [51], to generate pseudo masks, which is pre-trained on COCO with *box-only* annotations. And then we assign the pseudo masks to the ground-truth boxes through the assignment Alg. 1. As shown in Tab. 7, the improvements provided by self-training are much limited and the naive self-training even performs worse than the training with *box-only* annotations, *e.g.*, CondInst with R-50 and $3 \times$ schedule obtains 31.3 AP with pseudo masks, but inferior to the box-supervised version (31.8 AP). Though the self-training scheme enables the supervised training with pseudo masks and achieves comparable performance, we believe the high-quality pseudo masks are not well exploited. Significantly, BoxTeacher achieves higher mask AP compared to both self-training, in an end-to-end manner without complicated steps or procedures for label generation.

Effects of the Strong Data Augmentation. In this study, we explore the effect of strong augmentation on the proposed BoxTeacher, and apply augmentation to the input images of the student. Specifically, we defined two levels of data augmentation in A.1 (in Appendix), *i.e.*, strong augmentation and weak augmentation. As Tab. 8 shows, both strong and weak augmentation hurt the performance of CondInst and BoxTeacher under the $1 \times$ training schedule. Differently, BoxTeacher is more robust to the augmentations as the mask AP drops 0.4 compared to CondInst. However, BoxTeacher remarkably benefits more from the strong data augmentation when increasing the schedule to $3 \times$. In comparison to CondInst, BoxTeacher with strong augmentation will enforce the consistency between the student and teacher. Interestingly, Tab. 8 indicates that using strong augmentation is merely beneficial to the weakly supervised instance segmentation (+0.6 AP), but has no effect to the fully supervised object detection (+0.1AP), sug-

Table 4. **Pseudo Mask Loss.** We evaluate the effects of different loss for BoxTeacher.

$\mathcal{L}_{\text{pixel}}$	$\mathcal{L}_{\text{affinity}}$	AP	AP ₅₀	AP ₇₅
✗	-	30.7	52.5	31.2
bce	-	28.9	49.2	29.5
dice	-	31.8	53.1	32.8
dice	✓	32.2	53.5	33.2

Table 5. **Effect of the Weights of Pseudo Mask Loss.** We adopt $\lambda_{\text{pixel}} = 0.5$ and $\lambda_{\text{affinity}} = 0.1$ as the default setting.

λ_{pixel}	$\lambda_{\text{affinity}}$	AP	AP ₅₀	AP ₇₅
0.1	-	31.4	53.0	32.4
0.5	-	31.8	53.1	32.8
1.0	-	31.5	52.8	32.3
0.5	0.1	32.2	53.5	33.2
0.5	0.5	31.7	52.8	32.8

Table 6. **Effects of Mask Score.** We evaluate different mask scores, and it shows that the mask-aware confidence performs better.

Mask Score	AP	AP ₅₀	AP ₇₅
✗	32.2	53.5	33.2
cls	32.0	53.5	33.1
iou	32.2	53.5	33.4
mean-entropy	31.8	53.3	32.6
mask-aware	32.6	53.5	33.8

Table 7. **Comparison with Naive Self-Training.** As discussed in Sec. 3, we leverage the pre-trained BoxInst to generate pseudo mask labels and assign the pseudo masks to the ground-truth boxes. Then we adopt the pseudo masks and train the CondInst with different schedules and backbones. †: the mask AP achieved by the pseudo labeler, *i.e.*, BoxInst, with *box-only* annotations. ‡: the ideal mask AP could be achieved by CondInst if trained with box annotations following BoxInst.

Method	Backbone	Schedule	Pseudo Label	AP [†]	AP [‡]	AP	AP ₅₀	AP ₇₅
CondInst	R-50	1×	BoxInst, R-50	30.7	30.7	31.0	53.1	31.6
CondInst	R-50	3×	BoxInst, R-50	30.7	31.8	31.3	53.8	31.7
CondInst	R-50	3×	BoxInst, R-101	33.0	31.8	32.5	54.9	33.2
CondInst	R-101	3×	BoxInst, R-101	33.0	33.0	32.9	55.4	33.7
BoxTeacher	R-50	1×	End-to-End	-	-	32.6	53.5	33.8
BoxTeacher	R-50	3×	End-to-End	-	-	34.2	56.0	35.4
BoxTeacher	R-101	3×	End-to-End	-	-	35.2	57.1	36.8

Table 8. **The Effects of Data Augmentation.** We explore whether strong data augmentation will be beneficial to BoxTeacher, which has been widely exploited in semi-supervised methods. We apply weak and strong augmentation to both CondInst and the proposed BoxTeacher. AP^b and AP^m denote the AP for box and mask.

Method	Schd.	weak	strong	AP ^b	AP ^m
CondInst	1×			39.6	36.2
CondInst	1×	✓		39.6	35.6 ^{-0.6}
CondInst	1×		✓	39.2	35.3 ^{-0.9}
BoxTeacher	1×			39.4	32.6
BoxTeacher	1×	✓		39.1	32.4 ^{-0.2}
BoxTeacher	1×		✓	38.8	32.2 ^{-0.4}
CondInst	3×			41.9	37.5
CondInst	3×		✓	42.0	37.6 ^{+0.1}
BoxTeacher	3×			41.7	34.2
BoxTeacher	3×		✓	41.8	34.8 ^{+0.6}

Table 9. **Ablations on Exponential Moving Average.** We evaluate the performance of BoxInst *w/* or *w/o* EMA to make it clear whether the improvements are brought by EMA in BoxTeacher.

Method	w/ EMA	AP ^{bbox}	AP	AP ₅₀	AP ₇₅
BoxInst	✗	39.3	30.6	52.2	31.0
BoxInst	✓	39.4	30.7	52.5	31.2

gesting that consistency regularization might facilitate the learning from noisy pseudo masks.

Effects of Exponential Moving Average. To see whether EMA could partially bring some performance improve-

ments, we re-train BoxInst with EMA to obtain the averaged model to evaluate the performance. Tab. 9 shows that applying EMA has little impact to the final performance, proving that the improvements of BoxTeacher are mainly brought by the effects of pseudo masks.

Qualitative Comparisons. Fig. 3 provides visualization results of the proposed BoxTeacher on the COCO *test-dev*. Even with *box-only* annotations, BoxTeacher can output high-quality segmentation masks with fine boundaries.

6. Conclusions

In this paper, we explore the naive self-training with pseudo labeling for box-supervised instance segmentation, which is much limited by the noisy pseudo masks. To address this issue, we present an effective training framework, namely BoxTeacher, which contains a collaborative teacher and perturbed student for mutually generating high-quality masks and training with pseudo masks. We adopt mask-aware confidence scores to measure the quality of pseudo masks and noise-aware mask loss to train the student with pseudo masks. In the experiments, BoxTeacher achieves promising improvements on COCO, PASCAL VOC, and Cityscapes datasets, indicating that the proposed training framework is effective and can achieve higher level of weakly supervised instance segmentation.

Acknowledgement. This work was in part supported by NSFC (No. 62276108).

References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *CVPR*, 2019.
- [2] Aditya Arun, C. V. Jawahar, and M. Pawan Kumar. Weakly supervised instance segmentation by learning annotation consistent instances. In *ECCV*, 2020.
- [3] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *NeurIPS*, 2014.
- [4] David Berthelot, Nicholas Carlini, Ian J. Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019.
- [5] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT: real-time instance segmentation. In *ICCV*, 2019.
- [6] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: high quality object detection and instance segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1483–1498, 2021.
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [8] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *CVPR*, 2021.
- [9] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *CVPR*, 2020.
- [10] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2021.
- [11] Bowen Cheng, Omkar Parkhi, and Alexander Kirillov. Pointly-supervised instance segmentation. In *CVPR*, 2022.
- [12] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *NeurIPS*, 2021.
- [13] Tianheng Cheng, Xinggang Wang, Shaoyu Chen, Wenqiang Zhang, Qian Zhang, Chang Huang, Zhaoxiang Zhang, and Wenyu Liu. Sparse instance activation for real-time instance segmentation. In *CVPR*, 2022.
- [14] Tianheng Cheng, Xinggang Wang, Lichao Huang, and Wenyu Liu. Boundary-preserving Mask R-CNN. In *ECCV*, 2020.
- [15] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [16] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, 2017.
- [17] Bin Dong, Fangao Zeng, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. SOLQ: segmenting objects by learning queries. In *NeurIPS*, 2021.
- [18] 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 *ICLR*, 2021.
- [19] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 2010.
- [20] Yuxin Fang, Shusheng Yang, Xinggang Wang, Yu Li, Chen Fang, Ying Shan, Bin Feng, and Wenyu Liu. Instances as queries. In *ICCV*, 2021.
- [21] Stanley C. Fraclick. Learning to recognize patterns without a teacher. *IEEE Trans. Inf. Theory*, 1967.
- [22] Yongtao Ge, Qiang feng Zhou, Xinlong Wang, Chunhua Shen, Zhibin Wang, and Hao Li. Point-teaching: Weakly semi-supervised object detection with point annotations. *arXiv*, abs/2206.00274, 2022.
- [23] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *ICCV*, 2017.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [25] Cheng-Chun Hsu, Kuang-Jui Hsu, Chung-Chi Tsai, Yen-Yu Lin, and Yung-Yu Chuang. Weakly supervised instance segmentation using the bounding box tightness prior. In *NeurIPS*, pages 6582–6593, 2019.
- [26] Zitong Huang, Yiping Bao, Bowen Dong, Erjin Zhou, and Wangmeng Zuo. W2n: Switching from weak supervision to noisy supervision for object detection. In *ECCV*, 2022.
- [27] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring R-CNN. In *CVPR*, 2019.
- [28] Anna Khoreva, Rodrigo Benenson, Jan Hendrik Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, 2017.
- [29] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross B. Girshick. Pointrend: Image segmentation as rendering. In *CVPR*, 2020.
- [30] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017.
- [31] Shiyi Lan, Zhiding Yu, Christopher B. Choy, Subhashree Radhakrishnan, Guilin Liu, Yuke Zhu, Larry S. Davis, and Anima Anandkumar. Discobox: Weakly supervised instance segmentation and semantic correspondence from box supervision. In *ICCV*, 2021.
- [32] Jungbeom Lee, Jihun Yi, Chaehun Shin, and Sungroh Yoon. BBAM: bounding box attribution map for weakly supervised semantic and instance segmentation. In *CVPR*, 2021.
- [33] Wentong Li, Wenyu Liu, Jianke Zhu, Miaomiao Cui, Xiansheng Hua, and Lei Zhang. Box-supervised instance segmentation with level set evolution. In *ECCV*, 2022.
- [34] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, 2014.

- [35] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. In *ECCV*, 2016.
- [36] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. In *ICLR*, 2021.
- [37] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021.
- [38] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 2016.
- [39] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- [40] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017.
- [41] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S. Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *ICLR*, 2021.
- [42] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. "grabcut": interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 2004.
- [43] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NeurIPS*, pages 1163–1171, 2016.
- [44] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020.
- [45] Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *CoRR*, abs/2005.04757, 2020.
- [46] Russell Stewart, Mykhaylo Andriluka, and Andrew Y. Ng. End-to-end people detection in crowded scenes. In *CVPR*, 2016.
- [47] Lin Sui, Chen-Lin Zhang, and Jianxin Wu. Salvage of supervision in weakly supervised object detection. In *CVPR*, 2022.
- [48] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017.
- [49] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *ECCV*, 2020.
- [50] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: fully convolutional one-stage object detection. In *ICCV*, 2019.
- [51] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. Boxinst: High-performance instance segmentation with box annotations. In *CVPR*, 2021.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [53] Xinggang Wang, Jiawei Feng, Bin Hu, Qi Ding, Longjin Ran, Xiaoxin Chen, and Wenyu Liu. Weakly-supervised instance segmentation via class-agnostic learning with salient images. In *CVPR*, 2021.
- [54] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. SOLO: segmenting objects by locations. In *ECCV*, 2020.
- [55] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. In *NeurIPS*, 2020.
- [56] Yuchao Wang, Haochen Wang, Yujun Shen, Jingjing Fei, Wei Li, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xinyi Le. Semi-supervised semantic segmentation using unreliable pseudo-labels. In *CVPR*, 2022.
- [57] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [58] Enze Xie, Peize Sun, Xiaoge Song, Wenhui Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask: Single shot instance segmentation with polar representation. In *CVPR*, 2020.
- [59] Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, 2020.
- [60] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *ICCV*, 2021.
- [61] Jianlong Yuan, Yifan Liu, Chunhua Shen, Zhibin Wang, and Hao Li. A simple baseline for semi-supervised semantic segmentation with strong data augmentation*. In *ICCV*, 2021.
- [62] Rufeng Zhang, Zhi Tian, Chunhua Shen, Mingyu You, and Youliang Yan. Mask encoding for single shot instance segmentation. In *CVPR*, 2020.
- [63] Wenwei Zhang, Jiangmiao Pang, Kai Chen, and Chen Change Loy. K-net: Towards unified image segmentation. In *NeurIPS*, 2021.
- [64] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In *CVPR*, 2018.
- [65] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets V2: more deformable, better results. In *CVPR*, 2019.
- [66] Yi Zhu, Yanzhao Zhou, Huijuan Xu, Qixiang Ye, David S. Doermann, and Jianbin Jiao. Learning instance activation maps for weakly supervised instance segmentation. In *CVPR*, 2019.