

# Semantic-aware SAM for Point-Prompted Instance Segmentation

Zhaoyang Wei<sup>1\*</sup>, Pengfei Chen<sup>1\*</sup>, Xuehui Yu<sup>1\*</sup>, Guorong Li<sup>1</sup>,  
 Jianbin Jiao<sup>1</sup>, Zhenjun Han<sup>1†</sup>

<sup>1</sup>University of Chinese Academy of Sciences(UCAS)

## Abstract

Single-point annotation in visual tasks, with the goal of minimizing labelling costs, is becoming increasingly prominent in research. Recently, visual foundation models, such as Segment Anything (SAM), have gained widespread usage due to their robust zero-shot capabilities and exceptional annotation performance. However, SAM’s class-agnostic output and high confidence in local segmentation introduce **semantic ambiguity**, posing a challenge for precise category-specific segmentation. In this paper, we introduce a cost-effective category-specific segmenter using SAM. To tackle this challenge, we have devised a Semantic-Aware Instance Segmentation Network (SAPNet) that integrates Multiple Instance Learning (MIL) with matching capability and SAM with point prompts. SAPNet strategically selects the most representative mask proposals generated by SAM to supervise segmentation, with a specific focus on object category information. Moreover, we introduce the Point Distance Guidance and Box Mining Strategy to mitigate inherent challenges: **group** and **local** issues in weakly supervised segmentation. These strategies serve to further enhance the overall segmentation performance. The experimental results on Pascal VOC and COCO demonstrate the promising performance of our proposed SAPNet, emphasizing its semantic matching capabilities and its potential to advance point-prompted instance segmentation. The code is available at <https://github.com/zhaoyangwei123/SAPNet>.

## 1. Introduction

Instance segmentation seeks to discern pixel-level labels for both instances of interest and their semantic content in images, a crucial function in domains like autonomous driving, image editing, and human-computer interaction. Despite impressive results demonstrated by various studies [5, 11, 16, 40–42], the majority of these high-performing methods are trained in a fully supervised manner and heavily dependent on detailed pixel-level mask annotations,

\* Equal contribution.

† Corresponding authors. (hanzhj@ucas.ac.cn)

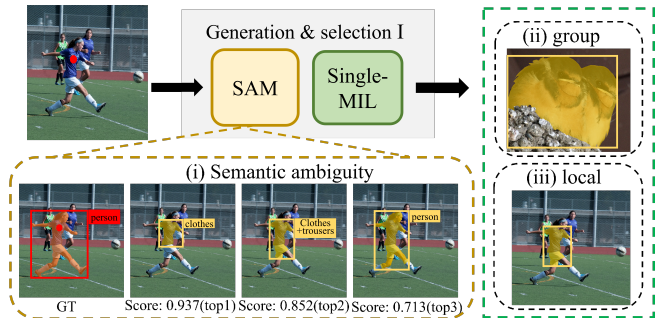


Figure 1. Three Challenges Brought by SAM and single-MIL. Orange dash box illustrates that **semantic ambiguity** in SAM-generated masks, where it erroneously assigns higher scores to non-object categories like clothes, despite the person being our desired target. Green dash box depicts a comparison between mask proposals using single-MIL and SAPNet. It illustrates two primary challenges: **‘group’**, where segmentation encounters difficulties in isolating individual targets among adjacent objects of the same category, and **‘local’**, where MIL favors foreground-dominant regions, resulting in overlooked local details.

thereby incurring significant labeling costs. To address this challenge, researchers are increasingly focusing on weakly supervised instance segmentation, leveraging cost-effective supervision methods, such as bounding boxes [23, 27, 39], points [14, 28], and image-level labels [21, 45].

Recently, visual foundation models, such as Segment Anything (SAM)[22], have been widely employed by researchers for their exceptional generalization capabilities and impressive annotation performance. Numerous studies based on SAM, such as [20, 44] have emerged, building upon the foundations of SAM to further enhance its generalization capabilities and efficiency. However, these efforts have predominantly focused on improving the annotation performance of SAM. One limitation arises from SAM’s lack of classification ability, resulting in class-agnostic segmentation results that fail to accurately segment specific categories as desired.

To tackle the inherent semantic ambiguity in SAM and achieve specific-category segmentation, we propose integrating weak annotations with SAM, employing point annotations as prompts to imbue semantic information into

SAM’s outputs. A straightforward approach involves leveraging SAM’s intrinsic scoring mechanism, selecting the top-scoring mask as the corresponding label for each category. However, when annotating object points are fed into the SAM, its category-agnostic characteristic tends to assign higher scores to parts of the object, resulting in generated mask annotations that fail to encompass the object as a whole. In Fig. 1 orange dashed box, we aim to obtain the ‘person’ mask annotation, but SAM predicts the proposals of ‘clothes’, ‘clothes+trousers’ and ‘person’. Relying solely on the score SAM provides is insufficient, as the highest score corresponds to ‘clothes’ (col-2), which does not meet our specific needs.

To address this challenge, we have proposed SAPNet, a semantically-aware instance segmentation network designed for high-quality, end-to-end segmentation. In this study, we design a proposal selection module (PSM) using the Multiple Instance Learning (MIL) paradigm to choose proposals that align closely with the specified semantic label. However, the MIL-based method relies on the classification score, often leading to group and local predictions [4, 21, 24]. In Fig. 1 green dashed box, the group issue is evident, where two objects of the same category are often both included when they are in close proximity. It also illustrates the local issue, where the MIL classifier frequently predicts the most discriminative region instead of the entire object. To overcome these limitations, we have introduced Point Distance Guidance (PDG) and Box Mining Strategies (BMS). Specifically, we penalize the selection results by calculating the Euclidean distances between the annotated points of identical categories enclosed within the proposals. Additionally, for more localized proposals, we filter out higher-quality proposals from their corresponding bags and dynamically merge them in scale. By fully exploiting the positional clues to prevent local and group prediction, we aim to select the proposal that most effectively represents the object category in refinement stage. The primary contributions of this work can be outlined as follows:

- 1) We introduce SAPNet, an end-to-end semantic-aware instance segmentation network based on point prompts. SAPNet combines the visual foundation model SAM with semantic information to address its inherent semantic ambiguity, facilitating the generation of semantically-aware proposal masks.

- 2) We incorporate Point Distance Guidance (PDG) and Box Mining Strategies (BMS) to prevent local and group predictions induced by MIL-based classifiers in both the proposal selection and refinement stages.

- 3) SAPNet achieves state-of-the-art performance in Point-Prompted Instance Segmentation (PPIS), significantly bridging the gap between point-prompted and fully-supervised segmentation methods on two challenging benchmarks (COCO and VOC2012).

## 2. Related Work

**Weakly-Supervised Instance Segmentation (WSIS)** offers a practical approach for accurate object masks using minimal supervision. It spans a range of annotations, from image labels to bounding boxes. Research has focused on narrowing the performance gap between weakly and fully-supervised methods, primarily through box-level [18, 25, 39] and image-level annotations [1, 21]. Box-based methods have explored structural constraints to guide the segmentation, as seen in BBTP [18], BoxInst [39], and Box2Mask [29], and applied structural constraints to drive segmentation, treating it as a multiple-instance learning task or enforcing color consistency based on CondInst [40]. These approaches, while innovative, can complicate training and sometimes neglect the object’s overall shape due to their focus on local features and proposal generation, like MCG [2]. Conversely, the proposal-free methods, like IRN [1], rely on class relationships for mask production but can falter in accurately separating instances. To preserve object integrity, recent methods such as Discobox [23] and BESTIE [21] integrate advanced semantic insights into instance segmentation using pairwise losses or saliency cues [30, 39, 42]. However, semantic drift remains an issue, with mislabeling or missed instances resulting in inferior pseudo labels [3] compromising segmentation quality.

**Pointly-Supervised Detection and Segmentation (PSDS)** cleverly balances minimal annotation costs with satisfactory localization accuracy. By introducing point annotations, WISE-Net [24], P2BNet [9] and BESTIE [21] improve upon weakly supervised methods that suffer from vague localizations. That only slightly increases the costs (by about 10%) and is almost as quick as the image-level annotation, but that is far speedier than more detailed bounding box or mask annotations. Such precision allows for tackling semantic bias, as seen in methods like PointRend [12], which utilize multiple points for improved accuracy, despite requiring additional bounding box supervision. Recent advancements in point-supervised instance segmentation, employed by WISE-Net and Point2Mask [28], show that even single-point annotations can yield precise mask proposals. WISE-Net skillfully localizes objects and selects masks, while BESTIE enhances accuracy using instance cues and self-correction to reduce semantic drift. Attnshift [31] advances this by extending single points to reconstruct entire objects. Apart from their complexity, these methods have yet to fully demonstrate their effectiveness, indicating ongoing challenges in harnessing single-point annotations for image segmentation and presenting clear avenues for further research.

**Prompting and Foundation Models.** Prompt-based learning enables pretrained foundation models to adapt to various tasks using well-crafted prompts. SAM [22], a prominent example in computer vision, exemplifies robust zero-

shot generalization and interactive segmentation across multiple applications. Additionally, SAM-based models like Fast-SAM [44] increases speed, HQ-SAM [20] improves segmentation quality, and Semantic-SAM [26] optimizes performance by training on diverse data granularities. Foundational models, pre-trained on large datasets, help improve generalization in downstream tasks, especially in data-scarce scenarios. Basing on SAM, Rsprompter [8] utilizes SAM-derived pseudo labels for improved remote sensing segmentation, meanwhile, adaptations for medical imaging and video tracking are explored in A-SAM [17] and Tracking Anything [43]. Further, [10] and [19] have integrated SAM with Weakly Supervised Semantic Segmentation networks to refine pseudo labels. Our research builds upon these innovations, transforming point annotations into mask proposals in instance segmentation to significantly enhancing performance.

### 3. Methodology

#### 3.1. Overview

The overview of our method is illustrated in Fig. 2, SAPNet comprises of two branches: one dedicated to the selection and refinement of mask proposals to generate pseudo-labels and the other employing solov2 head [42] for instance segmentation supervised by the generated pseudo labels. The central focus of our approach is the pseudo-label generation branch, exclusively utilized during the training phase, which includes the PSM, PNPG, and PRM modules. Following the initial proposal inputs, the PSM employs multi-instance learning and a point-distance penalty to identify semantically rich proposals. Subsequently, coupled with selected proposals from the PSM stage, the PNPG generates quality positive-negative bags to mitigate background and locality issues, emphasizing the primary regions of interest. Then, the PRM processes these bags, which selects refined proposals from positive bags to improve final box quality. Ultimately, the mask mappings derived from these box proposals are utilized to guide the segmentation branch. This guarantees the acquisition of high-quality category-specified mask proposals to supervise the segmentation branch.

#### 3.2. Proposal Selection Module

SAM’s limited semantic discernment causes category-agnostic labeling, leading to inconsistent proposal quality for the same objects. Employing these proposals directly for segmentation supervision could introduce noise and impair performance. Our goal is to design a category-specific segmenter, which needs to select the most semantically representative proposals for robust supervision.

Motivated by the insights from WSDDN [4] and P2BNet [9], our proposal selection module employs multi-instance

learning and leverages labeling information to prioritize high-confidence proposals for segmentation. In the training phase, we leverage SAM[22] solely to generate category-agnostic proposals. To avoid excessive memory use and slow training, we convert them into box proposals using the minimum bounding rectangle, and combine with depth features  $F \in \mathbb{R}^{H \times W \times D}$  from the image  $I \in \mathbb{R}^{H \times W}$ , serve as input to the PSM. Utilizing our designed MIL loss, PSM precisely predicts each proposal’s class and instance details. It selects the highest-scoring proposal as the semantically richest bounding box for each object, effectively choosing higher quality mask proposals.

Given an image  $I$  with  $N$  point annotations  $Y_n = \{(p_i, c_i)\}_{i=1}^N$ , where  $p_i$  is the coordinate of the annotated point and  $c_i$  is the class index. We transform each class-informative point  $p_i$  into  $M$  semantic mask proposals, which is further converted to a semantic proposal bag  $B_i \in \mathbb{R}^{M \times 4}$ . As illustrated in Fig. 2, after passing through a  $7 \times 7$  RoIAlign layer and two fully-connected layers, features  $F_i \in \mathbb{R}^{M \times H \times W \times D}$  are extracted from proposal bag  $B_i$ . Like in [4] and [37], the features  $F$  serve as input for the classification branch and instance branch, using fully-connected layer  $f$  and  $f'$  to generate  $\mathbf{W}_{cls} \in \mathbb{R}^{M \times K}$  and  $\mathbf{W}_{ins} \in \mathbb{R}^{M \times K}$ . A softmax activation function over  $K$  class and  $M$  instance dimensions yields the classification scores  $\mathbf{S}_{cls} \in \mathbb{R}^{M \times K}$  and instance scores  $\mathbf{S}_{ins} \in \mathbb{R}^{M \times K}$ .

$$\mathbf{W}_{cls} = f(\mathbf{F}); [\mathbf{S}_{cls}]_{mk} = e^{[\mathbf{W}_{cls}]_{mk}} / \sum_{k=1}^K e^{[\mathbf{W}_{cls}]_{mk}}.$$

$$\mathbf{W}_{ins} = f'(\mathbf{F}); [\mathbf{S}_{ins}]_{mk} = e^{[\mathbf{W}_{ins}]_{mk}} / \sum_{m=1}^M e^{[\mathbf{W}_{ins}]_{mk}}. \quad (1)$$

where  $[\cdot]_{mk}$  is the value in row  $m$  and column  $k$  of matrix.

**Point Distance Guidance.** SAM and MIL struggle with distinguishing adjacent objects of the same category, often merging two separate objects into one and giving high score. To combat this, we incorporate instance-level annotated point information and introduce a spatially aware selection with a point-distance penalty mechanism.

To address the challenge of overlapping objects and thereby enhance model optimization, we propose a strategy specifically aimed at penalizing instances of object overlap. For each  $m$ -th proposal within the set  $B_i$ , we define  $t_{mj} = 1$  to denote an overlap with any proposal in another identical class bag  $B_j$ ; otherwise,  $t_{mj} = 0$ . The penalty imposed increases in proportion to the distance of the overlapping objects from the proposal in question. This penalty,  $W_{dis}$ , is represented using the Euclidean distance between the annotated points of the overlapping proposals. Subsequently, the reciprocal of  $W_{dis}$  is then passed through a sigmoid function to compute the distance score  $\mathbf{S}_{dis}$  for the proposal.

$$[\mathbf{W}_{dis}]_{im} = \sum_{j=1, j \neq i}^N \|p_i - p_j\| * t_{mj}.$$

$$[\mathbf{S}_{dis}]_{im} = (1/e^{-(1/[\mathbf{W}_{dis}]_{im})})^d. \quad (2)$$

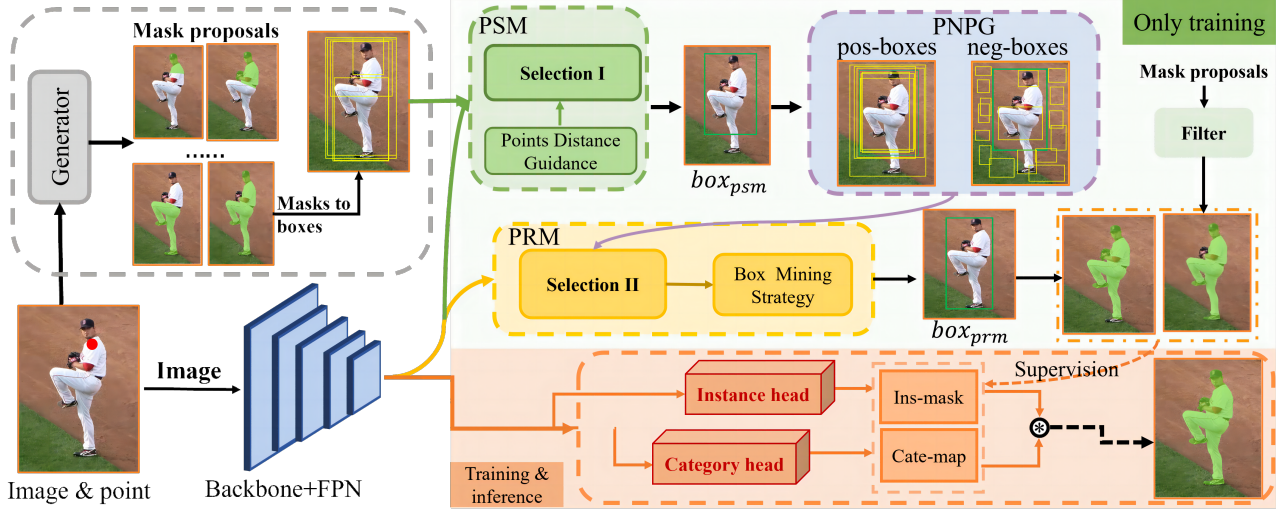


Figure 2. The framework of SAPNet comprises two components: one for generating mask proposals and another for their utilization in instance segmentation. The process starts with generating category-agnostic mask proposals using point prompts within a visual foundation model. That is followed by an initial proposal selection via MIL combined with PDG. Next, the PRM refines these proposals using positive and negative samples from PNPG, capturing global object semantics. Finally, augmented with the multi-mask proposal supervision, the segmentation branch aims to improve segmentation quality.




Figure 3. The mechanism of the proposal selection module. where  $[\cdot]_{im}$  is the value at the row  $i$  and column  $m$  in the matrix, and  $d$  is the exponential factor.

**PSM Loss.** The final score  $\mathbf{S}$  of each proposal is obtained by computing the Hadamard product of the classification score, the instance score, and the distance score, while the score  $\hat{\mathbf{S}}$  for each proposal bag  $B_i$  is obtained by summing the scores of the proposals in  $B_i$ . The MILloss of the PSM is constructed using the form of binary cross-entropy, and it is defined as follows:

$$\mathbf{S} = \mathbf{S}_{cls} \odot \mathbf{S}_{ins} \odot \mathbf{S}_{dis} \in \mathbb{R}^{M \times K}, \hat{\mathbf{S}} = \sum_{m=1}^M [\mathbf{S}]_m \in \mathbb{R}^K.$$

$$\mathcal{L}_{psm} = CE(\hat{\mathbf{S}}, \mathbf{c}) = -\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K c_k \log(\hat{\mathbf{S}}_k) + (1 - c_k) \log(1 - \hat{\mathbf{S}}_k) \quad (3)$$

where  $\mathbf{c} \in \{0, 1\}^K$  is the one-hot category's label.

Utilizing the MILloss, the PSM module skillfully identifies each proposal's category and instance. The module selects the proposal with the highest score, marked as  $\mathbf{S}$ , for a specific object and identifies a bounding box enriched with semantic information.

### 3.3. Positive and Negative Proposals Generator

To further refine the selection of more accurate bounding boxes, we employ PNPG based on  $box_{psm}$  selected via

PSM. That consists of two components: PPG and NPG. The PPG is designed to generate a richer set of positive samples, enhancing bag's quality. Concurrently, the NPG is responsible for generating negative samples, which are crucial for assisting model training. These negative samples, including background samples for all objects and part samples for each, are crucial in resolving part issues and ensuring high-quality bounding box selection. The positive sample set  $B^+$  produced by PPG and the negative sample set  $\mathcal{U}$  generated by NPG are utilized for training the subsequent PRM.

**Positive Proposals Generator (PPG).** Within this phase, to implement adaptive sampling for the identified bounding box, we capitalize on the  $box_{psm}$  derived from the PSM stage, coupled with the point distance penalty score  $\mathbf{S}_{dis}$  attributed to each proposal. To further elaborate, for each  $box_{psm}$  (denoted as  $b_x^*, b_y^*, b_w^*, b_h^*$ ) isolated during the PSM phase, its dimensions are meticulously recalibrated leveraging a scale factor  $v$  and its associated within-category inclusion score  $\mathbf{S}_{dis}$  to generate an augmented set of positive proposals  $(b_x, b_y, b_w, b_h)$ . The formulation is defined as follows:

$$\begin{aligned} b_w &= (1 \pm v/\mathbf{S}_{dis}) \cdot b_w^*, & b_h &= (1 \pm v/\mathbf{S}_{dis}) \cdot b_h^*, \\ b_x &= b_x^* \pm (b_w - b_w^*)/2, & b_y &= b_y^* \pm (b_h - b_h^*)/2. \end{aligned} \quad (4)$$

These newly cultivated positive proposals are carefully integrated into the existing set  $B_i$  to enhance the positive instances' pool. Such enhancements are pivotal in optimizing the training of the forthcoming PRM.

**Negative Proposals Generator(NPG).** MIL-based selection within a single positive bag may overemphasize the background noise, leading to inadequate focus on the object. To solve this, we create a negative bag from the back-

---

**Algorithm 1** Positive and Negative Proposals Generation

---

**Input:**  $T_{neg1}, T_{neg2}, box_{psm}$  from PSM stage, image  $I$ , positive bags  $B^+$ .

**Output:** Positive proposal bags  $B^+$ , Negative proposal set  $\mathcal{U}$ .

```
1: // Step1: positive proposals sampling
2: for  $i \in N, N$  is the number of object in image  $I$  do
3:    $B_i^+ \leftarrow B_i, B_i \in B$ ;
4:    $B_i^+ = B_i^+ \cup PPG(box_{psm}^i)$ ;
5: end for
6: // Step2: background negative proposals sampling
7:  $\mathcal{U} \leftarrow \{\}$ ;
8:  $proposals \leftarrow random\_sampling(I)$  for each image  $I$ ;
9:  $iou = IOU(proposals, B_i)$  for each  $B_i \in B$ ;
10: if  $iou < T_{neg1}$  then
11:    $\mathcal{U} = \mathcal{U} \cup proposals$ ;
12: end if
13: // Step3: part negative proposals sampling
14: for  $i \in N, N$  is the number of object in image  $I$  do
15:    $proposals \leftarrow part\_neg\_sampling(box_{psm}^i)$ ;
16:    $iou = IOU(proposals, box_{psm}^i)$ ;
17:   if  $iou < T_{neg2}$  then
18:      $\mathcal{U} = \mathcal{U} \cup proposals$ ;
19:   end if
20: end for
```

---

ground proposals post-positive bag training, which helps MIL maximize the attention towards the object.

Considering the image dimensions, we randomly sample proposals according to each image’s width and height, for negative instance sampling. We assess the Intersection over Union (IoU) between these negatives and the positive sets, filtering out those below a threshold  $T_{neg1}$ .

Additionally, to rectify MIL localization errors, we enforce the sampling of smaller proposals with an IoU under a second threshold,  $T_{neg2}$ , from inside  $box_{psm}$  based on its width and height, that is scored highest in PSM, as negative examples. These negative instances, partially capturing the object, drive the model to select high-quality bounding boxes that encompass the entire object. The PNPG is systematically elaborated upon in Algorithm 1.

### 3.4. Proposals Refinement Module

In the PSM phase, we employ MIL to select high-quality proposals from bag  $B^+$ . However, as shown in Fig. 2, the  $box_{psm}$  outcomes derived solely from a single-stage MIL are suboptimal and localized. Inspired by PCL [38], we consider refining the proposals in a second phase. However, in contrast to most WSOD methods which choose to continue refining using classification information in subsequent stages, we have established high-quality positive and negative bags, and further combined both classification and instance branches to introduce the PRM module to refine the proposals, aiming to obtain a high-quality bounding box.

The PRM module, extending beyond the scope of PSM,

focuses on both selection and refinement. It combines positive instances from the PPG with the initial set, forming an enriched  $B^+$ . Simultaneously, it incorporates the negative instance set  $\mathcal{U}$  from NPG, providing a comprehensive foundation for PRM. This integration leads to a restructured MIL loss in PRM, replacing the conventional CELoss with Focal Loss for positive instances. The modified positive loss function is as follows:

$$\mathcal{L}_{pos} = \frac{1}{N} \sum_{i=1}^N \langle \mathbf{c}_i^T, \widehat{\mathbf{S}}_i \rangle \cdot \text{FL}(\widehat{\mathbf{S}}_i^*, \mathbf{c}_i). \quad (5)$$

where FL is the focal loss [32],  $\widehat{\mathbf{S}}_i^*$  and  $\widehat{\mathbf{S}}_i$  represent the bag score predicted by PRM and PSM, respectively.  $\langle \mathbf{c}_i^T, \widehat{\mathbf{S}}_i \rangle$  represents the inner product of the two vectors, meaning the predicted bag score of the ground-truth category.

Enhancing background suppression, we use negative proposals and introduce a dedicated loss for these instances. Notably, these negative instances pass only through the classification branch for instance score computation, with their scores derived exclusively from classification. The specific formulation of this loss function is detailed below:

$$\beta = \frac{1}{N} \sum_{i=1}^N \langle \mathbf{c}_i^T, \widehat{\mathbf{S}}_i \rangle, \quad (6)$$

$$\mathcal{L}_{neg} = -\frac{1}{|\mathcal{U}|} \sum_{\mathcal{U}} \sum_{k=1}^K \beta \cdot ([\mathbf{S}_{neg}^{cls}]_k)^2 \log(1 - [\mathbf{S}_{neg}^{cls}]_k). \quad (7)$$

The PRM loss consists of the MIL loss  $\mathcal{L}_{pos}$  for positive bags and negative loss  $\mathcal{L}_{neg}$  for negative samples, *i.e.*,

$$\mathcal{L}_{prm} = \alpha \mathcal{L}_{pos} + (1 - \alpha) \mathcal{L}_{neg}, \quad (8)$$

where  $\alpha = 0.25$  by default.

**Box Mining Strategy.** MIL’s preference for segments with more foreground presence and SAM’s tendency to capture only parts of an object often bring to final bounding boxes,  $box_{prm}$ , the ‘local’ issue of MIL inadequately covers the instances. To improve the bounding box quality, we introduce a box mining strategy that adaptively expands  $box_{select}$  from proposal selection in PRM, by merging it with the original proposals filter, aiming to address MIL’s localization challenges.

The Box Mining Strategy (BMS) consists of two primary components: (i) We select the top  $k$  proposals from the positive proposal bag  $B^+$ , to create a set  $G$ . We evaluate the proposals in  $G$  against  $box_{select}$  based on IoU and size, using a threshold  $T_{min1}$ . Proposals larger than  $box_{select}$  and with an IoU above  $T_{min1}$  undergo dynamic expansion through IoU consideration, which allows for the adaptive integration with  $box_{select}$ . That mitigates the ‘local’ issue and maintains the bounding box’s consistency to the object’s true boundaries. (ii) Frequently, issues related to lo-

cality can lead to an exceedingly low IoU between proposals and  $box_{select}$ . Nonetheless, the ground truth box can fully encompass the  $box_{part}$ . Therefore, when component (i) conditions are unmet, if a proposal can entirely encapsulate  $box_{select}$ , we reset the threshold  $T_{min2}$ . Proposals surpassing this threshold adaptively merge with  $box_{select}$  to generate the final  $box_{prm}$ , used to yield  $Mask_{prm}$ . These two components collectively form our BMS strategy. A detailed procedure of this approach will be delineated in Algorithm 2 of the supplementary materials.

**Loss Function.** After acquiring the final supervision masks,  $Mask_{prm}$  and the filtered  $Mask_{sam}$  in Multi-mask Proposals Supervision (MPS) in Sec. 7 of supplementary, we use them together to guide the dynamic segmentation branch. To comprehensively train SAPNet, we integrate the loss functions from the PSM and PRM, culminating in the formulation of the total loss for our model, denoted as  $L_{total}$ . The aggregate loss function,  $L_{total}$  can be articulated as:

$$\mathcal{L}_{total} = \mathcal{L}_{mask} + \mathcal{L}_{cls} + \lambda \cdot \mathcal{L}_{psm} + \mathcal{L}_{prm} \quad (9)$$

where,  $\mathcal{L}_{Dice}$  is the Dice Loss [35],  $\mathcal{L}_{cls}$  is the Focal Loss [32], and  $\lambda$  is set as 0.25.

## 4. Experiment

### 4.1. Experimental Settings

**Datasets.** We use the publicly available MS COCO [33] and VOC2012SBD [13] datasets for experiments. COCO17 has 118k training and 5k validation images with 80 common object categories. VOC consists of 20 categories and contains 10,582 images for model training and 1,449 validation images for evaluation.

**Evaluation Metric.** We use mean average precision  $mAP@[.5,.95]$  for the MS-COCO. The  $\{AP, AP_{50}, AP_{75}, AP_{Small}, AP_{Middle}, AP_{Large}\}$  is reported for MS-COCO and for VOC12SBD segmentation, and we report  $AP_{25,50,75}$ . The  $mIoU_{box}$  is the average IoU between predicted pseudo-boxes and GT-boxes in the training set. It measures SAPNet’s ability to select mask proposals without using the segmentation branch.

**Implementation Details.** In our study, we employed the Stochastic Gradient Descent (SGD) optimizer, as detailed in [6]. Our experiments were conducted using the mmdetection toolbox [7], following standard training protocols for each dataset. We used the ResNet architecture [15], pretrained on ImageNet [36], as the backbone. For COCO, batch size was set at four images per GPU across eight GPUs, and for VOC2012, it was four GPUs. More details of the experiment are in Sec. 8 of the supplementary.

### 4.2. Experimental Comparisons

Tab. 1 shows the comparison results between our method and previous SOTA approaches [11, 16, 34, 40, 42] on

COCO. In our experiments, we provide SAM with both the labeled points and the annotations generated by the point annotation enhancer [9]. SAM then utilizes these inputs to generate subsequent mask proposals for selection and supervision. For fair comparison, we design two baselines: the top-1 scored mask from SAM and MIL-selected SAM mask proposals are used as SOLOv2 supervision, respectively. Tab. 1 shows our method substantially surpasses these baselines in performance.

**Comparison with point-annotated methods.** Our approach achieves a 31.2  $AP$  performance with a ResNet-50 backbone, surpassing all previous point-annotated methods, including BESTIE on HRNet-48 and AttnShift on Vit-B. Our model exhibits significant improvements under a 1x training schedule, with a 13.5  $AP$  increase when compared to the previous SOTA method, BESTIE. Furthermore, under a 3x training schedule, SAPNet outperforms AttnShift, which relies on large model training, with 13.4  $AP$ , improvements. Importantly, our method is trained end-to-end without needing post-processing, achieving SOTA performance in point-annotated instance segmentation.

**Comparison with other annotation-based methods.** Our SAPNet has significantly elevated point annotation, regardless of point annotation’s limitations in annotation time and quality compared to box annotation. Utilizing a ResNet-101 backbone and a 3x training schedule, SAPNet surpasses most box-annotated instance segmentation methods, achieving a 1.4  $AP$  improvement over BoxInst. Moreover, SAPNet’s segmentation performance nearly matches the mask-annotated methods, effectively bridging the gap between point-annotated and these techniques.

**Segmentation performance on VOC2012SBD.** Tab. 2 compares segmentation methods under different supervisions on the VOC2012 dataset. SAPNet reports an enhancement of 7.7  $AP$  over the AttnShift approach, evidencing a notable advancement in performance. Thereby, it significantly outstrips image-level supervised segmentation methods. Additionally, SAPNet surpasses box-annotated segmentation methods, such as BoxInst by 3.4  $AP_{50}$  and DiscoBox by 32.6  $AP_{50}$ . Further, our point-prompted method achieves 92.3% of the Mask-R-CNN.

### 4.3. Ablation Studies

More experiments have been conducted on COCO to further analyze SAPNet’s effectiveness and robustness.

**Training Stage in SAPNet.** The ablation study of the training stage is given in Tab. 3. We trained solov2 using the top-1 scored mask provided by SAM and compared it to the two training strategies of SAPNet. In the two-stage approach, the segmentation branch and multiple-mask supervision of SAPNet are removed. Instead, we use the selected mask to train a standalone instance segmentation model, as described by [42]. The end-to-end training method corre-

Method	Ann.	Backbone	sched.	Arch.	mAP	mAP <sub>50</sub>	mAP <sub>75</sub>	mAP <sub>s</sub>	mAP <sub>m</sub>	mAP <sub>l</sub>
<b>Fully-supervised instance segmentation models.</b>										
Mask R-CNN [16]	$\mathcal{M}$	ResNet-50	1x	Mask R-CNN	34.6	56.5	36.6	18.3	37.4	47.2
YOLOACT-700 [5]	$\mathcal{M}$	ResNet-101	4.5x	YOLOACT	31.2	54.0	32.8	12.1	33.3	47.
PolarMask [16]	$\mathcal{M}$	ResNet-101	2x	PolarMask	32.1	53.7	33.1	14.7	33.8	45.3
SOLOv2 [42]	$\mathcal{M}$	ResNet-50	1x	SOLOv2	34.8	54.9	36.9	13.4	37.8	53.7
CondInst [40]	$\mathcal{M}$	ResNet-50	1x	CondInst	35.3	56.4	37.4	18.0	39.4	50.4
SwinMR [34]	$\mathcal{M}$	Swin-S	50e	SwinMR	43.2	67.0	46.1	24.8	46.3	62.1
Mask2Former [11]	$\mathcal{M}$	Swin-S	50e	Mask2Former	46.1	69.4	52.8	25.4	49.7	68.5
<b>Weakly-supervised instance segmentation models.</b>										
IRNet [45]	$\mathcal{I}$	ResNet-50	1x	Mask R-CNN	6.1	11.7	5.5	-	-	-
BESTIE [21]	$\mathcal{I}$	HRNet-48	1x	Mask R-CNN	14.3	28.0	13.2	-	-	-
BBTP [18]	$\mathcal{B}$	ResNet-101	1x	Mask R-CNN	21.1	45.5	17.2	11.2	22.0	29.8
BoxInst [39]	$\mathcal{B}$	ResNet-101	3x	CondInst	33.2	56.5	33.6	16.2	35.3	45.1
DiscoBox [23]	$\mathcal{B}$	ResNet-50	3x	SOLOv2	32.0	53.6	32.6	11.7	33.7	48.4
Boxlevelset [27]	$\mathcal{B}$	ResNet-101	3x	SOLOv2	33.4	56.8	34.1	15.2	36.8	46.8
WISE-Net [24]	$\mathcal{P}$	ResNet-50	1x	Mask R-CNN	7.8	18.2	8.8	-	-	-
BESTIE <sup>†</sup> [21]	$\mathcal{P}$	HRNet-48	1x	Mask R-CNN	17.7	34.0	16.4	-	-	-
AttnShift [31]	$\mathcal{P}$	Vit-B	50e	Mask R-CNN	21.2	43.5	19.4	-	-	-
SAM-SOLOv2	$\mathcal{P}$	ResNet-50	1x	SOLOv2	24.6	41.9	25.3	9.3	28.6	38.1
MIL-SOLOv2	$\mathcal{P}$	ResNet-50	1x	SOLOv2	26.8	47.7	26.8	11.2	31.5	40.4
<b>SAPNet(ours)</b>	$\mathcal{P}$	ResNet-50	1x	SOLOv2	<b>31.2</b>	<b>51.8</b>	<b>32.3</b>	<b>12.6</b>	<b>35.1</b>	<b>47.8</b>
<b>SAPNet(ours)*</b>	$\mathcal{P}$	ResNet-101	3x	SOLOv2	<b>34.6</b>	<b>56.0</b>	<b>36.6</b>	<b>15.7</b>	<b>39.5</b>	<b>52.1</b>

Table 1. Mask annotation( $\mathcal{M}$ ), image annotation( $\mathcal{I}$ ), box annotation( $\mathcal{B}$ ) and point annotation( $\mathcal{P}$ ) performance on COCO-17 val. ‘Ann.’ is the type of the annotation and ‘sched.’ means schedule. \* is the multi-scale augment training for re-training segmentation methods, and other experiments are on single-scale training. SwinMR is Swin-Transformer-Mask R-CNN. SwinMR and Mask2Former use multi-scale data augment strategies for SOTA.

Method	Sup.	Backbone	$AP_{25}$	$AP_{50}$	$AP_{75}$
Mask R-CNN [16]	$\mathcal{M}$	R-50	78.0	68.8	43.3
Mask R-CNN [16]	$\mathcal{M}$	R-101	79.6	70.2	45.3
BoxInst [39]	$\mathcal{B}$	R-101	-	61.4	37.0
DiscoBox	$\mathcal{B}$	R-101	72.8	62.2	37.5
BESTIE [21]	$\mathcal{I}$	HRNet	53.5	41.7	24.2
IRNet [45]	$\mathcal{I}$	R-50	-	46.7	23.5
BESTIE <sup>†</sup> [21]	$\mathcal{I}$	HRNet	61.2	51.0	26.6
WISE-Net [24]	$\mathcal{P}$	R-50	53.5	43.0	25.9
BESTIE [21]	$\mathcal{P}$	HRNet	58.6	46.7	26.3
BESTIE <sup>†</sup> [21]	$\mathcal{P}$	HRNet	66.4	56.1	30.2
Attnshift [31]	$\mathcal{P}$	Vit-S	68.3	54.4	25.4
Attnshift <sup>†</sup> [31]	$\mathcal{P}$	Vit-S	70.3	57.1	30.4
<b>SAPNet(ours)</b>	$\mathcal{P}$	R-101	<b>76.5</b>	<b>64.8</b>	<b>58.7</b>

Table 2. Instance segmentation performance on the VOC2012 test set. <sup>†</sup> indicates applying MRCNN refinement.

sponds to the architecture illustrated in Fig. 2. Our findings indicate that our method is more competitive than directly employing SAM (31.2  $AP$  vs 24.6  $AP$ ), and the visualization of Fig. 4 shows us this enhancement. Moreover, the end-to-end training strategy boasts a more elegant model structure and outperforms the two-stage approach in overall efficiency (31.2  $AP$  vs 30.18  $AP$ ).

**Effect of Each Component.** Given the limited performance of SAM-top1, we opted for the single-MIL as our baseline. With a preliminary selection using MIL1, we




Figure 4. The comparative visualization between SAM-top1 and SAPNet is presented, showcasing SAM’s segmentation outcomes in green masks and our results in yellow. The orange and red bounding boxes highlight the respective mask boundaries.

train stage on coco	sched.	$AP$	$AP_{50}$	$AP_{75}$
SAM-top1	1x	24.6	41.9	25.3
Two stage	1x	30.2	49.8	31.5
End to end	1x	31.2	51.8	32.3

Table 3. The experimental comparisons of segmenters in COCO dataset, SAM-top1 is the highest scoring mask generated by SAM.

have achieved a segmentation performance of 26.8  $AP$ . **i) Point Distance Guidance.** We updated the proposal scores from the existing MIL by integrating the PDG module into the foundational MIL selection. This approach successfully segments adjacent objects of the same category, improving the segmentation performance by 0.7 points (27.5 vs 26.8).

mil1	PDG	mil2	PNPG	BMS	MPS	$mAP$
✓						26.8
✓	✓					27.5
✓	✓	✓				27.7
✓	✓	✓	✓			29.7
✓	✓	✓	✓	✓		30.8
✓	✓	✓	✓	✓	✓	31.2

Table 4. The effect of each component in SAPNet: proposal selection module(MIL1), point distance guidance(PDG), positive and negative proposals generator(PNPG), proposal selection module(MIL2), box mining strategy(BMS), and Multi-mask Proposals Supervision(MPS) in Sec. 7 of supplementary.

**ii) MIL2.** Building on the previous step, we incorporate a second MIL selection module to refine the initially selected boxes, resulting in a performance increment of 0.2 points. **iii) PNPG.** For MIL2, we devised the positive-negative sample sets, aiming to enhance the input quality for the PRM module and use the negative samples to suppress background. This adjustment leads to a segmentation performance boost of 2 points (29.7 vs 27.7). **iv) BMS.** Within the PRM, we refine the selected boxes using BMS, pushing the segmentation performance up by 1.1 points (30.8 vs 29.7). **v) MPS.** Utilizing MPS for segmentation branch supervision yields a 0.4-point performance improvement.

**Threshold of BMS.** For point refinement, there are two constraints (described in Sec. 3.4).  $T_{min1}$  and  $T_{min2}$  are thresholds of the Box Mining Strategy. In Tab. 5, it shows that the two constraints together to obtain performance gain. After multiple experiments, we have found that there is a significant performance improvement when  $T_{min1}$  and  $T_{min2}$  are set to 0.6 and 0.3, respectively.

**Components of PNPG.** Tab. 6 presents the results of a dissected ablation study on the Positive and Negative Proposals Generator(PNPG), illustrating the respective impacts of the positive and negative examples on the model’s performance. It is evident that the construction of negative examples plays a significant role in enhancing model efficacy. Furthermore, the beneficial effects of both positive and negative examples are observed to be cumulative.

**Performance Analysis.** As presented in Tab. 7, we conducted a statistical analysis to validate SAPNet’s capability to address ‘local’ issue and compare the outcomes selected by the single-MIL with those obtained by SAPNet in the absence of segmentation branch integration. Specifically, the part problem generated by the single-MIL, where MIL is inclined to select proposals with a higher proportion of foreground, is exemplified in Fig. 6 of supplementary. On this premise, we initially establish an evaluative criterion  $R_v = \frac{area_{mask}}{area_{box}}$ , which is the ratio of the mask area to the bounding box area. Subsequently, we compute  $R_{v_i}$  for each proposal within the proposal bag corresponding to every instance across the entire COCO dataset and select the maximum  $R_{v_{max}}$  to compute the mean value over the dataset,

$T_{min1}$	$T_{min2}$	$AP$	$AP_{50}$	$AP_{75}$	$AP^s$	$AP^m$	$AP^l$
0.5	0.3	30.9	51.3	32.0	12.2	34.7	47.4
0.5	0.4	30.7	51.2	31.8	11.9	34.7	47.1
0.6	0.3	<b>31.2</b>	<b>51.8</b>	<b>32.3</b>	<b>12.6</b>	<b>35.1</b>	<b>47.8</b>
0.6	0.4	30.8	51.1	32.0	12.1	34.7	47.3
0.7	0.3	31.0	51.5	32.2	12.6	34.9	47.3
0.7	0.4	30.7	51.1	31.9	12.0	34.6	47.2

Table 5. Constraints in box mining strategy.

PNPG		$AP$	$AP_{50}$	$AP_{75}$
$PPG$	$NPG$			
		29.3	49.7	30.0
✓		29.8	50.5	30.8
	✓	30.7	51.2	31.7
✓	✓	<b>31.2</b>	<b>51.8</b>	<b>32.3</b>

Table 6. Meticulous ablation experiments in PNPG

Method	$Gap$	$mIoU_{box}$
Single-MIL	0.199	63.8
SAPNet	0.131	69.1

Table 7. Experimental analysis with part problem.

which is then designated as the threshold  $T_{rv}$ . Ultimately, we identify the ground truth  $R_{v_{gt}}$  and objects where  $R_{v_{max}}$  exceeds  $T_{rv}$  and calculates the discrepancy between  $R_v$  values selected by single-MIL and SAPNet. The description is as follows:

$$Gap_{single} = R_{v_{single}} - R_{v_{gt}}, \quad Gap_{our} = R_{v_{our}} - R_{v_{gt}}. \quad (10)$$

Tab. 7 shows that the proposed SAPNet mitigates the locality issue faced by the single-MIL. Furthermore, the boxes selected via SAPNet exhibit a substantially higher IoU with GT than those selected by the single-MIL.

## 5. Conclusion

In this paper, we propose SAPNet, an innovative end-to-end point-prompted instance segmentation framework. SAPNet transforms point annotations into category-agnostic mask proposals and employs dual selection branches to elect the most semantic mask for each object, guiding the segmentation process. To address challenges such as indistinguishable adjacent objects of the same class and MIL’s locality bias, we integrate PDG and PNPG, complemented by a Box Mining Strategy for enhanced proposal refinement. SAPNet uniquely merges segmentation and selection branches under multi-mask supervision, significantly enhancing its segmentation performance. Extensive experimental comparisons on VOC and COCO datasets validate the SAPNet’s effectiveness in point-prompted instance segmentation.

## 6. Acknowledgements

This work was supported in part by the Youth Innovation Promotion Association CAS, the National Natural Science Foundation of China (NSFC) under Grant No. 61836012, 61771447 and 62272438, and the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant No.XDA27000000.



## References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2209–2218, 2019. 2
- [2] Pablo Andrés Arbeláez, Jordi Pont-Tuset, and Jonathan T. Barron *et al.* Multiscale combinatorial grouping. In *CVPR*, 2014. 2
- [3] Aditya Arun, CV Jawahar, and M Pawan Kumar. Weakly supervised instance segmentation by learning annotation consistent instances. In *European Conference on Computer Vision*, pages 254–270. Springer, 2020. 2
- [4] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *CVPR*, 2016. 2, 3
- [5] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact: Real-time instance segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9157–9166, 2019. 1, 7
- [6] Léon Bottou. Stochastic gradient descent tricks. In *Neural Networks: Tricks of the Trade: Second Edition*, pages 421–436. Springer, 2012. 6
- [7] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, *et al.* Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. <https://github.com/open-mmlab/mmdetection>. 6
- [8] Keyan Chen, Chenyang Liu, Hao Chen, Haotian Zhang, Wenyuan Li, Zhengxia Zou, and Zhenwei Shi. Rsprompter: Learning to prompt for remote sensing instance segmentation based on visual foundation model. *arXiv preprint arXiv:2306.16269*, 2023. 3
- [9] Pengfei Chen, Xuehui Yu, Xumeng Han, Najmul Hassan, Kai Wang, Jiachen Li, Jian Zhao, Humphrey Shi, Zhenjun Han, and Qixiang Ye. Point-to-box network for accurate object detection via single point supervision. In *European Conference on Computer Vision*, pages 51–67. Springer, 2022. 2, 3, 6
- [10] Tianle Chen, Zheda Mai, Ruiwen Li, and Wei-lun Chao. Segment anything model (sam) enhanced pseudo labels for weakly supervised semantic segmentation. *arXiv preprint arXiv:2305.05803*, 2023. 3
- [11] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2022. 1, 6, 7
- [12] Bowen Cheng, Omkar Parkhi, and Alexander Kirillov. Pointly-supervised instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2617–2626, 2022. 2
- [13] Mark Everingham, Luc Van Gool, and Christopher K. I. Williams *et al.* The pascal visual object classes (VOC) challenge. *IJCV*, 2010. <http://host.robots.ox.ac.uk/pascal/VOC/>. 6
- [14] Junsong Fan, Zhaoxiang Zhang, and Tieniu Tan. Pointly-supervised panoptic segmentation. In *European Conference on Computer Vision*, pages 319–336. Springer, 2022. 1
- [15] Kaiming He, Xiangyu Zhang, and Shaoqing Ren *et al.* Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [16] Kaiming He, Georgia Gkioxari, and Piotr Dollár *et al.* Mask R-CNN. In *ICCV*, 2017. 1, 6, 7
- [17] Sheng He, Rina Bao, Jingpeng Li, P Ellen Grant, and Yangming Ou. Accuracy of segment-anything model (sam) in medical image segmentation tasks. *arXiv preprint arXiv:2304.09324*, 2023. 3
- [18] 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*, 2019. 2, 7
- [19] Peng-Tao Jiang and Yuqi Yang. Segment anything is a good pseudo-label generator for weakly supervised semantic segmentation. *arXiv preprint arXiv:2305.01275*, 2023. 3
- [20] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality. *arXiv preprint arXiv:2306.01567*, 2023. 1, 3
- [21] Beomyoung Kim, Youngjoon Yoo, Chaeun Rhee, and Junmo Kim. Beyond semantic to instance segmentation: Weakly-supervised instance segmentation via semantic knowledge transfer and self-refinement. In *CVPR*, 2022. 1, 2, 7
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, *et al.* Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. <https://segment-anything.com/>. 1, 2, 3
- [23] Shiyi Lan, Zhiding Yu, Christopher 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3406–3416, 2021. 1, 2, 7
- [24] Issam H. Laradji, Negar Rostamzadeh, Pedro O. Pinheiro, David Vázquez, and Mark Schmidt. Proposal-based instance segmentation with point supervision. In *ICIP*, 2020. 2, 7
- [25] Jungbeom Lee, Jihun Yi, Chaehun Shin, and Sungroh Yoon. Bbam: Bounding box attribution map for weakly supervised semantic and instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2643–2652, 2021. 2
- [26] Feng Li, Hao Zhang, Peize Sun, Xueyan Zou, Shilong Liu, Jianwei Yang, Chunyuan Li, Lei Zhang, and Jianfeng Gao. Semantic-sam: Segment and recognize anything at any granularity. *arXiv preprint arXiv:2307.04767*, 2023. 3
- [27] Wentong Li, Wenyu Liu, Jianke Zhu, Miaomiao Cui, Xian-Sheng Hua, and Lei Zhang. Box-supervised instance segmentation with level set evolution. In *European conference on computer vision*, pages 1–18. Springer, 2022. 1, 7
- [28] Wentong Li, Yuqian Yuan, Song Wang, Jianke Zhu, Jianshu Li, Jian Liu, and Lei Zhang. Point2mask: Point-supervised panoptic segmentation via optimal transport. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 572–581, 2023. 1, 2

- [29] Wentong Li, Wenyu Liu, Jianke Zhu, Miaomiao Cui, Risheng Yu Xiansheng Hua, and Lei Zhang. Box2mask: Box-supervised instance segmentation via level-set evolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 2
- [30] Wentong Li, Yuqian Yuan, Song Wang, Wenyu Liu, Dongqi Tang, Jianke Zhu, Lei Zhang, et al. Label-efficient segmentation via affinity propagation. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [31] Mingxiang Liao, Zonghao Guo, , and Yuze Wang *et al.* Attentionshift: Iteratively estimated part-based attention map for pointly supervised instance segmentation. In *CVPR*, 2023. 2, 7
- [32] Tsung-Yi Lin, Priya Goyal, and Ross B. Girshick *et al.* Focal loss for dense object detection. In *ICCV*, 2017. 5, 6
- [33] Tsung-Yi Lin, Michael Maire, and Serge *et al.* Belongie. Microsoft coco: Common objects in context. In *ECCV*, 2014. <https://cocodataset.org/>. 6
- [34] Ze Liu, Yutong Lin, and Yue Cao *et al.* Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 6, 7
- [35] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016. 6
- [36] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015. 6
- [37] Peng Tang and Xinggang Wang *et al.* Multiple instance detection network with online instance classifier refinement. In *CVPR*, 2017. 3
- [38] Peng Tang, Xinggang Wang, and Song Bai *et al.* PCL: proposal cluster learning for weakly supervised object detection. *IEEE TPAMI*, 2020. 5
- [39] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. Boxinst: High-performance instance segmentation with box annotations. In *CVPR*, 2021. 1, 2, 7
- [40] Zhi Tian, Bowen Zhang, Hao Chen, and Chunhua Shen. Instance and panoptic segmentation using conditional convolutions. *IEEE TPAMI*, 2023. 1, 2, 6, 7
- [41] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. SOLO: segmenting objects by locations. In *ECCV*, 2020.
- [42] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020. <https://github.com/WXinlong/SOLO>. 1, 2, 3, 6, 7
- [43] Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos. *arXiv preprint arXiv:2304.11968*, 2023. 3
- [44] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything. *arXiv preprint arXiv:2306.12156*, 2023. 1, 3
- [45] Yanning Zhou, Hao Chen, Jiaqi Xu, Qi Dou, and Pheng-Ann Heng. Innet: Instance relation network for overlapping cervical cell segmentation. In *MICCAI*, 2019. 1, 7