

# S<sup>4</sup>M: Boosting Semi-Supervised Instance Segmentation with SAM

Heeji Yoon<sup>1\*</sup> Heeseong Shin<sup>1\*</sup> Eunbeen Hong<sup>1</sup> Hyunwook Choi<sup>2</sup>  
 Hansang Cho<sup>3</sup> Daun Jeong<sup>3</sup> Seungryong Kim<sup>1†</sup>

<sup>1</sup>KAIST AI <sup>2</sup>Korea University <sup>3</sup>Samsung Electro-Mechanics

<https://cvlab-kaist.github.io/S4M>

## Abstract

*Semi-supervised instance segmentation poses challenges due to limited labeled data, causing difficulties in accurately localizing distinct object instances. Current teacher-student frameworks still suffer from performance constraints due to unreliable pseudo-label quality stemming from limited labeled data. While the Segment Anything Model (SAM) offers robust segmentation capabilities at various granularities, directly applying SAM to this task introduces challenges such as class-agnostic predictions and potential over-segmentation. To address these complexities, we carefully integrate SAM into the semi-supervised instance segmentation framework, developing a novel distillation method that effectively captures the precise localization capabilities of SAM without compromising semantic recognition. Furthermore, we incorporate pseudo-label refinement as well as a specialized data augmentation with the refined pseudo-labels, resulting in superior performance. We establish state-of-the-art performance, and provide comprehensive experiments and ablation studies to validate the effectiveness of our proposed approach.*

## 1. Introduction

Instance segmentation—simultaneously detecting objects and delineating their pixel-level boundaries—is fundamental to applications such as autonomous driving and medical imaging [62, 64]. Although fully-supervised methods [6, 7, 19, 24] have achieved impressive accuracy, their dependence on extensive annotated datasets limits scalability due to the labor-intensive nature of pixel-level labeling.

Consequently, recent work [3, 5, 16, 28, 54] has explored semi-supervised learning (SSL) approaches that additionally leverage unlabeled images. Typically, these methods generate pseudo-labels using a teacher network trained on a limited labeled dataset, which subsequently guides the

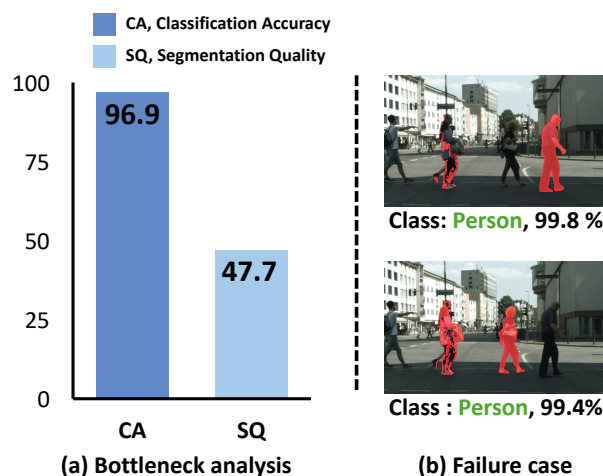


Figure 1. **Analysis on pseudo-labels by the teacher in a teacher-student framework for semi-supervised instance segmentation.** (a) Bottleneck analysis revealing that the primary limitation lies in mask quality rather than classification. Note that class accuracy (CA) is computed on matched pairs with IoU > 0.5, and segmentation quality (SQ) is measured by the standard segmentation quality metric from panoptic quality [32]. (b) Example failure cases with correct, confident class prediction but inaccurate masks.

training of a student network. However, the restricted availability of labeled data renders the pseudo-labels error-prone, thereby reducing the benefits of incorporating unlabeled images in the teacher-student framework.

To address this, recent studies have proposed several strategies, including noise filtering [54], the use of auxiliary information (e.g., depth maps [5]), and dedicated training stages to stabilize the teacher-student framework [3, 16]. Despite these advances, the inherent limitations imposed by the small-scale labeled data continue to produce pseudo-labels with significant errors, ultimately impeding the performance of the student network.

Recently, large-scale vision foundation models [20, 40, 42, 55], pretrained on web-scale datasets, have demon-

\*These authors contributed equally.

†Corresponding author.

strated exceptional performance across diverse tasks and exhibit strong generalization without task-specific fine-tuning [1, 2, 63]. In particular, the Segment Anything Model (SAM) [33, 45] has gathered significant attention as a prompt-driven segmentation foundation model, capable of predicting fine-grained masks at any granularity, from whole objects to parts and sub-parts, using geometric prompts such as points and bounding boxes. Trained on an unprecedented scale of images, SAM demonstrates efficacy and generalization across various domains [4, 29, 46] and in diverse tasks [44, 57].

These advancements motivate our exploration of SAM to enhance instance segmentation through knowledge distillation, pseudo-label enhancement, and data augmentation, which are fundamental elements in semi-supervised learning. However, it still faces challenges in directly applying SAM to instance segmentation due to its class-agnostic design [33, 52]: instance segmentation inherently requires both mask and class predictions, yet SAM is not designed to generate class-conditioned masks.

In this work, we integrate SAM into a semi-supervised instance segmentation framework to address challenges posed by limited labeled data. We first examine the deficiencies of existing semi-supervised approaches [3] by visualizing pseudo-labels generated by their teacher networks, thereby establishing a basis for incorporating SAM. As illustrated in Fig. 1, these networks reliably identify classes, but often fail in precise localization by grouping multiple instances into a single mask, resulting in *under-segmentation*. Although it might appear straightforward to use SAM to separate instances within these pseudo-labels, its class-agnostic design frequently produces masks that capture only fine-grained segments of an object rather than the object as a whole, resulting in *over-segmentation* [49]. This underscores the necessity for carefully balancing under- and over-segmentation when integrating SAM into the semi-supervised instance segmentation framework.

In this regard, we argue that it is crucial to identify what and what not to learn from SAM for tackling the under- and over-segmentation problem, and propose a novel framework for distilling SAM to improve the teacher and student networks. In specific, we first improve the teacher network trained on a small amount of label data by a novel knowledge distillation objective, which can effectively acquire the fine-grained localization capabilities of SAM while avoiding over-segmentation or hindering semantic recognition. We further enhance our framework by propagating the strong segmentation capability of SAM through pseudo-label refinement and an augmentation strategy designed for instance segmentation.

Our framework for boosting Semi-Supervised Instance Segmentation with SAM, or S<sup>4</sup>M, establishes state-of-the-art performance in all benchmarks, demonstrating the ef-

fectiveness of our approach. We further provide detailed ablations and analysis on our methodology, as well as qualitative comparison with baselines.

Our main contributions can be summarized as follows:

- We incorporate Segment Anything Model (SAM) into the semi-supervised instance segmentation framework, and present our explorations for improving the teacher-student framework with SAM.
- We carefully design our framework to fully leverage SAM through structural distillation, pseudo-label refinement and data augmentation while avoiding potential drawbacks of directly adopting SAM.
- We establish state-of-the-art performance across benchmarks, and provide thorough ablations and analysis to validate our approach.

## 2. Related Work

**Semi-supervised instance segmentation.** Dominant approaches to semi-supervised instance segmentation [3, 5, 16, 28, 53] have been based on student-teacher pseudo-labeling, where teacher model generates pseudo-labels for unlabeled images, which are then used by student model for training. In this framework, applying weak and strong data augmentations for generating and predicting pseudo-labels, respectively, effectively utilizes unlabeled data, as proposed in FixMatch [50]. The semi-supervised instance segmentation task was first introduced by Noisy Boundaries [53], which proposed a noise-tolerant mask head to filter noisy student predictions. Polite teacher [16] employs EMA teacher while filtering out pseudo-labels by confidence thresholding. PAIS [28] introduced a dynamically changing loss weight based on the quality of pseudo-labels, improving the utilization of unlabeled data by retaining low-confidence labels. More recently, GuidedDistillation [3] proposed an improved teacher-student distillation with a guided burn-in stage and integrated a vision transformer-based architecture into a semi-supervised instance segmentation framework, achieving substantial performance gains. Building on [3], Depth-Guided [5] integrates a depth foundation model for depth-aware cues. Yet both still struggle to produce noise-tolerant masks due to scarce labels. Here, we present the first semi-supervised instance segmentation framework to incorporate SAM, fully leveraging its generalization power.

**Segment Anything Model.** SAM [33, 45] is a foundation model for zero-shot image segmentation that generalizes across domains without task-specific fine-tuning. Supporting interactive prompts—points, boxes, or masks—it enables flexible, instance-aware segmentation in diverse scenarios. Building on this capability, several studies have explored its use in semi- [37, 56, 61] or weakly [34] supervised semantic segmentation. In this work, we present

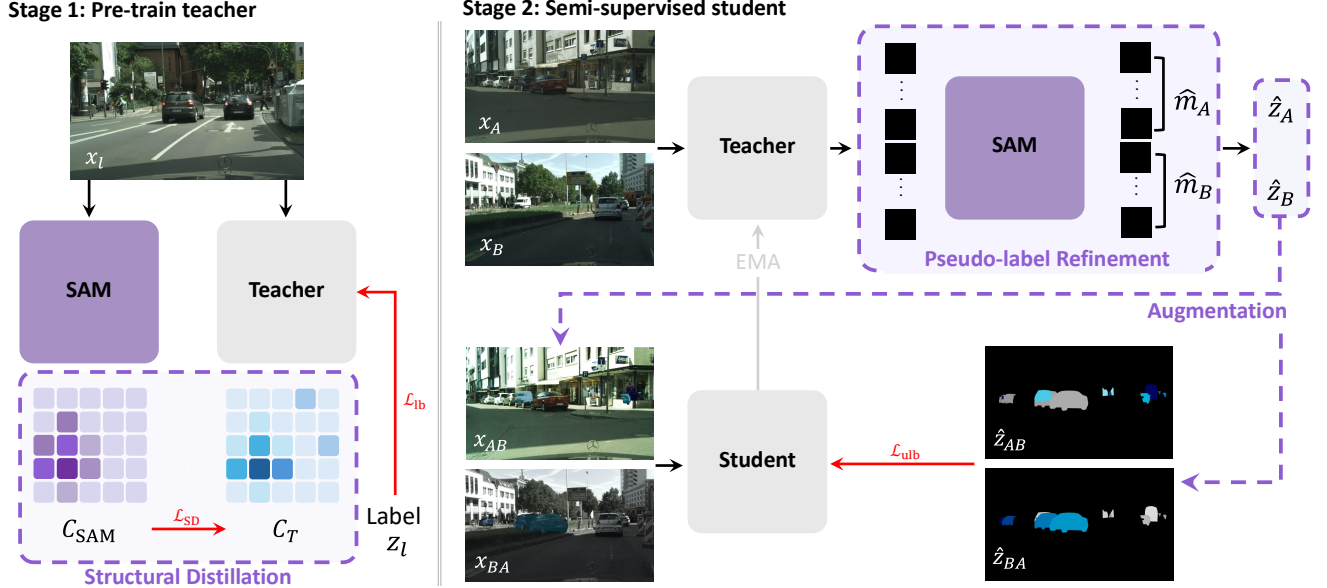


Figure 2. **Overall pipeline of the proposed framework,  $S^4M$ .** We propose  $S^4M$ , a semi-supervised instance segmentation framework that leverages SAM knowledge via three key components. First, we enhance the teacher network with structural distillation to transfer SAM’s spatial understanding. The student then learns from unlabeled data using SAM-refined pseudo-labels, further enhanced by augmentation with refined pseudo-labels.

the first approach to integrate SAM into semi-supervised instance segmentation through a carefully designed framework that addresses its potential limitations.

**Knowledge distillation.** Knowledge distillation (KD) is a widely known technique for transferring knowledge from a teacher to a student model [22]. Although conventional approaches in KD aim to train the student to mimic the teacher’s output, relational knowledge distillation (RKD) [41] introduces an effective approach to distill structural relations among learned representations. In the domain of semantic segmentation, several studies [21, 31, 36] investigated the effectiveness KD and RKD. In this work, we explore transferring the generalized structural knowledge embedded in the vision foundation model, SAM, into an instance segmentation model through knowledge distillation within a semi-supervised framework.

### 3. Preliminaries

**Problem formulation.** In semi-supervised instance segmentation, we leverage a large set of unlabeled data  $\mathcal{D}_U = \{x_u\}$  and a small labeled data set  $\mathcal{D}_L = \{(x_l, z_l)\}$ , where each image  $x \in \mathbb{R}^{3 \times H \times W}$  has a spatial resolution of height  $H$  and width  $W$ . The ground truth  $z_l = \{(c_l^k, m_l^k)\}$  for labeled image  $x_l$  consists of class labels  $c^k \in \{1, \dots, K\}$  and binary masks  $m^k \in \{0, 1\}^{H \times W}$ , where  $k$  indexes each instance in  $x_l$ . The goal is to improve model performance beyond what  $\mathcal{D}_L$  alone provides.

Our  $S^4M$  is built on the widely adopted teacher-student

framework with consistency regularization [3, 5] that formulates the training pipeline into two stages [3, 5]. In the first stage, we pre-train the teacher  $\mathcal{F}_T$  on labeled data  $\mathcal{D}_L$  with the objective  $\mathcal{L}_T = \mathcal{L}_{lb}$ . After obtaining the pre-trained teacher network, we then train the student  $\mathcal{F}_S$  utilizing both labeled data  $\mathcal{D}_L$  and unlabeled data  $\mathcal{D}_U$ . At the second stage, the teacher  $\mathcal{F}_T$  processes weakly augmented views  $weak(x_u)$  to generate pseudo-labels  $\hat{z}_u = \{(\hat{c}_u^k, \hat{m}_u^k)\}$ , retaining predictions where class confidence exceeds  $\tau_c$ . The student  $\mathcal{F}_S$  then learns from strongly augmented views  $strong(x_u)$  by matching its predictions to  $\hat{z}_u$ .

The total objective is defined as  $\mathcal{L}_S = \mathcal{L}_{lb} + \lambda_{ulb} \mathcal{L}_{ulb}$ . The network is jointly optimized by a cross entropy loss  $l_{cls}$  for classification and a mask loss  $l_{mask}$  consisted of dice loss [38] and binary cross entropy. Consequently, we define  $\mathcal{L}_{lb} = l_{cls}(\tilde{c}_l^k, c_l^k) + \lambda_{mask} l_{mask}(\tilde{m}_l^k, m_l^k)$  and  $\mathcal{L}_{ulb} = l_{cls}(\tilde{c}_u^k, \hat{c}_u^k) + \lambda_{mask} l_{mask}(\tilde{m}_u^k, \hat{m}_u^k)$ , where  $\{(\tilde{c}, \tilde{m})\}$  are model predictions.

**Baseline segmentation network.** Our framework builds on Mask2Former [7], a unified architecture for segmentation that we adapt for instance segmentation, following Guided Distillation [3]. The model comprises three core components. An image encoder extracts low-resolution features from the input image, and a pixel decoder progressively upsamples and refines these features to construct a multi-scale feature pyramid. A transformer decoder processes  $N$  learnable query embeddings, where iteratively interacting with the multi-scale features to generate class

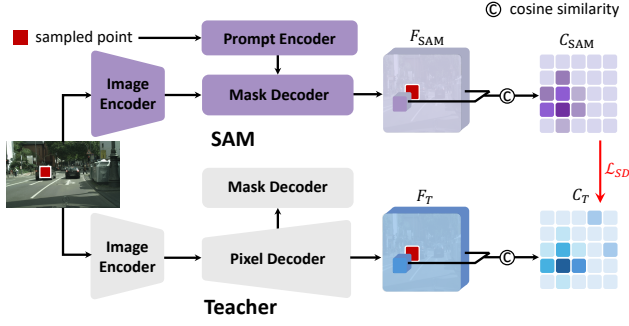


Figure 3. **Illustration of structural distillation with SAM for training the teacher.** We distill the self-similarity matrix extracted from the decoder feature of SAM to enhance the teacher for addressing under-segmentation.

embeddings for each of the  $N$  segments. The binary masks for each segment are generated by computing the dot product between the segment embeddings and the per-pixel embeddings, followed by a sigmoid activation.

**SAM.** The SAM [45, 52] is architecturally structured around three core components: an image encoder, a prompt encoder, and a mask decoder. Image encoder utilizes ViT-based backbone [15, 48] to extract image features and generate  $H' \times W'$  spatial embedding, where  $H'$  and  $W'$  denote the height and width of the feature map, respectively. The prompt encoder captures interactive positional cues from inputs such as points, boxes, and masks, transforming them into prompt embedding that inform the segmentation process. The final mask decoder fuses image and prompt embedding using a modified transformer block with bidirectional self-attention and cross-attention. It then upsamples the image embedding and applies an MLP-based dynamic classifier to compute the mask’s foreground probability at each location. Our approach utilizes two outputs from SAM that contains its rich segmentation knowledge.

## 4. Method

In this section, we detail our approach for integrating SAM into the semi-supervised instance segmentation framework. First, we analyze the limitations of the pseudo-labels produced by existing teacher networks and introduce a structural distillation strategy—leveraging SAM as a meta-teacher, to enhance teacher performance in Sec. 4.1. Next, we describes our approach to fully integrate SAM into the student network training via pseudo-label refinement in Sec. 4.2 and augmentation in Sec. 4.3.

### 4.1. Improving teacher with structural distillation

Recent works [3, 16] have shown that a robust teacher network is essential for effective semi-supervised learning. We observe that teacher networks in existing methods tend to suffer from under-segmentation, which is mainly caused by

the scarcity of labeled data and results in difficulty differentiating multiple instances, as illustrated in Fig. 1. Inspired by the strong localization capabilities of SAM, we propose a knowledge distillation strategy in which SAM functions as a meta-teacher, guiding the teacher network  $\mathcal{F}_T$  toward finer localization. Since the teacher network is initially trained on limited labeled data, the rich representations from SAM are particularly valuable in this label-scarce regime.

A major challenge in this approach is to avoid inheriting undesirable properties from the meta-teacher, such as over-segmentation. Although SAM effectively captures fine-grained regions, its limited semantic understanding stemming from training with geometric prompts rather than semantic labels [52] can lead to over-segmentation and sub-optimal classification performance [51, 52]. Since instance segmentation requires both accurate segmentation and robust classification, we refrain from directly minimizing the feature distance [47] between SAM and the teacher.

Instead, we design a distillation loss that focuses on imitating the structural layout of the image, which can be represented as a self-similarity matrix [8–10, 23–27] of the feature map. One important aspect is identifying where to distill from, in order to ensure that the self-similarity we learn from accurately captures the localized structure of the image, while avoiding over- and under-segmentation. In this regard, we further propose distilling the feature structures obtained from the decoder features of SAM, instead of encoder features [52]. Consequently, we randomly sample  $P$  points, yielding  $i \in \{1, \dots, P\}$  and prompt the decoder with the sampled point to obtain more localized features, as illustrated in Fig. 3.

Following this, we first extract the feature map from both SAM and the teacher model. The feature map extracted from the teacher model is interpolated to match the spatial dimensions of SAM, producing  $F_{\text{SAM}}, F_T \in \mathbb{R}^{H' \times W' \times d}$ . Using the sampled prompt locations, we then extract the spatially corresponding embeddings from both, resulting in  $F'_{\text{SAM}}, F'_T \in \mathbb{R}^{P \times d}$ . Finally, we compute the cosine similarity within the features, yielding self-similarity matrices  $C_{\text{SAM}}, C_T \in \mathbb{R}^{P \times H'W'}$ .

Formally, this is defined as:

$$C_{\text{SAM}} = \frac{F'_{\text{SAM}} \cdot F_{\text{SAM}}}{\|F'_{\text{SAM}}\| \|F_{\text{SAM}}\|}, \quad C_T = \frac{F'_T \cdot F_T}{\|F'_T\| \|F_T\|}. \quad (1)$$

We define structural distillation (SD) loss as:

$$\mathcal{L}_{\text{SD}} = \frac{1}{H'W'} \sum_i \rho(C_{\text{SAM}}(i) - C_T(i)), \quad (2)$$

where  $\rho$  is the Huber function [30], and  $i \in \{1, \dots, H'W'\}$  is the index along the first dimension representing the query, updating the objective for the teacher as  $\mathcal{L}_T = \mathcal{L}_{\text{lb}} + \mathcal{L}_{\text{SD}}$ .



## 4.2. Refining pseudo-labels with SAM

In addition to improving the teacher network, we propose to further boost the semi-supervised instance segmentation framework by refining the pseudo-label with SAM for minimizing the remaining error to mitigate error propagation stemming from noisy labels and to maximize the potential of the unlabeled data. Given a pseudo-label generated from the teacher, we can obtain geometric prompts to obtain refined labels from SAM. However, we find that naïve methods, such as selecting the center point of the mask [33] or obtaining the bounding box is prone to error and often results in over-segmentation or degenerate masks. To prevent this, we introduce a simple trick by stochastically sampling multiple points as prompts to SAM.

Given a pseudo-label mask  $\hat{m}_u^k \in \{0, 1\}^{H \times W}$ , we can also access the per-pixel probability of the mask  $\tilde{m}_u^k \in [0, 1]^{H \times W}$  before applying threshold to obtain binary masks. We then can obtain a probability distribution by normalizing the following:

$$p(a, b) = \begin{cases} \tilde{m}_u^k, & \text{if } \hat{m}_u^k(a, b) = 1, \\ 0, & \text{if } \hat{m}_u^k = 0, \end{cases} \quad (3)$$

where  $a, b$  represents the spatial locations. Consequently, we obtain refined pseudo-labels by prompting SAM with  $K$  points sampled from the distribution  $\tilde{p}(a, b) = \frac{p(a, b)}{\sum p(a, b)}$ , where we set  $K = 3$  in our method. As shown in Fig. 4, we can observe that SAM can effectively refine noisy pseudo-labels as high quality pseudo-labels. We note that although the refined pseudo-label may not always outperform the original due to stochasticity, the student network still benefits from improved samples and shows consistent gains.

## 4.3. Augmenting images with refined pseudo-labels

The motivation for leveraging weak-to-strong consistency [50] in semi-supervised learning is to enforce consistent predictions under challenging conditions using strong augmentations [12, 14, 58, 59]. This method has significantly improved tasks like semantic segmentation by utilizing techniques such as [39, 58], which boost robustness and generalizability. However, compared to semi-supervised semantic segmentation, instance segmentation has been less explored, often relying solely on photometric augmentations. Here, we introduce Augmentation with Refined Pseudo-label (ARP), an augmentation strategy inspired by prior work [17], addressing unreliable pseudo-labels through refinement to more effectively enhance performance. As illustrated in Figure 2, ARP generates synthetic images on the fly by leveraging pseudo-labels from a teacher network. Let  $x_A, x_B \in D_U$  denote a randomly sampled pair of weakly augmented images from a training batch. Their refined pseudo-masks,  $\{\hat{m}_A^k\}_{k=1}^{N_A}$  and




Figure 4. **Visualization of pseudo-labels before and after refinement.** We visualize pseudo-labels from the teacher network before (left) and after (right) refinement. With SAM, we can refine pseudo-labels with under-segmentation, often containing noisy parts of nearby instances, into high-quality pseudo-labels.

$\{\hat{m}_B^k\}_{k=1}^{N_B}$ , are aggregated into binary masks  $M_A$  and  $M_B$ , where  $N_A$  and  $N_B$  denote the number of pseudo-label instances for  $x_A$  and  $x_B$ , respectively. These masks are then used to bidirectionally paste detected instances between  $x_A$  and  $x_B$ , generating synthetic images  $x_{AB}, x_{BA}$  as follows:

$$\begin{aligned} x_{AB} &\leftarrow M_B \odot x_B + (1 - M_B) \odot x_A, \\ x_{BA} &\leftarrow M_A \odot x_A + (1 - M_A) \odot x_B. \end{aligned} \quad (4)$$

Here,  $\odot$  denotes the element-wise product between the binary mask and the image. The corresponding pseudo-labels  $\hat{z}_A$  and  $\hat{z}_B$  are also augmented accordingly into  $\hat{z}_{AB}$  and  $\hat{z}_{BA}$ . Subsequently, the student network  $\mathcal{F}_S$  is trained on the photometric augmented  $x_{AB}$  and  $x_{BA}$ . By placing instances from paired images into each other’s contexts, the method introduces diverse spatial variations, including novel scenes and potential occlusions. These transformations encourage consistent training under challenging conditions, enhancing robustness and generalization.

## 5. Experiments

### 5.1. Experimental setup

**Datasets and evaluation metric.** We conducted our experiments on two datasets. The **Cityscapes** dataset [11] contains 1024 x 2048 resolution driving scene images, comprising 2975 training images and 500 validation images, with pixel-level annotations for 8 semantic instance categories. For the semi-supervised setup, subsets comprising 5%, 10%, and 20% of the training images were sampled and used in our experiments. The **COCO** dataset [35] is a large-scale benchmark containing 118,287 images with instance segmentation annotations and is widely used in the field. For our experiments, we train on 1%, 2%, and 5% subsets of the training data and evaluate the performance on the 5000-image validation set. Following pre-



Figure 5. **Qualitative comparison on the Cityscapes dataset** [11] using 10% labeled data, comparing the baseline semi-supervised method GuidedDistillation [53] (top), and our approach (bottom). Compared to supervised training and the baseline method, our approach not only detects and segments instances more accurately but also exhibits higher discriminability between instances of the same class.

Methods	Cityscapes			
	5%	10%	20%	30%
DataDistillation [43]	13.7	19.2	24.6	27.4
NoisyBoundaries [54]	17.1	22.1	29.0	32.4
PAIS [28]	18.0	22.9	29.2	32.8
Guided Distillation [3]	23.0	30.8	33.1	35.6
Depth-Guided [5]	<u>23.2</u>	<u>30.9</u>	<u>34.1</u>	<u>36.7</u>
<b>S<sup>4</sup>M (Ours)</b>	<b>30.1</b>	<b>33.3</b>	<b>37.3</b>	<b>37.8</b>

Table 1. **Quantitative comparison on Cityscapes.** We provide comparison of Average Precision (AP) on Cityscapes under different label ratios with state-of-the-art methods. Results for DataDistillation is obtained from [54].

Methods	COCO			
	1%	2%	5%	10%
DataDistillation [43]	3.8	11.8	20.4	24.2
NoisyBoundaries [54]	7.7	16.3	24.9	29.2
PAIS [28]	21.1	-	29.3	31.0
Guided Distillation [3]	21.5	25.3	29.9	<u>35.0</u>
Depth-Guided [5]	<u>22.3</u>	<u>26.3</u>	<u>31.5</u>	<b>35.1</b>
<b>S<sup>4</sup>M (Ours)</b>	<b>24.2</b>	<b>28.1</b>	<b>32.1</b>	34.6

Table 2. **Quantitative comparison on COCO .** We provide comparison of Average Precision (AP) on COCO under different label ratios with state-of-the-art methods. Results for DataDistillation is obtained from [54].

vious works [3, 5, 19, 54], our results are evaluated using the mask-AP metric.

**Implementation details.** We implemented our approach based on the GuidedDistillation [3] codebase, using Mask2Former [7] with a ResNet-50 [18] backbone as the instance segmentation model. For the Cityscapes dataset, we trained the model with a batch size of 16 for 90K iterations on an RTX 3090 GPU with 24GB of RAM, while for the COCO dataset, the model was trained with a batch size of 12 for 368K iterations on an RTX A6000 GPU with 48GB of RAM. The models were optimized using the AdamW optimizer with a learning rate of  $10^{-4}$ , a weight decay of 0.05, and a multiplier of 0.1 applied to backbone updates. The thresholds for pseudo labels were set to 0.7 for class confidence and 5 for instance size. Additionally, the EMA weight decay rate  $\alpha$  was set to 0.9996, and the unsupervised loss weight  $\lambda_u$  was set to 2. For pseudo-label refinement, we utilized Segment Anything Model2 [45] with the Hiera-L [48] backbone. In our main experiments, this model was used without any additional fine-tuning on the datasets. The code will be made available upon acceptance.

## 5.2. Main results

**Quantitative results.** Table 1 presents the quantitative comparison results of our method with several existing works. We report the performance of models trained on the labeled data splits for the two datasets, Cityscapes and COCO, using their respective validation sets. The results demonstrate that our model outperforms previous methods, achieving state-of-the-art performance. On the **Cityscapes** dataset, our method achieves performance improvements of 6.9, 2.4, 3.2 and 1.1 points AP over the previous state-of-the-art methods for 5%, 10%, 20% and 30% labeled subsets, respectively. Compared to the teacher network, our method achieves improvements of 12.9, 10.8, 7.7, and 5.8 points AP for the same partitions. These results demonstrate the effectiveness of our proposed methods in achieving substantial performance gains even with a limited amount of labeled data. Notably, the results under the most challenging setting, 5%, are particularly remarkable. On the **COCO** dataset, **S<sup>4</sup>M** achieves state-of-the-art performance across the 1%, 2%, and 5% splits. In particular, it achieves improvements of 1.9 and 1.8 points AP over the previous best



Figure 6. **Qualitative comparison on the COCO dataset [35]** using 2% labeled data, comparing the baseline semi-supervised method GuidedDistillation [53] (top), and our approach (bottom).

methods for the 1% and 2% splits. We observe a slight drop in the 10% split, where we find that the stochasticity in the pseudo-label refinement could cause the drop given that the pseudo-labels are already in high-quality from having more labeled data. While could potentially address by adjusting  $K$  or calibrating  $\tilde{p}$  in the refinement, we highlight the substantial gains in lower label ratio splits, which aligning with the principle of semi-supervised learning for enhancing the framework with only small amount of labeled data.

**Qualitative results.** We provide qualitative examples for both Cityscapes and COCO in Figure 5 and Figure 6. When comparing the results of our method to those from comparative baseline, GuidedDistillation, we observe that our model is trained to align more closely with the goals of instance segmentation. In the case of GuidedDistillation, both datasets exhibit instances where multiple objects of the same semantic class are not properly separated, leading to the inclusion of multiple instances in a single mask proposal. In contrast, our model not only demonstrates higher performance in accurately distinguishing between instances but also achieves improved segmentation accuracy. We contribute our gains particularly to the careful adoption of SAM, effectively addressing under- and over-segmentation seen in compared baseline. We provide further qualitative results in the supplementary materials.

### 5.3. Ablation studies

We provide ablation studies for validating our approach and our design choices. All experiments were performed with the Mask2Former model with ResNet-50 backbone, consistent with the setup used in the main experiments. We report the results of the best model obtained over 45K training iterations from Cityscapes dataset using the 10% partition of the labeled data if not specified.

SD	Cityscapes				COCO			
	5%	10%	20%	30%	1%	2%	5%	10%
$\times$	14.9	19.0	27.6	29.9	13.5	19.5	25.8	30.1
$\checkmark$	17.2	22.5	29.6	32.0	15.1	20.3	26.1	30.3

Table 3. **Effects of the structural distillation (SD) loss for pre-training the teacher.** We compare our improved teacher with the structural distillation loss to the baseline teacher reproduced from [3], which is trained only with  $\mathcal{L}_{lb}$ .

**Effects of the SD loss for the teacher.** Tab 3 presents the performance of the teacher network enhanced with the proposed SD loss across varying label ratios on the Cityscapes and COCO datasets. The baseline refers to a Mask2Former [7] trained with  $\mathcal{L}_{lb}$  for different label ratios, following GuidedDistillation [3]. On Cityscapes with 5% labeled data, the SD loss yields an improvement of +2.3 AP over the baseline, with consistent gains of +3.5 AP, +2.0 AP, and +2.1 AP observed at 10%, 20%, and 30% label ratios, respectively. Similarly, on COCO, the SD loss achieves a notable gain of +1.6 AP under the 1% label setting, with improvements maintained across higher label ratios. These results indicate that incorporating the SD loss effectively enhances the pre-training process across all splits with especially higher gains for lower label ratio splits, demonstrating its efficacy in label-scarce settings.

**Ablation of SD loss in different training stages.** Table 4 presents an ablation study on the application of the SD loss, specifically examining the effect of employing the structural distillation loss,  $\mathcal{L}_{SD}$ , at different training stages. When  $\mathcal{L}_{SD}$  is applied exclusively during teacher training or solely during student training, both configurations yield improvements. However, the teacher-only configuration delivers a



$\mathcal{L}_{SD}$		AP
teacher	student	
×	×	30.1
✓	×	<b>32.8</b>
×	✓	30.7
✓	✓	29.4

Table 4. Ablation on  $\mathcal{L}_{SD}$  in different training stages.

$F_{SAM}$	distill. loss	AP
-	-	19.0
encoder	feature	21.1
encoder	structural	21.4
decoder	structural	<b>22.4</b>

Table 5. Ablation on  $F_{SAM}$  and distillation loss.

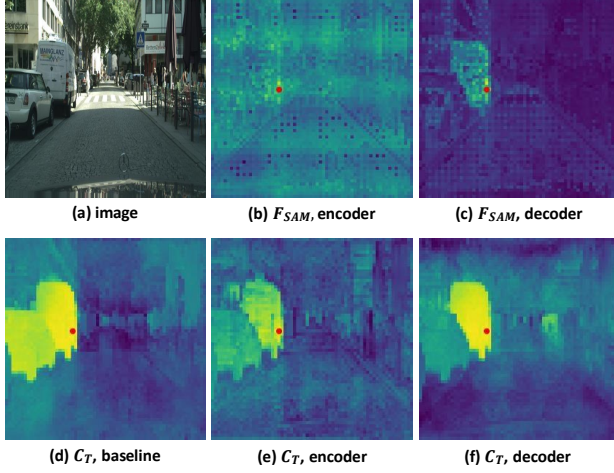


Figure 7. **Visualization of self-similarity matrices  $C$ .** Given an image (a), we visualize  $C_{SAM}$  with encoder (b) and decoder (c) features for  $F_{SAM}$ . We also visualize  $C_T$  for baseline teacher (d), as well as the corresponding teachers trained with  $\mathcal{L}_{SD}$  with  $C_{SAM}$  from encoder (e) and decoder (f).

substantially higher overall gain, underscoring the critical importance of a robust teacher network. Surprisingly, applying  $\mathcal{L}_{SD}$  to both the teacher and student stages results in a performance drop after the burn-in iterations [3], suggesting that the burn-in stage to be problematic.

**Ablation on design choices for SD loss.** In Table 5, we present ablations investigating the design choices of our SD loss by comparing different strategies for applying the structural distillation loss  $\mathcal{L}_{SD}$  in the teacher network. When using feature distillation [47], we minimize the Euclidean distance between  $F_{SAM}$  and  $F_T$ . Although all configurations yield noticeable improvements compared to the baseline without it, we observe that employing structural distillation leads to better teacher performance than feature distillation. Moreover, utilizing decoder-derived features for  $F_{SAM}$  results in an additional gain, verifying our choices.

Figure 7 provides visualizations of  $C_{SAM}$  and  $C_T$  in the context of the structural distillation loss. As shown in (b-c), the  $C_{SAM}$  obtained from the decoder demonstrates improved localization and maintains high similarity within the

	SD	PR	ARP	AP
(I)	×	×	×	28.4
(II)	✓	×	×	29.2
(III)	×	×	✓	27.8
(IV)	✓	✓	×	32.3
(V)	✓	×	✓	31.0
(VI)	✓	✓	✓	<b>32.8</b>

Table 6. **Component analysis.** We conduct ablation study on our key components structural distillation (SD), pseudo-label refinement (PR), and augmentation (ARP).

target instance at an appropriate level of granularity. Additionally, panels (e-f) illustrate the corresponding teacher self-similarity matrix  $C_T$  from a teacher trained with  $\mathcal{L}_{SD}$ , which, when compared to the baseline in panel (d) (i.e., without  $\mathcal{L}_{SD}$ ), clearly distinguishes the similarity across different instances within the same class.

**Component analysis.** We conduct ablation experiments to evaluate the effectiveness of the main components of our method: structural distillation loss (SD), pseudo-label refinement (PR) and the proposed augmentation strategy, ARP. As shown in Table 6, the performance improvement of the teacher model trained with  $\mathcal{L}_{SD}$  loss positively contributes to the learning of the student network (II). Furthermore, applying pseudo-label refinement leads to a significant performance boost (IV), demonstrating the effectiveness of incorporating SAM into a semi-supervised instance segmentation framework. While ARP alone can negatively affect student training (III)—likely due to noise from unrefined pseudo-labels—its combination with a teacher trained under SD still yields performance gains (V), suggesting improved pseudo-label generation. Moreover, combining ARP with pseudo-label refinement achieves the best performance (VI), indicating that these two methods work synergistically to enhance learning.

## 6. Conclusion

In this work, we propose **S<sup>4</sup>M**, a novel semi-supervised instance segmentation framework that integrates the SAM through structured distillation, pseudo-label refinement, and data augmentation. By selectively leveraging precise localization capability of SAM while mitigating its over-segmentation tendency, our approach significantly improves the teacher-student framework. Extensive experiments demonstrate state-of-the-art performance across benchmarks, highlighting the effectiveness of our method in enhancing semi-supervised instance segmentation.



## Acknowledgement

This research was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (RS-2019-II190075, RS-2024-00509279, RS-2025-II212068, RS-2023-00227592, RS-2025-02214479, RS-2024-00457882) and the Culture, Sports, and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism (RS-2024-00345025, RS-2024-00333068), and National Research Foundation of Korea (RS-2024-00346597).

## References

- [1] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2(3):4, 2021. 2
- [2] Honggyu An, Jinhyeon Kim, Seonghoon Park, Jaewoo Jung, Jisang Han, Sunghwan Hong, and Seungryong Kim. Cross-view completion models are zero-shot correspondence estimators. *arXiv preprint arXiv:2412.09072*, 2024. 2
- [3] Tariq Berrada, Camille Couprie, Karteek Alahari, and Jakob Verbeek. Guided distillation for semi-supervised instance segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 475–483, 2024. 1, 2, 3, 4, 6, 7, 8
- [4] Yunkang Cao, Xiaohao Xu, Chen Sun, Yuqi Cheng, Zongwei Du, Liang Gao, and Weiming Shen. Segment any anomaly without training via hybrid prompt regularization. *arXiv preprint arXiv:2305.10724*, 2023. 2
- [5] Xin Chen, Jie Hu, Xiawu Zheng, Jiangang Lin, Liujuan Cao, and Rongrong Ji. Depth-guided semi-supervised instance segmentation. *arXiv preprint arXiv:2406.17413*, 2024. 1, 2, 3, 6
- [6] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in neural information processing systems*, 34:17864–17875, 2021. 1
- [7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 1, 3, 6, 7
- [8] Seokju Cho, Sunghwan Hong, Sangryul Jeon, Yunsung Lee, Kwanghoon Sohn, and Seungryong Kim. Cats: Cost aggregation transformers for visual correspondence. *Advances in Neural Information Processing Systems*, 34:9011–9023, 2021. 4
- [9] Seokju Cho, Sunghwan Hong, and Seungryong Kim. Cats++: Boosting cost aggregation with convolutions and transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7174–7194, 2022.
- [10] Seokju Cho, Heeseong Shin, Sunghwan Hong, Anurag Arnab, Paul Hongsuck Seo, and Seungryong Kim. Catseg: Cost aggregation for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4113–4123, 2024. 4
- [11] 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 *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5, 6, 1
- [12] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 5
- [13] Haixing Dai, Chong Ma, Zhiling Yan, Zhengliang Liu, Enze Shi, Yiwei Li, Peng Shu, Xiaozheng Wei, Lin Zhao, Zihao Wu, et al. Samaug: Point prompt augmentation for segment anything model. *arXiv preprint arXiv:2307.01187*, 2023. 1
- [14] Terrance DeVries. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 5
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [16] Dominik Filipiak, Andrzej Zapala, Piotr Tempczyk, Anna Fensel, and Marek Cygan. Polite teacher: Semi-supervised instance segmentation with mutual learning and pseudo-label thresholding. *IEEE Access*, 12:37744–37756, 2024. 1, 2, 4
- [17] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2918–2928, 2021. 5
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1, 6
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 1
- [21] Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, and Youliang Yan. Knowledge adaptation for efficient semantic segmentation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 578–587, 2019. 3
- [22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015. 3
- [23] Sunghwan Hong and Seungryong Kim. Deep matching prior: Test-time optimization for dense correspondence. In

- Proceedings of the IEEE/CVF international conference on computer vision*, pages 9907–9917, 2021. 4
- [24] Sunghwan Hong, Seokju Cho, Jisu Nam, Stephen Lin, and Seungryong Kim. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *European Conference on Computer Vision*, pages 108–126. Springer, 2022. 1
- [25] Sunghwan Hong, Jisu Nam, Seokju Cho, Susung Hong, Sangryul Jeon, Dongbo Min, and Seungryong Kim. Neural matching fields: Implicit representation of matching fields for visual correspondence. *Advances in Neural Information Processing Systems*, 35:13512–13526, 2022.
- [26] Sunghwan Hong, Seokju Cho, Seungryong Kim, and Stephen Lin. Unifying feature and cost aggregation with transformers for semantic and visual correspondence. *arXiv preprint arXiv:2403.11120*, 2024.
- [27] Sunghwan Hong, Jaewoo Jung, Heeseong Shin, Jiaolong Yang, Seungryong Kim, and Chong Luo. Unifying correspondence pose and nerf for generalized pose-free novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20196–20206, 2024. 4
- [28] Jie Hu, Chen Chen, Liujuan Cao, Shengchuan Zhang, Annan Shu, Guannan Jiang, and Rongrong Ji. Pseudo-label alignment for semi-supervised instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16337–16347, 2023. 1, 2, 6
- [29] Yuhao Huang, Xin Yang, Lian Liu, Han Zhou, Ao Chang, Xinrui Zhou, Rusi Chen, Junxuan Yu, Jiongquan Chen, Chaoyu Chen, et al. Segment anything model for medical images? *Medical Image Analysis*, 92:103061, 2024. 2
- [30] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, pages 492–518. Springer, 1992. 4
- [31] Deyi Ji, Feng Zhao, Hongtao Lu, Feng Wu, and Jieping Ye. Structural and Statistical Texture Knowledge Distillation and Learning for Segmentation. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 5555. 3
- [32] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9404–9413, 2019. 1
- [33] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2, 5
- [34] Hyeokjun Kweon and Kuk-Jin Yoon. From sam to cams: Exploring segment anything model for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19499–19509, 2024. 2
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context, 2015. 5, 7, 1
- [36] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [37] Juzheng Miao, Cheng Chen, Keli Zhang, Jie Chuai, Quanzheng Li, and Pheng-Ann Heng. Cross prompting consistency with segment anything model for semi-supervised medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 167–177. Springer, 2024. 2
- [38] 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. 3
- [39] Viktor Olsson, Wilhelm Tranheden, Julianio Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1369–1378, 2021. 5
- [40] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1
- [41] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3967–3976, 2019. 3
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 1
- [43] Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omniscient supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4119–4128, 2018. 6
- [44] Frano Rajić, Lei Ke, Yu-Wing Tai, Chi-Keung Tang, Martin Danelljan, and Fisher Yu. Segment anything meets point tracking. *arXiv preprint arXiv:2307.01197*, 2023. 2
- [45] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 2, 4, 6
- [46] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024. 2
- [47] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 4, 8
- [48] Chaitanya Ryali, Yuan-Ting Hu, Daniel Bolya, Chen Wei, Haoqi Fan, Po-Yao Huang, Vaibhav Aggarwal, Arkabandhu

- Chowdhury, Omid Poursaeed, Judy Hoffman, et al. Hiera: A hierarchical vision transformer without the bells-and-whistles. In *International Conference on Machine Learning*, pages 29441–29454. PMLR, 2023. 4, 6
- [49] Heeseong Shin, Chaehyun Kim, Sunghwan Hong, Seokju Cho, Anurag Arnab, Paul Hongsuck Seo, and Seungryong Kim. Towards open-vocabulary semantic segmentation without semantic labels. *Advances in Neural Information Processing Systems*, 37:9153–9177, 2025. 2
- [50] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 2, 5
- [51] Vibashan VS, Shubhankar Borse, Hyojin Park, Debasmit Das, Vishal Patel, Munawar Hayat, and Fatih Porikli. Pos-sam: Panoptic open-vocabulary segment anything. *arXiv preprint arXiv:2403.09620*, 2024. 4
- [52] Haoxiang Wang, Pavan Kumar Anasosalu Vasu, Fartash Faghri, Raviteja Vemulapalli, Mehrdad Farajtabar, Sachin Mehta, Mohammad Rastegari, Oncel Tuzel, and Hadi Pouransari. Sam-clip: Merging vision foundation models towards semantic and spatial understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3635–3647, 2024. 2, 4
- [53] Zhenyu Wang, Yali Li, and Shengjin Wang. Noisy boundaries: Lemon or lemonade for semi-supervised instance segmentation? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16826–16835, 2022. 2, 6, 7
- [54] Zhenyu Wang, Yali Li, and Shengjin Wang. Noisy boundaries: Lemon or lemonade for semi-supervised instance segmentation? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16826–16835, 2022. 1, 6
- [55] Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain Brégier, Yohann Cabon, Vaibhav Arora, Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jérôme Revaud. Croco: Self-supervised pre-training for 3d vision tasks by cross-view completion. *Advances in Neural Information Processing Systems*, 35:3502–3516, 2022. 1
- [56] Guoping Xu, Xiaoxue Qian, Hua-Chieh Shao, Jax Luo, Weiguo Lu, and You Zhang. A segment anything model-guided and match-based semi-supervised segmentation framework for medical imaging. *Medical physics*, 2025. 2
- [57] 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. 2
- [58] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019. 5
- [59] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 5
- [60] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Xianzheng Ma, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023. 1
- [61] Yichi Zhang, Jin Yang, Yuchen Liu, Yuan Cheng, and Yuan Qi. Semisam: Enhancing semi-supervised medical image segmentation via sam-assisted consistency regularization. In *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 3982–3986. IEEE, 2024. 2
- [62] Ziyu Zhang, Sanja Fidler, and Raquel Urtasun. Instance-level segmentation for autonomous driving with deep densely connected mrfs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [63] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision*, pages 696–712. Springer, 2022. 2
- [64] Juan Carlos Ángeles Cerón, Gilberto Ochoa Ruiz, Leonardo Chang, and Sharib Ali. Real-time instance segmentation of surgical instruments using attention and multi-scale feature fusion. *Medical Image Analysis*, 81:102569, 2022. 1