

Lite-SAM Is Actually What You Need for Segment Everything

Jianhai Fu^{1,*}, Yuanjie Yu^{1,2,*}, Ningchuan Li^{1,†}, Yi Zhang¹,
Qichao Chen¹, Jianping Xiong¹, Jun Yin², and Zhiyu Xiang^{2,†}

¹ Zhejiang Dahua Technology Co., Ltd., Hangzhou, China

² Zhejiang University, Hangzhou, China

Abstract. The Segment Anything model (SAM) has brought significant changes to the segmentation field with its superior performance, but its extensive computational resource requirements remain a limiting factor. Many works, such as MobileSAM, Edge-SAM, and MobileSAM-v2, have explored lightweight solutions. However, their use of traditional Grid Search sampling strategies or two-stage concatenation methods, which do not allow for end-to-end training, severely limit the performance of segment everything (SegEvery).

This paper introduces Lite-SAM, an efficient end-to-end solution for the SegEvery task designed to reduce computational costs and redundancy. Lite-SAM is composed of four main components: a streamlined CNN-Transformer hybrid encoder (LiteViT), an automated prompt proposal network (AutoPPN), a traditional prompt encoder, and a mask decoder. All these components are integrated within the SAM framework. Our LiteViT, a high-performance lightweight backbone network, has only 1.16M parameters, which is a 23% reduction compared to the lightest existing backbone network Shufflenet. We also introduce AutoPPN, an innovative end-to-end method for prompt boxes and points generation. This is an improvement over traditional grid search sampling methods, and its unique design allows for easy integration into any SAM series algorithm, extending its usability.

we have thoroughly benchmarked Lite-SAM across a plethora of both public and private datasets. The evaluation encompassed a broad spectrum of universal metrics, including the number of parameters, SegEvery execution time, and accuracy. The findings reveal that Lite-SAM, operating with a lean 4.2M parameters, significantly outpaces its counterparts, demonstrating performance improvements of 43x, 31x, 20x, 21x, and 1.6x over SAM, MobileSAM, Edge-SAM, EfficientViT-SAM, and MobileSAM-v2 respectively, all the while maintaining competitive accuracy. This underscores Lite-SAM's prowess in achieving an optimal equilibrium between performance and precision, thereby setting a new state-of-the-art(SOTA) benchmark in the domain.

Keywords: SegEvery · AutoPPN · LiteViT · End-to-End

*Equal contribution.

†Corresponding author.

1 Introduction

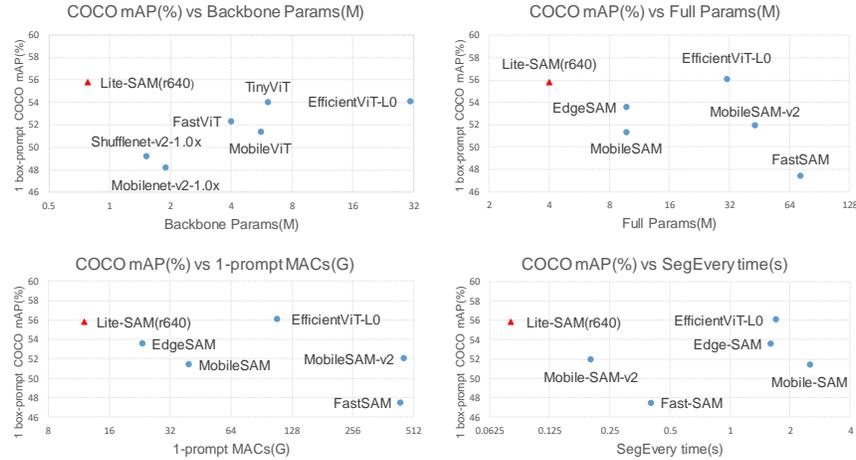


Fig. 1: The proposed Lite-SAM achieves SOTA performance in terms of Backbone Parameters (top left), Full Parameters (top right), Multiply-Accumulate Operations (bottom left), and SegEvery time (bottom right) tasks while maintaining computational efficiency. The metrics were evaluated on the zero-shot learning of the COCO dataset. Note that the comparison of backbone parameters is made against lightweight network structures (params ≤ 40 M), with MAE not falling within this scope.

Zhang et al. [41] have made a remarkable leap in the field of NLP, resulting in a significant breakthrough in generative AI (AIGC, also known as Artificial Intelligence Generated Content) [42]. This breakthrough has largely been enabled by the GPT-series models [3, 27], which are foundation models [2] trained on extensive text datasets. Capitalizing on the success of these foundation models in NLP, multiple studies [9, 26, 44] have explored the integration of image encoders and text encoders via contrastive learning [43].

The Meta Research team has recently introduced an ambitious program known as “segment everything” project called SAM [13]. SAM represents a crucial advancement for vision framework, drawing parallels to the impact of GPT in NLP. It comprises two key components: a ViT-based image encoder and a prompt-guided mask decoder, which work in conjunction. SAM is designed to handle two segmentation tasks: SegAny and SegEvery. Both tasks involve class-agnostic mask segmentation but differ in their objectives. SegAny uses a specific prompt, such as a point or box, to isolate and segment a particular item of interest within an image. In contrast, SegEvery’s objective is to delineate all discernible subjects in the image.

Chaoning Zhang et al. [39] proposed a “decoupled distillation” aimed at distilling the ViT-H decoder of SAM [13], which yielded a more efficient lightweight encoder that could integrate with SAM’s decoder. However, this algorithm model lacks robustness in platform adaptation and exhibits considerable accuracy loss during such translations, rendering it less suitable for deployment on mobile devices. Zhao et al. [46] introduced the Fast-SAM model, built upon YOLOv8 [12], that demonstrates remarkable segmentation capabilities. Its main limitation, however, is the absence of a full range of interactive modalities, notably lacking in dedicated box and point functionalities. Li et al. [15] engineered Semantic-SAM, a model that enhances the segmentation and recognition versatility of images across varying scales. It is imperative to highlight, though, that its substantial number of parameters contributes to longer inference times. Han Cai et al. [4] presented EfficientViT, introduced a novel lightweight algorithm called which achieved promising results. Chong Zhou et al. proposed Edge-SAM [47], an algorithm that accomplishes real-time execution for the SegAny task on an iPhone. All the aforementioned methods [4, 13, 15, 35, 39, 46, 47] are all evaluated for SegAny; however, the SegEvery continues to be highly time-demanding.

MobileSAM-v2 [40] proposed an innovative training approach for YOLOv8 [12] that uses pre-generated prompts (Object-Aware Prompt Sampling) in place of the traditional Gridsearch sampling strategy, enhancing the efficiency of the SegEvery process. Despite this improvement, this approach necessitates the use of separate models, which is considered a stopgap measure. Due to YOLOv8’s inherent inference and training demands, the overall time savings may be limited.

In order to address the aforementioned issues, our contributions can be summarized as follows:

- We Introduced LiteViT, a lightweight CNN-Transformer encoder, enhancing accuracy with reduced parameters, ideal for limited computational environments.
- The development of AutoPPN, an automated prompt proposal network, improving efficiency over grid search methods and integrating easily with SAM series algorithms.
- Validated Lite-SAM’s performance through experiments, as depicted in Fig. 1, showing accelerated results on SegEvery while preserving accuracy.

2 Related Works

2.1 Segment Anything

In the evolving field of image segmentation, the SAM [13] stands out as a significant progress. Its groundbreaking training methodology and exceptional performance on extensive visual datasets distinguish it. SAM is particularly adept in class-agnostic segmentation and shows impressive efficacy in zero-shot scenarios.

In addition to the work on lightweight versions of Segment Anything and its variants mentioned in [4, 13, 15, 35, 39, 40, 46, 47], a series of works combining SAM with various downstream tasks have also achieved impressive results.

Grounded SAM [28] integrates Grounding DINO’s open-set detection with SAM, enabling text-guided detection and segmentation in images. SegGPT [32] standardizes diverse segmentation data into a single image format, excelling in segmenting both in-domain and out-of-domain subjects with strong performance. Zou et al. [49] presented SEEM, featuring a versatile decoding mechanism for various segmentation tasks, aiming to create a universal interface akin to large language models. Inpaint Anything [37] introduces a novel “click and fill” method for mask-free image inpainting, blending SAM models with AIGC to create an efficient and user-friendly solution for inpainting tasks. SAM3D [36] advances 3D perception by mapping 2D segmentation to 3D spaces. It enables 3D point cloud mask prediction using RGB images with the SAM model, eliminating the need for additional training or fine-tuning.

As a multipurpose foundational model, SAM has greatly enhanced interactive segmentation techniques and demonstrated remarkable flexibility across diverse segmentation tasks. Its contributions has notably expanded the horizons for applications in open-world image understanding. However, a noteworthy limitation of SAM is its constrained real-time processing capabilities, which poses obstacles for time-sensitive applications.

2.2 Lightweight ViT and CNN

Historically, mobile vision applications have heavily relied on lightweight Convolutional Neural Networks (CNNs) like MobileNet [11] and ShuffleNet [21, 45]. The MobileNet series [10, 29] was pioneering in its segmentation of convolution blocks into depth-wise and point-wise convolutions, significantly reducing model size and computational demand. The emergence of Vision Transformers (ViTs) [7] has spurred efforts to streamline these architectures, resulting in more compact and efficient models such as Deit-Small (Deit-S) and Deit-Tiny (Deit-T) [30]. MobileViT [23] fused ViTs with conventional convolutions, outperforming MobileNet-v2 [29] by focusing on improved local feature extraction, a forte of CNNs. The trend toward computational economy is further advanced by subsequent models, including EfficientFormer [17], Efficient ViT [20], Next-ViT [16], TinyViT [33], and FastViT [31].

Through extensive experimentation, our Lite-SAM algorithm achieves an optimal balance between model complexity and inference speed. In our research, we introduce Lite-SAM, a lightweight algorithm that capitalizes on the LiteViT backbone and leverages a prompt-based network architecture, namely AutoPPN. Lite-SAM distinguishes itself by having a low parameter count and reduced computational costs, yet it is capable of attaining performance benchmarks similar to those of SAM-B. Our comprehensive testing indicates that Lite-SAM strikes an optimal balance, offering reduced model complexity while maintaining swift inference speeds.

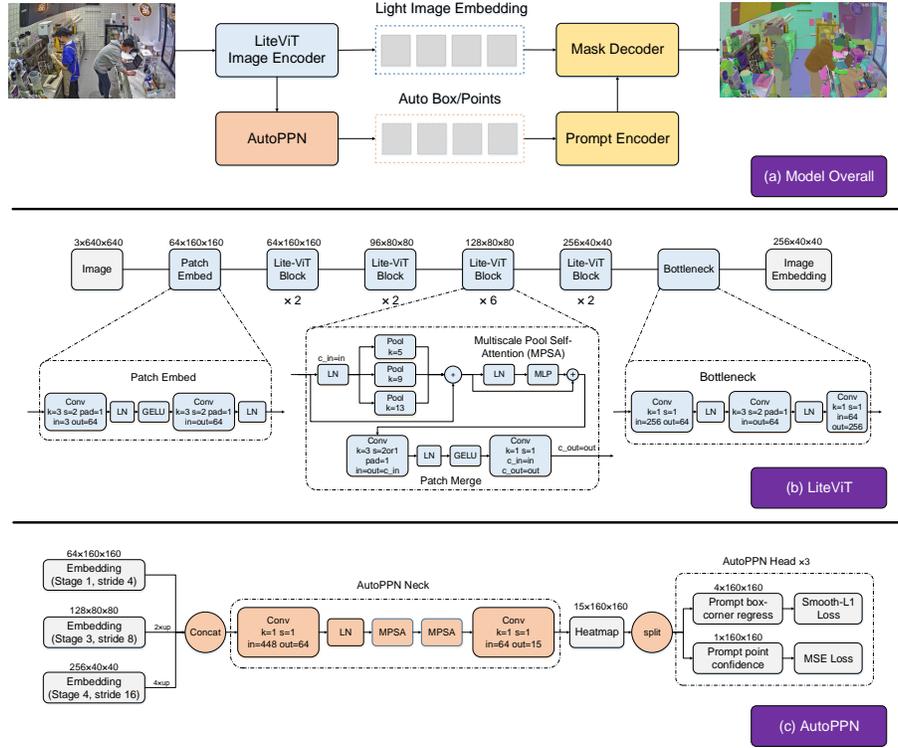


Fig. 2: (a) Overview of the proposed Lite-SAM. The architecture consists of two detachable blocks, namely the Lightweight ViT backbone (LiteViT), Automated Prompt Proposal Network (AutoPPN). (b) Macro Architecture of LiteViT. (c) Macro Architecture of AutoPPN.

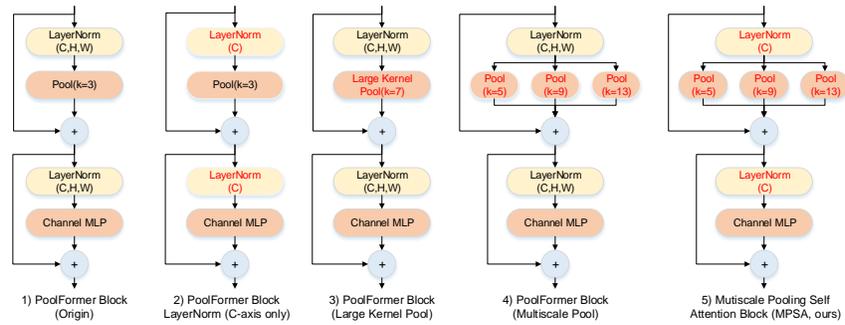


Fig. 3: Overview of architectural choice. (1) represents the original PoolFormer Block, (2),(3) and (4) show the modifications to PoolFormer Block, and (5) the final version of our Multiscale Pooling Self Attention module.

3 Method: Lite-SAM

3.1 Design motivation and choices

We present the Lite-SAM architecture, which consists of four main components: a LiteViT encoder, an AutoPPN network, a standard prompt encoder, and a mask decoder as delineated in the SAM framework [13]. This configuration is visualized in Fig. 2 (a). The novel AutoPPN module was specifically designed to streamline the automated prompt task. It simultaneously regresses both bounding box prompts and point prompts in an end-to-end fashion, which significantly cuts down the inference time for the SegEvery task when compared with dense positional encoding schemes from previous research. This advancement is key to achieving real-time segmentation. A comprehensive overview of the architecture and training methodologies will be provided in the forthcoming sections.

3.2 LiteViT Architecture

Standard self-attention token mixers [7] are known for their high computational cost. In contrast, the combined CNN-Transformer hybrid structure plays an essential role in crafting lightweight Vision Transformer (ViT) networks [20, 24, 31, 38]. This hybrid balances model accuracy with computational efficiency. Inspired by efficient variations of self-attention layers in existing research, we have developed our LiteViT image encoder, beginning with a PoolFormer-S12 [38] baseline. We closely examine our architectural decisions, as detailed in Tab. 1 and illustrated in Fig. 3. As a supplement, we have also demonstrated the excellent scalability of LiteViT in Tab. 1.

We base our image encoder model on a novel building block, referred to as the LiteViT Block. The detailed architectural specifications can be found in Fig. 2 (b). To overcome the challenge of capturing local features, we incorporate multiscale pooling into our lightweight attention module. Specifically, we introduce the Multi-Scale Pooling Module (MSPM) module to enhance the receptive field at each stage of the network architecture efficiently.

Within a LiteViT Block, the input is first processed by the MSPM module, followed by a convolutional MLP (Multilayer Perceptron) module; each stage is connected via skip connections. To facilitate downsampling and adjust the output channels at each stage, we employ a dedicated module known as the Patch Merge module, which effectively acts as a stem convolutional layer.

3.3 AutoPPN

The standard approach of using dense positional encoding for prompts may not be suitable for real-time segmentation tasks due to the processing time required. To enhance the inference performance of the SegEvery task, we introduce the AutoPPN module, the architecture is detailed in Fig. 2 (c).

It has been well-established that representing objects by a single point located at the center of their bounding box is a straightforward and efficient technique

Table 1: LiteViT Attention block Ablation Studies. All models are trained and benchmarked using the same settings described in Sec. 4.2, with unified input resolution 640×640 . As a supplementary addition, we have meticulously documented the performance metrics of LiteViT, specifically its floating-point operations (FLOPs), latency, and evaluation metrics, when scaled to 2 and 3 times the parameter volume of the baseline LiteViT network. Notably, this scaling achieves impressive mAP scores of 56.9% and 58.1% for 1-box prompt segmentation on the COCO dataset, respectively. These results underscore the commendable scalability of LiteViT.

Architectural	Attention Block Choices	Params ↓ (M)	MACs ↓ (G)	Backbone Latency ↓ (ms/1-batch)	COCO 1-box prompt mAP ↑ (%)	Stages 1-4 Embedding_dims
PoolFormer-S12 (Baseline [38])	Fig. 3 (1)	11.9	45.2	30	55.1	[64, 128, 320, 512]
PoolFormer-S12-Tiny (Embedding_dims pruned)	Fig. 3 (1)	0.54	4.2	7.4	50.9	[32, 64, 96, 128]
		1.15	10.8	8.4	52.7	
	Fig. 3 (2)	1.15	10.8	8.1	53.1	
	Fig. 3 (3)	1.15	10.8	8.6	54.0	[64, 96, 128, 256]
	Fig. 3 (4)	1.16	10.9	8.7	55.2	
LiteViT (ours, Sec. 3.2)	Fig. 3 (5)	1.16 (-90%)	10.9 (-76%)	8.6 (3.4x up ↑)	55.3 (+0.2% ↑)	
LiteViT (~2× parameters)	Fig. 3 (5)	2.19	22.3	12.4	56.9	[96, 128, 192, 384]
LiteViT (~3× parameters)	Fig. 3 (5)	3.63	38.0	15.9	58.1	[128, 160, 256, 512]

[14, 48]. Building on this concept, our AutoPPN framework predicts both prompt points and bounding boxes in an end-to-end manner from the output feature map. The corresponding loss is composed of two elements: confidence in the point prompt and accuracy in the bounding box regression. We have implemented three significant modifications to refine our approach, which are detailed below:

- (1) We have enhanced our network by replacing the basic stem convolution network with a more sophisticated stem MSPM network. This updated network effectively integrates multiscale spatial information, which significantly boosts the detection recall for large-scale objects or entities, such as the sky, buildings, and water bodies.
- (2) To estimate the confidence of point prompts, we have incorporated the use of distance transforms. This facilitates the calculation of the distance between a point and its corresponding mask, as depicted in Fig. 4. In cases where a point falls within multiple masks, we select the one with the smallest area. The pseudo code for this procedure is provided in supplementary materials. Unlike the Gaussian-based method referenced in [48], our technique enables the creation of a softened pointwise ground-truth distinction between foreground and background. Additionally, our method prioritizes the identification of the most central points of objects or entities rather than simply focusing on the center of their bounding boxes. This modification helps to alleviate the ambiguity present in scenarios involving unclear point prompts. When computing loss, we have opted for a hard mining Mean Squared Er-

Table 2: AutoPPN Ablation Studies. All models are trained and benchmarked using the same settings described in Sec. 4.2.

PPN Architectural Choices Sec. 3.3	Stem Conv → MSPM(Item 1)	New GT & Loss(Item 2)	Object Grouping(Item 3)	Mask AR@1000(%)
Baseline = Stem Conv + Focal-Loss/Smooth-L1 Loss + w/o Object Grouping	-	-	-	48.8
1 improvement strategy	✓	-	-	49.5
	-	✓	-	50.1
	-	-	✓	49.7
2 improvement strategies	✓	✓	-	51.4
	✓	-	✓	52.3
	-	✓	✓	51.1
AutoPPN (all improvement strategies)	✓	✓	✓	53.0

Table 3: Comparison of speed and accuracy acceleration of AutoPPN in SOTA models. To ensure a fair comparison, we conducted AutoPPN training on both SAM and MobileSAM using the same data and training parameters.

Model	Sampling Strategy	SegEvery Time↓ (ms)	COCO AR@1000 ↑ (%)
SAM-B [13]	Grid-Search (32 x 32)	2084	55.1
SAM-B + AutoPPN	AutoPPN(256 points)	120 (17.3x up)	54.7
MobileSAM [39]	Grid-Search (32 x 32)	2500	53.2
MobileSAM + AutoPPN	AutoPPN(256 points)	115 (21.7x up)	52.6
LiteViT(ours)	Grid-Search (32 x 32)	1320	53.4
LiteViT + AutoPPN (ours, Lite-SAM)	AutoPPN(256 points)	80 (16.5x up)	52.8

ror (MSE) Loss instead of the commonly used Focal-Loss for point prompt estimation.

Meanwhile, the Smooth-L1 Loss remains the same as stated in [14] for bounding box regression. It is also important to note that any unlabeled regions are excluded from the loss computation process. During inference, we only apply point-based non-maximum suppression (Point-NMS) and adhere to the practice of selecting the Top N points, as described in [48].

- (3) During the end-to-end regression stage, we divided the target masked regions into three groups based on the size of their bounding rectangles: large ($\max(\frac{h}{H_{img}}, \frac{w}{W_{img}}) \geq 0.25$), medium ($\max(\frac{h}{H_{img}}, \frac{w}{W_{img}}) \in (0.05, 0.25)$), and small ($\max(\frac{h}{H_{img}}, \frac{w}{W_{img}}) \leq 0.05$). Separate loss calculations were performed for each group. The three improvements greatly enhance the performance, see Tab. 2. We denote

$$L_{ppn} = L_{H-MSE} + L_{S-L1} \quad (1)$$

the total loss of PPN regression, where L_{H-MSE} refers to the hard mining MSE Loss and L_{S-L1} the Smooth L1 Loss for box regression.

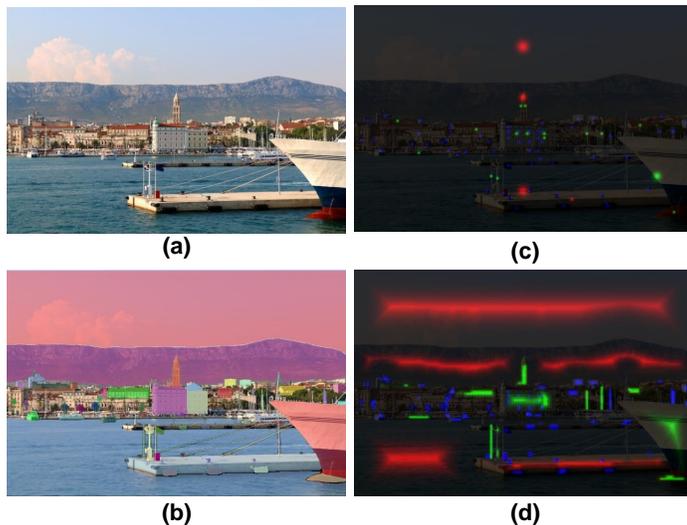


Fig. 4: We compare two methods of generating pointwise foreground/background labels within an image (sa_3196.jpg) from SA-1B [13] (a). All the masks are visualized as shown in (b). The pointwise labels generated by large, medium, small masks, are visualized with red, green and blue color, respectively. Comparing with bounding box center with gaussian kernel approach (c), distance transform approach (d) provides a more satisfactory result with less ambiguity.

3.4 Total loss

For the comprehensive training of Lite-SAM, we incorporate the mask loss, which combines the original Focal-Loss [18] and Dice-Loss [18] from SAM [13]. This combination quantifies the pixel-wise alignment between the predicted mask and the ground truth. Additionally, a mean squared error loss measures the discrepancy between the IoU prediction and the intersection of the predicted mask with the ground truth mask. The mask loss is formally expressed as:

$$L_{mask} = \lambda_f L_{focal} + \lambda_d L_{dice} + \lambda_i L_{iou} \quad (2)$$

With Eq. (1) and Eq. (2), the total training loss is defined by $L_{total} = L_{ppn} + L_{mask}$.

4 Experiments

In this section, we present a comprehensive evaluation of our proposed Lite-SAM framework. To ensure a rigorous and equitable comparison, we utilized the same evaluation protocol as employed by other SOTA methods.

4.1 Datasets

Public data. Lite-SAM was trained on SA-1B [13]. We selected three public datasets to assess the zero-shot capabilities of our model: MSCOCO 2017 [19], LVIS [8], and BSDS500 [22].

4.2 Implementation details

Hyperparameters. We developed Lite-SAM using the PyTorch framework and trained it on 128 NVIDIA A40 GPUs, achieving an aggregate batch size of 256. The model underwent training from scratch, without the use of any pre-existing weights. With just 18% of the SA-1B dataset [13] dataset, our model demonstrated impressive results. We utilized the Adam optimizer with an initial learning rate of $4e-5$ and completed the training process in 4 epochs, which took a total of 50 hours. Throughout the training, all images were resized to 640×640 pixels. Concurrently, it is essential to recognize that the choice of using 18% of the SA-1B data was based on a trade-off between training time and accuracy. The ablation study results regarding the selection of training data size and backbones, are presented in our supplementary material.

For supervising the guided prompt predictions, our loss function, AutoPPN-Loss, included a mix of hard mining MSE Loss for pointwise objectness and L1-Loss for prompt box regression, with a respective ratio of 2:1. For the mask prediction component, we employed a blended loss function combining Focal-Loss [18] and Dice-Loss [18] with a weighting of 10:1. In addition, a mean squared error loss was introduced to estimate the accuracy of the intersection over union (IoU) prediction compared to the ground truth mask alignment. The mean Intersection over Union (mIoU) metric was selected as our evaluation standard for segmentation performance.

4.3 Comparison of speed and accuracy acceleration of AutoPPN in SOTA models

As shown in Tab. 3, the integration of AutoPPN leads to appreciable improvements in SegEvery time, while preserving the recall rates. Specifically, with the SAM-B [13] model, AutoPPN achieves a speedup of 17.3-fold relative to that of a conventional Grid Search method. For MobileSAM [39], the speedup stands at 21.7-fold. When applied to our LiteViT model, AutoPPN manages a speedup of 16.5-fold, reducing the SegEvery time to less than 80 ms, a significant milestone. These experimental results compellingly illustrate the efficiency of AutoPPN in addressing the speed bottleneck typically associated with Grid Search.

4.4 Comparison with SOTA lightweight models on COCO 2017

In Tab. 4, we conducted detailed experimental comparisons and found that different backbone models exhibit varying levels of performance across each metric. Among these models, our proposed LiteViT (which serves as our backbone

Table 4: Comparison with SOTA lightweight backbone models on COCO.

(1) All experiments are conducted based on open-source models and trained from scratch with same implementation described in Sec. 4.2. (2) The time tests for SegAny were conducted on an A40 GPU, while keeping the same environment. (3) In supplementary material, we present the classification capability of LiteViT on the ImageNet dataset to serve as a pre-trained model.

Backbone Model	COCO				SegAny time (ms)	Params (M)	MACs (G)	Input
	1-box (mAP)	1-box (mIoU)	1-point (mAP)	1-point (mIoU)				
MobileNetv2 [10, 29]	48.2%	69.6%	23.5%	48.5%	5.1	1.89	4.0	640
ShuffleNetv2 [21, 45]	49.2%	70.6%	24.2%	49.6%	5.8	1.52	5.1	640
MobileViT [24]	51.4%	72.1%	26.2%	53.8%	11.9	5.57	13.7	640
EfficientViT [20]	54.1%	73.6%	28.4%	53.1%	18.0	30.73	106.5	640
Fast ViT [31]	52.3%	70.0%	28.0%	48.0%	7.5	3.98	7.0	640
TinyViT [33]	54.0%	73.4%	26.5%	52.6%	17	6.07	36.6	640
LiteViT(ours,backbone)	55.8%	74.8%	32.9%	55.3%	7.9	1.16	10.2	640

model) outperformed the other lightweight backbone models in all metrics and achieved the best results. Specifically, LiteViT reached a performance of 55.8% for 1-box mAP, 74.8% for 1-box mIoU, 32.9% for 1-point mAP, and 55.3% for 1-point mIoU. Furthermore, LiteViT demonstrated clear advantages in terms of inference time, model parameter count, and computational load compared to other models. Overall, our experimental results establish that LiteViT, serving as our backbone model, is an exceptional lightweight backbone option, achieving SOTA performance on the COCO dataset. It also offers faster inference times and a relatively smaller model size. These results confirm its effectiveness and its competitive edge.

4.5 Comparison with SOTA Algorithms on COCO and LVIS validation sets using AP and mIoU metric

In Tab. 5, we conducted detailed experimental comparisons among the latest algorithms from the SAM series on the COCO and LVIS datasets. The results show that the SAM-H [13] model achieved superior performance, obtaining the highest metrics on both datasets. Specifically, its 1-box mIoU reached 76.5% on the COCO dataset and 75.3% on the LVIS dataset. In contrast, the 1-point mIoU scored 57.4% for COCO and 56.4% for LVIS.

The experimental results classify the models into two categories based on their size: large models with a parameter count exceeding 10M and lightweight models with fewer than 10M parameters. Among the large models are SAM-B/L/H [13], Semantic-SAM-L [15], Fast-SAM [46], EfficientViT-L0-SAM [4], MobileSAM-v2 [40], and EfficientViT-L1-SAM [4]. The lightweight category includes Mobile-SAM [39], Edge-SAM [47], and Lite-SAM.

Lite-SAM, a lightweight model, achieved a 1-box mIoU performance that surpassed SAM-B [13] by 1.3%, with significantly fewer parameters and com-

Table 5: Zero-Shot Image Segmentation Results on MSCOCO 2017 and LVIS validation sets using mIoU and AP metric. (1) Note that the 1-box prompt result is not available in Semantic-SAM-L’s released code (similarly hereinafter). (2) “r640” means the input resolution is 640×640. (3) Note: We adopted the method in reference [4], which is entirely based on the ground truth (GT) box for predictions, instead of using ViTDet-H’s results as prompts [13].

Model	MSCOCO(mIoU) ↑		LVIS(mIoU) ↑		MSCOCO↑				LVIS ↑			
	1-box (%)	1-point (%)	1-box (%)	1-point (%)	AP (%)	AP ^S (%)	AP ^M (%)	AP ^L (%)	AP (%)	AP ^S (%)	AP ^M (%)	AP ^L (%)
SAM-B [13]	75.0	52.2	73.5	55.2	56.6	47.4	60.3	68.0	61.1	50.3	71.6	76.7
SAM-L [13]	76.4	56.8	75.0	55.8	59.4	48.8	65.3	72.1	64.9	53.5	76.1	81.9
SAM-H [13]	76.5	57.4	75.3	56.4	59.8	49.4	63.8	71.9	65.2	53.6	76.5	82.1
Semantic-SAM-L [15]	N/A	54.7	N/A	34.8	N/A							
Fast-SAM [46]	72.8	50.2	67.3	46.8	47.5	37.9	48.1	56.4	43.8	35.1	45.6	59.7
EfficientViT-L0-SAM [4]	74.5	51.3	73.1	52.9	56.1	44.3	59.7	70.8	59.8	46.8	70.2	80.1
EfficientViT-L1-SAM [4]	75.2	51.5	73.9	54.8	57.1	45.4	60.8	71.5	61.4	48.0	72.5	81.6
Mobile-SAM-v2* [40]	72.8	50.5	67.7	42.4	51.4	41.6	55.1	64.1	52.8	42.2	63.2	69.6
Mobile-SAM [39]	72.8	50.5	67.7	42.4	51.4	41.6	55.1	64.1	52.8	42.2	63.2	69.6
Edge-SAM [47]	74.0	51.9	69.4	43.8	52.5	42.7	56.0	65.3	54.1	43.5	63.9	70.7
Lite-SAM(ours, r640)	74.8	55.8	73.2	54.4	55.8	46.7	59.6	69.6	58.4	45.9	66.9	77.5
Lite-SAM(ours, r1024)	76.3	56.9	75.7	57.3	56.5	47.4	61.0	70.7	60.7	49.3	71.9	79.8

computational demands. Lite-SAM also outperformed Mobile-SAM [39] and Edge-SAM [47] in terms of mIoU metrics. Regarding the Average Precision (AP) metric, SAM-H [13] still recorded the highest values, with an AP of 59.8% on COCO and 65.2% on LVIS. Lite-SAM performed better than Mobile-SAM [39] and Edge-SAM [47] but slightly lower than EfficientViT-L1-SAM [4] in terms of AP.

Overall, these experiments highlight the outstanding performance of our Lite-SAM algorithm, confirming its effectiveness and competitive edge in the field.

4.6 Comparison with SOTA Algorithms Complexity and SegEvery Speed Evaluation

In Tab. 6, we have presented detailed experimental comparisons of the latest algorithms in the SAM series. The results reveal substantial variations in parameter size, Multiply-Accumulate Operations (MACs), and SegEvery runtime across the different algorithmic models. The Sampling Strategy is categorized into three types: Grid-Search, Object-aware, and AutoPPN.

SAM-B [13] boasts a parameter size of 90M, MACs of 371G, and a SegEvery runtime of 2.1s. The lightweight models, namely Mobile-SAM [39] and Edge-SAM [47], have parameter sizes of 9.7M and offer 39.6/23.4G MACs, respectively. Mobile-SAM-v2 [40] implements the Object-aware strategy, leveraging YOLOv8 [12] to perform box and point detection in advance, characterizing it as a two-stage algorithm.

Our newly developed Lite-SAM is designed as an end-to-end algorithm with a minimal parameter size of only 4.2M. Impressively, it has reduced the SegEvery runtime to a mere 80ms for the first time. This model not only demonstrates the best performance in regards to parameter size and MACs, but also in SegEvery inference time, which underlines its efficiency and competitive edge.

Table 6: Comparison with SOTA Algorithms: Model Complexity, SegEvery Speed, and Mask AR@1000 metric Evaluation on COCO2017. (1) “r640” means the input resolution is 640×640 . For specific calculation details, please refer to the code examples in the supplementary materials. (2) * The “Fast-SAM” does not have true interactive segmentation via point or box prompts. It employs heuristic rules for post-process object selection, a method that aligns marginally with the SAM principles. Therefore, the Fast-SAM algorithm and the other SAM series algorithms are completely different, making a comparison between them of no value. (3) ** The “Mobile-SAM-v2” paper does not include statistics on the parameter size and computational complexity of the Object-aware model, so the data has been re-estimated (43M = Yolov8(33.3M) + MobileSAM(9.7M)).

Model	Params ↓	MACs ↓	SegEvery Time ↓	Mask AR @1000 (%)	Sampling Strategy	Train Strategy	Year
SAM-B [13]	90M	371G	2.1s	55.1	Grid-Search (32 x 32)	pretrain on MAE	2023
SAM-L [13]	308M	1.3T	3.3s	56.6	Grid-Search (32 x 32)	pretrain on MAE	2023
SAM-H [13]	635M	2.7T	3.5s	58.7	Grid-Search (32 x 32)	pretrain on MAE	2023
Semantic-SAM [15]	202M	1.4T	2.6s	55.0	Grid-Search (32 x 32)	from scratch	2023
EfficientViT-L0-SAM [4]	31M	109G	1.7s	56.7	Grid-Search (32 x 32)	from scratch	2023
Fast-SAM [46] *	72.2M	443G	0.04s	53.3	Post-Process Object Selection	pretrain on YOLOv8	2023
Mobile-SAM [39]	9.7M	39.6G	2.5s	53.2	Grid-Search (32 x 32)	distillation	2023
Edge-SAM [47]	9.7M	23.4G	1.6s	51.9	Grid-Search (32 x 32)	distillation	2024
Mobile-SAM-v2 [40] **	43M	470G	0.13s	53.6	Object-Aware	distillation	2024
Lite-SAM(r640)	4.2M	12.7G	0.08s	52.8	AutoPPN(256 points)	from scratch	2024
Lite-SAM(r1024)	4.2M	32.5G	0.1s	54.1	AutoPPN(256 points)	from scratch	2024

To demonstrate that Lite-SAM delivers results on par with other SAM architectures, while also showcasing its exceptional performance relative to other lightweight SAMs, we have included visual qualitative assessments for the “SegEvery” and “SegAny” tasks as supplementary material. These illustrations underscore the effectiveness of the Lite-SAM approach.

Table 7: Zero-shot transfer to edge detection on BSDS500. Evaluation data of other methods is from [13].

Method	Year	ODS	OIS	AP	R50
HED [34]	2015	0.788	0.808	0.840	0.923
EDETR [25]	2022	0.840	0.858	0.896	0.930
<i>zero-shot transfer methods:</i>					
Sobel filter	1968	0.539	-	-	-
Canny [5]	1986	0.600	0.640	0.580	-
Felz-Hutt [6]	2004	0.610	0.640	0.560	-
SAM-H [13]	2023	0.768	0.786	0.794	0.928
Fast-SAM [46]	2023	0.750	0.790	0.793	0.903
Lite-SAM(ours)	2023	0.761	0.788	0.793	0.919

4.7 Zero-Shot Edge Detection

We assessed the zero-shot edge detection capability of Lite-SAM on the BSDS500 dataset [1, 22], following the experimental parameters established by SAM [13] and Fast-SAM [46]. As shown in Tab. 7, Lite-SAM R50 attains a metric score of 0.919, slightly behind SAM’s 0.928 and surpassing Fast-SAM’s 0.903.

5 Conclusion

In this paper, we propose an end-to-end lightweight algorithm called Lite-SAM, which aims to address the high computational complexity issue of the SegEvery model in the SAM series. Lite-SAM consists of the LiteViT module and the AutoPPN module, enabling modular deployment. Our algorithm achieves a 16-fold speedup in inference time while maintaining a minimal decrease in accuracy compared to the SegEvery mode. Through extensive experimental tests, we demonstrate that our approach satisfies the requirements of efficient and resource-friendly segmentation algorithms, providing possibilities for practical applications in various fields.

Acknowledgements

This work is supported by Zhejiang Dahua Technology Co., Ltd. and Zhejiang University.

References

1. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence* **33**(5), 898–916 (2010) [14](#)
2. Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., et al.: On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021) [2](#)
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020) [2](#)
4. Cai, H., Li, J., Hu, M., Gan, C., Han, S.: Efficientvit: Multi-scale linear attention for high-resolution dense prediction [3](#), [11](#), [12](#), [13](#)
5. Canny, J.: A computational approach to edge detection (1986) [13](#)
6. Dongcai, S.: Efficient graph based image segmentation. *Image Processing* [13](#)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020) [4](#), [6](#)

8. Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5356–5364 (2019) [10](#)
9. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020) [2](#)
10. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for mobilenetv3. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1314–1324 (2019) [4](#), [11](#)
11. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017) [4](#)
12. Jocher, G., Chaurasia, A., Qiu, J.: Yolo by ultralytics (2023), <https://github.com/ultralytics/ultralytics> [3](#), [12](#)
13. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023) [2](#), [3](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
14. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints (2019) [7](#), [8](#)
15. Li, F., Zhang, H., Sun, P., Zou, X., Liu, S., Yang, J., Li, C., Zhang, L., Gao, J.: Semantic-sam: Segment and recognize anything at any granularity. arXiv preprint arXiv:2307.04767 (2023) [3](#), [11](#), [12](#), [13](#)
16. Li, J., Xia, X., Li, W., Li, H., Wang, X., Xiao, X., Wang, R., Zheng, M., Pan, X.: Next-vit: Next generation vision transformer for efficient deployment in realistic industrial scenarios. arXiv preprint arXiv:2207.05501 (2022) [4](#)
17. Li, Y., Yuan, G., Wen, Y., Hu, J., Evangelidis, G., Tulyakov, S., Wang, Y., Ren, J.: Efficientformer: Vision transformers at mobilenet speed. Advances in Neural Information Processing Systems **35**, 12934–12949 (2022) [4](#)
18. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection (2018) [9](#), [10](#)
19. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014) [10](#)
20. Liu, X., Peng, H., Zheng, N., Yang, Y., Hu, H., Yuan, Y.: Efficientvit: Memory efficient vision transformer with cascaded group attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14420–14430 (2023) [4](#), [6](#), [11](#)
21. Ma, N., Zhang, X., Zheng, H.T., Sun, J.: Shufflenet v2: Practical guidelines for efficient cnn architecture design. In: Proceedings of the European conference on computer vision (ECCV). pp. 116–131 (2018) [4](#), [11](#)
22. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001. vol. 2, pp. 416–423. IEEE (2001) [10](#), [14](#)
23. Mehta, S., Rastegari, M.: Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. arXiv preprint arXiv:2110.02178 (2021) [4](#)
24. Mehta, S., Rastegari, M.: Mobilevit: Light-weight, general-purpose, and mobile-friendly vision transformer (2022) [6](#), [11](#)
25. Pu, M., Huang, Y., Liu, Y., Guan, Q., Ling, H.: Edter: Edge detection with transformer (2022) [13](#)

26. Qiao, Y., Munir, M.S., Adhikary, A., Le, H.Q., Raha, A.D., Zhang, C., Hong, C.S.: Mp-fedel: Multi-prototype federated contrastive learning for edge intelligence. arXiv preprint arXiv:2304.01950 (2023) [2](#)
27. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training (2018) [2](#)
28. Ren, T., Liu, S., Zeng, A., Lin, J., Li, K., Cao, H., Chen, J., Huang, X., Chen, Y., Yan, F., Zeng, Z., Zhang, H., Li, F., Yang, J., Li, H., Jiang, Q., Zhang, L.: Grounded sam: Assembling open-world models for diverse visual tasks (2024) [4](#)
29. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018) [4](#), [11](#)
30. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: International conference on machine learning. pp. 10347–10357. PMLR (2021) [4](#)
31. Vasu, P.K.A., Gabriel, J., Zhu, J., Tuzel, O., Ranjan, A.: Fastvit: A fast hybrid vision transformer using structural reparameterization. arXiv preprint arXiv:2303.14189 (2023) [4](#), [6](#), [11](#)
32. Wang, X., Zhang, X., Cao, Y., Wang, W., Shen, C., Huang, T.: Seggpt: Segmenting everything in context (2023), <https://arxiv.org/abs/2304.03284> [4](#)
33. Wu, K., Zhang, J., Peng, H., Liu, M., Xiao, B., Fu, J., Yuan, L.: Tinyvit: Fast pretraining distillation for small vision transformers. In: European Conference on Computer Vision. pp. 68–85. Springer (2022) [4](#), [11](#)
34. Xie, S., Tu, Z.: Holistically-nested edge detection (2015) [13](#)
35. Xiong, Y., Varadarajan, B., Wu, L., Xiang, X., Xiao, F., Zhu, C., Dai, X., Wang, D., Sun, F., Iandola, F., Krishnamoorthi, R., Chandra, V.: EfficientSAM: Leveraged masked image pretraining for efficient segment anything (2023) [3](#)
36. Yang, Y., Wu, X., He, T., Zhao, H., Liu, X.: Sam3d: Segment anything in 3d scenes (2023), <https://arxiv.org/abs/2306.03908> [4](#)
37. Yu, T., Feng, R., Feng, R., Liu, J., Jin, X., Zeng, W., Chen, Z.: Inpaint anything: Segment anything meets image inpainting (2023), <https://arxiv.org/abs/2304.06790> [4](#)
38. Yu, W., Luo, M., Zhou, P., Si, C., Zhou, Y., Wang, X., Feng, J., Yan, S.: Metaformer is actually what you need for vision (2022) [6](#), [7](#)
39. Zhang, C., Han, D., Qiao, Y., Kim, J.U., Bae, S.H., Lee, S., Hong, C.S.: Faster segment anything: Towards lightweight sam for mobile applications. arXiv preprint arXiv:2306.14289 (2023) [3](#), [8](#), [10](#), [11](#), [12](#), [13](#)
40. Zhang, C., Han, D., Zheng, S., Choi, J., Kim, T.H., Hong, C.S.: MobileSAMv2: Faster segment anything to everything (2023) [3](#), [11](#), [12](#), [13](#)
41. Zhang, C., Zhang, C., Li, C., Qiao, Y., Zheng, S., Dam, S.K., Zhang, M., Kim, J.U., Kim, S.T., Choi, J., et al.: One small step for generative ai, one giant leap for agi: A complete survey on chatgpt in aigc era. arXiv preprint arXiv:2304.06488 (2023) [2](#)
42. Zhang, C., Zhang, C., Zheng, S., Qiao, Y., Li, C., Zhang, M., Dam, S.K., Thwal, C.M., Tun, Y.L., Huy, L.L., et al.: A complete survey on generative ai (aigc): Is chatgpt from gpt-4 to gpt-5 all you need? arXiv preprint arXiv:2303.11717 (2023) [2](#)
43. Zhang, C., Zhang, K., Pham, T.X., Niu, A., Qiao, Z., Yoo, C.D., Kweon, I.S.: Dual temperature helps contrastive learning without many negative samples: Towards understanding and simplifying moco. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14441–14450 (2022) [2](#)

44. Zhang, C., Zhang, K., Zhang, C., Pham, T.X., Yoo, C.D., Kweon, I.S.: How does simsiam avoid collapse without negative samples? a unified understanding with self-supervised contrastive learning. arXiv preprint arXiv:2203.16262 (2022) [2](#)
45. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6848–6856 (2018) [4](#), [11](#)
46. Zhao, X., Ding, W., An, Y., Du, Y., Yu, T., Li, M., Tang, M., Wang, J.: Fast segment anything. arXiv preprint arXiv:2306.12156 (2023) [3](#), [11](#), [12](#), [13](#), [14](#)
47. Zhou, C., Li, X., Loy, C.C., Dai, B.: Edgesam: Prompt-in-the-loop distillation for on-device deployment of sam (2023) [3](#), [11](#), [12](#), [13](#)
48. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019) [7](#), [8](#)
49. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Wang, J., Wang, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once (2023), <https://arxiv.org/abs/2304.06718> [4](#)