

Self-Sampling Meta SAM: Enhancing Few-shot Medical Image Segmentation with Meta-Learning

Tianang Leng*

Huazhong University of Science and Technology
 Wuhan, China

tianangl@hust.edu.cn

Kun Han

University of California, Irvine
 Irvine, California, United States

khan7@uci.edu

Yiming Zhang*

Tokyo Institute of Technology
 Tokyo, Japan

zhang.y.bl@m.titech.ac.jp

Xiaohui Xie

University of California, Irvine
 Irvine, California, United States

xhx@ics.uci.edu

Abstract

While the Segment Anything Model (SAM) excels in semantic segmentation for general-purpose images, its performance significantly deteriorates when applied to medical images, primarily attributable to insufficient representation of medical images in its training dataset. Nonetheless, gathering comprehensive datasets and training models that are universally applicable is particularly challenging due to the long-tail problem common in medical images.

To address this gap, here we present a **Self-Sampling Meta SAM** (SSM-SAM) framework for few-shot medical image segmentation. Our innovation lies in the design of three key modules: 1) An online fast gradient descent optimizer, further optimized by a meta-learner, which ensures swift and robust adaptation to new tasks. 2) A Self-Sampling module designed to provide well-aligned visual prompts for improved attention allocation; and 3) A robust attention-based decoder specifically designed for few-shot medical image segmentation to capture relationship between different slices. Extensive experiments on a popular abdominal CT dataset and an MRI dataset demonstrate that the proposed method achieves significant improvements over state-of-the-art methods in few-shot segmentation, with an average improvements of 10.21% and 1.80% in terms of DSC, respectively. In conclusion, we present a novel approach for rapid online adaptation in interactive image segmentation, adapting to a new organ in just 0.83 minutes. Code is available at <https://github.com/DragonDescentZerotsu/SSM-SAM>

*equal contribution

1. Introduction

Medical image segmentation plays a pivotal role in clinical applications, including disease diagnosis, abnormality detection and treatment planning. Historically, the segmentation of anatomical structures was performed manually by experienced physicians, a laborious and time-consuming task.

Recent advancements in deep learning offer automated segmentation tools that deliver near-human accuracy swiftly. Nevertheless, these tools need extensive training on large annotated datasets for optimal performance, which are costly and time-intensive since they require inputs from experts with extensive clinical experience. Thus, in the field of medical imaging, few-shot learning [29, 49, 60] has gained significant interest from researchers because it is able to segment accurately without extensive labeled data. In fact, to achieve optimal performance on unseen classes, few-shot learning models must excel in extracting representative features from limited data. However, current few-shot learning frameworks for medical image segmentation [15, 50, 54] predominantly pretrain their models using data from a single medical domain, often with restricted datasets, leading to limited feature extraction ability. If we can leverage the strong feature extraction abilities obtained from extensive training data of diverse domains, the model can capture more distinctive features, enhancing its adaptability to new tasks.

Recently, Segment anything model (SAM) [28] has attracted significant attention. Trained on a large segmentation dataset of over 1 billion masks across various domains of natural images, SAM has strong feature extraction and generalization abilities and can segment any object on a certain image. Given SAM's excellent zero-shot

transferability, a natural idea is to directly apply SAM for medical image segmentation to address the issue of relatively scarce medical image data [20, 46]. However, recent studies [40, 47] have shown that the performance of SAM is overall moderate and varies significantly across different datasets and different cases, demonstrating the potential promise of SAM within the context of medical images but also shows that the model cannot be applied directly with high confidence. Such phenomenon is attributed to the vast differences between natural and medical images, as SAM was primarily trained on natural images. To better adapt SAM to medical images, most prior studies focused on integrating lightweight Adapters [10, 56] or on freezing the heavy image and prompt encoders, opting to fine-tune solely the mask decoder [23, 31, 38]. However, the limited data available for unseen classes will hinder the few-shot segmentation performance of these approaches.

Driven by the desire to maximize the powerful extracted features of SAM through prompts [61], and to leverage its potential zero-shot capabilities without vast training data, we introduce the Self-Sampling Meta SAM (SSM-SAM) framework, which utilizes a Meta-learning method MAML++. Our framework consists of two main parts. The first part is a redesigned backbone that removes the original SAM’s prompt encoder and mask decoder; instead, we employ a self-sampling prompt encoder and Flexible Mask Attention Decoder (FMAD). Also, we integrate adapters into the image encoder as previous works do, allowing SAM to learn features of unfamiliar tasks. This enables SAM to have better transferability for few-shot learning and is termed **SS-SAM** (Self-Sampling SAM without MAML++). The second part, which we refer to as **SSM-SAM** (with MAML++), serves as our final framework for few-shot learning. It is a meta-learning based optimizer layered on top of this backbone to further boost SAM’s few-shot performance on medical images.

We performed experiments for few-shot learning on an abdomen CT dataset and an MRI dataset. We also utilized a fully supervised medical image segmentation task to evaluate the performance of our SS-SAM backbone (without MAML++) in comparison to prior methods on CT dataset.

Our contributions are summarized as follows:

- An effective online parameter adaptation technique optimized by a MAML++ [1] based meta-learner to enhance SAM’s generalization and few-shot learning capacities.
- A positive-negative self-sampling module that can generate aligned visual prompts to better extract the contextual relationship.
- A novel Flexible Mask Attention Decoder specifically designed for medical image few-shot segmentation.

- Our method outperformed the SOTA framework for few-shot medical image segmentation by an average of 10.21% on MICCAI15 Multi-Atlas Abdomen Labeling challenge dataset [4] and 1.80% on ISBI 2019 Combined Healthy Abdominal Organ Segmentation Challenge [27] in terms of DSC.

Each module we use is plug-and-play, aiming to facilitate the deployment of fast and robust online image segmentation systems across various industries and making our model a baseline for improving performance where foundational models like SAM have struggled in the past.

2. Related Work

Foundation Models. A foundation model is essentially a large pre-trained model, typically developed using self-supervised learning across diverse datasets. This allows it to be fast adapted to specific tasks through mechanisms by fine-tuning or in-context learning. At present, foundation models have reached a level of maturity in the NLP domain, with models like BERT [13], GPT-3 [7]. In the realm of computer vision, the Segment Anything Model (SAM) [28], pre-trained on 11M images with 1B masks, first introduced a foundation model for image segmentation. This model can interpret a variety of prompts given by users, such as points, boxes, masks, and texts, to generate user-specified segmentations and is adept at segmenting any object within an image. However, despite its zero-shot generalization capabilities, SAM has demonstrated sub-optimal performance across multiple downstream tasks [26]. This observation has stimulated ongoing research efforts to enhance SAM’s performance in these tasks [18, 38, 56, 63].

Adapters. Adapters have emerged as a powerful tool in the realm of transfer learning for NLP and computer vision. Introduced as a lightweight alternative to full model fine-tuning, adapters allow for model customization while preserving the pre-trained parameters [22]. Instead of updating all the parameters in the model, adapter methods insert small bottleneck layers in the model architecture, which are then fine-tuned while the original model parameters remain unchanged. Recently, the ViT-Adapter [12] has been developed to equip a standard ViT [16] with the capability to perform a variety of downstream tasks. In a further development, an Explicit Visual Prompting (EVP) [36] technique has facilitated the integration of explicit visual cues into the Adapters. In this work, We integrate such Adapters [10] into the image encoder of SAM to avoid training a large amount of parameters.

Visual Prompts. The application of visual prompts in computer vision tasks has its roots in interactive segmentation, a methodology that necessitates user input, such as clicks [11, 25, 34, 52, 58], bounding boxes [43, 57], or scribbles [5, 6, 33], or convert spatial queries to masks and feed them

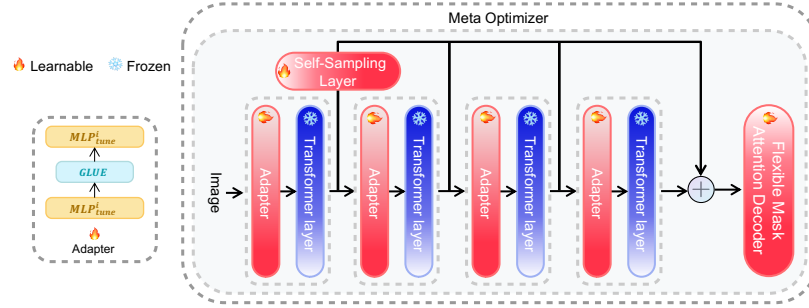


Figure 1. Overview of our proposed framework for SSM-SAM. SSM-SAM backbone consists of three main components: a) Injecting Adapters into the encoder of SAM to quickly acquire task-related information. b) Using the Self-Sampling Module to replace the position encoding based prompt encoder to strengthen the relationship between the feature and the prompt. c) Flexible Mask Attention Decoder (FMAD) to enhance boundary information and refine the final generated mask map. On top of the backbone, we employ a meta-learning based online optimizer.

into the image backbone [35] to assist the algorithm in accurately delineating object boundaries. These visual prompts augment the segmentation process, yielding more accurate and dependable outcomes [53]. However, such form of visual prompts either struggle to adapt to unseen prompts [28] or might be heavy in applications as each interaction necessitates processing the image through the feature extractor [35]. Recently, visual sampler has been introduced to transform all non-textual queries into visual prompts that reside within the same visual embedding space to address these limitations [65]. However, no prior research has endeavored to employ this approach for generating prompts using the SAM trained on an extensive image corpus. Here we seek to bridge this research gap.

Model Agnostic Meta Learning (MAML). One continuous challenge of few-shot medical image segmentation is the distribution mismatch between training and testing datasets, particularly the long tail problem. Fortunately, MAML [17] and MAML++ [1] offers meta-learning frameworks tailored to combat this issue. It is elegantly simple yet can find suitable model-agnostic initialization parameters that are trained through various tasks and can quickly adapt to new tasks. Many previous works like [39, 48, 59] use meta-learning to address different problems in medical image segmentation. Yet, MAML++ framework has been primarily focused on vision tasks like classification and recognition. All these previous works inspired us to utilize MAML++ to enhance foundational models’ few-shot learning ability.

3. Methodology

3.1. Overview of SSM-SAM architecture

Fig. 1 depicts the architecture of our few-shot learning framework. We keep the image encoder of SAM frozen and add a learnable **Adapter** layer [36] to each of the transformer layers for parameter-efficient training. After pass-

ing through the image encoder, we perform a self-sampling operation on the first image embedding to update it. Afterward, we feed the four output image embeddings into our Flexible Mask Attention Decoder (FMAD) to refine these embeddings and generate the predicted mask. On top of this backbone, we implement a MAML++ based meta-learner to search for optimal initialization parameters for rapid adaptation to different organs.

3.2. Few-shot Online Optimizer

Segmentation models’ clinical deployment has historically been hindered by a distribution mismatch between training and testing datasets. Given the impossibility of collecting comprehensive representative data, we have innovated in our training approach. Rather than creating a universal offline model, we designed our model to recognize and adapt to new data types, ensuring it remains relevant for unseen images. In our few-shot medical image segmentation scenario, we consider a distribution over organs, denoted as $p(O)$, to which we aim our model to adapt. We represent our model as f_{θ} with parameter θ , which is trained on images of different organs, I_{o_i} , following the distribution $p(I_{o_i})$, where o_i signifies the i^{th} organ. The online optimizer optimizes the parameters via back-propagating steps as follows:

$$\theta'_i \leftarrow \theta_i - \alpha \nabla_{\theta'_i} L_{o_i}(f_{\theta'_i}) \quad (1)$$

where α is the learning rate and θ'_i represents the model parameters adapted to fit the i^{th} organ. We employ balanced cross-entropy and IoU loss to supervise our network so L_{o_i} can be expressed as:

$$L_{o_i} = L_{bce} + L_{iou} \quad (2)$$

The overview of the online optimizer is outlined in Algorithm 1.

Algorithm 1 SSM-SAM Online Optimizer

Input: K image-mask pairs $P_i = \{(I_i, M_i)_k\}$ of the target organ o_i
Input: trained parameters: θ
Require: learning rate α ; number of online optimization steps S

- 1: $\theta'_i \leftarrow \theta$
- 2: **for** s from 1 to S **do**
- 3: **for** $(I_i, M_i)_k$ in P_i **do**
- 4: Update parameters with gradient descent:
 $\theta'_i \leftarrow \theta'_i - \alpha \nabla_{\theta'_i} L_{O_i}(f_{\theta'_i}((I_i, M_i)_k))$
- 5: **end for**
- 6: **end for**

Output: updated parameters: θ'_i

3.3. Meta-learning

We employ meta-learning for our model training to guarantee that the initial parameters θ exhibit robust generalization capabilities, facilitating rapid online optimization. MAML++ [1] is an ideal strategy for our purpose, as it is designed to learn suitable initial model parameters optimized for swift adaptation to new tasks. The complete algorithm is detailed in Algorithm 2. Note that the outermost for-loop

Algorithm 2 SSM-SAM Offline Meta-learner

Input: organ set O ; initial weight and learning rate θ, β
Require: number of epochs E ; number of within organ optimization steps S ; online learning rate α

- 1: **for** e from 1 to E **do**
- 2: Sample T organs $\{o_1, o_2, \dots, o_T\}$ from O
- 3: **for** i from 1 to T **do**
- 4: $\theta'_i \leftarrow \theta$
- 5: Sample K image-mask pairs $P_i = \{(I_i, M_i)_k\}$ of the target organ o_i
- 6: $\theta'_i \leftarrow$ **Online Optimizer**($S, P_i, \theta'_i, \alpha$)
- 7: **end for**
- 8: Resample K images-mask pairs P'_i for each o_i
- 9: $\beta \leftarrow$ **CosineAnnealingLR**(β, e)
- 10: $\theta \leftarrow \theta - \beta \nabla_{\theta} \frac{1}{T} \sum_i^T L_{o_i}(f_{\theta'_i}(P'_i))$
- 11: **end for**

Output: updated parameters: θ

is the meta-learner, which is defined as follows:

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \frac{1}{T} \sum_i^T L_{o_i}(f_{\theta'_i}) \quad (3)$$

where i is the i^{th} organ. T is the number of organs in a task batch for optimizing the meta-learner. We incorporate the cosine annealing learning rate as recommended in [1] to enhance our model’s efficiency in fitting the training set.

3.4. Adapted Image Encoder

The image encoder of SAM uses a Vision Transformer (ViT) [16], which is pre-trained with MAE [19]. When given an image of any size, it is essential to first resize it to a resolution of 1024 x 1024. The ViT then processes this resized image to generate an image embedding with dimensions $C \times H \times W$. For this research, we chose the ViT-B variant (SAM-b), (where $C = 256, H = 64, W = 64$). To enable efficient learning with faster updates and address the issue of excessive GPU memory usage, we keep the image encoder frozen and inject a trainable **Adapter** to each of the 12 transformer layers as mentioned before. We only train the parameters within **Adapters**, which are tasked with learning task-specific knowledge and low-level structural information from the features extracted from image embeddings. For the k -th Adapter layer, we take patch embedding F^k as input and obtain updated parameters P^k

$$P^k = \text{MLP}_{\text{up}}(\text{GELU}(\text{MLP}_{\text{tune}}^k(F^k))) \quad (4)$$

where $\text{MLP}_{\text{tune}}^k$ refers to a linear layer within each Adapter for generating distinct prompts. MLP_{up} is an up-projection layer shared across all Adapters designed to match the dimension of transformer features. P^k represents the output associated with each transformer layer.

Also, we divided image encoder to four sub-blocks, which is carried out with the intent of deriving multi-level information [32]. After passing through the image encoder, these four image embeddings of size $B \times 256 \times 64 \times 64$ are fed into the subsequent module.

3.5. Self-Sampling Module

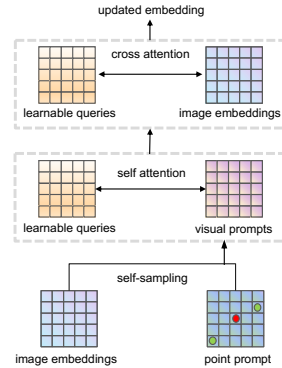


Figure 2. Structure of queries and prompt interaction during training and evaluation.

We employ a *positive-negative attention self-sampling* module inspired by [65] to convert all kinds of non-textual queries to visual prompts that lie in the same visual embedding space, as opposed to the conventional position encoding in SAM [28]. Our algorithm is illustrated in Fig. 2 and

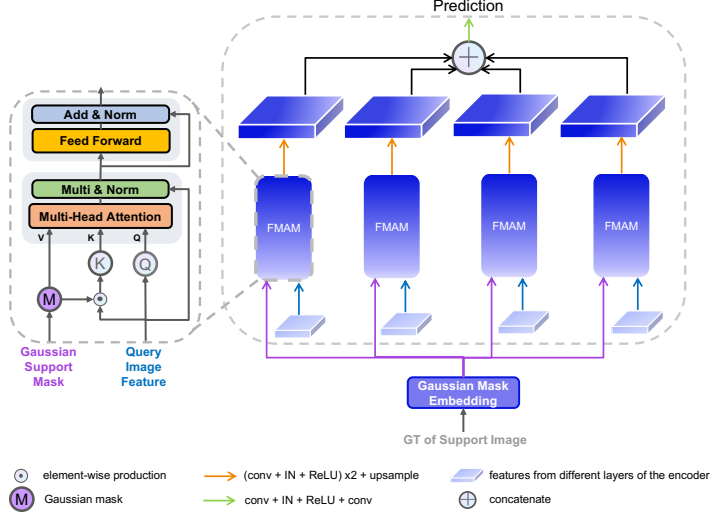


Figure 3. Process of Flexible Mask Attention Module (FMAM) and Flexible Mask Attention Decoder(FMAD).

can be summarized as follows:

$$P = \text{concatenate}([P_{positive}, P_{negative}]) \quad (5)$$

$$P_v = \text{Self-Sampler}(P, Z_h) \quad (6)$$

where $P_{positive}$, $P_{negative}$ stand for positive points (region we aim to segment) and negative points (background region) sampled from the image, respectively. Z_h is the feature maps extracted from the image and P denotes the point prompts.

Afterward, we concatenated these positive points and negative points to form the point prompts. This approach ensures that the model does not only focus on the positive parts, which could potentially lead to a higher false positive rate, but also adequately attends to the negative parts. With this method, the model can better discerns the organ boundaries, enhancing segmentation performance when we only give one point to prompt the model during inference.

On receiving the point prompts, we perform direct interpolation from the feature map to get the corresponding embeddings as visual prompts. This ensures that the prompts and the image embeddings share identical semantic features. We initialize a few tokens as learnable global queries and concatenate them with these visual prompts. After that, we apply self-attention [51] between a set of learnable queries and the visual prompts. Then, cross-attention is applied only between the image embeddings and these updated queries.

3.6. Flexible Mask Attention Decoder

Since we divided each 3D volume data into 12 chunks as in section 4.1, there is a certain pattern to how each organ changes across different slices. This observation inspired us to treat the segmentation task like a tracking task.

Consequently, we introduced the Flexible Mask Attention Decoder (FMAD). The core of this decoder is the Flexible Mask Attention Module (FMAM), as illustrated in Fig 3. By leveraging cross-attention, we can expand or contract the blurred support mask to predict the mask for the query image, mitigating the challenge of direct mask prediction. The Gaussian blurred support masks provide the model with ample room to explore potential mask positions while efficiently suppressing distractions from other parts of the image.

In decoding, the attention block bridges the Gaussian blurred support mask and queries images to update the mask. We first apply a Gaussian filter on original support mask $\mathbf{M} \in \mathbb{R}^{H \times W \times 1}$, repeat it for C times and reshape to $N \times C$ to get $\mathbf{M}' \in \mathbb{R}^{N \times C}$. Then we can easily cast it on query feature $\mathbf{F} \in \mathbb{R}^{N \times C}$ extracted from image encoder to get $\mathbf{K} \in \mathbb{R}^{N \times C}$ and $\mathbf{Q} \in \mathbb{R}^{N \times C}$ as follows:

$$\mathbf{K} = (\mathbf{M}' \odot \mathbf{F})\mathbf{W}^K \quad (7)$$

where \odot means element-wise production, $\mathbf{W}^K \in \mathbb{R}^{C \times C}$ means key projection matrix. Following [51], we also adopt the dot-product to compute the similarity matrix $\mathbf{A}_{K \rightarrow Q} \in \mathbb{R}^{N \times C}$ between the query and key as follows:

$$\mathbf{A}_{K \rightarrow Q} = \text{Atten}(\mathbf{Q}, \mathbf{K}) = \text{Softmax}_{\text{col}}(\bar{\mathbf{Q}}\bar{\mathbf{K}}^T / \tau) \quad (8)$$

where $\bar{\mathbf{Q}}$ and $\bar{\mathbf{K}}$ are l_2 -normalized features of \mathbf{Q} and \mathbf{K} across the channel dimension, and τ is a temperature parameter controlling the Softmax distribution, which is the same as in [55]. Then we convert \mathbf{M}' to $\mathbf{V} \in \mathbb{R}^{N \times C}$ though:

$$\mathbf{V} = (\mathbf{M}')\mathbf{W}^V \quad (9)$$

where $\mathbf{W}^V \in \mathbb{R}^{C \times C}$ denotes value projection matrix. With the attention matrix $\mathbf{A}_{K \rightarrow Q}$ from key to query, we can transform the value via $\mathbf{A}_{K \rightarrow Q} \mathbf{V} \in \mathbb{R}^{N \times C}$.

Unlike typical attention layers, at the first residual connection, we use *multiply & Norm* instead of *Add & Norm* to propagate query feature \mathbf{F} as follows:

$$\mathbf{F}' = \text{InsNorm}(\mathbf{A}_{K \rightarrow Q} \mathbf{V} \odot \mathbf{F}) \quad (10)$$

By using the updated Gaussian mask from the attention layer, we can further refine the predicted mask. After processing each of the four image embeddings, we upsample, concatenate, and put them through a convolutional layer. This approach facilitates improved information integration across layers to predict the final mask.

4. Experiments

4.1. Setup

Datasets To verify the generality and robustness of our approach, we perform experiments on two datasets, ABD-30 and ABD-MRI. We evaluate the complete few-shot segmentation framework (SSM-SAM) on both datasets, while the backbone (SS-SAM) is only assessed on ABD-30 to validate its efficiency.

- ABD-30 (from MICCAI15 Multi-Atlas Abdomen Labeling challenge [4]) contains 30 cases with 3779 axial abdominal clinical CT images.

- ABD-MRI (from Combined Healthy Abdominal Organ Segmentation (CHAOS) challenge [27] held in IEEE International Symposium on Biomedical Imaging (ISBI) 2019) contains 20 3D abdominal MRI scans with total four different labels.

Liver, spleen and left and right kidney are used as semantic classes following previous settings [14, 42, 50]. Within each experiment, one organ is considered as unseen semantic class for testing while the rest are used for training.

Evaluation Metrics We use the Dice Similarity Coefficient (DSC) to evaluate the prediction mask \mathbf{m} against the ground truth mask \mathbf{g} :

$$\text{DSC}(\mathbf{m}, \mathbf{g}) = \frac{2|m \cap g|}{|m| + |g|} \quad (11)$$

Implementation Details All the images extracted from 3D volume data are reshaped to 1024×1024 to fit into the SAM model. We follow the same protocol used in [42, 45, 50] to do 1-way 1-shot learning by dividing the 3D CT scans into 12 chunks and segmenting all the query slices in one chunk by using the center slice in the chunk as the support image, see Fig.4 for visual representation. We use ViT-B [16] version of SAM and supervise our network using balanced cross entropy loss and IoU loss as Eq. (2) between the predicted mask and ground truth mask. AdamW




Figure 4. Representative chunk

optimizer [37] is used for all the experiments with an initial learning rate of $2e-4$. Cosine decay [1] is applied to the learning rate. Few-shot segmentation task is trained for 50 epochs on a single NVIDIA A40 GPU and fully supervised segmentation task is trained for 50 epochs using PyTorch on a single NVIDIA RTX 3090 GPU.

4.2. Main Results

Results on Few-shot Medical Image Segmentation

Table 1 shows the few-shot segmentation performance of SSM-SAM with previous work on ABD-30 and ABD-MRI respectively. SE-Net [45] represents the first architecture designed explicitly for few-shot medical image segmentation. PANet [54] is an extended version of the widely-used prototypical network [49], tailored for natural image segmentation. SSL-ALPNet [42] integrates self-supervised learning with prototypical networks. Affine denotes the accuracy result after aligning the support and query images globally using an affine transformation. RP-Net [50] stands as the state-of-the-art framework for few-shot medical image segmentation, leveraging both a context relation encoder and a recurrent module. [50] reported performance for all the methods above, so these numbers are directly quoted.

First, compared to state-of-the-art method RP-Net, SSM-SAM outperforms RP-Net by 10.21% and 1.80% on ABD-30 and ABD-MRI respectively. Second, the use of FMAD results in improvements of 2.17% \sim 3.41% and 1.84% \sim 2.62% respectively. Additionally, employing the MAML++ based optimizer leads to enhancements of 5.98% \sim 7.22% and 4.47% \sim 5.25% respectively. This indicates that both our meta-optimizer and FMAD are highly effective.

Results on Fully Supervised Medical Image Segmentation

To evaluate the performance of our backbone (without MAML++), we compare SS-SAM with recent SOTA methods on ABD-30 dataset without Flexible Mask Attention Module, including U-Net [44], Att-UNet [41], TransUnet [9], Swin-Unet [8], MissFormer [24], TransDeepLab [3], HiFormer [21], DAE-Former [2] and SAMed [62] following [62]. For a fair comparison, we integrate the self-sampling module into some of the state-of-the-art models because the original SAM and our method utilize visual prompts, which incorporate ground truth information into the model while previous state-of-the-art models did




Figure 5. Testing process of MAML++ based model (in deep blue) and Fine-tuning of a model pretrained on the same distribution of tasks without MAML++ (in light blue)

Dataset	Method	Spleen \uparrow	Kidney L \uparrow	Kidney R \uparrow	Liver \uparrow	mean \uparrow
ABD-30	SE-Net [45]	0.23	32.83	14.34	0.27	11.91
	PANet [54]	25.59	32.34	17.37	38.42	29.42
	SSL-ALPNet [42]	60.25	63.34	54.82	73.65	63.02
	Affine	48.99	43.44	45.67	68.93	51.75
	RP-Net [50]	72.19	75.03	70.89	80.53	74.66
	SS-SAM (w/o MAML++, w/o FMAD)	77.24	72.89	79.78	72.00	75.48
	SS-SAM (w/o MAML++, w/ FMAD)	78.99	73.33	83.49	74.81	77.65
	SSM-SAM (w/ MAML++, w/o FMAD)	80.55	77.83	82.79	82.70	81.46
	SSM-SAM (w/ MAML++, w/ FMAD)	86.95	80.96	84.47	87.12	84.87
ABD-MRI	SE-Net [45]	51.80	62.11	61.32	27.43	50.66
	PANet [54]	50.90	53.45	38.64	42.26	46.33
	SSL-ALPNet [42]	67.02	73.63	78.39	73.05	73.02
	Affine	62.87	64.70	69.10	65.00	65.41
	RP-Net [50]	75.69	79.30	84.66	71.51	77.79
	SS-SAM (w/o MAML++, w/o FMAD)	71.99	76.52	72.13	69.36	72.50
	SS-SAM (w/o MAML++, w/ FMAD)	73.64	78.04	76.69	72.13	75.12
	SSM-SAM (w/ MAML++, w/o FMAD)	76.77	80.47	77.44	76.32	77.75
	SSM-SAM (w/ MAML++, w/ FMAD)	78.81	81.70	80.38	77.50	79.59

Table 1. DSC comparison with other methods on ABD-30 and ABD-MRI for few-shot learning (unit:%).

Method	DSC \uparrow	Aorta	Gallbladder	Kidney(L)	Kidney(R)	Liver	Pancreas	Spleen	Stomach
TransUNet [9]	77.48	87.23	63.13	81.87	77.02	94.08	55.86	85.08	75.62
SwinUnet [8]	79.13	85.47	66.53	83.28	79.61	94.29	56.58	90.66	76.60
MissFormer [24]	81.96	86.99	68.65	85.21	82.00	94.41	65.67	91.92	80.81
HiFormer [21]	80.39	86.21	65.69	85.23	79.77	94.61	59.52	90.99	81.08
DAE-Former [2]	82.43	88.96	72.30	86.08	80.88	94.98	65.12	91.94	79.19
PaNN [64]	90.8	92.5	72.9	95.3	92.0	95.3	-	96.8	-
<i>MissFormer</i> [24]	<i>78.84</i>	<i>85.89</i>	<i>64.75</i>	<i>84.19</i>	<i>76.98</i>	<i>94.18</i>	<i>55.19</i>	<i>89.07</i>	<i>80.47</i>
<i>HiFormer</i> [21]	<i>78.24</i>	<i>85.94</i>	<i>60.88</i>	<i>84.16</i>	<i>76.57</i>	<i>94.54</i>	<i>54.00</i>	<i>89.81</i>	<i>80.02</i>
<i>DAE-Former</i> [2]	<i>80.46</i>	<i>87.26</i>	<i>67.46</i>	<i>82.15</i>	<i>77.53</i>	<i>94.36</i>	<i>63.32</i>	<i>90.85</i>	<i>80.77</i>
SAMed [62]	81.88	87.77	69.11	80.45	79.95	94.80	72.17	88.72	82.06
SAM Adapter [10]	92.01	95.35	87.88	96.16	96.73	97.23	79.86	90.60	92.26
SS-SAM	93.09	95.96	89.77	96.09	96.79	97.61	79.15	95.15	94.23

Table 2. Comparison to state-of-the-art models using a fully supervised method on the ABD-30 dataset. Best results are highlighted in bold (unit:%). For a fair comparison, the results in italics are from the corresponding model with the same self-sampling module. Notably, our modified SAM (SS-SAM) outperforms other leading models.

not leverage such additional information. The results are highlighted in blue italics in Tab. 2. Surprisingly, since

these models are unaware of how to process this information, their performance slightly declined. We also observe

Prompt	Spleen	Kidney L	Kidney R	Liver	Mean
no prompt	91.83	93.65	92.55	96.21	93.56
w/ position encoding	93.05	95.75	96.87	97.27	95.74
w/ self-sampling	95.15	96.09	96.79	97.61	96.41

Table 3. Ablation Study on no prompt, position encoding (pe) of points and self-sampling

that SS-SAM achieves state-of-the-art performance. These experiments demonstrate that our models are capable of achieving high performance and also refute the potential claim that introducing prompts is a form of "cheating".

4.3. Ablation Study

Effect of self-sampling. As previously explained, the self-sampling module derives visual prompts by interpolating between the provided point coordinates and image embeddings. Subsequently, it conducts cross-attention with the image embeddings. On the other hand, SAM's point prompt employs positional encoding on the points before performing cross-attention with the image embeddings. These two approaches differ in their treatment of the sampled points. To evaluate the effectiveness of the proposed methods, we performed ablation studies on four distinct organs: the spleen, left kidney, right kidney, and liver, using a fully supervised approach. For each organ, we kept the image encoder and decoder the same, with three different prompt methods: a) no prompt, b) point prompt with position encoding, and c) point prompt with self-sampling. As presented in Tab. 3, our self-sampling method achieved the highest mean




Figure 6. Examples of predictions of SAM, SSM-SAM without FMAM and SSM-SAM with FMAM

Effects of Meta-learner. To highlight the few-shot learning advantages of our model achieved by integrating a meta-learner, apart from the experiment setting in Tab. 1, we set up a more rigorous experimental setting to compare performances with and without the meta-learner.

Following the configurations specified in [17, 30], we selected only 125 images for the four organs. In each experi-

ment, one organ is chosen as the test task (e.g., liver), while the remaining three organs (spleen, left kidney, right kidney) are used as training tasks. Only five images of each training task (spleen, left kidney, right kidney) are allocated during training. During testing, we adopt a 1-way, 5-shot approach, using five images of the test task (liver) as the support set for training. The remaining 120 images serve as the query set to evaluate the model's performance.

Our findings were illuminating. A visual representation of the online optimization process of each of the four organs is provided in Fig.5. With just a limited number of training images, models equipped with the meta-learner outperformed those without by a significant margin of 12.07% on average. Furthermore, models augmented by the meta-learner not only demonstrated swifter convergence in the initial epochs but also exhibited enhanced stability and performance in the concluding epochs.

4.4. Qualitative Results

In Fig.6, we present masks produced by SAM and variants of our algorithm. Notably, SAM struggles to segment smaller organs or those with indistinct boundaries. Conversely, SSM-SAM with FMAM outperforms its counterpart without FMAM, effectively minimizing the impact of similar distracting regions in the CT scans.

5. Conclusions

In this paper, we introduce a universal approach, SSM-SAM, designed to optimize and adapt a foundational model such as the Segment Anything Model (SAM), for few-shot medical image segmentation. Our method incorporates a positive-negative Self-Sampling prompt encoder and a Flexible Mask Attention Decoder to enhance the contextual relationship and tiny boundary information essential for mask generation. Moreover, our fast online meta-learning based optimizer facilitates high performance even without extensive training data and can be plugged into other frameworks effortlessly. Experiments demonstrate that SSM-SAM outperforms the previous state-of-the-art approach by as much as 10% in terms of DSC. Furthermore, the proposed SSM-SAM can produce segmentations in just 50 seconds per organ, indicating its potential for real-time applications in medical settings.

References

- [1] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your maml. *arXiv preprint arXiv:1810.09502*, 2018.
- [2] Reza Azad, René Arimond, Ehsan Khodapanah Aghdam, Amirhosein Kazerouni, and Dorit Merhof. Dae-former: Dual attention-guided efficient transformer for medical image segmentation. *arXiv preprint arXiv:2212.13504*, 2022.
- [3] Reza Azad, Moein Heidari, Moein Shariatnia, Ehsan Khodapanah Aghdam, Sanaz Karimijafarbigloo, Ehsan Adeli, and Dorit Merhof. Transdeeplab: Convolution-free transformer-based deeplab v3+ for medical image segmentation. In *International Workshop on PRedictive Intelligence In MEDicine*, pages 91–102. Springer, 2022.
- [4] J Igelsias M Styner T Langerak B Landman, Z Xu and A Klein. Miccai multi-atlas labeling beyond the cranial vault—workshop and challenge. *Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge*, 2015.
- [5] Junjie Bai and Xiaodong Wu. Error-tolerant scribbles based interactive image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 392–399, 2014.
- [6] Dhruv Batra, Adarsh Kowdle, Devi Parikh, Jiebo Luo, and Tsuhan Chen. icoseg: Interactive co-segmentation with intelligent scribble guidance. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3169–3176. IEEE, 2010.
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [8] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation. In *European conference on computer vision*, pages 205–218. Springer, 2022.
- [9] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021.
- [10] Tianrun Chen, Lanyun Zhu, Chaotao Ding, Runlong Cao, Shangzhan Zhang, Yan Wang, Zejian Li, Lingyun Sun, Papa Mao, and Ying Zang. Sam fails to segment anything?—sam-adapter: Adapting sam in underperformed scenes: Camouflage, shadow, and more. *arXiv preprint arXiv:2304.09148*, 2023.
- [11] Xi Chen, Zhiyan Zhao, Feiwu Yu, Yilei Zhang, and Manni Duan. Conditional diffusion for interactive segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7345–7354, 2021.
- [12] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*, 2022.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [14] Hao Ding, Changchang Sun, Hao Tang, Dawen Cai, and Yan Yan. Few-shot medical image segmentation with cycle-resemblance attention. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2488–2497, 2023.
- [15] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *BMVC*, volume 3, 2018.
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [17] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- [18] Shizhan Gong, Yuan Zhong, Wenao Ma, Jinpeng Li, Zhao Wang, Jingyang Zhang, Pheng-Ann Heng, and Qi Dou. 3dsam-adapter: Holistic adaptation of sam from 2d to 3d for promptable medical image segmentation, 2023.
- [19] 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.
- [20] Sheng He, Rina Bao, Jinpeng Li, P Ellen Grant, and Yangming Ou. Accuracy of segment-anything model (sam) in medical image segmentation tasks. *arXiv preprint arXiv:2304.09324*, 2023.
- [21] Moein Heidari, Amirhossein Kazerouni, Milad Soltany, Reza Azad, Ehsan Khodapanah Aghdam, Julien Cohen-Adad, and Dorit Merhof. Hiformer: Hierarchical multi-scale representations using transformers for medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6202–6212, 2023.
- [22] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019.
- [23] Mingzhe Hu, Yuheng Li, and Xiaofeng Yang. Skinsam: Empowering skin cancer segmentation with segment anything model. *arXiv preprint arXiv:2304.13973*, 2023.
- [24] Xiaohong Huang, Zhifang Deng, Dandan Li, Xueguang Yuan, and Ying Fu. Missformer: An effective transformer for 2d medical image segmentation. *IEEE Transactions on Medical Imaging*, 2022.
- [25] Won-Dong Jang and Chang-Su Kim. Interactive image segmentation via backpropagating refinement scheme. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5297–5306, 2019.
- [26] Wei Ji, Jingjing Li, Qi Bi, Tingwei Liu, Wenbo Li, and Li Cheng. Segment anything is not always perfect: An investigation of sam on different real-world applications, 2023.

- [27] A Emre Kavur, N Sinem Gezer, Mustafa Barış, Sinem Aslan, Pierre-Henri Conze, Vladimir Groza, Duc Duy Pham, Soumick Chatterjee, Philipp Ernst, Savaş Özkan, et al. Chaos challenge-combined (ct-mr) healthy abdominal organ segmentation. *Medical Image Analysis*, 69:101950, 2021.
- [28] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023.
- [29] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8334–8343, 2021.
- [30] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2869–2878, 2020.
- [31] Yuheng Li, Mingzhe Hu, and Xiaofeng Yang. Polypsam: Transfer sam for polyp segmentation. *arXiv preprint arXiv:2305.00293*, 2023.
- [32] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *European Conference on Computer Vision*, pages 280–296. Springer, 2022.
- [33] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3159–3167, 2016.
- [34] Zheng Lin, Zhao Zhang, Lin-Zhuo Chen, Ming-Ming Cheng, and Shao-Ping Lu. Interactive image segmentation with first click attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13339–13348, 2020.
- [35] Qin Liu, Zhenlin Xu, Gedas Bertasius, and Marc Niethammer. Simpleclick: Interactive image segmentation with simple vision transformers. *arXiv preprint arXiv:2210.11006*, 2022.
- [36] Weihuang Liu, Xi Shen, Chi-Man Pun, and Xiaodong Cun. Explicit visual prompting for low-level structure segmentations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19434–19445, 2023.
- [37] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [38] Jun Ma and Bo Wang. Segment anything in medical images. *arXiv preprint arXiv:2304.12306*, 2023.
- [39] Anastasia Makarevich, Azade Farshad, Vasileios Belagianis, and Nassir Navab. Metamedseg: Volumetric meta-learning for few-shot organ segmentation, 2021.
- [40] Maciej A Mazurowski, Haoyu Dong, Hanxue Gu, Jichen Yang, Nicholas Konz, and Yixin Zhang. Segment anything model for medical image analysis: an experimental study. *Medical Image Analysis*, page 102918, 2023.
- [41] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.
- [42] Cheng Ouyang, Carlo Biffi, Chen Chen, Turky Kart, Huaqi Qiu, and Daniel Rueckert. Self-supervision with superpixels: Training few-shot medical image segmentation without annotation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 762–780. Springer, 2020.
- [43] Martin Rajchl, Matthew CH Lee, Ozan Oktay, Konstantinos Kamnitsas, Jonathan Passerat-Palmbach, Wenjia Bai, Mellisa Damodaram, Mary A Rutherford, Joseph V Hajnal, Bernhard Kainz, et al. Deepcut: Object segmentation from bounding box annotations using convolutional neural networks. *IEEE transactions on medical imaging*, 36(2):674–683, 2016.
- [44] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [45] Abhijit Guha Roy, Shayan Siddiqui, Sebastian Pölsterl, Nassir Navab, and Christian Wachinger. ‘squeeze & excite’ guided few-shot segmentation of volumetric images. *Medical image analysis*, 59:101587, 2020.
- [46] Saikat Roy, Tassilo Wald, Gregor Koehler, Maximilian R Rokuss, Nico Disch, Julius Holzschuh, David Zimmerer, and Klaus H Maier-Hein. Sam. md: Zero-shot medical image segmentation capabilities of the segment anything model. *arXiv preprint arXiv:2304.05396*, 2023.
- [47] Peilun Shi, Jianing Qiu, Sai Mu Dalike Abaxi, Hao Wei, Frank P-W Lo, and Wu Yuan. Generalist vision foundation models for medical imaging: A case study of segment anything model on zero-shot medical segmentation. *Diagnostics*, 13(11):1947, 2023.
- [48] Rishav Singh, Vandana Bharti, Vishal Purohit, Abhinav Kumar, Amit Kumar Singh, and Sanjay Kumar Singh. Metamed: Few-shot medical image classification using gradient-based meta-learning. *Pattern Recognition*, 120:108111, 2021.
- [49] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.
- [50] Hao Tang, Xingwei Liu, Shanlin Sun, Xiangyi Yan, and Xiaohui Xie. Recurrent mask refinement for few-shot medical image segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3918–3928, 2021.
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [52] Guotai Wang, Wenqi Li, Maria A Zuluaga, Rosalind Pratt, Premal A Patel, Michael Aertsen, Tom Doel, Anna L David, Jan Deprest, Sébastien Ourselin, et al. Interactive medical image segmentation using deep learning with image-

- specific fine tuning. *IEEE transactions on medical imaging*, 37(7):1562–1573, 2018.
- [53] Jiaqi Wang, Zhengliang Liu, Lin Zhao, Zihao Wu, Chong Ma, Sigang Yu, Haixing Dai, Qiushi Yang, Yiheng Liu, Songyao Zhang, Enze Shi, Yi Pan, Tuo Zhang, Dajiang Zhu, Xiang Li, Xi Jiang, Bao Ge, Yixuan Yuan, Dinggang Shen, Tianming Liu, and Shu Zhang. Review of large vision models and visual prompt engineering, 2023.
- [54] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *proceedings of the IEEE/CVF international conference on computer vision*, pages 9197–9206, 2019.
- [55] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1571–1580, 2021.
- [56] Junde Wu, Rao Fu, Huihui Fang, Yuanpei Liu, Zhaowei Wang, Yanwu Xu, Yueming Jin, and Tal Arbel. Medical sam adapter: Adapting segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.12620*, 2023.
- [57] Jiajun Wu, Yibiao Zhao, Jun-Yan Zhu, Siwei Luo, and Zhuowen Tu. Milcut: A sweeping line multiple instance learning paradigm for interactive image segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 256–263, 2014.
- [58] Ning Xu, Brian Price, Scott Cohen, Jimei Yang, and Thomas S Huang. Deep interactive object selection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 373–381, 2016.
- [59] Hanchao Yu, Shanhui Sun, Haichao Yu, Xiao Chen, Honghui Shi, Thomas S Huang, and Terrence Chen. Foal: Fast online adaptive learning for cardiac motion estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4313–4323, 2020.
- [60] Bingfeng Zhang, Jimin Xiao, and Terry Qin. Self-guided and cross-guided learning for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8312–8321, 2021.
- [61] Chaoning Zhang, Sheng Zheng, Chenghao Li, Yu Qiao, Taegoo Kang, Xinru Shan, Chenshuang Zhang, Caiyan Qin, Francois Rameau, Sung-Ho Bae, et al. A survey on segment anything model (sam): Vision foundation model meets prompt engineering. *arXiv preprint arXiv:2306.06211*, 2023.
- [62] Kaidong Zhang and Dong Liu. Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785*, 2023.
- [63] Tao Zhou, Yizhe Zhang, Yi Zhou, Ye Wu, and Chen Gong. Can sam segment polyps?, 2023.
- [64] Yuyin Zhou, Zhe Li, Song Bai, Chong Wang, Xinlei Chen, Mei Han, Elliot Fishman, and Alan L Yuille. Prior-aware neural network for partially-supervised multi-organ segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10672–10681, 2019.
- [65] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *arXiv preprint arXiv:2304.06718*, 2023.