

# Unleashing the Potential of SAM for Medical Adaptation via Hierarchical Decoding

Zhiheng Cheng<sup>1</sup> Qingyue Wei<sup>2</sup> Hongru Zhu<sup>3</sup> Yan Wang<sup>1</sup> Liangqiong Qu<sup>4</sup>  
Wei Shao<sup>5</sup> Yuyin Zhou<sup>6</sup>

<sup>1</sup>Shanghai Key Laboratory of Multidimensional Information Processing, East China Normal University

<sup>2</sup>Stanford University <sup>3</sup>Johns Hopkins University <sup>4</sup>The University of Hong Kong

<sup>5</sup>University of Florida <sup>6</sup>UC Santa Cruz

## Abstract

The Segment Anything Model (SAM) has garnered significant attention for its versatile segmentation abilities and intuitive prompt-based interface. However, its application in medical imaging presents challenges, requiring either substantial training costs and extensive medical datasets for full model fine-tuning or high-quality prompts for optimal performance. This paper introduces H-SAM: a prompt-free adaptation of SAM tailored for efficient fine-tuning of medical images via a two-stage hierarchical decoding procedure. In the initial stage, H-SAM employs SAM's original decoder to generate a prior probabilistic mask, guiding a more intricate decoding process in the second stage. Specifically, we propose two key designs: 1) A class-balanced, mask-guided self-attention mechanism addressing the unbalanced label distribution, enhancing image embedding; 2) A learnable mask cross-attention mechanism spatially modulating the interplay among different image regions based on the prior mask. Moreover, the inclusion of a hierarchical pixel decoder in H-SAM enhances its proficiency in capturing fine-grained and localized details. This approach enables SAM to effectively integrate learned medical priors, facilitating enhanced adaptation for medical image segmentation with limited samples. Our H-SAM demonstrates a 4.78% improvement in average Dice compared to existing prompt-free SAM variants for multi-organ segmentation using only 10% of 2D slices. Notably, without using any unlabeled data, H-SAM even outperforms state-of-the-art semi-supervised models relying on extensive unlabeled training data across various medical datasets. Our code is available at <https://github.com/Cccccczh404/H-SAM>.

## 1. Introduction

Accurate delineation of tissues, organs, and regions of interest through medical image segmentation is pivotal in aid-

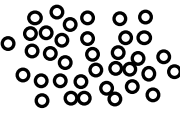



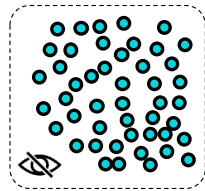


Unlabeled	Labeled	Mean Dice (%)	Methods
		78.27	SOTA Semi-Supervised Methods
		86.83 	H-SAM
		75.57	SOTA SAM Variants
		80.35 	H-SAM

Figure 1. H-SAM is advantageous in few-shot medical image segmentation. It achieves over 80% in average Dice using only 10% slices for multi-organ segmentation, outperforming existing prompt-free SAM adaptation methods. Without using any unlabeled data at all, it even outperforms state-of-the-art semi-supervised models that use extensive unlabeled training data for prostate segmentation.

ing medical professionals' diagnostic precision and treatment planning processes [13, 15]. Furthermore, it plays a fundamental role in propelling disease research and discovery [53]. Nonetheless, a significant challenge in this field lies in the demand for deep learning models to undergo extensive training on large annotated datasets, a resource often challenging to procure within the medical domain.

Recently, Segment Anything Model (SAM) [37], trained with over a billion masks from diverse natural images, demonstrates remarkable zero-shot learning capabilities. This breakthrough presents an avenue for significant advancements in medical image segmentation, especially considering the limited availability of extensive datasets in the medical realm. However, SAM's performance on medical

images diminishes notably in zero-shot settings, exhibiting reduced accuracy and robustness [25, 29, 35, 43, 50]. This decline can be attributed to SAM’s lack of exposure to medical images during training, as its extensive training revolves around natural images [32, 71].

While training SAM exclusively on medical datasets is a potential solution, it incurs substantial training costs and risks of over-fitting to single datasets [49]. Efforts to bridge the gap between medical and natural image domains involve adapting SAM to specific medical datasets [14, 23, 69]. Previous works primarily focus on inserting adapter layers into the image encoder with minimal decoder changes [23, 70]. Most of these efforts employ prompted SAM adaptation, generating prompts using point or bounding boxes from ground truth during testing [21, 64, 72]. However, creating accurate prompts demands domain knowledge from medical experts, which is often limited, time-consuming, and prone to noise, compromising segmentation accuracy. In response, prompt-free SAM adaptation methods [9, 56, 70] have emerged, yet they typically yield inferior results compared to prompted methods due to the lack of medical knowledge that prompts provide.

We present H-SAM, a prompt-free variant of the Segment Anything Model (SAM), aimed at integrating medical knowledge via a streamlined two-stage hierarchical mask decoder while maintaining the image encoder frozen. Initially, input images are processed by a LoRA-adapted image encoder. H-SAM employs SAM’s original lightweight mask decoder in the first stage to generate a prior probabilistic mask, guiding a more intricate second decoding stage. Two key designs underpin this process: 1) A **class-balanced, mask-guided self-attention** mechanism recalibrates the image embedding using self-attention from the prior mask, ensuring balanced representation across categories with noise augmentation. 2) A **learnable mask cross-attention** mechanism employs the prior mask to modulate cross-attention spatially within the subsequent Transformer decoder, attenuating less relevant background noise. Moreover, a hierarchical pixel decoder complements the hierarchical Transformer decoder, enhancing the model’s precision and ability to capture localized details. Figure 2 illustrates the overall pipeline.

In both fully-supervised and few-shot multi-organ segmentation tasks, our H-SAM surpasses existing prompt-free SAM variants. Specifically, using 10% and 100% of 2D slices, H-SAM achieves an average Dice score improvement of 4.78% and 3.48% on the Synapse dataset, respectively. This superior performance is also evident in other few-shot medical image segmentation tasks including the left atrial (LA) dataset and PROMISE2012 dataset. Notably, as illustrated in Figure 1, H-SAM excels without the need for any unlabeled data, surpassing state-of-the-art semi-supervised models that rely on extensive unlabeled

training datasets. The promising performance of 87.27% and 89.22% achieved on the prostate and left atrial segmentation using only 3 and 4 cases highlights H-SAM’s potential in medical imaging applications.

## 2. Related Work

**Medical Foundation Models** Foundation models are pre-trained, large-scale models that allow for rapid customization through fine-tuning or in-context learning, as exemplified in [18, 55]. Despite these notable signs of progress, challenges persist in complex tasks like image segmentation, primarily due to the difficulty of obtaining annotated masks. Segment Anything Model (SAM) [37], with extensive training on a dataset of over 1 billion natural images, showcases impressive performance in image segmentation. Particularly, in diverse real-world scenarios, SAM demonstrates powerful capabilities for zero-shot generalization, signifying its potential to address intricate computer vision challenges.

With the surge of Segment Anything Model [37], previous work seeks to apply SAM on medical images [5, 16, 73]. However, empirical studies of SAM’s zero-shot capability on medical images [12, 17, 27, 60] reveal a significant decline when dealing with unseen medical features [52, 61]. Consequently, recent work explores effective adaptations of prompted SAM on medical datasets. Due to the computational intensity associated with training all parameters of SAM, researchers primarily concentrate on updating a subset of SAM’s parameters. Many works require prompts to generalize SAM on medical datasets. For instance, MedSAM [49] curates a large medical image dataset to adapt bound box prompted SAM. Medical SAM Adapter [64] fine-tunes point-prompted SAM using Adaption modules. Several studies transfer SAM from 2D to 3D by adding layers to support volumetric inputs [24, 39, 40, 46]. Other prompted SAM adaptations employ exemplar learning [21] and new prompt mechanisms [72]. Apart from prompted SAM adaptations, which require prompts sampled from ground truth during testing, prompt-free SAM adaptation methods are proposed for medical segmentation where prompts are not necessarily available. AutoSAM [30] fine-tunes SAM without prompts by freezing the SAM encoder and adding prediction heads to generate segmentation masks. SAMed [70] achieves competitive prompt-free results by introducing LoRA [28] layers into the image encoder while using the original decoder.

This study also introduces a prompt-free version of SAM, designed to enhance efficient finetuning with limited medical data. It differentiates from existing research by introducing a novel hierarchical decoding process that incorporates medical prior knowledge.

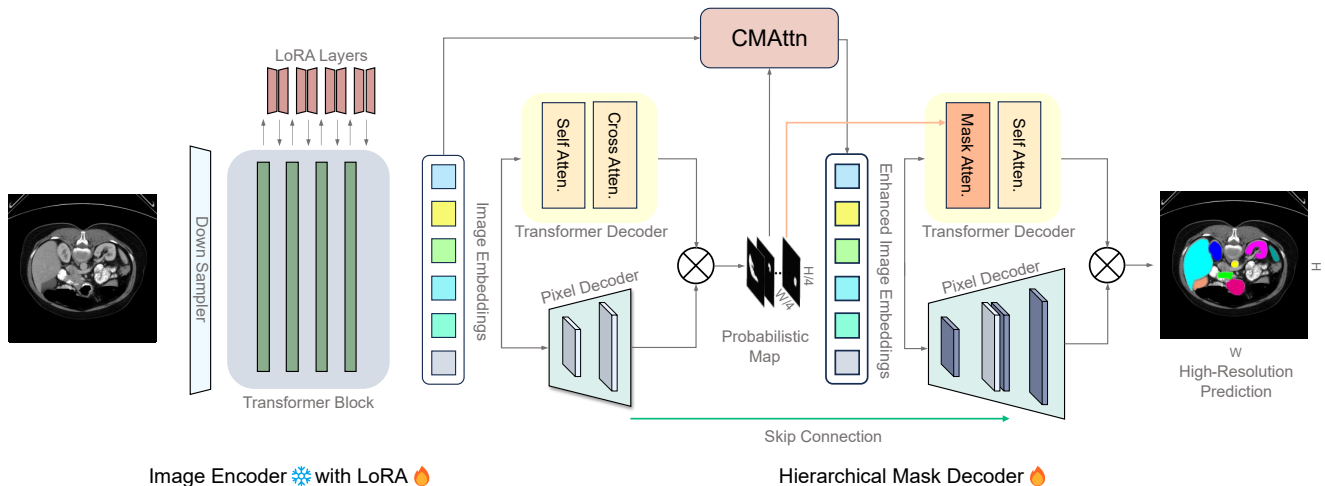


Figure 2. The H-SAM framework integrates a LoRA-adapted image encoder and a sophisticated 2-stage hierarchical decoder. We finetune the prompt encoder with default embeddings under a prompt-free setting. A key innovation lies in our hierarchical mask decoder, which strategically utilizes predictions from the stage-1 decoder as priors to achieve nuanced segmentation with 2 implementations: Class-Balanced Mask-Guided Self-Attention (CMAtn), and Learnable Mask Cross-Attention. And a hierarchical pixel decoder is employed to complement the enriched object queries derived from the transformer decoder.

**Model Fine-Tuning** Efficient fine-tuning of foundation models is crucial, with adapters becoming key for integrating new knowledge into pretrained models and providing a more efficient alternative to complete finetuning [10, 26]. Low Rank Adaptation (LoRA) [28] advocates for a gradual parameter update within transformer blocks, aiming for a low-rank approximation to refine large-scale models. In our proposed H-SAM, we incorporate LoRA adapters into SAM image encoder to avoid over-fitting in medical image dataset adaptation.

**Medical Image Segmentation** Medical image segmentation refers to partitioning or dividing a medical image into the dense prediction of pixels corresponding to lesions or organs in imaging modalities such as CT [22, 62, 74, 75] and MRI [36, 68]. With the rapid progress of deep learning, U-Net [58] with its elaborate network design ushers a new era for medical image segmentation. Building on U-Net’s foundation, a series of works emerge to enhance segmentation performance with U-shaped models [19, 33, 76]. Recent advancements of vision transformers in natural image analysis also prompt exploration in medical image segmentation [20]. A prevalent network design strategy involves integrating transformer blocks into the U-Net framework, resulting in novel architectures such as TransUnet [7] and Swin-Unet [4]. MISSFormer [31] is a u-shaped encoder-decoder network with enhanced transformer blocks. Recent works also explore medical image segmentation in a few-shot setting. RP-Net [59] proposed to iteratively capture context relationships using a U-shaped network. CAT-Net and 3D TransUNet [8, 44] design a cross masked attention

Transformer to focus only on foreground regions between the support image and query image.

Unlike existing studies, this study is focused on the fine-tuning of large foundational models, specifically SAM, to facilitate more efficient adaptation for medical image segmentation tasks. To achieve this, we introduce a hierarchical decoding strategy designed to optimize SAM’s capabilities, thereby unlocking its full potential for efficient and effective fine-tuning on medical tasks.

### 3. Methodology

#### 3.1. H-SAM Overview

Given an image  $I$  of size  $W \times H$ , our goal is to predict its corresponding segmentation map of  $W \times H$ . Each pixel in this map is assigned to a category from a predefined class list, aiming for maximal alignment with the ground truth  $gt$ . Our segmentation framework H-SAM is built up upon SAM, integrating a LoRA-adapted image encoder and a simple but effective 2-stage hierarchical decoder.

**LoRA-adapted Image Encoder** As illustrated in Figure 2, H-SAM utilizes the original image encoder of SAM and freezes all layers to preserve pre-learned knowledge. Then, we adopt the same LoRA implementation as SAMed [70] to add a smaller, trainable bypass composed of two low-rank matrices. In line with LoRA, these bypasses first compress the transformer features into a low-rank space. Subsequently, they reproject these condensed features to match the output feature channels of the frozen transformer blocks. Only these bypass matrices are updated

during training, allowing for minor yet effective model adjustments. For the prompt encoder, H-SAM does not need any prompt and simply updates a default embedding during training.

**Mask Decoder** The original SAM mask decoder consists of a Transformer decoder and a pixel decoder. The Transformer decoder processes image embeddings extracted from the image encoder, employing self-attention mechanisms to evaluate the significance of various image regions and cross-attention mechanisms to focus on relevant areas for segmentation. Subsequently, the pixel decoder refines this output, generating a detailed segmentation map, and assigning a class or category to each pixel.

**Hierarchical Decoding** Our H-SAM introduces a more intricate two-stage hierarchical decoding procedure. In the first stage, H-SAM employs SAM’s original decoder to create a prior (probabilistic) mask, which will be used to guide more intricate decoding in the second stage, as illustrated in Figure 2. This second stage mirrors the original one with both a Transformer decoder and a pixel decoder. To enhance image embedding input and optimize cross-attention in the second Transformer decoder, we introduce two novel modules. Firstly, a class-balanced, mask-guided self-attention mechanism is proposed to rectify the issue of unbalanced label distribution, thereby enhancing the image embedding for the second-stage Transformer decoder (Sec. 3.2). Secondly, we incorporate a learnable mask cross-attention mechanism within the second Transformer decoder. This mechanism adeptly modulates the spatial dynamics among various image regions, guided by the information from the prior mask, thereby enhancing the segmentation process (Sec. 3.3). Collectively, these decoders constitute a hierarchical Transformer decoder framework. Further, we propose a hierarchical pixel decoder, inspired by U-Net architecture, to supplement the hierarchical Transformer decoder and further refine the segmentation outcome. Concretely, the pixel decoder in the second stage integrates features from the first-stage pixel decoder through skip connections, enabling the generation of high-resolution predictions (Sec. 3.4).

### 3.2. Enhanced Image Embedding via Class-Balanced Mask-Guided Self-Attention

As shown in Figure 3, we incorporate a Class-Balanced Mask-Guided Self-Attention (CMAtn) block to enhance the image embeddings as input for the second-stage Transformer decoder. This is particularly helpful in the case where we have an imbalance between the abundant instances in the head categories and the scarcity of instances in the tail categories. We use a mask feature which is acquired by directly multiplying image embedding without

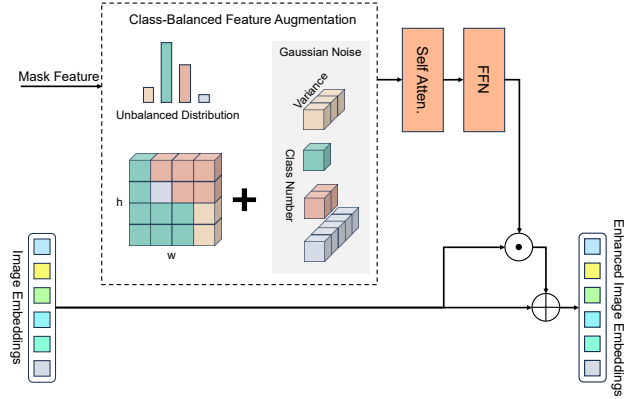


Figure 3. The illustration of Class-Balanced Mask-Guided Self-Attention (CMAtn) block.

upsampling from the first decoder as the input mask feature for CMAtn. Before the self-attention block, we adopt a class-balanced augmentation to introduce more variations to tail categories. Inspired by previous approaches utilizing logit adjustment in long-tail problems [41], we perturb the mask feature with Gaussian noise whose variance is inversely proportional to category sample frequencies:

$$P(gt = i) + \mathcal{N}(0, var(i)), \quad (1)$$

where  $P \in \mathbb{R}^{N \times C \times H \times W}$  is the normalized input mask feature.  $gt$  is the ground truth mask resized to the same size.  $\mathcal{N}$  is the added Gaussian noise. The variance list is calculated offline and stored as  $var$ .

After self-attention, we adopt a linear layer to compress the channel dimension, and incorporate the resulting mask feature to input image embeddings using Hadamard product  $\odot$ . A residual path is designed to retain information from the initial image embedding.

### 3.3. Learnable Mask Cross Attention

Mask-attention is a variant of cross-attention first proposed in Mask2Former [11]. Unlike cross-attention which attends to the global context, mask-attention operates only on the area within the predicted mask. Original mask-attention adds a transformed binary mask to the cross attention operation via

$$X = softmax(t(M) + KQ^T)V + X, \quad (2)$$

where  $X$  is the input query feature of the transformer block.  $K$ ,  $Q$ , and  $V$  are the key, query and value in cross attention.  $t(M)$  is a function mapping binarized input  $\{0, 1\}$  to  $\{-\infty, 0\}$ . This mask formulation has two limitations: (1) the gradient of mask  $M$  vanishes through  $t(M)$ ; (2) binarized mask  $M$  treats all foreground pixels without differentiation, limiting its ability to interpret further information from the mask prior.



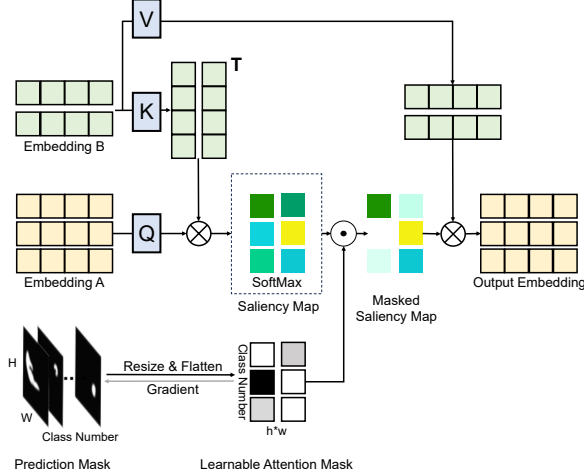


Figure 4. The illustration of learnable mask-attention.

To address these limitations, we propose to use learnable mask cross-attention in the second decoding stage as shown in Figure 4, which can be formulated as:

$$X = M \odot \text{softmax}(KQ^T)V + X, \quad (3)$$

which employs an untransformed probabilistic map  $M$  resized to the same spatial resolution as the saliency map in cross attention. With element-wise product between the mask and the saliency map, masked regions will be ignored by multiplying a near-zero probability. This new formulation mitigates the aforementioned limitations and facilitates rapid convergence and better performance. Our learnable mask cross-attention within the second-stage Transformer decoder leverages more information from the probabilistic map and can assign varying degrees of importance to different foreground regions.

### 3.4. Hierarchical pixel decoder

Complementing the Transformer decoder, SAM’s original pixel decoder directly upsamples the image embedding into a segmentation map of  $H/4 \times W/4$ . We argue that this resolution is not capable of capturing some intricate local details and small-scale medical objects in medical image segmentation necessitates the use of U-shaped networks with multiple skip connections. To enhance the details in the segmentation output, we adopt U-shaped architectures only in pixel decoders with skip connections for effectively handling multi-scale objects in medical images with acceptable computation cost. Within our H-SAM hierarchical decoding, the second-stage transformer decoder undergoes meticulously crafted mask-guided designs to propagate the prior mask from the first stage to the next stage (as shown in Sec. 3.2 and Sec. 3.3). Here we propose the hierarchical pixel decoder designed to complement the enriched object

queries derived from the transformer decoder. Similar to the hierarchical Transformer decoder, the hierarchical pixel decoder also consists of two successive pixel decoders, strategically incorporates skip connections to integrate features from the first pixel decoder to the second, and further up-sample the resolution from  $H/4 \times W/4$  to the full resolution  $H \times W$ . Benefiting from skip-connected localized features, the hierarchical pixel decoder will be supplemented with the Transformer decoder to output an enriched representation with enhanced resolution.

**Training Loss** The training loss combines pixel-wise classification loss and binary mask loss for each segmented prediction:

$$\mathcal{L} = \lambda_{ce}\mathcal{L}_{ce} + \lambda_{dice}\mathcal{L}_{dice}, \quad (4)$$

where the pixel-wise classification loss  $\mathcal{L}_{ce}$  and  $\mathcal{L}_{dice}$  denote binary cross-entropy loss and dice loss, respectively [51]. For our 2-stage hierarchical structure, there is also a  $\lambda_w$  for each loss in the 2 stage. The final loss  $\mathcal{L}_{total}$  a sum of  $\lambda_w\mathcal{L}_{stage1}$  and  $(1 - \lambda_w)\mathcal{L}_{stage2}$ . The parameter  $\lambda_w$  is set to gradually decrease from 0.8 in a way of exponential decay, at a decay coefficient of 0.005. The first decoder output is supervised by 1/4 resolution ground truth, and the second output by full resolution. The final output is ensembled by taking the average of probabilities from the two outputs.

**Deep Supervision** The training loss will be applied to every stage in our hierarchical decoding procedure. The supervisory signals for the stage-1 mask decoder stem from ground truth masks of  $H/4 \times W/4$ , while the stage-2 mask decoder is directly supervised using original high-resolution ground truth. This ensures thorough supervision and enhances the overall effectiveness of the model.

## 4. Experiments

### 4.1. Dataset and Evaluation

We conduct experiments of multi-organ semantic segmentation on three medical datasets, including Synapse multi-organ CT [38], the left atrial (LA) dataset [6], and PROMISE12 [45]. We utilize Dice coefficient and the average Hausdorff distance (HD) as evaluation metrics.

**Synapse Multi-Organ CT.** The Synapse dataset is from MICCAI 2015 Multi-Atlas Abdomen Labeling Challenge, which contains 3779 axial contrast enhanced abdominal CT images in total. The training set contains 2212 axial slices. We strictly follow TransUnet [7] and SAMed [70] for dataset split and data preprocessing. The dataset is split into 18 training cases and 12 test cases. The CT volumes

Training set	Method	Spleen	Right Kidney	Left Kidney	Gallbladder	Liver	Stomach	Aorta	Pancreas	Mean Dice $\uparrow$ (%)	HD $\downarrow$	
10%	AutoSAM [30]	68.80	77.44	76.53	24.87	88.06	52.70	75.19	34.58	55.69	31.67	
	SAM Adapter [9]	72.42	68.38	66.77	22.38	89.69	53.15	66.74	26.76	58.28	54.22	
	SAMed [70]	85.82	82.25	82.62	63.15	92.72	67.20	78.72	52.12	75.57	23.02	
	H-SAM (ours)	90.21	84.16	85.65	70.70	94.29	76.10	85.54	56.17	<b>80.35</b>	<b>15.54</b>	
Fully Supervised	TransUnet [7]	87.23	63.13	81.87	77.02	94.08	55.86	85.08	75.62	77.48	31.69	
	SwinUnet [4]	85.47	66.53	83.28	79.61	94.29	56.58	90.66	76.60	79.13	21.55	
	TransDeepLab [1]	86.04	69.16	84.08	79.88	93.53	61.19	89.00	78.40	80.16	21.25	
	DAE-Former [2]	88.96	72.30	86.08	80.88	94.98	65.12	91.94	79.19	82.43	17.46	
	MERIT [57]	92.01	84.85	87.79	74.40	95.26	85.38	87.71	71.81	84.90	13.22	
	AutoSAM [30]	80.54	80.02	79.60	41.37	89.24	61.14	82.56	44.22	62.08	27.56	
	SAM Adapter [9]	83.68	79.00	79.02	57.49	92.67	69.48	77.93	43.07	72.80	33.08	
	SAMed [70]	87.77	69.11	80.45	79.95	94.80	72.17	88.72	82.06	81.88	20.64	
	<b>H-SAM (ours)</b>	93.34(8.22)	89.93(9.59)	91.88(8.25)	73.49(23.01)	95.72(1.49)	87.10(7.14)	89.38(2.82)	71.11(10.27)		<b>86.49</b>	<b>8.18</b>

Table 1. Comparison to state-of-the-art models on Synapse multi-organ CT dataset with both few-shot and fully-supervised settings. Our model shows outstanding results in both of the training settings. The value in ( $\cdot$ ) is standard deviation.

for Synapse dataset of each volume contain 85 to 198 slices. The resolution of Synapse dataset is  $512 \times 512$  during few-shot training, and  $224 \times 224$  during fully-supervised training. We evaluate eight abdominal organs (aorta, gallbladder, spleen, left kidney, right kidney, liver, pancreas, stomach), following TransUnet [7].

**LA.** The left atrial (LA) dataset is from 2018 Atrial Segmentation Challenge [6]. We strictly follow UA-MT [67] and BCP [3] for data split and data preprocess. Specifically, LA dataset is split into 80 scans for training and 20 scans for evaluation. And we keep 4(5%) scans as labeled data while the rest scans in the training set are treated as unlabeled data for these two settings respectively. And we resize each 2D slice into  $512 \times 512$  during training. Note that we do not use the unlabeled data for training H-SAM. Instead, we only use these selected 4 labeled scans to train our model.

**PROMISE12.** PROMISE2012 dataset is from the Prostate MR Image Segmentation 2012 [45]. It contains 50 3D transversal T2-weighted MR images of the prostate with manual binary prostate gland segmentation and is obtained from multiple centers with different acquisition protocols. During the experiments, we strictly follow MLB-Seg [63] for data split and data preprocessing. Specifically, we split it into 40 / 10 cases for training / evaluation. 3 out of 40 are selected as data with labels while the rest 37 scans are used as unlabeled data. The resolution of PROMISE12 dataset is  $512 \times 512$ . Again, note that only the 3 labeled cases are used for training H-SAM.

## 4.2. Implementation details

All our implementation is in PyTorch and we train all our models on 4 NVIDIA RTX A5000 GPUs. During training, we adopt a data augmentation combination of elastic deformation, rotation, and scaling. The training loss is a combination of Cross-Entropy loss and Dice loss. We adopt the same LoRA settings as SAMed, in which the rank of LoRA

is set to 4. A ViT-B and a ViT-L backbone are adopted separately for few-shot and fully-supervised training. For a fair comparison, we utilize the same resolution of  $224 \times 224$  for fully supervised training on Synapse as other SAM variants and SOTA methods. The maximal training epoch is set to 300. The optimizer algorithm used for updating is based on the AdamW, and  $\beta_1$ ,  $\beta_2$ , and weight decay are set to 0.9, 0.999, and 0.1.

## 4.3. Results

**Synapse Multi-Organ CT** In Table 1, H-SAM shows outstanding few-shot transferability with limited seen medical images (10%). Compared to other SAM prompt-free variants: Auto SAM [30], SAM Adapter [9] and SAMed [70], our H-SAM reports results of 80.35% with limited training scans equal to only 1 volume, which outperforms other SAM adaptation variants by a large margin ( $\approx 5\%$ ). We also evaluate our H-SAM under a fully-supervised setting on Synapse multi-organ CT dataset. We evaluate our model in a fair comparison with the prompt-free setting with several state-of-the-art methods, including TransUnet [7], SwinUnet [4], TransDeepLab [1], DAE-Former [2] and MERIT[57], along with other SAM prompt-free variants: Auto SAM [30], SAM Adapter [9] and SAMed [70]. Our method achieves promising results in multi-organ segmentation with an 86.49% Mean Dice Coefficient, which is higher than newly released medical segmentation networks DAE-Former (82.43%) and MERIT (84.90%). H-SAM also easily outperforms other prompt-free SAM variants (86.24% vs. 81.88%).

**LA** As shown in Table 2, we conduct the few-shot semantic segmentation experiment on LA dataset. Here, we present results using 4 labeled scans (5% of the dataset) for training. We compare our method with two categories of approaches: 1) SAM efficient adaptation methods including AutoSAM [30], SAM Adapter, and SAMed [70] and 2) semi-supervised methods such as UA-MT [67], MC-Net [65], SS-Net [66], and BCP [3]. We also compare

Methods	Scans used		Mean Dice (%) $\uparrow$
	Labeled	Unlabeled	
UA-MT [67]			82.26
SASSNet [42]			81.60
DTC [47]			81.25
URPC [48]	4(5%)	76(95%)	82.48
MC-Net [65]			83.59
SS-Net [66]			86.33
BCP [3]			88.02
nnUnet [34]			64.02
AutoSAM [30]			74.73
SAM Adapter [70]	4(5%)	0(0%)	82.79
SAMed [70]			87.72
Ours			<b>89.22</b>

Table 2. Results of LA dataset under semi-supervision using 4 labeled scans.

our results with nnUnet [34], a well-known baseline for medical image segmentation. To ensure a fair comparison, we adhered to the data split protocols established in UA-MT [67] and BCP [3], employing identical sets of 4 labeled scans for our few-shot experiments. The semi-supervised methods also utilized the remaining 76 unlabeled scans during their training. And all the methods were evaluated on the same test dataset. Our proposed H-SAM method outperforms both the SAM efficient adaptation and semi-supervised methods under both split settings. This underscores the ability of H-SAM to achieve superior or comparable results with significantly less data compared to semi-supervised methods. Under the few-shot setting, H-SAM also demonstrates superior improvements in dice coefficients compared to SAMed [70] (89.22% vs. 87.72%), indicating the effectiveness of H-SAM in few-shot learning.

**PROMISE12** We also conducted the few-shot semantic segmentation experiment on PROMISE12 dataset. Similar to the LA dataset, we compared H-SAM against SAM efficient adaptation methods (AutoSAM [30], SAM Adapter [9], and SAMed [70]), nnUnet [34], and semi-supervised methods (UA-MT [67], MC-Net [65], SS-Net [66], and MLB-Seg [63]). For a fair comparison with the semi-supervised methods, we strictly follow MLB-Seg [63] and use the same 3 labeled cases as our training set under the few-shot setting, while the semi-supervised methods incorporated an additional 37 unlabeled cases. And all the methods are tested on the same test dataset. As shown in Table 3, H-SAM achieved a significant improvement in Dice coefficients (approximately 10.94%) over the semi-supervised method MLB-Seg, despite being trained on only three labeled cases. This result highlights the efficiency of H-SAM in leveraging limited labeled data. Compared to the SAM efficient adaptation method SAMed [70], our method also demonstrated superior performance (87.27% vs. 86.00%), further establishing its effectiveness in few-

Methods	Scans used		Mean Dice (%) $\uparrow$
	Labeled	Unlabeled	
UA-MT [67]			65.05
DTC [47]			63.44
SASSNet [42]			73.43
MC-Net [65]	3(7.5%)	37(92.5%)	72.66
SS-Net [66]			73.19
Self-Paced [54]			74.02
MLB-Seg [63]			78.27
nnUnet [34]			84.22
AutoSAM [30]			68.40
SAM Adapter [9]	3(7.5%)	0(0%)	75.45
SAMed [70]			86.00
Ours			<b>87.27</b>

Table 3. Results of PROMISE12 dataset under semi-supervision.

Learnable Mask-Attention	Hierarchical Pixel Decoder	CM Self-Attention	Mean Dice (%)
$\times$	$\times$	$\times$	75.57
$\checkmark$	$\times$	$\times$	77.68
$\checkmark$	$\checkmark$	$\times$	78.58
$\times$	$\checkmark$	$\times$	77.05
$\times$	$\checkmark$	$\checkmark$	79.03
$\times$	$\times$	$\checkmark$	77.71
$\checkmark$	$\times$	$\checkmark$	78.76
$\checkmark$	$\checkmark$	$\checkmark$	<b>80.35</b>

Table 4. Effectiveness of the key contributions in H-SAM: Learnable Mask-Attention, CMAttn and Hierarchical Pixel Decoder.

Methods	Mean Dice (%)
w/o. mask-attention	75.57
original mask-attention	75.61
learnable mask-attention (ours)	<b>77.68</b>

Table 5. Effectiveness of our learnable mask-attention against original mask-attention and baseline.

shot semantic segmentation tasks.

#### 4.4. Ablation study

**Effectiveness of Learnable Mask Cross Attention** As shown in Table 4, our learnable mask attention improves the baseline by 2.1%. Within the hierarchical decoding structure of H-SAM, our learnable mask cross attention shows no signal of performance saturation with other two key contributions: CMAttn and Hierarchical Pixel Decoder. We also compare learnable mask cross attention against normal cross attention and unlearnable mask attention in Table 5. Note that, unlearnable mask attention operation brings little improvement due to the lack of gradient backpropagation. Inversely, our proposed learnable mask cross attention brings an immediate performance increase of 2.1% with jointing training and inheritance of mask-guided prior.

**Effectiveness of CMAtn** In Table 4, we present the ablation of Class-Balanced Mask-Guided Self-Attention (CMAtn). The CMAtn alone brings a 1.2% improvement to the baseline, showing that SAM benefits from a more informed image embedding as input for the mask decoder. Combined with Learnable Mask-Attention, the two mask-guided implementations improve the baseline model by 3.2% in terms of Mean Dice.

**Effectiveness of Hierarchical Pixel Decoder** In Table 4, the hierarchical pixel decoder promotes the dice coefficient by 2.2%. Then we evaluate the effectiveness of a combination of mask-guided implementations with the hierarchical pixel decoder by adding or disabling one of them each time. Table 4 shows an improvement of 3.0% and 3.5%, respectively combined with Learnable Mask-Attention and CMAtn.

#### 4.5. Efficiency Analysis

To show that the advantages of our hierarchical decoding are not achieved by additional training cost in the decoding section, here we conduct ablation experiments between our model and other prompt-free SAM variants in terms of total parameters and performance. AutoSAM [30] is not listed in the table because it freezes the image encoder with no adapters. The original lightweight SAM mask decoder consists of only 2 transformer layers. SAMed [70] adds LoRA adapter in the image encoder, while remaining a default SAM mask decoder with 2 transformer layers. SAM Adapter [9] injects adapter layers in both the image encoder and mask decoder. In general, the training cost of our H-SAM is equal to SAMed with 1 additional lightweight mask decoder. For a fair comparison, we double and dribble the number of SAMed transformer layers into 4 and 6, to achieve a comparable and even larger parameter scale than the training cost in our H-SAM. As shown in Table 6, SAMed [70] benefits from increased transformer layers. However, the performance promotion comes with huge computation costs. H-SAM consistently outperforms SAMed. Compared with SAM Adapter, H-SAM shows a 7.5% performance improvement with a 20M lesser parameter scale. The results prove the superiority of our H-SAM in wisely generating finer medical segmentation with little extra computation cost.

#### 4.6. Qualitative Results

As shown in Figure 5, we compare our H-SAM to other prompt-free medical SAM variants, including AutoSAM [9], SAM Adapter [30] and SAMed [70]. As pointed out in the first row and the last row, compared to other SAM variants, H-SAM provides a precise mask prediction with lesser noise. In the second row, where other methods misrecognize Aorta and Pancreas to Stomach and

Methods	Transformer Layers	Total Parameter	Mean Dice (%)
SAMed [70]	2	108.8M	75.57
	4	112.5M	76.80
	6	116.2M	78.05
SAM Adapter [9]	2	131.5M	72.80
H-SAM(ours)	4	112.3M	<b>80.35</b>

Table 6. Efficiency analysis of our H-SAM against deeper default SAM mask decoder. H-SAM shows better performance with fewer parameters.

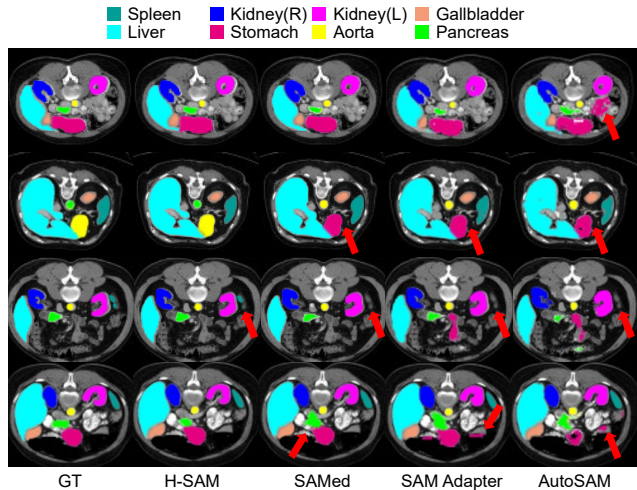


Figure 5. The qualitative results of H-SAM and other SAM variants, including SAMed, SAM Adapter, and AutoSAM.

Aorta, H-SAM provides correctly attributes each organ to their categories. H-SAM also performs superiorly with small-scale organs. In the third row, while all other variants miss Spleen, only H-SAM provides correct prediction for all of the organs.

## 5. Conclusion

We present H-SAM, which is a simple and efficient hierarchical mask decoder for adaptation of Segment Anything Model on medical image segmentation. Using a probabilistic map from a default decoder as prior to guide finer medical segmentation in the sequential decoding unit, our H-SAM puts forward a new direction of SAM adaptation. Notably, H-SAM achieves this superior performance without relying on any unlabeled data, surpassing even state-of-the-art semi-supervised models that use extensive unlabeled datasets in various medical imaging contexts. This underscores H-SAM’s significant potential in advancing the field of medical image segmentation, offering a robust, efficient, and data-economic solution.



## References

- [1] 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. 6
- [2] Reza Azad, René Arimond, Ehsan Khodapanah Aghdam, Amirhossein Kazerooni, and Dorit Merhof. Dae-former: Dual attention-guided efficient transformer for medical image segmentation. In *International Workshop on PRedictive Intelligence In MEdicine*, pages 83–95. Springer, 2023. 6
- [3] Yunhao Bai, Duowen Chen, Qingli Li, Wei Shen, and Yan Wang. Bidirectional copy-paste for semi-supervised medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11514–11524, 2023. 6, 7
- [4] 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. 3, 6
- [5] Shurong Chai, Rahul Kumar Jain, Shiyu Teng, Jiaqing Liu, Yinhao Li, Tomoko Tateyama, and Yen-wei Chen. Ladder fine-tuning approach for sam integrating complementary network. *arXiv preprint arXiv:2306.12737*, 2023. 2
- [6] Chen Chen, Wenjia Bai, and Daniel Rueckert. Multi-task learning for left atrial segmentation on ge-mri. In *Statistical Atlases and Computational Models of the Heart. Atrial Segmentation and LV Quantification Challenges: 9th International Workshop, STACOM 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers 9*, pages 292–301. Springer, 2019. 5, 6
- [7] 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. 3, 5, 6
- [8] Jieneng Chen, Jieru Mei, Xianhang Li, Yongyi Lu, Qihang Yu, Qingyue Wei, Xiangde Luo, Yutong Xie, Ehsan Adeli, Yan Wang, et al. 3d transunet: Advancing medical image segmentation through vision transformers. *arXiv preprint arXiv:2310.07781*, 2023. 3
- [9] 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, 2023. 2, 6, 7, 8
- [10] 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. 3
- [11] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 4
- [12] Dongjie Cheng, Ziyuan Qin, Zekun Jiang, Shaoting Zhang, Qicheng Lao, and Kang Li. Sam on medical images: A comprehensive study on three prompt modes. *arXiv preprint arXiv:2305.00035*, 2023. 2
- [13] Junlong Cheng, Shengwei Tian, Long Yu, Chengrui Gao, Xiaojing Kang, Xiang Ma, Weidong Wu, Shijia Liu, and Hongchun Lu. Resgnet: Residual group attention network for medical image classification and segmentation. *Medical Image Analysis*, 76:102313, 2022. 1
- [14] Can Cui, Ruining Deng, Quan Liu, Tianyuan Yao, Shunxing Bao, Lucas W Remedios, Yucheng Tang, and Yuankai Huo. All-in-sam: from weak annotation to pixel-wise nuclei segmentation with prompt-based finetuning. *arXiv preprint arXiv:2307.00290*, 2023. 2
- [15] Jeffrey De Fauw, Joseph R Ledsam, Bernardino Romera-Paredes, Stanislav Nikolov, Nenad Tomasev, Sam Blackwell, Harry Askham, Xavier Glorot, Brendan O’Donoghue, Daniel Visentin, et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature medicine*, 24(9):1342–1350, 2018. 1
- [16] Guoyao Deng, Ke Zou, Kai Ren, Meng Wang, Xuedong Yuan, Sancong Ying, and Huazhu Fu. Sam-u: Multi-box prompts triggered uncertainty estimation for reliable sam in medical image. *arXiv preprint arXiv:2307.04973*, 2023. 2
- [17] Ruining Deng, Can Cui, Quan Liu, Tianyuan Yao, Lucas W Remedios, Shunxing Bao, Bennett A Landman, Lee E Wheless, Lori A Coburn, Keith T Wilson, et al. Segment anything model (sam) for digital pathology: Assess zero-shot segmentation on whole slide imaging. *arXiv preprint arXiv:2304.04155*, 2023. 2
- [18] 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. 2
- [19] Foivos I Diakogiannis, François Waldner, Peter Caccetta, and Chen Wu. Resunet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 162:94–114, 2020. 3
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [21] Weijia Feng, Lingting Zhu, and Lequan Yu. Cheap lunch for medical image segmentation by fine-tuning sam on few exemplars. *arXiv preprint arXiv:2308.14133*, 2023. 2
- [22] Yabo Fu, Yang Lei, Tonghe Wang, Walter J Curran, Tian Liu, and Xiaofeng Yang. A review of deep learning based methods for medical image multi-organ segmentation. *Physica Medica*, 85:107–122, 2021. 3
- [23] Yifan Gao, Wei Xia, Dingdu Hu, and Xin Gao. Desam: Decoupling segment anything model for generalizable medical image segmentation. *arXiv preprint arXiv:2306.00499*, 2023. 2

- [24] 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. *arXiv preprint arXiv:2306.13465*, 2023. [2](#)
- [25] Sheng He, Rina Bao, Jingpeng Li, P Ellen Grant, and Yangming Ou. Accuracy of segment-anything model (sam) in medical image segmentation tasks. *arXiv preprint arXiv:2304.09324*, 2023. [2](#)
- [26] 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. [3](#)
- [27] Chuanfei Hu and Xinde Li. When sam meets medical images: An investigation of segment anything model (sam) on multi-phase liver tumor segmentation. *arXiv preprint arXiv:2304.08506*, 2023. [2](#)
- [28] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. [2](#), [3](#)
- [29] Mingzhe Hu, Yuheng Li, and Xiaofeng Yang. Skinsam: Empowering skin cancer segmentation with segment anything model. *arXiv preprint arXiv:2304.13973*, 2023. [2](#)
- [30] Xinrong Hu, Xiaowei Xu, and Yiyu Shi. How to efficiently adapt large segmentation model (sam) to medical images. *arXiv preprint arXiv:2306.13731*, 2023. [2](#), [6](#), [7](#), [8](#)
- [31] Xiaohong Huang, Zhifang Deng, Dandan Li, and Xueguang Yuan. Missformer: An effective medical image segmentation transformer. *arXiv preprint arXiv:2109.07162*, 2021. [3](#)
- [32] Yuhao Huang, Xin Yang, Lian Liu, Han Zhou, Ao Chang, Xinrui Zhou, Rusi Chen, Junxuan Yu, Jiongquan Chen, Chaoyu Chen, et al. Segment anything model for medical images? *arXiv preprint arXiv:2304.14660*, 2023. [2](#)
- [33] Fabian Isensee, Jens Petersen, Andre Klein, David Zimmerer, Paul F Jaeger, Simon Kohl, Jakob Wasserthal, Gregor Koehler, Tobias Norajitra, Sebastian Wirkert, et al. nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486*, 2018. [3](#)
- [34] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021. [7](#)
- [35] Ge-Peng Ji, Deng-Ping Fan, Peng Xu, Ming-Ming Cheng, Bowen Zhou, and Luc Van Gool. Sam struggles in concealed scenes—empirical study on “segment anything”. *arXiv preprint arXiv:2304.06022*, 2023. [2](#)
- [36] Yuanfeng Ji, Haotian Bai, Chongjian Ge, Jie Yang, Ye Zhu, Ruimao Zhang, Zhen Li, Lingyan Zhannng, Wanling Ma, Xiang Wan, et al. Amos: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation. *Advances in Neural Information Processing Systems*, 35: 36722–36732, 2022. [3](#)
- [37] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. [1](#), [2](#)
- [38] Bennett Landman, Zhoubing Xu, J Igelsias, Martin Styner, T Langerak, and Arno Klein. Miccai multi-atlas labeling beyond the cranial vault—workshop and challenge. In *Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge*, page 12, 2015. [5](#)
- [39] Wenhui Lei, Xu Wei, Xiaofan Zhang, Kang Li, and Shaoting Zhang. Medlsam: Localize and segment anything model for 3d medical images. *arXiv preprint arXiv:2306.14752*, 2023. [2](#)
- [40] Chengyin Li, Prashant Khanduri, Yao Qiang, Rafi Ibn Sultan, Indrin Chetty, and Dongxiao Zhu. Auto-prompting sam for mobile friendly 3d medical image segmentation. *arXiv preprint arXiv:2308.14936*, 2023. [2](#)
- [41] Mengke Li, Yiu-ming Cheung, and Yang Lu. Long-tailed visual recognition via gaussian clouded logit adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6929–6938, 2022. [4](#)
- [42] Shuailin Li, Chuyu Zhang, and Xuming He. Shape-aware semi-supervised 3d semantic segmentation for medical images. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I 23*, pages 552–561. Springer, 2020. [7](#)
- [43] Yuheng Li, Mingzhe Hu, and Xiaofeng Yang. Polypsam: Transfer sam for polyp segmentation. *arXiv preprint arXiv:2305.00293*, 2023. [2](#)
- [44] Yi Lin, Yufan Chen, Kwang-Ting Cheng, and Hao Chen. Few shot medical image segmentation with cross attention transformer. *arXiv preprint arXiv:2303.13867*, 2023. [3](#)
- [45] Geert Litjens, Robert Toth, Wendy Van De Ven, Caroline Hoeks, Sjoerd Kerkstra, Bram Van Ginneken, Graham Vincent, Gwenael Guillard, Neil Birbeck, Jindang Zhang, et al. Evaluation of prostate segmentation algorithms for mri: the promise12 challenge. *Medical image analysis*, 18(2):359–373, 2014. [5](#), [6](#)
- [46] Yihao Liu, Jiaming Zhang, Zhangcong She, Amir Kheradmand, and Mehran Armand. Sann (segment any medical model): A 3d slicer integration to sam. *arXiv preprint arXiv:2304.05622*, 2023. [2](#)
- [47] Xiangde Luo, Jieneng Chen, Tao Song, and Guotai Wang. Semi-supervised medical image segmentation through dual-task consistency. In *Proceedings of the AAAI conference on artificial intelligence*, pages 8801–8809, 2021. [7](#)
- [48] Xiangde Luo, Wenjun Liao, Jieneng Chen, Tao Song, Yinan Chen, Shichuan Zhang, Nianying Chen, Guotai Wang, and Shaoting Zhang. Efficient semi-supervised gross target volume of nasopharyngeal carcinoma segmentation via uncertainty rectified pyramid consistency. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part II 24*, pages 318–329. Springer, 2021. [7](#)
- [49] Jun Ma and Bo Wang. Segment anything in medical images. *arXiv preprint arXiv:2304.12306*, 2023. [2](#)
- [50] 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*, 89:102918, 2023. [2](#)

- [51] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016. [5](#)
- [52] S Mohapatra, A Gosai, and G Schlaug. Sam vs bet: A comparative study for brain extraction and segmentation of magnetic resonance images using deep learning. *arXiv preprint arXiv:2304.04738*, 2:4, 2023. [2](#)
- [53] David Ouyang, Bryan He, Amirata Ghorbani, Neal Yuan, Joseph Ebinger, Curtis P Langlotz, Paul A Heidenreich, Robert A Harrington, David H Liang, Euan A Ashley, et al. Video-based ai for beat-to-beat assessment of cardiac function. *Nature*, 580(7802):252–256, 2020. [1](#)
- [54] Jizong Peng, Ping Wang, Christian Desrosiers, and Marco Pedersoli. Self-paced contrastive learning for semi-supervised medical image segmentation with meta-labels. *Advances in Neural Information Processing Systems*, 34: 16686–16699, 2021. [7](#)
- [55] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#)
- [56] Md Mostafijur Rahman and Radu Marculescu. G-cascade: Efficient cascaded graph convolutional decoding for 2d medical image segmentation. *arXiv preprint arXiv:2310.16175*, 2023. [2](#)
- [57] Md Mostafijur Rahman and Radu Marculescu. Multi-scale hierarchical vision transformer with cascaded attention decoding for medical image segmentation. *arXiv preprint arXiv:2303.16892*, 2023. [6](#)
- [58] 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. [3](#)
- [59] 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. [3](#)
- [60] Tassilo Wald, Saikat Roy, Gregor Koehler, Nico Disch, Maximilian Rouven Rokuss, Julius Holzschuh, David Zimmerer, and Klaus Maier-Hein. Sam. md: Zero-shot medical image segmentation capabilities of the segment anything model. In *Medical Imaging with Deep Learning, short paper track*, 2023. [2](#)
- [61] An Wang, Mobarakol Islam, Mengya Xu, Yang Zhang, and Hongliang Ren. Sam meets robotic surgery: An empirical study in robustness perspective. *arXiv preprint arXiv:2304.14674*, 2023. [2](#)
- [62] Yan Wang, Yuyin Zhou, Wei Shen, Seyoun Park, Elliot K Fishman, and Alan L Yuille. Abdominal multi-organ segmentation with organ-attention networks and statistical fusion. *Medical image analysis*, 55:88–102, 2019. [3](#)
- [63] Qingyue Wei, Lequan Yu, Xianhang Li, Wei Shao, Cihang Xie, Lei Xing, and Yuyin Zhou. Consistency-guided meta-learning for bootstrapping semi-supervised medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 183–193. Springer, 2023. [6, 7](#)
- [64] 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. [2](#)
- [65] Yicheng Wu, Minfeng Xu, Zongyuan Ge, Jianfei Cai, and Lei Zhang. Semi-supervised left atrium segmentation with mutual consistency training. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part II 24*, pages 297–306. Springer, 2021. [6, 7](#)
- [66] Yicheng Wu, Zhonghua Wu, Qianyi Wu, Zongyuan Ge, and Jianfei Cai. Exploring smoothness and class-separation for semi-supervised medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 34–43. Springer, 2022. [6, 7](#)
- [67] Lequan Yu, Shujun Wang, Xiaomeng Li, Chi-Wing Fu, and Pheng-Ann Heng. Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*, pages 605–613. Springer, 2019. [6, 7](#)
- [68] Chenyi Zeng, Lin Gu, Zhenzhong Liu, and Shen Zhao. Review of deep learning approaches for the segmentation of multiple sclerosis lesions on brain mri. *Frontiers in Neuroinformatics*, 14:610967, 2020. [3](#)
- [69] Jingwei Zhang, Ke Ma, Saarthak Kapse, Joel Saltz, Maria Vakalopoulou, Prateek Prasanna, and Dimitris Samaras. Sam-path: A segment anything model for semantic segmentation in digital pathology. *arXiv preprint arXiv:2307.09570*, 2023. [2](#)
- [70] Kaidong Zhang and Dong Liu. Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785*, 2023. [2, 3, 5, 6, 7, 8](#)
- [71] Lian Zhang, Zhengliang Liu, Lu Zhang, Zihao Wu, Xiaowei Yu, Jason Holmes, Hongying Feng, Haixing Dai, Xiang Li, Quanzheng Li, et al. Segment anything model (sam) for radiation oncology. *arXiv preprint arXiv:2306.11730*, 2023. [2](#)
- [72] Yiming Zhang, Tianang Leng, Kun Han, and Xiaohui Xie. Self-sampling meta sam: Enhancing few-shot medical image segmentation with meta-learning. *arXiv preprint arXiv:2308.16466*, 2023. [2](#)
- [73] Yizhe Zhang, Tao Zhou, Peixian Liang, and Danny Z Chen. Input augmentation with sam: Boosting medical image segmentation with segmentation foundation model. *arXiv preprint arXiv:2304.11332*, 2023. [2](#)
- [74] 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. 3
- [75] Yuyin Zhou, Yan Wang, Peng Tang, Song Bai, Wei Shen, Elliot Fishman, and Alan Yuille. Semi-supervised 3d abdominal multi-organ segmentation via deep multi-planar co-training. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 121–140. IEEE, 2019. 3
- [76] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer, 2018. 3