

# How to build the best medical image segmentation algorithm using foundation models: a comprehensive empirical study with Segment Anything Model

Hanxue Gu <sup>1\*</sup>, Haoyu Dong <sup>1\*</sup>, Jichen Yang <sup>1</sup>, Maciej A. Mazurowski <sup>1,2,3,4</sup>

<sup>1</sup> Department of Electrical and Computer Engineering, Duke University, US

<sup>2</sup> Department of Biostatistics & Bioinformatics, Duke University, US

<sup>3</sup> Department of Computer Science, Duke University, US

<sup>4</sup> Department of Radiology, Duke University, US

\*These authors contributed equally to this work.

## Abstract

Automated segmentation is a fundamental medical image analysis task, which enjoys significant advances due to the advent of deep learning. While foundation models have been useful in natural language processing and some vision tasks for some time, the foundation model developed with image segmentation in mind - Segment Anything Model (SAM) - has been developed only recently and has shown similar promise. However, there are still no systematic analyses or “best-practice” guidelines for optimal fine-tuning of SAM for medical image segmentation. This work summarizes existing fine-tuning strategies with various backbone architectures, model components, and fine-tuning algorithms across 18 combinations and evaluates them on 17 datasets covering all common radiology modalities. Our study reveals that (1) fine-tuning SAM leads to slightly better performance than previous segmentation methods, (2) fine-tuning strategies that use parameter-efficient learning in both the encoder and decoder are superior to other strategies, (3) network architecture has a small impact on final performance, and (4) further training SAM with self-supervised learning can improve final model performance. We also demonstrate the ineffectiveness of some methods popular in the literature and further expand our experiments into few-shot and prompt-based settings. Lastly, we released our code and MRI-specific fine-tuned weights, which consistently obtained superior performance over the original SAM, at <https://github.com/mazurowski-lab/finetune-SAM>.

## Keywords

Machine Learning, Image Segmentation, Vision Foundation Model

## Article informations

<https://doi.org/doi.org/10.59275/j.melba.2025-86a6>

©2025 Gu, Dong, Yang and Mazurowski. License: CC-BY 4.0

Received: 2024-08-22, Published 2025-05-08

Corresponding author: [hanxue.gu@duke.edu](mailto:hanxue.gu@duke.edu)



## 1. Introduction

Medical image segmentation refers to identifying the region of interest, typically anatomical structures or tumors, in an image. This process is essential in the field of medical image analysis, contributing significantly to the effectiveness of computer-aided diagnosis (Al-Antari et al., 2018). The introduction of deep learning techniques has led to a variety of automated segmentation methods (Liu and Pang, 2020; Weng et al., 2023), further enhancing the precision and efficiency of analyzing medical images. Despite the progress made, it is still difficult to determine the most effective

method for developing an optimal segmentation model, especially with the recent introduction of vision foundation models dedicated to segmentation. Therefore, in our study, we aim to explore and define the optimal strategies for developing automated segmentation algorithms specifically designed for medical images and with a focus on foundation models.

In the era of deep learning, a common practice for automated medical image segmentation is to directly train UNet (Ronneberger et al., 2015) or its variants (Siddique et al., 2021; Wang et al., 2022) on a specific dataset and

task. Recently, there has been growing interest in customizing vision foundation models for specific segmentation tasks (Zhang and Liu, 2023; Gu et al., 2024). Foundation models are typically transformer-based (Vaswani et al., 2017) neural networks pre-trained on extremely large datasets, allowing them to generalize to a wide range of tasks and datasets. The Segment Anything Model (SAM) (Kirillov et al., 2023) was proposed as the first large model designed specifically for segmentation tasks. SAM employs a prompt-based training method, aiming to generate segmentation masks for target objects in response to user-provided prompts, e.g., points or bounding boxes. It has demonstrated comparable or even superior zero-shot performance to fully supervised models in processing natural images. However, SAM performance is less impressive when applied directly to medical images (Mazurowski et al., 2023; Deng et al., 2023; Huang et al., 2024), and the requirement of prompts makes its direct usage with medical imaging segmentation tasks challenging.

Removing the prompt requirement in SAM is relatively straightforward, as we can use “dummy” prompt embeddings as inputs during fine-tuning, a strategy verified by concurrent work (Zhang and Liu, 2023). However, the approaches proposed for adapting SAM to medical imaging tasks vary widely. Some works focused on adapting SAM’s decoder only (Yue et al., 2023), while others chose to modify the entire network (Zhang and Liu, 2023). Some works introduced an additional pre-training stage (Ma et al., 2024) using multiple medical images, while others usually do not. Though some prior survey work (Lee et al., 2024; Zhang and Jiao, 2023) have attempted to list different strategies and categorize them into various groups, questions remain regarding (1) which fine-tuned strategy is the most effective, (2) whether these models can achieve better performance than traditional UNets, and (3) whether using additional data is helpful.

In this work, we aim to answer these questions by considering three common dataset availability scenarios in the medical imaging field: (1) only a single labeled dataset; (2) multiple labeled datasets for different tasks; and (3) multiple labeled and unlabeled datasets. The first scenario reflects the most common deep learning scenario of adapting SAM to a dataset of interest, and the latter two scenarios offer the opportunity to first incorporate broad medical domain knowledge into SAM using supervised or self-supervised training before adapting it to a specific dataset. Figure 1 summarizes the strategies and experiment setups we include. To avoid terminology confusion, we name these strategies as **task-specific supervised training**, **task-expansive prompt-based learning**, and **task-agnostic self-supervised learning** respectively.

When considering the design choice during task-specific supervised training, our work aligns with several concurrent works (Lee et al., 2024; Zhang and Jiao, 2023; Zhang

et al., 2024b). However, our work differs from theirs in that we conduct comprehensive experiments over 17 medical imaging datasets across various modalities for a fair comparison, rather than simply compiling disparate results, and provide an organized and new catalog that could summarize existing approaches. The choice of each category is shown in Figure 2 and will be discussed in Section 2.2. Our category definition reflects the common design choices when building a segmentation algorithm from SAM, and we can categorize most existing works based on the proposed criteria, shown in Table 1.

When incorporating extra medical knowledge into SAM, we explore two main setups of pre-training approaches: task-expansive prompt-based learning and task-agnostic self-supervised learning, depending on the availability of labels. Note that our objective is not to propose a novel pre-training strategy, as such strategies are often constrained by the availability of computation resources and datasets. Instead, our focus is on determining if additional pre-training with more available medical data is helpful and, if so, what the best combination of pre-training and task-specific fine-tuning is. Therefore, we adopt the strategies used during SAM’s development, and one of the most commonly applied pre-training strategies in this field. Namely, we use the same interactive training pipeline where prompts are simulated from the masks for task-expansive prompt-based learning; and masked autoencoding (MAE) (He et al., 2022) as the objective for task-agnostic self-supervised learning. We evaluate these approaches on the same 17 datasets.

Our study is extended to two crucial scenarios when deploying a segmentation algorithm for annotation purposes: few-shot learning and interactive segmentation. The first scenario reflects the case of annotating a new dataset with only a limited number of annotations, while the latter examines the effectiveness of using tools to assist the annotation process. In both scenarios, we limit the evaluation datasets to MRI. When conducting few-shot learning, we limit the number of training samples to 5 and re-examine the effectiveness of setups 1-3 described in Figure 1. In the interactive segmentation scenario, we focus on task-specific supervised training with either boxes or points as input prompts.

In conclusion, we systematically study different approaches to fine-tune SAM in the medical imaging field. Our paper investigates an important question: What is the best practice to develop an automated segmentation algorithm based on foundation models? We find the answers depend on the availability of the datasets and consider three scenarios: (1) only a single labeled dataset; (2) multiple labeled datasets (with different objects of interest); and (3) multiple labeled and unlabeled datasets. We propose task-specific supervised training, task-expansive prompt-based learning followed by task-specific supervised training,

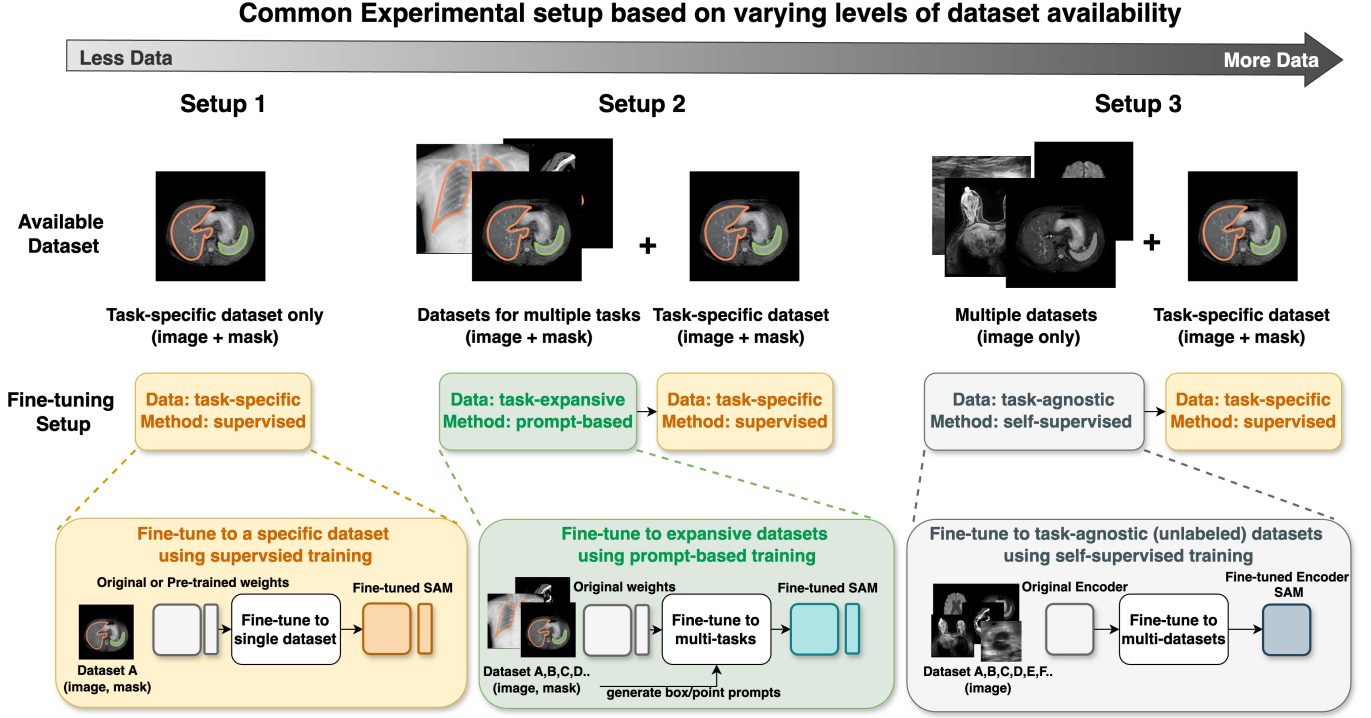


Figure 1: Overview of general fine-tuning strategies based on different levels of dataset availability.

and task-agnostic self-supervised learning followed by task-specific supervised training as the solution to each scenario, respectively. By conducting systematic experiments on each strategy, we observe several useful guidelines and summarize them in Section 5.

## 2. Method

In this section, we first recapitulate SAM’s structure. Next, we describe the selected methods under different data availability scenarios.

### 2.1 SAM and removal of prompts

To begin with, we first briefly present an overview of the Segment Anything Model (SAM) (Kirillov et al., 2023) before introducing the fine-tuning strategies. SAM consists of three components: (1) image encoder, (2) prompt encoder, and (3) mask decoder, as illustrated in Figure 2 of the overview of SAM’s architecture. The image encoder, built on a robust Vision Transformer (ViT) framework, effectively converts a 2D image into a latent feature representation. Notably, in SAM’s original architecture, the image encoder offers flexibility in model size, with options including ViT-H(uge), ViT-L(arge), and ViT-B(ase) (Dosovitskiy et al., 2020). The prompt encoder is flexible, accommodating different prompt types such as point, box, or mask prompts. Depending on the input type, it generates either sparse embeddings for point/box prompts or dense embeddings for mask prompts. Lastly, the mask decoder is a transformer

decoder using two-way cross-attention that integrates image and prompt embeddings to produce a multi-channel mask output, each channel representing varying confidence levels. SAM is trained on SA-1B, a dataset collected by the same team behind its development, which contains more than 1B masks from 11M natural images. During training, first, with equal probabilities, box and point prompts are generated. Point prompts are sampled uniformly from the ground truth mask, and box prompts are taken as the ground truth mask’s bounding box, with random noise added in each coordinate. The objective is the combination of focal loss (Ross and Dollár, 2017) and dice loss (Milletari et al., 2016).

#### 2.1.1 Interactive segmentation to automated segmentation

When converting SAM to auto-mode (without prompts), we can simply use the default embeddings generated by using “None” as inputs to the prompt encoder, i.e.

```
sparse_embeddings, dense_embeddings =
sam.prompt_encoder(points=None,
boxes=None, masks=None),
```

where the “sparse\_embeddings” and “dense\_embeddings” serve as inputs to SAM’s mask decoder. Empirical evidence shows the effectiveness of this simple approach (Zhang and Liu, 2023). More advanced methods, including adding a detection network to automatically generate box prompts or a self-prompt module into SAM’s structure (Wu et al., 2023b; Gao et al., 2023; Shaharabany et al., 2023), are

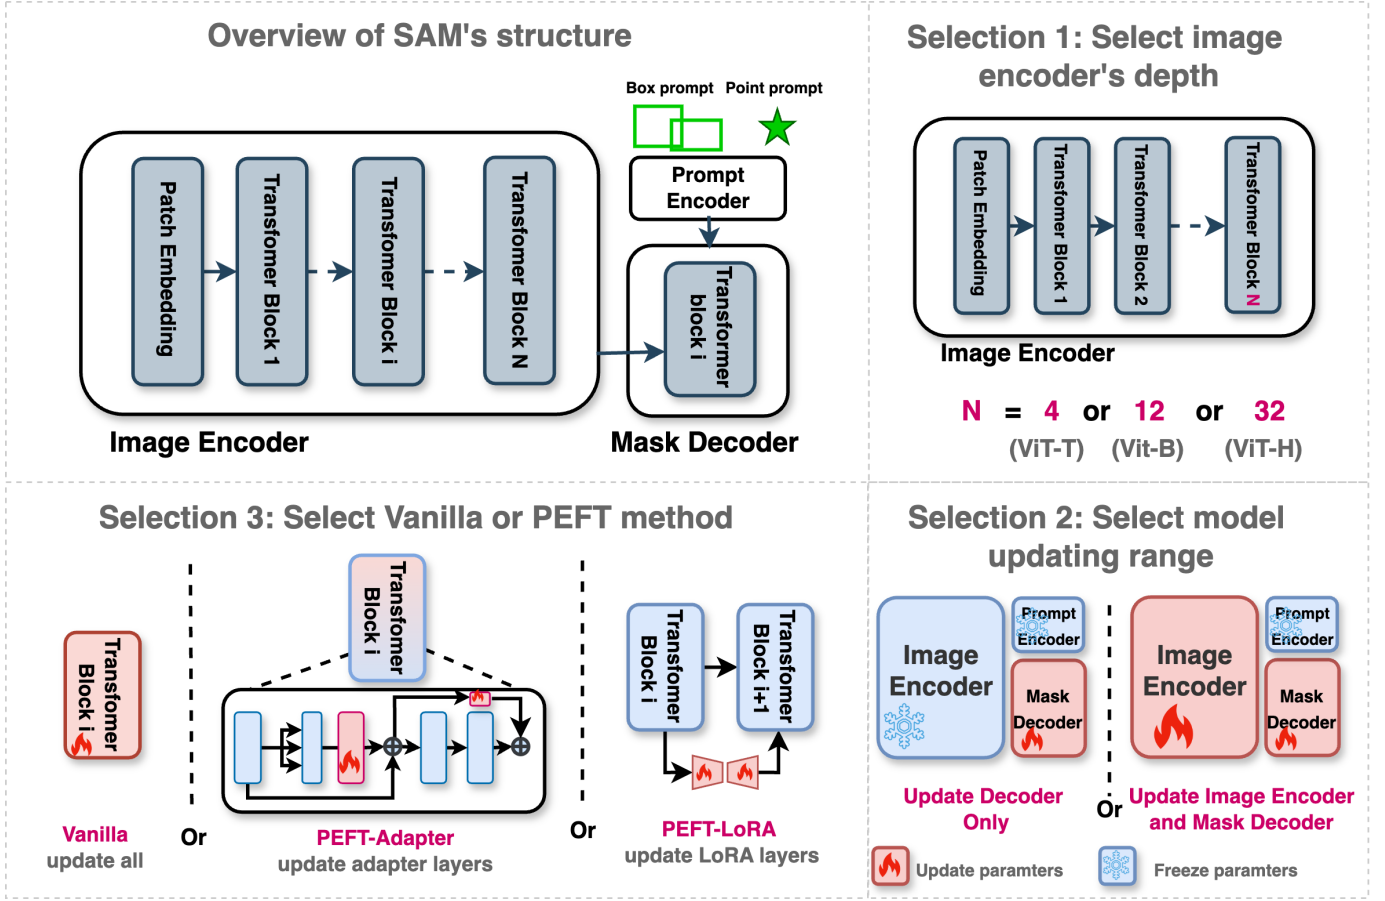


Figure 2: Visualization of the overview of the SAM's architecture and task-specific fine-tuning architectures selected in our study: including 3 encoder architecture  $\times$  2 model updating range  $\times$  3 vanilla/PEFT methods = 18 choices.

feasible but will introduce extra complexity into SAM's structure, hampering our analysis on the effectiveness of fine-tuning strategies. Moreover, we believe our work offers a potentially effective strategy, i.e., feeding the generated prompts to the fine-tuned SAM instead of the original one, for this research direction.

## 2.2 Variables in task-specific supervised training

We consider three aspects in task-specific supervised training that can affect the performance, as shown in Figure 2. The choice of each category is discussed next.

### 2.2.1 Image Encoder Architecture

In this work, we select ViT-B and ViT-H from SAM's provided checkpoints because they represent the most efficient and effective versions of SAM's encoder, respectively. Additionally, we explore a more compact option: ViT-T(iny). This backbone is introduced in MobileSAM (Zhang et al., 2023a), where the authors distill the knowledge from SAM's pre-trained ViT-H to ViT-T, a customized vision transformer with only 1% of the network parameters, and achieve an inference speed that is 50 times faster compared to the original SAM. The lightweight nature can widen the appli-

cation of SAM, ease the computational requirements, and, more importantly, raise interesting research questions, such as whether distillation can maintain or even lead to the superior performance of SAM during fine-tuning.

### 2.2.2 Model Component

When selecting which set(s) of SAM's parameters to update, the most straightforward way is to select one or multiple components from SAM. As discussed in Section 2.1.1, we use the default prompt embeddings and thus can keep the prompt encoder as is.

Some methods state the effectiveness of maintaining the encoder's weights and only adapting the lightweight mask decoder of SAM's architecture, as shown in Table 1. These approaches are grounded in the understanding that SAM's encoder, trained on a comprehensive collection of natural images, is sufficiently powerful and knowledgeable in the general imaging domain, with zero-shot ability to perform feature extraction. Updating the lightweight decoder can enhance performance in specific applications. Also, it is suggested that fine-tuning the encoder of SAM would demand considerably higher computational resources. However, due to the significant appeal difference between



natural images and medical images, it remains unclear if SAM's encoder could effectively capture the distinctive features of medical images without any modifications to the image encoder. Therefore, we consider these two choices: *mask decoder only* and *mask decoder plus image encoder*.

### 2.2.3 Vanilla or Parameter-efficient Fine-tuning Method

When conducting model fine-tuning or transfer learning, the easiest and most intuitive approach is to update all model parameters involved when feeding medical domain data, which is known as vanilla fine-tuning (Azizi et al., 2022).

Recently, parameter-efficient fine-tuning (PEFT) techniques have gained increasing attention. Compared with vanilla fine-tuning, PEFT keeps most of the network parameters fixed and only updates a small percentage, i.e. less than 5%, of them. Such a strategy significantly reduces computational cost and is beneficial in preventing both catastrophic forgetting and overfitting (Dutt et al., 2023). The simplest PEFT techniques involve selecting and fine-tuning a subset of the parameters from the original pre-trained model. These selective methods concentrate on identifying the optimal parameters for fine-tuning, aiming to strike an ideal balance between parameter efficiency and overall model performance. For instance, BitFit (Zaken et al., 2021) proposed fine-tuning only the bias term in a transformer-based network, and Tinyt1 (Cai et al., 2020) shares a similar idea but focuses on CNN-based networks. Touvron et al. (Touvron et al., 2022) have demonstrated the effectiveness of selectively fine-tuning only the attention layers in Vision Transformers (ViTs), whereas alternative approaches suggest that exclusively optimizing the normalization layers (Basu et al., 2023).

Rather than updating certain selected layers, recent research (Li et al., 2022) indicates that incorporating extra modules into the large models appears to be more beneficial because it reconfigures models towards a downstream task rather than solely updating network parameters. Among these additional module methods, Adapters, which were originally introduced for learning multi-domain representations (Rebuffi et al., 2017, 2018), have emerged as a particularly prevalent choice. In addition to the conventional Adapter block (Houlsby et al., 2019), Low-Rank Adaptation (LoRA), known for its efficient low-rank matrices added into self-attention layers that require even fewer parameters, has also gained popularity (Hu et al., 2021). Additionally, AdaptFormer (Chen et al., 2022), which integrates a scalable Adapter into the multi-layer perceptron (MLP) layer of the attention block specifically for vision transformers, has shown that adding just 2% extra parameters can surpass the performance of fully fine-tuned networks across various natural imaging benchmarks. In our study, we aim to inves-

tigate the effectiveness of applying PEFT for fine-tuning SAM within the medical imaging domain. Considering that Adapter and LoRA are the most popular PEFT methods adopted in LLMs and vision foundation transformers (Xu et al., 2023), we choose to include them, as well as the simplest Vanilla, in our experiments. The implementation details of Adapter and LoRA are presented in Figure 2, Selection 3. Specifically, Adapter introduces two additional layers in each transformer block, and LoRA adds trainable pairs of rank decomposition matrices in parallel to existing blocks. Following (Hu et al., 2021), we set the rank of the additional matrices to 4, which provides a good balance between performance and computational efficiency. In summary, we select the following methods:

1. Vanilla - a "baseline" approach that updates all parameters.
2. Adapter / LoRA - these approaches involve adding extra layers with original parameters fixed.

### 2.2.4 Summary

To sum up, we choose the following combinations to fine-tune SAM's performance for specific datasets or tasks, as illustrated in Figure 2: (1) Variations in the image encoder's architecture, including ViT-H, ViT-B, and ViT-T from MobileSAM (Zhang et al., 2023a). (2) Different approaches to the fine-tuning scope of the network, encompassing both the encoder and decoder updates, as well as the decoder only. (3) Diverse parameter updating settings, comprising Vanilla, PEFT with Adapter blocks, and PEFT with LoRA blocks. By integrating these three variables, we present a total of  $3 \times 2 \times 3 = 18$  distinct architectures for fine-tuning SAM on task-specific supervised training setup.

## 2.3 Task-general Prompt-based Learning

When multiple **labeled** datasets are available, an approach to adapting SAM to the medical imaging field involves additional pre-training, i.e. starting with SAM's pre-trained weights and continuing the learning process with a compilation of the medical datasets, followed by task-specific supervised training, as is similar in (Cheng et al., 2023; Ma et al., 2024). The procedure is illustrated in Figure 1 setup 2. The motivation for adopting a two-step training approach, rather than utilizing automated fine-tuning on multiple datasets, lies in the fact that various datasets, despite featuring medical images of the same anatomical area and being captured using the same imaging technique, may concentrate on entirely distinct objectives. For example, two brain MRI datasets might each focus on segmenting different entities, such as the hippocampus in one and tumors in the other. Applying automated segmentation to these datasets could introduce ambiguity and make the training

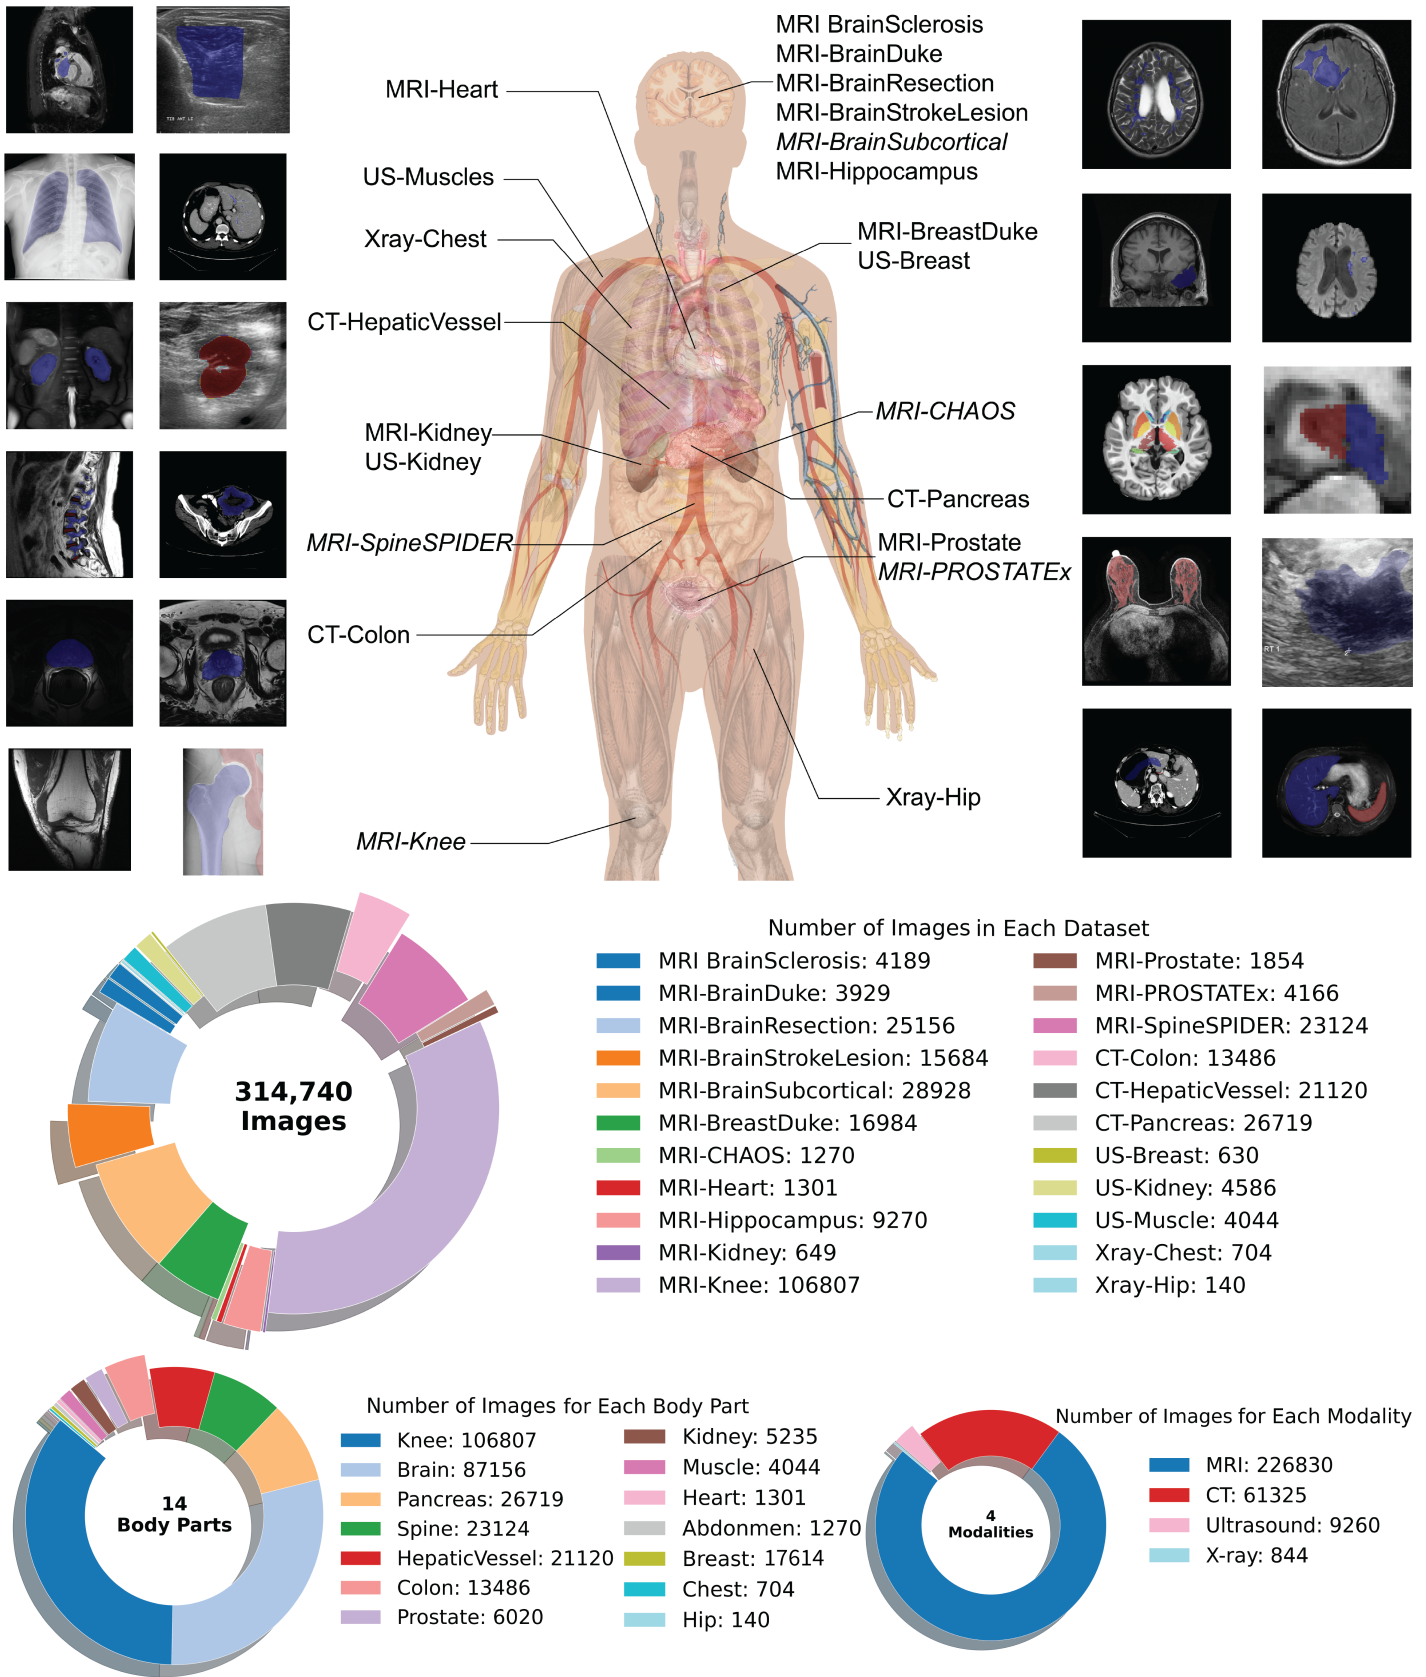


Figure 3: A visual summary of the collected datasets, including the visual example from each dataset and the number of images under different criteria.

Category	Methods	Dataset	Backbone
Decoder + Vanilla	SAM for medical images? (Huang et al., 2024)	COSMOS 1050K	ViT-H / ViT-B
	Generalist Vision Foundation (Shi et al., 2023a)	Multiple	ViT-H
	How to efficiently (Hu et al., 2023b)	ACDC (Bernard et al., 2018)	ViT-H
	All-in-SA (Cui et al., 2023)	Monuseg (Kumar et al., 2019)	ViT-H
	AdaptiveSAM (Paranjape et al., 2023)	Multiple	ViT-B
	SkinSAM (Hu et al., 2023a)	HAM10000 (Tschandl et al., 2018)	ViT-B
	DeSAM (Gao et al., 2023)	Multi-Prostate	ViT-H
	SurgicalSAM (Yue et al., 2023)	EndoVis (Allan et al., 2020)	ViT-H
	SAM3D (Bui et al., 2023)	Multiple	ViT-B
	SAM Fewshot (Xie et al., 2024)	Multiple	N/A
Decoder + LoRA	BLO-SAM (Zhang et al., 2024a)	Multiple	ViT-B
Encoder + Vanilla	AutoSAM (Shaharabany et al., 2023)	Multiple	ViT-H
	SAMUS (Lin et al., 2023)	Multiple	N/A
En/Decoder + Vanilla	Polyp-SAM (Li et al., 2023b)	Multi-Colonoscopy	ViT-L / ViT-B
	CellViT (Hörst et al., 2023)	Multi-nuclei	ViT-H
En/Decoder + LoRA	Customized SAM (Zhang and Liu, 2023)	Multi-Organ	ViT-H / ViT-B
	MA-SAM (Chen et al., 2023a)	BTVC (Fang and Yan, 2020)	ViT-H
	Cheap Lunch (Feng et al., 2023)	BraTS 18 (Menze et al., 2014)	ViT-B
En/Decoder + Adapter	Medical SAM Adapter (Wu et al., 2023a)	BTVC (Fang and Yan, 2020)	ViT-H
	SAM Fails to Segment? (Chen et al., 2023b)	kvasir-SEG (Jha et al., 2020)	ViT-H
	Learnable Ophthalmology SAM (Qiu et al., 2023)	Multiple	N/A
	Auto-Prompting SAM (Li et al., 2023a)	Multiple	LightWeight ViT
	SAM-Med2D (Cheng et al., 2023)	SAM-Med2D (Cheng et al., 2023)	ViT-B
	3DSAM-adapter (Gong et al., 2023)	Multi-CT	ViT-B
	SAM-Path (Zhang et al., 2023b)	Multiple	LightWeight ViT
	Self-Sampling Meta (Leng et al., 2024)	ABD-30 (Landman et al., 2015)	ViT-B
	Cross-modality (Shi et al., 2023b)	MRI	ViT-B

Table 1: Summary of existing fine-tuning methods, grouped by the component-method combinations. En/Decoder means both the encoder and decoder are optimized.

challenging.

Therefore, we use a prompt-based segmentation setup during task-expansive prompt-based learning, which involves providing prompts to identify the targeted area of interest interactively. This approach addresses the issue of the above-mentioned ambiguity. Furthermore, this choice conceptually aligns with the second phase of SAM’s second-stage training process (Kirillov et al., 2023), allowing it to be considered as an extra step of supervised training specifically tailored for medical images within the SAM framework.

To simulate a prompt, we first randomly selected a type of object if the ground truth mask contains multiple objects. Then, for each contiguous region of the selected object, we generate a box or a point prompt with equal probability. The generation of box prompts follows SAM’s strategy: a tight box containing the target box is created first, followed by some random noises added to both locations. The random noises are generated with the standard deviation equal to 10% of the box coordinates and to a maximum of 20 pixels. For point generation, we follow (Mazurowski et al., 2023) to first compute the distance of each pixel to its closest edge within the object region and sample a pixel with a distance above the median distance. This operation allows us to sample a point prompt around the center of the object. Such a design choice differs from that of the SAM, which

samples the first point prompt uniformly and subsequent points from the error region between the previous mask prediction and the ground-truth mask. Since our ultimate goal is to learn an automated segmentation algorithm, we don’t need an iterative process to refine predictions based on prompts.

Considering the differences across various medical imaging modalities, alongside the observation that images within the same modality often share similar visual features (Konz and Mazurowski, 2024), we emphasize the value of modality-specific dataset knowledge in pre-training, a practice also common in the literature (Tang et al., 2022). Specifically, we focus on supervised pre-training on MRI, which comprises most of the collected datasets.

#### 2.4 Task-agnostic Self-supervised Learning

When using multiple **unlabeled** medical image datasets, we can utilize the same idea as in the last section to conduct additional pre-training before applying task-specific supervised training. As the label information is unavailable in this scenario, we adapt the self-supervised learning (SSL) techniques, referred to as task-agnostic self-supervised learning. SSL refers to training a network without any human-collected labels. This is an effective strategy since the number of available datasets can increase dramatically. Popular



Alias in our paper	Description and citation	# of images	# of masks
<b>Datasets used for dataset-specific fine-tuning and pre-training:</b>			
MRI-BrainSclerosis (MRI-Scler.)	Brain Sclerosis with Lesion Segmentation (Muslim et al., 2022)	4189	1813
MRI-BrainDuke (MRI-Brain)	Brain with FLAIR Abnormality Segmentation Masks (Buda, 2019)	3929	1372
MRI-BrainResection (MRI-Resect.)	EPISURG Segmentation on the Resection CaViTy (Pérez-García et al., 2020)	25156	5377
MRI-BrainStrokeLesion (MRI-Stroke.)	ISLES 2022: Stroke Lesion Segmentation (Hernandez Petzsche et al., 2022)	15684	4827
MRI-BreastDuke (MRI-Breast)	Breast Cancer Screening (Saha et al., 2018; Lew et al., 2024)	16984	16984
MRI-Heart	Medical Segmentation Decathlon Cardiac (Antonelli et al., 2022)	1301	1301
MRI-Hippocampus (MRI-Hippo.)	Medical Segmentation Decathlon Hippocampus (Antonelli et al., 2022)	9270	6622
MRI-Kidney	T2-weighted Kidney MRI (Daniel et al., 2021)	649	647
MRI-Prostate (MRI-Prost.)	Original Multi-Parametric MRI Images of Prostate (Lemaitre et al., 2016)	1854	1854
CT-Colon	Medical Segmentation Decathlon Colon (Antonelli et al., 2022)	13486	1285
CT-HepaticVessel (CT-vessel)	Medical Segmentation Decathlon Hepatic Vessel (Antonelli et al., 2022)	21120	13042
CT-Pancreas (CT-Panc.)	Medical Segmentation Decathlon Pancreas (Antonelli et al., 2022)	26719	8792
US-Breast	Breast Ultrasound Images Dataset (Al-Dhabyani et al., 2019)	630	630
US-Kidney	Kidney Segmentation in Ultrasound Images (Song et al., 2022)	4586	4585
US-Muscle	Transverse Musculoskeletal Ultrasound images (Marzola et al., 2021)	4044	4043
Xray-Chest	Montgomery County and Shenzhen Chest X-ray Dataset (Jaeger et al., 2014)	704	704
Xray-Hip	X-ray Images of the Hip Joints (Gut, 2021)	140	140
<b>Datasets used for pre-training only:</b>			
MRI-BrainSubcortical	14 Brain Sub-cortical Structures (Sivaswamy et al., 2021)	28928	5891
MRI-CHAOS	Segmentation of Liver, Kidneys and Spleen from MRI data (Kavur et al., 2021)	1270	992
MRI-Knee	MRNet Knee MRI Examinations (Bien et al., 2018)	106807	Not used
MRI-ProstateX	ProstateX Dataset of Prostate MR Studies (Armato III et al., 2018)	4166	3195
MRI-SpineSPIDER	SPIDER - Lumbar Spine Segmentation in MRI (van der Graaf et al., 2023)	23124	12763

Table 2: Description of the collected datasets. The first 17 datasets are used throughout the paper for training and evaluation, and the remaining 5 are used exclusively for supervised or self-supervised pre-training.

SSL methods generally fall into two categories: contrastive learning (He et al., 2020; Chen et al., 2020; Park et al., 2020) or masked image modeling (He et al., 2022; Xie et al., 2022b). In the medical image field, further methods are tailored to fit the characteristics of medical images (Tang et al., 2022), including dimension discrepancy (Xie et al., 2022a), structure similarity (Jiang et al., 2023), etc.

Among these methods, we select the Masked Autoencoder (MAE) (He et al., 2022) as our self-supervised pre-training strategy for three reasons: (1) The computation of CL is infeasible given the input size of SAM, i.e.  $1024 \times 1024$ ; (2) SAM also utilized an MAE pre-trained vision transformer as the initial model; (3) There is increasing evidence showing the effectiveness of using MAE during medical image segmentation (Zhou et al., 2023; Zhuang et al., 2023; Wang et al., 2024). MAE first divides the input images into non-overlapping patches and removes most of them. Its goal is to reconstruct the original inputs given the remaining ones. For the same reason as in task-expansive prompt-based learning, we also conduct training with MRIs only and evaluate all datasets.

### 3. Datasets

We select 22 datasets from publicly available sources online, including TCIA (The Cancer Imaging Archive), Kaggle, Zendo, Grand Challenge, and Mendeley Data. Figure 3 presents a visual summary of the datasets used in this paper. In total, we collect 14 magnetic resonance images (MRI) datasets, 3 computed tomography (CT) datasets, 3 ultrasound (US) datasets, and 2 X-ray datasets. Our

datasets cover images from 14 main human anatomical locations. Examples of all 22 datasets are also included in the same figure, in addition to the image count of each dataset, each anatomical location, and each modality. The details of each dataset are listed in Table 2, including the citation, description, number of images, and number of masks. All datasets use expert-annotated segmentation masks, except for the MRI-Knee dataset, which is used exclusively for self-supervised pre-training. Therefore, the number of masks for this dataset is indicated as “not used”.

Since SAM is a 2D-based segmentation network, 3D volumes are first extracted to 2D slices. We convert NIfTI, MHA, and DICOM volumes by normalizing all pixel values to 0-255 and further upsampling each slice to  $1024 \times 1024$  using nearest-neighbor interpolation to meet SAM’s required input size. The pixel values are repeated three times across the three channels for all grayscale images. We split each dataset into training, validation, and test sets by dividing patients into 7 : 1 : 2 splits, accordingly.

In our study, we utilize these curated medical imaging datasets across various levels of data availability as follows:

1. *Specific Dataset Fine-tuning*: This involves datasets applied for fine-tuning on a single labeled dataset, including all datasets listed in the first section of Table 2.
2. *Prompt-based Supervised Training with Multiple Labeled Datasets*: This category includes all MRI datasets that come with available masks, utilized for supervised pre-training.
3. *Self-Supervised Learning with Both Labeled and Unlabeled*



*beled Datasets*: This level pertains to the use of all MRI datasets mentioned in our table for self-supervised learning, incorporating both labeled and unlabeled datasets.

## 4. Experiments

In this section, we conduct various experiments to assess the effectiveness of fine-tuning strategies across multiple datasets and configurations. First and foremost, we analyze the behaviors of these strategies when conducting task-specific supervised training. Next, we examine the effectiveness of applying task-expansive prompt-based learning or task-agnostic self-supervised learning before task-specific supervised training. These experiments are further extended to the few-shot learning and interactive segmentation setting. Throughout the paper, we use the Dice similarity score (DSC%) and normalized surface distance (NSD%) as the evaluation metric.

### 4.1 Which fine-tuning combination is the best if applying task-specific supervised training only?

#### 4.1.1 Implementation details

When conducting task-specific supervised training, in addition to finding the most effective fine-tuning strategy, we also want to compare the performance against conventional UNet-based algorithms. Therefore, we have chosen three well-known and frequently used automated segmentation models for comparison: UNet (Ronneberger et al., 2015), nnUNet (Isensee et al., 2021), and SwinUNet (Cao et al., 2022):

1. *UNet* serves as the baseline of the convolutional neural network-based segmentation algorithm.
2. *nnUNet* configures a UNet-based algorithm according to the specific attributes of the input datasets and approximates the performance of UNet with the optimal design choice.
3. *SwinUNet* replaces the encoder of UNet with a swin-transformer (Liu et al., 2021).

For PEFT with Adapter blocks, we implemented it based on (Wu et al., 2023b), which added adapter blocks to SAM's attention blocks. In the image encoder, adapter blocks were added into transformer blocks right after the multi-attention head or within the MLP residual block with a scalable parameter (Chen et al., 2022). In the default setting, we added Adapter blocks at the first two and last two transformer blocks (layers) in the image encoder. In the mask decoder, Adapters were added after the multi-head attention. For PEFT with LoRA blocks, we implemented it based on SAMed (Zhang and Liu, 2023), which inserted

LoRA layers into the "query" and "value" projection layers of each transformer block in the image encoder and mask decoder.

In the pre-processing step for our datasets, we removed all data lacking target objects, specifically images where the corresponding masks only consisted of the background, for both training and testing purposes. We also followed the same normalization mean and standard deviation as the original SAM. We utilized dice loss and cross-entropy loss with equal weights. For all experiments across various datasets, we adopted a basic learning rate of  $1e-4$  with a warmup of 200 iterations and applied AdamW with 0.1 weight decay for the remaining iterations. The batch size was set to 4. Histogram equalization and color jitter were used as augmentations during training. The maximum epochs were set to 200 epochs, incorporating an early stopping mechanism that halted training if there was no improvement in the highest validation DSC for 20 consecutive epochs. For settings that require heavy computing resources, such as ViT-H with the vanilla setting, we used Distributed Data Parallel (DDP) combined with model parallelism to split batches or the model across 2-3 GPUs. The specifics regarding the total number of parameters in the model, the number of parameters subject to updates for each combination, and the computing resources needed are detailed in Table 3.

#### 4.1.2 Results

Figure 4 shows the performance of various task-specific fine-tuning approaches across each dataset and the performance of various UNet-based algorithms. The detailed results of each architecture-component-method fine-tuning strategy over 17 datasets can be found in the Appendix. Note that we provide the numerical results for all figure-based results in the Appendix. By comparing the red and blue bars, we first observe an almost definite improvement (except in the ViT-B + Vanilla setting) in average performance when updating both the encoder and decoder parameters. This discovery indicates that SAM's original image encoder, initially trained on a vast collection of natural images, is not well-suited for efficiently extracting features from medical images, aligned with previous studies (Mazurowski et al., 2023). Therefore, freezing it and directly adapting the mask decoder for downstream tasks in the medical domain yields suboptimal performance.

When comparing the performances of Vanilla, Adapter, and LoRA, we observe little difference between the two PEFT methods (Adapter and LoRA). The rank between Vanilla and the PEFT methods depends on the choices of backbones. When using the standard-size backbone (ViT-H or ViT-B), we observe an absolute advantage of using PEFT techniques. For example, when comparing the middle three blue bars in Figure 4, we observe that




Figure 4: Left: Illustration of the dataset and training setting: task-specific supervised training; Middle: Visualization of the performance of various task-specific supervised training strategies, including various Image-Encoder depth: 4 (ViT-T) or 12 (ViT-B) or 32 (ViT-H); updating all parameters (V) or use PEFT-Adapter (A) or PEFT-LoRA (L); fine-tuning Mask Decoder only (De) or Image Encoder + Mask Decoder (En/De), and UNet and its variances. Right: Difference between the best fine-tuned SAM model (ViT-H + En/Decoder + LoRA) and the best UNet-based model (nnUNet) across individual datasets.

using PEFT with Adapter or LoRA achieves 77.35% and 78.93% DSC, respectively, while vanilla fine-tuning yields only 65.39% DSC. When comparing Adapter and LoRA, we observe that LoRA performs better when applied to a model with a larger size, *i.e.*, ViT-H. This can be caused by the superior parameter efficiency of LoRA (Hu et al., 2021). We also conducted an additional study to insert the adapter block after each layer and observed little difference. However, the difference between vanilla fine-tuning and PEFT is significantly reduced when using ViT-T, the small backbone introduced by MobileSAM (Zhang et al., 2023a). We hypothesize that the large gap between Vanilla and PEFT methods with a standard-size backbone is caused by the model collapse and knowledge-forgotten issues when fine-tuning all parameters of a large model (Dutt et al., 2023; Fu et al., 2023). To support our hypothesis, we track the change in validation DSC on all datasets. Examples of each modality are shown in Figure 5. It reveals that ViT-H + En/Decoder + Vanilla would cause a model collapse at an early stage. For instance, when applied to the MRI-BrainDuke dataset, this configuration results in a significant performance decrease, with a drop of approximately 35% in DSC.

Next, our results show that the choice of encoder network architecture has a minimal impact on performance when all other decisions remain the same. For instance, under the same settings of updating both encoder and decoder with Adapter (EN/Decoder + Adapter), we observe similar DSC scores across different encoders: ViT-H, ViT-B, and

ViT-T. This similarity in performance occurs despite ViT-H having a significantly larger encoder with 6.8 times network parameters compared with ViT-B and 64 times parameters compared with ViT-T, as listed in Table 3. We deduce that while ViT-H excels in foundational settings for multiple segmentation tasks, it may be too parameter-rich for singular medical imaging segmentation tasks, given the typically smaller dataset sizes in the medical imaging field. Such property can also make the training difficult. As shown in Figure 5 top right, it takes 140 epochs to train the ViT-H + En/Decoder + Adapter setting, which is approximately 30 times longer when conducting the same setting with ViT-B, yet the latter gives slightly better performance (80.53% DSC vs. 81.21% DSC when using ViT-H and ViT-B as the backbone respectively). Given the uniform results across our 17 datasets, we deduce that using PEFT (Adapter or LoRA) to fine-tune SAM within a smaller framework offers more benefits. This strategy strikes an optimal balance between efficiency and performance, especially for medical image segmentation tasks where datasets typically range from a couple of hundred to tens of thousands of images.

Lastly, when comparing the results of automated segmentation by directly training on UNets, we observe that fine-tuning SAM with the En/Decoder generally yields better average performance. To further study the optimal performance across individual datasets, we select the best SAM-based method (ViT-H+LoRA+En/Decoder) vs. the best UNet-based method (nnUNet) in Figure 4, right block. It appears that SAM excels in tasks where the objects of

Model			# of param.	# of trainable param.	GPU allocated (MB)	Single iteration time (s)
ViT-H	Decoder	Vanilla	640.95M	3.92M (0.61%)	22206	1.82
		Adapter	641.08M	0.13M (0.021%)	17210	1.88
		LoRA	640.97M	0.65M (0.10%)	17214	1.81
	En/Decoder	Vanilla	640.95M	640.95M (100.0%)	99,208	5.2
		Adapter	650.93M	9.98M (1.53%)	69,144	3.8
		LoRA	641.63M	1.31M (0.20%)	68,540	3.82
ViT-B	Decoder	Vanilla	93.60M	3.92M (4.19%)	11158	0.44
		Adapter	93.73M	0.13M (0.14%)	11532	0.45
		LoRA	93.62M	0.65M (0.69%)	11530	0.44
	En/Decoder	Vanilla	93.60M	93.60M (100.0%)	37306	1.16
		Adapter	97.28M	3.68M (3.79%)	33454	1
		LoRA	93.77M	0.80M (0.86%)	30578	0.95
ViT-T	Decoder	Vanilla	9.99M	3.92M (39.22%)	1404	0.08
		Adapter	10.12M	0.13M (1.13%)	1398	0.09
		LoRA	10.01M	0.65M (6.49%)	1400	0.08
	En/Decoder	Vanilla	9.99M	9.99M (100.0%)	10102	0.25
		Adapter	10.31M	0.32M (3.12%)	4018	0.16
		LoRA	10.04M	0.79M (7.82%)	4036	0.15

Table 3: Specifics of parameters and computer resources occupied for 18 combinations of fine-tuning choices. GPU allocations are measured when the batch size is set to 4, and the time for forward and backward passes is measured on a single A6000.

interest are complicated, such as MRI-Brain Sclerosis or CT-Colon. Conversely, nnUNet demonstrates higher precision in tasks with simpler targets, including Xray-Hip or MRI-Kidney.

#### 4.2 Is task-expansive prompt-based learning helpful for further automated segmentation?

As discussed in Section 2.3, we utilize an interactive segmentation setting for task-expansive prompt-based learning. In this section, we provide the implementation details and analysis of its performance.

##### 4.2.1 Implementation details

When implementing task-expansive prompt-based learning, we utilized a collection of datasets and their corresponding masks. Our experimental setup focused on training with all MRI datasets listed in Table 2 except “MRI-Knee”. In this collection criteria, we include 13 MRI datasets, including 45k training images and non-empty mask pairs, and 7.9k evaluation images and mask pairs. We selected ViT-B as the encoder architecture since (1) it achieved similar performance to ViT-H in the data-specific fine-tuning setting; (2) MobileSAM was a distilled version of SAM, which was not provided by the initial SAM and thus less representative.

Under this multi-dataset supervised training setting, we configured a batch size of 3, utilizing DDP across 3 A6000 GPUs for training. The training was stopped when the validation performance did not improve for 5 epochs. To train the network, we fixed the prompt encoder and updated the image encoder and mask decoder in a vanilla setting

(same as SAM’s supervised training stage). We generated a box or point prompt for each image as described in Section 2.3. Additionally, in this setting, we incorporated MedSAM (Ma et al., 2024) as a parallel approach since its training algorithm was similar to the proposed task-expansive prompt-based learning but with a different database.

In this experimental setup, our methodology follows a two-phase procedure, depicted in Figure 1. Setup 2: (1) Phase 1: task-expansive prompt-based learning. (2) Phase 2: task-specific supervised training for each dataset, beginning from pre-trained weights. We name the weights after phase 1 as **SupSAM**. The implementation details in the task-specific setting were identical to those of the previous section (Section 4.1). In phase 2, we had 2 model component choices (En/Decoder or Decoder only) and 3 Vanilla/PEFT choices in the automated setting. We evaluated the final automated segmentation performance after phase 2.

##### 4.2.2 Results

The results are shown in Figure 6 and Figure 7, in which we list the performance of the original method when starting from SAM’s initial weights in gray. We further illustrate the difference, highlighted in green (when positive) or red (if negative), when changing the initial weights from the original SAM to those acquired from task-expansive prompt-based learning or MedSAM, respectively. When task-specific supervised training is applied solely to the decoder (the upper 3 blocks in both figures), we observe a consistent performance decrease in both scenarios. Since

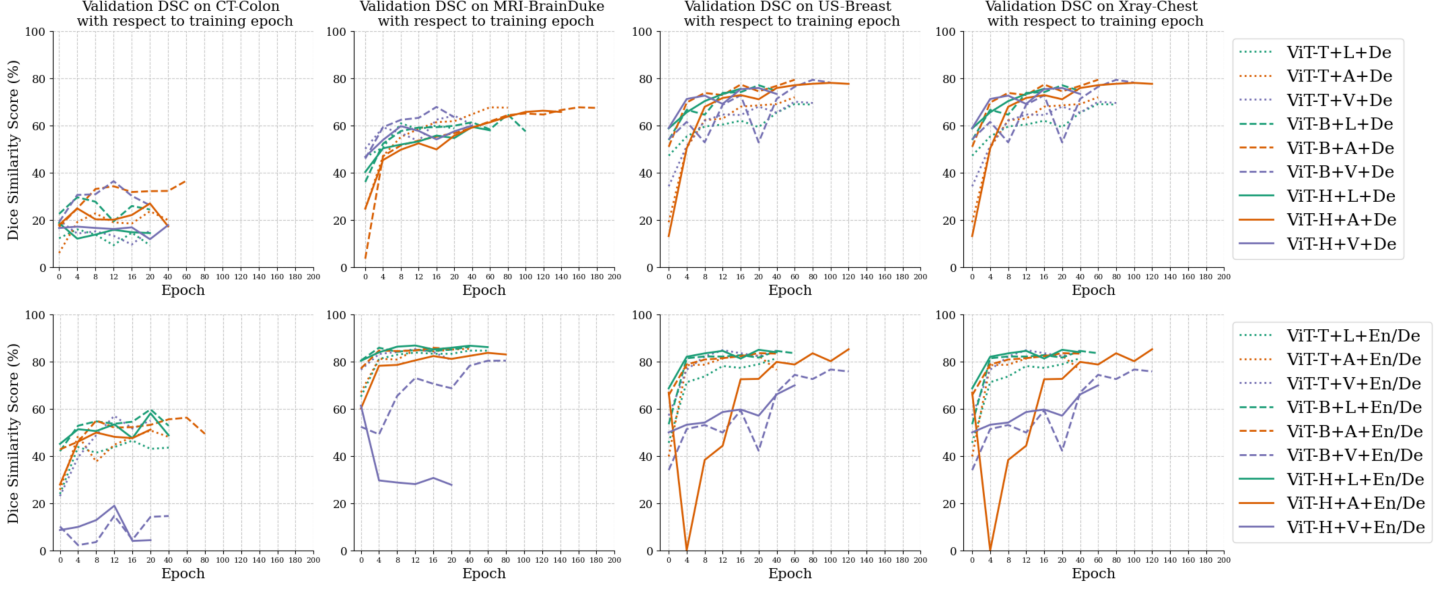


Figure 5: Validation Dice Coefficient Score (DSC) with respect to training epoch for different combinations of fine-tuning strategies. We select one representative dataset from each modality for visualization, and the behavior is consistent across different datasets. Note that the total number of epochs is not consistent because the training stops when the validation DSC is not improved for 20 epochs.

in this case, phase 2 fine-tuning is applied exclusively to the decoder, it appears that weights learning through supervised training fail to adequately prepare the image encoder with the medical imaging domain knowledge necessary for automated feature extraction. Thus, attempting to conserve computational resources while starting with a publicly available, prompt-based medical model, such as MedSAM, becomes unnecessary if the goal is to update only the small decoder for customized datasets.

When fine-tuning both the encoder and decoder during phase 2 task-specific supervised training (Figure 6 and 7, bottom 3 blocks), we observe either a negative or no effect when utilizing a different pre-trained weights when using PEFT methods and a positive effect when using Vanilla fine-tuning. However, the latter setting leads to worse overall performance and thus offers little practical value. Despite the decrement being alleviated when using MedSAM, the overall trend shows the ineffectiveness of prompt-based learning when the ultimate goal is automated segmentation, as MedSAM has a training set around 34 times larger yet results in no performance change at best. We infer that the observed performance enhancement in En/Decoder-Vanilla is attributed to supervised pre-training, potentially nudging the weights closer to the medical domain, thereby reducing the likelihood of model collapse. Nonetheless, since this configuration still yields the lowest segmentation performance compared to the PEFT setting, this improvement does not offer a practical benefit.

Our research indicates that supervised pre-training across

various labeled medical images, despite working well in the original setting (MedSAM (Ma et al., 2024) is verified in its paper, and ours is verified in Section 4.2), fails to enhance automated segmentation performance. We believe the ineffectiveness of additional supervised pre-training is due to the difference in the tasks: (1) The weights learned from interactive (prompt-based) mode vs. automated mode might be different and could not communicate with each other effectively. Additionally, extended pre-training may disrupt the transferability of the original weights. (2) Furthermore, training the network on multiple segmentation tasks simultaneously could lead to conflicts within the training process, thus making the training more challenging. These observations suggest that hybrid training, *i.e.*, fine-tuning SAM for both the automatic segmentation task and the prompt-based segmentation task simultaneously, may not be effective.

#### 4.3 Is task-agnostic self-supervised learning helpful for further automated segmentation?

As discussed in Section 2.4, we select masked autoencoder (MAE) as the method for task-agnostic self-supervised learning. In this section, we provide the implementation details and analyze its performance.

##### 4.3.1 Implementation details

In this stage, all MRI datasets in Table 2 were used during training. The mask information was ignored, and only the



## The impact of using task-expansive prompt-based training weights: SAM vs. SupSAM

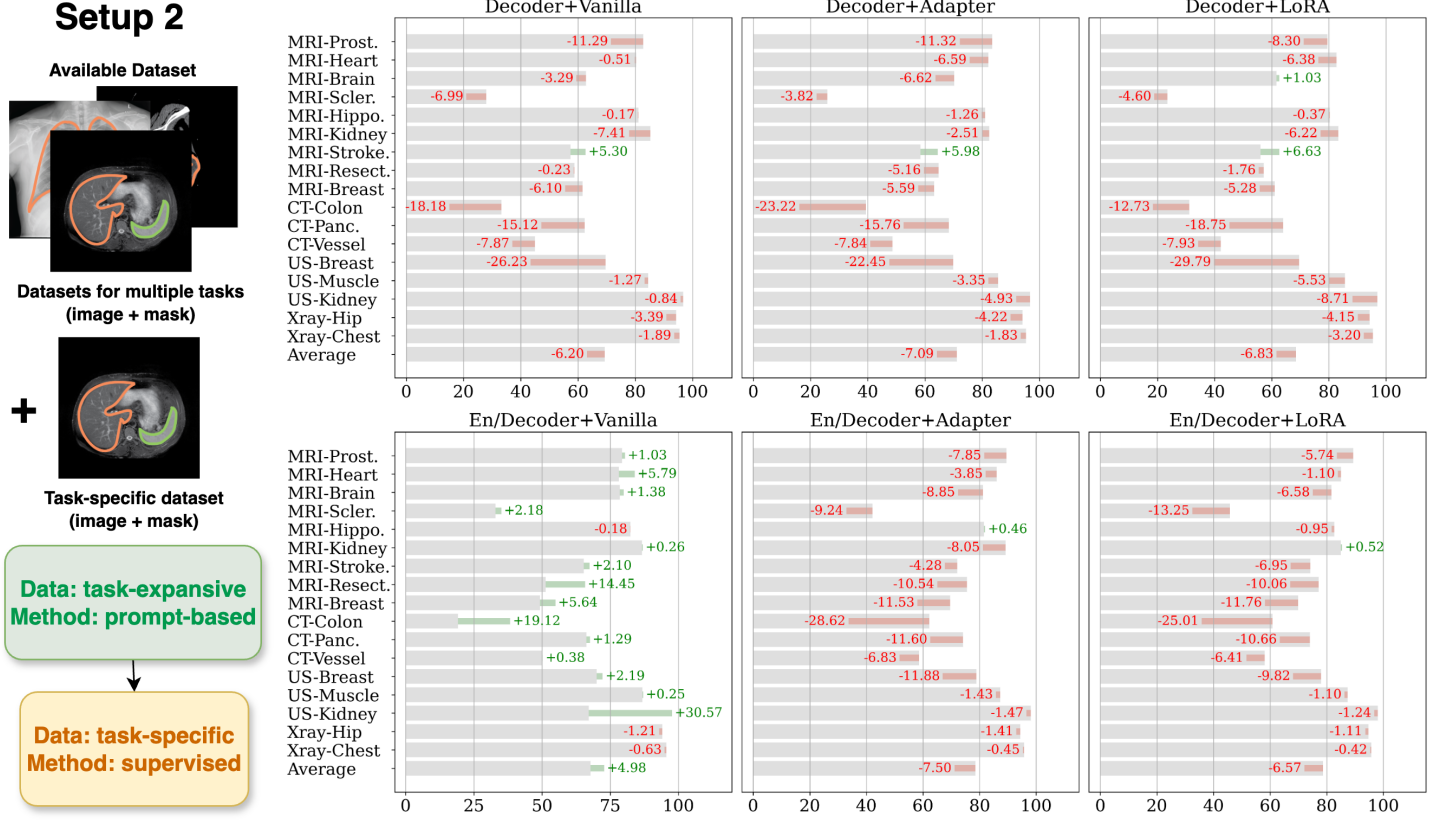


Figure 6: The comparison performance of applying various task-specific supervised training strategies to SAM directly (in gray), and the approach of using task-expansive prompt-based learning first when additional labeled datasets are available. The differences are represented as SupSAM - SAM, where green indicates a positive difference ( $>0$ ) and red indicates a negative difference ( $<0$ ), meaning that using SupSAM to initiate fine-tuning led to a performance increase or drop, respectively.

images from the training set were utilized. Blank slices were further removed since they were not informative for the reconstruction task. This resulted in 159,104 images for training.

Following the previous section, we only select ViT-B as the encoder architecture. To employ MAE’s objective, we introduced an additional decoder to reconstruct the inputs, which was discarded after training. The additional decoder’s structure, mask ratio, and learning rate schedule all follow that of the original paper. Since the shape of SAM’s input was restricted to  $1024 \times 1024$ , we set the batch size to 2. To accommodate for the small batch size, we set the cumulative iteration, i.e. the number of iterations to only compute gradients without updating the network parameters, to 16. The training was conducted on 8 A6000 for 50 epochs.

In this experimental setup, our methodology follows a similar two-phase procedure: (1) Phase 1: task-agnostic

self-supervised learning. (2) Phase 2: task-specific supervised training for each dataset, beginning from pre-trained weights. We name the weights after phase 1 as **SSLSAM**, which is shared publicly at <https://github.com/mazurowski-lab/finetune-SAM>.

### 4.3.2 Results

Figure 8 presents the results of using SSLSAM as the weights before task-specific supervised training. We observe an improvement in the average performance of all MRIs across all architecture combinations besides En/Decoder + LoRA. Noticeably, when selecting En/Decoder + Adapter as the fine-tuning strategy, we observe consistent improvements over all MRI datasets, and the model achieves an average performance of 78.09%, which surpasses the leading ViT-H based method by 1.08% and the best UNet-based method (nnUNet) by 3.58%. Contradictory to task-expansive prompt-based learning, the features learned during

## The impact of using task-expansive prompt-based training weights: SAM vs. MedSAM




Figure 7: The comparison performance of applying various task-specific supervised training strategies to SAM directly (in gray), and the approach of using MedSAM (Ma et al., 2024), a variation of task-expansive prompt-based learning first when additional labeled datasets are available. The differences are represented as MedSAM - SAM, where green indicates a positive difference ( $>0$ ) and red indicates a negative difference ( $<0$ ), meaning that using MedSAM to initiate fine-tuning led to a performance increase or drop, respectively.

SSL are task agnostic, and since we only pre-train the encoder, there is no dependence on the task condition, which we have shown to hurt the performance. In the non-MRI data sets base, we observe a minor decrease in performance. This may be due to the MRI's specificity.

### 4.4 Which fine-tuning setup works the best in the few-shot learning setting?

#### 4.4.1 Implementation details

One of the popular scenarios for adopting fine-tuning strategies is when the target dataset has a limited number of annotated images, known as few-shot fine-tuning. Few-shot fine-tuning gets even more attention in the field of medical imaging segmentation, given the difficulty in securing a substantial quantity of expert annotations. We explore the effectiveness of different fine-tuning procedures in the following setting.

Our study primarily investigates 5-shot learning, where

we randomly select five images and their corresponding annotations from the training set and validation set, respectively. In the case of datasets derived from 3D volumes, these samples were chosen from distinct volumes to ensure diversity. In addition to the constraint of using only 5 training and validation images, we also modified the training hyperparameters to suit the few-shot setting better. Specifically, we set a maximum of 1000 epochs for training and a batch size of 2, and did not include any early stopping criteria. We increased the number of epochs to ensure sufficient updates for learning medical image patterns, used a smaller batch size to have diverse updates, and removed early stopping to prevent premature termination in the few-shot setting. We chose this strategy based on preliminary results indicating better performance compared to the default setting for full-set fine-tuning.

In this setting, we also investigated the previous three setups within a few-shot framework, including (1) task-specific fine-tuning exclusively and (2) two pre-trained setups: task-

## The impact of using task-agnostic self-supervised training weights: SAM vs. SSSLAM




Figure 8: The comparison performance of applying various task-specific supervised training strategies to SAM directly (in gray) and the approach of using task-agnostic self-supervised learning first when additional labeled datasets are available. The differences are represented as SSSLAM - SAM, where green indicates a positive difference ( $>0$ ) and red indicates a negative difference ( $<0$ ), meaning that using SSSLAM to initiate fine-tuning led to a performance increase or drop, respectively.

expansive prompt-based learning and task-agnostic self-supervised learning, each followed by additional task-specific fine-tuning. Notably, we cannot use the weights learned from task-expansive prompt-based learning as its training pipeline requires all labeled images. Thus, we consider the MedSAM weights as an alternative. In this section, we investigate three key questions: (1) whether additional pre-training offers superior initial weights for the few-shot setting, (2) which fine-tuning strategy proves most effective in a few-shot context, (3) whether fine-tuning foundation models still works better than UNets in the few-shot setting, and (4) how does performance vary with different levels of few-shot data availability, such as 10-shot or 20-shot?

### 4.4.2 Results

Figure 9 shows the results for the 5-shot setting, comparing the fine-tuning of SAM’s models with direct training on UNets. The results demonstrate that using ViT-B +

En/Decoder + Adapter on the initial SAM yields optimal performance across all configurations. Comparing the results with UNets’, especially with the best-performing nnUNet, fine-tuned SAM still manages to achieve a slight improvement. This observation aligns with the findings presented in Section 4.1. The performance of three selected datasets with different levels of few-shot data availability is shown in Appendix Figure 12.

Consistently, the observation that MedSAM is ineffective persists in this few-shot scenario. While SSSLAM has shown effectiveness in the full-dataset setting, its efficacy diminished and even became detrimental in the few-shot scenario. This observation aligns with the understanding that the task-agnostic training disrupts the model’s connection with the segmentation tasks. While beneficial as initial weights for full-dataset training, this approach may not adequately rebuild these connections with only five images, resulting in lower performance. In summary, we observe the ineffectiveness of additional pre-training (either using




Figure 9: Left: Visualization of the performance of fine-tuned models compared with UNets in 5-shot setting (Section 4.4). Right: Difference between the best fine-tuned SAM model (ViT-B-A-En/Decoder) and the best UNet-based model (nnUnet) across individual datasets.

task-expansive prompt-based learning or task-agnostic self-supervised learning) with limited training data. This can be contradictory to the common belief that imposing medical knowledge on the network in advance will be beneficial in the few-shot learning scenario.

#### 4.5 What if I can provide prompts to SAM?

Despite the primary goal of this paper being to find the best practice to build an automated segmentation algorithm from the foundation models, we conduct a non-trivial extension that studies the effectiveness of various task-specific supervised training strategies in the interactive segmentation setting.

##### 4.5.1 Implementation details

We considered two types of prompts: point or box. To generate a prompt, we utilized the same strategy as in task-expansive prompt-based learning. The generated prompts were fed to the prompt encoder to obtain the prompt embeddings during both training and inference. The prompt encoder was fixed. We also limited the evaluation data to MRI only since this task was not our main focus.

We included the performance of applying prompts to the original SAM and its oracle performance; the latter reflects the hypothetical performance when an extra step is taken to determine the best prediction from the network. Additionally, we include weights acquired from task-expansive prompt-based learning and MedSAM in the box-prompt setting. When applying task-specific supervised training to SAM, we followed Section 4.1 to include all strategies.

##### 4.5.2 Results

Figure 10 shows the performance of all task-specific supervised training strategies in the interactive mode, as well as that of the original SAM as a reference. First, we observe that the performance of SAM with point prompts is not satisfactory. Replacing the points with boxes significantly improves performance. In this setting, we also examine the weights acquired from task-expansive prompt-based learning and observe a similar performance, with a difference of only 0.92% in DSC compared to the initial SAM. This observation confirms that the inefficiency of task-expansive prompt-based learning is inherently due to having different tasks in mind.

To be noted, the best automated task-specific supervised training strategy yields an average performance of 77.01% on all MRI datasets, which is better than the original SAM with point prompts and on par with that with box prompts.

Conducting task-specific supervised training is also shown to be effective in the interactive segmentation scenario, especially when points are used as the prompts. Moreover, we find that the best task-specific supervised training strategy is ViT-T + En/Decoder + Adapter in both scenarios, achieving 83.83%/88.17% DSC when using a point/box as a prompt. This difference also suggests that providing box prompts can lead to better segmentation performance with the additional annotation cost.

## 5. Conclusion

In summary, we systematically study different strategies to fine-tune the Segment Anything Model (SAM) in the medical imaging field. Our extensive experiments reveal several interesting observations and provide a guide for developing






Figure 10: Performance of the task-specific fine-tuning strategies and initial SAM when using point/box prompts.

Dataset Setting	Condition	Optimal Strategy
Single Dataset	NA	ViT-H + En/Decoder + LoRA
Multiple Datasets	NA	SSLSAM + En/Decoder + Adapter
Single Dataset	Few-shot	ViT-B + En/Decoder + Adapter
Single Dataset	Prompt-based	ViT-T + En/Decoder + Adapter

Table 4: The optimal fine-tuning strategy under each scenario. SSLSAM refers to the weights obtained after self-supervised learning, introduced in Figure 1, Step 3. LoRA and Adapter are two parameter-efficient fine-tuning strategies, and En/Decoder means the adaption layers in both the encoder and decoder will be updated.

an automated segmentation algorithm from SAM. Our findings are based on dataset availability and are summarized as follows:

When only a single dataset is available for fine-tuning, we find that:

1. Fine-tuning SAM, if with proper strategies, leads to better segmentation performance than training from scratch with nnUNet (Isensee et al., 2021).
2. Increasing the size of network architecture leads to a marginal increase in segmentation performance, while exponentially longer training time.
3. Using parameter-efficient fine-tuning (PEFT) techniques has a clear advantage over fine-tuning all parameters on medium-to-large-sized architectures (ViT-H or ViT-B), but not on small architectures (ViT-T).

4. Optimizing both the encoder and decoder (En/De) has clear advantages over only the decoder, with the latter being a popular choice in the literature.
5. In summary, we recommend using **ViT-B + En/Decoder + PEFT** as the fine-tuning framework because of its leading performance and moderate computation cost.

When having multiple datasets, we can incorporate general medical knowledge into SAM through additional pre-training, and find that:

1. Using additional labeled data for other tasks with prompt-based learning is an inefficient strategy to further improve SAM's performance.
2. Using additional unlabeled data with self-supervised learning can improve SAM's performance when the test data

is from the same modality as the pre-trained data, and has no effect otherwise.

When we adjust the fine-tuning task into few-shot learning, we find that:

1. Fine-tuning SAM significantly outperforms UNet, and still maintains a slight improvement compared with nn-UNet in the 5-shot setting.
2. Additional task-expansive prompt-based learning or task-agnostic self-supervised learning do not contribute to improvements in the final few-shot learning task.

Lastly, in the interactive segmentation scenario, we observe that:

1. Fine-tuned SAM brings substantial performance improvements compared to SAM and MedSAM in the interactive segmentation setting.
2. Our recommendation framework is **ViT-T + En/Decoder + PEFT** given its superior performance and low computation cost.

In addition to these findings, we also observe several limitations when fine-tuning SAM. Namely, the computational cost of fine-tuning SAM can be high when using the default backbones. As shown in Table 3, when using ViT-H or ViT-B as the backbone, the number of parameters is 640.95M and 93.60 M. Using parameter-efficient fine-tuning strategies leads to better segmentation performance but does not improve the training efficiency. A potential solution is to utilize smaller backbones, such as ViT-T, *i.e.*, a lightweight encoder distilled from pre-trained ViT-H (Zhang and Jiao, 2023), to balance the performance and computation cost. Another limitation is that our study only focuses on the 2D segmentation tasks. To alleviate this issue, we have conducted a preliminary study that utilizes nnUNet 3D as a strong 3D segmentation baseline. We observed that the performance of DSC(%) is 87.14 vs. 87.35, 40.56 vs. 29.96, 47.58 vs. 46.95 on MRI-Heart, MRI-Scler., and CT-Colon, respectively, when comparing between nnUNet 3D vs. 2D. This observation suggests that the improvements are limited when the baseline methods are switched from 2D to 3D. On the other hand, adapting SAM to 3D segmentation is beyond the scope of this paper. We believe some potential directions include adapting SAM directly to 3D (Gong et al., 2024) or fine-tuning the recently proposed SAM 2 (Ravi et al., 2024), which is designed specifically for video segmentation.

## Acknowledgments

Research reported in this publication was supported by the National Institute Of Biomedical Imaging And Bioengineering of the National Institutes of Health under Award Number R01EB031575. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

## Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects.

## Conflicts of Interest

The conflicts of interest have not been entered yet.

## Data availability

All data used in this study are publicly available. We provide detailed descriptions and proper citations for each dataset used to ensure full transparency and reproducibility. Additionally, all code for data preprocessing and preparation is publicly released and can be used to fully reproduce our results. The code is available at: <https://github.com/mazurowski-lab/finetune-SAM>.

## References

- Mugahed A Al-Antari, Mohammed A Al-Masni, Mun-Taek Choi, Seung-Moo Han, and Tae-Seong Kim. A fully integrated computer-aided diagnosis system for digital x-ray mammograms via deep learning detection, segmentation, and classification. *International journal of medical informatics*, 117:44–54, 2018.
- Walid Al-Dhabyani, Mohammed Gomaa, Hussien Khaled, and Fahmy Aly. Deep learning approaches for data augmentation and classification of breast masses using ultrasound images. *Int. J. Adv. Comput. Sci. Appl*, 10(5): 1–11, 2019.
- Max Allan, Satoshi Kondo, Sebastian Bodenstedt, Stefan Leger, Rahim Kadkhodamohammadi, Imanol Luengo, Felix Fuentes, Evangello Flouty, Ahmed Mohammed, Marius Pedersen, et al. 2018 robotic scene segmentation challenge. *arXiv preprint arXiv:2001.11190*, 2020.
- Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger,

- Ronald M Summers, et al. The medical segmentation decathlon. *Nature communications*, 13(1):4128, 2022.
- Samuel G Armato III, Henkjan Huisman, Karen Drukker, Lubomir Hadjiiski, Justin S Kirby, Nicholas Petrick, George Redmond, Maryellen L Giger, Kenny Cha, Artem Mamonov, et al. Prostatex challenges for computerized classification of prostate lesions from multiparametric magnetic resonance images. *Journal of Medical Imaging*, 5(4):044501–044501, 2018.
- Shekoofeh Azizi, Laura Culp, Jan Freyberg, Basil Mustafa, Sebastien Baur, Simon Kornblith, Ting Chen, Patricia MacWilliams, S Sara Mahdavi, Ellery Wulczyn, et al. Robust and efficient medical imaging with self-supervision. *arXiv preprint arXiv:2205.09723*, 2022.
- Samyadeep Basu, Daniela Massiceti, Shell Xu Hu, and Soheil Feizi. Strong baselines for parameter efficient few-shot fine-tuning. *arXiv preprint arXiv:2304.01917*, 2023.
- Olivier Bernard, Alain Lalande, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging*, 37(11):2514–2525, 2018.
- Nicholas Bien, Pranav Rajpurkar, Robyn L Ball, Jeremy Irvin, Allison Park, Erik Jones, Michael Bereket, Bhavik N Patel, Kristen W Yeom, Katie Shpanskaya, et al. Deep-learning-assisted diagnosis for knee magnetic resonance imaging: development and retrospective validation of mrnet. *PLoS medicine*, 15(11):e1002699, 2018.
- Mateusz Buda. Brain mri segmentation: Brain mri images together with manual flair abnormality segmentation masks. *Kaggle*, 2019.
- Nhat-Tan Bui, Dinh-Hieu Hoang, Minh-Triet Tran, and Ngan Le. Sam3d: Segment anything model in volumetric medical images. *arXiv preprint arXiv:2309.03493*, 2023.
- Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. *Advances in Neural Information Processing Systems*, 33:11285–11297, 2020.
- 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.
- Cheng Chen, Juzheng Miao, Dufan Wu, Zhiling Yan, Sekeun Kim, Jiang Hu, Aoxiao Zhong, Zhengliang Liu, Lichao Sun, Xiang Li, et al. Ma-sam: Modality-agnostic sam adaptation for 3d medical image segmentation. *arXiv preprint arXiv:2309.08842*, 2023a.
- Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems*, 35:16664–16678, 2022.
- 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*, 2023b.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/chen20j.html>.
- Junlong Cheng, Jin Ye, Zhongying Deng, Jianpin Chen, Tianbin Li, Haoyu Wang, Yanzhou Su, Ziyang Huang, Jilong Chen, Lei Jiang, et al. Sam-med2d. *arXiv preprint arXiv:2308.16184*, 2023.
- 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.
- Alexander J Daniel, Charlotte E Buchanan, Thomas Allcock, Daniel Scerri, Eleanor F Cox, Benjamin L Prestwich, and Susan T Francis. Automated renal segmentation in healthy and chronic kidney disease subjects using a convolutional neural network. *Magnetic resonance in medicine*, 86(2):1125–1136, 2021.
- Ruining Deng, Can Cui, Quan Liu, Tianyuan Yao, Lucas W. Remedios, Shunxing Bao, Bennett A. Landman, Lee Wheless, Lori A. Coburn, Keith T. Wilson, Yaohong Wang, Shilin Zhao, Agnes B. Fogo, Haichun Yang, Yucheng Tang, and Yuankai Huo. Segment anything model (sam) for digital pathology: Assess zero-shot segmentation on whole slide imaging. *ArXiv*, abs/2304.04155, 2023. URL <https://api.semanticscholar.org/CorpusID:258049163>.

- 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.
- Raman Dutt, Linus Ericsson, Pedro Sanchez, Sotirios A Tsaftaris, and Timothy Hospedales. Parameter-efficient fine-tuning for medical image analysis: The missed opportunity. *arXiv preprint arXiv:2305.08252*, 2023.
- Xi Fang and Pingkun Yan. Multi-organ segmentation over partially labeled datasets with multi-scale feature abstraction. *IEEE Transactions on Medical Imaging*, 39(11): 3619–3629, 2020.
- 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.
- Zihao Fu, Haoran Yang, Anthony Man-Cho So, Wai Lam, Lidong Bing, and Nigel Collier. On the effectiveness of parameter-efficient fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 12799–12807, 2023.
- 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.
- 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.
- 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 tumor segmentation. *Medical Image Analysis*, 98:103324, 2024.
- Hanxue Gu, Roy Colglazier, Haoyu Dong, Jikai Zhang, Yaqian Chen, Zafer Yildiz, Yuwen Chen, Lin Li, Jichen Yang, Jay Willhite, Alex M. Meyer, Brian Guo, Yashvi Atul Shah, Emily Luo, Shipra Rajput, Sally Kuehn, Clark Bulleit, Kevin A. Wu, Jisoo Lee, Brandon Ramirez, Darui Lu, Jay M. Levin, and Maciej A. Mazurowski. Segmentanybone: A universal model that segments any bone at any location on mri, 2024.
- Daniel Gut. X-ray images of the hip joints. 1, July 2021. . URL <https://data.mendeley.com/datasets/zm6bxzhmfz>. Publisher: Mendeley Data.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735, 2020.
- 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.
- Moritz R Hernandez Petzsche, Ezequiel de la Rosa, Uta Han-nig, Roland Wiest, Waldo Valenzuela, Mauricio Reyes, Maria Meyer, Sook-Lei Liew, Florian Kofler, Ivan Ezhov, et al. Isles 2022: A multi-center magnetic resonance imaging stroke lesion segmentation dataset. *Scientific data*, 9(1):762, 2022.
- Fabian Hörst, Moritz Rempe, Lukas Heine, Constantin Seibold, Julius Keyl, Giulia Baldini, Selma Ugurel, Jens Siveke, Barbara Grünwald, Jan Egger, et al. Cellvit: Vision transformers for precise cell segmentation and classification. *arXiv preprint arXiv:2306.15350*, 2023.
- 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.
- 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.
- Mingzhe Hu, Yuheng Li, and Xiaofeng Yang. Skinsam: Empowering skin cancer segmentation with segment anything model. *arXiv preprint arXiv:2304.13973*, 2023a.
- Xinrong Hu, Xiaowei Xu, and Yiyu Shi. How to efficiently adapt large segmentation model (sam) to medical images. *arXiv preprint arXiv:2306.13731*, 2023b.
- Yuhao Huang, Xin Yang, Lian Liu, Han Zhou, Ao Chang, Xinrui Zhou, Rusi Chen, Junxuan Yu, Jiongquan Chen, Chaoyu Chen, et al. Segment anything model for medical images? *Medical Image Analysis*, 92:103061, 2024.
- 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.
- Stefan Jaeger, Sema Candemir, Sameer Antani, Yi-Xiáng J Wáng, Pu-Xuan Lu, and George Thoma. Two public chest



- x-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative imaging in medicine and surgery*, 4(6):475, 2014.
- Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D Johansen. Kvasir-seg: A segmented polyp dataset. In *MultiMedia Modeling: 26th International Conference, MMM 2020, Daejeon, South Korea, January 5–8, 2020, Proceedings, Part II 26*, pages 451–462. Springer, 2020.
- Yankai Jiang, Mingze Sun, Heng Guo, Xiaoyu Bai, Ke Yan, Le Lu, and Minfeng Xu. Anatomical invariance modeling and semantic alignment for self-supervised learning in 3d medical image analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15859–15869, 2023.
- 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.
- 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.
- Nicholas Konz and Maciej A Mazurowski. The effect of intrinsic dataset properties on generalization: Unraveling learning differences between natural and medical images. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024. URL <https://openreview.net/forum?id=ixP76Y33y1>.
- Neeraj Kumar, Ruchika Verma, Deepak Anand, Yanning Zhou, Omer Fahri Onder, Efstratios Tsougenis, Hao Chen, Pheng-Ann Heng, Jiahui Li, Zhiqiang Hu, et al. A multi-organ nucleus segmentation challenge. *IEEE transactions on medical imaging*, 39(5):1380–1391, 2019.
- 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*, volume 5, page 12, 2015.
- Ho Hin Lee, Yu Gu, Theodore Zhao, Yanbo Xu, Jianwei Yang, Naoto Usuyama, Cliff Wong, Mu Wei, Bennett A Landman, Yuankai Huo, et al. Foundation models for biomedical image segmentation: A survey. *arXiv preprint arXiv:2401.07654*, 2024.
- Guillaume Lemaitre, Robert Martí Marly, and Fabrice Meriaudeau. Original multi-parametric mri images of prostate. 2016.
- Tianang Leng, Yiming Zhang, Kun Han, and Xiaohui Xie. Self-sampling meta sam: Enhancing few-shot medical image segmentation with meta-learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7925–7935, 2024.
- Christopher O Lew, Majid Harouni, Ella R Kirksey, Elianne J Kang, Haoyu Dong, Han Gu, Lars J Grimm, Ruth Walsh, Dorothy A Lowell, and Maciej A. Mazurowski. A publicly available deep learning model and dataset for segmentation of breast, fibroglandular tissue, and vessels in breast mri. *Scientific reports*, 14 1:5383, 2024. URL <https://api.semanticscholar.org/CorpusID:268251677>.
- 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*, 2023a.
- Wei-Hong Li, Xialei Liu, and Hakan Bilen. Cross-domain few-shot learning with task-specific adapters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7161–7170, 2022.
- Yuheng Li, Mingzhe Hu, and Xiaofeng Yang. Polyp-sam: Transfer sam for polyp segmentation. *arXiv preprint arXiv:2305.00293*, 2023b.
- Xian Lin, Yangyang Xiang, Li Zhang, Xin Yang, Zengqiang Yan, and Li Yu. Samus: Adapting segment anything model for clinically-friendly and generalizable ultrasound image segmentation. *arXiv preprint arXiv:2309.06824*, 2023.
- Caixia Liu and Mingyong Pang. Extracting lungs from ct images via deep convolutional neural network based segmentation and two-pass contour refinement. *Journal of Digital Imaging*, 33:1465–1478, 2020.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15:1–9, 2024.
- Francesco Marzola, Nens van Alfen, Jonne Doorduyn, and Kristen M Meiburger. Deep learning segmentation of

- p>transverse musculoskeletal ultrasound images for neuromuscular disease assessment.
- Computers in Biology and Medicine*
- , 135:104623, 2021.
- 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.
- Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- 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.
- Ali M Muslim, Syamsiah Mashohor, Gheyath Al Gawwam, Rozi Mahmud, Marsyita binti Hanafi, Osama Alnuaimi, Raad Josephine, and Abdullah Dhaifallah Almutairi. Brain mri dataset of multiple sclerosis with consensus manual lesion segmentation and patient meta information. *Data in Brief*, 42:108139, 2022.
- Jay N Paranjape, Nithin Gopalakrishnan Nair, Shameema Sikder, S Swaroop Vedula, and Vishal M Patel. Adaptivesam: Towards efficient tuning of sam for surgical scene segmentation. *arXiv preprint arXiv:2308.03726*, 2023.
- Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, pages 319–345. Springer, 2020.
- F Pérez-García, R Rodionov, A Alim-Marvasti, R Sparks, J Duncan, and S Ourselin. Episurg: a dataset of postoperative magnetic resonance images (mri) for quantitative analysis of resection neurosurgery for refractory epilepsy. university college london. *DOI*, 1(0.5522):04, 2020.
- Zhongxi Qiu, Yan Hu, Heng Li, and Jiang Liu. Learnable ophthalmology sam. *arXiv preprint arXiv:2304.13425*, 2023.
- Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *Advances in neural information processing systems*, 30, 2017.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Efficient parametrization of multi-domain deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8119–8127, 2018.
- 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.
- T-YLPG Ross and GKHP Dollár. Focal loss for dense object detection. In *proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2980–2988, 2017.
- Ashirbani Saha, Michael R Harowicz, Lars J Grimm, Connie E Kim, Sujata V Ghate, Ruth Walsh, and Maciej A Mazurowski. A machine learning approach to radio-genomics of breast cancer: a study of 922 subjects and 529 dce-mri features. *British journal of cancer*, 119(4): 508–516, 2018.
- Tal Shaharabany, Aviad Dahan, Raja Giryes, and Lior Wolf. Autosam: Adapting sam to medical images by overloading the prompt encoder, 2023.
- 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. *Diagnosics*, 13(11):1947, 2023a.
- Xiaoyu Shi, Shurong Chai, Yinhao Li, Jingliang Cheng, Jie Bai, Guohua Zhao, and Yen-Wei Chen. Cross-modality attention adapter: A glioma segmentation fine-tuning method for sam using multimodal brain mr images. *arXiv preprint arXiv:2307.01124*, 2023b.
- Nahian Siddique, Sidike Paheding, Colin P Elkin, and Vijay Devabhaktuni. U-net and its variants for medical image segmentation: A review of theory and applications. *IEEE Access*, 9:82031–82057, 2021.
- Jayanthi Sivaswamy, Alphin J Thottupattu, Raghav Mehta, R Sheelakumari, Chandrasekharan Kesavadas, et al. Subcortical structure segmentation database for young population. *arXiv preprint arXiv:2111.01561*, 2021.

- Yuxin Song, Jing Zheng, Long Lei, Zhipeng Ni, Baoliang Zhao, and Ying Hu. Ct2us: Cross-modal transfer learning for kidney segmentation in ultrasound images with synthesized data. *Ultrasonics*, 122:106706, 2022.
- Yucheng Tang, Dong Yang, Wenqi Li, Holger R Roth, Bennett Landman, Daguang Xu, Vishwesh Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20730–20740, 2022.
- Hugo Touvron, Matthieu Cord, Alaaeldin El-Nouby, Jakob Verbeek, and Hervé Jégou. Three things everyone should know about vision transformers. In *European Conference on Computer Vision*, pages 497–515. Springer, 2022.
- Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018.
- Jasper W van der Graaf, Miranda L van Hooff, Constantinus FM Buckens, Matthieu Rutten, Job LC van Susante, Robert Jan Kroeze, Marinus de Kleuver, Bram van Ginneken, and Nikolas Lessmann. Lumbar spine segmentation in mr images: a dataset and a public benchmark. *arXiv preprint arXiv:2306.12217*, 2023.
- 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.
- Risheng Wang, Tao Lei, Ruixia Cui, Bingtao Zhang, Hongying Meng, and Asoke K Nandi. Medical image segmentation using deep learning: A survey. *IET Image Processing*, 16(5):1243–1267, 2022.
- Wenxuan Wang, Jing Wang, Chen Chen, Jianbo Jiao, Yuanxiu Cai, Shanshan Song, and Jiangyun Li. Fremim: Fourier transform meets masked image modeling for medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7860–7870, 2024.
- Jingxi Weng, Benjamin Wildman-Tobriner, Mateusz Buda, Jichen Yang, Lisa M Ho, Brian C Allen, Wendy L Ehieli, Chad M Miller, Jikai Zhang, and Maciej A Mazurowski. Deep learning for classification of thyroid nodules on ultrasound: validation on an independent dataset. *Clinical Imaging*, 99:60–66, 2023.
- 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*, 2023a.
- Qi Wu, Yuyao Zhang, and Marawan Elbatel. Self-prompting large vision models for few-shot medical image segmentation. In *MICCAI Workshop on Domain Adaptation and Representation Transfer*, pages 156–167. Springer, 2023b.
- Weiyi Xie, Nathalie Willems, Shubham Patil, Yang Li, and Mayank Kumar. Sam fewshot finetuning for anatomical segmentation in medical images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3253–3261, 2024.
- Yutong Xie, Jianpeng Zhang, Yong Xia, and Qi Wu. Unimiss: Universal medical self-supervised learning via breaking dimensionality barrier. In *European Conference on Computer Vision*, pages 558–575. Springer, 2022a.
- Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022b.
- Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui Tao, and Fu Lee Wang. Parameter-efficient fine-tuning methods for pretrained language models: A critical review and assessment. *arXiv preprint arXiv:2312.12148*, 2023.
- Wenxi Yue, Jing Zhang, Kun Hu, Yong Xia, Jiebo Luo, and Zhiyong Wang. Surgicalsam: Efficient class promptable surgical instrument segmentation. *arXiv preprint arXiv:2308.08746*, 2023.
- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*, 2021.
- Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023a.
- 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. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 161–170. Springer, 2023b.

Kaidong Zhang and Dong Liu. Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785*, 2023.

Li Zhang, Youwei Liang, and Pengtao Xie. Blo-sam: Bi-level optimization based overfitting-preventing finetuning of sam. *arXiv preprint arXiv:2402.16338*, 2024a.

Yichi Zhang and Rushi Jiao. Towards segment anything model (sam) for medical image segmentation: a survey. *arXiv [Preprint]*, 2023.

Yichi Zhang, Zhenrong Shen, and Rushi Jiao. Segment anything model for medical image segmentation: Current applications and future directions. *Computers in Biology and Medicine*, page 108238, 2024b.

Lei Zhou, Huidong Liu, Joseph Bae, Junjun He, Dimitris Samaras, and Prateek Prasanna. Self pre-training with masked autoencoders for medical image classification and segmentation. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–6. IEEE, 2023.

Jia-Xin Zhuang, Luyang Luo, and Hao Chen. Advancing volumetric medical image segmentation via global-local masked autoencoder. *arXiv preprint arXiv:2306.08913*, 2023.



## Appendix A. Proof of the central theorem

In this appendix we prove the central theorem and present additional experimental results.

*Remainder omitted in this sample.*

## Appendix B. Appendix section

In the appendix, we provide the numerical results of all experiments.

Tables 5, 6, 7 present the results of applying task-specific supervised training to ViT-H, ViT-B, and ViT-T respectively, and table 8 shows the performance of UNets. Figure 4 summarizes the results of these 4 tables, and Figure 11 visualizes some examples of segmentation results of using foundation model vs. using nnUNet.

Tables 9, 10, 11 show the performance of replacing SAM's weights with SupSAM (acquired from task-expansive prompt-based learning), MedSAM (from Ma et al. (2024)), and SSLSAM (acquired from task-agnostic self-supervised learning), respectively. Since all weights utilize the same encoder architecture (ViT-B), we show the difference by conducting Table 6 minus the corresponding table in Figure 6, 7, and 8, respectively.

Tables 13 and 12 show the performance of applying task-specific supervised training in the few-shot learning setting, which is summarized in Figure 9. The results for the selected three datasets of expanded few-shot learning: 10-shot, 20-shot, are shown in Figure 12.

Table 14 shows the performance of initial SAM in the point-based segmentation setting, with tables 15, 16, and 17 showing applying task-specific supervised training in the same setting. Tables 18, 19, 20, and 21 are conducted in the same setting but with the box as the prompt. All results are summarized in Figure 10.

Dataset	ViT-H													
	Decoder						En/Decoder							
	V		A		L		V		A		A*		L	
	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	82.58	31.37	82.42	33.33	78.60	35.08	73.79	18.47	88.97	42.61	88.80	47.09	89.64	50.01
MRI-Heart	80.86	64.16	80.61	62.69	80.53	58.43	73.69	48.07	85.07	75.47	86.22	76.76	85.07	74.35
MRI-Brain	45.32	19.37	60.19	26.92	50.87	23.94	53.89	15.06	80.53	42.41	82.65	54.43	84.09	62.38
MRI-Scler.	21.16	21.03	22.01	21.18	22.51	11.04	27.81	22.36	42.74	20.04	48.44	48.32	47.24	47.96
MRI-Hippo.	81.25	29.43	80.30	26.83	79.01	22.84	81.22	27.22	81.61	20.61	82.56	29.08	83.33	43.17
MRI-Kidney	82.34	57.39	84.08	58.66	80.60	48.60	87.65	62.25	88.81	68.05	90.29	68.47	83.02	58.77
MRI-Stroke.	47.52	35.99	53.06	40.32	47.22	34.59	56.51	52.79	71.64	60.21	67.82	55.95	72.39	60.86
MRI-Resect.	43.96	26.17	60.09	42.77	54.23	31.38	57.26	38.33	75.03	61.56	67.77	35.28	76.62	65.12
MRI-Breast	59.48	66.62	61.10	67.89	59.05	65.72	60.41	66.50	69.00	74.91	69.11	74.05	71.71	77.69
CT-Colon	22.67	17.00	35.08	26.53	26.77	17.40	24.12	12.65	55.44	42.78	54.30	43.06	60.91	55.61
CT-Panc.	61.78	41.37	67.43	46.85	64.06	43.96	64.45	40.67	73.92	55.41	73.01	53.38	75.74	59.90
CT-vessel.	40.94	45.68	44.75	49.84	42.15	46.23	40.96	54.58	58.16	55.81	54.67	50.22	58.53	60.53
US-Breast	71.27	26.74	74.17	28.43	70.22	25.86	64.91	18.29	77.33	32.43	78.18	33.09	77.36	36.41
US-Muscle	84.50	19.28	85.52	20.39	85.05	20.50	86.66	23.12	86.16	19.29	86.92	20.15	87.49	24.13
US-Kidney	96.97	68.83	97.08	69.05	96.53	62.97	71.34	21.92	98.13	81.29	98.34	82.05	98.18	81.88
Xray-Hip	93.78	73.71	94.35	72.17	93.00	74.47	92.02	70.84	93.96	75.77	94.33	73.78	94.73	73.44
Xray-Chest	95.29	63.29	95.35	64.03	95.22	62.69	94.85	56.99	95.81	69.49	95.45	70.11	95.74	70.56
Average	65.36	41.86	69.24	44.81	66.34	40.57	65.39	38.24	77.35	52.83	77.58	53.84	78.93	58.99

Table 5: The performance of SAM under different fine-tuning strategies in the automated mode across different datasets. V = vanilla, A = Adapter, L = LoRA. A\* is a variant of A in which we applied Adapter after each network layer.

Dataset	ViT-B											
	Decoder						En/Decoder					
	V		A		L		V		A		L	
	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	82.77	33.71	83.44	36.50	79.51	29.78	79.36	25.69	89.50	49.88	89.37	42.66
MRI-Heart	80.29	63.53	82.16	66.50	80.90	62.60	78.16	51.79	86.01	76.92	85.04	24.47
MRI-Brain	62.73	29.46	70.23	36.56	65.35	31.54	78.57	42.86	81.21	47.26	81.65	57.61
MRI-Scler.	27.95	28.34	25.92	25.54	27.91	26.47	32.90	28.67	42.12	37.57	45.83	32.28
MRI-Hippo.	81.08	31.21	81.09	27.50	80.35	30.07	82.45	28.69	81.52	28.67	82.69	26.36
MRI-Kidney	85.29	60.88	82.50	53.40	79.58	42.45	86.70	59.03	89.19	68.82	84.93	72.13
MRI-Stroke.	57.33	43.97	58.47	45.58	52.67	35.47	65.32	55.95	72.05	61.61	74.24	63.65
MRI-Resect.	58.74	39.84	64.76	48.44	57.23	7.85	51.28	22.87	75.53	63.25	71.23	65.18
MRI-Breast	61.60	68.69	63.16	70.46	61.02	68.56	49.30	64.10	69.54	74.92	69.95	76.68
CT-Colon	33.25	22.30	39.29	30.17	31.12	21.64	19.18	11.77	62.30	54.53	60.87	53.95
CT-Panc.	62.28	38.55	68.34	48.62	63.97	42.83	66.25	45.00	74.16	57.06	74.07	57.58
CT-vessel.	44.95	48.46	48.71	51.39	42.16	48.28	49.89	53.71	58.56	60.45	58.08	59.97
US-Breast	69.64	27.40	69.93	23.73	69.56	24.00	70.02	25.38	78.87	34.87	77.97	32.92
US-Muscle	84.53	19.51	85.53	19.19	85.61	21.12	86.80	23.43	87.24	24.20	87.37	23.90
US-Kidney	96.70	66.50	96.76	66.41	96.99	67.93	67.12	12.29	98.08	80.80	98.05	80.35
Xray-Hip	94.29	70.04	94.11	68.23	97.68	71.44	94.14	71.53	94.44	79.32	94.73	81.01
Xray-Chest	95.45	65.18	95.29	62.85	95.45	65.48	95.58	66.31	95.73	69.87	95.78	69.47
Average	69.80	44.81	71.16	46.22	64.77	40.11	67.82	40.53	78.59	57.05	78.69	54.12

Table 6: The performance of SAM under different fine-tuning strategies in the automated mode across different datasets. V = vanilla, A = Adapter, L = LoRA.

Dataset	ViT-T											
	Decoder						En/Decoder					
	V		A		L		V		A		L	
	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	81.36	32.07	80.01	30.26	77.17	27.22	88.90	49.55	89.37	48.57	86.70	41.30
MRI-Heart	80.75	64.09	80.39	60.28	76.24	50.87	85.15	74.76	84.93	73.79	84.35	72.54
MRI-Brain	52.68	23.88	58.43	26.55	28.12	8.20	78.32	53.18	79.42	56.09	77.14	52.36
MRI-Scler.	21.20	20.38	18.22	15.84	15.29	14.47	42.60	43.84	41.05	40.98	41.92	42.54
MRI-Hippo.	78.97	28.95	78.39	23.73	76.00	23.74	82.90	44.45	82.96	43.85	82.56	39.60
MRI-Kidney	82.49	52.68	79.93	46.70	67.54	36.20	89.85	70.32	88.39	68.55	88.38	68.96
MRI-Stroke.	42.13	30.31	46.35	34.00	34.08	19.68	71.02	59.46	72.14	60.31	70.37	58.82
MRI-Resect.	40.35	27.88	40.56	21.47	37.10	19.96	74.34	61.28	75.74	63.86	72.93	60.12
MRI-Breast	53.06	61.55	55.10	62.58	45.23	54.33	66.92	74.75	68.28	76.27	66.18	72.84
CT-Colon	14.97	11.16	23.09	14.35	23.16	15.50	55.99	48.99	53.76	47.94	49.63	41.11
CT-Panc.	59.01	38.52	64.54	42.06	55.43	33.58	72.08	54.74	73.21	57.13	70.36	52.21
CT-vessel.	37.25	41.47	36.55	37.00	32.25	37.85	56.62	59.39	57.25	59.08	55.09	56.54
US-Breast	61.42	18.59	67.38	21.88	49.64	13.28	77.55	35.52	74.50	32.32	71.15	30.16
US-Muscle	84.68	19.37	85.50	19.87	59.94	7.15	86.72	23.67	86.71	22.52	87.18	23.94
US-Kidney	96.11	61.66	96.37	59.30	92.60	44.55	97.45	81.29	98.06	80.41	97.88	78.02
Xray-Hip	97.58	68.87	93.83	64.80	91.04	46.13	92.33	76.82	94.33	76.34	94.61	76.76
Xray-Chest	95.08	60.87	95.26	62.78	91.50	43.07	95.87	70.25	95.75	68.62	95.72	68.49
Average	63.27	39.22	64.57	38.10	56.02	29.34	77.33	58.11	77.40	57.78	76.01	55.36

Table 7: The performance of SAM under different fine-tuning strategies in the automated mode across different datasets. V = Vanilla, A = Adapter, L = LoRA.

Dataset	UNet		nnUNet		SwinUNet	
	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	82.59	11.64	81.98	46.39	82.03	10.78
MRI-Heart	82.85	25.12	87.14	79.51	84.10	32.13
MRI-Brain	79.00	16.97	74.13	50.37	75.83	18.20
MRI-Scler.	36.71	12.38	40.56	39.93	36.73	11.69
MRI-Hippo.	80.59	14.53	76.18	37.62	83.25	26.12
MRI-Kidney	87.03	22.60	91.69	72.90	83.83	20.85
MRI-Stroke.	69.67	26.84	72.27	55.86	70.45	35.97
MRI-Resect.	71.07	27.88	74.15	63.60	65.91	21.41
MRI-Breast	68.77	56.02	72.49	75.52	70.33	57.44
CT-Colon	43.64	12.04	47.58	43.25	46.39	12.00
CT-Panc.	70.06	17.38	74.78	58.07	68.54	15.96
CT-vessel.	57.06	31.49	58.44	58.50	56.06	31.89
US-Breast	64.83	5.26	76.03	33.56	72.02	7.80
US-Muscle	83.99	4.99	87.30	25.77	86.62	7.32
US-Kidney	96.74	43.96	98.37	83.16	97.75	51.09
Xray-Hip	97.86	33.56	98.46	83.52	94.29	33.41
Xray-Chest	95.51	23.82	96.40	70.29	95.62	23.89
Average	74.44	22.86	76.94	57.52	74.69	24.71

Table 8: The performance of UNet, nnUNet, and SwinUNet across different datasets.




Figure 11: Comparison of Segmentation Results: Examples showcasing the performance of the foundation model (ViT-H + L + Encoder/Decoder) versus the non-foundation-based method (nnUNet).




Figure 12: Visualization of three selected datasets' model (ViT-B, EN/Decoder, Adapter) performance (DSC) vs. different dataset availability: 5-shot, 10-shot, 20-shot, and the full set.



Dataset	SupSAM											
	Decoder						En/Decoder					
	V		A		L		V		A		L	
	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	71.48	18.20	72.12	19.67	81.98	31.59	82.01	29.76	81.65	29.06	83.63	31.45
MRI-Heart	79.78	58.22	75.57	48.18	76.34	51.57	83.95	69.98	82.16	65.93	83.94	66.78
MRI-Brain	59.44	27.04	63.61	29.37	62.68	30.42	79.95	42.81	76.89	41.55	75.07	39.52
MRI-Scler.	20.96	18.96	22.10	19.53	18.84	16.20	34.32	29.70	32.88	31.32	32.58	31.24
MRI-Hippo.	80.91	27.00	79.83	25.63	79.98	26.10	82.27	28.68	81.98	29.73	81.74	28.03
MRI-Kidney	77.88	46.81	79.99	49.42	77.10	46.21	87.99	62.83	85.71	59.58	85.45	59.88
MRI-Stroke.	62.63	51.19	64.45	53.61	62.68	50.60	67.42	59.07	67.77	58.47	67.29	56.92
MRI-Resect.	58.51	36.59	59.60	38.20	55.47	34.02	69.20	50.09	64.99	46.51	67.07	46.93
MRI-Breast	55.50	63.82	57.57	65.76	55.74	63.65	58.06	66.64	58.01	67.31	58.19	67.81
CT-Colon	15.07	10.17	16.07	10.24	18.39	10.70	38.30	25.62	33.68	24.19	35.86	26.17
CT-Panc.	47.16	23.66	52.58	27.69	45.22	23.81	66.34	43.13	62.56	39.01	63.41	41.59
CT-Vessel	37.08	48.31	40.87	50.24	34.23	43.51	50.27	52.09	51.73	54.59	51.67	54.21
US-Breast	43.41	8.53	47.48	8.89	39.77	7.50	73.44	28.24	66.99	21.35	68.15	21.98
US-Muscle	83.26	16.17	82.18	15.89	80.08	13.26	87.05	23.04	85.81	19.11	86.27	20.46
US-Kidney	95.86	55.78	91.83	43.45	88.28	25.50	97.69	74.37	96.61	65.56	96.81	66.91
Xray-Hip	90.90	38.34	89.89	45.61	90.12	36.53	95.64	41.32	93.03	61.71	96.32	66.39
Xray-Chest	93.56	48.37	93.46	47.15	92.25	41.28	94.95	62.11	95.28	63.15	95.36	63.78
Average	63.14	35.29	64.07	32.66	62.93	32.66	73.59	46.66	72.06	46.03	72.12	46.75

Table 9: The performance of applying various task-specific supervised training to SSLSAM, weights acquired from task-expansive prompt-based learning.

Dataset	MedSAM											
	Decoder						En/Decoder					
	V		A		L		V		A		L	
	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	71.40	18.05	75.78	20.54	70.82	15.44	87.42	40.69	89.76	48.84	89.02	47.40
MRI-Heart	80.13	60.41	81.90	64.03	80.27	63.26	84.97	73.29	85.17	76.56	85.59	75.33
MRI-Brain	61.76	25.07	69.62	32.20	60.33	23.47	69.54	33.80	80.60	54.07	78.98	51.51
MRI-Scler.	22.20	18.29	25.23	22.33	25.16	22.49	36.28	34.08	43.49	43.14	40.43	39.61
MRI-Hippo.	79.81	25.76	81.69	28.10	80.67	25.44	82.70	30.19	82.71	30.31	82.67	29.24
MRI-Kidney	81.04	49.73	82.28	49.14	79.52	46.28	87.95	61.80	88.38	68.38	70.84	69.14
MRI-Stroke.	60.21	46.98	64.02	52.55	62.38	48.93	67.30	58.11	72.09	60.98	70.84	63.02
MRI-Resect.	51.53	26.41	54.56	30.64	49.31	24.98	67.37	49.06	76.41	63.31	70.10	55.80
MRI-Breast	53.95	63.03	54.80	62.91	53.02	60.66	61.93	70.06	67.37	74.47	66.64	74.32
CT-Colon	21.84	13.58	26.25	15.34	22.54	14.78	40.74	31.36	55.23	44.76	45.30	35.31
CT-Panc.	51.22	29.20	59.16	35.69	53.69	30.07	66.07	47.49	72.63	55.42	72.26	52.42
CT-Vessel	40.19	43.35	44.54	49.16	39.74	41.14	53.63	56.56	57.83	59.06	56.22	58.27
US-Breast	57.55	14.65	67.14	20.98	62.34	18.36	60.06	15.43	76.88	33.15	78.48	33.60
US-Muscle	84.22	19.43	85.64	20.56	84.11	18.53	86.85	22.87	86.90	24.33	87.25	24.51
US-Kidney	96.05	58.25	97.03	68.55	97.18	61.07	97.89	63.74	98.18	81.37	98.12	80.62
Xray-Hip	95.52	62.12	93.49	62.18	93.27	59.53	94.74	78.87	94.73	77.58	94.55	75.22
Xray-Chest	95.06	60.54	95.28	61.65	95.25	61.73	95.79	69.20	95.86	69.86	95.77	69.92
Average	64.80	37.58	68.14	41.22	65.27	37.65	73.14	49.52	77.90	57.13	75.47	55.33

Table 10: The performance of applying various task-specific supervised training to MedSAM, weights acquired from Ma et al. (2024).

Dataset	SSLSAM											
	Decoder						En/Decoder					
	V		A		L		V		A		L	
	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	82.86	33.48	84.38	34.08	83.55	34.52	88.82	44.53	90.29	50.10	86.74	39.23
MRI-Heart	78.25	58.08	83.67	68.59	84.08	69.83	84.45	73.93	87.41	78.75	85.52	73.81
MRI-Brain	63.26	29.01	71.22	36.88	64.77	31.13	76.71	50.13	83.31	55.51	80.33	51.39
MRI-Scler.	31.96	29.36	32.32	29.82	31.32	27.77	39.80	40.00	47.03	41.60	37.74	35.98
MRI-Hippo.	80.63	30.00	81.95	29.57	81.90	31.25	83.13	40.52	82.83	34.75	82.28	32.71
MRI-Kidney	85.54	58.16	87.10	61.05	84.91	56.20	89.80	64.29	89.57	68.57	89.69	67.00
MRI-Stroke.	65.10	52.10	66.71	55.73	63.80	48.91	70.34	59.74	73.89	62.33	71.11	59.37
MRI-Resect.	67.92	49.55	69.26	52.37	69.98	51.64	71.37	57.10	78.73	66.47	75.64	61.82
MRI-Breast	62.30	70.34	62.15	70.77	60.97	69.03	66.23	73.68	69.72	75.92	67.65	73.66
CT-Colon	35.31	23.78	37.70	8.23	23.44	10.42	45.29	41.50	55.56	45.50	45.49	35.51
CT-Panc.	58.99	35.69	66.19	42.16	61.29	38.58	72.07	53.84	74.50	57.05	70.66	50.63
CT-Vessel	48.49	49.78	47.76	51.46	49.97	49.68	42.14	53.52	57.67	59.99	57.51	57.86
US-Breast	56.43	15.21	68.46	22.30	63.22	18.58	78.15	33.97	77.84	33.31	75.16	28.54
US-Muscle	84.26	19.79	85.76	19.31	85.33	20.29	85.67	24.60	87.70	22.64	86.99	23.00
US-Kidney	96.64	64.48	97.34	70.84	96.78	65.26	98.02	81.94	98.24	82.39	97.96	78.80
Xray-Hip	93.73	65.77	93.18	59.00	93.59	63.39	94.14	23.03	94.73	76.89	93.96	67.01
Xray-Chest	95.32	63.68	95.04	58.85	95.23	62.57	94.93	52.72	95.72	66.90	95.80	68.38
Average	69.82	44.28	72.36	45.61	70.24	44.32	75.36	51.32	79.10	57.95	76.48	53.51

Table 11: The performance of applying various task-specific supervised training to SSLSAM, weights acquired from task-agnostic self-supervised learning.

Dataset	ViT-B				MedSAM				SSLSAM			
	Decoder		En/Decoder		Decoder		En/Decoder		Decoder		En/Decoder	
	V	A	V	A	V	A	V	A	V	A	V	A
DSC												
MRI-Prost.	66.77	68.15	45.47	73.11	58.73	57.93	73.04	70.20	62.69	67.46	69.67	73.88
MRI-Heart	56.94	62.55	19.14	72.34	64.70	59.65	69.21	66.28	69.00	68.90	35.30	61.30
MRI-Brain	31.54	38.29	3.28	60.31	17.10	43.01	41.33	64.50	31.62	50.19	0.09	55.09
MRI-Scler.	2.34	1.49	2.26	16.17	4.17	3.58	13.08	12.91	2.59	2.39	0.98	12.70
MRI-Hippo.	60.12	51.86	70.45	67.79	46.23	51.93	65.05	60.49	42.91	51.42	0.00	64.73
MRI-Kidney	74.46	74.45	67.70	76.21	68.40	66.98	79.27	67.56	73.90	66.70	0.00	77.10
MRI-Stroke.	8.59	15.46	25.20	25.54	11.60	15.79	27.09	34.02	14.03	26.60	31.60	25.50
MRI-Resect.	27.01	35.33	11.30	51.30	13.90	17.75	29.38	33.00	36.81	46.40	4.44	50.30
MRI-Breast	45.00	43.83	33.61	52.25	36.30	37.60	38.12	42.27	40.17	42.50	21.00	48.80
Average	41.42	43.49	36.05	55.01	35.83	39.36	48.37	50.35	41.52	46.95	18.12	52.16
NSD												
MRI-Prost.	19.37	18.60	1.21	23.61	7.96	9.52	21.55	19.49	13.99	14.48	6.08	20.06
MRI-Heart	31.09	41.58	0.08	50.52	41.53	27.22	45.22	42.02	42.70	40.51	0.00	34.04
MRI-Brain	7.76	11.94	3.95	21.71	5.30	13.17	10.57	32.38	8.00	15.63	1.27	19.34
MRI-Scler.	1.75	1.62	3.33	17.00	2.92	2.87	11.14	14.34	1.16	1.89	7.56	8.76
MRI-Hippo.	12.03	10.64	7.80	23.79	8.31	10.24	13.90	12.49	5.77	10.84	0.76	13.63
MRI-Kidney	44.17	42.86	0.00	51.76	23.77	25.00	43.70	28.31	39.61	32.26	0.00	45.29
MRI-Stroke.	7.25	10.52	18.93	27.77	8.64	10.39	18.03	26.07	9.92	18.27	0.00	16.55
MRI-Resect.	12.64	14.36	2.52	27.63	6.18	8.89	7.08	11.89	13.95	19.36	1.01	21.38
MRI-Breast	50.72	49.82	31.22	58.87	38.63	41.03	39.66	46.58	45.33	48.69	1.16	52.98
Average	20.83	22.52	7.68	33.63	15.97	16.53	23.52	26.03	20.12	22.51	1.99	25.87

Table 12: The performance of applying task-specific supervised training to the ViT-B architecture on the few-shot setting 5-shot with different initial weights.

Dataset	UNet		nnUNet		SwinUNet	
	DSC	NSD	DSC	NSD	DSC	NSD
MRI-Prost.	6.34	1.43	70.13	22.75	2.67	1.26
MRI-Heart	1.26	4.14	67.69	55.95	59.94	34.23
MRI-Brain	6.49	1.32	46.60	25.78	53.69	18.20
MRI-Scler.	0.77	1.25	10.40	9.35	1.64	0.37
MRI-Hippo.	34.05	4.39	71.29	22.34	46.58	7.76
MRI-Kidney	47.95	21.63	75.81	48.88	47.51	18.00
MRI-Stroke.	2.01	4.67	37.30	26.20	38.99	23.36
MRI-Resect.	2.66	2.50	54.34	38.29	22.16	4.63
MRI-Breast	32.40	35.14	54.06	58.97	37.39	37.79
Average	14.88	8.50	54.18	34.28	34.51	16.18

Table 13: The performance of UNet, nnUNet, and SwinUNet across different datasets under a few-shot setting (5-shot).

Dataset	ViT-H		ViT-B		ViT-T	
	H	O	H	O	H	O
MRI-Prost.	52.63	67.65	35.52	68.94	41.06	65.71
MRI-Heart	49.20	58.75	48.39	65.24	43.50	63.42
MRI-Brain	34.20	54.57	29.23	58.55	30.33	58.38
MRI-Scler.	1.80	3.91	1.58	5.28	1.72	4.50
MRI-Hippo.	34.78	55.15	17.34	59.33	22.20	59.94
MRI-Kidney	41.41	73.31	44.17	76.22	46.07	70.85
MRI-Stroke.	21.38	39.85	21.75	49.98	20.17	44.43
MRI-Resect.	27.86	45.84	17.38	60.82	26.07	55.89
MRI-Breast	5.00	21.61	12.83	21.81	5.00	28.57
Average	29.81	46.74	25.35	51.80	26.24	50.19

Table 14: The performance of SAM with different choices of encoder architectures with point prompts. H = Highest confidence, O = Oracle mode.

Dataset	ViT-H					
	Decoder			En/Decoder		
	V	A	L	V	A	L
MRI-Prost.	87.74	86.89	86.82	79.84	89.55	88.10
MRI-Heart	84.77	85.07	86.46	83.77	85.60	86.40
MRI-Brain	74.44	72.87	72.66	74.15	80.76	75.09
MRI-Scler.	62.95	62.64	62.64	51.42	60.77	64.26
MRI-Hippo.	83.13	82.03	82.95	83.99	81.88	83.26
MRI-Kidney	88.49	87.37	89.06	89.31	86.87	92.01
MRI-Stroke.	61.18	66.09	60.47	73.87	78.22	81.52
MRI-Resect.	73.38	75.20	68.06	66.57	79.96	77.09
MRI-Breast	64.84	69.56	67.61	70.12	77.14	78.13
Average	75.66	76.41	75.19	74.78	80.08	80.65

Table 15: The performance of applying task-specific supervised training during point-based segmentation with ViT-H as the encoder architecture.

Dataset	ViT-B					
	Decoder			En/Decoder		
	V	A	L	V	A	L
MRI-Prost.	87.23	87.22	87.10	88.54	89.57	87.48
MRI-Heart	84.90	85.20	84.98	86.04	87.54	85.41
MRI-Brain	76.05	76.48	77.71	73.82	83.70	77.61
MRI-Scler.	58.41	59.18	62.78	59.42	66.95	60.98
MRI-Hippo.	82.06	82.98	83.66	84.23	84.04	83.43
MRI-Kidney	89.34	89.04	86.93	88.81	90.80	89.77
MRI-Stroke.	69.08	69.67	52.63	72.62	79.19	78.12
MRI-Resect.	73.11	77.15	73.02	71.36	81.56	75.39
MRI-Breast	64.57	67.79	69.79	56.86	73.38	76.28
Average	76.08	77.19	75.40	75.74	81.86	79.39

Table 16: The performance of applying task-specific supervised training during point-based segmentation with ViT-B as the encoder architecture.

Dataset	ViT-T					
	Decoder			En/Decoder		
	V	A	L	V	A	L
MRI-Prost.	87.76	84.06	78.06	90.46	90.81	84.38
MRI-Heart	85.19	85.50	79.28	87.72	88.03	86.60
MRI-Brain	74.38	72.63	59.00	84.12	85.78	83.98
MRI-Scler.	61.29	63.85	50.53	64.78	72.56	71.04
MRI-Hippo.	82.97	80.71	79.20	84.95	87.36	77.90
MRI-Kidney	88.30	87.60	81.25	91.12	91.48	89.69
MRI-Stroke.	63.80	65.78	64.84	76.22	81.61	81.21
MRI-Resect.	74.08	74.48	68.94	80.66	81.62	79.82
MRI-Breast	59.92	64.80	59.12	69.55	75.23	73.61
Average	75.30	75.49	68.91	81.06	83.83	80.91

Table 17: The performance of applying task-specific supervised training during point-based segmentation with ViT-T as the encoder architecture.



Dataset	ViT-H		ViT-B		ViT-T		MedSAM		SupSAM
	H	O	H	O	H	O	H	O	-
MRI-Prost.	87.09	91.60	89.75	92.20	88.50	90.99	87.02	92.51	91.82
MRI-Heart	67.29	85.62	83.03	85.96	83.49	86.82	73.99	75.73	83.74
MRI-Brain	73.44	84.51	81.74	85.48	81.70	84.50	77.13	78.32	80.57
MRI-Scler.	39.13	78.58	75.15	79.91	76.43	79.72	38.52	41.28	56.26
MRI-Hippo.	71.19	81.37	78.48	82.31	78.52	81.75	70.76	72.69	84.75
MRI-Kidney	87.51	89.58	87.84	90.02	87.40	89.31	70.36	72.44	87.46
MRI-Stroke.	55.97	81.09	78.14	82.21	77.50	82.15	55.88	57.95	78.06
MRI-Resect.	71.97	85.68	83.58	87.09	81.58	84.85	66.92	68.20	78.90
MRI-Breast	54.69	63.38	56.87	66.96	54.09	63.06	19.52	21.12	64.70
Average	67.59	82.38	79.40	83.57	78.80	82.57	62.23	64.47	78.47

Table 18: The performance of SAM with different choices of encoder architectures with box prompts. H = Highest confidence, O = Oracle mode. MedSAM and SupSAM is applicable since both utilize box during training.

Dataset	ViT-H					
	Decoder			En/Decoder		
	V	A	L	V	A	L
MRI-Prost.	91.61	94.18	94.52	91.08	94.82	95.16
MRI-Heart	64.24	90.66	88.91	86.21	89.17	91.88
MRI-Brain	77.84	82.93	83.21	73.03	87.09	88.14
MRI-Scler.	62.29	73.70	74.06	62.53	74.84	72.18
MRI-Hippo.	84.95	86.02	88.32	88.14	88.61	89.45
MRI-Kidney	89.44	89.59	90.06	89.51	90.06	90.15
MRI-Stroke.	81.05	81.14	81.95	81.66	87.70	86.42
MRI-Resect.	81.66	81.29	81.47	83.23	86.93	87.70
MRI-Breast	72.91	74.71	72.89	69.55	79.70	80.68
Average	78.44	83.80	83.93	80.55	86.55	86.86

Table 19: The performance of applying task-specific supervised training during box-based segmentation with ViT-H as the encoder architecture.

Dataset	ViT-B					
	Decoder			En/Decoder		
	V	A	L	V	A	L
MRI-Prost.	94.13	94.43	94.76	92.76	94.78	94.58
MRI-Heart	90.27	90.52	90.35	90.18	90.48	90.05
MRI-Brain	82.68	82.85	84.71	76.32	87.98	85.27
MRI-Scler.	68.58	71.80	69.12	61.74	69.43	72.00
MRI-Hippo.	86.38	87.02	89.04	85.89	88.24	88.68
MRI-Kidney	91.37	90.33	89.44	91.63	91.53	90.89
MRI-Stroke.	72.45	81.19	75.46	67.33	85.84	82.76
MRI-Resect.	77.15	77.35	83.83	76.86	84.91	85.49
MRI-Breast	64.57	69.88	73.70	59.66	74.00	73.14
Average	80.84	82.82	83.38	78.04	85.24	84.76

Table 20: The performance of applying task-specific supervised training during box-based segmentation with ViT-B as the encoder architecture.

Dataset	ViT-T					
	Decoder			En/Decoder		
	V	A	L	V	A	L
MRI-Prost.	93.37	94.45	94.07	93.51	94.98	94.13
MRI-Heart	88.75	91.26	90.10	90.50	91.02	91.05
MRI-Brain	78.92	82.23	79.38	84.15	88.69	88.01
MRI-Scler.	75.41	76.24	74.59	74.29	79.49	78.49
MRI-Hippo.	86.64	88.01	87.59	87.60	90.38	86.85
MRI-Kidney	90.27	89.37	87.13	92.15	92.33	91.22
MRI-Stroke.	70.94	81.77	81.40	81.44	89.33	88.13
MRI-Resect.	76.44	85.68	82.91	81.00	86.77	87.62
MRI-Breast	62.81	70.90	63.12	74.09	80.54	74.74
Average	80.39	84.43	82.25	84.30	88.17	86.69

Table 21: The performance of applying task-specific supervised training during box-based segmentation with ViT-T as the encoder architecture.