ACS 2025 Surgeons and Engineers: A Dialogue on Surgical Simulation Meeting

P-D-07

Research Abstracts

S-SAM: SVD-based Fine-Tuning of Segment Anything Model for Medical Image Segmentation

Jay Nitin Paranjape; Shameema Sikder, MD, FACS; S. Swaroop Vedula, MBBS, PhD, MPH; and Vishal M. Patel

Johns Hopkins University, Baltimore, MD; Johns Hopkins University School of Medicine, Baltimore, MD

Introduction: Medical image segmentation has traditionally been approached by fine-tuning entire models for each new modality or dataset, a process that is often computationally intensive and time-consuming. While the Segment Anything Model (SAM) has been introduced to facilitate efficient segmentation in natural images, its adaptation to medical imaging still demands expert annotations—such as point or bounding box prompts—during both training and inference. This dependency on expert input presents a significant challenge for the practical implementation of SAM in clinical settings.

Methods: In this work, we introduce S-SAM, a novel adaptation technique that addresses these challenges by significantly reducing the number of trainable parameters required for tuning—specifically, to only 0.4% of SAM's parameters. S-SAM distinguishes itself by eliminating the need for expert prompts during both training and inference. Instead, it utilizes label names as prompts to generate accurate segmentation masks. This approach enhances the efficiency of the tuning process while reducing the reliance on time-consuming expert annotations

Results: We evaluated S-SAM across five distinct medical imaging modalities: endoscopic images, x-rays, ultrasound, CT, and histology images. Our experimental results demonstrate that S-SAM not only achieves superior performance compared to state-of-the-art methods but also surpasses existing SAM adaptation techniques like LoRA or Adaptive. Notably, S-SAM accomplishes this while tuning a significantly smaller subset of parameters, underscoring its effectiveness and efficiency in diverse medical imaging tasks.

Conclusions: S-SAM addresses key limitations associated with existing adaptation methods for SAM. By reducing the need for extensive parameter tuning and expert prompts, S-SAM offers a more practical and resource-efficient solution for clinical applications. Our findings suggest that S-SAM is well-suited for adapting SAM to a variety of medical imaging challenges, with the potential to outperform current state-of-the-art techniques. The code for S-SAM is publicly available.