

Efficient MedSAMs: Segment Anything in Medical Images on Laptop

Jun Ma, Feifei Li*, Sumin Kim*, Reza Asakereh, Bao-Hiep Le, Dang-Khoa Nguyen-Vu, Alexander Pfefferle, Muxin Wei, Ruochen Gao, Donghang Lyu, Songxiao Yang, Lennart Purucker, Zdravko Marinov, Marius Staring, Haisheng Lu, Thuy Thanh Dao, Xincheng Ye, Zhi Li, Gianluca Brugnara, Philipp Vollmuth, Martha Foltyn-Dumitru, Jaeyoung Cho, Mustafa Ahmed Mahmutoglu, Martin Bendszus, Irada Pflüger, Aditya Rastogi, Dong Ni, Xin Yang, Guang-Quan Zhou, Kaini Wang, Nicholas Heller, Nikolaos Papanikolopoulos, Christopher Weight, Yubing Tong, Jayaram K Udupa, Cahill J. Patrick, Yaqi Wang, Yifan Zhang, Francisco Contijoch, Elliot McVeigh, Xin Ye, Shucheng He, Robert Haase, Thomas Pinetz, Alexander Radbruch, Inga Krause, Erich Kobler, Jian He, Yucheng Tang, Haichun Yang, Yuankai Huo, Gongning Luo, Kaisar Kushibar, Jandos Amankulov, Dias Toleshbayev, Amangeldi Mukhamejan, Jan Egger, Antonio Pepe, Christina Gsaxner, Gijs Luijten, Shohei Fujita, Tomohiro Kikuchi, Benedikt Wiestler, Jan S. Kirschke, Ezequiel de la Rosa, Federico Bolelli, Luca Lumetti, Costantino Grana, Kunpeng Xie, Guomin Wu, Behrus Puladi, Carlos Martín-Isla, Karim Lekadir, Victor M. Campello, Wei Shao, Wayne Brisbane, Hongxu Jiang, Hao Wei, Wu Yuan, Shuangli Li, Yuyin Zhou, and Bo Wang

Abstract

Promptable segmentation foundation models have emerged as a transformative approach to addressing the diverse needs in medical images, but most existing models require expensive computing, posing a big barrier to their adoption in clinical practice. In this work, we organized the first international competition dedicated to promptable medical image segmentation, featuring a large-scale dataset spanning nine common imaging modalities from over 20 different institutions. The top teams developed lightweight segmentation foundation models and implemented an efficient inference pipeline that substantially reduced computational requirements while maintaining state-of-the-art segmentation accuracy. Moreover, the post-challenge phase advanced the algorithms through the design of performance booster and reproducibility tasks, resulting in improved algorithms and validated reproducibility of the winning solution. Furthermore, the best-performing algorithms have been incorporated into the open-source software with a user-friendly interface to facilitate clinical adoption. The data and code are publicly available to foster the further development of medical image segmentation foundation models and pave the way for impactful real-world applications.



INTRODUCTION

Segmentation is a fundamental task in medical image analysis, aiming to provide accurate boundaries for anatomies and pathologies, which plays an important role in many clinical tasks, such as disease detection and diagnosis, surgical planning, and treatment monitoring [1], [2]. Over the past decade, deep learning-based models have revolutionized medical image segmentation by providing state-of-the-art accuracy and unprecedented automation [3]. More recently, the emergence of foundation models, such as Segment Anything Model (SAM) [4], [5], has introduced a new paradigm for segmentation tasks [6], [7]. These models are trained on large-scale annotated images and capable of being generalized across multiple domains and transferred to unseen domains [8]. For example, MedSAM [9] and BiomedParse [10] have demonstrated strong zero-shot segmentation ability across common medical imaging modalities and a wide range of anatomies by fine-tuning SAM on over one million medical image-mask pairs.

However, despite these advances, adopting foundation models for segmentation in clinical practice remains an ongoing challenge. One key bottleneck preventing the widespread use of these foundation models is the high demand for computing resources. Many of these segmentation foundation models are resource-intensive because of large model weights, making them difficult to deploy on commonly available hardware, such as laptops, creating a barrier to democratizing the use of advanced segmentation tools, especially in resource-constrained clinical environments. Additionally, as segmentation foundation models have been developed and validated under different settings, it is difficult to assess their performance rigorously. This calls for standardized benchmarks, yet their absence hinders the identification of optimal models [11].

International competitions have proven to be effective in driving advancements in various domains, such as ImageNet competition [12] for deep convolutional neural networks [13] and the critical assessment of protein structure prediction competition for AlphaFold [14]. By providing a standard platform for researchers to compare their models, these challenges foster collaborative efforts of researchers all over the world to identify novel approaches while ensuring transparent evaluation and reproducible results [15]–[17]. Thus, an international benchmark for promptable medical image segmentation can greatly boost the progress in creating lightweight, efficient models that are also suitable for clinical deployment.

In this work, we present the competition results for promptable medical image segmentation, specifically designed to encourage the development of more efficient segmentation foundation models for diverse modalities and segmentation targets. Importantly, we collaborate with 24 institutions worldwide to curate a new testing set with over 4000 cases, allowing for fair model evaluation while minimizing the risk of data leakage. Based on this comprehensive benchmark dataset, the top-performing algorithms demonstrate substantial improvements in both segmentation accuracy and efficiency, with models achieving segmentation results over ten times faster than existing SAM-based foundation models [4], [9]. Moreover, we work with the two best-performing algorithm developers and integrated the models into the open-source image computing platform 3D Slicer with a user-friendly interface, making state-of-the-art segmentation models more accessible and practical for a wider range of clinical applications.




Fig. 1. **Competition design.** **a.** The task is to develop universal segmentation foundation models that can accept various medical image inputs with target bounding box prompts and generate the corresponding segmentation masks. The model should be lightweight and deployable on laptops without reliance on graphics processing unit (GPU). **b.** Three phases in the competition. During the development phase, participants train their models on the training set and obtain performance metrics on the online validation set on Codabench. The top 20 teams on the validation leaderboard are invited to submit their algorithm dockers and we manually evaluate them on the hidden testing set. After that, we release all the technical details and code of the top ten teams and launch a post-challenge phase to invite participants to further boost their model performance and reproduce the winning algorithms.

RESULTS

Competition Design: Universal and Lightweight Medical Image Segmentation Foundation Models

The competition task is designed for promptable medical image segmentation. Specifically, the input contains various medical images and bounding boxes for segmentation targets, and the output is the corresponding segmentation mask. The bounding box is selected as the prompt because it can precisely specify the target, which has lower ambiguity than point prompts. The challenge consists of three distinct phases: a development phase lasting 122 days, a testing phase spanning 35 days, and a post-challenge phase of 35 days. These phases focus respectively on model training, evaluation on a hidden testing set, and subsequent improvements and reproducibility analysis.

During the development phase, participants were provided with a large-scale and diverse training dataset, serving as the foundation for participants to design, develop, and train their models. Moreover, an online leaderboard was made available on Codabench [18], allowing for automatic evaluation of participants' results. Participants could leverage this feature to test their models against a tuning set, obtain performance feedback, and iteratively tune and refine their approaches. This phase was instrumental in enabling participants to optimize their models and prepare them for subsequent testing.

During the testing phase, the top 20 teams from the validation leaderboard were invited to make testing submissions. However, submissions were not limited to these top teams and other participants were also welcome to submit their solutions for evaluation. Participants were required to encapsulate their algorithms in Docker containers, ensuring a standardized and portable format for evaluation. Each Docker container was manually executed on a hidden testing set, with all evaluations conducted on the same workstation to ensure fairness and consistency. The teams were ranked based on both the accuracy and efficiency of their models. Specifically, the evaluation criteria included the Dice Similarity Coefficient (DSC), Normalized Surface Distance (NSD), and runtime performance. This rigorous and equitable process determined the testing leaderboard rankings.

The post-challenge was designed with two primary objectives: further improving model performance by integrating successful strategies from the top-performing algorithms (performance booster) and evaluating the reproducibility of the winning solutions (reproducibility analysis). This phase aimed to push the boundaries of segmentation performance using collective knowledge and assess whether the winning solution could be easily implemented by others. In the performance-booster subtask, participants were encouraged to enhance their models by incorporating two new datasets, fast inference methods, and other optimization strategies contributed by participants. In the reproducibility subtask, all participants were invited to reproduce the results of the top-performing teams. We made the code and corresponding methodologies of top solutions publicly available, allowing participants to validate them on the testing set. These collaborative efforts not only fostered the development of a refined solution capable of outperforming the original winning algorithms but also verified the reproducibility of the top solutions.

Competition Dataset: A Global Collaboration

We provided participants with a diverse set of training and testing images to facilitate robust model development and evaluation. Figure 2a illustrates several visual examples from the training and testing sets across various imaging modalities, accompanied by their corresponding bounding box prompts and reference masks. The training dataset comprises 1,809,644 image-mask pairs spanning 10 imaging modalities, including Computed Tomography (CT), Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET), ultrasound, X-ray, Optical Coherence Tomography (OCT), endoscopy, fundus, and microscopy images. Figure 2b visualizes the distribution of image-mask pairs across these modalities, with CT and MRI images dominating the dataset due to their widespread use in clinical diagnostics. All images were sourced from publicly available datasets with proper licensing permissions and were preprocessed to ensure standardization in intensity range and format. However, some public datasets prohibit redistribution. To address this limitation, we compiled a comprehensive list of over 100 medical image segmentation datasets, detailing information such as modality, segmentation targets, number of cases, and download links. Participants were also encouraged to contribute by adding new public datasets to the list. To ensure a fair comparison, participants were restricted to using only the officially preprocessed data and datasets included in the curated list for model development.

Since the majority of the public medical image segmentation datasets have been included in this competition, we opted to create a completely new testing set rather than annotating existing datasets to mitigate the potential risk of data leakage. To ensure the dataset was both diverse and representative, we launched a "Call for Dataset" initiative, inviting researchers from around the globe to contribute de-identified medical images in accordance with predefined criteria and licensing requirements (Methods). As a result, we successfully collected 124,004 image-mask pairs from 4,414 cases (Fig. 2c, Supplementary Table 3). While the testing set spans a wide range of imaging modalities, CT and MRI images continue to dominate, reflecting their prevalence in clinical diagnostics.

Fig. 2d illustrates the geographical distribution of participants and testing data sources in the competition. The blue circles represent the locations of participant teams, with their size corresponding to the relative number of teams in each region. The challenge attracted participation from over 200 individuals and 61 data contributors across 24 institutions, which highlights the international interest in advancing medical image segmentation and demonstrates the competition's success in engaging a wide array of contributors from different regions around the world.




Fig. 2. **Competition dataset.** **a**, Example images and the annotations in the training and testing set. **b**, The number of image-mask pairs in the public training set. **c**, The number of image-mask pairs in the hidden testing set. All testing data has been newly collected specifically for this challenge and was not publicly available beforehand. **d**, Geographical distribution of participants and data contributors in the testing set.

Algorithm Overview

We first developed a baseline model, LiteMedSAM, by distilling the large vision transformer encoder [19] in MedSAM to a TinyViT [20] (Methods). The code and model were publicly available at the beginning of the competition to reduce the entrance barrier for participants. In addition to the baseline model, we received 22 algorithm submissions during the testing phase where 20 algorithms were from the top 20 teams on the validation leaderboard.

Most of the top-performing teams employed a SAM-like promptable segmentation model, including an image encoder to extract features from input images, a prompt encoder to process the bounding box coordinates, and a mask decoder to generate segmentation masks. The key to standing out is to design lightweight image encoders and implement fast inference strategies. We have summarized the key components of the top five algorithms, focusing on their network architectures and efficient inference strategies, in Supplementary Table 4. In this section, we specifically highlight the top three best-performing algorithms.

Best-performing algorithm. Le et al. (T1-seno [21]) proposed MedficientSAM, which used the EfficientViT-L1 model as the image encoder and conducted knowledge distillation from the MedSAM encoder [9] on the training set. Then, the whole model was fine-tuned in an end-to-end way where the prompt encoder and the mask decoder were initialized with MedSAM’s pretrained weights. For faster inference, the Python-based pipeline was ported to C++, and the trained model was exported to OpenVINO format. Moreover, an embedding caching mechanism was implemented for 3D images, which only needed to compute image embeddings once for multiple box prompts. Furthermore, all the modules were compiled from source code to take advantage of runtime optimizations such as Advanced Vector Extensions and Link Time Optimization.

Second-best-performing algorithm. Pfefferle et al. (T2-automlfreiburg [22]) introduced a data-aware fine-tuning framework to efficiently produce tailored models for specific data or modality types. The image encoder was the lightweight EfficientViT-L0 model with a three-stage fine-tuning pipeline: knowledge distillation, general fine-tuning on the whole training set, and data-aware fine-tuning on a subset. In addition to the above OpenVINO and embedding caching

mechanism for inference optimization, model caching and slim packages (e.g., Python 3.11 and OpenCV headless version) were used to speed up loading models and dockers, respectively.

Third-best-performing algorithm. Wei et al. (T3-skippinglegday [23]) proposed RepMedSAM, which replaced MedSAM’s image encoder with a pure CNN and lightweight RepVit-M model [24]. The prompt encoder and mask decoder were from MedSAM, while the image encoder was trained by knowledge distillation. No specific runtime optimization was implemented during inference because the designed model already had relatively better latency compared to the baseline model.

Performance analysis on the hidden testing set

All the submitted algorithms were independently executed on the same computing platform (Methods) for a fair comparison. Fig 3a and b show the average DSC and NSD scores across the 2D and 3D datasets, respectively. The 2D segmentation results reveal generally high average DSC and NSD scores for most algorithms, with several teams, such as seno and hmi306, achieving median scores over 0.9. Those teams also have relatively narrow interquartile ranges (IQRs), suggesting that their algorithms demonstrate both high accuracy and consistent performance across different cases.

The overall range of scores for 3D segmentation appears broader than in the 2D case, reflecting the added complexity of volumetric segmentation tasks. 3D segmentation results also exhibit more variability with substantially lower median scores compared to 2D results. While teams such as skippinglegday and waterlooviplab maintain strong performance with high average DSC and NSD scores, many teams display wider IQRs and longer whiskers, indicating greater inconsistency in their segmentation results.

In addition to segmentation accuracy, the efficiency of the segmentation algorithm is another crucial aspect, as it substantially impacts the overall user experience. The bubble plots in Fig 3c and d present the trade-off between segmentation accuracy and efficiency, measured as runtime in seconds, for the 2D and 3D datasets, respectively. The size of the circles is proportional to the NSD score. For 2D segmentation, most algorithms achieve high DSC scores (above 0.8) while maintaining a runtime of less than 10 seconds per image. Notably, teams seno and automlfburg achieve exceptional accuracy and efficiency, with a NSD near 0.9 and a runtime less than 2 seconds. In contrast, 3D segmentation results show a wider spread in runtimes, ranging from approximately 10 seconds to over 70 seconds per case, reflecting the greater computational complexity of volumetric data processing. Nevertheless, teams seno and automlfburg still obtain the fastest inference with competitive DSC and NSD scores, demonstrating a great trade-off between segmentation accuracy and efficiency.

Next, we conducted a fine-grained analysis of the three overall best-performing teams and baseline across the nine modalities (Fig 3e, Supplementary Table 5-9). For the 3D modalities CT, MRI, and PET, T2-automlfburg and T3-skippinglegday achieve comparable performances, which are significantly better than T1 ($p < 0.01$). Across most 2D modalities, T1-seno consistently outperforms the others, achieving the highest accuracy with relatively narrow interquartile ranges, reflecting its robustness and consistency. T2-automlfburg, while competitive on endoscopy, fundus, and microscopy, shows inferior performance on the other modalities, particularly in ultrasound and X-ray. T3-skippinglegday approaches the performance of T1-seno for most modalities, especially excelling in microscopy. The LiteMedSAM baseline, while competitive for some modalities, such as PET and X-ray, lags substantially behind for most modalities, with lower scores and wider performance ranges.

Furthermore, we compare the runtime of the four algorithms across 3D and 2D modalities. For 3D modalities (Fig 3f), T1-seno and T2-automlfburg demonstrate consistently lower runtimes with minimal variance than other algorithms, indicating their efficiency and stability across all 3D images. In contrast, the baseline exhibits higher runtimes and substantial variability, particularly for CT, with runtimes exceeding 300 seconds in large CT scans. Similar trends are observed across 2D modalities (Fig 3g). Notably, T1-seno consistently consumes around one to two seconds for all 2D images, which is five to ten times faster than the baseline model.

Finally, we visualized some segmentation examples of the three best-performing algorithms for each modality (Fig 3h, Supplementary Fig. 1). The top algorithms demonstrate superior versatility, achieving accurate results for most of the images even if the object boundaries are not clear. However, certain challenging cases highlight areas where further refinement is needed. For instance, in 3D modalities like CT and MR, the segmentation of heterogeneous lesions, characterized by irregular shapes and varying intensities, leads to some over- or under-segmentations. Additionally, unseen or less common targets, such as teeth segmentation in X-ray images, pose unique challenges where the algorithms struggle to generalize effectively, resulting in suboptimal predictions. These challenges could be addressed by enriching the training dataset through community-driven efforts to include more diverse and rare cases. Moreover, integrating interactive refinement mechanisms, which allow users to provide real-time corrections, could also help improve segmentation results.

Post-challenge analysis

We found that none of the algorithms consistently achieved the best performance across all modalities, indicating that these algorithms could be complementary. This motivated us to design a post-challenge to invite participants to enhance their algorithms by using the strategies from top solutions. Although all the top-three algorithms only used the provided training set, using external public datasets from the predefined dataset list was allowed and encouraged. To mitigate the effects of training datasets, we reached out to the teams with the highest DSC and NSD scores in each modality and




Fig. 3. **Evaluation results of 23 algorithms on the hidden testing set.** Dot and box plot of the Average DSC and NSD scores on the **a**, 2D ($n = 2,309$ images) and **b**, 3D ($n = 2,105$ scans) testing set. The box plots display descriptive statistics across all testing cases, with the median value represented by the horizontal line within the box, the lower and upper quartiles delineating the borders of the box and the vertical black lines indicating $1.5 \times \text{IQR}$. The algorithms are organized on the x-axis based on their corresponding ranks. The bubble plots show the trade-off between segmentation accuracy (DSC and NSD) and efficiency on the **c**, 2D and **d**, 3D testing set. The circle size is proportional to the NSD score. **e**, Modality-wise segmentation accuracy performance (Average DSC and NSD) of three best-performing teams and LiteMedSAM baseline. Modality-wise segmentation efficiency (runtime) of three best-performing teams and LiteMedSAM baseline on 3D **f**, and 2D **g**, modalities. **h**, Visualized segmentation results of nine modalities. The blue bounding box and green overlay denote prompt and reference standard, respectively. For each image, the best DSC score from the three algorithms and the corresponding contour are presented.




Fig. 4. **Results of the post-challenge.** **a**, The new best-performing algorithm (T2-automlfreiburg-post) and comparison to its predecessor (T2-automlfreiburg) and the previous best-performing algorithm (T1-seno) across all modalities. **b**, Runtime comparison. **c**, T1-reproduce successfully replicated the performance of T1-seno across all modalities

identified the employed external datasets. Only one X-Ray [25] and one PET dataset [26] were used as additional datasets by the modality-wise winning teams. Thus, we also added those two datasets to the training set.

We received six submissions for the performance-booster task. One algorithm (T2-automlfreiburg-post) [22] surpassed the previous best-performing algorithm (T1-seno) and its previous algorithm (T2-automlfreiburg) in terms of the overall rank. The main technical improvements include using a shared EfficientViT model for all 3D modalities, early stopping to avoid over-fitting, and speeding up the inference pipeline with C++ implementations from T1-seno [21]. Fig 4a shows the modality-wise performance of the three algorithms. T2-automlfreiburg-post substantially improves the previous algorithm on PET, ultrasound, X-ray, and microscopy modalities while preserving the performance for the other modalities. Importantly, the runtime is reduced by 2x, matching the efficiency of the fastest algorithm T1-seno (Fig 4b).

In addition, we received one submission (T1-reproduce) for the reproducibility analysis task, aiming to achieve similar performance to T1-seno using the released code [27]. The results, as shown in Fig 4c, indicate that T1-reproduce successfully replicated the performance of T1-seno across all modalities, achieving nearly identical average DSC and NSD scores, although the training batch size was smaller than the original setting because of the computing resource limitation.

DISCUSSION

This work presents the first international competition dedicated to promptable medical image segmentation foundation models. One of the key achievements of the competition was the curation of a large-scale, unpublished dataset for benchmarking, consisting of diverse imaging modalities and newly annotated cases. This dataset not only mitigated the risk of data leakage from existing public sources but also provided a robust and comprehensive foundation for evaluating segmentation algorithms. The challenge attracted substantial global participation, with thousands of submissions on the online leaderboard from over 200 participants, showcasing the high level of interest in the community.

This competition also introduced a novel task focused on efficient segmentation, aiming to develop deployable segmentation models that balance high accuracy with practical runtime performance without reliance on expensive GPUs. Although all the top teams employed SAM-like architectures, two key designs made their algorithms stand out: efficient network architecture and fast inference strategies. In particular, all the three best-performing algorithms replaced the heavy-weight ViT image encoder [19] in MedSAM with a lightweight image encoder, such as EfficientViT [28] and RepViT [24], which reduced the parameters by 50%. Moreover, knowledge distillation was an effective way to transfer knowledge from larger, more accurate models to compact models. Notably, using lightweight image encoders with knowledge distillation preserved segmentation accuracy, which achieved comparable to or even better accuracy than MedSAM while substantially reducing computational overhead (Supplementary Fig. 2).

In addition to lightweight image encoders, we also identified several practical strategies for fast inference on CPU, including embedding caching, the use of a C++ inference pipeline, and integration with OpenVINO combined with optimized Docker deployment. Embedding caching minimizes redundant computations by storing pre-computed 2D slice embeddings for 3D image segmentation. The C++ inference pipeline improves execution speed by leveraging a compiled language’s low-level optimizations, enabling faster runtimes compared to Python. Additionally, OpenVINO and optimized Docker deployment streamline inference by utilizing a highly efficient runtime environment and lightweight containerization, reducing resource consumption and startup latency. Together, these strategies enable faster and more resource-efficient deployments on edge devices without sacrificing segmentation accuracy.

Furthermore, our post-challenge phase also introduced two innovative tasks: the performance booster and the reproducibility analysis, aimed at advancing algorithm performance and ensuring scientific rigor. One of the new submissions set a new state-of-the-art by combining elements of existing successful solutions, demonstrating the power of collaborative advancements. The reproducibility analysis further validated the robustness of the competition’s outcomes, as the winning solution during the testing phase was successfully reproduced using the released code and documentation. This highlights the transparency and reliability of the methods developed during the challenge, ensuring that the results can serve as a solid foundation for future research.

While the competition showcased substantial algorithm advancements for universal medical image segmentation models, there remains a substantial gap between these cutting-edge algorithms and their integration into clinical practice. The complexity of deploying advanced models, combined with the need for user-friendly interfaces, often hinders their adoption in real-world settings. To bridge this gap, we collaborated with the two best-performing teams to incorporate their algorithms into 3D Slicer [29], a widely-used open-source platform for medical image analysis. This integration allows clinicians and researchers to access state-of-the-art segmentation algorithms without requiring any coding expertise, making these powerful tools more accessible and practical for routine use.

This work also has several limitations. First, while we provided one of the largest and most diverse datasets to date, the majority of curated testing images originated from North America, Europe, and Asia, with a noticeable lack of representation from South America, Africa, and Oceania. Second, all the top-performing algorithms utilized 2D models for 3D data segmentation, which overlooked the slice-wise or 3D contextual information inherent in volumetric medical images [5]. This limitation may impact segmentation performance, particularly for tasks requiring spatial consistency across slices. Finally, the prompts were based on the bounding box because of its effectiveness and less ambiguity. However, this setting did not consider more flexible interactive segmentation approaches, allowing for iterative user feedback refinement.

Future iterations of the competition will aim to address current limitations and push the boundaries further. In particular, we will expand the dataset by incorporating more images from recent public datasets [30] and under-represented regions to enhance the global generalizability and fairness of the algorithms [31]. Additionally, we will design a new task to benchmark interactive segmentation [11] and text-based segmentation methods, reflecting the evolving research directions [10], [32]. This benchmark will assess how effectively algorithms can integrate user feedback through interactive processes and respond to natural language descriptions for segmentation tasks.

In conclusion, this competition marked a pivotal step forward in advancing promptable medical image segmentation, introducing a novel task that emphasized efficiency alongside accuracy and a large and diverse testing set from 24 different institutions. The competition attracted over 200 participants worldwide, and the two best-performing algorithms, featuring a transformer-based lightweight image encoder and fast inference implementations, achieved over 10 times faster inference speed than the previous foundation model while maintaining comparable or even superior accuracy. The innovative post-challenge phase further validated the reproducibility of top-performing solutions and set a new state-of-the-art through integrated strategies, highlighting the importance of transparency and collaboration in driving progress. The two best-performing algorithms have been incorporated into 3D Slicer to bridge the gap between algorithmic innovation and clinical deployment. We believe these efforts will pave the way for more robust, accessible, and impactful medical image segmentation technologies, ultimately benefiting global healthcare practices.

Affiliations

- Jun Ma is with AI Collaborative Centre, University Health Network; Department of Laboratory Medicine and Pathobiology, University of Toronto; Vector Institute, Toronto, Canada
- Feifei Li is with Peter Munk Cardiac Centre, University Health Network, Toronto, Canada. (*Equal contribution)
- Sumin Kim is Toronto General Hospital Research Institute, University Health Network; Department of Computer Science, University of Toronto; University Health Network; Vector Institute, Toronto, Canada. (*Equal contribution)
- Reza Asakereh is with Peter Munk Cardiac Centre, University Health Network, Toronto, Canada.
- Bao-Hiep Le is with University of Science, Vietnam National University, Ho Chi Minh City, Vietnam.
- Dang-Khoa Nguyen-Vu is with University of Science, Vietnam National University, Ho Chi Minh City, Vietnam.
- Alexander Pfefferle is with Institute of Computer Science, University of Freiburg, Freiburg, Germany.
- Muxin Wei is with the School of Medicine and Health, Harbin Institute of Technology, Harbin, China.
- Ruochen Gao is with the Division of Image Processing, Department of Radiology, Leiden University Medical Center, Leiden, the Netherlands.
- Donghang Lyu is with the Division of Image Processing, Department of Radiology, Leiden University Medical Center, Leiden, the Netherlands.
- Songxiao Yang is with the Department of System and Control Engineering, School of Engineering, Institute of Science Tokyo (formerly Tokyo Institute of Technology), Tokyo, Japan.
- Lennart Purucker is with the Institute of Computer Science, University of Freiburg, Freiburg, Germany.
- Zdravko Marinov is with Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Karlsruhe, Germany.
- Marius Staring is with the Division of Image Processing, Department of Radiology, Leiden University Medical Center, Leiden, the Netherlands.
- Haisheng Lu is with the School of Information and Communication Engineering, University of Electronic Science and Technology of China, Chengdu, China.
- Thuy Thanh Dao is with the School of Electrical Engineering and Computer Science, University of Queensland, Brisbane, Australia.
- Xincheng Ye is with the School of Electrical Engineering and Computer Science, University of Queensland, Brisbane, Australia.
- Zhi Li is with the School of Cyberspace, Hangzhou Dianzi University, Hangzhou, China.
- Gianluca Brugnara is with the Division for Computational Radiology and Clinical AI and the Department of Neuroradiology, University Hospital Bonn, Germany.
- Philipp Vollmuth is with the Division for Computational Radiology and Clinical AI and the Department of Neuroradiology, University Hospital Bonn, Germany.
- Martha Foltyn-Dumitru is with the Division for Computational Radiology and Clinical AI and the Department of Neuroradiology, University Hospital Bonn, Germany.
- Jaeyoung Cho is with the Division for Computational Radiology and Clinical AI, the Department of Neuroradiology, University Hospital Bonn, Germany.
- Mustafa Ahmed Mahmutoglu is with the Department of Neuroradiology, Heidelberg University Hospital, Heidelberg, Germany.
- Martin Bendszus is with the Department of Neuroradiology, Heidelberg University Hospital, Heidelberg, Germany.
- Irada Pflüger is with the Department of Neuroradiology, Heidelberg University Hospital, Heidelberg, Germany.
- Aditya Rastogi is with the Division for Computational Radiology and Clinical AI, Department of Neuroradiology, University Hospital Bonn, Germany.

- Dong Ni is with the School of Biomedical Engineering, Shenzhen University, Shenzhen, China.
- Xin Yang is with the School of Biomedical Engineering, Shenzhen University, Shenzhen, China.
- Guang-Quan Zhou is with the School of Biological Science and Medical Engineering, Southeast University, Nanjing, China.
- Kaini Wang is with the School of Biological Science and Medical Engineering, Southeast University, Nanjing, China.
- Nicholas Heller is with the Department of Urology, Cleveland Clinic, Cleveland, United States.
- Nikolaos Papanikolopoulos is with the Department of Computer Science, University of Minnesota, Minneapolis, United States.
- Christopher Weight is with the Department of Urology, Cleveland Clinic, Cleveland, United States.
- Yubing Tong is with the Department of Radiology, University of Pennsylvania, Philadelphia, United States.
- Jayaram K Udupa is with the Department of Radiology, University of Pennsylvania, Philadelphia, United States.
- Cahill J. Patrick is with the Department of Orthopaedic Surgery, Children's Hospital of Philadelphia, Philadelphia, United States.
- Yaqi Wang is with the College of Media Engineering, Communication University of Zhejiang, Hangzhou, China.
- Yifan Zhang is with Hangzhou Dental Hospital Group, Hangzhou, China.
- Francisco Contijoch is with Bioengineering, Radiology, UC San Diego, La Jolla, California, United States.
- Elliot McVeigh is with Bioengineering, Radiology, and Cardiology, UC San Diego, La Jolla, California, United States.
- Xin Ye is with the Department of Ophthalmology, Hangzhou, China.
- Shucheng He is with the Department of Ophthalmology, Hangzhou, China.
- Robert Haase is with the Department of Neuroradiology, University Hospital Bonn, Bonn, Germany.
- Thomas Pinetz is with the Institute for Applied Mathematics, University of Bonn, Bonn, Germany.
- Alexander Radbruch is with the Department of Neuroradiology, University Hospital Bonn, Bonn, Germany.
- Inga Krause is with the Department of Neuroradiology, University Hospital Bonn, Bonn, Germany.
- Erich Kobler is with the Department of Neuroradiology, University Hospital Bonn, Bonn, Germany.
- Jian He is with the Department of Nuclear Medicine, Nanjing Drum Tower Hospital, Nanjing, China.
- Yucheng Tang is with the Healthcare and Life Science, NVIDIA, Redmond, United States.
- Haichun Yang is with the Department of Pathology, Microbiology and Immunology, Vanderbilt University Medical Center, Nashville, United States.
- Yuankai Huo is with the Department of Computer Science, Vanderbilt University, Nashville, United States.
- Gongning Luo is with the Department of Computer Science, King Abdullah University of Science and Technology, Thuwal, Kingdom of Saudi Arabia.
- Kaisar Kushibar is with the Faculty of Mathematics and Computer Science, University of Barcelona, Barcelona, Spain.
- Jandos Amankulov is with the Radiology and Nuclear Medicine, Kazakh Institute of Oncology and Radiology, Almaty, Kazakhstan.
- Dias Toleshbayev is with the Radiology and Nuclear Medicine, Kazakh Institute of Oncology and Radiology, Almaty, Kazakhstan.
- Amangeldi Mukhamejan is with the Radiology and Nuclear Medicine, Kazakh Institute of Oncology and Radiology, Almaty, Kazakhstan.
- Jan Egger is with the Institute for AI in Medicine, University Hospital Essen, Essen, Germany.
- Antonio Pepe is with the Institute of Computer Graphics and Vision, Graz University of Technology, Graz, Austria.
- Christina Gsaxner is with the Institute of Computer Graphics and Vision, Graz University of Technology, Graz, Austria.
- Gijs Luijten is with the Institute for AI in Medicine, University Hospital Essen, Essen, Germany.
- Shohei Fujita is with the Department of Radiology, the University of Tokyo, Tokyo, Japan.
- Tomohiro Kikuchi is with the Data Science Center, Jichi Medical University, Tochigi, Japan.
- Benedikt Wiestler is with the AI for Image-Guided Diagnosis and Therapy, Technical University of Munich, Munich, Germany.
- Jan S. Kirschke is with the Department of Diagnostic and Interventional Neuroradiology, Technical University of Munich, Munich, Germany.
- Ezequiel de la Rosa is with the Department of Quantitative Biomedicine, University of Zurich, Zurich, Switzerland.
- Federico Bolelli is with the Department of Engineering, University of Modena and Reggio Emilia, Modena, Italy.
- Luca Lumetti is with the Department of Engineering, University of Modena and Reggio Emilia, Modena, Italy.
- Costantino Grana is with the Department of Engineering, University of Modena and Reggio Emilia, Modena, Italy.
- Kunpeng Xie is with the Department of Oral and Maxillofacial Surgery and Institute of Medical Informatics, University Hospital RWTH Aachen, Aachen, Germany.
- Guomin Wu is with the Department of Oral, Plastic, and Aesthetic Surgery, School and Hospital of Stomatology, Jilin University, Changchun, China.
- Behrus Puladi is with the Department of Oral and Maxillofacial Surgery and Institute of Medical Informatics, University Hospital RWTH Aachen, Aachen, Germany.
- Carlos Martín-Isla is with the Department of Mathematics and Computer Science, University of Barcelona, Barcelona,

Spain.

- Karim Lekadir is with the Catalan Institution for Research and Advanced Studies, University of Barcelona, Barcelona, Spain.
- Victor M. Campello is with the Department of Mathematics and Computer Science, University of Barcelona, Barcelona, Spain.
- Wei Shao is with the Department of Medicine, University of Florida, Gainesville, United States.
- Wayne Brisbane is with the Department of Urology, University of California, Los Angeles, Los Angeles, United States.
- Hongxu Jiang is with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, United States.
- Hao Wei is with the Department of Biomedical Engineering, the Chinese University of Hong Kong, Hong Kong SAR, China.
- Wu Yuan is with the Department of Biomedical Engineering, the Chinese University of Hong Kong, Hong Kong SAR, China.
- Shuangli Li is with the Department of Ophthalmology, Zigong First People’s Hospital, Zigong, China.
- Yuyin Zhou is with the Department of Computer Science and Engineering, University of California, Santa Cruz, USA
- Bo Wang (Corresponding Author) is with Peter Munk Cardiac Centre, University Health Network; Department of Laboratory Medicine and Pathobiology and Department of Computer Science, University of Toronto; Vector Institute, Toronto, Canada. E-mail: bowang@vectorinstitute.ai

METHODS

Dataset curation and pre-processing

All the training images were from publicly available datasets with a license for re-distribution (Supplementary Table 1-2). The original images have a wide range of format, such as nifti, dicom, mhd, nrrd, jpg, and png. We normalized them to the same npz format with image and reference standard inside it, allowing participants to get rid of tedious data cleaning and focus on model development.

The pre-processing followed common practice [9], [33]. Specifically, for CT images, we first adjusted the intensity to the proper window level and width followed by rescaling to $[0, 255]$. For MRI and PET images, we applied intensity cut-off with the lower-bound and upper-bound of 0.5% and 99.5% percentile of foreground intensity and then rescaled the intensity to $[0, 255]$. For the 2D modalities, we applied the same intensity preprocessing as MRI if their intensity ranges are not in $[0, 255]$. Otherwise, no preprocessing was applied. For the reference standard, we converted the lesion semantic mask to instance mask, allowing the generation of lesion-wise bounding box prompts.

Baseline: LiteMedSAM

We provided an out-of-the-box baseline model, LiteMedSAM, with a detailed tutorial to reduce the entry barriers. The development of LiteMedSAM involves a two-stage process: distillation and fine-tuning. In the first stage, we distill the extensive medical imaging knowledge from MedSAM’s heavyweight ViT-B [19] image encoder into Tiny-ViT [20], a compact hybrid architecture combining Transformer and convolution layers. In the second stage, we fine-tune the entire LiteMedSAM model.

To prepare the input images, for MedSAM’s ViT-b encoder, the images are first resized to 1024×1024 using bi-cubic interpolation, while for Tiny-ViT, the longest edge of the input images is first resized to 256 using bilinear interpolation, and then the images are padded to 256×256 with zeros on the right and bottom to maintain a consistent size. After resizing, min-max normalization is applied to scale the intensity values from the range of $[0, 255]$ to $[0, 1]$. Data augmentation is then performed, which includes random horizontal and vertical flips with a probability of 0.5. These preprocessing and augmentation steps are consistent across both the distillation and fine-tuning stages.

During the distillation process, MedSAM’s image encoder serves as the teacher model, while Tiny-ViT acts as the student model. The teacher model remains frozen throughout the distillation. The output feature embeddings of both image encoders have a dimensionality of $64 \times 64 \times 256$. We adopt Mean Squared Error (MSE) loss to encourage the feature embeddings of the student to match those of the teacher. The distillation is performed using the AdamW optimizer [34] with a learning rate of $5e-5$, a weight decay of 0.01, and a batch size of 8. The process continues until either 1000 epochs are reached or the loss no longer decreases.

In the second stage, we plug in the distilled Tiny-ViT as the image encoder and fine-tune the entire LiteMedSAM model, including the prompt encoder and mask decoder. During the fine-tuning process, a loss function comprised by the unweighted sum of the Dice loss and the binary cross-entropy (BCE) loss is adopted. The same AdamW optimizer settings as in the distillation stage are used. For fine-tuning, the input images are resized to 256×256 using the same image processing steps as in the distillation stage for Tiny-ViT. The fine-tuning process stops after 1000 epochs or when the training loss falls below 0.005, whichever occurs first. Both the distillation and fine-tuning of LiteMedSAM are performed on the same dataset used to train MedSAM.

Deployment: 3D Slicer plugin

Despite the open sourcing of these advanced algorithms, a significant challenge remains in integrating them seamlessly into clinical workflows, as this often demands basic coding skills. To address this, we incorporated the two best-performing algorithms as a plugin for 3D Slicer, an open-source software platform for medical image analysis and three-dimensional visualization. The plugin was implemented using the “loadable script module” type, allowing for the easy integration of Python scripts and seamless interaction with the Slicer framework.

The module provides a user-friendly graphical interface where users can specify preprocessing options and define the region of interest (ROI) for segmentation. Its architecture is based on a generic abstract class to facilitate the easy integration of new segmentation models. Importantly, all these models can be executed on laptops without reliance on GPUs.

Evaluation metrics and platform

The evaluation metrics contained two segmentation accuracy metrics and one efficiency metric. We followed the recommendations in Metrics Reloaded [35] to evaluate the segmentation accuracy. Specifically, we used Dice Similarity Coefficient (DSC) and Normalized Surface Distance (NSD) to quantitatively evaluate the region overlap and boundary similarity, respectively. The efficiency was measured by runtime (seconds) and we recorded the Docker container execution time for each case. The submitted Docker algorithms were evaluated on the same platform with Ubuntu 20.04 system, which contains one Intel CPU (Xeon W-2133 3.60GHz) and an upperbound of 8GB RAM. Docker version 20.10.13 The evaluation code and platform information were released at the beginning of the competition for a transparent evaluation.

Ranking scheme and statistical analysis

We employed the commonly used rank-then-aggregate scheme to obtain the final rank [16], [36], which contains three steps. First, we computed the three metrics (DSC, NSD, and runtime) for each testing case and each algorithm. Then, we ranked the 23 algorithms teams for each modality and each metric. Finally, we averaged all the rankings to obtain the final rank. Wilcoxon signed-rank test was used to statistically compare the performance of the algorithms. We also used the bootstrapping (N=1000) to analyze the ranking stability. The analysis was based on the widely used toolkit ChallengeR [37].

Data availability

All the datasets in this study have been publicly available on the codabench [18] competition website (Dataset page) <https://www.codabench.org/competitions/1847/>.

Code availability

The code and trained models are publicly available at <https://github.com/bowang-lab/MedSAM/tree/LiteMedSAM>. 3D slicer plugin can be accessed at <https://github.com/bowang-lab/MedSAMslicer>. The code and model weights of best-performing teams are available at the competition website (Ranking page) <https://www.codabench.org/competitions/1847/>.

Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council of Canada (RGPIN-2020-06189 and DGECR-2020-00294) and CIFAR AI Chair programs. This research was enabled, in part, by computing resources provided by the Digital Research Alliance of Canada. B.L. and D.N. were supported by Vingroup Innovation Foundation (VINIF) in project code VINIF2019.DA19. S.Y. was supported by JST SPRING, Japan Grant Number JPMJSP2106. L.P. was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 499552394 – SFB 1597. L.Z. was supported by Zhejiang Provincial Natural Science Foundation of China under Grant No. LDT23F01015F01. M.A.M. was supported by Else Kröner Research College for young physicians (ref. number: 2023_EKFK.02) N.H. was supported by DoD HT9425-23-1-0918. Y.T., J.K.U., and C.J.P. were supported by NIH R01HL150147. F.C. was supported by NIH K01HL143113. E.M. was supported by NIH R01HL144678. E.K. was supported by DFG KO 7162/1-1 543939932 and FWF 10.55776/COE12. H.Y. was supported by DOD HT9425-23-1-0003. K.K. was supported by Juan de la Cierva fellowship by the Ministry of Science and Innovation of Spain with reference number FJC2021-047659-I. F.B., L.L., and C.G. were supported by the University of Modena and Reggio Emilia and Fondazione di Modena, through the FAR 2024 and FARD-2024 funds (Fondo di Ateneo per la Ricerca). We also appreciate all the data owners for providing public medical images to the community.

REFERENCES

- [1] B. He, A. C. Kwan, J. H. Cho, N. Yuan, C. Pollick, T. Shiota, J. Ebinger, N. A. Bello, J. Wei, K. Josan, G. Duffy, M. Jujavarapu, R. Siegel, S. Cheng, J. Y. Zou, and D. Ouyang, "Blinded, randomized trial of sonographer versus AI cardiac function assessment," *Nature*, vol. 616, no. 7957, pp. 520–524, 2023.
- [2] K. Cao, Y. Xia, J. Yao, X. Han, L. Lambert, T. Zhang, W. Tang, G. Jin, H. Jiang, X. Fang *et al.*, "Large-scale pancreatic cancer detection via non-contrast ct and deep learning," *Nature medicine*, vol. 29, no. 12, pp. 3033–3043, 2023.
- [3] S. Minaee, Y. Boykov, F. Porikli, A. Plaza, N. Kehtarnavaz, and D. Terzopoulos, "Image segmentation using deep learning: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, pp. 3523–3542, 2021.
- [4] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, "Segment anything," *arXiv preprint arXiv:2304.02643*, 2023.
- [5] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár, and C. Feichtenhofer, "Sam 2: Segment anything in images and videos," *arXiv preprint arXiv:2408.00714*, 2024. [Online]. Available: <https://arxiv.org/abs/2408.00714>
- [6] M. A. Mazurowski, H. Dong, H. Gu, J. Yang, N. Konz, and Y. Zhang, "Segment anything model for medical image analysis: an experimental study," *Medical Image Analysis*, vol. 89, p. 102918, 2023.
- [7] Y. Huang, X. Yang, L. Liu, H. Zhou, A. Chang, X. Zhou, R. Chen, J. Yu, J. Chen, C. Chen *et al.*, "Segment anything model for medical images?" *Medical Image Analysis*, vol. 92, p. 103061, 2024.
- [8] X. Zou, J. Yang, H. Zhang, F. Li, L. Li, J. Gao, and Y. J. Lee, "Segment everything everywhere all at once," *arXiv preprint arXiv:2304.06718*, 2023.
- [9] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang, "Segment anything in medical images," *Nature Communications*, vol. 15, p. 654, 2024.
- [10] T. Zhao, Y. Gu, J. Yang, N. Usuyama, H. H. Lee, S. Kiblawi, T. Naumann, J. Gao, A. Crabtree, J. Abel, C. Moung-Wen, B. Piening, C. Bifulco, M. Wei, H. Poon, and S. Wang, "A foundation model for joint segmentation, detection and recognition of biomedical objects across nine modalities," *Nature Methods*, 2024.
- [11] Z. Marinov, P. F. Jäger, J. Egger, J. Kleesiek, and R. Stiefelhagen, "Deep interactive segmentation of medical images: A systematic review and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 12, pp. 10998–11018, 2024.
- [12] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems*, vol. 25, 2012.
- [14] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko *et al.*, "Highly accurate protein structure prediction with alphafold," *nature*, vol. 596, no. 7873, pp. 583–589, 2021.
- [15] W. Bulten, K. Kartasalo, P.-H. C. Chen, P. Ström, H. Pinckaers, K. Nagpal, Y. Cai, D. F. Steiner, H. van Boven, R. Vink *et al.*, "Artificial intelligence for diagnosis and gleason grading of prostate cancer: the panda challenge," *Nature Medicine*, vol. 28, no. 1, pp. 154–163, 2022.
- [16] J. Ma, Y. Zhang, S. Gu, C. Ge, S. Mae, A. Young, C. Zhu, X. Yang, K. Meng, Z. Huang, F. Zhang, Y. Pan, S. Huang, J. Wang, M. Sun, R. Zhang, D. Jia, J. W. Choi, N. Alves, B. de Wilde, G. Koehler, H. Lai, E. Wang, M. Wiesenfarth, Q. Zhu, G. Dong, J. He, J. He, H. Yang, B. Huang, M. Lyu, Y. Ma, H. Guo, W. Xu, K. Maier-Hein, Y. Wu, and B. Wang, "Unleashing the strengths of unlabelled data in deep learning-assisted pan-cancer abdominal organ quantification: the flare22 challenge," *The Lancet Digital Health*, vol. 6, no. 11, p. e815–e826, 2024.
- [17] P. R. A. S. Bassi, W. Li, Y. Tang, F. Isensee, Z. Wang, J. Chen, Y.-C. Chou, Y. Kirchhoff, M. R. Rokuss, Z. Huang, J. Ye, J. He, T. Wald, C. Ulrich, M. Baumgartner, S. Roy, K. Maier-Hein, P. F. Jaeger, Y. Ye, Y. Xie, J. Zhang, Z. Chen, Y. Xia, Z. Xing, L. Zhu, Y. Sadegheih, A. Bozorgpour, P. Kumari, R. Azad, D. Merhof, P. Shi, T. Ma, Y. Du, F. BAI, T. Huang, B. Zhao, H. Wang, X. Li, H. Gu, H. Dong, J. Yang, M. A. Mazurowski, S. Gupta, L. Wu, J.-X. Zhuang, H. Chen, H. R. Roth, D. Xu, M. B. Blaschko, S. Decherchi, A. Cavalli, A. Yuille, and Z. Zhou, "Touchstone benchmark: Are we on the right way for evaluating AI algorithms for medical segmentation?" in *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.
- [18] Z. Xu, S. Escalera, A. Pavão, M. Richard, W.-W. Tu, Q. Yao, H. Zhao, and I. Guyon, "Codabench: Flexible, easy-to-use, and reproducible meta-benchmark platform," *Patterns*, vol. 3, no. 7, p. 100543, 2022.
- [19] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2020.
- [20] K. Wu, J. Zhang, H. Peng, M. Liu, B. Xiao, J. Fu, and L. Yuan, "Tinyvit: Fast pretraining distillation for small vision transformers," in *European Conference on Computer Vision*, 2022, pp. 68–85.
- [21] B.-H. Le, D.-K. Nguyen-Vu, T.-H. Nguyen-Mau, H.-D. Nguyen, and M.-T. Tran, "Medficientsam: A robust medical segmentation model with optimized inference pipeline for limited clinical settings," in *CVPR 2024: Segment Anything In Medical Images On Laptop*, 2024.
- [22] A. Pfefferle, L. Purucker, and F. Hutter, "Daft: Data-aware fine-tuning of foundation models for efficient and effective medical image segmentation," in *CVPR 2024: Segment Anything In Medical Images On Laptop*, 2024.
- [23] M. Wei, S. Chen, S. Wu, and D. Xu, "Rep-medsam: Towards real-time and universal medical image segmentation," in *CVPR 2024: Segment Anything In Medical Images On Laptop*, 2024.
- [24] A. Wang, H. Chen, Z. Lin, J. Han, and G. Ding, "Repvit: Revisiting mobile cnn from vit perspective," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15909–15920.
- [25] Y. Zhang, F. Ye, L. Chen, F. Xu, X. Chen, H. Wu, M. Cao, Y. Li, Y. Wang, and X. Huang, "Children's dental panoramic radiographs dataset for caries segmentation and dental disease detection," *Scientific Data*, vol. 10, no. 1, p. 380, 2023.
- [26] S. Gatidis, M. Früh, M. P. Fabritius, S. Gu, K. Nikolaou, C. L. Fougère, J. Ye, J. He, Y. Peng, L. Bi, J. Ma, B. Wang, J. Zhang, Y. Huang, L. Heiliger, Z. Marinov, R. Stiefelhagen, J. Egger, J. Kleesiek, L. Sibille, L. Xiang, S. Bendazzoli, M. Astaraki, M. Ingrisch, C. C. Cyran, and T. Küstner, "Results from the autopet challenge on fully automated lesion segmentation in oncologic pet/ct imaging," *Nature Machine Intelligence*, vol. 6, no. 11, p. 1396–1405, 2024.
- [27] Z. Li, Y. WANG, and S. Wang, "Reproducibility analysis: Reproduce the top one team results," in *CVPR 2024: Segment Anything In Medical Images On Laptop*, 2024.
- [28] H. Cai, J. Li, M. Hu, C. Gan, and S. Han, "Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 17302–17313.
- [29] R. Kikinis, S. D. Pieper, and K. G. Vosburgh, "3d slicer: a platform for subject-specific image analysis, visualization, and clinical support," in *Intraoperative imaging and image-guided therapy*. Springer, 2013, pp. 277–289.
- [30] W. Li, C. Qu, X. Chen, P. R. Bassi, Y. Shi, Y. Lai, Q. Yu, H. Xue, Y. Chen, X. Lin, Y. Tang, Y. Cao, H. Han, Z. Zhang, J. Liu, T. Zhang, Y. Ma, J. Wang, G. Zhang, A. Yuille, and Z. Zhou, "Abdomenatlas: A large-scale, detailed-annotated, & multi-center dataset for efficient transfer learning and open algorithmic benchmarking," *Medical Image Analysis*, vol. 97, p. 103285, 2024.
- [31] C. Shi, R. Rezai, J. Yang, Q. Dou, and X. Li, "A survey on trustworthiness in foundation models for medical image analysis," *arXiv preprint arXiv:2407.15851*, 2024.
- [32] Z. Zhao, Y. Zhang, C. Wu, X. Zhang, Y. Zhang, Y. Wang, and W. Xie, "One model to rule them all: Towards universal segmentation for medical images with text prompts," *arXiv preprint arXiv:2312.17183*, 2023.

- [33] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, “nnu-net: a self-configuring method for deep learning-based biomedical image segmentation,” *Nature Methods*, vol. 18, no. 2, pp. 203–211, 2021.
- [34] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*, 2019.
- [35] L. Maier-Hein, A. Reinke, P. Godau, M. D. Tizabi, F. Buettner, E. Christodoulou, B. Glocker, F. Isensee, J. Kleesiek, M. Kozubek, M. Reyes, M. A. Riegler, M. Wiesenfarth, A. E. Kavur, C. H. Sudre, M. Baumgartner, M. Eisenmann, D. Heckmann-Nötzel, T. Radsch, L. Acion, M. Antonelli, T. Arbel, S. Bakas, A. Benis, M. B. Blaschko, M. J. Cardoso, V. Cheplygina, B. A. Cimini, G. S. Collins, K. Farahani, L. Ferrer, A. Galdan, B. van Ginneken, R. Haase, D. A. Hashimoto, M. M. Hoffman, M. Huisman, P. Jannin, C. E. Kahn, D. Kainmueller, B. Kainz, A. Karargyris, A. Karthikesalingam, F. Kofler, A. Kopp-Schneider, A. Kreshuk, T. Kurc, B. A. Landman, G. Litjens, A. Madani, K. Maier-Hein, A. L. Martel, P. Mattson, E. Meijering, B. Menze, K. G. M. Moons, H. Müller, B. Nichyporuk, F. Nickel, J. Petersen, N. Rajpoot, N. Rieke, J. Saez-Rodriguez, C. I. Sánchez, S. Shetty, M. van Smeden, R. M. Summers, A. A. Taha, A. Tiulpin, S. A. Tsiftaris, B. Van Calster, G. Varoquaux, and P. F. Jäger, “Metrics reloaded: recommendations for image analysis validation,” *Nature Methods*, vol. 21, no. 2, p. 195–212, 2024.
- [36] S. Bakas, M. Reyes, A. Jakab, S. Bauer, M. Rempfler, A. Crimi *et al.*, “Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge,” *arXiv preprint arXiv:1811.02629*, 2018.
- [37] M. Wiesenfarth, A. Reinke, B. A. Landman, M. Eisenmann, L. A. Saiz, M. J. Cardoso *et al.*, “Methods and open-source toolkit for analyzing and visualizing challenge results,” *Scientific Reports*, vol. 11, no. 1, pp. 1–15, 2021.