

WSI-SAM: Multi-resolution Segment Anything Model (SAM) for histopathology whole-slide images

Hong Liu

*Eindhoven University of Technology
Eindhoven, The Netherlands*

H.LIU2@TUE.NL

Haosen Yang

*University of Surrey
London, UK*

HAOSEN.YANG.6@GMAIL.COM

Paul J. van Diest

*University Utrecht
Utrecht, The Netherlands*

PDIEST@UMCUTRECHT.NL

Josien P.W. Pluim

*Eindhoven University of Technology
Eindhoven, The Netherlands*

J.PLUIM@TUE.NL

Mitko Veta

*Eindhoven University of Technology
Eindhoven, The Netherlands*

M.VETA@TUE.NL

Editor: My editor

Abstract

The Segment Anything Model (SAM) marks a significant advancement in segmentation models, offering robust zero-shot abilities and dynamic prompting. However, existing medical SAMs are not suitable for the multi-scale nature of whole-slide images (WSIs), restricting their effectiveness. To resolve this drawback, we present WSI-SAM, enhancing SAM with precise object segmentation capabilities for histopathology images using multi-resolution patches, while preserving its efficient, prompt-driven design, and zero-shot abilities. To fully exploit pretrained knowledge while minimizing training overhead, we keep SAM frozen, introducing only minimal extra parameters and computational overhead. In particular, we introduce High-Resolution (HR) token, Low-Resolution (LR) token and dual mask decoder. This decoder integrates the original SAM mask decoder with a lightweight fusion module that integrates features at multiple scales. Instead of predicting a mask independently, we integrate HR and LR token at intermediate layer to jointly learn features of the same object across multiple resolutions. Experiments show that our WSI-SAM outperforms state-of-the-art SAM and its variants. In particular, our model outperforms SAM by 4.1 and 2.5 percent points on a ductal carcinoma in situ (DCIS) segmentation tasks and breast cancer metastasis segmentation task (CAMELYON16 dataset). The code will be available at <https://github.com/HongLiuuuuu/WSI-SAM>.

Keywords: Foundation Models, Computational pathology, Whole-slide images.

1 Introduction

Segmentation is crucial in histopathology images, enabling pathologists to analyze tissue regions such as tumor and stroma and in turn use those results for a number of tasks such as disease diagnosis, treatment planning, and monitoring progression. Current deep learning-based models Feng et al. (2021a); Wang et al. (2021); Xu et al. (2019); Feng et al. (2021b); Ni et al. (2019); Qaiser et al. (2019); Han et al. (2022) have shown great promise in histopathology image segmentation, which can significantly reduce time, labor, and expertise required, and enable large-scale dataset analysis. However, these models are typically designed and trained for a specific segmentation task, which greatly limits the generalization to wider applications in clinical practice. Therefore, it is essential to develop universal models with zero-shot ability that can be trained once and then applied to a wide range of histopathology segmentation tasks.

Recently, the segment anything model (SAM) Kirillov et al. (2023) was released as a segmentation foundation model for natural images, showcasing remarkable zero-shot abilities across various scenarios. Enabling its application across a wide range of tasks Mazurowski et al. (2023); Cheng et al. (2023); Wang et al. (2023); Yue et al. (2023); Lei et al. (2023); Fazekas et al. (2023); Chen et al. (2023); Zhang et al. (2023b) through simple prompting, this breakthrough has catalyzed a significant paradigm shift. Some attempts have been made to apply SAM to histopathology images to boost the performance on various segmentation tasks. Med-SAM Ma et al. (2024) fine-tunes the mask decoder of SAM on various medical datasets, while Medical SAM Adapter Wu et al. (2023) incorporates an adapter trained on diverse medical datasets. Although these methods show strong performance on particular tasks, we hypothesize that they are suboptimal in processing histopathology WSIs that possess a pyramid structure of multiple resolutions. For instance, to capture detailed information of ductal carcinoma in situ (DCIS) lesions, patches must be extracted from images as large as $10,000 \times 10,000$ pixels at $10\times$ magnification to manage costs van Rijthoven et al. (2021); Gu et al. (2018) and ensure compatibility with SAM, which has an input resolution of 1024×1024 . A single resolution SAM model might fail because more contextual information is needed in order to "understand" the object of interest is a lesion that should also include the lesion interior, as shown in Figure 1. More visualization results can be found in the supplementary materials.

To overcome the aforementioned limitation, in this work, we propose WSI-SAM, a model that leverages multi-resolution patches (i.e., subregions from a WSI) to perform segmentation in a zero-shot manner. Similar to HookNet van Rijthoven et al. (2021), we extract a patch from a position within the WSI and couple it with a concentric lower-resolution patch to capture essential contextual information, as illustrated in Figure 1 (right). Instead of fusing the contextual information directly at the pixel level, we execute aggregation at the token level due to token represent the segmented object information in the SAM mask generation mechanism. Given that re-training the SAM model can significantly degrades its general zero-shot performance Ke et al. (2023), we propose the WSI-SAM architecture. This design tightly integrates with and re-utilizes the existing learned structure of SAM to fully retain its zero-shot capability. Specifically, we firstly introduce High-Resolution (HR) Token and Low-Resolution (LR) Token. Unlike the original output tokens, our HR and LR Token and their associated MLP layers are trained to predict masks of high-resolution and

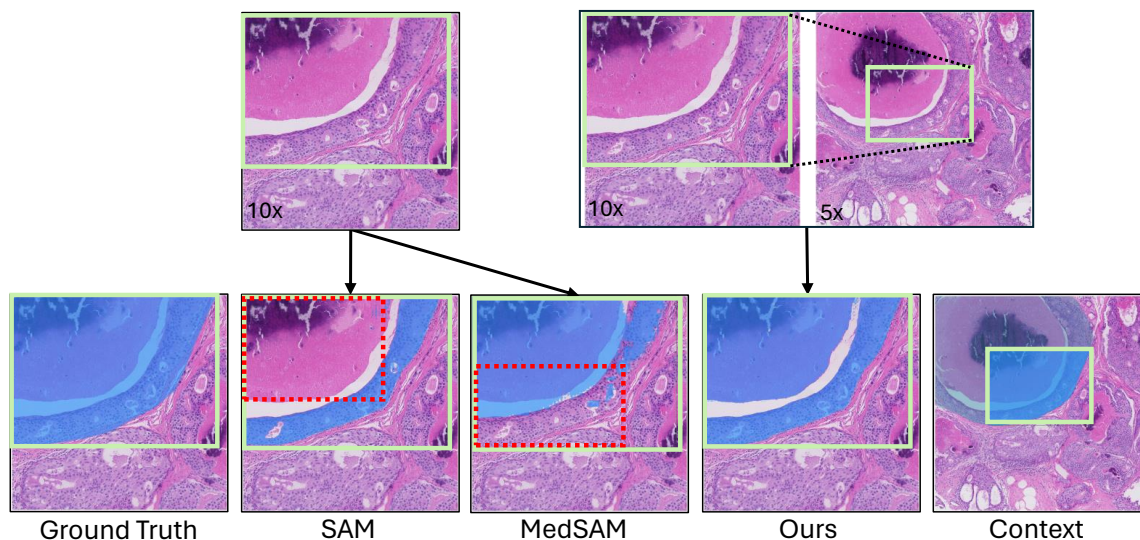


Figure 1: Comparison of DCIS segmentation in H&E-stained breast tissue by SAM, MedSAM, and our WSI-SAM. Using the same green box as input prompt on a 10 \times magnification patch, SAM erroneously does not segment the interior wall of DCIS lesion. This error is compounded by the presence of calcifications and necrosis in the interior of the duct. MedSAM overlooks the ductal region beneath the lumen. Leveraging additional context (right), our WSI-SAM predict more accurate entire DCIS area, despite the intervening background and dark region.

low-resolution patches. Secondly, we propose a dual mask decoder that integrates the original SAM mask decoder with a fusion module. The fusion module enables the integration of global semantic context with local detailed features by combining SAM’s mask decoder features with early and late features from its ViT encoder. Finally, instead of predicting masks independently, we integrate the HR and LR tokens at the intermediate decoder layer for contextual information aggregation across both resolutions to generate accurate mask details.

Our contributions are summarized as follows: (1) This paper introduces WSI-SAM, a segmentation framework building upon SAM, designed to seamlessly integrate contextual cues with high-resolution details, enabling the prediction of the detailed segmentation mask for histopathology images in a zero-shot manner. (2) To achieve this objective, we introduce the HR and LR Token, Dual Mask Decoder, and Token Aggregation, which together facilitate enhanced segmentation without the need for training from scratch. (3) The effectiveness of our method has been validated by two benchmarks, DCIS and CAMELYON16.

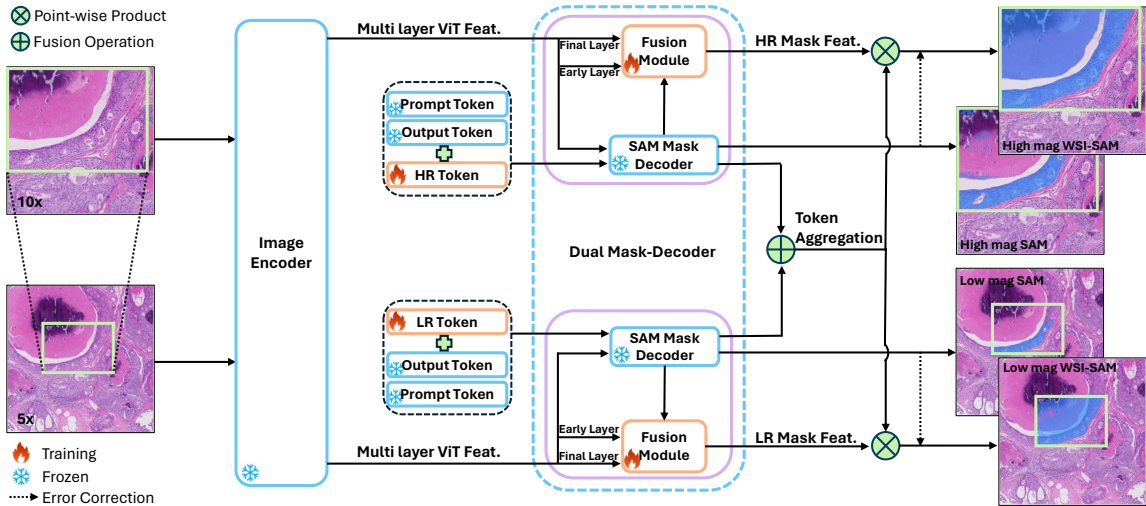


Figure 2: WSI-SAM model architecture, which introduces HR and LR Tokens, Dual Mask Decoder and Token Aggregation to SAM for for improving the mask quality in histopathology WSIs.

2 Method

2.1 Preliminary: SAM architecture

SAM consists of three fundamental components. The image encoder, utilizing a ViT-based backbone, extracts image features to produce image embeddings. The prompt encoder captures positional information from input points, boxes, or masks, aiding in the mask decoding process. Lastly, the mask decoder leverages both image embeddings and prompt tokens to generate final mask predictions. One of SAM’s remarkable features is its robust zero-shot capability to novel scenarios, thanks to extensive training on a vast repository of prompt-mask pairs.

2.2 Ours: WSI-SAM

To maintain SAM’s zero-shot transfer ability and avoid model overfitting or catastrophic forgetting, we opt for minimal adaptation rather than direct fine-tuning SAM or incorporating a new, extensive decoder network. To this end, we introduce three novel components in our WSI-SAM, as illustrated in Figure 2.

High and Low Resolution Tokens We introduce an efficient adaptation method to enhance mask quality in histopathology WSIs. As shown in Figure 2, we utilize the output token for mask prediction Kirillov et al. (2023), predicting dynamic MLP weights and subsequently applying a point-wise product with the mask features. To improve mask quality in WSIs using SAM, we avoid directly utilizing SAM’s coarse masks as input. Instead, we introduce HR and LR Tokens along with a new mask prediction layer, enabling refined mask prediction at multiple resolutions. As depicted in Figure 2, we preserve SAM’s mask decoder unchanged while introducing two new learnable tokens, HR and LR Tokens, each

with dimensions 1×256 . These tokens are then merged with SAM’s output tokens (sized 4×256) and prompt tokens (sized $N_{prompt} \times 256$), serving as input to SAM’s mask decoder. The HR and LR Tokens interact with their corresponding resolution image features for its feature updating.

Dual Mask Decoder To achieve the best interaction between token and ViT features, we propose dual mask decoder, which include the SAM mask decoder and a lightweight fusion module. HR and LR Tokens employ the point-wise MLP, which is shared with other tokens. Once processed through the mask decoder layers, the updated HR and LR Tokens acquire comprehensive insights into the global image context, crucial geometric/type information conveyed by prompt tokens, and hidden mask details inherent in other output tokens. To enhance mask quality, we augment the mask decoder features of SAM with both high-level object context and low-level boundary/edge details at Fusion modules. Specifically, we integrate new features by merging features from various stages of the SAM model, which include features from both early and late layers of ViT encoder, as well as mask features obtained from SAM’s mask decoder.

Token Aggregation Rather than predicting the mask independently, we introduce Token Aggregation to capture contextual information across multiple resolutions, thereby enhancing mask detail accuracy. We merge detailed features from the updated HR Token with broader contextual features from the LR Token, as these tokens represent features of the same object at different resolutions. As illustrated in Figure 2, the merge operation is formed by averaging the updated HR and LR Tokens. This aggregation method is both simple and effective, yielding segmentation results that preserve details while requiring minimal memory and computational resources. Subsequently, a spatial point-wise product is applied to the HR and LR mask features to facilitate mask generation.

2.3 Training Objective

While training WSI-SAM, a separate loss is computed for each resolution. We propose a loss function

$$L = \lambda L_{high} + (1 - \lambda)L_{low},$$

where L_{high} and L_{low} are the combination of dice loss and cross-entropy loss for different resolutions, respectively, and λ controls the importance of each resolution.

3 Experiments and Results

Training data To train WSI-SAM in a data-efficient manner, we train on the **CATCH** Wilm et al. (2022) dataset. This dataset contain 350 WSIs of seven distinct subtypes of canine cutaneous tumors, augmented with 12,424 polygon annotations across 13 histological classes. Following the official division, the dataset is split into a training set with 245 WSIs and a validation set encompassing 35 WSIs.

Benchmark settings We assess our model’s performance on two Whole Slide Image (WSI) datasets in a zero-shot manner. The first is a dataset of Ductal Carcinoma in Situ (**DCIS** Wetstein et al. (2021)), a non-invasive breast cancer collection comprising 116 WSIs.

From this, 50 WSIs were randomly selected for our test set. The second is the **CAMELYON16**¹ dataset, where we incorporated its official test set into our own. This set includes a total of 47 WSIs.

3.1 Implementation Details

During training, we maintain the model parameters of the pre-trained SAM model unchanged, focusing solely on making the proposed WSI-SAM learnable. Given SAM’s capability for handling flexible segmentation prompts, we train WSI-SAM using a variety of prompt types, such as bounding boxes, randomly sampled points, and coarse masks. These degraded masks are created by incorporating random Gaussian noise into the boundary areas of the GT masks. Given hardware constraints, we randomly select a limited number of 1024×1024 patches at $10\times$ magnification and their corresponding concentric patches at $5\times$ magnification from each WSI for training. We adopt TinyViT Zhang et al. (2023a) as its backbone. Our training employs a mini-batch size of 1 and a learning rate 0.001. We conducted all experiments using PyTorch on a setup equipped with an NVIDIA TITAN Xp GPU.




Figure 3: Example of predicted tissue segmentation on DCIS. Yellow boxes indicate the incorrect predictions.

3.2 Zero-shot Comparison with SAMs

Setup In our evaluation process, we conduct zero-shot mask prediction using prompts. Our investigation includes two distinct types of prompts: box and point. To precisely measure enhancements in mask quality, we utilize the dice score as our reporting metric.

Zero shot segmentation with box prompt We compare with SAM, fine-tuned SAM and MedSAM on both DCIS and CAMELYON16 datasets. We simulate realistic human-annotated bounding boxes by introducing noise to the GT object boxes. As shown in Table 1, we observe a marked decrease in performance when fine-tuning SAM on histopathology

1. <https://camelyon17.grand-challenge.org/Data/>

Table 1: Comparison of zero-shot segmentation results on the DCIS and CAMELYON16 test sets using bounding boxes as input prompts. * indicate fine-tune SAM’s mask decoder on CATCH.

Method	DCIS	CAMELYON16			
	DCIS	IDC	ILC	DUC	avg
SAM Kirillov et al. (2023)	73.35	86.06	81.36	72.86	80.09
SAM*	61.22	74.17	72.88	40.00	62.35
MedSAM Ma et al. (2024)	64.71	77.91	75.24	71.64	74.93
WSI-SAM (Ours)	77.50	89.76	84.30	73.55	82.54

Table 2: Comparison of zero-shot segmentation results on the test set of DCIS. We use nnU-Net Isensee et al. (2021) trained on DCIS as box prompt generator.

Method	SAM	MedSAM	WSI-SAM
Dice	56.00	56.43	57.37

WSIs, attributed to the disruption of pre-trained knowledge through fine-tuning. In contrast, WSI-SAM outperforms SAM by 4.1 and 2.5 percent points on the DCIS and CAMELYON16 datasets, respectively. Furthermore, when compared to MedSAM, which was trained on a variety of medical datasets, WSI-SAM demonstrates a significant performance enhancement, as evidenced by the increase from **64.71%** to **77.50%**. Moreover, to simulate real-world scenarios, we further evaluated our method using prompts generated from nnU-Net Isensee et al. (2021). WSI-SAM surpasses both SAM and MedSAM (e.g., **56.0**, **56.73** vs **59.9**), showcasing the resilience and robustness of our approach. To provide additional validation, we conduct a segmentation qualitative analysis on DICS, as illustrated in Figure 3.

Zero shot segmentation with point prompt To investigate the segmentation capabilities of WSI-SAM using interactive point prompts, we conducted a comparison with SAM, evaluating their performances across a range of input point quantities on both the DCIS and CAMELYON16 datasets. Noting that MedSAM is limited to box prompts. For the CAMELYON16 dataset, we present an average performance across the three tumor classes, with comprehensive metrics for each class detailed in the supplementary materials. WSI-SAM consistently outperforms SAM on both datasets, regardless of the number of point prompts utilized, as shown in the Figure 4.

3.3 Ablation Study

We conducted an ablation study for WSI-SAM on DCIS dataset with bounding boxes as input prompts.

Effect of the HR and LR Tokens. WSI-SAM leverages HR and LR Tokens to fuse more contextual information, enhancing mask prediction accuracy. We examined various aggregation target, specifically fuse the HR and LR feature by average-pooling, fuse with expand HR feature region on the HR feature, and HR and LR token aggregation. Results presented in Table 3b indicate that utilizing HR and LR Tokens, outperforms these alternatives, achieving performance improvements of 5.3 percent points on DCIS dataset.




Figure 4: Comparison of zero-shot interactive segmentation results using a varying number of input points on the DCIS and CAMELYON16 dataset.

Table 3: **Ablation experiments on DCIS dataset.** (a) Ablation study on aggregation ways of HR and LR features or Tokens using box prompts; (b) Ablation study on aggregation target for integrating contextual information using box prompts; (c) Ablation study on different values of λ in loss function.

(a)		(b)		(c)	
Aggregation	Dice	Aggregation Target	Dice	λ	Dice
Concat.-FC	77.24	HR feat. and LR feat.	72.20	0.0	74.16
Max.	75.91	HR feat. and expand HR feat.	59.04	0.5	77.50
Avg.	77.50	HR Token and LR Token	77.50	1.0	75.20

Ablation on Tokens Aggregation We evaluate three ways of fusing HR and LR Tokens: (1) Concatenating followed by one learnable FC layer (Concat-FC); (2) Max pooling (Max); (3) Average pooling (Avg). Table 3a shows that the average pooling turns out to be the best way for aggregation.

Ablation on λ in loss function. The influence of losses from HR and LR Tokens is modulated by λ . Experiments were conducted with $\lambda = 0.0$ to ignore the HR Token, $\lambda = 0.5$ to balance both tokens equally, and $\lambda = 1.0$ to ignore the LR Token. According to Table 3c, setting $\lambda = 0.5$ yields the best performance enhancement.

4 Conclusion

We introduce WSI-SAM, the zero-shot segmentation model tailored for WSIs, incorporating innovative HR and LR Tokens to refine the mask prediction of SAM’s output token and enhancing the original SAM with minimal additional computational cost. Our zero-shot transfer evaluations on the DCIS and CAMELYON16 datasets showcase WSI-SAM’s superior performance, marking a substantial advancement in zero-shot segmentation for WSIs.

Appendix A.

Table 4: Comparison of zero-shot segmentation results on CAMELYON16 test sets using different numbers of points as input prompts. (a) 3 positive points; (b) 5 positive points; (c) 10 positive points.

(a)			
Method	IDC	ILC	DUC
SAM	85.26	80.76	60.13
WSI-SAM	85.63	81.66	61.73

(b)			
Method	IDC	ILC	DUC
SAM	87.53	82.09	66.48
WSI-SAM	89.53	84.90	68.16

(c)			
Method	IDC	ILC	DUC
SAM	89.43	83.97	74.61
WSI-SAM	91.53	86.23	75.22

WSI-SAM

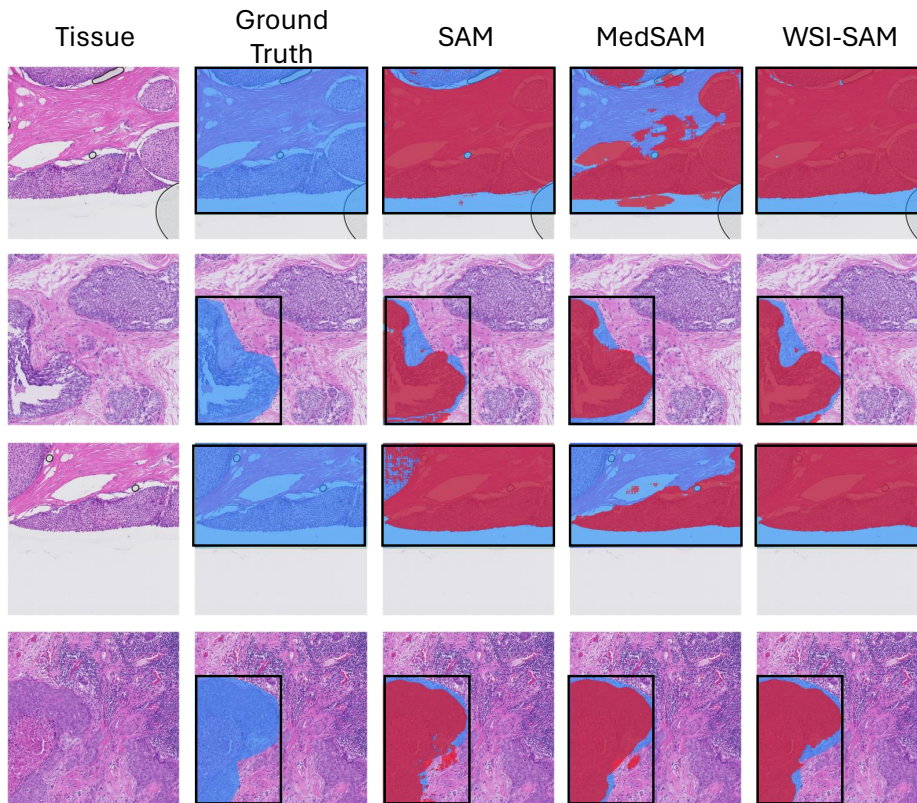


Figure 5: Comparison of segmentation mask predictions in DCIS dataset.

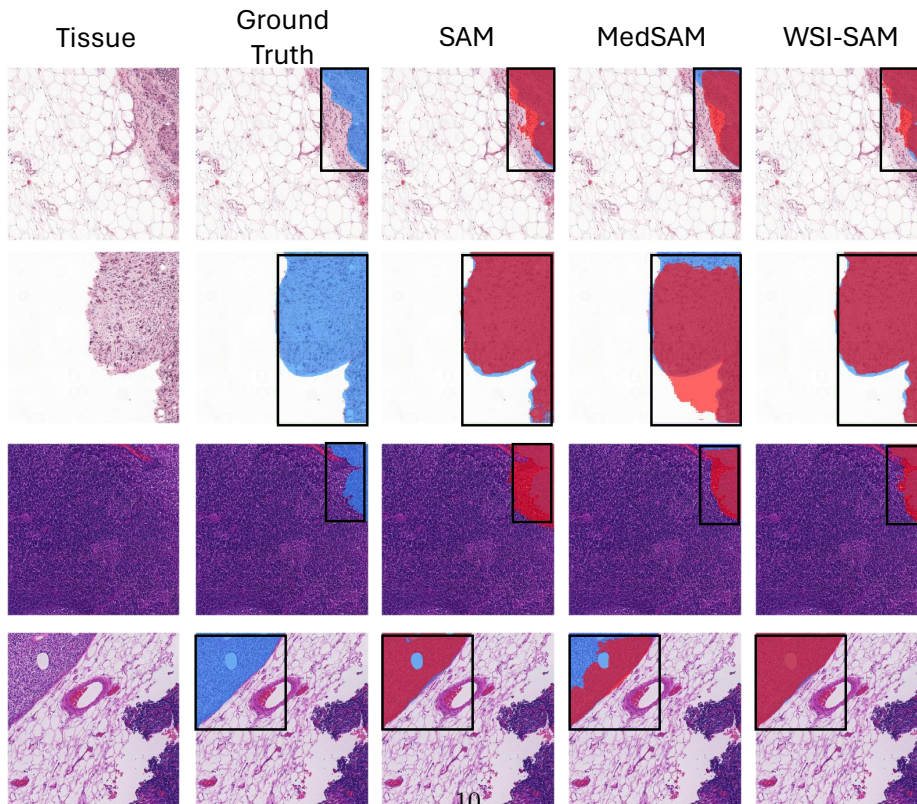


Figure 6: Comparison of segmentation mask predictions in CAMELYON16 dataset.

References

- Cheng Chen, Juzheng Miao, Dufan Wu, Zhiling Yan, Sekeun Kim, Jiang Hu, Aoxiao Zhong, Zhengliang Liu, Lichao Sun, Xiang Li, Tianming Liu, Pheng-Ann Heng, and Quanzheng Li. Ma-sam: Modality-agnostic sam adaptation for 3d medical image segmentation. arXiv:2309.08842, 2023. arXiv preprint.
- Junlong Cheng, Jin Ye, Zhongying Deng, Jianpin Chen, Tianbin Li, Haoyu Wang, Yanzhou Su, Ziyang Huang, Jilong Chen, Lei Jiang, Hui Sun, Junjun He, Shaoting Zhang, Min Zhu, and Yu Qiao. Sam-med2d. arXiv:2308.16184, 2023. arXiv preprint.
- Botond Fazekas, José Morano, Dmitrii Lachinov, Guilherme Aresta, and Hrvoje Bogunović. Samedoct: Adapting segment anything model (sam) for retinal oct. arXiv:2308.09331, 2023. arXiv preprint.
- Ruiwei Feng, Xuechen Liu, Jintai Chen, Danny Z. Chen, Honghao Gao, and Jian Wu. A deep learning approach for colonoscopy pathology wsi analysis: Accurate segmentation and classification. *IEEE Journal of Biomedical and Health Informatics*, 25(10):3700–3708, 2021a. doi: 10.1109/JBHI.2020.3040269.
- Yanbo Feng, Adel Hafiane, and Hélène Laurent. A deep learning based multiscale approach to segment the areas of interest in whole slide images. *Computerized Medical Imaging and Graphics*, 90:101923, 2021b. ISSN 0895-6111. doi: <https://doi.org/10.1016/j.compmedimag.2021.101923>.
- Feng Gu, Nikolay Burlutskiy, Mats Andersson, and Lena Kajland Wilén. Multi-resolution networks for semantic segmentation in whole slide images. In *Computational Pathology and Ophthalmic Medical Image Analysis*, pages 11–18. Springer International Publishing, 2018.
- Chu Han, Jiatai Lin, Jinhai Mai, Yi Wang, Qingling Zhang, Bingchao Zhao, Xin Chen, Xipeng Pan, Zhenwei Shi, Zeyan Xu, Su Yao, Lixu Yan, Huan Lin, Xiaomei Huang, Changhong Liang, Guoqiang Han, and Zaiyi Liu. Multi-layer pseudo-supervision for histopathology tissue semantic segmentation using patch-level classification labels. *Medical Image Analysis*, 80:102487, 2022. doi: <https://doi.org/10.1016/j.media.2022.102487>.
- Fabian Isensee, Paul F. Jaeger, Simon A. A. 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. doi: 10.1038/s41592-020-01008-z.
- Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality. In *Advances in Neural Information Processing Systems*, volume 36, pages 29914–29934, 2023.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3992–4003, 2023. doi: 10.1109/ICCV51070.2023.00371.

- Wenhui Lei, Xu Wei, Xiaofan Zhang, Kang Li, and Shaoting Zhang. Medlsam: Localize and segment anything model for 3d ct images. arXiv:2306.14752, 2023. arXiv preprint.
- Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15:1–9, 2024.
- 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. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2023.102918>.
- Haomiao Ni, Hong Liu, Kuansong Wang, Xiangdong Wang, Xunjian Zhou, and Yueliang Qian. Wsi-net: Branch-based and hierarchy-aware network for segmentation and classification of breast histopathological whole-slide images. In *Machine Learning in Medical Imaging*, pages 36–44. Springer International Publishing, 2019.
- Talha Qaiser, Yee-Wah Tsang, Daiki Taniyama, Naoya Sakamoto, Kazuaki Nakane, David Epstein, and Nasir Rajpoot. Fast and accurate tumor segmentation of histology images using persistent homology and deep convolutional features. *Medical Image Analysis*, 55: 1–14, 2019. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2019.03.014>.
- Mart van Rijthoven, Maschenka Balkenhol, Karina Siliņa, Jeroen van der Laak, and Francesco Ciompi. Hooknet: Multi-resolution convolutional neural networks for semantic segmentation in histopathology whole-slide images. *Medical Image Analysis*, 68:101890, 2021. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2020.101890>.
- Haoyu Wang, Sizheng Guo, Jin Ye, Zhongying Deng, Junlong Cheng, Tianbin Li, Jianpin Chen, Yanzhou Su, Ziyang Huang, Yiqing Shen, Bin Fu, Shaoting Zhang, Junjun He, and Yu Qiao. Sam-med3d. arXiv:2310.15161, 2023. arXiv preprint.
- Xiyue Wang, Yuqi Fang, Sen Yang, Delong Zhu, Minghui Wang, Jing Zhang, Kai yu Tong, and Xiao Han. A hybrid network for automatic hepatocellular carcinoma segmentation in h&e-stained whole slide images. *Medical Image Analysis*, 68:101914, 2021. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2020.101914>.
- Suzanne C. Wetstein, Nikolas Stathonikos, Josien P.W. Pluim, Yujing J. Heng, Natalie D. ter Hoeve, Celien P.H. Vreuls, Paul J. van Diest, and Mitko Veta. Deep learning-based grading of ductal carcinoma in situ in breast histopathology images. *Laboratory Investigation*, 101(4):525–533, 2021. ISSN 0023-6837. doi: <https://doi.org/10.1038/s41374-021-00540-6>.
- Frauke Wilm, Marco Fragoso, Christian Marzahl, Jingna Qiu, Chloé Puget, Laura Diehl, Christof A. Bertram, Robert Klopffleisch, Andreas Maier, Katharina Breininger, and Marc Aubreville. Pan-tumor canine cutaneous cancer histology (catch) dataset. *Scientific Data*, 9(1):588, 2022. doi: [10.1038/s41597-022-01692-w](https://doi.org/10.1038/s41597-022-01692-w).
- Junde Wu, Wei Ji, Yuanpei Liu, Huazhu Fu, Min Xu, Yanwu Xu, and Yueming Jin. Medical sam adapter: Adapting segment anything model for medical image segmentation. arXiv:2304.12620, 2023. arXiv preprint.

Gang Xu, Zhigang Song, Zhuo Sun, Calvin Ku, Zhe Yang, Cancheng Liu, Shuhao Wang, Jianpeng Ma, and Wei Xu. Camel: A weakly supervised learning framework for histopathology image segmentation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10681–10690, 2019. doi: 10.1109/ICCV.2019.01078.

Wenxi Yue, Jing Zhang, Kun Hu, Yong Xia, Jiebo Luo, and Zhiyong Wang. Surgicalsam: Efficient class promptable surgical instrument segmentation. arXiv:2308.08746, 2023. arXiv preprint.

Chaoning Zhang, Dongshen Han, Sheng Zheng, Jinwoo Choi, Tae-Ho Kim, and Choong Seon Hong. Mobilesamv2: Faster segment anything to everything. arXiv:2312.09579, 2023a. arXiv preprint.

Yichi Zhang, Yuan Cheng, and Yuan Qi. Semisam: Exploring sam for enhancing semi-supervised medical image segmentation with extremely limited annotations. arXiv:2312.06316, 2023b. arXiv preprint.