



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

SimTxtSeg: Weakly-Supervised Medical Image Segmentation with Simple Text Cues

Yuxin Xie^{1,2}, Tao Zhou³, Yi Zhou^{1,2*}, and Geng Chen⁴

¹ School of Computer Science and Engineering, Southeast University, China

² Key Laboratory of New Generation Artificial Intelligence Technology and Its Interdisciplinary Applications, Ministry of Education, China

³ Nanjing University of Science and Technology, China

⁴ Northwestern Polytechnical University, China

{silver_iris@163.com, yizhou.szcn@gmail.com}

Abstract. Weakly-supervised medical image segmentation is a challenging task that aims to reduce the annotation cost while keep the segmentation performance. In this paper, we present a novel framework, SimTxtSeg, that leverages simple text cues to generate high-quality pseudo-labels and study the cross-modal fusion in training segmentation models, simultaneously. Our contribution consists of two key components: an effective Textual-to-Visual Cue Converter that produces visual prompts from text prompts on medical images, and a text-guided segmentation model with Text-Vision Hybrid Attention that fuses text and image features. We evaluate our framework on two medical image segmentation tasks: colonic polyp segmentation and MRI brain tumor segmentation, and achieve consistent state-of-the-art performance. Source code is available at: <https://github.com/xyx1024/SimTxtSeg>.

Keywords: Weakly-supervised medical image segmentation · Textual-to-visual cue converter · Text-vision hybrid attention.

1 Introduction

Medical image segmentation [24] plays a crucial role in medical image analysis, which are usually trained in a fully-supervised manner. However, this kind of approach heavily suffers expensive annotation cost of providing pixel-level labels, impeding practical clinical application. In recent years, a wave of weakly-supervised segmentation models has emerged, which operate with different label levels, such as image-level[10,23,29], point-level[8,19], scribble-level[14,22,25], and bounding box-level[9,26,30] methods. By leveraging techniques such as reinforcement learning and active learning, they bridge the gap between pseudo-labels and ground truths, enabling pixel-level segmentation for medical images. Despite their innovative approaches, the current challenge lies in the fact that the results achieved by these methods still fall short of the performance exhibited by fully-supervised learning ways. Therefore, we aim to study lower-cost and higher-quality pseudo-labels for weakly-supervised medical image segmentation.

* Corresponding author: Yi Zhou

The Segment Anything Model (SAM)[12], a general visual foundation segmentation model, has garnered widespread attention due to its remarkable segmentation and robust zero-shot generalization capabilities. Although SAM has been trained using large-scale data with pixel-level labels, its performance for medical image segmentation is unsatisfactory due to the lack of reliable clinical training data. Consequently, many researchers have fine-tuned SAM specifically in medical domains[17,31,13,4], including full fine-tuning and parameter-efficient tuning, achieving promising performance. Nonetheless, models based on SAM still require providing manual visual prompts (e.g. point and box prompts) for each image, increasing the difficulty and time required for expert physicians to make annotations. Therefore, we aim to explore a novel and automatic approach only using simple text cues to accomplish weakly-supervised medical image segmentation, through equipping SAM with a language-to-vision prompt converter. Moreover, to further conduct cross-modal fusion, we seek to better integrate text cues into the target visual segmentation model, ultimately enhancing the effect of language-driven segmentation performance.

In this paper, we propose a weakly-supervised medical image segmentation pipeline, **SimTxtSeg**. After the establishment of a domain-specific pre-training framework, the text prompt can easily be converted into a visual prompt and a pseudo-mask. Hence, only with a simple text cue, a target segmentation model can be trained in a weakly-supervised manner, eliminating the need of repeatedly providing pixel-level annotations. The most significant problem we investigated is how to effectively integrate information from simple text cues into the visual segmentation task model, such as transforming textual prompts into visual ones. Overall, we put forth SimTxtSeg, consisting of two key components: a Textual-to-Visual Cue Converter and a Text-Vision Hybrid Attention module.

We highlight our contributions as follows: **1)** We propose to address weakly-supervised medical image segmentation using simple textual prompts, by extending the zero-shot generalization capability of SAM, thereby reducing the burden of pixel-level annotation on medical images. **2)** The proposed SimTxtSeg includes a Textual-to-Visual Cue Converter(TVCC) and a Text-Vision Hybrid Attention(TVHA) module, promoting the integration of textual cues into the visual medical image segmentation task. **3)** Through extensive comparison and ablation experiments, we validate the effectiveness of our approach, demonstrating state-of-the-art performance across multiple datasets, including scenarios such as intestinal polyp segmentation and MRI brain tumor segmentation.

2 Proposed Method

2.1 Problem Formulation

Given a dataset with N image-text pairs, *i.e.*, $D = \{(I_1, T_1), \dots, (I_N, T_N)\}$, where $I_i \in \mathbb{R}^{H \times W \times 3}$ represents the i -th images, H, W represent height,width, and $T_i \in \mathbb{R}^L$ is a brief description of the image, L is the length of the sentence. First, we aim to train a textual-to-visual cue converter capable of directly localizing regions of interest in an image using simple descriptive text, then obtain the

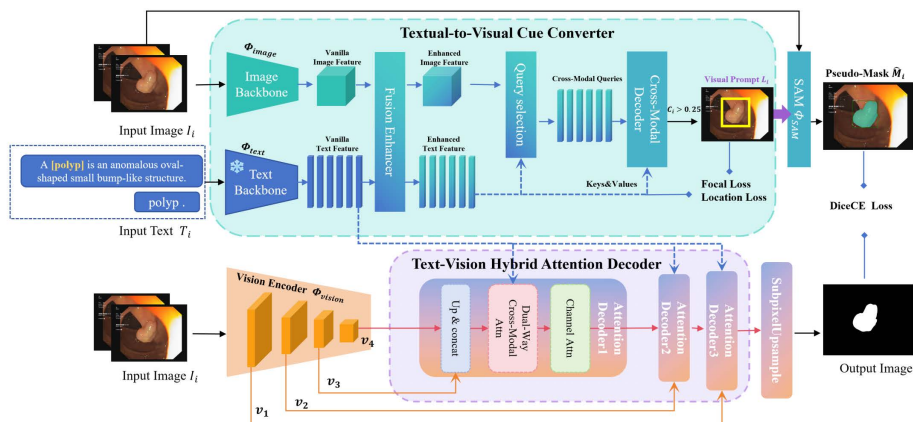


Fig. 1: **The framework of SimTxtSeg.** The textual-to-visual cue converter enables SAM to generate pseudo masks via text cues. Then, the weakly-supervised segmentation model is enhanced by text-vision hybrid attention.

pseudo-masks $\hat{M} \in \mathbb{R}^{H \times W \times 1}$ by using SAM’s zero-shot capability, thus eliminating the need for pixel-level annotations from doctors.

$$B_i, S_i = \Phi_{converter}(\Phi_{image}(I_i), \Phi_{text}(T_i)), \quad (1)$$

$$\hat{M}_i = \Phi_{SAM}(I_i, B_i), \quad (2)$$

where B_i, S_i denote the bounding boxes and their confidence predicted by our textual-to-visual cue converter ($\Phi_{converter}$), which has an image backbone (Φ_{image}) and a text backbone (Φ_{text}). Φ_{SAM} represents Segment Anything Model. Second, we propose a text-guided medical image segmentation model incorporated with a text-vision hybrid attention module in the decoder. To demonstrate the effectiveness of this weakly-supervised manner, we train it with image-text pairs D_{train} and the pseudo-masks \hat{M}_{train} . The overall pipeline is shown in Fig. 1.

2.2 Textual-to-Visual Cue Converter

Inspired by GroundingDINO[15], the construction of our Textual-to-Visual Cue Converter consists of image and text backbones for feature extraction, a feature enhancer for image and text feature fusion, a language-guided query selection module for query initialization, and a cross-modality decoder for box refinement. We utilize the mmdetection framework [18] to fine-tune the Textual-to-Visual Cue Converter based on Swin-T [16], employing a domain-specific medical dataset. Our training data are transformed to the ODVG format for precise alignment of regions and phrases, *i.e.*, $D_{train} = \{(I_1, T_1, G_1), \dots, (I_N, T_N, G_N)\}$, where $G_i = \{\{Bbox_1, Phrase_1\}, \dots, \{Bbox_N, Phrase_N\}\}$ contains the bounding boxes and their corresponding phrases, N is the number of the lesions in the

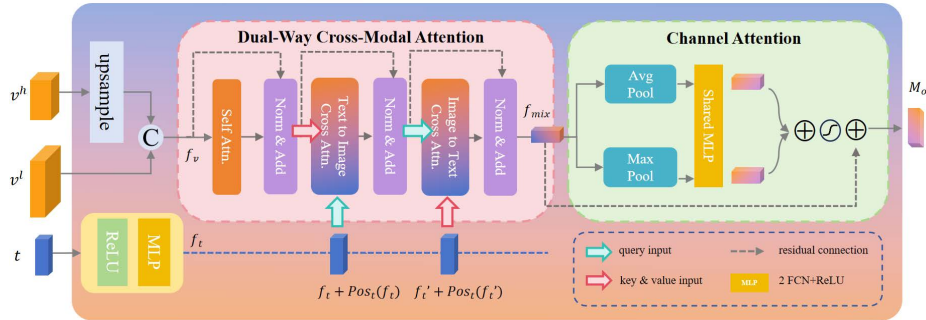


Fig. 2: Detailed structure of our text-vision hybrid attention decoder layer, containing the essential dual-way cross-modal attention and channel attention.

image. During training, we keep the weights of position embedding, backbone, and the language model (BERT-BASE) fixed, focusing solely on training the feature enhancer and cross-modality decoder. Consequently, it accurately generates precise annotation boxes based on textual cues.

Once the Textual-to-Visual Cue Converter is pre-trained, we can straightforwardly transfer the textual prompts into visual prompts for any new dataset within the same medical domain. Then we employ SAM as our pseudo-masks generator with the visual prompts B , configuring the confidence threshold at 0.25. In our study, both the vanilla SAM and SAM-med2d are experimented.

2.3 Text-Guided Segmentation with Text-Vision Hybrid Attention

The objective of our work is to train a medical image segmentation model based on weakly-supervised text cues. Notably, these simple text cues serve a dual purpose: to generate pseudo-labels for supervision and to be directly integrated into the target segmentation model, effectively infusing intricate semantic details into the visual task model.

Vision Encoder & Text Encoder: Given an image I , we choose ConvNext-Tiny[28] as our vision encoder Φ_{vision} : $v_i = \Phi_{vision}(I)$, where i refers to the i -th layer in the backbone. We extract its first four layers of output for feature fusion, which are defined as $v_1 \in \mathbb{R}^{H/4 \times W/4 \times C1}$, $v_2 \in \mathbb{R}^{H/8 \times W/8 \times C2}$, $v_3 \in \mathbb{R}^{H/16 \times W/16 \times C3}$, $v_4 \in \mathbb{R}^{H/32 \times W/32 \times C4}$. Given a sentence T , we take BERT-BASE[5] as our tokenizer and text backbone, and take its last embedding $t \in \mathbb{R}^{l \times C}$, where l is the length of token, and C refers to the feature dimension. $t = \Phi_{text}(tokenize(T))$.

Text-Vision Hybrid Attention Decoder: We employ three Text-Vision Hybrid Attention decoder layers, and a subpixel-upsample layer in our decoder. The details of Text-Vision Hybrid Attention decoder layer are illustrated in Fig. 2, which consists of a dual-way cross-modal attention and a channel attention.

Let v^h represents the high-level feature from the previous decoder layer ($v^h = v_4$ in the first decoder), and v^l represents the low-level feature from the

corresponding encoder layer, we upsample v^h and concatenate it with v^l , then obtain the output f_v .

As a cue, the text embedding t is aligned with visual feature dimensions through a projection layer, which is shown in the following equation:

$$f_t = \Phi_{projection}(t) = ReLU(MLP(t)), \quad (3)$$

where MLP contains a 1×1 convolution layer, a GELU activation function and a linear layer. $f_t \in \mathbb{R}^{L_j \times C_j}$ is the output text embedding, L_j, C_j represent the length and channel number of the output token in the j -th decoder layer. Also, the image embedding's shape is projected into $HW \times C_j$, consistent with f_t .

For a more fine-grained integration of text and visual features, we propose a **dual-way cross-modal attention** in Fig. 2. Given f_v, f_t representing aligned image and text embeddings, the dual-way cross-modal attention module performs three steps. First, we compute self-attention on the f_v , using image position embedding as the query and key, f_v as value. Residual connection is employed to preserve the vision feature. The self attention is processed as:

$$f'_v = LayerNorm(MHSA(f_v)) + f_v, \quad (4)$$

where $MHSA(\cdot)$ refers to Multi-Head Self Attention. Second, a text-to-vision attention is applied, which means cross-attention from text (text position embedding as query) to the image embedding (image position as key, f_v as value). Then, a norm-and-add layer is applied. The text-to-vision cross-attention process is shown in Eq. 5:

$$f'_t = LayerNorm(MHCA_{tv}(f_t, f'_v)) + f_t, \quad (5)$$

where $MHCA_{tv}(\cdot)$ represents text-to-vision Multi-Head Cross-Attention. Finally, we employ a vision-to-text attention, with image position embedding as query, text position as key, and f_t as value, followed by the add-norm function, to get the fused feature f_{mix} :

$$f_{mix} = LayerNorm(MHCA_{vt}(f'_v, f'_t)) + f'_v, \quad (6)$$

where $MHCA_{vt}(\cdot)$ represents vision-to-text Multi-Head Cross-Attention.

To further exploit the most useful feature channels, we introduce channel attention to automatically highlight the relevant feature channels while suppressing irrelevant channels. As is shown in Fig. 2, the mixed feature f_{mix} undergoes global max pooling and global average pooling based on its width and height to fuse the spatial information across the entire feature map. The pooled features are individually processed through an MLP, which learns channel-specific weights and biases to enhance or suppress certain features. Then, the MLP outputs are element-wisely summed together and passed through a sigmoid activation, to obtain the decoder layer's output feature map M_o as:

$$M_o = Sigmoid(MLP(AvgPool(f_{mix})) + MLP(MaxPool(f_{mix}))) + f_{mix}. \quad (7)$$

3 Experiments and Results

3.1 Experiment Setup

Colonic Polyp Dataset: We utilize the following datasets for colonic polyp segmentation: CVC-ClinicDB[2], CVC-ColonDB[21], ETIS-LaribPolypDB[20], Kvasir[11], PolypGen[1]. In total, there are 3,784 images of colonic polyps, including both images appearing polyps and normal cases. We randomly split these datasets into training (3190 images), validation (299 images), and testing (295 images) set as the ratio of 8:1:1. The image size is reshaped to (384×384) .

MRI Brain Tumor Dataset: For brain tumor segmentation, we utilize LGG Segmentation Dataset[3] from The Cancer Imaging Archive, which comprises 3,929 brain MRI images. Other settings remain consistent with the Polyp datasets.

Text Cues: We have designed two kind of text prompt granularities for each task: individual words and descriptive sentences. To avoid handcrafted prompting cost, we use GPT-4 to generate a concise sentence within 20 words. In the subsequent analysis, we will evaluate the effectiveness of these different granularities for SimTxtSeg.

Evaluation Metrics: We adopt mean Intersection over Union (mIoU) and mean Dice coefficient to evaluate the medical image segmentation performance.

Implementation Details: To pre-train the Textual-to-Visual Cue Converter, we employed Adam optimizer with an initial learning rate of 2×10^{-4} , weight decay of 1×10^{-4} , a batch size of 8, and trained it for 100 epochs. As for the parameter scheduler, we adopted both LinearLR and MultiStepLR. To train the text-guided segmentation model with Text-Vision Hybrid Attention, we freeze the text branch parameters and employ ConvNeXt as the vision backbone, with an input image size of 384. The learning rate adjustment strategy is ReduceLROnPlateau. All the methods are implemented using PyTorch, accelerated by 4 NVIDIA 4090 Ti GPUs.

3.2 Comparisons with the State-of-the-Art Methods

Comparison results against seven state-of-the-art methods are reported in Table 1. These methods fall into two categories: three fully-supervised models (ResUNet[6], PraNet[7], and Ariadne’s Thread[32]) and four weakly-supervised models with different label levels (WeakPolyp[26], BoxPolyp[27], Boxshrink[9], and S²ME[22]). We compared the segmentation performance of these SOTAs with our proposed pseudo-label generator(Pseudo-L:TVCC+SAM) and the final weakly-supervised model(SimTxtSeg-w-TVHA). It is observed that the generated pseudo-label quality is roughly on par with that of the SOTA fully-supervised models, with even a slight edge on the polyp dataset, and our final segmentation performance surpasses other SOTA weakly-supervised models. Among different kinds of weak supervision cue, our text-based cue is the weakest annotation without any spatial labeling and has the lowest cost compared to visual cues like boxes and scribbles which still cost a lot. On the polyp dataset, we achieve a +1.38% improvement in mDice and a +3.36% improvement in mIoU.

Table 1: Comparisons with state-of-the-arts on polyp and brain tumor segmentation, containing 3 fully-supervised and 4 weakly-supervised models. The results marked in gray means our generated pseudo-mask quality used for supervision.

Method			Polyp		Brain Tumor		
			mIoU(%)	mDice(%)	mIoU(%)	mDice(%)	
Fully-Supervised			ResUNet(2020)	75.31	82.60	58.42	71.27
			PraNet(2020)	81.32	87.30	74.14	82.49
			Ariadne’s Thread(2023)	80.65	87.14	71.55	81.20
Weakly-Supervised	box	WeakPolyp(2023)	79.40	85.61	63.43	74.82	
	scribble	S2ME(2023)	49.62	66.33	15.38	26.66	
	box+half anno.	BoxPolyp(2022)	79.11	86.86	67.40	77.64	
	box	boxshrink(2023)	64.22	78.21	57.02	66.36	
	-	Pseudo-L:TVCC+SAM	81.06	87.46	72.38	81.69	
	-	SimTxtSeg-w/o-TVHA	74.92	83.15	66.57	77.86	
	-	SimTxtSeg-w/o-CMA	80.83	87.22	71.16	81.57	
	text	SimTxtSeg-w/o-CA	80.64	86.87	70.42	80.97	
text	SimTxtSeg-w-TVHA	82.47	88.24	71.34	81.74		

Table 2: Eval. of prompt types and SAM variants for pseudo-mask generation.

Prompt Type	Polyp		Brain Tumor		SAM Variant	Polyp		Brain Tumor	
	mIoU	mDice	mIoU	mDice		mIoU	mDice	mIoU	mDice
[15]-w-class	22.15	29.29	10.89	13.02	SAM-base	76.87	84.32	72.38	81.69
Class name	80.84	87.30	68.00	78.15	SAM-huge	81.06	87.46	68.30	78.36
Sentence	81.06	87.46	68.30	78.36	SAM-med2d-base	70.62	79.07	67.20	77.34

Moreover, on the brain tumor dataset, our method achieves a +4.1% improvement in mDice and a +3.94% improvement in mIoU. Qualitative comparison of segmentation performance is visualized in Fig. 3.

3.3 Ablation Study

Impact of prompt types. We evaluated class name and sentence as text cues during training textual-to-visual cue converter and compared their effectiveness for pseudo-mask generation on SAM-huge. Also, we tested the performance by the original GroundingDINO[15] with class name prompt. As Table 2 shows, training the textual-to-visual cue converter with sentences (e.g. A polyp is an anomalous oval-shaped small bump-like structure.) tends to yield slightly better results than training it with class names (e.g. polyp), since the converter generates better box cues, achieving 80.1% mAP for polyp and 74.8% mAP for brain tumor. The original GroundingDINO fails to generate useful pseudo masks.

Impact of SAM variants. We compared three pre-trained SAM models: SAM-huge, SAM-base, and SAM-Med2d-base[31] for pseudo-label generation, which differ in model parameters and pretraining dataset. As seen in Table 2, the SAM-huge performs better for polyp images while SAM-base yields superior results for brain tumor dataset. Due to extensive pretraining of SAM-Med2d-

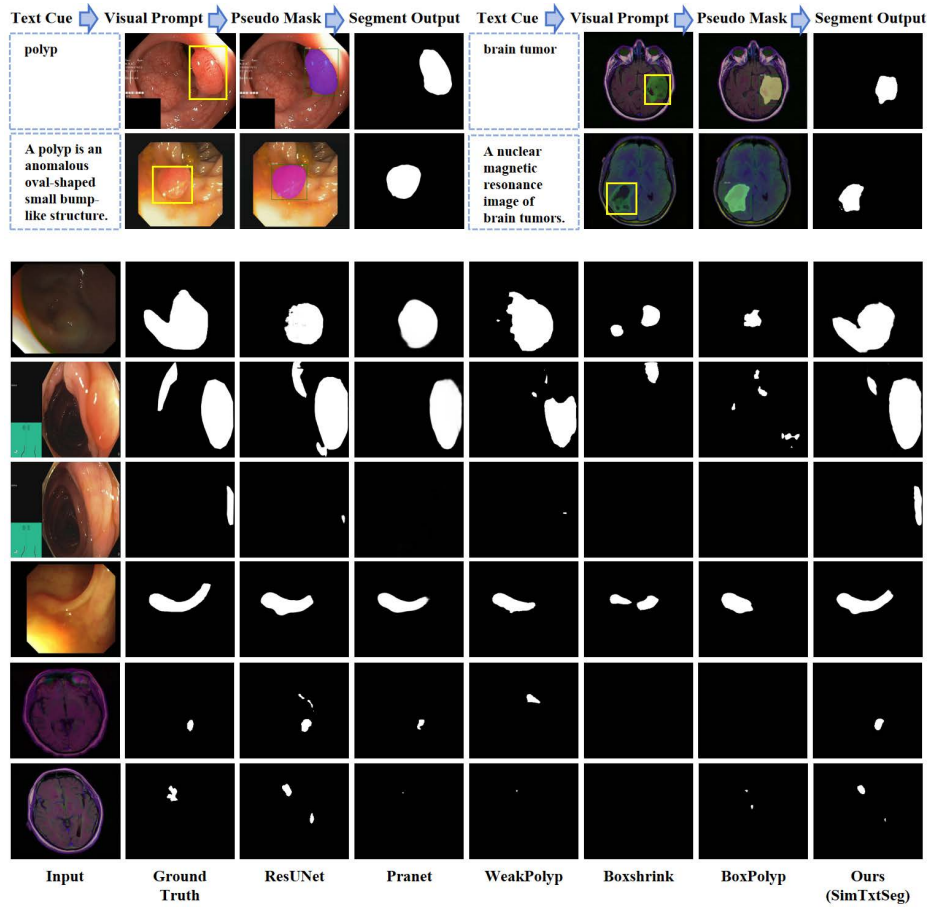


Fig. 3: Qualitative visualization on polyp and brain tumor segmentation.

base specifically with CT and MRI data, it exhibits significant bias when applied to polyp data, resulting in poor generalizability compared to the general SAM.

Impact of our TVHA. SimTxtSeg-w/o-TVHA denotes the model without Text-Vision Hybrid Attention, using UNet decoder instead. SimTxtSeg-w/o-CMA denotes the model without Dual-Way Cross-Modal Attention, SimTxtSeg-w/o-CA denotes the model without Channel Attention. From Table 1, it is observed that, after incorporating the TVHA, our model’s performance has significantly improved. Specifically, on the polyp dataset, we observe a +5.09% increase in mDice and a +7.55% increase in mIoU. The contribution of both modules to model performance improvement is roughly equal, but using them together achieves the best results. Also, we surprisingly find that the results by SimTxtSeg even surpass the pseudo masks by TVCC+SAM used for weakly supervision.

4 Conclusion

This paper proposes an effective SimTxtSeg for weakly-supervised medical image segmentation via inputting simple text cues, which contains a textual-to-visual cue converter and a text-vision hybrid attention mechanism. Extensive experiments are conducted to prove that, using simple text cues, our approach achieves state-of-the-art performance with minimal supervision. In the future, we will extend our method to more medical image analysis areas and fuse the TVCC and SAM into an end-to-end fashion for improvement.

Acknowledgments. This work was partially supported by the National Natural Science Foundation of China (Grants No 62106043, 62172228), and the Natural Science Foundation of Jiangsu Province (Grants No BK20210225).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ali, S., Jha, D., Ghatwary, N., Realdon, S., Cannizzaro, R., Salem, O.E., Lamarque, D., Daul, C., Riegler, M.A., Anonsen, K.V., et al.: A multi-centre polyp detection and segmentation dataset for generalisability assessment. *Scientific Data* **10**(1), 75 (2023)
2. Bernal, J., Sánchez, F.J., Fernández-Esparrach, G., Gil, D., Rodríguez, C., Vilaríño, F.: Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized medical imaging and graphics* **43**, 99–111 (2015)
3. Buda, M., Saha, A., Mazurowski, M.A.: Association of genomic subtypes of lower-grade gliomas with shape features automatically extracted by a deep learning algorithm. *Computers in biology and medicine* **109**, 218–225 (2019)
4. Deng, G., Zou, K., Ren, K., Wang, M., Yuan, X., Ying, S., Fu, H.: Sam-u: Multi-box prompts triggered uncertainty estimation for reliable sam in medical image. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 368–377. Springer (2023)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
6. Diakogiannis, F.I., Waldner, F., Caccetta, P., Wu, C.: Resunet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing* **162**, 94–114 (2020)
7. Fan, D.P., Ji, G.P., Zhou, T., Chen, G., Fu, H., Shen, J., Shao, L.: Pranel: Parallel reverse attention network for polyp segmentation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 263–273. Springer (2020)
8. Gama, P.H., Oliveira, H., dos Santos, J.A.: Learning to segment medical images from few-shot sparse labels. In: *2021 34th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*. pp. 89–96. IEEE (2021)

9. Gröger, M., Borisov, V., Kasneci, G.: Boxshrink: From bounding boxes to segmentation masks. In: Workshop on Medical Image Learning with Limited and Noisy Data. pp. 65–75. Springer (2022)
10. Hu, X., Chen, Y.J., Ho, T.Y., Shi, Y.: Conditional diffusion models for weakly supervised medical image segmentation. arXiv preprint arXiv:2306.03878 (2023)
11. Jha, D., Smedsrud, P.H., Riegler, M.A., Halvorsen, P., de Lange, T., Johansen, D., Johansen, H.D.: 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. pp. 451–462. Springer (2020)
12. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
13. Lei, W., Wei, X., Zhang, X., Li, K., Zhang, S.: Medlsam: Localize and segment anything model for 3d medical images. arXiv preprint arXiv:2306.14752 (2023)
14. Li, Z., Zheng, Y., Luo, X., Shan, D., Hong, Q.: Scribblevc: Scribble-supervised medical image segmentation with vision-class embedding (2023)
15. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023)
16. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
17. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024)
18. MMDetection Contributors: OpenMMLab Detection Toolbox and Benchmark (Aug 2018), <https://github.com/open-mmlab/mmdetection>
19. Roth, H.R., Yang, D., Xu, Z., Wang, X., Xu, D.: Going to extremes: weakly supervised medical image segmentation. *Machine Learning and Knowledge Extraction* **3**(2), 507–524 (2021)
20. Silva, J., Histace, A., Romain, O., Dray, X., Granado, B.: Toward embedded detection of polyps in wce images for early diagnosis of colorectal cancer. *International journal of computer assisted radiology and surgery* **9**, 283–293 (2014)
21. Tajbakhsh, N., Gurudu, S.R., Liang, J.: Automated polyp detection in colonoscopy videos using shape and context information. *IEEE transactions on medical imaging* **35**(2), 630–644 (2015)
22. Wang, A., Xu, M., Zhang, Y., Islam, M., Ren, H.: S²me: Spatial-spectral mutual teaching and ensemble learning for scribble-supervised polyp segmentation. arXiv preprint arXiv:2306.00451 (2023)
23. Wang, C., Zhang, D., Yan, R.: Boosting weakly-supervised image segmentation via representation, transform, and compensator. arXiv preprint arXiv:2309.00871 (2023)
24. Wang, R., Lei, T., Cui, R., Zhang, B., Meng, H., Nandi, A.K.: Medical image segmentation using deep learning: A survey. *IET Image Processing* **16**(5), 1243–1267 (2022)
25. Wang, Z., Voiculescu, I.: Weakly supervised medical image segmentation through dense combinations of dense pseudo-labels. In: MICCAI Workshop on Data Engineering in Medical Imaging. pp. 1–10. Springer (2023)
26. Wei, J., Hu, Y., Cui, S., Zhou, S.K., Li, Z.: Weakpolyp: You only look bounding box for polyp segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 757–766. Springer (2023)

27. Wei, J., Hu, Y., Li, G., Cui, S., Kevin Zhou, S., Li, Z.: Boxpolyp: Boost generalized polyp segmentation using extra coarse bounding box annotations. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 67–77. Springer (2022)
28. Woo, S., Debnath, S., Hu, R., Chen, X., Liu, Z., Kweon, I.S., Xie, S.: Convnext v2: Co-designing and scaling convnets with masked autoencoders. In: CVPR. pp. 16133–16142 (June 2023)
29. Xie, X., Fan, H., Yu, Z., Bai, H., Tang, Y.: Weakly-supervised medical image segmentation based on multi-task learning. In: International Conference on Intelligent Robotics and Applications. pp. 395–404. Springer (2022)
30. Xu, Y., Gong, M., Xie, S., Batmanghelich, K.: Box-adapt: Domain-adaptive medical image segmentation using bounding boxsupervision (2021)
31. Ye, J., Cheng, J., Chen, J., Deng, Z., Li, T., Wang, H., Su, Y., Huang, Z., Chen, J., Jiang, L., et al.: Sa-med2d-20m dataset: Segment anything in 2d medical imaging with 20 million masks. arXiv preprint arXiv:2311.11969 (2023)
32. Zhong, Y., Xu, M., Liang, K., Chen, K., Wu, M.: Ariadne’s thread: Using text prompts to improve segmentation of infected areas from chest x-ray images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 724–733. Springer (2023)