

Towards Accurate and Efficient Pan-cancer Screening

Linshan Wu¹[0000–0002–0486–184X] and Jiaxin Zhuang¹[0000–0001–9287–4263]

The Hong Kong University of Science and Technology, Hong kong, China
{wulinshan503503}@gmail.com

Abstract. Pan-cancer screening is challenging due to the difficulty of localizing tiny tumors within large CT volumes. The efficiency also plays an important role in large-scale routine physical examination. To this end, we propose an accurate and efficient framework for pan-cancer screening. Our model employs a classification model to localize the diseased regions and a segmentation model to precisely segment the lesions. Specifically, we first crop a group of sub-volumes from the whole CT volume and select those with lesions for segmentation. We lead the MICCAI FLARE25 pan-cancer challenge, with 59% DSC and 52% NSD on the validation leaderboard.

Keywords: Pan-cancer Screening

1 Introduction

Cancer is the leading cause of death. AI-driven screening using large-scale computed tomography (CT) scans [14,32,31,28,37,17] has gained growing clinical interest, owing to the wide availability and low cost of CT in routine health examinations. In particular, pan-cancer screening seeks to develop a universal model capable of detecting and segmenting diverse lesions across large CT volumes, offering substantial potential for clinical impact.

Although recent studies have achieved promising results in cancer detection [14,4,13,21,33,34], developing effective pan-cancer screening models remains challenging [19,2]. Current approaches typically train segmentation models to identify cancerous regions. However, lesions often occupy only minute regions within large CT volumes, leading to severe foreground-background imbalance. This issue complicates the accurate localization of diverse lesion types across multiple organs.

In contrast to AI models, radiologists typically scan an entire CT volume before concentrating on specific regions for detailed analysis. Inspired by this diagnostic strategy, we propose a novel two-stage framework for pan-cancer screening. Our approach first uses a classification model to coarsely localize potentially diseased regions, then employs a segmentation model for precise lesion delineation, with both components working synergistically. Specifically, candidate sub-volumes are extracted from the full CT scan, and the system learns to identify those containing lesions for further segmentation. This design ensures

both high accuracy and efficiency in pan-cancer screening. Notably, our method achieved leading results in the MICCAI FLARE25 challenge, with scores of 59% DSC and 52% NSD, significantly outperforming the FLARE 2024 champion solution.

2 Method

2.1 Preprocessing

All experiments were implemented using PyTorch [22] and MONAI [5]. We used the AdamW optimizer with an initial learning rate of $3e-4$ and a cosine decay scheduler. A batch size of 4 was used, with 4 sub-volumes cropped from each CT volume. The model is trained for 500 epochs.

For pre-processing, we followed the FLARE 2025 guidelines by splitting the dataset into chest and abdomen CT scans. Hounsfield Unit (HU) values were windowed and normalized to the range $[0, 1]$. The clipping ranges were set to $[-900, 650]$ for chest CT and $[-175, 250]$ for abdomen CT. Each sub-volume was cropped to a size of $[96, 96, 64]$. During inference, we employed a sliding-window strategy with an overlap ratio of 0.25 to partition each CT volume into sub-volumes, in line with common practice.

Table 1. Pre-processing details and Training settings.

Clipped HU for chest CT	$[-900, 650]$
Clipped HU for abdomen CT	$[-175, 250]$
Crop Size	$[96, 96, 64]$
Spacing	$[1.0, 1.0, 3.0]$
Normalized value	$[0, 1]$
Segmentation model	SwinUNETR [8]
Network Parameters	104M
Segmentation model Loss	Dice-CE
Optimizer	AdamW
Batch size	4
Scheduler	Cosine
Learning rate	$3e-4$
Training epochs	500

2.2 Proposed Method

We did not use the partial labels. Unlabeled images were not used. We did not use the pseudo labels generated by the FLARE23 winning algorithm.

Segmentation [14,17,25,29,30,27,26,16,9,18] plays a vital role in medical image analysis. Nevertheless, in whole-body CT scans, lesions often occupy only




Fig. 1. Network architecture. We use SwinUNETR [8] as the segmentation model.

tiny regions, resulting in extreme foreground-background imbalance and making it difficult to accurately localize diverse lesions across multiple organs.

To address these challenges, we propose a two-stage framework that combines efficiency with accuracy. Our approach consists of a classification module that identifies whether a sub-volume contains lesions, and a segmentation module that precisely delineates lesion boundaries—both of which can incorporate state-of-the-art network architectures. In this work, we use a lightweight 3D ResNet-18 as the classifier and SwinUNETR [8] as the segmentation model.

During training, randomly cropped sub-volumes are processed by both models. The segmentation model is trained with a combination of binary cross-entropy and Dice loss, supervised by ground-truth masks. The classifier is trained using segmentation results as pseudo-labels. During inference, the input volume is divided into sub-volumes via a sliding window. The classifier first filters out healthy regions, and only candidate sub-volumes containing lesions are forwarded to the segmentation module. This strategy introduces minimal overhead but significantly reduces computation by avoiding redundant processing of healthy areas. Consequently, our method not only improves inference efficiency but also reduces false positives by focusing primarily on diseased regions.

2.3 Post-processing

No post-processing technique is used.

3 Experiments

3.1 Dataset and evaluation measures

The segmentation targets cover various lesions. The training dataset is curated from more than 50 medical centers under the license permission, including TCIA [6], LiTS [3], MSD [24], KiTS [10,12,11], and FLARE 2023 [20], COVID-19-20 [23], CHOS [15], and LIDC [1]. The training set includes more than 5,000+ abdomen and chest CT scans. The validation and testing sets cover various cancer types. The lesion annotation process used ITK-SNAP [36], nnU-Net [14], MedSAM [17,18], and Slicer Plugins [7,18].

The evaluation metrics encompass two accuracy measures—Dice Similarity Coefficient (DSC) and Normalized Surface Dice (NSD)—alongside two efficiency measures—running time and area under the GPU memory-time curve. These metrics collectively contribute to the ranking computation. Furthermore, the running time and GPU memory consumption are considered within tolerances of 45 seconds and 4 GB, respectively.

3.2 Implementation details

Environment settings The development environments and requirements are presented in Table 2.

Table 2. Development environments and requirements.

System	e.g., Ubuntu 18.04.5 LTS or Windows 11
CPU	e.g., Intel(R) Core(TM) i9-7900X CPU@3.30GHz
RAM	16×4GB; 2.67MT/s
GPU (number and type)	Four NVIDIA V100 16G
CUDA version	11.8
Programming language	e.g., Python 3.10
Deep learning framework	torch 2.0, torchvision 0.2.2

4 Results and discussion

Table 3. Quantitative evaluation results.)

Methods	Public Validation		Online Validation		Testing	
	DSC(%)	NSD(%)	DSC(%)	NSD(%)	DSC(%)	NSD (%)
Algorithm1	56.4±1.0	48.1±1.2	59	52		

4.1 Quantitative results on validation set

The results are shown in Table 3. We achieve 59% DSC and 52% NSD in the validation leaderboard.

4.2 Qualitative results on validation set

4.3 Results on final testing set

This is a placeholder. We will send you the testing results during MICCAI.




Fig. 2. Visualization results.

Table 4. Quantitative evaluation of segmentation efficiency in terms of the running them and GPU memory consumption. Total GPU denotes the area under GPU Memory-Time curve. Evaluation GPU platform: NVIDIA QUADRO RTX5000 (16G). Note: if you didn’t make validation docker submissions during the challenge, you can obtain these metrics on your local GPU. Our evaluation code has been released at <https://github.com/JunMa11/FLARE/tree/main/FLARE23>. You could use a similar table format for ablation studies. Please don’t change the case IDs. The organizer provided the results of these two cases to us.

Case ID	Image Size	Running Time (s)	Max GPU (MB)	Total GPU (MB)
0001	(512, 512, 55)	31	14044	-
0051	(512, 512, 100)	23	14724	-
0017	(512, 512, 150)			
0019	(512, 512, 215)			
0099	(512, 512, 334)			
0063	(512, 512, 448)			
0048	(512, 512, 499)			
0029	(512, 512, 554)			

4.4 Limitation and future work

This is a brief technical report for the FLARE 2025 challenge, and we will extend it to a better method with more thorough discussions.

Acknowledgements The authors of this paper declare that the segmentation method they implemented for participation in the FLARE 2025 challenge has not used any pre-trained models nor additional datasets other than those provided by the organizers. The proposed solution is fully automatic without any manual intervention. We thank all data owners for making the CT scans publicly available and CodaLab [35] for hosting the challenge platform.

Disclosure of Interests

The authors declare no competing interests.

References

1. Armato III, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Zhao, B., Aberle, D.R., Henschke, C.I., Hoffman, E.A., et al.: The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics* **38**(2), 915–931 (2011) [3](#)
2. Bassi, P.R., Li, W., Tang, Y., Isensee, F., Wang, Z., Chen, J., Chou, Y.C., Kirchoff, Y., Rokuss, M.R., Huang, Z., et al.: Touchstone benchmark: Are we on the right way for evaluating ai algorithms for medical segmentation? *Advances in Neural Information Processing Systems* **37**, 15184–15201 (2024) [1](#)

3. Bilic, P., Christ, P., Li, H.B., Vorontsov, E., Ben-Cohen, A., Kaissis, G., Szeskin, A., Jacobs, C., Mamani, G.E.H., Chartrand, G., Lohöfer, F., Holch, J.W., Sommer, W., Hofmann, F., Hostettler, A., Lev-Cohain, N., Drozdal, M., Amitai, M.M., Vivanti, R., Sosna, J., Ezhov, I., Sekuboyina, A., Navarro, F., Kofler, F., Paetzold, J.C., Shit, S., Hu, X., Lipková, J., Rempfler, M., Piraud, M., Kirschke, J., Wiestler, B., Zhang, Z., Hülsemeyer, C., Beetz, M., Ettlinger, F., Antonelli, M., Bae, W., Bellver, M., Bi, L., Chen, H., Chlebus, G., Dam, E.B., Dou, Q., Fu, C.W., Georgescu, B., i Nieto, X.G., Gruen, F., Han, X., Heng, P.A., Hesser, J., Moltz, J.H., Igel, C., Isensee, F., Jäger, P., Jia, F., Kaluva, K.C., Khened, M., Kim, I., Kim, J.H., Kim, S., Kohl, S., Konopczynski, T., Kori, A., Krishnamurthi, G., Li, F., Li, H., Li, J., Li, X., Lowengrub, J., Ma, J., Maier-Hein, K., Maninis, K.K., Meine, H., Merhof, D., Pai, A., Perslev, M., Petersen, J., Pont-Tuset, J., Qi, J., Qi, X., Rippel, O., Roth, K., Sarasua, I., Schenk, A., Shen, Z., Torres, J., Wachinger, C., Wang, C., Weninger, L., Wu, J., Xu, D., Yang, X., Yu, S.C.H., Yuan, Y., Yue, M., Zhang, L., Cardoso, J., Bakas, S., Braren, R., Heinemann, V., Pal, C., Tang, A., Kadoury, S., Soler, L., van Ginneken, B., Greenspan, H., Joskowicz, L., Menze, B.: The liver tumor segmentation benchmark (lits). *Medical Image Analysis* **84**, 102680 (2023) [3](#)
4. Cao, K., Xia, Y., Yao, J., Han, X., Lambert, L., Zhang, T., Tang, W., Jin, G., Jiang, H., Fang, X., et al.: Large-scale pancreatic cancer detection via non-contrast ct and deep learning. *Nature Medicine* **29**(12), 3033–3043 (2023) [1](#)
5. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701* (2022) [2](#)
6. Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., Tarbox, L., Prior, F.: The cancer imaging archive (tcia): maintaining and operating a public information repository. *Journal of Digital Imaging* **26**(6), 1045–1057 (2013) [3](#)
7. Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., et al.: 3d slicer as an image computing platform for the quantitative imaging network. *Magnetic Resonance Imaging* **30**(9), 1323–1341 (2012) [3](#)
8. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: *International MICCAI brainlesion workshop*. pp. 272–284. Springer (2021) [2](#), [3](#)
9. He, S., Nie, Y., Wang, H., Yang, S., Wang, Y., Cai, Z., Chen, Z., Xu, Y., Luo, L., Xiang, H., et al.: Gsco: Towards generalizable ai in medicine via generalist-specialist collaboration. *arXiv preprint arXiv:2404.15127* (2024) [2](#)
10. Heller, N., Isensee, F., Maier-Hein, K.H., Hou, X., Xie, C., Li, F., Nan, Y., Mu, G., Lin, Z., Han, M., Yao, G., Gao, Y., Zhang, Y., Wang, Y., Hou, F., Yang, J., Xiong, G., Tian, J., Zhong, C., Ma, J., Rickman, J., Dean, J., Stai, B., Tejpaul, R., Oestreich, M., Blake, P., Kaluzniak, H., Raza, S., Rosenberg, J., Moore, K., Walczak, E., Rengel, Z., Edgerton, Z., Vasdev, R., Peterson, M., McSweeney, S., Peterson, S., Kalapara, A., Sathianathen, N., Papanikolopoulos, N., Weight, C.: The state of the art in kidney and kidney tumor segmentation in contrast-enhanced ct imaging: Results of the kits19 challenge. *Medical Image Analysis* **67**, 101821 (2021) [3](#)
11. Heller, N., Isensee, F., Trofimova, D., Tejpaul, R., Zhao, Z., Chen, H., Wang, L., Golts, A., Khapun, D., Shats, D., Shoshan, Y., Gilboa-Solomon, F., George, Y., Yang, X., Zhang, J., Zhang, J., Xia, Y., Wu, M., Liu, Z., Walczak, E., McSweeney,

- S., Vasdev, R., Hornung, C., Solaiman, R., Schoepfoerster, J., Abernathy, B., Wu, D., Abdulkadir, S., Byun, B., Spriggs, J., Struyk, G., Austin, A., Simpson, B., Hagstrom, M., Virnig, S., French, J., Venkatesh, N., Chan, S., Moore, K., Jacobsen, A., Austin, S., Austin, M., Regmi, S., Papanikolopoulos, N., Weight, C.: The kits21 challenge: Automatic segmentation of kidneys, renal tumors, and renal cysts in corticomedullary-phase ct. *arXiv preprint arXiv:2307.01984* (2023) [3](#)
12. Heller, N., McSweeney, S., Peterson, M.T., Peterson, S., Rickman, J., Stai, B., Tejapaul, R., Oestreich, M., Blake, P., Rosenberg, J., et al.: An international challenge to use artificial intelligence to define the state-of-the-art in kidney and kidney tumor segmentation in ct imaging. *American Society of Clinical Oncology* **38**(6), 626–626 (2020) [3](#)
13. Hu, C., Xia, Y., Zheng, Z., Cao, M., Zheng, G., Chen, S., Sun, J., Chen, W., Zheng, Q., Pan, S., et al.: Ai-based large-scale screening of gastric cancer from noncontrast ct imaging. *Nature Medicine* pp. 1–9 (2025) [1](#)
14. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021) [1](#), [2](#), [3](#)
15. Kavur, A.E., Gezer, N.S., Barış, M., Aslan, S., Conze, P.H., Groza, V., Pham, D.D., Chatterjee, S., Ernst, P., Özkan, S., Baydar, B., Lachinov, D., Han, S., Pauli, J., Isensee, F., Perkonig, M., Sathish, R., Rajan, R., Sheet, D., Dovletov, G., Speck, O., Nürnberger, A., Maier-Hein, K.H., Bozdağı Akar, G., Ünal, G., Dicle, O., Selver, M.A.: CHAOS Challenge - combined (CT-MR) healthy abdominal organ segmentation. *Medical Image Analysis* **69**, 101950 (2021) [3](#)
16. Liu, Q., He, M., Kuang, Y., Wu, L., Yue, J., Fang, L.: A multi-level label-aware semi-supervised framework for remote sensing scene classification. *IEEE Transactions on Geoscience and Remote Sensing* **61**, 1–12 (2023) [2](#)
17. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024) [1](#), [2](#), [3](#)
18. Ma, J., Yang, Z., Kim, S., Chen, B., Baharoon, M., Fallahpour, A., Asakereh, R., Lyu, H., Wang, B.: Medsam2: Segment anything in 3d medical images and videos. *arXiv preprint arXiv:2504.03600* (2025) [2](#), [3](#)
19. Ma, J., Zhang, Y., Gu, S., Ge, C., Mae, S., Young, A., Zhu, C., Yang, X., Meng, K., Huang, Z., et al.: Unleashing the strengths of unlabelled data in deep learning-assisted pan-cancer abdominal organ quantification: the flare22 challenge. *The Lancet Digital Health* **6**(11), e815–e826 (2024) [1](#)
20. Ma, J., Zhang, Y., Gu, S., Ge, C., Wang, E., Zhou, Q., Huang, Z., Lyu, P., He, J., Wang, B.: Automatic organ and pan-cancer segmentation in abdomen ct: the flare 2023 challenge. *arXiv preprint arXiv:2408.12534* (2024) [3](#)
21. Ni, X., Wu, L., Zhuang, J., Wang, Q., Wu, M., Vardhanabhuti, V., Zhang, L., Gao, H., Chen, H.: Mg-3d: Multi-grained knowledge-enhanced 3d medical vision-language pre-training. *arXiv preprint arXiv:2412.05876* (2024) [1](#)
22. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* **32** (2019) [2](#)
23. Roth, H.R., Xu, Z., Tor-Díez, C., Jacob, R.S., Zember, J., Molto, J., Li, W., Xu, S., Turkbey, B., Turkbey, E., et al.: Rapid artificial intelligence solutions in a pandemic—the covid-19-20 lung ct lesion segmentation challenge. *Medical Image Analysis* **82**, 102605 (2022) [3](#)

24. Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., van Ginneken, B., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., Bilic, P., Christ, P.F., Do, R.K.G., Gollub, M., Golia-Pernicka, J., Heckers, S.H., Jarnagin, W.R., McHugo, M.K., Napel, S., Vorontsov, E., Maier-Hein, L., Cardoso, M.J.: A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *arXiv preprint arXiv:1902.09063* (2019) [3](#)
25. Wu, L., Fang, L., He, X., He, M., Ma, J., Zhong, Z.: Querying labeled for unlabeled: Cross-image semantic consistency guided semi-supervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(7), 8827–8844 (2023) [2](#)
26. Wu, L., Fang, L., Yue, J., Zhang, B., Ghamisi, P., He, M.: Deep bilateral filtering network for point-supervised semantic segmentation in remote sensing images. *IEEE Transactions on Image Processing* **31**, 7419–7434 (2022) [2](#)
27. Wu, L., Lu, M., Fang, L.: Deep covariance alignment for domain adaptive remote sensing image segmentation. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–11 (2022) [2](#)
28. Wu, L., Nie, Y., He, S., Zhuang, J., Luo, L., Mahboobani, N., Vardhanabhuti, V., Chan, R.C.K., Peng, Y., Rajpurkar, P., et al.: Unibiomed: A universal foundation model for grounded biomedical image interpretation. *arXiv preprint arXiv:2504.21336* (2025) [1](#)
29. Wu, L., Zhong, Z., Fang, L., He, X., Liu, Q., Ma, J., Chen, H.: Sparsely annotated semantic segmentation with adaptive gaussian mixtures. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15454–15464 (2023) [2](#)
30. Wu, L., Zhong, Z., Ma, J., Wei, Y., Chen, H., Fang, L., Li, S.: Modeling the label distributions for weakly-supervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025) [2](#)
31. Wu, L., Zhuang, J., Chen, H.: Large-scale 3d medical image pre-training with geometric context priors. *arXiv preprint arXiv:2410.09890* (2024) [1](#)
32. Wu, L., Zhuang, J., Chen, H.: Voco: A simple-yet-effective volume contrastive learning framework for 3d medical image analysis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 22873–22882 (2024) [1](#)
33. Wu, L., Zhuang, J., Ni, X., Chen, H.: Freetumor: Advance tumor segmentation via large-scale tumor synthesis. *arXiv preprint arXiv:2406.01264* (2024) [1](#)
34. Wu, L., Zhuang, J., Zhou, Y., He, S., Ma, J., Luo, L., Wang, X., Ni, X., Zhong, X., Wu, M., et al.: Freetumor: Large-scale generative tumor synthesis in computed tomography images for improving tumor recognition. *arXiv preprint arXiv:2502.18519* (2025) [1](#)
35. Xu, Z., Escalera, S., Pavão, A., Richard, M., Tu, W.W., Yao, Q., Zhao, H., Guyon, I.: Codabench: Flexible, easy-to-use, and reproducible meta-benchmark platform. *Patterns* **3**(7), 100543 (2022) [6](#)
36. Yushkevich, P.A., Gao, Y., Gerig, G.: Itk-snap: An interactive tool for semi-automatic segmentation of multi-modality biomedical images. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. pp. 3342–3345 (2016) [3](#)
37. Zhuang, J., Wu, L., Wang, Q., Vardhanabhuti, V., Luo, L., Chen, H.: Mim: Mask in mask self-supervised pre-training for 3d medical image analysis. *arXiv preprint arXiv:2404.15580* (2024) [1](#)

Table 5. Checklist Table. Please fill out this checklist table in the answer column.

Requirements	Answer
A meaningful title	Yes
The number of authors (≤ 6)	2
Author affiliations and ORCID	Yes
Corresponding author email is presented	Yes
Validation scores are presented in the abstract	Yes
Introduction includes at least three parts: background, related work, and motivation	Yes
A pipeline/network figure is provided	Figure 1
Pre-processing	Page 2
Strategies to use the partial label	Did not use
Strategies to use the unlabeled images.	Did not use
Strategies to improve model inference	Page 3
Post-processing	Did not use
The dataset and evaluation metric section are presented	Page 4
Environment setting table is provided	Table 1
Training protocol table is provided	Table 1
Ablation study	Page 6
Efficiency evaluation results are provided	Table 4
Visualized segmentation example is provided	Figure 3
Limitation and future work are presented	Yes
Reference format is consistent.	Yes