

Efficient Whole-Body Tumor Segmentation with a 5.6M Parameter 3D U-Net

Ziyan Huang¹[0000–0002–1533–5239]

Shanghai Jiao Tong University {ziyanhuang}@sjtu.edu.cn

Abstract. Accurate organ and cancer segmentation in medical imaging, especially in 3D CT scans, is essential for precise diagnosis, treatment planning, and disease monitoring. The MICCAI FLARE24 challenge aims to advance pan-cancer segmentation algorithms, with Task 1 focusing on whole-body cancer segmentation in CT scans. In this paper, we present an efficient approach utilizing a lightweight 3D U-Net architecture to address this challenge. Our model comprises only four resolution stages and approximately 5.6 million parameters, significantly reducing computational demands while maintaining performance. Our method processed the 279 CT images in the public validation set in just 18 minutes, averaging under 4 seconds per scan. For individual cases, the average prediction time per scan was approximately 20 seconds. We achieved an average Dice Similarity Coefficient (DSC) of 25.34% and a Normalized Surface Dice (NSD) of 24.40% on the public validation set.

1 Introduction

Whole-body tumor segmentation from CT scans is a critical yet challenging task in medical image analysis. The MICCAI FLARE24 challenge, specifically Task 1, aims to advance the field of pan-cancer segmentation by addressing this complex problem. The difficulty of this challenge arises from several factors:

- **Anatomical Diversity:** The vast anatomical diversity across different patients and the varied appearance of tumors in different body regions make it challenging to develop a universally effective segmentation model.
- **Computational Demands:** Whole-body CT scans are large, often containing thousands of slices. Processing these scans in real-time or near-real-time settings poses significant computational challenges.
- **Partial Labeling:** The dataset’s partial labeling nature, where not all tumors are annotated in every scan, adds another layer of complexity to the learning process.

Recent advancements in deep learning have led to significant progress in medical image segmentation. In the context of semi-supervised and partial-label segmentation, several approaches have shown promise:

- **Self-training Methods:** These methods iteratively use the model’s own predictions on unlabeled data to augment the training set, gradually improving performance [11].

- **Consistency-based Approaches:** These leverage the principle that model predictions should remain consistent under different perturbations of the input or model [25].
- **Adversarial Learning:** Some methods employ adversarial training to align the distributions of labeled and unlabeled data, enhancing generalization [23].
- **Multi-task Learning:** By combining segmentation with auxiliary tasks like reconstruction or classification, these methods can better utilize partially labeled data [3].

While these approaches have achieved success in various medical imaging tasks, they often require significant computational resources, limiting their applicability in real-world clinical scenarios.

Motivated by the need for an efficient yet accurate solution to whole-body tumor segmentation, we propose a lightweight 3D U-Net architecture. Our approach is guided by the following key considerations:

- **Efficiency:** In clinical settings, fast processing times are crucial. We aim to develop a model that can segment whole-body CT scans in near real-time while maintaining high accuracy.
- **Resource Constraints:** Many healthcare facilities may lack access to high-end GPU resources. Our lightweight model addresses this by minimizing computational requirements.
- **Generalizability:** Given the diverse nature of whole-body tumors, we focus on creating a model that can generalize well across different anatomical regions and tumor types.

Our main contributions are:

- (1) A compact 3D U-Net architecture with only 5.6 million parameters, significantly reducing the model size compared to standard implementations.
- (2) An efficient segmentation approach that balances accuracy and speed, processing whole-body CT scans in approximately 10 seconds.
- (3) Comprehensive evaluation on a large-scale dataset of over 10,000 CT scans, demonstrating the model’s effectiveness across diverse tumor types and anatomical locations.
- (4) Insights into the trade-offs between model complexity and performance in the context of whole-body tumor segmentation, potentially guiding future research in this area.

By addressing the challenges of whole-body tumor segmentation with a focus on efficiency and practicality, our work aims to bridge the gap between advanced AI models and real-world clinical applications.

2 Method

2.1 Preprocessing

We utilized the default preprocessing pipeline of **nnU-Net** [12]. The preprocessing steps included:

- **Dataset Selection:** Only images with tumor annotations were used for training.
- **Resampling:** All CT scans were resampled to an isotropic voxel spacing of $2 \times 2 \times 3 \text{ mm}^3$.
- **Intensity Normalization:** Hounsfield Unit (HU) values were clipped at the 0.05% and 99.5% percentiles and then normalized by subtracting the mean and dividing by the standard deviation, based on statistics computed from the entire dataset.

2.2 Proposed Method

Figure 1 illustrates the architecture of our proposed model in comparison to deeper variants. Unlike conventional U-Net or nnU-Net structures, which typically employ 5 or even 6 resolution stages, our model adopts a shallower 4-stage design. As the later stages in U-Net architectures generally involve high channel dimensions and incur substantial parameter costs, our streamlined structure significantly reduces the total number of parameters without compromising segmentation performance. Apart from the number of stages, the rest of our model architecture remains aligned with the standard nnU-Net framework.




Fig. 1. Architectures with different stage numbers: (a) our 4-stage model, (b) a 5-stage variant, and (c) the 6-stage nnU-Net. More stages lead to deeper models and higher complexity.

Loss function: we use the summation between Dice loss and cross-entropy loss because compound loss functions have proven robust in various medical image segmentation tasks [14].

We did not employ specific strategies to reduce false positives on CT scans from healthy patients, handle partial labels, or use unlabeled images. We only used the provided labeled data without distinguishing whether it was partially labeled, and unlabeled images were not used. Additionally, we did not use the pseudo labels generated by the FLARE23 winning algorithm.

We used a smaller model architecture and resampled the images to a relatively low resolution of $2 \times 2 \times 3 \text{ mm}^3$, which helped to improve inference speed and reduce resource consumption during the inference phase.

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 [4], LiTS [2], MSD [22], KiTS [8,10,9], autoPET [7,6], TotalSegmentator [24], and AbdomenCT-1K [19], FLARE 2023 [18], DeepLesion [27], COVID-19-CT-Seg-Benchmark [17], COVID-19-20 [21], CHOS [13], LNDB [20], and LIDC [1]. The training set includes 4000 abdomen CT scans where 2200 CT scans with partial labels and 1800 CT scans without labels. The validation and testing sets include 100 and 400 CT scans, respectively, which cover various abdominal cancer types, such as liver cancer, kidney cancer, pancreas cancer, colon cancer, gastric cancer, and so on. The lesion annotation process used ITK-SNAP [28], nnU-Net [12], MedSAM [15], and Slicer Plugins [5,16].

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 1.

Table 1. Development environments and requirements.

System version	CentOS Linux release 7.6.1810
CPU	Dual AMD Rome 7742@3.4GHz
RAM	32×32GB; 3200MT/s
GPU (number and type)	1x NVIDIA A100 80GB Tensor Core GPUs
CUDA version	11.2
Programming language	Python 3.8.0
Deep learning framework	Pytorch (Torch 1.10.1)
Specific dependencies	nnU-Net 2.2.0
Code	https://github.com/Ziyan-Huang/FLARE24

Training Protocols Our training protocols followed the default settings of nnU-Net.

1. Specifically, no specific strategies were applied for the processing of unlabeled images and partial labels; we used only the provided labeled data without distinguishing between partially labeled or unlabeled data.

2. For data augmentation, we employed the default extensive augmentation techniques as implemented in nnU-Net.

3. The patch sampling strategy was also based on the default configuration of nnU-Net.

4. Additionally, we did not perform specific model selection, and training continued until the final epoch without early stopping or optimal model selection.

Table 2. Training protocols.

Network initialization	He
Batch size	2
Patch size	$128 \times 128 \times 96$
Total epochs	2000
Optimizer	SGD
Initial learning rate (lr)	0.001
Lr decay schedule	poly
Training time	20 hours
Loss function	Dice plus CE
Number of model parameters	5.6M ¹
Number of flops	59.32G ²
CO ₂ eq	1 Kg ³

4 Results and discussion

Table 3. Quantitative evaluation results.

Methods	Public Validation		Online Validation		Testing	
	DSC(%)	NSD(%)	DSC(%)	NSD(%)	DSC(%)	NSD (%)
Ours	25.34 \pm 31.56	24.40 \pm 27.80	–	–	43.11 \pm 39.98	35.67 \pm 35.76

4.1 Quantitative results on validation set

Quantitative results are shown in Table 3. On the public validation set, our method achieved a mean Dice of 25.34% and NSD of 24.40%. On the hidden test set, performance improves to 43.11% Dice and 35.67% NSD, showing good generalization. Among 40 healthy cases, 16 showed false positives, mostly in ambiguous soft tissue regions.

4.2 Qualitative results on validation set

Figure 2 shows two successful and two failed tumor segmentation cases. In the good cases (top two rows), the model accurately captures the tumor region, achieving high Dice scores (e.g., 0.9451). In contrast, the failed cases (bottom two rows) show missed detections or incomplete coverage of tumor areas, with Dice scores below 0.3. These failures are often due to small tumor size, low contrast, or surrounding complex textures.

4.3 Segmentation efficiency results on validation set

Table 4 shows inference efficiency on representative cases. Average runtime is around 22 seconds per scan, with peak GPU usage below 5.2 GB, showing good scalability.

Table 4. Quantitative evaluation of segmentation efficiency in terms of the running time and GPU memory consumption. Total GPU denotes the area under GPU Memory-Time curve.

Case ID	Image Size	Running Time (s)	Max GPU (MB)	Total GPU (MB)
0001	(512, 512, 55)	28.51	4696	141667
0051	(512, 512, 100)	18.85	4916	92644
0017	(512, 512, 150)	18.87	5170	97570
0019	(512, 512, 215)	19.47	4930	95986
0099	(512, 512, 334)	21.82	5177	112956
0063	(512, 512, 448)	18.16	4776	86711
0048	(512, 512, 499)	25.17	5017	126278
0029	(512, 512, 554)	24.57	4639	113978

4.4 Results on final testing set

Table 5 summarizes the final testing results of our method submitted to the FLARE23 challenge. The metrics include Dice Similarity Coefficient (DSC), Normalized Surface Dice (NSD), inference time, and GPU memory usage. Our method achieved competitive segmentation accuracy while maintaining relatively efficient runtime and resource consumption.

4.5 Limitation and Future Work

Despite achieving efficiency, our model has several limitations:

- (1) **Accuracy:** The DSC and NSD scores suggest there is room to improve segmentation accuracy, particularly for small or irregular tumors.

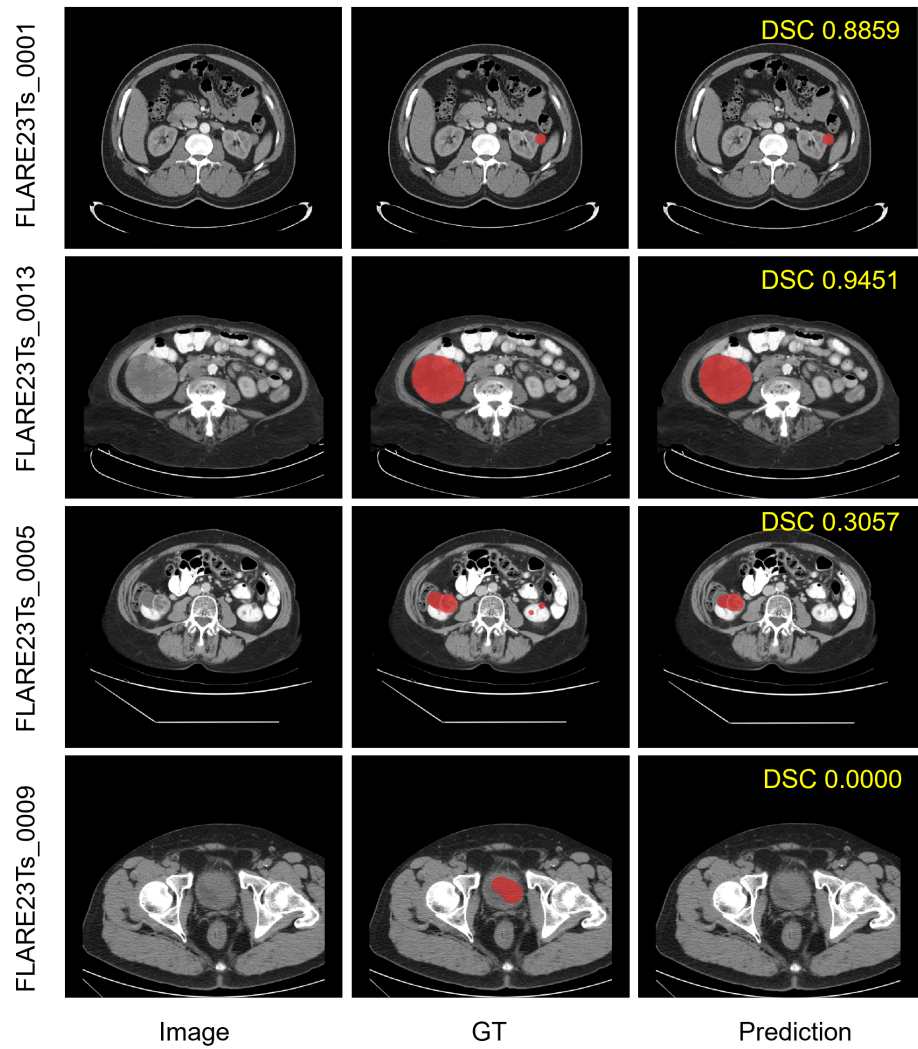


Fig. 2. Two examples with good segmentation results and two examples with bad segmentation results in the validation set.

Table 5. Final testing results on the FLARE23 challenge (team: gmai).

Avg DSC	43.11 ± 39.98
Median DSC	47.11 (0.00, 85.15)
Avg NSD	35.67 ± 35.76
Median NSD	30.68 (0.00, 70.06)
Avg Time (s)	32.46 ± 14.32
Median Time (s)	28.39 (20.62, 40.23)
Avg GPU (MB)	43451.6 ± 19941.9
Median GPU (MB)	39401.5 (28946.0, 57435.5)

- (2) **Partial Labels:** We did not utilize unlabeled or partially labeled data, limiting generalization. Future work will incorporate semi-supervised learning or pseudo-labeling.
- (3) **False Positives:** High false positive rates in healthy cases reduce clinical reliability. Advanced post-processing methods will be explored.

5 Conclusion

We presented a 5.6M parameter 3D U-Net model for efficient whole-body tumor segmentation. The model achieved a DSC of 25.34% and NSD of 24.40%, with an inference time of less than 4 seconds per scan, highlighting its suitability for near real-time applications. Key contributions include reducing model complexity while maintaining a balance between speed and accuracy. Future work will focus on improving segmentation precision and expanding the model’s applicability across diverse clinical settings.

Acknowledgements The authors of this paper declare that the segmentation method they implemented for participation in the FLARE 2024 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 [26] 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) [4](#)
2. 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) [4](#)
 3. Chen, S., Bortsova, G., García-Uceda Juárez, A., Van Tulder, G., De Bruijne, M.: Multi-task attention-based semi-supervised learning for medical image segmentation. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III* 22. pp. 457–465. Springer (2019) [2](#)
 4. 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) [4](#)
 5. 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) [4](#)
 6. Gatidis, S., Früh, M., Fabritius, M., Gu, S., Nikolaou, K., Fougère, C.L., Ye, J., He, J., Peng, Y., Bi, L., Ma, J., Wang, B., Zhang, J., Huang, Y., Heiliger, L., Marinov, Z., Stiefelhagen, R., Egger, J., Kleesiek, J., Sibille, L., Xiang, L., Bendazolli, S., Astaraki, M., Schölkopf, B., Ingrisch, M., Cyran, C., Küstner, T.: The autopet challenge: towards fully automated lesion segmentation in oncologic pet/ct imaging. *Nature Machine Intelligence* (2023) [4](#)
 7. Gatidis, S., Hepp, T., Früh, M., La Fougère, C., Nikolaou, K., Pfannenberger, C., Schölkopf, B., Küstner, T., Cyran, C., Rubin, D.: A whole-body fdg-pet/ct dataset with manually annotated tumor lesions. *Scientific Data* **9**(1), 601 (2022) [4](#)
 8. 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., Tejapaul, 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., Sathianathan, 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) [4](#)
9. 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) [4](#)
 10. Heller, N., McSweeney, S., Peterson, M.T., Peterson, S., Rickman, J., Stai, B., Tejpaul, 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) [4](#)
 11. Huang, Z., Wang, H., Ye, J., Niu, J., Tu, C., Yang, Y., Du, S., Deng, Z., Gu, L., He, J.: Revisiting nnu-net for iterative pseudo labeling and efficient sliding window inference. In: *MICCAI Challenge on Fast and Low-Resource Semi-supervised Abdominal Organ Segmentation*. pp. 178–189. Springer (2022) [1](#)
 12. 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) [2](#), [4](#)
 13. 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) [4](#)
 14. Ma, J., Chen, J., Ng, M., Huang, R., Li, Y., Li, C., Yang, X., Martel, A.L.: Loss odyssey in medical image segmentation. *Medical Image Analysis* **71**, 102035 (2021) [3](#)
 15. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**, 654 (2024) [4](#)
 16. Ma, J., Kim, S., Li, F., Baharoon, M., Asakereh, R., Lyu, H., Wang, B.: Segment anything in medical images and videos: Benchmark and deployment. *arXiv preprint arXiv:2408.03322* (2024) [4](#)
 17. Ma, J., Wang, Y., An, X., Ge, C., Yu, Z., Chen, J., Zhu, Q., Dong, G., He, J., He, Z., Cao, T., Zhu, Y., Nie, Z., Yang, X.: Towards data-efficient learning: A benchmark for covid-19 ct lung and infection segmentation. *Medical Physics* **48**(3), 1197–1210 (2021) [4](#)
 18. 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) [4](#)
 19. Ma, J., Zhang, Y., Gu, S., Zhu, C., Ge, C., Zhang, Y., An, X., Wang, C., Wang, Q., Liu, X., Cao, S., Zhang, Q., Liu, S., Wang, Y., Li, Y., He, J., Yang, X.: Abdomenct-1k: Is abdominal organ segmentation a solved problem? *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(10), 6695–6714 (2022) [4](#)
 20. Pedrosa, J., Aresta, G., Ferreira, C., Atwal, G., Phoulady, H.A., Chen, X., Chen, R., Li, J., Wang, L., Galdran, A., et al.: Lndb challenge on automatic lung cancer patient management. *Medical Image Analysis* **70**, 102027 (2021) [4](#)

21. 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) [4](#)
22. 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) [4](#)
23. Tang, Y., Wang, S., Qu, Y., Cui, Z., Zhang, W.: Consistency and adversarial semi-supervised learning for medical image segmentation. *Computers in Biology and Medicine* **161**, 107018 (2023) [2](#)
24. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., Bach, M., Segeroth, M.: Totalsegmentator: Robust segmentation of 104 anatomic structures in ct images. *Radiology: Artificial Intelligence* **5**(5), e230024 (2023) [4](#)
25. Xu, Z., Wang, Y., Lu, D., Yu, L., Yan, J., Luo, J., Ma, K., Zheng, Y., Tong, R.K.y.: All-around real label supervision: Cyclic prototype consistency learning for semi-supervised medical image segmentation. *IEEE Journal of Biomedical and Health Informatics* **26**(7), 3174–3184 (2022) [2](#)
26. 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) [8](#)
27. Yan, K., Wang, X., Lu, L., Summers, R.M.: Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *Journal of Medical Imaging* **5**(3), 036501–036501 (2018) [4](#)
28. 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) [4](#)