



Model Ensemble for Brain Tumor Segmentation in Magnetic Resonance Imaging

Daniel Capellán-Martín^{1,2}, Zhifan Jiang¹, Abhijeet Parida¹, Xinyang Liu¹,
Van Lam¹, Hareem Nisar¹, Austin Tapp¹, Sarah Elsharkawi^{1,3},
María J. Ledesma-Carbayo², Syed Muhammad Anwar^{1,4},
and Marius George Linguraru^{1,4}(✉)

¹ Sheikh Zayed Institute for Pediatric Surgical Innovation, Children's National Hospital, Washington, DC 20010, USA
mlingura@childrensnational.org

² Biomedical Image Technologies, ETSI Telecomunicación, Universidad Politécnica de Madrid, Madrid 28040, Spain and CIBER-BBN, ISCIII, Madrid, Spain
³ Princeton University, Princeton, NJ 08544, USA

⁴ School of Medicine and Health Sciences, George Washington University, Washington, DC 20052, USA

Abstract. Segmenting brain tumors in multi-parametric magnetic resonance imaging enables performing quantitative analysis in support of clinical trials and personalized patient care. This analysis provides the potential to impact clinical decision-making processes, including diagnosis and prognosis. In 2023, the well-established Brain Tumor Segmentation (BraTS) challenge presented a substantial expansion with eight tasks and 4,500 brain tumor cases. In this paper, we present a deep learning-based ensemble strategy that is evaluated for newly included tumor cases in three tasks: pediatric brain tumors (PED), intracranial meningioma (MEN), and brain metastases (MET). In particular, we ensemble outputs from state-of-the-art nnU-Net and Swin UNETR models on a region-wise basis. Furthermore, we implemented a targeted post-processing strategy based on a cross-validated threshold search to improve the segmentation results for tumor sub-regions. The evaluation of our proposed method on unseen test cases for the three tasks resulted in lesion-wise Dice scores for PED: 0.653, 0.809, 0.826; MEN: 0.876, 0.867, 0.849; and MET: 0.555, 0.6, 0.58; for the enhancing tumor, tumor core, and whole tumor, respectively. Our method was ranked first for PED, third for MEN, and fourth for MET, respectively.

Keywords: Brain tumor segmentation · MRI · Deep learning · Pediatric brain tumors · Meningioma · Metastases

D. Capellán-Martín, Z. Jiang and A. Parida—These authors contributed equally.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2024
U. Baid et al. (Eds.): crossMoDA 2023/BraTS 2023, LNCS 14669, pp. 221–232, 2024.
https://doi.org/10.1007/978-3-031-76163-8_20

1 Introduction

The brain tumor segmentation (BraTS) challenge held in conjunction with the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) conference, established in 2012, has generated a benchmark dataset for the segmentation of adult brain gliomas [4–7, 21]. The BraTS 2023 challenge has expanded to a cluster of challenges, encompassing a variety of tumor types alongside augmentation tasks [1, 10, 15, 17, 19, 25]. Herein, we propose a segmentation technique for newly introduced tasks featuring smaller datasets or new types of tumors. Particularly, these tasks include pediatric brain tumors (PED) [16], intracranial meningioma (MEN) [18], and brain metastasis (MET) [22]. This paper presents the methodology primarily developed for PED tumor segmentation, which was also adapted for MEN and MET segmentation tasks.

Brain cancer has become the leading cause of cancer death among children in the United States [8]. Although rare, pediatric high-grade brain tumors can be aggressive. For example, the median overall survival was reported to be less than one year for pediatric diffuse mid-line gliomas (DMGs, which replaced the formerly known diffuse intrinsic pontine gliomas (DIPGs) to emphasize that the disease may affect areas other than the pons) [24]. Multi-parametric magnetic resonance imaging (mpMRI) is essential in pediatric brain tumor diagnosis and monitoring of tumor progression. While adult and pediatric brain tumors share certain similarities, their locations in the brain and imaging characteristics may be different. For DMGs, necrosis is rare or unclear and the tumor may or may not present enhancement on post-gadolinium T1-weighted MRI. Hence, imaging tools specially designed to analyze pediatric brain tumors are necessary to improve clinical management. Automatic tumor segmentation is usually the first and most important step for the success of such analysis. Over the years, several methods have been presented for adult brain tumor segmentation using the BraTS dataset, however, efforts are needed to develop methods more suitable for pediatric brain tumors. The recently introduced PED task in BraTS 2023 provided an opportunity to develop such methods.

Although, a few previous works addressed pediatric brain tumors [10, 14, 20], their automatic segmentation remains challenging. This is due to a lack of training data and variation in heterogeneous histologic sub-regions including peritumoral edematous/invaded tissue, necrotic core, and enhancing tumor. BraTS-PED 2023 provides these types of data for the first benchmarking initiative on pediatric brain tumor segmentation. The task provides the largest annotated publicly-available retrospective cohort of high-grade gliomas including astrocytomas and DMG/DIPG in children.

In this work, we developed an ensemble approach involving two state-of-the-art deep learning models. Our approach to segmenting pediatric tumors is tested on two additional tasks at BraTS 2023: the segmentation of meningiomas (MEN) and brain metastases (MET). Meningioma is the most common primary intracranial tumor in adults and can result in significant morbidity and mortality for affected patients. Brain metastases are the most common form of

central nervous system (CNS) malignancy in adults. Accurate detection of small metastatic lesions is essential for patient prognosis, as missing even one lesion can lead to repeated interventions and treatment delays. For each task, training was performed only on the dataset provided by each sub-challenge.

2 Methods

The model ensemble technique is a widely recognized strategy used with machine learning methods and is aimed at increasing the stability and accuracy of model predictions. To harness the benefits inherent to both convolutional neural networks and vision transformers, we adopted an ensemble approach (see Fig. 1) involving two state-of-the-art models: nnU-Net (“no new U-Net”, winner of BraTS 2020) [13] and Swin UNETR (“Swin U-Net transformers”, top-performing model of BraTS 2021) [27]. Given that all tasks address the segmentation of multiple tumor sub-regions, our analysis and ensemble approach were conducted in a label-wise manner.

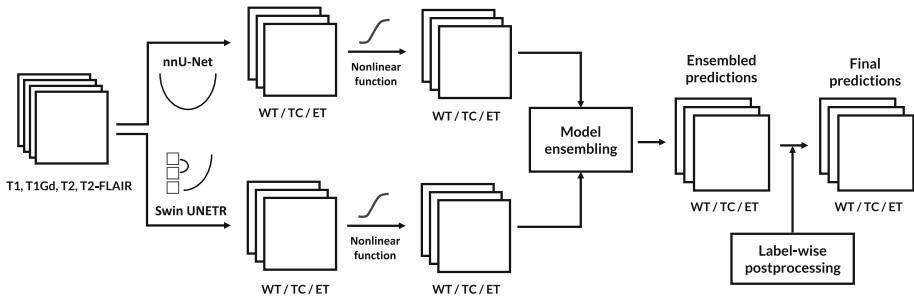


Fig. 1. Proposed method: model ensembling and post-processing pipeline. Outputs are obtained from two state-of-the-art deep learning models. These outputs are subjected to nonlinear activation functions and ensembling strategies. Finally, the ensembled predictions are subjected to a specifically tailored post-processing step.

2.1 Data Description

For all three sub-challenges, mpMRI data included pre- and post-gadolinium T1-weighted (T1 and T1CE), T2-weighted (T2), and T2-weighted fluid attenuated inversion recovery (T2-FLAIR) MRI. The mpMRI scans were pre-processed in a standardized fashion, including co-registration to the same anatomical template [25], re-sampling to an isotropic resolution, and skull-stripping. Each task provided three manual segmentation labels with slightly different definitions. For example, the labels for PED included enhancing tumor (ET), non-enhancing component (NCR - a combination of non-enhancing tumor, cystic component, and necrosis), and peritumoral edematous area (ED). The labels of MEN and

MET included enhancing tumor (ET), non-enhancing tumor core (NETC), and surrounding non-enhancing FLAIR hyperintensity (SNFH). Based on the three labels, three tumor sub-regions were defined to evaluate algorithm performance: enhancing tumor (ET), tumor core (TC) and whole tumor (WT). The TC region included ET and NCR labels, and WT included all three labels. Figure 2 shows samples of training examples for the three tasks.

Imaging data for PED are split into training ($n = 99$), validation ($n = 45$), and additional testing subsets. Imaging data for MEN are split into training ($n = 1,000$), validation ($n = 141$), and testing datasets [18]. Imaging data for MET are split into training ($n = 165$), validation ($n = 31$), and testing datasets [22].

2.2 Deep Learning Models and Experimental Setup

nnU-Net. The nnU-Net, which is based on the U-Net architecture [26], is a self-configuring deep learning framework for semantic segmentation. According to the specific imaging modality and unique attributes of each dataset, the framework autonomously adjusts its internal configurations [13]. This results in an improved segmentation performance and generalization when compared to other state-of-the-art methods for biomedical image segmentation.

For each of the three tasks, we trained a full-resolution 3D nnU-Net (v2) model using a five-fold cross-validation approach. A preprocessing consisting




Fig. 2. Training examples in the PED, MEN, and MET tasks (from top to bottom) with the following tumor subregions: enhancing tumor ET (blue), a combination of nonenhancing tumor, cystic component, and necrosis NCR (red), and peritumoral edematous area ED in PED or surrounding nonenhancing FLAIR hyperintensity (SNFH) in MEN and MET (green). (Color figure online)

of a zero mean unit variance normalization was applied to input images. Each input image was divided into patches of $128 \times 128 \times 128$ voxels for PED and $128 \times 160 \times 112$ for MEN and MET. The model output consisted on three channels corresponding to the three tumor sub-regions. Region-based training was employed and the patch size was determined by the GPU memory allocation, favoring larger patches while remaining within the GPU's capacity [13]. We used a class-weight loss function that combined Dice loss and cross entropy loss. To optimize the loss function, we used the stochastic gradient descent (SGD) optimizer with Nesterov momentum with the following parameters: initial learning rate of 0.01, momentum of 0.99, and weight decay of $3e-05$. Each of the five folds was trained for 100 epochs on an NVIDIA A100 (40 GB) GPU. At inference time, images were predicted using a sliding window approach. The window size matched the patch size used during training. The nnU-Net implementation is available in an open-source repository: <https://github.com/MIC-DKFZ/nnUNet>.

Swin UNETR. The Swin UNETR framework employs a vision transformer (ViT)-based [9] hierarchical structure for localized self-attention using non-overlapping windows [11, 12, 27]. Swin UNETR's innovative local window self-attention outperforms traditional ViT, which is well suited to multiscale tasks. The framework includes a 3D Swin transformer encoder with window-shifting for extended receptive fields, and it connects to a multiscale residual U-Net-like decoder to perform tasks like 3D medical image segmentation.

For each of the three tasks, we trained a full-resolution 3D Swin UNETR model using a five-fold cross-validation approach. A preprocessing consisting of a zero mean unit variance normalization was applied to input images. Each input image was sampled four times using patches of $96 \times 96 \times 96$ voxels to fully utilize the GPU's memory. The model output was 4 channels corresponding to the three labels and background. We used a class-weight loss function that combined Dice loss and focal loss. To optimize the loss function, we used the AdamW optimizer with an initial learning rate of 0.0001, momentum of 0.99, and weight decay of $3e-05$. Each of the folds were trained for 600 epochs on an NVIDIA A5000 (24 GB) GPU and NVIDIA A6000 (48GB) GPU. The Swin UNETR implementation is part of the PyTorch-based framework MONAI:¹. Hyper-parameter optimization was carried out using Optuna [2]:².

2.3 Model Ensembling

To enhance the accuracy and robustness of the segmentation outcome, we propose a model ensembling strategy. This approach (see Fig. 1) involves harnessing the complementary strengths of the two models described, nnU-Net and Swin UNETR, to collectively address the task of pixel classification.

¹ <https://monai.io>.

² optuna.readthedocs.io/.

The ensembling strategy was optimized for each task based on each model’s performance. For the PED task, we trained an nnU-Net solely on the ET region and ensembled its predictions with a Swin UNETR, which was trained on all labels, for ET region prediction. For TC and WT, we ensembled predictions from nnU-Net and Swin UNETR (as described previously), both trained on all labels. For MEN, predictions from both implementations, trained on all labels, were combined to generate ET, TC, and WT regions. Finally, for the MET task, only nnU-Net, trained on all labels, was used after experimental evaluation, as Swin UNETR showed inferior performance in this scenario.

When using a combination of nnU-Net and Swin UNETR, we ensembled the outputs (after applying the corresponding nonlinear activation functions) from each fold obtained during cross-validation ($k = 5$) training of both models. This allowed us to leverage the advantages of both approaches.

2.4 Post-processing

New performance metrics were introduced in this year’s BraTS challenge to assess segmentation models at a lesion-wise level rather than over the entire/multiple tumor region(s). We developed a post-processing strategy to adapt to the lesion-wise scores. This post-processing (see Fig. 3) was applied on the ensembled predictions and, first, removed disconnected regions (which contributed to undesired noise) smaller than 130 voxels for PED, 110 voxels for MEN, and 15 voxels for MET (Fig. 4). These optimal threshold values were determined by experimenting with the cross-validation data.




Fig. 3. Post-processing strategy. The ensemble predictions were first cleaned of small disconnected regions. Then, for the PED task, ET and ED labels were redefined based on ET/WT and ED/WT thresholds, respectively.

Subsequently, given the nature of pediatric brain tumors, in the PED task, numerous cases had empty ground truth annotations for the ET and the ED labels. Therefore, redefining these labels in cases where they fell below a certain threshold with respect to the WT volume was particularly useful. For example, if the ET/WT ratio fell below the threshold, the ET label would be redefined to either NCR or ED, which corresponds to the TC region. Figure 5a displays the lesion-wise Dice score *vs.* ET/WT threshold curve obtained from the optimal threshold search process performed on the cross-validation sets. For ET, we




Fig. 4. Threshold search on the cross-validation set for identifying small disconnected regions in PED, MEN and MET, respectively. LW refers to lesion-wise metrics.

obtained an optimal threshold of 0.04. On the other hand, Fig. 5b displays the lesion-wise Dice score vs ED/WT threshold curve obtained from the same process but applied to the ED label, which yielded an optimal threshold of 1.00. These optimal threshold values were also determined by experimenting with the cross-validation data.

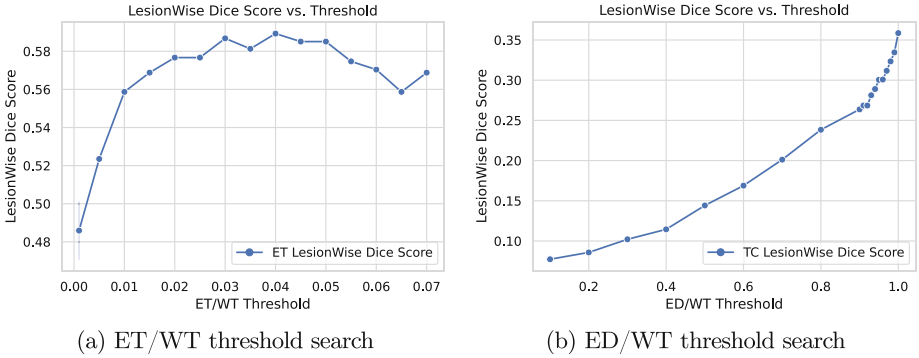


Fig. 5. Threshold search on the cross-validation set. Abbreviations: ED, Peritumoral edematous; ET, enhancing tumor; WT, whole tumor.

By carrying out these post-processing steps, we aimed for a balance between maintaining the integrity of the original predictions and enhancing the final scores by mitigating potential misclassifications.

3 Results

Table 1 provides a comprehensive overview of the performance evaluation of our models across the validation datasets for each task. This performance evaluation was performed automatically by the challenge’s digital platform, with no access to the validation ground truth data. Additionally, Fig. 6 shows qualitative results on some cases of the validation datasets for PED, MEN, and MET tasks.

Table 2 provides quantitative evaluation of our proposed solution across the test datasets for each task.

Table 1. Quantitative results on the validation datasets of PED, MEN and MET. Lesion-wise (LW) Dice coefficients and 95% Hausdorff distances (HD95) were computed for enhancing tumor (ET), tumor core (TC), and whole tumor (WT), respectively. Numbers in bold indicate the best results for each task.

Task	Model	LW Dice			LW HD95 (mm)		
		ET	TC	WT	ET	TC	WT
PED	nnU-Net	0.462	0.731	0.798	224.71	30.41	26.08
	Swin UNETR	0.362	0.641	0.716	191.47	82.51	63.33
	Ensemble	0.466	0.73	0.797	158.89	39.34	35.11
	Post-processing	0.733	0.782	0.817	75.93	25.54	24.18
MEN	nnU-Net	0.818	0.799	0.794	51.06	55.74	55.59
	Swin UNETR	0.64	0.644	0.59	117.85	114.43	135.63
	Ensemble	0.833	0.832	0.804	45.94	43.68	53.26
	Post-processing	0.852	0.846	0.832	37.98	38.68	42.9
MET	nnU-Net	0.565	0.614	0.545	117.53	110.88	115.5
	Swin UNETR	0.341	0.378	0.343	204.33	191.12	200.05
	Ensemble	0.533	0.594	0.539	123.18	110.28	115.38
	Post-processing	0.559	0.604	0.565	110.07	105.12	102.77
	Without Swin	0.608	0.649	0.587	91.62	91.29	95.7

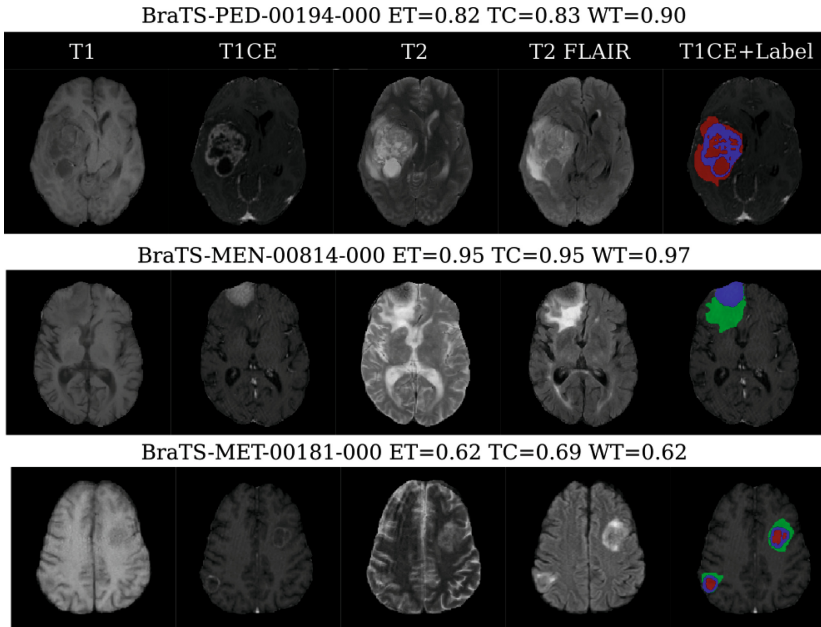


Fig. 6. Qualitative results on the validation datasets of PED, MEN, and MET. The selected cases show the median of the averaged lesion-wise Dice over three tumor regions for each task (NCR-red, ED-green, ET-blue). (Color figure online)

Table 2. Quantitative results on the testing datasets of PED, MEN and MET and ranks obtained in the competition. Lesion-wise (LW) Dice coefficients and 95% Hausdorff distances (HD95) were computed for enhancing tumor (ET), tumor core (TC), and whole tumor (WT), respectively.

Task	Statistic	LW Dice			LW HD95 (mm)		
		ET	TC	WT	ET	TC	WT
PED	Mean	0.653	0.809	0.826	43.89	21.82	20.86
	(Standard deviation)	(0.32)	(0.185)	(0.183)	(108.59)	(75.15)	(75.31)
	1st rank 25th quantile	0.569	0.809	0.840	1.41	2.96	2.96
	Median	0.741	0.856	0.875	3.67	5.26	4.24
MEN	75th quantile	0.916	0.888	0.895	8.98	9.17	6.76
	Mean	0.876	0.867	0.849	30.04	31.69	35.17
	(Standard deviation)	(0.217)	(0.227)	(0.231)	(80.9)	(83.53)	(86.77)
	3rd rank 25th quantile	0.907	0.885	0.863	1.00	1.00	1.00
MET	Median	0.968	0.968	0.953	1.00	1.00	1.62
	75th quantile	0.985	0.985	0.975	2.91	3.61	4.36
	Mean	0.555	0.6	0.58	113.96	112.84	108.17
	(Standard deviation)	(0.279)	(0.3)	(0.289)	(122.53)	(121.6)	(122.13)
4th rank	25th quantile	0.312	0.370	0.342	1.66	1.29	2.45
	Median	0.638	0.691	0.644	75.78	75.60	9.49
	75th quantile	0.813	0.863	0.826	202.39	189.43	189.03

4 Discussion

Overall, the combination of ensemble models and post-processing procedures yielded clear improvements in the performance for the PED, MEN, and MET tasks. With our approach, we are able to synergize the capabilities of nnU-Net and Swin UNETR models based on the validation scores to effectively generate more accurate and robust segmentation. Our strategy minimized undesired smaller contributions, which would negatively impact lesion-wise performance.

Despite the consistent labeling of the three tumor sub-regions (ET, TC, WT) across all tasks, the lesions' morphology and spatial location are variable across subjects and tasks. Among these three sub-regions, ET identification was the most challenging in the PED task, as evidenced by the lowest Dice coefficient for this region. For the MEN and MET tasks, a better balance between label segmentation was achieved. This confirms the importance of applying label-wise post-processing techniques tailored to each task. Consequently, binary metrics disproportionately penalize false positive predictions. Moreover, the novel metrics centered on lesions result in severe penalties for isolated volumes that do not correspond to actual lesions. All of this accentuates the necessity for post-processing strategies to address disconnected and small components.

Each task’s training process was restricted to the usage of task-specific data, adhering to the predefined guidelines of the BraTS 2023 challenge. Outside of the scope of this challenge, it would be beneficial to enhance the model’s capacity with additional pre-training to ensure it generalizes effectively.

Our intention was to tackle the diversity in the challenge data from both the perspective of the model and that of the data. Towards this, we argue that self supervised learning strategies and the concept of foundation model [3, 23] holds the potential to augment model’s generalizability, thus making it suitable for a more extensive range of clinical applications. The PED and MET tasks were conducted with limited available data, which is not surprising given the rarity of brain tumors. To overcome this limitation, the incorporation of generative artificial intelligence could also be a potential solution to augment the training set with synthetic data.

5 Conclusion

The segmentation of rare tumors from radiology images using deep learning techniques is challenging given the limited available data. We developed a model ensemble technique to address the segmentation of three types of brain tumors included in the BraTS 2023 challenge: i.e., pediatric, meningioma, and metastases. The ensemble prediction from nnU-Net and Swin UNETR provided better segmentation accuracy for all tumor types. Furthermore, our targeted post-processing strategy based on cross-validated thresholding search improved the performance in all tasks, being more noticeable in pediatric tumor segmentation due to the label-wise targeted post-processing steps. Through this work, we explored ensembling techniques that leverage the generalizability of deep learning models and strategies across diverse data distributions and tasks. The success of our method is demonstrated by the challenge ranking positions on unseen test cases, being the winner in PED task, ranked third for MEN, and fourth for MET task.

Acknowledgements. Partial support for this work was provided by the National Cancer Institute (UG3 CA236536) and by the Spanish Ministerio de Ciencia e Innovación, the Agencia Estatal de Investigación and NextGenerationEU funds, under grants PDC2022-133865-I00 and PID2022-141493OB-I00. The authors gratefully acknowledge the Universidad Politécnica de Madrid (www.upm.es) for providing computing resources on Magerit Supercomputer.

References

1. Adewole, M., Rudie, J.D., Gbadamosi, A., et al.: The Brain Tumor Segmentation (BraTS) Challenge 2023: Glioma Segmentation in Sub-Saharan Africa Patient Population (BraTS-Africa) (2023)
2. Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M.: Optuna: a next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2019)

3. Anwar, S.M., Parida, A., Atito, S., et al.: SS-CXR: multitask representation learning using self supervised pre-training from chest x-rays. [arXiv:2211.12944](https://arxiv.org/abs/2211.12944) (2022)
4. Baid, U., Ghodasara, S., Mohan, S., et al.: The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv preprint [arXiv:2107.02314](https://arxiv.org/abs/2107.02314) (2021)
5. Bakas, S., Akbari, H., Sotiras, A., et al.: Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Sci. Data* **4**(1), 170117 (2017). <https://doi.org/10.1038/sdata.2017.117>
6. Bakas, S., Akbari, H., Sotiras, A., et al.: Segmentation labels and radiomic features for the pre-operative scans of the TCGA-GBM collection. *The Cancer Imaging Archive* (2017). <https://doi.org/10.7937/K9/TCIA.2017.KLXWJJ1Q>
7. Bakas, S., Akbari, H., Sotiras, A., et al.: Segmentation labels and radiomic features for the pre-operative scans of the TCGA-LGG collection. *The Cancer Imaging Archive* (2017). <https://doi.org/10.7937/K9/TCIA.2017.GJQ7R0EF>
8. Curtin, S., Minino, A., Anderson, R.: Declines in cancer death rates among children and adolescents in the united states, 1999-2014. *National Center for Health Statistics Data Brief* (2016)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al.: An image is worth 16x16 words: transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) (2020)
10. Fathi Kazerooni, A., Arif, S., Madhogarhia, R., et al.: Automated tumor segmentation and brain tissue extraction from multiparametric MRI of pediatric brain tumors: a multi-institutional study. *Neuro-Oncol. Adv.* **5**(1), vdad027 (2023)
11. Hatamizadeh, A., Nath, V., Tang, Y., et al.: Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images. In: Crimi, A., Bakas, S. (eds.) *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, pp. 272–284. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-08999-2_22
12. Hatamizadeh, A., Tang, Y., Nath, V., et al.: UNETR: transformers for 3D medical image segmentation. In: *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 574–584 (2022)
13. Isensee, F., Jaeger, P.F., Kohl, S.A., et al.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat. Methods* **18**(2), 203–211 (2021)
14. Jiang, Z., Parida, A., Anwar, S.M., et al.: Automatic visual acuity loss prediction in children with optic pathway gliomas using magnetic resonance imaging. In: *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 1–5 (2023)
15. Karargyris, A., Umeton, R., Sheller, M., et al.: Federated benchmarking of medical artificial intelligence with MedPerf. *Nat. Mach. Intell.* **5**, 799–810 (2023)
16. Kazerooni, A.F., Khalili, N., Liu, X., et al.: The Brain Tumor Segmentation (BraTS) Challenge 2023: Focus on Pediatrics (CBTN-CONNECT-DIPGR-ASNR-MICCAI BraTS-PEDs) (2023)
17. Kofler, F., Meissen, F., Steinbauer, F., et al.: The Brain Tumor Segmentation (BraTS) Challenge 2023: Local Synthesis of Healthy Brain Tissue via Inpainting (2023)
18. LaBella, D., Adewole, M., Alonso-Basanta, M., et al.: The ASNR-MICCAI Brain Tumor Segmentation (BraTS) Challenge 2023: Intracranial Meningioma (2023)
19. Li, H.B., Conte, G.M., Anwar, S.M., et al.: The Brain Tumor Segmentation (BraTS) Challenge 2023: Brain MR Image Synthesis for Tumor Segmentation (BraSyn) (2023)

20. Liu, X., Bonner, E., Jiang, Z., et al.: Automatic segmentation of rare pediatric brain tumors using knowledge transfer from adult data. In: 20th IEEE International Symposium on Biomedical Imaging (2023)
21. Menze, B.H., Jakab, A., Bauer, S., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* **34**(10), 1993–2024 (2015). <https://doi.org/10.1109/TMI.2014.2377694>
22. Moawad, A.W., Janas, A., Baid, U., et al.: The Brain Tumor Segmentation (BraTS-METS) Challenge 2023: Brain Metastasis Segmentation on Pre-treatment MRI (2023)
23. Parida, A., Capellan-Martin, D., Atito, S., et al.: DiCoM—diverse concept modeling towards enhancing generalizability in chest X-ray studies. [arXiv:2402.15534](https://arxiv.org/abs/2402.15534) (2024)
24. Rashed, W.M., Maher, E., Adel, M., et al.: Pediatric diffuse intrinsic pontine glioma: where do we stand? *Cancer Metastasis Rev.* **38**(4), 759–770 (2019)
25. Rohlfing, T., Zahr, N.M., Sullivan, E.V., Pfefferbaum, A.: The SRI24 multichannel atlas of normal adult human brain structure. *Hum. Brain Mapp.* **31**(5), 798–819 (2010)
26. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
27. Tang, Y., Yang, D., Li, W., Roth, H.R., et al.: Self-supervised pre-training of swin transformers for 3D medical image analysis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20730–20740 (2022)