Clinically Informed, Dataset-Adaptive Lightweight Segmentation of Brain Metastases within the nnU-Net Framework

Willem P.E. Boonzaier[1,2] [0009-0007-7216-6543], Izak S. Pretorius[1,2] [0009-0003-0200-0602], Aondona M, Iorumbur[2,3][0009-0006-5621-0981], Udunna C. Anazodo[2,4,5] [0000-0001-8864-035X]

[1] Department of Medical Physics, University of the Free State, Bloemfontein, South Africa
[2] SPRINT AI Training for African Medical Imaging Knowledge Translation (SPARK) Program, Montreal, Canada
[3] Department of Physics, Federal University of Technology, Minna, Nigeria
[4] Department of Neurology and Neurosurgery, McGill University, Montreal, QC, Canada
[5] The Montreal Neurological Institute-Hospital, Montreal, QC, Canada

boonzaierwpe@ufs.ac.za

Abstract. We present a clinically informed, dataset-adaptive lightweight pipeline for brain metastasis sub-region segmentation within the nnU-Net v2 framework. Targeting the four BraTS 2026 METS sub-regions — enhancing tumour (ET), non-enhancing tumour core (NETC), surrounding non-enhancing FLAIR hyperintensity (SNFH), and resection cavity (RC) — our approach addresses key barriers to clinical deployment through five complementary contributions: (i) RC-balanced dataset curation using stratified K-means clustering to counter severe post-treatment class imbalance; (ii) Connected Component Dice (CCDice) loss to explicitly supervise multi-lesion spatially-disconnected structures; (iii) Jacobian sensitivity-guided XTiny model compression reducing model size by 80% relative to the standard nnU-Net ResEncM baseline; (iv) an on-the-fly 3D clinical prior attention map derived from T1ce spatial gradients encoding grey-white matter interface, vascular junction, and Rolandic region predilections; and (v) always-active modality dropout enforcing robustness to missing MRI sequences. Independent ablation shows CCDice loss and RC-balanced curation as the highest-impact modifications, with RC DSC gains of +29.1% and +21.9% respectively over the full-dataset nnU-Net baseline. The resulting model is compact, modality-robust, and anatomically informed, advancing the case for clinically deployable brain metastasis segmentation.

Keywords: Clinical attention, Auto segmentation, Brain metastasis, Lightweight, Dataset curation.

1 Introduction

Brain metastases present a significant clinical challenge, occurring in 20–40% of advanced primary cancers and predominantly colonizing the gray-white matter junction (GWMJ), vascular zone boundaries, and the Rolandic region via hematogenous

spread.[1, 2]. Accurate automated segmentation is crucial for diagnosis, prognosis, treatment planning, and response evaluation. [3–6].

Deep learning models, particularly U-Net architectures, have shown promise in automating the detection and segmentation of brain tumors, including metastases[6–8]. These automated approaches aim to overcome the limitations of manual segmentation, which is often time-consuming, subjective, and prone to inter-observer variability[6, 9]. While convolutional neural networks have set benchmarks for performance, they face two critical hurdles for clinical deployment: architectural complexity and modality dependence[10]. Consequently, there is an urgent need for lightweight architectures that maintain high sensitivity in resource-constrained environments and robust pipelines that can handle missing input modalities through training-time regularization, such as modality dropout[11, 12]. By providing a standardized, high-quality dataset, the BraTS-METS challenge creates an ideal environment for advancing model performance and generalizability[13]. Top performing models in this challenge have in the past included highly complex, self-configuring pipelines such as nnU-Net[14] and advanced hybrid architectures utilizing deformable convolutions or state-space models to capture the subtle, irregular morphology of brain metastases. While these frameworks consistently achieve top-tier segmentation accuracy, they frequently rely on extensive ensemble strategies and high-capacity architectures that prioritize performance at the expense of computational efficiency[15–17]. Modern research argues against the traditional assumption that computational efficiency requires a compromise in performance, the XtinyU-Net architecture demonstrates that significant reductions in parameter count and memory footprint can be achieved without sacrificing diagnostic accuracy[18]. Furthermore, multi-encoder head configurations—such as separating structural sequences (T1/T1ce versus T2/T2-FLAIR) or spatial and spectral information and fusing modalities later in the architecture can improve overall accuracy and eliminate severe outliers[19].

To bridge the gap between high-performance modelling and clinical deployment, we present a lightweight architecture designed for robustness across modalities and informed by typical anatomical presentation. Our principal contributions are: (i) an XTiny type Residual Encoder U-Net integrated within the nnU-Net self-configuring framework; (ii) a Connected Component Dice loss that rewards complete lesion-level detection; (iii) always-active modality dropout for robustness to missing sequences; (iv) a differentiable clinical prior attention map modulating spatial feature importance at training time; (v) a principled dataset curatgion strategy based on unsupervised clustering of tumour and image quality features and adjustment of foreground sampling probability during training, (vi) multi-encoder heads for splitting modalities and spatial and spectral information, (vii) adaptive Tversky loss to optimize recall and precision ratios and (viii) small lesion oversampling to increase F1 accuracy for small lesions.

2 Methods

2.1 Dataset and evaluation methods

We trained on fold 0 of the default nnU-Net 5-fold split, as little differences were observed between folds, using the BraTS 2026 Brain Metastases dataset comprising mpMRI scans (T1, T1ce, T2, T2-FLAIR) annotated for NETC (label 1), SNFH (label 2), ET (label 3), and RC (label 4), with RC present only in post-treatment cases [13]. Blinded validation was performed on the BraTS 2026 challenge validation dataset and platform (Synapse) using lesion-wise Dice (DSC) and lesion-wise normalized surface distance (NSD) as well as F1 scores for lesion detection accuracy for large and small lesions per class[13].

2.2 Architecture and training setup

We adopted the nnU-Net v2 3D full-resolution ResEncUNet with patch size and spacing determined automatically by the ResEncM planner. All experimental models trained for 200 epochs (SGD, Nesterov momentum 0.99, polynomial LR from 0.01 and standard Dice & Cross entropy loss) on an NVIDIA RTX 5000 16GB GPU.

2.3 Experimental optimization

A baseline out-of-the-box nnU-Net trained on the full BraTS METS 2026 training set served as the control. Eight modifications were evaluated independently (Table 1). Retention of modifications in the final model was informed by their observed effect on Synapse validation scores during iterative challenge submissions; therefore, the ablation gains are development-set estimates rather than independent held-out estimates.

Table 1. Experimental modifications evaluated independently against the baseline nnU-Net.

Experiment name	Description
E1-Population reduction	Dataset exploration, feature clustering and sub-population extraction
E2-Loss function change	Dice-Cross entropy → Connected component Dice-Cross entropy.
E3-Model compression	XTiny U-Net style model compression of features and number of residual blocks per convolutional layer.
E4-Clinical attention mapping	Custom prior clinically informed spatial attention map.
E5-Modality dropout	Probabilistic modality dropout training.
E6-Foreground probability sampling adjustment	0.33 → 0.6
E7- Epoch adaptive Tversky loss	Tversky loss with independent alpha and beta control per region, adapting every epoch.
E8-Triple-Encoder head configuration	Triple-encoder design fused at bottleneck.

To counter severe class imbalance, with RC present in only 12.7% of cases, we constructed a balanced 300-case subset (150 RC-positive, 150 RC-negative) using K-means clustering over tumour morphology and image-quality features. With batch size 2, this produces approximately 75% per-batch RC exposure. Within each pool, clustering preserved phenotypic diversity using tumour volumes, lesion counts, vol-ume ratios, per-modality coefficient of variation and foreground entropy, and T1ce contrast-to-noise ratio.

K-means was applied independently within each pool, with cluster number select-ed by joint inspection of within-cluster sum of squares and silhouette score over k=2–15.

Standard voxel-level Dice loss was replaced with Connected Component Dice (CCDice), which computes overlap for individual connected components and penalises fragmented predictions and false-positive blobs independently of volume. Components were identified by 26-connectivity on the argmax prediction; this partition was treated as fixed for the training step, while Dice was computed from the underlying soft probabilities. Predicted and ground-truth components were greedily matched by maximum voxel-wise IoU; unmatched components were scored as false positives or false negatives. CCDice loss was combined with cross-entropy. An additional Tversky loss term was used with independently adaptive α and β weights for each region. Following each epoch, α was updated from the validation precision and recall according to equation 1,

$$\alpha_t = R_t /(R_t + P_t + \varepsilon) \tag{1}$$

With

$$\beta_t = 1- \alpha_t \tag{2}$$

where (R_t) and (P_t) denote recall and precision at epoch (t).

This increases the false-positive penalty when recall exceeds precision and increases the false-negative penalty when precision exceeds recall, thereby dynamically balancing precision and recall across epochs. To prevent unstable changes early in training, α was constrained within predefined bounds and its change per epoch was limited. Multi-encoder heads (T1+T1ce, T2+T2-FLAIR, and spectral) were fused at the bottleneck, and fore-ground sampling probability was increased from 0.33 to 0.6.

A 3D clinical prior attention map α was generated on-the-fly in normalized per-axis voxel coordinates rather than registered atlas space. It combined: (i) a grey-white matter interface term from the absolute first difference of T1ce along the x-axis; (ii) bilateral vascular-junction Gaussians centred at x=0.33 and 0.66; and (iii) a Rolandic-region anisotropic Gaussian centred at (0.5,0.5,0.85). The components were combined as α=wGWMI+0.5wvasc+0.5wRol, normalized to [0,1], and used to scale input features (fig. 1).

$$x' = x \times (1 + \alpha) \tag{3}$$

To reduce model complexity, we used the training-free XTinyU-Net selection strategy of Kimbowa et al. [18]. Importantly, the final “tiny” architecture was not chosen

by an arbitrary parameter target or by repeatedly training candidate models. Instead, nnU-Net's planned topology was retained, and an ordered family of smaller candidate configurations was generated by progressively capping convolutional channel widths; depth (number of residual blocks per stage) was evaluated separately. Each candidate was randomly initialized and evaluated on a small set of unlabeled training images using an input-output Jacobian sensitivity metric. The resulting sensitivity values were normalized and their changes between consecutive candidates were used to identify the transition from a stable sensitivity plateau to the abrupt variation associated with representational collapse. The smallest candidate still on the stable side of this boundary was selected as the XTiny configuration. For this dataset, the selected width configuration was [16, 32, 64, 128, 256, 256], while the selected depth configuration was [1, 2, 2, 2, 2, 2] (11 residual blocks). Thus, mod-el size was determined from the dataset-specific initialization-time sensitivity curve rather than from a manually imposed compression ratio. The procedure was implemented in the nnU-Net initialization pipeline. Modality dropout was applied independently to the four input channels with p=0.25, while requiring at least one channel to remain.

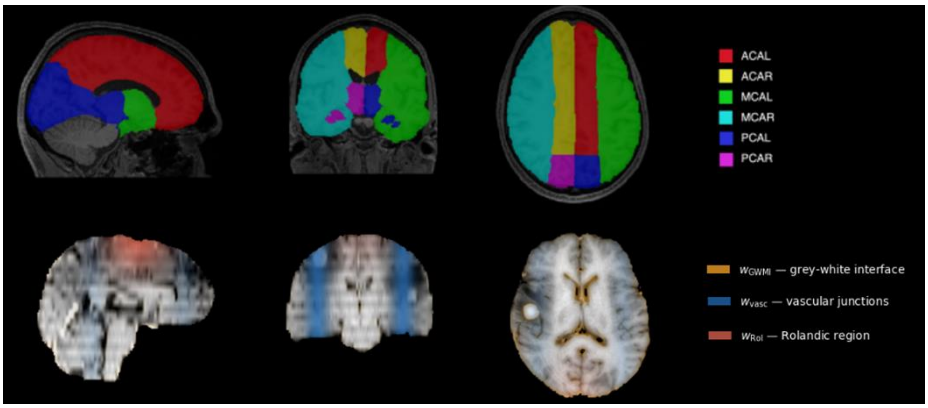

Fig. 1. Clinical prior attention map overlaid on MRI. Top row: cerebrovascular territory atlas (ACAL, ACAR, MCAL, MCAR, PCAL, PCAR) on T1w MRI (sagittal, coronal, transverse), illustrating anatomical metastatic predilection zones[20]. Bottom row: three-component prior attention map α on T1ce MRI. Amber: grey-white matter interface (w_{GWMI}) from T1ce spatial gradient; blue: bilateral vascular junction paths (w_{vasc}) at x=0.33 and x=0.66; coral: Rolandic region prior (w_{RoI}) centred on the superior central sulcus applied as $x' = x \times (1 + \alpha)$.

Final model composition. The final model submitted for the BraTS 2026 METS challenge was constructed by implementing all the optimization changes that led to metric improvement during the model optimization phase. The model was trained for 400 epochs in total, and the final checkpoint was used for inference. Individual ablation experiments (E1–E8) were each compared against a 200-epoch baseline; the final combined model was instead trained for 400 epochs to ensure full convergence, as compute constraints precluded also training a matched 400-epoch ablation baseline. Part of the final model's improvement over the 200-epoch baseline may therefore reflect

additional training rather than the combined effect of E1–E8 alone; we report this as a limitation rather than an isolated, epoch-matched effect size.

3 Results and discussion

The silhouette score for the feature clustering peaked at k=2 then dropped steeply after k=3, indicating that k=3 captured the first meaningful phenotypic subdivision beyond a binary split without fragmenting coherent groups (Fig. 2a). k=3 was therefore selected for both pools. Proportional subsampling from each cluster yielded the final 150 cases per pool ensuring rare phenotypes were represented at higher rates than their natural prevalence (Fig. 2b).

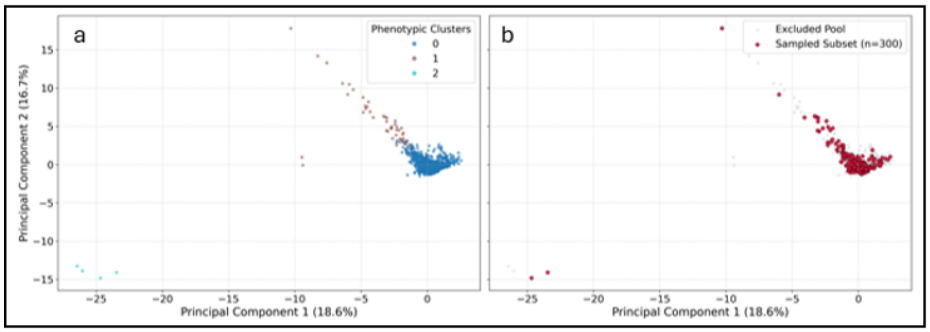

Fig. 2. PCA projection of the BraTS 2026 METS training population (PC1: 18.6%, PC2: 16.7%). (a) K-means clustering (k=3) identifies a dominant typical metastatic phenotype (cluster 0, blue), an atypical large-lesion subgroup (cluster 1, brown), and rare extreme cases (cluster 2, cyan). (b) Final 300-case sampled subset (red) against the excluded pool (grey), demonstrating broad phenotypic coverage across all three clusters.

The width sensitivity analysis showed a low-sensitivity plateau across the 5–12M-parameter range, with the configuration [16, 32, 64, 128, 256, 256] selected at the stable side boundary. Depth sensitivity identified [1, 2, 2, 2, 2, 2] with 11 residual blocks. The resulting XTiny-style single-encoder model contained approximately 20.7M trainable parameters versus 101.5M for standard nnU-Net ResEncM (Fig. 3b).

Each experimental modification was evaluated independently against a baseline nnU-Net trained on the full BraTS 2026 METS training set, with results expressed as percentage gain in lesion-wise DSC and NSD across all four sub-regions (Fig. 4).

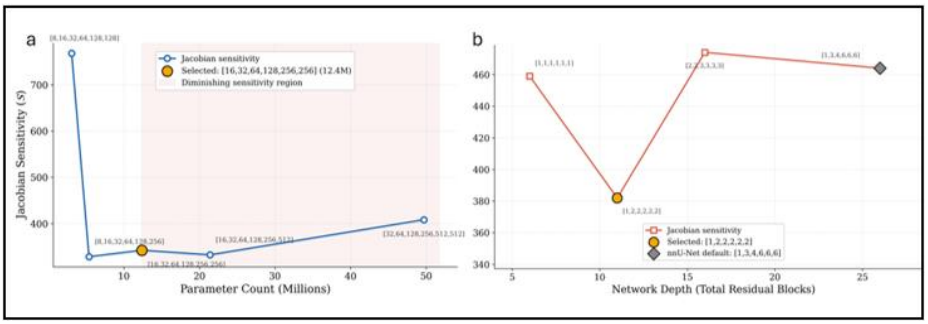

Fig. 3. Jacobian sensitivity analysis for XTiny U-Net style model compression. (a) Width sensitivity as a function of parameter count across five convolutional layer width configurations. The shaded region indicates configurations where additional parameters produce diminishing or worsening gradient signal. (b) Depth sensitivity as a function of total residual blocks across four configurations.

CCDice (E2) produced the largest individual gains, particularly for RC (DSC +29.1%, NSD +25.2%), consistent with the multi-lesion nature of brain metastases. RC-balanced curation (E1) also substantially improved RC (DSC +21.9%, NSD +21.6%) and produced approximately 6–8% gains across ET, NETC and SNFH. Model compression (E3) and clinical attention (E4) produced moderate, consistent gains; E4 improved RC DSC by 6.8% and NSD by 9.3%. Modality dropout (E5) pro-duced smaller but consistently positive gains.

Increasing foreground sampling probability to 0.6 produced negative SNFH changes (DSC −3.2%, NSD −4.5%) but small positive gains for ET, RC and NETC, consistent with greater emphasis on focal lesions at the expense of diffuse SNFH.

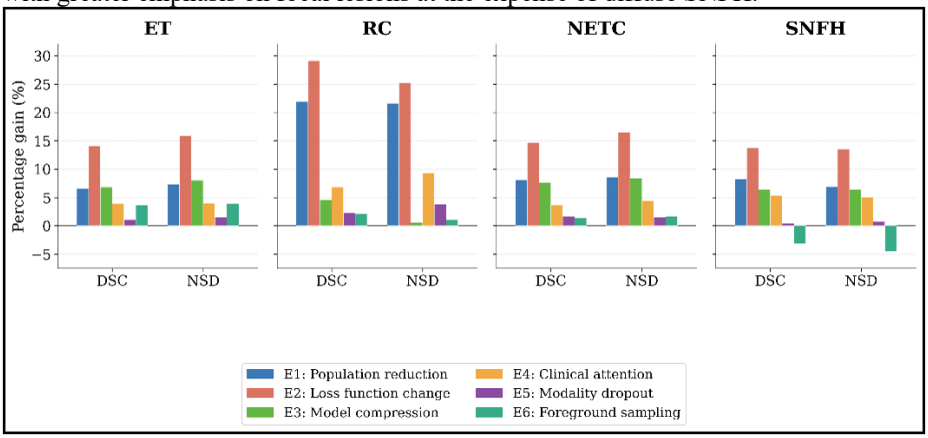

Fig. 4. Percentage gain in lesion-wise DSC and NSD relative to the full nnU-Net baseline for each experimental modification, evaluated independently across ET, RC, NETC, and SNFH subregions.

Final model performance. After E1–E6 (Best Dice model), lesion-wise DSC/NSD were 0.600/0.652 (ET), 0.501/0.383 (RC), 0.623/0.665 (NETC), and 0.554/0.562 (SNFH). Large-lesion F1 scores were 0.676, 0.084, 0.684 and 0.677, respectively, while

small-lesion F1 remained lower at 0.354, 0.000, 0.356 and 0.270. E7 and E8 were then applied to improve small-lesion F1. Tversky loss increased mean small-lesion F1 by 6.8% but reduced large-lesion F1 by 2.8%. Adding the triple encoder produced a 7.0% small-lesion F1 increase and 0.23% large-lesion increase, with RC showing the main Dice/NSD trade-off (fig. 5).

3.1 Limitations

Several limitations should be noted. First, modifications retained in the final model were selected using iterative Synapse validation submissions, so ablation gains are development-set estimates rather than fully independent estimates. Second, the final model used 400 epochs whereas ablations used 200, so some improvement may re-flect additional training. Lastly, evaluation remained within the BraTS-METS ecosystem; external validation, missing-modality inference experiments, and direct inference-time/memory benchmarking remain necessary as well as comparison to current state of the art models.

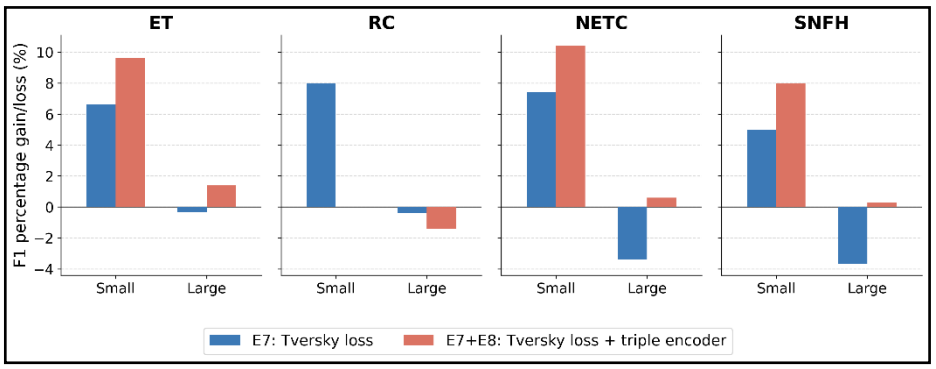

Fig. 5. Percentage gain in lesion-wise F1 score for small and large lesion relative to the best dice model employing E1-6 implementations for each experimental modification, evaluated independently across ET, RC, NETC, and SNFH sub-regions. Results for added adaptive Tversky loss and triple encoder architecture design.

4 Conclusion

We presented a clinically informed, dataset-adaptive lightweight pipeline for brain metastasis sub-region segmentation within nnU-Net. The largest gains arose from RC-balanced curation and Connected Component Dice loss, addressing class imbalance and the spatially disconnected nature of metastases.

Jacobian sensitivity-guided XTiny compression reduced model size by 80% relative to nnU-Net ResEncM while maintaining consistent performance. The clinical prior attention map further improved performance, particularly for RC, without addi-tional annotations.

Taken together, these contributions advance the case for a clinically-grounded, computationally efficient approach to brain metastasis segmentation — one that is small enough to run on a single mid-range GPU, robust to real-world acquisition variability, and informed by the disease biology that defines where and how metastases present. Future work will focus on optimization of the clinical prior design, ex-tending the model compression to additional tasks and clinical problems, quantifying inference-time gains due to model reduction strategies and quantification of modality dropout training on missing-modality inference experimentation.

Acknowledgments. This work was part of the Sprint AI Training for African Medical Imaging Knowledge Translation (SPARK) Academy 2026 summer school on deep learning in medical imaging. The authors would like to thank the instructors of the summer for providing insightful background knowledge on brain tumours that informed the research presented here, most notably: Maruf Adewole, Mohannad Barakat, Craig Jones, Noha Magdy, Amal Saleh, Aondona Iorumbur, Bernes Atabonfack, Confidence Raymond, Craig Jones, Ilerioluwakiiye Abolade, Jeremiah Fadugba, Juampablo Heras Rivera, Kenneth Aguh, Lukman E. Ismaila, Maruf Adewole, Mohtady Barakat, Nourou Dine Bankole, Saikat Roy, Toufiq Musah, Willem Boonzaier.. The authors acknowledge the funding support provided to SPARK through the Lacuna Fund for Health and Equity (PI: Udunna Anazodo), the RSNA R & E Foundation (PI: Farouk Dako), the University of Washington Population Health Initiative Tier2 Grant (PI: Mehmet Kurt), McGill University Healthy Brain and Healthy Lives (HBHL, Anazodo), and the National Science and Engineering Research Council of Canada (NSERC) Discovery Launch Supplement (PI:UdunnaAnazodo, DGECR-2022-00136).

Code Availability. The source code used in this study is publicly available at the GitHub repository provided: https://github.com/Williamwhy/brats-mets-2026.

Artificial Intelligence Declaration. Artificial intelligence tools were used for language editing, proofreading and graphical editing of the manuscript. All study methodology, experiments, data analysis, interpretation, and scientific conclusions were performed by the authors.

References

1. Mehrabian, H., Detsky, J., Soliman, H., Sahgal, A., Stanisz, G.J.: Advanced magnetic resonance imaging techniques in management of brain metastases. Front. Oncol. 9, 1–16 (2019). https://doi.org/10.3389/fonc.2019.00440.
2. Schneider, T., Kemmling, A., Schroeder, J., Pantel, K., Glatzel, M., Schoen, G., Mohme, M., Fiehler, J., Gellißen, S.: Inverse perfusion requirements of supra- and infratentorial brain metastases formation. Front. Neurol. 9, 1–7 (2018). https://doi.org/10.3389/fneur.2018.00391.
3. Liu, Y., Stojadinovic, S., Hrycushko, B., Wardak, Z., Lau, S., Lu, W., Yan, Y., Jiang, S.B., Zhen, X., Timmerman, R., Nedzi, L., Gu, X.: A deep convolutional neural network-based automatic delineation strategy for multiple brain metastases stereotactic radiosurgery. PLoS One. 12, 1–17 (2017). https://doi.org/10.1371/journal.pone.0185844.
4. Heyn, C., Moody, A.R., Tseng, C.L., Wong, E., Kang, T., Kapadia, A., Howard, P., Maralani, P., Symons, S., Goubran, M., Martel, A., Chen, H., Myrehaug, S., Detsky, J., Sahgal, A.,

Soliman, H.: Segmentation of Brain Metastases Using Background Layer Statistics (BLAST). Am. J. Neuroradiol. 44, 1135–1143 (2023). https://doi.org/10.3174/ajnr.A7998.
5. Li, B., Sun, Q., Fang, X., Yang, Y., Li, X.: A novel metastatic tumor segmentation method with a new evaluation metric in clinic study. Front. Med. 11, 1–12 (2024). https://doi.org/10.3389/fmed.2024.1375851.
6. Swapna, J., Sundeep Reddy, G., Krishna Chaitanya Reddy, S., Nagi Reddy, N.: Enhanced Brain Tumor Segmentation using Convolutional Neural Network. (2025). https://doi.org/10.4108/eai.28-4-2025.2358087.
7. Pennig, L., Shahzad, R., Caldeira, L., Lennartz, S., Thiele, F., Goertz, L., Zopfs, D., Meißner, A.K., Fürtjes, G., Perkuhn, M., Kabbasch, C., Grau, S., Borggrefe, J., Laukamp, K.R.: Automated detection and segmentation of brain metastases in malignant melanoma: Evaluation of a dedicated deep learning model. Am. J. Neuroradiol. 42, 655–662 (2021). https://doi.org/10.3174/AJNR.A6982.
8. Verma, A., Yadav, A.K.: Brain tumor segmentation with deep learning: Current approaches and future perspectives. J. Neurosci. Methods. 418, 110424 (2025). https://doi.org/10.1016/j.jneumeth.2025.110424.
9. Islam, R., Hossain, S.: Enhanced Brain Tumor Segmentation Using CBAM-Integrated Deep Learning and Area Quantification. Int. J. Biomed. Imaging. 2025, (2025). https://doi.org/10.1155/ijbi/2149042.
10. Müller, S., Weickert, J., Graf, N.: Robustness of brain tumor segmentation. J. Med. Imaging. 7, 1–18 (2020). https://doi.org/10.1117/1.jmi.7.6.064006.
11. Rahman, M.M., Marculescu, R.: MK-UNet: Multi-kernel Lightweight CNN for Medical Image Segmentation. (2025). https://doi.org/10.1109/ICCVW69036.2025.00114.
12. Feng, X., Ghimire, K., Kim, D.D., Chandra, R.S., Zhang, H., Peng, J., Han, B., Huang, G., Chen, Q., Patel, S., Bettagowda, C., Sair, H.I., Jones, C., Jiao, Z., Yang, L., Bai, H.: Brain Tumor Segmentation for Multi-Modal MRI with Missing Information. J. Digit. Imaging. 36, 2075–2087 (2023). https://doi.org/10.1007/s10278-023-00860-7.
13. Maleki, N., Amiruddin, R., Moawad, A.W., Yordanov, N., Gkampenis, A., Fehringer, P., Umeh, F., Chukwurah, C., Memon, F., Petrovic, B., Albalooshy, B., Cramer, J., Krycia, M., Shrickel, E.B., Ikuta, I., Thompson, G., Vidal, L., Kosovic, V., Goldman-Yassen, A.E., Hill, V., So, T.Y., Mhana, S., Alotaibi, A., Page, N., Bhatia, P., Guelen, M.S., Sharifi, Y., Jakovljevic, M., Abosabie, S.A.S., Abosabie, S.A., Ghonim, M., Ghonim, M., Manteghinejad, A., Baid, U., Janas, A., Krantchev, K., Ramakrishnan, D., Saluja, R., Jekel, L., Adewole, M., Albrecht, J., Anazodo, U., Aneja, S., Anwar, S.M., Bergquist, T., Chiang, V., Chung, V., Conte, G.M., Dako, F., Eddy, J., Ezhov, I., Khalili, N., Farahani, K., Iglesias, J.E., Jiang, Z., Johanson, E., Kazerooni, A.F., Kofler, F., Labella, D., Leemput, K. Van, Li, H.B., Linguraru, M.G., Liu, X., Meier, Z., Menze, B.H., Moy, H., Osenberg, K., Piraud, M., Reitman, Z., Shinohara, R.T., Wang, C., Wiestler, B., Wiggins, W., Shafique, U., Willms, K., Avesta, A., Bousabarah, K., Chakrabarty, S., Gennaro, N., Holler, W., Kaur, M., Lamontagne, P., Lin, M., Lost, J., Marcus, D.S., Maresca, R., Merkaj, S., Pedersen, G.C., Von Reppert, M., Sotiras, A., Teytelboym, O., Tillmans, N., Westerhoff, M., Youssef, A., Godfrey, D., Floyd, S., Rauschecker, A.M., Villanueva-Meyer, J., Pflüger, I., Cho, J., Bendszus, M., Brugnara, G., Guzman Perez-Carillo, G.J., Johnson, D.R., Kam, A., Yin, B., Kwan, M., Lai, L., Lall, N.U., Narayana Patro, S., Wu, L., Bansal, A., Barkhof, F., Besada, C., Chu, S., Druzgal, J., Dusoi, A., Farage, L., Feltrin, F., Fong, A., Fung, S.H., Gray, R.I.,

Iv, M., Postma, A.A., Mahajan, A., Joyner, D., Krumpelman, C., Letourneau-Guillon, L., Lincoln, C.M., Maros, E., Miller, E., Morón, F., Nimchinsky, E.A.: Analysis of the MICCAI Brain Tumor Segmentation - Metastases (BraTS-METS) 2025 Lighthouse Challenge: Brain Metastasis Segmentation on Pre- and Post-treatment MRI. 53, (2025).

14. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nat. Methods. 18, 203–211 (2021). https://doi.org/10.1038/s41592-020-01008-z.
15. Bancerek, M., Rudzki, P., Nalepa, J.: Segmentation of Pre and Post-treatment Brain Metastases Using nnU-Nets. Lect. Notes Comput. Sci. 16376 LNCS, 161–172 (2026). https://doi.org/10.1007/978-3-032-16365-3_15.
16. Umeh, F., Yordanov, N.Y., Maleki, N., Amiruddin, R., Moawad, A., Pytlarz, M., Chukwurah, C., Aboian, M.: Taking Advantage of MONAI and DiNTS Frameworks to Develop a State-of-the-Art Algorithm for Automatic Segmentation of Brain Metastases. Lect. Notes Comput. Sci. 16376 LNCS, 182–189 (2026). https://doi.org/10.1007/978-3-032-16365-3_17.
17. Krikorian, W., Purwar, A.: Automated Segmentation for the Brain Tumor Segmentation (BraTS) Metastases 2025 Challenge Using Multi-Architectural Deep Learning. Lect. Notes Comput. Sci. 16376 LNCS, 173–181 (2026). https://doi.org/10.1007/978-3-032-16365-3_16.
18. Kimbowa, A., Heidari, M., Liu, D., Hacihaliloglu, I.: XTinyU-Net: Training-Free U-Net Scaling via Initialization-Time Sensitivity. 1–11 (2026).
19. Pan, Y., Yong, H., Lu, W., Li, G., Cong, J.: Brain tumor segmentation by combining MultiEncoder UNet with wavelet fusion. J. Appl. Clin. Med. Phys. 25, 1–12 (2024). https://doi.org/10.1002/acm2.14527.
20. Rowsthorn, E., Cribb, L., Sinclair, B., Pham, W., Chong, T.T.J., Yiallourou, S., Cavuoto, M., Vivash, L., O'Brien, T.J., Shao, X., Wang, D.J.J., Law, M., Pase, M.P., Harding, I.H.: Relationships between measures of neurovascular integrity and fluid transport in aging: a multi-modal neuroimaging study. Fluids Barriers CNS. 22, (2025). https://doi.org/10.1186/s12987-025-00664-7.