

Biophysics Informed Pathological Regularisation for Brain Tumour Segmentation

Lipei Zhang, Yanqi Cheng, Lihao Liu,
Carola-Bibiane Schönlieb, and Angelica I Aviles-Rivero

Department of Applied Mathematics and Theoretical Physics,
University of Cambridge, UK
{l1z452,yc443,l1610,cbs31,ai323}@cam.ac.uk

Abstract. Recent advancements in deep learning have significantly improved brain tumour segmentation techniques; however, the results still lack confidence and robustness as they solely consider image data without biophysical priors or pathological information. Integrating biophysics-informed regularisation is one effective way to change this situation, as it provides an prior regularisation for automated end-to-end learning. In this paper, we propose a novel approach that designs brain tumour growth Partial Differential Equation (PDE) models as a regularisation with deep learning, operational with any network model. Our method introduces tumour growth PDE models directly into the segmentation process, improving accuracy and robustness, especially in data-scarce scenarios. This system estimates tumour cell density using a periodic activation function. By effectively integrating this estimation with biophysical models, we achieve a better capture of tumour characteristics. This approach not only aligns the segmentation closer to actual biological behaviour but also strengthens the model’s performance under limited data conditions. We demonstrate the effectiveness of our framework through extensive experiments on the BraTS 2023 dataset, showcasing significant improvements in both precision and reliability of tumour segmentation.

Keywords: Glioma · Segmentation · Partial Differential Equations · Deep Learning · Representation Regularisation.

1 Introduction

Glioblastoma, the most aggressive brain cancer, represents 14.3% of primary malignant central nervous system tumours. It exhibits rapid, heterogeneous growth that complicates detection and treatment. Thus often leading to poor prognosis despite advanced interventions [16, 17]. Magnetic Resonance Imaging (MRI) plays a crucial role in diagnosing brain tumours. It provides high-resolution images that are essential for accurate tumour margin segmentation. These images also guide treatment decisions across various modalities, such as T1, T1-weighted with contrast (T1w-gd), T2-weighted (T2w), and T2 Fluid Attenuated Inversion Recovery (FLAIR) [4]. Advances in deep learning, particularly with introduction of the U-Net architecture [24], have significantly improved segmentation in

different imaging modalities [1, 6, 11, 13, 20, 22]. The integration of U-Net and Transformer models [8, 9, 14], utilising long-range attention mechanisms, show improved accuracy in the Brain Tumour Segmentation (BraTS) Challenge [3].

However, the limited data availability in biology and medicine highlights the challenges of the generalisability issues arising from the limited availability of medical data. This underscores the importance of incorporating domain-specific knowledge into deep learning models to enhance their utility in medical imaging. Physics-Informed Neural Networks (PINNs) offer a promising approach to overcome generalisability issues. They integrate boundary conditions and partial differential equations (PDEs) into the learning process. It has enabled their application across various fields. For example, it is used in elasticity reconstruction [23], Alzheimer’s disease analysis [28], and glioma progression estimation [19]. However, the effort in glioma has primarily focused on parameter estimation. Previous efforts have not succeeded in directly integrating physics-informed learning with automated segmentation models. The integration could leverage the inherent structure and dynamics of biological systems. This approach has the potential to yield more accurate and biologically plausible segmentation results. This method could connect data-driven models with model-driven models that include specific expert knowledge in their learning.

In this work, we propose a novel method that designs biophysics-informed regularisation to improve brain tumor segmentation through deep learning. This approach embeds pathological insights into tumour growth dynamics, enhancing segmentation accuracy and robustness. Our contributions are:

1. A unified framework that integrate deep learning with brain tumour growth PDE models through a biophysics-informed regularisation, enhancing segmentation precision and robustness across various scenarios.
2. Introducing MRI-driven tumour cell density estimator estimates along with biophysics-informed regularisation, achieving segmentations that reflect real tumour growth mechanism. In this estimator, a designed periodic activation function in the high-level feature extracting layers captures the non-linear dynamics of the biophysical model, enabling detailed signal and spatial derivative learning for calculation of biophysics-informed regularisation term.
3. These innovations have yielded exceptional segmentation results on the BraTS 2023 dataset, proving effective across various architectures, dataset sizes, missing modalities and combinations with other losses.

2 Proposed Method

Brain tumour segmentation. In this problem, we aim to segment 4 regions: normal region (y_0), tumour core (TC, y_1), and whole tumour (WT, y_2), enhancing tumour (ET, y_3), in a given MRI data with T1 (i_0), T1Gd (i_1), T2 (i_2), and FLAIR (i_3) volumes. The segmentation model can be defined as f_θ such that the probability of each class of P can be obtained as $P = f_\theta(I)$, where

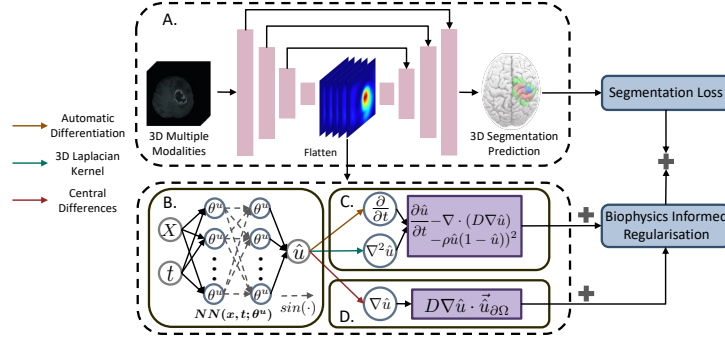


Fig. 1. The Biophysics-informed optimisation for segmentation. A. Main structure for brain tumour segmentation. B. Tumour cell density estimator (the flattened feature map will be concatenated with assumed time matrix T). C. Calculation of PDE loss. D. Calculation of boundary loss.

$I = \{i_0, i_1, i_2, i_3\}$, $P = \{p_0, p_1, p_2, p_3\}$, and $Y = \{y_0, y_1, y_2, y_3\}$ indicates one-hot mask of each subregion. The function f_θ is usually nonlinear and θ is a large vector of parameters. The learning phase selects θ in order to minimise a loss function \mathcal{L} that measures the accuracy of the predicted segmentation $f_\theta(I)$. This overall training flow is shown in Fig. 1(a). The parameters θ are obtained by minimising a loss function plus a regularisation term. Consider a dice loss with supervised manner, we optimise the network by:

$$\arg \min_{\theta} \mathcal{L}(I, Y, \theta) = \arg \min_{\theta} \left(\underbrace{\sum_{n \in \{0,1,2,3\}} \left(1 - \frac{2y_n p_n + 1}{y_n + p_n + 1} \right)}_{\text{Segmentation Loss}} + \underbrace{R_\theta(f_\theta(I))}_{\text{Regularisation}} \right) \quad (1)$$

Biophysics-informed optimisation. Exploring effective regularisation terms $R_\theta(\cdot)$ is essential to improve model robustness. Unlike current regularisation approaches based on VAE [21], we propose that incorporating biophysical models of tumour growth model into the training process as an more explainable learning bias. It enables the segmentation model to capture more accurate pathological features representative of tumour areas. To this end, we introduce a biophysics-informed regularisation module designed as a plug-and-play component. This module integrates two key elements: additional fully connected layers with periodic activation functions for pathological information extraction, and a combination of proliferation and diffusion PDEs alongside boundary conditions. This integration not only improves the segmentation’s ability to capture tumour details by embedding biophysical regularisation directly into the end-to-end optimisation process but also ensures the model’s outputs align more closely with actual biological behaviours. The overall objective loss, balanced by weights λ_1 and λ_2 , is formulated as:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{Dice}} + \underbrace{\lambda_1 \mathcal{L}_{\text{PDE}} + \lambda_2 \mathcal{L}_{\text{BC}}}_{\text{Biophysics Informed Regularisation}} \quad (2)$$

Tumour cell density estimator. In this first part of this module, we introduce a neural network component, $NN(x, t; \theta^u)$, which employs multiple fully connected layers with periodic activation functions to disentangle high-level feature maps from middle stage to potential tumour cell density \hat{u} . At these stages, the feature maps primarily highlight pathological features, focusing on tumour regions. This architecture offers a significant advantage, that it learns the non-linear aspects of governing PDEs indirectly, thus endowing the model with the capability to capture detailed signals and spatial derivatives without specifying them explicitly.

The process begins with a set of feature maps (B, C, H, W, D) , where B represents batch size, C is channel number and H, W, D denote height, width and depth, respectively. These maps are flattened across H, W, D dimensions to form a matrix X with dimensions $(B, C, H \times W \times D)$. This procedure ensures efficient sequential processing and maintains the independence of the information of each channel. This approach, different from SIREN's coordinate consideration [27], incorporates an assumed temporal dimension T , viewing each segmentation optimisation step as a discrete moment in time. The temporal dimension T , a matrix of the same shape as X , but filled with the assumed time step t , is concatenated with X . Time embedding is crucial for simulating tumour dynamics during training, ensuring the model's output aligns with expected growth dynamics. It also facilitates the calculation of the first-order time derivative, integrating temporal dynamics into our model. The concatenated features $y = (X, T) = ((x_1, t_1)(x_2, t_2), \dots, (x_t, t_C))$ with size of $(B, C, H \times W \times D \times 2)$, with x_c and t_c represent dimensional and temporal matrix on the c^{th} channel respectively, undergo transformation as follows:

$$\hat{u} = \Gamma(y) = W_n(\gamma_{n-1} \circ \gamma_{n-2} \circ \dots \circ \gamma_0)(y) + b_n, \quad y_i \mapsto \gamma_i(y_i) = \sin(W_i y_i + b_i) \quad (3)$$

where, $\gamma_i : \mathbb{R}^{M_i} \rightarrow \mathbb{R}^{N_i}$ denotes the i^{th} layer of the transformation. It comprises the affine transforms defined by the weight matrix $W_i \in \mathbb{R}^{N_i \times M_i}$, and the biases $b_i \in \mathbb{R}^{N_i}$ utilised on the input $y_i \in \mathbb{R}^{M_i}$, followed by the sine nonlinearity. This process extracts the tumour cell density \hat{u} , facilitating the calculation of tumour proliferation and diffusion PDEs. During the inference phase, this module can be pruned to avoid any increase in processing time.

Calculation of biophysics-informed regularisation. Tumour cell density proliferation and diffusion across the brain is predominantly based on the reaction-diffusion equation [12], the model incorporates logistic growth to represent cell proliferation and enforces Neumann boundary conditions to model tumour boundaries within the brain domain Ω ,

$$\begin{aligned} \frac{\partial u}{\partial t} &= \underbrace{\nabla \cdot (d \nabla u)}_{\text{Diffusion}} + \underbrace{f(u)}_{\text{Reaction}} \\ f(u) &= \rho u (1 - u), \quad d \nabla u \cdot \mathbf{u}_{\partial \Omega} = 0 \end{aligned} \quad (4)$$

These equations describes the evolution of the tumour cell density u , which infiltrates neighbouring tissues with a diffusion tensor d , and proliferates ac-

cording to the law defined with $f(u)$. We consider the reaction term with a logistic growth term with a net proliferation rate ρ . Here, d and ρ have same dimension with u . The $\frac{\partial u}{\partial t}$ with respect to time t is computed through automatic differentiation due to t embedded with input x . The predicted cell density vectors u are reshaped back to 3D format (B, C, H, W, D) . $F_\Delta(u) = \nabla \cdot (d\nabla u)$, is approximated using a 3D Laplacian kernel (K). This kernel is more efficient than traditional central finite difference method for preserving spatial information when computing the second derivative and boundary condition related to location.

$$K = \left[\begin{array}{ccc|ccc} 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & -6 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{array} \right] \quad (5)$$

After considering all voxels with total number of $N_{biophy} = H \times W \times D$ in feature maps, the PDE loss term takes the following form:

$$\mathcal{L}_{\text{PDE}} = \left\| \frac{1}{N_{biophy}} \sum_{i=1}^{N_{biophy}} \left(\frac{\partial u_i}{\partial t_i} - \nabla \cdot (D\nabla u_i) - \rho u_i(1 - u_i) \right)^2 \right\| \quad (6)$$

$\|\cdot\|$ represents the norm across channels and batches. Additionally, boundary constraints are specified to model the diffusion process accurately, ensuring the diffusion remains non-negative and aligns with the spatial domain of the brain within the feature maps. Here, this can be describe by the following formula:

$$\begin{aligned} 0 \leq x \leq H, \quad 0 \leq y \leq W, \quad 0 \leq z \leq D \\ u_x(0, y, z) = 0, \quad u_x(H, y, z) = 0 \\ u_y(x, 0, z) = 0, \quad u_y(x, W, z) = 0 \\ u_z(x, y, 0) = 0, \quad u_z(x, y, D) = 0 \end{aligned} \quad (7)$$

The corresponding boundary constraints take the form, with $\mathbb{1}$ indicates the indicator function:

$$\mathcal{L}_{\text{BC}} = \left\| d \sum_{x=0}^H \sum_{y=0}^W \sum_{z=0}^D \left[\frac{\mathbb{1}_{\{0,H\}}(x)}{WD} u_x^2(y, z) + \frac{\mathbb{1}_{\{0,W\}}(y)}{HD} u_y^2(x, z) + \frac{\mathbb{1}_{\{0,D\}}(z)}{HW} u_z^2(x, y) \right] \right\| \quad (8)$$

3 Experiment

Datasets. Our evaluations were applied on the BraTS 2023 dataset [3], consisting of 1251 cases within the official training partition. Each case comprises mpMRI scans with four modalities: T1, T1c, T2, and FLAIR, each having dimensions of $240 \times 240 \times 155$. Three tumor subregions—enhancing tumour (ET), tumour core (TC), and whole tumour (WT)—were labelled by neuroradiologists. Given the absence of labels in the official validation and testing sets, the 1251 cases was split into training, validation, and test sets with a ratio of 7:1:2. During training, the 4-channel MRI volumes centred on the tumour were cropped to $128 \times 128 \times 128$ voxel patches. To mitigate overfitting, augmentation techniques

were employed. Z-score standardisation was performed on non-zero voxels across each channel, with outlier values being clipped. For inference, same centred cropping on the tumour of 4-channel MRI volumes was conducted, complemented by test-time augmentation (TTA) techniques [26] and overlapped sliding window inference [11]. Other preprocessing setups aligned with [5].

Main experimental Details and compared methods. All experiments, developed on MONAI v1.3.0 [18] and Pytorch 2.1.0, was executed on an Nvidia A10 GPU with 24GB memory with automatic mixed precision training. The number of epoch was 175 epochs and batch size was set to 1. We used the Ranger 2020 optimiser [15, 29] with a starting learning rate of 3×10^{-4} and a cosine decay schedule. Dice Loss [20] was computed both batch-wise and channel-wise, unweighted, alongside a consistent loss weight for PDE and BC (λ_1 and λ_2 set to 1). Final models were used for testing. Baseline comparisons included 3D UNet [24], R2-UNet [1], nn-UNet [11], UNETR [9], SegResNet, and SegResNet-VAE [2], all employing single dice loss. We compared these with our proposed biophysics-informed loss, excluding SegResNetVAE for regularization comparison. More configurations are shown in appendix. The diffusion coefficient d and proliferation rate ρ are in the range of $[0.02, 1.5] \text{ mm}^2/\text{day}$ and ρ between $[0.002, 0.2]/\text{day}$ for tumour tissues respectively [12]. High-level feature maps, targeting tumour region, assume isotropic diffusion across voxels due to gliomas' tendency to disrupt neural pathways, leading to uniform diffusion. The input size of feature map for the tumour cell estimator was $(B \times C \times 16 \times 16 \times 16)$ to ensure computational efficiency. The d and ρ values randomly sampled from the specified ranges for each voxel. This random sampling introduces regularisation, ensuring biophysical priors fall within a valid distribution.


Evaluation Metrics. The results were compared with 2 quantitative criteria, Dice similarity coefficient (Dice) [7] and Hausdorff Distance (HD) [10], on each of the three sub-regions. We take the norm as Euclidean norm for HD. Outliers were removed using the 95th percentile (HD95) approach [5].

Results and Discussion. The segmentation results for the BraTS2023 dataset, shown in Table. 1, indicate that integrating biophysics-informed regularisation with UNet architectures leads to better accuracy than standard methods. Our approach not only improved the Dice coefficients for all tumour regions but also decreased Hausdorff distances, proving its effectiveness in enhancing segmentation precision and better geometric accuracy in tumour margin prediction. This method also showed consistency across different models, suggesting its robustness. By integrating biophysics into training, we can achieve more reliable segmentations, potentially improving clinical outcomes.

The additional studies depicted in Fig 2 present an comprehensive experiments of the biophysics-informed UNet model. In Fig 2 (a), the employment of Sine over ReLU activation functions is scrutinised, revealing superior performance in identifying multiple regions, and the inclusion of boundary conditions

Table 1. Comparison table of UNet family with and without the biophysics-informed regularisation on mean and standard deviation of Dice score and Hausdorff distance.

Method	Dice \uparrow	Hausdorff Distance (mm) \downarrow
	TC/WT/ET	TC/WT/ET
UNet	90.68/92.29/87.28	6.16/7.85/10.99
UNet (Biophy)	91.83/92.34/88.26	3.99/6.94/7.47
R2-UNet	90.75/91.86/87.29	5.82/7.34/11.24
R2-UNet (Biophy)	91.02/91.91/87.53	4.83/7.14/9.25
nn-UNet	91.24/92.48/87.38	4.44/6.95/10.85
nn-UNet (Biophy)	91.70/92.69/87.55	4.00/6.51/10.39
UNet-TR	88.20/90.73/85.41	6.75/11.36/12.62
UNet-TR (Biophy)	89.06/91.10/85.96	5.97/8.98/11.12
SegResNet	91.96/91.99/87.28	3.91/7.53/10.80
SegResNet-VAE	92.00/91.02/87.02	3.88/6.80/10.73
SegResNet (Biophy)	92.13/92.19/87.46	3.62/6.68/10.56

**Fig. 2.** (a) shows the impact of activation functions and boundary conditions on the biophysics-informed UNet performance on Dice score. (b)-(d) shows comparisons of UNet, with and without the biophysics-informed module with mean on Dice score over the 3 regions. (b) explores performance variations using two or three modalities, (c) presents comparisons on different training set sizes, (d) compares different loss combinations.

mitigates excessive predictions in the potential infiltrated area (ET). These findings demonstrate that the biophysics-informed regularisation effectively offers a prior understanding of growth and marginal constraints. Fig 2 (b) illustrates the model's robustness even with reduced imaging modalities, where using only T1 and FLAIR yields comparable results to including an additional T2 modality. Training size variations in Fig 2 (c) show our efficiency in data-scarce scenarios. Integrating our module into UNet models even outperform standard UNets on larger datasets. Lastly, Fig 2 (d) explores different loss functions, including Dice+Cross Entropy, Focal, and Jaccard Loss. The consistent enhancement of UNet's performance, due to biophysics informed regularisation, showcases the module's flexibility. This adaptability enhances network optimisation and performance of segmentation.

Main Comparison Visualisation in Fig. 3 demonstrates that biophysics informed regularisation notably enhances tumour segmentation in various networks, yielding results that align more closely with ground truth. This method

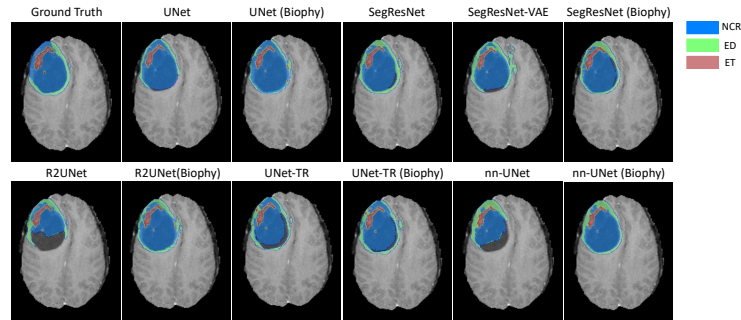


Fig. 3. MRI tumour segmentation comparison: All enhanced by biophysics-informed regularisation. The GD-enhancing tumour (ET), the peritumoral edematous tissue (ED), and the necrotic tumour core (NCR). The WT combines red, blue, and green; the TC merges green and blue; and the ET is represented by red.




Fig. 4. Localisation in Brain Tumour Segmentation: The images compare the UNet with its biophysics-informed version using GradCAM, revealing more precise tumour localisation. The ground truth scan is shown first, followed by GradCAM visualisations of last encoder and decoder layers, with and without biophysics-informed regularisation.

surpasses traditional dice loss and VAE regularisation by capturing tumour boundaries and structures more effectively, showcasing its potential to improve segmentation across different neural architectures.

Interpretability Visualisation in Fig. 4, utilising Grad-CAM [25], shows how biophysics-informed regularisation refines UNet performance. The focused activation on tumour areas, guided by biophysical tumour characteristics, leads to precise localisation and enhanced interpretability, ensuring model outputs mirror real tumour pathology for accurate medical segmentation.

4 Conclusion

Our study showcases the effectiveness of biophysics-informed regularisation in improving brain tumour segmentation accuracy and robustness. Integrating tumour growth dynamics with deep learning, we enhance traditional models significantly. Tested on the BraTS 2023 dataset, our method boosts various networks performance, proves robust against data scarcity and varying training losses, and sets a foundation for future exploration of domain-specific knowledge in deep

learning for medical imaging, particularly in weakly supervised or unsupervised settings.

Acknowledgments. AAR gratefully acknowledges funding from the Cambridge Centre for Data-Driven Discovery and Accelerate Programme for Scientific Discovery, made possible by a donation from Schmidt Futures, EPSRC Digital Core Capability Award, and CMIH and CCIMI, University of Cambridge. This project also was supported with funding from the Oracle for Research Project Award. CBS acknowledges support from the Philip Leverhulme Prize, the Royal Society Wolfson Fellowship, the EPSRC advanced career fellowship EP/V029428/1, EPSRC grants EP/S026045/1 and EP/T003553/1, EP/N014588/1, EP/T017961/1, the Wellcome Innovator Awards 215733/Z/19/Z and 221633/Z/20/Z, CCMI and the Alan Turing Institute.

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

References

1. Alom, M.Z., Hasan, M., Yakopcic, C., Taha, T.M., Asari, V.K.: Recurrent residual convolutional neural network based on u-net (r2u-net) for medical image segmentation. arXiv preprint arXiv:1802.06955 (2018)
2. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* **39**(12), 2481–2495 (2017)
3. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F.C., Pati, S., et al.: The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv preprint arXiv:2107.02314 (2021)
4. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J.S., Freymann, J.B., Farahani, K., Davatzikos, C.: Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data* **4**(1), 1–13 (2017)
5. Carré, A., Deutsch, E., Robert, C.: Automatic brain tumor segmentation with a bridge-unet deeply supervised enhanced with downsampling pooling combination, atrous spatial pyramid pooling, squeeze-and-excitation and evonorm. In: *International MICCAI Brainlesion Workshop*. pp. 253–266. Springer (2021)
6. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II* 19. pp. 424–432. Springer (2016)
7. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945)
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)
9. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. pp. 574–584 (2022)

10. Huttenlocher, D.P., Klanderman, G.A., Rucklidge, W.J.: Comparing images using the hausdorff distance. *IEEE Transactions on pattern analysis and machine intelligence* **15**(9), 850–863 (1993)
11. 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)
12. Lê, M., Delingette, H., Kalpathy-Cramer, J., Gerstner, E.R., Batchelor, T., Unkelbach, J., Ayache, N.: Bayesian personalization of brain tumor growth model. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part II* 18. pp. 424–432. Springer (2015)
13. Liu, L., Hu, X., Zhu, L., Fu, C.W., Qin, J., Heng, P.A.: ψ -net: Stacking densely convolutional lstms for sub-cortical brain structure segmentation. *IEEE transactions on medical imaging* **39**(9), 2806–2817 (2020)
14. Liu, L., Huang, Z., Liò, P., Schönlieb, C.B., Aviles-Rivero, A.I.: Pc-swinmorph: Patch representation for unsupervised medical image registration and segmentation. *arXiv preprint arXiv:2203.05684* (2022)
15. Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., Han, J.: On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265* (2019)
16. Louis, D.N., Perry, A., Reifenberger, G., Von Deimling, A., Figarella-Branger, D., Cavenee, W.K., Ohgaki, H., Wiestler, O.D., Kleihues, P., Ellison, D.W.: The 2016 world health organization classification of tumors of the central nervous system: a summary. *Acta neuropathologica* **131**(6), 803–820 (2016)
17. Low, J.T., Ostrom, Q.T., Cioffi, G., Neff, C., Waite, K.A., Kruchko, C., Barnholtz-Sloan, J.S.: Primary brain and other central nervous system tumors in the united states (2014-2018): A summary of the cbtrus statistical report for clinicians. *Neuro-oncology practice* **9**(3), 165–182 (2022)
18. Ma, N., et al.: Project-monai/monai: 1.3. 0. zenodo (2023)
19. Meaney, C., Das, S., Colak, E., Kohandel, M.: Deep learning characterization of brain tumours with diffusion weighted imaging. *Journal of Theoretical Biology* **557**, 111342 (2023)
20. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *2016 fourth international conference on 3D vision (3DV)*. pp. 565–571. Ieee (2016)
21. Myronenko, A.: 3d mri brain tumor segmentation using autoencoder regularization. In: *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 4th International Workshop, BrainLes 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers, Part II* 4. pp. 311–320. Springer (2019)
22. Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., et al.: Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999* (2018)
23. Ragoza, M., Batmanghelich, K.: Physics-informed neural networks for tissue elasticity reconstruction in magnetic resonance elastography. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 333–343. Springer (2023)
24. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)

25. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
26. Shanmugam, D., Blalock, D., Balakrishnan, G., Gutttag, J.: Better aggregation in test-time augmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1214–1223 (2021)
27. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* **33**, 7462–7473 (2020)
28. Song, T.A., Chowdhury, S.R., Yang, F., Jacobs, H.I., Sepulcre, J., Wedeen, V.J., Johnson, K.A., Dutta, J.: A physics-informed geometric learning model for pathological tau spread in alzheimer’s disease. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VII 23. pp. 418–427. Springer (2020)
29. Zhang, M., Lucas, J., Ba, J., Hinton, G.E.: Lookahead optimizer: k steps forward, 1 step back. *Advances in neural information processing systems* **32** (2019)