



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Spatiotemporal Graph Neural Network Modelling Perfusion MRI

Ruodan Yan¹, Carola-Bibiane Schönlieb¹, and Chao Li^{1,2,3}✉

¹ Department of Applied Mathematics and Theoretical Physics, University of Cambridge
Cambridge

² Department of Clinical Neurosciences, University of Cambridge

³ School of Science and Engineering, University of Dundee
c1647@cam.ac.uk

Abstract. Perfusion MRI (pMRI) offers valuable insights into tumor vascularity and promises to predict tumor genotypes, thus benefiting prognosis for glioma patients, yet effective models tailored to 4D pMRI are still lacking. This study presents the first attempt to model 4D pMRI using a GNN-based spatiotemporal model (**PerfGAT**), integrating spatial information and temporal kinetics to predict Isocitrate DeHydrogenase (IDH) mutation status in glioma patients. Specifically, we propose a graph structure learning approach based on edge attention and negative graphs to optimize temporal correlations modeling. In addition, we design a dual-attention feature fusion module to integrate spatiotemporal features while addressing tumor-related brain regions. Further, we develop a class-balanced augmentation methods tailored to spatiotemporal data, which could mitigate the common label imbalance issue in clinical datasets. Our experimental results demonstrate that the proposed method outperforms other state-of-the-art approaches, promising to model pMRI effectively for patient characterization.

Keywords: Dynamic Susceptibility Contrast MRI · Glioma · Isocitrate Dehydrogenase · Spatiotemporal Learning · Graph Neural Network.

1 Introduction

Glioma is the most common adult malignant brain tumor, characterized by dismal prognosis [12]. Isocitrate DeHydrogenase (IDH) mutation is one of the most important genetic markers for gliomas, indicating tumor malignancy. Clinically, determining IDH mutation status relies on invasive procedures, i.e., immunohistochemistry and gene sequencing of surgical samples, which are often time-consuming and less feasible for some patients. In parallel, MRI offers a non-invasive characterization of brain tumors. Previous MRI studies [18, 19] show that the genotype of glioma can be predicted using machine learning, offering valuable clinical decision support.

Perfusion MRI (pMRI) is a technique to reflect brain vascularity and micro-circulation. By capturing the dynamics of contrast agents that pass through the

vasculature, pMRI acquires sensitive spatiotemporal information to investigate neovascularization caused by angiogenesis, which could indicate tumor aggressiveness and better predict the IDH mutations over conventional magnetic resonance imaging, as IDH mutation status indicates glioma malignancy grading. However, it remains a challenge to develop effective spatiotemporal modelling for pMRI.

Previous studies rely predominantly on predefined pharmacokinetic models [9], leveraging mathematical representations to calculate several perfusion parameters from the kinetic curve, e.g., relative cerebral blood volume (rCBV). However, the derived parametric maps are inherently simplified and limited in modelling complex physiological kinetics. Although various attempts have been made to develop machine learning/deep learning approaches [3,13] for extracting more comprehensive features from parametric maps, sophisticated spatiotemporal modelling of pMRI remains an unmet need to decode perfusion kinetics effectively.

Scanty work developed spatiotemporal models to learn kinetics in perfusion imaging. Zhang et al. [22] proposed 3D-ResNet based ConvLSTM [15] framework, to predict patient survival. However, since this model is based on contrast-enhanced CT, it cannot include the detailed tissue structure and regional perfusion available from pMRI. Another spatiotemporal model, SwiFT [6], is proposed for functional MRI (fMRI) based on the Swin Transformer [11]. This model employs a window attention mechanism using the image patches defined by the Euclidean space, which limits its ability to model the perfusion patterns of anatomical regions.

Brain networks promise to comprehensively characterise the global and regional brain, where brain regions from the prior atlas are treated as nodes and their connectivities as edges [20]. Based on fMRI, graph neural networks (GNNs) have been shown to characterize brain function and disorders by modelling brain networks [10]. In particular, GNNs can model spatiotemporal graphs, e.g. STGCN for traffic forecasting [23], which promise to model the kinetics of perfusion networks. However, existing spatiotemporal graph models are not designed for medical images and, therefore, may be unable to capture the complex spatial structures and intricate physiological processes from the perfusion graph.

Further, both function and perfusion networks typically have noisy and dense connections, leading to overfitting risk in GNN modelling. Graph structure learning promises to refine the connectivities in functional graphs [24]. However, this method may be less effective in perfusion graphs due to the intrinsic differences between functional and perfusion graphs in topological properties and temporal resolution of raw MRI. An unmet need remains to develop approaches to optimize the structure of perfusion graphs reflecting vasculature networks.

In this study, we propose a spatiotemporal GNN framework with specific model developments tailored for tumor-bearing brain, and to learn the kinetics of pMRI and predict the genotype of glioma. Our contributions include:

1. **Perfusion Graph Structure Learning.** We propose a graph structural learning module utilizing edge attention and a negative graph to optimize

the temporal graph structure. Specifically, edge attention allows the model to adaptively assign attention to different edges based on their importance in capturing temporal correlations and identifying the most relevant graph structure. Moreover, a negative graph provides complementary information to identify potential missing important connections. This allows the model to prioritize relevant connections and refine noisy temporal correlations.

2. **Dual-attention Guided Spatiotemporal Feature Fusion.** To optimize the encoder’s capability in extracting pertinent spatiotemporal features while mitigating the influence of concurrent pathologies, we introduce a dual-attention mechanism, which systematically prioritizes brain regions by node attention and fuse spatiotemporal features with semantic attention, enabling the identification of most tumor-specific spatiotemporal features at both regional tumor and global brain levels.
3. **Class-balanced feature augmentation.** Class imbalance is a significant challenge in predicting genotypes due to natural incidence. Common methods for mitigating class imbalance, such as oversampling and cost-sensitive learning, are limited by overfitting and introducing noise or distortion to the data distribution [1, 2]. Here we propose a spatiotemporal augmentation method which recombines local and graph features from the minority class during the retraining stage, and incorporates a class-balanced classifier retraining approach [5]. Thereby generating data augmentation with more diverse representations reflecting patient brains and improving model generalizability and robustness.

As far as we know, this is the first work to develop spatiotemporal graph learning for characterizing pMRI. Our experimental results show that our method outperforms other state-of-the-art models in predicting the genotype of glioma.

2 Methodology

Overview. The proposed PerfGAT (Fig. 1) firstly constructs a spatiotemporal graph using DSC-MRI, brain atlas and tumor masks to incorporate focal tumors. Subsequently, spatial and temporal graph features and local tumor features are generated using distinct encoders. To reduce noisy temporal connections, we employ a graph structure learning approach based on edge attentions (details in 2.1). The dual-attention feature fusion mechanism (details in 2.2) strategically integrates features across local tumor and spatial and temporal graphs, yielding comprehensive representations from pMRI. Finally, to tackle the class-imbalance issue while avoiding data distortion, we employ a recombining augmentation mechanism tailored to spatiotemporal graph (details in 2.3).

Spatiotemporal graph generation. We construct a spatiotemporal graph \mathbf{G} from the raw DSC-MRI using the SRI24/TZO atlas [14] and tumor masks. Specifically, the nodes are defined as the brain regions of the atlas and the tumor region is segmented based on the tumor mask and taken as a separate

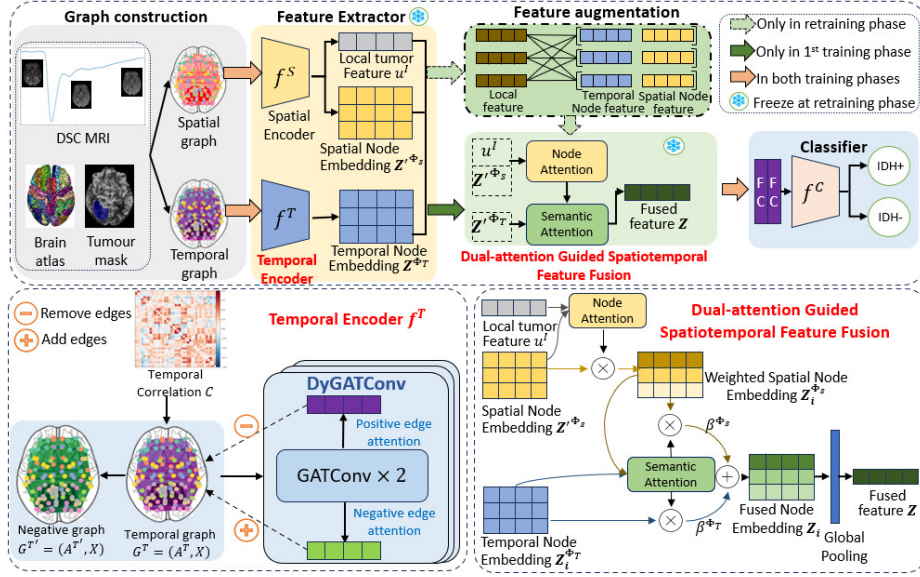


Fig. 1. Model Design. DSC-MRI, brain atlas, and tumor masks are employed to construct the input graphs, with features extracted and augmented before dual-attention fusion. Temporal graphs with their negative graphs are adjusted via DyGATConv, and features are fused using node and semantic attention to derive the final features.

node. The average time-signal curve of brain regions is extracted as node features $\mathbf{X}^B \in \mathbb{R}^{N \times T}$, where N is the node number, and T is the time length.

The perfusion graph consists of temporal and spatial connections. 1) Temporal connections represent the temporal correlations of perfusion kinetics across brain regions. We first derive a temporal correlation matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$, where C_{ij} represent the correlation coefficient between node pairs. Thresholding this matrix yields the temporal adjacency matrix $\mathbf{A}^T \in \mathbb{R}^{N \times N}$, where $A_{ij}^T = 1$ if C_{ij} exceeds a predefined threshold, and 0 otherwise. 2) Spatial connections are defined by the Euclidean distance of brain regions. Each node is connected to its k nearest neighbours, forming spatial adjacency matrix $\mathbf{A}^S \in \mathbb{R}^{N \times N}$, where $A_{ij}^S = 1$ indicates connected edges and $A_{ij}^S = 0$ indicates unconnected edges. This results in temporal graphs $\mathbf{G}^T = (\mathbf{X}^B, \mathbf{A}^T)$ and spatial graphs $\mathbf{G}^S = (\mathbf{X}^B, \mathbf{A}^S)$.

2.1 Graph Structure Learning

Our approach to graph structure learning is aimed at optimizing temporal graph connectivity to mitigate noisy temporal connections and enhance adaptability to task-specific nuances (Alg. 1). The graph structure learning module operates within each layer to dynamically adjust the graph structure. It begins by generating a negative graph, denoted as $G^{T^-} = (A^{T^-}, X^B)$, capturing the edges absent in the original graph while maintaining the same nodes. The edge

Algorithm 1 Graph Structure Learning

Input: $G^T = (X^B, A^T)$, $\alpha, \beta, l \leftarrow 0$

repeat

$l \leftarrow l + 1$

$G^{T^-} = (X^B, A^{T^-}) \leftarrow \text{Negative_graph}(G^T)$

 Update E^+ from G^T based on Equation (1)

 Update E^- from G^{T^-} based on Equation (1)

 edges_to_delete \leftarrow indices of lowest αe_{ij}^+ in E^+

 edges_to_add \leftarrow indices of highest βe_{ij}^- in E^-

$G^T \leftarrow \text{remove_edges}(\text{edges_to_delete})$

$G^T \leftarrow \text{add_edges}(\text{edges_to_add})$

 Update X^B based on Equation (1)

until $l = \text{max_layer}$

return G^T

attentions of the original graph and the negative graph are computed through the graph attention layer (GATConv [16]) based on Eq.1, where e_{ij}^+ and e_{ij}^- denote the attention scores for the edges in each graph, respectively. Subsequently, edges with low e_{ij}^+ are considered noisy edges to remove. Conversely, edges with high e_{ij}^- representing informative connections absent in the original graph are added.

$$e_{ij} = \text{softmax}_i(\text{LeakyReLU}(\mathbf{w}_a^T [\mathbf{W}_A * x_i \| \mathbf{W}_A * x_j])) \quad (1)$$

where \mathbf{W}_A is a weight matrix; x_i and x_j are node feature i and j ; \mathbf{w}_a is a weight vector; e_{ij} is edge attention score of connecting nodes i and j .

After acquiring the edge attention, the node feature is updated by Eq.2.

$$x_i^{(l+1)} = \sum_{j \in \mathcal{N}(i)} e_{ij} * \mathbf{W}_n^{(l)} * x_j^{(l)} \quad (2)$$

where $x_i^{(l+1)}$ is the feature of node i in layer $(l+1)$; $x_j^{(l)}$ is the feature of node j in layer l ; e_{ij} is edge attention score connecting node i, j ; \mathbf{W}_n is a weight matrix.

2.2 Dual-attention guided spatiotemporal feature fusion

In parallel to temporal encoder ($f^T(\cdot)$), spatial graph is encoded by GAT layers ($f^{\text{GAT}}(\cdot)$) and local tumor is encoded by 3DResNet + ConvLSTM [22] ($f^I(\cdot)$). Embeddings of temporal node $\mathbf{Z}^{\phi_T} = f^T(X^B, A^T)$, spatial node $\mathbf{Z}^{\phi_S} = f^{\text{GAT}}(X^B, A^S)$, and local tumor $u^I = f^I(X^I)$ (\mathbf{X}^I : local tumor image) are acquired.

The proposed dual-attention mechanism is structured by node attention and semantic attention. It first focuses on spatial features through node attention to emphasize tumor-related node embedding, consequently mitigating the influence of noise inherent in complex high-dimensional perfusion data and confounding factors. After project spatial node embeddings and tumor features into the

shared latent space, The node attention is computed as follows:

$$a_i = S(f_N^{\text{MLP}}(z_i^{\Phi'_s}), f_F^{\text{MLP}}(u^I)) \quad (3)$$

where a_i denotes the attention of the i th node; $z_i^{\Phi'_s}$ is the spatial feature of i th node; $S(\cdot)$ is the cosine similarity function, and $f_N^{\text{MLP}}(\cdot)$ and $f_F^{\text{MLP}}(\cdot)$ are projection heads for node embeddings and tumor features.

Subsequently, inspired by the HAN model [17], we use semantic attention to further fuse spatial and temporal node embeddings. The spatial and temporal node embeddings are taken as input and projected to a latent space ($f^{\text{MLP}}(\cdot)$). The transformed embeddings are measured using a semantic-level attention vector \mathbf{q} . We then average all the node embeddings as the significance of each dimension (spatial or temporal). The learned weights of each type β_{Φ_d} ($d \in (S, T)$) are then obtained by softmax normalization to compute the final node embeddings \mathbf{Z} :

$$\beta_{\Phi_d} = \text{Softmax}\left(\frac{1}{|\mathcal{N}|} \sum_{i \in \mathcal{N}} \mathbf{q} \cdot \tanh(f^{\text{MLP}}(\mathbf{Z}^{\Phi_d}))\right) \quad (4)$$

$$\mathbf{Z} = \beta^{\Phi_s} \cdot \mathbf{Z}^{\Phi_s} + \beta^{\Phi_T} \cdot \mathbf{Z}^{\Phi_T} \quad (5)$$

Table 1. The results for IDH classification on the meta-cohort dataset.

Model Name	ACC (%)	B-ACC (%)	SPEC (%)	SENS (%)	F1 (%)	AUC (%)
T1+CNN	92.5	50.0	100	00.0	00.0	78.4
rCBV+CNN	92.5	50.0	100	00.0	00.0	76.5
ResNet + ConvLSTM [15]	80.6	71.1	82.3	60.0	31.6	81.0
SwiFT [6]	77.6	69.5	79.0	60.0	19.5	80.0
GCN [8]	57.3	69.4	55.4	83.3	20.8	77.5
GIN [21]	92.1	57.1	97.6	16.7	22.2	57.4
GAT [16]	69.7	68.3	69.9	66.7	22.9	75.1
EGAT [4]	65.2	65.9	65.1	66.7	20.5	77.1
STGCN [23]	52.2	55.8	51.6	60.0	15.8	63.2
BrainGNN [10]	85.1	55.2	90.3	20.0	16.7	58.7
PerfGAT	88.8	86.2	89.2	83.3	50.0	94.4

2.3 Class-balanced Recombining Augmentation

To address the limitations of random resampling methods and enhance the robustness of our spatiotemporal GNN model, we propose a class-balanced recombining augmentation technique tailored for multimodal graph learning. Here, we randomly select pairs of minority class samples $h_{\text{minority}} = (u^I, u^B)$ where $u^B = (Z^{\Phi_s}, Z^{\Phi_T})$, and $h'_{\text{minority}} = (u^I, u^{B'})$. By recombining features within the minority class, we generate augmented samples $h_{\text{aug1}} = (u^I, u^{B'})$ and $h_{\text{aug2}} =$

(u^I, u^B) . This augmentation process ensures the synthesis of sufficient minority samples to achieve class balance within a highly biased dataset. By encouraging the model to learn more diverse representations of the minority class, our innovative recombining method significantly enhances its generalizability to unseen data, further strengthening the effectiveness of our spatiotemporal GNN model.




Fig. 2. ROC curves of comparison models (left), and ablation models (right)

3 Experiments and Results

3.1 Datasets and Image Pre-processing

We compiled a meta-cohort with Dynamic Susceptibility Contrast (DSC) MRI of 444 glioma patients, including the UPenn-GBM dataset available from The Cancer Imaging Archive (TCIA), and an in-house dataset. Among patients, 30 (6.8%) are IDH mutants, while 414 (93.2%) are wildtype IDH. Patients were divided into training, validation, and testing sets in a 7:1:2 ratio.

The DSC images were first co-registered with post-contrast T1 using the FMRIB Linear Image Registration Tool (FLIRT) from FSL. Subsequently, skull stripping was performed using the Brain Extraction Tool (BET) in FSL. Normalization included histogram matching and voxel smoothing with SUSAN noise reduction according to BraTS preprocessing procedures. Finally, all MRIs were non-linearly transformed to the standard space by co-registering them to the SRI atlas [14], using the Advanced Normalization Tools (ANTs).

Tumor masks, delineating the tumor core, were generated following standard procedures outlined in the brain tumor segmentation benchmark. Initially, an initial tumor segmentation (automask) was generated using a pretrained nnUNet, which was then manually inspected and corrected by two clinical experts.

3.2 Experimental Settings and Implementations

We used two 24GB NVIDIA RTX5000 GPU and PyTorch. The main training stage had a maximum of 200 epochs, and an early stopping was used to avoid

overfitting when validation loss no longer decreased over 10 epochs. For retraining, we set the number of epochs as 20. The batch size for graph and 3D image data was 10, while for 4D image data was 1 to reduce computational costs. The optimizer was Adam [7], and the learning rate was 1e-4 for all trainings.

Our code is publicly available at <https://github.com/DaisyYan2000/PerfGAT>.

Table 2. The ablation results for IDH classification on the meta-cohort dataset.

Model Config	ACC (%)	B-ACC (%)	SPEC (%)	SENS (%)	F1 (%)	AUC (%)
w/o Graph structure learning	86.5	77.3	88.0	66.7	40.0	82.3
w/o Dual-attention	88.8	78.5	90.4	66.7	44.4	88.8
w/o Recombine augmentation*	83.1	83.2	83.1	83.3	40.0	89.4
w/o Spatial graph [†]	77.5	80.2	77.1	83.3	33.3	93.2
w/o Tumor local [‡]	78.7	80.8	78.3	83.3	34.5	92.0
PerfGAT	88.8	86.2	89.2	83.3	50.0	94.4

* Random resampling is used for mitigating imbalance datasets; † Only node attention for feature fusion; ‡ Only semantic attention for feature fusion.

3.3 Performance evaluation

We compare our PerfGAT with other state-of-the-art models (Table 1). The comparison methods are grouped into three categories: CNN-based spatial model for T1 and rCBV images, non-GNN spatiotemporal models (ResNet+ConvLSTM [22], SwiFT [6]), and GNN-based spatiotemporal models (GCN [8], GIN [21], GAT [16], EGAT [4], STGCN [23] and BrainGNN [10]). For all comparative models, we use the random resampling methods to mitigate imbalance datasets problem for fair comparisons.

The results show that PerfGAT outperforms all the comparison models in most metrics, achieving at least 21.2%, 58.2% and 16.5% over other models in balanced accuracy, F1 score, and AUC, respectively. Notably, the achievement in superior balanced-accuracy and sensitivity indicating the model performance in learning the minority class. Figure 2 plots the ROC curves of all models. This underscores the effectiveness of our method in classifying IDH mutations even in unbalanced datasets. Note that sensitivities are similar because of limited negative samples in the test sets.

3.4 Ablation Experiments

Table 2 shows that graph structure learning effectively reduces noise, improving at least 2.7% in all metrics. Regarding feature fusion, dual-attention mechanisms outperform simple concatenation by 9.8%, 24.9%, 12.6% and 6.3% in balanced accuracy, sensitivity, F1 and AUC, respectively. Compared to random resampling, the recombine augmentation method achieved at least 3.6% improvement in most metrics except sensitivity, indicating less data distortion. Additionally,

when removing the spatial graph or local tumor information, the model performance decreases in most metrics, suggesting that informative features can be identified from both local lesions and the global brain.

4 Conclusion

In this study, we propose a spatiotemporal GNN framework, PerfGAT, for predicting the genotype of glioma using pMRI data. Our approach optimizes temporal connections using graph structure learning, and fuses spatiotemporal features through a dual-attention mechanism. To tackle the class imbalance challenge, we propose a class-balanced augmentation technique tailored to spatiotemporal features. Our proposed method addresses key challenges in pMRI analysis, demonstrating superior performance compared to state-of-the-art approaches. Future research may include multi-modal graph analysis and clinical validation.

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

References

1. Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research* **16**, 321–357 (Jun 2002). <https://doi.org/10.1613/jair.953>, <http://dx.doi.org/10.1613/jair.953>
2. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples (2019)
3. Heo, D., Lee, J., Yoo, R.E., Choi, S.H., Kim, T.M., Park, C.K., Park, S.H., Won, J.K., Lee, J.H., Lee, S.T., Choi, K.S., Lee, J.Y., Hwang, I., Kang, K.M., Yun, T.J.: Deep learning based on dynamic susceptibility contrast MR imaging for prediction of local progression in adult-type diffuse glioma (grade 4). *Sci Rep* **13**(1), 13864 (Aug 2023). <https://doi.org/10.1038/s41598-023-41171-9>, <https://www.nature.com/articles/s41598-023-41171-9>, number: 1 Publisher: Nature Publishing Group
4. Kamiński, K., Ludwiczak, J., Jasiński, M., Bukala, A., Madaj, R., Szczepaniak, K., Dunin-Horkawicz, S.: Rossmann-toolbox: a deep learning-based protocol for the prediction and design of cofactor specificity in Rossmann fold proteins. *Brief Bioinform* **23**(1), bbab371 (Jan 2022). <https://doi.org/10.1093/bib/bbab371>
5. Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition (2020)
6. Kim, P.Y., Kwon, J., Joo, S., Bae, S., Lee, D., Jung, Y., Yoo, S., Cha, J., Moon, T.: Swift: Swin 4d fmri transformer (2023)
7. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017)
8. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks (2017)
9. Li, C., Wang, S., Liu, P., Torheim, T., Boonzaier, N.R., van Dijken, B.R., Schönlieb, C.B., Markowetz, F., Price, S.J.: Decoding

- the Interdependence of Multiparametric Magnetic Resonance Imaging to Reveal Patient Subgroups Correlated with Survivals. *Neoplasia* **21**(5), 442–449 (May 2019). <https://doi.org/10.1016/j.neo.2019.03.005>, <https://www.sciencedirect.com/science/article/pii/S1476558618306377>
10. Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., Scheinost, D., Staib, L.H., Ventola, P., Duncan, J.S.: BrainGNN: Interpretable Brain Graph Neural Network for fMRI Analysis. *Medical Image Analysis* **74**, 102233 (Dec 2021). <https://doi.org/10.1016/j.media.2021.102233>, <https://www.sciencedirect.com/science/article/pii/S1361841521002784>
 11. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows (2021)
 12. Louis, D.N., Perry, A., Wesseling, P., Brat, D.J., Cree, I.A., Figarella-Branger, D., Hawkins, C., Ng, H.K., Pfister, S.M., Reifenberger, G., Soffietti, R., von Deimling, A., Ellison, D.W.: The 2021 WHO Classification of Tumors of the Central Nervous System: a summary. *Neuro Oncol* **23**(8), 1231–1251 (Jun 2021). <https://doi.org/10.1093/neuonc/noab106>, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8328013/>
 13. Pinto, A., Pereira, S., Meier, R., Alves, V., Wiest, R., Silva, C.A., Reyes, M.: Enhancing Clinical MRI Perfusion Maps with Data-Driven Maps of Complementary Nature for Lesion Outcome Prediction, p. 107–115. Springer International Publishing (2018). https://doi.org/10.1007/978-3-030-00931-1_13, http://dx.doi.org/10.1007/978-3-030-00931-1_13
 14. 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 (May 2010). <https://doi.org/10.1002/hbm.20906>
 15. SHI, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.k., WOO, W.c.: Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In: *Advances in Neural Information Processing Systems*. vol. 28. Curran Associates, Inc. (2015), https://papers.nips.cc/paper_files/paper/2015/hash/07563a3fe3bbe7e3ba84431ad9d055af-Abstract.html
 16. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks (2018)
 17. Wang, X., Ji, H., Shi, C., Wang, B., Cui, P., Yu, P., Ye, Y.: Heterogeneous graph attention network (2021)
 18. Wang, X., Price, S., Li, C.: Multi-task Learning of Histology and Molecular Markers for Classifying Diffuse Glioma (Jun 2023). <https://doi.org/10.48550/arXiv.2303.14845>, <http://arxiv.org/abs/2303.14845>, arXiv:2303.14845 [cs, eess]
 19. Wei, Y., Chen, X., Zhu, L., Zhang, L., Schönlieb, C.B., Price, S., Li, C.: Multi-Modal Learning for Predicting the Genotype of Glioma. *IEEE Transactions on Medical Imaging* **42**(11), 3167–3178 (Nov 2023). <https://doi.org/10.1109/TMI.2023.3244038>, <https://ieeexplore.ieee.org/abstract/document/10042035>, conference Name: IEEE Transactions on Medical Imaging
 20. Wei, Y., Li, C., Price, S.J.: Quantifying Structural Connectivity in Brain Tumor Patients. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 519–529. Lecture Notes in Computer Science, Springer International Publishing, Cham (2021). https://doi.org/10.1007/978-3-030-87234-2_49

21. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? (2019)
22. Yao, J., Shi, Y., Cao, K., Lu, L., Lu, J., Song, Q., Jin, G., Xiao, J., Hou, Y., Zhang, L.: DeepPrognosis: Preoperative prediction of pancreatic cancer survival and surgical margin via comprehensive understanding of dynamic contrast-enhanced CT imaging and tumor-vascular contact parsing. *Medical Image Analysis* **73**, 102150 (Oct 2021). <https://doi.org/10.1016/j.media.2021.102150>, <https://www.sciencedirect.com/science/article/pii/S1361841521001961>
23. Yu, B., Yin, H., Zhu, Z.: Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. pp. 3634–3640 (Jul 2018). <https://doi.org/10.24963/ijcai.2018/505>, <http://arxiv.org/abs/1709.04875>, arXiv:1709.04875 [cs, stat]
24. Zhang, P., Wen, G., Cao, P., Yang, J., Zhang, J., Zhang, X., Zhu, X., Zaiane, O.R., Wang, F.: BrainUSL: Unsupervised Graph Structure Learning for Functional Brain Network Analysis. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 205–214. *Lecture Notes in Computer Science*, Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-43993-3_20