

A Lesion-aware Edge-based Graph Neural Network for Predicting Language Ability in Patients with Post-stroke Aphasia

Zijian Chen¹, Maria Varkanitsa², Prakash Ishwar¹, Janusz Konrad¹,
Margrit Betke³, Swathi Kiran² and Archana Venkataraman¹

¹ Department of Electrical and Computer Engineering, Boston University

² Center for Brain Recovery, Boston University

³ Department of Computer Science, Boston University

{zijianc,mvarkan,pi,jkonrad,betke,skiran,archanav}@bu.edu

Abstract. We propose a lesion-aware graph neural network (LEGNet) to predict language ability from resting-state fMRI (rs-fMRI) connectivity in patients with post-stroke aphasia. Our model integrates three components: an edge-based learning module that encodes functional connectivity between brain regions, a lesion encoding module, and a subgraph learning module that leverages functional similarities for prediction. We use synthetic data derived from the Human Connectome Project (HCP) for hyperparameter tuning and model pretraining. We then evaluate the performance using repeated 10-fold cross-validation on an in-house neuroimaging dataset of post-stroke aphasia. Our results demonstrate that LEGNet outperforms baseline deep learning methods in predicting language ability. LEGNet also exhibits superior generalization ability when tested on a second in-house dataset that was acquired under a slightly different neuroimaging protocol. Taken together, the results of this study highlight the potential of LEGNet in effectively learning the relationships between rs-fMRI connectivity and language ability in a patient cohort with brain lesions for improved post-stroke aphasia evaluation.

Keywords: Lesion-aware modeling · Graph neural networks · Functional connectivity · Data augmentation · Aphasia prediction.

1 Introduction

Stroke is one of the major causes for disability worldwide [7], with approximately one-third of stroke survivors affected by speech and language impairments, known as aphasia [1]. Resting-state fMRI (rs-fMRI) captures steady-state patterns of co-activation in the brain and provides a unique glimpse into the altered brain network organization due to the stroke [5]. Exploring this relationship is crucial for understanding the mechanisms underlying aphasia and for developing effective, personalized treatment strategies. However, developing models that can simultaneously accommodate patient-specific changes in functional connectivity due to a lesion (i.e., the stroke area) and use this information to predict generalized language impairments remains an open challenge.

Prior studies have attempted to predict language ability using neuroimaging data. The earliest work [17] developed a stacked random forest (RF) model that performed feature selection across multiple modalities and then used these features to predict the composite Aphasia Quotient scored from the revised Western Aphasia Battery, i.e., WAB-AQ [10]. Another study [11] used support vector regression (SVR) to predict WAB-AQ by stacking features from functional MRI, structural MRI, and cerebral blood flow data. Recent work [3] proposed a supervised learning method for feature selection and fusion methods to integrate features from different modalities and predicted WAB-AQ using RF and SVR. An earlier study [2] used similar multimodal ML methods to predict treatment response, rather than baseline functionality. Finally, Wang et al. [21] used persistent diagrams derived from patient rs-fMRI to identify aphasia subtypes [21]. While these studies represent seminal contributions, they largely treat the data as a "bag of features" and do not fully capitalize on network-level information.

We propose to address this gap with Graph neural networks (GNN), which represent the brain as a graph, where nodes correspond to regions of interest (ROIs) and edges represent functional connections between ROIs. Convolutions on the graph aggregate information from neighboring nodes or edges. They can be node-based, as seen in models like BrainGNN [12], GAT [20], and GIN [23], or edge-based, as formulated in the BrainNetCNN [9] and the HGCNN [8] models. GNNs have shown superior performance compared to traditional machine learning techniques in predicting cognitive outcomes related to autism [4,12], aging and intelligence [8], Alzheimer’s Disease [25], and ADHD [26]. However, these applications revolve around intact brain networks, which is not the case for a large lesion caused by stroke. Previous work [15,16] took the approach of masking out the lesioned ROIs from the input data. However, this strategy ignores the possibility of informative brain signals from around the lesion boundary. Another challenge is the limited availability of rs-fMRI data from stroke patients. One approach is to reduce the number of features in the analysis [3,11,17]. However, feature selection may inadvertently remove key information in the data, and prior studies have not been diligent about cleanly separating data used for feature selection from that used for performance evaluation [2,17].

In this paper, we introduce a novel lesion-aware edge-based GNN model, which we call LEGNet, that uses rs-fMRI connectivity to predict language ability in patients with post-stroke aphasia. LEGNet is designed to aggregate information from neighboring edges of the brain graph, thus aligning with both the nature of rs-fMRI connectivity and the distributed interactions that contribute to language performance. We incorporate lesion information into LEGNet by encoding the stroke size and position into the model and by using this encoding to constrain the graph convolution process. To address data scarcity, we draw from the approach of [14] and develop a comprehensive data augmentation strategy that inserts an "artificial lesion" into healthy neuroimaging data and simulates the corresponding impact on rs-fMRI connectivity and language ability. We demonstrate that LEGNet outperforms baseline deep learning methods on two in-house datasets of patients with post-stroke aphasia.




Fig. 1. The Lesion-aware Edge-based BrainGNN Model. **Top:** Edge-to-edge message passing and edge-to-node aggregation. **Bottom Left:** Patient-specific lesion size and position encoding. **Bottom Right:** Subgraph updating and language prediction.

2 Lesion-Aware GNN with Simulated Training Data

2.1 LEGNet Model Architecture

An overview of our LEGNet model architecture is shown in Fig. 1. LEGNet is designed to bridge the gap between the region- or node-based characterization of a lesion and rs-fMRI connectivity, which is defined on edges. As seen, our model includes three components: an edge-based learning module, a lesion encoding module, and a subgraph learning module that connects the two viewpoints.

Formally, let N be the number of ROIs in the brain. The input to LEGNet is the patient rs-fMRI connectivity $\mathbf{X} \in \mathbb{R}^{N \times N}$, which is obtained by exponentiating the correlation matrix computed from the mean time series of non-lesioned voxels within each ROI, as introduced in [15]. If the entire ROI lies within the lesion, then the time series is zero. The entries of \mathbf{X} can be viewed as *edge features* in the underlying brain graph defined on the ROIs.

Edge-Based Learning: From the input \mathbf{X} , LEGNet first performs an edge-to-edge convolution [9] given by the following relationship:

$$\mathbf{H}_{ij} = \phi \left(\sum_{n \in \mathcal{N}(i)} \mathbf{r}_n \mathbf{X}_{in} + \sum_{n \in \mathcal{N}(j)} \mathbf{c}_n \mathbf{X}_{nj} \right), \quad (1)$$

where $\mathbf{H}_{ij} \in \mathbb{R}^{d_0}$ is the feature map of edge (i, j) , $\mathcal{N}(i)$ is the set of neighboring nodes to ROI i , including i itself, $\mathbf{r}_n \in \mathbb{R}^{d_0}$ and $\mathbf{c}_n \in \mathbb{R}^{d_0}$ are the learnable filters for each node n , and ϕ is an activation function that is applied element-wise. Intuitively, Eq. (1) aggregates the connectivity information along neighboring edges that share the same end-nodes and updates the edge features accordingly.

Following this step, LEGNet maps the edge features back into the node space:

$$\mathbf{h}_i^{(1)} = \phi \left(\sum_{n \in \mathcal{N}(i)} \mathbf{g}_n \mathbf{H}_{in} + \mathbf{b}_1 \right), \quad i = 1, 2, \dots, N, \quad (2)$$

where $\mathbf{h}_i^{(1)} \in \mathbb{R}^{d_1}$ is the feature map of node i , $\mathbf{g}_n \in \mathbb{R}^{d_1 \times d_0}$ is the learnable filter, and $\mathbf{b}_1 \in \mathbb{R}^{d_1}$ is the learnable bias term from [9].

Lesion Encoding: The LEGNet lesion encoding module captures the size and position of the stroke for downstream processing. This is done by computing the percentage of spared gray matter p_i in each ROI i . We use this information to construct the diagonal lesion embedding matrix $\mathbf{L} \in \mathbb{R}^{N \times N}$ for each patient:

$$\mathbf{L} = \begin{bmatrix} | & | & | & \cdots & | \\ \mathbf{L}_1 & \mathbf{L}_2 & \cdots & \mathbf{L}_N \\ | & | & | & \cdots & | \end{bmatrix} = \begin{bmatrix} p_1 & & & & \\ & \ddots & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & p_N \end{bmatrix}, \quad (3)$$

If an ROI is intact, then $p_i = 1$ to indicate no lesion; otherwise, $0 \leq p_i < 1$.

Subgraph Learning: At a high level, the subgraph learning module divides the nodes/ROIs into k subgroups based on their lesion encoding information and their (learned) contributions to the final prediction. First, LEGNet uses the lesion encoding \mathbf{L} to update the node representations via:

$$\mathbf{h}_i^{(2)} = \phi \left(\sum_{j \in \mathcal{N}(i)} \mathbf{W}_j \mathbf{h}_j^{(1)} \right), \quad i = 1, 2, \dots, N, \quad (4)$$

where $\mathbf{h}_i^{(2)} \in \mathbb{R}^{d_2}$ is the updated representation for node i , and, inspired by [12], the filters $\mathbf{W}_j \in \mathbb{R}^{d_2 \times d_1}$ are parameterized using the lesion matrix \mathbf{L} as follows:

$$\text{vec}(\mathbf{W}_j) = \Theta_2 \cdot \psi(\Theta_1 \mathbf{L}_j) + \mathbf{b}_2. \quad (5)$$

The learnable parameters $\Theta_2 \in \mathbb{R}^{d_2 d_1 \times k}$ and $\Theta_1 \in \mathbb{R}^{k \times N}$ are shared across all regions and all subjects. The bias term is \mathbf{b}_2 , and ψ an activation function.

The assignment score for each node j is computed as $\psi(\Theta_1 \mathbf{L}_j)$ and depends on its lesion embedding. The score indicates the involvement of node j in each subgraph. In this way, ROIs with similar lesion information and functionality are grouped together and updated with similar filters. Following the subgraph learning, the updated node features $\mathbf{h}_i^{(2)}$ are fed into a fully connected layer, with dimension d_3 , to predict the scalar WAB-AQ, which quantifies language ability.

Training Loss: We train LEGNet using the mean squared error between the actual y_m and predicted \hat{y}_m language performance for each subject m , together with a ridge regularization term on the network filters:

$$\ell = \frac{1}{M} \sum_{i=1}^M (\hat{y}_i - y_i)^2 + \lambda R(\Theta_1, \Theta_2, \mathbf{b}_1, \mathbf{b}_2, \mathbf{r}, \mathbf{c}, \mathbf{g}), \quad (6)$$

where M is the total number of subjects and R is an L^2 -norm.

2.2 Synthetic Data Generation for Model Pre-Training

Given the heterogeneity of stroke, we pre-train LEGNet using a large simulated dataset. This pre-trained model is then fine-tuned using our small patient dataset. Our strategy is to insert "artificial lesions" into the neuroimaging data of healthy subjects and simulate its impact on rs-fMRI connectivity and language.

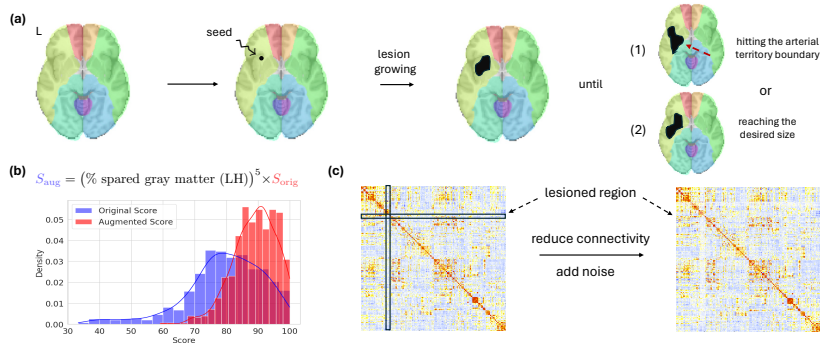


Fig. 2. Synthetic data generation workflow. **(a)** Artificial lesions are created to lie within a single arterial territory. **(b)** Language score density is adjusted after the lesion augmentation. **(c)** Functional connectivity is corrupted in the lesioned area.

Our pipeline for generating synthetic data is shown in Fig. 2. We first simulate a unique structural lesion for each subject based on the following rules: (1) lesions are left-hemisphere only; (2) lesions are placed randomly but do not cross arterial territories [13]; (3) lesion sizes range from 5% to 20% of one arterial territory; (4) lesions are spatially continuous and simply-connected (i.e., without holes in the inside). Next, the artificial lesion is used to mask out voxels when computing ROI mean time series. We also diminish and add Gaussian noise to the connectivity represented in \mathbf{X} between the lesioned region and the rest of the brain, followed by clipping the values to lie within the original connectivity range. Finally, the language performance score is re-scaled proportional to the percentage spared gray matter (< 1) to simulate the negative impact of the lesion on functionality.

2.3 Implementation Details

Simulated-Lesion HCP (HCP-SL): We use rs-fMRI data from 700 randomly selected subjects in the Human Connectome Project (HCP) S1200 database [19] as the foundation for generating synthetic data. Following the standard HCP minimal preprocessing pipeline [18], we parcellate the brain into 246 ROIs using the Brainnetome atlas [6]. The subject language score is accuracy in answering simple math and story-related questions during an fMRI language task. Artificial lesions are inserted and modify the data as described in Section. 2.2.

Pre-training: Pre-training is done via 10-fold cross validation (CV). We use a two-stage grid search to fix the model hyperparameters $\{\lambda, k, d_0, d_1, d_2, d_3\}$, with a coarse stage used to select a suitable power of 2 from 2^1 to 2^6 , followed by a fine stage with increments of 1-2. The regularizer λ is swept across $[10^{-4}, 1]$. The final values are $\{\lambda = 0.005, k = 8, d_0 = 4, d_1 = 8, d_2 = 2, d_3 = 8\}$. We use the Adam optimizer with learning rate starting from 0.01 and decaying by a factor of 0.95 every 20 steps. Early stopping is also applied based on the validation loss.

Once the hyperparameters are selected, we pre-train the LEGNet architecture in order to provide a better model initialization for our stroke datasets.⁴

Application to Post-Stroke Aphasia: We use repeated 10-fold CV on our larger in-house dataset (see Section 3.1) to evaluate the performance of LEGNet in a real-world setting. The hyperparameters and optimizations are fixed at the values determined during the pre-training phase on synthetic data. The same procedure is applied to all baseline methods. We compare two scenarios: (1) training the models (LEGNet and baselines) from scratch, and (2) using the pre-trained model as the initialization for our repeated CV experiment.

Cross-Dataset Generalization: As further validation, we quantify the language ability prediction performances when the models (LEGNet and baselines) are trained on DS-1 and applied to our second in-house dataset (DS-2) which has slightly different patient characteristics than DS-1.

2.4 Baseline Models

We compare LEGNet with four baseline approaches. The first baseline is a modified BrainGNN model (BrainGNN†) [12], which uses the same subgraph learning modules but does not perform edge-based learning or incorporate lesion information. The second baseline is BrainNetCNN model [9] with ROIs masked from the rs-fMRI connectivity input if the percentage of spared gray matter is less than 0.3 (BNC-masked). The third baseline is the BrainNetCNN with a two-channel input, with one channel being the unaltered rs-fMRI connectivity and the second channel being a lesion mask (BNC-2channel). The final baseline is support vector regression (SVR) with the lower-triangle of the rs-fMRI connectivity matrix used as the input feature vector. The baseline models inherit the appropriate subset of hyperparameters from LEGnet. We used the default radial basis function kernel with $C = 100$, $\gamma = 2/(N(N - 1))$, and $\epsilon = 0.1$.

3 Experimental Results

3.1 Datasets of Post-Stroke Aphasia

In-House Dataset 1 (DS-1): This dataset consists of 52 patients with chronic post-stroke aphasia with left-hemisphere lesions and aged between 35–80 years. Structural MRI (T1-weighted; TE=2.98ms, TR=2300ms, TI=900ms, res=1mm isotropic) and rs-fMRI (EPI; TE=20ms, TR=2/2.4s, res= $1.72 \times 1.72 \times 3\text{mm}^3$) were acquired on a Siemens 3T scanner. Both scans are pre-processed using the CONN toolbox [22]. Lesion boundaries were delineated manually by trained professionals and normalized to the MNI space. We use the Brainnetome atlas [6] to delineate 246 ROIs for the input rs-fMRI connectivity. Finally, all patients were evaluated using the WAB test [10] to obtain a measure of overall language

⁴ All code and synthetic data will be made public upon paper acceptance.

Table 1. Language prediction on DS-1 using repeated 10-fold CV for LEGNet and the baselines in Section 2.4. The asterisk * indicates statistically worse performance ($p < 0.05$) compared to the best performing model in highlighted in bold.

Methods	RMSE	MAE	R^2	Correlation Coeff.
LEGNet	17.38 ± 0.44	12.56 ± 3.02	0.35 ± 0.03	0.61 ± 0.03
BrainGNN†	$19.46 \pm 0.30^*$	$14.02 \pm 2.59^*$	$0.29 \pm 0.04^*$	$0.59 \pm 0.02^*$
BNC-mask	$19.18 \pm 1.82^*$	11.61 ± 2.33	$0.22 \pm 0.05^*$	$0.55 \pm 0.03^*$
BNC-2channel	$22.78 \pm 0.53^*$	$19.41 \pm 0.89^*$	$0.21 \pm 0.10^*$	$0.53 \pm 0.05^*$
SVR	$20.45 \pm 0.35^*$	$16.85 \pm 0.26^*$	$0.13 \pm 0.05^*$	$0.55 \pm 0.03^*$
LEGNet (w/o HCP-SL)	$18.39 \pm 0.68^*$	$15.33 \pm 0.57^*$	$0.29 \pm 0.05^*$	0.58 ± 0.04

ability (i.e., WAB-AQ). This value ranges from 0-100 with lower scores indicating severe aphasia, and higher scores indicating mild aphasia.

In-House Dataset 2 (DS-2): This dataset consists of 18 patients with chronic post-stroke aphasia that were recruited separately from DS-1. While the inclusion criteria and neuroimaging acquisition protocols are the same as for DS-1, the distribution of WAB-AQ scores is different. This provides an ideal scenario to evaluate cross-dataset generalization of LEGNet and the baseline models.

3.2 Performance Characterization and Model Interpretation

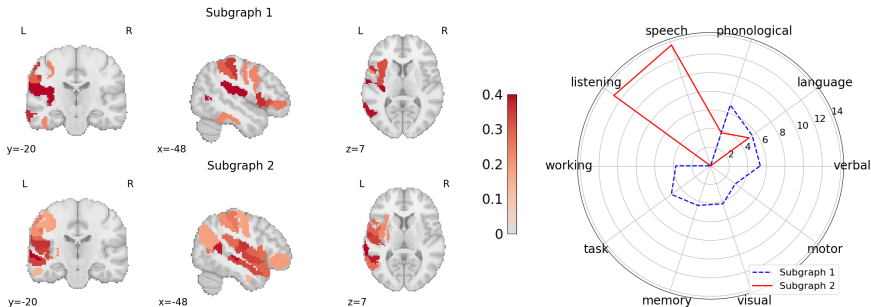
Table 1 reports the predictive performance of each method using repeated 10-fold CV on DS-1. LEGNet achieves the best performance in RMSE, R^2 , and correlation coefficient. While it is second-best to BNC-masked in MAE, the difference is not statistically significant. As a baseline, we applied the LEGNet architecture to DS-1 from a random initialization, i.e., without having access to synthetic data. To avoid data leakage, we selected the hyperparameters based on the corresponding modules used in previous studies [12,15]. We note a statistically significant decrease in performance w/o HCP, which underscores the importance of using synthetic data to design and initialize the deep network.

Fig. 3 (left) illustrates the top two subgraphs identified by LEGNet for the best-performing model during repeated 10-fold CV. The top subgraphs are identified by averaging the subgraph assignment scores for each ROI ($\psi(\Theta_1 \mathbf{L}_j)$ from Section 2.1) across all 52 patients in DS-1. We use Neurosynth [24] to decode the functionality associated with the ROIs assigned to each of the top two subgraphs, as shown in Fig. 3 (right). We note that LEGNet assigns high scores to regions that are related to the language ability. Intuitively, these regions also influence the prediction of language ability, as described in Section 2.1.

Finally, we tested generalization performance by applying the model that performs best on DS-1 to DS-2 without any fine-tuning (Table 2). In terms of R^2 , LEGNet maintains a leading position but, along with the other baselines, also shows a decrease compared to the validation performance in Table 1. While BrainGNN and SVR also show a decrease, BrainNetCNN-based models exhibit

Table 2. Language prediction on DS-2 with the model from repeated 10-fold CV that generalizes the best. Top performance is highlighted in bold.

Methods	RMSE	MAE	R^2	Correlation Coeff.
LEGNet	17.71	8.74	0.19	0.44
BrainGNN†	18.52	12.64	0.11	0.34
BNC-mask	19.24	12.68	0.04	0.31
BNC-2channel	19.70	13.47	0.01	0.28
SVR	18.43	12.65	0.12	0.35
LEGNet (w/o HCP-SL)	18.36	11.26	0.12	0.40

**Fig. 3.** **Left:** Top two subgraphs in in DS-1. **Right:** The z -score of each subgraph from Neurosynth [24] indicating the association strength with a particular function.

a sharper drop, indicating their reduced robustness on unseen data. The other three metrics follow a similar trend. This is expected due to the slight distribution shift between DS-1 and DS-2. Nevertheless, LEGNet still outperforms all baseline methods, indicating superior generalization ability.

4 Conclusion

We have introduced LEGNet, a novel lesion-aware edge-based graph neural network model designed to predict language performance in post-stroke aphasia patients from rs-fMRI connectivity. LEGNet bridges the gap between the lesion boundary defined on nodes and rs-fMRI connectivity defined on edges, while simultaneously using the lesion size and position to guide both the graph convolution and subgraph identification processes. Our synthetic data generation procedure addresses the challenge of limited patient data by simulating lesioned brain networks in healthy subjects. Pretraining on the augmented HCP dataset allows for unbiased hyperparameter selection and a reliable model initialization for fine-tuning on patient data. We demonstrate that LEGNet outperforms state-of-the-art methods in predictive accuracy and generalization ability, thus highlighting its potential as a reliable tool for post-stroke aphasia evaluation.

Acknowledgements

This work was supported by the National Institutes of Health R01 HD108790 (PI Venkataraman), the National Institutes of Health R01 EB029977 (PI Caffo), the National Institutes of Health R21 CA263804 (PI Venkataraman), the National Institutes of Health P50DC012283 (BU Site PI Kiran) and the National Institutes of Health R01 DC016950 (PI Kiran).

References

- Berthier, M.L.: Poststroke aphasia: epidemiology, pathophysiology and treatment. *Drugs & aging* **22**, 163–182 (2005)
- Billot, A., Lai, S., Varkanitsa, M., Braun, E.J., Rapp, B., Parrish, T.B., Higgins, J., Kurani, A.S., Caplan, D., Thompson, C.K., et al.: Multimodal neural and behavioral data predict response to rehabilitation in chronic poststroke aphasia. *Stroke* **53**(5), 1606–1614 (2022)
- Chennuri, S., Lai, S., Billot, A., Varkanitsa, M., Braun, E.J., Kiran, S., Venkataraman, A., Konrad, J., Ishwar, P., Betke, M.: Fusion approaches to predict post-stroke aphasia severity from multimodal neuroimaging data. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2644–2653 (2023)
- Dsouza, N.S., Nebel, M.B., Crocetti, D., Robinson, J., Mostofsky, S., Venkataraman, A.: M-GCN: A multimodal graph convolutional network to integrate functional and structural connectomics data to predict multidimensional phenotypic characterizations. In: *Medical Imaging with Deep Learning*. pp. 119–130. PMLR (2021)
- Falconer, I., Varkanitsa, M., Kiran, S.: Resting-state brain network connectivity is an independent predictor of responsiveness to language therapy in chronic post-stroke aphasia. *Cortex* (2024)
- Fan, L., Li, H., Zhuo, J., Zhang, Y., Wang, J., Chen, L., Yang, Z., Chu, C., Xie, S., Laird, A.R., et al.: The human brainnetome atlas: a new brain atlas based on connectural architecture. *Cerebral cortex* **26**(8), 3508–3526 (2016)
- Feigin, V.L., Brainin, M., Norrving, B., Martins, S., Sacco, R.L., Hacke, W., Fisher, M., Pandian, J., Lindsay, P.: World stroke organization (WSO): global stroke fact sheet 2022. *International Journal of Stroke* **17**(1), 18–29 (2022)
- Huang, J., Chung, M.K., Qiu, A.: Heterogeneous graph convolutional neural network via hodge-laplacian for brain functional data. In: *International Conference on Information Processing in Medical Imaging*. pp. 278–290. Springer (2023)
- Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017)
- Kertesz, A.: Western aphasia battery–revised (2007), <https://doi.org/10.1037/t15168-000>
- Kristinsson, S., Zhang, W., Rorden, C., Newman-Norlund, R., Basilakos, A., Bonilha, L., Yourganov, G., Xiao, F., Hillis, A., Fridriksson, J.: Machine learning-based multimodal prediction of language outcomes in chronic aphasia. *Human brain mapping* **42**(6), 1682–1698 (2021)
- 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 (2021)

13. Liu, C.F., Hsu, J., Xu, X., Kim, G., Sheppard, S.M., Meier, E.L., Miller, M.I., Hillis, A.E., Faria, A.V.: Digital 3D brain MRI arterial territories atlas. *Scientific Data* **10**(1), 74 (2023)
14. Nandakumar, N., Hsu, D., Ahmed, R., Venkataraman, A.: A deep learning framework to localize the epileptogenic zone from dynamic functional connectivity using a combined graph convolutional and transformer network. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2023)
15. Nandakumar, N., Manzoor, K., Agarwal, S., Pillai, J.J., Gujar, S.K., Sair, H.I., Venkataraman, A.: Automated eloquent cortex localization in brain tumor patients using multi-task graph neural networks. *Medical image analysis* **74**, 102203 (2021)
16. Nandakumar, N., Manzoor, K., Agarwal, S., Pillai, J.J., Gujar, S.K., Sair, H.I., Venkataraman, A.: A multi-scale spatial and temporal attention network on dynamic connectivity to localize the eloquent cortex in brain tumor patients. In: *Information Processing in Medical Imaging: 27th International Conference, IPMI 2021, Virtual Event, June 28–June 30, 2021, Proceedings* 27. pp. 241–252. Springer (2021)
17. Pustina, D., Coslett, H.B., Ungar, L., Faseyitan, O.K., Medaglia, J.D., Avants, B., Schwartz, M.F.: Enhanced estimations of post-stroke aphasia severity using stacked multimodal predictions. *Human brain mapping* **38**(11), 5603–5615 (2017)
18. Smith, S.M., Beckmann, C.F., Andersson, J., Auerbach, E.J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., et al.: Resting-state fMRI in the human connectome project. *Neuroimage* **80**, 144–168 (2013)
19. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al.: The WU-Minn human connectome project: an overview. *Neuroimage* **80**, 62–79 (2013)
20. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y., et al.: Graph attention networks. *stat* **1050**(20), 10–48550 (2017)
21. Wang, Y., Yin, J., Desai, R.H.: Topological inference on brain networks across subtypes of post-stroke aphasia. *ArXiv* (2023)
22. Whitfield-Gabrieli, S., Nieto-Castanon, A.: Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain connectivity* **2**(3), 125–141 (2012)
23. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018)
24. Yarkoni, T., Poldrack, R.A., Nichols, T.E., Van Essen, D.C., Wager, T.D.: Large-scale automated synthesis of human functional neuroimaging data. *Nature methods* **8**(8), 665–670 (2011)
25. Zhang, J., Guo, Y., Zhou, L., Wang, L., Wu, W., Shen, D.: Constructing hierarchical attentive functional brain networks for early AD diagnosis. *Medical Image Analysis* **94**, 103137 (2024)
26. Zhao, K., Duka, B., Xie, H., Oathes, D.J., Calhoun, V., Zhang, Y.: A dynamic graph convolutional neural network framework reveals new insights into connectome dysfunctions in ADHD. *Neuroimage* **246**, 118774 (2022)