A Federated Graph Learning Framework for Brain Connectome

Chaomeng Chen (gianluigi-chen@bupt.edu.cn) [1]
Sen Su (susen@bupt.edu.cn)
School of Computer Science, Beijing University of Posts and Telecommunications,
Beijing, China

Abstract

Neuroimaging, especially through Functional Magnetic Resonance Imaging (fMRI), plays a pivotal role in understanding brain activity by leveraging blood-oxygen level dependent (BOLD) signals to estimate neural activities across the brain. The interpretation of these signals through functional connectivity (FC) matrices facilitates the application of Graph Neural Networks (GNN) for analyzing brain network structures, offering insights into both normal and abnormal brain functions. Despite the potential of centralized learning methods in this domain, challenges related to data privacy and the feasibility of sharing sensitive medical datasets across institutions limit their application. This study introduces the Federated Graph Learning Framework for Brain Connectome (FGLBC), addressing these concerns. This novel approach enables the collaborative training of GNN models across multiple entities, such as hospitals, without compromising data privacy. The FGLBC framework implements a privacy-preserving local GNN training (PPGT) algorithm that incorporates Differential Privacy (DP) to safeguard sensitive information during model training. Furthermore, we introduce a unique similarity-weighted aggregation (SWA) algorithm that enhances the aggregation process, thereby boosting the global model's utility and performance. Our comprehensive evaluation across benchmark datasets demonstrates that the FGLBC not only preserves user privacy but also achieves or surpasses the performance of existing methods.

Keywords: Graph Neural Networks; Cognitive Science; Federated Learning; Brain Connectome

Introduction

Neuroimaging techniques are essential for capturing neural signals and measuring brain activity. Among these, Functional Magnetic Resonance Imaging (fMRI) is a prominent non-invasive method that measures whole-brain neural activity by tracking blood-oxygen level dependent (BOLD) signals at specific intervals (Matthews & Jezzard, 2004).. fMRI employs statistical metrics like Pearson correlation and mutual information to determine functional connectivity (FC) across different brain regions. These FC matrices are crucial for graph-based network analysis (Sporns, 2018; Wang, Zuo, & He, 2010), providing vital data for Graph Neural Network (GNN), which offer deep insights into both standard and abnormal brain functions (He et al., 2020).

GNN has become a dominant paradigm in deep learning, especially for graph-structured data (Wu et al., 2020). The structural similarity between brain networks and graphs has notably boosted the application of GNNs to analyze brain FC networks. Such analyses are crucial to pinpointing specific features or states in brain signals, greatly aiding the study of how cognitive functions rely on the integrated activities of various brain regions. However, most existing studies rely on centralized learning approaches (Y. Zhang, Tetrel, Thirion, & Bellec, 2021; Kim & Ye, 2020; Azevedo et al., 2022; Huang, Xia, Xu, & Qiu, 2022), which may not be viable where hospitals each maintain private brain imaging datasets due to privacy concerns.

In situations where multiple parties hold localized data, such as hospitals with individual fMRI datasets, a significant challenge is enabling collaborative training without risking data privacy. The Federated Graph Learning (FGL) (Lalitha, Kilinc, Javidi, & Koushanfar, 2019) framework provides a promising solution, allowing GNN models to be trained independently on local datasets while a central server manages weight aggregation and updates. This framework effectively handles data isolation and improves model generalizability. Despite the development of various FGL approaches, challenges persist in managing spatio-temporal graph data and identifying community structures within networks (Jiang, Jung, Karl, & Zhao, 2022; Luo et al., 2021; Baek, Jeong, Jin, Yoon, & Hwang, 2023). The main concerns include the vulnerability of gradients to inference attacks and the need for more precise FL aggregation techniques.

In response to these challenges, we have developed the Federated Graph Learning Framework for Brain Connectome (FGLBC), a novel approach designed for cross-institutional collaborative training of GNN models. FGLBC directly tackles privacy issues by integrating Differential Privacy (DP) (Dwork, 2006) into the gradient updates of GNN models through a privacy-preserving local GNN training (PPGT) algorithm. Additionally, our newly devised Weight Similarity-Weighted Aggregation (SWA) algorithm minimizes the negative impacts of uncooperative participants

[1]Corresponding Author. Email: gianluigi-chen@bupt.edu.cn

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*

on the global model's performance. The contributions of our research are summarized as follows:

- We introduce the Federated Graph Learning for Brain Connectome (FGLBC), optimized for collaborative GNN model training across multiple hospitals.
- We advance the field by proposing a novel privacy-preserving local GNN training (PPGT) and similarity-weighted aggregation (SWA) algorithm that enhances data privacy and aggregation efficacy.
- Through comprehensive experiments and analysis of benchmark datasets, our methods have demonstrated the capability to match or exceed the efficacy of existing methods while rigorously maintaining user privacy.

Related Work

Federated Graph Learning

Federated Graph Learning (FGL) has become a cornerstone in distributed Graph Neural Network training, significantly expanding GNN applications. FGL research is divided into three primary categories: inter-graph FGL, intra-graph FGL, and graph-structured FGL.

In inter-graph FGL, each participant utilizes a distinct graph dataset to enhance the efficiency of GNNs in modeling local data. This model's effectiveness has been demonstrated in various contexts, e.g., (Xie, Ma, Xiong, & Yang, 2021), and its adaptability for more generalized models has been explored, e.g., (Zhu et al., 2021). Additionally, it effectively handles the complexities of spatial-temporal graph data, as explored in recent studies (Jiang et al., 2022; Luo et al., 2021). Conversely, intra-graph FGL operates under a different premise where each client manages a segment of a larger graph. This configuration effectively resolves issues with missing links in subgraphs by supplementing neighboring data voids, as indicated in (K. Zhang, Yang, Li, Sun, & Yiu, 2021), and facilitates the identification of community structures (Baek et al., 2023). Finally, graph-structured FGL leverages graphical representations to elucidate the relationships between clients. This approach has shown versatility in handling various types of data, such as in image processing (Chen, Long, Wu, Zhou, & Jiang, 2022) and traffic data analysis (Meng, Rambhatla, & Liu, 2021).

Our research focuses on enhancing inter-graph FGL by developing a privacy-preserving FL framework tailored for brain connectome graphs.

GNN For Brain Connection

In brain network analysis, various methodologies for measuring functional connectivity (FC), such as Pearson correlation (Y. Zhang et al., 2021), (Kim & Ye, 2020), (Azevedo et al., 2022), (Huang et al., 2022), and partial correlation (Li et al., 2021), are fundamental for constructing brain graphs. These methods facilitate a deeper understanding of brain functions through the analysis of interactions between brain regions. However, they tend to overlook the dynamic nature of FC, remarkably its temporal variability, a crucial aspect highlighted by (Calhoun, Miller, Pearlson, & Adalı, 2014). Moreover, the selection of FC metrics profoundly influences the efficacy of subsequent analytical tasks, as noted by (Korhonen, Zanin, & Papo, 2021). In response to these challenges, recent advancements have been made in modeling dynamic brain graphs to more accurately reflect the non-stationary characteristics of brain activity. For instance, (Gadgil et al., 2020) introduced a technique utilizing spatio-temporal Graph Neural Networks, building upon the foundational work in spatio-temporal GNN by (Yan, Xiong, & Lin, 2018). This approach significantly enhances the integration of temporal dynamics, offering novel insights into the evolving nature of the brain. Concurrently, (Kim, Ye, & Kim, 2021) has applied the spatio-temporal attention mechanism from the transformer model, initially developed by (Vaswani et al., 2017), to effectively process and analyze dynamic brain graphs.

Preliminaries

Differential Privacy

Differential privacy, as established by (Dwork, 2006), provides a structured framework that offers provable safeguards against various privacy breaches. In this framework, two datasets, D and D', are considered adjacent, denoted as $D \simeq D'$, if they differ by only a single record.

Definition 1 (ε, δ)-Differential Privacy. Let $\mathcal{M} : D \to \mathcal{R}$ be a randomized mechanism that ensures (ε, δ)-differential privacy (ε-DP). This mechanism operates over a domain D and range $\mathcal{R}$, and guarantees that for all pairs of adjacent datasets D and D', and for any subset $\mathcal{Z}$ of $\mathcal{R}$:

$$\Pr(\mathcal{M}(D) \in \mathcal{Z}) \leq e^{\varepsilon} \times \Pr(\mathcal{M}(D') \in \mathcal{Z}) + \delta \tag{1}$$

Here, ε represents the privacy budget, which quantifies the allowable difference in outcomes between D and D', whereas δ denotes the probability of the mechanism failing to maintain ε-DP. Lower values of ε and δ enhance the rigor and strength of privacy protection.

Federated Learning

In a FL framework with M clients, each client m possesses a unique dataset D_m. The primary goal is to minimize the following objective function:

$$\min_{\mathbf{w}=\{w_1,\ldots,w_M\}} \frac{1}{M} \sum_{m=1}^{M} \frac{|D_m|}{N} L_m(w_m; D_m), \tag{2}$$

where N is the total number of data points across all clients, L_m denotes the loss function, and w_m are the model parameters for client m.

The server aggregates the model parameters from all clients as follows:

$$\bar{w}_{global} \leftarrow w_{global} + \frac{1}{M} \sum_{m=1}^{M} \frac{|D_m|}{N} w_m, \tag{3}$$

and then redistributes the updated global model to all clients.

Despite its potential, these methods can only perform with heterogeneous data distributions among clients. Recent innovations improve this by integrating personalized strategies into the local training and aggregation process, thus enhancing the performance of w_m on client m's dataset.

Spatio-Temporal Graph Convolutional Network (ST-GCN)

The ST-GCN (Gadgil et al., 2020) processes graph-structured data over time, making it highly suitable for applications like action recognition and brain activity prediction.

Notation Consider a graph $G = (V, E)$ with vertices V and edges E. Each vertex $v_i \in V$ is associated with a feature vector $x_i \in \mathbb{R}^F$, where F denotes the feature count. For dynamic scenarios, we analyze a series of graphs $\{G_t\}_{t=1}^{T}$, with T representing the timeline.

ST-GCN Model ST-GCN adapts convolution operations to graphs, defined as:

$$f_{out}(v_i) = \frac{1}{Z_{ii}} \sum_{v_j \in B(v_i)} f_{in}(v_j) \cdot \omega(v_j), \tag{4}$$

where $f_{in}(v_j)$ and $f_{out}(v_i)$ are the input and output features at node v_i, $B(v_i)$ represents the node's spatial-temporal neighborhood, ω denotes the convolutional kernel, and Z_{ii} is the normalization factor.

The graph convolution is split into spatial and temporal phases, each handled by distinct kernels. The spatial convolution at time t is:

$$f_i^{'} = \Lambda^{-\frac{1}{2}}(A+I)\Lambda^{-\frac{1}{2}} f_t W_{SG}, \tag{5}$$

with A as the adjacency matrix, Λ as the degree matrix, I as the identity matrix, f_t as the temporal feature matrix, and W_{SG} as the spatial kernel.

Classifying BOLD Time Series with ST-GCN ST-GCN excels in classifying Blood Oxygen Level Dependent (BOLD) time series, utilizing the graph's structural properties. The classification is achieved by directing ST-GCN's outputs to a fully connected layer with a sigmoid activation, generating class probabilities.

Methodology

This section presents a comprehensive overview of the FGLBC framework. We then present two core components of FGLBC: Privacy-Preserving Local GNN Training and the Similarity-Weighted Aggregation algorithm.

Overview of FGLBC

This section introduces the FGLBC framework, illustrated in Figure 1. FGLBC operates through a central server collaborating with multiple hospitals. Each hospital contributes by training models on local subgraphs, followed by model updates through weight adjustments. The process involves the following steps:

1. **Privacy-Preserving Local GNN Training:** Hospitals execute the PPGT algorithm, details of which are provided later.

2. **Weight Upload:** Hospitals upload their model weights to the central server.

3. **Similarity-Weighted Aggregation:** The server performs aggregation using the SWA algorithm.

4. **Global Weight Update:** The server distributes the global weights back to the hospitals.

Privacy-Preserving Local GNN Training (PPGT)

In the PPGT, we enhance the application of the ST-GCN tailored for fMRI data analysis. This approach involves transforming functional brain networks into detailed spatio-temporal graphs. Central to this is our comprehensive Spatio-Temporal Graph Convolution (ST-GC) framework, detailed previously.

Our model represents brain networks using nodes that correspond to BOLD signals from different brain regions. The edges are designed to capture the intricate connectivity that encompasses both temporal dynamics and spatial relationships, thus facilitating a deeper understanding of the brain connectome's activity.

To protect the privacy of sensitive medical data, we integrate the Differential Privacy-Stochastic Gradient Descent (DP-SGD) algorithm into the ST-GCN classifier. A critical component is the precise adjustment of privacy controls to ensure patient data protection while maintaining the model's utility.

The critical element of our methodology is computing the gradient of the ST-GCN's loss function, $\mathcal{L}(\theta)$, with respect to the model parameters w, for each mini-batch of graph data x_i. We adhere to stringent privacy norms by rigorously clipping this gradient to ensure that the L_2 norm remains below a predefined threshold C. This step is crucial for balancing data utility with privacy.

After gradient clipping, we add a calculated amount of random Gaussian noise, $\mathcal{N}(0, \sigma^2 C^2 \mathbf{I})$, where $\sigma = \sqrt{2 \log \frac{1.25}{\delta}} / \varepsilon$. This is not just noise injection; it is a strategic enhancement of our privacy framework, meticulously calibrated to evaluate its impact on privacy protection.

In the classifier's final stage, we implement a fully connected layer that transitions from analyzing spatio-temporal features to producing actionable classification probabilities. This phase marks the progression towards practical predictive decision-making.

The Similarity Weighted Aggregation Algorithm (SWA)

To improve the aggregation effectiveness in FGL, we propose the SWA. The initial challenge of diverse model

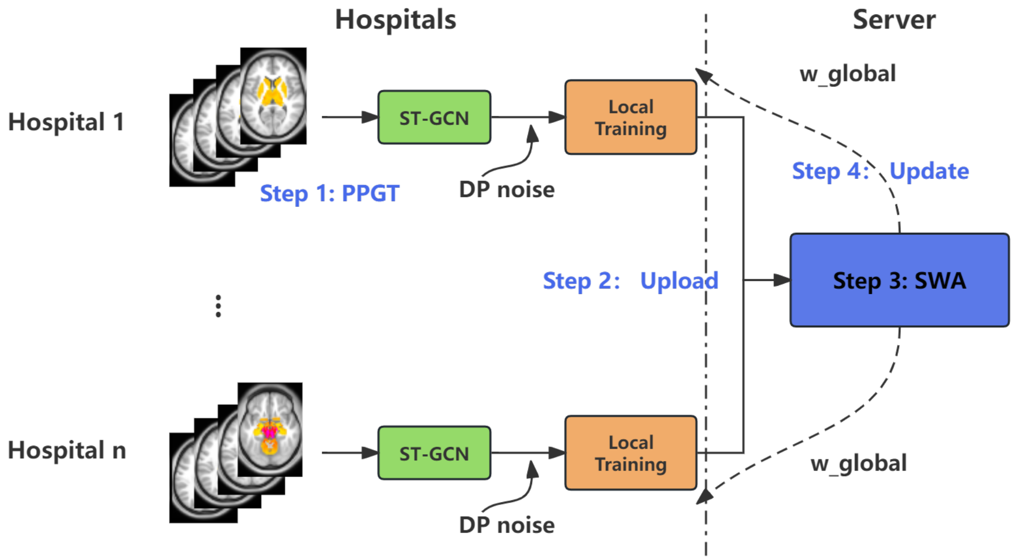

Figure 1: Framework of FGLBC

dimensions in GNNs is addressed by employing Principal Component Analysis (PCA) (Wold, Esbensen, & Geladi, 1987), which projects model weights w_j into a consistent, lower-dimensional space h_j. We then compute the cosine similarity between these projections and the global model weights h_{global}, defined as $S(j) = \frac{h_j \cdot h_{global}}{\|h_j\| \|h_{global}\|}$, to determine their similarity $S(i,j)$.

Once $S(i,j)$ is established, we update the global model by incorporating weights from various clients based on this similarity metric, adjusted to avoid negative values: $S(i,j) = \max(S(i,j),0)$. The aggregation method averages the local GNNs according to their functional similarity, utilizing the following formula:

$$w_{global} \leftarrow w_{global} + \sum_{j=1}^{K} \alpha_j \cdot w_j, \quad \alpha_j = \frac{\exp(\tau \cdot S(j))}{\sum_{j=1}^{K} \exp(\tau \cdot S(j))} \tag{6}$$

Here, α_j denotes the normalized similarity between client j and the global model, effectively adjusting the influence of each client's model. The hyperparameter τ, which can be increased to enhance community-specific model aggregation, regulates the sensitivity of similarity measures, emphasizing the contributions of similar subgraphs.

Experiments and Results

Datasets

We make use of two datasets: NCANDA and HCP in our experiments:

- **NCANDA** (Brown et al., 2015): The dataset includes 773 rs-fMRI scans from adolescents aged 12–21. It features scans from 388 individuals younger and 385 older than 16, with a gender distribution of 376 males and 397 females. Among them, 638 adolescents adhered to NIAAA guidelines for minimal alcohol consumption. Processing involved aligning each mean BOLD image with individual T1 MRI and MNI space and normalizing BOLD signals in 34 cortical ROIs to z-scores.
- **HCP** (Van Essen et al., 2013): This dataset includes rs-fMRI scans of 1,091 young adults aged 22–35, with 498 females and 593 males. After excluding five scans with insufficient frames, 1,200-frame scans from the first session (15 min, TR = 0.72 s) were used. The scans were processed through the HCP minimal pipeline and fMRISurface, aligning volumes to standard CIFTI gray coordinates space. Two subsets were created, HCP-Rest and HCP-Task, based on the activity status during scanning.

Baselines

We evaluate FGLBC against two baselines: (1) **FedAvg+DP**, which integrates a standard FL algorithm with a differential privacy (DP) mechanism, and (2) **FedGNN +DP**, a benchmark system for graph neural networks in federated learning, also incorporating standard DP.

Experimental Settings

The study involved four hospitals, with each client's privacy budget set at 10. The dataset was evenly distributed among the clients. The primary loss function employed was cross-entropy. All experiments were conducted on GPU servers equipped with 10 NVIDIA Tesla V100 GPUs, and the default τ value was set at 10.

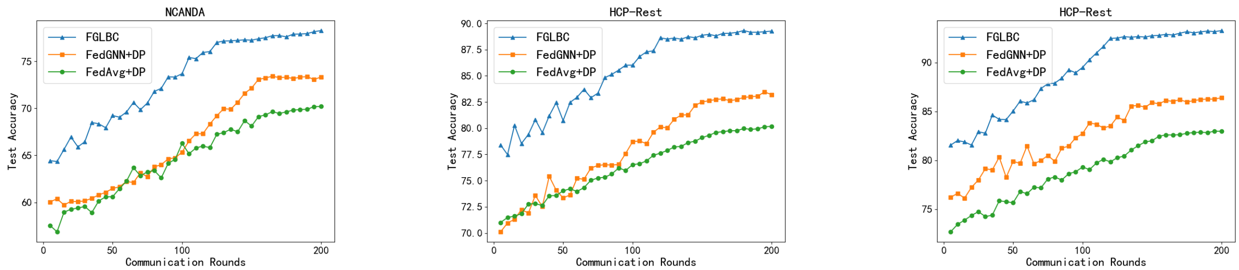

(a) **NCANDA** (b) **HCP-Rest** (c) **HCP-Task**

Figure 2: Test accuracy curves of our proposed FGLBC and other baselines methods along the communication rounds.

Performance and Analysis

Table 1: Accuracy on NCANDA, HCP-Rest and HCP-Task dataset.

Methods	NCANDA	HCP-Rest	HCP-Task
FGLBC	78.16	89.34	93.25
FedGNN+DP	74.22	83.12	86.28
FedAvg+DP	70.31	80.15	82.97

Performance Comparison

We conducted a comprehensive comparison of three methods: FGLBC, FedGNN +DP, and FedAvg+DP, assessing their accuracy on the NCANDA, HCP-Rest, and HCP-Task datasets. As demonstrated in Table 1, FGLBC consistently outperformed the other methods, notably achieving a remarkable accuracy of 93.25% on the challenging HCP-Task dataset. This performance underscores FGLBC's ability to simultaneously preserve model privacy and enhance effectiveness through an optimized weighted aggregation process. Conversely, FedGNN +DP and FedAvg+DP showed inferior performance across all datasets, with FedAvg+DP notably lagging at only 70.31% accuracy on the NCANDA dataset. These results underscore the superior adaptability and efficiency of FGLBC in managing diverse types of datasets.

Table 2: Effects of Different Numbers of Hospitals on HCP-Task Dataset.

Methods	2	4	6	8	10
FGLBC	91.02	93.24	93.10	93.04	92.98
FedGNN+DP	84.21	86.23	83.52	81.74	79.83
FedAvg+DP	80.07	82.97	80.17	77.93	72.57

Convergence Analysis

Figure 2 illustrates the average testing accuracy curves from five random training iterations across all methods. PPFGL consistently achieves higher average testing accuracy across all three datasets and shows faster convergence. This performance indicates that our SWA algorithm significantly improves the model's aggregation efficiency.

Effects of Different Numbers of Hospitals

We evaluated the model's performance within an expanded FGL framework by varying the number of participating hospitals from 2 to 10, as shown in Table 2. FGLBC maintained the highest and most consistent accuracy, experiencing only slight decreases as the number of hospitals increased, showcasing its robustness and scalability. Conversely, FedGNN+DP showed optimal performance with four hospitals but suffered significant accuracy losses with larger numbers, indicating its susceptibility to data heterogeneity. FedAvg+DP started with lower accuracy, which declined sharply with more than four hospitals, underscoring its challenges in managing larger graph datasets. These results demonstrate that our SWA effectively addresses data heterogeneity from multiple hospitals, maintaining model utility in federated learning.

Effects of Different Privacy Budgets

Table 3 presents the effects of varying privacy budgets (ε) on the performance of FGLBC, FedGNN+DP, and FedAvg+DP across three datasets: NCANDA, HCP-Rest, and HCP-Task. The data reveal a general trend where performance improves with increased privacy budgets. FGLBC consistently outperforms other methods in all scenarios, particularly at lower privacy budgets. Performance enhancements tend to plateau when the privacy budget exceeds 10, highlighting the efficacy of our methods in preserving accuracy despite the addition of varying levels of DP noise.

Conclusion

In this study, we introduce the FGLBC, designed to facilitate the collaborative training of GNN while prioritizing data privacy. Within FGLBC, each hospital independently trains its local GNN model. Model weights are then securely uploaded to a central server for aggregation before being redistributed to update the local models. The FGLBC framework incorporates a PPGT algorithm integrating DP to protect sensitive data during training. Additionally, we

Table 3: Effects of Different Privacy Budgets on Cora, CiteSeer and Pubmed datasets.

Mathods	NCANADA				HCP-Rest				HCP-Task			
	2	5	10	20	2	5	10	20	2	5	10	20
FGLBC	70.98	73.65	78.16	80.54	79.27	83.98	89.34	91.65	80.12	84.47	93.25	93.73
FedGNN+DP	65.48	68.99	74.22	76.15	72.67	76.05	83.12	84.65	73.12	80.54	86.28	88.81
FedAvg+DP	59.60	63.48	70.31	72.33	65.78	72.16	80.15	82.92	70.79	74.68	82.97	85.02

introduce a novel Similarity-Weighted Aggregation SWA algorithm that enhances the efficiency of the aggregation process, thus improving the utility and performance of the global model. Extensive experiments confirm that FGLBC significantly outperforms existing methods, establishing its efficacy in safeguarding privacy and enhancing model performance.

Acknowledgments

This work is supported by the National Natural ScienceFoundation of China (62072052) and the Innovation Research Group Project of NSFC (61921003).

References

Azevedo, T., Campbell, A., Romero-Garcia, R., Passamonti, L., Bethlehem, R. A., Lio, P., & Toschi, N. (2022). A deep graph neural network architecture for modelling spatio-temporal dynamics in resting-state functional mri data. *Medical Image Analysis*, *79*, 102471.

Baek, J., Jeong, W., Jin, J., Yoon, J., & Hwang, S. J. (2023). Personalized subgraph federated learning. In *International conference on machine learning* (pp. 1396–1415).

Brown, S. A., Brumback, T., Tomlinson, K., Cummins, K., Thompson, W. K., Nagel, B. J., … others (2015). The national consortium on alcohol and neurodevelopment in adolescence (ncanda): a multisite study of adolescent development and substance use. *Journal of studies on alcohol and drugs*, *76*(6), 895–908.

Calhoun, V. D., Miller, R., Pearlson, G., & Adalı, T. (2014). The chronnectome: time-varying connectivity networks as the next frontier in fmri data discovery. *Neuron*, *84*(2), 262–274.

Chalnick, A., & Billman, D. (1988). Unsupervised learning of correlational structure. In *Proceedings of the tenth annual conference of the cognitive science society* (pp. 510–516). Hillsdale, NJ: Lawrence Erlbaum Associates.

Chen, F., Long, G., Wu, Z., Zhou, T., & Jiang, J. (2022). Personalized federated learning with graph. *arXiv preprint arXiv:2203.00829*.

Dwork, C. (2006). Differential privacy. In *International colloquium on automata, languages, and programming* (pp. 1–12).

Feigenbaum, E. A. (1963). The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought.* New York: McGraw-Hill.

Gadgil, S., Zhao, Q., Pfefferbaum, A., Sullivan, E. V., Adeli, E., & Pohl, K. M. (2020). Spatio-temporal graph convolution for resting-state fmri analysis. In *Medical image computing and computer assisted intervention–miccai 2020: 23rd international conference, lima, peru, october 4–8, 2020, proceedings, part vii 23* (pp. 528–538).

He, T., Kong, R., Holmes, A. J., Nguyen, M., Sabuncu, M. R., Eickhoff, S. B., … Yeo, B. T. (2020). Deep neural networks and kernel regression achieve comparable accuracies for functional connectivity prediction of behavior and demographics. *NeuroImage*, *206*, 116276.

Hill, J. A. C. (1983). A computational model of language acquisition in the two-year old. *Cognition and Brain Theory*, *6*, 287–317.

Huang, S.-G., Xia, J., Xu, L., & Qiu, A. (2022). Spatio-temporal directed acyclic graph learning with attention mechanisms on brain functional time series and connectivity. *Medical Image Analysis*, *77*, 102370.

Jiang, M., Jung, T., Karl, R., & Zhao, T. (2022). Federated dynamic graph neural networks with secure aggregation for video-based distributed surveillance. *ACM Transactions on Intelligent Systems and Technology (TIST)*, *13*(4), 1–23.

Kim, B.-H., & Ye, J. C. (2020). Understanding graph isomorphism network for rs-fmri functional connectivity analysis. *Frontiers in neuroscience*, *14*, 630.

Kim, B.-H., Ye, J. C., & Kim, J.-J. (2021). Learning dynamic graph representation of brain connectome with spatio-temporal attention. *Advances in Neural Information Processing Systems*, *34*, 4314–4327.

Korhonen, O., Zanin, M., & Papo, D. (2021). Principles and open questions in functional brain network reconstruction. *Human Brain Mapping*, *42*(11), 3680–3711.

Lalitha, A., Kilinc, O. C., Javidi, T., & Koushanfar, F. (2019). Peer-to-peer federated learning on graphs. *arXiv preprint arXiv:1901.11173*.

Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., … Duncan, J. S. (2021). Braingnn: Interpretable brain graph neural network for fmri analysis. *Medical Image Analysis*, *74*, 102233.

Luo, M., Chen, F., Hu, D., Zhang, Y., Liang, J., & Feng, J. (2021). No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *Advances in Neural Information Processing Systems*, *34*, 5972–5984.

Matlock, T. (2001). *How real is fictive motion?* Doctoral dissertation, Psychology Department, University of California, Santa Cruz.

Matthews, P. M., & Jezzard, P. (2004). Functional magnetic resonance imaging. *Journal of Neurology, Neurosurgery &*

Psychiatry, *75*(1), 6–12.
Meng, C., Rambhatla, S., & Liu, Y. (2021). Cross-node federated graph neural network for spatio-temporal data modeling. In *Proceedings of the 27th acm sigkdd conference on knowledge discovery & data mining* (pp. 1202–1211).
Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
Ohlsson, S., & Langley, P. (1985). *Identifying solution paths in cognitive diagnosis* (Tech. Rep. No. CMU-RI-TR-85-2). Pittsburgh, PA: Carnegie Mellon University, The Robotics Institute.
Shrager, J., & Langley, P. (Eds.). (1990). *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
Sporns, O. (2018). Graph theory methods: applications in brain networks. *Dialogues in clinical neuroscience*, *20*(2), 111–121.
Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., Ugurbil, K., . . . others (2013). The wu-minn human connectome project: an overview. *Neuroimage*, *80*, 62–79.
Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., . . . Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.
Wang, J., Zuo, X., & He, Y. (2010). Graph-based network analysis of resting-state functional mri. *Frontiers in systems neuroscience*, *4*, 1419.
Wold, S., Esbensen, K., & Geladi, P. (1987). Principal component analysis. *Chemometrics and intelligent laboratory systems*, *2*(1-3), 37–52.
Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, *32*(1), 4–24.
Xie, H., Ma, J., Xiong, L., & Yang, C. (2021). Federated graph classification over non-iid graphs. *Advances in Neural Information Processing Systems*, *34*, 18839–18852.
Yan, S., Xiong, Y., & Lin, D. (2018). Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 32).
Zhang, K., Yang, C., Li, X., Sun, L., & Yiu, S. M. (2021). Subgraph federated learning with missing neighbor generation. *Advances in Neural Information Processing Systems*, *34*, 6671–6682.
Zhang, Y., Tetrel, L., Thirion, B., & Bellec, P. (2021). Functional annotation of human cognitive states using deep graph convolution. *NeuroImage*, *231*, 117847.
Zhu, Q., Yang, C., Xu, Y., Wang, H., Zhang, C., & Han, J. (2021). Transfer learning of graph neural networks with ego-graph information maximization. *Advances in Neural Information Processing Systems*, *34*, 1766–1779.