
Joint Embedding of Structural and Functional Brain Networks with Graph Neural Networks for Mental Illness Diagnosis

Yanqiao Zhu^{1,2} Hejie Cui³ Lifang He⁴ Lichao Sun⁴ Carl Yang^{3†}

Abstract

Multimodal brain networks characterize complex connectivities among different brain regions from both structural and functional aspects and provide a new means for mental disease analysis. Recently, Graph Neural Networks (GNNs) have become a de facto model for analyzing graph-structured data. However, how to employ GNNs to extract effective representations from brain networks in multiple modalities remains rarely explored. Moreover, as brain networks provide no initial node features, how to design informative node attributes and leverage edge weights for GNNs to learn is left unsolved. To this end, we develop a novel multiview GNN for multimodal brain networks. In particular, we regard each modality as a view for brain networks and employ contrastive learning for multimodal fusion. Then, we propose a GNN model which takes advantage of the message passing scheme by propagating messages based on degree statistics and brain region connectivities. Extensive experiments on two real-world disease datasets (HIV and Bipolar) demonstrate the effectiveness of our proposed method over state-of-the-art baselines.

1. Introduction

Mental illness is nowadays highly prevalent and has shown to be impactful for people’s physical health. Recent years have witnessed a growing academic interest in brain network analysis, which has demonstrates its effectiveness in mental health analysis (Su et al., 2020). In neuroscience, brain networks are often represented in different modalities from structural (e.g., Diffusion Tensor Imaging (DTI))

and functional aspects (e.g., functional Magnetic Resonance Imaging (fMRI)) (Zimmermann et al., 2018). These networked data represent complex structures of human brain connectivities. For example, in fMRI networks, edge connections represent correlations among brain regions with functional stimulations. Therefore, they are of paramount research values to understanding biologically mechanisms of brain functions. Moreover, the existing body of research on brain network analysis suggests that different modalities of brain networks convey complementary information to each other and the fusion of multiple modalities could lead to consistent improvements for brain analysis (Sun et al., 2017; Ma et al., 2017a; Zhang et al., 2018).

Recently, Graph Neural Networks (GNNs) have emerged as a powerful tool for understanding graph-structured data in many domains (Kipf & Welling, 2017; Veličković et al., 2018; Yang et al., 2021), where graph structures (i.e., adjacency matrices) and node features are embedded into a low-dimensional space for downstream machine learning. Unlike previous shallow network embedding models that can be regarded as a certain case of matrix factorization, GNN is more powerful in terms of representation ability (Xu et al., 2019; Qiu et al., 2018), which makes it suitable for analyzing brain networks usually of high nonlinearity (Zhang et al., 2018).

To date, there remains a paucity of studies on applying deep GNNs in analyzing multiview brain networks and we identify two obstacles for effectively learning embeddings of brain regions in multimodal data. Firstly, due to the multimodal nature of brain network data, different modalities encode different biomedical semantics of brain regions. How to learn effective node embeddings in such a multiview setting with GNNs remains rarely explored. Secondly, unlike other conventional graphs such as social networks, most brain networks could be expressed in a form of weighted adjacency matrix describing connections among brain regions without initial node features. How to design informative node attributes and corresponding edge weights for GNNs to learn is left unsolved.

To address these aforementioned challenges, in this work, we develop a novel multiview framework for multimodal brain network analysis, which we refer to BrainNN for brevity. Unlike previous work which presumes the existence of a common underlying graph structure beneath different

¹Center for Research on Intelligent Perception and Computing, Institute of Automation, Chinese Academy of Sciences ²School of Artificial Intelligence, University of Chinese Academy of Sciences ³Department of Computer Science, Emory University ⁴Department of Computer Science and Engineering, Lehigh University. Correspondence to: Carl Yang <j.carlyang@emory.edu>.

modalities, we first treat the brain networks under different modalities as multiple views of the brain and resort to contrastive learning for jointly embedding structural and functional brain networks. Particularly, we ensure cross-view consistency by imposing a contrastive objective on node embeddings across different views. In this way, we are able to adaptively distill discriminative knowledge from each modality without the need of defining or learning a common brain network structure.

In addition, due to the lack of node and edge features in brain network analysis, we propose to derive informative attributes from the original multimodal data. For incorporating node attributes, we take advantage of existing structural descriptors, such as degree profiles (Cai & Wang, 2018) to extract features for each node. To make full use of edge connectivity signals, contrary to directly utilizing edge weights for aggregating neighborhoods, we propose a general message passing scheme for brain networks such that we properly embed edge weights into learned node representations. Through extensive experiments on two real-world brain disease classification datasets (HIV and Bipolar), we find that our BrainNN employing degree features and the message passing scheme with edge embeddings can consistently achieve better performance across different datasets.

2. The Proposed Method: BrainNN

2.1. Preliminaries

Problem definition. We consider the problem of multiview brain network analysis, where each brain network describes connectivities between brain regions in multiple modalities. Suppose we are given a dataset $\mathcal{M} = \{(\{\mathcal{G}_i^s, \mathcal{G}_i^f\}, y_i)\}_{i=1}^S$ consisting of S subjects, where \mathcal{G}_i^s and \mathcal{G}_i^f represent the structural and functional modalities of the i^{th} subject described in brain networks respectively and y_i is its corresponding disease label. Each modality can be described in a weighted graph $\mathcal{G}_i^* = (\mathcal{V}, \mathcal{E}_i^*, \mathbf{W}_i^*)$, where $\mathcal{V} = \{v_i\}_{i=1}^N$ is the node set of size N defined by region of interest (ROI) (same across subjects), $\mathcal{E}_i^* = \mathcal{V} \times \mathcal{V}$ is the edge set, and $\mathbf{W}_i^* \in \mathbb{R}^{N \times N}$ is the weighted adjacency matrix describing interconnections between ROIs. The purpose of multiview brain network analysis is to learn a low-dimensional representation for each subject. The learned subject embeddings can be used to facilitate disease diagnosis and treatment.

Message passing graph neural networks. GNNs are widely used as a backbone for extracting features of graph-structured data. For the modeling of a node v_i , a GNN involves two key components: (1) aggregating messages from its neighborhood and (2) updating its representation in the previous layer with the aggregated message. The two operations can be formulated as

$$\mathbf{a}_i^{(l)} = \text{agg}^{(l)} \left(\text{msg}^{(l)} \left(\left\{ \mathbf{h}_j^{(l-1)} \mid v_j \in \mathcal{N}_i \cup \{v_i\} \right\} \right) \right), \quad (1)$$

$$\mathbf{h}_i^{(l)} = \text{upd}^{(l)} \left(\mathbf{h}_i^{(l-1)}, \mathbf{a}_i^{(l)} \right). \quad (2)$$

For simplicity, we denote a GNN model containing L




Figure 1: Our proposed BrainNN composed of two stages: graph representation learning and multimodal fusion.

stacked layers as $f^{(L)}$ thereafter. For a graph \mathcal{G} , an extra readout function g is required to obtain a graph-level embedding:

$$\mathbf{z} = g \left(\left\{ \mathbf{h}_i^{(L)} \mid v_i \in \mathcal{V} \right\} \right), \quad (3)$$

where $\mathbf{h}_i^{(L)} = f^{(L)}(\mathcal{G}_i)$ is the node embedding for node i .

2.2. The Overall Framework

Figure 1 summarizes the overall framework of our proposed BrainNN method. We employ contrastive learning for aggregating information from multiple modalities (§2.3), where each modality is modeled as a single view network by a specialized message passing GNN (§2.4).

Initially, we construct node features via local statistics for each view of brain network and feed each modality graph into a GNN model; the resulting node-level embeddings are aggregated using a shared multilayer perception (MLP) model to jointly learn structural and functional embeddings. After that, we enforce cross-view consistency to adaptively integrate information from the two views by contrasting local and global representations.

2.3. Multimodal Fusion across Multiple Views

We compute graph embeddings for each subject and denote the embeddings for functional and structural views as $\mathbf{z}_i^f = g(f^{(L)}(\mathcal{G}_i^f))$ and $\mathbf{z}_i^s = g(f^{(L)}(\mathcal{G}_i^s))$ respectively. At the inference time, we aggregate the representations from the two views and take the average of them $\mathbf{z}_i = (\mathbf{z}_i^f + \mathbf{z}_i^s)/2$ as the graph representations for downstream tasks.

After that, inspired by recent success of graph contrastive learning (Veličković et al., 2019; Zhu et al., 2020; 2021a;b), we employ a contrastive objective that distinguishes node representations of one view with graph representations of the other and vice versa:

$$\mathcal{J}_{\text{con}} = \frac{1}{2S} \sum_{\mathcal{G}_i \in \mathcal{M}} \left[\frac{1}{N} \sum_{v_i \in \mathcal{V}} \left(I(\mathbf{h}_i^f; \mathbf{z}_i^s) + I(\mathbf{h}_i^s; \mathbf{z}_i^f) \right) \right]. \quad (4)$$

This global-local contrasting mode encourages multi-scale consistency of graph representations (Hassani & Khasahmadi, 2020) and is shown to be effective in graph-oriented end tasks. For optimization efficiency, we estimate the mutual information $I(X; Y)$ in Eq. (4) using Jason-Shannon Divergence (JSD) (Hjelm et al., 2019) described as

$$I(\mathbf{h}_i; \mathbf{z}_i) = -\text{sp}(-d(\mathbf{h}_i, \mathbf{z}_i)) - \frac{1}{N-1} \sum_{v_j \in \mathcal{V} \setminus \{v_i\}} \text{sp}(d(\mathbf{h}_i, \mathbf{z}_j)), \quad (5)$$

where $\text{sp}(x) = \log(1+e^x)$ and d is a discriminator function, which takes the inner product with a sigmoid activation. Finally, we train the model with the unsupervised objective \mathcal{J}_{con} along with the supervised cross-entropy loss.

2.4. Message Passing GNNs for Brain Networks

Constructing node features for modality graphs. Non-attributed brain networks for graph classification bring challenges for applying graph neural network techniques. We take advantages of existing local statistical measures such as degree profiles. In particular, we study local degree profiles (LDP) (Cai & Wang, 2018) for each brain modality, where each feature \mathbf{x}_n of modality graph \mathcal{G}_{ij} is computed as

$$\mathbf{x}_n = [\text{deg}(n); \min(\mathcal{D}_n); \max(\mathcal{D}_n); \text{mean}(\mathcal{D}_n); \text{std}(\mathcal{D}_n)], \quad (6)$$

where $\mathcal{D}_n = \{\text{deg}(m) \mid (n, m) \in \mathcal{E}_{ij}\}$ describes the degree statistics of node n 's one-hop neighborhood, and $[\cdot; \cdot]$ denotes concatenation. Such computation can be done in $\mathcal{O}(\mathcal{E})$ time, which is computational friendly.

Implementing message passing GNN for handling edge weights. After obtained node features for each modality graph, we feed it into a message passing GNNs, with parameters shared across all modalities. For notation simplicity, we focus on one modality graph in this section and thus omit the subscript referring to specific graphs. Since the brain region connectivity is expressed in edge weights, we first construct a message vector $\mathbf{m}_{ij} \in \mathbb{R}^D$ composed of node embeddings of a node i , its neighborhood j , and edge weights w_{ij} :

$$\mathbf{m}_{ij}^{(l)} = t_{\Theta} \left(\left[\mathbf{h}_i^{(l)}; \mathbf{h}_j^{(l)}; w_{ij} \right] \right), \quad (7)$$

where t_{Θ} denotes a MLP layer parameterized by Θ , and l is the index of the current GNN layer.

Then, we aggregate messages from all neighborhoods followed by a non-linear transformation; the node-wise propagation rule can be written as:

$$\mathbf{h}_i^{(l)} = \sigma \left(\sum_{j \in \mathcal{N}_i \cup \{i\}} \mathbf{m}_{ij}^{(l-1)} \right), \quad (8)$$

where σ is a non-linear activation function such as

$$\text{ReLU}(\cdot) = \max(0, \cdot).$$

Finally, we employ another MLP with residual connections (He et al., 2016) for summarizing all node embeddings to compute graph-level embeddings $\mathbf{z} \in \mathbb{R}^D$:

$$\mathbf{z}' = \sum_{i \in \mathcal{V}} \mathbf{h}_i^{(k)}, \quad \mathbf{z} = t_{\Phi}(\mathbf{z}') + \mathbf{z}'. \quad (9)$$

3. Experiments

3.1. Experimental Protocols

Datasets. We use two real-world datasets (see details in (Ma et al., 2017b)) to evaluate the effectiveness of our method: **Human Immunodeficiency Virus Infection (HIV)** contains 35 patients (positive) and 35 seronegative controls (negative). Each modality graph is with 90 nodes and edge weights are calculated as correlations between brain regions. **Bipolar Disorder (BP)** consists of 52 bipolar subjects in euthymia and 45 healthy controls with matched age and gender. It stimulates 82 brain regions, according to Freesurfer-generated cortical/subcortical gray matter regions. In our experiments, we randomly use 80% of data for training, 10% for validation, and the remaining 10% for test.

Metrics. In experiments, we report the classification performance in terms of Accuracy and Area Under the ROC Curve (AUC), which are widely used for disease identification. Larger values indicate better performance.

Baselines. We include a broad range of baseline methods for comprehensive evaluation. For shallow embeddings methods, our experiments include M2E (Liu et al., 2018), MIC (Shao et al., 2015), MPCA (Lu et al., 2008), and MK-SVM (Dyrba et al., 2015). For these shallow baselines, we feed modality graphs as input and obtain the embeddings for all subjects. These embeddings are then fed into logistic regression classifiers.

We also consider the following up-to-date deep learning models: 3D-CNN (Gupta et al., 2013), GAT (Veličković et al., 2019), GCN (Kipf & Welling, 2017), DiffPool (Ying et al., 2018), and MVGCN (Zhang et al., 2018). Among them, 3D-CNN and MVGCN are designed for multiview learning and we perform classification in an end-to-end manner. For the other three baselines, we apply it with each single modality and report the best performance.

3.2. Results and Analysis

The overall performance is presented in Table 1. It is apparently seen that our proposed method achieves the best in two datasets. To be specific, BrainNN outperforms previous graph-based baselines MVGCN and DiffPool by large margins, up to 11% absolute improvements, which demonstrates the effectiveness of multiview fusion and the applicability of message passing GNN on brain networks. Compared to traditional tensor-based methods, such as M2E and MPCA,

Table 1: Comparison of different models on HIV and BP datasets. The highest performance is highlighted in boldface and passed significant tests with $p = 0.01$.

Method	HIV		BP	
	Accuracy	AUC	Accuracy	AUC
M2E	50.61	51.53	57.78	53.63
MIC	55.63	56.61	51.21	50.12
MPCA	67.24	66.92	56.92	56.86
MK-SVM	65.71	68.89	60.12	56.78
3D-CNN	74.31	73.53	63.33	61.62
GAT	68.58	67.31	61.31	59.93
GCN	70.16	69.94	64.44	64.24
DiffPool	71.42	71.08	62.22	62.54
MVGCN	74.29	73.75	62.22	62.64
BrainNN	77.14	79.79	73.64	67.54

Table 2: Performance of ablated models on HIV and BP datasets. Results passed significant tests with $p = 0.01$.

Method	HIV		BP	
	Accuracy	AUC	Accuracy	AUC
V-GCN	70.00	75.83	67.14	61.17
CONCAT	66.36	72.39	67.27	61.13
BrainNN	77.14	79.79	73.64	67.54

our work achieves significantly better results, which verifies the necessity of incorporating deep GNN for learning informative graph representations. The rationale of jointly embedding structural and functional brain networks can be further supported by the superiority of our work compared with deep models designed for single view such as GCN and GAT. Moreover, we notice that 3D-CNN obtains promising results, owing to the ability of modeling locality of input features. However, it computes graph representations with mere pooling layers, which fails to model the interaction among different views. Our work, on the contrary, leverages contrastive learning to model interaction of structural and functional views of brain networks, leading to discriminative representations in an adaptive manner.

3.3. Ablation Studies

We further verify the effectiveness of two key components in our model: message passing GNN and multiview fusion based on contrastive learning. Firstly, to examine the impact of the message passing GNN, we compare it with a vanilla Graph Convolutional Network model (**V-GCN**). The V-GCN directly treats the correlation weights as the adjacency matrix and performs weighted spectral convolution, which does not properly handle edge weights. Secondly, to validate the effectiveness of our multimodal fusion scheme, we further compare it against a simple baseline model named **CONCAT** that concatenates embeddings of all two views without multiview fusion. For fair comparison,

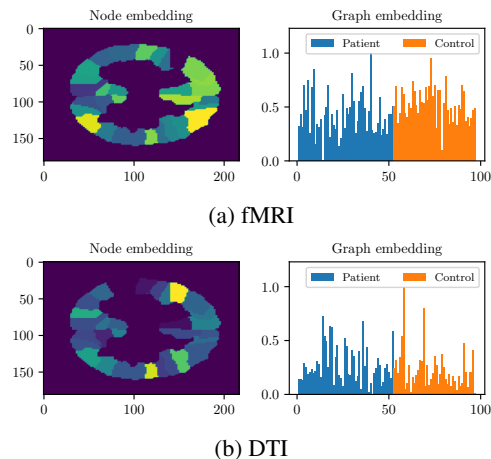


Figure 2: Visualization of embedded features H and Z from fMRI and DTI on the BP dataset.

all other experimental configurations are kept the same as in previous section. The results are summarized in Table 2. We observe that on the two datasets, our BrainNN method outperforms its two downgraded versions. It is worthy noting that the V-GCN still achieves close-to-optimal performance, demonstrating the effectiveness of applying deep GNNs with multiview interaction.

3.4. Visualization

To qualitatively examine the effectiveness of our method, we visualize the embeddings learned by our model from clinical perspective. In Figure 2, we show the learned node embeddings in the left panel, where the coordinate system represents neuroanatomy and color denotes the intensity of brain activities. It is seen from the figure that node embeddings from structural and functional perspectives are widely different to each other. We also present the visualization for graph embeddings representing the factor strengths for both patients and health controls. We observe a relatively positive correlation in the controls' embeddings, while the patients have a relatively negative correlation with the factors. It is thus evidently necessary to properly model and combine both views to deliver ideal clinical diagnoses.

4. Conclusion

In this work, we have proposed a novel BrainNN framework that jointly embeds multimodal brain networks with GNNs for mental illness diagnosis. Extensive experiments on two real-world datasets demonstrate the effectiveness of our proposed method. The study of applying GNNs for brain networks remains widely open with many challenges left to solve. For example, the current databases of small scales greatly confine the training of deep GNN models, for which we plan to investigate transfer learning and pre-training techniques (Hu et al., 2020a;b) in the near future.

Acknowledgements

Lifang He was supported by NSF ONR N00014-18-1-2009 and Lehigh's accelerator grant S00010293.

References

- Cai, C. and Wang, Y. A Simple Yet Effective Baseline for Non-Attributed Graph Classification. *arXiv.org*, November 2018.
- Dyrba, M., Grothe, M., Kirste, T., and Teipel, S. J. Multi-modal Analysis of Functional and Structural Disconnection in Alzheimer's Disease Using Multiple Kernel SVM. *Hum. Brain Mapp.*, 36(6):2118–2131, April 2015.
- Gupta, A., Ayhan, M., and Maida, A. Natural Image Bases to Represent Neuroimaging Data. In *ICML*, pp. 987–994, 2013.
- Hassani, K. and Khasahmadi, A. H. Contrastive Multi-View Representation Learning on Graphs. In *ICML*, pp. 4116–4126, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep Residual Learning for Image Recognition. In *CVPR*, pp. 770–778, 2016.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning Deep Representations by Mutual Information Estimation and Maximization. In *ICLR*, 2019.
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. Strategies for Pre-training Graph Neural Networks. In *ICLR*, 2020a.
- Hu, Z., Dong, Y., Wang, K., Chang, K.-W., and Sun, Y. GPT-GNN: Generative Pre-Training of Graph Neural Networks. In *KDD*, pp. 1857–1867, 2020b.
- Kipf, T. N. and Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*, 2017.
- Liu, Y., He, L., Cao, B., Yu, P. S., Ragin, A. B., and Leow, A. D. Multi-View Multi-Graph Embedding for Brain Network Clustering Analysis. In *AAAI*, pp. 117–124, 2018.
- Lu, H., Plataniotis, K. N., and Venetsanopoulos, A. N. MPCA: Multilinear Principal Component Analysis of Tensor Objects. *TNN*, 19(1):18–39, 2008.
- Ma, G., He, L., Lu, C.-T., Shao, W., Yu, P. S., Leow, A. D., and Ragin, A. B. Multi-view Clustering with Graph Embedding for Connectome Analysis. In *CIKM*, pp. 127–136, 2017a.
- Ma, G., Lu, C.-T., He, L., Yu, P. S., and Ragin, A. B. Multi-view Graph Embedding with Hub Detection for Brain Network Analysis. In *ICDM*, pp. 967–972, 2017b.
- Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., and Tang, J. Network Embedding as Matrix Factorization: Unifying DeepWalk, LINE, PTE, and node2vec. In *WSDM*, pp. 459–467, 2018.
- Shao, W., He, L., and Yu, P. S. Clustering on Multi-source Incomplete Data via Tensor Modeling and Factorization. In *PAKDD*, pp. 485–497, 2015.
- Su, C., Xu, Z., Pathak, J., and Wang, F. Deep Learning in Mental Health Outcome Research: A Scoping Review. *Transl. Psychiatry*, 10(116):1–26, April 2020.
- Sun, L., Wang, Y., Cao, B., Yu, P. S., Srisa-an, W., and Leow, A. D. Sequential Keystroke Behavioral Biometrics for Mobile User Identification via Multi-view Deep Learning. In *ECML-PKDD*, pp. 228–240, 2017.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph Attention Networks. In *ICLR*, 2018.
- Veličković, P., Fedus, W., Hamilton, W. L., Liò, P., Bengio, Y., and Hjelm, R. D. Deep Graph Infomax. In *ICLR*, 2019.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How Powerful are Graph Neural Networks? In *ICLR*, 2019.
- Yang, C., Xiao, Y., Zhang, Y., Sun, Y., and Han, J. Heterogeneous Network Representation Learning: A Unified Framework with Survey and Benchmark. *TKDE*, 2021.
- Ying, R., You, J., Morris, C., Ren, X., Hamilton, W. L., and Leskovec, J. Hierarchical Graph Representation Learning with Differentiable Pooling. In *NeurIPS*, pp. 4801–4811, 2018.
- Zhang, X., He, L., Chen, K., Luo, Y., Zhou, J., and Wang, F. Multi-View Graph Convolutional Network and Its Applications on Neuroimage Analysis for Parkinson's Disease. *AMIA Annu. Symp. Proc.*, 2018:1147–1156, December 2018.
- Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. Deep Graph Contrastive Representation Learning. In *GRL+@ICML*, June 2020.
- Zhu, Y., Xu, Y., Liu, Q., and Wu, S. An Empirical Study of Graph Contrastive Learning. *arXiv.org*, 2021a.
- Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. Graph Contrastive Learning with Adaptive Augmentation. In *WWW*, pp. 2069–2080, 2021b.
- Zimmermann, J., Griffiths, J. D., and McIntosh, A. R. Unique Mapping of Structural and Functional Connectivity on Cognition. *J. Neurosci.*, 38(45):9658–9667, November 2018.