A Novel GNN Framework Integrating Neuroimaging and Behavioral Information to Understand Adolescent Psychiatric Disorders

Weifeng Yu[1] DAG9WJ@VIRGINIA.EDU
Gang Qu[2] GQU1@TULANE.EDU
Young-geun Kim[3] KIMYO145@MSU.EDU
Lei Xu[1] LX557@NYU.EDU
Aiying Zhang[1] XSA5DB@VIRGINIA.EDU
[1] *School of Data Science, University of Virginia, Charlottesville, VA*
[2] *Department of Biomedical Engineering, Tulane University, New Orleans, LA*
[3] *Department of Statistics and Probability, Michigan State University, East Lansing, MI*

Editors: Under Review for MIDL 2025

Abstract

Functional connectivity (FC) is widely used to study various psychiatric disorders, but its consistency is often undermined by significant inter-subject variability. While these differences can be reflected by behavioral characteristics, few studies have combined them with FC. To this end, we propose a novel graph learning framework that enhances differentiation of psychiatric disorders by integrating FC with behavioral characteristics. Additionally, we apply Grad-CAM to enhance model interpretability by identifying key regions of interest involved in distinguishing individuals with psychiatric disorders from healthy controls. Experiments with the Adolescent Brain Cognitive Development dataset highlighted two critical insights: the thalamus and specific ROIs within the somatomotor and cingulo-opercular networks are vital for identifying psychiatric disorders. Additionally, visualization of latent representations indicated that individuals with externalizing disorders, specifically Oppositional Defiant Disorder, are distinguishable from healthy controls. These findings highlight the potential of our graph learning framework in discerning psychiatric disorders, offering potential for enhanced diagnostic accuracy.

Keywords: fMRI, adolescent psychiatric disorder, neurobehavior, graph autoencoder, interpretability.

1. Introduction

Late childhood and early adolescence are critical stages for brain functional development, often accompanied by the onset and development of multiple psychiatric problems, including anxiety disorders (ANX) (Siegel and Dickstein, 2011), obsessive–compulsive disorder (OCD), oppositional defiant disorder (ODD) (Ghosh et al., 2017), conduct disorder (Cond) (Fairchild et al., 2019; Stein et al., 2019), and attention-deficit hyperactivity disorder (ADHD) (Swanson et al., 1998; Sun et al., 2022), that affect cognitive development, social functioning, and overall quality of life, potentially leading to long-term impairments and increased risk for persistent psychiatric disorders in adulthood (Costello et al., 2003). It is crucial to understand the underlying neurobehavioral mechanisms of these disorders

at early stages and to identify biomarkers that could potentially inform the development of effective prevention and intervention strategies.

Functional magnetic resonance imaging (fMRI) offers a non-invasive, high-resolution method for capturing brain activity by detecting fluctuations in blood-oxygenation-level-dependent (BOLD) signals. Though BOLD signals provide substantial information about neuroactivities, their temporal structures make it challenging to reveal synchronous activity for interregional brain communication (Yan et al., 2022; Wang et al., 2024b). To alleviate this issue, functional connectivity (FC) (Smitha et al., 2017), estimated as the temporal association between different regions of interest (ROIs) derived from the fMRI time series data, has become a crucial tool for phenotype association study (Orlichenko et al., 2022) and psychiatric disorders research (Zhou et al., 2020; Zhang et al., 2019).

Graph Neural Networks (GNNs) are powerful tools for embedding graph-structured data, which is essential for integrating complex brain networks (Zhang et al., 2022; Wang et al., 2023; Zhu et al., 2022). This capability is particularly valuable for neuroimaging studies, as it facilitates comprehensive analysis and visualization of brain structures and the functional interactions between ROIs. Prior studies have demonstrated that representations from population-level graphs in graph learning are effective for demographic classifications, such as brain cognition (Qu et al., 2021a; Xiao et al., 2020) and development (Xiao et al., 2022; Chen et al., 2024). However, psychiatric disorder classification is inherently more challenging than intelligence classification due to high inter-subject variability and FC heterogeneity (Langhammer et al., 2024; Wang et al., 2024a), which often obscure condition-specific patterns. Accordingly, FCP-GNN (Gu et al., 2024) leveraged demographic information to construct edges and used FC as node features in a population graph, effectively distinguishing autism spectrum disorder. However, their approach did not incorporate behavioral characteristics and focused on a single type of psychiatric disorder, highlighting the necessity to enhance representations and broaden applications to the classification of psychiatric disorders.

To this end, we propose **B**ehavioral **E**dge **G**eneration **G**raph **A**uto**E**ncoder (BEG-GAE), a novel GNN framework that integrates relevant behavioral characteristics with FC data to enhance the brain network representation underlying psychiatric disorders. In this approach, node features are derived from FC data, while edge features are informed by behavioral characteristics, enabling the model to capture subtle connectivity changes linked to psychiatric disorders. To further enhance the interpretability of the model, we adopt the gradient-weighted class activation mapping (Grad-CAM) (Selvaraju et al., 2017; Qu et al., 2021b) to highlight the ROIs that are most critical to the classification of psychiatric disorders.

2. Methodology

As shown in Figure 1, the BEG-GAE consists of four steps: 1) **Node (i.e., subject) feature extraction**: For each subject, node embeddings are generated from individual FC using a GAE (Kipf and Welling, 2016; Noman et al., 2024); 2) **Edge generation**: A weighted adjacency matrix representing subject similarities is estimated based on behavioral characteristics; 3) **Latent space visualization**: Extracting latent representations from the integrated graph using another GAE to help distinguish psychiatric disorders; 4) **Perfor-**

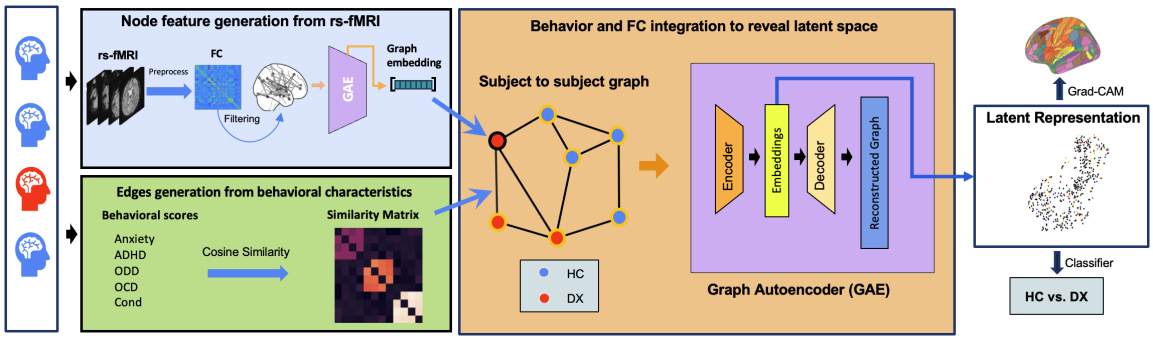

Figure 1: Schematic diagram for the BEG-GAE.

mance Evaluation and Feature Analysis: Performing classification for validation and feature analysis using logistic regression and Grad-CAM.

2.1. Embedding Extraction using Graph Autoencoder

We extract node embeddings using a GAE. For an input graph, all node features are concatenated into the feature matrix $\boldsymbol{X} \in \mathbb{R}^{n \times d}$, where n is the number of nodes and d is the feature dimensionality. During the encoding phase, the graph convolutional layer processes $\boldsymbol{X}$ and produces a latent representation $\boldsymbol{H}$:

$$\boldsymbol{H} = \sigma\big(\tilde{\boldsymbol{D}}^{-\frac{1}{2}}\,\tilde{\boldsymbol{A}}\,\tilde{\boldsymbol{D}}^{-\frac{1}{2}}\,\boldsymbol{X}\,\boldsymbol{W}\big) \quad \text{and} \quad \tilde{\boldsymbol{A}} = \boldsymbol{A} + \boldsymbol{I},$$

where $\tilde{\boldsymbol{A}}$ is the adjacency matrix with self-loop ($\boldsymbol{I}$ is the identity matrix), $\tilde{\boldsymbol{D}}$ is the diagonal matrix whose diagonal entries are the node degrees of $\tilde{\boldsymbol{A}}$, $\boldsymbol{W}$ corresponds to the learnable weight matrix, and $\sigma(\cdot)$ is the nonlinear activation function. The operation yields resulting $\boldsymbol{H}$, which amalgamates both the structural and attribute-based data from the graph.

In the decoding step, the latent representation $\boldsymbol{H}$ is used to reconstruct an approximation $\boldsymbol{X'}$ of the original feature matrix:

$$\boldsymbol{X}' = \sigma\big(\tilde{\boldsymbol{D}}^{-\frac{1}{2}}\,\tilde{\boldsymbol{A}}\,\tilde{\boldsymbol{D}}^{-\frac{1}{2}}\,\boldsymbol{H}\,\boldsymbol{W}'\big),$$

where $\boldsymbol{W'}$ is the reconstruction weight matrix. The optimization objective minimizes the Mean Squared Error (MSE) loss between $\boldsymbol{X}'_{\boldsymbol{i}}$ and $\boldsymbol{X}_{\boldsymbol{i}}$, defined as:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{nd}\sum_{i=1}^{n}\sum_{j=1}^{d}(\boldsymbol{X}_{ij} - \boldsymbol{X}'_{ij})^2$$

This optimization process is designed to encourage the model to learn embeddings $\boldsymbol{H}$ that retain essential information from $\boldsymbol{X}$, effectively capturing information for population-level graph embeddings.

2.2. Population Graph Generation

A population graph integrates all samples, where each node represents a subject. In our approach, we first construct an FC graph for each subject, with nodes representing ROIs.

The features of these nodes capture the functional relationships between ROIs, derived from the subject's FC matrix. We then generate graph embeddings for each subject-specific FC graph using a GAE during the **node feature generation** stage. These embeddings serve as the node features in the population graph. Finally, we construct edges of the population graph based on the cosine similarity between behavioral score vectors associated with each subject. For subjects i and j with behavioral score vectors $\boldsymbol{b}_i$ and $\boldsymbol{b}_j$, the cosine similarity is computed as $\mathcal{S}_{ij} = \frac{\boldsymbol{b}_i^\top \boldsymbol{b}_j}{\|\boldsymbol{b}_i\|\|\boldsymbol{b}_j\|}$. To ensure the symmetry of the similarity matrix, each element is defined as the maximum between corresponding elements across the matrix diagonal. Edges are established between nodes when the similarity values exceed a predefined threshold, thus controlling the sparsity of the graph and ensuring that the connections accurately reflect substantive feature similarities.

2.3. Model Interpretability with Grad-CAM

Grad-CAM is applied for model interpretability, leveraging the gradient information flowing to compute the importance of each node with respect to the predicted class scores. Specifically, it computes the gradients of the predicted class score $\boldsymbol{y}^c$ with respect to the node embeddings $\boldsymbol{h}^k$ of a graph convolutional layer. The gradient $\boldsymbol{\alpha}_k^c$ for each node embedding k with respect to class c is calculated as $\boldsymbol{\alpha}_k^c = \frac{1}{Z}\sum_i \frac{\partial \boldsymbol{y}^c}{\partial \boldsymbol{h}_i^k}$, where Z is the number of nodes in the layer, and $\frac{\partial \boldsymbol{y}^c}{\partial \boldsymbol{h}_i^k}$ denotes the partial derivative of the score $\boldsymbol{y}^c$ with respect to each node i in the embedding $\boldsymbol{h}^k$.

The Grad-CAM heatmap $\boldsymbol{L}^c$ is then generated by aggregating the node embeddings weighted by their calculated importance, followed by a ReLU activation to ensure only positive contributions are retained: $\boldsymbol{L}^c = \text{ReLU}\left(\sum_k \boldsymbol{\alpha}_k^c \boldsymbol{h}^k\right)$. This method enables precise tracking of ROIs' contributions to the model decisions, enhancing interpretability by visually identifying key influential features.

3. Experiment and Result

3.1. Datasets

We investigated subjects from the University of Pittsburgh site of the Adolescent Brain Cognitive Development (ABCD) study, which is designed to explore brain development and mental health for children aged 9–10. Resting-state fMRI (rs-fMRI) and behaviors related to five primary psychiatric disorders were explored, including ANX, OCD, ADHD, ODD, and Conduct Disorder. Participants with less prevalent conditions were excluded, resulting in a final sample of 440 participants (188 female, 252 male, 334 healthy controls (HC), 106 all diagnosed disorders (DX), as shown in Figure 2).

Preprocessed rs-fMRI data from the ABCD study were analyzed following the standardized ABCD pipeline,including motion correction, B0 distortion correction, and gradient nonlinearity adjustments (Hagler Jr et al., 2019). We extracted 379 ROIs using the Glasser atlas (Glasser et al., 2016) for cortical parcellations and the Aseg atlas (Fischl et al., 2002) for subcortical parcellations. Behavioral characteristics were assessed using the Child Behavior Checklist (CBCL) (Thompson et al., 2019), which includes syndrome scales that evaluate the overall symptoms associated with each of the five primary psychiatric disorders.

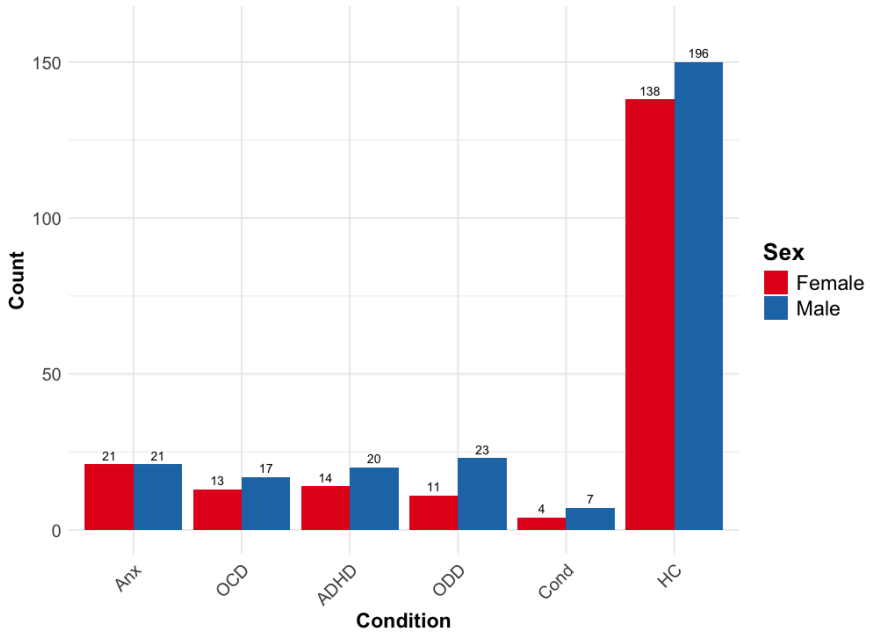

Figure 2: Diagnostic distribution by sex of subjects included in the analysis

3.2. Experimental Setup

To validate the effectiveness of various latent representations in distinguishing between HC and disorder groups, we employed two principal experimental approaches: t-SNE visualization (Van der Maaten and Hinton, 2008) and binary classification tasks. t-SNE visualization was employed to qualitatively demonstrate the separation of categories within the latent space, revealing clustering tendencies and distribution patterns among different groups. For quantitative evaluation, binary classification tasks were conducted using logistic regression with $L1$ regularization. Each classification task incorporated label balancing and 5-fold cross-validation to ensure robust and unbiased evaluation. The performance of the models was assessed using four metrics: accuracy, F1-scores, recall, and area under the receiver operating characteristic curve (AUC). The classification tasks included HC vs. internalizing disorders (i.e., ANX, OCD), HC vs. externalizing disorders (i.e., ADHD, ODD, Cond), and HC vs. DX. Building on these evaluations, we employed pairwise t-tests on the metrics from each cross-validation fold to statistically assess model performance. Metrics derived from the BEG-GAE model served as the baseline for comparison against other latent representations, allowing us to determine whether the observed differences were statistically significant.

To facilitate these experimental assessments, three baseline methods were compared against the proposed method, progressing from simple raw features to unsupervised learning, and finally to a graph-based approach.

- **Functional Connectivity Only (Flattened FC):** Utilizes raw flattened FC features directly, without embedding.

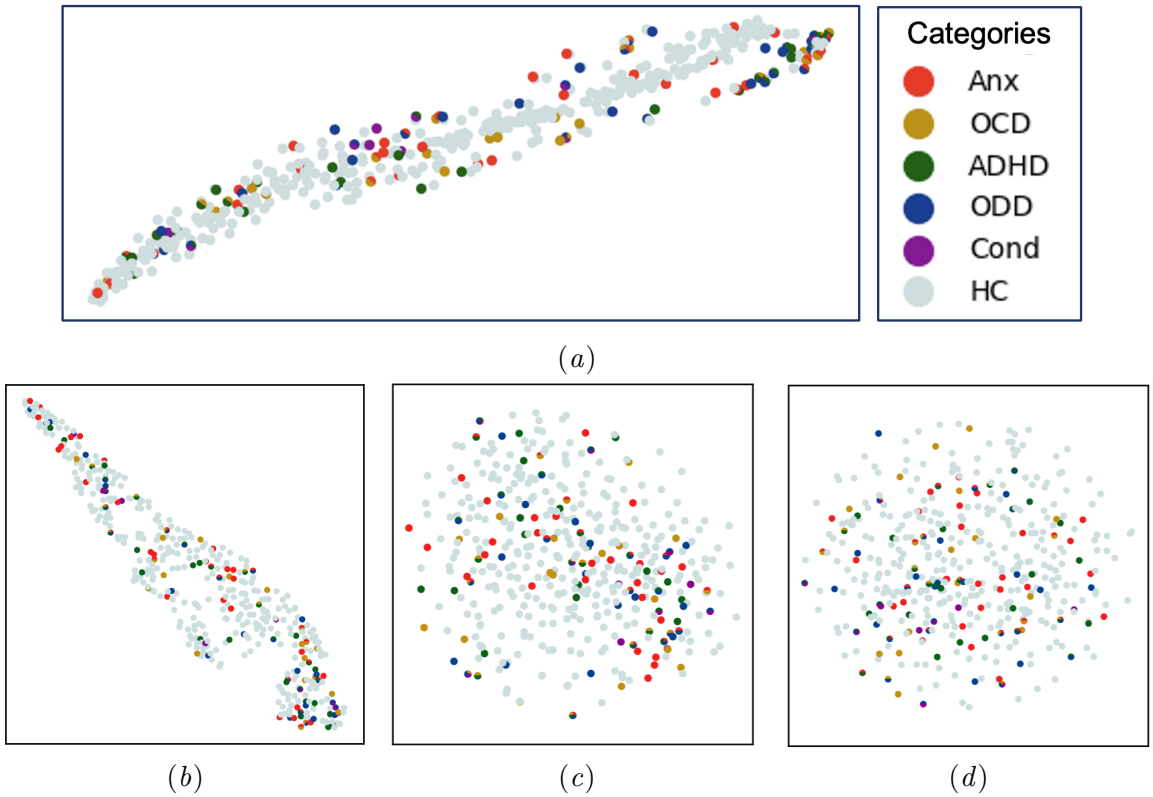

Figure 3: Comparison of t-SNE Dimensionality Reduction Across Different Frameworks: (a) BEG-GAE (**ours**); (b) GAE+FC; (c) MLP Autoencoder+FC; (d) Flattened FC.

- **Multilayer Perceptron Autoencoder (AE+FC):** Utilizes an MLP-based autoencoder to flattened FC matrices to learn latent embeddings.
- **Graph Autoencoder (GAE+FC):** Utilizes a Graph Autoencoder to flattened FC matrices to capture the underlying graph structure of brain connectivity.

3.3. Experimental Result

Distribution of Multilabel Embeddings: As shown in Figure 3, embeddings generated by other competing methods show minimal separation between psychiatric disorders and HC, with significant overlap in the latent space. BEG-GAE, through none of the five conditions, exhibit highly concentrated distributions. However, we observe that ODD—including cases of comorbidity with ADHD and Cond—tends to deviate more from the primary distribution. In contrast, ANX and OCD are more frequently aligned with the main distribution pattern.

Binary Classification: As shown in Table 1, BEG-GAE outperformed all other latent representations across all classification tasks. This indicates that combining behavioral scores with FC yields a more comprehensive representation, enhancing the discriminative capacity of our model, effectively differentiating HCs from those with psychiatric disorders. Additionally, the model achieved better performance in distinguishing the externalizing group from HC compared to the internalizing versus HC classification. Notably, t-tests

HC vs. Internalizing Disorders								
Framework	**Accuracy**	**p-value**	**F1**	**p-value**	**Recall**	**p-value**	**AUC**	**p-value**
Flattened FC	0.52 (0.03)	0.08	0.54 (0.04)	0.38	0.57 (0.09)	0.76	0.54 (0.05)	0.72
AE+FC	0.52 (0.16)	0.44	0.50 (0.21)	0.58	0.51 (0.24)	0.80	**0.60 (0.17)**	0.83
GAE+FC	0.56 (0.09)	0.33	0.56 (0.10)	0.66	**0.58 (0.14)**	0.50	0.57 (0.13)	0.96
BEG-GAE (Ours)	**0.62 (0.09)**	-	**0.59 (0.11)**	-	0.56 (0.13)	-	0.57 (0.12)	-
HC vs. Externalizing Disorders								
Framework	**Accuracy**	**p-value**	**F1**	**p-value**	**Recall**	**p-value**	**AUC**	**p-value**
Flattened FC	0.56 (0.08)	0.001	0.55 (0.11)	0.002	0.54 (0.12)	0.035	0.61 (0.09)	0.0009
AE+FC	0.52 (0.05)	0.0008	0.53 (0.03)	0.004	0.54 (0.07)	0.05	0.55 (0.07)	0.001
GAE+FC	0.55 (0.13)	0.04	0.56 (0.11)	0.04	0.56 (0.10)	0.08	0.51 (0.19)	0.02
BEG-GAE (Ours)	**0.79 (0.07)**	-	**0.78 (0.08)**	-	**0.79 (0.13)**	-	**0.82 (0.09)**	-
HC vs. DX								
Framework	**Accuracy**	**p-value**	**F1**	**p-value**	**Recall**	**p-value**	**AUC**	**p-value**
Flattened FC	0.53 (0.05)	0.001	0.52 (0.08)	0.005	0.52 (0.14)	0.021	0.53 (0.05)	0.004
AE+FC	0.55 (0.09)	0.05	0.55 (0.08)	0.028	0.57 (0.11)	0.019	0.57 (0.10)	0.077
GAE+FC	0.50 (0.03)	0.001	0.48 (0.07)	0.001	0.47 (0.11)	0.003	0.56 (0.02)	0.01
BEG-GAE (Ours)	**0.73 (0.06)**	-	**0.74 (0.06)**	-	**0.78 (0.10)**	-	**0.75 (0.07)**	-

* The mean and standard deviation (in parentheses) are reported.

Table 1: Classification Performance Comparison

conducted on the performance metrics revealed statistically significant differences for performance comparisons between our method and competing approaches, except for the HC versus internalizing disorders classification, where the results were not significant. We will further discuss the implications of the non-significant t-test results for the HC versus internalizing disorders group in the discussion section.

Identifying Key ROIs Relevant to Psychiatric Disorders: As shown in Figure 4, Grad-CAM analysis identifies several ROIs that help differentiate between HC and DX groups. Notably, the thalamus stands out for its role in sensory relay and regulation of consciousness. Additionally, regions within the somatomotor network, essential for voluntary motor control and coordination, and the cingulo-opercular network, crucial for cognitive control and emotional regulation, show strong associations with psychiatric disorders. These findings suggest that these networks and regions may play a significant role in understanding psychiatric disorders.

4. Discussion

4.1. Challenges in Differentiating Internalizing Disorders from Healthy Control

Our models for internalizing disorders, particularly ANX, demonstrated limited performance, consistent with prior research (Ipser et al., 2013) indicating that individuals with Specific Phobia—a subtype of ANX—exhibit significant variability in their neural signals. The young age of our cohort (9–10 years) introduces an additional challenge, as anxiety disorders typically evolve and manifest differently over time, complicating early detection efforts. As illustrated in Figure 3(a), the ANX group largely overlaps with healthy controls, further obstructing the identification of distinct diagnostic patterns. The classification outcomes, coupled with non-significant t-test results, suggest that individual variability and cohort-specific factors obscure the disorder-specific features we aimed to detect. Additionally, although fronto-parietal network regions are frequently implicated in psychiatric disorders, our model failed to identify these regions. Given their strong association with ANX (Ma et al., 2019), it is possible that their effects are more subtle or context-dependent, making them particularly challenging to isolate within our dataset.

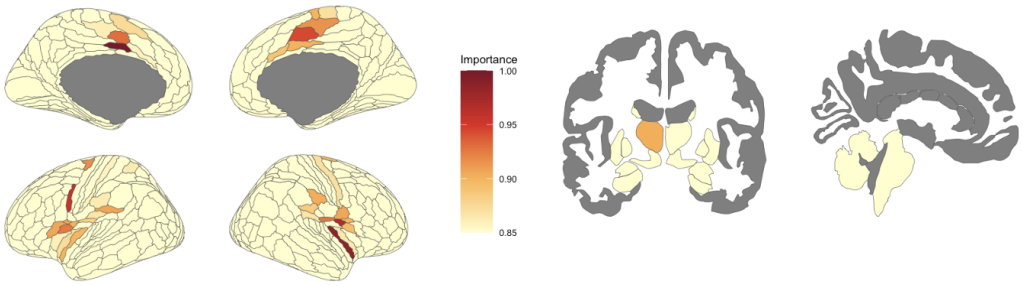

Figure 4: Grad-CAM–identified regions of interest (ROIs) related to psychiatric disorders. Left: Cortical regions; Right: Subcortical regions. Node importance values ranging from 0 to 1, representing each ROI's importance in distinguishing psychiatric disorders.

4.2. Limitations and Future Directions

This study exclusively focuses on data from a single site. While pooling data from multiple sites can increase sample size and improve statistical power, it introduces additional variability that risks obscuring the biological or functional patterns of interest. In future research, we plan to apply site-effect removal techniques, such as ComBat (Yu et al., 2018), to harmonize multi-site data and mitigate scanner-related variability. These approaches will allow us to utilize more samples from the dataset, facilitating more generalizable findings and broader applicability.

Multimodal fusion, as demonstrated in prior research, has been shown to enhance the richness and interpretability of learned representations across applications such as intelligence (Qu et al., 2024), sex classification (Patel et al., 2024), and brain cognition (Hu et al., 2021). However, this study is centered on FC, which provides valuable insights into neural interactions but overlooks other critical dimensions of brain organization. To address this limitation, future work will incorporate additional modalities, such as structural MRI (sMRI) and diffusion tensor imaging (DTI). By integrating these modalities with fMRI, the resulting graph representations are expected to capture complementary and diverse features of brain organization, thereby enriching the representation space and advancing our understanding of complex neural patterns.

5. Conclusion

We introduce BEG-GAE, an innovative framework that combines resting-state fMRI data with behavioral characteristics to advance the representation of psychiatric disorders. Our findings reveal that the BEG-GAE model generates representations that surpass traditional methods, including Autoencoders, Graph Autoencoders (GAEs), and raw functional connectivity features. Additionally, our analysis identifies key brain regions, particularly within the somatomotor and cingulo-opercular networks, as critical for classifying psychiatric disorders. These results underscore the potential of BEG-GAE to improve psychiatric diagnostics in late childhood and early adolescence by elucidating the intricate associations between brain connectivity and psychiatric disorders.

6. Acknowledgments

This work is partly supported by the UVA Brain Institute Transformative Neuroscience Pilot Grants. Data used in this study were obtained from the Adolescent Brain Cognitive Development[SM] (ABCD) Study (https://abcdstudy.org), held in the NIMH Data Archive (NDA).

References

Longyun Chen et al. Explainable spatio-temporal graph evolution learning with applications to dynamic brain network analysis during development. *NeuroImage*, 298:120771, 2024.

E Jane Costello et al. Prevalence and development of psychiatric disorders in childhood and adolescence. *Archives of general psychiatry*, 60(8):837–844, 2003.

Graeme Fairchild et al. Conduct disorder. *Nature Reviews Disease Primers*, 5(1):43, 2019.

Bruce Fischl et al. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, 2002.

Abhishek Ghosh, Anirban Ray, and Aniruddha Basu. Oppositional defiant disorder: current insight. *Psychology research and behavior management*, pages 353–367, 2017.

Matthew F Glasser et al. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615): 171–178, 2016.

Yuheng Gu et al. A novel population graph neural network based on functional connectivity for mental disorders detection. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 221–233. Springer, 2024.

Donald J Hagler Jr et al. Image processing and analysis methods for the adolescent brain cognitive development study. *Neuroimage*, 202:116091, 2019.

Wenxing Hu et al. Interpretable multimodal fusion networks reveal mechanisms of brain cognition. *IEEE transactions on medical imaging*, 40(5):1474–1483, 2021.

Jonathan C Ipser, Leesha Singh, and Dan J Stein. Meta-analysis of functional brain imaging in specific phobia. *Psychiatry and clinical neurosciences*, 67(5):311–322, 2013.

Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.

Till Langhammer et al. Resting-state functional connectivity in anxiety disorders: a multicenter fmri study. *Molecular Psychiatry*, pages 1–10, 2024.

Zijuan Ma et al. Frontoparietal network abnormalities of gray matter volume and functional connectivity in patients with generalized anxiety disorder. *Psychiatry Research: Neuroimaging*, 286: 24–30, 2019.

Fuad Noman et al. Graph autoencoders for embedding learning in brain networks and major depressive disorder identification. *IEEE Journal of Biomedical and Health Informatics*, 2024.

Anton Orlichenko et al. Latent similarity identifies important functional connections for phenotype prediction. *IEEE Transactions on Biomedical Engineering*, 70(6):1979–1989, 2022.

Binish Patel et al. Explainable multimodal graph isomorphism network for interpreting sex differences in adolescent neurodevelopment. *Applied Sciences*, 14(10):4144, 2024.

Gang Qu et al. Brain functional connectivity analysis via graphical deep learning. *IEEE Transactions on Biomedical Engineering*, 69(5):1696–1706, 2021a.

Gang Qu et al. Ensemble manifold regularized multi-modal graph convolutional network for cognitive ability prediction. *IEEE Transactions on Biomedical Engineering*, 68(12):3564–3573, 2021b.

Gang Qu et al. Integrated brain connectivity analysis with fmri, dti, and smri powered by interpretable graph neural networks. *ArXiv*, pages arXiv–2408, 2024.

Ramprasaath R Selvaraju et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.

Rebecca S Siegel and Daniel P Dickstein. Anxiety in adolescents: Update on its diagnosis and treatment for primary care providers. *Adolescent health, medicine and therapeutics*, pages 1–16, 2011.

KA Smitha et al. Resting state fmri: A review on methods in resting state connectivity analysis and resting state networks. *The neuroradiology journal*, 30(4):305–317, 2017.

Dan J Stein et al. Obsessive–compulsive disorder. *Nature reviews Disease primers*, 5(1):52, 2019.

Wenxin Sun, Mingxuan Yu, and Xiaojing Zhou. Effects of physical exercise on attention deficit and other major symptoms in children with adhd: A meta-analysis. *Psychiatry research*, 311:114509, 2022.

James M Swanson et al. Attention-deficit hyperactivity disorder and hyperkinetic disorder. *The Lancet*, 351(9100):429–433, 1998.

Wesley K Thompson et al. The structure of cognition in 9 and 10 year-old children and associations with problem behaviors: Findings from the abcd study's baseline neurocognitive battery. *Developmental cognitive neuroscience*, 36:100606, 2019.

Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

Junqi Wang et al. Dynamic weighted hypergraph convolutional network for brain functional connectome analysis. *Medical Image Analysis*, 87:102828, 2023.

Wei Wang et al. Multiview hyperedge-aware hypergraph embedding learning for multisite, multiatlas fmri based functional connectivity network analysis. *Medical Image Analysis*, 94:103144, 2024a.

Yingying Wang et al. A deep dynamic causal learning model to study changes in dynamic effective connectivity during brain development. *IEEE Transactions on Biomedical Engineering*, 2024b.

Li Xiao et al. A manifold regularized multi-task learning model for iq prediction from two fmri paradigms. *IEEE Transactions on Biomedical Engineering*, 67(3):796–806, 2020. doi: 10.1109/TBME.2019.2921207.

Li Xiao et al. Distance correlation-based brain functional connectivity estimation and non-convex multi-task learning for developmental fmri studies. *IEEE Transactions on Biomedical Engineering*, 69(10):3039–3050, 2022.

Weizheng Yan et al. Deep learning in neuroimaging: Promises and challenges. *IEEE Signal Processing Magazine*, 39(2):87–98, 2022.

Meichen Yu et al. Statistical harmonization corrects site effects in functional connectivity measurements from multi-site fmri data. *Human brain mapping*, 39(11):4213–4227, 2018.

Aiying Zhang et al. Aberrant brain connectivity in schizophrenia detected via a fast gaussian graphical model. *IEEE Journal of Biomedical and Health Informatics*, 23(4):1479–1489, 2019. doi: 10.1109/JBHI.2018.2854659.

Hao Zhang et al. Classification of brain disorders in rs-fmri via local-to-global graph neural networks. *IEEE transactions on medical imaging*, 42(2):444–455, 2022.

Zhongxing Zhou et al. Prediction and classification of sleep quality based on phase synchronization related whole-brain dynamic connectivity using resting state fmri. *NeuroImage*, 221:117190, 2020.

Yanqiao Zhu et al. Joint embedding of structural and functional brain networks with graph neural networks for mental illness diagnosis. In *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 272–276. IEEE, 2022.