Adaptive Graph Learning with Multi-Graph Convolutions for Brain Disorder Classification

Fuad Noman[1 [0000−0002−1756−8239]], Raphaël C.-W. Phan[1 [0000−0001−7448−4595]], Hernando Ombao[2 [0000−0001−7020−8091]], and Chee-Ming Ting[1 [0000−0002−6037−3728]]

[1] School of Information Technology, Monash University Malaysia
{fuad.noman,raphael.phan,ting.cheeming}@monash.edu
[2] Statistics Program, King Abdullah University of Science and Technology
hernando.ombao@kaust.edu.sa

Abstract. Functional Magnetic Resonance Imaging (fMRI) provides crucial insights into brain activity but presents challenges due to its high-dimensional, dynamic, and noisy nature. Traditional graph-based approaches for fMRI analysis often rely on predefined correlation structures, which may not accurately reflect the true underlying functional connectivity. To address this limitation, we propose a graph learning framework that dynamically constructs brain graphs and leverages Spline Convolutional Neural Networks (SplineCNN) for localized spatial feature extraction. Our model introduces a Learner Graph module, which infers graph structures in a data-driven manner, mitigating the reliance on predefined connectivity measures. The SplineCNN and Multi-Graph Convolution modules capture fine-grained spatial dependencies, offering improved adaptability to the heterogeneous nature of fMRI data. Additionally, we incorporate contrastive learning to align learned representations with domain-specific priors to improve generalization. Experimental results demonstrate that our approach outperforms traditional correlation-based methods in neurological disorder classification. The proposed framework provides a principled, adaptive solution for learning graph representations from fMRI, enhancing generalizability and robustness in brain network analysis.

Keywords: Adaptive graph convolution · Multi-level · Brain disorder.

1 Introduction

Functional Magnetic Resonance Imaging (fMRI) has emerged as a cornerstone for studying brain function, offering a non-invasive window into the dynamics of neural activity through Blood Oxygen Level Dependent (BOLD) signals [3]. This modality has been instrumental in uncovering aberrant functional connectivity patterns associated with neurological disorders such as Autism Spectrum Disorder (ASD) and Major Depressive Disorder (MDD) [4,5]. However, the high-dimensional, noisy, and temporally dynamic nature of fMRI data poses significant challenges for accurate brain network modeling and disorder diagnosis [6].

Traditional approaches often construct Functional Connectivity (FC) graphs using predefined measures like Pearson correlation or partial correlation, which may fail to capture the complex, subject-specific topological structures inherent in brain networks [7,8]. Recent advances in graph-based deep learning have shown promise in addressing these limitations by modeling brain networks as graphs, where nodes represent regions of interest (ROIs) and edges denote functional connections [9, 10]. Graph Convolutional Networks (GCNs), for instance, have been employed to leverage the non-Euclidean structure of brain connectomes, outperforming conventional Convolutional Neural Networks (CNNs) that assume regular grid-like data [11, 12].

Zeng et al. [13] introduced the Knowledge-driven Multi-Graph Convolutional Network (KMGCN), which integrates individual and population-level graphs to enhance brain network analysis, demonstrating superior performance in identifying biomarkers for Alzheimer's Disease and ASD. Similarly, Noman et al. [14] proposed a Graph Autoencoder (GAE) framework with GCNs to embed topological features of brain networks for MDD classification, highlighting the potential of unsupervised graph learning. Despite these advances, existing methods often rely on static or prior-driven graph structures, limiting their adaptability to the heterogeneous and dynamic nature of fMRI data across individuals and disorders. Also, transductive population-based GCN approaches, which often integrate multi-graph convolutions, face critical limitations during inference that compromise their generalizability. These methods commonly rely on population-level graph structures that incorporate information from both training and test sets, introducing a bias that inflates performance metrics and undermines fair evaluation [13]. Such dependence on test-time access to training data distributions is impractical for real-world clinical applications, where unseen data must be processed independently, and increases the risk of overfitting to specific training cohorts [16].

To overcome these challenges, we propose an Adaptive Graph Learning with Multi-Graph Convolutions (AGMGC) framework for fMRI-based brain disorder diagnosis. Our approach departs from predefined connectivity assumptions by introducing a Learner Graph module, which dynamically infers graph structures from raw fMRI time series in a data-driven manner. This is complemented by a Spline Convolutional Neural Network (SplineCNN) module [22], which extracts localized spatial features from the learned graphs, and a Multi-Graph Convolution module (MGCNet) that captures fine-grained dependencies for robust classification. Inspired by contrastive learning paradigms [15], we incorporate a contrastive loss to align the learned representations with domain-specific priors, enhancing generalization across datasets. We evaluate our framework on two prominent datasets: the Autism Brain Imaging Data Exchange (ABIDE) for ASD classification and the REST-meta-MDD Consortium database for MDD diagnosis, targeting the identification of disorder-specific connectivity patterns.

Our contributions are threefold: 1.) A novel graph learning framework that adaptively constructs brain graphs from fMRI, reducing reliance on static correlation measures. 2.) Integration of SplineCNN and multi-graph convolutions to

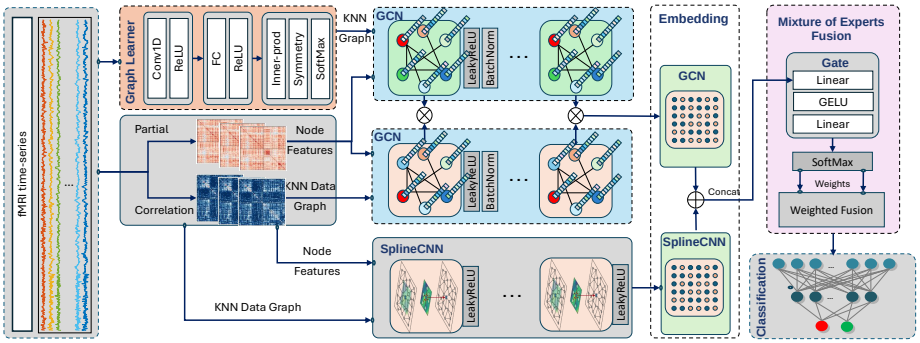

Fig. 1. Overall framework of the proposed AGMGC approach. The model employs a Learner Graph module to dynamically construct brain graphs from fMRI time series, augmented by data-driven graphs using partial correlation. SplineCNN and Multi-Graph Convolution modules process these multi-graph representations, capturing localized and global spatial features. Feature fusion and contrastive learning are applied to enhance generalization.

capture localized and global spatial dependencies, improving diagnostic accuracy. 3.) A contrastive learning strategy to enhance the robustness and generalizability of the learned representations.

2 Methods

As illustrated in Fig. 1, our AGMGC framework integrates adaptive graph construction, multi-graph convolutions, and contrastive learning to model fMRI-derived brain networks.

2.1 Data Preprocessing and Graph Representation

For each subject, we consider fMRI-derived graph $G = (V, E)$, where $V = \{v_1, \ldots, v_N\}$ denotes N ROIs (e.g., 200 using CC200 atlas), and E represents functional connections with corresponding adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. Let $\mathbf{T} \in \mathbb{R}^{N \times T}$ be the BOLD time series matrix, where T is the number of time points (truncated to the minimum T across subjects, e.g., 78 for ABIDE, 140 for REST-meta-MDD). We standardize $\mathbf{T}$ by subtracting the mean and dividing by the standard deviation computed across all subjects.

2.2 Learner Graph Module

Traditional methods define the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ using Pearson correlation or partial correlation [13]. Instead, our **Learner Graph module** infers $\mathbf{A}$ dynamically from $\mathbf{T}$. We employ a convolutional encoder inspired by [14]:

$$\mathbf{Z}^{(1)} = \text{ReLU}(\text{Conv1d}(\mathbf{T}; \mathbf{W}_c)) \tag{1}$$

where Conv1d is a 1D convolution with kernel size 3, stride 2, and weights $\mathbf{W}_c$, reducing the temporal dimension to $T' = \lfloor (T + 2 - 3)/2 \rfloor + 1$. The output $\mathbf{Z}^{(1)} \in \mathbb{R}^{N \times T'}$ is then transformed:

$$\mathbf{Z}^{(2)} = \text{ReLU}(\mathbf{W}_1 \cdot \text{mean}(\mathbf{Z}^{(1)}, \text{dim} = 1)) \tag{2}$$

$$\mathbf{A}_L = \text{softmax}(\mathbf{W}_2 \cdot \mathbf{Z}^{(2)} \cdot (\mathbf{W}_2 \cdot \mathbf{Z}^{(2)})^\top) \tag{3}$$

where $\mathbf{W}_1 \in \mathbb{R}^{T' \times 64}$ and $\mathbf{W}_2 \in \mathbb{R}^{64 \times N}$ are learnable parameters, and $\mathbf{A}_L \in \mathbb{R}^{N \times N}$ is the learned symmetric adjacency matrix normalized via softmax. To ensure sparsity, we apply kNN thresholding:

$$\mathbf{A}'_L = \text{kNN}(\mathbf{A}_L, k) \tag{4}$$

where $k = \lfloor 0.15N \rfloor$ retains the top 15% strongest connections.

2.3 Multi-Graph Construction

We augment $\mathbf{A}'_L$ with a data-driven graph $\mathbf{A}_K$, computed as the partial correlation of $\mathbf{T}$. Both $\mathbf{A}'_L$ and $\mathbf{A}_K$ are thresholded to ensure consistent sparsity, forming a multi-graph tuple $\{\mathbf{A}'_L, \mathbf{A}_K\}$.

2.4 SplineCNN and Multi-Graph Convolution

The **SplineCNN** module extracts localized features from $\mathbf{A}_K$ using a spline-based convolution:

$$\mathbf{H}_S = \text{SplineConv}(\mathbf{X}, \mathbf{A}_K; \mathbf{W}_S) \tag{5}$$

where $\mathbf{X} = \mathbf{A}_K + \text{diag}(\mathbf{1}_N)$ (pseudo-features), $\mathbf{W}_S$ are spline weights, and:

$$\mathbf{H}_S = \sum_{k=1}^{K} (\mathbf{B}_k \mathbf{W}_S^{(k)} \mathbf{X}) \tag{6}$$

with $\mathbf{B}_k$ as the spline basis computed from $\mathbf{A}_K$, and $K = 4$ basis functions. The output $\mathbf{H}_S \in \mathbb{R}^{N \times D}$ (e.g., $D = 32$) captures spatial dependencies.

The **Multi-Graph Convolution** (MGCNet) processes multi-graph features. First, multi-graph feature aggregation is performed as:

$$\mathbf{H_G}^{(0)} = (\mathbf{A}_K \mathbf{X}) \odot (\mathbf{A}'_L \mathbf{X}). \tag{7}$$

Then, graph convolutions are applied sequentially:

$$\mathbf{H_G}^{(1)} = \text{BN}(\text{LeakyReLU}(\mathbf{H_G}^{(0)} \Theta_1) \Theta_2), \tag{8}$$

$$\mathbf{H_G}^{(2)} = (\mathbf{A}_K \mathbf{H_G}^{(1)}) \odot (\mathbf{A}'_L \mathbf{H_G}^{(1)}). \tag{9}$$

Here, Θ_i are learnable weight matrices, $\odot$ denotes element-wise multiplication, and BN represents batch normalization.

The embeddings from SplineCNN $\mathbf{H}_S \in \mathbb{R}^{N\times D}$ and the Multi-Graph Convolution module $\mathbf{H}_G \in \mathbb{R}^{N\times D}$ are fused using a dynamic gated Mixture of Experts (MoE) approach to adaptively combine their contributions. The MoE module concatenates the two embeddings along the feature dimension to form $\mathbf{H}_{\text{cat}} = [\mathbf{H}_S, \mathbf{H}_G] \in \mathbb{R}^{N\times 2D}$. A gating network then computes weights for each expert:

$$\mathbf{W} = \text{softmax}(\text{FCN}_2(\text{GELU}(\text{FCN}_1(\mathbf{H}_{\text{cat}})))) \tag{10}$$

where $\text{FCN}_1 : \mathbb{R}^{2D} \to \mathbb{R}^4$ and $\text{FCN}_2 : \mathbb{R}^4 \to \mathbb{R}^2$ are fully connected layers, and $\mathbf{W} \in \mathbb{R}^{N\times 2}$ represents the softmax-normalized weights for $\mathbf{H}_S$ and $\mathbf{H}_G$ at each node. The final fused embedding is computed as:

$$\mathbf{H} = \mathbf{W}_{:,0:1} \cdot \mathbf{H}_S + \mathbf{W}_{:,1:2} \cdot \mathbf{H}_G \tag{11}$$

where $\mathbf{W}_{:,0:1}$ and $\mathbf{W}_{:,1:2}$ are the weight slices for $\mathbf{H}_S$ and $\mathbf{H}_G$, respectively, and $(\cdot)$ denotes element-wise multiplication with broadcasting. The fused embedding $\mathbf{H} \in \mathbb{R}^{N\times D}$ is then flattened and passed through a fully connected network for classification:

$$\mathbf{y} = \text{softmax}(\text{FCN}(\text{flatten}(\mathbf{H}))) \tag{12}$$

2.5 Training Loss

The model is trained using a composite loss function that combines classification, clustering, and contrastive objectives to enhance both predictive accuracy and graph representation quality. We employ a mixup-augmented cross-entropy loss to improve generalization. Given predictions $\mathbf{y}$, true labels $\mathbf{y}_a$ and $\mathbf{y}_b$ for two samples, and a mixing coefficient $\lambda \in [0, 1]$ sampled from a Beta distribution, the mixup classification loss is defined as:

$$\mathcal{L}_{\text{mixup}} = w_{\text{mixup-weight}} \cdot [\lambda \mathcal{L}_{\text{CE}}(\mathbf{y}, \mathbf{y}_a) + (1-\lambda)\mathcal{L}_{\text{CE}}(\mathbf{y}, \mathbf{y}_b)] \tag{13}$$

where $\mathcal{L}_{\text{CE}}$ is the cross-entropy loss, and mixup-weight $= 2$ amplifies the contribution of the mixup term. To ensure consistency in the learned graph topology, a mixup cluster loss $\mathcal{L}_{\text{cluster}}$ is applied to the learner topology $\mathbf{A}'_L$, aligning it with the interpolated targets $\mathbf{y}_a$ and $\mathbf{y}_b$ under the same λ:

$$\mathcal{L}_{\text{cluster}} = \text{cluster}_{\text{loss}}(\mathbf{A}'_L, \mathbf{y}_a, \mathbf{y}_b, \lambda) \tag{14}$$

Additionally, a contrastive loss aligns the learned adjacency matrix $\mathbf{A}'_L$ with the data-driven adjacency $\mathbf{A}_K$ (e.g., partial correlation $\mathbf{p}_{\text{corr}}$):

$$\mathcal{L}_{\text{cont}} = \frac{1}{N}\sum_{i=1}^{N} \|\mathbf{A}'_{L,i} - \mathbf{A}_{K,i}\|_2^2 \tag{15}$$

where $\mathbf{A}'_{L,i}$ and $\mathbf{A}_{K,i}$ are row vectors of the respective adjacency matrices. The total loss combines these terms:

$$\mathcal{L} = \mathcal{L}_{\text{mixup}} + \mathcal{L}_{\text{mixup-cluster}} + \alpha \mathcal{L}_{\text{cont}} \tag{16}$$

with α (set as $\lambda_{\text{contrastive}}$ in training) controlling the weight of the contrastive term, empirically tuned to balance the objectives.

3 Experiments and results

3.1 Datasets

We evaluate the proposed framework on two widely-used fMRI datasets: the ABIDE and the REST-meta-MDD Consortium database. ABIDE [1] comprises 1,035 subjects, including 530 individuals with ASD and 505 healthy controls (NC), with 878 males and 157 females, a mean age of 16.95 years, and a standard deviation of 8 years. The REST-meta-MDD database [2] includes 1,601 subjects, with 830 diagnosed with MDD and 771 NC, consisting of 622 males and 979 females, a mean age of 34.47 years, and a standard deviation of 12.36 years. These datasets provide diverse demographic and clinical profiles, enabling robust assessment of our model across ASD and MDD classification tasks.

3.2 Experimental Setup

Training and Evaluation We train AGMGC using AdamW optimizer (lr = 0.000761, wd = 5.3e-06) over 100 epochs with 5-fold stratified cross-validation on ABIDE and REST-meta-MDD. Performance metrics include accuracy (ACC), sensitivity (SEN), specificity (SPE), precision (PRE), F1-score (F1), and area under the receiver operating characteristic curve (AUC).

Comparison with Baseline Methods: To assess the performance of our AGMGC framework, we benchmark it against a range of baseline methods, including conventional machine learning techniques, Support Vector Machines (SVM) and Multi-Layer Perceptrons (MLP), and several prominent graph neural network models: PopulationGCN [18], BrainGNN [19], GATE [20], PLSNet [21], and KMGCN [13]. For fair evaluation, we utilized publicly available implementations of these competing methods, ensuring all models were trained and tested under identical computational conditions.

Ablation Study: To elucidate the individual contributions of our AGMGC framework's components, we conducted an ablation study by systematically evaluating the performance of key modules in isolation. Specifically, we assessed the efficacy of SplineCNN module, which captures localized spatial features from the learned graph topology, by training and testing the model with only this component active. Subsequently, we evaluated the MGCNet module, responsible for integrating multi-graph representations from both adaptive and data-driven adjacency matrices, in a separate configuration.

Table 1. Classification performance on the ABIDE dataset for NC vs. ASD.

	ACC	PRE	SEN	SPE	F1	AUC
SVM	$67.34_{\pm 2.26}$	$67.30_{\pm 3.03}$	$70.78_{\pm 02.70}$	$63.80_{\pm 4.47}$	$68.94_{\pm 1.87}$	$67.29_{\pm 2.29}$
MLP	$65.12_{\pm 6.13}$	$66.50_{\pm 5.55}$	$64.43_{\pm 08.20}$	$65.93_{\pm 6.66}$	$65.33_{\pm 6.36}$	$65.18_{\pm 6.10}$
PopulationGCN [18]	$80.73_{\pm 3.04}$	$79.55_{\pm 3.66}$	$85.59_{\pm 05.52}$	$75.34_{\pm 4.98}$	$82.35_{\pm 3.09}$	$80.47_{\pm 3.00}$
BrainGNN [19]	$70.74_{\pm 4.45}$	$71.91_{\pm 4.95}$	$73.66_{\pm 05.33}$	$67.54_{\pm 7.76}$	$72.65_{\pm 3.96}$	$70.60_{\pm 4.53}$
PLSNet [21]	$74.69_{\pm 3.73}$	$74.96_{\pm 1.57}$	$77.95_{\pm 10.52}$	$70.87_{\pm 5.41}$	$76.12_{\pm 5.51}$	$83.88_{\pm 3.46}$
KMGCN [13]	$77.68_{\pm 2.28}$	$80.99_{\pm 7.13}$	$75.99_{\pm 08.05}$	$79.44_{\pm 11.37}$	$77.94_{\pm 3.12}$	$\mathbf{85.88_{\pm 4.82}}$
GATE [20]	$75.10_{\pm 6.91}$	$74.71_{\pm 4.99}$	$79.49_{\pm 9.27}$	$70.21_{\pm 4.87}$	$76.97_{\pm 6.92}$	$74.85_{\pm 6.78}$
SplineCNN (Ours)	$79.53_{\pm 4.00}$	$83.11_{\pm 6.09}$	$77.38_{\pm 08.32}$	$81.85_{\pm 08.74}$	$79.78_{\pm 4.76}$	$79.61_{\pm 4.02}$
MGCNet (Ours)	$82.36_{\pm 2.37}$	$80.46_{\pm 3.37}$	$\mathbf{87.65_{\pm 2.15}}$	$76.58_{\pm 6.15}$	$83.84_{\pm 1.55}$	$82.11_{\pm 2.62}$
AGMGC (Ours)	$\mathbf{84.65_{\pm 3.36}}$	$\mathbf{86.14_{\pm 6.99}}$	$85.43_{\pm 06.15}$	$\mathbf{84.08_{\pm 9.16}}$	$\mathbf{85.44_{\pm 2.86}}$	$84.75_{\pm 3.28}$

Table 2. Classification performance on the REST-meta-MDD dataset for NC vs. MDD.

	ACC	PRE	SEN	SPE	F1	AUC
SVM	$64.71_{\pm 1.59}$	$64.70_{\pm 1.47}$	$70.36_{\pm 3.02}$	$58.62_{\pm 03.20}$	$67.38_{\pm 1.68}$	$64.49_{\pm 1.60}$
MLP	$62.71_{\pm 0.88}$	$64.69_{\pm 0.81}$	$61.93_{\pm 4.39}$	$63.56_{\pm 03.38}$	$63.20_{\pm 2.05}$	$62.74_{\pm 0.79}$
PopulationGCN [18]	$72.30_{\pm 1.55}$	$72.66_{\pm 1.37}$	$\mathbf{89.34_{\pm 1.90}}$	$43.93_{\pm 04.12}$	$80.12_{\pm 1.07}$	$66.63_{\pm 1.92}$
BrainGNN [19]	$63.63_{\pm 1.16}$	$66.60_{\pm 2.63}$	$84.84_{\pm 8.00}$	$28.23_{\pm 13.97}$	$74.35_{\pm 2.05}$	$56.53_{\pm 3.19}$
PLSNet [21]	$71.25_{\pm 1.16}$	$75.37_{\pm 3.86}$	$82.34_{\pm 7.75}$	$51.87_{\pm 13.48}$	$78.39_{\pm 1.79}$	$75.94_{\pm 1.15}$
KMGCN [13]	$68.75_{\pm 5.28}$	$74.65_{\pm 3.65}$	$77.12_{\pm 7.51}$	$54.05_{\pm 07.75}$	$75.74_{\pm 4.67}$	$72.19_{\pm 5.98}$
GATE [20]	$73.79_{\pm 1.43}$	$78.3_{\pm 1.59}$	$80.37_{\pm 0.45}$	$62.85_{\pm 3.32}$	$79.31_{\pm 0.96}$	$71.61_{\pm 1.8}$
SplineCNN (Ours)	$74.55_{\pm 2.02}$	$78.96_{\pm 3.46}$	$82.18_{\pm 2.82}$	$61.17_{\pm 9.22}$	$80.45_{\pm 1.11}$	$71.67_{\pm 3.47}$
MGCNet (Ours)	$79.59_{\pm 4.38}$	$79.87_{\pm 4.96}$	$82.95_{\pm 3.36}$	$75.71_{\pm 6.85}$	$81.34_{\pm 3.81}$	$79.33_{\pm 4.51}$
AGMGC (Ours)	$\mathbf{82.36_{\pm 2.84}}$	$\mathbf{82.11_{\pm 3.2}}$	$80.94_{\pm 6.11}$	$\mathbf{83.64_{\pm 03.90}}$	$\mathbf{81.40_{\pm 3.44}}$	$\mathbf{82.29_{\pm 2.93}}$

3.3 Results on ABIDE Dataset

Table 1 summarizes the classification performance on the ABIDE dataset. Among baseline methods, traditional machine learning approaches (SVM and MLP) achieved modest results, with SVM recording an accuracy of 67.34% and MLP at 65.12%, reflecting their limited capacity to model complex graph structures in fMRI data. State-of-the-art graph neural network models exhibited stronger performance, with PopulationGCN [18] achieving an accuracy of 80.73% and KMGCN [13] reaching an AUC of 85.88, highlighting the advantage of leveraging population-level graph information. However, these methods trailed our full AGMGC model, which attained the highest accuracy (84.65%), precision (86.14%), specificity (84.08%), and F1-score (85.44%), alongside a competitive AUC (84.75).

The ablation study reveals the complementary strengths of our framework's components. The SplineCNN module achieved an accuracy of 79.53% and a high specificity of 81.85%, underscoring its effectiveness in extracting localized spatial features. The MGCNet module outperformed SplineCNN in accuracy (82.36%) and sensitivity (87.65%), demonstrating its capability to integrate multi-graph representations for robust classification. The full AGMGC model, combining both modules with contrastive learning, consistently outperformed its individual components across most metrics, indicating a synergistic enhancement in capturing both local and global connectivity patterns. Notably, AGMGC's lower standard deviations in key metrics (e.g., F1-score: ±2.86) compared to base-

lines like GATE (±6.92) and PLSNet (±5.51) suggest improved stability, likely attributable to the adaptive graph learning and dynamic fusion mechanisms.

3.4 Results on REST-meta-MDD Dataset

Table 2 presents the classification performance on the REST-meta-MDD dataset for NC vs. MDD. Conventional methods (SVM: 64.71%, MLP: 62.71%) again underperformed, struggling with the high-dimensional fMRI data. Among GNN baselines, PopulationGCN [18] excelled in sensitivity (89.34%) but faltered in specificity (43.93%), indicating a bias toward MDD detection at the expense of NC classification. GATE [20] achieved a balanced accuracy of 73.79% and precision of 78.30%, while KMGCN [13] showed moderate performance (accuracy: 68.75%) with higher variability across folds.

Our AGMGC framework outperformed all baselines, achieving the highest accuracy (82.36%), specificity (83.64%), F1-score (81.40%), and AUC (82.29%), with competitive precision (82.11%). The ablation results highlight distinct module contributions: SplineCNN delivered a strong F1-score (80.45%) and sensitivity (82.18%), reflecting its focus on local feature extraction, whereas MGCNet improved specificity (75.71%) and overall accuracy (79.59%), leveraging multi-graph integration. The full AGMGC model balanced these strengths, achieving superior specificity and stability (e.g., AUC variance: ±2.93 vs. KMGCN's ±5.98), likely due to the adaptive graph construction and Mixture of Experts fusion mitigating biases seen in population-based methods like PopulationGCN.

4 Conclusion

We developed the Adaptive Graph Learning with Multi-Graph Convolutions (AGMGC) framework for fMRI-based brain disorder diagnosis, addressing the limitations of static connectivity measures and biased inference processes prevalent in population-based graph neural network approaches. By integrating a data-driven Learner Graph module, SplineCNN for localized feature extraction, and a Multi-Graph Convolution Network with dynamic Mixture of Experts fusion, AGMGC achieved superior classification performance on the ABIDE and REST-meta-MDD datasets. On ABIDE, AGMGC attained an accuracy of 84.65% and an F1-score of 85.44%, outperforming state-of-the-art baselines like PopulationGCN and KMGCN. For REST-meta-MDD, it recorded an accuracy of 82.36% and a specificity of 83.64%, demonstrating robust generalization across disorders. Ablation studies confirmed the complementary roles of SplineCNN and MGCNet, with their synergy enhancing both predictive accuracy and stability.

Our work advances fMRI analysis by offering a scalable and generalizable solution that mitigates reliance on predefined graph structures and test-time training data dependencies. Future research could extend AGMGC to dynamic fMRI data to capture temporal connectivity fluctuations and explore multimodal integration with structural imaging for a comprehensive brain network understanding. Additionally, future efforts will focus on developing interpretability analyses,

such as validating ROI importance scores and visualizing learned graphs, to align with established neuroscience findings. These enhancements could further refine its diagnostic potential, positioning AGMGC as a valuable tool for precision neuroscience and clinical applications.

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

References

1. Di Martino, A., Yan, C.-G., Li, Q., Denio, E., Castellanos, F.X., Alaerts, K., Anderson, J.S., Assaf, M., Bookheimer, S.Y., Dapretto, M., et al.: The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism. Mol. Psychiatry 19(6), 659–667 (2014).
2. Yan, C.-G., et al.: Reduced default mode network functional connectivity in patients with recurrent major depressive disorder. Proc. Natl. Acad. Sci. U.S.A. 116(18), 9078–9083 (2019).
3. Buxton, R.B.: Dynamic models of BOLD contrast. NeuroImage 62(1), 953–961 (2012).
4. Bassett, D.S., Bullmore, E.T.: Human brain networks in health and disease. Curr. Opin. Neurol. 22(4), 340–347 (2009).
5. Woodward, N.D., Cascio, C.J.: Resting-state functional connectivity in psychiatric disorders. JAMA Psychiatry 72(8), 743–744 (2015).
6. Wang, J., et al.: Functional MRI in brain research. NeuroImage 231, 117872 (2021).
7. Zhu, X., et al.: Cross-network interaction for diagnosis of major depressive disorder. Brain Imaging Behav. 15(3), 1279–1289 (2021).
8. Qu, G., et al.: Challenges in fMRI connectivity analysis. NeuroImage 235, 118013 (2021).
9. Warren, S.L., Moustafa, A.A.: Graph neural networks in neuroscience. Neurocomputing 512, 161–175 (2023).
10. Bessadok, A., et al.: Graph neural networks for brain imaging. Med. Image Anal. 78, 102427 (2022).
11. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: Proc. 5th Int. Conf. on Learning Representations (ICLR). OpenReview.net, Toulon (2017).
12. Kawahara, J., et al.: BrainNetCNN: Convolutional neural networks for brain connectomes. Med. Image Anal. 36, 1038–1049 (2017).
13. Zeng, X., et al.: Knowledge-driven multi-graph convolutional network. Med. Image Anal. 99, 103368 (2025).
14. Noman, F., et al.: Graph autoencoders for brain network embedding. IEEE J. Biomed. Health Inform. 28(3), 1644–1655 (2024).
15. Chen, T., et al.: A simple framework for contrastive learning. In: Proc. 37th Int. Conf. on Machine Learning (ICML), PMLR, vol. 119, pp. 1597–1607 (2020).
16. Li, X., et al.: Recent advances in graph neural networks for neuroimaging. NeuroImage 285, 120483 (2024).
17. Zhang, H., et al.: Dynamic graph convolutional networks for brain disorder diagnosis. IEEE Trans. Med. Imaging 42(5), 1335–1345 (2023).

18. Parisot, S., Ktena, S.I., Ferrante, E., Lee, M., Guerrero, R., Glocker, B., Rueckert, D.: Disease prediction using graph convolutional networks: Application to autism spectrum disorder and Alzheimer's disease. Med. Image Anal. 48, 117–130 (2018).
19. 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. Med. Image Anal. 74, 102233 (2021).
20. Peng, L., Wang, N., Xu, J., Zhu, X., Li, X.: GATE: Graph CCA for temporal self-supervised learning for label-efficient fMRI analysis. IEEE Trans. Med. Imaging 42(2), 391–402 (2022).
21. Wang, Y., Long, H., Zhou, Q., Bo, T., Zheng, J.: PLSNet: Position-aware graph convolutional network-based autism spectrum disorder diagnosis via functional connectivity learning and ROIs sifting. Comput. Biol. Med. 163, 107184 (2023).
22. Fey, M., Lenssen, J.E., Weichert, F., Müller, H.: SplineCNN: Fast geometric deep learning with continuous B-spline kernels. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), pp. 869–877. IEEE, Salt Lake City (2018).