CROSS-SCALE META-PATH LEARNING FOR HETEROGENEOUS BRAIN NETWORK ANALYSIS

Anonymous authors
Paper under double-blind review

ABSTRACT

The brain is a complex system of interconnected neurons that can be modeled as a brain network. Brain disorders often involve abnormal functional connectivity, making it important to jointly capture local interactions among brain regions and high-order coordination across functional subnetworks. However, existing brain graph learning methods mainly rely on direct pairwise connections or entangled multi-hop propagation, limiting their ability to distinguish heterogeneous connection semantics and explicitly model high-order cross-subnetwork communication. To address these issues, we propose a meta-path-guided cross-scale heterogeneous brain network learning framework. Specifically, we design a low-order heterogeneous communication module that separately models intra-subnetwork and inter-subnetwork connections through relation-specific message passing and subnetwork-aware attention. We further introduce a high-order meta-path communication module that constructs high-order communication structures within a unified functional-subnetwork meta-path space. By adaptively selecting disease-discriminative meta-paths, the module captures high-order cross-subnetwork interactions. In addition, node-level cross-scale contrastive learning is employed to align low-order and high-order representations and promote their complementarity. Experimental results on the ADNI and ABIDE datasets demonstrate that the proposed method outperforms existing approaches and identifies informative brain regions and meta-paths, validating its effectiveness and interpretability for brain disease diagnosis.

1 INTRODUCTION

Brain disorders are often associated with complex and subtle alterations in functional connectivity Yang et al. (2026). Effectively identifying disease-related brain network patterns is therefore essential for facilitating clinical diagnosis and elucidating potential pathological mechanisms Li et al. (2026); Ma et al. (2026b). Resting-state functional magnetic resonance imaging (rs-fMRI) provides a non-invasive means of measuring spontaneous brain activity and enables the construction of functional brain networks based on inter-regional signal correlations, thereby offering a valuable tool for characterizing disease-associated disruptions in functional brain organization Elorette et al. (2024). Given the inherently non-Euclidean structure of functional brain networks, graph neural networks (GNNs) have been widely adopted to learn discriminative representations of neurological disorders from complex network topologies Cui et al. (2022); Zheng et al. (2024b).

Existing brain graph learning methods primarily rely on direct functional connections between brain regions for message passing, thereby emphasizing local pairwise interactions Huang et al. (2024). However, brain function is not merely an aggregation of isolated brain regions and their direct connections. Instead, it relies on coordinated communication across distinct functional subnetworks Wang et al. (2025); Zhang et al. (2024). Disease-related abnormalities may therefore manifest not only as changes in local connection strengths, but also as high-order interaction patterns mediated by intermediate functional subnetworks Zhou et al. (2024); Wen et al. (2024). Although stacking graph convolutional layers or directly performing multi-hop adjacency propagation can enlarge the receptive field, such strategies tend to entangle propagation paths associated with different functional subnetworks, making it difficult to distinguish the communication semantics induced by distinct intermediate subnetworks Xue et al. (2025). Consequently, modeling disease-related high-order

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

communication while preserving explicit functional subnetwork semantics remains a key challenge in brain network learning.

Beyond high-order communication, direct functional connections in brain networks also exhibit substantial relational heterogeneity Chen et al. (2024). According to the functional subnetworks to which brain regions belong, these connections can be categorized as intra-subnetwork and inter-subnetwork relations. They respectively characterize local coordination within the same functional system and direct interactions across distinct functional systems, thereby preserving the node and relation semantics encoded by the functional subnetwork partition Bannadabhavi et al. (2023). Recent studies have begun to incorporate functional community priors or heterogeneous graph structures to model the modular organization of brain networks and the interactions among functional subnetworks Pei et al. (2025). For example, BrainHGL treats functional subnetworks as node types and learns high-order cross-subnetwork representations through meta-path encoding and inter-path interaction Wen et al. (2024). Nevertheless, existing methods do not explicitly disentangle relation-specific low-order communication from high-order communication along functional-subnetwork meta-paths, nor do they exploit their cross-scale correspondences for collaborative representation learning.

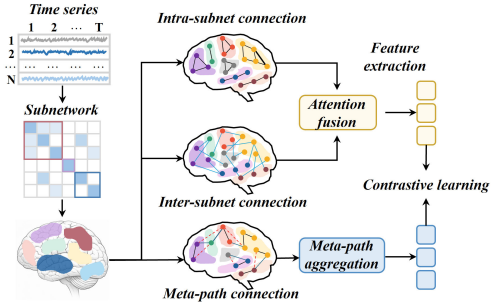

Figure 1: Motivation of functional-subnetwork-aware cross-scale brain network learning.

To address these challenges, we propose a meta-path-guided cross-scale framework for heterogeneous brain network representation learning, as illustrated in Figure 1. First, we develop a high-order meta-path communication (HMC) module, which constructs high-order communication patterns mediated by intermediate functional subnetworks within a unified functional-subnetwork meta-path space. A global competitive selection mechanism is further employed to estimate the task-specific importance of individual meta-paths, thereby yielding high-order communication structures with explicit subnetwork semantics. Then, we introduce a low-order heterogeneous communication (LHC) module, which categorizes direct functional connections into intra-subnetwork and inter-subnetwork relations. Relation-specific message passing is used to separately extract local connectivity information from the two relation types, followed by a subnetwork-aware relational attention mechanism for adaptive fusion. Finally, we incorporate node-level cross-scale contrastive learning Xu et al. (2024) to promote semantic consistency between representations at different topological scales, and fuse their graph-level representations for disease diagnosis. The main contributions of this work are summarized as follows:

- We propose a meta-path-guided cross-scale framework for heterogeneous brain network representation learning, which jointly captures local functional interactions and high-order coordination across functional subnetworks.
- We develop the HMC and LHC modules. HMC constructs and selects discriminative high-order communication structures in a unified subnetwork meta-path space, while LHC models intra-subnetwork and inter-subnetwork functional connections through relation-specific propagation and subnetwork-aware attention.
- We exploit node correspondences across different topological scales to perform node-level cross-scale contrastive learning. Diagnostic and ablation experiments on the ADNI and ABIDE datasets demonstrate the effectiveness of the proposed method.

2 Related Work

2.1 Heterogeneous Brain Network

Heterogeneous brain network learning aims to exploit the diverse semantics of brain regions and functional connections to enhance representation learning for brain disease diagnosis. Most existing brain graph methods treat brain regions as homogeneous nodes and process all functional

connections uniformly, making it difficult to adequately characterize the modular organization and relational heterogeneity of brain networks Makhlouf et al. (2025); Shi et al. (2025). To address this limitation, several studies have incorporated anatomical or functional priors. For example, ComBrainTF rearranges functional connectivity matrices according to predefined functional communities to learn intra-community and inter-community connectivity patterns Bannadabhavi et al. (2023). BrainGNN learns soft assignment scores for brain regions to construct task-relevant latent functional communities Li et al. (2021). M2CDCA combines spectral clustering with structural network priors to facilitate community detection. Another line of work directly models brain networks as heterogeneous graphs by defining node types according to anatomical regions or functional subnetworks Zhu et al. (2022). For example, HyperDiag constructs region-specific functional hypergraphs to capture high-order dependencies within individual regions and topological interactions across regions Ma et al. (2026b). Nevertheless, existing methods mainly focus on node types, community assignments, or region-level structures, while rarely explicitly distinguishing different relation types among direct functional connections.1

2.2 High-order Brain Network Learning

High-order brain network learning has shown considerable potential for brain disease diagnosis by moving beyond conventional pairwise functional connectivity to characterize complex dependencies among multiple brain regions. For example, DHGNN dynamically updates hypergraph structures to capture evolving high-order relationships Jiang et al. (2019). I^2HBN jointly models intra-subject and inter-subject high-order correlations for ASD diagnosis Han et al. (2024). HGFM employs large-scale self-supervised pretraining to learn transferable high-order representations across brain disease datasets Han et al. (2025). Beyond hypergraph-based methods, hierarchical modeling and long-range dependency learning have also been explored to characterize the multiscale organization of brain networks. For example, BrainHGT integrates long-short range attention with prior-guided community clustering to jointly model region-level, region-to-community, and inter-community interactions Ma et al. (2026a). However, hypergraphs generally represent brain regions as unordered sets and thus cannot explicitly encode ordered communication semantics across functional subnetworks. Conventional multi-hop propagation may also entangle information conveyed along distinct paths. In contrast, meta-paths explicitly characterize cross-subnetwork communication patterns through sequences of functional subnetwork types Shi et al. (2023).

3 Method

The proposed framework is illustrated in Figure 2, which jointly models low-order heterogeneous connectivity and high-order meta-path communication via cross-scale contrastive learning for brain disease diagnosis.

3.1 High-Order Meta-Path Communication

For each subject, the mean rs-fMRI time series was extracted from each ROI. Functional connectivity between every pair of ROIs was then estimated using Pearson correlation, resulting in a weighted 90×90 functional connectivity matrix. Each ROI was treated as a graph node, and the resulting functional connectivity matrix served as the adjacency matrix of the brain network. Existing brain graph methods primarily rely on direct functional connections between brain regions, limiting their ability to capture high-order cross-subnetwork coordination mediated by intermediate functional subnetworks. To address this limitation, we introduce a high-order meta-path communication module that constructs semantically explicit high-order communication patterns within a unified functional-subnetwork meta-path space and models high-order interactions among brain regions.

Given the functional subnetwork label s_i of brain region i, we first extract inter-subnetwork functional connections from the functional connectivity matrix A:

$$A_{ij}^{\text{inter}} = A_{ij} \cdot I(s_i \neq s_j) \tag{1}$$

where $I(\cdot)$ denotes the indicator function. Based on these connections, we construct a unified functional-subnetwork meta-path space. Each candidate meta-path $\mathcal{P}_m = (a, b, c)$ represents information propagation from a source functional subnetwork a to a target functional subnetwork c

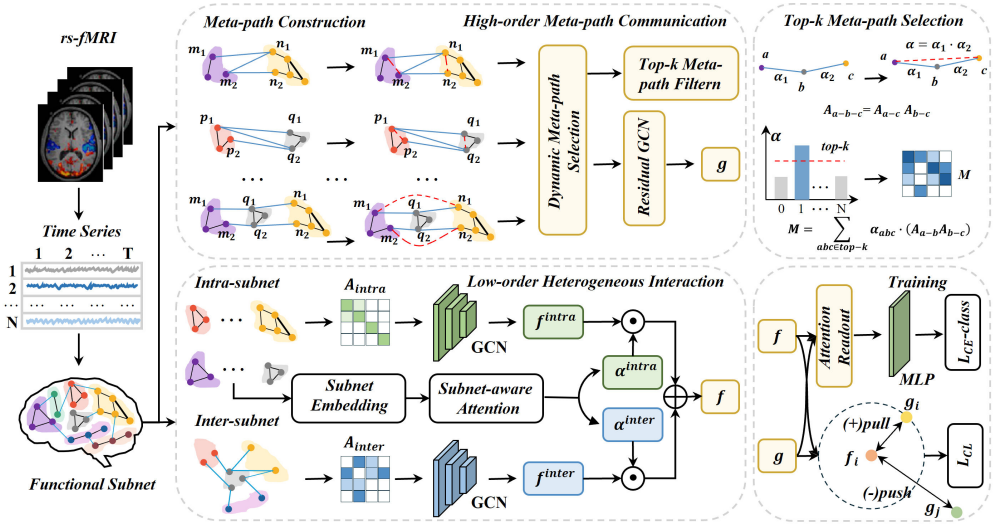

Figure 2: Overview of the proposed framework. The LHC module models intra-subnetwork and inter-subnetwork direct connections through relation-specific message passing, while the HMC module constructs and selects disease-discriminative functional-subnetwork meta-paths to capture high-order communication. Cross-scale contrastive learning aligns the low-order and high-order node representations, which are subsequently aggregated and fused for brain disease diagnosis.

through a subnetwork b. Since each step corresponds to cross-subnetwork propagation, adjacent subnetwork types are required to be different, i.e., $a \neq b$ and $b \neq c$. Accordingly, the candidate meta-path space is defined as:

$$\mathcal{P} = \{(a, b, c) \in \mathcal{S}^3 \mid a \neq b,\ b \neq c\} \tag{2}$$

where, $\mathcal{S} = \{1, \ldots, 7\}$ denotes the set of functional subnetwork types. We include all cross-subnetwork multi-hop communication patterns in a unified candidate space, enabling the model to compare different functional-subnetwork communication sequences within a shared semantic space rather than relying on a small set of manually predefined meta-paths.

To map subnetwork-level meta-paths to high-order communication structures at the brain-region level, let $D_a \in R^{N \times N}$ denote the diagonal indicator matrix of functional subnetwork a. Specifically, $(D_a)_{ii} = 1$ if brain region i belongs to functional subnetwork a, and 0 otherwise. For a candidate meta-path $\mathcal{P}_m = (a, b, c)$, its high-order adjacency matrix is defined as:

$$A_{\mathcal{P}_m} = D_a A^{\text{inter}} D_b A^{\text{inter}} D_c \tag{3}$$

where $D_a A^{\text{inter}} D_b$ extracts the connections from the source functional subnetwork a to the intermediate functional subnetwork b, while $D_b A^{\text{inter}} D_c$ describes the connections from the intermediate functional subnetwork b to the target functional subnetwork c.

Not all candidate meta-paths are relevant to disease classification, and directly aggregating them may introduce redundant communication and task-irrelevant noise. Therefore, we adopt a global competitive meta-path selection mechanism to adaptively identify disease-discriminative functional-subnetwork communication paths. Let M_p denote the number of meta-paths and ω_m the learnable weight of the m-th meta-path. We rank all candidate meta-paths by their weights and retain the top-K ones. The normalized importance of the m-th selected meta-path is defined as:

$$\mathcal{P}_K = \{\mathcal{P}_m \mid m \in \text{TopK}\left(\{\omega_1, \omega_2, \ldots, \omega_{M_p}\}, K\right)\} \tag{4}$$

$$\beta_m = \frac{\exp(\omega_m)}{\sum_{\mathcal{P}_n \in \mathcal{P}_K} \exp(\omega_n)} \tag{5}$$

Based on the selected meta-paths, the high-order communication matrix is computed as:

$$M = \sum_{\mathcal{P}_m \in \mathcal{P}_K} \beta_m A_{\mathcal{P}_m} \tag{6}$$

where $A_{\mathcal{P}_m}$ denotes the high-order adjacency matrix induced by meta-path $\mathcal{P}_m$, encoding long-range communication across functional subnetworks. We normalize M to obtain $\widetilde{M}$ for stable propagation. The high-order representations are then updated using residual graph convolution:

$$H_g^{(l)} = \sigma\left(H_g^{(l-1)} + \widetilde{M} H_g^{(l-1)} W_g^{(l)}\right) \tag{7}$$

where $H_g^{(l-1)}$ denotes the high-order representation at layer $l-1$, and $W_g^{(l)}$ is a learnable weight matrix. After multiple layers, the resulting representation H_g captures high-order communication across functional subnetworks and is jointly optimized with the low-order heterogeneous representation through cross-scale contrastive learning.

3.2 Low-Order Heterogeneous Communication

To model direct functional connections and their relational heterogeneity, we represent the brain network as a functional-subnetwork-aware heterogeneous graph. Each brain region is treated as a node whose type is determined by its functional subnetwork. Accordingly, the nodes are divided into seven types. Let s_i denote the functional subnetwork label of brain region i, which specifies the type of node v_i. Based on whether the connected regions belong to the same functional subnetwork, functional connections are divided into two types:

$$r \in \mathcal{R} = \{r_{\text{intra}}, r_{\text{inter}}\} \tag{8}$$

where r_{intra} denotes an intra-subnetwork connection with $s_i = s_j$, and r_{inter} denotes an inter-subnetwork connection with $s_i \neq s_j$. Conventional homogeneous graph neural networks treat all nodes and edges uniformly, failing to capture node and relation heterogeneity induced by functional subnetworks. We therefore adopt relation-specific message passing to separately model intra- and inter-subnetwork functional connections. At layer l, the relation-specific message received by node i under relation r is defined as:

$$m_{i,r}^{(l)} = \frac{1}{|\mathcal{N}_i^r|} \sum_{j \in \mathcal{N}_i^r} f_r^{(l-1)}(A_{ij}) \odot W_r^{(l-1)} h_j^{(l-1)} \tag{9}$$

where $\mathcal{N}_i^r$ denotes the neighbors of node i under relation r, $h_j^{(l-1)}$ is the representation of node j at layer $l-1$, and $W_r^{(l-1)}$ is a learnable relation-specific transformation. A_{ij} denotes the functional connectivity strength between nodes i and j, $f_r^{(l-1)}$ is a weighting function, and $\odot$ denotes element-wise multiplication. This formulation separately aggregates intra- and inter-subnetwork messages according to their relation types.

To enhance sensitivity to functional-subnetwork semantics, we introduce a subnetwork-aware relation attention mechanism. For node i and relation type r, the attention score is computed as:

$$e_{i,r}^{(l)} = a_r^\top \sigma\left(W_h m_{i,r}^{(l)} + W_s q_{s_i}\right) \tag{10}$$

where q_{s_i} is the learnable embedding of the functional subnetwork containing node i, W_h and W_s are linear transformations for the message representation and subnetwork embedding, respectively, $a_r^\top$ is the attention parameter for relation r, and $\sigma(\cdot)$ is a nonlinear activation function. The relation importance for node i is then normalized as:

$$\alpha_{i,r}^{(l)} = \frac{\exp\left(e_{i,r}^{(l)}\right)}{\sum_{r' \in \mathcal{R}_i} \exp\left(e_{i,r'}^{(l)}\right)} \tag{11}$$

Finally, messages from different relation types are fused by attention, while a residual connection preserves the node representation:

$$h_i^{(l)} = \sigma\left[\frac{1}{C(i)} \sum_{r \in \mathcal{R}_i} \alpha_{i,r}^{(l)} m_{i,r}^{(l)} + W_o^{(l-1)} h_i^{(l-1)}\right] \tag{12}$$

where $C(i)$ denotes the number of relation types associated with node i, and $W_o^{(l-1)}$ is the residual transformation matrix. Through this process, the LHC module jointly models intra-subnetwork and inter-subnetwork direct communication, yielding a functional-subnetwork-aware low-order representation H_f.

Table 1: Disease prediction performance of different methods on ADNI dataset (Mean/Std) (%)

Method	NC vs. EMCI			NC vs. LMCI			NC vs. AD			NC vs. EMCI vs. LMCI vs. AD		
	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC
STIGR	70.46/07.51	69.65/09.42	69.69/10.26	74.72/06.65	68.18/10.58	73.02/09.52	75.81/04.67	70.46/08.19	71.75/07.33	48.42/07.56	46.37/08.16	67.28/06.37
STCAL	69.81/07.50	72.78/08.22	66.43/08.57	72.45/05.56	67.79/07.98	71.92/07.93	74.22/07.01	67.37/13.22	74.63/09.92	45.98/07.03	42.56/08.14	64.22/06.77
HyperDiag	70.13/08.04	71.23/06.96	69.49/10.55	75.95/05.57	71.69/06.36	74.73/09.74	75.80/07.16	71.01/10.56	75.07/10.28	58.68/08.73	56.57/09.88	79.63/05.98
BrainHGT	71.10/02.62	72.07/04.22	70.92/05.39	76.33/07.58	70.66/09.12	75.73/07.98	76.21/06.28	72.17/08.24	76.90/08.62	57.21/06.63	54.80/07.83	78.55/06.87
FORGE	69.84/05.36	70.73/07.55	68.13/05.77	73.98/05.37	68.01/09.38	69.28/08.75	75.00/07.11	69.56/10.98	73.62/07.21	49.55/05.56	48.45/06.22	68.33/03.57
HeCo	66.01/06.02	68.70/08.55	64.35/09.73	71.67/07.55	67.65/08.17	64.19/09.44	73.40/07.38	68.06/07.53	76.36/06.65	44.87/05.19	42.53/05.77	64.07/05.28
MS-GNN	68.55/06.18	67.06/11.74	68.37/07.56	72.03/04.72	68.39/06.20	73.39/08.68	74.16/06.37	68.63/11.07	71.86/08.56	47.26/05.10	43.82/07.04	66.77/04.92
Async-ST	68.90/05.43	69.40/05.60	69.32/07.73	72.40/05.41	65.93/09.17	70.30/10.02	73.80/06.20	69.87/07.07	74.41/05.13	44.09/04.94	42.51/05.43	65.35/05.63
BrainHGL	70.18/07.48	69.99/10.43	71.01/09.21	75.17/07.47	69.23/12.93	73.07/10.81	75.78/06.02	70.42/08.22	73.84/08.96	46.15/04.99	45.13/06.00	64.52/06.62
DHGFormer	70.75/05.05	71.94/06.06	67.81/09.21	74.82/07.18	69.54/10.86	72.68/08.35	75.38/05.86	70.87/07.70	72.11/10.45	47.09/06.60	45.33/07.73	67.45/06.49
Ours	**73.66/04.86**	**74.67/04.53**	**73.56/06.42**	**80.29/07.94**	**76.50/10.33**	**80.59/10.20**	**82.23/06.56**	**79.58/07.33**	**79.52/10.28**	**64.12/08.38**	**62.40/10.17**	**83.00/07.76**

Table 2: Disease prediction performance of different methods on ABIDE dataset (Mean/Std) (%)

Method	NC vs. ASD		
	ACC	F1	AUC
STIGR	75.71/09.83	73.83/10.20	78.61/10.78
STCAL	73.33/06.61	69.98/08.64	79.12/09.02
HyperDiag	75.56/05.12	72.69/05.41	81.57/05.78
BrainHGT	74.92/10.10	70.45/12.83	80.93/08.99
FORGE	72.06/05.94	67.88/08.43	79.55/07.84
HeCo	70.00/08.04	64.95/09.37	78.84/08.81
MS-GNN	71.74/09.59	69.15/11.64	78.62/09.06
Async-ST	72.06/10.21	66.95/15.11	77.55/14.04
BrainHGL	72.53/03.33	69.59/03.36	79.56/05.55
DHGFormer	73.01/06.26	68.96/08.93	77.32/06.45
Ours	**79.99/03.49**	**76.35/05.06**	**86.15/03.57**

3.3 Cross-Scale Contrastive Learning

Low-order heterogeneous communication and high-order meta-path communication capture direct heterogeneous connections and cross-subnetwork long-range coordination, respectively. To promote their complementary interaction, we introduce cross-scale contrastive learning. Let H_f and H_g denote the low- and high-order node representations produced by LHC and HMC, respectively. Two projection heads map them into a shared contrastive space, yielding Z_f and Z_g. Cross-scale representations of the same brain region within a subject form positive pairs, while all others form negative pairs. The cross-scale contrastive loss is defined as:

$$\ell_{i,n} = -\log \frac{\exp\left(\cos\left(z_{f,i,n}, z_{g,i,n}\right)/\tau\right)}{\sum_{j=1}^{B}\sum_{m=1}^{N}\exp\left(\cos\left(z_{f,i,n}, z_{g,j,m}\right)/\tau\right)} \tag{13}$$

$$\mathcal{L}_{CL} = \frac{1}{BN}\sum_{i=1}^{B}\sum_{n=1}^{N}\ell_{i,n} \tag{14}$$

where B denotes the batch size, N the number of brain regions, and $z_{f,i,n}$ and $z_{g,i,n}$ the low- and high-order projected representations of the n-th brain region in the i-th subject, respectively. τ is the temperature parameter. The two node representations are then aggregated by attention to obtain graph-level representations, which are concatenated and fed into a classifier for disease diagnosis. The overall training objective is defined as:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda\mathcal{L}_{CL} \tag{15}$$

where $\mathcal{L}_{CE}$ is the cross-entropy classification loss and λ controls the contribution of the contrastive loss. This framework jointly models low-order short-range heterogeneous communication and high-order long-range meta-path communication, while cross-scale contrastive learning promotes complementary representations of direct local connectivity and high-order multi-hop coordination.

Table 3: Ablation study on ADNI dataset (Mean/Std) (%)

Method	NC vs. EMCI			NC vs. LMCI			NC vs. AD			NC vs. EMCI vs. LMCI vs. AD		
	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC
w/o LHC	67.24/07.68	67.14/08.44	67.40/07.17	73.25/09.42	69.96/08.04	74.34/10.24	76.18/06.15	70.93/07.01	76.34/07.77	57.76/06.92	54.88/08.30	80.08/05.87
w/o HMC	70.50/05.54	70.25/08.17	70.07/09.07	76.80/07.57	73.72/07.87	73.71/09.17	77.81/07.00	73.81/07.71	75.62/09.45	60.91/09.58	58.76/10.62	82.10/07.80
w/o CL	71.40/06.41	70.90/07.65	68.17/08.75	78.67/06.32	73.87/08.55	79.62/07.08	79.01/06.44	73.69/08.91	77.23/07.93	61.92/06.99	60.34/07.48	82.57/04.94
Ours	**73.66/04.86**	**74.67/04.53**	**73.56/06.42**	**80.29/07.94**	**76.50/10.33**	**80.59/10.20**	**82.23/06.56**	**79.58/07.33**	**79.52/10.28**	**64.12/08.38**	**62.40/10.17**	**83.00/07.76**

Table 4: Ablation study on ABIDE dataset (Mean/Std) (%)

Method	NC vs. ASD		
	ACC	F1	AUC
w/o LHC	73.01/06.95	71.92/07.07	78.64/07.22
w/o HMC	76.50/06.68	73.37/07.16	83.84/07.15
w/o CL	77.77/06.30	75.77/07.28	83.97/05.92
Ours	**79.99/03.49**	**76.35/05.06**	**86.15/03.57**

4 EXPERIMENT

4.1 DATA ACQUISITION AND PROCESSING

We evaluate the proposed method on two publicly available resting-state functional magnetic resonance imaging (rs-fMRI) datasets covering both Alzheimer's disease and autism spectrum disorder. ADNI: This dataset comprises 533 subjects, including 142 normal controls (NC), 173 early mild cognitive impairment (EMCI) subjects, 112 late mild cognitive impairment (LMCI) subjects, and 106 Alzheimer's disease (AD) subjects. ABIDE: This dataset comprises 630 subjects, including 337 normal controls (NC) and 293 autism spectrum disorder (ASD) subjects.

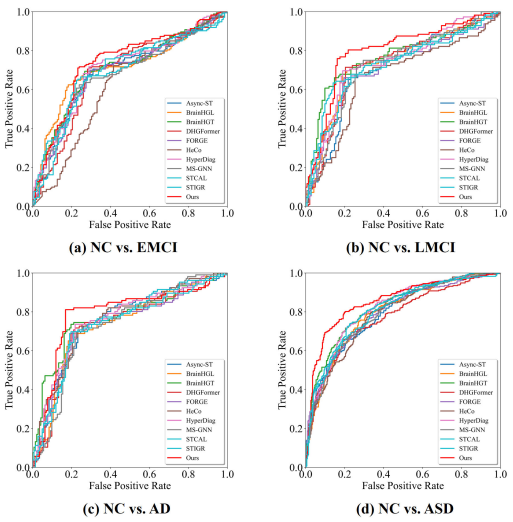

Figure 3: ROC curves of all methods under different classification tasks.

We employed the DPARSF toolbox based on SPM8 to perform identical preprocessing procedures on both fMRI datasets. Slice timing correction, spatial normalization to the standard EPI template, linear detrending, and nuisance regression were performed to reduce the effects of head motion, cerebrospinal fluid, and white matter signals. Finally, the preprocessed fMRI data were parcellated into 90 regions of interest (ROIs) using the Automated Anatomical Labeling (AAL) atlas, and each ROI was assigned to one of seven functional subnetworks.

4.2 EXPERIMENTAL SETTINGS

The proposed method is implemented in PyTorch using an NVIDIA GeForce RTX 3090 GPU. We use the Adam optimizer for end-to-end optimization, with an initial learning rate of 5×10^{-3}, a weight decay of 5×10^{-4}, and a batch size of 32. The maximum number of training epochs is set to 300, and early stopping is adopted to prevent overfitting. The contrastive loss weight λ is selected from $\{0.001, 0.01, 0.05, 0.1, 0.3, 0.5, 1.0\}$. We evaluate the model using 10-fold cross-validation and report the mean and standard deviation over the 10 folds. On the ADNI dataset, we conduct three binary classification tasks and one four-class classification task: NC vs. EMCI, NC vs. LMCI, NC vs. AD, and NC vs. EMCI vs. LMCI vs. AD. On the ABIDE dataset, we further

evaluate the proposed method on the NC vs. ASD binary classification task. For all tasks, we use accuracy (ACC), F1-score (F1), and the area under the receiver operating characteristic curve (AUC) as evaluation metrics.

4.3 Comparison with Existing Methods

We compare our method with ten competing methods, including STIGR Liu et al. (2025), STCAL Liu et al. (2024), HyperDiag Ma et al. (2026b), BrainHGT Ma et al. (2026a), FORGE Chen & Yu (2026), HeCo Wang et al. (2021), MS-GNN Hwang et al. (2024), Async-ST Zhang et al. (2025), BrainHGL Xia et al. (2024), and DHGFormer Xue et al. (2025). We conduct a comprehensive comparison with recent state-of-the-art methods on two brain disease datasets. As shown in Table 1, our model achieves the highest accuracy on the ADNI dataset. Specifically, it outperforms the second-best method by 2.56%, 3.96%, 6.02%, and 5.44% on the NC vs. EMCI, NC vs. LMCI, NC vs. AD, and NC vs. EMCI vs. LMCI vs. AD tasks, respectively. As shown in Table 2, our model also achieves the highest accuracy of 79.99% on the ABIDE dataset, exceeding previous methods by at least 4.28%. Moreover, the ROC curves in Figure 3 consistently lie above those of the competing methods, demonstrating strong performance in terms of AUC. The consistent improvements in ACC, F1-score, and AUC indicate that the proposed model learns more discriminative brain network representations and generalizes effectively across different brain disease diagnosis tasks.

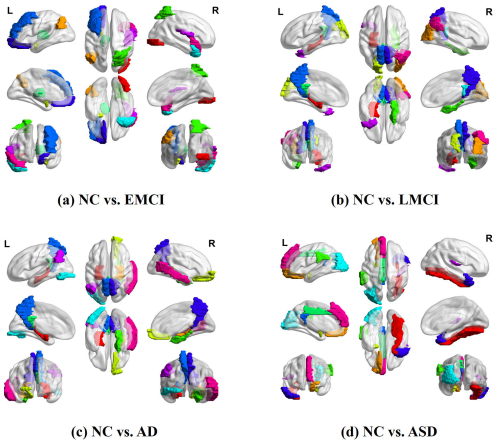

Figure 4: Most discriminative brain regions across different diagnostic tasks.

4.4 Ablation Studies

To evaluate the effectiveness of each core component, we conduct ablation studies on the ADNI and ABIDE datasets, with the results reported in Tables 3 and 4. Specifically, w/o LHC replaces the low-order heterogeneous communication module with a homogeneous graph convolutional network, w/o HMC removes the high-order meta-path communication module, and w/o CL removes the cross-scale contrastive learning objective while retaining only the classification loss.

The results show that replacing LHC causes the largest performance degradation, demonstrating the importance of explicitly modeling heterogeneous intra- and inter-subnetwork connections for learning discriminative local brain network representations. Removing HMC also consistently reduces performance, indicating that meta-path-based multi-hop propagation effectively captures high-order long-range interactions across functional subnetworks. Moreover, removing cross-scale contrastive learning leads to inferior results, confirming the benefit of aligning low- and high-order representations to capture complementary information across scales. The full model achieves the best overall performance on both datasets, validating the effectiveness and complementarity of the three components.

4.5 Discriminative Brain Regions

To interpret the model predictions, we rank the ROIs by their average importance scores and visualize the top ten regions for each task in Figure 4. For NC vs. EMCI, the selected regions are broadly distributed across the frontal, parietal, temporal, visual, limbic, and subcortical systems, including the angular gyrus, amygdala, superior parietal lobule, thalamus, and temporal pole Li et al. (2026; 2025a). For NC vs. LMCI, the most discriminative regions are mainly located in the hippocampus, parahippocampal gyrus, bilateral posterior cingulate gyri, bilateral precuneus, angular gyrus, and occipital cortex Sperling et al. (2011). A similar but more pronounced pattern is observed for NC vs.

AD, where the bilateral hippocampi and several default-mode-related regions, including the posterior cingulate gyrus, precuneus, and angular gyrus, receive high importance scores Li et al. (2025b). For NC vs. ASD, the identified regions span sensorimotor, cingulo-limbic, auditory, visual, and association areas, including the postcentral gyrus, amygdala, supramarginal gyrus, cingulate cortex, Heschl's gyrus, and medial superior frontal gyrus Zheng et al. (2024a). Overall, the model identifies both disease-specific and shared discriminative patterns, providing anatomically interpretable evidence for its predictions Qiu et al. (2024).

4.6 Discriminative Meta-Path Analysis

To interpret the high-order long-range meta-path communication module, we visualize the most discriminative meta-paths selected by our model across different diagnostic tasks, as shown in Figure 5. For NC vs. EMCI, the selected meta-paths mainly involve communications among temporal and high-order association regions. These results suggest that abnormal interactions among sensory, attention, and association networks may already exist at the early stage of cognitive impairment. For NC vs. LMCI, a similar pattern is observed, where DMN-mediated communication pathways play an important role, indicating that large-scale functional integration across brain networks is progressively disrupted with cognitive decline. For NC vs. AD, the model identifies multiple DMN-centered communication patterns. These pathways involve key regions such as the angular gyrus, temporal association cortex, and parietal regions, suggesting that AD-related abnormalities are mainly characterized by impaired long-range coordination between sensory processing regions and high-order cognitive networks. For NC vs. ASD, the discriminative meta-paths mainly capture interactions among the frontoparietal network (FPN), default mode network (DMN), sensorimotor network (SMN), and attention network (ATN). The corresponding important regions include the angular gyrus, precuneus, inferior parietal lobule, and temporal regions, which are consistent with previous findings that ASD is associated with abnormal large-scale functional network integration. Overall, the selected meta-paths demonstrate that our model can identify disease-related high-order cross-subnetwork communication patterns rather than relying solely on isolated functional connections.

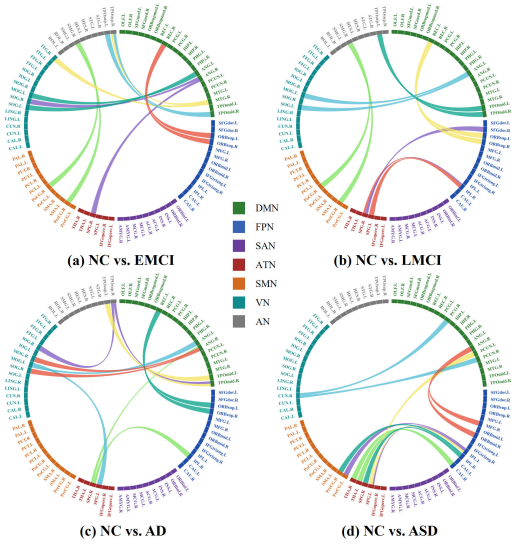

Figure 5: Most discriminative brain regions across different diagnostic tasks.

5 Conclusion

In this paper, we propose a meta-path-guided cross-scale heterogeneous brain network learning framework for brain disease classification. The framework jointly models direct heterogeneous functional connections and high-order communication across functional subnetworks through the LHC and HMC modules. Furthermore, node-level cross-scale contrastive learning is introduced to enhance the coordination of low-order and high-order representations. Extensive experiments on the ADNI and ABIDE datasets demonstrate that the proposed method achieves superior performance compared with existing approaches. It also identifies disease-relevant brain regions and discriminative meta-paths, validating its effectiveness and interpretability for brain disease diagnosis.

AI use statement

In this work, we used generative AI tools for manuscript language polishing, formula typesetting and figure caption refinement. We have not used generative AI tools for core model design, experimental design and result analysis, and data collection are not applicable to this work. Addition-

ally, we used generative AI tools for preliminary literature-search assistance. We have reviewed all AI-assisted work. We manually cross-checked AI-polished text against our original arguments, and all experiment-related descriptions are verified by two authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

All experiments in this paper are based on publicly released ADNI and ABIDE datasets. The original data collection has completed corresponding ethical review and IRB approval in prior publications. We do not collect new human-subject data in this work. All analyses follow the original dataset usage agreement. No privacy-releasing or discriminatory processing is performed on subject data.

REPRODUCIBILITY STATEMENT

To support reproducibility, anonymous source code will be released as supplementary material upon submission. All baseline implementations and evaluation metrics follow descriptions in existing published papers.

REFERENCES

Anushree Bannadabhavi, Soojin Lee, Wenlong Deng, Rex Ying, and Xiaoxiao Li. Community-aware transformer for autism prediction in fmri connectome. In *International conference on medical image computing and computer-assisted intervention*, pp. 287–297. Springer, 2023.

Dongdong Chen, Mengjun Liu, Zhenrong Shen, Linlin Yao, Xiangyu Zhao, Zhiyun Song, Haolei Yuan, Qian Wang, and Lichi Zhang. Exploring multiconnectivity and subdivision functions of brain network via heterogeneous graph network for cognitive disorder identification. *IEEE Transactions on Neural Networks and Learning Systems*, 36(7):12400–12414, 2024.

Qianyu Chen and Shujian Yu. Continual learning for fmri-based brain disorder diagnosis via functional connectivity matrices generative replay. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25099–25109, 2026.

Hejie Cui, Wei Dai, Yanqiao Zhu, Xuan Kan, Antonio Aodong Chen Gu, Joshua Lukemire, Liang Zhan, Lifang He, Ying Guo, and Carl Yang. Braingb: a benchmark for brain network analysis with graph neural networks. *IEEE transactions on medical imaging*, 42(2):493–506, 2022.

Catherine Elorette, Atsushi Fujimoto, Frederic M Stoll, Satoka H Fujimoto, Niranjana Bienkowska, Liza London, Lazar Fleysher, Brian E Russ, and Peter H Rudebeck. The neural basis of resting-state fmri functional connectivity in fronto-limbic circuits revealed by chemogenetic manipulation. *Nature communications*, 15(1):4669, 2024.

Xiangmin Han, Rundong Xue, Shaoyi Du, and Yue Gao. Inter-intra high-order brain network for asd diagnosis via functional mris. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 216–226. Springer, 2024.

Xiangmin Han, Rundong Xue, Jingxi Feng, Yifan Feng, Shaoyi Du, Jun Shi, and Yue Gao. Hypergraph foundation model for brain disease diagnosis. *IEEE Transactions on Neural Networks and Learning Systems*, 36(10):17702–17716, 2025. doi: 10.1109/TNNLS.2025.3554755.

Jinghan Huang, Nanguang Chen, and Anqi Qiu. Topological cycle graph attention network for brain functional connectivity. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 723–732. Springer, 2024.

Yechan Hwang, Soojin Hwang, Guorong Wu, and Won Hwa Kim. Multi-order simplex-based graph neural network for brain network analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 532–541. Springer, 2024.

Jianwen Jiang, Yuxuan Wei, Yifan Feng, Jingxuan Cao, and Yue Gao. Dynamic hypergraph neural networks. In *Ijcai*, pp. 2635–2641, 2019.

Shengrong Li, Qi Zhu, Chunwei Tian, Wei Shao, and Daoqiang Zhang. Interpretable dynamic brain network analysis with functional and structural priors. *IEEE Transactions on Medical Imaging*, 44(12):4878–4889, 2025a. doi: 10.1109/TMI.2025.3584231.

Shengrong Li, Qi Zhu, Chunwei Tian, Li Zhang, Bo Shen, Chuhang Zheng, Daoqiang Zhang, and Wei Shao. Spatio-temporal evolutionary graph learning for brain network analysis using medical imaging. *IEEE Transactions on Image Processing*, 34:5860–5872, 2025b. doi: 10.1109/TIP.2025.3607633.

Shengrong Li, Qi Zhu, Chunwei Tian, Xinyang Zhang, Wei Shao, Jie Wen, and Daoqiang Zhang. Neuroh-tgl: Neuro-heterogeneity guided temporal graph learning strategy for brain disease diagnosis. *Advances in Neural Information Processing Systems*, 38:109349–109371, 2026.

Xiaoxiao Li, Yuan Zhou, Nicha Dvornek, Muhan Zhang, Siyuan Gao, Juntang Zhuang, Dustin Scheinost, Lawrence H Staib, Pamela Ventola, and James S Duncan. Braingnn: Interpretable brain graph neural network for fmri analysis. *Medical image analysis*, 74:102233, 2021.

Rui Liu, Zhi-An Huang, Yao Hu, Zexuan Zhu, Ka-Chun Wong, and Kay Chen Tan. Spatial–temporal co-attention learning for diagnosis of mental disorders from resting-state fmri data. *IEEE Transactions on Neural Networks and Learning Systems*, 35(8):10591–10605, 2024. doi: 10.1109/TNNLS.2023.3243000.

Rui Liu, Yao Hu, Jibin Wu, Ka-Chun Wong, Zhi-An Huang, Yu-An Huang, and Kay Chen Tan. Dynamic graph representation learning for spatio-temporal neuroimaging analysis. *IEEE Transactions on Cybernetics*, 55(3):1121–1134, 2025. doi: 10.1109/TCYB.2025.3531657.

Jiajun Ma, Yongchao Zhang, Chao Zhang, Zhao Lv, and Shengbing Pei. Brainhgt: A hierarchical graph transformer for interpretable brain network analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 17617–17625, 2026a.

Yulan Ma, Fangkun Li, Wenchao Yang, Qian Si, Chenglong Yu, and Yang Li. Hyperdiag: Temporal–regional hypergraph learning via topology-enhanced state propagation for brain disease diagnosis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 7874–7882, 2026b.

Ahmed T Makhlouf, William Drew, Jacob L Stubbs, Joseph J Taylor, Donato Liloia, Jordan Grafman, David Silbersweig, Michael D Fox, and Shan H Siddiqi. Heterogeneous patterns of brain atrophy in schizophrenia localize to a common brain network. *Nature Mental Health*, 3(1):19–30, 2025.

Shengbing Pei, Jiajun Ma, Zhao Lv, Chao Zhang, and Jihong Guan. Community-aware graph transformer for brain disorder identification. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pp. 4191–4199, 2025.

Xinmei Qiu, Fan Wang, Yongheng Sun, Chunfeng Lian, and Jianhua Ma. Towards graph neural networks with domain-generalizable explainability for fmri-based brain disorder diagnosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 454–464. Springer, 2024.

Gen Shi, Yifan Zhu, Jian K Liu, and Xuesong Li. Hegcl: Advance self-supervised learning in heterogeneous graph-level representation. *IEEE Transactions on Neural Networks and Learning Systems*, 35(10):13914–13925, 2023.

Gen Shi, Yifan Zhu, Wenjin Liu, Quanming Yao, and Xuesong Li. Heterogeneous graph-based multimodal brain network learning. *IEEE Transactions on Knowledge and Data Engineering*, 37 (8):4664–4676, 2025. doi: 10.1109/TKDE.2025.3569648.

Reisa A Sperling, Paul S Aisen, Laurel A Beckett, David A Bennett, Suzanne Craft, Anne M Fagan, Takeshi Iwatsubo, Clifford R Jack Jr, Jeffrey Kaye, Thomas J Montine, et al. Toward defining the preclinical stages of alzheimer's disease: recommendations from the national institute on aging-alzheimer's association workgroups on diagnostic guidelines for alzheimer's disease. *Alzheimer's & dementia*, 7(3):280–292, 2011.

Song Wang, Zhenyu Lei, Zhen Tan, Jiaqi Ding, Xinyu Zhao, Yushun Dong, Guorong Wu, Tianlong Chen, Chen Chen, Aiying Zhang, et al. Brainmap: Learning multiple activation pathways in brain networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 14432–14440, 2025.

Xiao Wang, Nian Liu, Hui Han, and Chuan Shi. Self-supervised heterogeneous graph neural network with co-contrastive learning. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 1726–1736, 2021.

Guangqi Wen, Peng Cao, Lingwen Liu, Maochun Hao, Siyu Liu, Junjie Zheng, Jinzhu Yang, Osmar R Zaiane, and Fei Wang. Heterogeneous graph representation learning framework for resting-state functional connectivity analysis. *IEEE Transactions on Medical Imaging*, 44(3):1581–1595, 2024.

Zhengwang Xia, Huan Wang, Tao Zhou, Zhuqing Jiao, and Jianfeng Lu. Customized relationship graph neural network for brain disorder identification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 109–118. Springer, 2024.

Jiaxing Xu, Qingtian Bian, Xinhang Li, Aihu Zhang, Yiping Ke, Miao Qiao, Wei Zhang, Wei Khang Jeremy Sim, and Balázs Gulyás. Contrastive graph pooling for explainable classification of brain networks. *IEEE Transactions on Medical Imaging*, 43(9):3292–3305, 2024.

Rundong Xue, Hao Hu, Zeyu Zhang, Xiangmin Han, Juan Wang, Yue Gao, and Shaoyi Du. Dhgformer: Dynamic hierarchical graph transformer for disorder brain disease diagnosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 268–278. Springer, 2025.

Meimei Yang, Yongheng Sun, Qianqian Wang, Andrea Bozoki, Maureen Kohi, and Mingxia Liu. Hyperbolic kernel graph neural networks for neurocognitive decline analysis from multimodal brain imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(4):4827–4841, 2026. doi: 10.1109/TPAMI.2025.3646188.

Jianjia Zhang, Yunan Guo, Luping Zhou, Lei Wang, Weiwen Wu, and Dinggang Shen. Constructing hierarchical attentive functional brain networks for early ad diagnosis. *Medical Image Analysis*, 94:103137, 2024.

Jianjia Zhang, Xiaotong Wu, Xiang Tang, Luping Zhou, Lei Wang, Weiwen Wu, and Dinggang Shen. Asynchronous functional brain network construction with spatiotemporal transformer for mci classification. *IEEE Transactions on Medical Imaging*, 44(3):1168–1180, 2025. doi: 10.1109/TMI.2024.3486086.

Kaizhong Zheng, Shujian Yu, Liangjun Chen, Lujuan Dang, and Badong Chen. Bpi-gnn: Interpretable brain network-based psychiatric diagnosis and subtyping. *NeuroImage*, 292:120594, 2024a. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2024.120594. URL https://www.sciencedirect.com/science/article/pii/S1053811924000892.

Kaizhong Zheng, Shujian Yu, Baojuan Li, Robert Jenssen, and Badong Chen. Brainib: Interpretable brain network-based psychiatric diagnosis with graph information bottleneck. *IEEE Transactions on Neural Networks and Learning Systems*, 36(7):13066–13079, 2024b.

Jie Zhou, Biao Jie, Zhengdong Wang, Zhixiang Zhang, Tongchun Du, Weixin Bian, Yang Yang, and Jun Jia. Lcgnet: Local sequential feature coupling global representation learning for functional connectivity network analysis with fmri. *IEEE Transactions on Medical Imaging*, 43(12):4319–4330, 2024.

Qi Zhu, Jing Yang, Shuihua Wang, Daoqiang Zhang, and Zheng Zhang. Multi-modal non-euclidean brain network analysis with community detection and convolutional autoencoder. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 7(2):436–446, 2022.