HI-BRAIN: BRAIN-INSPIRED BIDIRECTIONAL HIERARCHICAL LEARNING FOR DYNAMIC BRAIN CONNECTIVITY ANALYSIS

Anonymous authors
Paper under double-blind review

ABSTRACT

Modeling the brain's dynamic information-processing mechanisms is a core challenge in neuroscience and artificial intelligence. Existing methods typically first construct dynamic functional brain connectivity (DFBC), and further use topological modeling or prior constraints to characterize such dynamic information transmission. However, they fail to comprehensively model the hierarchical organization of the brain, leading to learned features with limited biological plausibility. To address this issue, we propose a brain-inspired bidirectional hierarchical network (Hi-Brain) for DFBC analysis. Hi-Brain effectively simulates bottom-up functional aggregation and top-down functional regulation in the brain. Specifically, we first segment fMRI signals into multiple subsequences to model the brain's dynamic patterns. For each subsequence, we design a neuro-inspired node activation learner to capture the adaptive activation of neural nodes under different states. Building on this, we further introduce a high-order subnetwork learner to describe the functional differentiation and global collaboration of the sub-networks, enabling bottom-up high-order functional aggregation. Finally, the aggregated information is used to guide feature optimization in lower-level neural organizations, thereby simulating top-down functional regulation in the brain. Experiments on the ADNI and Parkinson's disease datasets demonstrate that Hi-Brain not only achieves superior predictive performance, but also effectively explains the neural mechanisms underlying its predictions. More importantly, Hi-Brain provides a new perspective for developing brain-inspired machine learning models.

1 INTRODUCTION

The brain dynamically integrates neural information to support diverse cognitive functions (Wei et al., 2026). To characterize this dynamic pattern, researchers commonly use functional magnetic resonance imaging (fMRI) to construct dynamic functional brain connectivity (DFBC) (Kim et al., 2021). DFBC models the dynamic information interactions among neural nodes, serving as an important bridge between neural activity and complex cognition (Li et al., 2025a). Therefore, analyzing DFBC is of great significance for both neuroscience research and clinical applications.

Recently, the most popular way for analyzing DFBC is to model it as a sequence of dynamic brain graphs, and use graph neural networks to sequentially extract spatio-temporal features from each graph (Liu et al., 2025; Zhu et al., 2024). For example, STAGIN (Kim et al., 2021) first uses multiple graph isomorphism networks to capture the topological information, and then employs a transformer to model temporal dependencies across brain graphs. Although this data-driven strategy captures the time-varying structure of DFBC, it ignores valuable prior knowledge in the brain. Therefore, several prior-driven methods have been proposed (Chen et al., 2023; Wang et al., 2023; Shehzad et al., 2025). These methods typically incorporate structural or functional priors of brain connectivity into model design, and impose biological constraints on feature learning (Gao et al., 2024; Li et al., 2025b). For example, ST2GCN (Li et al., 2025a) incorporates the modular property of the brain by introducing functional subnetwork constraints, aiming to encourage feature similarity within subnetworks and dissimilarity across subnetworks. However, these methods mainly adopt data-driven or prior-constrained modeling paradigms, lacking network architectures guided by the brain's

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.

intrinsic information-processing mechanisms. As a result, the learned neural representations fail to accurately reflect brain activity patterns.

In fact, brain information processing relies on bottom-up functional integration and top-down functional regulation (Katsuki & Constantinidis, 2012; Barrett & Miller, 2026). As shown in Fig. 1, during neural information processing, the brain first activates task-relevant neural nodes (Wang et al., 2025). These nodes then interact with each other and differentiate into subnetworks responsible for specific functions (Li et al., 2025b). In particular, complex functions are not independently determined by isolated pairs of brain regions, but rely on coordinated interactions among multiple brain regions. Different functional subnetworks further achieve information collaboration through high-order associations, thereby integrating local information into more comprehensive whole-brain representations (Huang et al., 2025). Meanwhile, whole-brain representation is fed back to guide lower-level neural activity, which enables adaptive optimization of information processing (Chen et al., 2025; Ju & Bassett, 2020). The optimized information is used for final decision-making. Clearly, designing network architectures based on this bidirectional regulatory mechanism can lead to more biologically plausible representations.

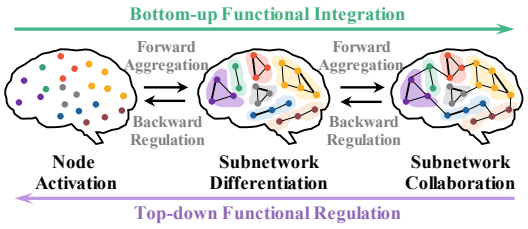

Figure 1: Bidirectional functional aggregation and regulation mechanisms of the brain.

To this end, we propose a brain-inspired bidirectional hierarchical network (Hi-Brain) for DFBC analysis, aiming to jointly characterize the dynamic aggregation and hierarchical regulation of neural information. Specifically, we first apply a sliding window to segment fMRI signals into multiple subsequences to capture brain dynamics. For each subsequence, we design a neuro-inspired node activation learner (N2AL) to model the adaptive activation of neural nodes under different states. Building on this, we devise a high-order subnetwork learner (HOSL) to simulate the complex differentiation and cooperation among multiple subnetworks, enabling bottom-up functional aggregation in the brain. Subsequently, the aggregated functional information is further used to optimize information extraction in lower-level neural organizations, which corresponds to top-down functional regulation (TDFR). Finally, we integrate whole-brain representations across different windows for downstream tasks. Our contributions are summarized as follows:

- To emulate the intrinsic information processing mechanism of the brain, we propose Hi-Brain to capture cross-level information interactions at the node, subnetwork, and whole-brain levels, thereby simulating bottom-up functional aggregation in the brain.
- We design a TDFR module to mimic the top-down regulation of brain function, in which whole-brain information guides lower-level neural feature learning.
- Experiments show that Hi-Brain not only achieves strong predictive performance, but also reveals the neural mechanisms underlying its predictions. In addition, it provides the machine learning community with a more biologically plausible computational framework.

2 Related Work

Brain Connectivity Analysis. Brain connectivity analysis provides an effective modeling framework for characterizing information interactions among brain regions (Zhan et al., 2026). Existing analysis methods typically extract structural features directly from whole-brain connectivity. For example, BNT (Kan et al., 2022b) employs a transformer to learn discriminative brain embeddings. OT-STGCN (Zhu et al., 2024) introduces a hubness propagation model to capture high-order evolution between adjacent DFBCs. In addition, several methods focus on modeling the modular organization of the brain. For instance, LSGNN (Chen et al., 2023) maps brain into multiple subnetworks using a divisible graph neural network and extracts local features from each subnetwork. CAGT (Pei et al., 2025) introduces a community-aware transformer to model intra- and inter-subnetwork interactions. However, brain information processing is not merely a single-level topological modeling process. It involves the collaboration of multi-level neural organizations and is regulated by feedback from whole-brain information (Katsuki & Constantinidis, 2012; Barrett & Miller, 2026).

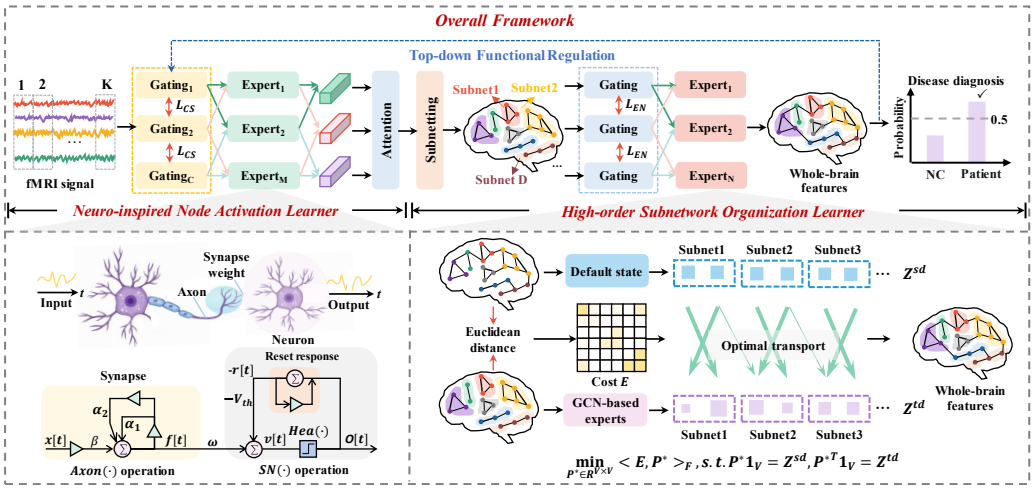

Figure 2: The proposed brain-inspired bidirectional hierarchical network for DFBC analysis.

Inspired by this mechanism, we propose Hi-Brain, which aims to unify bottom-up functional aggregation and top-down functional regulation within a single framework.

Mixture of Experts. A mixture-of-experts (MoE) system consists of a set of neural network-based experts and a learnable gating mechanism (Berisha et al., 2025). Owing to its strong feature extraction capability, MoE has been widely used in brain network analysis. For example, BrainMoE (Zhang et al., 2025a) designs an MoE framework to learn the structural connectivity of different nodes, and introduces a disease-specific MoE to improve disease classification performance. BrainMAP (Wang et al., 2025) proposes a sequential model to identify the activation pathways in brain connections, and then combines a MoE-based aggregation module to capture parallel information across pathways. Existing MoE methods can learn discriminative brain features, but they lack a clear correspondence with the intrinsic mechanisms of the brain. To address this limitation, this paper does not simply introduce MoE as a generic module. Instead, it redesigns MoE in a brain-inspired manner at both the node and subnetwork levels. This design effectively simulates node activation under different functional states, as well as the differentiation and collaboration of subnetworks.

3 Proposed Method

As shown in Fig. 2, we propose a brain-inspired bidirectional hierarchical network. It couples node activation, subnetwork organization, and whole-brain regulation into a closed-loop learning process, thereby jointly modeling bottom-up functional aggregation and top-down functional regulation.

3.1 Neuro-inspired Node Activation Learner

In this work, the fMRI signal is represented as $X \in \mathbb{R}^{V \times L}$, where V denotes the number of nodes and L is the signal length. To represent the dynamic patterns of the brain, we divide X into K sliding windows to generate a series of subsequences $\{X_k\}_{k=1}^{K}$. In the bottom-up functional aggregation pathway, neural nodes are the fundamental units of information aggregation, and their activation states reflect selective responses to input information. Although fMRI does not directly measure neuronal spikes, its BOLD signals exhibit temporally smooth and history-dependent dynamics. This property requires modeling information dependencies across successive observations, which can be naturally captured by the recurrent state integration of spiking neural networks (SNNs). Therefore, we use an SNN as the node activation learner. Specifically, we introduce an SNN based on an infinite impulse response (IIR) filter (Fang et al., 2020), which leverages a recurrent architecture to enable nodes to adaptively respond by integrating current inputs and historical states. The membrane potential $v[t]$ is computed by aggregating temporally filtered synaptic inputs, while spike generation follows a threshold-triggering mechanism with threshold V_{th}:

$$v[t] = -V_{th} r[t] + \sum_{i=1}^{S} \omega_i f_i[t] \tag{1}$$

$$r[t] = e^{-\frac{1}{\tau_r}} r[t-1] + O[t-1] \tag{2}$$

$$O[t] = Hea\left(v[t] - V_{th}\right) \tag{3}$$

where t denotes time, $r[t]$ is the reset filter, S represents the number of synapses, ω_i denotes the synaptic weight, $f_i[t]$ is the second-order IIR filter response corresponding to the ith synapse, τ_r controls the decay speed of the reset impulse, and $Hea(\cdot)$ represents the Heaviside step function. Next, the temporal response of each input channel is modeled using a second-order IIR filter:

$$f_i[t] = \alpha_1 f_i[t-1] + \alpha_2 f_i[t-2] + \beta x_i[t-1] \tag{4}$$

where $x_i[t-1]$ denotes the input signal of the ith channel, and α_1, α_2, and β are predefined filter coefficients controlling temporal dynamics, whose detailed definitions are provided in Appendix A. From a computational perspective, the model can be decomposed into two fundamental operations. $Axon(\cdot)$ represents the propagation of spikes to the next neuron, which can be modeled using the second-order exponential IIR filter in Eq. (4). $SN(\cdot)$ denotes the neuronal process of updating membrane potentials and generating spike outputs, as described in Eq. (1)-Eq. (3). Based on this formulation, each node adaptively encodes its activation through recurrent temporal dynamics. Therefore, the node-level expert network can be formulated as: $SNN(\cdot) = SN\left(Axon(\cdot)\right)$.

In fact, the brain dynamically switches among multiple functional states, exhibiting distinct regional activation patterns (Barttfeld et al., 2015). To capture this property, we design a MoE with C gating networks and M SNNs. Each gating network models the selective modulation of neural nodes under different states. Given a subsequence X_k, the cth gating network produces expert weights via:

$$P_{k,m}^{c} = \frac{\exp\left(\mathrm{Gate}_{k,m}^{c}(X_k)\right)}{\sum_{m'=1}^{M} \exp\left(\mathrm{Gate}_{k,m'}^{c}(X_k)\right)} \tag{5}$$

where $P_{k,m}^{c}$ is the contribution of the mth expert under the cth gating network. To encourage the diversity of node activation across states, we introduce a similarity constraint:

$$\mathcal{L}_{CS}^{k} = \frac{2}{C(C-1)} \sum\nolimits_{1 \le c_1 < c_2 \le C} \mathrm{sim}\left(P_k^{c_1}, P_k^{c_2}\right) \tag{6}$$

where $sim(\cdot)$ denotes the cosine similarity. Then, SNN-based expert networks are used to adaptively learn node activation features: $H_{k,m} = SNN_m(X_k)$. The output of each expert is weighted by its corresponding contribution to obtain the node activation pattern under the given state: $H_k^c = \sum_{m=1}^{M} P_{k,m}^{c} H_{k,m}$. On this basis, we apply an attention mechanism to fuse the C activation features, obtaining more expressive node representations $\bar{H}_k$.

3.2 High-order Subnetwork Organization Learner

Activated nodes are organized into multiple subnetworks that encode functional specialization and collaboration. Moreover, brain information processing relies on coordinated interactions among multiple regions (Momi et al., 2025). To this end, we design a high-order subnetwork organization learner (HOSL). Specifically, given the node features $\bar{H}_k$, we first construct brain networks A_k using Pearson correlation. Then, we incorporate functional priors and partition each A_k into D subnetworks $\{A_k^d \in \mathbb{R}^{V_d \times V_d}\}_{d=1}^{D}$ based on a predefined template (Long et al., 2019), where V_d is the number of nodes in the dth subnetwork. To model subnetwork differentiation, we introduce an entropy-regularized gating mechanism that learns expert assignment weights. For subnetwork A_k^d, the expert weights are computed as:

$$P_{k,n}^{d} = \frac{\exp\left(\mathrm{Gate}_{k,n}^{d}\left(A_k^d\right)\right)}{\sum_{n'=1}^{N} \exp\left(\mathrm{Gate}_{k,n'}^{d}\left(A_k^d\right)\right)} \tag{7}$$

where $P_{k,n}^{d}$ denotes the contribution of the nth expert to A_k^d, and N is the number of subnetwork-level experts. Then, we impose an entropy regularization on the expert weights: $\mathcal{L}_{EN}^{k} = -\sum_{d=1}^{D}\sum_{n=1}^{N} P_{k,n}^{d} \log\left(P_{k,n}^{d}\right)$. The entropy constraint encourages sparse expert assignments, promoting subnetwork-level functional differentiation. Here, we adopt GCN-based experts to learn functional representations of subnetworks: $Z_{k,n}^{d} = \sigma\left(A_k^d \bar{H}_k^d W_{k,n}^{d}\right)$, where $\bar{H}_k^d$ is the activated

node representation of the dth subnetwork, and $W_{k,n}^{d}$ is the learnable weight matrix. The expert outputs are aggregated using the learned weights, yielding the subnetwork-level representation Z_k^d. Subsequently, all subnetwork representations Z_k^d are concatenated along the node dimension, and their feature dimension is reduced to $V \times 1$ using a node-wise readout function. Finally, we apply a global softmax over all nodes to obtain the normalized distribution Z_k^{td}.

To model high-order functional aggregation across subnetworks, we introduce optimal transport. In optimal transport, the transport relations among all nodes are jointly optimized under global constraints, thereby enabling the characterization of high-order dependencies in the brain. Let $Z_k^{td} \in \mathbb{R}^{V\times 1}$ denote the target distribution obtained by concatenating normalized subnetwork features, and let $Z_k^{sd} = [1/V, 1/V, \cdots, 1/V] \in \mathbb{R}^{V\times 1}$ denote a uniform source distribution. Z_k^{sd} reflects that the brain does not exhibit functional preferences at initialization. High-order subnetwork collaboration can be formulated as follows:

$$\min_{P_k^* \in \mathbb{R}^{V\times V}} \langle E_k, P_k^* \rangle_F, \quad \text{s.t.} \quad P_k^* \mathbf{1}_V = Z_k^{sd}, \quad (P_k^*)^T \mathbf{1}_V = Z_k^{td} \tag{8}$$

where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product, E_k is the cost matrix defined by the pairwise Euclidean distances between node features, P_k^* denotes the optimal transport plan, and $\mathbf{1}_V$ is an all-ones vector. The problem is solved via the Sinkhorn algorithm (Cuturi, 2013). The resulting P_k^* integrates information from node- and subnetwork-level processing, yielding a whole-brain representation in a bottom-up manner.

3.3 Top-Down Functional Regulation

The brain uses the aggregated information to regulate lower-level neural organizations. Therefore, we further model top-down functional regulation (TDFR) by using global information to guide node activation. Specifically, we project P_k^* to match the dimension of X_k via a linear layer $\phi(\cdot)$, and inject it into the node-level gating mechanism defined in Eq.(5):

$$P_{k,m}^c = \frac{\exp\left(Gate_{k,m}^c\left(X_k + \phi(P_k^*)\right)\right)}{\sum_{m'=1}^{M} \exp\left(Gate_{k,m'}^c\left(X_k + \phi(P_k^*)\right)\right)} \tag{9}$$

We stack multiple Hi-Brain layers to model hierarchical brain organization. Let $X_k^{(0)} = X_k$ denote the original fMRI subsequence of the kth window. The lth layer Hi-Brain layer is formulated as:

$$F_k^{(l)} = \text{N2AL}^{(l)}\left(X_k^{(l-1)}\right), \quad P_k^{*(l)} = \text{HOSL}^{(l)}\left(F_k^{(l)}\right), \quad X_k^{(l)} = F_k^{(l)} + \phi^{(l)}\left(P_k^{*(l)}\right) \tag{10}$$

where $F_k^{(l)}$ is the node-level activation representation learned by N2AL at the lth layer, and $P_k^{*(l)}$ denotes the whole-brain collaboration representation obtained by HOSL. The regulated representation $X_k^{(l)}$ is then used as the input to the next layer. In this way, each layer first performs bottom-up functional aggregation through N2AL and HOSL, and then uses TDFR to feed the aggregated whole-brain information back to lower-level representations. Finally, $P_k^{(l_h)}$ from all windows is concatenated for classification, where l_h denotes the number of stacked Hi-Brain layers. The overall training objective is: $\mathcal{L} = \mathcal{L}_{CE} + \lambda_1 \mathcal{L}_{CS} + \lambda_2 \mathcal{L}_{EN}$, where $\mathcal{L}_{CE}$ denotes the cross-entropy loss, and $\mathcal{L}_{CS}$ and $\mathcal{L}_{EN}$ represent the similarity constraint and entropy constraint aggregated over all windows, respectively, λ_1 and λ_2 are used to control the relative importance of different losses.

4 Experiments

In this section, we aim to address the following research questions (RQs): **RQ1:** How does Hi-Brain perform compared with state-of-the-art methods? **RQ2:** What is the contribution of each module in Hi-Brain to the overall prediction performance? **RQ3:** Can Hi-Brain effectively explain the mechanisms underlying its prediction results? **RQ4:** Does Hi-Brain realize the brain's bidirectional hierarchical mechanism?

4.1 Experimental Setup

Data Description and Preprocessing. We conduct experiments on the publicly available ADNI dataset and a Parkinson's disease (PD) dataset collected from the Affiliated Hospital of Nanjing

Table 1: Classification performance of different methods on the ADNI and PD datasets (%).

Dataset	Method	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC
		NC vs MCI vs AD			NC vs MCI			NC vs AD			MCI vs AD		
ADNI	FBNetGen	57.50	56.97	67.75	61.25	66.46	58.47	63.33	58.50	57.53	65.83	73.71	64.07
	BNT	60.21	40.94	61.96	69.75	79.65	57.58	74.00	71.72	75.89	74.50	81.56	57.64
	STAGIN	56.46	27.75	54.76	69.25	80.97	56.44	67.08	47.99	55.49	74.44	85.35	53.71
	ACIFBN	62.35	57.88	70.42	71.25	78.96	64.49	79.17	71.53	78.44	80.54	87.13	75.82
	OT-STGCN	59.02	41.31	58.52	76.08	80.14	60.89	73.75	57.61	66.41	76.26	85.40	68.15
	LSGNN	58.75	36.96	58.20	60.27	68.62	54.04	70.42	53.07	66.26	64.05	73.72	56.11
	CAGT	60.75	36.25	59.29	75.75	83.81	62.68	80.83	76.90	83.29	75.75	83.45	67.40
	BrainMAP	57.65	50.53	63.08	70.39	77.75	60.82	76.08	67.37	71.64	79.46	86.29	69.52
	ST2GCN	60.39	54.44	63.91	75.49	83.67	65.65	78.10	73.34	76.43	82.97	88.66	77.10
	Hi-Brain	**67.60**	**61.26**	**74.93**	**79.00**	**85.61**	**72.11**	**85.00**	**81.08**	**84.93**	**85.28**	**90.62**	**77.52**
	Method	NC vs TPD vs PPD			NC vs TPD			NC vs PPD			TPD vs PPD		
PD	FBNetGen	62.50	57.65	69.23	74.56	52.80	65.65	74.55	78.39	51.82	72.00	73.05	66.14
	BNT	63.75	56.47	71.75	82.50	70.88	74.83	79.00	77.18	78.62	76.73	79.50	70.20
	STAGIN	55.63	43.52	63.71	81.89	74.20	76.55	82.20	80.89	79.30	78.00	63.00	59.85
	ACIFBN	61.25	48.72	66.33	82.89	77.91	79.80	76.97	73.90	68.59	74.18	73.36	63.79
	OT-STGCN	59.38	38.24	62.37	81.56	79.31	81.59	82.20	83.59	75.50	75.82	73.39	63.05
	LSGNN	56.88	46.44	64.46	73.75	74.84	77.08	74.17	76.93	68.47	72.00	76.61	58.62
	CAGT	59.38	57.06	66.25	81.25	75.88	80.42	78.33	80.06	73.99	80.00	85.43	60.00
	BrainMAP	57.24	53.04	64.75	82.50	77.85	81.63	80.90	81.79	81.57	77.50	81.40	76.45
	ST2GCN	60.63	53.74	65.30	86.25	83.57	77.75	83.75	85.30	81.62	81.00	82.69	80.82
	Hi-Brain	**66.14**	**59.21**	**74.09**	**92.50**	**90.03**	**87.54**	**86.36**	**87.04**	**84.08**	**85.75**	**85.47**	**82.35**

Medical University. The ADNI dataset includes 510 subjects: 140 normal controls (NC), 268 mild cognitive impairment (MCI), and 102 Alzheimer's disease (AD). The PD dataset includes 162 subjects: 54 NC, 44 tremor-dominant Parkinson's disease (TPD), and 64 postural instability and gait disorder Parkinson's disease (PPD). We use the DPARSF (Yan & Zang, 2010) toolbox for fMRI preprocessing, including realignment to the EPI template and detrending to reduce head motion and physiological noise. Then, we use the AAL atlas to define nodes and extract the fMRI signal of each node. For the ADNI dataset, V = 90 and L = 140. For the PD dataset, V = 90 and L = 220.

Metrics. We evaluate the performance of Hi-Brain on brain disease diagnosis tasks. For the ADNI dataset, the classification tasks include: NC vs MCI vs AD, NC vs MCI, NC vs AD, and MCI vs AD. For the PD dataset, the classification tasks include: NC vs TPD vs PPD, NC vs TPD, NC vs PPD, and TPD vs PPD. We use 10-fold cross-validation to evaluate the performance of all methods. The evaluation metrics include accuracy (ACC), F1-score (F1), and AUC.

Implementation Details. Experiments are trained on an NVIDIA GeForce RTX 4090 GPU. The training details are provided in Appendix B. The number of windows K is varied within {2, 3, 4, 5, 6, 7, 8}. Following prior work (Barttfeld et al., 2015), the number of states is set to C=7. The number of node-level experts M and subnetwork-level experts N are both varied within {1, 2, 3, 4, 5, 6}. The brain is partitioned into D=8 subnetworks (Long et al., 2019): default mode network (DMN), fronto-parietal network (FPN), salience network (SAN), attention network (ATN), somatosensory and motor network (SMN), visual network (VN), auditory network (AN), and other networks (ON). λ_1 and λ_2 are selected from {1e-6, 1e-5, 1e-4, 1e-3, 1e-2, 1e-1}. All hyperparameters are determined via grid search. Our code has been uploaded to the supplementary materials.

4.2 Comparison Experiment (RQ1)

Comparison Methods. To evaluate the performance of Hi-Brain, we compare it with nine advanced brain network analysis methods. Data-driven methods include FBNetGen (Kan et al., 2022a), BNT (Kan et al., 2022b), STAGIN (Kim et al., 2021), ACIFBN (Zhang et al., 2024a), and OT-STGCN (Zhu et al., 2024), while prior-driven methods include LSGNN (Chen et al., 2023), CAGT (Pei et al., 2025), BrainMAP (Wang et al., 2025), and ST2GCN (Li et al., 2025a).

Classification Performance. Table 1 presents the diagnostic performance of different methods on the ADNI and PD datasets. The standard deviations are reported in Appendix C. The proposed method consistently achieves the best performance. On the ADNI dataset, Hi-Brain improves the accuracy by 5.25%, 2.92%, 4.17%, and 2.31% across the four tasks, respectively. On the PD dataset,

Hi-Brain achieves accuracy improvements of 2.39%, 6.25%, 2.61%, and 4.75% on the four tasks, respectively. The superior performance of Hi-Brain can be attributed to its ability to simulate both bottom-up information aggregation and top-down functional regulation in the brain, thereby learning more biologically plausible representations. Data-driven methods such as BNT, STAGIN, and OT-STGCN learn topological features of brain networks, but they overlook the valuable neurological priors. Prior-driven methods such as LSGNN, BrainMAP, and ST2GCN can learn biologically meaningful representations, but they typically focus on shallow functional priors. In contrast, the proposed method is inspired by the brain's bidirectional hierarchical mechanisms, enabling unified modeling of cross-level information flow and regulation. In addition, to evaluate the robustness of Hi-Brain, we conduct several parameter sensitivity analyses in Appendix D.

4.3 Ablation Experiment (RQ2)

To validate the effectiveness of each module, we conduct extensive ablation experiments with three variants: (1) w/o N2AL, where SNN is replaced with GRU. (2) w/o HOSL, where attention mechanism is used to learn inter-subnetwork associations. (3) w/o TDFR, where TDFR is removed. As shown in Table 2, removing any module leads to a decline in performance. For example, in the NC vs MCI task, removing the three modules results in performance drops of 6.50%, 4.50%, and 3.75%, respectively. Replacing N2AL with GRU leads to performance degradation because SNNs are beneficial for capturing the temporal dynamics, thereby enabling more effective node activation modeling. HOSL outperforms attention by jointly optimizing node interactions under global constraints, facilitating subnetwork differentiation and high-order functional collaboration. Removing TDFR leads to performance degradation. This suggests that bidirectional regulation better reflects brain functional organization than unidirectional aggregation. Therefore, these modules are all effective and jointly simulate the bidirectional hierarchical regulation mechanism in the brain. Appendix E presents ablation studies on $\mathcal{L}_{EN}$ and $\mathcal{L}_{CS}$, demonstrating that these two losses promote the functional specificity and collaboration of subnetworks, respectively.

Table 2: Ablation results of the proposed method on the ADNI and PD datasets (%).

Dataset	Method	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC
		NC vs MCI vs AD			NC vs MCI			NC vs AD			MCI vs AD		
ADNI	w/o N2AL	59.58	45.85	65.84	72.50	81.33	62.87	79.58	73.47	79.52	77.75	84.87	73.84
	w/o HOSL	61.00	44.62	62.81	74.50	83.32	65.43	81.67	79.40	83.39	79.17	86.13	72.45
	w/o TDFR	62.08	54.43	66.86	75.25	83.58	66.91	82.92	77.99	84.67	80.00	85.47	76.34
	Hi-Brain	**67.60**	**61.26**	**74.93**	**79.00**	**85.61**	**72.11**	**85.00**	**81.08**	**84.93**	**85.28**	**90.62**	**77.52**
	Method	NC vs TPD vs PPD			NC vs TPD			NC vs PPD			TPD vs PPD		
PD	w/o N2AL	56.25	51.06	62.27	86.25	77.88	78.63	77.08	73.15	69.40	76.25	79.87	70.67
	w/o HOSL	59.38	56.03	67.02	87.50	83.50	85.50	77.50	81.22	75.42	80.00	85.82	63.95
	w/o TDFR	60.63	56.39	67.19	88.75	79.57	82.04	80.42	82.11	75.75	81.25	85.10	69.67
	Hi-Brain	**66.14**	**59.21**	**74.09**	**92.50**	**90.03**	**87.54**	**86.36**	**87.04**	**84.08**	**85.75**	**85.47**	**82.35**

4.4 Interpretability Analysis (RQ3)

To reveal the predictive mechanism of Hi-Brain, we identify the brain regions most important for prediction. Specifically, we treat the bidirectionally regulated node activation as the contribution weight of each brain region, and then employ a t-test to identify key regions in different tasks. The results are shown in Fig. 3(a). For the NC vs MCI task, the key brain regions include PCG.L, PHG.R, ITG.L, and MTG.R. These regions are mainly involved in memory and semantic processing, which are commonly impaired in MCI (Billette et al., 2022; Zhang et al., 2024b). For the NC vs AD task, the important brain regions include PCUN.R, HIP.L, AMYG.L, and PreCG.R. This is because AD is often accompanied by emotional dysregulation and memory impairment, which may cause significant alterations in the brain regions responsible for these functions (Tang et al., 2023). For the NC vs TPD task, THA.L, PreCG.R, SMA.R, and CAU.L are the key brain regions. These regions are major components of the thalamo-cortical motor circuit, which serves as the basis for tremor generation and modulation (Mendonça et al., 2024). Therefore, excitation or inhibition in these regions reflects the motor disturbances caused by TPD. For the NC vs PPD task, the key brain regions include SMA.L, PreCG.L, MOG.L, and LING.R. These regions are mainly involved in motor execution and visual perception, and their abnormalities reflect the neural functional alterations

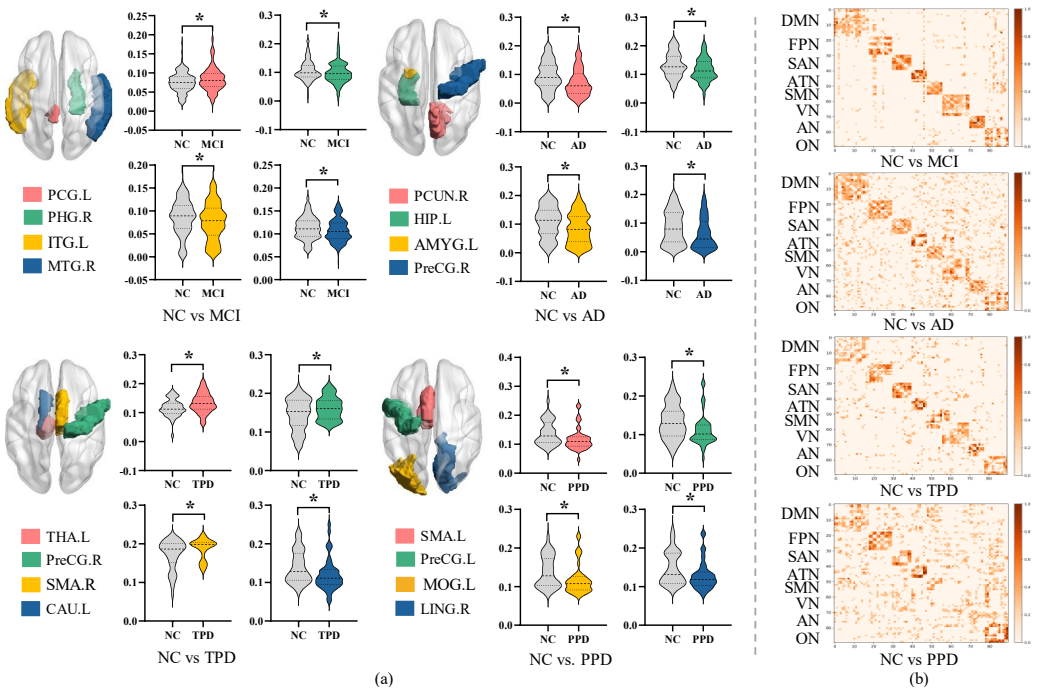

Figure 3: (a) Distribution of important brain regions. * indicates p-value < 0.05. (b) Visualization of functional differentiation and collaboration among subnetworks across different tasks.

in PPD patients with impaired gait regulation Wu et al. (2026). Overall, the key brain regions identified by Hi-Brain are closely associated with brain disorders. In addition, we further explore the functional differentiation and collaboration among different subnetworks. The learned whole-brain feature matrices P^* for different prediction tasks are shown in Fig. 3(b). First, we observe that each brain region interacts with multiple other regions, indicating that Hi-Brain can capture high-order interactions among brain regions. In addition, the elements in the diagonal blocks exhibit significant aggregation, forming a modular structure. This indicates that Hi-Brain can effectively simulate the functional differentiation of subnetworks. Meanwhile, substantial connections are observed in the off-diagonal blocks, suggesting that the method captures functional collaboration among different subnetworks. For NC vs MCI, the connectivity strengths of both DMN and FPN decrease. This may be because MCI patients exhibit impairments in memory processing and cognitive control, leading to abnormal interactions among the corresponding subnetworks (Zhang et al., 2024b). For NC vs AD, the connections between ATN and VN show a clear reduction. This may be because AD impairs patients' attention control and causes visual dysfunction (Li et al., 2025a). Meanwhile, we also observe that compensatory mechanisms in the brain are activated, with AD patients enhancing inter-subnetwork collaboration to compensate for functional deficits (Corriveau-Lecavalier et al., 2026). For NC vs TPD, SMN and DMN are the most affected subnetworks. These subnetworks mainly govern motor control and self-regulation, and abnormal motor regulation is a typical characteristic of TPD patients (Zhang et al., 2025b). Therefore, SMN and DMN can be regarded as key subnetworks for distinguishing NC from TPD. For NC vs PPD, the important subnetworks are SMN and VN. This is because postural and gait control in PPD depends on visuomotor integration, and the abnormalities in SMN and VN jointly reflect the underlying neural impairments in PPD patients (Nishida et al., 2024). Clearly, Hi-Brain not only achieves superior prediction performance, but also explains the neural mechanisms underlying its predictions.

4.5 Brain-inspired Hierarchical Mechanism Validation (RQ4)

To verify whether Hi-Brain effectively realizes a brain-inspired bidirectional hierarchical information-processing mechanism, we design a hierarchical temporal prediction experiment. For bottom-up functional aggregation, we construct two prediction tasks. First, the fMRI signal of a single brain region is used to predict the average signal of the remaining regions within the same

subnetwork, evaluating node to subnetwork information aggregation. Second, the average fMRI signal of one subnetwork is used to predict the average signal of the remaining subnetworks, evaluating subnetwork to whole brain information integration. For top-down functional regulation, the signals of all remaining brain regions are used to predict the target brain region, evaluating whole-brain regulation of local brain activity. For comparison, we select three baselines: LSTM (Greff et al., 2016) as a classical temporal modeling method, BrainMass (Yang et al., 2024) as an advanced fMRI foundation model, and Spike-DB (Wei et al., 2026) as an advanced brain-inspired model. We use MSE, R^2, and Pearson correlation (PC) as evaluation metrics, and the results are reported in Table 3. Notably, Hi-Brain consistently performs best across the node to subnetwork, subnetwork to whole brain, and whole brain to node information-flow directions. For example, its MSE values on the three tasks are 0.1046, 0.1049, and 0.1015, respectively. These results indicate that Hi-Brain can more effectively model cross-level information transmission among nodes, subnetworks, and the whole brain. In contrast, LSTM mainly models general temporal dependencies, while BrainMass and Spike-DB provide stronger brain activity modeling capabilities but lack an explicit bidirectional hierarchical mechanism. Moreover, Fig. 4 presents representative prediction sequences for three tasks. The temporal consistency between the predicted and observed signals indicates that Hi-Brain can infer the dynamic changes of one brain level from the activity of another, suggesting that the model realizes a brain-inspired bidirectional hierarchical interaction mechanism.

Table 3: Validation of the brain-inspired mechanism via temporal prediction.

Method	Node→Subnetwork			Subnetwork→Whole Brain			Whole Brain→Node		
	MSE	R^2	PC	MSE	R^2	PC	MSE	R^2	PC
LSTM	0.2015	0.6088	0.7643	0.1493	0.6757	0.8222	0.1343	0.8244	0.9081
BrainMass	0.1364	0.7441	0.8655	0.1347	0.7073	0.8614	0.1239	0.8270	0.9095
Spike-DB	0.1472	0.7250	0.8528	0.1244	0.7298	0.8554	0.1166	0.8283	0.9103
Hi-Brain	**0.1046**	**0.8014**	**0.9032**	**0.1049**	**0.7722**	**0.8830**	**0.1015**	**0.8385**	**0.9161**

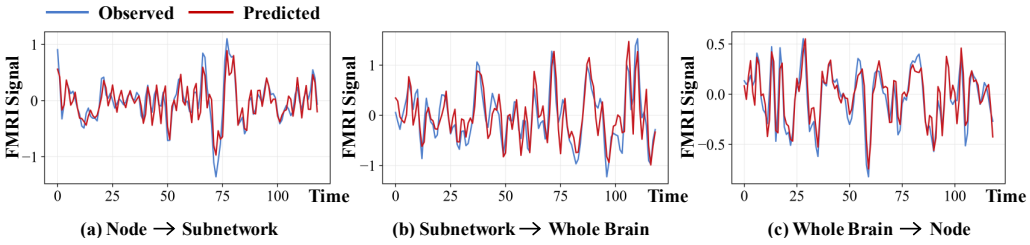

Figure 4: Visualization of predicted and observed fMRI signals in the temporal prediction tasks.

5 Conclusion and Future Work

Inspired by the human nervous system, we propose a bidirectional hierarchical neural network for dynamic brain activity analysis. It can effectively simulate bottom-up functional aggregation and top-down functional regulation in the brain. Specifically, we first divide the fMRI signals into multiple subsequences to characterize the dynamics of the brain. For each subsequence, we design a neuro-inspired node activator to capture the adaptive activation of neural nodes under different states. On this basis, we further introduce a high-order subnetwork learner to characterize efficient differentiation and collaboration among subnetworks, enabling bottom-up functional aggregation. Subsequently, the aggregated functional information is further used to optimize information extraction in lower-level neural organizations, thereby realizing top-down functional regulation. Experimental results on the ADNI and PD datasets show that Hi-Brain not only achieves strong predictive performance, but also reveals the neural mechanisms underlying its predictions. Notably, Hi-Brain provides the machine learning community with a more biologically plausible computational framework. Unfortunately, Hi-Brain analyzes brain activity solely from fMRI signals, without explicitly considering anatomical structure. Structural connectivity is the basis of functional interactions. In future work, we will explore functional regulation under structural constraints, thereby further improving the biological plausibility of the model.

AI USE STATEMENT

In this work, we used generative AI tools for language translation and polishing. We did not use generative AI tools for research idea generation, methodology design, experiment implementation, data analysis, or result interpretation. All scientific content was produced by the authors. We have reviewed all AI-assisted text and take full responsibility for the final content of this work.

ETHICS STATEMENT

The ADNI dataset used in this study is publicly available and was collected under protocols approved by the relevant institutional review boards. The PD dataset was collected at an affiliated hospital of Nanjing Medical University in accordance with relevant ethical guidelines and regulations. All participants provided written informed consent before participation.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of the datasets and preprocessing procedures, experimental settings, and data partitioning strategies in the experimental setup section. In addition, the hyperparameter settings of the proposed model are provided in Appendix B. The source code is also provided in the supplementary materials. These details provide sufficient information to support the reproducibility of the proposed method.

REFERENCES

Lisa Feldman Barrett and Earl K Miller. Categorization is 'baked'into the brain. *Nature Reviews Neuroscience*, pp. 1–22, 2026.

Pablo Barttfeld, Lynn Uhrig, Jacobo D Sitt, Mariano Sigman, Béchir Jarraya, and Stanislas Dehaene. Signature of consciousness in the dynamics of resting-state brain activity. *Proceedings of the National Academy of Sciences*, 112(3):887–892, 2015.

Uranik Berisha, Jens Mehnert, and Alexandru Paul Condurache. Efficient data driven mixture-of-expert extraction from trained networks. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 20082–20091, 2025.

Ornella V Billette, Gabriel Ziegler, Merita Aruci, Hartmut Schütze, Jasmin M Kizilirmak, Anni Richter, Slawek Altenstein, Claudia Bartels, Frederic Brosseron, Arturo Cardenas-Blanco, et al. Novelty-related fmri responses of precuneus and medial temporal regions in individuals at risk for alzheimer disease. *Neurology*, 99(8):e775–e788, 2022.

Dongdong Chen, Mengjun Liu, Zhenrong Shen, Xiangyu Zhao, Qian Wang, and Lichi Zhang. Learnable subdivision graph neural network for functional brain network analysis and interpretable cognitive disorder diagnosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 56–66. Springer, 2023.

Zhaonan Chen, Yanmei Liu, Yunqi Yang, Lizhao Wang, Meiling Qin, Zhishan Jiang, Min Xu, and Siyu Zhang. Whole-brain mapping of basal forebrain cholinergic neurons reveals a long-range reciprocal input-output loop between distinct subtypes. *Science Advances*, 11(22):eadt1617, 2025.

Nick Corriveau-Lecavalier, Ellen Dicks, Peter R Martin, Heather J Wiste, Jeffrey L Gunter, Michael G Kamykowski, Matthew L Senjem, Christopher G Schwarz, Hugo Botha, Jonathan Graff-Radford, et al. Default mode network failure across the alzheimer's disease spectrum. *Brain*, pp. awag146, 2026.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems*, 26, 2013.

Haowen Fang, Amar Shrestha, Ziyi Zhao, and Qinru Qiu. Exploiting neuron and synapse filter dynamics in spatial temporal learning of deep spiking neural network. *arXiv preprint arXiv:2003.02944*, 2020.

Bin Gao, Aiju Yu, Chen Qiao, Vince D Calhoun, Julia M Stephen, Tony W Wilson, and Yu-Ping Wang. An explainable unified framework of spatio-temporal coupling learning with application to dynamic brain functional connectivity analysis. *IEEE Transactions on Medical Imaging*, 44 (2):941–951, 2024.

Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28(10): 2222–2232, 2016.

Yuye Huang, David O'Connor, Bailey B Harris, Rajita Sinha, R Todd Constable, and Elizabeth V Goldfarb. Dynamic brain mechanisms supporting salient memories under cortisol. *Science Advances*, 11(50):eadz4143, 2025.

Harang Ju and Danielle S Bassett. Dynamic representations in networked neural systems. *Nature Neuroscience*, 23(8):908–917, 2020.

Xuan Kan, Hejie Cui, Joshua Lukemire, Ying Guo, and Carl Yang. Fbnetgen: Task-aware gnn-based fmri analysis via functional brain network generation. In *International Conference on Medical Imaging with Deep Learning*, pp. 618–637. PMLR, 2022a.

Xuan Kan, Wei Dai, Hejie Cui, Zilong Zhang, Ying Guo, and Carl Yang. Brain network transformer. *Advances in Neural Information Processing Systems*, 35:25586–25599, 2022b.

Fumi Katsuki and Christos Constantinidis. Early involvement of prefrontal cortex in visual bottom-up attention. *Nature Neuroscience*, 15(8):1160–1166, 2012.

Byung-Hoon Kim, Jong Chul Ye, and Jae-Jin Kim. Learning dynamic graph representation of brain connectome with spatio-temporal attention. *Advances in Neural Information Processing Systems*, 34:4314–4327, 2021.

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.

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.

Yicheng Long, Chujun Chen, Mengjie Deng, Xiaojun Huang, Wenjian Tan, Li Zhang, Zebin Fan, and Zhening Liu. Psychological resilience negatively correlates with resting-state brain network flexibility in young healthy adults: a dynamic functional magnetic resonance imaging study. *Annals of Translational Medicine*, 7(24):809, 2019.

Marcelo D Mendonça, Pedro C Ferreira, Francisco Oliveira, Raquel Barbosa, Bruna Meira, Durval C Costa, Albino J Oliveira-Maia, and Joaquim Alves da Silva. Relative sparing of dopaminergic terminals in the caudate nucleus is a feature of rest tremor in parkinson's disease. *npj Parkinson's Disease*, 10(1):209, 2024.

Davide Momi, Zheng Wang, Sara Parmigiani, Ezequiel Mikulan, Sorenza P Bastiaens, Mohammad P Oveisi, Kevin Kadak, Gianluca Gaglioti, Allison C Waters, Sean Hill, et al. Stimulation mapping and whole-brain modeling reveal gradients of excitability and recurrence in cortical networks. *Nature Communications*, 16(1):3222, 2025.

Akira Nishida, Atsushi Shima, Daisuke Kambe, Koji Furukawa, Haruhi Sakamaki-Tsukita, Kenji Yoshimura, Ikko Wada, Yusuke Sakato, Yuta Terada, Masanori Sawamura, et al. Frontoparietal-striatal network and nucleus basalis modulation in patients with parkinson disease and gait disturbance. *Neurology*, 103(3):e209606, 2024.

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.

Ahsan Shehzad, Dongyu Zhang, Shuo Yu, Shagufta Abid, and Feng Xia. Dynamic graph transformer for brain disorder diagnosis. *IEEE Journal of Biomedical and Health Informatics*, 29(6):4388–4400, 2025.

Yuanhong Tang, Yan Yan, Jian Mao, Junjun Ni, and Hong Qing. The hippocampus associated gabaergic neural network impairment in early-stage of alzheimer's disease. *Ageing Research Reviews*, 86:101865, 2023.

Qianqian Wang, Mengqi Wu, Yuqi Fang, Wei Wang, Lishan Qiao, and Mingxia Liu. Modularity-constrained dynamic representation learning for interpretable brain disorder analysis with functional mri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 46–56. Springer, 2023.

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.

Shaolong Wei, Qiyu Sun, Mingliang Wang, Liang Sun, Weiping Ding, and Jiashuang Huang. Spike-based digital brain: a novel fundamental model for brain activity analysis. In *International Conference on Learning Representations*, volume 2026, pp. 108842–108858, 2026.

Zhuang Wu, Sha Zhu, Ronghua Hong, Zhuoyu Zhang, Yanzi Peng, Jingxing Zhang, Lizhen Pan, Qiang Guan, Yuhui Chen, and Lingjing Jin. A network centered on the visual-motor cortex is critically involved in postural abnormality in parkinson's disease. *Movement Disorders*, 41(1): 156–167, 2026.

Chaogan Yan and Yufeng Zang. Dparsf: a matlab toolbox for" pipeline" data analysis of resting-state fmri. *Frontiers in Systems Neuroscience*, 4:1377, 2010.

Yanwu Yang, Chenfei Ye, Guinan Su, Ziyao Zhang, Zhikai Chang, Hairui Chen, Piu Chan, Yue Yu, and Ting Ma. Brainmass: Advancing brain network analysis for diagnosis with large-scale self-supervised learning. *IEEE Transactions on Medical Imaging*, 43(11):4004–4016, 2024.

Ling Zhan, Zhen Li, Junjie Huang, and Tao Jia. Accelerating benchmarking of functional connectivity modeling via structure-aware core-set selection. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=0RYazbfSzW.

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, 2024a.

Jing Zhang, Xiaowei Yu, Tong Chen, Chao Cao, Minheng Chen, Yan Zhuang, Yanjun Lyu, Lu Zhang, Li Su, Tianming Liu, et al. Brainnet-moe: Brain-inspired mixture-of-experts learning for neurological disease identification. In *2025 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 1995–2001. IEEE, 2025a.

Moxuan Zhang, Siyu Zhou, Pengda Yang, Huizhi Wang, Jinli Ding, Xiaobo Wang, Xuzhu Chen, Chaonan Zhang, Anni Wang, Yuan Gao, et al. The precuneus region drives brain network changes in tremor-dominant parkinson's disease: Insights from a morphological causal analysis. *MedComm*, 6(11):e70441, 2025b.

Ying Zhang, Le Xue, Shuoyan Zhang, Jiacheng Yang, Qi Zhang, Min Wang, Luyao Wang, Mingkai Zhang, Jiehui Jiang, Yunxia Li, et al. A novel spatiotemporal graph convolutional network framework for functional connectivity biomarkers identification of alzheimer's disease. *Alzheimer's Research & Therapy*, 16(1):60, 2024b.

Qi Zhu, Shengrong Li, Xiangshui Meng, Qiang Xu, Zhiqiang Zhang, Wei Shao, and Daoqiang Zhang. Spatio-temporal graph hubness propagation model for dynamic brain network classification. *IEEE Transactions on Medical Imaging*, 43(6):2381–2394, 2024.

A Definitions of Filter Coefficients

In Eq. (4), α_1, α_2, and β are predefined filter coefficients used to control temporal dynamics. Their detailed derivations are as follows:

$$\alpha_1 = e^{-\frac{1}{\tau_m}} + e^{-\frac{1}{\tau_s}} \tag{11}$$

$$\alpha_2 = -e^{-\frac{\tau_m+\tau_s}{\tau_m\tau_s}} \tag{12}$$

$$\beta = e^{-\frac{1}{\tau_m}} - e^{-\frac{1}{\tau_s}} \tag{13}$$

where τ_m and τ_s denote time constants.

B Parameter Settings

Hi-Brain is implemented based on PyTorch and snnTorch. The settings of the model parameters and hyperparameters are shown in Table 4.

Table 4: Different parameter settings.

Config	Value
Training epoch	300
Early stopping patience	25
Threshold potential V_{th}	1
Decay speed τ_r	0.9
Time constant τ_m	4
Time constant τ_s	1
Dropout	0.5
Optimizer	Adam
Learning rate	0.003
Hi-Brain layers	2

C Comparison Results with Standard Deviation

Due to text layout and page limitations, only the mean values from 10 tests are presented in the main text. To ensure a comprehensive presentation of the experimental results, we list the mean values and standard deviations (std) for each evaluation metric in Table 5 and Table 6.

Table 5: Classification results (mean/std) of different methods on the ADNI dataset (%).

Method	ACC	F1	AUC	ACC	F1	AUC
	NC vs MCI vs AD			NC vs MCI		
FBNetGen	57.50/06.88	56.97/04.69	67.75/06.11	61.25/09.29	66.46/11.22	58.47/12.03
BNT	60.21/03.78	40.94/07.51	61.96/06.90	69.75/06.17	79.65/04.68	57.58/09.83
STAGIN	56.46/05.62	27.75/05.83	54.76/05.34	69.25/03.54	80.97/01.53	56.44/07.40
ACIFBN	62.35/05.02	57.88/06.12	70.42/03.40	71.25/04.37	78.96/05.37	64.49/08.50
OT-STGCN	59.02/06.82	41.31/10.78	58.52/09.15	76.08/07.28	80.14/03.83	60.89/08.05
LSGNN	58.75/02.91	36.96/09.56	58.20/04.99	60.27/08.83	68.62/13.00	54.04/08.45
CAGT	60.75/02.97	36.25/10.65	59.29/06.91	75.75/05.37	83.81/02.88	62.68/14.58
BrainMAP	57.65/04.31	50.53/06.11	63.08/04.32	70.39/08.22	77.75/07.70	60.82/14.26
ST2GCN	60.39/05.93	54.44/07.19	63.91/04.36	75.49/04.21	83.67/02.36	65.65/08.15
Hi-Brain	**67.60/04.18**	**61.26/07.89**	**74.93/07.78**	**79.00/07.52**	**85.61/04.75**	**72.11/13.95**
Method		NC vs AD			MCI vs AD	
FBNetGen	63.33/07.86	58.50/10.03	57.53/11.74	65.83/10.40	73.71/08.32	64.07/12.72
BNT	74.00/10.68	71.72/07.97	75.89/09.99	74.50/09.07	81.56/08.43	57.64/18.27
STAGIN	67.08/04.73	47.99/12.31	55.49/15.81	74.44/01.11	85.35/00.73	53.71/11.38
ACIFBN	79.17/05.27	71.53/10.66	78.44/08.56	80.54/05.90	87.13/03.65	75.82/11.19
OT-STGCN	73.75/08.34	57.61/23.67	66.41/12.45	76.26/06.02	85.40/03.93	68.15/06.89
LSGNN	70.42/06.83	53.07/18.88	66.26/13.00	64.05/09.21	73.72/10.06	56.11/07.38
CAGT	80.83/04.64	76.90/05.47	83.29/06.59	75.75/04.48	83.45/02.17	67.40/12.03
BrainMAP	76.08/07.54	67.37/15.48	71.64/12.83	79.46/03.67	86.29/02.92	69.52/06.02
ST2GCN	78.10/07.91	73.34/10.60	76.43/11.61	82.97/03.64	88.66/03.02	77.10/08.08
Hi-Brain	**85.00/03.82**	**81.08/05.06**	**84.93/05.13**	**85.28/03.52**	**90.62/01.88**	**77.52/11.62**

Table 6: Classification results (mean/std) of different methods on the PD dataset (%).

Method	ACC	F1	AUC	ACC	F1	AUC
	NC vs TPD vs PPD			NC vs TPD		
FBNetGen	62.50/13.11	57.65/16.24	69.23/13.49	74.56/08.13	52.80/31.54	65.65/23.42
BNT	63.75/05.48	56.47/08.10	71.75/03.40	82.50/06.12	70.88/24.91	74.83/12.12
STAGIN	55.63/07.63	43.52/10.77	63.71/04.28	81.89/09.71	74.20/26.10	76.55/21.18
ACIFBN	61.25/10.75	48.72/17.28	66.33/16.23	82.89/10.93	77.91/16.31	79.80/18.10
OT-STGCN	59.38/08.95	38.24/08.67	62.37/14.32	81.56/10.89	79.31/13.10	81.59/13.40
LSGNN	56.88/08.59	46.44/11.69	64.46/10.71	73.75/27.07	74.84/17.88	77.08/16.43
CAGT	59.38/05.76	57.06/08.42	66.25/08.37	81.25/06.25	75.88/08.15	80.42/13.62
BrainMAP	57.24/09.07	53.04/10.34	64.75/09.84	82.50/06.12	77.85/08.38	81.63/10.27
ST2GCN	60.63/09.70	53.74/13.57	65.30/12.33	86.25/08.75	83.57/11.01	77.75/14.76
Hi-Brain	**66.14/06.93**	**59.21/12.07**	**74.09/11.87**	**92.50/06.12**	**90.03/08.59**	**87.54/13.36**
Method		NC vs PPD			TPD vs PPD	
FBNetGen	74.55/07.51	78.39/04.73	51.82/25.05	72.00/12.49	73.05/14.63	66.14/19.14
BNT	79.00/09.60	77.18/15.50	78.62/11.08	76.73/09.82	79.50/09.28	70.20/14.43
STAGIN	82.20/06.04	80.89/10.86	79.30/09.40	78.00/13.27	63.00/36.08	59.85/21.41
ACIFBN	76.97/13.27	73.90/26.74	68.59/24.62	74.18/14.29	73.36/26.45	63.79/10.20
OT-STGCN	82.20/08.69	83.59/09.09	75.50/15.85	75.82/11.14	73.39/25.59	63.05/18.53
LSGNN	74.17/08.98	76.93/09.53	68.47/16.28	72.00/07.48	76.61/08.45	58.62/13.83
CAGT	78.33/11.76	80.06/11.19	73.99/16.68	80.00/08.29	85.43/08.51	60.00/22.27
BrainMAP	80.90/06.36	81.79/05.89	81.57/11.23	77.50/12.25	81.40/10.02	76.45/25.56
ST2GCN	83.75/08.00	85.30/08.81	81.62/13.08	81.00/10.44	82.69/10.31	80.82/14.22
Hi-Brain	**86.36/08.79**	**87.04/08.73**	**84.08/11.48**	**85.75/09.09**	**85.47/08.60**	**82.35/13.98**

D PARAMETER SENSITIVITY ANALYSIS

We first explore the effect of the number of sliding windows K on accuracy. As shown in Fig. 5(a), the performance is optimal when K=5 or 6. This is because a small K fails to capture the dynamic evolution, whereas an excessively large K provides limited temporal information within each window, making it difficult to accurately reflect information interactions among brain regions. Therefore, K=5 or 6 provides a better trade-off between capturing temporal dynamics and effective interaction.

We also conduct a sensitivity analysis on λ_1 and λ_2. In Fig. 5(b), the accuracy first increases and then decreases as λ_1 and λ_2 grow. This is because a small λ_1 makes it difficult for the model to distinguish different neural states, while a large λ_1 may overemphasize state separation and lead to distorted representations. Meanwhile, a small λ_2 results in unclear expert specialization, whereas a large λ_2 causes the model to rely on only a few experts for feature learning, thereby weakening its representation capacity. Therefore, moderate values of λ_1 and λ_2 are more beneficial for achieving superior performance.

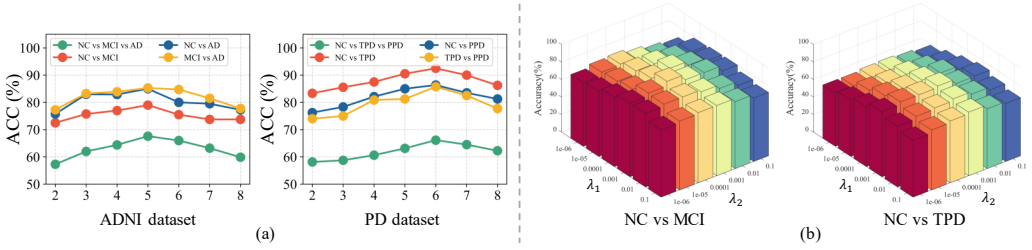

Figure 5: (a) Impact of K on accuracy. (b) Impact of λ_1 and λ_2 on accuracy.

Number of node-level experts M and subnetwork-level experts N are both varied within $\{1, 2, 3, 4, 5, 6\}$. As shown in Fig. 6, it can be observed that the accuracy rate does not significantly change with the variations of M and N. This indicates that the proposed method has excellent robustness.

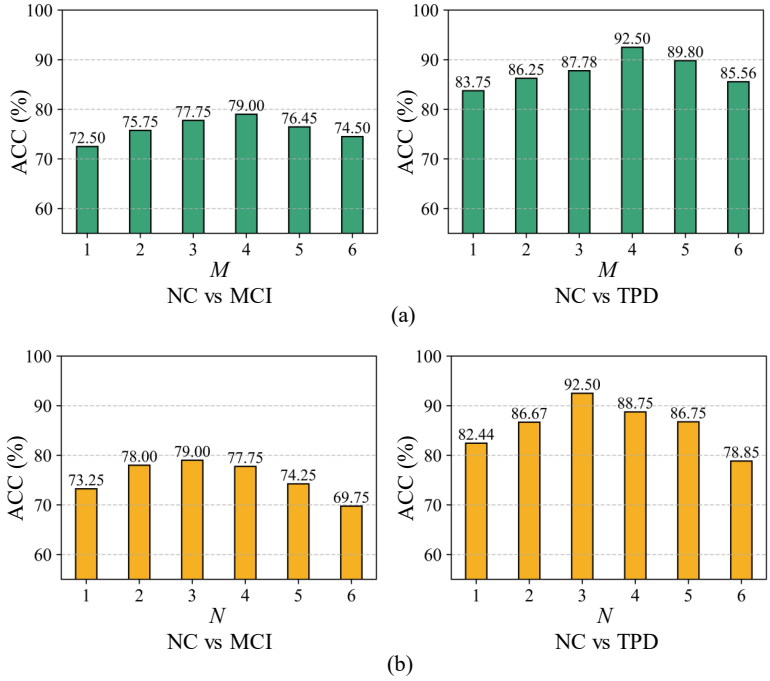

Figure 6: Impact of M and N on accuracy across diagnostic tasks.

E Ablation on Losses

To verify the effectiveness of similarity constraint $\mathcal{L}_{CS}$ and entropy constraint $\mathcal{L}_{EN}$, we conduct additional ablation experiments. As shown in Table 7, removing either the $\mathcal{L}_{CS}$ or $\mathcal{L}_{EN}$ decreases the prediction performance, while combining both achieves the best performance. This phenomenon indicates that these losses can capture the functional specificity and collaboration of subnetworks, thereby improving diagnostic performance.

Table 7: Ablation results of $\mathcal{L}_{CS}$ and $\mathcal{L}_{EN}$ on the ADNI and PD datasets (%).

Dataset	$\mathcal{L}_{CS}$	$\mathcal{L}_{EN}$	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC
			NC vs MCI vs AD			NC vs MCI			NC vs AD			MCI vs AD		
ADNI	✓	✗	62.08	53.82	67.60	73.25	81.73	68.25	81.25	76.34	80.45	78.61	84.31	74.61
	✗	✓	63.33	52.73	67.47	75.25	83.52	67.44	82.50	77.77	83.00	79.44	85.16	77.03
	✓	✓	**67.60**	**61.26**	**74.93**	**79.00**	**85.61**	**72.11**	**85.00**	**81.08**	**84.93**	**85.28**	**90.62**	**77.52**
			NC vs TPD vs PPD			NC vs TPD			NC vs PPD			TPD vs PPD		
PD	✓	✗	56.88	52.64	63.36	87.50	77.37	87.38	77.92	80.29	72.48	77.78	80.91	72.06
	✗	✓	64.12	55.80	71.89	90.00	85.12	87.00	81.25	85.40	72.46	82.22	84.17	77.21
	✓	✓	**66.14**	**59.21**	**74.09**	**92.50**	**90.03**	**87.54**	**86.36**	**87.04**	**84.08**	**85.75**	**85.47**	**82.35**