

HyperDiag: Temporal-Regional Hypergraph Learning via Topology-Enhanced State Propagation for Brain Disease Diagnosis

Yulan Ma¹, Fangkun Li¹, Wenchao Yang¹, Qian Si², Chenglong Yu¹, Yang Li^{1*}

¹School of Automation Science and Electrical Engineering, Beihang University, China

²School of Cyber Science and Technology, Beihang University, China

{mayulan, fangkunli, yangwenchao, siqian, yuyubuaa, liyang}@buaa.edu.cn

Abstract

Dynamic brain networks provide a powerful representation for capturing temporal variations in functional brain connectivity and have gained increasing attention in brain disease diagnosis. However, most existing methods extract features from isolated time windows, making it difficult to capture the high-order dynamic evolution of brain activity. Moreover, these methods often neglect the functional heterogeneity among brain regions, thereby limiting diagnostic performance. To address these limitations, we propose **HyperDiag**, a novel temporal-regional **Hypergraph** learning via topology-enhanced state propagation for brain disease **Diagnosis**. Specifically, we first design a dual-level hypergraph learning strategy: a temporally-evolving hypergraph message passing strategy to capture dynamic high-order dependencies within and across time windows, and meanwhile, a region-wise functional hypergraph learning strategy to capture regional dependencies. Subsequently, we construct a topology-enhanced selective state-space propagation network to integrate complementary information from both the temporally-evolving and region-wise features. Extensive experiments on four brain disorder datasets (ABIDE-I, ADNI, REST-meta-MDD, and Epilepsy) demonstrate that HyperDiag not only outperforms state-of-the-art methods but also identifies biologically meaningful abnormal connections, offering potential biomarkers for clinical interpretation.

Introduction

Accurate and early diagnosis of brain disorders, such as autism spectrum disorder (ASD), Alzheimer’s disease (AD), major depressive disorder (MDD) and epilepsy, is essential for enabling timely intervention and improving patient outcomes (Lord et al. 2018). However, reliable diagnosis remains challenging due to subtle, dynamic, and heterogeneous brain alterations that are difficult to detect using standard clinical assessments (Ma et al. 2024). Recent studies have shown that dynamic brain networks derived from resting-state functional magnetic resonance imaging (rs-fMRI) are both reliable and sensitive for detecting abnormal brain patterns associated with various neurological and psychiatric disorders (Wen et al. 2025).

Conventional dynamic brain networks primarily model pairwise brain regions of interest (ROIs) correlations from isolated time windows, which are insufficient to capture the complex high-order interactions of brain activity (Yang et al. 2019; Cui et al. 2023b). To improve model performance, hyper-connectivity networks (HCNs) have been introduced to characterize high-order functional dependencies across multiple ROIs, offering a more comprehensive representation (Liu et al. 2024; Han, Lei, and Li 2025). However, most existing HCNs only capture spatial topologies, while neglecting the intrinsic temporal dynamics embedded in rs-fMRI signals. Furthermore, growing evidence suggests that brain activity is inherently non-stationary and continuously evolves over time (Xu et al. 2025). Therefore, a novel high-order brain network construction approach that integrates temporal dynamics is required to facilitate the discovery of meaningful functional biomarkers for brain disorder.

In addition to temporal dynamics, the human brain exhibits a modular anatomical organization, including regions like the limbic region, frontal region, and subcortical nuclei, each of which supports specific cognitive and emotional processes (Jiao et al. 2024; Van Overwalle 2024; Shao, Li, and Wu 2025). Many existing brain network models treat the brain as a homogeneous network, neglecting regional modularity and anatomical priors, which reduces diagnostic specificity and limits their ability to detect localized pathological disruptions (Makhlouf et al. 2025; Han, Lei, and Li 2025). Moreover, brain disorders affect distinct brain regions, and the functional interactions between these regions are crucial for characterizing disease manifestations (Chen, Dang, and Zhang 2021). Therefore, a feature learning network is required that not only captures connectivity patterns within individual brain regions but also models functional interactions across regions.

After analyzing brain activity patterns, effectively integrating discriminative features is critical for robust brain disorder diagnosis. For example, MMTGCN (Yao et al. 2021) adopts a mutual multi-scale triplet graph convolutional network to integrate deep features of brain connectivity. STGHP (Zhu et al. 2024) develops a multi-channel spatio-temporal graph convolutional network that collaboratively extracts both temporal and spatial features. DCLNet (Zhou et al. 2025) leverages both collaborative encoders and contrastive learning to capture complementary information

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

from static and dynamic functional connectivity. Despite improvements in performance, integrating different connectivity patterns remains challenging due to the spatial modularity and dynamic nature of brain activity, which require capturing long-range dynamic dependencies and more robust modeling (Béna and Goodman 2025). Recently, structured state-space models (SSMs), exemplified by Mamba (Gu and Dao 2023), have demonstrated remarkable performance in long-range context modeling tasks within natural language processing and computer vision by enabling efficient and selective information flow (Yu and Wang 2025). However, their application to neuroimaging, particularly in non-Euclidean data like rs-fMRI, remains underexplored, since current SSMs ignore the underlying topological dependencies, which are crucial for accurately modeling the spatial modularity and dynamic interactions in brain networks.

To address the aforementioned challenges, we propose **HyperDiag**, a novel temporal-regional **Hypergraph** learning via topology-enhanced state propagation for brain disease **Diagnosis**. Specifically, we first design a temporally-evolving hypergraph message passing strategy to capture dynamic high-order interactions within and across time windows. Next, we construct region-wise functional hypergraph learning strategy, which leverages location prior information to explore regional dependencies through region-specific hypergraphs and a spectral-aware diffusion mechanism. Finally, we develop a topology-enhanced selective state-space propagation network to deeply explore complementary information within evolving and regional features. Extensive evaluations on the ABIDE-I, ADNI, REST-meta-MDD, and Epilepsy datasets demonstrate that HyperDiag not only achieves superior diagnostic performance but also uncovers biologically meaningful abnormal connections, offering potential biomarkers for clinical interpretation. Our contributions are:

- We propose HyperDiag, a novel temporal-regional hypergraph learning via topology-enhanced state propagation for brain disease diagnosis, achieving superior performance and identifying potential biomarkers for clinical interpretation.
- We design a dual-level hypergraph learning strategy: a temporally-evolving hypergraph message passing strategy to capture dynamic high-order dependencies within and across time windows, and a region-wise functional hypergraph learning strategy to explore anatomical modularity and fuse them via spectral diffusion.
- We introduce a topology-enhanced selective state-space propagation network that incorporates a novel mechanism to integrate topological cues from both temporal dynamics and anatomical structures, facilitating selective and long-range fusion of complementary information.

Related Work

High-order Brain Network for Disease Analysis

High-order brain network modeling, has shown great potential for brain disorder diagnosis. For example, NHCM (Du, Niu, and Calhoun 2021) introduced a hypergraph clustering framework to identify latent biotypes in mental illnesses.

DHGI (Zhu et al. 2019) proposed a dynamic hypergraph inference model; however, the term "dynamic" refers to iterative structural optimization rather than the explicit modeling of temporal dynamics. More recently, I²HBN (Han et al. 2024) developed an inter-intra high-order brain network framework for ASD diagnosis. HGFM (Han et al. 2025) presented a hypergraph foundation model to improve diagnostic performance. Despite these advances, existing methods fail to capture the dynamic high-order dependencies among three or more brain regions that evolve over time (Liu et al. 2024). In addition, most current models treat the brain as a homogeneous network, neglecting region-specific high-order interactions shaped by anatomical and functional modularity (Makhlouf et al. 2025; Shi et al. 2025).

State Spaces for Sequence Modeling

State space models (SSMs) offer a principled framework for modeling dynamical systems by representing latent temporal processes. While classical SSMs are limited in modeling long-range dependencies due to time-invariant transition mechanisms, recent developments have revisited SSMs with data-dependent and structured enhancements. Among these, Mamba (Gu and Dao 2023) introduces a selective scan-based state space formulation that dynamically modulates the information flow via input-conditioned transition matrices, achieving superior performance in long-sequence modeling. Mamba and its variants have demonstrated remarkable success in natural language processing (Zhang et al. 2024), vision (Ahamed and Cheng 2024; Liu et al. 2025), and content-based reasoning (Sarem et al. 2024), yet their application to non-Euclidean data.

Method

The proposed framework is presented in Fig. 1, further details of each step are described in the following sections.

Temporally-Evolving Hypergraph Message Passing

Evolutionary Hypergraph Construction Understanding the temporal evolution of brain connectivity patterns is essential for accurately modeling dynamic functional brain networks. To this end, we propose temporally-evolving hypergraph message passing that jointly captures intra-window high-order functional relationships and inter-window semantic transitions. Given the fMRI time series $\mathbf{X} \in \mathbb{R}^{T \times N}$ for a subject with T time points and N brain ROIs, we first segment the data into W non-overlapping temporal windows of fixed length L . This results in a set of sub-sequences $\{\mathbf{X}^w\}_{w=1}^W$, where each segment $\mathbf{X}^w \in \mathbb{R}^{L \times N}$ corresponds to the signal within the w -th window. For each temporal window, we construct a window-specific sparse hypergraph $\mathcal{H}^w = (\mathcal{V}, \mathcal{E}^w)$, where the vertex set \mathcal{V} corresponds to brain ROIs, and the hyperedge set \mathcal{E}^w captures high-order interactions among them. To capture functional relationships, we express the time series $\mathbf{x}_n^w \in \mathbb{R}^L$ of n -th ROI as a sparse linear combination of the remaining ROIs $\hat{\mathbf{X}}^w \in \mathbb{R}^{L \times (N-1)}$:

$$\min_{\alpha_n^w} \frac{1}{2} \left\| \mathbf{x}_n^w - \hat{\mathbf{X}}^w \alpha_n^w \right\|_2^2 + \lambda_1 \|\alpha_n^w\|_1, \quad (1)$$

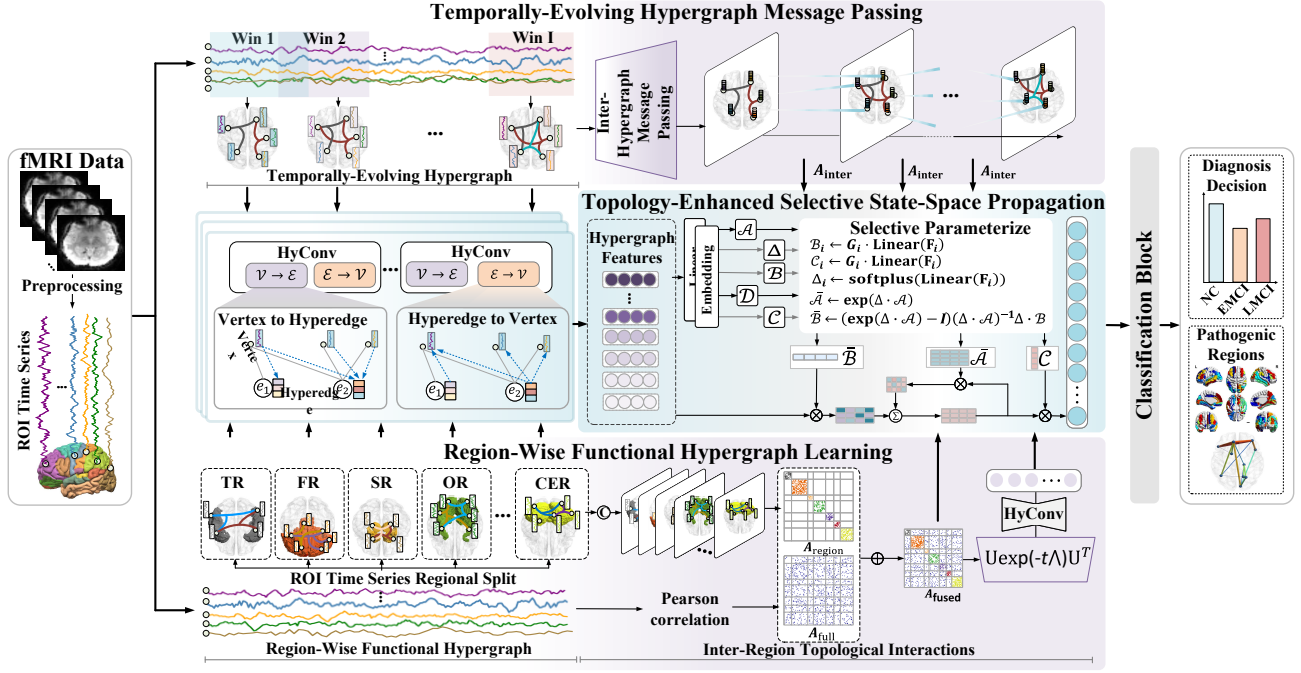


Figure 1: The schematic diagram of the proposed HyperDiag.

where $\alpha_n^w \in \mathbb{R}^{N-1}$ is the sparse coefficient vector quantifying the influence of other ROIs on the n -th ROI, and λ is a regularization parameter that controls the sparsity. Higher values in α_n^w indicate stronger functional interactions. By applying sparse representation to each ROI, we obtain the interactive relationships between the centroid ROI and other ROIs that serve as hyperedges. Repeating this process for all ROIs results in a set of M hyperedges $\mathcal{E}^w = \{e_1^w, e_2^w, \dots, e_M^w\}$. The hypergraph is then represented by incidence matrix $\mathbf{H}^w \in \{0, 1\}$, where $\mathbf{H}^w(n, e) = 1$ if ROI n participates in hyperedge e^w .

Intra-Inter Hypergraph Message Passing Subsequently, we extract high-order semantic representations from each hypergraph via a stack of hypergraph convolution (HyConv) layers, which propagate information across ROIs through shared hyperedges. The l -th HyConv layer is formulated as:

$$\mathbf{Z}_l^w = \sigma(\mathbf{H}^w \mathbf{W}_e \mathbf{D}_e^{-1} (\mathbf{H}^w)^\top \mathbf{Z}_{l-1}^w \Theta_l), \quad (2)$$

where $\mathbf{Z}_0^w = \mathbf{X}^w$, Θ_l is a learnable projection matrix, and \mathbf{D}_e denotes the hyperedge degree matrix. To enhance generalization and mitigate overfitting to overly dominant nodes, we introduce DropMax regularization after the second HyConv layer:

$$\mathbf{Z}_2^{w'} = \text{DropMax}(\mathbf{Z}_2^w) = \mathbf{Z}_2^w \odot \mathbf{M}, \quad (3)$$

where \mathbf{M} is a binary mask applied to suppress highly activated nodes. Ultimately, three HyConv layers are applied to each window, yielding intra-window hypergraph features $\{\mathbf{Z}^w\}_{w=1}^W$.

To further capture temporal evolution and semantic transitions across adjacent brain states, inter-hypergraph kernelized attention is introduced to propagate information across

windows. For a sequence of intra-window hypergraph features $\{\mathbf{Z}^w\}_{w=1}^W$, we aim to propagate information from \mathbf{Z}^{w+1} to \mathbf{Z}^w for all $w \in \{1, \dots, W-1\}$. Each \mathbf{Z}^w is projected into queries $\mathbf{Q}^w = \mathbf{Z}^w \mathbf{W}_q$, keys $\mathbf{K}^{w+1} = \mathbf{Z}^{w+1} \mathbf{W}_k$, and values $\mathbf{V}^{w+1} = \mathbf{Z}^{w+1} \mathbf{W}_v$ using shared linear projections. To enable efficient inter-graph propagation, we apply a random feature approximation approach. Each query vector \mathbf{q}_i^w (i.e., the i -th row of \mathbf{Q}^w) and each key vector \mathbf{k}_j^{w+1} (i.e., the j -th row of \mathbf{K}^{w+1}) are independently mapped into a shared kernel space using a learnable projection matrix \mathcal{M} . Moreover, to enhance stochasticity and exploration in learning inter-graph dependencies, Gumbel noise is injected into the kernelized transformation:

$$\phi(\mathbf{p}) = \frac{1}{\sqrt{m}} \exp\left(\frac{\mathcal{M}^\top \mathbf{p} + \mathcal{G}}{\tau} - \frac{1}{2} \|\mathbf{p}\|^2\right), \quad (4)$$

where $\mathbf{p} \in \{\mathbf{q}_i^w, \mathbf{k}_j^{w+1}\}$ denotes a node-level feature vector from either the current or next window, $\mathcal{G} \sim \text{Gumbel}(0, 1)$ is sampled per entry, and τ is a temperature parameter. The message from \mathbf{Z}^{w+1} to \mathbf{Z}^w is aggregated as:

$$\hat{\mathbf{Z}}^w = \frac{\phi(\mathbf{Q}^w) \cdot (\phi(\mathbf{K}^{w+1})^\top \mathbf{V}^{w+1})}{\phi(\mathbf{Q}^w) \cdot \left(\sum_{j=1}^N \phi(\mathbf{k}_j^{w+1})^\top\right) + \epsilon}, \quad (5)$$

where the denominator ensures row-wise normalization, and ϵ is a small constant added for numerical stability.

In addition to updating node representations, we also extract an inter-hypergraph affinity matrix for each transition to capture temporal structural cues:

$$\mathbf{A}_{\text{inter}}^{w \rightarrow w+1} = \text{softmax}\left(\frac{\phi(\mathbf{Q}^w) \cdot \phi(\mathbf{K}^{w+1})^\top}{\sqrt{d_h}}\right), \quad (6)$$

which reflects node-level influence across adjacent time windows and provides a differentiable structure for modeling inter-hypergraph transitions. In summary, this module yields a sequence of interactive features $\{\hat{\mathbf{Z}}^w\}_{w=1}^{W-1}$, which capture dynamic high-order dependencies within and across time windows. Meanwhile, it also produces the corresponding inter-hypergraph affinity matrices $\{\mathbf{A}_{\text{inter}}^{w \rightarrow w+1}\}_{w=1}^{W-1}$.

Region-Wise Functional Hypergraph Learning

Region-Wise Functional Hypergraph Construction Considering that brain activity involves both temporal dynamics and region-specific dependencies, where functional patterns are often similar within the same anatomical region, we further propose a region-wise functional hypergraph learning strategy to capture these localized interactions. Specifically, we divide the whole-brain fMRI time series for a subject into eight macro-regions according to the AAL atlas (Tzourio-Mazoyer et al. 2002; Cui et al. 2023b): the frontal region (FR), central region (CR), limbic region (LR), temporal region (TR), occipital region (OR), parietal region (PR), subcortical region (SR), and cerebellar region (CER). Accordingly, the global fMRI time series $\mathbf{X} \in \mathbb{R}^{T \times N}$ is decomposed into a collection of regional sub-series $\{\mathbf{X}^r\}$ ($r \in \{\text{FR, CR, LR, TR, OR, PR, SR, CER}\}$), where $\mathbf{X}^r \in \mathbb{R}^{T \times N_r}$ denotes ROI time series in region r . For each region, a functional sub-hypergraph is constructed to capture latent high-order dependencies among its ROIs. In particular, each ROI time series $\mathbf{x}_i^r \in \mathbb{R}^T$ is represented as a sparse linear combination of remaining ROIs within the same region:

$$\min_{\beta_i^r} \frac{1}{2} \left\| \mathbf{x}_i^r - \hat{\mathbf{X}}^r \beta_i^r \right\|_2^2 + \lambda_2 \|\beta_i^r\|_1, \quad (7)$$

where $\hat{\mathbf{X}}^r \in \mathbb{R}^{T \times (N_r - 1)}$ is obtained by removing \mathbf{x}_i^r from \mathbf{X}^r , and $\beta_i^r \in \mathbb{R}^{N_r - 1}$ is a sparse coefficient vector that reflects the contribution of other ROIs to the i -th ROI. The regularization parameter λ_2 controls the sparsity level. Larger values in β_i^r indicate stronger functional associations. By applying sparse representation to each ROI, we obtain the interactive relationships between the centroid ROI and other ROIs in each region, which are regarded as hyperedges. Repeating this process for all ROIs results in \mathcal{E}^r hyperedges. These define a region-specific hypergraph $\mathcal{H}^r = (\mathcal{V}^r, \mathcal{E}^r)$, where \mathcal{V}^r is the vertex set of ROIs and \mathcal{E}^r is the set of hyperedges in region r . The incidence matrix $\mathbf{H}^r \in \{0, 1\}$ is defined $\mathbf{H}^r(i, e) = 1$ if ROI i participates in hyperedge e .

After constructing the hypergraph Laplacian-inspired adjacency matrix to model local high-order dependencies:

$$\mathbf{A}^r = \mathbf{H}^r \mathbf{W}_e \mathbf{D}_e^{-1} (\mathbf{H}^r)^\top, \quad (8)$$

where \mathbf{W}_e is a diagonal hyperedge weight matrix and \mathbf{D}_e is the diagonal hyperedge degree matrix. We then apply HyConv with DropMax regularization to model local high-order dependencies among ROIs. The region-level hypergraph convolution is formulated as:

$$\mathbf{Z}^r = \text{DropMax}(\sigma(\mathbf{A}^r \sigma(\mathbf{A}^r \mathbf{X}^r \Theta_1) \Theta_2)), \quad (9)$$

where Θ_1, Θ_2 are learnable projection matrices and $\sigma(\cdot)$ denotes a non-linear activation (e.g., ReLU). This intra-hypergraph encoding is independently applied to each

of the eight anatomical regions. The resulting region-specific features \mathbf{Z}^r are concatenated along the ROI dimension to obtain the unified representation $\mathbf{Z}_{\text{fused}} = \text{Concat}(\mathbf{Z}^{\text{FR}}, \mathbf{Z}^{\text{CR}}, \dots, \mathbf{Z}^{\text{CER}})$.

Inter-Region Topological Interactions Next, we enhance inter-region topological interactions using a spectral-aware diffusion mechanism. Specifically, we define a fused region adjacency matrix $\mathbf{A}_{\text{fused}} = \mathbf{A}_{\text{full}} + \mathbf{A}_{\text{region}}$, which integrates both global and region-specific high-order functional topologies. Here, $\mathbf{A}_{\text{region}} = \text{BlockDiag}(\mathbf{A}^{\text{FR}}, \dots, \mathbf{A}^{\text{CER}})$ is a block-diagonal functional connectivity matrix constructed from region-wise hypergraphs, while \mathbf{A}_{full} is the global functional connectivity matrix computed using Pearson correlation across all time series, with ROIs arranged in the same order as in $\mathbf{A}_{\text{region}}$. We then perform eigendecomposition of the normalized Laplacian:

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A}_{\text{fused}} \mathbf{D}^{-1/2} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top, \quad (10)$$

where \mathbf{U} contains the eigenvectors and $\mathbf{\Lambda}$ holds the eigenvalues. Based on this spectrum, the heat diffusion kernel is defined as:

$$\mathbf{K}_t = \mathbf{U} \exp(-t\mathbf{\Lambda}) \mathbf{U}^\top, \quad (11)$$

where $t \in \mathbb{R}^{C \times N}$ controls the diffusion scale, is class-specific and learnable. The operation $\mathbf{K}_t \mathbf{Z}^r$ produces a topologically smoothed representation, which is concatenated to obtain $\mathbf{Z}_{\text{cat}} = \text{Concat}(\mathbf{Z}_{\text{fused}}, \mathbf{K}_t \mathbf{Z}^r)$. Finally, a two-layer HyConv is used to integrate the concatenated features and perform global hypergraph learning:

$$\mathbf{Z}_{\text{region}} = \sigma(\hat{\mathbf{A}} \sigma(\hat{\mathbf{A}} \mathbf{Z}_{\text{cat}} \Theta_1) \Theta_2), \quad (12)$$

where $\mathbf{Z}_{\text{region}}$ denotes the fused region-wise hypergraph features, $\hat{\mathbf{A}} = \tilde{\mathbf{D}}^{-1/2} (\mathbf{A}_{\text{fused}} + \mathbf{I}) \tilde{\mathbf{D}}^{-1/2}$ is the normalized adjacency matrix with self-loops, and $\tilde{\mathbf{D}}$ is its degree matrix.

Topology-Enhanced Selective State-Space Propagation

After analyzing brain activity patterns, effectively integrating discriminative features is critical for robust brain disorder diagnosis. Therefore, we propose a topology-enhanced selective state-space propagation network to effectively integrate the aforementioned temporally-evolving and region-wise hypergraph features. Specifically, given the temporally-evolved representations $\{\hat{\mathbf{Z}}^w\}_{w=1}^{W-1}$ and the fused region-wise hypergraph features $\mathbf{Z}_{\text{region}}$, we first construct a unified dynamic input sequence: $\mathbf{F} = \text{Concat}(\mathbf{Z}_{\text{region}}, \hat{\mathbf{Z}}^1, \dots, \hat{\mathbf{Z}}^{W-1})$, which serves as input to a hypergraph-aware selective state-space system. We initialize the modeling with a linear continuous-time state-space formulation:

$$h'(t) = \mathcal{A}h(t) + \mathcal{B}x(t), \quad y(t) = \mathcal{C}^\top h(t), \quad (13)$$

where \mathcal{A} is a fixed diagonal matrix, initialized with non-positive real numbers, ensuring numerical stability in the matrix exponential. \mathcal{B}, \mathcal{C} are input and output projection matrices, respectively. To adapt this formulation to discrete se-

quences over brain nodes, we apply zero-order hold discretization with a learnable timescale parameter Δ :

$$\bar{\mathbf{A}} = \exp(\Delta \cdot \mathbf{A}), \quad (14)$$

$$\bar{\mathbf{B}} = (\exp(\Delta \cdot \mathbf{A}) - \mathbf{I}) (\Delta \cdot \mathbf{A})^{-1} (\Delta \cdot \mathbf{B}). \quad (15)$$

To incorporate both anatomical modularity and temporal transition structures, we define a composite topological operator:

$$\mathbf{G} = \mathbf{A}_{\text{fused}} + \sum_{w=1}^{W-1} \mathbf{A}_{\text{inter}}^{w \rightarrow w+1}. \quad (16)$$

where $\mathbf{A}_{\text{inter}}^{w \rightarrow w+1}$ denotes the inter-hypergraph affinity matrix that models semantic transitions across adjacent hypergraphs, $\mathbf{A}_{\text{fused}}$ represents the fused region adjacency matrix that encodes anatomical modularity. Before integrating them, we align the brain regions of the two topologies and apply linear layers to dynamically rescale and balance the contributions of the two topological structures.

Each brain node i is associated with a latent state \mathbf{L}_i that evolves under a topology-informed dynamic process:

$$\mathbf{L}_i = \bar{\mathbf{A}}_i \mathbf{L}_{i-1} + \bar{\mathbf{B}}_i \mathbf{F}_i, \quad (17)$$

$$\mathcal{F}_i = \mathbf{C}_i \mathbf{L}_i^\top, \quad (18)$$

where the node-specific parameters $\bar{\mathbf{A}}_i$, $\bar{\mathbf{B}}_i$, and \mathbf{C}_i are modulated by the corresponding topological structure \mathbf{G}_i :

$$\mathbf{B}_i = \mathbf{G}_i \cdot \text{Linear}(\mathbf{F}_i), \quad (19)$$

$$\mathbf{C}_i = \mathbf{G}_i \cdot \text{Linear}(\mathbf{F}_i), \quad (20)$$

$$\Delta_i = \text{softplus}(\text{Linear}(\mathbf{F}_i)). \quad (21)$$

The final latent outputs $\mathcal{F}_{\text{out}} = [\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_N]$ are passed through two fully connected layers followed by a softmax function to produce classification probabilities.

Experiment

Datasets

We evaluate our method on four datasets, all based on rs-fMRI modality, including three public datasets and one clinical epilepsy dataset. **ABIDE-I**: 1035 subjects (530 NC, 505 ASD) were selected from the Preprocessed Connectomes Project after resampling. **ADNI-2**: 640 samples (219 NC, 257 EMCI, 164 LMCI) were retained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database. **REST-meta-MDD**: Following (Chen et al. 2022; Yan et al. 2019), 2034 subjects (917 NC, 1117 MDD) were included. **Epilepsy**: The rs-fMRI data were acquired from Xuanwu Hospital, including 23 patients with temporal lobe epilepsy (TLE) and 26 patients diagnosed with frontal lobe epilepsy (FLE).

Implementation Details and Evaluation Metrics

The experiments are performed on the Intel Core i9-13900KF 3.0GHz CPU and NVIDIA RTX 3090 Ti GPU. The software environment includes Matlab R2021a, Python 3.10.11, Pytorch 2.0.1, and CUDA 11.8. We train HyperDiag¹ by setting the batch size as 16, the learning rate as

¹Code: <https://github.com/mylbuaa/HyperDiag>.

Method	ACC	SEN	SPE
ST-Transformer (Deng et al. 2022)	71.0	72.0	70.0
MAHGCN (Liu et al. 2023a)	72.7	69.8	75.4
DG-DMSGCN (Cui et al. 2023a)	72.0	73.3	70.5
STCAL (Liu et al. 2023b)	73.0	79.8	65.9
MGCA-RAFFNet (Ma et al. 2024)	75.1	76.0	74.1
Ours	76.8	77.2	76.4

where bold fonts indicate the best performance.

Table 1: Performance comparison of different methods on the ABIDE-I dataset (NC vs. ASD).

0.001. The model is optimized by using the Adam optimizer and the cross-entropy loss function. To ensure reliable evaluation results, we employ a nested 10-fold cross-validation strategy (Li et al. 2022). We repeat the nested scheme 10 times by changing the sample division to obtain reliable evaluation results. Accuracy (ACC), sensitivity (SEN) and specificity (SPE) are adopted for performance evaluation.

Comparison with State-of-the-art Methods

We conduct comprehensive comparisons with recent state-of-the-art methods on four brain disorder datasets. The results are summarized in Tables 1–4. As shown in Table 1, our model achieves the highest ACC of 76.8% on the ABIDE-I dataset, outperforming existing state-of-the-art methods. On the ADNI-2 dataset (Table 2), our method achieves an ACC of 84.7%, demonstrating its strength in multi-class diagnosis. Table 3 reports an ACC of 67.0% on the REST-meta-MDD dataset, with at least a 1.9% gain over prior methods, demonstrating superior generalization to depressive heterogeneity. As shown in Table 4, our model also achieves the highest ACC of 73.5% on the Epilepsy dataset, with an ACC gain of at least 2.1% compared to previous methods. We employ two-sample t-tests, and the results indicate that the differences between our model and the compared methods are statistically significant (p-values < 0.05 and < 0.001). These consistent improvements across all datasets highlight the robustness and generalizability of our method. Notably, the superior performance stems from task-driven innovations specifically tailored to neuroimaging analysis. Each component is purposefully designed to address a specific issue and is tightly integrated into a cohesive architecture, thereby achieving strong performance and facilitating the discovery of clinically meaningful biomarkers.

Ablation Studies

We conduct ablation studies to evaluate the contribution of each component in our proposed method. As shown in Table 5, the ablation analysis includes the following settings:

- Removing the temporally-evolving hypergraph message passing (w/o TEHMP), including its variants without intra-window (w/o TEHMP_{intra}) and inter-window message passing (w/o TEHMP_{inter}).
- Removing the region-wise functional hypergraph learning (w/o RWFHL), along with its variants that exclude

Method	ACC	SEN _N	SEN _E	SEN _L
MMTGCN † (Yao et al. 2021)	80.3	77.9	83.3	79.1
STA-BiGRU (Huang et al. 2022)	81.5	-	-	-
PFC-STAA (Cui et al. 2023b)	83.3	80.5	89.3	78.3
STW-HCN (Liu et al. 2024)	83.0	82.8	83.9	81.8
ACI-FBN * (Zhang et al. 2025)	83.6	84.0	84.4	81.7
Ours	84.7	84.9	85.2	83.5

where † means the result is from MGCA-RAFFNet (Ma et al. 2024) and * means the method is reproduced by ourselves.

Table 2: Performance comparison of different methods on the ADNI-2 dataset (NC vs. EMCI vs. LMCI).

Method	ACC	SEN	SPE
MDGL (Ma et al. 2024)	58.0	65.0	-
BrainGSLs † (Wen et al. 2023)	62.5	62.5	62.2
BrainMass † (Yang et al. 2024)	64.0	65.7	61.4
HGFM (Han et al. 2025)	65.1	67.5	61.5
Ours	67.0	67.9	66.2

where † means the result is from HGFM (Han et al. 2025).

Table 3: Performance comparison of different methods on the REST-meta-MDD dataset (NC vs. MDD).

Method	ACC	SEN	SPE
PFC-STAA * (Cui et al. 2023b)	67.4	65.2	69.2
MGCA-RAFFNet * (Ma et al. 2024)	69.4	69.6	69.2
ACI-FBN * (Zhang et al. 2025)	71.4	69.6	73.1
Ours	73.5	73.9	73.1

where * means the method is reproduced by ourselves.

Table 4: Performance comparison of different methods on the Epilepsy dataset (TLE vs. FLE).

intra-region (w/o RWFHL_{intra}) and inter-region hypergraph learning (w/o RWFHL_{inter}).

- Removing the topology-enhanced selective state-space propagation (w/o TESSP).

Effectiveness of Temporally-Evolving Hypergraph Message Passing (TEHMP). From Table 5, removing TEHMP (w/o TEHMP) leads to at least a 2.8% drop in ACC across all four datasets, indicating its effectiveness in capturing high-order dynamic dependencies both within and across time windows. Further ablation reveals that removing intra-window message passing (w/o TEHMP_{intra}) causes an additional decrease in ACC, highlighting the importance of intra-window interactions. Similarly, removing inter-window message passing (w/o TEHMP_{inter}) results in a smaller performance decline, underscoring the role of temporal relationships across windows.

Effectiveness of Region-Wise Functional Hypergraph Learning (RWFHL). Table 5 shows that removing RWFHL leads to ACC drops of 3.7%, 4.1%, 3.4%, and 8.2% on four datasets, respectively. Removing intra-region learning (w/o RWFHL_{intra}) results in a larger performance decline than removing inter-region learning (w/o RWFHL_{inter}),




Figure 2: Comparison of different fusion strategies.

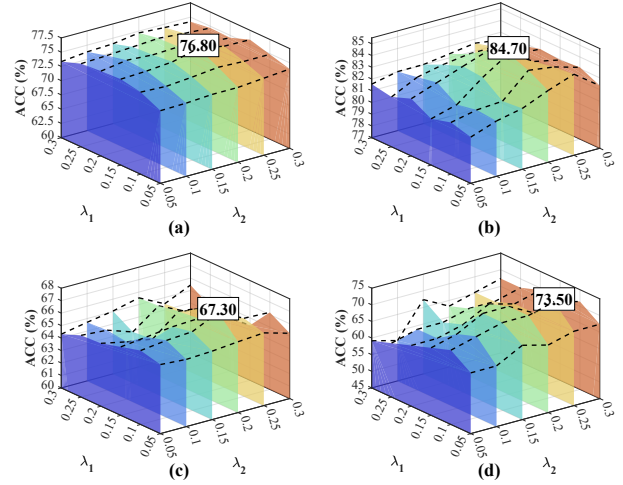


Figure 3: The influence of the hyper-parameters on classification performance. (a) NC vs. ASD, (b) NC vs. EMCI vs. LMCI, (c) NC vs. MDD, and (d) TLE vs. FLE.

highlighting the importance of region-specific hypergraphs for modeling functional heterogeneity, while inter-region modeling enhances cross-regional integration.

Effectiveness of Topology-Enhanced Selective State-Space Propagation (TESSP). TESSP is designed to aggregate temporally-evolving hypergraph features and topology-enhanced region-wise hypergraph features. As shown in Table 5, removing TESSP consistently leads to performance drops of 1.2%, 1.2%, 1.1%, and 2.1% across the four datasets, highlighting its crucial role in aggregating discriminative features. In addition, from Fig. 2, we compare TESSP with two alternative fusion strategies: feature concatenation and feature selective aggregation. The results show that TESSP outperforms these alternatives.

Hyperparameter Sensitivity Analysis

The hyperparameters λ_1 and λ_2 control the sparsity of the temporally-evolving and region-wise hypergraphs, respectively. We vary them from 0.05 to 0.3 with a step of 0.05, and report the ACC values in Fig. 3. The optimal settings are $\lambda_1 = 0.15$, $\lambda_2 = 0.2$ for NC vs. ASD and NC vs. EMCI vs. LMCI; $\lambda_1 = 0.1$, $\lambda_2 = 0.2$ for NC vs. MDD; and $\lambda_1 = 0.1$, $\lambda_2 = 0.25$ for TLE vs. FLE. Both overly large and small values degrade performance by over-penalizing the topology or retaining redundant connections.

Method	NC vs. ASD			NC vs. MDD			NC vs. EMCI vs. LMCI				TLE vs. FLE		
	ACC	SEN	SPE	ACC	SEN	SPE	ACC	SEN _N	SEN _E	SEN _L	ACC	SEN	SPE
w/o TEHMP	74.0	74.5	73.5	64.3	65.0	63.7	80.9	81.7	80.9	79.9	61.2	60.9	61.5
w/o TEHMP _{intra}	75.1	75.3	74.9	65.9	66.5	65.4	82.2	82.2	82.1	82.3	63.3	69.6	57.7
w/o TEHMP _{inter}	75.7	76.0	75.3	65.2	65.8	64.8	83.4	84.5	84.1	81.1	63.3	65.2	61.5
w/o RWFHL	73.1	73.8	72.5	63.2	63.9	62.7	81.3	81.3	81.7	80.5	65.3	69.6	61.5
w/o RWFHL _{intra}	74.2	74.7	73.7	64.4	64.8	64.0	82.5	82.2	82.5	82.9	67.4	69.6	65.4
w/o RWFHL _{inter}	74.9	75.5	74.3	64.9	65.4	64.5	82.8	83.1	82.9	82.3	69.4	73.9	65.4
w/o TESSP	75.6	76.2	74.9	66.1	66.0	66.3	83.6	84.5	84.1	81.7	71.4	73.9	69.2
Ours (complete)	76.8	77.2	76.4	67.3	68.7	66.2	84.7	84.9	85.2	83.5	73.5	73.9	73.1

Table 5: Experimental results of the ablation study on different brain disease tasks.

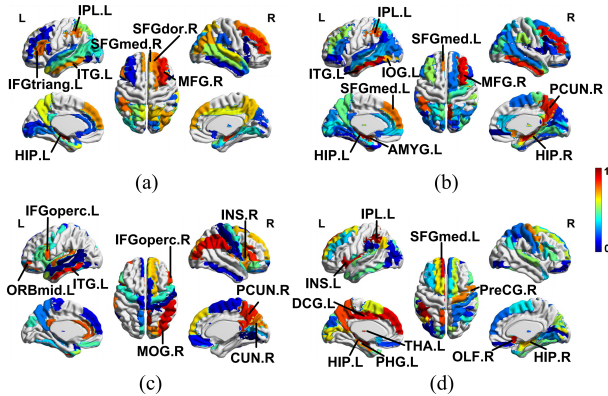


Figure 4: Most discriminative brain regions across different diagnostic tasks: (a) NC vs. ASD, (b) NC vs. EMCI vs. LMCI, (c) NC vs. MDD, and (d) TLE vs. FLE.

Discriminative Brain Regions and Connections

Discriminative Brain Regions. To identify potential neurobiomarkers, we compute ROI importance weights and project them onto the AAL template. Each diagnostic task exhibits distinct discriminative regions. For instance, in Fig. 4 (a), MFG.R, SFGdor.R, IPL.L, and IFGtriang.L are highly weighted, aligning with ASD-related impairments in social cognition and executive function (Wen et al. 2023). In Fig. 4 (b), regions such as HIP.R and PCUN.R, implicated in memory and cognitive decline (Raine and Rao 2022). In Fig. 4 (c), discriminative areas include IFGoperc.L and INS.R, associated with affective processing. In Fig. 4 (d), prominent regions like HIP.R and THA.L reflect abnormal motor and memory pathways.

Discriminative Brain Connections. To further investigate abnormal connectivity patterns, we perform interpretability analysis on the confidence matrix. As shown in Fig. 5, each diagnostic task exhibits distinct discriminative connections, where chords with the same color indicate membership in the same hyperedge. For example, in Fig. 5(a), temporal-frontal hyperedges (e.g., TPOmid.L–ITG.L–IFGtriang.L) suggest social-executive dysfunction. Fig. 5(b) highlights parietal-limbic-frontal interactions (e.g., IPL.L–OLF.R–PHG.L) indicative of MCI-related memory decline. In Fig. 5(c), frontal-limbic couplings dominate, linked to emotional dysregulation. For Fig. 5(d),

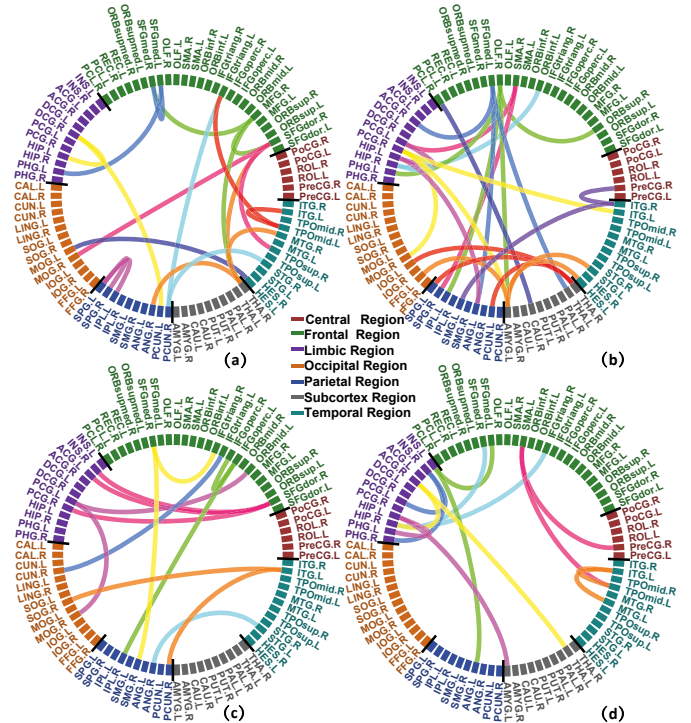


Figure 5: Most discriminative brain connections across different diagnostic tasks: (a) NC vs. ASD, (b) NC vs. EMCI vs. LMCI, (c) NC vs. MDD, and (d) TLE vs. FLE.

discriminative hyperedges include HIP.L–PHG.L–AMYG.L and PHG.L–INS.R–PCG.L, which reflect both mesial temporal and motor-related abnormalities. These findings not only corroborate prior evidence but also reveal previously underexplored abnormal connections, demonstrating the model’s ability to capture high-order interactions.

Conclusion

We propose HyperDiag, a novel and robust framework that integrates temporally-evolving and region-wise hypergraph modeling with topology-enhanced selective state-space propagation. Experiments on four datasets demonstrate superior performance and the ability to identify meaningful biomarkers.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 623B2011, Grant 62325301, and Grant U24B20186; in part by the Beijing Natural Science Foundation under Grant Z220017, Grant L256008; in part by the National Key Research and Development Program of China under Grant 2023YFC2416600; and in part by Natural Science Foundation Key Project of Zhejiang Province, China under Grant LZ23F030001. It was also supported by the Academic Excellence Foundation of BUAA for PhD students. In part by the Chongqing Municipal Health Commission, China under Grant 2025GGXM005.

We gratefully acknowledge the DIRECT Consortium and the International Big-Data Center for Depression Research, Institute of Psychology, Chinese Academy of Sciences, for sharing the REST-meta-MDD dataset. The DIRECT Consortium is the data-generating consortium and the group author of the REST-meta-MDD project in its original publication. We also thank the following team members of XuanWu Hospital Functional neurosurgery department Epilepsy Database (XWH-FED) for their support in data collection and management: Runshi Gao, Xiumei Wang, Tao Yu, and Liang Qiao.

References

- Ahamed, M. A.; and Cheng, Q. 2024. MambaTab: A plug-and-play model for learning tabular data. In *2024 IEEE 7th International Conference on Multimedia Information Processing and Retrieval (MIPR)*, 369–375. IEEE.
- Béna, G.; and Goodman, D. F. 2025. Dynamics of specialization in neural modules under resource constraints. *Nature Communications*, 16(1): 187.
- Chen, X.; Lu, B.; Li, H. X.; Li, X. Y.; Wang, Y. W.; Castellanos, F. X.; Cao, L. P.; Chen, N. X.; Chen, W.; Cheng, Y. Q.; et al. 2022. The DIRECT consortium and the REST-meta-MDD project: towards neuroimaging biomarkers of major depressive disorder. *Psychoradiology*, 2(1): 32–42.
- Chen, Y.; Dang, M.; and Zhang, Z. 2021. Brain mechanisms underlying neuropsychiatric symptoms in Alzheimer’s disease: a systematic review of symptom-general and-specific lesion patterns. *Molecular Neurodegeneration*, 16(1): 38.
- Cui, W.; Du, J.; Sun, M.; Zhu, S.; Zhao, S.; Peng, Z.; Tan, L.; and Li, Y. 2023a. Dynamic multi-site graph convolutional network for autism spectrum disorder identification. *Computers in Biology and Medicine*, 157: 106749.
- Cui, W.; Ma, Y.; Ren, J.; Liu, J.; Ma, G.; Liu, H.; and Li, Y. 2023b. Personalized functional connectivity based spatio-temporal aggregated attention network for MCI identification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31: 2257–2267.
- Deng, X.; Zhang, J.; Liu, R.; and Liu, K. 2022. Classifying ASD based on time-series fMRI using spatial-temporal transformer. *Computers in biology and medicine*, 151: 106320.
- Du, Y.; Niu, J.; and Calhoun, V. D. 2021. A New Hypergraph Clustering Method For Exploring Transdiagnostic Biotypes In Mental Illnesses: Application To Schizophrenia And Psychotic Bipolar Disorder. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 971–974.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Han, X.; Lei, M.; and Li, J. 2025. Hypergraph-based semantic and topological self-supervised learning for brain disease diagnosis. *Pattern Recognition*, 111921.
- Han, X.; Xue, R.; Du, S.; and Gao, Y. 2024. Inter-intra high-order brain network for ASD diagnosis via functional MRIs. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 216–226. Springer.
- Han, X.; Xue, R.; Feng, J.; Feng, Y.; Du, S.; Shi, J.; and Gao, Y. 2025. Hypergraph Foundation Model for Brain Disease Diagnosis. *IEEE Transactions on Neural Networks and Learning Systems*, 1–15.
- Huang, H.; Liu, Q.; Jiang, Y.; Yang, Q.; Zhu, X.; and Li, Y. 2022. Deep Spatio-Temporal Attention-Based Recurrent Network From Dynamic Adaptive Functional Connectivity for MCI Identification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 30: 2600–2612.
- Jiao, L.; Ma, M.; He, P.; Geng, X.; Liu, X.; Liu, F.; Ma, W.; Yang, S.; Hou, B.; and Tang, X. 2024. Brain-inspired learning, perception, and cognition: A comprehensive review. *IEEE Transactions on Neural Networks and Learning Systems*.
- Li, Y.; Liu, J.; Jiang, Y.; Liu, Y.; and Lei, B. 2022. Virtual Adversarial Training-Based Deep Feature Aggregation Network From Dynamic Effective Connectivity for MCI Identification. *IEEE Transactions on Medical Imaging*, 41(1): 237–251.
- Liu, J.; Cui, W.; Chen, Y.; Ma, Y.; Dong, Q.; Cai, R.; Li, Y.; and Hu, B. 2024. Deep Fusion of Multi-Template Using Spatio-Temporal Weighted Multi-Hypergraph Convolutional Networks for Brain Disease Analysis. *IEEE Transactions on Medical Imaging*, 43(2): 860–873.
- Liu, M.; Zhang, H.; Shi, F.; and Shen, D. 2023a. Hierarchical graph convolutional network built by multiscale atlases for brain disorder diagnosis using functional connectivity. *IEEE Transactions on Neural Networks and Learning Systems*, 35(11): 15182–15194.
- Liu, R.; Huang, Z.-A.; Hu, Y.; Zhu, Z.; Wong, K.-C.; and Tan, K. C. 2023b. 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.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; Jiao, J.; and Liu, Y. 2025. Vmamba: Visual state space model. *Advances in neural information processing systems*, 37: 103031–103063.
- Lord, C.; Elsabbagh, M.; Baird, G.; and Veenstra-Vanderweele, J. 2018. Autism spectrum disorder. *The lancet*, 392(10146): 508–520.

- Ma, Y.; Cui, W.; Liu, J.; Guo, Y.; Chen, H.; and Li, Y. 2024. A Multi-Graph Cross-Attention-Based Region-Aware Feature Fusion Network Using Multi-Template for Brain Disorder Diagnosis. *IEEE Transactions on Medical Imaging*, 43(3): 1045–1059.
- Makhlouf, A. T.; Drew, W.; Stubbs, J. L.; Taylor, J. J.; Liloia, D.; Grafman, J.; Silbersweig, D.; Fox, M. D.; and Siddiqi, S. H. 2025. Heterogeneous patterns of brain atrophy in schizophrenia localize to a common brain network. *Nature Mental Health*, 3(1): 19–30.
- Raine, P. J.; and Rao, H. 2022. Volume, density, and thickness brain abnormalities in mild cognitive impairment: an ALE meta-analysis controlling for age and education. *Brain imaging and behavior*, 16(5): 2335–2352.
- Sarem, M.; Jurdi, T.; Alshlawy, L.; and Massrie, E. 2024. Improving long text classification based on Selective State Space model (Mamba). In *2024 IEEE 17th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc)*, 32–38. IEEE.
- Shao, Y.; Li, Y.; and Wu, B. 2025. A Directional Attention Fusion and Multi-Head Spatial-Channel Attention Network for Facial Expression Recognition. *Journal of Machine Learning and Information Security*, 1(1): 7.
- Shi, G.; Zhu, Y.; Liu, W.; Yao, Q.; and Li, X. 2025. Heterogeneous Graph-Based Multimodal Brain Network Learning. *IEEE Transactions on Knowledge and Data Engineering*, 37(8): 4664–4676.
- Tzourio-Mazoyer, N.; Landeau, B.; Papathanassiou, D.; Crivello, F.; Etard, O.; Delcroix, N.; Mazoyer, B.; and Joliot, M. 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage*, 15(1): 273–289.
- Van Overwalle, F. 2024. Social and emotional learning in the cerebellum. *Nature Reviews Neuroscience*, 1–16.
- Wen, G.; Cao, P.; Liu, L.; Hao, M.; Liu, S.; Zheng, J.; Yang, J.; Zaiane, O. R.; and Wang, F. 2025. Heterogeneous Graph Representation Learning Framework for Resting-State Functional Connectivity Analysis. *IEEE Transactions on Medical Imaging*, 44(3): 1581–1595.
- Wen, G.; Cao, P.; Liu, L.; Yang, J.; Zhang, X.; Wang, F.; and Zaiane, O. R. 2023. Graph self-supervised learning with application to brain networks analysis. *IEEE Journal of Biomedical and Health Informatics*, 27(8): 4154–4165.
- Xu, Z.; Ma, C.; Wang, C.; Guo, F.; Zheng, M.; Fang, P.; and Zhu, Y. 2025. Dynamic changes in brain function during sleep deprivation: Increased occurrence of non-stationary states indicates the extent of cognitive impairment. *NeuroImage*, 309: 121099.
- Yan, C. G.; Chen, X.; Li, L.; Castellanos, F. X.; Bai, T. J.; Bo, Q. J.; Cao, J.; Chen, G. M.; Chen, N. X.; Chen, W.; et al. 2019. Reduced default mode network functional connectivity in patients with recurrent major depressive disorder. *Proceedings of the National Academy of Sciences*, 116(18): 9078–9083.
- Yang, P.; Zhou, F.; Ni, D.; Xu, Y.; Chen, S.; Wang, T.; and Lei, B. 2019. Fused sparse network learning for longitudinal analysis of mild cognitive impairment. *IEEE transactions on cybernetics*, 51(1): 233–246.
- Yang, Y.; Ye, C.; Su, G.; Zhang, Z.; Chang, Z.; Chen, H.; Chan, P.; Yu, Y.; and Ma, T. 2024. Brainmass: Advancing brain network analysis for diagnosis with large-scale self-supervised learning. *IEEE transactions on medical imaging*, 43(11): 4004–4016.
- Yao, D.; Sui, J.; Wang, M.; Yang, E.; Jiaerken, Y.; Luo, N.; Yap, P.-T.; Liu, M.; and Shen, D. 2021. A mutual multi-scale triplet graph convolutional network for classification of brain disorders using functional or structural connectivity. *IEEE transactions on medical imaging*, 40(4): 1279–1289.
- Yu, W.; and Wang, X. 2025. Mambaout: Do we really need mamba for vision? In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 4484–4496.
- Zhang, H.; Chen, C.; Mei, L.; Liu, Q.; and Mao, J. 2024. Mamba Retriever: Utilizing Mamba for Effective and Efficient Dense Retrieval. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 4268–4272.
- Zhang, J.; Wu, X.; Tang, X.; Zhou, L.; Wang, L.; Wu, W.; and Shen, D. 2025. Asynchronous Functional Brain Network Construction With Spatiotemporal Transformer for MCI Classification. *IEEE Transactions on Medical Imaging*, 44(3): 1168–1180.
- Zhou, J.; Jie, B.; Wang, Z.; Zhang, Z.; Bian, W.; Yang, Y.; Li, H.; Sun, F.; and Liu, M. 2025. DCLNet: Double Collaborative Learning Network on Stationary-Dynamic Functional Brain Network for Brain Disease Classification. *IEEE Transactions on Image Processing*, 34: 4026–4039.
- Zhu, Q.; Li, S.; Meng, X.; Xu, Q.; Zhang, Z.; Shao, W.; and Zhang, D. 2024. Spatio-Temporal Graph Hubness Propagation Model for Dynamic Brain Network Classification. *IEEE Transactions on Medical Imaging*, 43(6): 2381–2394.
- Zhu, Y.; Zhu, X.; Kim, M.; Yan, J.; Kaufer, D.; and Wu, G. 2019. Dynamic Hyper-Graph Inference Framework for Computer-Assisted Diagnosis of Neurodegenerative Diseases. *IEEE Transactions on Medical Imaging*, 38(2): 608–616.