Learning Geometry-Aware Bundles with Spectral Heat Kernel Diffusion for fMRI-based Brain Disease Diagnosis

Feiyu Yin[1], Jie Xu[1], Qi Zhang[2], Yipin Zhang[3], Junhu Fu[1], Ziqi Yu[2], Shengjie Zhang[2(✉)], and Jinhua Yu[1(✉)]

[1] College of Biomedical Engineering, Fudan University, Shanghai, 200433, China.
jhyu@fudan.edu.cn

[2] Department of Radiology, Ruijin Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai, 200025, China.
zsjxll@gmail.com

[3] Institute of Science and Technology for Brain-Inspired Intelligence, Fudan University, Shanghai, 200433, China

Abstract. Accurate diagnosis of neurodegenerative disorders using functional magnetic resonance imaging (fMRI) remains challenging due to the intricate functional connectivity patterns and inherent imaging noise. Conventional Graph Neural Networks (GNNs) typically project the brain into a shared global space, overlooking that different brain regions reside on heterogeneous functional manifolds. While existing attention or heterogeneous graph methods merely adjust edge weights, they fail to capture the complex geometric misalignment among regional manifolds. To address this limitation, we propose Vector Bundle Graph Contrastive Learning (VB-GCL), a novel dual-stream framework grounded in vector bundle theory for brain disease classification. Unlike standard GNNs, vector bundles provide a natural framework for coordinating locally distinct geometries. VB-GCL lifts the brain connectivity graph to a vector bundle, where each region of interest (ROI) is associated with a learnable orthogonal subspace, enabling the model to explicitly capture region-specific functional geometry. Based on this structure, we introduce a Bundle Laplacian that jointly encodes topological connectivity and local geometric orientation. Our framework performs multi-hop geometric diffusion on the bundle manifold for structural feature aggregation and classification, while simultaneously employing Bundle-Laplacian-driven spectral perturbations to generate contrastive views that encourage noise-robust representations. Extensive experiments on the ADNI and an in-house clinical dataset demonstrate that VB-GCL achieves state-of-the-art performance across multiple diagnostic tasks. Ablation studies and interpretability analysis further confirm the effectiveness of each component and reveal clinically meaningful pathological biomarkers, providing a promising tool for precision brain disease diagnosis.

Keywords: Brain Disease Diagnosis · fMRI Analysis · Graph Neural Network · Vector Bundle Theory

1 Introduction

With the accelerating aging of the global population, the accurate diagnosis of dementia, primarily Alzheimer's disease (AD) and vascular diseases represented by Moyamoya disease (MMD), has emerged as a critical clinical imperative [4,28,6]. Functional Magnetic Resonance Imaging (fMRI), with its capability to map dynamic blood-oxygen-level-dependent (BOLD) signals, serves as a pivotal tool for unraveling the functional connectome underlying cognitive decline [3,18]. In recent years, Graph Neural Networks (GNNs) have emerged as a powerful framework for fMRI analysis. By conceptualizing brain regions as nodes and functional associations as edges, GNNs have achieved state-of-the-art performance in disease classification tasks [24,7].

Despite their promising performance, most existing GNNs share a critical geometric limitation: they implicitly project the complex brain connectome into a single shared Euclidean feature space [2,11]. This homogeneous assumption overlooks the heterogeneity of brain regions, which may reside on distinct functional manifolds with region-specific local geometric properties [1,23]. Such heterogeneity implies that features from different region of interests (ROIs) are not necessarily expressed in a globally consistent coordinate system, leading to geometric misalignment during message passing. This type of coordinate misalignment is closely related to the connection Laplacian and its learnable counterparts in recent sheaf and bundle neural networks, which model message passing through transformations between ROI-specific local feature spaces. However, conventional GNNs lack such geometric alignment mechanisms and instead aggregate regional features directly in a shared Euclidean space. Forcing heterogeneous local manifolds to collapse into this shared space can induce representational distortion and feature entanglement. As a result, the learned representations become more vulnerable to the inherent noise of fMRI networks. To mitigate this vulnerability and extract noise-invariant features, Graph Contrastive Learning (GCL) is widely adopted [27,25]. However, conventional GCL strategies (e.g., random edge or node dropping) paradoxically violate the geometric nature of the brain from another dimension [29]. Since different brain regions possess independent local coordinate systems, their information exchange relies on the intact network topology to achieve smooth geometric alignment. Violent random edge dropping blindly severs the geometric coherence between adjacent local spaces, causing the generated contrastive views to deviate from the underlying manifold structure. Such geometrically blind augmentations fail to provide meaningful noise regularization; instead, they destroy the discriminative manifold invariants essential for generalization in small-sample clinical scenarios [26].

To address these limitations, we propose a novel Vector Bundle Graph Contrastive Learning (VB-GCL), a novel dual-stream framework grounded in vector bundle theory. Departing from the homogeneous assumption of a single shared global space, VB-GCL elevates the brain connectivity graph to a vector bundle manifold. By assigning a learnable independent orthogonal subspace to each ROI, this bundle structure explicitly captures local geometric heterogeneity and prevents representational distortion caused by forced dimensionality reduction.

To enable information exchange across these independent local coordinate systems, we employ the Bundle Laplacian for geometric feature alignment between adjacent ROIs. Furthermore, within the dual-stream contrastive network, to prevent heuristic data augmentations from corrupting the manifold topology, we design a spectral perturbation module based on the heat kernel. By modulating the diffusion time scales, this module yields contrastive views by smoothing features along the Bundle Laplacian at different scales, so that augmentations remain on the connectivity topology rather than randomly perturbing it. Through this unified architecture, Laplacian smoothing effectively filters out the intrinsic noise of fMRI signals, while preserving the structural invariants crucial for the generalization capability of brain networks.

The main contributions of this work are summarized as follows:

1. We lift the brain connectome to a vector bundle manifold via learnable orthogonal subspaces and the Bundle Laplacian, assigning each ROI an independent local geometry beyond the shared-space assumption of traditional GNNs.

2. We propose a dual-stream framework coupling bundle geometric diffusion with contrastive learning. Replacing heuristic edge dropping with Bundle Laplacian-driven spectral perturbations, we generate topologically coherent views that enhance robustness against fMRI noise.

3. Evaluations on Alzheimer's Disease Neuroimaging Initiative (ADNI) and an in-house dataset demonstrate that VB-GCL outperforms advanced methods. Furthermore, interpretability analysis localizes disease-specific ROIs that are statistically separable between groups.

2 Method

In this section, we detail the architecture of the proposed VB-GCL algorithm. As illustrated in Fig. 1, VB-GCL integrates three core modules alongside a downstream classification path: (1) **Data Input & Preprocessing**, which extracts ROI-wise time series to construct the initial brain connectivity graph; (2) **The Bundle Transformation**, which lifts the original graph space onto a bundle subspace by assigning local orthogonal coordinate systems; and (3) **Dual-Stream Contrastive Network**, which employs a heat kernel-based spectral perturbation module to generate weakly and strongly smoothed representations for robust contrastive learning. Ultimately, the representations are fed into the classification path for the final clinical prediction.

2.1 Vector Bundle and Bundle Laplacian Construction

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ be the fMRI-derived brain network with $N = |\mathcal{V}|$ nodes, and let $\mathbf{L}$ denote its original graph Laplacian matrix. We associate each node $i \in \mathcal{V}$ with a local orthogonal subspace $\mathcal{F}_i \in \mathbb{R}^d$, where d is the bundle dimension. We define a global orthogonal transport operator $\mathcal{O} = \mathrm{diag}(\mathbf{O}_1, \ldots, \mathbf{O}_N)$, where $\mathbf{O}_i$ is a learnable orthogonal matrix constrained by stabilized QR decomposition

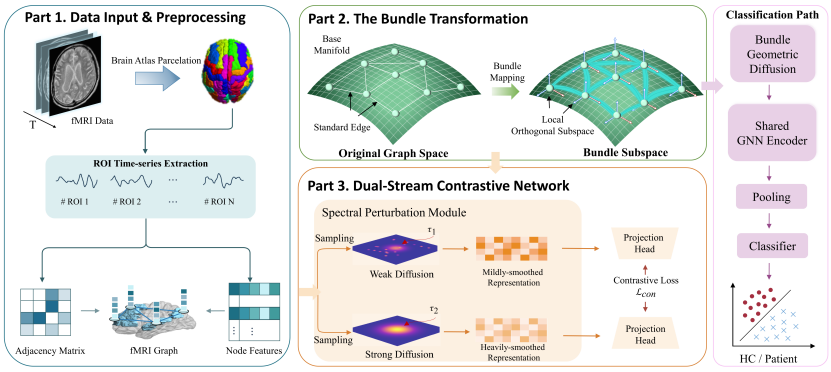

Fig. 1. Overall architecture of the proposed VB-GCL. The model constructs a vector bundle subspace to encode local brain network geometry. A shared GNN encoder is optimized jointly by a primary bundle diffusion classification path and an auxiliary spectral contrastive module, which leverages heat kernel perturbations to enforce multi-scale representation learning.

to align adjacent coordinate systems. To jointly encode topological connectivity and subspace geometry, we formulate the Bundle Laplacian $\mathbf{L}_{\mathcal{B}}$:

$$\mathbf{L}_{\mathcal{B}} = \mathcal{O}^{\top}(\mathbf{L} \otimes \mathbf{I}_d)\mathcal{O}, \tag{1}$$

where $\otimes$ denotes the Kronecker product and $\mathbf{I}_d$ is the identity matrix. For numerical stability, we apply symmetric normalization $\hat{\mathbf{L}}_{\mathcal{B}} = \mathbf{D}^{-1/2}\mathbf{L}_{\mathcal{B}}\mathbf{D}^{-1/2}$, with the degree matrix $\mathbf{D}_{ii} = \sum_j |(\mathbf{L}_{\mathcal{B}})_{ij}|$, where j indexes the adjacent nodes of i.

2.2 Geometry-Aware Bundle Diffusion

To capture the manifold geometry encoded in $\hat{\mathbf{L}}_{\mathcal{B}}$, the bundle diffusion layer employs a K-order polynomial filter to simulate physical geometric diffusion across multi-hop neighborhoods. The updated features $\mathbf{H}_{\mathcal{B}}^{(l)}$ at layer l are computed as:

$$\mathbf{H}_{\mathcal{B}}^{(l)} = \sigma\left(\mathrm{LN}\left(\sum_{k=0}^{K}(\hat{\mathbf{L}}_{\mathcal{B}}^{k}\mathbf{H}_{\mathcal{B}}^{(l-1)})\mathbf{W}_k + \mathbf{H}_{\mathcal{B}}^{(l-1)}\right)\right), \tag{2}$$

where $\mathbf{W}_k$ denotes layer-specific learnable projection weights, $\mathrm{LN}(\cdot)$ is Layer Normalization, and σ is the LeakyReLU activation. A residual connection is added to mitigate over-smoothing during deep message passing.

2.3 Dual-Stream Spectral View Generation and Contrastive Learning

To regularize representations against fMRI noise and individual variability, we propose a spectral contrastive augmentation strategy based on heat diffusion.

Avoiding computationally prohibitive matrix exponentials, we approximate the heat kernel via a first-order linear expansion. Specifically, the heat kernel $e^{-\tau_t \hat{\mathbf{L}}_{\mathcal{B}}}$ admits the first-order Taylor expansion $e^{-\tau_t \hat{\mathbf{L}}_{\mathcal{B}}} \approx I - \tau_t \hat{\mathbf{L}}_{\mathcal{B}}$ for small diffusion time τ_t, which yields the closed-form spectral view in Eq. (3). Two topology-preserving spectral views are generated using two different diffusion time scales, denoted by τ_1 and τ_2, respectively:

$$\mathbf{X}_{\text{spec}}^{(t)} \approx \left(\mathbf{I} - \tau_t \hat{\mathbf{L}}_{\mathcal{B}}\right) \mathbf{X}_{\mathcal{B}}, t \in \{1, 2\}. \tag{3}$$

Here $\mathbf{X}_{\mathcal{B}} = \mathbf{H}_{\mathcal{B}}^{(l)}$ denotes the bundle features from the diffusion path in Eq. (2), and $\mathbf{X}_{\text{spec}}^{(t)}$ are the two spectral views.

To prevent overfitting, we apply a random feature masking strategy to both generated views. These two spectral views are then processed by a shared backbone and a projection head to yield L_2-normalized embeddings z_m in a contrastive latent space. For a mini-batch of B graphs, let $I \equiv \{1, \ldots, 2B\}$ denote the indices of the concatenated augmented batch. For an anchor $m \in I$, we define $A(m) \equiv I \backslash \{m\}$ as the set of all other indices, and $P(m) \equiv \{p \in A(m) : y_p = y_m\}$ as the subset of positive samples sharing the same class label y_m. To explicitly maximize margin separability and noise robustness, we optimize the network using a supervised contrastive loss:

$$\mathcal{L}_{con} = \sum_{m \in I} \frac{-1}{|P(m)|} \sum_{p \in P(m)} \log \frac{\exp(z_m^\top z_p / \gamma)}{\sum_{a \in A(m)} \exp(z_m^\top z_a / \gamma)}, \tag{4}$$

where γ is a temperature scaling parameter.

3 Experiments and Results

Datasets. We utilized fMRI data from the public ADNI database, including 238 normal controls (NCs) and 213 patients, and an in-house cohort (98 NCs, 120 MMD patients). Both datasets provide demographics and scanner metadata. The in-house dataset was approved by the local Institutional Review Board (IRB), with all patient records fully anonymized. fMRI preprocessing via DPABI [21] included discarding the first ten volumes for magnetic equilibration, head motion correction, spatial normalization, and Gaussian smoothing. Finally, regional mean BOLD signals were extracted using the AAL90 atlas.

Implementation Details. Our method was implemented in PyTorch and trained on 8 NVIDIA V100 GPUs. All experiments were conducted using a standard 10-fold cross-validation strategy. We configured the model with a learning rate of 0.001 and a batch size of $B = 32$. Regarding the classification module, we employed a Multi-Layer Perceptron (MLP) with a 32-dimensional hidden layer. To stabilize training and mitigate overfitting, this classifier incorporates batch normalization, a LeakyReLU activation function, and a dropout layer with a rate of 0.7 before projecting to the final number of classes. To instantiate the joint

objective mentioned earlier, the total loss is formulated as $\mathcal{L}_{total} = \mathcal{L}_{CE} + \mathcal{L}_{con}$, where $\mathcal{L}_{CE}$ is the standard cross-entropy loss for the primary diagnostic task. The network is optimized using the Adam optimizer. Our code is publicly available at https://github.com/CodeMedYY/VB-GCL.

Table 1. Classification performance denoted by mean and standard deviation of the proposed method and all comparison methods on the ADNI and in-house dataset (%).

Method	ADNI				In-house			
	Acc.	Sen.	Spe.	AUC	Acc.	Sen.	Spe.	AUC
GIN	$73.39_{3.62}$	$73.94_{2.13}$	$72.82_{5.51}$	$72.92_{4.61}$	$67.88_{3.31}$	$68.34_{6.46}$	$67.33_{6.03}$	$67.77_{3.71}$
BNT	$76.05_{2.63}$	$74.97_{3.10}$	$76.94_{3.45}$	$76.25_{2.72}$	$70.65_{4.13}$	$71.67_{4.08}$	$69.45_{7.54}$	$69.74_{2.32}$
FNETGen	$75.38_{4.20}$	$74.54_{3.68}$	$76.13_{5.78}$	$74.93_{7.44}$	$69.70_{4.34}$	$68.33_{5.00}$	$71.33_{7.77}$	$68.60_{3.53}$
BrainGNN	$74.94_{3.49}$	$72.99_{4.18}$	$76.60_{4.41}$	$73.95_{5.70}$	$69.72_{4.18}$	$70.84_{4.16}$	$68.33_{7.03}$	$68.56_{2.81}$
LSGNN	$75.61_{4.64}$	$75.11_{4.24}$	$76.10_{7.58}$	$75.26_{5.98}$	$71.56_{4.47}$	$73.33_{3.33}$	$69.45_{7.54}$	$69.99_{1.86}$
Com-BrainTF	$75.39_{3.62}$	$74.51_{3.77}$	$76.15_{6.60}$	$75.50_{5.22}$	$70.20_{4.60}$	$71.67_{5.53}$	$68.45_{6.69}$	$68.91_{2.93}$
A-GCL	$77.38_{3.45}$	$75.49_{3.97}$	$79.13_{6.82}$	$78.01_{3.21}$	$72.47_{5.44}$	$74.17_{4.49}$	$70.45_{8.18}$	$71.57_{4.21}$
SPAN	$76.71_{3.40}$	$76.40_{5.14}$	$77.43_{5.78}$	$76.83_{3.45}$	$71.10_{4.09}$	$72.50_{5.33}$	$69.33_{4.67}$	$70.53_{3.90}$
Ours	$79.82_{2.92}$	$78.42_{4.62}$	$81.14_{6.96}$	$80.91_{3.93}$	$74.78_{4.61}$	$75.00_{5.27}$	$74.45_{5.23}$	$74.86_{5.53}$

Classification performance. To comprehensively evaluate our approach, we benchmarked it against eight representative methods: classical graph networks (GIN [20]), brain-specific deep learning models (BrainGNN [13], LSGNN [5], FNETGen [10]), transformer-based architectures (BNT [11], Com-BrainTF [1]), and recent state-of-the-art graph contrastive learning frameworks (A-GCL[27], SPAN [14]). Table 1 summarizes the classification performance. As observed, our proposed VB-GCL consistently outperforms all competing methods across all four evaluation metrics on both datasets. Notably, compared to the second-best method (A-GCL), our model achieves an absolute improvement of 2.44% in accuracy and 2.90% in AUC on the ADNI dataset, with similar margins observed on the in-house dataset. This consistent superiority demonstrates the effectiveness of lifting the homogeneous graph to a vector bundle to capture heterogeneous local geometric properties. Furthermore, it validates the strong generalization capability of our dual-stream framework across diverse clinical fMRI datasets.

Ablation studies. To investigate the individual contributions of the key components in our framework, we conducted extensive ablation studies, as shown in Table 2. The base model, which uses a standard GNN backbone without geometric enhancements, yields the poorest performance. When the vector bundle formulation is removed ("w/o Bundle", degrading to a standard homogeneous graph), the accuracy decreases by 3.10% on ADNI. This heavily substantiates the necessity of modeling region-specific functional geometry. Similarly, the removal of spectral contrastive learning ("w/o SCL") leads to a clear performance degradation, highlighting its critical role in providing noise regularization and

mitigating overfitting in small-sample regimes. Finally, removing the geometric diffusion ("w/o Diff") or attention pooling ("w/o Attn") also impairs performance. Together, these results confirm that all proposed modules are highly complementary, and their optimal synergy yields the best diagnostic outcome.

Table 2. Ablation study results of the proposed VB-GCL on two fMRI datasets (%).

Method	ADNI				In-house			
	Acc.	Sen.	Spe.	AUC	Acc.	Sen.	Spe.	AUC
Base	$73.16_{3.26}$	$72.63_{3.66}$	$73.59_{4.07}$	$72.74_{4.85}$	$68.34_{3.28}$	$66.67_{3.73}$	$70.33_{7.22}$	$67.77_{3.59}$
w/o Bundle	$76.72_{3.90}$	$76.02_{5.02}$	$77.87_{6.03}$	$77.37_{5.07}$	$71.60_{5.13}$	$72.50_{5.53}$	$69.33_{6.46}$	$70.72_{2.94}$
w/o Diff	$78.27_{3.55}$	$76.97_{6.54}$	$79.48_{5.91}$	$78.94_{4.38}$	$72.06_{5.40}$	$73.33_{5.00}$	$70.45_{8.18}$	$72.29_{5.70}$
w/o Attn	$79.38_{2.79}$	$78.87_{4.25}$	$79.89_{6.02}$	$79.77_{4.58}$	$74.33_{4.57}$	$75.83_{4.49}$	$72.45_{7.82}$	$73.04_{5.34}$
w/o SCL	$78.49_{3.14}$	$77.94_{4.16}$	$79.04_{6.39}$	$78.90_{4.00}$	$73.42_{4.34}$	$75.00_{3.73}$	$71.45_{8.66}$	$73.12_{5.72}$
Ours	$79.82_{2.92}$	$78.42_{4.62}$	$81.14_{6.96}$	$80.91_{3.93}$	$74.78_{4.61}$	$75.00_{5.27}$	$74.45_{5.23}$	$74.86_{5.53}$

Hyperparameter sensitivity. We investigate the influence of the bundle dimension and node embedding size on model performance, as shown in Fig. 2(a) and (b). The diagnostic accuracy on the ADNI dataset remains remarkably robust across varying dimensions. Conversely, for the smaller in-house dataset, performance peaks at moderate values (e.g., an embedding dimension of 64) and exhibits slight degradation at higher dimensions. This trend suggests that excessive parameterization may induce overfitting on limited clinical cohorts, and selecting a moderate dimensionality strikes an optimal balance between expressive capacity and generalizability.

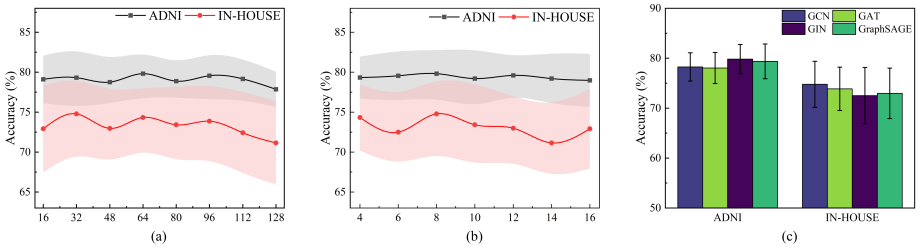

Fig. 2. Hyperparameter sensitivity and backbone analysis. (a) and (b) illustrate the impact of varying bundle and node embedding dimensions on classification accuracy, respectively. Shaded areas represent standard deviation. (c) Performance comparison of the proposed framework using different standard GNN backbones.

Backbone encoders. To evaluate the adaptability of our framework, we compare different mainstream GNNs (GCN [12], GAT [16], GIN, and GraphSAGE [9]) as the backbone encoder, as illustrated in Fig. 2(c). While GIN peaks on ADNI via strong topological expressivity, it degrades on the In-house dataset due to potential overfitting. Conversely, GCN yields the most robust generalization across both datasets. This stability arises because GCN's isotropic smoothing naturally aligns with our bundle geometric diffusion, effectively regularizing topology-aware representations without over-complicating message passing.

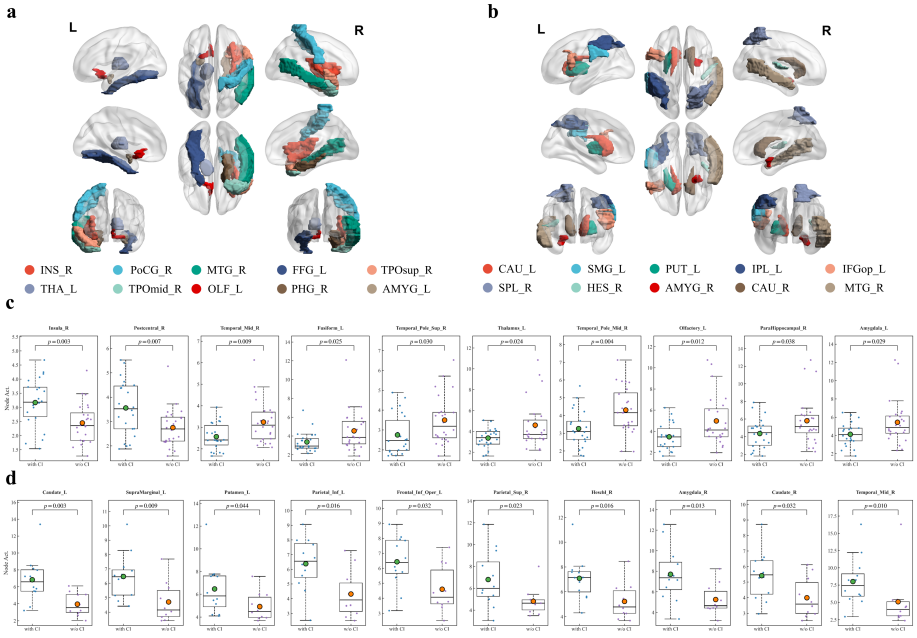

Fig. 3. Discriminative ROIs and node activity analysis. (a, b) Top 10 most discriminative ROIs identified by VB-GCL on the ADNI and in-house datasets, respectively. (c, d) Node activity comparisons of the corresponding ROIs between the with-CI and w/o-CI groups, where p-values further verify the rationality of the ROI selection.

Interpretability analysis. Node importance was assessed using global saliency maps, generated by averaging the gradients of the output logits with respect to the adjacency matrices across all subjects [2,22]. The top 10 most critical regions, ranked by these importance scores, were subsequently visualized using BrainNet Viewer [19]. To further validate the rationality of these selected ROIs, we performed node activity analysis between the with-cognitive impairment (with-CI) and without-cognitive impairment (w/o-CI) groups. This analysis demonstrated significant differences (Fig. 3c, d). For ADNI (Fig. 3a, c), the top discriminative ROIs localize in the limbic and temporal regions (e.g., parahippocampal

gyrus, $p = 0.038$; amygdala, $p = 0.029$), perfectly aligning with established AD neurodegeneration patterns [15,17]. Conversely, for MMD (Fig. 3b, d), critical ROIs shift to the basal ganglia (caudate, $p = 0.003$) and frontoparietal networks, accurately reflecting chronic hypoperfusion territories driving vascular cognitive impairment [8]. All identified ROIs show significant CI vs. w/o CI separability ($p < 0.05$), confirming that VB-GCL transcends black-box predictions to provide transparent, clinically actionable neuropathological insights.

4 Conclusion

This paper presents VB-GCL, a novel dual-stream framework grounded in vector bundle theory, to address the limitations of conventional GNNs in modeling complex geometric misalignments across heterogeneous functional brain manifolds. By lifting the fMRI connectivity graph into a vector bundle, VB-GCL associates each ROI with a learnable orthogonal subspace, effectively aligning locally distinct geometries. Furthermore, the proposed Bundle Laplacian elegantly couples multi-hop geometric diffusion for structural feature aggregation with spectral contrastive perturbations for robust noise regularization. Extensive evaluations on the ADNI and an in-house dataset demonstrate that VB-GCL achieves state-of-the-art diagnostic performance while identifying clinically meaningful pathological biomarkers. Future work will explore extending this framework to multimodal neuroimaging analysis.

Acknowledgments. This work was supported in part by the National Key Research and Development Program (2025YFC2423000), and in part by the National natural science foundation of China (82372096).

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Bannadabhavi, A., Lee, S., Deng, W., Ying, R., Li, X.: Community-aware transformer for autism prediction in fmri connectome. In: International conference on medical image computing and computer-assisted intervention. pp. 287–297. Springer (2023)
2. Bessadok, A., Mahjoub, M.A., Rekik, I.: Graph neural networks in network neuroscience. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(5), 5833–5848 (2022)
3. Biswal, B.B., Uddin, L.Q.: The history and future of resting-state functional magnetic resonance imaging. Nature **641**(8065), 1121–1131 (2025)
4. Cai, H., Sheng, X., Wu, G., Hu, B., Cheung, Y.M., Chen, J.: Brain network classification for accurate detection of alzheimer's disease via manifold harmonic discriminant analysis. IEEE transactions on neural networks and learning systems **35**(12), 17266–17280 (2023)

5. Chen, D., Liu, M., Shen, Z., Zhao, X., Wang, Q., Zhang, L.: 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)
6. Chen, X., Zeng, W., Wu, G., Lei, Y., Ni, W., Wang, Y., Gu, Y., Yu, J.: Identification of vascular cognitive impairment in adult moyamoya disease via integrated graph convolutional network. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 669–678. Springer (2022)
7. Cui, H., Dai, W., Zhu, Y., Kan, X., Gu, A.A.C., Lukemire, J., Zhan, L., He, L., Guo, Y., Yang, C.: Braingb: a benchmark for brain network analysis with graph neural networks. IEEE Transactions on Medical Imaging (2022)
8. Gupta, N., Muthusami, P., Pai, V., Shroff, M.: Arterial ischemic stroke in children: navigating diagnostic and management pathways. RadioGraphics **45**(10), e250017 (2025)
9. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. Advances in neural information processing systems **30** (2017)
10. Kan, X., Cui, H., Lukemire, J., Guo, Y., Yang, C.: 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 (2022)
11. Kan, X., Dai, W., Cui, H., Zhang, Z., Guo, Y., Yang, C.: Brain network transformer. Advances in Neural Information Processing Systems **35**, 25586–25599 (2022)
12. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
13. Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., Scheinost, D., Staib, L.H., Ventola, P., Duncan, J.S.: Braingnn: Interpretable brain graph neural network for fmri analysis. Medical image analysis **74**, 102233 (2021)
14. Lin, L., Chen, J., Wang, H.: Spectral augmentation for self-supervised learning on graphs. arXiv preprint arXiv:2210.00643 (2022)
15. Lu, W., Gong, Q., Chen, Y., Qiu, S., An, J.: Hippocampus-centered structural covariance network reorganization in alzheimer's disease: An individualized graph-based biomarker validated by machine learning. Neural Networks p. 108542 (2026)
16. Petar, V., Guillem, C., Arantxa, C., Adriana, R., Pietro, L., Yoshua, B., et al.: Graph attention networks. In: International conference on learning representations. vol. 8 (2018)
17. Wang, M., Shao, W., Zhang, D., Liu, Q.: Identification of genetic risk factors based on disease progression derived from modeling longitudinal phenotype latent pattern representation. IEEE Transactions on Medical Imaging (2025)
18. Wang, X., Yao, L., Rekik, I., Zhang, Y.: Contrastive functional connectivity graph learning for population-based fmri classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 221–230. Springer (2022)
19. Xia, M., Wang, J., He, Y.: Brainnet viewer: a network visualization tool for human brain connectomics. PloS one **8**(7), e68910 (2013)
20. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? In: International Conference on Learning Representations (2018)
21. Yan, C.G., Wang, X.D., Zuo, X.N., Zang, Y.F.: Dpabi: data processing & analysis for (resting-state) brain imaging. Neuroinformatics **14**(3), 339–351 (2016)
22. Ye, K., Tang, H., Dai, S., Guo, L., Liu, J.Y., Wang, Y., Leow, A., Thompson, P.M., Huang, H., Zhan, L.: Bidirectional mapping with contrastive learning on

multimodal neuroimaging data. In: International conference on medical image computing and computer-assisted intervention. pp. 138–148. Springer (2023)
23. Yin, F., Lei, Y., Dai, S., Zeng, W., Wu, G., Zhan, L., Yu, J.: A heterogeneous graph neural network fusing functional and structural connectivity for mci diagnosis. arXiv preprint arXiv:2411.08424 (2024)
24. Zeng, W., Yin, F., Song, P., Wu, Y., Zhao, C., Wu, G., Yu, J.: Adaptive multi-scale dynamic graph representation learning with overlapping community-awareness for asd classification. IEEE Journal of Biomedical and Health Informatics **29**(12), 8711–8718 (2025)
25. Zhang, H., Liu, X., Ma, Y., Yuan, Y., Zhang, D., Zhang, L., et al.: Graph-level contrastive learning with self-aware and cross-sample topology augmentation for brain disorder diagnosis using rs-fmri. Neural Networks p. 108379 (2025)
26. Zhang, P., Wen, G., Cao, P., Yang, J., Zhang, J., Zhang, X., Zhu, X., Zaiane, O.R., Wang, F.: Brainusl: U nsupervised graph s tructure l earning for functional brain network analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 205–214. Springer (2023)
27. Zhang, S., Chen, X., Shen, X., Ren, B., Yu, Z., Yang, H., Jiang, X., Shen, D., Zhou, Y., Zhang, X.Y.: A-gcl: Adversarial graph contrastive learning for fmri analysis to diagnose neurodevelopmental disorders. Medical Image Analysis **90**, 102932 (2023)
28. Zhou, H., He, L., Chen, B.Y., Shen, L., Zhang, Y.: Multi-modal diagnosis of alzheimer's disease using interpretable graph convolutional networks. IEEE Transactions on Medical Imaging **44**(1), 142–153 (2024)
29. Zong, Y., Zuo, Q., Ng, M.K.P., Lei, B., Wang, S.: A new brain network construction paradigm for brain disorder via diffusion-based graph contrastive learning. IEEE transactions on pattern analysis and machine intelligence **46**(12), 10389–10403 (2024)