BrainACU: Asynchronous Coordination Unit Modeling for Dynamic Brain Network Analysis

Guiliang Guo[1], Lanting Li[2], Qi Liu[1], Guangqi Wen[3], Lingwen Liu[4], Ruoxian Song[1], Peng Cao[1(✉)], Jinzhu Yang[1], Fei Wang[5(✉)], Xiaoli Liu[6], and Osmar R. Zaiane[7]

[1] Computer Science and Engineering, Northeastern University, Shenyang, China
caopeng@cse.neu.edu.cn
[2] College of Artificial Intelligence, Shenyang Aerospace University, Shenyang, China
[3] School of Computer Science and Artificial Intelligence, Shandong Normal University, Jinan, China
[4] Graduate School of Information, Production and Systems, Waseda University, Kitakyushu, Fukuoka, Japan
[5] Nanjing Medical University, Nanjing, China
[6] AiShiWeiLai AI Research, Beijing, China
[7] University of Alberta, Edmonton, Alberta, Canada

Abstract. Resting-state functional MRI (rs-fMRI) provides crucial insights into dynamic brain network mechanisms underlying psychiatric disorders. While deep learning methods have advanced rs-fMRI-based disease recognition and dynamic brain network modeling, existing approaches are often limited by two key assumptions: **(1) fixed subnetwork partitioning**, where predefined functional templates or partition rules impose a fixed region–subnetwork mapping that, although interpretable, restricts the discovery of complex coordination patterns; and **(2) synchronization assumption**, where inter-subnetwork interactions are modeled as strictly time-aligned, overlooking physiologically plausible delays and weakening diagnostic signals. To address these limitations, we propose **BrainACU**, an asynchronous coordination unit modeling framework for dynamic brain network analysis. BrainACU removes fixed partitioning by learning coordination units (CUs) directly from data, capturing both synchronous and asynchronous coordination. Using conventional sliding windows, we segment BOLD sequences into temporally consistent patches for subsequent CU learning and interaction modeling. BrainACU then learns stable ROI-to-CU mappings via data-driven clustering with cross-sample consensus, and projects variable-size CU signals into a unified latent space for dimensional alignment. Finally, an asynchronous interaction module models both time-aligned and time-lagged unit-level dependencies to form dynamic brain network representations. Experiments on ASD, MDD, and BD demonstrate that BrainACU outperforms existing methods and provides more discriminative and interpretable dynamic brain network evidence. Our code: https://github.com/guoguiliang111/BrainACU.

G. Guo and L. Li contributed equally to this work.

Keywords: Dynamic brain network analysis · Coordination unit learning · Asynchronous interactions · Interpretability

1 Introduction

Resting-state functional MRI (rs-fMRI) has become a key non-invasive approach for mechanistic studies and computer-aided diagnosis of psychiatric disorders (e.g., ASD, MDD, and BD) [2, 7]. Increasing evidence suggests that these disorders are better explained by disrupted, time-varying coordination among functional networks rather than isolated regional abnormalities, motivating functional connectivity (FC) analysis and dynamic brain network modeling [23, 5]. Despite recent advances of CNN-, graph neural network-, and Transformer-based methods for functional connectivity analysis and dynamic modeling [11, 14, 24, 31, 9, 13], existing approaches remain constrained by two major limitations: **(1) Fixed subnetwork partitioning.** Many methods rely on predefined functional templates or fixed partition rules to establish an ROI–subnetwork correspondence [30, 17].This strong structural prior may obscure atypical cross-template coordination and restrict the discovery of non-canonical, disease-relevant coordination groupings, especially in multi-disorder settings (e.g., ASD, MDD, and BD) where different pathologies can induce distinct organizational patterns [21, 27]. **(2) Synchronization assumption.** Inter-subnetwork interactions are often operationalized using zero-lag dependence measures under strict temporal alignment, implicitly assuming that coordination occurs synchronously [15, 18]. This neglects physiologically plausible temporal lags and asynchronous coupling in real neural systems, causing delayed-propagation disease signals to be weakened into noise or misaligned correlations, and causing interpretability to overemphasize zero-lag (time-aligned) interactions while underestimating delayed couplings [20, 19].

To address these limitations, we propose **BrainACU**, an *asynchronous coordination unit* framework for dynamic brain network analysis. BrainACU learns data-driven *Coordination Units* (CUs) from ROIs and models both time-aligned and time-lagged cross-CU interactions within a unified attention-based architecture, shifting from fixed templates to data-driven organization and relaxing strict synchrony assumptions. Specifically, given a subject's BOLD sequence, we use a standard fixed-length sliding window to form temporally normalized patches, yielding consistently indexed temporal segments that provide a unified reference for modeling time-aligned and time-lagged CU interactions. Based on these patches, we learn a stable **assignment of ROIs to CUs** via per-subject ROI clustering followed by cross-subject consensus aggregation, and project CU-specific signals of varying sizes into a shared latent space to obtain dimension-aligned unit representations. An asynchronous interaction module then models both time-aligned and time-lagged cross-CU dependencies via directed attention, producing a discriminative dynamic coordination representation, while L_{lag} suppresses near-diagonal (zero-lag) attention to encourage delayed coupling and L_{div} discourages CU embedding similarity to avoid collapse and promote com-

plementary coordination patterns. Experiments on multi-disorder settings show consistent gains, e.g., 76.8/77.9 (ACC/AUC) on HC vs. MDD, 78.6/77.9 on HC vs. BD, and 72.3/74.1 on HC vs. ASD (NYU), with ablations and delay-statistics analyses confirming the benefit of asynchronous coordination modeling. Our contributions are summarized as follows:

1. We propose **BrainACU**, a dynamic brain network analysis framework that learns *data-driven Coordination Units* (CUs), reducing reliance on fixed subnetwork partitioning and providing a more flexible basis for characterizing disease-relevant coordination beyond canonical templates and zero-lag synchrony.
2. We unify *synchronous* and *asynchronous* coordination modeling at the CU level, explicitly capturing time-shifted dependencies to avoid zero-lag synchronization bias and to improve interpretability.
3. Extensive experiments on **ASD**, **MDD**, and **BD** demonstrate that BrainACU achieves consistent gains over representative baselines, while ablation studies verify the necessity of CU discovery, latent alignment, and asynchronous interaction; moreover, lag-step statistics reveals disease-related delayed coordination patterns, providing mechanistically meaningful interpretability.

2 Method

Overview. We propose **BrainACU**, an asynchronous coordination unit (CU) modeling framework for dynamic brain network analysis. Given an ROI-level BOLD sequence segmented into fixed-length sliding-window patches, BrainACU (i) discovers *data-driven* CUs, (ii) builds *dimension-aligned* CU representations, and (iii) explicitly models *time-aligned* and *time-lagged* cross-CU dependencies via directed cross-CU attention.

2.1 Problem Setup and Sliding-Window Patching

For each subject, the rs-fMRI signal is represented as:

$$\mathbf{X} \in \mathbb{R}^{T\times C}, \tag{1}$$

where T is the number of time points and C is the number of ROIs. To obtain a consistently indexed token timeline, we apply a fixed-length sliding window:

$$\mathbf{P}^{(m)} = \mathbf{X}[t_m : t_m{+}L{-}1, :] \in \mathbb{R}^{L\times C}, \quad t_m = (m-1)s+1, \quad m = 1,\ldots,M, \tag{2}$$

where L is the window length, s is the stride, and $M = \left\lfloor \frac{T-L}{s} \right\rfloor + 1$ is the number of patches. In our implementation, we use L=30 and s=10 to produce patch tensors of shape $\mathbf{P} \in \mathbb{R}^{B\times L\times M\times C}$ for a minibatch of size B (with B=32).

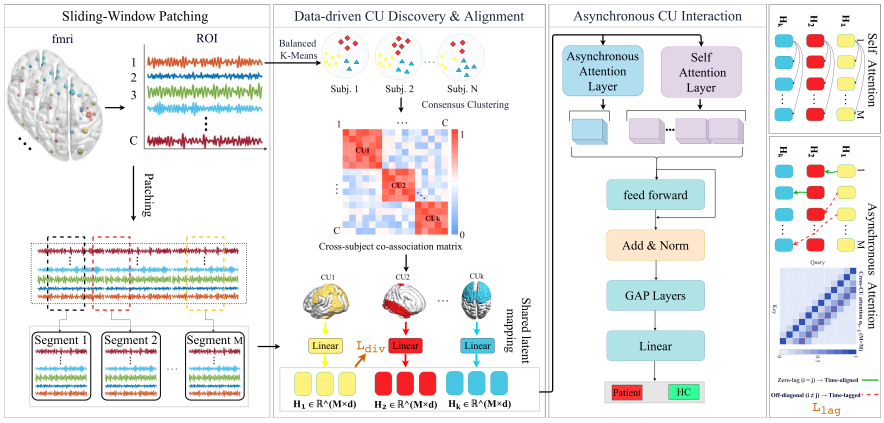

Fig. 1. Overview of BrainACU. Given an ROI-level BOLD sequence $\mathbf{X} \in \mathbb{R}^{T\times C}$, we segment it into fixed-length temporal patches to form a token timeline. BrainACU discovers *data-driven* coordination units (CUs) by learning a stable ROI-to-CU assignment via per-subject clustering followed by cross-subject consensus. Each CU is mapped to a shared latent space to obtain dimension-aligned CU tokens. An asynchronous interaction module models both time-aligned and time-lagged dependencies across CUs via directed cross-CU attention, while intra-CU self-attention captures within-unit temporal dynamics. All CU streams are fused for diagnosis prediction.

2.2 Data-Driven CU Discovery via Consensus Clustering

BrainACU learns a stable ROI-to-CU assignment from data, avoiding predefined functional templates. Let K be the CU number (we use K=3 throughout). We first perform *per-subject ROI clustering* into K groups using a balanced k-means variant to prevent degenerate CU sizes, producing subject-wise labels $\mathbf{z}^{(n)} \in \{1,\dots,K\}^C$. To obtain a *stable* grouping across subjects, we construct a cross-subject co-association matrix and apply consensus clustering:

$$\mathbf{A}_{ij} = \frac{1}{N}\sum_{n=1}^{N} \mathbb{1}\left[z_i^{(n)} = z_j^{(n)}\right], \qquad i,j \in \{1,\dots,C\}, \tag{3}$$

where N is the number of training subjects and $\mathbf{A}_{ij}$ measures how often ROI i and ROI j fall into the same cluster across subjects. We then run K-means on $\mathbf{A}$ to obtain final CU labels $\hat{\mathbf{z}} \in \{1,\dots,K\}^C$.

2.3 CU Tokens via Shared Latent Mapping

For each patch index m, we treat the ROI activations as $\mathbf{x}_m \in \mathbb{R}^C$ and group them into CU-specific vectors $\mathbf{x}_{m,k} \in \mathbb{R}^{|\mathcal{V}_k|}$ using the learned ROI sets $\mathcal{V}_k = \{i \mid \hat{z}_i = k\}$. Because $|\mathcal{V}_k|$ varies across CUs, we map each CU to a shared latent dimension d via CU-specific linear projections:

$$\mathbf{h}_{m,k} = \mathbf{W}_k \mathbf{x}_{m,k} \in \mathbb{R}^d, \qquad \mathbf{W}_k \in \mathbb{R}^{d\times|\mathcal{V}_k|}. \tag{4}$$

Here $\mathbf{h}_{m,k}$ is the CU token for CU k at patch m, and d is the shared embedding size (we use d=36). This yields K CU token sequences $\{\mathbf{h}_{m,k}\}_{m=1}^{M}$ that are dimension-aligned and ready for attention-based interaction modeling.

Diversity regularization for CU specialization. A practical risk of data-driven grouping is CU collapse (different CUs becoming redundant). We regularize CU specialization by penalizing similarity among CU mean embeddings:

$$\mathcal{L}_{\text{div}} = \sum_{i<j} \left(\frac{\bar{\mathbf{h}}_i^\top \bar{\mathbf{h}}_j}{\|\bar{\mathbf{h}}_i\|_2 \, \|\bar{\mathbf{h}}_j\|_2} \right)^2, \quad \bar{\mathbf{h}}_k = \frac{1}{M} \sum_{m=1}^{M} \mathbf{h}_{m,k}, \tag{5}$$

where $\bar{\mathbf{h}}_k$ is the temporal mean token of CU k (before cross-CU interaction). This encourages different CUs to encode complementary coordination patterns.

2.4 Asynchronous Interaction Module

Given CU token sequences $\{\mathbf{h}_{m,k}\}_{m=1}^{M}$ from Eq. (4), we form $\mathbf{H}_k = [\mathbf{h}_{1,k}; \ldots; \mathbf{h}_{M,k}] \in \mathbb{R}^{M \times d}$ ($K = 3$, $d = 36$). BrainACU applies (i) intra-CU self-attention over the patch timeline and (ii) directed cross-CU attention to capture lagged coordination; each attention block is followed by a lightweight FFN with residual connection and LayerNorm (as in standard Transformer layers). For each CU k, we update $\mathbf{H}_k$ by:

$$\mathbf{H}_k^{\text{in}} = \text{Attn}_k\big(\mathbf{H}_k \mathbf{W}_Q^{\text{in}}, \mathbf{H}_k \mathbf{W}_K^{\text{in}}, \mathbf{H}_k \mathbf{W}_V^{\text{in}}\big) = \text{Softmax}\left(\frac{\mathbf{Q}_k \mathbf{K}_k^\top}{\sqrt{d_h}} \right) \mathbf{V}_k \in \mathbb{R}^{M \times d}, \tag{6}$$

where $\mathbf{W}_{\{\cdot\}}^{\text{in}} \in \mathbb{R}^{d \times d_h}$ and d_h is the head dimension.

Directed cross-CU attention. For each ordered pair $(j \to i)$ with $i \neq j$, CU i queries CU j:

$$\boldsymbol{\alpha}_{i \leftarrow j} = \text{Softmax}\left(\frac{\mathbf{Q}_{i \leftarrow j} \mathbf{K}_j^\top}{\sqrt{d_h}} \right) \in \mathbb{R}^{M \times M}, \quad \mathbf{U}_{i \leftarrow j} = \boldsymbol{\alpha}_{i \leftarrow j} \mathbf{V}_j \in \mathbb{R}^{M \times d}, \tag{7}$$

with $\mathbf{Q}_{i \leftarrow j} = \mathbf{H}_i^{\text{in}} \mathbf{W}_{Q,ij}^{\text{cr}}$, $\mathbf{K}_j = \mathbf{H}_j^{\text{in}} \mathbf{W}_{K,ij}^{\text{cr}}$, $\mathbf{V}_j = \mathbf{H}_j^{\text{in}} \mathbf{W}_{V,ij}^{\text{cr}}$. Here $\boldsymbol{\alpha}_{i \leftarrow j}(p, q)$ denotes the attention weight from the q-th token (patch index) of CU j to the p-th token of CU i; off-diagonal mass corresponds to lagged coupling. We aggregate incoming messages by $\mathbf{U}_i = \sum_{j \neq i} \mathbf{U}_{i \leftarrow j}$ and return $\{\boldsymbol{\alpha}_{i \leftarrow j}\}$ for lag-step statistics and interpretability.

Lag-promoting attention regularization. Without constraints, the directed cross-CU attention may degenerate to near-diagonal (near zero-lag) alignment. We penalize near-diagonal attention mass using a Gaussian kernel:

$$\mathcal{L}_{\text{lag}} = \sum_{i \neq j} \sum_{p=1}^{M} \sum_{q=1}^{M} \exp\left(-\frac{(p-q)^2}{2\kappa^2} \right) \cdot \boldsymbol{\alpha}_{i \leftarrow j}(p, q), \tag{8}$$

where κ controls the width of the small-lag band. Minimizing $\mathcal{L}_{\text{lag}}$ suppresses trivial zero-lag alignment and encourages lag-aware cross-CU coupling.

2.5 Prediction and Training Objective

All CU streams (intra-CU and cross-CU) are fused by concatenation and a position-wise MLP, then temporally pooled for binary classification. We optimize BrainACU end-to-end with binary cross-entropy plus the two method-aligned regularizers:

$$\mathcal{L} = \mathcal{L}_{\text{cls}} + \lambda_{\text{lag}}\mathcal{L}_{\text{lag}} + \lambda_{\text{div}}\mathcal{L}_{\text{div}}, \tag{9}$$

where $\mathcal{L}_{\text{cls}}$ is the binary cross-entropy, λ_{lag}=0.1 weights lag-aware coordination, and λ_{div}=0.01 weights CU specialization.

3 Experiments and Results

3.1 Datasets

We evaluate BrainACU on two rs-fMRI benchmarks for three diagnoses (**MDD**, **BD**, **ASD**). The private cohort constructed from Nanjing Medical University (NMU)[8] includes **246** controls, **151** MDD, and **126** BD subjects, preprocessed via **DPABI** [28]. From the public **ABIDE** dataset, we use only the **NYU** site (**74** ASD, **98** controls) to reduce multi-center variability. We used AAL-based pre-processed functional data for evaluation. Five-fold cross-validation is applied, with hyperparameters tuned on internal validation set.

3.2 Experimental Results and Discussion

Comparison with state-of-the-art. We compare with (i) **traditional FC-based classifiers** on handcrafted connectivity features, (ii) **static-graph** neural models on aggregated FC, and (iii) **dynamic-graph** models for dFC. Table 1 reports mean±std over 5-fold CV. Overall, BrainACU achieves the best performance on Center-NMU and remains competitive on ABIDE (NYU), suggesting that *data-driven coordination units* and *asynchronous (lagged) coordination modeling* provide complementary cues beyond sync-only interaction modeling. On Center-NMU, BrainACU achieves **76.8±2.3**/**77.9±2.7** (ACC/AUC) for HC vs. MDD and **78.6±2.1**/**77.9±2.5** for HC vs. BD (Table 1), outperforming strong dynamic baselines (e.g., BrainDGT) by clear margins. On ABIDE (NYU), BrainACU achieves the best overall performance in our table, supporting the value of modeling *lagged cross-unit coordination* across disorders.

Ablation analysis. We conduct single-factor ablations by removing one component from the full BrainACU (Table 1). Replacing data-driven CU discovery with a fixed partition reduces performance on all tasks, indicating the benefit of cohort-adaptive coordination grouping in multi-disorder settings. Removing the shared latent-space mapping also causes consistent drops, showing that dimension alignment is important for stable cross-CU interaction modeling.

[8] Approved by the Medical Science Ethics Committee of Nanjing Medical University (Ref. 2020-KY027-01).

Table 1. Comparison with state-of-the-art methods and ablations on **MDD**, **BD**, and **ASD** (ABIDE NYU). The best results are in bold. The average and std of ACC and AUC are reported.

Category	Method	HC vs. MDD		HC vs. BD		HC vs. ASD (NYU)	
		ACC(%)	AUC(%)	ACC(%)	AUC(%)	ACC(%)	AUC(%)
Traditional	FC + SVM	68.4 ± 2.7	67.2 ± 3.2	65.6 ± 3.1	67.3 ± 3.5	68.1 ± 2.9	69.7 ± 3.6
	FC + RF	64.7 ± 2.9	61.2 ± 3.4	65.2 ± 3.1	63.6 ± 4.0	66.2 ± 2.8	65.1 ± 3.7
	SVM-RFE [4]	70.3 ± 3.8	70.1 ± 4.0	66.9 ± 3.7	67.6 ± 3.9	65.2 ± 4.1	65.9 ± 4.3
Static graph	GroupINN [29]	67.1 ± 3.9	65.7 ± 4.2	68.3 ± 3.8	65.2 ± 4.0	63.3 ± 4.3	64.3 ± 4.4
	BrainGNN [14]	65.2 ± 3.6	63.7 ± 3.9	69.1 ± 3.4	68.7 ± 3.6	65.2 ± 3.8	65.4 ± 4.0
	MVS-GCN [26]	68.1 ± 3.5	68.1 ± 3.7	67.1 ± 3.6	65.0 ± 3.8	69.3 ± 3.7	68.3 ± 3.9
	ASD-DiagNet [1]	68.2 ± 3.4	68.0 ± 3.6	72.1 ± 3.2	72.1 ± 3.4	67.4 ± 3.5	67.5 ± 3.7
	FRL [10]	68.0 ± 3.3	68.9 ± 3.5	71.1 ± 3.2	70.7 ± 4.1	67.1 ± 3.5	69.4 ± 3.7
	EAG-RS [8]	66.5 ± 3.6	70.9 ± 3.8	73.5 ± 3.1	72.1 ± 4.3	66.8 ± 3.7	68.1 ± 3.9
	D-CoRP [6]	71.7 ± 2.7	71.0 ± 2.9	72.6 ± 2.7	71.8 ± 2.9	67.9 ± 1.7	65.3 ± 2.1
Dynamic graph	MDGL [16]	67.4 ± 3.7	66.7 ± 3.9	69.1 ± 3.5	70.3 ± 3.7	66.9 ± 3.9	64.1 ± 4.1
	STA-GIN [12]	73.1 ± 3.2	66.0 ± 3.5	72.4 ± 3.1	68.5 ± 3.3	68.9 ± 3.4	67.2 ± 3.7
	BrainDGT [22]	74.3 ± 2.6	71.3 ± 2.8	73.7 ± 3.4	74.7 ± 2.7	69.1 ± 3.1	69.7 ± 4.2
	MCDGLN [25]	72.6 ± 2.5	71.1 ± 3.7	71.6 ± 3.2	72.1 ± 2.6	67.2 ± 2.7	69.3 ± 2.9
	ST-GCN [3]	60.3 ± 4.8	62.9 ± 5.1	61.3 ± 4.7	60.9 ± 4.9	56.7 ± 5.2	52.6 ± 5.6
Ours	w/o CU Discovery	74.6 ± 2.8	75.2 ± 3.1	76.1 ± 2.6	75.4 ± 2.9	69.5 ± 3.3	71.1 ± 3.6
	w/o CU Alignment	75.3 ± 2.6	76.0 ± 2.9	76.8 ± 2.4	76.0 ± 2.7	70.0 ± 3.1	71.7 ± 3.4
	w/o $\mathcal{L}_{\text{lag}}$	76.1 ± 2.4	77.0 ± 2.7	77.6 ± 2.1	76.8 ± 2.5	71.2 ± 2.9	72.8 ± 3.2
	w/o $\mathcal{L}_{\text{div}}$	76.4 ± 2.3	77.3 ± 2.6	77.9 ± 2.0	77.1 ± 2.3	71.6 ± 2.8	73.2 ± 3.1
	w/o Async Interaction	73.1 ± 3.0	74.0 ± 3.4	75.0 ± 2.7	74.2 ± 3.2	69.0 ± 3.5	70.6 ± 3.8
Full Model	**BrainACU**	**76.8±2.3**	**77.9±2.7**	**78.6±2.1**	**77.9±2.5**	**72.3±3.1**	**74.1±3.4**

Disabling the asynchronous module yields the largest degradation, confirming that time-lagged cross-CU dependencies carry key diagnostic cues beyond sync-only modeling, especially for MDD and BD. Finally, removing $\mathcal{L}_{\text{lag}}$ leads to a smaller but consistent decrease, suggesting it prevents attention from collapsing to near-diagonal alignments and strengthens lag-aware behavior. Removing $\mathcal{L}_{\text{div}}$ also slightly degrades performance, implying that encouraging CU specialization helps avoid redundant units and yields more complementary coordination patterns.

Hyper-parameter analysis on CU number. Table 2 reports mean±std over 5-fold CV. Overall, performance peaks at a moderate granularity (K=3), while larger K tends to over-partition ROIs and introduce redundant CU streams, making cross-CU interaction modeling harder to optimize and less robust. Conversely, too small K may underfit fine-grained coordination patterns.

3.3 Interpretability Analysis of Asynchronous Coordination

BrainACU provides interpretable evidence for time-lagged coordination via its asynchronous interaction module. We analyze **lag-step statistics** on the BD cohort. Quantitatively, from head-averaged attention matrices we extract attention

Table 2. Effect of the CU number K on three diagnosis tasks (5-fold CV).

K	HC vs. MDD		HC vs. BD		HC vs. ASD (NYU)	
	ACC(%)	AUC(%)	ACC(%)	AUC(%)	ACC(%)	AUC(%)
2	75.1 ± 2.4	76.2 ± 2.8	76.9 ± 2.0	76.1 ± 2.2	70.8 ± 2.9	72.5 ± 3.2
3	$\mathbf{76.8 \pm 2.3}$	$\mathbf{77.9 \pm 2.7}$	$\mathbf{78.6 \pm 2.1}$	$\mathbf{77.9 \pm 2.5}$	$\mathbf{72.3 \pm 3.1}$	$\mathbf{74.1 \pm 3.4}$
4	74.5 ± 2.5	75.6 ± 2.9	76.2 ± 2.3	75.3 ± 2.6	70.1 ± 3.2	71.8 ± 3.5
5	72.8 ± 2.7	74.0 ± 3.1	74.5 ± 2.5	73.6 ± 2.8	68.5 ± 3.4	69.9 ± 3.7

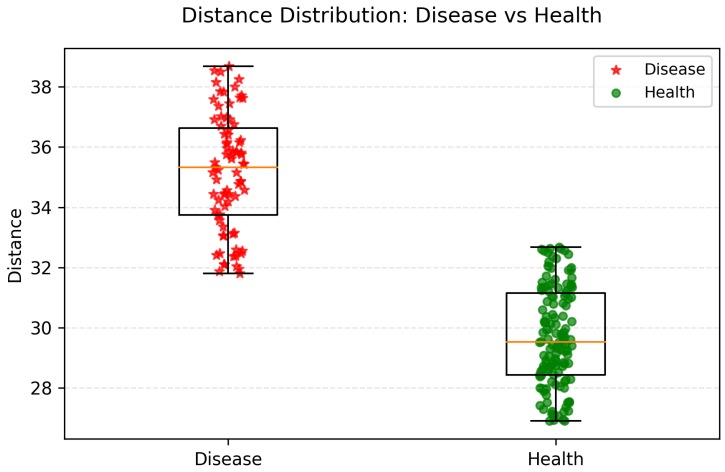

Fig. 2. Lag-step statistics derived from cross-CU attention. Disease groups exhibit larger lag steps than controls, suggesting abnormal delayed coordination patterns captured by BrainACU.

extrema and define lag steps as $d_{ij} = |i - j|$ (with i/j the query/key temporal indices); summary statistics (mean/median/variance) show consistently larger lag steps in the disease group, suggesting disrupted and delayed coordination.

4 Conclusion

We presented **BrainACU**, a Transformer-based framework for dynamic fMRI diagnosis that explicitly models asynchronous interactions among brain subnetworks. BrainACU combines temporal patching, subnetwork partition-and-mapping, and asynchronous interaction modeling in an end-to-end architecture. Experiments demonstrate clear improvements over representative baselines, and interpretability analyses reveal disease-related delayed coordination patterns. These results suggest that subnetwork-asynchronous modeling is an effective and interpretable direction for rs-fMRI-based psychiatric diagnosis.

Acknowledgments. This research was supported by the National Natural Science Foundation of China (62076059), the Science and Technology Joint Project of Liaoning Province (2023JH2/101700367 and ZX20240193) and the Fundamental Research Funds for the Central Universities (N25BJD005).

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

References

1. Eslami, T., Mirjalili, V., Fong, A., Laird, A.R., Saeed, F.: Asd-diagnet: a hybrid learning approach for detection of autism spectrum disorder using fmri data. Frontiers in neuroinformatics **13**, 70 (2019)
2. Fox, M.D., Raichle, M.E.: Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nature reviews neuroscience **8**(9), 700–711 (2007)
3. Gadgil, S., Zhao, Q., Pfefferbaum, A., Sullivan, E.V., Adeli, E., Pohl, K.M.: Spatio-temporal graph convolution for resting-state fmri analysis. In: International conference on medical image computing and computer-assisted intervention. pp. 528–538. Springer (2020)
4. Guyon, I., Weston, J., Barnhill, S., Vapnik, V.: Gene selection for cancer classification using support vector machines. Machine learning **46**(1), 389–422 (2002)
5. Hoheisel, L., Kambeitz-Ilankovic, L., Wenzel, J., Haas, S.S., Antonucci, L.A., Ruef, A., Penzel, N., Schultze-Lutter, F., Lichtenstein, T., Rosen, M., et al.: Alterations of functional connectivity dynamics in affective and psychotic disorders. Biological Psychiatry: Cognitive Neuroscience and Neuroimaging **9**(8), 765–776 (2024)
6. Hu, H., Zhang, H., Li, C.: D-corp: Differentiable connectivity refinement for functional brain networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 119–129. Springer (2024)
7. Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Della Penna, S., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., et al.: Dynamic functional connectivity: promise, issues, and interpretations. Neuroimage **80**, 360–378 (2013)
8. Jung, W., Jeon, E., Kang, E., Suk, H.I.: Eag-rs: a novel explainability-guided roi-selection framework for asd diagnosis via inter-regional relation learning. IEEE Transactions on Medical Imaging **43**(4), 1400–1411 (2023)
9. 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)
10. Kang, E., Heo, D.W., Lee, J., Suk, H.I.: A learnable counter-condition analysis framework for functional connectivity-based neurological disorder diagnosis. IEEE Transactions on Medical Imaging **43**(4), 1377–1387 (2023)
11. Kawahara, J., Brown, C.J., Miller, S.P., Booth, B.G., Chau, V., Grunau, R.E., Zwicker, J.G., Hamarneh, G.: Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. NeuroImage **146**, 1038–1049 (2017)
12. Kim, B.H., Ye, J.C., Kim, J.J.: Learning dynamic graph representation of brain connectome with spatio-temporal attention. Advances in Neural Information Processing Systems **34**, 4314–4327 (2021)

13. Kim, P., Kwon, J., Joo, S., Bae, S., Lee, D., Jung, Y., Yoo, S., Cha, J., Moon, T.: Swift: Swin 4d fmri transformer. Advances in Neural Information Processing Systems **36**, 42015–42037 (2023)
14. 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)
15. Liu, Z.Q., Luppi, A.I., Hansen, J.Y., Tian, Y.E., Zalesky, A., Yeo, B.T., Fulcher, B.D., Misic, B.: Benchmarking methods for mapping functional connectivity in the brain. Nature Methods **22**(7), 1593–1602 (2025)
16. Ma, Y., Wang, Q., Cao, L., Li, L., Zhang, C., Qiao, L., Liu, M.: Multi-scale dynamic graph learning for brain disorder detection with functional mri. IEEE Transactions on Neural Systems and Rehabilitation Engineering **31**, 3501–3512 (2023)
17. Menon, V.: Large-scale brain networks and psychopathology: a unifying triple network model. Trends in cognitive sciences **15**(10), 483–506 (2011)
18. Meszlényi, R.J., Hermann, P., Buza, K., Gál, V., Vidnyánszky, Z.: Resting state fmri functional connectivity analysis using dynamic time warping. Frontiers in neuroscience **11**, 75 (2017)
19. Mitra, A., Snyder, A.Z., Blazey, T., Raichle, M.E.: Lag threads organize the brain's intrinsic activity. Proceedings of the National Academy of Sciences **112**(17), E2235–E2244 (2015)
20. Mitra, A., Snyder, A.Z., Hacker, C.D., Raichle, M.E.: Lag structure in resting-state fmri. Journal of neurophysiology **111**(11), 2374–2391 (2014)
21. Salehi, M., Greene, A.S., Karbasi, A., Shen, X., Scheinost, D., Constable, R.T.: There is no single functional atlas even for a single individual: Functional parcel definitions change with task. NeuroImage **208**, 116366 (2020)
22. Shehzad, A., Zhang, D., Yu, S., Abid, S., Xia, F.: Dynamic graph transformer for brain disorder diagnosis. IEEE Journal of Biomedical and Health Informatics (2025)
23. Song, X., Shu, K., Yang, P., Zhao, C., Zhou, F., Frangi, A.F., Xiao, X., Dong, L., Wang, T., Wang, S., et al.: Knowledge-aware multisite adaptive graph transformer for brain disorder diagnosis. IEEE transactions on medical imaging **44**(6), 2370–2383 (2024)
24. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
25. Wang, P., Wen, X., Lei, Y., Guo, Y., Li, J., Hao, Y., Cao, R., Gao, C., Cao, R.: Mcdgln: Masked connection-based dynamic graph learning network for autism spectrum disorder. Brain Research Bulletin **224**, 111290 (2025)
26. Wen, G., Cao, P., Bao, H., Yang, W., Zheng, T., Zaiane, O.: Mvs-gcn: A prior brain structure learning-guided multi-view graph convolution network for autism spectrum disorder diagnosis. Computers in biology and medicine **142**, 105239 (2022)
27. Wu, X., Liang, C., Bustillo, J., Kochunov, P., Wen, X., Sui, J., Jiang, R., Yang, X., Fu, Z., Zhang, D., et al.: The impact of atlas parcellation on functional connectivity analysis across six psychiatric disorders. Human Brain Mapping **46**(5), e70206 (2025)
28. 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)
29. Yan, Y., Zhu, J., Duda, M., Solarz, E., Sripada, C., Koutra, D.: Groupinn: Grouping-based interpretable neural network for classification of limited, noisy brain data. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 772–782 (2019)

30. Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., et al.: The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of neurophysiology (2011)
31. Zhang, H., Song, R., Wang, L., Zhang, L., Wang, D., Wang, C., Zhang, W.: Classification of brain disorders in rs-fmri via local-to-global graph neural networks. IEEE transactions on medical imaging **42**(2), 444–455 (2022)