# Local-Global Coupling Spiking Graph Transformer for Brain Disorders Diagnosis from Two Perspectives

**Geng Zhang**[1,2,3], **Jiangrong Shen**[1,2,4]*, **Kaizhong Zheng**[1,2,3], **Liangjun Chen**[1,2,3], **Badong Chen**[1,2,3]*

[1]National Key Laboratory of Human-Machine Hybrid Augmented Intelligence
[2]National Engineering Research Center for Visual Information and Applications
[3]Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University
[4]Faculty of Electronic and Information Engineering, Xi'an Jiaotong University

## Abstract

Brain disorders have been consistently associated with abnormalities in specific brain regions or neural circuits. Identifying key brain regional activities and functional connectivity patterns is essential for discovering more precise neurobiological biomarkers. However, previous studies have primarily emphasized alterations in functional connectivity while overlooking abnormal neuronal population activity within brain regions. To bridge this gap, we propose a novel Local-Global Coupling Spiking Graph Transformer (LGC-SGT) that jointly models both inter-regional connectivity differences and deviations in neuronal population firing rates within brain regions, enabling a dual-perspective neuropathological analysis. The global pathway leverages spike-based computation in LGC-SGT to model biologically plausible aberrant neural firing dynamics, while the local pathway adaptively captures abnormal graph-based representations of brain connectivity learned by local plasticity in the liquid state machine module. Furthermore, we design a shortcut-enhanced output strategy in LGC-SGT with the hybrid loss function to suppress outlier interference caused by inter-individual and inter-center variability, enabling a more robust decision boundary. Extensive experiments on three brain disorder datasets demonstrate that our model consistently outperforms state-of-the-art graph methods in brain disorder diagnosis. Moreover, it facilitates the extraction of interpretable neurobiological biomarkers by jointly analyzing regional neural activity and functional connectivity, offering a more comprehensive framework for brain disorder understanding and diagnosis.

## 1 Introduction

Brain disorders are complex, multi-level conditions characterized by alterations at both the neuronal and network levels [1]. Evidence suggests that excessive neuronal activity in the hippocampus and amygdala, along with reorganization of functional connections within the prefrontal cortex and the default mode network, likely contributes to cognitive impairments [2]. Connectivity patterns within these regions are particularly crucial, as they mediate the integration of neural circuits that underpin cognitive function [3]. Moreover, stress-related disorders, such as major depressive disorder (MDD) and schizophrenia, have been shown to be associated with neural activity in the amygdala [4]. Thus, abnormal activity in specific neuronal populations can disrupt neural oscillations, also leading to cognitive dysfunction [5]. Spiking neural networks (SNNs), which biologically model neuronal firing behaviors in a biologically plausible manner, offer a unique opportunity to integrate

*Corresponding author: jrshen@zju.edu.cn, chenbd@mail.xjtu.edu.cn

39th Conference on Neural Information Processing Systems (NeurIPS 2025).

connectivity analysis with population-level neural dynamics. Therefore, this study aims to distinguish patients from healthy individuals using SNNs, considering both neuronal activity and connectivity reorganization at the network level.

Recent advances in SNNs have demonstrated their capability to characterize static functional connectivity patterns[6, 7]. Building on this foundation, recent studies have extended SNN frameworks to model dynamic functional and structural connectivity [8, 9]. However, their analytical paradigms are limited to functional connectivity changes between brain regions, overlooking the representation of neuronal dynamics within brain regions. This limits their completeness in modeling brains. For example, Alzheimer's pathological changes often begin with micro-disruptions in synaptic plasticity within the hippocampus, while abnormalities in macro-level functional connectivity may not become apparent for several years [10]. Therefore, by modeling the dynamic characteristics of neuronal clusters within brain regions alongside functional connectivity differences, we can capture brain changes from both microscopic and macroscopic perspectives, providing more accurate biomarkers. However, existing SNN frameworks seldom integrate these two perspectives, *i.e.* network topology and neuronal population responses, within a single computational model, limiting their ability to disentangle multi-scale biomarkers.

To bridge this gap, we propose a Local-Global Coupling Spiking Graph Transformer (LGC-SGT) that jointly models inter-regional connectivity discrepancies and neuronal population firing rate deviations across brain areas, enabling a dual-perspective analysis of neuropathology. Two critical perspectives are provided by our model: 1) disorder-sensitive alterations in inter-regional connectivity strength and topology, and 2) abnormal firing rate responses of neuronal populations within specific brain regions. Therefore, our framework dynamically links macroscale connectivity patterns with microscale-inspired spiking population dynamics. Our spike-based transformer integrates the local-global coupling module, the spiking transformer feature extraction block, and the shortcut-enhanced output strategy with a hybrid loss function. The local-global coupling module captures the spatiotemporal interactions of both aberrant firing rate and abnormal brain connectivity. The local liquid state machine (LSM) pathway guided by functional magnetic resonance imaging (fMRI) initializes region-specific neuronal populations with random connectivity, then updates synaptic weights via spike-timing-dependent plasticity (STDP) to emulate local plasticity, constructing the functional connection pattern. The global adaptive spiking graph convolution pathway refines inter-regional connectivity weights via Backpropagation and integrates global information to capture the neural firing dynamics. Besides, the shortcut-enhanced output module with the hybrid loss function is designed to obtain a more robust decision boundary by suppressing the interference caused by the outliers since the differences between inter-individual and inter-center.

The contributions of our paper could be summarized as follows:

- We propose the local-global coupling spiking graph transformer framework for brain disorder diagnosis. The spike-based transformer integrates local and global pathways, attending to spatiotemporal interactions of both aberrant firing rate dynamics and abnormal brain connectivity.
- The shortcut-enhanced output strategy with the hybrid loss function is designed to obtain a more robust decision boundary by suppressing the interference caused by the outliers since the variety of inter-individual and inter-center.
- The experiments on three brain disorder datasets show that the proposed LGC-SGT model outperforms existing methods and provides the dual-perspective biomarker discovery of interregional connectivity discrepancies and neuronal population firing rate deviations.

## 2 Related Work

### 2.1 Function-based Brain Networks Analysis

With the rapid advancement of neuroimaging technologies, researchers have progressively established a systematic neuroimaging-based paradigm for investigating brain networks [11], with fMRI emerging as a crucial data foundation for constructing brain networks through its unique integration of whole-brain coverage and high spatial resolution [12]. The complexity of modeling the brain's topological relationships has sparked researchers' interest, leading to the development of methods using Graph Neural Networks (GNN) for brain network analysis [13]. Additionally, approaches such as adaptive

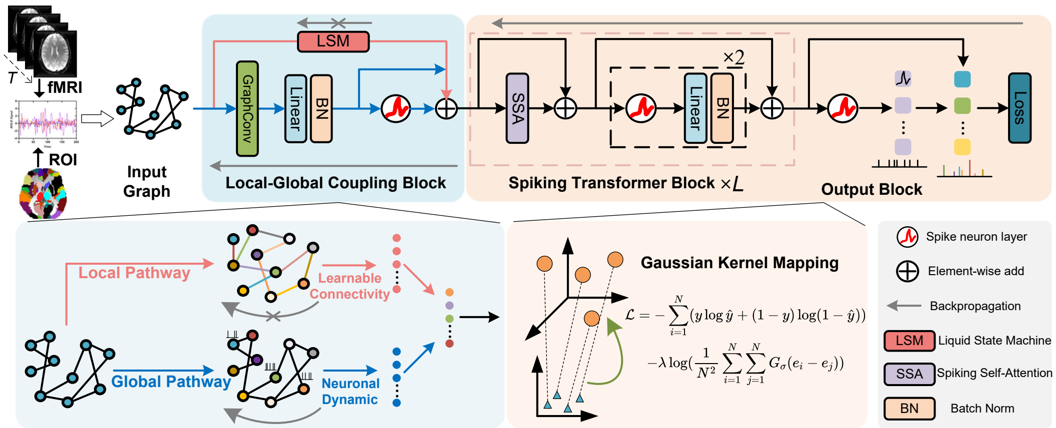


Figure 1: The overall framework of the proposed LGC-SGT model.

brain region discrimination [14], information-theory-driven connectivity selection [15], and long-range dependency modeling [16] have been proposed. However, these methods typically focus on modeling functional connectivity relationships and often overlook the representation of neuronal population dynamics within the brain.

### 2.2 SNN-based Brain Networks Analysis

Recent years have seen growing applications of SNNs in diverse fields [17–19], owing to their biological plausibility and low-power characteristics [20]. In brain cognitive modeling, SNNs enable biologically-based brain network modeling by simulating the dynamic processes of biological neurons. Three-dimensional spiking neural network architectures facilitate the analysis of spatiotemporal neural data through biomimetic topological design [21]. By spatio-temporal associative memories in SNNs, Kasabov et al. [22] achieved brain disorder classification based on fMRI data. The biologically-inspired characteristics of SNNs reveal how mindfulness training reorganizes brain networks [23] and allow for data analysis of the coupling between brain structure and functional connectivity [8, 9]. Additionally, SNNs constrained by fMRI topology uncover auditory coding mechanisms from the perspective of information transfer efficiency [24]. Anatomically-guided spiking networks leverage brain anatomical features to guide SNNs in autonomously learning brain network characteristics [25]. Despite these advancements, current SNN frameworks still have limitations in brain network modeling, particularly in their lack of collaborative modeling that integrates inter-region network topology with the responses of neuronal populations within brain regions.

## 3 Method

### 3.1 Problem Definition

In brain network analysis, the core objective of our model is to predict the presence of a brain disorder or to discover new biomarkers. Here, the network is constructed based on graph-structured data derived from resting-state fMRI data. Formally, we derive a set of graph-structured samples $\left\{\mathcal{G}^1, \cdots, \mathcal{G}^N\right\}$ for $N$ participants and $n$ ROIs, where each sample $\mathcal{G}^i = (A^i, X^i, \mathrm{y}^i)$ denotes the functional architecture of the $i$-th objects. In each functional graph of $\mathcal{G}^i$, $A^i \in \{0, 1\}^{n\times n}$ is the graph adjacency matrix that represents the functional connection relationship between brain ROIs. $X^i \in \mathbb{R}^{n\times n}$ denotes the node feature matrix that captures functional activity intensity of the brain ROIs, with the corresponding label of $\mathrm{y}^i$ that distinguishes brain disorders or healthy controls (HCs). Our following model aims to learn a map $f(A^i, X^i) \rightarrow \mathrm{y}^{\mathrm{i}}$ and discover the biomarkers in the process.

### 3.2 Local-Global Coupling Spiking Graph Transformer

As shown in Fig. 1, our LGC-SGT framework contains the local-global coupling module, the spiking transformer feature extraction block, and the shortcut-enhanced output module with a hybrid loss function. The model begins by constructing a graph embedding that transforms graph-structured brain ROIs feature matrices and functional connectivity matrices into feature embeddings. Establish a mapping of elementary connection patterns to higher-order pathological features. Specifically, we propose a local-global coupling block that synergistically integrates locally constrained biologically inspired features with globally optimized graph convolutional representations. This integration facilitates precise localization of abnormal brain regions and identification of pathology-relevant biomarkers. Next, we construct the spiking transformer block to effectively capture dependencies among features while strengthening their interactive relationships through dynamic attention weighting. The features are processed through a spike layer, after which the model combines the features before the spike layer with the binarized features from the spike layer to enhance the output used for binary classification between HCs and patients with brain disorders.

**Local-Global Coupling Module.** Building upon the multiple level learning plasticity in the human brain, we propose a local-global coupling (LGC) module designed to simultaneously capture local and global features of brain networks. Specifically, the LGC module integrates two parallel pathways: (1) local pathway learned by unsupervised learning via STDP, and (2) global pathway via supervised learning by backpropagation. This dual-path architecture enables synergistic fusion of neuronal dynamics and network organization.

***Local Pathway.*** This pathway mainly consists of three components: an input layer, a reservoir network, and a readout layer. The input layer encodes external signals into spike trains for the reservoir network. The liquid state machine layer comprises spiking neurons with sparse random connectivity, maintaining a biologically plausible excitatory-to-inhibitory ratio of 4:1, which mimics the cortical microcircuit organization [26]. The readout layer, typically implemented as a linear layer, decodes the states of neurons within the reservoir network. The calculation process of LSM can be described as follows:

$$y(t) = f^M((L^M x)(t)), \tag{1}$$

where $x$ is input of LSM, $L^M$ is the nonlinear dynamic system (reservoir network), $f^M$ is the readout function (readout layer), and $y(t)$ is the output of the readout layer at time $t$. The nonlinear dynamic system $L^M$ here is realized with LIF neurons (all spiking neurons in this paper are LIF neurons), given by

$$u^{(t),\,\text{pre}} = \tau u^{(t)} + W x^t, u^{(t)} = u^{(t),\,\text{pre}} \cdot \left(1 - o^{(t)}\right). \tag{2}$$

$$o^{(t)} = \begin{cases} 1 & \text{if } u^{(t),\,\text{pre}} > V_{\text{th}} \\ 0 & \text{otherwise.} \end{cases}, \tag{3}$$

where $u^{(t),\,\text{pre}}$ is the pre-synaptic input at time step $t$, $\tau$ is the leaky factor, $W$ is the weight parameter, $V_{\text{th}}$ is the firing threshold, and $o^{(t)}$ is the spike output. The STDP rule is only used in the local pathway forward propagation, which can be mathematically represented as follows:

$$\Delta W = \begin{cases} A_{\text{LTP}} \cdot e^{-\frac{\Delta t}{\tau_{\text{LTP}}}}, & \text{if } \Delta t > 0 \\ \\ -A_{\text{LTD}} \cdot e^{\frac{\Delta t}{\tau_{\text{LTD}}}}, & \text{if } \Delta t < 0 \end{cases} \tag{4}$$

where $\Delta W$ represents the change in synaptic weight, $\Delta t$ is the time interval between the pre-synaptic and post-synaptic spikes, $A_{\text{LTP}}$ and $A_{\text{LTD}}$ are control constants of long-term potentiation and long-term depression, respectively, and $\tau_{\text{LTP}}$ and $\tau_{\text{LTD}}$ are time constants.

***Global Pathway.*** The local pathways may lead to the loss of original graph structural details. To mitigate this information loss, we introduce a global pathway designed to preserve holistic network attributes. This global pathway employs a spiking graph convolutional network with backpropagation to capture whole-brain dynamics, formally defined by the following equation:

$$H = \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} X. \tag{5}$$

$$O_{Global} = SN(BN(Linear(H))) + H, \tag{6}$$

where $\tilde{A} = A + I$ is the adjacent matrix with self-connection, $\tilde{D}$ is the degree matrix of $\tilde{A}$, $H$ is the new node representation with connected aggregation, and $O_{Global}$ is the output of the global pathway. $SN$ represents the spiking layer and $BN$ represents the Batch Norm layer.

**Spiking Transformer Blocks.** We employ a spiking transformer [27] as the backbone network for learning features. The core of the spiking transformer is spiking self-attention (SSA). Specifically, given the feature, it first computes the SSA, given by

$$Q(X) = SN(Linear_q(X)), K(X) = SN(Linear_k(X)), V(X) = SN(Linear_v(X)). \tag{7}$$

$$SSA(X) = MLP(Q(X) \cdot SN(K^T(X) \odot V(X))). \tag{8}$$

After determining the attention, a shortcut message aggregation scheme is applied from the features to the attention, followed by another SNN-enhanced MLP block with the shortcut. The following equations can describe the overall calculation process:

$$Attn^l = SSA(X^l) + X^l, l = 1 \cdots L. \tag{9}$$

$$O^l = MLP(SN(Attn^l))) + Attn^l, l = 1 \cdots L. \tag{10}$$

**Shortcut-enhanced Output Strategy.** Due to the discrete spike firing mechanism employed by SNNs, the step function is non-differentiable. Currently, surrogate gradient methods are widely adopted to address this challenge. However, as the network depth increases, employing surrogate gradients inevitably leads to approximation errors and gradient vanishing. Typically, shortcuts can be introduced to alleviate gradient vanishing and reduce the information loss [28], but this causes a conflict between the real-valued features in intermediate layers and the binarization of SNNs. To address this, researchers have proposed placing spiking activation functions at the beginning of each block to maintain local binarization [29], while enforcing a final spiking layer at the end of the network to achieve global binarization, thus ensuring the network remains purely event-driven. However, binary spike outputs result in information loss. To mitigate the information loss caused by binarization, particularly its impact on classification performance, we propose a shortcut-enhanced output strategy by mixing the features before the spiking layer with their subsequent binarized features for the final output:

$$O_{\text{final}} = O^l \cdot SN_{\text{final}}(O^l), \tag{11}$$

where $O^l$represents the continuous features before the spiking layer, and $SN_{\text{final}}(O^l)$ converts the continuous features into spiking features. The final output is obtained by multiplying the two, where the spiking signals represent important features, and the real-valued coefficients determine the intensity of the features.

### 3.3 Minimum Error Entropy Criterion

Significant individual differences exist between patients and HCs, and inter-site variability caused by differences in equipment and operational procedures further complicates the data. Traditional loss functions based on cross-entropy (CE) may lack robustness when dealing with such complex data distributions. To address this limitation, we introduce minimum error entropy (MEE) [30] as the extra optimization objective of the model. By minimizing the information entropy of the classification errors, MEE effectively suppresses the interference caused by outliers arising from individual and site differences while reducing reliance on assumptions about specific data distributions.

In cross-entropy loss, the optimization objective essentially measures the difference between two probability distributions, with the optimization performed independently on a point-by-point basis. In contrast, MEE takes into account the statistical dependencies among the error samples. The error information $e$ can be measured by Rényi's entropy [31]

$$H_\alpha(e) = \frac{1}{1-\alpha} \log V_\alpha(e), \tag{12}$$

where error $e = \hat{y} - y$ between target $\hat{y}$ and model output $y$, $\alpha$ is the order of entropy, usually $\alpha = 2$ for the convenience of calculation, and $V_\alpha(e)$ is the information potential, given by

$$\hat{V}_2(e) = \int \hat{p}^2(e) = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} G_\sigma(e_i - e_j). \tag{13}$$

$$\hat{p}(x) = \frac{1}{N} \sum_{i=1}^{N} G_\sigma(x - e_i),\ G_\sigma(e) = \exp\left(-\frac{e^2}{2\sigma^2}\right), \tag{14}$$

where $G_\sigma$ is Gaussian kernel with bandwidth $\sigma$, and $p(:)$ is the probability density function of error $e$. Obviously, minimizing the error entropy criterion $H_2(e)$ is the same as maximum information potential $\hat{V}_2(e)$. Hereby, we define the MEE-Loss as:

$$\mathcal{L}_{MEE} = \hat{H}_2(e) = -\log(\frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} G_\sigma(e_i - e_j)). \tag{15}$$

Then the overall loss can be calculated as follows:

$$\mathcal{L}_{Total} = \mathcal{L}_{CE} + \lambda\mathcal{L}_{MEE}. \tag{16}$$

**Theoretical Analysis** To validate the robustness of MEE-loss, we conducted a theoretical analysis focusing on gradient perturbations caused by outliers. The CE-loss is defined as follows:

$$\mathcal{L}_{\mathrm{CE}} = -\sum_{i=1}^{C} p_i \log q_i,\ q_i = \frac{e^{z_i}}{\sum_{j=1}^{C} e^{z_j}},\ z_i = f(x;\theta). \tag{17}$$

where $p_i$ is the target distribution, $q_i$ is the predict distribution, $f$ is the network mapping, $z_i$ is the network output for category $i$, $x$ is input data, $\theta$ is the learnable network parameter. Next, we compute the gradient:

$$\nabla_\theta \mathcal{L}_{\mathrm{CE}} = -\sum_{i=1}^{C} \frac{\partial \mathcal{L}_{\mathrm{CE}}}{\partial z_i} \cdot \nabla_\theta z_i = -\sum_{i=1}^{C} (p_i - q_i) \cdot \nabla_\theta z_i. \tag{18}$$

Similarly, the gradient of MEE-loss is given by:

$$\nabla_\theta J_{\mathrm{MEE}} = -\log(\frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} G'_\sigma\,(e_i - e_j) \cdot (\nabla_\theta e_i - \nabla_\theta e_j)). \tag{19}$$

If there are outliers in the input, the corresponding network will produce abnormal outputs. The gradient of CE is linearly related to $z_i$, such that $\nabla_\theta z_i \uparrow$ as $|z_i| \uparrow$. As outliers increase, their influence on parameter updates grows linearly, directly affecting the gradient update. In contrast, for MEE, the Gaussian kernel mapping results in a decrease in the contribution of the kernel function as outliers increase: $G'_\sigma\,(e_i - e_j) \downarrow$ as $|e_i - e_j| \uparrow$. This leads to the contribution of outliers to the gradient approaching zero.

# 4 Experiments

In this section, we analyze the following aspects to demonstrate the effectiveness of our proposed LGC-SGT model and its dual-perspective biomarker discovery capability.
**Q1.** How does the performance of the model compare to other state-of-the-art baseline models?
**Q2.** Does the model detect abnormal firing rate responses of specific neuronal populations?
**Q3.** Does the model identify the functional connectivity changes in disorder-sensitive regions?

## 4.1 Experimental Settings

**Datasets, Data Preprocessing, and Implementation Details.** We evaluate the proposed LGC-SGT model using three brain network analysis-related fMRI datasets. (1) The ABIDE dataset [1], which contains 528 patients with autism spectrum disorder (ASD) and 571 HCs. (2) The REST-meta-MDD dataset[2], which contains 848 patients with MDD and 794 HCs. (3) The SRPBS dataset[3], which contains 92 patients with schizophrenia and 92 HCs. The details of datasets, data preprocessing, and experimental implementation can be found in Appendix A.

[1]`http://preprocessed-connectomes-project.org/abide/`
[2]`https://rfmri.org/maps`
[3]`https://bicr-resource.atr.jp/srpbsfc/`

**Baseline Models.** The selected baselines correspond to two categories. The first category is ANNs, including generalized graph networks and specialized brain networks. Specifically, these include GCN [32], GAT [33], GIN [34], SIB [35], DIR-GNN [36], ProtGNN [37], BrainGNN [13], IBGNN [38], CI-GNN [39], BrainIB [15], ContrastPool [14], and ALTER [16]. The second category is the SNNs, including SpikingGCN [40], MSG [41], SiGNN [42], and SpikingGT [27].

### 4.2 Performance Comparison (Q1)

**Tenfold Cross-Validation.** Table 1 presents a comparison between the proposed LGC-SGT model and the baseline models. The LGC-SGT significantly outperforms both categories of baseline methods across all three datasets. ***(a) Compared to ANNs.*** For generalized graph networks, our model demonstrates notable improvements in the accuracy metric, achieving a 6.2% increase on the REST-meta-MDD dataset, a 5.1% increase on the ABIDE dataset, and a 9.2% increase on the SRPBS dataset. For specialized brain networks, we also observe superior performance, with ACC improvements of 0.9%, 2.5%, and 3.2% on REST-meta-MDD, ABIDE, and SRPBS, respectively. ***(b) Compared to SNNs.*** Our model achieves improvements of 4.9%, 4.3%, and 4.8% on REST-meta-MDD, ABIDE, and SRPBS datasets, respectively. These experimental results indicate that LGC-SGT consistently outperforms the existing baseline methods across all evaluated datasets.

Table 1: Tenfold cross-validation performance with ANNs and SNNs on three datasets (REST-meta-MDD, ABIDE, and SRPBS) (%). The best results are marked in **bold**.

| Category | | Methods | REST-meta-MDD | ABIDE | SRPBS |
|---|---|---|---|---|---|
| ANN | Generalized | GCN [32] | $60.8 \pm 1.3$ | $64.4 \pm 3.8$ | $82.7 \pm 8.6$ |
| | | GAT [33] | $64.7 \pm 1.7$ | $67.9 \pm 3.9$ | $84.2 \pm 8.1$ |
| | | GIN [34] | $65.4 \pm 3.2$ | $67.9 \pm 3.5$ | $84.3 \pm 8.2$ |
| | | SIB [35] | $57.7 \pm 3.2$ | $62.7 \pm 4.4$ | $81.1 \pm 9.0$ |
| | | DIR-GNN [36] | $64.4 \pm 1.8$ | $67.0 \pm 2.7$ | $83.9 \pm 6.5$ |
| | | ProtGNN [37] | $61.0 \pm 2.1$ | $65.3 \pm 3.1$ | $84.3 \pm 5.0$ |
| | Specialized | BrainGNN [13] | $60.2 \pm 3.0$ | $62.7 \pm 2.1$ | $82.2 \pm 8.5$ |
| | | IBGNN [38] | $63.0 \pm 2.7$ | $64.0 \pm 3.1$ | $84.8 \pm 8.6$ |
| | | CI-GNN [39] | $66.8 \pm 4.0$ | $67.5 \pm 3.3$ | $87.5 \pm 6.8$ |
| | | BrainIB [15] | $70.0 \pm 2.2$ | $70.2 \pm 2.0$ | $90.3 \pm 4.6$ |
| | | ContrastPooL [14] | $65.1 \pm 1.9$ | $68.6 \pm 2.7$ | $89.2 \pm 3.4$ |
| | | ALTER [16] | $65.8 \pm 3.3$ | $70.5 \pm 1.4$ | $89.5 \pm 5.9$ |
| SNN | Generalized | SpikeGCN [40] | $66.0 \pm 2.1$ | $68.7 \pm 1.6$ | $86.5 \pm 2.5$ |
| | | MSG [41] | $60.4 \pm 2.4$ | $62.9 \pm 3.7$ | $81.9 \pm 2.7$ |
| | | SiGNN [42] | $66.2 \pm 2.5$ | $69.1 \pm 1.8$ | $87.6 \pm 7.2$ |
| | | SpikeGT [27] | $61.8 \pm 3.5$ | $65.8 \pm 1.4$ | $88.7 \pm 3.6$ |
| | Specialized | LGC-SGT (Ours) | $\mathbf{70.9 \pm 1.7}$ | $\mathbf{73.0 \pm 2.6}$ | $\mathbf{93.5 \pm 3.0}$ |

Table 2: Leave-one-site-out cross-validation performance on REST-meta-MDD and ABIDE dataset (%).

| Dataset | Methods | S1 | S2 | S3 | S4 | S5 | S6 | S7 | S8 | S9 | S10 | S11 | S12 | S13 | S14 | S15 | S16 | S17 | Mean |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| REST-meta-MDD | CI-GNN | 63.3 | **83.0** | 85.2 | 76.3 | 70.0 | **68.8** | 75.4 | 73.2 | 72.2 | 68.3 | 81.1 | 62.3 | 67.3 | **63.2** | 68.3 | **75.2** | 64.4 | 71.6 |
| | BrainIB | 63.3 | 73.0 | **85.4** | **77.8** | 71.3 | **68.8** | 73.2 | 75.7 | **80.6** | 72.0 | **82.1** | **67.1** | **69.4** | **63.2** | 70.1 | 71.1 | 68.9 | 72.5 |
| | SpikeGT | 55.5 | 56.7 | 75.6 | 62.5 | 70.1 | 62.5 | 70.4 | 62.2 | 61.1 | 65.6 | 73.1 | 57.3 | 63.3 | 58.1 | 63.9 | 68.4 | 64.4 | 64.2 |
| | LGC-SGT | **63.7** | 70.0 | **85.4** | **77.8** | **74.7** | 68.4 | **76.1** | **78.4** | **80.6** | **75.3** | **82.1** | 62.2 | **69.4** | 60.4 | **70.8** | 68.9 | **78.9** | **73.1** |
| ABIDE | CI-GNN | 83.3 | 71.1 | 67.4 | 67.8 | 71.7 | 64.3 | **67.9** | 77.8 | **78.1** | 69.2 | **83.3** | **75.0** | 64.2 | 63.4 | 61.2 | 66.3 | 75.1 | 71.0 |
| | BrainIB | 83.3 | 71.1 | **72.7** | **73.4** | 66.7 | **70.1** | **67.9** | 77.8 | 66.7 | **83.3** | 75.0 | **75.0** | 65.3 | 74.7 | 64.8 | 73.3 | 82.1 | 73.1 |
| | SpikeGT | 79.2 | 57.9 | **72.7** | 67.2 | 61.4 | 69.0 | 53.6 | 72.2 | 56.1 | 70.0 | 72.2 | 70.0 | 69.4 | 68.7 | 63.5 | 74.3 | 78.6 | 68.0 |
| | LGC-SGT | **87.5** | **76.3** | 70.9 | **73.4** | **73.7** | **70.1** | 64.3 | **80.6** | 70.2 | 73.3 | 77.8 | 72.5 | **73.5** | **74.8** | **67.6** | **75.3** | **85.7** | **74.6** |

**Leave-One-Site-Out Cross-Validation.** To further assess the performance of our model, we conduct leave-one-site-out cross-validation. ABIDE and REST-meta-MDD datasets contain 17

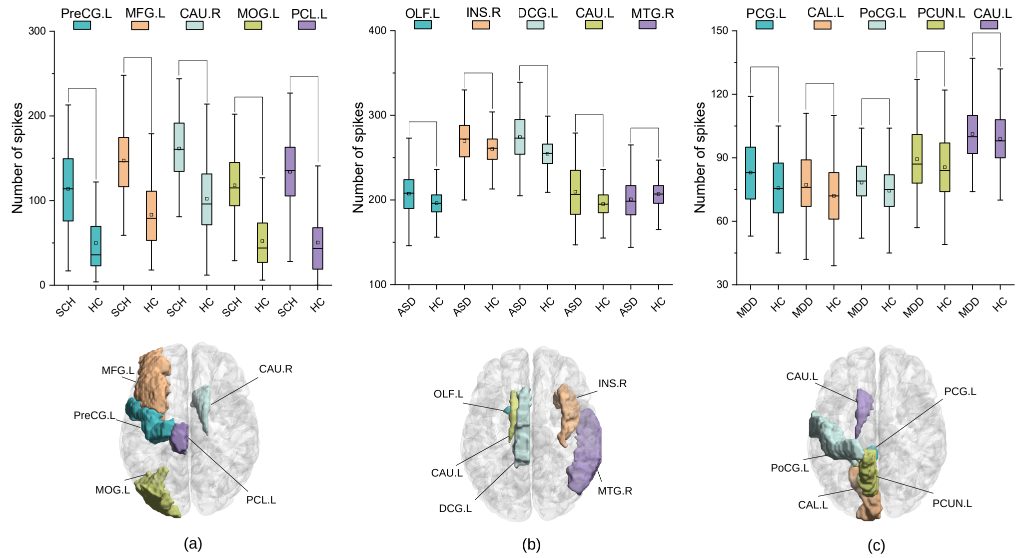


Figure 2: Comparison of brain region activity of HCs and patients on schizophrenia, ASD, and MDD. The top 5 different regions are shown. The two-sample $t$-test with Welch's correction was employed to evaluate the significance of intergroup differences. Statistical significance was defined as $p < 0.05$. (a) Schizophrenia. (b) ASD. (c) MDD.

independent sites. In detail, we divide each dataset into the training set (16 sites out of 17 sites) to train the model and the testing set (remaining site out of 17 sites) for testing the model. We compared our model with CI-GNN, BrainIB, and SpikeGT. CI-GNN and BrainIB were chosen due to their performance better in tenfold cross-validation compared to other models, while SpikeGT represents a typical transformer-based SNN model. The experimental results are summarized in Table 2. Compared to the baselines, our model achieved the highest site-average accuracy, reaching 73.1% and 74.6% on REST-meta-MDD and ABIDE datasets, respectively. Moreover, it demonstrates significant improvements across most sites, indicating that the strength of our model lies in its ability to generalize well across completely different multi-site datasets despite site-specific differences.

Table 3: Ablation study for the modules we design.

| Dataset | Method | Accuracy |
|---|---|---|
| REST-meta-MDD | Baseline | **70.94%** |
| | w/o MEE | 70.77% |
| | w/o LGC | 69.99% |
| | w/o SEO | 69.83% |
| ABIDE | Baseline | **73.03%** |
| | w/o MEE | 72.91% |
| | w/o LGC | 72.55% |
| | w/o SEO | 71.73% |
| SRPBS | Baseline | **93.51%** |
| | w/o MEE | 90.50% |
| | w/o LGC | 89.77% |
| | w/o SEO | 89.74% |

**Ablation Study.** To rigorously evaluate the contribution of each proposed component to our model's performance, we conduct comprehensive ablation studies across three benchmark datasets. The experimental design systematically isolates three critical elements: (i) the MEE-loss function, (ii) the LGC module, and (iii) the shortcut-enhanced output strategy (SEO). As shown in Table 3, the LGC-SGT achieves the highest classification accuracy on three brain disorders classification tasks, demonstrating the effectiveness of the proposed modules.

## 5 In-depth Analysis of brain neuronal population dynamics (Q2)

To analyze the neuronal population dynamics of our model, we conducted rigorous statistical comparisons. The results show that the differences in spiking activity between patients with brain disorders and HCs are statistically significant. Fig. 2 illustrates the five ROIs exhibiting maximal inter-group differences for each disorder, accompanied by their neuroanatomical schematics. ***Schizophrenia***: Patients demonstrated aberrant activity in several brain regions, including the PreCG.L, MFG.L, PCL.L, CAU.R, and MOG.L. These findings align with prior studies [43–45] that abnormal activity in

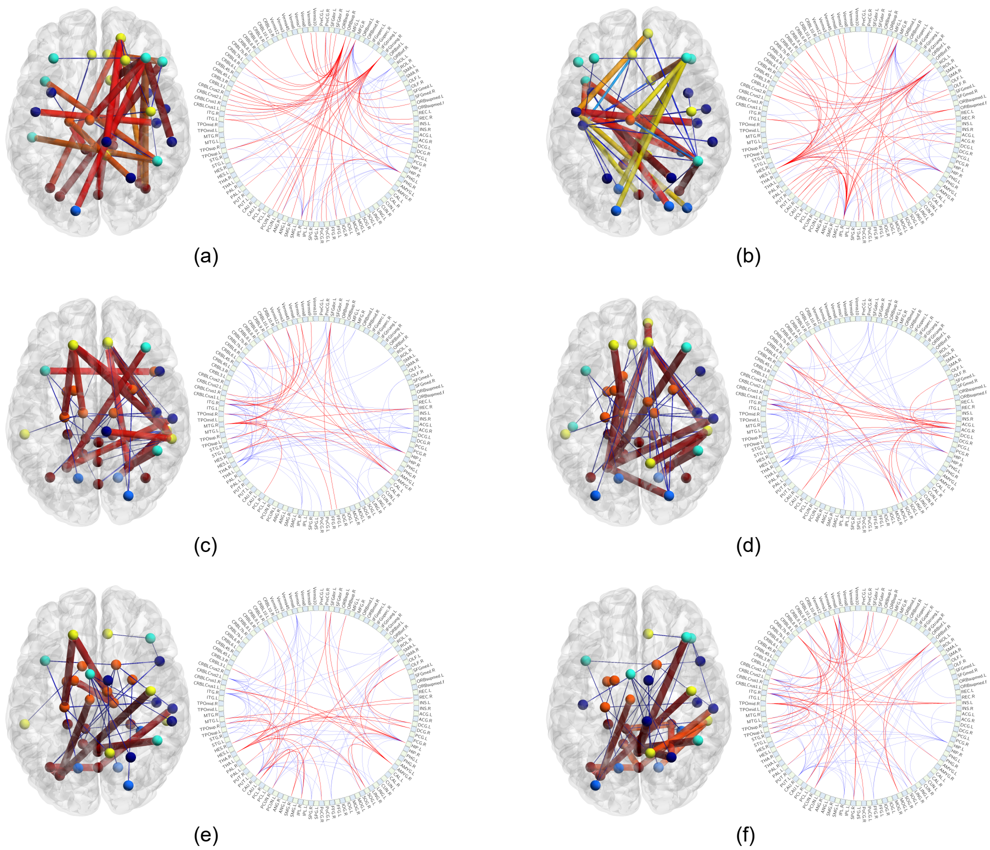


Figure 3: Comparison of brain functional connections of HCs and patients with schizophrenia. In the anatomical brain map, the thickness of the connections represents the strength of functional connectivity; in the circular layout chord diagram, blue lines indicate shared functional connections between the patient and HCs, while red lines represent connections unique to either the patient or HCs. (a) Patients with schizophrenia. (b) HCs of schizophrenia. (c) Patients with ASD. (d) HCs of ASD. (e) Patients with MDD. (f) HCs of MDD.

these regions serves as a potential cause of schizophrenia. These regions are primarily involved in the maintenance of working memory, conflict monitoring, and the regulation of goal-directed behavior. ***ASD***: significant activity differences emerge in the OLF.L, INS.R, DCG.L, CAU.L, and MTG.R. Supporting literature [46–48] further suggests that abnormalities in these regions are closely related to the core symptoms characteristic of ASD. These regions are responsible for detecting internal sensory signals and coordinating the allocation of attentional resources. ***MDD***: patients showed significant activity alterations in the PCG.L, CAL.L, PoCG.L, PCUN.L, and CAU.L. Consistent with these observations, previous studies [49, 50] have confirmed that abnormal activity in these brain regions may underlie the development of major depressive disorder. Abnormalities in these brain regions may lead to a bias in the processing of negative information.

Comparative analysis across three brain disorders revealed distinct neural activity patterns of aberrant brain regions when contrasted with HCs. These results indicate that our model effectively simulates brain activity and identifies regions exhibiting abnormal neural activity in patients compared to HCs. Clinical and neuroimaging results also validate the rationale of the biomarkers we discovered. The detailed experimental setting can be found in Appendix B.1.

## 6 In-depth Analysis of brain functional connectivity (Q3)

To analyze the functional pattern under different brain disorders, we visualize functional connectivity. The results show that the functional connectivity between patients with brain disorders and HCs is significantly different. As shown in Fig. 3, our model learned distinct connectivity patterns for schizophrenia patients and HCs, and obtained some functional connections that could distinguish patients from HCs. ***Schizophrenia***: functional connectivity was significantly increased between the MFG.R and IFGtriang.R (frontoparietal network). This finding is consistent with previous studies [51] reporting that the connectivity between the sensorimotor network and the frontoparietal network is enhanced in schizophrenia. ***ASD***: the increased connectivity pointing to the HIP.L region, which aligns with prior research indicating functional abnormalities of the hippocampus in ASD [52]. ***MDD***: connectivity between the THA.R and HIP.L has been increased, consistent with findings from [53] that connectivity change in these regions is part of the pathophysiology of MDD.

The experimental results demonstrate that our model accurately reproduces the abnormal functional connectivity patterns observed clinically, effectively simulates the physiological mechanisms underlying brain functional connectivity, and reveals distinct connectivity patterns characteristic of brain disorders. Clinical and neuroimaging results also validate the rationale of the biomarkers we discovered. The detail of the experiment is provided in the Appendix B.2.

## 7 Conclusion and Discussions

**Conclusion.** In summary, we propose the LGC-SGT model that jointly models inter-regional connectivity discrepancies and neuronal population firing rate deviations, enabling a dual-perspective analysis of neuropathology. By integrating a shortcut-enhanced output strategy with a hybrid loss function, our framework achieves robust decision boundaries for brain disorder diagnosis. Extensive experiments on three brain disorder datasets demonstrate that the LGC-SGT model consistently outperforms existing graph learning methods, successfully identifying dual-perspective biomarkers from both inter-regional connectivity and neuronal firing rate abnormalities. This approach offers novel insights for clinical research and therapeutic strategies.

**Limitations.** Although our current work effectively simulates brain functional states using resting-state fMRI data, incorporating multimodal data, such as diffusion tensor imaging (DTI) for structural connectivity and electroencephalogram (EEG) for neural oscillations, would provide additional anatomical constraints and temporal resolution, enabling more accurate brain modeling grounded in biological and clinical relevance.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant No. 62436005, U21A20485, 62306274), China Postdoctoral Science Foundation under Grant Number (No. GZB20250394), Young Talent Fund of Xi'an Association for Science and Technology (0959202513037).

## References

[1] Alex Fornito, Andrew Zalesky, and Michael Breakspear. The connectomics of brain disorders. *Nature Reviews Neuroscience*, 16(3):159–172, 2015.

[2] Nick Corriveau-Lecavalier, Jenna N Adams, Larissa Fischer, Eóin N Molloy, and Anne Maass. Cerebral hyperactivation across the alzheimer's disease pathological cascade. *Brain Communications*, 6(6):fcae376, 2024.

[3] Kenneth D Harris and Thomas D Mrsic-Flogel. Cortical connectivity and sensory coding. *Nature*, 503(7474):51–58, 2013.

[4] Sonja Sudimac, Vera Sale, and Simone Kühn. How nature nurtures: Amygdala activity decreases as the result of a one-hour walk in nature. *Molecular psychiatry*, 27(11):4446–4452, 2022.

[5] Thomas Prévot and Etienne Sibille. Altered gaba-mediated information processing and cognitive dysfunctions in depression and other brain disorders. *Molecular psychiatry*, 26(1):151–167, 2021.

[6] Nikola Kasabov, Lei Zhou, Maryam Gholami Doborjeh, Zohreh Gholami Doborjeh, and Jie Yang. New algorithms for encoding, learning and classification of fmri data in a spiking neural network architecture: A case on modeling and understanding of dynamic cognitive processes. *IEEE Transactions on Cognitive and Developmental Systems*, 9(4):293–303, 2016.

[7] Norhanifah Murli, Nikola Kasabov, and Bana Handaga. Classification of fmri data in the neucube evolving spiking neural network architecture. In *Neural Information Processing: 21st International Conference, ICONIP 2014, Kuching, Malaysia, November 3-6, 2014. Proceedings, Part I 21*, pages 421–428. Springer, 2014.

[8] Nikola K Kasabov, Maryam Gholami Doborjeh, and Zohreh Gholami Doborjeh. Mapping, learning, visualization, classification, and understanding of fmri data in the neucube evolving spatiotemporal data machine of spiking neural networks. *IEEE transactions on neural networks and learning systems*, 28(4):887–899, 2016.

[9] Shaolong Wei, Jiashuang Huang, Mingliang Wang, Shu Jiang, and Weiping Ding. Scnn: Spike coupling neural network for multimodal brain network analysis. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2025.

[10] Jorge J Palop and Lennart Mucke. Network abnormalities and interneuron dysfunction in alzheimer disease. *Nature Reviews Neuroscience*, 17(12):777–792, 2016.

[11] Alisa M Loosen, Ayaka Kato, and Xiaosi Gu. Revisiting the role of computational neuroimaging in the era of integrative neuroscience. *Neuropsychopharmacology*, 50(1):103–113, 2025.

[12] Michael D Fox and Marcus E Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8(12):700–711, 2007.

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

[14] Jiaxing Xu, Qingtian Bian, Xinhang Li, Aihu Zhang, Yiping Ke, Miao Qiao, Wei Zhang, Wei Khang Jeremy Sim, and Balázs Gulyás. Contrastive graph pooling for explainable classification of brain networks. *IEEE Transactions on Medical Imaging*, 2024.

[15] Kaizhong Zheng, Shujian Yu, Baojuan Li, Robert Jenssen, and Badong Chen. Brainib: Interpretable brain network-based psychiatric diagnosis with graph information bottleneck. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.

[16] Shuo Yu, Shan Jin, Ming Li, Tabinda Sarwar, and Feng Xia. Long-range brain graph transformer. *Advances in Neural Information Processing Systems*, 37:24472–24495, 2024.

[17] Tong Bu, Maohua Li, and Zhaofei Yu. Inference-scale complexity in ann-snn conversion for high-performance and low-power applications. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24387–24397, 2025.

[18] Jiangrong Shen, Yulin Xie, Qi Xu, Gang Pan, Huajin Tang, and Badong Chen. Spiking neural networks with temporal attention-guided adaptive fusion for imbalanced multi-modal learning. *arXiv preprint arXiv:2505.14535*, 2025.

[19] Qi Xu, Jie Deng, Jiangrong Shen, Biwu Chen, Huajin Tang, and Gang Pan. Hybrid spiking vision transformer for object detection with event cameras. In *Forty-second International Conference on Machine Learning*, 2025.

[20] De Ma, Xiaofei Jin, Shichun Sun, Yitao Li, Xundong Wu, Youneng Hu, Fangchao Yang, Huajin Tang, Xiaolei Zhu, Peng Lin, et al. Darwin3: a large-scale neuromorphic chip with a novel isa and on-chip learning. *National Science Review*, 11(5):nwae102, 2024.

[21] Nikola K Kasabov. Neucube: A spiking neural network architecture for mapping, learning and understanding of spatio-temporal brain data. *Neural networks*, 52:62–76, 2014.

[22] Nikola K Kasabov, Helena Bahrami, Maryam Doborjeh, and Alan Wang. Brain-inspired spatio-temporal associative memories for neuroimaging data classification: Eeg and fmri. *Bioengineering*, 10(12):1341, 2023.

[23] Zohreh Doborjeh, Maryam Doborjeh, Tamasin Taylor, Nikola Kasabov, Grace Y Wang, Richard Siegert, and Alex Sumich. Spiking neural network modelling approach reveals how mindfulness training rewires the brain. *Scientific reports*, 9(1):6367, 2019.

[24] Yihua Song, Lei Guo, Menghua Man, and Youxi Wu. The spiking neural network based on fmri for speech recognition. *Pattern Recognition*, 155:110672, 2024.

[25] Jiashuang Huang, Mingliang Wang, Hengrong Ju, Weiping Ding, and Daoqiang Zhang. Agbn-transformer: Anatomy-guided brain network transformer for schizophrenia diagnosis. *Biomedical Signal Processing and Control*, 102:107226, 2025.

[26] Lionel G Nowak, Maria V Sanchez-Vives, and David A McCormick. Lack of orientation and direction selectivity in a subgroup of fast-spiking inhibitory interneurons: cellular and synaptic mechanisms and comparison with other electrophysiological cell types. *Cerebral cortex*, 18(5): 1058–1078, 2008.

[27] Yundong Sun, Dongjie Zhu, Yansong Wang, Zhaoshuo Tian, Ning Cao, and Gregory O'Hared. Spikegraphormer: A high-performance graph transformer with spiking graph attention. *arXiv preprint arXiv:2403.15480*, 2024.

[28] Wei Fang, Zhaofei Yu, Yanqi Chen, Tiejun Huang, Timothée Masquelier, and Yonghong Tian. Deep residual learning in spiking neural networks. *Advances in Neural Information Processing Systems*, 34:21056–21069, 2021.

[29] Yifan Hu, Lei Deng, Yujie Wu, Man Yao, and Guoqi Li. Advancing spiking neural networks toward deep residual learning. *IEEE transactions on neural networks and learning systems*, 36 (2):2353–2367, 2024.

[30] Yuqing Xie, Yingsong Li, Yuantao Gu, Jiuwen Cao, and Badong Chen. Fixed-point minimum error entropy with fiducial points. *IEEE Transactions on Signal Processing*, 68:3824–3833, 2020.

[31] PA Bromiley, NA Thacker, and E Bouhova-Thacker. Shannon entropy, renyi entropy, and information. *Statistics and Inf. Series (2004-004)*, 9(2004):2–8, 2004.

[32] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[33] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. Graph attention networks. *stat*, 1050(20):10–48550, 2017.

[34] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.

[35] Junchi Yu, Tingyang Xu, Yu Rong, Yatao Bian, Junzhou Huang, and Ran He. Recognizing predictive substructures with subgraph information bottleneck. *IEEE transactions on pattern analysis and machine intelligence*, 46(3):1650–1663, 2021.

[36] Ying-Xin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. Discovering invariant rationales for graph neural networks. *arXiv preprint arXiv:2201.12872*, 2022.

[37] Zaixi Zhang, Qi Liu, Hao Wang, Chengqiang Lu, and Cheekong Lee. Protgnn: Towards self-explaining graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 9127–9135, 2022.

[38] Hejie Cui, Wei Dai, Yanqiao Zhu, Xiaoxiao Li, Lifang He, and Carl Yang. Interpretable graph neural networks for connectome-based brain disorder analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 375–385. Springer, 2022.

[39] Kaizhong Zheng, Shujian Yu, and Badong Chen. Ci-gnn: A granger causality-inspired graph neural network for interpretable brain network-based psychiatric diagnosis. *Neural Networks*, 172:106147, 2024.

[40] Zulun Zhu, Jiaying Peng, Jintang Li, Liang Chen, Qi Yu, and Siqiang Luo. Spiking graph convolutional networks. *arXiv preprint arXiv:2205.02767*, 2022.

[41] Li Sun, Zhenhao Huang, Qiqi Wan, Hao Peng, and Philip S Yu. Spiking graph neural network on riemannian manifolds. *Advances in Neural Information Processing Systems*, 37:34025–34055, 2024.

[42] Dong Chen, Shuai Zheng, Muhao Xu, Zhenfeng Zhu, and Yao Zhao. Signn: A spike-induced graph neural network for dynamic graph representation learning. *Pattern Recognition*, 158: 111026, 2025.

[43] Debo Dong, Dezhong Yao, Yulin Wang, Seok-Jun Hong, Sarah Genon, Fei Xin, Kyesam Jung, Hui He, Xuebin Chang, Mingjun Duan, et al. Compressed sensorimotor-to-transmodal hierarchical organization in schizophrenia. *Psychological medicine*, 53(3):771–784, 2023.

[44] Kynon JM Benjamin, Ria Arora, Arthur S Feltrin, Geo Pertea, Hunter H Giles, Joshua M Stolz, Laura D'Ignazio, Leonardo Collado-Torres, Joo Heon Shin, William S Ulrich, et al. Sex affects transcriptional associations with schizophrenia across the dorsolateral prefrontal cortex, hippocampus, and caudate nucleus. *Nature communications*, 15(1):3980, 2024.

[45] Long-Biao Cui, Shu-Wan Zhao, Ya-Hong Zhang, Kun Chen, Yu-Fei Fu, Ting Qi, Mengya Wang, Jing-Wen Fan, Yue-Wen Gu, Xiao-Fan Liu, et al. Associated transcriptional, brain and clinical variations in schizophrenia. *Nature Mental Health*, 2(10):1239–1249, 2024.

[46] Liron Rozenkrantz, Ditza Zachor, Iris Heller, Anton Plotkin, Aharon Weissbrod, Kobi Snitz, Lavi Secundo, and Noam Sobel. A mechanistic link between olfaction and autism spectrum disorder. *Current biology*, 25(14):1904–1910, 2015.

[47] Nadine Gogolla. The insular cortex. *Current Biology*, 27(12):R580–R586, 2017.

[48] Jinping Xu, Chao Wang, Ziyun Xu, Tian Li, Fangfang Chen, Kai Chen, Jingjing Gao, Jiaojian Wang, and Qingmao Hu. Specific functional connectivity patterns of middle temporal gyrus subregions in children and adults with autism spectrum disorder. *Autism Research*, 13(3): 410–422, 2020.

[49] Sugai Liang, Wei Deng, Xiaojing Li, Andrew J Greenshaw, Qiang Wang, Mingli Li, Xiaohong Ma, Tong-Jian Bai, Qi-Jing Bo, Jun Cao, et al. Biotypes of major depressive disorder: Neuroimaging evidence from resting-state default mode network patterns. *NeuroImage: Clinical*, 28:102514, 2020.

[50] Brett L Foster, Seth R Koslov, Lyndsey Aponik-Gremillion, Megan E Monko, Benjamin Y Hayden, and Sarah R Heilbronner. A tripartite view of the posterior cingulate cortex. *Nature Reviews Neuroscience*, 24(3):173–189, 2023.

[51] Linlin Fan, Miao Yu, Amy Pinkham, Yiyi Zhu, Xiaowei Tang, Xiang Wang, Xiaobin Zhang, Junji Ma, Jinbo Zhang, Xiangrong Zhang, et al. Aberrant large-scale brain modules in deficit and non-deficit schizophrenia. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, 113:110461, 2022.

[52] Sarah M Banker, Xiaosi Gu, Daniela Schiller, and Jennifer H Foss-Feig. Hippocampal contributions to social and cognitive deficits in autism spectrum disorder. *Trends in neurosciences*, 44 (10):793–807, 2021.

[53] Weijian Liu, Jurjen Heij, Shu Liu, Luka Liebrand, Matthan Caan, Wietske van der Zwaag, Dick J Veltman, Lin Lu, Moji Aghajani, and Guido van Wingen. Hippocampal, thalamic, and amygdala subfield morphology in major depressive disorder: an ultra-high resolution mri study at 7-tesla. *European archives of psychiatry and clinical neuroscience*, pages 1–17, 2024.

## NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Accurately for sure.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The limitation is described in the conclusion and discussion section.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: The detailed theoretical analysis is provided in the paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The detailed experimental settings are provided in the paper.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We provide the dataset citation and the experimental settings.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes] ,

   Justification: All the training and test details are provided.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: We ensure the experimental results are convincing by conducting several runs.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Yes.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Yes.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

    Justification: There is no societal impact of the work performed.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: the paper poses no such risks.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The citation has been provided in the paper.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The new model is provided in the paper.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer:[NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.

# A Datasets and Implementation Details

## A.1 Datasets and Data Preprocessing

In the construction of brain connectomes, we first preprocessed the raw resting-state fMRI data using Statistical Parametric Mapping (SPM) software, including slice timing correction and head motion correction. Subsequently, the resting-state fMRI images were normalized to a standard space using deformation parameters that map the fMRI images to the Montreal Neurological Institute (MNI) template. In addition, a Gaussian filter with a half maximum width of 6 mm is used to smooth the functional images. The resulting fMRI images are filtered with a bandpass filter (0.01–0.08 Hz). After preprocessing, the average time series for each participant is extracted from each ROIs using the AAL atlas, which comprises 90 cerebral regions and 26 cerebellar regions. Then, functional connectivity was computed between all ROIs pairs using Pearson's correlation coefficient, followed by Fisher's r-to-z transformation, resulting in a 116 × 116 symmetric matrix representing the brain network. Their demographic information is summarized in Table 4.

Table 4: Description of datasets. Demographic and clinical characteristics.

| Characteristic | ABIDE | | Rest-meta-MDD | | SRPBS | |
|---|---|---|---|---|---|---|
| | ASD | HC | MDD | HC | Schizophrenia | HC |
| Sample Size | 528 | 571 | 828 | 776 | 92 | 92 |
| Age | 17.0 ± 8.4 | 17.1 ± 7.7 | 34.3 ± 11.5 | 34.4 ± 13.0 | 39.6 ± 10.4 | 38.0 ± 12.4 |
| Gender(M/F) | 464/64 | 471/100 | 301/527 | 318/458 | 47/45 | 60/32 |

## A.2 Implementation Details

In the proposed LGC-SGT model, we had the following hyperparameters and the range of them in Table 5. And all experiments were performed on a single NVIDIA RTX 3090 GPU.

Table 5: Ranges of different hyperparameters.

| Parameter | Rest-meta-MDD | ABIDE | SRPBS |
|---|---|---|---|
| number of spiking transformer blocks | 3 | 2 | 1 |
| timesteps | 8 | 2 | 4 |
| lambda | 0.1 | 0.2 | 0.2 |
| MEE kernel | 0.5 | 0.1 | 0.2 |
| batch size | 32 | | |
| learning rate | $1e-4$ | | |
| weight decay | $1e-4$ | | |
| attention heads | 2 | | |
| embedding dimension | 116 | | |

# B Two Perspective Biomarker Discovery

The differences in the association mechanisms between abnormal activity in brain regions and abnormal functional connectivity in pathological contexts primarily reflect their impact on different levels of neural networks: the former focuses more on the activity levels of local brain regions, while the latter emphasizes the functional connectivity relationships between brain regions.

## B.1 Brain Neuronal Population Dynamics

The analysis employs the Automated Anatomical Labeling (AAL) atlas as the criteria for brain region division, establishing a precise one-to-one mapping between 116 spiking neurons and neuroanatomical regions (90 cerebral and 26 cerebellar).

### B.2 Brain Functional Connectivity

To analyze brain functional connectivity under different brain disorders, we performed separate STDP training of the model's local pathways using samples from each disorder group and the HCs. To further investigate disorder-specific differences in functional connectivity patterns, we conducted the following analyses on the trained networks for each group: first, we established a mapping between the model's reservoir layer neurons and specific brain regions; second, we constructed functional connectivity maps based on the synaptic weight matrices obtained after training, quantifying the interaction strength between different brain regions; finally, by comparing network characteristics between patient and HCs, we identified abnormal connectivity patterns specific to each disorder, which may help elucidate the underlying neural circuit mechanisms of brain disorders. For visualization, we selected the top 100 edges with the highest weights along with their connected nodes.