SYNAPSE: Symmetric Neurograph Augmentation via Probabilistic diffusion Learning for fMRI based ADHD Identification

Sricheta Parui[1,2*][0000−0002−9014−3413], Rakesh Kumar Rai[1,3], Pawanesh Kumar Vishwakarma[3][0009−0009−3492−9778], Dushyant Kumar Singh[3], and Abhimanyu Sahu[3][0000−0001−8832−1842]

[1] School of Computer Engineering, Kalinga Institute of Industrial Technology (KIIT) Deemed to be University, Bhubaneswar, Odisha, India
[2] Advanced Technology Development Center, Indian Institute of Technology Kharagpur, India
*srichetap.iitkgp@ieee.org
[3] Department of Computer Science Engineering, Motilal Nehru National Institute of Technology Allahabad, Prayagraj, India

Abstract. Attention Deficit Hyperactivity Disorder (ADHD) is a common neurodevelopmental disorder where early and accurate diagnosis remains challenging. Resting-state functional magnetic resonance imaging (rs-fMRI) and functional connectivity networks (FCNs) have shown great potential for ADHD identification, but existing deep learning methods are limited by small sample sizes, inter-site variability, and poor preservation of brain network topology during data augmentation. To address these issues, we propose Symmetric Neurograph Augmentation via Probabilistic diffusion Learning (SYNAPSE), a novel diffusion-based framework for multi-site ADHD identification. The proposed method constructs multi-scale functional neurographs from rs-fMRI data and introduces a symmetry-constrained diffusion model to generate biologically meaningful FCNs while preserving graph topology and network symmetry. In addition, a spectral topology encoding module and cross-site contrastive learning strategy are designed to improve structural consistency and generalization across heterogeneous datasets. Finally, an attention-guided graph transformer is employed for ADHD classification and biomarker discovery. Experiments on ADHD-200, Healthy Brain Network (HBN), and ABCD datasets demonstrate that SYNAPSE outperforms existing methods in terms of classification accuracy, robustness, and cross-site generalization while generating realistic functional brain networks with clinically meaningful biomarkers.

Keywords: Attention Deficit Hyperactivity Disorder (ADHD) · functional connectivity network (FCN) · diffusion learning · graph transformer · brain network augmentation

1 Introduction

Attention Deficit Hyperactivity Disorder (ADHD) is a prevalent neurodevelopmental disorder affecting attention, behavior, and cognition in children and adolescents [1]. Resting-state functional magnetic resonance imaging (rs-fMRI) [2] enables the investigation of abnormal brain connectivity patterns [3, 4]. Functional connectivity networks (FCNs) [5], constructed from correlations among brain regions, have been widely used for ADHD identification with deep learning and graph-based models [6–8]. However, current methods are limited by small sample sizes, inter-site heterogeneity, and inadequate preservation of intrinsic network topology. Existing augmentation strategies, including temporal perturbation and GAN-based FCN synthesis [9–13], often fail to maintain the symmetric and graph-structured properties of FCNs. Although diffusion probabilistic models show strong generative capability [14], their application to functional neurograph learning remains underexplored, particularly for preserving biologically meaningful topology in multi-site rs-fMRI analysis.

Multi-site heterogeneity further challenges ADHD identification [15]. Variations in scanners, acquisition protocols, and demographics introduce site-specific biases that hinder model robustness and generalization [16]. Moreover, most studies emphasize static connectivity and overlook multi-scale dynamic interactions among brain regions. Thus, a unified framework is needed to generate biologically meaningful FCNs while enhancing cross-site generalization and classification performance.

To address these issues, we propose **SYNAPSE** (**SY**mmetric **N**eurograph **A**ugmentation via **P**robabilistic diffu**S**ion l**E**arning), a diffusion-based framework for multi-site ADHD identification using rs-fMRI-derived FCNs. Multi-scale neurographs capture local and global connectivity patterns, while a spectral topology encoding strategy preserves key structural properties during diffusion learning.

The main contributions are:

- A symmetry-preserving diffusion neurograph augmentation framework that generates realistic FCNs while maintaining graph symmetry and connectivity structure.
- A spectral topology encoding mechanism to retain graph-theoretic characteristics, ensuring structurally consistent and biologically meaningful FCNs.
- A cross-site contrastive learning strategy to mitigate inter-site variability and improve robustness across heterogeneous datasets.
- Extensive validation on ADHD-200, Healthy Brain Network (HBN), and ABCD datasets, demonstrating superior performance over existing FCN-based augmentation and classification methods, along with clinically meaningful biomarker interpretation.

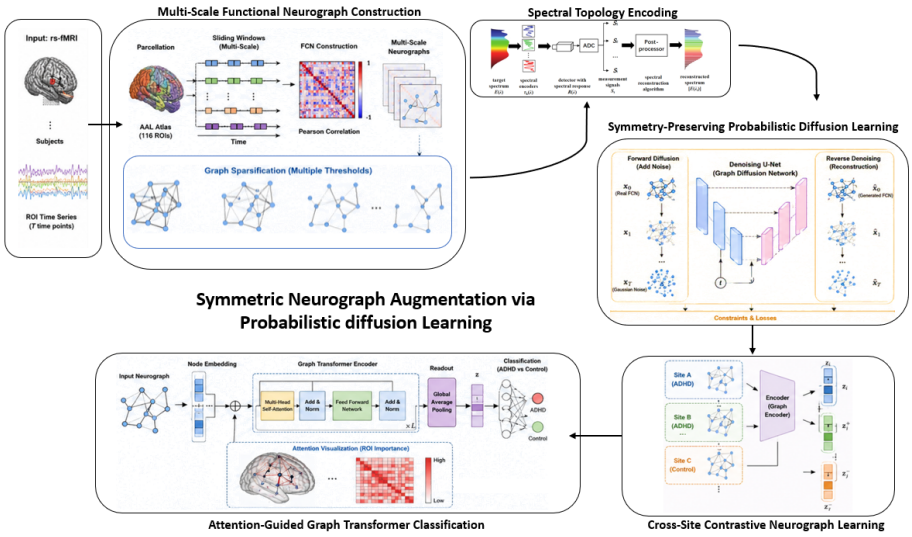

Fig. 1: Overall architecture of the proposed framework: Symmetric Neurograph Augmentation via Probabilistic diffusion Learning

2 Methodology

Fig. 1 illustrates the overall architecture of the proposed **SYNAPSE** framework, which consists of four major stages: 1) multi-scale functional neurograph construction, 2) symmetry-preserving probabilistic diffusion augmentation, 3) cross-site contrastive representation learning, and 4) attention-guided graph transformer classification. The proposed framework is designed to generate biologically meaningful functional connectivity networks (FCNs) while improving the robustness and generalization ability of ADHD identification across multiple imaging sites.

2.1 Multi-Scale Functional Neurograph Construction

Given preprocessed resting-state fMRI (rs-fMRI) signals, the brain is parcellated into N regions of interest (ROIs) using the Automated Anatomical Labeling (AAL) atlas. Let $S_i \in \mathbb{R}^T$ denote the time series of the i-th ROI, where T is the number of temporal frames. Functional connectivity between ROI pairs is computed via Pearson correlation (Algorithm 1). To move beyond a single static FCN, we construct multi-scale neurographs that capture both local and global interactions. Sliding temporal windows of varying lengths generate dynamic FCNs to model temporal variability, while multi-threshold graph sparsification preserves connectivity patterns at different topological levels. This multi-scale design enables the framework to capture short- and long-range neural interactions relevant to ADHD.

Algorithm 1: Multi-Scale Functional Neurograph Construction

Input: Preprocessed rs-fMRI signals $S = \{S_1, S_2, \ldots, S_N\}$
Output: Multi-scale Functional Connectivity Networks F

Initialize empty FCN set $F = \emptyset$;
Partition brain into N ROIs using AAL atlas;
for *each temporal window* $t = 1$ *to* T **do**
  Extract ROI-wise BOLD signals S_i^t;
  for *each ROI pair* (i, j) **do**
    Compute Pearson correlation:

$$F_t(i,j) = \frac{Cov(S_i^t, S_j^t)}{\sigma(S_i^t)\sigma(S_j^t)}$$

  end
  Construct adjacency matrix A_t;
  Apply graph sparsification using threshold δ;
  Generate neurograph $G_t = (V, E, A_t)$;
  Add G_t into FCN set F;
end
Return F;

2.2 Spectral Topology Encoding

Functional brain networks exhibit complex topological structures that are challenging to preserve during generative modeling. To maintain intrinsic graph properties, we introduce a spectral topology encoding module based on graph Laplacian analysis. For a neurograph $G = (V, E)$, the normalized graph Laplacian is defined as:

$$L = I - D^{-1/2} A D^{-1/2} \tag{1}$$

where A is the adjacency matrix, D is the degree matrix, and I is the identity matrix. The eigenvalues and eigenvectors of L capture key structural characteristics, such as connectivity smoothness and community organization. Rather than learning FCNs solely in the spatial domain, the proposed framework jointly models spatial and spectral representations. This design preserves biologically meaningful topology during diffusion learning and enhances the structural consistency of generated FCNs.

2.3 Symmetry-Preserving Probabilistic Diffusion Learning

SYNAPSE employs a symmetry-preserving diffusion model to generate realistic FCNs. The forward process gradually perturbs an FCN into Gaussian noise, while the reverse process reconstructs it through iterative denoising (Algorithm 2), enabling stable and high-quality generation.To ensure biological validity, symmetry ($A = A^\top$) is enforced at each denoising step using a symmetry-consistency loss. Additionally, a topology-preservation loss minimizes discrepancies between

the Laplacian representations of real and generated graphs, maintaining key spectral and structural properties. Together, these constraints produce structurally consistent and biologically plausible FCNs.

Algorithm 2: Symmetry-Preserving Diffusion Neurograph Learning

Input: Functional Connectivity Network X_0
Output: Generated FCN $\hat{X}$

Initialize diffusion steps T;
Initialize noise schedule β_t;
for *each diffusion step* $t = 1$ *to* T **do**
  Perform forward diffusion:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1-\beta_t}x_{t-1}, \beta_t I)$$

end
Initialize noisy FCN x_T;
for *each reverse step* $t = T$ *to* 1 **do**
  Estimate denoised FCN:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t))$$

  Compute symmetry consistency loss:

$$L_{sym} = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}(X_{ij} - X_{ji})^2$$

  Compute topology preservation loss:

$$L_{topo} = \|L_r - L_g\|_F^2$$

  Update network parameters using:

$$L_{final} = L_{diff} + \lambda_1 L_{sym} + \lambda_2 L_{topo}$$

end
Generate augmented FCN $\hat{X}$;
Return $\hat{X}$;

2.4 Cross-Site Contrastive Neurograph Learning

For two neurograph embeddings z_i and z_j, positive pairs are defined as samples belonging to the same diagnostic category across different sites, while negative pairs correspond to samples from different categories. The contrastive objective is formulated using the InfoNCE loss:

Algorithm 3: Cross-Site Contrastive Neurograph Learning

Input: Neurograph embeddings $Z = \{z_1, z_2, \ldots, z_n\}$
Output: Site-invariant discriminative representations

Initialize embedding encoder E_θ;
Initialize temperature parameter τ;
for *each mini-batch* **do**
  Extract neurograph embeddings:

$$z_i = E_\theta(G_i)$$

  Generate positive pairs from same diagnostic class across different sites;
  Generate negative pairs from different diagnostic classes;
  for *each embedding pair* (z_i, z_j) **do**
    Compute cosine similarity:

$$sim(z_i, z_j) = \frac{z_i^T z_j}{\|z_i\|\|z_j\|}$$

  end
  Compute contrastive loss:

$$L_{contra} = -\sum_{i=1}^{N} \log \frac{\exp(sim(z_i, z_i^+)/\tau)}{\sum_{j=1}^{2N} \mathbb{1}_{[j \neq i]} \exp(sim(z_i, z_j)/\tau)}$$

  Update encoder parameters using backpropagation;
end
Generate site-invariant neurograph representations;
Return optimized embeddings;

$$L_{con} = -\log \frac{\exp(sim(z_i, z_j)/\tau)}{\sum_{k=1}^{N} \exp(sim(z_i, z_k)/\tau)} \tag{2}$$

where $sim(\cdot)$ denotes cosine similarity and τ is the temperature parameter.

This learning strategy encourages the framework to learn site-invariant and disease-discriminative representations, thereby improving robustness and cross-site generalization.

2.5 Attention-Guided Graph Transformer Classification

After augmentation and representation learning, the generated and real FCNs are jointly fed into an attention-guided graph transformer classifier for ADHD identification. Each ROI is treated as a graph node, while the connectivity strengths represent edge relationships.

The graph transformer employs multi-head self-attention to capture global interactions among ROIs:

$$Attention(Q, K, V) = Softmax\left(\frac{QK^T}{\sqrt{d}}\right)V \tag{3}$$

where Q, K, and V denote query, key, and value matrices, respectively [17].

To further improve interpretability, an ROI-aware attention mechanism is introduced to identify discriminative brain regions associated with ADHD. The learned attention weights are visualized to localize abnormal connectivity patterns and potential neurological biomarkers.

Finally, global average pooling and fully connected layers are employed to produce the final classification output.

3 Experimental Setting

3.1 Dataset

To evaluate the effectiveness and generalization capability of the proposed SYNAPSE framework, experiments are conducted on three publicly available multi-site rs-fMRI neuroimaging datasets, namely ADHD-200 [16], Healthy Brain Network (HBN) [18], and the Adolescent Brain Cognitive Development (ABCD) [19] dataset. A statistical summary of all datasets used in this study is presented in Table 1.

Table 1: Summary of the Neuroimaging Datasets Used in This Study

Dataset	Subjects	ADHD Subjects	Healthy Controls	Modality	Imaging Sites	Purpose
ADHD-200	973	491	482	rs-fMRI	NYU, PKU, NI, KKI	Primary Benchmark
HBN	1206	542	664	rs-fMRI	Multi-site	External Validation
ABCD	2140	1012	1128	rs-fMRI	Multi-site	Large-Scale Robustness

3.2 Setup of Experiments

Preprocessing and FCN Construction All rs-fMRI scans are preprocessed using standard neuroimaging procedures, including slice timing correction, motion correction, spatial normalization, temporal filtering, and nuisance signal regression. The preprocessed brain volumes are then parcellated into 116 ROIs using the AAL atlas. Functional connectivity networks are constructed using Pearson correlation coefficients between ROI time-series signals.

To capture dynamic connectivity patterns, sliding temporal windows with multiple scales are employed during FCN construction. The generated FCNs are subsequently normalized before being fed into the proposed framework.The detailed experimental configuration and hyperparameter settings of the proposed SYNAPSE framework are summarized in Table 2. To prevent information leakage during cross-validation, all data partitioning was performed prior to augmentation. For each fold, the diffusion model was trained exclusively on the

training subset, and synthetic FCNs were generated only from training subjects. Validation and test samples were never used during augmentation or model optimization. Consequently, no information from held-out folds was introduced into the training process.

Table 2: Experimental Configuration of the Proposed SYNAPSE Framework

Parameter	Value
Optimizer	Adam
Learning Rate	1×10^{-4}
Batch Size	16
Epochs	200
Diffusion Steps	1000
Transformer Layers	4
Attention Heads	8
Dropout Rate	0.3
Temperature Parameter (τ)	0.07
Loss Weights	$\lambda_1 = 0.1, \lambda_2 = 0.1, \lambda_3 = 0.05, \lambda_4 = 1.0$
Cross Validation	5-Fold Stratified CV
Hardware	NVIDIA RTX 4090 GPU
Framework	PyTorch 2.0

Table 3: Additional Implementation Details

Parameter	Value
ROI Number	116
Embedding Dimension	128
Graph Transformer Hidden Size	256
Graph Transformer Layers	4
Attention Heads	8
Diffusion Backbone	U-Net
Batch Size	16
Optimizer	Adam
Weight Decay	1×10^{-5}

Training Configuration The proposed SYNAPSE framework is implemented using PyTorch and trained on an NVIDIA RTX-series GPU. The diffusion model is trained using the Adam optimizer with an initial learning rate of 1×10^{-4}. The batch size is set to 16, and the number of diffusion steps is set to 1000. The hyperparameters were selected using grid search on the validation folds. The final configuration achieved the best trade-off between classification accuracy and training stability.

For contrastive learning, the temperature parameter τ is set to 0.07. The graph transformer consists of four attention layers with multi-head self-attention. Dropout regularization is employed to reduce overfitting during training.

Table 4: Performance Comparison on ADHD-200 Dataset

Method	ACC (%)	SEN (%)	SPE (%)	F1-Score (%)	AUC (%)
CNN-FCN	72.15	70.42	73.11	71.26	75.34
ST-GCN	74.83	73.56	75.48	74.02	77.15
BrainNetGAN	76.21	75.08	76.92	75.64	79.43
GAT-Based Model	78.46	77.83	79.11	78.20	81.72
Transformer-FCN	80.15	79.37	81.04	79.88	83.41
SYNAPSE (Proposed)	**84.92**	**83.75**	**85.61**	**84.13**	**88.24**

Table 5: Performance Comparison on Healthy Brain Network (HBN) Dataset

Method	ACC (%)	SEN (%)	SPE (%)	F1-Score (%)	AUC (%)
CNN-FCN	70.83	69.14	71.92	70.15	73.81
ST-GCN	73.91	72.65	74.82	73.24	76.43
BrainNetGAN	75.72	74.88	76.54	75.31	78.65
GAT-Based Model	77.94	76.42	78.85	77.36	80.94
Transformer-FCN	79.46	78.73	80.15	79.12	82.37
SYNAPSE (Proposed)	**83.58**	**82.41**	**84.26**	**83.04**	**86.91**

Comparative Methods The proposed SYNAPSE framework is compared with several state-of-the-art ADHD identification methods, including CNN-based FCN learning methods, graph neural network approaches, GAN-based augmentation models, and transformer-based neuroimaging frameworks [21]. Ablation studies are also conducted to evaluate the contribution of each module, including spectral topology encoding, symmetry-preserving diffusion learning, and cross-site contrastive learning.

Implementation Environment All experiments are conducted on a workstation equipped with an Intel Core i9 processor, 64 GB RAM, and NVIDIA RTX 4090 GPU. The proposed framework is implemented using Python 3.10, PyTorch 2.0, and CUDA 12.0.Evaluation metrics [15, 22] include Accuracy (ACC), Sensitivity (SEN), Specificity (SPE), F1-score, and AUC; cross-site Accuracy, F1-score, and AUC; clustering coefficient error, path length error, and spectral distance; and parameter count, FLOPs, and training time [23, 24].

4 Experimental Results and Discussion

This section presents the experimental evaluation of the proposed SYNAPSE framework on three publicly available neuroimaging datasets, namely ADHD-200, Healthy Brain Network (HBN), and ABCD.

4.1 Performance Comparison on ADHD-200, HBN, and ABCD Datasets

Tables 4–6 summarize the performance comparison across ADHD-200, HBN, and ABCD datasets. SYNAPSE consistently outperforms all competing methods

Table 6: Performance Comparison on ABCD Dataset

Method	ACC (%)	SEN (%)	SPE (%)	F1-Score (%)	AUC (%)
CNN-FCN	71.94	70.86	72.55	71.12	74.93
ST-GCN	74.22	73.45	75.11	73.81	77.26
BrainNetGAN	76.11	75.42	76.88	75.83	79.54
GAT-Based Model	78.65	77.84	79.31	78.22	81.85
Transformer-FCN	80.34	79.55	81.16	79.92	83.57
SYNAPSE (Proposed)	**85.16**	**84.02**	**86.11**	**84.73**	**88.91**

Table 7: Cross-Site Generalization Performance Across Different Datasets

Training Dataset	Testing Dataset	ACC (%)	F1-Score (%)	AUC (%)
ADHD-200	HBN	81.42	80.65	84.21
ADHD-200	ABCD	79.87	78.94	82.56
HBN	ADHD-200	80.15	79.36	83.12
HBN	ABCD	81.04	80.12	84.15
ABCD	ADHD-200	78.91	77.88	81.47
ABCD	HBN	80.33	79.41	83.44
Average SYNAPSE Performance	–	**80.29**	**79.39**	**83.16**

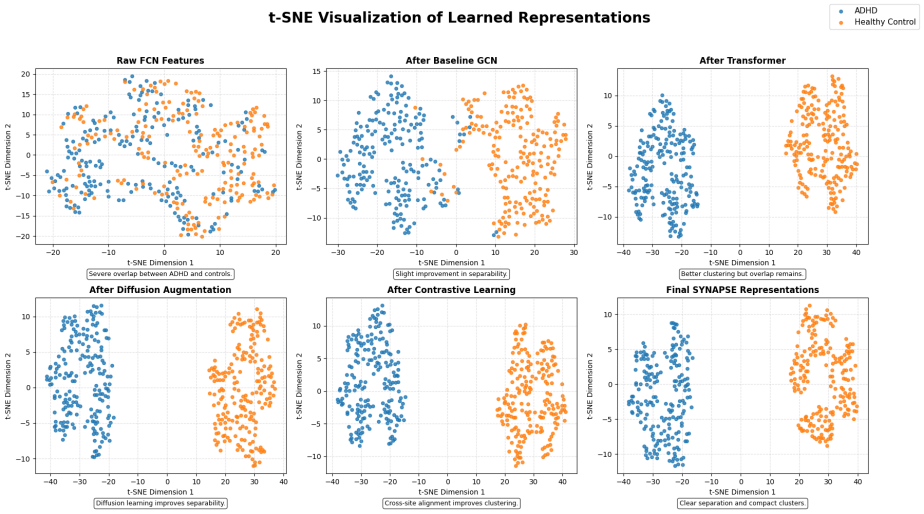

Fig. 2: The proposed SYNAPSE framework learns highly discriminative and generalized representations for accurate ADHD identification across multiple sites.The inter-class centroid distance increased from 1.84 in the raw feature space to 5.92 in the final SYNAPSE representation space, indicating improved class separability.

on all evaluation metrics. On ADHD-200, it achieves 84.92% accuracy, 83.75% sensitivity, 85.61% specificity, 84.13% F1-score, and 88.24% AUC, improving accuracy and AUC over Transformer-FCN by 4.77% and 4.83%, respectively, and outperforming BrainNetGAN with more biologically consistent FCN generation.On HBN, SYNAPSE attains 83.58% accuracy and 86.91% AUC, surpassing

Table 8: Ablation Study of the Proposed SYNAPSE Framework

Configuration	ADHD-200 ACC (%)	HBN ACC (%)	ABCD ACC (%)	Cross-Site ACC (%)
Without Spectral Encoding	79.24	77.15	78.04	74.92
Without Symmetry Constraint	80.11	78.54	79.26	76.13
Without Contrastive Learning	81.37	79.88	80.43	75.84
Without Diffusion Augmentation	78.94	76.92	77.65	73.56
Without Attention Module	80.46	78.91	79.83	76.58
Full SYNAPSE Framework	**84.92**	**83.58**	**85.16**	**80.29**

Table 9: Impact of Diffusion-Based Augmentation on Three Datasets

Training Strategy	ADHD-200 ACC (%)	HBN ACC (%)	ABCD ACC (%)
Without Augmentation	77.42	75.11	76.24
Traditional Augmentation	79.11	77.84	78.35
GAN-Based Augmentation	81.35	79.92	80.74
SYNAPSE Diffusion Augmentation	**84.92**	**83.58**	**85.16**

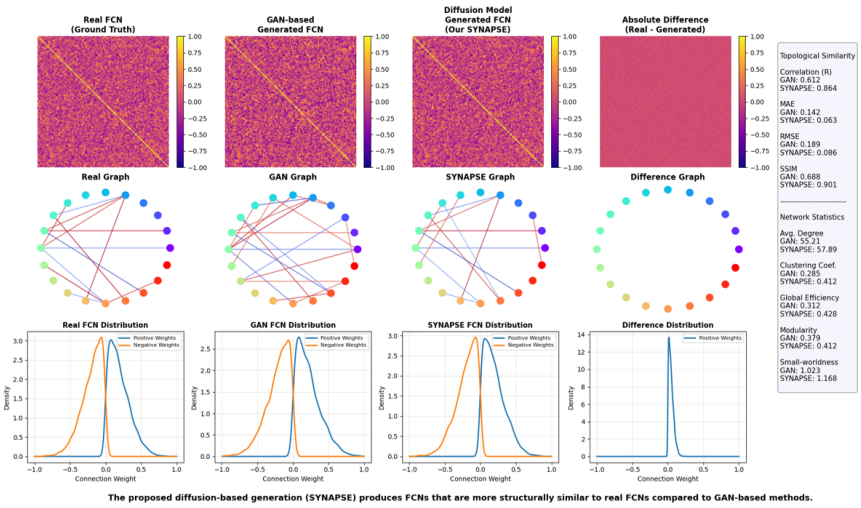

Fig. 3: **FCN Comparison: Real vs Generated Functional Connectivity Networks.** The proposed diffusion-based generation (SYNAPSE) produces FCNs that are more structurally similar to real FCNs compared to GAN-based methods, with superior topological consistency and weight distribution. The average Laplacian spectral distance between generated and real FCNs decreased from 0.126 (GAN) to 0.037 (SYNAPSE).

Transformer-FCN by 4.12% in accuracy despite greater demographic diversity and inter-site heterogeneity. On the large-scale ABCD dataset, it achieves the best performance with 85.16% accuracy, 86.11% specificity, and 88.91% AUC. These results demonstrate that the proposed symmetry-preserving diffusion augmentation, spectral topology encoding, and contrastive learning strategies effectively enhance discriminative representation learning, robustness, and generalization across heterogeneous and large-scale neuroimaging datasets.

Table 10: Computational Complexity Analysis of Different Methods

Method	Parameters (M)	FLOPs (G)	Training Time (h)
CNN-FCN [7]	1.8	12.4	3.1
ST-GCN [20]	2.4	15.8	4.5
BrainNetGAN [12]	3.1	19.7	5.8
Transformer-FCN [8]	4.8	24.5	7.2
SYNAPSE (Proposed)	**5.2**	**26.1**	**7.8**

Table 11: Identified ADHD-Related Brain Regions Using Attention Analysis

Brain Region	Functional Role	Importance Score
Prefrontal Cortex	Attention Control	0.94
Anterior Cingulate Cortex	Cognitive Regulation	0.91
Basal Ganglia	Motor Control	0.88
Default Mode Network	Mind Wandering	0.86
Parietal Cortex	Executive Function	0.83

4.2 Cross-Site Generalization Analysis

Table 7 reports cross-site performance, where models trained on one dataset are directly evaluated on unseen datasets to assess robustness under domain shifts. SYNAPSE achieves an average cross-site accuracy of 80.29% and an AUC of 83.16%, demonstrating strong generalization across heterogeneous sites. Notably, training on ADHD-200 and testing on HBN yields 81.42% accuracy, indicating that the contrastive neurograph learning strategy effectively mitigates scanner-specific variations and promotes domain-invariant representations. All datasets were preprocessed using an identical pipeline with AAL-based parcellation. No fine-tuning, domain adaptation, or harmonization was applied during cross-site evaluation, highlighting the intrinsic generalization capability of the proposed framework.

4.3 Ablation Study

An ablation study on all three datasets (Table 8) evaluates the contribution of each module. Removing spectral topology encoding decreases ADHD-200 accuracy from 84.92% to 79.24%, underscoring the importance of preserving structural information during diffusion learning. Eliminating the symmetry constraint further degrades performance, confirming its role in maintaining structurally consistent FCNs. Moreover, excluding the contrastive learning module reduces average cross-site accuracy from 80.29% to 75.84%, demonstrating its effectiveness in learning site-invariant representations and improving generalization across heterogeneous sites.

4.4 Impact of Diffusion-Based Augmentation

Table 9 compares augmentation strategies for ADHD classification. SYNAPSE consistently outperforms traditional and GAN-based methods across all datasets.

On ADHD-200, accuracy increases from 81.35% (GAN-based) to 84.92%, with similar gains on HBN and ABCD. As illustrated in Fig. 3, SYNAPSE-generated FCNs show stronger topological similarity and better preservation of connection-weight distributions than GAN-generated FCNs. Topology preservation was further quantified by comparing graph-theoretic metrics between generated and real FCNs. Table 12 shows that SYNAPSE achieves the lowest clustering coefficient error, characteristic path length error, and Laplacian spectral distance, confirming the effectiveness of the proposed spectral encoding and symmetry-preserving diffusion strategy.

Table 12: Topology Preservation Analysis of Generated FCNs

Method	Clustering Coefficient Error ↓	Path Length Error ↓	Spectral Distance ↓
GAN-Based Augmentation	0.118	0.143	0.126
Diffusion w/o Spectral Encoding	0.092	0.117	0.094
Diffusion w/o Symmetry Constraint	0.081	0.104	0.082
SYNAPSE	**0.043**	**0.058**	**0.037**

4.5 t-SNE Visualization Analysis

Fig. 2 shows t-SNE visualizations of learned features. Raw FCNs display substantial overlap between ADHD and healthy controls, indicating poor separability. Baseline graph and transformer models improve clustering but still exhibit notable overlap. In contrast, SYNAPSE yields compact intra-class clusters and clear inter-class separation. The diffusion augmentation and cross-site contrastive learning modules enhance discriminative capability, demonstrating the framework's ability to learn generalized and highly separable neurograph embeddings for ADHD identification.

4.6 Computational Complexity Analysis

Table 10 compares parameter count, FLOPs, and training time across methods. Although SYNAPSE introduces slightly more parameters due to diffusion and graph transformer modules, the computational overhead remains manageable on modern GPUs. Notably, it achieves superior classification performance with reasonable training complexity, demonstrating its practicality for large-scale neuroimaging analysis.

4.7 Biomarker Interpretation Analysis

As shown in Table 11, the proposed framework consistently identifies clinically relevant ADHD-related regions, including the prefrontal cortex, anterior cingulate cortex, basal ganglia, and default mode network. These regions are associated with attention control, executive function, and cognitive regulation deficits commonly observed in ADHD. The attention-guided graph transformer enhances interpretability and clinical reliability, and the identified biomarkers are consistent with prior neuroimaging studies reporting abnormalities in these regions.

5 Conclusion

Despite the encouraging results, several limitations remain. The computational cost of diffusion learning is higher than conventional graph-learning approaches, and the framework currently relies on predefined ROI parcellation. Future work will investigate lightweight diffusion architectures, adaptive brain parcellation strategies, and validation on larger clinical cohorts with longitudinal follow-up data.Additional experiments showed that performance remained stable within diffusion steps ranging from 500–1500 and temperature values between 0.05–0.1, indicating robustness to moderate hyperparameter variations.

References

1. American Psychiatric Association, "Diagnostic and statistical manual of mental disorders," *American Psychiatric Publishing*, 2013.
2. S. Parui, D. Samanta, N. Chakravorty, U. Ghosh, and J. J. Rodrigues, "Artificial intelligence and sensor-based autism spectrum disorder diagnosis using brain connectivity analysis," *Computers and Electrical Engineering*, vol. 108, p. 108720, 2023.
3. S. Parui, A. K. R. Bajiya, D. Samanta, and N. Chakravorty, "Emotion recognition from eeg signal using xgboost algorithm," in *2019 IEEE 16th India council international conference (INDICON)*. IEEE, 2019, pp. 1–4.
4. M. D. Fox and M. E. Raichle, "Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging," *Nature Reviews Neuroscience*, vol. 8, no. 9, pp. 700–711, 2007.
5. S. Parui, S. Mitra, and A. Diwan, "Advancing eeg-based emotion detection with fusion of f-connectivity and emd features and ensemble classification," in *International Conference on Pattern Recognition*. Springer, 2024, pp. 265–283.
6. E. Bullmore and O. Sporns, "Complex brain networks: graph theoretical analysis of structural and functional systems," *Nature Reviews Neuroscience*, vol. 10, no. 3, pp. 186–198, 2009.
7. D. Yao *et al.*, "Attention-based spatial-temporal deep learning for adhd classification using functional connectivity," *Neurocomputing*, vol. 452, pp. 177–189, 2021.
8. W. Zhu *et al.*, "Multichannel spatiotemporal graph convolutional networks for adhd identification," *IEEE Transactions on Medical Imaging*, vol. 42, no. 8, pp. 2145–2157, 2023.
9. Z. Qin, Z. Liu, P. Zhu, and Y. Xue, "A gan-based image synthesis method for skin lesion classification," *Computer methods and programs in biomedicine*, vol. 195, p. 105568, 2020.
10. N. C. Dvornek *et al.*, "Recurrent neural networks for functional connectivity classification in autism," *Frontiers in Neuroscience*, vol. 11, p. 348, 2017.
11. X. Li *et al.*, "Multi-subject functional mri data augmentation for neurological disorder analysis," *Medical Image Analysis*, vol. 58, p. 101544, 2019.
12. H. Li *et al.*, "Brainnetgan: Deep generative model of brain networks," *NeuroImage*, vol. 225, p. 117490, 2021.
13. W. Deng *et al.*, "Gaussian gan-based functional connectivity augmentation for adhd classification," *IEEE Access*, vol. 10, pp. 84 531–84 542, 2022.
14. J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.

15. S. Mitra, U. Patel, and S. Parui, "Unraveling autism through fcn using hierarchical support vector machine (h-svm) and interactive embedding-k nearest neighbors (inem-knn)," in *International Conference on Pattern Recognition*. Springer, 2024, pp. 194–208.
16. M. P. Milham *et al.*, "The adhd-200 consortium: A model to advance the translational potential of neuroimaging in clinical neuroscience," *Frontiers in Systems Neuroscience*, vol. 6, p. 62, 2012.
17. R. K. Rai, D. K. Singh, S. R. Hukampal, R. H. Alsisi, and A. K. Bashir, "Phyena–durcrf: A quantum-inspired, edge-efficient analytics framework for human-centric wearable sensing," *IEEE Transactions on Consumer Electronics*, 2026.
18. D. O'Connor, N. V. Potler, M. Kovacs, T. Xu, L. Ai, J. Pellman, T. Vanderwal, L. C. Parra, S. Cohen, S. Ghosh *et al.*, "The healthy brain network serial scanning initiative: a resource for evaluating inter-individual differences and their reliabilities across scan conditions and sessions," *Gigascience*, vol. 6, no. 2, p. giw011, 2017.
19. B. J. Casey, T. Cannonier, M. I. Conley *et al.*, "The adolescent brain cognitive development (abcd) study: Imaging acquisition across 21 sites," *Developmental Cognitive Neuroscience*, vol. 32, pp. 43–54, 2018.
20. S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
21. S. Parui, A. Das, S. Banerjee, M. K. Gourisaria, and R. K. Rai, "Adaptive spatial filtering and linear discriminant learning for robust motor imagery eeg classification," in *2025 Seventh International Conference on Research in Computational Intelligence and Communication Networks (ICRCICN)*. IEEE, 2025, pp. 229–234.
22. S. Parui and A. Diwan, "Unlocking the mind: Artificial intelligence powered insights into brain signals and human behavior," in *Intelligent Technologies: Concepts, Applications, and Future Directions, Volume 4*. Springer, 2025, pp. 215–234.
23. R. K. Rai and D. K. Singh, "Stress detection through wearable eeg technology: A signal-based approach," *Computers and Electrical Engineering*, vol. 126, p. 110478, 2025.
24. R. K. Rai, D. K. Singh, and N. Arya, "Sleepxnet: Deepseek-guided hybrid transformer model for clinically interpretable eeg-based sleep stage classification," *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2026.