Biologically Informed Graph Attention Network for Interpretable Brain Connectivity Analysis

1st Roger Zhu
Independent Researcher
Millburn High School
Millburn, NJ, USA
27zhur@millburn.org

2nd Rishi Upadhyay
Department of Computer Science
University of California, Los Angeles
Los Angeles, CA, USA
rishiu@ucla.edu

***Abstract*—Brain connectome analysis has long been a focus in neuroscience and neuroimaging research. However, many recent approaches focus heavily on fMRI and ROI-based modeling approaches, often overlooking key structural features and connectivity patterns of white matter tracts. In this work, we present a novel heterogeneous graph framework that incorporates structural features of white matter tracts, representing both white matter clusters and anatomical regions as nodes with edges capturing their spatial intersection. This approach allows us to model both structural integrity and connectivity, along with regional context for more biologically meaningful patterns. Using GraphSAGE for feature aggregation and multi-head Graph Attention Networks (GAT) for interpretability, our proposed architecture achieves an AUC of 0.891 for predicting sex on the Human Connectome Project 1200 dataset, matching the performance of other existing models and validating our graph construction while using significantly fewer training data. Furthermore, we extract attention scores, identifying key white matter tracts and anatomical regions involved in sex differentiation, grounding model predictions in neuroanatomical relevance. Our results highlight the importance of structural information and pave the way for more biologically informed models.**

***Index Terms*—heterogeneous graph neural networks, diffusion MRI, white matter tracts, sex classification.**

I. Introduction

The human brain is a complex network of interconnected regions. Understanding the brain's wiring, the connectome, has long been a central goal in neuroscience [1], [2]. Advances in magnetic resonance imaging (MRI) techniques, including diffusion MRI (dMRI) and functional MRI (fMRI), have enabled the rise of network neuroscience, where the brain's structural and functional features can be modeled as complex brain graphs [3], [4]. These connectomes are powerful tools for exploring individual differences in cognition, behavior, and even identifying neurological disorders at an early stage [5].

However, most prior work in brain connectome analysis focuses heavily on functional connectivity derived from fMRI, representing the brain as a graph of anatomical regions (ROIs) connected by a statistical relationship in activity [6], [7]. While effective at times, this approach overlooks the rich microstructural detail information found in white matter fiber tracts. Even structural connectome studies simplify tractography by reducing it to region-to-region connection weights, discarding valuable information regarding the integrity, density, and shape of the underlying fibers [5], [8]. Understanding the structural properties of white matter fiber tracts is crucial, as they serve as the physical wiring of the brain, enabling communication and forming the foundation for higher-level function. To address this, we propose a novel graph representation framework allowing for accurate and interpretable brain connectome analysis to predict sex. Specifically, our contributions are as follows:

1) We introduce a new graph formulation that separates the brain into distinct node types, allowing us to incorporate the microstructural details.
2) A novel GNN framework that combines GraphSAGE and multi-head GAT to combine performance and interpretability.
3) Our model achieves a competitive AUC (0.891) on sex classification while using up to three times fewer training data.

II. Related Work

A. Machine Learning in Brain Connectomics

Classical machine learning approaches treat the connectome as high-dimensional feature data for prediction. Typically, one would compute a connectivity matrix for each subject (often using Pearson correlations from resting-state fMRI or tractography-based edges from dMRI) and then either apply feature extractions or use the raw connectivity values in classifiers such as SVM, Logistic Regression, and Random Forest [9]. These connectomes have been used to diagnose disorders and predict patient outcomes. For example, Zeraati & Davoodi (2025) present ASD-GraphNet, where functional connectomes are converted into node-level, edge-level, and graph-level features such as centralities and node degree. These features are then fed into classic classifiers to distinguish autism from control, getting an accuracy of around 75% [10]. Other approaches of feature extraction include using tensor decomposition on DTI-derived structural networks, or applying SVMs or kernel methods for disease diagnosis [11] [12]. Major limitations of extracting handcrafted features include their reduced expressivity, as they often flatten the data.

B. Graph-based Deep Learning

More recently, graph neural networks (GNNs) have been used to model the connectome directly as a graph, where most often, nodes would represent regions of interest. In these approaches, each subject's brain graph is input to a GNN that learns representations for classification or regression. Architecturally, recent GNNs have employed graph convolutions for brain graphs, in which features are aggregated from all neighbouring nodes equally. Kazi et al. (2023) proposed a multi-head Graph Convolutional Network (GCN) that applies parallel graph convolutions focusing on nodes and edges, achieving greater results compared to classical ML approaches [13]. Kazi et al. (2025) further developed a GCN-inspired model with a connectivity attention module that combines deep learning approaches with extracted feature matrices to predict age [14].

Another common direction in connectomic analysis is the fusion of multiple imaging modalities, particularly combining structural (dMRI) and functional (fMRI) connectivity. These approaches aim to capture a more complete picture of the brain by leveraging complementary information. Qu et al. (2025) proposed a MaskGNN that combines derived functional and structural connectivity matrices, along with a cortical atlas for intelligence score predictions [15]. Other approaches incorporate genetic or clinical data alongside connectomic data to capture unique patterns [16].

III. Materials and Methods

A. Dataset

Diffusion-weighted MRI data used within this study were acquired from the Human Connectome Project (HCP) 1200, a large multi-modal dataset comprised of data from young healthy adults (age 22-37) [17]. Out of the total 1206 subjects available, 233 participants who did not complete the full 3T diffusion MRI protocol were excluded, resulting in a final sample of 973 subjects. The dataset includes a balanced distribution of sexes, with 451 males and 522 females used for downstream modeling. For each subject, we extracted the following files:

1) Diffusion-weighted images.
2) Diffusion brain masks.
3) Anatomical parcellation label maps: brain masks segmented into anatomical regions of interest obtained from the Desikan-Killiany Atlas [18].

All diffusion MRI (dMRI) data were acquired on a Siemens Skyra Connectome Scanner with the following parameters: b-values of 1000, 2000, and 3000 s/mm^2; 90 diffusion directions per shell plus 6 b=0 acquisitions interspersed throughout each run; 1.25 mm isotropic resolution; and a multiband factor of 3. These high-resolution scans allow for detailed modeling of white matter pathways and structural features.

B. Data Preprocessing

All dMRI data were processed with the HCP MR Diffusion Pipeline, which normalizes and corrects for distortions [17]. To reduce computational load while preserving angular coverage, the diffusion MRI data were downsampled from 288 to 90 diffusion-weighted directions. We implemented a farthest point sampling strategy in cosine space to select directions while retaining b=0 volumes in order to provide anatomical context. After normalizing all nonzero b-vectors to unit length, the algorithm iteratively selected vectors that were most dissimilar from the current set to ensure directional diversity as demonstrated in Fig. 1. The final subset combined these 90 directions with the original b=0 images, and the corresponding diffusion volumes, b-values, and b-vectors were extracted to produce a reduced DWI dataset, preserving angular resolution critical for tractography.

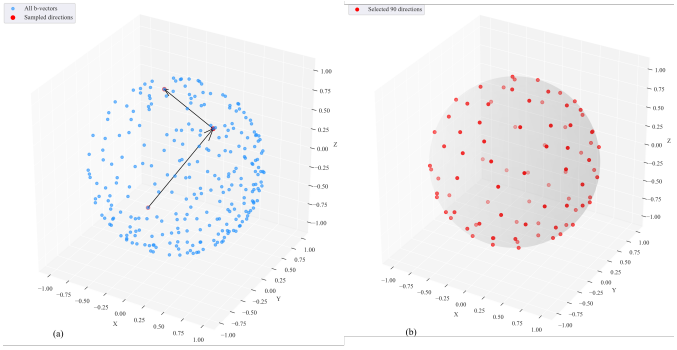

Fig. 1. Visualization of diffusion sampling strategy on the unit sphere. (a) Depicts the Farthest Point Sampling Algorithm: the initial randomly selected direction is highlighted in red, and subsequent arrows indicate directions chosen to maximize angular diversity. (b) Final subset of 90 directions evenly distributed.

Whole-brain tractography was then performed using the DIPY library on the downsampled MRI volumes [19]. First, a diffusion tensor model (DTI), which quantifies the direction and magnitude of water diffusion at each voxel, was fit to the data [20]. From this model, we computed voxel-wise diffusion properties such as fractional anisotropy (FA), a scalar measure commonly used to measure white matter structural integrity. Brain white matter masks were then constructed where FA> 0.2 to ensure our tractography algorithm remains within the bounds of white matter tracts. While a DTI provides a basic model of diffusion, it struggles with resolving crossing fibers. Therefore, to better capture the brain's complex fiber structure, we used a constrained spherical deconvolution (CSD) model that estimates the fiber orientation distribution function (fODF), which is a continuous function describing the probability of fiber directions within each voxel [21]. CSD achieves this by deconvolving the observed diffusion signal with a response function, which characterizes the expected signal from a single coherent fiber population. We used the single-shell single-tissue (SSST) response function, derived from FA> 0.7 voxels. The tractography is then conducted with a step size of 0.5mm and a maximum angle of 30 degrees to prevent sharp, unrealistic turns. The number of streamlines generated is capped at 100,000 for computational reasons.

To reduce redundancy and simplify the tractography data,

we clustered whole-brain streamlines into anatomically similar fiber bundles using the QuickBundles algorithm [22]. QuickBundles performs clustering based on the minimum average direct-flip (MDF) distance between streamlines, accounting for streamline geometry regardless of orientation. The algorithm was chosen for its algorithmic simplicity, allowing for fast and computationally efficient clustering. Prior to clustering, streamlines were filtered to remove those shorter than 10 mm or longer than 250 mm to reduce spurious or non-biological tracts. The clustering threshold was set to 10 mm. Clustered tracts were further filtered to exclude those with fewer than 20 streamlines in order to reduce noise for downstream tasks. Each resulting cluster represents a distinct white matter fiber tract.

Anatomical parcellations were obtained from the Desikan-Killiany Atlas for cortical region of interests (ROIs), and subcortical segmentations from FreeSurfer for a total of 114 ROIs [18]. To integrate these ROIs with the diffusion data, the parcellation mask was aligned to the subject's diffusion space using center-of-mass alignment. Nearest-neighbor interpolation was used during resampling to preserve integer label values. This coarse alignment strategy is sufficient for the purposes of graph construction, as the goal is to determine which streamlines spatially intersect each ROI rather than to achieve precise boundary-level registration; voxel-level label identity is preserved via nearest-neighbor interpolation.

C. Graph Construction

The proposed graph construction scheme was designed to model brain connectivity while incorporating important structural information. For each subject, we constructed a heterogeneous graph $\mathcal{G} = (V, E)$ for each subject. The node set V is composed of two types: streamline cluster nodes V_c and region-of-interest (ROI) nodes V_r. Streamline cluster nodes encode two features: mean FA across all streamlines, reflecting the structural integrity of the white matter tract; and the proportion of streamlines, defined as the size of the cluster relative to the total number of streamlines per subject. ROI nodes are encoded with a one-hot vector indicating the anatomical region label and the mean FA value within the ROI mask. Edges $E \subseteq V_c \times V_r$ were created between cluster and ROI nodes if there was an intersection between the streamlines in a cluster and the voxels labeled as part of an ROI Fig 2. Intersection is defined as binary with no weighting by the number of intersecting points. The resulting graph is an undirected, bipartite graph. To quantify key characteristics of the constructed graph dataset, summarized statistics are provided in Table I.

D. Model Architecture

The proposed heterogeneous graph representation is utilized to predict the biological sex of an individual subject, which we framed as a graph classification task. A schematic overview of the model architecture is provided in Fig 3. Heterogeneous graphs are passed through node feature encoders, various GNN layers, and a pooling layer where embeddings are aggregated before being handed off to the classification head.

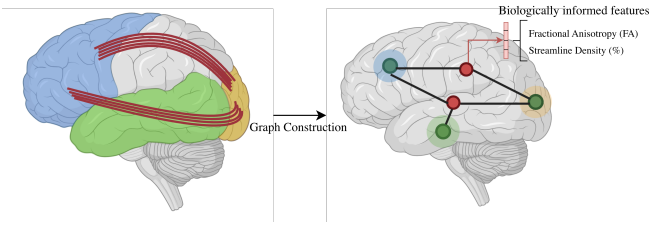

Fig. 2. Schematic illustration of graph construction from dMRI tractography. Clustered streamlines intersect ROIs in anatomical space, resulting in edges being drawn. This effectively captures important structural information and brain connectivity.

TABLE I
KEY STATISTICS OF THE CONSTRUCTED GRAPH DATASET. REPORTED VALUES ARE MEANS ± STANDARD DEVIATIONS.

Statistic	Value (Mean ± SD)
Number of Cluster Nodes	1012 ± 43.7
Number of ROI Nodes	112 ± 0.58
Number of Edges	8064 ± 481
Edge Density	0.07
Average Node Degree (Cluster)	7.96 ± 4.99
Average Node Degree (ROI)	71.77 ± 72.02

To initialize the node representation, separate encoders are used for each node type due to the difference in feature dimensionality. Each node type is passed through an independent linear projection followed by a ReLU activation and a batch normalization layer. These encoders map both cluster and ROI features into a shared hidden space of dimension 64, enabling subsequent message-passing layers to operate in a unified feature space.

To effectively capture local graph structure and inter-node interactions, we applied one layer of GraphSAGE followed by a Graph Attention Network (GAT) layer to the heterogeneous graph [23]. The GraphSAGE layer aggregates neighbourhood features through a mean-based message-passing scheme. For each node v, the updated embedding is computed as:

$$\mathbf{h}_v^{(k)} = \sigma\Bigg(\mathbf{W} \cdot \text{AGGREGATE}\Big(\Big\{\mathbf{h}_v^{(k-1)}\Big\} \cup \Big\{\mathbf{h}_u^{(k-1)} : u \in \mathcal{N}(v)\Big\}\Big)\Bigg), \tag{1}$$

Where $h_v^{(k)}$ is the hidden representation at layer $k, N(v)$ denotes the neighbours of node v, and σ is a non-linear activation function (ReLU). This allows for the model to learn informative node feature representations while incorporating local graph structure.

Following the GraphSAGE layer, the GAT layer is applied to assign learnable importance weights to edges [24]. These weights can later be extracted for downstream analysis and model interpretability. For each edge (v, u), attention coeffi-

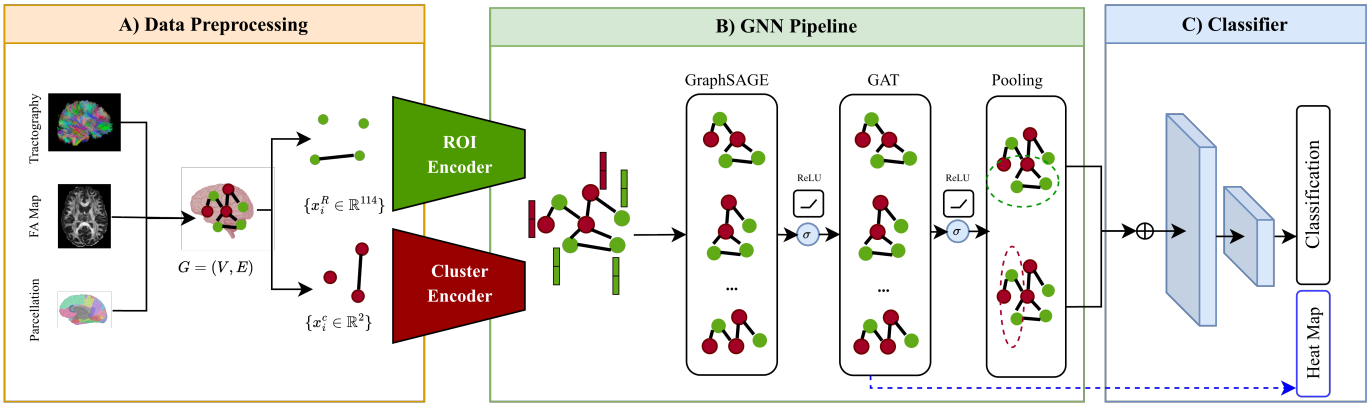

Fig. 3. Overview of the proposed heterogeneous graph neural network architecture. Node features are encoded separately for cluster and ROI nodes, followed by GraphSAGE and GAT layers for message passing. Global mean pooling is applied to obtain a graph-level embedding, which is passed to a classification head for sex prediction.

cients α_{vu} are computed using a shared attention mechanism across all node types:

$$\alpha_{vu} = \frac{\exp\left(\text{LeakyReLU}\left(\mathbf{a}^\top \left[\mathbf{W}\mathbf{h}_v \,\|\, \mathbf{W}\mathbf{h}_u\right]\right)\right)}{\sum_{k \in \mathcal{N}(v)} \exp\left(\text{LeakyReLU}\left(\mathbf{a}^\top \left[\mathbf{W}\mathbf{h}_v \,\|\, \mathbf{W}\mathbf{h}_k\right]\right)\right)}, \quad (2)$$

Where $\|$ denotes concatenation and α is a learnable weight vector. To further enhance our model, multi-head attention with 2 heads, where outputs were concatenated, was implemented.

Finally, a global mean pooling layer is applied to obtain a graph-level embedding for classification. The pooling is applied separately to both ROI and cluster node embeddings and is concatenated before being passed on to the final classification head. The classification head is made up of one fully connected linear layer that maps the pooled features to a binary prediction indicating the subject's sex. A sigmoid activation function is applied to map the logits into probabilities.

E. Training Details

All models were trained on a Google Colab instance with an NVIDIA T4 GPU. The Adam optimizer was used in training with a learning rate of $1\text{x}10^{-4}$ and weight decay of $1\text{x}10^{-5}$. Models were trained for 200 epochs with a batch size of 8. Early stopping was implemented with a patience of 15 to prevent overfitting. The dataset was split into 70% training, 15% validation, and 15% testing with model performance measured on the held-out test set.

IV. Results

A. Quantitative Results

Quantitative results on the test dataset are reported in Table II and compared contexually to existing models of varying datasets and modalities. Note that this is not a direct benchmark comparison but rather an indication that our approach achieves competitive performance despite limited data. Our method achieved an average AUC of 0.874 ± 0.01, with a highest AUC of 0.891.

TABLE II
Comparison of Sex Classification Performance Across Models and Imaging Modalities.

Model	Sample Size	Modality	AUC (↑)
Logistic Regression [25]	$> 5500^a$	fMRI	0.76 ± 0.15
BraiNN [26]	3308	sMRI	0.927
Proposed Model	973	dMRI	0.874 ± 0.01

[a] Dataset pooled from ABIDE, HCP, ACPI, COBRE-MIND, and T1000.

B. Ablation Study

To evaluate the contribution of each architectural component, an ablation study was conducted on identical data splits differing only in random initialization seeds, and summarized in Table III. Starting from a baseline model using only one GAT layer and classifier head, the model achieved an AUC of 0.599 ± 0.01. Introducing the node encoder led to the greatest improvement, with an increase in AUC by 0.248, highlighting the importance of initial feature transformation. Subsequent additions include GraphSAGE and multi-head attention. Ultimately, the final model yielded the best result with an AUC of 0.874 ± 0.01, which was statistically significant compared to all ablated variants, determined by a paired t-test at the 95% confidence level.

C. Qualitative Results and Interpretability

To gain insight into the model's decision-making process, we qualitatively analyze attention scores learned by the model. From the GAT layer, we averaged attention scores across

TABLE III
ABLATION STUDY ON THE HUMAN CONNECTOME PROJECT 1200 DATASET. MEAN AND STANDARD DEVIATION OVER TEN RUNS ARE REPORTED.

Model Variant	Acc. (↑)	AUC (↑)	F1 (↑)
Baseline	0.517 (< 0.01)	0.599 (< 0.01)	0.643 (< 0.01)
+ Node Encoder	0.754 (0.01)	0.847 (0.003)	0.781 (0.02)
+ GraphSAGE	0.781 (0.34)	0.858 (0.01)	0.800 (0.44)
+ Multi-head Attention	0.790 ± 0.01	0.874 ± 0.01	0.808 ± 0.02

Parentheses indicate p-values from paired t-tests compared to the proposed model.

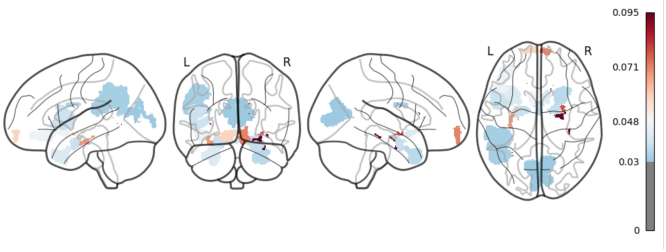

Fig. 4. Heatmap depicting the top anatomical regions of the brain with the highest attention scores derived from the GAT layer. Attention scores are averaged across all subjects.

the two attention heads and then aggregated across all edges incident to each node. The per-node scores are averaged across all subjects to produce heatmaps of the top anatomical regions in Figure 4 and the top edges in Fig 5. We identified key regions including the lateral ventricles and the frontal pole cortex that were consistent across all ten runs, indicating that the findings are stable with respect to random initialization rather than artifacts of a single training seed.

V. DISCUSSION

A. Summary of Findings

Our model achieved an AUC of 0.891, matching the results of state-of-the-art models using our novel heterogeneous graph that jointly models white fiber clusters and anatomical regions, allowing for the model to incorporate both structural integrity and regional context. Despite using significantly fewer training samples, up to 3 times fewer, we achieved competitive results. Furthermore, the use of GraphSAGE for node representation and multi-head GAT for message passing enabled not only strong performance but also offered interpretability through learned attention scores. These scores allow us to identify anatomical regions and fiber tracts relevant to sex differentiation, demonstrating biologically grounded predictions.

B. Interpretation and Biological Relevance

As described earlier, we extracted and analyzed attention scores from our multi-head GAT layer, allowing us to identify consistent patterns of focus. Most notably, regions such as the lateral ventricles and the right frontal pole consistently received high attention across all subjects. The lateral ventricles are fluid-filled cavities in the brain that contain cerebrospinal fluid [27]. Multiple studies have found that, on average, males exhibit larger overall brain volume [28]. Consequently, this

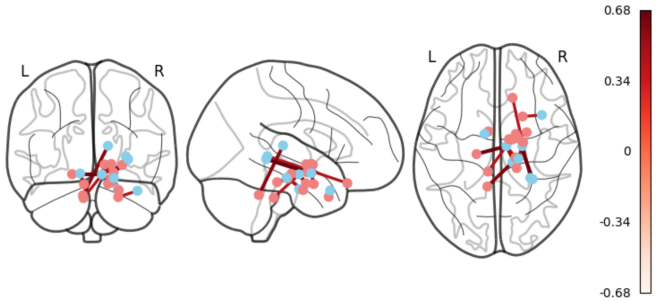

Fig. 5. Diagram depicting the top 15 edges with the highest attention weights. Cluster nodes are in blue, and ROI nodes are in red.

results in a larger ventricular volume, which is likely what our model has identified. On the other hand, the right frontal pole is a part of the prefrontal cortex, which is involved in many important functions such as emotions and strategic decision-making [29]. Prior neuroimaging studies have suggested differences in cortical thickness between sexes in the frontal areas, along with greater hemispheric connectivity shown in males over females [30] [31].

These findings are not only consistent with known neurobiological sex differences but also highlight our model's ability to focus on meaningful structural features without guidance. Furthermore, our heterogeneous graph approach allows the model to preserve fine-grained detail while simultaneously incorporating broader anatomical context, capturing richer patterns and relationships that are often lost in traditional machine learning approaches.

C. Limitations

While demonstrating promising results, our study does have potential limitations. The study was conducted exclusively on the HCP 1200 dataset exclusively of young healthy adults, in which only one type of scanner was utilized. Additionally, the data split did not explicitly account for family structure within HCP, as related participants may appear across training and test sets, which could lead to optimistic performance estimates. The lack of diversity may reduce generalizability to other older populations or acquisition protocols. Furthermore, although attention weights provide an interpretable view into model behavior, they do not necessarily indicate a cause-and-effect relationship. Moreover, while averaging attention scores across subjects reveals consistent trends, it may smooth over individual-specific patterns that could be important for subgroup analysis.

D. Future Directions

Building on this work, several directions for development exist. First, applying the proposed framework to more diverse datasets, including multi-site and multi-scanner acquisitions, would allow for evaluation of generalizability across different populations. Second, access to greater computational resources would enable the use of full-resolution diffusion MRI without downsampling, allowing for more accurate tractographies and

potentially more meaningful graphs. Finally, this study can be extended to investigate neurodegenerative diseases such as Alzheimer's Disease or Parkinson's, where structural change may be predictive of disease progression.

VI. Conclusion

In this work, we proposed a novel biologically informed heterogeneous graph framework for machine learning sex prediction. The framework jointly models both white matter fiber clusters and anatomical regions. Unlike traditional approaches, our method preserves rich microstructural detail. Our structural-aware graph formulation, along with our interpretable GNN architecture, outperformed classical machine learning approaches while using significantly fewer training samples. Importantly, our attention analysis revealed biologically meaningful regions, further validating our model's focus and potential for biological insight. Our work paves the way for more interpretable and biologically grounded models of brain connectivity.

Acknowledgment

This work used publicly available data from the Human Connectome Project.

References

[1] D. S. Bassett and O. Sporns, "Network neuroscience," Nature Neuroscience, vol. 20, no. 3, pp. 353–364, Mar. 2017, doi: 10.1038/nn.4502.

[2] D. S. Bassett, E. Bullmore, B. A. Verchinski, V. S. Mattay, D. R. Weinberger, and A. Meyer-Lindenberg, "Hierarchical Organization of Human Cortical Networks in Health and Schizophrenia," J. Neurosci., vol. 28, no. 37, pp. 9239–9248, Sep. 2008.

[3] P. Sorrentino et al., "The structural connectome constrains fast brain dynamics," eLife, vol. 10, p. e67400, Jul. 2021, doi: 10.7554/eLife.67400.

[4] D. M. Barch et al., "Function in the human connectome: Task-fMRI and individual differences in behavior," NeuroImage, vol. 80, pp. 169–189, Oct. 2013, doi: https://doi.org/10.1016/j.neuroimage.2013.05.033.

[5] Y. Sang and W. Li, "Classification Study of Alzheimer's Disease Based on Self-Attention Mechanism and DTI Imaging Using GCN," IEEE Access, vol. 12, pp. 24387–24395, 2024, doi: 10.1109/ACCESS.2024.3364545.

[6] Q. Li, "Functional connectivity inference from fMRI data using multivariate information measures," Neural Networks, Nov. 2021, doi: https://doi.org/10.1016/j.neunet.2021.11.016.

[7] L. Xiao et al., "Correlation Guided Graph Learning to Estimate Functional Connectivity Patterns from fMRI Data," IEEE Trans Biomed Eng, vol. 68, no. 4, pp. 1154–1165, Apr. 2021, doi: 10.1109/TBME.2020.3022335.

[8] M. Froeling, P. Pullens, and A. Leemans, "DTI Analysis Methods: Region of Interest Analysis." doi: 10.1007/978-1-4939-3118-7-9.

[9] J. Anbarasi, R. Kumari, M. Ganesh, and R. Agrawal, "Translational Connectomics: overview of machine learning in macroscale Connectomics for clinical insights," BMC Neurology, vol. 24, no. 1, p. 364, Sep. 2024, doi: 10.1186/s12883-024-03864-0.

[10] M. Zeraati and A. Davoodi, "ASD-GraphNet: A novel graph learning approach for Autism Spectrum Disorder diagnosis using fMRI data," Computers in Biology and Medicine, vol. 196, p. 110723, Sep. 2025, doi: 10.1016/j.compbiomed.2025.110723.

[11] L. Zhan et al., "Frontiers — Boosting brain connectome classification accuracy in Alzheimer's disease using higher-order singular value decomposition", Accessed: Sep. 24, 2025. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC4513242/

[12] L. Dodero, H. Q. Minh, M. S. Biagio, V. Murino, and D. Sona, "Kernel-based classification for brain connectivity graphs on the Riemannian manifold of positive definite matrices," in 2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI), Apr. 2015, pp. 42–45. doi: 10.1109/ISBI.2015.7163812.

[13] A. Kazi, J. Mora, B. Fischl, A. V. Dalca, and I. Aganj, "Multi-Head Graph Convolutional Network for Structural Connectome Classification," ArXiv, p. arXiv:2305.02199v2, 2023, Available: https://pubmed.ncbi.nlm.nih.gov/37205262/

[14] A. Kazi, J. Mora, B. Fischl, A. V. Dalca, and I. Aganj, "Structural Connectome Analysis using a Graph-based Deep Model for Age and Dementia Prediction," Mar. 2025, doi: 10.1101/2025.03.09.642165.

[15] G. Qu, Z. Zhou, V. D. Calhoun, A. Zhang, and Y.-P. Wang, "Integrated brain connectivity analysis with fMRI, DTI, and sMRI powered by interpretable graph neural networks," Medical Image Analysis, vol. 103, p. 103570, Apr. 2025, doi: https://doi.org/10.1016/j.media.2025.103570.

[16] S. Oh, S. Kim, J.-E. Lee, B.-Y. Park, J. H. Won, and H. Park, "Multimodal analysis of disease onset in Alzheimer's disease using Connectome, Molecular, and genetics data," NeuroImage Clinical, vol. 43, pp. 103660–103660, Jan. 2024, doi: https://doi.org/10.1016/j.nicl.2024.103660.

[17] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. J. Behrens, E. Yacoub, and K. Ugurbil, "The WU-Minn Human Connectome Project: An overview," NeuroImage, vol. 80, pp. 62–79, Oct. 2013, doi: https://doi.org/10.1016/j.neuroimage.2013.05.041.

[18] R. S. Desikan et al., "An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest," NeuroImage, vol. 31, no. 3, pp. 968–980, Jul. 2006, doi: https://doi.org/10.1016/j.neuroimage.2006.01.021.

[19] E. Garyfallidis et al., "Frontiers — Dipy, a library for the analysis of diffusion MRI data", Accessed: Sep. 24, 2025. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC3931231/

[20] P. J. Basser, J. Mattiello, and D. LeBihan, "MR diffusion tensor spectroscopy and imaging," Biophysical Journal, vol. 66, no. 1, pp. 259–267, Jan. 1994, doi: https://doi.org/10.1016/S0006-3495(94)80775-1.

[21] R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly, "Anatomically-constrained tractography: Improved diffusion MRI streamlines tractography through effective use of anatomical information," NeuroImage, vol. 62, no. 3, pp. 1924–1938, Sep. 2012, doi: https://doi.org/10.1016/j.neuroimage.2012.06.005.

[22] E. Garyfallidis, M. Brett, M. M. Correia, G. B. Williams, and I. Nimmo-Smith, "QuickBundles, a Method for Tractography Simplification," Frontiers in Neuroscience, vol. 6, 2012, doi: https://doi.org/10.3389/fnins.2012.00175.

[23] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive Representation Learning on Large Graphs".

[24] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph Attention Networks," Feb. 04, 2018, arXiv: arXiv:1710.10903. doi: 10.48550/arXiv.1710.10903.

[25] O. Al Zoubi et al., "Machine Learning Evidence for Sex Differences Consistently Influences Resting-State Functional Magnetic Resonance Imaging Fluctuations Across Multiple Independently Acquired Data Sets," Brain Connect, vol. 12, no. 4, pp. 348–361, May 2022, doi: 10.1089/brain.2020.0878.

[26] M. Ebel, M. Domin, N. Neumann, C. O. Schmidt, M. Lotze, and M. Stanke, "Classifying sex with volume-matched brain MRI," Neuroimage Rep, vol. 3, no. 3, p. 100181, Jul. 2023, doi: 10.1016/j.ynirp.2023.100181.

[27] C. L. Scelsi, T. A. Rahim, J. A. Morris, G. J. Kramer, B. C. Gilbert, and S. E. Forseen, "The Lateral Ventricles: A Detailed Review of Anatomy, Development, and Anatomic Variations," American Journal of Neuroradiology, vol. 41, no. 4, Feb. 2020, doi: https://doi.org/10.3174/ajnr.A6456.

[28] A. R. Erdogan, S. Dane, M. D. Aydin, M. Özdikici, and S. Diyarbakirl, "Sex and Handedness Differences In Size of Cerebral Ventricles of Normal Subjects," International Journal of Neuroscience, vol. 114, no. 1, pp. 67–73, Jan. 2004, doi: https://doi.org/10.1080/00207450490249428.

[29] W. R. Hathaway and B. W. Newton, "Neuroanatomy, Prefrontal Cortex," in StatPearls, Treasure Island (FL): StatPearls Publishing, 2025. Accessed: Sep. 24, 2025. [Online]. Available: http://www.ncbi.nlm.nih.gov/books/NBK499919/

[30] A. N. V. Ruigrok et al., "A meta-analysis of sex differences in human brain structure," Neurosci Biobehav Rev, vol. 39, no. 100, pp. 34–50, Feb. 2014, doi: 10.1016/j.neubiorev.2013.12.004.

[31] M. Ingalhalikar et al., "Sex differences in the structural connectome of the human brain," Proceedings of the National Academy of Sciences, vol. 111, no. 2, pp. 823–828, Dec. 2013, doi: https://doi.org/10.1073/pnas.1316909110.