

BrainMAP: Learning Multiple Activation Pathways in Brain Networks

Song Wang^{1*}, Zhenyu Lei^{1*}, Zhen Tan², Jiaqi Ding³, Xinyu Zhao³, Yushun Dong⁴,
Guorong Wu³, Tianlong Chen³, Chen Chen¹, Aiyong Zhang¹, Jundong Li¹

¹University of Virginia

²Arizona State University

³University of North Carolina at Chapel Hill

⁴Florida State University

{sw3wv, vjd5zr, zrh6du, xsa5db, jundong}@virginia.edu, ztan36@asu.edu, {jiaqid, xinyuzh, tianlong}@cs.unc.edu, yd24f@fsu.edu, guorong_wu@med.unc.edu

Abstract

Functional Magnetic Resonance Image (fMRI) is commonly employed to study human brain activity, since it offers insight into the relationship between functional fluctuations and human behavior. To enhance analysis and comprehension of brain activity, Graph Neural Networks (GNNs) have been widely applied to the analysis of functional connectivities (FC) derived from fMRI data, due to their ability to capture the synergistic interactions among brain regions. However, in the human brain, performing complex tasks typically involves the activation of certain *pathways*, which could be represented as paths across graphs. As such, conventional GNNs struggle to learn from these pathways due to the long-range dependencies of multiple pathways. To address these challenges, we introduce a novel framework **BrainMAP** to learn **M**ultiple **A**ctivation **P**athways in **B**rain networks. BrainMAP leverages sequential models to identify long-range correlations among sequentialized brain regions and incorporates an aggregation module based on Mixture of Experts (MoE) to learn from multiple pathways. Our comprehensive experiments highlight BrainMAP’s superior performance. Furthermore, our framework enables explanatory analyses of crucial brain regions involved in tasks.

Code — <https://github.com/LzyFischer/BrainMAP>

1 Introduction

Recently, significant research has focused on learning complex patterns in brain activities, which has promoted tasks such as cognitive process decoding (Li and Fan 2019; Thomas, Ré, and Poldrack 2022; Finn, Poldrack, and Shine 2023) and the diagnosis of mental health disorders (Jo, Nho, and Saykin 2019; Eslami et al. 2019). Generally, brain activities could be represented as functional magnetic resonance imaging (fMRI) data (Fox and Raichle 2007; Zhang, Ji, and Liu 2024), which measures blood-oxygen-level-dependent (BOLD) responses and reflects changes in metabolic demand associated with neural activity (Kohoutová et al. 2020; Davis et al. 2020). By leveraging fMRI’s unique blend of spatial and temporal characteristics, researchers can delve into the complexities of cognitive processes in the human

*Equal contribution.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

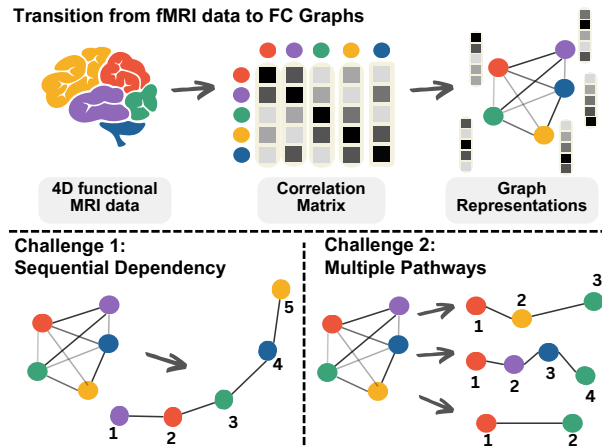


Figure 1: An illustration of the transition from fMRI data to FC graphs, along with two challenges for learning from pathways in FC graphs: (1) The sequential dependency, a fundamental feature of human brain activity, is not naturally presented in FC graphs; (2) Multiple pathways exist in FC graphs, making the extraction of them more difficult.

brain (Bassett and Sporns 2017). More specifically, BOLD signals are commonly used to construct networks of brain regions from fMRI data, where the functional connectivities (FC) among distinct brain regions are associated with various normal and pathological states (Kawahara et al. 2017), as shown in Fig. 1. Studying the FC features that correspond to the different brain states enables the identification of specific behavioral traits and neurological disorders linked to particular FC patterns (Morris et al. 2019).

To extract patterns in FC features, they are generally modeled as FC graphs, where nodes represent brain Regions of Interest (ROIs), and edges represent their relationships (Cui et al. 2022b,a). In this way, the correlations among brain regions could be explicitly represented (Said et al. 2023). With the development of Graph Machine Learning (GML) techniques, Graph Neural Networks (GNNs) are widely applied to FC graphs (Wang et al. 2022; Zhou et al. 2020). By capitalizing on the structured nature of the FC graphs and integrating local information, GNNs fa-

cilitate learning from patterns in functional connectivities and informative features (Li et al. 2021). While FC graphs offer valuable connectivity insights by depicting correlations among brain regions, existing works often overlook the *activation pathways* that are inherently present in these graphs. Specifically, in the human brain, performing tasks typically involves the activation of certain pathways (Sporns 2011), which could be represented as paths across the FC graphs (Sankar et al. 2018), as shown in Fig. 1. These pathways indicate the transmission of neural signals to a particular brain region. By incorporating these pathways into analysis, we could capture the complex interactions that might be overlooked when only considering pairwise correlations. Furthermore, these pathways provide insights into how different brain regions segregate into functional modules and collaborate to perform complex tasks.

However, despite the benefits of considering activation pathways on FC graphs, learning from these pathways is challenging. Since they are not explicitly represented in the graphs, without ground truth, models struggle to accurately learn and interpret them. Generally, pathways exhibit two crucial properties as shown in Fig. 1: (1) **Sequential dependency** is a fundamental feature of brain networks, where multiple regions co-activate and interact over long distances (Dahan et al. 2021). For example, in emotional memory processing, the hippocampus encodes memories, the amygdala assesses their emotional significance, and the prefrontal cortex uses this information for decision-making, exemplifying the long-range dependencies across multiple brain regions (Said et al. 2023). Nevertheless, while capturing these sequential dependencies is crucial for understanding the information flow in the brain, the structural nature of FC graphs makes it challenging to effectively model such dependencies. (2) **Multiple pathways** are generally necessary for the brain to process different behaviors and perform complex tasks. For example, in visual processing, the brain utilizes two parallel pathways: one along the dorsal visual cortex, which handles fast but coarse information, and the “what” stream along the ventral visual cortex, which processes slower but more detailed information (Lee et al. 2016). These distinct pathways correspond to different aspects of visual stimuli, emphasizing the need for multiple pathways in visual processing. Nevertheless, it is especially challenging to capture multiple pathways with existing GNN architectures. Due to the inherent limitations of the message-passing mechanism (Kipf and Welling 2017; Veličković et al. 2018; Tan et al. 2022), which focuses on aggregating information from neighboring nodes, GNNs struggle to effectively model the complex, long-range interactions in multiple pathways (Kim, Ye, and Kim 2021). Moreover, the interpretability of functional connectivity patterns is underexplored in current GNN-based approaches. Existing interpretable GNN models (Ying et al. 2019; Luo et al. 2020), which are typically designed to explain the importance of individual nodes and edges rather than considering their relationships within an activation path, struggle to provide explanations for interactions in long-range paths.

In this work, we propose **BrainMAP** to effectively learn from and interpret **M**ultiple **A**ctivation **P**athways present

in FC (functional connectivities) graphs while tackling the challenges posed by long-range dependencies and pathway correlations. To achieve this: (1) We propose an **Adaptive Graph Sequentialization** module to transform each FC graph into a node sequence that reflects the order of information flow, which enables the extraction of the hidden pathways that are crucial for modeling long-range interactions. (2) We design a **Hierarchical Pathway Integration** strategy that analyzes correlations among multiple pathways. Inspired by the human brain’s use of parallel pathways in complex tasks, we propose to integrate insights from diverse pathways, which captures complementary information contributed by each pathway. More importantly, our design improves interpretability by identifying the crucial brain regions in pathways that work together to support brain functions. Such interpretability offers deeper insights into the functional co-activated pattern of the brain. To evaluate our framework, We conduct experiments on five real-world fMRI datasets. The results demonstrate that our framework outperforms existing models in various prediction tasks on FC graphs while also offering comprehensive explanations for pathways. In summary, our contributions are as follows:

- **Innovation.** We present a novel framework for predictive tasks on FC graphs while providing comprehensive explanations to identify crucial brain regions—an area that has been underexplored in prior research.
- **Architecture.** We design an Adaptive Sequentialization module to transform FC graphs into node sequences for pathway learning, and a Pathway Integration module to aggregate and analyze correlations across multiple pathways on FC graphs.
- **Validation.** We conduct extensive experiments on various real-world FC datasets, and the results demonstrate the superior performance of our framework in both predictions and explanations.

2 Related Work

2.1 Brain Network Analysis

Brain network analysis involves understanding the intricate patterns of connectivity within the brain (Cui et al. 2022a; Kan et al. 2022; Zhang et al. 2022), which has gained increasing popularity recently due to its various applications, including identifying biomarkers for neurological diseases (Yang et al. 2022), understanding cognitive processes (Liu et al. 2023; Chen et al. 2024), and distinguishing different types of brain networks (Liao, Wan, and Du 2024). Among these, one of the most important tasks is the prediction of brain-related attributes, such as demographics and task states (Said et al. 2023; He et al. 2020). Recently, GNNs have significantly evolved as a major field of exploration for these tasks (Li et al. 2022; Cui et al. 2022a), due to their extraordinary ability to leverage the structured data (Li et al. 2021; Xu et al. 2024; Wang, Chen, and Li 2022). Nevertheless, GNN-based approaches often struggle to fully exploit the useful knowledge in brain networks, particularly the activation pathways that are inherently present in brains (Keller, Taube, and Lauber 2018) are neglected. To address this limitation, we propose to extract multiple underlying activation

pathways with adaptive structure sequentialization and Mixture of Experts (MoE), which thus enables a more comprehensive understanding of the brain connection.

2.2 Mixture of Experts

The Mixture of Experts (MoE) approach involves deploying a collection of expert networks, each designed to specialize in a particular task or a subset of the input space (Shazeer et al. 2017; Wang et al. 2024b). Originally derived from traditional machine learning models (Jacobs et al. 1991; Jordan and Jacobs 1994), MoE has since been adapted for deep learning, significantly enhancing its ability to handle complex vision and language tasks (Jiang et al. 2024). In addition to the strategy of interesting MoE layers with conventional neural networks (Vaswani et al. 2017; Dauphin et al. 2017), the concept of MoE is also extended to large and independent modules, e.g., language models as agents (Wang et al. 2024). In this work, we extend the MoE framework to address the challenge of multiple pathways in brain networks, focusing on learning the correlations across pathways. As a result, our framework is able to extract multiple pathways within and across different orders while learning from them.

3 Preliminary

We define an FC graph $G = (\mathcal{V}, \mathcal{E}, \mathbf{A}, \mathbf{X})$, where \mathcal{V} represents the set of nodes that indicate brain regions, \mathcal{E} denotes the edges that illustrate functional connections between these regions, $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the adjacency matrix capturing the connectivity structure, and $\mathbf{X} \in \mathbb{R}^{N \times d}$ denotes node features that may include various biological markers or other relevant attributes. $|\mathcal{V}| = N$ is the size of G , and d is the feature dimension size. We use Y to denote the prediction target of each graph in classification or regression tasks.

4 Methodology

As shown in Fig 2, BrainMAP is composed of two components: (1) *Adaptive Graph Sequentialization*, which learns the optimal sequence of brain regions by transforming the FC graph structure into a meaningful order that captures key dependencies, and (2) *Hierarchical Pathways Aggregation*, which utilizes multiple experts to extract diverse pathways from different orders and then aggregates them to capture complex interactions across multiple pathways. Each expert is instantiated as a sequential model, in order to extract long-range dependencies within potential pathways.

4.1 Adaptive Graph Sequentialization

When performing tasks such as visual or motor activities, research has shown that multiple brain regions often collaborate over long distances rather than functioning in isolation, which means cognitive processes emerge from the sequential activation of these regions (Thiebaut de Schotten and Forkel 2022). Consequently, capturing the order of sequential activation paths is crucial for accurate prediction in brain networks. Nevertheless, due to the complex structure of FC graphs, it is difficult to identify and extract such pathways. To address this, we propose an adaptive sequentialization strategy that transforms each FC graph into a sequence,

in order to preserve key pathway information and facilitate more effective modeling of the brain’s dynamic processes.

Learning Orders for FC Graphs. A significant obstacle in converting FC graphs into node sequences lies in the permutation invariance of brain network regions (i.e., nodes) (Said et al. 2023). This invariance contrasts with the inherently sequential nature of activation pathways, which do not naturally account for such invariance.

To tackle this, we introduce a learning-based strategy that utilizes an order-learning GNN to adaptively determine the node order for each input FC graph. With the order-learning GNN, we aim to learn the optimal sequence of nodes by arranging them based on their learned ordering scores in an ascending order. The benefit of using the learned scores to describe the order is that it avoids the massive search space of possible node orders (i.e., $N!$ for a graph of size N), which would otherwise make exhaustive search infeasible. In the following, we describe the process of learning the ordering scores with order-learning GNN. Given an input graph G , the ordering score $s_i \in \mathbb{R}$ of node v_i in G is learned as:

$$s_i = \text{GNN}_i(\mathcal{V}_i, \mathcal{E}_i, \mathbf{X}_i), \text{ where } \mathcal{V}_i = \mathcal{N}_i \cup \{v_i\}. \quad (1)$$

Here \mathbf{X}_i is the feature matrix of \mathcal{V}_i , which is the set of neighboring nodes of v_i . \mathcal{E}_i is the set of edges for nodes in \mathcal{V}_i . GNN_i is the order-learning GNN.

With the ordering scores $\{s_1, s_2, \dots, s_N\}$ of nodes in G , calculated in Eq. (1), we obtain the order $\hat{\phi}$ of N nodes $\{v_1, v_2, \dots, v_N\}$ as follows:

$$\hat{\phi} = (v_{\pi(1)}, v_{\pi(2)}, \dots, v_{\pi(N)}), \quad (2)$$

$$\text{where } \pi(i) = \underset{j \notin \{\pi(1), \pi(2), \dots, \pi(i-1)\}}{\text{argmin}} s_j.$$

Here π is a permutation of indices that sorts the scores $\{s_1, s_2, \dots, s_N\}$ in an ascending order. $\hat{\phi}$ denotes the obtained order of N nodes in the input graph.

Optimization of Order-Learning GNNs. To optimize the order-learning GNN, an obvious challenge is the lack of ground-truth orders. That being said, the optimal node order that consists of sufficient pathway information remains unavailable. Therefore, we propose to use the loss of BrainMAP output to select good and bad orders as the supervision signal. Intuitively, the orders that could provide smaller losses regarding the correct label (or ground-truth values in regression tasks) should be more similar to the optimal orders. In concrete, within each training step, we first randomly sample a batch of orders and compute their corresponding output. Then we select N_p orders with the smallest losses as positive samples (denoted as Φ_p), and select N_d orders with the largest losses as negative samples (denoted as Φ_n). Based on the concept of contrastive learning (Oord, Li, and Vinyals 2018; You et al. 2020; Xu et al. 2023, 2024), our optimization aims to increase the similarity between the learned order and the positive orders, while decreasing the similarity between the learned order and the negative orders. In this manner, we manage to gradually make the learned order approach the better orders during training.

$$\max \sum_{\phi \in \Phi_p} s(\hat{\phi}, \phi) - \lambda \sum_{\phi \in \Phi_n} s(\hat{\phi}, \phi), \quad (3)$$

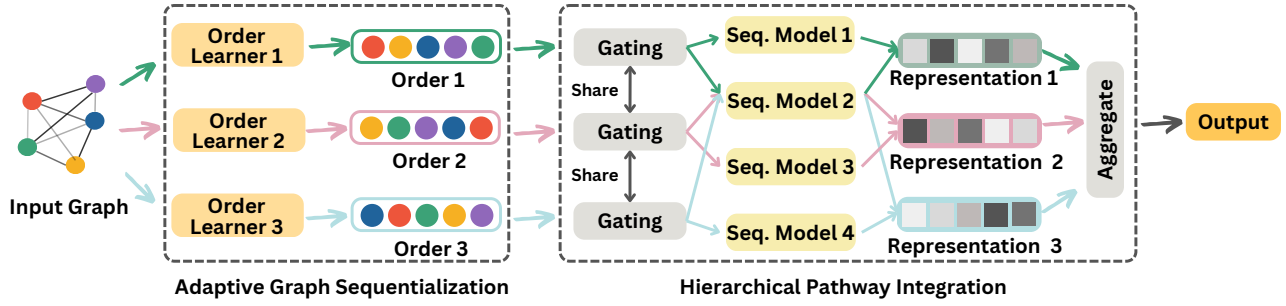


Figure 2: The overall process of BrainMAP. We first adaptively learn M ($M = 3$ in the figure) orders with three order-learners (GNNs). Then these orders are input into the gating function to select K ($K = 2$ in the figure) experts from a total number of P ($P = 4$ in the figure) experts. Each expert is implemented as a sequential model. The output of these experts will be aggregated into a representation for each order. Finally, the representations from all orders are aggregated into the output.

where λ determines the relative importance of the two objectives. $s(\hat{\phi}, \phi)$ denotes the similarity between $\hat{\phi}$ and ϕ .

To estimate the similarity between any two orders, a straightforward strategy is to leverage Spearman’s rank correlation coefficient (Spearman 1904). However, this coefficient is computed between two real ranks, each consisting of distinct values from 1 to N , which contrast the ordering scores we obtain. Moreover, directly converting the ordering scores into real ranks (i.e., from 1 to N) would prevent the flow of gradients, making optimization through gradient descent infeasible. To deal with this issue, we propose to calculate approximate rank scores as a substitute for the original ordering scores (i.e., s_i). These rank scores are differentiable and can be optimized using gradient descent.

Denoting any learned order $\hat{\phi}$, along with its ordering scores $\{s_1, s_2, \dots, s_N\}$, we calculate the approximate rank score $S_i \in \mathbb{R}$ of node v_i as follows:

$$S_i = \frac{s_i - \mathbb{E}[s]}{\sqrt{\mathbb{E}[s^2] - (\mathbb{E}[s])^2}} \cdot \sqrt{\frac{N^2 - 1}{12}} + \frac{(N + 1)(2N + 1)}{6},$$

where $\mathbb{E}[s] = \frac{1}{N} \sum_{i=1}^N s_i$, $\mathbb{E}[s^2] = \frac{1}{N} \sum_{i=1}^N s_i^2$.

(4)

In Eq. (4), the ordering score s_i is linearly transformed to S_i , based on the mean and variance of all N ordering scores, i.e., $\{s_1, s_2, \dots, s_N\}$. We perform such transformation to ensure that the mean and variance of $\{S_1, S_2, \dots, S_N\}$ are the same as those of a real rank variable of size N , which results in a similar distribution. Moreover, it also aligns with our loss design of enhancing the Spearman’s rank correlation coefficient (Spearman 1904), as introduced later. The consistency of mean and variance is verified in the following theorem.

Theorem 4.1. *The mean and standard deviation of S_i are the same as those of any real rank variable R for a sample size of N , i.e.,*

$$\mu(S_i) = \mu(R), \quad \sigma(S_i) = \sigma(R). \quad (5)$$

We provide the proof in Appendix A (in the Arxiv version of this work). According to Theorem 4.1, rank scores $\{S_1, S_2, \dots, S_N\}$ could represent an approximate rank as they share the same mean and variance.

Optimization Loss. For any randomly sampled (real) order ϕ , we use S_i^ϕ to represent the rank of v_i in ϕ . Since ϕ is a real order, we know S_i^ϕ is an integer and $1 \leq S_i^\phi \leq N$. Moreover, $S_i^\phi \neq S_j^\phi$ if $i \neq j$. Given the approximate rank scores $\{S_1, S_2, \dots, S_N\}$ of a learned order $\hat{\phi}$, we optimize it according to the following loss:

$$\mathcal{L}(\hat{\phi}, \Phi_p, \Phi_n) = \frac{\sum_{\phi \in \Phi_p} \sum_{i=1}^N (S_i - S_i^\phi)^2}{\sum_{i=1}^N \left(\sum_{\phi \in \Phi_p} (S_i - S_i^\phi)^2 + \sum_{\phi \in \Phi_n} (S_i - S_i^\phi)^2 \right)}, \quad (6)$$

Here Φ_p and Φ_n are the set of sampled positive orders and negative orders, respectively. $|\Phi_p| = N_p$ and $|\Phi_n| = N_n$. To validate the effectiveness of using loss $\mathcal{L}(\hat{\phi}, \Phi_p, \Phi_n)$ for optimization, we propose the following theorem.

Theorem 4.2. *Minimizing the loss \mathcal{L} described in Eq. (6) equals maximizing the Spearman’s rank correlation coefficient (Spearman 1904) between learned orders and good orders, while minimizing the coefficient between learned orders and bad orders.*

$$\min \mathcal{L}(\hat{\phi}, \Phi_p, \Phi_n) \equiv \max_{\phi \in \Phi_n} \frac{\sum_{\phi \in \Phi_n} (1 - r(\hat{\phi}, \phi))}{\sum_{\phi \in \Phi_p} (1 - r(\hat{\phi}, \phi))}, \quad (7)$$

where the coefficient is calculated as:

$$r(\hat{\phi}, \phi) = \rho_{S, S^\phi} = \frac{\text{cov}(S, S^\phi)}{\sigma_S \sigma_{S^\phi}}. \quad (8)$$

The proof of Theorem 4.2 is provided in Appendix B (in the Arxiv version of this work). According to Theorem 4.2, we know that optimizing the loss \mathcal{L} as described in Eq. (6) can increase the Spearman’s rank correlation coefficient between learned orders and good orders. Moreover, the objective also decreases the coefficient between learned orders and bad orders. In concrete, with the theoretical support from Theorem 4.2, we manage to optimize the order-learner with sampled good and bad orders.

Dataset	$ G $	$ N _{\text{avg}}$	$ E _{\text{avg}}$	d_{max}	d_{avg}	K_{avg}	d_x	#Classes	Prediction Task
HCP-Task	7,443	360	7,029.18	153	17.572	0.410	360	7	Graph Classification
HCP-Gender	1,078	1,000	45,578.61	413	45.579	0.466	1,000	2	Graph Classification
HCP-Age	1,065	1,000	45,588.40	413	45.588	0.466	1,000	3	Graph Classification
HCP-FI	1,071	1,000	45,573.67	413	45.574	0.466	1,000	-	Graph Regression
HCP-WM	1,078	1,000	45,578.61	413	45.579	0.466	1,000	-	Graph Regression

Table 1: The detailed statistics of datasets used in our experiments. $|G|$ denotes the number of graphs in each dataset, $|N|_{\text{avg}}$ and $|E|_{\text{avg}}$ represent the average number of nodes and edges, respectively. d signifies the degree, and K_{avg} represents the global clustering coefficient. The datasets encompass two types of prediction tasks: graph classification and graph regression.

4.2 Hierarchical Pathway Integration

Although sequential models can extract long-range pathways, they are inherently limited to identifying a single pathway at a time. In contrast, the human brain typically relies on multiple pathways to process various behaviors and perform complex tasks, as different pathways often contribute unique and complementary information (Morris et al. 2019). For instance, in visual processing, the brain employs two parallel pathways: one along the dorsal visual cortex, which quickly processes broad, less detailed information, and another one along the ventral visual cortex, which handles slower but more detailed information (Lee et al. 2016).

To deal with the challenge of multiple pathways, we propose to learn numerous activation pathways from each order of brain regions with multiple sequential models. Moreover, the activation pathways can be present in different orders. To effectively learn from these diverse pathways, we propose a two-level hierarchical integration approach, across and within different orders. (1) We first utilize the Mixture of Experts (MoE) strategy to integrate multiple pathways within each sequential order of brain regions. (2) Next, we aggregate the representations across different orders to obtain a comprehensive representation of brain activity.

► Step 1: Pathway Aggregation within Each Order.

Within each order, multiple sub-sequences may connect different sets of brain regions that appear as activation pathways, while they can be hard to extract with only one sequential model, due to the potential heterogeneity among pathways. Thus, we propose to learn multiple pathways simultaneously based on the MoE architecture, with each expert capturing different underlying pathways. To be specific, BrainMAP consists of multiple experts, each utilizing a different sequential model. To dynamically determine which experts are most suitable for a specific order, we design a gating function that ensures the similar pathways are consistently assigned to the same expert. In this manner, each expert specializes in capturing a specific type of pathway.

Formally, considering an input order $\hat{\phi}$ and P experts, the aggregation is performed as follows:

$$z' = \sigma \left(\sum_{i=1}^P G_i(\hat{\phi}) F_i(\hat{\phi}) \right), \quad (9)$$

where F_i is the sequential model of the i -th expert. G is the gating function that generates multiple decision scores with

the input as sequentialized brain regions $\hat{\phi}$, and $G(\hat{\phi}) \in \mathbb{R}^P$ denotes the scores to choose P experts for the graph. We employ an attention-based top-k gating design for G , which can be formalized with

$$\begin{aligned} G(\hat{\phi}) &= \text{Softmax}(\text{TopK}(Q(\hat{\phi}), K)), \\ Q(\hat{\phi}) &= \text{MLP}(\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V})), \\ \mathbf{Q} &= \mathbf{W}_Q \mathbf{h}, \quad \mathbf{K} = \mathbf{W}_K \mathbf{h}, \quad \mathbf{V} = \mathbf{W}_V \mathbf{h}, \\ \mathbf{h} &= \mathbf{W}_I \hat{\phi} + \text{PE}(\hat{\phi}), \end{aligned} \quad (10)$$

where PE denotes the sinusoidal positional encoding, which is utilized to inform the gating function with the order information of sequentialized graph representations. \mathbf{W}_I , \mathbf{W}_Q , \mathbf{W}_K , and \mathbf{W}_V are learnable parameters, and Attention denotes self-attention mechanism. Besides, K denotes the number of selected experts ($K \leq P$). $\text{TopK}(Q(\hat{\phi}), K)$ denotes that we keep the top K values in $Q(\hat{\phi})$, i.e.,

$$\begin{aligned} &\text{TopK}(Q(\hat{\phi}), K)_j \\ &= \begin{cases} Q(\hat{\phi})_j & \text{if } Q(\hat{\phi})_j \text{ is in the top } K \text{ values of } Q(\hat{\phi}), \\ -\infty & \text{otherwise.} \end{cases} \end{aligned} \quad (11)$$

► Step 2: Pathway Aggregation across Different Orders.

After the aggregation in Step 1, we obtain an output representation from each order. To aggregate the pathway information across different orders, we compute the weighted sum over representations learned from these orders, and the weights are the maximum value of $Q(\hat{\phi})$. In this manner, we achieve a final embedding for the input FC graph, i.e.,

$$z = \sum_{i=1}^M \text{Max}(Q(\hat{\phi}_i)) \cdot z'_i \quad (12)$$

where z'_i is the representation learned from the i -th order.

For the training of gating functions and experts in BrainMAP, we adopt the cross-entropy (CE) loss for classification and the mean absolute error (MAE) for regression tasks.

5 Experiments

In this section, we aim to answer the following research questions (RQs). **RQ1.** How well can BrainMAP perform on brain-related tasks compared to other alternatives? **RQ2.** How does each component contribute to the overall predictive performance? **RQ3.** How effectively can BrainMAP elucidate the rationale behind its predictive outcomes? **RQ4.** What impact does the design of MoE have on performance?

Dataset	HCP-Task \uparrow	HCP-Gender \uparrow	HCP-Age \uparrow	HCP-FI \downarrow	HCP-WM \downarrow
GCN	86.29 (± 0.98)	76.03 (± 2.40)	44.27 (± 2.69)	11.49 (± 0.15)	3.95 (± 0.05)
GAT	85.60 (± 1.26)	75.62 (± 2.22)	44.48 (± 2.35)	13.69 (± 0.52)	4.06 (± 0.11)
SAGE	84.49 (± 0.57)	74.69 (± 3.50)	45.83 (± 1.78)	<u>11.34</u> (± 0.12)	3.99 (± 0.06)
ResGCN	93.75 (± 0.35)	76.75 (± 0.65)	43.54 (± 0.90)	11.48 (± 0.29)	<u>3.92</u> (± 0.04)
GraphGPS	92.13 (± 2.00)	76.85 (± 1.54)	45.84 (± 3.21)	11.37 (± 0.86)	<u>3.98</u> (± 0.04)
Graph-Mamba	<u>94.17</u> (± 0.86)	<u>77.16</u> (± 3.13)	<u>46.35</u> (± 2.73)	11.51 (± 0.88)	3.94 (± 0.14)
BrainMAP	94.74 (± 0.07)	78.92 (± 0.49)	48.44 (± 1.65)	10.75 (± 0.61)	3.81 (± 0.03)

Table 2: Performance comparison of different models across various datasets. The best performance and the second-best performance are in bold and underlined, respectively. All experiments are repeated with 3 different random seeds.

Dataset	BrainMAP	w/o LR	w/o MoE	w/o LB
HCP-Task \uparrow	94.74	<u>94.56</u>	94.45	94.45
HCP-Gender \uparrow	78.92	78.09	77.58	<u>78.50</u>
HCP-Age \uparrow	48.44	<u>48.13</u>	47.81	47.54
HCP-FI \downarrow	3.81	3.84	3.91	3.91
HCP-WM \downarrow	10.75	<u>10.96</u>	11.19	11.01

Table 3: Ablation study of BrainMAP on various datasets.

5.1 Experimental Settings

We provide a brief introduction to the experimental settings. For the sequential model in our framework, we utilize the Mamba (Gu and Dao 2023), which is particularly effective in capturing long-range dependencies. The implementation details are explained in Appendix C (in the Arxiv version of this work).

Datasets. In our experiments, we consider the Human Connectome Project (HCP) dataset (Van Essen et al. 2013), which is a comprehensive publicly available neuroimaging dataset that includes both imaging data and a wide range of behavioral and cognitive data. We process the HCP-Task dataset by parcellating it into 360 distinct brain regions. With respect to other datasets, we use the processed ones from the NeuroGraph benchmark (Said et al. 2023).

Baselines. We compare our framework with baselines leveraged by the NeuroGraph benchmark and two state-of-the-art models GraphGPS (Rampásek et al. 2022) and Graph-Mamba (Wang et al. 2024a) that can extract long-range dependencies within the graph data.

5.2 Main Results

To answer **RQ1**, we first evaluate the performance of BrainMAP in comparison to all baselines on the HCP datasets. We make the following observations from empirical results in Table 2. (i) BrainMAP outperforms all baselines across various benchmarks with improvements up to 4.09% over the state-of-the-art, which demonstrates its ability to extract long-range dependencies between brain regions along multiple pathways. (ii) BrainMAP and Graph-Mamba surpass traditional GNN models by a large margin, with BrainMAP showing improvements of up to 12.13% on HCP-Task. The




Figure 3: Interpretation results of BrainMAP for the task MOTOR in HCP-Task. The average salient regions from random samples. The color bar ranges from 0.4 to 1. The bright-yellow color indicates a high score, while dark-red color indicates a relatively lower score. The ground-truth brain regions given by domain experts are circled in blue.

observation corroborates the benefit of extracting activation pathways for brain-related tasks. (iii) BrainMAP consistently outperforms Graph-Mamba, which demonstrates the effectiveness of learning multiple pathways.

5.3 Ablation Studies

To address **RQ2**, we conduct ablation studies on BrainMAP by removing different components, where *w/o LR* refers to the removal of the structure sequentializer and *w/o LB* indicates the exclusion of the load balancing loss for the MoE. The empirical results in Table 3 lead to the following observations. (i) Both the MoE and the structure sequentializer contribute to the overall performance, which suggests the importance of effective structure sequentialization and the extraction of multiple pathways. (ii) The MoE appears to be the most critical component for overall performance, indicating the significant role of learning multiple pathways. (iii) The removal of load balancing loss results in a reduction in the overall performance, which illustrates the importance of broad and balanced activation of experts.

Model	Hit@10	Hit@30	MRR
ResGCN	6.25	21.88	3.07
GraphGPS	8.75	24.38	3.13
Graph-Mamba	<u>15.00</u>	<u>31.25</u>	<u>6.27</u>
BrainMAP	19.38	33.75	9.26

Table 4: Interpretation results of BrainMAP for the task MOTOR from HCP-Task, where the alignment between salient brain regions obtained from different models and the ground truth given by domain experts is measured by three metrics.




Figure 4: The results of varying the number of experts in the MoE on four HCP benchmark datasets.

5.4 Explanation Study

To answer **RQ3** and better comprehend the prediction decisions made by different models, we aim to identify the salient brain regions that contribute the most to the predictions. To be more specific, we seek to identify the activated brain regions during a specific task MOTOR from the HCP-Task dataset. We first adopt explanation models to calculate the importance scores of brain regions during the task MOTOR of several random samples from the HCP-Task dataset, where the scores are then averaged to assess the interpretation ability. We adopt the commonly used GNNExplainer (Ying et al. 2019) for ResGCN and GraphGPS, and a Mamba-specific explanation method for Graph-Mamba and BrainMAP to calculate the importance scores. We select the salient brain regions, which are then compared to the ground-truth activated brain region of the HCP-Task given by domain experts, where the correspondence is measured with Hit@10, Hit@30, and Mean Reciprocal Rank (MRR). The results in Table 4 showcase that BrainMAP achieves higher precision in locating the activated brain regions for the MOTOR task, which demonstrates its reliability and effectiveness. Apart from the quantitative analysis of the interpretation ability of the BrainMAP, we also visualize the interpretation results in Fig. 3, where top-ranked brain regions of BrainMAP are highlighted with different colors, and ground-truth activated brain regions given by domain experts are circled. We could make the observation that BrainMAP is able to identify several ground-truth regions, which further demonstrates its effectiveness.

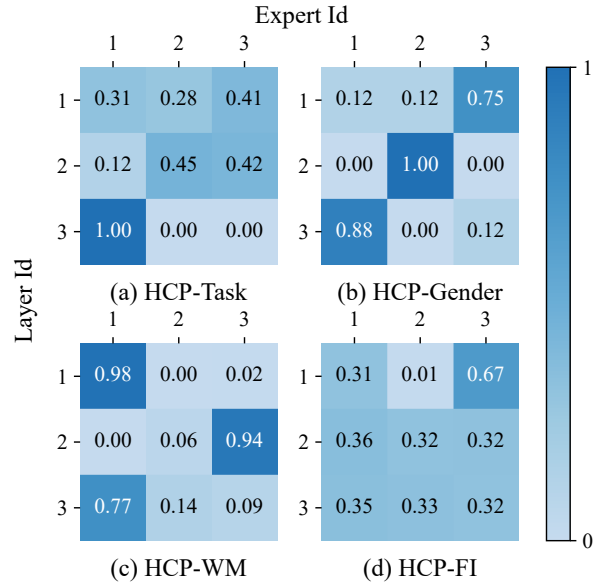


Figure 5: The activation distribution of the experts across different model layers. BrainMAP consistently maintains high activation rates on various HCP datasets.

5.5 MoE Analysis

The MoE is critical in extracting multiple pathways. To answer **RQ4**, we evaluate the impact of varying the number of experts in the MoE on the model’s performance. We could make the following observations from Fig. 4. (i) The performance of BrainMAP improves as the number of experts increases up to 3. It can be attributed to the fact that experts might be insufficient to extract diverse pathways necessary for comprehensive prediction. (ii) The accuracy remains nearly constant once the number of experts exceeds 4, which can be attributed to the limited number of potential pathways in the brain. To gain a deeper understanding of the MoE component, we analyze the activation distribution across different layers of BrainMAP, as shown in Fig. 5. The results illustrated that BrainMAP consistently maintains high activation rates, with a minimum of 66.7% activation rate for HCP-Gender. The findings further suggest that BrainMAP effectively extracts multiple pathways, as evidenced by the activation of diverse experts.

6 Conclusion

Despite significant progress has been made in understanding brain activity through functional connectivity (FC) graphs, challenges persist in effectively capturing and interpreting the complex, long-range dependencies and multiple pathways that are inherent in these graphs. In this work, we introduce BrainMAP, a novel framework designed to extract multiple long-range activation pathways with adaptive sequentialization and pathway aggregation. Experiments demonstrate the effectiveness of BrainMAP in extracting underlying activation pathways for predictions tasks.

Acknowledgments

This work is supported in part by the National Science Foundation under grants IIS-2006844, IIS-2144209, IIS-2223769, IIS-2331315, CNS-2154962, BCS-2228534, and CMMI-2411248, the Commonwealth Cyber Initiative Awards under grants VV-1Q24-011, VV-1Q25-004, and the research gift funding from Netflix and Snap. Xinyu Zhao and Tianlong Chen are supported by NIH OT2OD038045-01 and UNC SDSS Seed Grant.

References

- Bassett, D. S.; and Sporns, O. 2017. Network neuroscience. *Nature neuroscience*, 20(3): 353–364.
- Chen, J.; Qi, Y.; Wang, Y.; and Pan, G. 2024. Bridging the Semantic Latent Space between Brain and Machine: Similarity Is All You Need. *AAAI*, 38: 11302–11310.
- Cui, H.; Dai, W.; Zhu, Y.; Kan, X.; Gu, A. A. C.; Lukemire, J.; Zhan, L.; He, L.; Guo, Y.; and Yang, C. 2022a. Braingb: a benchmark for brain network analysis with graph neural networks. *IEEE transactions on medical imaging*.
- Cui, H.; Lu, Z.; Li, P.; and Yang, C. 2022b. On positional and structural node features for graph neural networks on non-attributed graphs. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 3898–3902.
- Dahan, S.; Williams, L. Z.; Rueckert, D.; and Robinson, E. C. 2021. Improving phenotype prediction using long-range spatio-temporal dynamics of functional connectivity. In *Machine Learning in Clinical Neuroimaging: 4th International Workshop*.
- Dauphin, Y. N.; Fan, A.; Auli, M.; and Grangier, D. 2017. Language modeling with gated convolutional networks. In *International conference on machine learning*.
- Davis, K. D.; Aghaeepour, N.; Ahn, A. H.; Angst, M. S.; Borsook, D.; Brenton, A.; Burczynski, M. E.; Crean, C.; Edwards, R.; Gaudilliere, B.; et al. 2020. Discovery and validation of biomarkers to aid the development of safe and effective pain therapeutics: challenges and opportunities. *Nature Reviews Neurology*, 16(7): 381–400.
- Eslami, T.; Mirjalili, V.; Fong, A.; Laird, A. R.; and Saeed, F. 2019. ASD-DiagNet: a hybrid learning approach for detection of autism spectrum disorder using fMRI data. *Frontiers in neuroinformatics*, 13: 70.
- Finn, E. S.; Poldrack, R. A.; and Shine, J. M. 2023. Functional neuroimaging as a catalyst for integrated neuroscience. *Nature*, 623(7986): 263–273.
- Fox, M. D.; and Raichle, M. E. 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature reviews neuroscience*.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- He, T.; Kong, R.; Holmes, A. J.; Nguyen, M.; Sabuncu, M. R.; Eickhoff, S. B.; Bzdok, D.; Feng, J.; and Yeo, B. T. 2020. Deep neural networks and kernel regression achieve comparable accuracies for functional connectivity prediction of behavior and demographics. *NeuroImage*.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1): 79–87.
- Jiang, A. Q.; Sablayrolles, A.; Roux, A.; Mensch, A.; Savary, B.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Hanna, E. B.; Bressand, F.; et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Jo, T.; Nho, K.; and Saykin, A. J. 2019. Deep learning in Alzheimer’s disease: diagnostic classification and prognostic prediction using neuroimaging data. *Frontiers in aging neuroscience*, 11: 220.
- Jordan, M. I.; and Jacobs, R. A. 1994. Hierarchical mixtures of experts and the EM algorithm. *Neural computation*.
- Kan, X.; Dai, W.; Cui, H.; Zhang, Z.; Guo, Y.; and Yang, C. 2022. Brain network transformer. *NeurIPS*, 35: 25586–25599.
- Kawahara, J.; Brown, C. J.; Miller, S. P.; Booth, B. G.; Chau, V.; Grunau, R. E.; Zwicker, J. G.; and Hamarneh, G. 2017. BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage*, 146: 1038–1049.
- Keller, M.; Taube, W.; and Lauber, B. 2018. Task-dependent activation of distinct fast and slow (er) motor pathways during motor imagery. *Brain stimulation*, 11(4): 782–788.
- Kim, B.-H.; Ye, J. C.; and Kim, J.-J. 2021. Learning dynamic graph representation of brain connectome with spatio-temporal attention. *NeurIPS*, 34: 4314–4327.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Kohoutová, L.; Heo, J.; Cha, S.; Lee, S.; Moon, T.; Wager, T. D.; and Woo, C.-W. 2020. Toward a unified framework for interpreting machine-learning models in neuroimaging. *Nature protocols*, 15(4): 1399–1435.
- Lee, W.-C. A.; Bonin, V.; Reed, M.; Graham, B. J.; Hood, G.; Glatfelter, K.; and Reid, R. C. 2016. Anatomy and function of an excitatory network in the visual cortex. *Nature*, 532(7599): 370–374.
- Li, H.; and Fan, Y. 2019. Interpretable, highly accurate brain decoding of subtly distinct brain states from functional MRI using intrinsic functional networks and long short-term memory recurrent neural networks. *NeuroImage*.
- Li, X.; Zhou, Y.; Dvornek, N.; Zhang, M.; Gao, S.; Zhuang, J.; Scheinost, D.; Staib, L. H.; Ventola, P.; and Duncan, J. S. 2021. Braingnn: Interpretable brain graph neural network for fmri analysis. *Medical Image Analysis*, 74: 102233.
- Li, Y.; Zhang, X.; Nie, J.; Zhang, G.; Fang, R.; Xu, X.; Wu, Z.; Hu, D.; Wang, L.; Zhang, H.; et al. 2022. Brain connectivity based graph convolutional networks and its application to infant age prediction. *IEEE transactions on medical imaging*, 41(10): 2764–2776.
- Liao, M.; Wan, G.; and Du, B. 2024. Joint Learning Neuronal Skeleton and Brain Circuit Topology with Permutation Invariant Encoders for Neuron Classification. In *AAAI*.
- Liu, X.; Zhou, M.; Shi, G.; Du, Y.; Zhao, L.; Wu, Z.; Liu, D.; Liu, T.; and Hu, X. 2023. Coupling artificial neurons in bert and biological neurons in the human brain. In *AAAI*.

- Luo, D.; Cheng, W.; Xu, D.; Yu, W.; Zong, B.; Chen, H.; and Zhang, X. 2020. Parameterized explainer for graph neural network. *NeurIPS*, 33: 19620–19631.
- Morris, C.; Ritzert, M.; Fey, M.; Hamilton, W. L.; Lenssen, J. E.; Rattan, G.; and Grohe, M. 2019. Weisfeiler and leman go neural: Higher-order graph neural networks. In *AAAI*, volume 33, 4602–4609.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv*.
- Rampásek, L.; Galkin, M.; Dwivedi, V. P.; Luu, A. T.; Wolf, G.; and Beaini, D. 2022. Recipe for a general, powerful, scalable graph transformer. *NeurIPS*, 35: 14501–14515.
- Said, A.; Bayrak, R.; Derr, T.; Shabbir, M.; Moyer, D.; Chang, C.; and Koutsoukos, X. 2023. Neurograph: Benchmarks for graph machine learning in brain connectomics. *NeurIPS*.
- Sankar, A.; Wu, Y.; Gou, L.; Zhang, W.; and Yang, H. 2018. Dynamic graph representation learning via self-attention networks. *arXiv preprint arXiv:1812.09430*.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Spearman, C. 1904. The Proof and Measurement of Association between Two Things. *The American Journal of Psychology*.
- Sporns, O. 2011. The human connectome: a complex network. *Annals of the new York Academy of Sciences*.
- Tan, Z.; Wang, S.; Ding, K.; Li, J.; and Liu, H. 2022. Transductive Linear Probing: A Novel Framework for Few-Shot Node Classification. In *LoG*.
- Thiebaut de Schotten, M.; and Forkel, S. J. 2022. The emergent properties of the connected brain. *Science*.
- Thomas, A. W.; Ré, C.; and Poldrack, R. A. 2022. Interpreting mental state decoding with deep learning models. *Trends in Cognitive Sciences*, 26(11): 972–986.
- Van Essen, D. C.; Smith, S. M.; Barch, D. M.; Behrens, T. E.; Yacoub, E.; Ugurbil, K.; Consortium, W.-M. H.; et al. 2013. The WU-Minn human connectome project: an overview. *Neuroimage*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2018. Graph attention networks. In *ICLR*.
- Wang, C.; Tsepa, O.; Ma, J.; and Wang, B. 2024a. Graphmamba: Towards long-range graph sequence modeling with selective state spaces. *arXiv preprint arXiv:2402.00789*.
- Wang, R.; An, S.; Cheng, M.; Zhou, T.; Hwang, S. J.; and Hsieh, C.-J. 2024b. One Prompt is not Enough: Automated Construction of a Mixture-of-Expert Prompts. *arXiv preprint arXiv:2407.00256*.
- Wang, S.; Chen, C.; and Li, J. 2022. Graph Few-shot Learning with Task-specific Structures. In *NeurIPS*.
- Wang, S.; Chen, Z.; Shi, C.; Shen, C.; and Li, J. 2024. Mixture of Demonstrations for In-Context Learning. In *NeurIPS*.
- Wang, X.; Yao, L.; Reikik, I.; and Zhang, Y. 2022. Contrastive functional connectivity graph learning for population-based fMRI classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 221–230. Springer.
- Xu, Y.; Peng, Z.; Shi, B.; Hua, X.; and Dong, B. 2024. Learning dynamic graph representations through timespan view contrasts. *Neural Networks*, 176: 106384.
- Xu, Y.; Shi, B.; Ma, T.; Dong, B.; Zhou, H.; and Zheng, Q. 2023. CLDG: Contrastive learning on dynamic graphs. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 696–707. IEEE.
- Yang, Y.; Zhu, Y.; Cui, H.; Kan, X.; He, L.; Guo, Y.; and Yang, C. 2022. Data-efficient brain connectome analysis via multi-task meta-learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4743–4751.
- Ying, Z.; Bourgeois, D.; You, J.; Zitnik, M.; and Leskovec, J. 2019. Gnnexplainer: Generating explanations for graph neural networks. *NeurIPS*.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *NeurIPS*.
- Zhang, X.; Wang, S.; Lin, N.; Zhang, J.; and Zong, C. 2022. Probing word syntactic representations in the brain by a feature elimination method. In *AAAI*, volume 36, 11721–11729.
- Zhang, Z.; Ji, J.; and Liu, J. 2024. MetaRLEC: Meta-Reinforcement Learning for Discovery of Brain Effective Connectivity. In *AAAI*.
- Zhou, J.; Cui, G.; Hu, S.; Zhang, Z.; Yang, C.; Liu, Z.; Wang, L.; Li, C.; and Sun, M. 2020. Graph neural networks: A review of methods and applications. *AI open*, 1: 57–81.