

Distribution Agreement

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Tianyi Zhang

April 17, 2025

Causal Brain Connectivity: Integrating Granger Directed Graphs in fMRI Analysis

By

Tianyi Zhang

Carl Yang
Advisor

The Department of Quantitative Theory & Methods

Carl Yang
Advisor

Jinho D. Choi
Committee Member

Wei Jin
Committee Member

2025

Causal Brain Connectivity: Integrating Granger Directed Graphs in fMRI Analysis

By

Tianyi Zhang

Carl Yang
Advisor

An abstract of
A thesis submitted to the Faculty of the Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements for the degree of
Bachelor of Science with Honors

The Department of Quantitative Theory & Methods

2025

Abstract

Causal Brain Connectivity: Integrating Granger Directed Graphs in fMRI Analysis By Tianyi Zhang

Functional Magnetic Resonance Imaging (fMRI) has significantly advanced our understanding of human brains by capturing dynamic neural activities, providing basis for causal analysis between brain regions. However, conventional correlation-based analyses often fail to account for the directionality and complexity of neural interactions. We propose an approach that integrates Granger causality with graph-based deep learning to better capture *effective connectivity* between brain regions. Specifically, we compare three methods: MLP-based approaches on flattened time series, Graph Convolutional Networks (GCNs) using undirected connectivity, and a GCN framework incorporating directed Granger-causal influences into brain graph construction. Through the optimization of Granger parameters such as the lag order via Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC), we investigate the impact of different graph construction methods on connectome-based outcome prediction. The directed graph framework demonstrates robustness to hyperparameter variations, while also providing biologically plausible insights into brain functionalities that complement undirected correlation-based graphs. Evaluations on classification and regression tasks using large-scale fMRI datasets reveal that directionality preserves predictive performance while offering additional understanding of information flow within brain networks. These findings emphasize the potential of Granger-causality-informed graphs for robust, nuanced, and causality-aware fMRI analyses.

Causal Brain Connectivity: Integrating Granger Directed Graphs in fMRI Analysis

By

Tianyi Zhang

Carl Yang
Advisor

A thesis submitted to the Faculty of the Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements for the degree of
Bachelor of Science with Honors

The Department of Quantitative Theory & Methods

2025

Acknowledgments

The authors gratefully acknowledge the support for the collection of the Philadelphia Neurodevelopmental Cohort (PNC) dataset, provided by Grant RC2MH089983 awarded to Raquel Gur and RC2MH089924 awarded to Hakon Hakorson.

The Adolescent Brain Cognitive Development (ABCD) Study is supported by the National Institutes of Health and additional federal partners. A full list of supporters is available at <https://abcdstudy.org/federal-partners.html>. Data used in this work were obtained from the ABCD Study (<https://abcdstudy.org>), held in the NIMH Data Archive. This multisite, longitudinal study is designed to recruit more than 10,000 children aged 9–10 and follow them over 10 years into early adulthood. ABCD consortium investigators designed and implemented the study and/or provided data but did not necessarily participate in the analysis or writing of this report.

This research was supported in part by Emory Brain Network Lab and GPU servers provided by the Computer Science Department of Emory University. The authors also gratefully acknowledge support from the NIH under award numbers R01MH105561 and R01MH118771. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

We gratefully acknowledge the contributions of the following individuals: Keqi Han, Jiawei Nie, and Carl Yang (Department of Computer Science, Emory University, Atlanta, GA, USA); Chenyu You (Department of Computer Science, Stony Brook University, Stony Brook, NY, USA); and Sanne van Rooij, Jennifer Stevens, Boadie Dunlop, and Charles Gillespie (Department of Psychiatry and Behavioral Sciences, Emory University, Atlanta, GA, USA).

Contents

1	Introduction	1
2	Related Works	3
2.1	Graph Neural Networks for Brain Connectome Analysis	3
2.2	Functional and Effective Connectivity	4
2.3	Causality and Directionality in Brain Networks	5
3	Problem Formulation	7
4	Method	10
4.1	MLP-Based Methods	10
4.2	Graph-Based Models	11
4.2.1	Undirected Graphs: Functional Connectivity.	11
4.2.2	Directed Graphs: Effective Connectivity.	12
4.3	Rationale for Selecting Granger Causality	14
4.3.1	Overview of Connectivity Inference Methods	15
5	Experiments	16
5.1	Experimental Settings	17
5.2	Model Performance (RQ1)	19
5.2.1	Graph-Based Models Outperform MLP.	19
5.2.2	Compatibility of Directed and Undirected Graphs.	19

5.3	Hyperparameter Study for Directed Graphs (RQ2)	20
5.3.1	Effects of Window Size, Step, and Lag on Directed Connectivity Patterns	22
5.4	Neurological Insights (RQ3)	27
5.4.1	Comparison of Undirected and Directed Connectivity Matrices	27
5.4.2	Consistent Directed Connectivity Patterns in the ABCD Cohort	32
6	Analysis	34
7	Conclusion	37
A	Granger Causality Algorithms	39
A.1	Data Partitioning	39
A.2	VAR Modeling and Hypothesis Testing	40
A.3	Aggregation of Test Results	41
A.4	Granger Causality Graph Construction	42
B	Model Description	44
B.1	Multilayer Perceptron (MLP)	44
B.2	Graph Convolutional Network (GCN)	44
	Bibliography	46

List of Figures

5.1	<i>Directed Connectivity Under Different Hyperparameters.</i> Each matrix is a Granger-causal adjacency (rescaled to $[0,1]$) using different window sizes (W), step sizes (S), and lag orders (L). Overall connectivity structure remains similar, despite parameter changes.	23
5.2	<i>Both matrices are derived from the ABCD dataset (HCP 360 atlas).</i> The undirected matrix (e) is correlation-based; the directed matrix (f) stems from Granger causality. Each is normalized to $[0,1]$. White grid lines outline functional networks (SM, DMN, VS, CE, DS, Vis).	27
5.3	<i>Chord Plots Showing the Top 2% of Edges.</i> Each segment represents a network; arcs link strongly connected or causally influencing ROIs.	28
5.4	<i>Chord Diagrams for Three Randomly Selected Participants in the ABCD Dataset.</i> (i), (j), and (k) show correlation-based undirected connectivity (top row) and the corresponding Granger-causal directed connectivity (bottom row). Edges are thresholded at the top 2% for visualization.	32

List of Tables

4.1	Comparison of connectivity inference methods for fMRI analysis. . . .	15
5.1	Summary of Datasets and Tasks	17
5.2	Performance Comparison Across Models and Tasks	19
5.3	Grid Search Results for Granger Causality Hyperparameters	21

Chapter 1

Introduction

Functional magnetic resonance imaging (fMRI) has significantly advanced our understanding of brain function by non-invasively measuring blood oxygen level-dependent signals (BOLD), thus capturing *dynamic* aspects of neural activity. However, many existing analyses continue to rely on *static, undirected pairwise correlations* between regions of interest (ROIs), overlooking the *directional* and *causal* nature of neural interactions. This omission limits our ability to determine how brain regions *causally influence* each other and restricts our understanding of the underlying mechanisms of neural communication.

Recent advances in *effective connectivity* emphasize modeling how activity causally propagates between brain regions, rather than merely identifying co-activation patterns [11, 61]. *Granger causality (GC)* is a well-established technique for estimating effective connectivity that provides a robust statistical framework to infer directed influences from time-series data, offering deeper insights into the flow and directionality of neural interactions than correlation-based methods [14, 26, 58]. In recent years, Graph Neural Networks (GNNs) have attracted broad interest for modeling graph-structured interactions between brain regions [12, 30, 31, 73, 52, 72]. Several pioneering methods have been developed for cognition assessment and neural disorder detection

[12, 31, 39, 75]. However, to the best of our knowledge, most existing neuroimaging applications of GNNs are built on correlation-based brain connectivity, which overlooks the causal directionality inherent in brain activity.

In this work, we propose a novel framework that integrates Granger-causal directed graphs with Graph Convolutional Networks (GCNs) to rigorously model *effective connectivity* in fMRI data. We validate our approach on the ABCD and PNC datasets, large neurodevelopmental cohorts well-suited for studying how functional networks evolve over development. Our approach centers on four insights. First, while correlation-based connectivity captures co-activation patterns, it overlooks the *causal directionality* essential for understanding how neural signal propagate. We address this by constructing *directed adjacency matrices* via Granger causality, capturing asymmetrical information flow across brain regions. Second, we systematically tune Granger parameters (e.g., lag order) using the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC), ensuring both statistical validity and scalability to large fMRI datasets. Third, integrating these Granger-causal graphs enables GCNs to explicitly model directed causal influences in brain networks, leveraging both local and global connectivity patterns for more accurate clinical outcome prediction. Finally, our Granger-causality-informed GCN directly models *directed, causal influences*, offering a complementary perspective to traditional correlation-based methods. Comprehensive validation shows that incorporating causality-guided connections delivers robust, competitive performance across diverse clinical tasks, while revealing deeper insights into neural information flow. Overall, our contributions bridge the gap between statistical causal inference and graph-based deep learning, addressing limitations of static functional connectivity analysis and paving the way for a new generation of *causality-aware* fMRI studies.

Chapter 2

Related Works

Understanding the complex interactions within the brain has been a central objective in neuroscience. Recent advances in deep learning and brain analysis methods now enable more powerful modeling and interpretation of neural connectivity. This section reviews three important research areas that underlie our study: *Graph Neural Networks for brain connectome analysis*, *functional and effective connectivity*, and *causality and Directionality in brain networks*.

2.1 Graph Neural Networks for Brain Connectome Analysis

Graph-based representations have become increasingly prominent in neuroscience for capturing the complexity of inter-regional communications in the brain [12, 31, 32, 39, 67, 77]. In these representations, each brain region is modeled as a node, while edges capture statistical or computational relationships (e.g., correlations or causal influences) between regions of interest (ROIs) [32, 77, 12, 31]. Conventional machine learning techniques often process neuroimaging data as high-dimensional vectors, thereby neglecting the inherent topological structure and interdependencies

among brain regions. In contrast, Graph Neural Networks (GNNs) explicitly treat the connectome as a graph, allowing them to leverage both local and global connectivity patterns.

Recent GNN-based approaches have demonstrated promise in both classification tasks, such as diagnostic prediction, and regression tasks, such as predicting cognitive scores, by exploiting these topological features [12, 31, 39]. Typically, these models learn node and graph embeddings via iterative message-passing and aggregation operations. This process not only preserves the overall structure of the connectome, but also highlights critical local connectivity patterns. For example, BrainNetCNN [34] introduces specialized convolutional filters designed specifically for brain connectivity matrices, while FCNet [48] processes pairwise correlation maps to capture functional connections between ROIs. More recent methodologies, such as BrainGB [12], have built on these ideas by integrating domain-specific anatomical or functional constraints within scalable GNN frameworks. However, many of these approaches predominantly utilize *undirected* functional connectivity graphs, thereby overlooking the potential directional nature of neural information flow.

2.2 Functional and Effective Connectivity

In neuroimaging, *functional connectivity* is typically defined through correlation-based measures that capture synchronized or coactivated neural activity patterns across different regions of the brain [4, 6, 19, 50, 61, 65]. This is often mathematically instantiated using the Pearson correlation coefficient or other similar metrics applied to time series data. Functional connectivity has been instrumental in the identification of large-scale networks such as the Default Mode Network (DMN), Visual Network, and Sensorimotor Network [2, 27, 29, 57, 60, 64]. However, while such correlation-based approaches offer valuable insights into co-activation patterns during resting-state or

task-driven scenarios, they fall short in elucidating the temporal order or directional causal influence between brain regions.

In contrast, *effective connectivity* focuses on uncovering the directional and causal driving forces between ROIs [11, 61, 75]. This approach models the dynamic influence that the activity in one brain region exerts over another. Granger causality is one of the most widely adopted methods for this purpose [14, 26, 58, 69]. It employs a Vector Autoregressive (VAR) framework to test whether past activity in one region can enhance the prediction of future activity in another, thus offering directed edges that represent putative causal influences [15, 56, 58, 49]. This distinction is crucial, as it empowers our approach to transcend simple correlational analyses and rigorously infer the causal mechanisms underpinning neural communication.

2.3 Causality and Directionality in Brain Networks

Shifting the focus to the neural foundations, the inherent directionality of brain connectivity is deeply rooted in the brain’s anatomical and physiological architecture [3, 10, 71]. At the core of this dynamic is the interplay between feedforward and feedback pathways [5, 8, 54, 78]. In the brain, sensory information is typically transmitted from primary sensory areas to higher-order association cortices through feedforward connections, while feedback pathways modulate and refine this information by projecting signals back to earlier stages [5, 8, 29, 54, 78]. This hierarchical organization supports complex cognitive functions and adaptive behavior.

At the microscopic level, the structural connectivity—formed by white matter tracts and synaptic networks—sets the stage for these directional interactions [13, 28, 35, 74]. Neurons form circuits that are not only spatially distributed but also functionally specialized, where the direction of information flow is determined by synaptic strengths, inhibitory and excitatory balances, and the intrinsic properties of neural populations

[46, 62]. Advanced neuroimaging techniques, such as network-level fMRI, enable us to capture these dynamics at a macroscopic scale, revealing how these microcircuit principles manifest in large-scale brain networks [18, 78].

By integrating causal inference methods into these analyses, researchers can disentangle the directional influences between brain regions. For example, combining noise-diffusion models or structural equation modeling with graph neural networks (GNNs) can help trace the flow of information, and map out how disruptions in these directed networks might contribute to neurological disorders or psychopathology [24, 75]. This approach not only provides a more refined understanding of neural communication but also bridges the gap between neurobiological mechanisms and their observable impact on cognition and behavior.

Thus, moving beyond traditional correlation analyses, modern causal inference techniques provide a powerful framework to decode the brain’s directional connectivity. They offer critical insights into both the flow of neural information and the alterations that may underpin various clinical conditions.

In this work, we build upon these advances to propose a *direction-aware* approach for fMRI analysis, incorporating Granger-causal edges into GNN-based models. By fusing effective connectivity estimates with the representational power of GNNs, our framework aims to preserve the predictive strengths of graph-based learning while uncovering unique insights into the causal orchestration of neural activity across the human connectome.

Chapter 3

Problem Formulation

This study addresses the challenge of downstream prediction using resting-state fMRI (rs-fMRI) data, taking into account the *directional* and *causal* nature of neural interactions. Traditional functional connectivity measures, such as Pearson’s correlation, often capture only undirected relationships, potentially overlooking the temporal dynamics that govern how neural signals propagate. By incorporating *directionality* and *causal inference* into graph-based connectivity representations [9, 16], we aim to more accurately model the functional organization of the brain and thereby improve predictive performance on tasks ranging from classification (e.g., diagnostic labels or gender) to regression (e.g., cognitive or clinical scores).

Data Representation and Notation

In rs-fMRI experiments, Blood Oxygen Level-Dependent (BOLD) signals are sampled over time from N distinct Regions of Interest (ROIs). Thus, for each subject m , the raw data can be represented as a matrix

$$\mathbf{X}^{(m)} \in \mathbb{R}^{N \times T}, \tag{3.1}$$

where $x_{i,t}^{(m)}$ denotes the BOLD intensity in ROI i at time step t , and T is the total number of time points. Each subject is associated with a label $y^{(m)}$, which can take the form of either a continuous value (e.g., a cognitive score) or a categorical label (e.g., diagnostic class or behavioral trait). The supervised learning objective is to predict:

$$\hat{y}^{(m)} = f(\mathbf{X}^{(m)}), \quad (3.2)$$

where f is a predictive function learned from a training set of $(\mathbf{X}^{(m)}, y^{(m)})$ pairs.

Graph-Based Connectivity Representations

To better capture the interplay among ROIs, we transform each subject’s rs-fMRI data into a graph structure. Specifically, we consider a graph

$$\mathcal{G}^{(m)} = (\mathcal{V}, \mathcal{E}^{(m)}),$$

where each node $v_i \in \mathcal{V}$ corresponds to an ROI, and edges $\mathcal{E}^{(m)}$ encode relationships between ROIs for subject m . These edges may be derived in two primary ways:

Undirected Functional Connectivity: Often obtained through measures like Pearson’s correlation, capturing how strongly two ROIs co-activate over the entire time series. While useful for identifying global co-activation networks, this approach ignores the temporal ordering of signals.

Directed Effective Connectivity: Derived via statistical methods such as Granger causality, providing edges that reflect the temporal direction of influences between ROIs. This approach attempts to unveil how activity in one ROI at a previous time point may predict future activity in another ROI.

By modeling brain networks in a graph form, we can apply *graph neural networks* to capture not only the local interactions among connected nodes but also the global

topology of the brain network.

Study Focus

This study aims to perform downstream prediction using resting-state fMRI (rs-fMRI) data by modeling the brain as a network of interacting regions. To enhance predictive accuracy and interpretability, we incorporate directionality and causal inference into graph-based connectivity representations. For a given subject m , the resulting graph-based representation—whether undirected or directed—serves as an input to a predictive model:

$$\hat{y}^{(m)} = \mathcal{F}(\mathbf{X}^{(m)}, \mathbf{A}^{(m)}), \quad (3.3)$$

where $\mathbf{A}^{(m)}$ denotes the adjacency matrix (undirected or directed) encoding the subject’s brain connectivity, and \mathcal{F} is typically a learnable function (e.g., a GNN). The desired output $\hat{y}^{(m)}$ may correspond to a continuous cognitive score or a categorical behavioral classification, depending on the prediction task.

Overall, the central proposition of this work is to construct brain network representations that *retain the directional causal signals* inherent in time-series data. By embedding these properties in a graph model, we aim to advance both the understanding of brain function and the performance of machine learning tasks in large-scale neuroimaging studies.

Chapter 4

Method

In this section, we first introduce two existing fMRI-based brain analysis paradigms: MLP-Based approaches and graph-based models utilizing undirected functional connectivity. We then present the proposed method based on Granger causality, which can capture directional interactions between brain regions, providing insights into causal relationships and temporal dependencies that are not detectable in undirected correlation-based networks.

4.1 MLP-Based Methods

A baseline approach to fMRI-based prediction is to apply a Multilayer Perceptron (MLP) directly to the BOLD signals, treating the data as a high-dimensional input. Instead of explicitly modeling connectivity between brain regions, this method learns feature representations through fully connected layers, allowing it to capture straightforward but potentially effective patterns in the data.

For a given subject m , the fMRI data $\mathbf{X}^{(m)} \in \mathbb{R}^{N \times T}$ is first flattened into a feature vector $\mathbf{z}^{(m)} \in \mathbb{R}^{N \cdot T}$, which serves as the input to an MLP:

$$\hat{y}^{(m)} = f_{\theta}(\mathbf{z}^{(m)}),$$

where f_θ represents the MLP model with learnable parameters θ . The MLP consists of multiple fully connected (dense) layers, each followed by a non-linear activation function (e.g., ReLU). While this approach can learn strong predictive mappings in some cases, it does not exploit the inherent spatial or temporal structure of the brain, and as a result, its interpretability may be limited.

For additional details on the MLP formulation and hyperparameters, we refer the reader to Appendix B.

4.2 Graph-Based Models

For each subject m , the node feature matrix $\mathbf{X}^{(m)} \in \mathbb{R}^{N \times T}$ is derived from time-series BOLD signals, where N represents the number of ROIs and T represents the number of time steps. Each row of $\mathbf{X}^{(m)}$ encapsulates the temporal activity of an ROI. Adjacency matrices encode ROI relationships in two forms. Undirected graphs represent functional connectivity based on Pearson correlation coefficients, producing symmetric adjacency matrices $\mathbf{C}^{(m)} \in \mathbb{R}^{N \times N}$. Directed graphs capture effective connectivity using Granger causality, resulting in directed adjacency matrices $\mathbf{A}^{(m)} \in \mathbb{R}^{N \times N}$.

4.2.1 Undirected Graphs: Functional Connectivity.

Functional connectivity between ROIs is computed using Pearson correlation coefficients. For each subject m , the adjacency matrix $\mathbf{C}^{(m)}$ is defined as:

$$C_{ij}^{(m)} = \frac{\sum_t (x_{i,t}^{(m)} - \mu_i^{(m)}) (x_{j,t}^{(m)} - \mu_j^{(m)})}{\sqrt{\sum_t (x_{i,t}^{(m)} - \mu_i^{(m)})^2} \sqrt{\sum_t (x_{j,t}^{(m)} - \mu_j^{(m)})^2}},$$

where $\mu_i^{(m)}$ and $\mu_j^{(m)}$ are the mean values of ROI i and ROI j time-series, respectively. The resulting graph $\mathcal{G}^{(m)}$ models symmetric relationships between ROIs, with edge

weights reflecting the strength of functional connectivity.

4.2.2 Directed Graphs: Effective Connectivity.

Effective connectivity is modeled using Granger causality [26, 14], allowing us to infer putative directional influences among ROIs. Below, we outline the main steps to construct the directed adjacency matrix $\mathbf{A}^{(m)}$.

Sliding Windows: We segment each ROI’s time-series into overlapping windows of size W , shifting by S . For each window w , the mean signal is computed:

$$\bar{x}_{i,w}^{(m)} = \frac{1}{W} \sum_{t \in w} x_{i,t}^{(m)},$$

yielding a mean-based time-series $\bar{\mathbf{X}}^{(m)} \in \mathbb{R}^{N \times W_{\text{windows}}}$. This step helps reduce high-frequency noise and captures short-term dynamics.

VAR Modeling: Within each window w , for every pair (i, j) , we fit a Vector Autoregressive (VAR) model of order L :

$$x_j(t) = \alpha_{j,0} + \sum_{k=1}^L \alpha_{j,k} x_j(t-k) + \sum_{k=1}^L \beta_{j,k} x_i(t-k) + \epsilon_j(t).$$

The intuition is to assess whether past values of ROI i improve the prediction of ROI j , beyond what j ’s own past can explain [58, 59].

Hypothesis Testing: We conduct a statistical test (commonly an F -test) to examine whether the lag coefficients $\beta_{j,k}$ associated with ROI i are jointly non-zero:

$$H_0 : \beta_{j,1} = \beta_{j,2} = \dots = \beta_{j,L} = 0 \quad \text{vs.} \quad H_1 : \exists k \text{ such that } \beta_{j,k} \neq 0.$$

If H_0 is rejected at a specified significance level ($\alpha = 0.05$), we conclude that i “Granger-causes” j in this window [26, 58].

Aggregation Across Windows: We repeat the above tests for each overlapping window. If ROI i is found to Granger-cause ROI j in a sufficient fraction of those windows (e.g., exceeding a threshold τ), we set

$$A_{ij}^{(m)} = 1 \quad (\text{directed edge from } i \text{ to } j),$$

otherwise $A_{ij}^{(m)} = 0$.

Through this procedure, the final adjacency matrix $\mathbf{A}^{(m)} \in \{0, 1\}^{N \times N}$ encodes the direction and flow of information among ROIs. In contrast to undirected correlation-based graphs, these directed edges highlight potential causal relationships and temporal ordering of brain activity. For additional technical details (including pseudocode for VAR fitting and statistical tests), see Appendix A.

We evaluate the constructed graphs using a Graph Convolutional Networks (GCNs) [12, 37]. Specifically, GCNs are applied separately to undirected $\mathbf{C}^{(m)}$ and directed $\mathbf{A}^{(m)}$. In undirected GCNs, $\mathbf{C}^{(m)}$ serves as the adjacency matrix for symmetric feature propagation, while directed GCNs use $\mathbf{A}^{(m)}$ to incorporate directional dependencies. The model generates embeddings, which are evaluated on downstream classification and regression tasks.

Overall, combining causal directionality with graph neural networks offers richer insights into how signals move across brain circuits. We refer interested readers to Appendix A for algorithms on the Granger causality steps and to Appendix B for additional details on the MLP and GCN formulations used in this work.

4.3 Rationale for Selecting Granger Causality

In our study, we evaluated multiple methods for constructing connectivity graphs from fMRI data. Traditional measures such as Pearson correlation and time cross-correlation are well-known for their computational efficiency and scalability; however, these methods inherently yield undirected graphs. As a result, they capture only co-activation or temporal synchrony and cannot resolve the causal direction of interactions among brain regions [20, 47].

In contrast, methods capable of inferring directionality—such as Transfer Entropy, Dynamic Causal Modeling (DCM), and Granger causality—offer deeper insights into effective connectivity by estimating causal relationships. Transfer Entropy, for instance, provides a non-linear measure of information transfer that can theoretically capture the direction of interactions. Yet, its high computational demand and sensitivity to noise limit its feasibility for large-scale fMRI datasets [68].

Dynamic Causal Modeling (DCM) adopts a biophysical approach, fitting a generative model to the data to estimate directional and signed influences between brain regions [21, 22]. Although DCM is robust and provides detailed insights into neural dynamics, it requires a priori specification of the network structure and becomes computationally prohibitive as network size increases. For example, on the ABCD dataset—which contains 360 nodes and 512 time series—DCM requires approximately 120 seconds per sample, and with over 7,000 samples, the processing time becomes impractical [47, 53].

Granger causality, on the other hand, infers directional influences by testing whether past activity in one region improves the prediction of future activity in another [25]. Despite its underlying assumptions of linearity and stationarity, our pre-processed fMRI data satisfy these conditions well enough to yield robust estimates. Importantly, Granger causality achieves a significant reduction in computation time—averaging about 11 seconds per sample on the ABCD dataset. Moreover, by employing a

chunking strategy, we further improved scalability without degrading performance, reducing the processing time to approximately 3 seconds per sample on average. Additionally, the use of parallelization further enhances the efficiency of iterative computations. Experiments on the smaller PNC dataset demonstrate that chunking not only maintains consistent performance but also effectively smooths noise, which does not inherently contribute to improved training.

In summary, the efficiency and scalability of Granger causality make it ideally suited for large-scale fMRI studies. Its ability to efficiently infer directional influences, combined with our practical improvements through chunking, offers a compelling balance between computational feasibility and the extraction of meaningful, causal interactions from complex neuroimaging data.

4.3.1 Overview of Connectivity Inference Methods

Table 4.1: Comparison of connectivity inference methods for fMRI analysis.

Method	Directed	Complexity	Scalability	Time (in ms)
Pearson Corr.	No	Low	High	152
Time Cross-Corr.	No	Low-Mod	High	456
Transfer Entropy	Yes	High	Limited	88,000
DCM	Yes	High/Mod	Limited	120,000
Granger Causality	Yes	Moderate	Moderate	11,000

Note: All timing measurements are based on the ABCD dataset, comprising 360 nodes and 512 time steps per sample. The time (in milliseconds) reflects the average duration required to process a single sample for graph construction.

Based on our experiments and practical considerations, we adopted Granger causality for constructing directed graphs. While DCM provides detailed biophysical insights, its computational intensity and requirement for pre-specified network models limit its applicability to large networks. Similarly, Transfer Entropy, though theoretically robust, is less practical for large-scale datasets due to its high complexity. In contrast, Granger causality offers a favorable compromise by efficiently inferring directional interactions while scaling well to large fMRI studies.

Chapter 5

Experiments

Our experimental design is structured to address three central research questions that aim to evaluate both predictive performance and the neurobiological interpretability of our approaches:

RQ1 *Does model performance vary across different graph construction methods?* We compare models based on flattened time-series data, undirected correlation graphs, and directed Granger-causal graphs to determine whether directional information offers advantages in terms of predictive accuracy and robustness.

RQ2 *What parameters best optimize Granger-causal graphs?* We systematically tune hyperparameters such as window size, step size, and lag order using model selection criteria (e.g., AIC and BIC) to identify configurations that effectively balance model complexity and temporal resolution in directed connectivity.

RQ3 *How do directed connectivity insights complement undirected graphs?* We focus on how causal (directed) edges can provide deeper neurobiological insights, particularly by highlighting potential top-down vs. bottom-up pathways that remain obscure in undirected correlation-based methods.

5.1 Experimental Settings

Datasets. We use two large-scale neuroimaging datasets: the Adolescent Brain Cognitive Development Study (ABCD) and the Philadelphia Neurodevelopmental Cohort (PNC) [9, 51].

Table 5.1: Summary of Datasets and Tasks

Dataset	Task	Atlas	# Subjects	# Time Steps	# Nodes	Response	# Classes
ABCD [9]	Classification	HCP 360	7,901	512	360	Gender	2
ABCD [9]	Regression	HCP 360	4,613	1,024	360	Cognitive Score	–
PNC [51]	Classification	Power 264	503	120	264	Gender	2

ABCD Dataset. The ABCD dataset tracks 9-10-year-olds through early adulthood with repeated MRI scans [9]. It includes 7,901 subjects, parcellated using the HCP 360 atlas [16], with a balanced gender distribution. For gender prediction tasks, 512 time steps are used, while regression tasks predicting the Cognition Summary Score utilize 1,024 time steps. Samples with fewer than 1,024 time steps were excluded, resulting in 4,613 samples for regression.

PNC Dataset. The PNC dataset comprises 503 subjects from the University of Pennsylvania and Children’s Hospital of Philadelphia, also with a balanced gender distribution [51]. Each subject provides 120 time steps of rs-fMRI data from 264 ROIs, with preprocessing steps including motion correction, normalization, and bandpass filtering [44].

Metrics. For binary gender classification, we use the Area Under the Receiver Operating Characteristic (AUC), Accuracy, and F1 Score, reflecting the model’s ability to distinguish between classes, overall correctness, and balance of precision and recall. The classification threshold is set at 0.5. For cognitive score regression, Mean Squared Error (MSE) is employed to measure the average squared difference between predicted and actual scores.

Implementation Details. All models were trained using the Adam optimizer (learning rate = 1×10^{-4} , weight decay = 1×10^{-4}) [36]. We used binary cross-

entropy loss for classification tasks and mean squared error (MSE) for regression tasks, training each model for up to 100 epochs with early stopping (patience = 10 epochs). Batch size was set to 16, with validation loss checked every 5 epochs, and mixup regularization (mixup=1) was applied to enhance generalization. Prior to training, each time series was z-score normalized on a per-ROI basis. To ensure robustness, we repeated experiments 5 times with different random seeds and report averaged results.

We evaluated two main architectures. The **MLP** baseline flattens the time-series data into a single feature vector and processes it through two fully connected layers (512 and 256 units, both with ReLU activation and dropout = 0.5). A final output layer uses sigmoid activation for classification tasks or a linear activation for regression tasks. The **GCNs**, on the other hand, incorporate node features and either undirected or directed adjacency matrices. Specifically, two graph convolution layers (256 units, ReLU) are followed by a fully connected layer (256 units, dropout = 0.5) and an output layer (sigmoid or linear, depending on the task). Except for the graph-specific operations, training procedures (e.g., optimizer settings, early stopping, batch size) remain consistent across both architectures.

Baselines. We compare our proposed model with a range of baselines: (i) *MLP Approach*—a time-series MLP that directly encodes BOLD data without network modeling; (ii) *Direction-Free Deep Learning Approaches*—including BrainNetCNN, FCNet, and BrainGB with functional connectivity (BrainGB w/FC), which exploit correlation-based connectivity features via BrainGB’s GCN architecture to model ROI relationships without directionality [12, 34, 48]; (iii) *Effective Connectivity Method*—which estimates directed interactions using the Noise-Diffusion Network (NDNetwork) followed by a GCN module [12, 24, 37].

Table 5.2: Performance Comparison Across Models and Tasks

Type	Method	Dataset: PNC			Dataset: ABCD			
		AUC ↑	Acc ↑	F1 ↑	AUC ↑	Acc ↑	F1 ↑	MSE ↓
Time-series	MLP	53.80	53.28	52.48	51.29	50.38	49.17	106.43
Direction-Free	BrainNetCNN [34]	52.73	54.16	53.33	54.36	53.29	54.10	85.81
	FCNet [48]	52.95	51.64	50.82	50.36	51.29	50.10	103.65
Effective Connectivity	BrainGB w/FC [12, 37]	<u>56.39</u>	56.74	55.32	<u>63.32</u>	59.74	<u>59.33</u>	79.93
	NDNetwork [24, 37]	54.46	53.58	50.95	TLE	TLE	TLE	TLE
Directed (Ours)	BrainGB w/GC [12, 14, 37]	56.74	<u>56.35</u>	<u>54.93</u>	63.76	<u>59.65</u>	59.56	<u>81.32</u>

TLE (Time Limit Exceeded): Indicates that the training exceeds 72 hours.

5.2 Model Performance (RQ1)

5.2.1 Graph-Based Models Outperform MLP.

Table 5.2 shows that graph-based models consistently outperform the MLP baseline. By incorporating functional connectivity (whether correlation-based or Granger-based), these models capture the spatial and temporal organization of neural systems more effectively than the MLP, which treats each voxel or ROI time series as an independent feature. In particular, *BrainGB w/FC* and *BrainGB w/GC* achieve higher AUC and F1 scores for classification tasks on the PNC dataset, and lower MSE for regression on the ABCD dataset. The MLP’s lower performance highlights the importance of leveraging inter-regional relationships in fMRI analysis.

5.2.2 Compatibility of Directed and Undirected Graphs.

Our findings further indicate that introducing directionality does not compromise predictive performance. Specifically, *BrainGB w/GC* achieves results comparable to, and occasionally better than, *BrainGB w/FC*, suggesting that Granger-causal edges can be as effective as correlation-based edges when used in GCNs. Moreover, the directed approach delivers additional insights into *causal* interactions between brain regions that cannot be captured by undirected methods. Notably, the dispersed connectivity in directed graphs enhances neurological interpretability, offering plausible pathways

of information flow throughout the connectome. In contrast, the Noise-Diffusion Network (NDNetwork) approach times out (TLE) when scaling to the large ABCD dataset, underscoring potential computational inefficiencies relative to correlation- or Granger-based methods. This highlights the practical advantage of constructing directed Granger-causal graphs, which remain both computationally tractable and neurobiologically meaningful. Overall, these results demonstrate that directionality can be introduced without sacrificing performance, while providing more nuanced understandings of functional relationships in the brain.

5.3 Hyperparameter Study for Directed Graphs (RQ2)

The construction of directed Granger-causal graphs involves optimizing three key hyperparameters—**window size** (W), **step size** (S), and **lag order** (L)—within a Vector Autoregressive (VAR) framework [58, 59, 75]. These parameters crucially determine how effectively the model captures temporal dependencies and causal influences in fMRI time-series data.

The **window size** (W) controls how many consecutive time steps are grouped together, with larger windows capturing more extended dependencies but risking the smoothing-over of finer-scale dynamics. Conversely, smaller windows provide sharper temporal resolution but may fail to capture longer-lag effects. The **step size** (S) specifies how far the sliding window shifts at each step; small shifts allow for denser sampling of the time series, while larger shifts reduce computation at the potential cost of missing fast-evolving phenomena. Finally, the **lag order** (L) defines how many past time points from each ROI are used to predict future activity in another ROI, balancing the trade-off between model complexity and its ability to capture relevant temporal delays.

To select the optimal L , we employ the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) [1, 42]. AIC is computed as $AIC = 2k - 2\ln(L)$, where k denotes the number of parameters in the VAR model, and $\ln(L)$ is the log-likelihood [1]. BIC uses a stricter penalty for model complexity: $BIC = k\ln(n) - 2\ln(L)$, where n is the number of observations [42]. The lag order that jointly minimizes AIC and BIC is typically regarded as optimal.

Using the ABCD dataset for a gender prediction task, we conducted a grid search over combinations of W , S , and L (see Table 5.3) to empirically assess how each hyperparameter setting influences both model fit (AIC/BIC) and predictive performance (AUC, Accuracy, F1).

Table 5.3: Grid Search Results for Granger Causality Hyperparameters

Window Size (W)	Step Size (S)	Lag Order (L)	AIC [1]	BIC [42]	AUC	Accuracy	F1
32	16	1	79.59	<u>83.90</u>	63.32	59.44	59.33
32	16	2	78.10	85.11	62.82	58.93	58.75
32	16	3	<u>77.34</u>	86.92	63.03	59.19	59.08
32	32	1	78.40	82.70	<u>63.62</u>	<u>59.61</u>	59.71
32	32	2	77.72	84.73	63.12	59.21	59.01
32	32	3	76.96	86.53	63.76	59.65	<u>59.56</u>
64	32	1	167.37	173.79	62.53	58.02	58.07
64	32	2	166.26	176.90	63.04	58.58	58.56
64	32	3	165.43	180.21	63.58	59.05	59.07
64	64	1	161.80	168.23	62.89	58.20	58.24
64	64	2	160.97	171.61	63.20	58.64	58.62
64	64	3	160.44	175.21	63.55	59.00	59.03
128	64	1	352.23	360.76	61.82	57.94	57.83
128	64	2	350.05	364.24	62.19	58.21	58.19
128	64	3	348.88	368.68	62.40	58.54	58.32
128	128	1	339.90	348.43	62.71	58.86	58.60
128	128	2	337.63	351.81	63.05	59.11	58.99
128	128	3	336.39	356.18	63.36	59.33	59.22

Findings

Despite the assumption that lower AIC and BIC would yield better results, performance remained robust under various hyperparameter settings. Table 5.3 highlights a nuanced picture of Granger-causal modeling under various window sizes, step sizes, and lag orders. While theory often suggests that minimizing AIC and BIC should yield the best predictive model, our results show that even configurations not strictly optimal

in these criteria can still produce robust outcomes. Notably, $(W = 32, S = 32, L = 3)$ achieves the lowest AIC (76.96) along with the highest AUC (63.76%) and Accuracy (59.65%), whereas $(W = 32, S = 32, L = 1)$ obtains the lowest BIC (82.70) and the best F1 score (59.71%). These differences illustrate that balancing AIC against BIC may lead to slightly different optimal settings, yet both configurations deliver comparable performance.

Beyond individual metrics, the grid search underscores that small modifications to W , S , or L rarely compromise predictive accuracy in a significant way. This suggests that fMRI signals contain sufficient temporal structure to be modeled effectively across a range of window and lag parameters. From a practical standpoint, researchers can afford to choose hyperparameters that best fit their computational budget or domain-specific needs (e.g., emphasizing fine-grained temporal resolution versus reducing runtime) without sacrificing much predictive power. In particular, smaller window sizes and step sizes can capture more granular dynamics, albeit at higher computational cost, while larger windows and fewer lags can simplify training without severely degrading performance. Overall, these results confirm that Granger-based causal modeling is robust to moderate variations in hyperparameter choices, reinforcing its applicability to diverse fMRI datasets and experimental protocols.

5.3.1 Effects of Window Size, Step, and Lag on Directed Connectivity Patterns

Figure 5.1 illustrates a series of directed connectivity matrices generated by varying the hyperparameters in our Granger-causal estimation. Although changes in window size, step size, and lag order modulate the edge density and fine structure of the matrices, the overall pattern of inter-regional relationships remains consistent. This indicates that the Granger-causal approach is robust to moderate variations in these settings. Smaller window sizes and step sizes tend to capture more fine-grained temporal

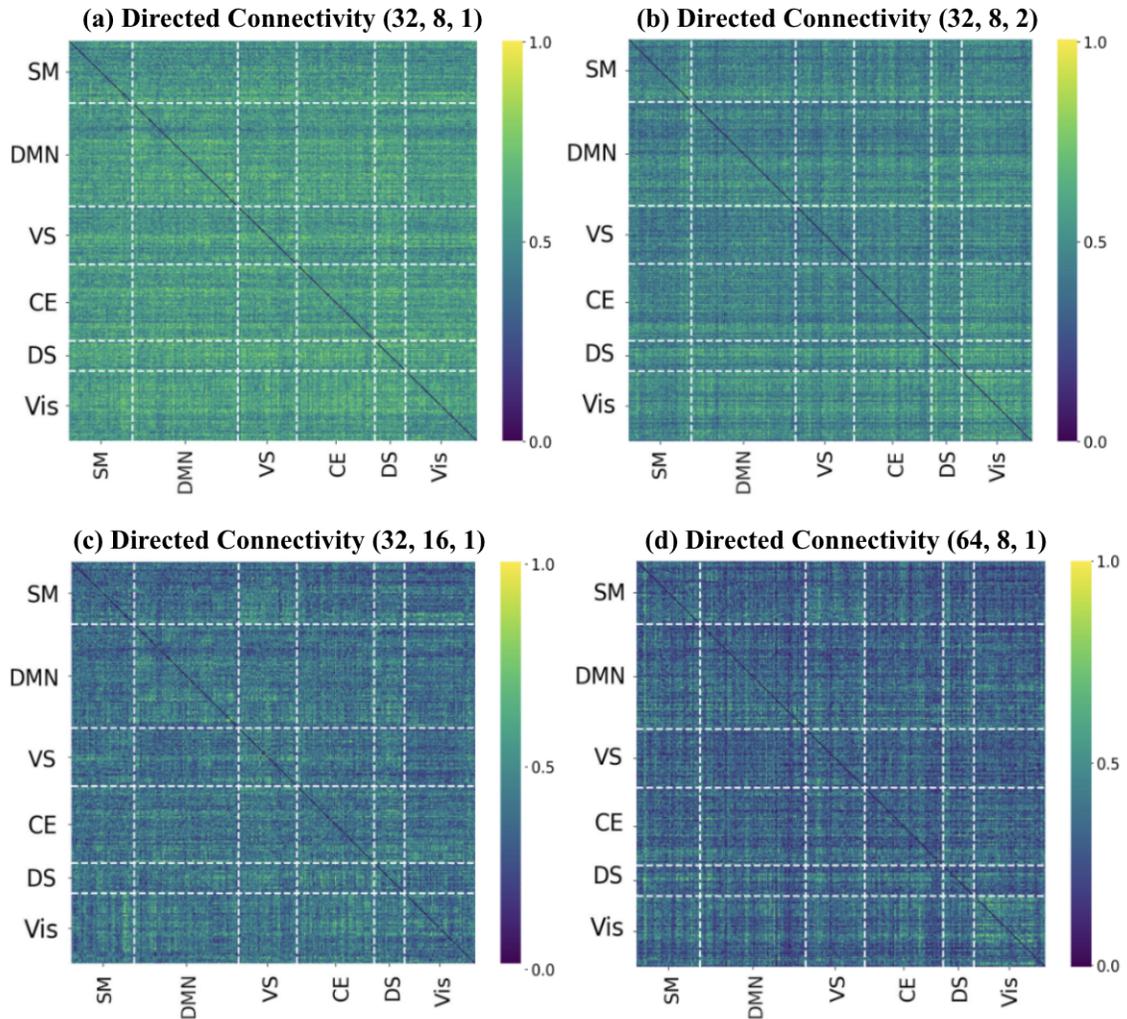


Figure 5.1: *Directed Connectivity Under Different Hyperparameters*. Each matrix is a Granger-causal adjacency (rescaled to $[0,1]$) using different window sizes (W), step sizes (S), and lag orders (L). Overall connectivity structure remains similar, despite parameter changes.

dynamics, resulting in matrices with higher resolution but increased computational cost. Larger windows, conversely, average over longer periods, potentially smoothing out transient effects while preserving the general connectivity pattern. Similarly, the lag order (L) is crucial: higher lag orders can model extended temporal dependencies but may also introduce noise and risk overfitting, while lower lag orders might overlook subtle delayed interactions. Our grid search (see Table 5.3) demonstrates that both lower and higher lag orders can yield competitive performance metrics. Neurologically, these parameter adjustments provide additional insights. For example, configurations

with lower lag orders might be more sensitive to immediate interactions among ROIs, whereas higher lag orders could capture more complex, hierarchical signal propagation—such as how sensorimotor activity may initiate cascades that influence higher-order regions like the DMN. Such directional insights, not discernible from correlation-based methods, highlight the potential of Granger-causal graphs to capture the brain’s hierarchical organization and feedback loops.

Overall, while tuning hyperparameters can optimize model fit (as reflected by lower AIC and BIC values), the stability of the inferred connectivity structure across a range of settings reinforces the utility of directed Granger-causal graphs. This robustness underscores their applicability to diverse fMRI datasets and experimental protocols, providing a reliable means of capturing meaningful neurobiological dynamics.

Sliding Window Length. The choice of window length strongly influences the detected connectivity dynamics. Shorter windows (e.g., 16 - 32 time steps) allow rapid fluctuations to be captured but risk higher variability and spurious connections due to limited data in each window [38, 45, 66, 76]. In contrast, longer windows (e.g., 64 - 128 time steps) produce smoother, more stable connectivity estimates by averaging out brief transients, albeit at the cost of temporal precision [38, 45, 66, 76]. In our analysis, shorter windows tended to produce denser effective connectivity matrices, revealing more significant directed edges, whereas longer windows yielded sparser matrices, capturing only the strongest, most persistent connections. Increasing window length also helps reduce spurious fluctuations in connectivity, strengthening confidence that the remaining directed links are robust [23, 38]. However, excessively long windows may assume stationarity and overlook meaningful dynamic shifts.

Step Size (Window Offset). The step size (i.e., overlap between successive windows) determines how finely one tracks changes in connectivity over time. Smaller steps (high overlap) yield a more continuous time course of connectivity, at the expense

of greater computational load and potential serial correlations. Larger steps (low overlap or non-overlapping windows) reduce redundancy but may skip over rapid transitions. Empirical evaluations indicate that step size, second only to window length, affects the detection of connectivity-state durations [?]. In practice, a moderate overlap (e.g., 50–80%) often balances smoothness with computational efficiency. In our runs, high overlap produced more gradual changes in directed edge strengths, whereas a large step sometimes caused abrupt transitions between windows. Checking for consistency across overlapping windows can also help validate truly robust directed influences (e.g., a persistent DMN→SM link should appear in consecutive windows rather than isolated ones).

Lag Order of the VAR Model. Another critical hyperparameter is the autoregressive lag order. Too few lags can bias Granger causality tests by failing to capture the actual delay structure of neural interactions [55, 63], potentially missing directed links with longer neural or hemodynamic latencies. Conversely, using too many lags consumes degrees of freedom and can introduce noise, diminishing statistical power [55, 63]. In our resting-state analyses, the appropriate lag order depends on the temporal resolution determined by the repetition time (TR) of the dataset. For the ABCD dataset, which has a TR of 0.8 s (yielding a sampling rate of approximately 1.25 Hz), a lag order of 1 (i.e. one TR, or 0.8 s) sometimes yielded sparse connectivity because one TR may be shorter than the typical hemodynamic delay [9, 16, 23]. Increasing the lag (e.g., to 3–5 TRs, corresponding to approximately 2.4–4.0 s) allowed detection of slower, multi-TR influences; however, very high lags (e.g., 10 TRs, or about 8 s) risked overfitting, resulting in nearly fully connected networks. In contrast, for the PNC dataset, where the TR is 3.0 s (sampling rate of approximately 0.33 Hz), one TR represents a 3-second interval [7, 44, 51]. Given this slower temporal resolution, the choice of lag order is even more critical to capture delayed neural interactions without

overestimating connectivity. This bias–variance tradeoff is especially salient considering that rapid neural interactions occurring on a sub-TR scale may not be detectable, particularly in datasets with slower sampling rates [63]. Therefore, selecting a lag order that aligns with the temporal scale of interest (and considering methods such as pre-whitening or deconvolution to better account for the hemodynamic response) is crucial for reliable inference in both the ABCD and PNC datasets.

Edge Density and Sparsity. Finally, these hyperparameter factors collectively affect the overall sparsity of the directed connectivity matrix. More conservative configurations (long windows, high overlap, low lag) often yield sparser networks, retaining only the strongest directed edges, whereas more aggressive settings (short windows, minimal overlap, higher lag) produce denser networks. The resulting sparsity or density can greatly influence interpretations of brain organization: a sparse directed network might suggest a hierarchical system with a few causal hubs, while a denser network implies more distributed, bidirectional information flow. Certain edges (e.g., within tightly coupled subsystems like Vis↔SM) appeared consistently across many configurations, suggesting they represent stable causal pathways. Other edges emerged only under specific parameter choices, indicating potential context- or parameter-dependent interactions. By identifying robust edges that persist across multiple configurations, researchers can be more confident in the neurobiological relevance of those directed influences, while edges sensitive to small parameter changes warrant careful scrutiny.

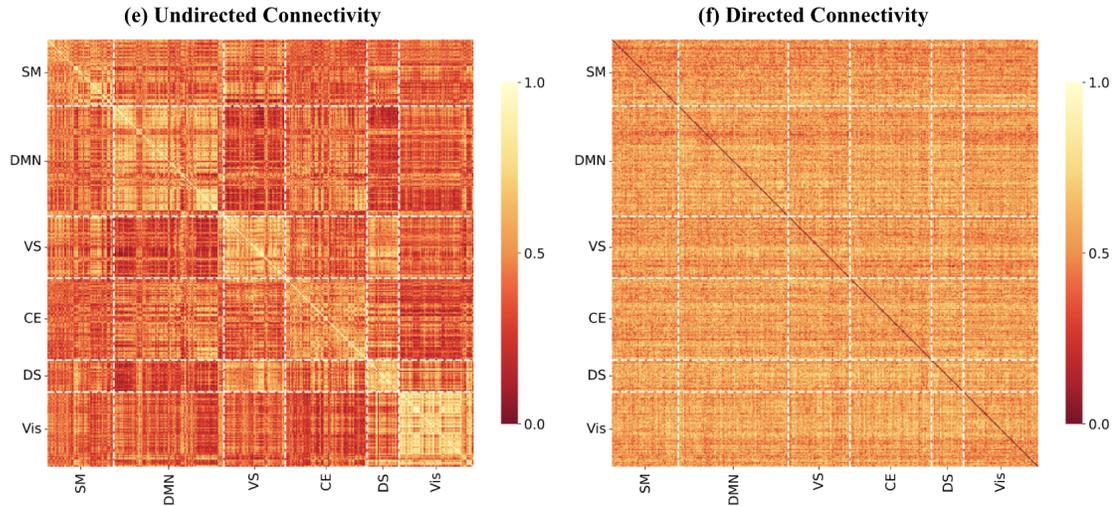


Figure 5.2: *Both matrices are derived from the ABCD dataset (HCP 360 atlas). The undirected matrix (e) is correlation-based; the directed matrix (f) stems from Granger causality. Each is normalized to [0,1]. White grid lines outline functional networks (SM, DMN, VS, CE, DS, Vis).*

5.4 Neurological Insights (RQ3)

5.4.1 Comparison of Undirected and Directed Connectivity Matrices

Figure 5.2 presents the correlation-based (e) and Granger-causal (f) connectivity matrices derived from the ABCD dataset using the HCP 360 atlas, each normalized to the range $[0,1]$. The undirected matrix (left) exhibits a prominent block structure, with large-scale networks (e.g., Sensorimotor (SM), Default Mode Network (DMN), and Visual (Vis)) forming visible diagonal blocks of high correlation. Such block-like patterns suggest that regions within each network tend to co-activate strongly and relatively synchronously.

By contrast, the directed matrix (right) has more diffuse or asymmetrical connectivity, reflecting unidirectional influences among regions. The diagonal blocks appear less pronounced, indicating that while intra-network connections remain important, inter-network edges become more relevant when directionality is taken into account.

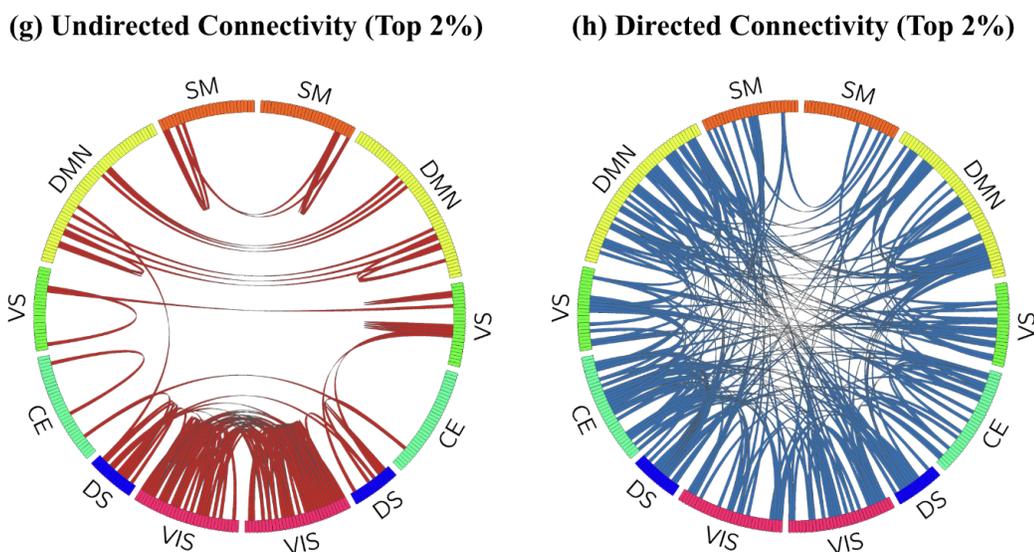


Figure 5.3: *Chord Plots Showing the Top 2% of Edges.* Each segment represents a network; arcs link strongly connected or causally influencing ROIs.

In particular, sensorimotor areas often emerge as ‘sources’ that drive activity in higher-order networks (e.g., DMN), rather than merely co-activating with them. The ability of Granger causality to capture temporal precedence and predictive relationships underscores potential causal flow, hierarchical organization, and feedback loops that can be obscured in correlation-based approaches.

Chord Plots of the Top 2% of Edges Figure 5.3 shows chord diagrams of the top 2% of connections for both undirected (left) and directed (right) connectivity. Each colored segment corresponds to a major functional network (e.g., SM, DMN, Vis, CE, DS), while the arcs represent the strongest pairwise links. In the undirected chord plot, arcs cluster heavily within the same functional modules (e.g., DMN-to-DMN), mirroring the block-like structure seen in the correlation matrix. This pattern suggests high intra-network synchrony.

In contrast, the directed chord diagram reveals a greater number of inter-network arcs, highlighting the potential for causal influences between distant regions. For instance, SM regions may project to DMN or other association networks, consistent

with a feedforward hierarchy in which sensorimotor signals propagate to higher-order systems. Moreover, some arcs appear unidirectional (e.g., SM \rightarrow DMN) rather than reciprocal, suggesting that while sensorimotor activity may drive default-mode regions, the reverse influence is comparatively weaker. Because directed connectivity does not double-count symmetric links, the overall set of edges can be larger, yet sparser in a node-to-node sense—reflecting the fact that many pairs exhibit only one dominant direction of influence.

Implications and Observations: The directed (Granger-causal) vs. undirected connectivity patterns highlight key neurodevelopmental dynamics in the ABCD cohort. Notably, strong directed influences from sensorimotor regions to the default mode network (DMN) suggest that in children and adolescents, lower-order sensorimotor signals heavily drive activity in higher-order self-referential networks. This bottom-up bias aligns with neurodevelopmental models positing that cortical hierarchies are initially dominated by feed-forward influences, with top-down feedback pathways maturing later in adolescence [41, 43]. Indeed, recent work shows that top-down propagations (e.g., frontoparietal-to-sensory cortical activity flows) become more prevalent with age during youth, indicating a strengthening of bidirectional (feedback) circuits over development [43]. Thus, the observed unidirectional sensorimotor \rightarrow DMN connectivity in younger participants likely marks an earlier developmental stage when external bodily inputs can intrude on or drive internal mentation. Consistent with this, information flow from visual regions to associative networks (including the DMN) tends to decrease as cognitive abilities improve [40], suggesting that maturation involves greater top-down regulation of sensory-driven activity. The pronounced directed links from visual networks to attention networks in the ABCD data similarly underscore children’s reliance on bottom-up visual inputs to engage attention orienting systems. As the dorsal attention network matures, more balanced reciprocal connectivity (i.e.,

increased feedback from attention-control regions to sensory cortex) supports selective attention and goal-directed behavior.

In parallel, the emergence of executive control loops—reciprocal connections among frontoparietal and other control regions—reflects the refinement of cognitive control circuitry. For example, the segregation (even anti-correlation) of the DMN from lateral frontoparietal networks becomes more pronounced from childhood to adulthood, and in 9–10 year-olds, weaker DMN–executive coupling (i.e., more decoupling) correlates with better cognitive performance [17]. This developmental shift from diffuse, positively coupled networks toward more specialized, anticorrelated networks suggests that bidirectional executive-DMN interactions are tuned to enable focus on external tasks by suppressing self-referential activity. In short, unidirectional influences in childhood may be a hallmark of immature hierarchical organization, whereas the emergence of robust bidirectional influences by adolescence marks a brain that can flexibly alternate between bottom-up and top-down modes of information flow [41, 43]. Such changes are biologically grounded in processes like synaptic pruning and myelination that sharpen network specificity, and they align with theories of gradual integration and segregation of brain networks during development.

These network-level developmental differences have important cognitive and clinical implications. Improved top-down connectivity within executive circuits and from salience/attention networks to sensory regions is crucial for developing self-regulation skills such as sustained attention, working memory, and inhibitory control. Conversely, if connectivity remains overly bottom-up or unbalanced, a child may struggle with regulating internal versus external focus. For instance, heightened sensorimotor–DMN coupling could reflect difficulties suppressing motor-sensory impulses during internally focused thought, potentially contributing to attention lapses or impulsivity. Indeed, children with more neurodevelopmental symptoms (e.g., inattention, hyperactivity) show atypical connectivity patterns involving the DMN—specifically, a trend toward

lower within-DMN coherence and stronger anti-correlation between the DMN and other networks—indicative of immature or dysregulated integration of internal and external attention systems [33]. Similarly, a recent large-scale study found that youth with better cognitive performance exhibited weaker coupling between DMN and executive networks at rest (and stronger within-network executive connectivity), supporting the idea that appropriate segregation of these networks facilitates executive function [17].

On the mental health front, the directionality of influences between networks provides insight into pathophysiology: for example, adolescent depression has been associated with a reduction in top-down inhibitory influence from the salience network to the DMN, resulting in an overactive DMN and excessive self-focused rumination when regulatory control fails [70]. Such a hierarchical imbalance—a predominance of intrinsic DMN activity unchecked by salience/executive circuits—aligns with the directed connectivity findings and could help explain internalizing symptoms. Likewise, aberrant executive-loop development and DMN connectivity have been implicated in ADHD and related conditions, where insufficient top-down control and delayed network segregation may lead to poor impulse control and distractibility [33].

In summary, the observed directed vs. undirected connectivity differences (e.g., sensorimotor \rightarrow DMN, visual \rightarrow attention, evolving executive network reciprocity) are developmentally meaningful: they illustrate a brain transitioning from predominantly unidirectional, bottom-up information flow in childhood to more bidirectional, integrative network communication by adolescence. This trajectory supports the maturation of attention, memory, and self-regulatory capacities, while deviations from it (e.g., persistent unidirectional or uncoupled networks) can foreshadow cognitive and mental health challenges. These interpretations are supported by recent large-scale developmental fMRI studies and underscore how directed connectivity analyses in youth can illuminate the biological underpinnings of cognitive development and psychopathology [17, 33, 40, 70].

5.4.2 Consistent Directed Connectivity Patterns in the ABCD Cohort

Cohort

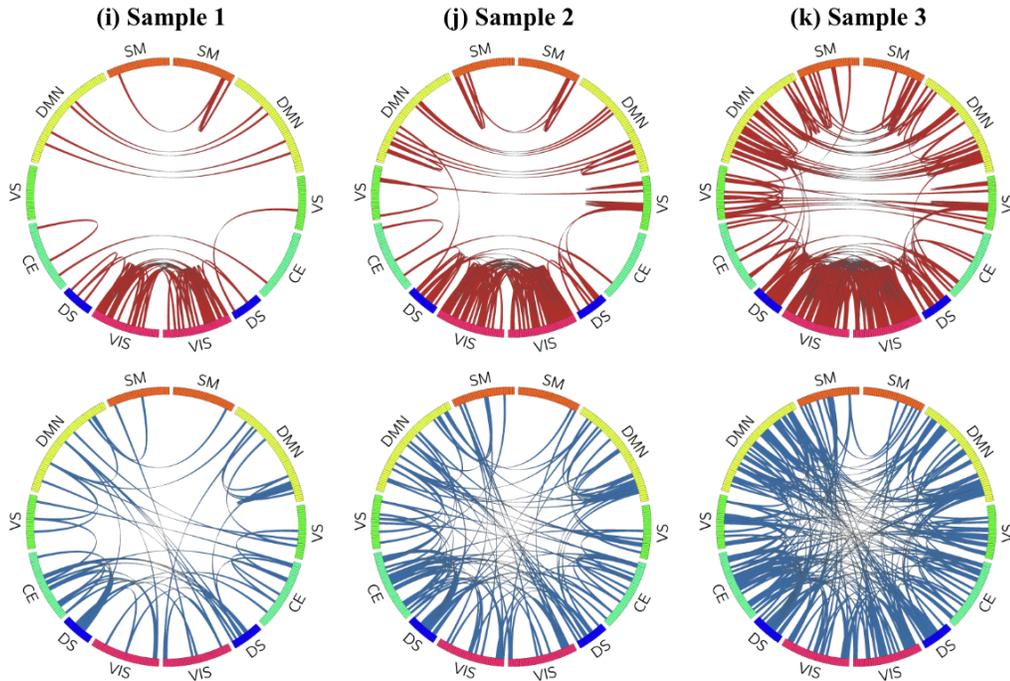


Figure 5.4: *Chord Diagrams for Three Randomly Selected Participants in the ABCD Dataset.* (i), (j), and (k) show correlation-based undirected connectivity (top row) and the corresponding Granger-causal directed connectivity (bottom row). Edges are thresholded at the top 2% for visualization.

Figure 5.4 illustrates chord diagrams from three randomly selected participants in the ABCD cohort. The top row displays undirected connectivity (red edges), while the bottom row shows the corresponding Granger-causal directed connectivity (blue edges). Strikingly, despite variations in individual-level dynamics, certain inter-network patterns consistently emerge across samples. For example, strong directed edges from the Sensorimotor (SM) network to both the Default Mode Network (DMN) and Visual (VIS) regions appear repeatedly. These recurring $SM \rightarrow DMN$ and $SM \rightarrow VIS$ influences support the notion of stable, large-scale information flow patterns, particularly from lower-order sensorimotor regions toward higher-order cognitive networks.

While individual differences in edge density, lateralization, and specific ROI con-

nections are expected due to noise and subject variability, the preservation of major hierarchical pathways indicates that the Granger-causal approach captures meaningful, generalizable aspects of brain functional architecture. Notably, directed connectivity maps are sparser and highlight predominant directional influences, thereby filtering out symmetric, bidirectional couplings that are present in undirected correlations. Recent developmental studies have demonstrated that such consistent directed patterns, especially the SM→DMN drive, are characteristic of typical maturation in youth [41, 40, 43]. Moreover, the reproducibility of these directed edges across samples reinforces their biological validity. As top-down connectivity from higher-order networks gradually strengthens with age, the persistent bottom-up influences observed here likely reflect a normative, developmental state where external sensorimotor inputs dominate internal cognitive processing [17].

Importantly, while these consistent pathways provide a robust template of neural communication, subtle individual differences remain. Variations in overall edge density, hemispheric lateralization, and the presence of intra-network loops suggest that personal factors or developmental stage nuances contribute to the observed connectivity profiles. Nevertheless, the stability of core directional links across diverse individuals highlights the potential of directed connectivity measures as reliable markers of typical brain development and as potential clinical biomarkers for neurodevelopmental disorders [33, 70]. Overall, the consistent emergence of these directed pathways underscores their representativeness and provides compelling evidence that the observed connectivity patterns are not artifacts of hyperparameter choices or isolated samples, but rather reflect systematic, biologically grounded neural dynamics shared across the population.

Chapter 6

Analysis

In this chapter, we synthesize our experimental findings and explore their broader implications for understanding neural connectivity and brain dynamics. Our results clearly demonstrate that graph-based models, particularly those leveraging Granger-causal directed graphs, maintain competitive predictive accuracy while providing valuable insights into the directional flow of neural signals. This dual advantage—robust performance coupled with richer interpretability—underscores the promise of causality-aware approaches in neuroscience research.

A key contribution of our work lies in revealing how directed edges capture asymmetrical and hierarchical interactions among brain regions. In particular, we observe that the Sensorimotor (SM) network frequently emerges as a driving influence on higher-order networks such as the Default Mode Network (DMN) and Visual (Vis) regions. This finding is consistent with neurobiological theories that propose SM regions play an active role in orchestrating and modulating cognitive processes, rather than merely co-activating alongside them. By highlighting such directional influences, our model goes beyond conventional undirected (correlation-based) methods, shedding light on plausible feedback loops and causal pathways that shape large-scale neural dynamics.

Our investigation into hyperparameter settings—namely window size (W), step size (S), and lag order (L)—further reveals that while specific choices can influence the density of inferred connections, the core connectivity structure remains consistent across a range of configurations. This robustness is important in both clinical and research contexts, where fMRI signals are typically noisy and subject to various artifacts. The ability to extract stable directed graphs even under different parameter choices indicates that Granger-based methods can capture fundamental temporal features of neural activity without being overly sensitive to minor methodological adjustments.

One particularly notable distinction is between undirected, correlation-based graphs and the directed graphs produced by Granger causality. Whereas correlation captures broad co-activation patterns, it does not illuminate the possible flow of information. By contrast, our directed approach reveals a richer picture of network organization, identifying potential upstream and downstream neural hubs and pointing to hierarchical ordering within and across major networks. These findings not only enhance predictive modeling but also open avenues for deeper neuroscientific inquiry, such as investigating how brain regions might shift roles from being drivers to receivers under different cognitive or developmental contexts.

Despite these strengths, practical considerations remain. First, constructing Granger-causal graphs can be computationally intensive, particularly for high temporal resolution datasets or when exploring larger lag orders. Continued efforts in algorithmic optimization, parallelization, or adaptive lag selection may help alleviate these challenges. Second, integrating complementary data modalities—such as structural connectivity from diffusion MRI or electrophysiological measures—could enable more comprehensive models and richer neurobiological interpretations. Third, coupling directed graph inference with emerging neural architectures, including attention-based or reinforcement-learning-inspired mechanisms, might further refine both predictive

performance and the discovery of meaningful causal insights. Finally, exploring end-to-end adjacency learning—where graph connectivity is jointly optimized alongside model parameters—offers a promising avenue to capture context-dependent interactions more flexibly than a strictly deterministic pipeline.

In summary, our findings emphasize the power of causality-aware graph representations in capturing the spatiotemporal complexity of the human connectome. By bridging the gap between traditional correlation-based analyses and more sophisticated causal inferences, these directed models pave the way for a deeper understanding of both typical and atypical brain functioning. This approach lays a solid foundation for future endeavors in cognitive neuroscience, clinical diagnostics, and any domain where elucidating the direction and magnitude of inter-regional influences is paramount.

Chapter 7

Conclusion

This study explores the integration of directed graph structures into brain connectivity analysis for gender classification, focusing on three key research questions that balance predictive performance and interpretability.

RQ1: Our experiments show that graph-based models substantially outperform MLPs, primarily by capturing complex, inter-regional interactions through functional connectivity. Additionally, Directed Graph Convolutional Networks (e.g., *BrainGB w/GC*) perform on par with their undirected counterparts (*BrainGB w/FC*), while also illuminating the causal flow of signals—a capability absent in standard correlation-based methods.

RQ2: We examine how hyperparameters such as window size, step size, and lag order affect the stability of Granger-causal graphs. Our results confirm that, despite the computational overhead of varying these parameters, the essential structure of directed connectivity remains consistent across a wide range of configurations. This robustness is valuable for both large-scale datasets and diverse experimental protocols.

RQ3: Comparative analyses reveal that undirected, correlation-based graphs offer broad co-activation patterns yet fail to capture directional influences among brain regions. By contrast, directed (Granger-based) networks highlight asymmetrical causal

pathways, clarifying how regions such as the Sensorimotor (SM) network may drive higher-order systems like the Default Mode Network (DMN) and Visual (Vis) regions. This directionality furnishes complementary insights into the hierarchical organization of the connectome.

Although these findings are promising, we note some limitations. Granger-causal graph construction can be computationally intensive—particularly for higher lag orders—and real-time or large-scale applications may require further optimization. Additionally, our two-stage deterministic pipeline does not facilitate end-to-end learning of the adjacency structure and model parameters, potentially restricting the discovery of subtler or context-dependent neural interactions. Future efforts could explore *learnable* adjacency frameworks and multi-modal data integration (e.g., structural MRI, electrophysiological signals) to enrich both the biological grounding and adaptability of the resulting networks.

Overall, by integrating causality-aware graph structures, this work enhances both predictive performance and our understanding of neural information flow in cognitive and clinical contexts. Our approach paves the way for more nuanced analyses of fMRI data, setting the stage for applications in developmental neuroscience, psychiatric research, and beyond.

Appendix A

Granger Causality Algorithms

In this appendix, we detail the algorithms for constructing a directed Granger causality graph from fMRI data. The process comprises: (1) data partitioning into overlapping time windows, (2) Vector Autoregressive (VAR) modeling with hypothesis testing for Granger causality, and (3) aggregation of significant causal influences. The final algorithm integrates these steps to output a directed adjacency matrix representing effective connectivity among brain regions.

A.1 Data Partitioning

Description. This step segments the time series data into overlapping windows, each of size W . The mean signal within each window is computed to reduce high-frequency noise and computational overhead in later steps. By controlling both window size W and step size S , this procedure captures short-term temporal dependencies while providing sufficient data for robust VAR modeling.

Algorithm 1: Data Partitioning

Input: $\mathbf{X} \in \mathbb{R}^{N \times T}$: fMRI data, $W, S \in$: window size and step size

Output: $\{\bar{\mathbf{X}}_w\}_{w=1}^{W_{\text{num}}}$: collection of window-averaged signals

1 **Compute number of windows:**

$$W_{\text{num}} \leftarrow \left\lfloor \frac{T - W}{S} \right\rfloor + 1$$

2 **for** $w = 1$ **to** W_{num} **do**

3 **for** $i = 1$ **to** N **do**

4

$$\bar{x}_{i,w} \leftarrow \frac{1}{W} \sum_{t=(w-1)S+1}^{(w-1)S+W} x_{i,t}$$

5 **end**

6 **end**

7 **return** $\{\bar{\mathbf{X}}_w\}_{w=1}^{W_{\text{num}}}$

A.2 VAR Modeling and Hypothesis Testing

Description. For each window, a VAR(L) model is fitted to pairs of ROI time series. The null hypothesis H_0 states that past values of ROI i do not explain additional variance in ROI j . If H_0 is rejected (based on a significance level α), we conclude that ROI i Granger-causes ROI j . These results are stored in an indicator matrix $\mathbf{I}^{(w)}$, one for each window w .

Algorithm 2: VAR Modeling and Hypothesis Testing

Input: Window-averaged data $\bar{\mathbf{X}}_w \in \mathbb{R}^N$,

VAR lag order $L \in$,

Significance level $\alpha \in (0, 1)$

Output: $\mathbf{I}^{(w)} \in \{0, 1\}^{N \times N}$: indicator matrix for window w

- 1 Initialize $\mathbf{I}^{(w)}$ with zeros;
- 2 **for** $i = 1$ to N **do**
- 3 **for** $j = 1$ to N **do**
- 4 **if** $i \neq j$ **then**
- 5 **Fit VAR(L) model:**
- $$x_j(t) = \alpha_{j,0}^{(w)} + \sum_{k=1}^L \alpha_{j,k}^{(w)} x_j(t-k) + \sum_{k=1}^L \beta_{ij,k}^{(w)} x_i(t-k) + \epsilon_j^{(w)}(t)$$
- Test null hypothesis** $H_0: \beta_{ij,1}^{(w)} = \dots = \beta_{ij,L}^{(w)} = 0$ *vs.* alternative;
- 6 **if** H_0 is rejected at level α **then**
- 7 $I_{ij}^{(w)} \leftarrow 1$
- 8 **end**
- 9 **end**
- 10 **end**
- 11 **end**
- 12 **return** $\mathbf{I}^{(w)}$

A.3 Aggregation of Test Results

Description. Causal inferences are aggregated across windows to capture consistent directional influences. Each pair (i, j) accrues a count of how many windows reject H_0 . If this count surpasses a threshold τ , edge $(i \rightarrow j)$ is deemed significant in the

final adjacency matrix.

Algorithm 3: Aggregation of Test Results

Input: $\{\mathbf{I}^{(w)}\}_{w=1}^{W_{\text{num}}}$: indicator matrices,
 Aggregation threshold $\tau \in$
Output: $\mathbf{A} \in \{0, 1\}^{N \times N}$: directed adjacency matrix

```

1 Initialize  $\mathbf{C} \in^{N \times N}$  with zeros;
2 for  $w = 1$  to  $W_{\text{num}}$  do
3   |  $\mathbf{C} \leftarrow \mathbf{C} + \mathbf{I}^{(w)}$ 
4 end
5 for  $i = 1$  to  $N$  do
6   | for  $j = 1$  to  $N$  do
7     | if  $i \neq j$  then
8       |
9     | end
10  | end
11 end
12 return  $\mathbf{A}$ 

```

$$\mathbf{A}_{ij} \leftarrow \begin{cases} 1, & \text{if } C_{ij} \geq \tau \\ 0, & \text{otherwise} \end{cases}$$

A.4 Granger Causality Graph Construction

Description. The final algorithm orchestrates the entire process. First, it partitions the input data using Algorithm 1, then performs Granger causality testing (Algorithm 2) for each window. Finally, it aggregates repeated causal detections using Algorithm 3 to yield a directed adjacency matrix \mathbf{A} .

Algorithm 4: Directed Granger Causality Graph Construction

Input: $\mathbf{X} \in \mathbb{R}^{N \times T}$: fMRI data,
 $W, S, L \in$: window size, step size, VAR lag order,
 $\alpha \in (0, 1)$: significance level,
 $\tau \in$: aggregation threshold
Output: $\mathbf{A} \in \{0, 1\}^{N \times N}$: directed adjacency matrix

1 Step 1: Data Partitioning

2 $\{\bar{\mathbf{X}}_w\}_{w=1}^{W_{\text{num}}} \leftarrow$ Algorithm 1 (\mathbf{X}, W, S);

3 Step 2: VAR Modeling and Testing

4 for $w = 1$ **to** W_{num} **do**

5 $\mathbf{I}^{(w)} \leftarrow$ Algorithm 2 ($\bar{\mathbf{X}}_w, L, \alpha$);

6 end

7 Step 3: Aggregation

8 $\mathbf{A} \leftarrow$ Algorithm 3 ($\{\mathbf{I}^{(w)}\}_{w=1}^{W_{\text{num}}}, \tau$);

9 return \mathbf{A}

Appendix B

Model Description

B.1 Multilayer Perceptron (MLP)

Given an input vector $\mathbf{x} \in \mathbb{R}^d$, a Multilayer Perceptron (MLP) computes a mapping $f_\theta(\mathbf{x})$ by applying a sequence of linear transformations and element-wise nonlinearities. Formally, let $\mathbf{h}^{(0)} = \mathbf{x}$, and for each layer $l = 1, \dots, L$:

$$\mathbf{h}^{(l)} = \sigma\left(\mathbf{W}^{(l)} \mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right),$$

where $\mathbf{W}^{(l)} \in \mathbb{R}^{d_l \times d_{l-1}}$ and $\mathbf{b}^{(l)} \in \mathbb{R}^{d_l}$ are trainable weights and biases, and $\sigma(\cdot)$ is a pointwise nonlinear activation (e.g., ReLU). After L layers, we obtain the final output $\hat{y} = \mathbf{W}^{(\text{out})} \mathbf{h}^{(L)} + \mathbf{b}^{(\text{out})}$ for either classification or regression tasks.

B.2 Graph Convolutional Network (GCN)

A Graph Convolutional Network (GCN) operates on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with N nodes. Let $\mathbf{X} \in \mathbb{R}^{N \times d}$ be the node feature matrix and $\mathbf{A} \in \{0, 1\}^{N \times N}$ the adjacency matrix (possibly directed or undirected, depending on the application). We often define $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$, and let $\tilde{\mathbf{D}}_{ii} = \sum_j \tilde{\mathbf{A}}_{ij}$ be the corresponding degree matrix [37].

For each GCN layer $l = 1, \dots, L$, let $\mathbf{H}^{(l)} \in \mathbb{R}^{N \times d_l}$ be the node embedding matrix (with $\mathbf{H}^{(0)} = \mathbf{X}$). The graph convolution updates the node embeddings as [37]:

$$\mathbf{H}^{(l+1)} = \sigma\left(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^{(l)} \mathbf{W}^{(l)}\right),$$

where $\mathbf{W}^{(l)} \in \mathbb{R}^{d_l \times d_{l+1}}$ are trainable weights, and $\sigma(\cdot)$ is an element-wise activation [37]. After L layers, a readout function aggregates the node embeddings into a single vector \mathbf{r} (e.g., via mean pooling), and the final prediction \hat{y} is obtained through a linear transformation of \mathbf{r} [37].

Whether the adjacency matrix \mathbf{A} is symmetric or not depends on whether undirected or directed edges are used.

Bibliography

- [1] H. Akaike. Akaike's information criterion. In *International Encyclopedia of Statistical Science*, pages 25–25. Springer, Berlin, Heidelberg, 2011.
- [2] Cecilia Algarin, Keerthana Deepti Karunakaran, Sussanne Reyes, Cristian Morales, Betsy Lozoff, Patricio Peirano, and Bharat Biswal. Differences on brain connectivity in adulthood are present in subjects with iron deficiency anemia in infancy. *Frontiers in aging neuroscience*, 9:54, 2017.
- [3] Abdul Rauf Anwar, Makii Muthalib, Stephane Perrey, Andreas Galka, Oliver Granert, Stephan Wolff, Günther Deuschl, Jan Raethjen, Ulrich Heute, and Muthuraman Muthuraman. Directionality analysis on functional magnetic resonance imaging during motor task using granger causality. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2287–2290. IEEE, 2012.
- [4] André M Bastos and Jan-Mathijs Schoffelen. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in systems neuroscience*, 9:175, 2016.
- [5] André Moraes Bastos, Julien Vezoli, Conrado Arturo Bosman, Jan-Mathijs Schoffelen, Robert Oostenveld, Jarrod Robert Dowdall, Peter De Weerd, Henry Kennedy, and Pascal Fries. Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85(2):390–401, 2015.

- [6] Janine Bijsterbosch, Stephen M Smith, and Christian Beckmann. *An introduction to resting state fMRI functional connectivity*. Oxford University Press, 2017.
- [7] Biao Cai, Pascal Zille, Julia M. Stephen, Tony W. Wilson, Vince D. Calhoun, and Yu-Ping Wang. Estimation of dynamic sparse connectivity patterns from resting state fmri. *IEEE Transactions on Medical Imaging*, 37(5):1224–1234, 2018. doi: 10.1109/TMI.2017.2786553.
- [8] Edward M Callaway. Feedforward, feedback and inhibitory connections in primate visual cortex. *Neural Networks*, 17(5-6):625–632, 2004.
- [9] B. J. Casey et al. The adolescent brain cognitive development (abcd) study: Imaging acquisition across 21 sites. *Developmental Cognitive Neuroscience*, 32: 43–54, 2018. doi: 10.1016/J.DCN.2018.03.001.
- [10] Daniel Chicharro and Anders Ledberg. When two become one: the limits of causality analysis of brain dynamics. *PloS one*, 7(3):e32466, 2012.
- [11] K.C. Chuang et al. Brain effective connectivity and functional connectivity as markers of lifespan vascular exposures in middle-aged adults: The bogalusa heart study. *Frontiers in Aging Neuroscience*, 15, 2023. doi: 10.3389/fnagi.2023.1110434.
- [12] H. Cui et al. Braingb: A benchmark for brain network analysis with graph neural networks (extended abstract). In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 4968–4969, 2022. doi: 10.1109/TMI.2022.3218745.
- [13] Jessica S Damoiseaux and Michael D Greicius. Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity. *Brain structure and function*, 213:525–533, 2009.

- [14] G. Deshpande et al. Effective connectivity during haptic perception: A study using granger causality analysis of functional magnetic resonance imaging data. *NeuroImage*, 40(4):1807–1814, 2008. doi: 10.1016/J.NEUROIMAGE.2008.01.044.
- [15] Mingzhou Ding, Yonghong Chen, and Steven L Bressler. Granger causality: basic theory and application to neuroscience. *Handbook of time series analysis: recent theoretical developments and applications*, pages 437–460, 2006.
- [16] E. Earl et al. Human connectome project minimal preprocessing pipelines to nipy. *GigaScience*, 5, 2016. doi: 10.1186/S13742-016-0147-0-H.
- [17] Monica E. Ellwood-Lowe, Susan Whitfield-Gabrieli, and Silvia A. Bunge. Brain network coupling associated with cognitive performance varies as a function of a child’s environment in the abcd study. *Nature Communications*, 12(1):7183, 2021. doi: 10.1038/s41467-021-27336-y.
- [18] Monique Ernst, Salvatore Torrisi, Nicholas Balderston, Christian Grillon, and Elizabeth A Hale. fmri functional connectivity applied to adolescent neurodevelopment. *Annual review of clinical psychology*, 11(1):361–377, 2015.
- [19] Andrew A Fingelkurts, Alexander A Fingelkurts, and Seppo Kähkönen. Functional connectivity in the brain—is it an elusive concept? *Neuroscience & Biobehavioral Reviews*, 28(8):827–836, 2005.
- [20] K. J. Friston, C. D. Frith, P. F. Liddle, and R. S. J. Frackowiak. Functional connectivity: The principal-component analysis of large (pet) data sets. *Journal of Cerebral Blood Flow and Metabolism*, 13(1):5–14, 1993.
- [21] K. J. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *NeuroImage*, 19(4):1273–1302, 2003.

- [22] K. J. Friston, J. Kahan, B. Biswal, and A. Razi. A dcm for resting state fmri. *NeuroImage*, 94:396–407, 2014.
- [23] Zening Fu, Jing Sui, Armin Iraj, Jingyu Liu, and Vince D. Calhoun. Cognitive and psychiatric relevance of dynamic functional connectivity states in a large ($n = 10,000$) children population. *Molecular Psychiatry*, 30(2):402–413, 2025. doi: 10.1038/s41380-024-02683-6.
- [24] M. Gilson et al. Estimation of directed effective connectivity from fmri functional connectivity hints at asymmetries of cortical connectome. *PLOS Computational Biology*, 12(3), 2016. doi: 10.1371/JOURNAL.PCBI.1004762.
- [25] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.
- [26] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969. doi: 10.2307/1912791.
- [27] Michael D Greicius, Kaustubh Supekar, Vinod Menon, and Robert F Dougherty. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral cortex*, 19(1):72–78, 2009.
- [28] Michael D Greicius, Kaustubh Supekar, Vinod Menon, and Robert F Dougherty. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral cortex*, 19(1):72–78, 2009.
- [29] T. S. Hale et al. Visual network asymmetry and default mode network function in adhd: An fmri study. *Frontiers in Psychiatry*, 5:81, 2014. doi: 10.3389/FPSYT.2014.00081.
- [30] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

- [31] X. Kan et al. Dynamic brain transformer with multi-level attention for functional brain network analysis. arXiv:2309.01941, 2023. Preprint.
- [32] Xuan Kan, Hejie Cui, Joshua Lukemire, Ying Guo, and Carl Yang. Fbnetgen: Task-aware gnn-based fmri analysis via functional brain network generation. In *International Conference on Medical Imaging with Deep Learning*, pages 618–637. PMLR, 2022.
- [33] Nicole R. Karcher, Kevin J. O’Brien, Sripada Kandala, and Deanna M. Barch. Associations between resting-state functional connectivity and a hierarchical dimensional structure of psychopathology in middle childhood. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 6(5):508–517, 2021. doi: 10.1016/j.bpsc.2020.09.008.
- [34] J. Kawahara et al. Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage*, 146:1038–1049, 2017. doi: 10.1016/J.NEUROIMAGE.2016.09.046.
- [35] Dae-Jin Kim, Patrick D Skosnik, Hu Cheng, Ben J Pruce, Margaret S Brumbaugh, Jennifer M Vollmer, William P Hetrick, Brian F O’Donnell, Olaf Sporns, Aina Puce, et al. Structural network topology revealed by white matter tractography in cannabis users: a graph theoretical analysis. *Brain connectivity*, 1(6):473–483, 2011.
- [36] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [37] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.

- [38] Nora Leonardi and Dimitri Van De Ville. On spurious and real fluctuations of dynamic functional connectivity during rest. *NeuroImage*, 104:430–436, 2015. doi: 10.1016/j.neuroimage.2014.09.007.
- [39] X. Li et al. Braingnn: Interpretable brain graph neural network for fmri analysis. *Medical Image Analysis*, 74:102233, 2021. doi: 10.1016/J.MEDIA.2021.102233.
- [40] Martina J. Lund, Dag Alnæs, Jaroslav Rokicki, Simon Schwab, Ole A. Andreassen, Lars T. Westlye, and Tobias Kaufmann. Functional connectivity directionality between large-scale resting-state networks across typical and non-typical trajectories in children and adolescence. *PLOS ONE*, 17(12):e0276221, 2022. doi: 10.1371/journal.pone.0276221.
- [41] Audrey C. Luo, Valerie J. Sydnor, Adam Pines, Bart Larsen, Antonia N. Kaczurkin, R. Ciric, and *et al.* Functional connectivity development along the sensorimotor-association axis enhances the cortical hierarchy. *Nature Communications*, 15(1):3511, 2024. doi: 10.1038/s41467-024-47748-w.
- [42] A. A. Neath and J. E. Cavanaugh. The bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(2):199–203, 2012.
- [43] Adam Pines, Arielle S. Keller, Bart Larsen, Maxwell A. Bertolero, Arian Ashourvan, and Danielle S. *et al.* Bassett. Development of top-down cortical propagations in youth. *Neuron*, 111(8):1316–1330.e5, 2023. doi: 10.1016/j.neuron.2023.01.014.
- [44] J. D. Power et al. Functional network organization of the human brain. *Neuron*, 72(4):665–678, 2011. doi: 10.1016/J.NEURON.2011.09.006.
- [45] Maria Giulia Preti, Thomas A. W. Bolton, and Dimitri Van De Ville. The

- dynamic functional connectome: state-of-the-art and perspectives. *NeuroImage*, 160:41–54, 2017. doi: 10.1016/j.neuroimage.2016.12.061.
- [46] Victor Ramirez-Amaya, Diano F Marrone, Fred H Gage, Paul F Worley, and Carol A Barnes. Integration of new neurons into functional neural networks. *Journal of Neuroscience*, 26(47):12237–12241, 2006.
- [47] A. Razi and K. J. Friston. The connected brain: Causality, models, and intrinsic dynamics. *IEEE Signal Processing Magazine*, 33(3):14–22, 2016.
- [48] A. Riaz et al. Fcnet: a convolutional neural network for calculating functional connectivity from functional mri. In *Connectomics in NeuroImaging: First International Workshop, CNI 2017, Held in Conjunction with MICCAI 2017*, pages 70–78, 2017.
- [49] Alard Roebroeck, Elia Formisano, and Rainer Goebel. Mapping directed influence over the brain using granger causality and fmri. *Neuroimage*, 25(1):230–242, 2005.
- [50] Baxter P Rogers, Victoria L Morgan, Allen T Newton, and John C Gore. Assessing functional connectivity in the human brain by fmri. *Magnetic resonance imaging*, 25(10):1347–1357, 2007.
- [51] T. D. Satterthwaite et al. Neuroimaging of the philadelphia neurodevelopmental cohort. *NeuroImage*, 86:544–553, 2014. doi: 10.1016/J.NEUROIMAGE.2013.07.064.
- [52] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
- [53] M. L. Seghier and K. J. Friston. Network discovery with large dcms. *NeuroImage*, 68:181–191, 2013.

- [54] Rachael D Seidler, Douglas C Noll, and G Thiers. Feedforward and feedback processes in motor control. *Neuroimage*, 22(4):1775–1783, 2004.
- [55] Anil K. Seth, Adam B. Barrett, and Lionel Barnett. Granger causality analysis in neuroscience and neuroimaging. *Journal of Neuroscience*, 35(8):3293–3297, 2015. doi: 10.1523/JNEUROSCI.4399-14.2015.
- [56] Anil K Seth, Adam B Barrett, and Lionel Barnett. Granger causality analysis in neuroscience and neuroimaging. *Journal of Neuroscience*, 35(8):3293–3297, 2015.
- [57] Maksim G Sharaev, Viktoria V Zavyalova, Vadim L Ushakov, Sergey I Kartashov, and Boris M Velichkovsky. Effective connectivity within the default mode network: dynamic causal modeling of resting-state fmri data. *Frontiers in human neuroscience*, 10:14, 2016.
- [58] Ali Shojaie and Emily B Fox. Granger causality: A review and recent advances. *Annual Review of Statistics and Its Application*, 9(1):289–319, 2022.
- [59] C. A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980. doi: 10.2307/1912017.
- [60] J. Smallwood et al. The default mode network in cognition: a topographical perspective. *Nature Reviews Neuroscience*, 22(8):503–513, 2021. doi: 10.1038/S41583-021-00474-4.
- [61] K.E. Stephan et al. Analyzing effective connectivity with functional magnetic resonance imaging. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(3): 446–459, 2010. doi: 10.1002/WCS.58.
- [62] Ian H Stevenson, James M Rebesco, Lee E Miller, and Konrad P Körding. Inferring functional connections between neurons. *Current opinion in neurobiology*, 18(6): 582–588, 2008.

- [63] Patrick A. Stokes and Patrick L. Purdon. A study of problems encountered in granger causality analysis from a neuroscience perspective. *Proceedings of the National Academy of Sciences*, 114(34):E7063–E7072, 2017. doi: 10.1073/pnas.1704663114.
- [64] Lucina Q Uddin, AM Clare Kelly, Bharat B Biswal, F Xavier Castellanos, and Michael P Milham. Functional connectivity of default mode network components: correlation, anticorrelation, and causality. *Human brain mapping*, 30(2):625–637, 2009.
- [65] Martijn P Van Den Heuvel and Hilleke E Hulshoff Pol. Exploring the brain network: a review on resting-state fmri functional connectivity. *European neuropsychopharmacology*, 20(8):519–534, 2010.
- [66] Victor M Vergara, Anees Abrol, and Vince D Calhoun. An average sliding window correlation method for dynamic functional connectivity. *Human Brain Mapping*, 40(7):2089–2103, 2019.
- [67] Simon Wein, Wilhelm M Malloni, Ana Maria Tomé, Sebastian M Frank, G-I Henze, Stefan Wüst, Mark W Greenlee, and Elmar W Lang. A graph neural network framework for causal inference in brain networks. *Scientific reports*, 11(1):8061, 2021.
- [68] X. Wen, G. Rangarajan, and M. Ding. Is granger causality a viable technique for analyzing fmri data? *PLoS ONE*, 8(7):e67428, 2013.
- [69] Xiaotong Wen, Govindan Rangarajan, and Mingzhou Ding. Is granger causality a viable technique for analyzing fmri data? *PloS one*, 8(7):e67428, 2013.
- [70] David Willinger, Isabelle Häberling, Iva Ilioska, Gregor Berger, Susanne Walitza, and Silvia Brem. Weakened effective connectivity between salience network and

- default mode network during resting state in adolescent depression. *Frontiers in Psychiatry*, 15:1386984, 2024. doi: 10.3389/fpsyt.2024.1386984.
- [71] Adam J Woods, Roy H Hamilton, Alexander Kranjec, Preet Minhaus, Marom Bikson, Jonathan Yu, and Anjan Chatterjee. Space, time, and causality in the human brain. *Neuroimage*, 92:285–297, 2014.
- [72] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.
- [73] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.
- [74] Chun-Hung Yeh, Derek K Jones, Xiaoyun Liang, Maxime Descoteaux, and Alan Connelly. Mapping structural connectivity using diffusion mri: challenges and opportunities. *Journal of Magnetic Resonance Imaging*, 53(6):1666–1682, 2021.
- [75] Yue Yu, Xuan Kan, Hejie Cui, Ran Xu, Yujia Zheng, Xiangchen Song, Yanqiao Zhu, Kun Zhang, Razieh Nabi, Ying Guo, Chao Zhang, and Carl Yang. Deep dag learning of effective brain connectivity for fmri analysis. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, 2023. doi: 10.1109/ISBI53787.2023.10230429.
- [76] Andrew Zalesky and Michael Breakspear. Towards a statistical test for functional connectivity dynamics. *Neuroimage*, 114:466–470, 2015.
- [77] Hao Zhang, Ran Song, Liping Wang, Lin Zhang, Dawei Wang, Cong Wang, and Wei Zhang. Classification of brain disorders in rs-fmri via local-to-global graph neural networks. *IEEE transactions on medical imaging*, 42(2):444–455, 2022.

- [78] Xingye Zhang, Shaoqian Wang, Jesse B Hoagg, and T Michael Seigler. The roles of feedback and feedforward as humans learn to control unknown dynamic systems. *IEEE transactions on cybernetics*, 48(2):543–555, 2017.