

Causal Functional Connectivity from Neural Dynamics

Rahul Biswas

A dissertation

submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington
2024

Reading Committee:

Eli Shlizerman, Chair

Chet Moritz

Andrea Stocco

Program Authorized to Offer Degree:
Electrical and Computer Engineering

© Copyright 2024

Rahul Biswas

University of Washington

Abstract

Causal Functional Connectivity from Neural Dynamics

Rahul Biswas

Chair of the Supervisory Committee:

Eli Shlizerman

Electrical and Computer Engineering

Representation of brain network interactions is fundamental to the translation of neural structure to brain function. As such, methodologies for mapping neural interactions into structural models, i.e., inference of functional connectome from neural recordings, are key for the study of brain networks. While multiple approaches have been proposed for functional connectomics based on statistical associations between neural activity, association does not necessarily incorporate causation. Additional approaches have been proposed to incorporate aspects of causality to turn functional connectomes into causal functional connectomes, however, these methodologies typically focus on specific aspects of causality. This warrants a systematic statistical framework for causal functional connectomics that defines the foundations of common aspects of causality, to assist in contrasting existing approaches and guide development of further causal methodologies. In this research, we first develop such a statistical guide and perform a comparative study of existing approaches of causal modeling for finding the causal functional connectome (CFC).

Thereafter, we consider a popular framework for inferring causal connectivity from observations - *Directed Probabilistic Graphical Models*. Its common formulation is not suitable for neural time series since it was developed for variables with independent and identically distributed static samples. In this work, we propose to model and estimate the causal functional connectivity from neural time series using a novel approach that adapts directed probabilistic graphical modeling to the time series scenario. In particular,

we develop the *Time-Aware PC* (TPC) algorithm for estimating the causal functional connectivity, which adapts the Peter-Clark (PC) algorithm—a state-of-the-art method for statistical causal inference. We show that the model outcome of TPC has the properties of reflecting causality of neural interactions such as being non-parametric, exhibits the *directed Markov* property in a time-series setting, and is predictive of the consequence of counterfactual interventions on the time series. We demonstrate the utility of the methodology to obtain the causal functional connectome for several datasets including simulations, benchmark datasets, and recent multi-array electro-physiological recordings from the mouse visual cortex.

We establish mathematical guarantees of estimation by the TPC algorithm under a standard set of assumptions on the underlying time series involving ρ -mixing that encompasses popular time series models such as vector auto-regressive moving average and linear processes. These findings are supported by simulations and benchmark real data analyses with TPC algorithm using different candidates of conditional dependence tests and with/without subsampling.

Finally, we compute the CFC from time series resting state *functional magnetic resonance imaging* (rs-fMRI) data recorded in studies of *Alzheimer's disease*. We apply the TPC algorithm to infer the CFC for the whole brain from rs-fMRI recordings for subjects from three clinical groups: cognitively normal, mild cognitive impairment, and Alzheimer's disease. We compare the CFC outcome of TPC with that of other related approaches in the literature. Then, we use the CFC outcomes of TPC and perform an exploratory analysis of the difference in strengths of CFC edges between Alzheimer's and cognitively normal groups, based on edge-wise p-values obtained by Welch's t-test. The brain regions thus identified are found to be in agreement with literature on brain regions impacted by Alzheimer's disease, published by researchers from clinical/medical institutions.

DEDICATION

Dedicated to the Supreme, to my parents Srikrishna and Shyamali Biswas, to the Brahma Kumaris,
and to all my teachers.

ACKNOWLEDGEMENTS

I express my heartfelt gratitude to my advisor, Dr. Eli Shlizerman, for his help, cooperation and support in this research. In addition, his guidance in producing high quality research articles and presentations is a valuable asset for my future career. Thanks also to my committee members for their help in my progress. I am fortunate to have such a wonderful supervisory committee.

Thanks to our graduate program coordinator, Dr. Anant Anantram, for supportive interactions and research discussions. Thanks also to all the advising and administrative staff at the Department of Electrical and Computer Engineering for their kind assistance.

Thanks to Dr. Thomas Richardson from the Department of Statistics for introducing me to causal modeling in his course. Thanks also to Dr. Ranjini Grove and Dr. Abel Rodriguez for their support.

I express my deepest gratitude to Dr. SuryaNarayana Sripada for his inspiration to pursue this direction of research. His invaluable guidance helped me to enter and navigate in the field of neuroscience. I am also thankful to his collaboration in research papers. Furthermore, his advice was consequential to navigate through turning points during my PhD journey.

I am thankful to all my friends and colleagues at UW, Mingfei Chen, Jie Gao, Sheridan Grant, Holly Hake, Jimin Kim, Nilanjana Laha, Astitwa Lathe, Trung Le, Jingyuan Li, Xiulong Liu, Anwasha Pan, Yandi Shen, Kun Su, Jinlin Xiang, Yang Zheng, for their camaraderie and encouragement.

Thanks to Dr. Somabha Mukherjee from the National University of Singapore for his authentic friendship and collaboration.

I am also thankful to the collaboration and mentorship of Dr. Xiangrong Kong of Johns Hopkins University since the time of my internship in 2014.

I express thanks to my professors at Indian Statistical Institute for introducing me to the profound beauty of mathematics and its statistical applications. In particular, I would like to thank Dr. Sashi Mohan Srivastava, Dr. Alok Goswami, Dr. Bimal Roy, Dr. Pradipta Bandyopadhyay, Dr. Sumitra Purkayastha, and Dr. Kiranmoy Das.

At a personal level, I express my deepest gratitude to the Brahma Kumaris for teaching a meditation practice, which is open to people of all backgrounds, free of charge. This meditation practice has helped me

to keep my mind stable despite challenges during my PhD that include a pandemic. In particular, thanks to my Brahma Kumaris Seattle family for their care and affection during my stay in Seattle.

I express my heartfelt gratitude to my parents, Srikrishna and Shyamali Biswas, for their unwavering support and sacrifices since my childhood till date. They always inspired me to pursue truth and humility in terms of values. My father, a high school physics teacher, helped to clear my basic concepts in science that fueled my progress forward. My mother is always caring and supports my endeavors in every field.

I am thankful for the support of my larger family. In particular, my grandfather, Ranajit Mandal, who is a high school mathematics teacher, supported my mathematics education in early childhood. My maternal uncle, Sandip Mandal, also supported my education in general in my early childhood. I am also thankful to my maternal uncle, Dr. Parthasarathi Mandal, professor at the University of Manchester, for his inspirations.

I am thankful to all my teachers in high school for imparting education to me. In particular, thanks to Dr. Zahedul Islam and Mrs. Swapna Islam for teaching me all through primary and secondary school not only on curricular topics but also laying a foundation of values and service to mankind.

Finally, I express heartfelt thanks to the Supreme for his eternal company and guidance to pursue truth.

Contents

1	Introduction	1
1.1	Outline	5
2	Functional and Causal Neural Connectomics: A Comparative Study	9
2.1	Neural Connectomics: Anatomical and Functional	9
2.1.1	Anatomical Connectome	10
2.1.2	Functional Connectome (Associative)	11
2.2	From Association to Causation	15
2.2.1	Causal Modeling	16
2.2.2	Granger Causality	19
2.2.3	Dynamic Causal Model	20
2.2.4	Directed Probabilistic Graphical Models	22
2.2.5	Comparitive Study of Approaches to Causal Functional Connectivity	27
2.3	Discussion	32
3	Functional and Causal Neural Connectomics: The Time-Aware PC Algorithm	39
3.1	Causal Functional Connectivity for Time Series	39
3.1.1	Unrolled Graphical Modeling of Time Series	39
3.1.2	Rolled CFC-DPGM	41
3.1.3	Estimation from data: Time-Aware PC (TPC) Algorithm	42
3.2	Results	45
3.2.1	Comparison Study of TPC with other Approaches to Causal Functional Connectivity	45

3.2.2	Application to Benchmark Data	50
3.2.3	Application to Neurobiological Data	51
3.3	Discussion	56
3.4	Materials and Methods	61
3.4.1	Causal Functional Connectivity for Static Variables - Review	61
3.4.2	Choice of gap between time-advanced samples in TPC algorithm	63
3.4.3	Connectivity weights in the CFC for time series	64
3.4.4	Properties of Rolled CFC-DPGM	66
4	Mathematical Guarantees	71
4.1	Causal Structure Learning	75
4.1.1	Conditions for Causal Inference	75
4.1.2	Estimation of the CPDAG: The PC algorithm	77
4.1.3	Causal recovery in time series: The Time-aware PC algorithm	77
4.1.4	Conditional Dependence Test in the Gaussian Regime: Pearson's Partial Correlations	79
4.1.5	Conditional Dependence Test in the Non-Gaussian Regime: The Hilbert Schmidt Criterion	80
4.2	Consistency of the PC Algorithm under temporal dependence	81
4.2.1	The General Consistency Framework	82
4.2.2	Consistency of Tests for Conditional Dependence	84
4.3	Analogous Results for Strongly Mixing Processes	88
4.4	Common Time Series Models	89
4.4.1	VARMA Process	89
4.4.2	Linear Process	90
4.5	Consistency of the Time-Aware PC algorithm	91
4.5.1	Main Result	92
4.6	Simulation Studies	93
4.7	Application to River Runoff Benchmark Data	102
4.8	Discussion	104

5	Causal Functional Connectivity in Alzheimer’s Disease Computed from Time Series fMRI data	105
5.1	Materials and Methods	108
5.1.1	Participants	108
5.1.2	Image Acquisition	108
5.1.3	fMRI Preprocessing	109
5.1.4	Inference of causal functional connectivity: Time-aware PC algorithm	109
5.1.5	Comparison with functional connectivity using other approaches	111
5.1.6	Alterations of CFC edges in Alzheimer’s disease	111
5.2	Results	112
5.2.1	Subject-specific Causal Functional Connectivity	112
5.2.2	Comparison with functional connectivity using other approaches	113
5.2.3	Alterations of CFC edges in Alzheimer’s disease	114
5.3	Discussion	116
6	Conclusion	119
6.1	Code Availability	122
A	Appendix to the Time-Aware PC Algorithm	159
A.1	Proof of Theorem 1	159
A.2	Proof of Corollary 7.1	159
A.3	Simulation Study Details	160
A.4	Benchmark Datasets	162
A.5	Visual Coding Neuropixels Dataset	163
B	Appendix to Mathematical Guarantees	165
B.1	Facts about the VARMA Process (4.8)	165
B.2	Simulation Study Details	167
B.3	Benchmark Datasets	169

C Appendix to Causal Functional Connectivity in Alzheimer’s Disease	171
C.1 Automated Anatomical Labeling (AAL) Atlas	171

List of Figures

2.1	An undirected PGM and its Markov Property (Figure from [1])	14
2.2	Various connectome representations for the same network of neurons (Figure from [1])	16
2.3	Causal graphs for electrical activity of brain nuclei for fear perception from visual stimulus [1]	17
2.4	Schematic of the haemodynamic model used by DCM for fMRI (Figure from [1])	21
2.5	Directed Markov Property in the context of fear stimulus (Figure from [1])	23
2.6	The PC algorithm (Figure from [1])	26
2.7	Comparative study of CFC inference (Figure from [1])	29
3.1	Causal modeling of neural time series (Figure from [2])	40
3.2	The TPC algorithm (Figure from [2])	42
3.3	Comparative study of CFC inference (Figure from [2])	46
3.4	Comparative study over levels of noise and thresholding parameter (Figure from [2])	48
3.5	Interventional connectivity weights (Figure from [2])	49
3.6	Application to Neuropixels dataset (Figure from [2])	52
3.7	Graphical comparison of estimated CFC over stimuli (Figure from [2])	54
3.8	Comparison of performance with Graph Neural Networks	60
3.9	Impact of intervention (Figure from [2])	69
4.1	Schematic depiction of causal modeling of time series by the Unrolled and Rolled causal graphs (Figure from [3])	73
4.2	Comparative study in simulations with TPC using different conditional dependence tests and with/without subsampling (Figure from [3])	94

4.3 Comparative study of performance based on TPR and TFPR over different noise levels when using different variations of TPC (Figure from [3]) 100

4.4 Comparative study of performance based on Combined Score over different levels of noise and significance level when using different variations of TPC (Figure from [3]) 101

5.1 Steps conveying the concept of the TPC algorithm to infer CFC from observed neural time series data (Figure from [4]) 110

5.2 CFC for an example subject who is CN, estimated by TPC algorithm (Figure from [4]) . . . 112

5.3 Comparison and demonstration of FC inferred by three methods: Associative FC using SPC, and Causal FC using TPC and GC (Figure from [4]) 113

5.4 Causal functional connections with edge-weights differing between clinical groups with edge-wise p-values ranging in 0 – 0.05 based on t-test (Figure from [4]) 115

List of Tables

2.1	Comparative summary of different approaches for causal modeling in functional connectomics (Table from [1])	33
3.1	Comparison of CFC inference on benchmark datasets (Table from [2])	50
3.2	Comparative summary of different approaches for causal modeling (Table from [2])	70
4.1	Percentage of occurrence of edges among simulation trials (Table from [3])	98
4.2	Comparison of performance on benchmarking datasets when using different variations of TPC (Table from [3])	103
5.1	Summary of demographic information and Mini Mental State Examination (MMSE) for CN, MCI and AD subjects (Table from [4])	108
5.2	CFC edges with lowest edge-wise p-values for CN vs. AD group comparison (Table from [4])	114
C.1	Names of regions in the AAL Atlas (Table from [4])	171

List of Abbreviations

A Accuracy. 29

AFC Associative Functional Connectivity. 11

CeA Central nucleus of the amygdala. 16

CFC Causal Functional Connectivity. 2

CTRNN Continuous Time Recurrent Neural Network. 28

DAG Directed Acyclic Graph. 17

DCM Dynamic Causal Model. 20

DMP Directed Markov Property. 24

DPGM Directed Probabilistic Graphical Model. 22

EEG Electroencephalography. 20

FC Functional Connectivity. 1

fMRI Functional Magnetic Resonance Imaging. 10

FN False Negative. 30

FP False Positive. 30

GC Granger Causality. 19

GES Greedy Equivalence Search. 27

GIES Greedy Interventional Equivalence Search. 27

HPA hypothalamus-pituitary-adrenal axis. 16

LGN Lateral geniculate nucleus. 16

MEG Magnetoencephalography. 20

ML Machine Learning. 3

MP Markov Property. 24

P Pulvinar. 16

PC Peter Clark. xiii

PGM Probabilistic Graphical Model. 12

PVN paraventricular nucleus. 16

S Sensitivity. 29

SC Superior colliculus. 16

TN True Negative. 30

TP True Positive. 29

VC Visual cortex. 16

Chapter 1

Introduction

This introduction is partially adapted from published work in Biswas and Shlizerman [1, 2].

The term “connectome” typically refers to a network of neurons and their anatomical links, such as chemical and electrical synapses. The connectome represents the anatomical map of the neural circuitry of the brain [5]. Connectome mapping can be achieved with the help of imaging techniques and computer vision methods at different scales [6; 7; 8]. The aim of finding the connectome is to provide insight into how neurons are connected and how they interact to form brain function.

While the anatomical connectome includes the backbone information on possible ways of how neurons could interact, it does not fully reflect the “wiring diagram of the brain”, which is expected to incorporate the dynamic nature of neurons’ activity and their interactions [9; 10; 11; 12]. In particular, the anatomical connectome does not correspond to how the anatomical structure relates to brain function since each anatomical connectivity map can encode several functional outcomes of the brain [13]. Thereby, the term connectome has been extended beyond the anatomical meaning. In particular, a map that reflects neurons’ functions is named Functional Connectome (FC) and it represents the network of interactions between neurons with respect to their activity over time [14]. Inferring FC is expected to lead to more fundamental understanding of brain function and dysfunction [15]. Indeed, FC is expected to include and facilitate inference of the governing neuronal pathways essential for brain functioning and behavior [16]. Two neurons are said to be functionally connected if there is a significant relationship between their activity over time where the activity can be recorded from neurons over time and measured with various measures [17]. In

contrast to the anatomical connectome, the functional connectome needs to be inferred, as it cannot be directly observed or mapped, since the transformation from activity to associations is intricate.

Several approaches have been introduced to infer the FC. These include approaches based on measuring correlations, such as pairwise correlation [18; 19], or sparse covariance matrix that is comparatively better than correlations given limited time points [20; 21]. Furthermore, for such scenarios, regularized precision matrix approaches were proposed to better incorporate conditional dependencies between neural time courses, where the precision matrix is inferred by a penalized maximum likelihood to promote sparsity [22; 23; 24]. While there is a wide variety of methods, there is still a lack of unification as to what defines the “functional connectome”. A taxonomy which provides a systematic treatment grounded from definitions and followed with algorithmic properties is currently unavailable. An existing taxonomy for FC considers generic aspects, such as, undirected and directed, model-based and model-free, time and frequency-domains [25]. Here, we consider the angle of association and causation, such as pairwise association vs. non-pairwise graphical associations.

Moreover, the prevalent research on FC, outlined above, deals with finding associations between neural signals in a non-causal manner. That is, in such a mapping we would know that a neuron A and a neuron B are active in a correlated manner, however, we would not know whether the activity in neuron A causes neuron B to be active ($A \rightarrow B$), or is it the other way around ($B \rightarrow A$)? Or, is there a neuron C which intermediates the correlation between A and B ($A \leftarrow C \rightarrow B$)? In short, questions of this nature distinguish causation from association. In this report, we provide a statistical framework for FC and investigate the aspect of causality in the notion of Causal Functional Connectome (CFC) which would answer the aforementioned causal questions.

While FC refers to the network of interactions between units of the brain, such as individual neurons or brain regions, with respect to their activity over time [14], the interactions can be associative or causal. When the interactions are causal, the network is termed as *Causal Functional Connectome* (CFC). The CFC maps how neural activity flows within neural circuits, and provides the possibility for inference of neural pathways essential for brain functioning and behavior, such as sensory-motor-behavioral pathways [16].

Several approaches aim to infer CFC, such as Granger Causality (GC), Dynamic Causal Modeling (DCM), and Directed Probabilistic Graphical Models (DPGM), each having their applicability and challenges, as

surveyed in [1]. GC obtains the directed functional connectivity from observed neural activity in a way that tells whether a neuron's *past* is predictive of another neuron's future, however it is unclear whether the prediction implies causation. In contrast, DCM compares *specific* mechanistic biological models based on data evidence, in which, model parameters represent causal influences between hidden neural states [26]. On the other hand, DPGM is a generic procedure to obtain causal relationships between nodes of a network from observations, using the *directed Markov* property. DPGM is non-parametric in the sense of capturing arbitrary functional relationships among the nodes and is predictive of the consequence of counterfactual interventions to the network [27; 28]. Such properties make DPGM a popular approach for causal modeling in various disciplines such as genomics and econometrics [29; 30; 31; 32; 33; 34; 35; 36; 37; 38; 39; 40].

In comparison, the landscape of machine learning (ML) today is predominantly dominated by predictive models that excel in identifying patterns within massive datasets [41; 42]. These models, powered by advancements in computational power and algorithmic sophistication, have achieved remarkable success in a wide range of applications, from image recognition to natural language processing [43; 44]. However, this success largely hinges on correlational relationships within data, which, while powerful for prediction in stable environments, often falter when applied to situations where the underlying data distribution changes or when decisions require understanding the effect of interventions [45; 46]. Traditional ML approaches are adept at prediction tasks where the goal is to forecast outcomes based on historical data. These models rely on the assumption that the future will resemble the past, making them less equipped to handle scenarios that demand an understanding of causal relationships—i.e., the mechanisms through which one variable influences another [47].

In contrast, causal inference seeks to understand and model these mechanisms, enabling predictions about the effects of potential interventions or changes. Causal models, such as Directed Probabilistic Graphical Models (DPGMs) and Structural Causal Models (SCMs), go beyond mere associations to describe how systems behave under manipulation. This capability is critical for tasks like policy evaluation, where understanding the impact of a proposed action is necessary, or in medicine, where predicting the outcome of a treatment on a new patient group is essential [48; 49; 50; 45]. Bernhard Schölkopf's work on "Causality for Machine Learning" represents a pivotal step towards integrating causal reasoning into ML [51]. Schölkopf argues that for ML to advance further, it must move beyond correlation to causation, enabling models to make

reliable predictions in novel situations and under interventions. Causal models such as DPGMs and SCMs offer frameworks for this integration, providing a systematic way to represent causal relationships. These models allow for predictability in interventions to distributions, counterfactual reasoning, and the estimation of causal effects from observational data, addressing the limitations of traditional ML in making predictions under changing conditions [48; 52; 53; 50].

The shift towards causality in ML is not just technical but philosophical, demanding a reconsideration of how we model the world and make decisions. It emphasizes the importance of understanding the underlying processes that generate data, which is essential for tasks ranging from scientific discovery to social policy and ethics [54]. The integration of causality into ML represents a paradigm shift, promising models that are not only more predictive but also more interpretable, stable, and capable of reasoning about interventions [55]. This holistic approach to ML, blending predictive power with causal understanding, opens new horizons for research and application, ensuring that future technologies are equipped to address not only what is happening but also why it happens and what might happen under different circumstances [45; 51]. In this thesis, we focus on such a task for neural data: incorporating causality into the inference of FC from neural time series.

We investigated the utility of DPGM in obtaining CFC from neural data in Biswas and Shlizerman [1]. DPGM, applicable to i.i.d. observations, can model causal relations between whole time series of different neurons in sense of average over time or at a specific time. Thereby, standard DPGM does not explicitly model the inter-temporal causal relations between neural activity in the time series. For inference of the DPGM, the PC algorithm is one of the most popular causal inference algorithms that assumes independent and identically distributed (i.i.d.) sampling of the nodes of the network and absence of latent confounders [52; 48]. The PC algorithm estimates the completed partially directed acyclic graph (CPDAG) which represents the class of all DAGs which are Markov Equivalent and thereby indistinguishable from observational data. In terms of computational time complexity, the PC algorithm is exponential (as a function of the number of nodes) in the worst case, but if the true DAG is sparse, which is often a reasonable assumption, it reduces to polynomial complexity. The question of asymptotic guarantees in estimation of the CPDAG by the PC algorithm has been studied in [52; 56; 57]. In particular, assuming only faithfulness, it has been shown that pointwise consistency of the PC algorithm can be achieved based on i.i.d. observations [52]. Furthermore, high dimensional uniform consistency of the PC algorithm based on i.i.d. samples has been shown to be

achieved with a further Gaussian distribution assumption [57], which popularized the use of PC algorithm in high dimensional settings.

However, in neural time series scenario, causal relations are between neural activity at different times. Assuming independent sampling of nodes of the common DPGM is not suitable as the nodes correspond to time series with samples which are dependent over time. Despite its asymptotic guarantees in estimation from i.i.d. data, in the dependent sampling i.e. time series scenario, the consistency of the PC algorithm is still unclear. Moreover, DPGM typically generate a Directed Acyclic Graph (DAG), while neural activity is often comprised of feedback loops over time [58; 59; 60; 61]. Though adaptations aim to include cycles in the DPGM they have a more complicated output [62]. In this thesis we address these limitations to improve the utility of DPGM for finding CFC in the neural time series scenario and then apply to real neurobiological data. The following section describes the outline of the report.

1.1 Outline

In Chapter 2, we provide a statistical framework for FC and investigate the aspect of causality in the notion of Causal Functional Connectome (CFC) which would answer the aforementioned causal questions [63; 64]. Specifically, we introduce the directed Markov graphical models as a framework for the representation of functional connectome and define the Directed Markov Property – an essential criterion for examining the causality of proposed functional connectomes. The framework that we introduce allows us to delineate the following properties for a statistical description of causal modeling

1. Format of causality.
2. Inclusion of temporal relationships in the model.
3. Generalization of the statistical model.
4. Dependence on parametric equations.
5. Estimation-based vs. hypothesis test-based inference of CFC from recorded data.
6. Inclusion of cycles and self-loops.

7. Incorporation of intervention queries in a counterfactual manner.
8. Ability to recover relationships between neurons when ground-truth dynamic equation of neural activity are given.

We discuss the applicability and the challenges of existing approaches for CFC inference with respect to these statistical properties. In particular, we compare existing approaches for causal functional connectome inference, such as Granger Causality (GC), Dynamic Causal Modeling (DCM) and Directed Probabilistic Graphical Models (DPGM) based on these properties. The comparative study provides a taxonomy of causal properties that existing methods address and an outlook of the properties for future extensions to address. This chapter is joint work with Eli Shlizerman and published in *Frontiers in Systems Neuroscience* [1].

In Chapter 3, we develop a novel approach for modeling and estimating causal functional connectivity by adapting directed probabilistic graphical models to the time series scenario. We introduce the *Time-Aware PC* (TPC) algorithm. It uses the PC algorithm as a starting point and adapts it to the neural time series setup by following processes such as time-delay, bootstrapping and pruning. These ensure that the inferred CFC is robust and well suited to the setup. The proposed CFC graphical model incorporates feedback loops in functional connectivity and is non-parametric, yet we show that the CFC graphical model accurately represents the causal relationships in the unknown dynamical process of neural activity. Furthermore, the proposed CFC graphical model is predictive of the consequence of counterfactual interventions, such as the alteration in the CFC due to ablation or external control of certain neurons. We apply the proposed methodology on neural signals simulated from different paradigms and CFC motifs and demonstrate the utility of TPC in recovering the generating motifs. We also apply the TPC on public benchmark datasets and compare the performance in recovery of the ground truth CFC with other approaches. We further demonstrate the use of TPC to obtain the CFC among sampled neurons in mice brain from electrophysiological neural signals. This chapter is joint work with Eli Shlizerman. It has been presented in its early form at the 29th Annual Computational Neuroscience Meeting (CNS 2020) [65] and its complete version is published at *PLoS Computational Biology: Methods* [2].

In Chapter 4, we show the mathematical guarantees of the TPC algorithm in estimating the causal graph from time series data. For clarity on the mathematical guarantees of the TPC algorithm which builds upon the PC algorithm in time series, it is first essential to have clarity on the mathematical guarantees of the PC

algorithm under temporal dependence. Thereby, we firstly show that under a standard set of assumptions on the underlying time series involving ρ -mixing, the PC algorithm is consistent under temporal dependence. Further, we show that for the popular time series models such as vector auto-regressive moving average and linear processes, consistency of the PC algorithm holds. We then prove the consistency for the Time-Aware PC algorithm. The findings are supported by simulations and benchmark real data analyses with TPC algorithm using different conditional dependence tests and with/without subsampling. This chapter is joint work with Somabha Mukherjee and is published in *Statistics and Computing* [3].

In Chapter 5, we compute the CFC from time series resting-state functional magnetic resonance imaging (rs-fMRI) data recorded for study of Alzheimer's Disease. Functional connectivity between brain regions is known to be altered in Alzheimer's disease and promises to be a biomarker for early diagnosis. Several approaches for functional connectivity obtain an un-directed network representing stochastic associations (correlations) between brain regions. However, association does not necessarily imply causation. In contrast, Causal Functional Connectivity (CFC) is more informative, providing a directed network representing causal relationships between brain regions. In this chapter, we obtained the causal functional connectome for the whole brain from resting-state functional magnetic resonance imaging (rs-fMRI) recordings for subjects from three clinical groups: cognitively normal, mild cognitive impairment, and Alzheimer's disease. We applied the recently developed Time-aware PC (TPC) algorithm to infer the causal functional connectome for the whole brain. TPC supports model-free estimation of whole brain CFC based on directed graphical modeling in a time series setting. We compared the CFC outcome of TPC with that of other related approaches in the literature. Then, we used the CFC outcomes of TPC and applied network-based statistics to identify CFC sub-networks that are altered in subjects with Alzheimer's disease compared to cognitively normal subjects. We also performed an exploratory analysis of the difference in strengths of CFC edges between Alzheimer's and cognitively normal groups, based on edge-wise p-values obtained by Welch's t-test. The obtained brain regions are found to be in agreement with literature on brain regions impacted by Alzheimer's disease, published by researchers from clinical/medical institutions. This chapter is joint work with SuryaNarayana Sripada, and is published in *Frontiers in Computational Neuroscience* [4].

Chapter 2

Functional and Causal Neural Connectomics: A Comparative Study

This chapter is a reproduction of published work in Biswas and Shlizerman [1], with minor adaptations to fit the thesis structure and formatting.

2.1 Neural Connectomics: Anatomical and Functional

Recent advances in neuro-imaging has made it possible to examine brain connectivity at micro and macro scales. These, for example, include electron microscopy reconstruction of the full nervous system with neurons and synapses of *C. elegans* in the mid-1980s [66]. Recent non-invasive diffusion tensor imaging, followed by computational tractography, allow to recover fiber tract pathways in human brain [67; 68; 69]. Also, two-photon tomography facilitates imaging axonal projections in the brain of mice [70]. Although the anatomical reconstruction gives insights into the building blocks and wiring of the brain, how that leads to function remains unresolved. A representative example is *C. elegans*, in which connections and neurons were fairly rapidly mapped and some neurons were associated with functions, it is still unclear what most of the connections “do” [71].

It became increasingly clear that anatomical wiring diagram of the brain could generate hypotheses for testing, but it is far from the resolution of how the anatomical structure relates to function and behavior.

This is because each wiring diagram can encode several functional and behavioral outcomes [13]. In both vertebrate and invertebrate brains, pairs of neurons can influence each other through several parallel pathways consisting of chemical and electrical synapses, where the different pathways can either result in similar or dissimilar behavior. Furthermore, neuromodulators re-configure fast chemical and electrical synapses and have been shown to act as key mediators for various brain states. Given this background, it is generally agreed upon that the anatomical synaptic connectivity does not provide adequate information to predict the physiological output of neural circuits. To understand the latter, one needs to understand the flow of information in the network, and for that, there is no substitute for recording neuronal activity and inferring functional associations from it. This is what the functional connectome aims to achieve.

The functional connectome is the network of information flow between neurons based on their activity and incorporates the dynamic nature of neuronal activity and interactions between them. To obtain neuronal activity and dynamics, the neuronal circuit needs to be monitored and/or manipulated. Recent approaches to record such activity include brain wide two-photon calcium single neuron imaging in vivo of *C. elegans* [72], wide-field calcium imaging of regions of the brain of behaving mice [73], two-photon calcium imaging [74], Neuropixel recordings of single neurons in the brain of behaving mice [75; 76], and functional Magnetic Resonance Imaging (fMRI) recordings of voxels in the human brain as part of the Human Connectome Project [77; 78]. It is noteworthy that FC aims to capture statistical dependencies based on neural recordings and does not rely on the underlying anatomical connectivity, thereby FC methods are applicable for different scales of neural data - micro [79], meso [80] and macro [81]. Yet, when the underlying anatomical connectivity is charted, e.g. macro-scale anatomical connectivity in Diffusion Tensor Imaging [82], it can be used to further constrain the inference of FC [83].

In this section, we set the notions of the neural connectome with regards to anatomical and functional aspects. These notions are from a statistical perspective and will assist to connect connectomics in further sections with causation.

2.1.1 Anatomical Connectome

Let us consider a brain network $V = \{v_1, \dots, v_N\}$ with N neurons labeled as v_1, \dots, v_N . We will denote the edges as $E_a \subset V \times V$ between pairs of neurons that correspond to anatomical connectivity between the

neurons [5]. We will refer to the graph $G = (V, E_a)$ as the anatomical connectome between the neurons in V . Each edge $(v, u) \in E_a$ will be marked by a weight $w_{vu} \in \mathbb{R}$ that quantifies the strength of the anatomical connection from v to u .

Examples

2.1.1) Binary Gap Connectome. If there is a gap junction connection from neuron v to neuron u , the set E_a will include the edges $v \rightarrow u$ and $u \rightarrow v$ and they will be marked with weight 1 [84]. The resulting graph is undirected in the sense that $(v, u) \in E_a$ iff $(u, v) \in E_a$. The resulting weight matrix with entries w_{vu} is symmetric.

2.1.2) Weighted Synaptic Connectome. If there is a synaptic connection from neuron v to neuron u , the set E_a will include the edge $v \rightarrow u$ and it will be marked with a weight which is equal to the number of synaptic connections starting from neuron v to neuron u . The resulting graph is a directed graph in the sense that $(v, u) \in E_a$ does not imply $(u, v) \in E_a$, and the weight matrix (w_{vu}) is asymmetric [85].

2.1.2 Functional Connectome (Associative)

Functional Connectome (FC) is a graph with nodes being the neurons in V and pairwise edges representing a stochastic relationship between the activity of the neurons. Weights of the edges describe the strength of the relationship. Let $X_v(t) \in \mathbb{R}$ denote a random variable measuring the activity of neuron v at time t . Examples for such variables are instantaneous membrane potential, instantaneous firing rate, etc. *Associative FC* (AFC) is an undirected graph, $G = (V, E)$, where $E_{v,u} \in E$, an undirected edge between v and u , represents stochastic association between neuron v and u . Edge weights signify the strength of the associations. Different approaches define stochastic associations leading to different candidates for AFC, as follows.

Pairwise Associative Connectivity

We first describe pairwise stochastic associations (*pairwise AFC*). Let us consider recordings at time points $0, \dots, T$, with activities $X_v = \{X_v(t) : t \in 0, \dots, T\}$ for neuron v . The following measures of pairwise association will correspond to pairwise AFC.

Pearson's Correlation The Pearson's correlation between X_v and X_u is defined as

$$r(X_v, X_u) = \frac{\sum_{t=0}^T (X_v(t) - \bar{X}_v)(X_u(t) - \bar{X}_u)}{\sqrt{\sum_{t=0}^T (X_v(t) - \bar{X}_v)^2 \sum_{t=0}^T (X_u(t) - \bar{X}_u)^2}}$$

where $\bar{X}_v = \frac{1}{T+1} \sum_{t=0}^T X_v(t)$, for $v \in V$. Pearson's correlation takes a value between -1 and 1 and the further its value is from 0, the larger is the degree of pairwise association. Neurons v and u are connected by pairwise AFC with respect to Pearson's correlation if $r(X_v, X_u)$ is greater than a threshold in absolute value and the value $r(X_v, X_u)$ would be the weight of the connection, i.e. $E_{v,u} = \text{thresh}(r(X_v, X_u))$. Correlation coefficient and the AFC based on it are sensitive to indirect third party effects such as an intermediary neuron, poly-synaptic influences, indirect influences, and noise.

Partial Correlation Partial correlation is an additional measure of pairwise stochastic association between random variables defined as follows. Let the covariance between X_v and X_u be $C_{vu} = \frac{1}{T+1} \sum_{t=0}^T (X_v(t) - \bar{X}_v)(X_u(t) - \bar{X}_u)$. The matrix of covariances, $\Sigma = (C_{vu})_{1 \leq v, u \leq N}$ is called the covariance matrix for the variables X_1, \dots, X_N . Let the v, u -th entry of the inverse covariance matrix Σ^{-1} be γ_{vu} . The partial correlation between X_v and X_u is defined as

$$\rho(X_v, X_u) = -\frac{\gamma_{vu}}{\sqrt{\gamma_{vv}\gamma_{uu}}}$$

Partial correlation rectifies the problem of correlation coefficient being sensitive to third party effects since it estimates the correlation between two nodes after removing the variance shared with other system elements. Partial correlation takes a value between -1 and 1 and the further its value is from 0, the larger is the degree of pairwise association. Neurons v and u are connected by pairwise AFC with respect to partial correlation, if $\rho(X_v, X_u)$ is greater than a threshold in absolute value and the value $\rho(X_v, X_u)$ would be the weight of the connection, i.e. $E_{v,u} = \text{thresh}(\rho(X_v, X_u))$.

Undirected Probabilistic Graphical Model

Undirected Probabilistic Graphical Models (PGM) allow for modeling and inferring stochastic associations while considering multi-nodal interactions beyond pairwise manner through a graphical model. Let $G =$

(V, E) be an undirected graph with neuron labels $V = \{v_1, \dots, v_N\}$ and edges E (Figure 2.1). Let Y_v denote a scalar-valued random variable corresponding to $v \in V$, for example, Y_v can be the value of X_v at recording time t : $Y_v = X_v(t)$, or average of recordings over time: $Y_v = \bar{X}_v$, etc. For a set of neurons $A \subset V$, let \mathbf{Y}_A denote the random vector $(Y_v, v \in A)$. The random vector \mathbf{Y}_V , is said to satisfy the *undirected graphical model* with graph G , if, Y_v is conditionally independent of Y_u given $\mathbf{Y}_{V \setminus \{v,u\}}$ for $(v, u) \notin E$, denoted as

$$Y_v \perp\!\!\!\perp Y_u | \mathbf{Y}_{V \setminus \{v,u\}} \text{ for } (v, u) \notin E. \quad (2.1)$$

When Eq. (2.1) holds, Y_v is said to satisfy the *Markov property* with the undirected graph G [86]. The Markov property with undirected graph G translates each *absent edge* between a pair of nodes v and u into *conditional independence* of Y_v and Y_u given all other nodes. In other words, nodes in the undirected PGM that are not connected by an edge appear as non-interacting when all the nodes are examined.

Definition 1. The AFC for neurons in V is the undirected graph $G = (V, E)$ such that $(Y_v, v \in V)$ satisfies the *Markov Property* with respect to G . When $Y_v = X_v(t)$, the associative FC is contemporaneous at each time t .

According to this definition, the graph edges $E_{v,u}$ are well-defined. With respect to their weights, there are no unique candidates. When $(Y_v, v \in V)$ follows a multivariate Gaussian distribution, an edge is present if and only if the partial correlation between them is non-zero. Thereby, in practice, partial correlation $\rho(X_v, X_u)$ is typically used for weight of edge $E_{v,u} \in E$ (Figure 2.1).

In the above we have defined AFC to be the undirected graphical model that has the Markov Property. A natural question that arises is what kind of probability distributions for the graphical model allow it to follow the Markov Property? The Factorization Theorem by Hammersley, Clifford and Besag prescribes conditions on the probability distributions of $Y_v, v \in V$ for which undirected graphical model has the Markov property [87].

Theorem 2.1.1 (Factorization Theorem). If $Y_v, v \in V$ has a positive and continuous density f with respect to the Lebesgue measure or is discrete with positive joint probabilities, then it satisfies the Markov property (Eq. 2.1) with respect to $G = (V, E)$ if and only if the distribution of $Y_v, v \in V$ *factorizes* with respect to G ,

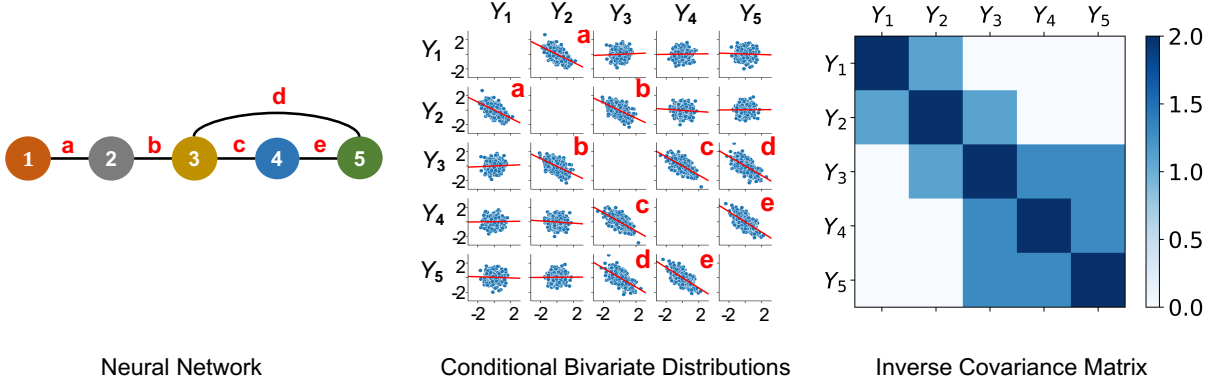


Figure 2.1 An undirected PGM and its Markov Property. Left: Network of 5 neurons that is defined in Example 2.1.1. Neurons are labeled as 1 – 5 and edges are labeled as a-e. Middle: Y_1, \dots, Y_5 following centered multivariate Gaussian distribution with entries of inverse covariance matrix γ_{ij} such that $\gamma_{13} = \gamma_{14} = \gamma_{15} = \gamma_{24} = \gamma_{25} = 0$, when it factorizes with respect to the Neural Network. This plot shows conditional bivariate distributions given other variables $\in (-0.2, 0.2)$ and demonstrates Eq. (2.1) where Y_v and Y_u are not correlated conditional on other variables (seen by nearly flat red trend-lines) for (v, u) not an edge of the Neural Network. This indicates that Y_1, \dots, Y_5 satisfies the Markov Property with the Neural Network. Right: Due to a Gaussian distribution, the non-zero entries of the Inverse Covariance Matrix of Y_1, \dots, Y_5 correspond to the edges of the Neural Network.

which means,

$$f(y) = \prod_{C \subset G: C \text{ is complete}} \phi_C(y_C), y \in \mathbb{R}^V \quad (2.2)$$

where, f is the density of $Y_v, v \in V$, ϕ_C is an arbitrary function with domain \mathbb{R}^C , and y_C is the sub-vector $(y_v : v \in C)$ and $C \subset G$ is complete, i.e. $E_{v,u} \in E$ for all $v \neq u \in C$, are connected by an edge.

Under the multivariate Gaussian assumption, Theorem 2.1.1 yields a simple prescription for obtaining the undirected graph with the Markov Property. When $Y_v, v \in V$ are distributed as multivariate Gaussian with positive definite covariance matrix Σ , then G is determined by the zeroes of the inverse covariance matrix Σ^{-1} , i.e., $E_{v,w} \in E$ if and only if $\Sigma_{ij}^{-1} \neq 0$. This is illustrated in Example 2.1.1 and Figure 2.1. This has been used for inferring undirected PGMs in several applications [88; 89]. In such a case, estimation of G is tantamount to estimation of Σ^{-1} . Methods for estimation of Σ^{-1} include Maximum Likelihood Estimation (MLE) which provides non-sparse estimates of Σ^{-1} [90] and penalized MLE, e.g. by Graphical Lasso, which provides sparse estimates of Σ^{-1} , with the estimates being statistically consistent under assumptions [91; 92; 93; 94].

Example 2.1.1 (Markov Property and Multivariate Gaussian Assumption). Let us consider an example with 5 neurons. Suppose $(X_1(t), \dots, X_5(t))$ follow a centered multivariate Gaussian distribution with positive

definite covariance matrix Σ and independent copies over time t . Let $\Sigma^{-1} = (\gamma_{ij})_{1 \leq i, j \leq N}$ be the inverse covariance matrix. The probability density of $(X_1(t), \dots, X_5(t))$ is given by

$$f(x^t) \propto \exp \left(-\frac{1}{2} \sum_{i,j=1}^5 \gamma_{ij} x_i^t x_j^t \right), x^t \in \mathbb{R}^5. \quad (2.3)$$

The neural network graph is illustrated in Figure 2.1-left and we note that the graph has the following complete subsets

$$\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{3, 5\}, \{3, 4, 5\}.$$

Thereby, for each t , it is observed that $\gamma_{13} = \gamma_{14} = \gamma_{15} = \gamma_{24} = \gamma_{25} = 0$ iff the density in Eq. (2.3) factorizes as

$$f(x^t) \propto \exp \left(-\frac{1}{2} (\gamma_{11} x_1^t{}^2 + \gamma_{12} x_1^t x_2^t + \gamma_{22} x_2^t{}^2) \right) \exp \left(-\frac{1}{2} \gamma_{23} x_2^t x_3^t \right) \exp \left(-\frac{1}{2} (\gamma_{33} x_3^t{}^2 + \gamma_{34} x_3^t x_4^t + \gamma_{44} x_4^t{}^2) \right) \\ \exp \left(-\frac{1}{2} (\gamma_{45} x_4^t x_5^t + \gamma_{55} x_5^t{}^2) \right) \exp \left(-\frac{1}{2} \gamma_{35} x_3^t x_5^t \right)$$

That is, the density in Eq. (2.3) factorizes according to Eq. (2.2) with respect to the graph. Hence, when $\gamma_{13} = \gamma_{14} = \gamma_{15} = \gamma_{24} = \gamma_{25} = 0$, according to Theorem 2.1.1 it follows that $(X_1(t), \dots, X_5(t))$ satisfies the Markov Property with respect to the graph in Figure 2.1-left.

2.2 From Association to Causation

In Section 2.1.2 we provided a systematic exposition of AFC, however, the ultimate phenomenon of functional connectivity is to capture the causal interaction between neural entities [63; 64; 95; 96]. Indeed, FC research is already targeting addition of causal statements to associative FC by identifying correlated brain regions as causal entities, suggesting the need to incorporate functional connectivity beyond association [97; 98; 99; 100]. From associations alone, causal inference is ambiguous with possible parallel, bidirectional and spurious causal pathways. Narrowing the space of causal pathways inferred from brain signals can significantly progress FC towards its aims of finding causal neural interactions (See Figure 2.2). For example, in a neural

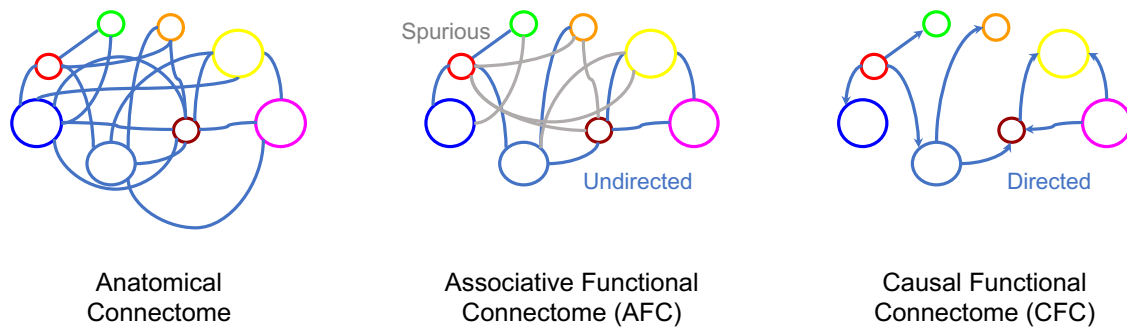


Figure 2.2 Various connectome representations for the same network of neurons. Left to right: Anatomical connectivity between brain regions, an undirected graph representing AFC where gray indicates spurious edges, and a directed graph with directions representing causal relationships.

circuit, an edge between neuron A and neuron B could mean either neuron A influences neuron B or vice versa. This directionality of influence is unclear from the association. In the instance when neuron C influences activity of both A and B neurons, while they do not influence each other, a spurious association would be found between A and B in the AFC (See example in Figure 2.2). Recent progress in causal inference in statistics literature has made it possible to define causality and make causal inferences. Given this background, we review causal modeling in statistics to aid the development of a novel framework for *Causal Functional Connectivity* (CFC) from neural dynamics.

2.2.1 Causal Modeling

Causation is a relation between two events, in which one event is the cause and the other is the effect. In the context of neuroscience, causality is a major factor. For example, in fear perception in the human brain, the causal relationships among the activity of retina, lateral geniculate nucleus (LGN), visual cortex (VC), superior colliculus (SC), pulvinar (P), central nucleus of the amygdala (CeA), paraventricular nucleus (PVN) and hypothalamus-pituitary-adrenal axis (HPA) can be considered. While the information flows from the Retina to HPA, there is no direct link between them. Fear stimulus of the retina causes triggering of the HPA (response region) mediated by the activity of intermediate brain nuclei in a specific sequence of activation through two merging pathways [101; 102; 103].

A causal model relates the cause and effect, rather than recording correlation in the data and allows the investigator to answer a variety of queries such as *associational queries* (e.g. having observed activity in LGN, what activity can we expect in CeA?), *abductive queries* (e.g. what are highly plausible explanations for




Figure 2.3 Causal graphs for electrical activity of brain nuclei for fear perception from visual stimulus. Left: Literature describes two routes for fear stimulus propagation from Retina to HPA: Retina → LGN → VC → CeA → PVN → HPA (cortical route) and Retina → SC → P → CeA → PVN → HPA (subcortical route). Right: It is demonstrated that even if there is intervention by ablation or lesion in the striate cortex of VC, thereby blindness, yet fear response to visual stimuli ("blindsight") is yielded through the subcortical route [104; 101].

active CeA during fear stimulus?), and *interventional queries* (e.g. what will happen to the causal pathways if there is an ablation of VC?). Often interventional queries are especially of interest in neural connectomics, with interventions including neural ablation (see Figure 2.3b) and external neuromodulation [13; 105]. In such cases, causal modeling aims to correctly predict the effect of an intervention in a counterfactual manner, that is, without necessarily performing the intervention but from observational data [48].

Representing Causal Relations with a Directed Graph

Directed graphs provide convenient means for expressing causal relations among the variables. The vertices correspond to variables, and directed edges between the vertices represent a causal relationship that holds between pairs of variables. Formally, a graph consists of a set V of vertices (nodes) and a set $E \subset V \times V$ of edges that connect some pairs of vertices. Each edge can be either directed or undirected. In a directed graph, all edges are directed. A directed acyclic graph (DAG) is a directed graph without cycles. The directed graph representing causal relationships is called a causal graph. Figure 2.3 is a causal graph among eight variables

representing electrical activity in eight brain regions. If G is a causal graph and there is a directed path from nodes a to b it implies that, the variable corresponding to node a is a cause for the variable corresponding to node b . For example, the electrical activity in Retina (node 1) is a common cause for electrical activity in LGN (node 2) and SC (node 3).

If there is any intervention of one of the variables such as ablation or neuromodulation of nodes, the network topology can be adapted with minor changes. In Figure 2.3 (right), to represent an ablation of visual cortex (VC), one would delete from the network all links incident to and originating from the node VC. To represent control of activity of VC by neuromodulation, one would delete all edges only incident to VC as then VC is not causally influenced by its parent region's activity but by external neuromodulation.

Statistical Properties of CFC Modeling

In the following we outline several statistical properties that are relevant in the context of causal modeling of functional connectivity.

1. *Format of causality.* This specifies how causation is defined in the model with respect to parameters or properties satisfied by the model.
2. *Inclusion of temporal relationships in the model.* Since the activity of neurons are related over time, this condition specifies whether such temporal relationships are incorporated in defining causal relationships among neural activity in the CFC model.
3. *Generalization of the statistical model.* This condition specifies model restrictions, such as linear or non-linear modeling, and informs whether such restrictions can be generalized.
4. *Parametric or non-parametric model.* This specifies whether the model is parametric (i.e. consisting of a parametric equation and needing estimation of parameters) or non-parametric (i.e. free of a parametric equation) [106]. Non-parametric models have the advantage of not requiring assumptions on specific dynamical equations for the neural dynamics.
5. *Estimation-based vs. hypothesis test-based inference of CFC from recorded data.* Approaches for inferring CFC either support estimation of the CFC from the data or test of significance of hypothetical

CFC models based on data. This condition specifies which category among these does the CFC model belong to.

6. *Inclusion of cycles and self-loops in the model.* Neural activity often consist of feedback loops and cycles [107]. This condition specifies whether such cycles and self-loops are represented in the CFC model.
7. *Incorporation of intervention queries in a counterfactual manner.* This condition specifies whether interventional queries are answered directly by the CFC model from observational data without performing the experimental intervention.
8. *Ability to recover relationships between neurons when ground-truth dynamic equation of neural activity are given.* It is often desirable that causal relationships between neural activity in ground truth dynamical equations are accurately represented by the inferred CFC [108; 109]. This condition sheds light into the performance of the CFC approach to recover such ground truth relationships from dynamical equations.

We proceed to delineate causation from statistical associations by surveying existing approaches, and describe relation of each approach to above statistical properties [110; 111; 48].

2.2.2 Granger Causality

Granger causality (GC) is a statistical methodology for determining whether one time series is useful in predicting another [112; 113]. Denoting $X_v(t)$ to be the state of neuron v at time t , the main idea of GC is that, X_j “Granger-causes” X_i if past of X_j contains information that helps predict the future of X_i better than using the information in the past of X_i or past of other conditioning variables X_k [114]. There are several variations of Granger Causality (also called Wiener-Granger Causality), based on linear vector auto-regressive model [115], nonlinear vector auto-regressive model [116], and non-parametric approaches [117]; conditional and pairwise approaches [23]. More formally, typical approach considers a linear Gaussian vector auto-regressive (VAR) linear model between the variables, in this case, states of neurons X_i [115],

$$X_i(t) = \sum_{j=1}^N \sum_{k=1}^K A_{ji}(k) X_j(t-k) + \epsilon_i(t)$$

where K is the maximum number of lags (model order) and $A_{ji}(K)$ are real-valued linear regression coefficients, and $\epsilon_i(t) \sim N(0, \sigma^2 I)$. Neuron j is said to *Granger-cause* neuron i if at least one of the coefficients $A_{ji}(k) \neq 0$ for $k = 1, \dots, K$. In practice, $A_{ji}(k)$ are estimated by minimizing the squared prediction error or by maximizing the likelihood or sparsity-inducing penalized likelihood [118; 113]. Granger Causality has been applied towards inference of CFC in the linear Gaussian VAR model setting [119; 108; 118]. Extensions of GC to categorical random variables, non-linear auto-regressive models and non-parametric models exist [120; 117; 121; 116].

GC was first introduced within econometrics and later has been used to find directed functional connectivity in electrophysiological studies [112; 122], in EEG or MEG datasets, either at source or sensor level [123; 124; 125; 126]. The slow dynamics and regional variability of the haemodynamic response to underlying neuronal activity in fMRI were shown to be able to confound the temporal precedence assumptions of GC [127; 128; 129; 130].

While Granger Causality provides a powerful tool for understanding which neural time series have a key role in predicting the future of other neural time series [131; 132; 133], studies express concern since prediction is not a formal setting to answer causal questions related to the consequence of interventions and counterfactuals [134; 135; 136]. Furthermore, in practice, GC uses a model assumption between the variables, e.g. a linear Gaussian VAR model, and results could differ when this assumption does not hold [115]. Notwithstanding the limitations, GC has been a well-known method in the neural time series scenarios, and applications [119; 137]. GC based on linear VAR model is equivalent to Transfer Entropy for Gaussian variables, while the latter is a non-linear method in its general formulation [138]. Transfer Entropy has been explored as a tool to explore connectomics at different scales [139], and also applied to the retina circuit [140].

2.2.3 Dynamic Causal Model

The Dynamic Causal Model (DCM) is an approach for modeling brain activity and causal interaction between brain regions (See Figure 2.4). DCM was first introduced for fMRI time series and treats the brain as a deterministic non-linear dynamical system network in which neural activity is considered to be unobserved and propagates in an input-state-output system [141]. The system is attached to a forward model that maps

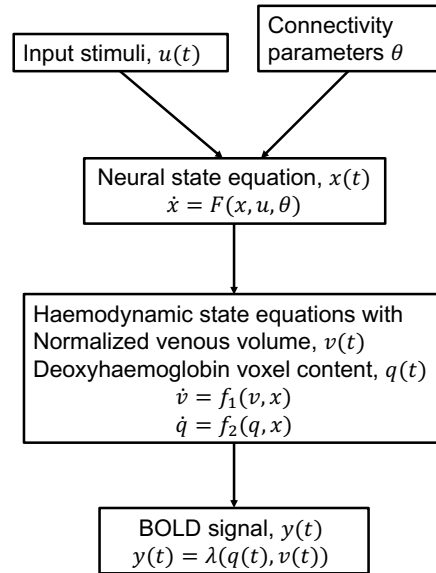


Figure 2.4 Schematic of the haemodynamic model used by DCM for fMRI. Input stimuli $u(t)$ lead to neural state $x(t)$ subject to connectivity parameters θ modeled by the neural state equation. Neuronal activity leads to increase in blood flow that changes in venous volume $v(t)$ and deoxyhaemoglobin content $q(t)$ modeled by the haemodynamic state equations. These haemodynamic states are fed to a forward non-linear model λ which results in the Blood Oxygenation Level Dependent (BOLD) fMRI response [142].

the neural activity to observed Blood Oxygenation Level Dependent (BOLD) fMRI signal [142; 143; 144]. The neural and observational models are specified by a particular parametric form that depends on the data and imaging modality and they together form a fully generative model [114]. DCM outputs evidence for different neural and/or observational models and posterior parameter estimates optimized by variational Bayesian techniques [145]. The coupling parameters between hidden states for different brain regions, in the differential equations that specify the model, constitute the CFC. The DCM model incorporates interactions due to the experimental equipment and changes due to external perturbations [146].

DCM compares the evidence for different hypothesized models, and thereby is a tool for testing hypothesis for model and experimental design. DCMs are not restricted to linear systems and include intricate models with large number of parameters for neural dynamics and experimental perturbations. The model is typically constructed to be biologically plausible. Constraints are exercised to the model through priors. The priors are useful to specify which connections are believed to be more likely. A Bayesian likelihood with the prior distribution is maximized to obtain the parameter estimates. The approach relies on precise modeling and aims to specify a biologically plausible detailed mechanistic model of the neuronal dynamics and imaging

equipment perturbation [147]. While DCM was originally formulated to be deterministic, recent advances can include stochastic fluctuations in the neural activity as well [148; 149]. The DCM framework has also been extended beyond fMRI and established in the magneto/encephalography domain [150], and in local field potentials [151].

2.2.4 Directed Probabilistic Graphical Models

Directed Probabilistic Graphical Models (DPGMs) provide a probabilistic foundation to causality in a manner that answers causal queries through a generalized model without requiring specific modeling of the neural dynamics. An important aspect to take into account in inference of CFC is stochasticity. Neural signals are stochastic due to noise and intrinsic nature of neuron firing [152; 153]. The variability and noise in neural dynamics is well known to challenge the determination of neural phenomenon, e.g. the detection of onset of epileptic seizures [154; 155]. So, when we say “spike in activity of neuron A is a cause of the spike in activity of neuron B”, the cause makes the effect more likely and not certain due to other factors like incoming inhibitory projections from other neurons and extracellular ion concentration [156; 157]. Moreover, additional arbitrary functional relationships between each cause and effect could exist such that these introduce arbitrary disturbances following some undisclosed probability distribution. For example, it is widely known that diet and stress in humans change the levels of neurotransmitters in the brain [158; 159]. Thereby, the strength of causal relationships between neurons can be perturbed by daily variability in diet and/or stress. Also, uncertainties in effect can occur from unobserved causes which is especially true in the context of neural signals as extraneous variables such as diet and stress are often not observed or due to recording from only a fraction of the units in the brain [160]. With these considerations, it is elaborated that values of exogenous variables do not imply values for the remaining variables in a deterministic manner. This motivates the need to consider a probabilistic foundation to causation and causal graphs provided by DPGM. We outline three conditions of DPGM that connect probabilities with causal graphs: The Directed Markov Property, the Causal Minimality Condition, and the Faithfulness Condition.

Let $G = (V, E)$ be a DAG over neuron labels $V = (v_1, \dots, v_N)$ with directed edges E (e.g., Figure 2.3). DPGM typically considers the graph to be a DAG because of coherent probability semantics with a DAG and challenges with directed cycles, while there could be more general extensions [27; 28]. Nodes v




Figure 2.5 Directed Markov Property in the context of fear stimulus. The DAG in the example of Figure 2.3 is annotated to illustrate the Directed Markov Property (Eq. 2.4). PVN is selected as node v , its random variable is hence Y_v (red). The parents of v denoted as $pa_G(v)$, and corresponding random variables are $Y_{pa_G(v)}$ (green). The non-descendants of v before parents, denoted as $nd_G(v) \setminus pa_G(v)$, and corresponding random variables $Y_{nd_G(v) \setminus pa_G(v)}$ (blue). Directed Markov Property holds with the true causal edges (black), as for them parents and children are functionally related (Eq. 2.6). Causal Minimality Condition ensures that potential edges (blue dotted) between $nd_G(v) \setminus pa_G(v)$ nodes and v are absent from the DAG.

and $u \in V$ are said to be *adjacent* if $v \rightarrow u \in E$ or $u \rightarrow v \in E$. A *path* is a sequence of distinct nodes in which successive nodes are adjacent. If $\pi = (v_0, \dots, v_k)$ is a path then v_0 and v_k are the end-points of the path. If every edge of π is of the form $v_{i-1} \rightarrow v_i$ then v_0 is an *ancestor* of v_k and v_k is a *descendant* of v_0 . We use the convention that v is an ancestor and descendant of itself. The set of *non-descendants* of v , denoted $nd_G(v)$, contains nodes $u \in V$ that are not descendants of v . The set of *parents* of $v \in V$ is denoted as $pa_G(v) = \{u \in V : u \rightarrow v \in E\}$. We mark the set $nd_G(v) \setminus pa_G(v)$ as the set that contains all nodes which are older ancestors of v before its parents (Figure 2.5). Let Y_v denote a scalar-valued random variable corresponding to $v \in V$, e.g., the neural recording at time t : $Y_v = X_v(t)$, average of recordings over time $Y_v = \bar{X}(v)$, and for a set of neurons $A \subset V$, \mathbf{Y}_A denotes the random vector $(Y_v, v \in A)$. With these notations, we outline the three conditions of DPGM.

1) Directed Markov Property $(Y_v, v \in V)$ is said to satisfy the *Directed Markov Property* with respect to the DAG G if and only if,

$$Y_v \perp\!\!\!\perp \mathbf{Y}_{nd_G(v) \setminus pa_G(v)} \mid \mathbf{Y}_{pa_G(v)} \quad (2.4)$$

for every $v \in V$. The Directed Markov Property translates the edges in the DAG into conditional independencies, such that each node Y_v and its older ancestors $\mathbf{Y}_{nd_G(v) \setminus pa_G(v)}$ are conditionally independent given its parents $\mathbf{Y}_{pa_G(v)}$. In other words, the influence of each node's ancestors beyond parents reaches to the node exclusively via its parents. In this way, the Directed Markov Property connects probabilistic conditional independencies with relationships of causal influence between nodes of a directed graph. For example, under the Directed Markov Property, in Figure 2.5, the assertion that the activity of VC and P are conditionally independent of the activity of PVN, given the activity of CeA at time t corresponds to the causal relationship that the influence of the activity of VC and P on the activity of PVN is mediated by the activity of CeA, represented in the DAG as CeA a parent node of PVN, and VC and P are non-descendant nodes beyond parents of CeA.

The Directed Markov Property for DPGM (DMP) is different from the Markov Property for undirected PGM (MP) in that DMP relates conditional independencies between random variables in a directed graph to causal relationships in the directed graph, whereas MP relates conditional independencies between random variables in an undirected graph to edges of association in the undirected graph. Yet, both incorporate multi-nodal interactions in the graphs beyond a pairwise manner. Furthermore, we had seen in Theorem 2.1.1 that MP yields a factorization of the probability density of the random variables comprising the PGM. The DMP also yields a factorization of the joint probability density for DPGM in an adapted manner as follows [161].

Theorem 2.2.1 (Factorization Theorem for DPGM). For $(Y_v : v \in V)$ real random variables with density f with respect to a product measure, it satisfies the Directed Markov Property (Eq. 2.4) with respect to the DAG G if and only if their distribution factorizes according to the G , which means,

$$f(y) = \prod_{v \in V} f(y_v | y_{pa_G(v)}), y \in \mathbb{R}^V \quad (2.5)$$

where f is the density of Y_v , and $f(y_v | y_{pa_G(v)})$ are conditional probability densities.

The Directed Markov Property can be equivalently represented with functional relationships between parents and child instead of conditional independencies, which is described in the following theorem [162].

Theorem 2.2.2 (Functional Equivalence of DPGM). If Y_v satisfies

$$Y_v = g_v(Y_{pa_{\tilde{G}}(v)}, \epsilon_v), v \in V \quad (2.6)$$

where ϵ_v are independent random variables and g_v are measurable functions for $v \in V$ and \tilde{G} is a DAG with vertices V , then $Y_v, v \in V$ satisfies the Directed Markov Property with respect to \tilde{G} . Conversely, if $Y_v, v \in V$ satisfies the Directed Markov Property with respect to a DAG \tilde{G} , then there are independent random variables ϵ_v and measurable functions g_v for which (Eq. 2.6) holds.

Since the functional equations admit a natural causal interpretation, so do DPGMs satisfying the Directed Markov Property [87]. The variables in $Y_{pa_{\tilde{G}}(v)}$ are direct causes of Y_v , meaning that changes in $Y_{pa_{\tilde{G}}(v)}$ lead to changes in distribution Y_v , but not necessarily the other way around. Furthermore, when $Y_v, v \in V$ satisfies the Directed Markov Property, then (Eq. 2.6) holds for some choice of functions $\{g_v\}$ and error distributions $\{\epsilon_v\}$, which implies causal relationships among $Y_v, v \in V$.

DPGMs can predict the consequence of a counterfactual intervention on the random variables [48]. Using Theorem 2.2.2 we show in the following that we only need to remove the edges pointing to the intervened random variables in the DPGM to incorporate the impact of a counterfactual intervention. More precisely, if before the intervention, G is DPGM and $Y_v, v \in V$ satisfies the DMP with respect to G , an intervention to the random variables will modify Eq. (2.6) only those variables that are impacted by the intervention. For example, let us consider the intervention forcing Y_{v_0} to take the value 0 or 1 regardless of value at other nodes. This intervention will change (Eq. 2.6) by excluding the equations in which Y_{v_0} is a function of other nodes. This corresponds to replacing $pa_G(v_0)$ by an empty set in equation (Eq. 2.6), and in other words, removing the edges pointing to node v_0 in G . That is, after the intervention, (Eq. 2.6) holds with a different graph G' that is obtained by removing the edges incident upon the intervened nodes in G . Equivalently, by Theorem 2.2.2, after the intervention $Y_v, v \in V$ satisfies DMP with respect to G' , and thus G' is the DPGM after the intervention.

2) Causal Minimality Condition Let $(Y_v, v \in V)$ satisfy the DMP with respect to the DAG G . G satisfies the *Causal Minimality Condition* if and only if for every proper subgraph H of G with vertex set V , $(Y_v, v \in V)$ does not satisfy the DMP with respect to H . In other words, if adding any edge on to G can also




Figure 2.6 The PC algorithm. Steps of the PC algorithm to infer the DPGM from observed data are summarized by five diagrams (left to right). Data for variables $Y_1 - Y_7$ is visualized in the context of neural recordings. Graph with nodes 1 – 7 corresponding to variables $Y_1 - Y_7$ has no edges. Then, an edge introduced between Y_i and Y_j if they are not independent or conditionally independent given any other variable(s) determined by statistical tests, which results in the undirected Skeleton Graph. Using rules of converting undirected to directed edges as depicted in the figure - Collider Detection rule, Known Non-Colliders rule and Cycle Avoidance rule, as depicted in Figure 6, the skeleton graph is converted to a DAG.

satisfy the DMP, we do not add such an edge, and consider the minimal graph G with respect to which DMP is satisfied to be the DPGM of $Y_v, v \in V$. In principle, since the complete set of causes is unknown, there can be multiple graphs that would fit a given distribution of random variables for DMP, each connecting the observed variables through different causal relationships. Among all those possible graphs satisfying DMP, the causal minimality condition considers the simplest one and ensures a unique causal DPGM.

3) Faithfulness Condition The Directed Markov Property with respect to a DAG G prescribes a set of conditional independence relations on the random variables comprising the graph. However, in general, a probability distribution P of random variables in DAG G that has the independence relations given by the DMP may also include other independence relations. If that does not occur such that all the conditional independence relations by the probability distribution P are encompassed by G , then we denote P and G as *faithful* to one another.

Inference of DPGM

Several methods have been developed for inferring DPGM with the Directed Markov Property, Causal Minimality and Faithfulness conditions for stationary observed data. These include the PC algorithm

(constraint-based, acyclic graph, no latent confounders, no selection bias) [52], FCI algorithm (constraint-based, acyclic graph, latent confounders, selection bias) [163], GES (score-based equivalence class search) [164] and GIES (score-based equivalence class search from data with multiple interventions) [165].

For example, we describe here the PC algorithm which is a popular statistical tool to infer causal connections between stationary random variables under independent and identically distributed sampling. It is a constraint-based method in which a consistent statistical test for conditional independence is used to select the connectivity graph among the random variables of interest. For Gaussian random variables and linear relationships, a standard choice for such a conditional independence test is constructed using the Fisher’s Z-transform [57]. For non-Gaussian variables and non-linear relationships, kernel and distance based tests of conditional independence are used (for e.g., the Kernel PC algorithm [166]). The algorithm first represents the observed variables by nodes of a graph and starts with an empty set of edges and decides whether to put an undirected edge between each pair of nodes, e.g. node 2 and 5 in Figure 2.6. In order to determine whether to have an edge between the pair of nodes, it performs consistent statistical tests for independence of the random variables for the pair of nodes or conditional independence given the random variables for other node(s). If any of the tests finds evidence for independence or conditional independence, an edge is drawn between the pair of nodes and otherwise no edge is drawn between the pair of nodes. The process is followed for each pair of nodes to result in an undirected graph, called the skeleton graph. Using rules such as the Collider Detection rule, Known Non-Colliders rule and Cycle Avoidance rule, the undirected edges are directed to convert the skeleton graph into a DAG. Figure 2.6 provides a schematic of the algorithm with the context of neural recordings. The PC algorithm is shown to be consistent for finding the true causal graph in the absence of latent confounders [52].

2.2.5 Comparative Study of Approaches to Causal Functional Connectivity

We compare the performance of exemplary approaches of CFC inference discussed above to recover relationships in ground truth dynamical equations by generating synthetic data from three simulation paradigms and estimate their CFC using the methods of GC and DPGM (see Figure 2.7). The simulation paradigms correspond to specific model assumptions to assess the impact of model assumptions on the performance of the approaches.

1. Linear Gaussian Time Series (Figure 2.7 top-row). Let $N(0, 1)$ denote a standard Normal random variable. We define $X_v(t)$ as a linear Gaussian time series for $v = 1, \dots, 4$ whose true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Let $X_v(0) = N(0, 1)$ for $v = 1, \dots, 4$ and for $t = 1, 2, \dots, 10000$,

$$\begin{aligned} X_1(t) &= 1 + N(0, 1), & X_2(t) &= -1 + N(0, 1), \\ X_3(t+1) &= 2X_1(t) + X_2(t) + N(0, 1), & X_4(t+1) &= 2X_3(t) + N(0, 1) \end{aligned}$$

2. Non-linear Non-Gaussian Time Series (Figure 2.7 middle-row). Let $U(0, 1)$ denote a *Uniformly* distributed random variable on the interval $(0, 1)$. We define $X_v(t)$ as a non-linear non-Gaussian time series for $v = 1, \dots, 4$ whose true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Let $X_v(0) = U(0, 1)$ for $v = 1, \dots, 4$ and for $t = 1, 2, \dots, 10000$,

$$\begin{aligned} X_1(t) &= U(0, 1), & X_2(t) &= U(0, 1), \\ X_3(t+1) &= 4 \sin(X_1(t)) + 3 \cos(X_2(t)) + U(0, 1), & X_4(t+1) &= 2 \sin(X_3(t)) + U(0, 1) \end{aligned}$$

3. Continuous Time Recurrent Neural Network (CTRNN) (Figure 2.7 bottom-row). We simulate neural dynamics by Continuous Time Recurrent Neural Networks, Eq. (2.7). $u_j(t)$ is the instantaneous firing rate at time t for a post-synaptic neuron j , w_{ij} is the linear coefficient to pre-synaptic neuron i 's input on the post-synaptic neuron j , $I_j(t)$ is the input current on neuron j at time t , τ_j is the time constant of the post-synaptic neuron j , with i, j being indices for neurons with m being the total number of neurons. Such a model is typically used to simulate neurons as firing rate units

$$\tau_j \frac{du_j(t)}{dt} = -u_j(t) + \sum_{i=1}^m w_{ij} \sigma(u_i(t)) + I_j(t), j = 1, \dots, m. \quad (2.7)$$

We consider a motif consisting of 4 neurons with $w_{13} = w_{23} = w_{34} = 10$ and $w_{ij} = 0$ otherwise. We also note that in Eq. 2.7, activity of each neuron $u_j(t)$ depends on its own past. Therefore, the true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4, 1 \rightarrow 1, 2 \rightarrow 2, 3 \rightarrow 3, 4 \rightarrow 4$. The time constant τ_i is set to 10 msec for each neuron i . We consider $I_i(t)$ to be distributed as independent Gaussian process with the mean of 1 and the standard deviation of 1. The signals are sampled at a time gap of $e \approx 2.72$ msec for a total duration of 10




Figure 2.7 Comparative study of CFC inference. CFC inference of GC, DPGM 1 and DPGM 2 methods is compared on three examples of motifs and simulation paradigms; from top to bottom: Linear Gaussian, Non-linear Non-Gaussian, CTRNN. Table: 4 neurons motifs that define the Ground Truth CFC (left) are depicted side by side with inferred CFC over several simulation instances according to the three different methods (right). An edge $v \rightarrow w$ in each inferred CFC corresponds to an edge detected in any of the inference instances. The percentage (blue) next to each edge indicates the number of times out of all instances that the edge was detected. Right: For each motif and simulation paradigm, Sensitivity (green) and Accuracy (orange) of each method is shown.

secs.

For these network motifs we compare the methods GC, DPGM 1 and DPGM 2. We compute the GC graph using the *Nitime* Python library which fits an MVAR model followed by computation of the Granger Causality by the *GrangerAnalyzer* [167].

We compute DPGM using the PC algorithm which requires several samples of a scalar-valued random variable Y_v (measured activity) for neurons $v \in V$. We consider two of such Y_v possibilities

DPGM 1: Neural recordings at time t : $Y_v = X_v(t), v \in V$. Different t gives different samples of Y_v .

DPGM 2: Windowed Average of recordings over a duration of 50 msec: $Y_v = \bar{X}_v, v \in V$, and averaging over different windows of 50 msec with a gap of 50 msec in between consecutive windows gives different samples of Y_v .

We quantified the performance of the algorithms by inference of CFC for 25 different simulations and summarize the performance by two metrics, Accuracy (A) and Sensitivity (S). Let True Positive (TP) be the

number of correctly identified edges, True Negative (TN) be the number of missing edges that were correctly identified, False Positive (FP) be the number of incorrectly identified edges and False Negative (FN) be the number of missing edges incorrectly identified across simulations. We define the Accuracy as

$$A = \frac{TP+TN}{TP+TN+FP+FN},$$

which measures the ratio of the count of correctly identified edges or missing edges to the count of all possible edges across simulations. In the motifs and simulation paradigms we consider, there are 4 neurons and 16 possible edges (including self-loops) per simulation resulting with total of 400 possible edges across 25 simulations. We also define the Sensitivity as

$$S = \frac{TP}{TP + FN}$$

the ratio of the count of true edges that were correctly identified to the total count of the true edges across simulations. In this comparative study, Sensitivity is more relevant than Accuracy since it focuses on the detection of the true edges. Indeed, in the extreme case of having the estimated CFC to be an empty set of edges across simulations, the linear Gaussian paradigm will still have 70% neuron pairs correctly identified to be not connected by an edge, thereby resulting in $A = 70\%$. Whereas, there will be 0% of true edges detected correctly resulting in $S = 0\%$, which reflects the undesirability of the empty graph estimate. We report both Accuracy and Sensitivity for a comprehensive summary of performance. We also report the percentage of the simulations that has each estimated edge present. Higher percentage indicates higher confidence in the detection of that edge. Figure 2.7 compares the results for GC, DPGM 1 and DPGM 2 methods in inference of the true CFC.

In *Linear Gaussian scenario* (top row in Fig. 2.7), the connections between neurons in the Ground Truth CFC are excitatory due to positive coefficients in the linear dynamical equation for neural activity. GC generates a sparse set of edges in which it correctly detects a single edge $3 \rightarrow 4$ among the three edges of the true CFC. DPGM 1 generates a large set of edges (9 out of 16 possible) with several of them being spurious. Indeed, each edge is present in less than 25% of the simulations. DPGM 2 generates the same number of edges as DPGM 1, however it has less spurious edges indicated by higher percentages for the expected edges in the

Ground Truth CFC ($1 \rightarrow 3, 3 \rightarrow 4$ with 92% and 100%). Overall, all methods result in $A > 80\%$ while sensitivity for GC, DPGM 1 and DPGM 2 varied significantly $S = 33.3\%, 8.7\%, 76.3\%$ respectively. We thereby conclude that among the three methods, GC is the most accurate but since it did not detect two out of three edges, it is not as sensitive as DPGM 2.

In *Non-linear Non-Gaussian scenario (middle row)*, in the Ground Truth CFC for the Non-linear Non-Gaussian scenario, $1 \rightarrow 3, 3 \rightarrow 4$ are excitatory due to $\sin(x)$ being an increasing function while $2 \rightarrow 3$ is an inhibitory connection due to $\cos(x)$ being a decreasing function for $x \in [0, 1]$ in their dynamical equation. As previously, GC consistently detects a sparse set of edges (single edge $1 \rightarrow 3$ with 100%) which is one of the three true edges. DPGM 1 obtains a large set of edges with several of them are spurious edges and all edges appear in less than 25% of the trials. For DPGM 2, the number of spurious edges obtained is more than GC and less than DPGM 1. DPGM 2 obtains correctly two out of the three true edges $1 \rightarrow 3$ and $2 \rightarrow 3$ in most of the trials (96% and 88% respectively). In summary, GC, DPGM 1 and DPGM 2 resulted in an accuracy of $A = 87.5\%, 79.4\%, 90.8\%$ respectively and sensitivity of $S = 33.3\%, 4.7\%, 63.7\%$ respectively. For this scenario, DPGM 2 has the highest accuracy and sensitivity among the methods.

In *CTRNN scenario (bottom row)*, self-loops are present for each neuron, and due to positive weights and increasing activation function $\sigma(\cdot)$ in their dynamical equation, the connections in the Ground Truth CFC are excitatory. GC obtains two of the three true non-self edges $1 \rightarrow 3, 2 \rightarrow 3$ for 60%, 40% of the trials. DPGM 1 detects spurious edges, but also infers the true edges $1 \rightarrow 3, 2 \rightarrow 3$ for 72% of the trials. DPGM 2 obtains less number of spurious edges compared to DPGM 1 and obtains all of the non-self true edges $1 \rightarrow 2$ and $2 \rightarrow 3$ for 100% of the trials. In summary, all methods result in a lower accuracy of $A \approx 70\%$ compared to other scenarios since they do not include self-loops and sensitivity is $S = 16.7\%, 28\%, 33.2\%$ for GC, DPGM 1 and DPGM 2, respectively, which is considerably lower than other scenarios. Among all methods, DPGM 2 obtains the highest accuracy followed by GC and lastly DPGM 1. DPGM 2 had the highest sensitivity compared to the other methods.

The choice of thresholds tunes the decision whether a connection exists in the CFC. For DPGM-based approaches, increasing the p-value cut-off for conditional independence tests increases the rate of detecting edges while also increasing the rate of detecting false positives. We use a p-value cut-off of 0.1, after trial and

error, for DPGM 1 and DPGM 2. For GC, a likelihood ratio statistic L_{uv} is obtained for testing $A_{uv}(k) = 0$ for $k = 1, \dots, K$. An edge $u \rightarrow v$ is outputted if L_{uv} has a value greater than a threshold. We use a percentile-based threshold [108], and output an edge $u \rightarrow v$ if L_{uv} is greater than 95 percentile of L_{ij} 's over all pairs of neurons (i, j) in the graph. We also trialed with different percentile thresholds of 90 percentile and higher, none of which outperformed DPGM.

2.3 Discussion

In this chapter, we establish a statistical guide to neural connectomics with regards to mapping neural interaction anatomy, function and causality with the means of graphical models. We first describe possibilities of mapping neural anatomical connections with a graph, i.e., anatomical connectome, and demonstrate the difference between such a mapping and a graph that captures functional interactions, i.e., associative functional connectome (AFC). Recognizing that the ultimate goal of functional connectomics is to infer causal interactions between neurons, we define the graphical tools and properties needed to distill AFC into a directional graph, i.e., causal functional connectome (CFC), which represents flows of cause and effect in the interaction between neurons. We then compare exemplary common approaches having the ultimate goal of finding "causation", such as Granger Causality (GC), Dynamic Causal Model (DCM) and Directed Probabilistic Graphical Model (DPGM), in the context of functional connectomics. In particular, we introduce the developments in statistical theory of DPGM to the subject of CFC inference, and define the Directed Markov Property that guarantees consideration of cause and effect in graph settings. We show that this property is key in the definition of probabilistic graphical models that could constitute neural CFC. We then describe the PC algorithm, a common statistical approach for inference of such graphs. Based on these notions and the outcomes of the Directed Markov Property we formulate criteria based on which CFC models can be compared.

We conclude by performing a holistic comparison, in Table 2.1, of several common approaches that do not obey the Directed Markov Property, such as Granger Causality (GC) and Dynamic Causal Model (DCM), with variants of the PC algorithm (DPGM), comparing them with respect to the criteria that we have outlined. We demonstrate the applicability and the challenges for inference of CFC from measured neural activity for each of the approaches on simulated motifs in Figure 2.7.

Table 2.1 Comparative summary of different approaches for causal modeling in functional connectomics.

	GC	DCM	DPGM
Format of Causality	Non-zero parameters in VAR model	Coupling parameters in biological model	Directed Markov Graph
Inclusion of temporal relationships	Yes	Yes	No , formulation for stationary variables
Generalizable statistical model	Yes	No	Yes
Non-parametric Model	Yes , parametric and non-parametric approaches exist.	No , biologically mechanistic non-linear model.	Yes , equivalent to an arbitrary functional relationship between nodes.
Supports CFC estimation	Yes	No , suitable for comparing model hypotheses	Yes
Occurrence of cycles (including self-loops) in the model	Yes (neuron $i \rightarrow i$ when $A_{ii}(k) \neq 0$ for some k)	Yes ($i \rightarrow i$ when $\theta_{ii} \neq 0$)	No , it is a DAG
Incorporation of interventional and counterfactual queries	No	No	Yes but for stationary variables.

Functional Connectivity between neurons describes statistical dependency between observed neural signals and is descriptive in nature without requiring accurate modeling. FC can either describe undirected statistical associations (AFC) or directed causal interdependencies among the observed neural signals (CFC) [25]. In contrast, there is another concept of Effective Connectivity in literature which aims to quantify the causal influence between hidden neural states and corresponds to the parameter of a mechanistic model that aims to explain the observed directed dependencies [95; 168]. In practice, the analysis of Effective Connectivity involves comparison of generative models with coupling among hidden brain states, based on evidence from observed data. Whereas, analysis of CFC is predictive in nature, that estimates the presence and/or strength of causal dependencies from the observed recordings and does not require generative modeling.

In this work, our aim is to formulate statistical properties and criteria related to causality of functional

connectomics, rather than propose a new approach for causal functional connectome modeling. Such formulation is expected to identify existing gaps in causal modeling and guide extensions of causal functional connectome models ideally satisfying all criteria that we have outlined. Indeed, capturing as many causal criteria is fundamental to any approach from statistical and application points of view. For example, one such property of importance is the ability to uncover directed relationships in ground truth dynamical equations [108; 109], which we include in our comparative study. We have compared the approaches to find CFC in simulations from linear gaussian, non-linear non-gaussian, and CTRNN to demonstrate how specific model assumptions in ground truth dynamical equations are impacting the utilization of the approaches in recovery of relationships between the neurons. For the methods that we have tested, our simulated comparison shows that GC output typically results in a sparse graph with inferred edges that indeed represent causal connections, but we find that it also misses multiple edges that represent causal connections (high accuracy; low sensitivity). For DPGM, we find that the output depends on the choice of measured activity. DPGM 1, which uses full neural time series, results in comparatively low accuracy and low sensitivity since does not capture dependencies across time. DPGM 2 uses neural activity averaged over time and results in comparatively high accuracy as well as sensitivity for detection of most of the true edges, but not all of them. These results are not surprising, since the PC algorithm guarantees causality for independent samples per time point thus guarantees the Directed Markov Property only for a single time point, however here we aim to infer a single CFC for the whole recording over time (as in DPGM 1) [48]. In such a case, the Directed Markov Property and causality are not guaranteed. Averaging over time and separating by time gaps to reduce interdependence between samples leads to improved performance of DPGM 2, but this does not necessarily reflect the time dependent causes and effects between and within neural time series, thus again does not guarantee causality.

Our exposition of properties that each approach is based upon and the comparative study show that each of the methods address different aspects of modeling causality of neural interaction and mapping them in the form of a graph. In particular, GC aims to obtain the directed functional connectivity from observed neural states, in a way that tells whether a variable's *past* is predictive of another's future, without requiring a detailed model. It consists of a framework based on auto-regression models which is relatively easy to compute. In contrast, DCM enables *specific* model hypotheses to be compared based on evidence from data and provides insights on causal connectivity between hidden neural states based on those models. DPGM is a

generic procedure that represents causal functional connectivity from observed neural states in a way that is predictive of the consequence of counterfactual interventions while answering queries of causal dependency in both interpretable and mathematically precise manner. To summarize these differences, we outline in Table 2.1 the strengths and weaknesses of each of the approaches with respect to applicability to various criteria of causality. Although these approaches are three distinct pathways popular for causal modeling, there have been attempts to combine them under a single framework [169], and to quantify causal effects that are applicable to either framework although under stringent conditions [170].

The comparative table demonstrates that with respect to the model that each approach is assuming, GC requires a linear model in its common use, though has recent non-linear and non-parametric extensions. DCM requires a strict well defined mechanistic biological model and thus can only compare different models based on evidence from data. In comparison, DPGM has an advantage that it does not require modeling of the neural dynamics using a parametric equation or assumption of a linear model. Furthermore, the Directed Markov Condition of the DPGM implies the existence of a functional relationship (Eq. 2.6) between parent and children connections in the graph, thus doing away with the need for modeling by specific linear or non-linear functions. In regards to guarantee of causality, GC can provide useful insights into a system's dynamical interactions in different conditions, however its causal interpretation is not guaranteed as it focuses on the predictability of future based on past observations of variables. DCM uses the parameters for coupling between hidden neural states in competing biological models to indicate CFC, however it compares hypothetical models based on evidence from data which relevance to causality is not guaranteed [141]. In comparison, DPGM provides a probabilistic foundation for causality which is predictive of the consequence of possible intervention like neuron ablation and counterfactual queries. Inference of CFC is possible with several causal graph inference algorithms such as the PC algorithm.

Such properties of causal interaction between entities are what makes DPGM popular in various disciplines such as genomics and econometrics [29; 30; 31; 32; 33; 34; 35; 36; 37; 38; 39; 40]. However, such an inference by DPGM typically produces a DAG, though adaptations exist which aim to include cycles in the CFC while having a more complicated output [62]. Furthermore, DPGM-based inference guarantees causality for independent measurements only. There is no such guarantee when considering whole neural time series due to temporal dependence, and this in practice leads to a decline in performance due to the

degree of temporal dependence in the time series. Despite these limitations, DPGM substantially adds to the causal interpretation of CFC by being predictive of counterfactual interventions, unlike GC and DCM. DPGM is model-free, unlike popular GC variations relying on auto-regression and DCM. DPGM estimates the CFC graph from observed data, unlike DCM. Furthermore, using average activity over time instead of the entire time series, DPGM 2 performs well at the recovery of ground truth connectivity in simulation studies over different motifs compared to GC (See Fig. 2.7).

In the context of fMRI datasets, several studies aim to evaluate the approaches. In simulated fMRI data from a DCM-based forward model, GC and its variations have suffered in performance [23; 64], suspected to be due to inter-regional differences in the haemodynamic response function and causal processes occurring faster than the sampling rate. Later studies on more detailed fMRI simulations based on spiking neuron models coupled to biophysically realistic haemodynamic observation models revealed that, GC is largely invariant to changes in haemodynamic response function properties, however, in the presence of severe downsampling and/or high measurement noise, which can be typical of fMRI data, GC suffers in performance [171; 172; 173]. In an attempt to address these challenges, novel approaches to GC have recently been proposed [174; 175; 176]. DCM exhibits good performance in comparing few model hypotheses for coupling between hidden neural states, however, DCM has the limitation that it is in general not mathematically and computationally feasible to search across all possible functional connectivity graphs and estimate from observed fMRI data [141]. In contrast, DPGM-based approaches have been studied to perform relatively well in searching for the CFC and estimate from observed data in the fMRI setting [64; 81], such as, PC algorithm used on the whole time series [23], Independent Multisample Greedy Equivalence Search (IMaGES) which addresses Simpson’s paradox in multi-subject fMRI data by assigning a penalized score for the CFC for each subject, combining them across individual subjects, and optimizing the score [64], and Fast Adjacency Skewness (FASK) which incorporates feedforward and feedback connections in the CFC graph in fMRI setting [177].

In conclusion, DPGM provides a probabilistic and interpretable formulation for CFC. We have established the statistical properties that the inferred DPGM should possess as well as demonstrated its performance in inference of CFC. While DPGM is a powerful causal framework, existing DPGM algorithms do not reflect the inter-temporal causal dependencies within and between the neural time series. Yet in the neural time

series setting, nodes of the connectivity graph are neurons that correspond to an entire time series of neural activity and comprise inter-temporal dependence. Thus the remaining challenge is the adaptation of the DPGM based formulation of CFC to incorporate inter-temporal dependencies in the neural time series. Such an adaptation will further increase the strength of using DPGM for CFC inference from neural dynamics.

Chapter 3

Functional and Causal Neural Connectomics: The Time-Aware PC Algorithm

This chapter is a reproduction of published work in Biswas and Shlizerman [2], with the exception of Figure 3.8 and its discussion in section 3.3, and minor adaptations to fit the thesis structure and formatting.

3.1 Causal Functional Connectivity for Time Series

In this section, we propose a novel methodology for modeling and estimation of the CFC for time series. This methodology is generic and applicable to various time series including neural recordings. The CFC is represented by a graph with nodes as neurons, each corresponds to a time series of neural activity, and edges indicating causal connectivity. In the following, we show we can explicitly model causal relations within and between time series in a DPGM framework and use it to define the CFC.

3.1.1 Unrolled Graphical Modeling of Time Series

We aim to incorporate the causal influence of the activity of neuron u at time t_1 upon the activity of neuron v at time t_2 in our proposed model. To do that, we *unroll* the time series $X_v(t), v \in V, t \in \{0, \dots, T\}$ into nodes (v, t) for neuron v and time t where the node (v, t) corresponds to the variable $X_v(t)$. We use directed edges between the nodes (v, t) to represent causal relations between $X_v(t)$. For example, the edge

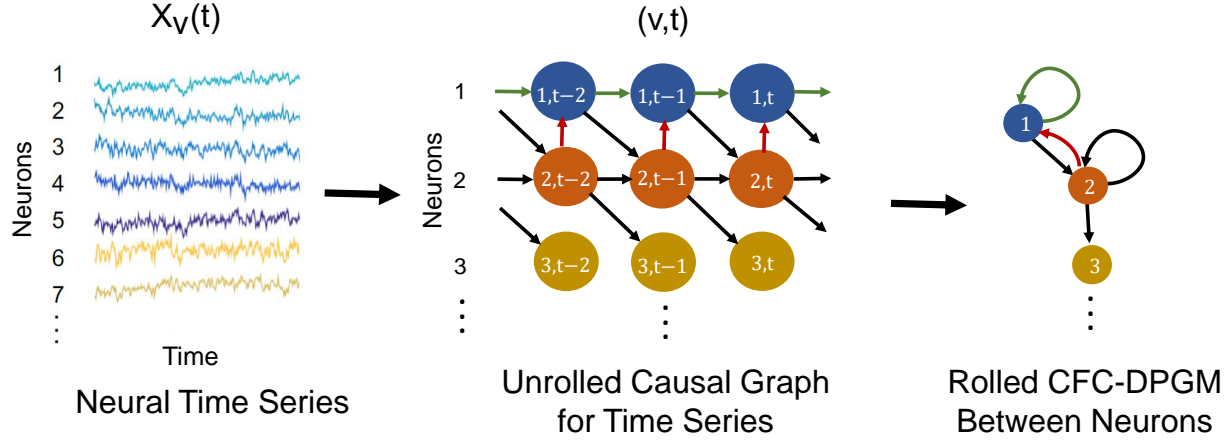


Figure 3.1 Causal modeling of neural time series. Schematic depiction of causal modeling of the neural time series (left) by the *Unrolled Causal Graph* (middle), and then rolling back its edges (colored) to define the *Rolled CFC-DPGM* (right).

$(u, t_1) \rightarrow (v, t_2)$ represents the causal influence of the activity of neuron u at time t_1 , $X_u(t_1)$ upon the activity of neuron v at time t_2 , $X_v(t_2)$. Let $\mathbf{V} = \{(v, t), v \in V, t \in \{0, \dots, T\}\}$ be the set of nodes in the unrolled time series, \mathbf{E} be the set of directed edges between the nodes and $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ be the *Unrolled Graph* for the time series (See Figure 3.1-middle). We assume that \mathbf{G} is a DAG in which there are no cycles in the edges in \mathbf{E} .

Causal relationships between neurons are either forward in time: (1) from neurons u at time t_1 to neuron v at time t_2 for $t_1 < t_2$, represented by $(u, t_1) \rightarrow (v, t_2)$ for $t_1 < t_2$ in \mathbf{G} ; or, (2) contemporaneous: when causal influences occur more rapidly than the sampling interval of the time series, represented in \mathbf{G} by $(u, t) \rightarrow (v, t)$ at time t [178; 179]. Causal relations cannot direct backward in time, that is, \mathbf{G} will not contain $(u, t_1) \rightarrow (v, t_2)$ for $t_1 > t_2$. Furthermore, for the contemporaneous causal influences, we do not allow the activity of u at time t to have causal influence on itself at time t , that is $(u, t) \rightarrow (u, t)$ is not allowed, while $(u, t) \rightarrow (v, t)$ is allowed. These considerations imply the absence of cycles in \mathbf{G} , thereby justifying the assumption for \mathbf{G} to be a DAG to model the causal interactions in the unrolled time series.

In practice, the true causal interactions between $X_v(t)$ are unknown. Yet, when $\mathbf{X} = \{X_v(t) : (v, t) \in \mathbf{V}\}$ satisfies the Directed Markov Property (DMP) with respect to DAG \mathbf{G} (See Section 3.4.1 for details about the DMP), then it implies that \mathbf{G} captures the causal functional interactions among $X_v(t)$, as we show in the functional equivalence of DMP in Section 3.4.1. We refer to such a DAG \mathbf{G} as the *Unrolled Causal*

Graph for the time series $X_v(t), v \in V, t \in \{0, \dots, T\}$.

3.1.2 Rolled CFC-DPGM

Typically for signals, effective CFC representation refers to relations between neurons rather than relations between different signals' times. Thereby, we propose to roll back the unrolled causal graph \mathbf{G} to define CFC between neurons (See Figure 3.1-right). The rolled graph is based on the principle that the existence of a causal relationship from neuron u at time t_1 to neuron v at time t_2 , $(u, t_1) \rightarrow (v, t_2) \in \mathbf{G}$ for $t_1 < t_2$, would imply that u is connected to v in the rolled CFC. In practice, the causal interactions weaken as the time-gap $t_2 - t_1$ grows. Thereby, we consider a maximum time-delay of interaction, τ , so that (u, t_1) and (v, t_2) would not share a significant influence between them if the time gap $t_2 - t_1 > \tau$ for all neurons u, v . Such a consideration of the maximum time-delay aids in making statistical inference from the time series data. Thus, if $(u, t_1) \rightarrow (v, t_2)$ in \mathbf{G} for some $t_1 \leq t_2 \leq t_1 + \tau$, then the CFC graph between neurons should include $u \rightarrow v$. We consolidate these concepts to define causal functional connectivity between neurons based on DPGM, in the following.

Definition 2 (Rolled CFC-DPGM). Let \mathbf{X} satisfy the DMP with respect to DAG $\mathbf{G} = (\mathbf{V}, \mathbf{E})$. The *Rolled CFC-DPGM* for neurons in V with maximum time-delay of interaction τ , is defined as the directed graph F_τ having edge $u \rightarrow v$ if $(u, t_1) \rightarrow (v, t_2) \in \mathbf{E}$ for some $t_1 \leq t_2 \leq t_1 + \tau$, $t_1, t_2 \in \{0, 1, \dots, T\}$. $u, v \in V$ could be either the same neuron or distinct neurons.

We show an example in Figure 3.1, where for neurons $V = \{1, 2, 3\}$, their unrolled DAG \mathbf{G} is represented by Figure 3.1-middle. When the neural data \mathbf{X} satisfies the directed Markov Property with respect to \mathbf{G} , the Rolled CFC graph with maximum time delay of interaction $\tau = 1$ is given by Figure 3.1-right. Note that the same CFC graph would be obtained by taking any value of $\tau \geq 1$.

Contemporaneous and Feedback Interactions We highlight that the Rolled CFC-DPGM in Definition 2 incorporates contemporaneous interactions, which can arise if the causal influences occur more rapidly than the sampling interval of the time series or the aggregation interval for aggregated time series. Such a scenario can arise for example in spiking neural datasets where peri-stimulus time histograms aggregate the spike trains over time intervals [180; 181], and in Functional Magnetic Resonance Imaging (fMRI)

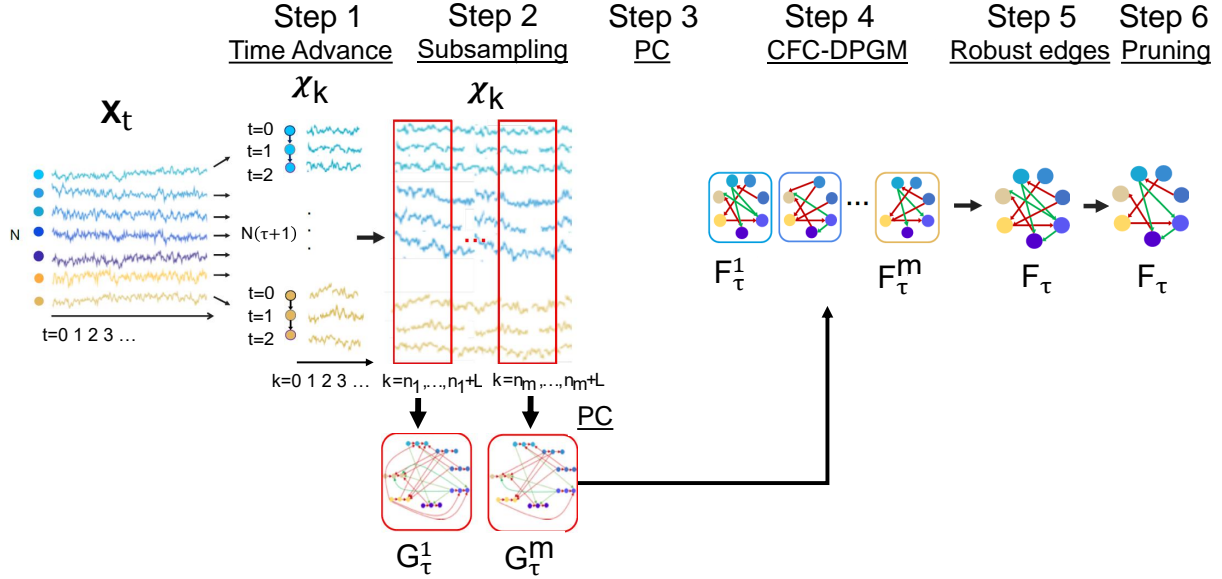


Figure 3.2 The TPC algorithm. Illustration of Steps 1-6 in the TPC Algorithm: Time Advance, Subsampling, PC, Rolled CFC-DPGM, Robust edges and Pruning.

datasets where low sampling rates is typical [142]. However, if the time span in sampling and aggregation is expected to be less than the time scale of causal interactions, then one can impose $t_1 < t_2$ to exclude contemporaneous interactions. Additionally, the Rolled CFC-DPGM accommodates self-loops in neural interactions [182; 183], by checking whether $(u, t_1) \rightarrow (u, t_2) \in \mathbf{E}$ for some $0 \leq t_1 < t_2 \leq t_1 + \tau$ in determining whether $u \rightarrow u \in F_\tau$. Longer feedback loops are also incorporated. For example, the existence of $u \rightarrow v \rightarrow u \in F_\tau$ is determined by checking whether $(u, t_1) \rightarrow (v, t_2)$ and $(v, t_2) \rightarrow (u, t_3) \in \mathbf{E}$ for some $0 \leq t_1 \leq t_2 \leq t_3 \leq t_1 + \tau$. By virtue of the technique of unrolling and rolling back, the Rolled CFC-DPGM captures causality while including cycles, since the Unrolled Graph is still a DAG and thereby meets the requirement for satisfying DMP.

3.1.3 Estimation from data: Time-Aware PC (TPC) Algorithm

In this section we outline the steps to estimate the Rolled CFC-DPGM from time series $X_v(t)$, $v \in V$, $t \in \{0, 1, \dots, T\}$, $\dim V = N$, which constitute the Time-Aware PC (TPC) Algorithm. Let us

denote the recordings from the N neurons at time t by $\mathbf{X}_t = (X_{v_1}(t), \dots, X_{v_N}(t))$, where $\dim \mathbf{X}_t = N$, for $t \in \{0, 1, \dots, T\}$. The output of TPC is a Rolled CFC that contains edges found significant in the duration of recording. The TPC algorithm first aims to estimate the unrolled DAG with set of nodes $\{(v, t) : v \in V, t \in \{0, 1, \dots, \tau\}\}$, for a fixed $\tau \geq 1$. For the estimation, TPC constructs samples $\chi_k = (\mathbf{X}_{2(\tau+1)k}, \dots, \mathbf{X}_{\tau+2(\tau+1)k})$, where $\dim \chi_k = N(\tau + 1)$, $k = 0, 1, \dots, K := \lfloor \frac{T-\tau}{2(\tau+1)} \rfloor$. Note that χ_k is a concatenation of the vectors $\mathbf{X}_{2(\tau+1)k}, \dots, \mathbf{X}_{\tau+2(\tau+1)k}$, each of dimension N , into a resulting vector of dimension $N(\tau + 1)$. We consider that the entry of χ_k corresponding to node v and time $t + 2(\tau + 1)k$, i.e. $X_v(t + 2(\tau + 1)k)$, for $k = 0, \dots, K$ are samples for node (v, t) , $v \in V, t \in \{0, 1, \dots, \tau\}$. In other words, samples for a node (v, t) consists of $X_v(t + 2(\tau + 1)k)$ for $k = 0, 1, \dots, K$, which are time advanced instances of neuron v starting at time t with time advances of multiples of $2(\tau + 1)$. Such a time-advance of multiples of $2(\tau + 1)$ ensures a substantial time-gap of $2(\tau + 1)$ units between samples, that reduces interdependence between the samples, considering the maximum time-delay of interaction τ (See Section 3.4.2). This operation increases the dimension of the considered time-series at each time by a multiple of $\tau + 1$ such that for original recordings of \mathbf{X}_t of dimension N , χ_k would have a dimension of $N(\tau + 1)$.

In Step 2, TPC selects m windows of length L starting from n_b chosen at random, i.e. $\{n_b + 1, n_b + 2, \dots, n_b + L\} \subset \{0, 1, \dots, N\}$, to obtain $\{\chi_k\}_{k=n_b+1}^{n_b+L}$, $b = 1, \dots, m$. The process is called *subsampling* since in the next step, the PC algorithm is applied on each subsample $\{\chi_k\}_{k=n_b+1}^{n_b+L}$ to output an estimate of the unrolled DAG with nodes (v, t) , $v \in V, t \in \{0, 1, \dots, \tau\}$, thus yielding a set of such estimates.

In Step 3, on each subsample $\{\chi_k\}_{k=n_b+1}^{n_b+L}$, the PC algorithm estimates the unrolled DAG with nodes (v, t) , $v \in V, t \in \{0, 1, \dots, \tau\}$ as a completed partially directed acyclic graph (CPDAG), defined as the graph union of DAGs that satisfy DMP with respect to the subsample, and denoted by \mathbf{G}_τ^b [111].

In Step 4, the Unrolled graph \mathbf{G}_τ^b is transformed to give the Rolled CFC-DPGM denoted F_τ^b . At this step, weights for edges in F_τ^b are also obtained using interventional connectivity weights (3.4.3), that quantifies the causal effect of intervention on each neuron to its connected neurons.

In Step 5, a single CFC consensus F_τ is obtained by keeping only those edges which have relative frequency of occurrence among F_τ^b , $b = 1, \dots, m$ to be greater than the cut-off γ . A single connectivity weight consensus for an edge in F_τ is achieved by averaging over the weights for the same edge, whenever present, over the m iterates. The resampling procedure promotes detection of stable edges [184].

Algorithm 1: TPC

Input : Recordings of activity of neurons in $V = \{v_1, \dots, v_N\}$ over time:
 $\mathbf{X}_t = (X_{v_1}(t), \dots, X_{v_N}(t)), t \in \{0, \dots, T\}$; maximum delay of interaction between neurons τ ; significance level α ;
For subsampling: window width L ; m iterations; subsampling stability cutoff γ ;

Output : CFC estimate, denoted F_τ .

1 begin

- Step 1. Time-advanced samples. Let $\chi_k = (\mathbf{X}_{2(\tau+1)k}, \dots, \mathbf{X}_{\tau+2(\tau+1)k})$, where $\dim \chi_k = N(\tau + 1)$, $k = 0, 1, \dots, K := \lfloor \frac{T-\tau}{2(\tau+1)} \rfloor$. Construct $\{\chi_k\}_{k=0}^K$.
- Step 2. Subsampling. Select a random integer $n_b \in \{0, \dots, K - L\}$ and choose the window of length L given by $\{n_b + 1, \dots, n_b + L\} \subseteq \{0, \dots, K\}$, to obtain $\{\chi_k\}_{k=n_b+1}^{n_b+L}$.
- Step 3. PC. Use the PC algorithm based on samples $\{\chi_k\}_{k=n_b+1}^{n_b+L}$ to estimate the unrolled DAG with nodes $(v, t), v \in V, t \in \{0, \dots, \tau\}$. Denote the output as \mathbf{G}_τ^b .
- Step 4. Rolled CFC-DPGM.
- (a) Convert \mathbf{G}_τ^b to rolled CFC-DPGM F_τ^b with max. delay of interaction τ using Def. 2.
 - (b) Find the connectivity weights $w_\tau^b(u, v)$ for connections $u \rightarrow v \in F_\tau^b$ by Def. 3.4.3.
- Step 5. Robust edges.
- (a) Repeat Steps 2-5 to obtain m iterates of the rolled CFC-DPGM: F_τ^b , and connectivity weights $w_\tau^b(u, v), b = 1, \dots, m$.
 - (b) Output a single CFC, F_τ , with only those edges whose relative frequency of occurrence among F_τ^b is above γ .
 - (c) Output single Connectivity Weights, $w_\tau(u, v)$, for connections $u \rightarrow v \in F_\tau$, as the average of $\{w_\tau^b(u, v) : u \rightarrow v \in F_\tau, u \rightarrow v \in F_\tau^b, i = 1, \dots, m\}$ when the set is non-empty and 0 otherwise.
- Step 6. Pruning. Remove from F_τ those edges $u \rightarrow v \in F_\tau$ with $|w_\tau(u, v)| < w_0$, where, $w_0 = \frac{1}{10} \max\{|w_\tau(u, v)| : u \rightarrow v \in F_\tau\}$.

2 end

Finally, in Step 6, F_τ is pruned to further reduce spurious edges, by removing the edges which have exceedingly low connectivity weights, determined by those edges in F_τ whose weights are less than a tenth in magnitude compared to the maximum magnitude for edge weights in F_τ . The TPC algorithm is outlined in Alg. 1 and Figure 3.2.

3.2 Results

3.2.1 Comparison Study of TPC with other Approaches to Causal Functional Connectivity

We compare the performance of TPC with different existing CFC inference approaches to recover relationships in ground truth dynamical equations by generating synthetic data from three simulation paradigms. In particular, we estimate their CFC using GC, DPGM and TPC. The simulation paradigms correspond to specific model assumptions to assess the impact of model assumptions on the performance of the approaches (See Appendix A.3).

We measured the algorithms' performance using CFC inference for 25 different simulations and summarized the results using three metrics: (1) Combined Score (CS), (2) True Positive Rate (TPR), (3) 1 - False Positive Rate (IFPR). Let True Positive (TP) represent the number of correctly identified edges, True Negative (TN) represent the number of correctly identified missing edges, False Positive (FP) represent the number of incorrectly identified edges, and False Negative (FN) represent the number of incorrectly identified missing edges across simulations. Therefore, the total number of ground-truth missing edges is $FP+TN$, and the total number of ground-truth existing edges is $TP+FN$. IFPR is defined as: $IFPR = \left(1 - \frac{FP}{FP+TN}\right) \cdot 100$, which measures the ratio of the number of correctly identified missing edges by the algorithm to the total number of true missing edges. Note that the rate is reported such that 100% corresponds to no falsely detected edges. TPR is defined $TPR = \left(\frac{TP}{TP+FP}\right) \cdot 100$ as the ratio of the number of correctly identified edges by the algorithm to the total number of true edges in percent. The Combined Score (CS) is given by Youden's Index [185; 186], as follows, $CS = TPR - FPR$.

In the motifs and simulation paradigms that we consider, there are 4 neurons and 16 possible edges (including self-loops) per simulation resulting with total of 400 possible edges across 25 simulations. Figure 3.3 compares in detail the results for GC, DPGM and TPC in inference of true CFC for noise level $\eta = 1$ and thresholding parameter $\alpha = 0.05$. Here we also report the percentage of the simulations that has each estimated edge present. Higher percentage indicates higher confidence in the detection of that edge. Figure 3.4 compares the Combined Score of the approaches over different values of noise level η and thresholding parameter α for each simulation paradigm.

In *Linear Gaussian scenario* (left column in Figure 3.3), the connections between neurons in the Ground

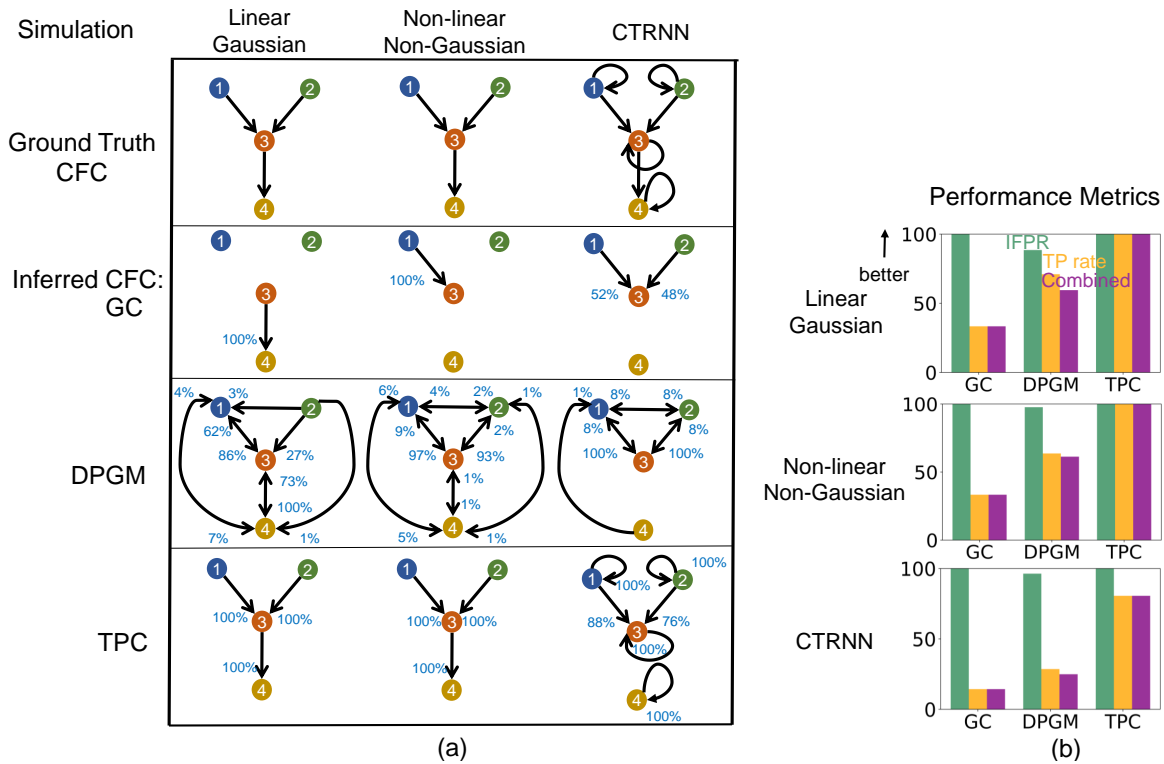


Figure 3.3 Comparative study of CFC inference. (a) CFC inference by GC, DPGM, and TPC, is compared on three examples of motifs and simulation paradigms; from left to right: Linear Gaussian, Non-linear Non-Gaussian, CTRNN. Table: 4-neurons motifs that define the Ground Truth CFC (row 1) are depicted along with inferred CFC over several simulation instances according to the three different methods (row 2-4). Each inferred CFC has an edge $v \rightarrow w$ that corresponds to an edge detected in any of the inference instances. The percentage (blue) next to each edge indicates the number of times the edge was detected out of all instances. (b) IFPR (green), TP rate (orange) and Combined Score (purple) of each method are shown for each motif.

Truth CFC are excitatory due to positive coefficients in the linear dynamical equation for neural activity. GC generates a sparse set of edges in which it correctly detects a single edge $3 \rightarrow 4$ among the three edges of the true CFC but misses two other edges. DPGM generates a large number of edges (9 out of 16), many of which are spurious, though it has a high percentage for expected edges in the Ground Truth CFC ($1 \rightarrow 3, 3 \rightarrow 4$ with 87% and 100% respectively). TPC obtains the Ground Truth CFC, with no spurious edges and obtains the expected edges in all of the trials ($1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$ with 100%, 100% and 100% respectively). Overall, GC, DPGM and TPC produce IFPR = 100%, 88.5%, 100%, TPR = 33.3%, 71.0%, and 100%, and CS = 33%, 59%, 100% respectively. Thereby, among the three methods, we conclude that TPC detects the edges perfectly, while GC is highly specific to correct edges, but since it does not detect two out of three edges it is not as sensitive as DPGM.

In the *Non-linear Non-Gaussian scenario (second column)*, in the Ground Truth CFC consists of $1 \rightarrow 3, 3 \rightarrow 4$ excitatory due to $\sin(x)$ being an increasing function, while $2 \rightarrow 3$ is an inhibitory connection due to $\cos(x)$ being a decreasing function for $x \in [0, 1]$ in the dynamical equation. As previously, GC consistently detects a sparse set of edges (single edge $1 \rightarrow 3$ with 100%) which is one of the three true edges. DPGM again generates a large number of edges, some of which are spurious. In the majority of trials (97% and 93%, respectively), DPGM correctly obtains two of the three true edges $1 \rightarrow 3$ and $2 \rightarrow 3$. In contrast, TPC obtains no spurious edges and the true edges were detected for all the trials ($1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$ with 100%, 100%, 100%). In summary, GC, DPGM and TPC yielded IFPR = 100%, 97.6%, 100% and TPR = 33.3%, 63.7%, 100% and CS = 33%, 61%, 100%. For this scenario, TPC again has the highest performance among the methods.

In *CTRNN scenario (third column)*, self-loops are present for each neuron, and due to positive weights and increasing activation function $\sigma(\cdot)$ in their dynamical equation, the connections in the Ground Truth CFC are excitatory. GC obtains two of the three true non-self edges $1 \rightarrow 3, 2 \rightarrow 3$ for 52%, 48% of the trials. DPGM detects spurious edges, but also infers the non-self true edges $1 \rightarrow 3, 2 \rightarrow 3$ for 100% of the trials. In comparison, TPC infers no spurious edges and all the self true edges for 100% of the trials and non-self true edges $1 \rightarrow 3$ and $2 \rightarrow 3$ for 88%, 76% of the trials. It is noteworthy that none of these approaches report the $3 \rightarrow 4$ edge in the CTRNN setting. This is suspected to be because of the dynamics of the process in which the strength of the self-inhibition of neuron 4 far exceeds the damped influence (via tanh activation

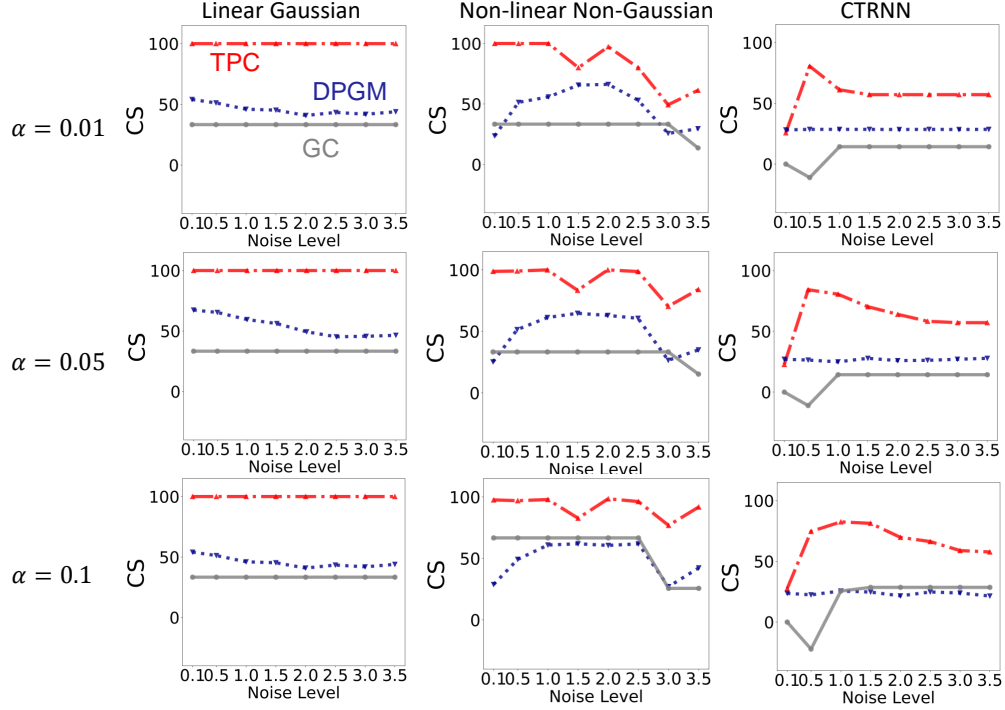


Figure 3.4 Comparative study over levels of noise and thresholding parameter. Combined Score of the three methods of CFC inference - TPC (red), DPGM (blue), GC (gray), over varying noise levels in simulation $\eta = 0.1, 0.5, 1.0, \dots, 3.5$, for simulated motifs from Linear Gaussian, Non-linear Non-Gaussian and CTRNN paradigms (left to right), with thresholding parameter $\alpha = 0.01, 0.05, 0.1$ (top to bottom).

function) from neuron 3 to 4. In summary, IFPR of GC, DPGM and TPC is 100%, 96.3%, 100% and TPR is 14.3%, 28.6%, 80.6% and CS is 14%, 25%, 81% respectively. Among all methods, TPC has the highest TPR, followed by DPGM and lastly GC. TPC and GC have the highest IFPR having not detected any false edges, followed by DPGM. In terms of the CS, TPC has the highest performance compared to other methods.

We compare Combined Score of TPC and other approaches across varying levels of simulation noise η from 0.1 to 3.5 and thresholding parameter $\alpha = 0.01, 0.05, 0.1$ in Figure 3.4. In the Linear Gaussian scenario, we note that TPC has a CS of $\approx 100\%$ across all levels of simulation noise and thresholding parameter α , and is followed by DPGM in performance and lastly GC. In the Non-linear Non-Gaussian scenario, TPC has the highest CS compared to other methods across levels of noise and α . In the CTRNN scenario, the performance of all the three approaches is lower compared to the other simulation paradigms for different level of η and α , yet TPC has higher CS compared to the other methods over the different parameter values.

We demonstrate the connectivity weights obtained by TPC and inferred nature of connections, whether

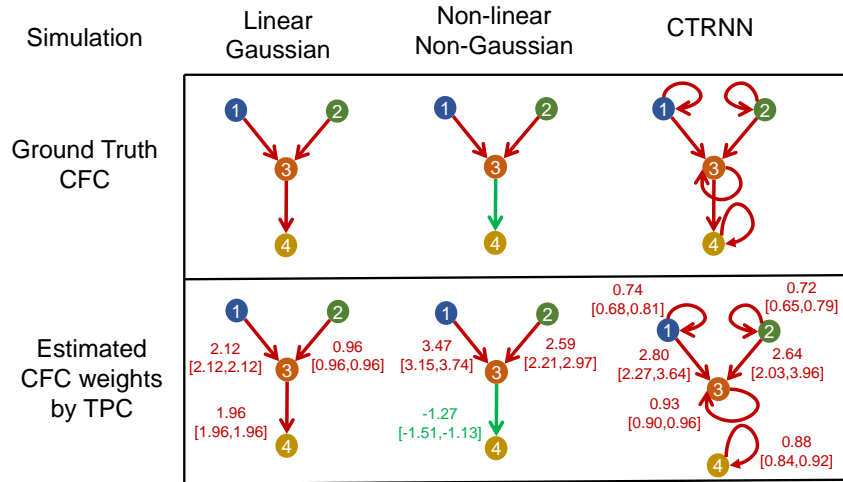


Figure 3.5 Interventional connectivity weights. Inference of interventional connectivity weights by the TPC algorithm with max delay 1 msec for the example motifs from the three simulation paradigms: Linear Gaussian VAR, Non-linear Non-Gaussian VAR, CTRNN (left to right). Top row: Ground Truth CFC with excitatory (red) and inhibitory (green) connections; Bottom row: Estimated CFC labeled with edge weights (median [min,max] over all instances) and inferred nature whether excitatory (red) or inhibitory (green).

excitatory or inhibitory, across simulations for noise level $\eta = 1$ and thresholding parameter $\alpha = 0.05$ in Figure 3.5. In the *Linear Gaussian scenario*, the estimated connectivity weight of $1 \rightarrow 3$, $2 \rightarrow 3$, $3 \rightarrow 4$ are 2.12, 0.96, 1.96, across simulation trials. Since the weights are positive, thereby the connections are labeled to be excitatory in all simulation trials, which agrees with the Ground Truth. In the *Non-linear Non-Gaussian scenario*, the estimated weight for the connection $2 \rightarrow 3$ is -1.27 in median and ranges between $-1.51, -1.13$. Therefore the weight is always negative and labeled inhibitory in the simulation trials. The weight for $1 \rightarrow 3$, $2 \rightarrow 3$ are 3.47, 2.57 in median and ranges between 3.15, 3.74 and 2.21, 2.97 respectively. Their weights are always positive and labeled excitatory in all the simulation trials. These labels for the nature of connections obtained by TPC agrees with the ground truth. In the *CTRNN scenario*, the estimated connectivity weight for $1 \rightarrow 1$, $1 \rightarrow 3$, $2 \rightarrow 2$, $2 \rightarrow 3$, $3 \rightarrow 3$, $4 \rightarrow 4$ are 0.74, 2.80, 0.72, 2.64, 0.93, 0.88 in median respectively and ranges over positive values in all the simulation trials. Thereby the connections are labeled excitatory in all the simulation trials, which agrees with the Ground Truth.

Combined [True, 1-False] Rates (%)		
Algorithm	River-Runoff (Real)	Logistic Map (Synthetic)
GC	37 [45, 92]	79 [86, 93]
PCMCI	45 [<u>100</u> , 45]	86 [89, 97]
selVAR	54 [91, 63]	87 [88, <u>99</u>]
DPGM	60 [64, <u>96</u>]	19 [27, 92]
TPC (Our)	72(+12%) [<u>100</u> , 72]	84(-3%) [<u>90</u> , 94]

Table 3.1 Comparison of CFC inference by GC, DPGM, PCMCI-GPDC and selVAR, and TPC on benchmarking datasets. For each dataset, each method’s Combined Score, True Positive Rate, and 1-False Positive Rate are reported (Higher value is better).

3.2.2 Application to Benchmark Data

We applied TPC to find the CFC from datasets in the public benchmarking platform - *CauseMe* [187], and compared with benchmarked approaches in the platform in their performance to recover causal interactions present in the datasets. We used the *River Runoff* (real data) and *Logistic Map* (synthetic data) benchmarking datasets (See Appendix A.4). We compare the approaches PCMCI-GPDC [188], selVAR [189] - which are among the top of the leaderboard for performance on the benchmarking datasets, GC, DPGM (estimated by the PC algorithm), and TPC (Our).

As previously, we measured the performance of the algorithms using 1 - False Positive Rate (IFPR), True Positive Rate (TPR) and Combined Score given by Youden’s Index ($CS = TPR - FPR$) (See Table 3.1). The River Runoff dataset is comprised of contemporaneous interactions and is expected to demonstrate the performance of the methods in an empirical setting. The Logistic Map dataset is synthetic and excludes contemporaneous interactions and shows the performance of the methods when the ground truth connectivities are specifically controlled for.

In terms of CS, TPC has recorded the best performance with a score of 72%, followed by DPGM, selVAR, PCMCI-GPDC and GC at 60%, 54%, 45%, 37% respectively. TPC exceeds the second best approach by **12%**. In terms of TPR, TPC and PCMCI have the highest scores at 100%, followed by selVAR, DPGM and GC at 91%, 64%, 45% respectively. In terms of IFPR, DPGM has the best performance with a score of 96% closely followed by GC at 92%, TPC at 72%, and selVAR and PCMCI with 63%, 45%. DPGM and GC turn out to have higher IFPR than TPC because they detect fewer false edges, but that is achieved by their detection of fewer edges altogether including fewer true edges leading to a low TPR. In contrast, TPC has

greater sensitivity in detecting edges, whose benefit is that TPC detects all the true edges correctly leading to 100% TPR. In terms of both TPR and IFPR, TPC maintains a better trade-off, and thereby a better CS compared to other methods.

For the logistic map dataset, in terms of CS, selVAR, PCMCI-GPDC and TPC have scores of 87%, 86%, 84% respectively, followed by GC with 79% and lastly DPGM with 19%. In terms of TPR, TPC has the highest score of 90%, followed by PCMCI, selVAR, GC and TPC with TPR of 89%, 88%, 86%, 79% respectively, with DPGM having comparatively lowest TPR of 27%. In terms of IFPR, all the approaches have a score of at least 90%. Thereby, TPC achieves a high CS of 84%, short of 3% from the best CS by selVAR of 87%.

The results indicate that in the real benchmark dataset of River-Runoff, TPC outperforms all methods by a substantial gap, whereas, in the synthetic benchmark dataset of Logistic Map, TPC has Combined Score of 84%, being in the top group of 84 – 87% CS performance. While selVAR and PCMCI achieve a CS of 87% and 86% respectively in the synthetic dataset, they achieve a low CS of 54% and 45% in the real dataset. Since the synthetic dataset is generated by a model controlling coupling strength between variables, low noise and devoid of contemporaneous interactions, thereby most of the methods including TPC perform fairly well in the range of 80%. In contrast, in the real dataset, the coupling between variables as well as noise are not controlled and contemporaneous interactions are expected to be present as the sampling resolution is greater than the time taken for interactions between the variables. Thereby, the real dataset provides a challenge for the methods where TPC outperforms other approaches with a CS and shows significant improvement in performance than the other methods. Presence of TPC in top group of performance for both benchmarks indicates the generality and applicability of TPC to various scenarios.

3.2.3 Application to Neurobiological Data

We proceed and test the methods on neural data consisting of electrophysiological recordings in the Visual Coding Neuropixels dataset of the Allen Brain Observatory [190; 191]. We compare TPC with Granger Causality (GC) and Sparse Partial Correlation, that are popular methods for obtaining CFC and Associative Functional Connectivity (AFC) from electrophysiological neural recordings. The dataset consists of sorted spike trains and local field potentials recorded simultaneously from up to six cortical visual areas, hippocampus, thalamus, and other adjacent structures of mice, while the mice passively view a stimuli shown

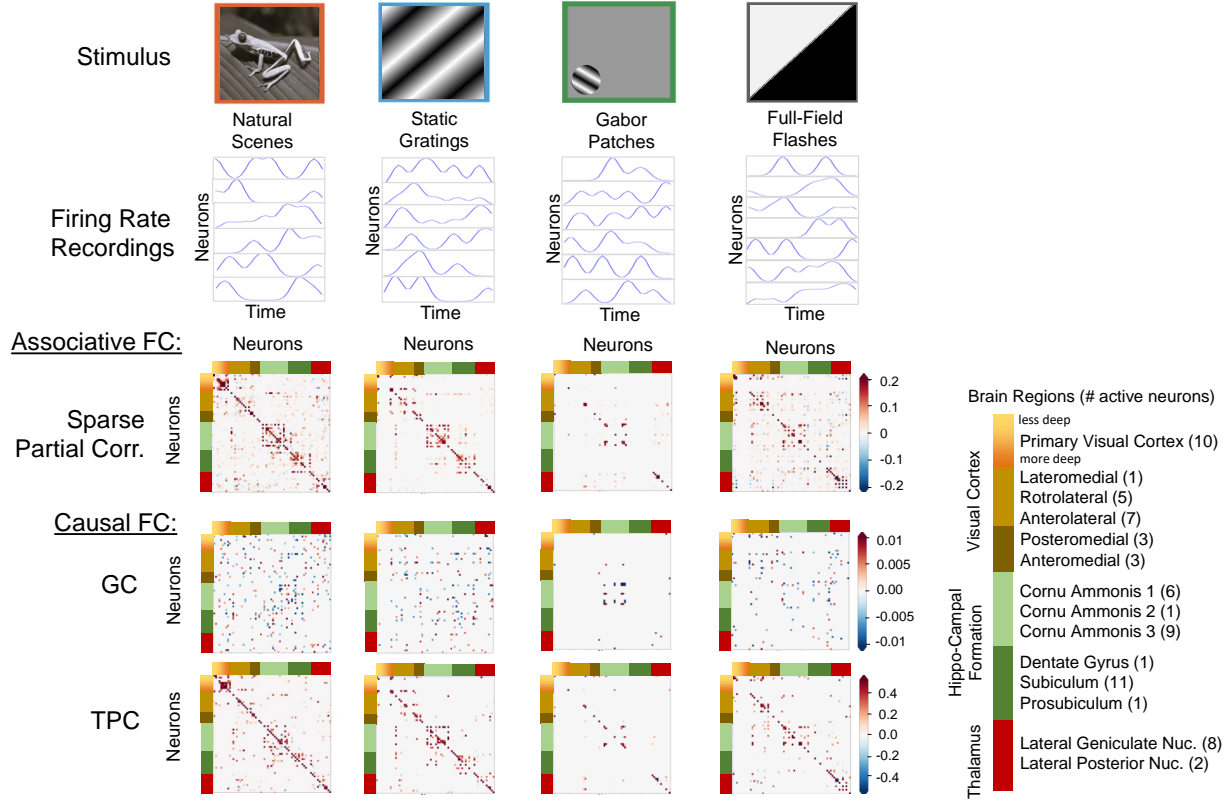


Figure 3.6 Application to Neuropixels dataset. Comparison and demonstration of the FC inferred for a benchmark of mice brain data from the Allen Institute’s Neuropixels dataset, by three methods for FC inference: Associative FC using Sparse Partial Correlation, and Causal FC using GC and TPC. The estimated FC is represented by its adjacency matrix with edge weights, which is symmetric for Associative FC and asymmetric for Causal FC. The mice were subject to different stimuli, among which we selected four stimuli categories with distinct characteristics: Natural Scenes, Static Gratings, Gabor Patches and Full-Field Flashes. The neurons are clustered by the region of brain: Visual Cortex, Hippo-Campal Formation, and Thalamus, which are further divided into sub-regions. In the adjacency matrices, a non-zero entry in (i, j) represents the connection of neuron $i \rightarrow j$.

to them. The stimuli include static gratings, drifting gratings, natural scenes/images and natural movies, which are shown to the mice with repetitions. The data has been recorded from the neurons with the recently developed technology of Neuropixels which allows real-time recording from hundreds of neurons across the brain simultaneously by inserting multiple probes into the brain [192]. Details of the dataset are in Appendix A.5.

Figure 3.6 shows the adjacency matrices for the FC obtained by the methods for one trial in each of the stimuli categories. For each stimuli categories, the AFC constitutes a distinct pattern of connectivity among the neurons. We quantify the similarity between a pair of adjacency matrices A and B by element-wise Pearson’s correlation between A and B , i.e. Pearson’s correlation between the vectorized form of the matrices

A and *B*. The CFC obtained by GC has a correlation of -0.02, -0.01, -0.03, -0.02 with the AFC for natural scenes, static gratings, Gabor patches and flashes respectively. In contrast, the CFC obtained by TPC has a correlation of 0.89, 0.88, 0.92, 0.90 with the AFC for natural scenes, static gratings, Gabor patches, and flashes respectively. It is expected that the CFC will be a directed subgraph of the AFC and be consistent with the overall patterns present in the AFC [193; 194]. However, the patterns present in the CFC obtained by GC do not match with the AFC as indicated by low element-wise correlations between the matrices. In contrast, the overall patterns present in the CFC obtained by TPC indeed match with the AFC as indicated by high element-wise correlation between the matrices. On a detailed level, there are differences between TPC-CFC and AFC: TPC results in a directed graph thereby its adjacency matrix is asymmetric while AFC is an undirected graph with symmetric adjacency matrix. Furthermore, the CFC obtained by TPC includes self-loops represented by the diagonals of the adjacency matrix and results in a sparse matrix devoid of noise since the connections passed conditional independence tests and subsampling stability thresholds. In the CFC obtained by TPC (see Figure 3.6), a greater extent of connectivity within the active neurons in Primary Visual Cortex is evoked by natural scenes, in Posteromedial and Anteromedial Visual Cortex by static gratings, in Anterolateral Visual Cortex and Thalamus by full-field flashes, compared to other stimuli. All four stimuli exhibit distinct patterns of connectivity in the Cornu Ammonis regions of the Hippo-Campal Formation. Natural Scenes and Static gratings evoke more prominent connectivity within the Subiculum compared to other stimuli.

Graphical Comparison of Estimated CFC over Stimuli

To study the differences in functional connectivity between the stimuli categories, we investigate the topological patterns in the CFC estimated above. The topological patterns can be summarized by graph theoretic measures [195; 196; 197], as follows. The graph measures were computed using the *Networkx* Python library [198], over different trials of each stimuli.

- Betweenness centrality: the fraction of all shortest paths that pass through a node, averaged over nodes, indicating the average effect of individual nodes on information flow among the remaining network's nodes.
- Transitivity: the fraction of all possible triangles present in the graph, indicating prevalence of clustered

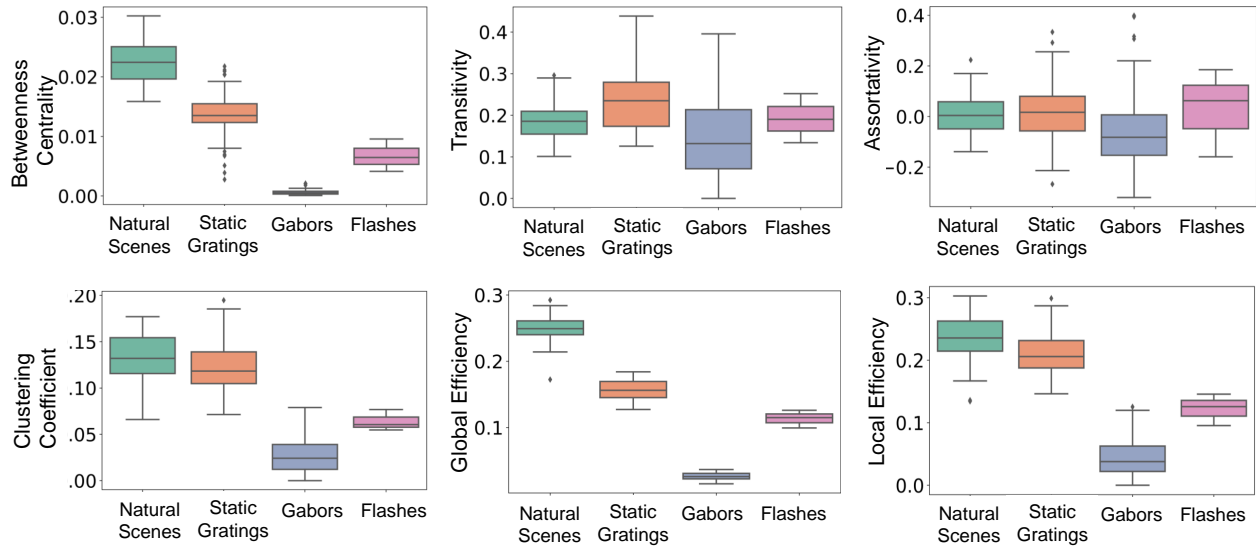


Figure 3.7 Graphical comparison of estimated CFC over stimuli. This figure compares the distribution of graph measures of CFC obtained by TPC over different stimuli: natural scenes, static gratings, Gabor patches and flashes. The distribution for each graph measure and stimuli is shown by a boxplot.

connectivity.

- Assortativity: measures the similarity of connections in the graph with respect to the node degree.
- Clustering Coefficient: the average of all clustering coefficients in the network, reflecting the tightness of connections between nodes.
- Global Efficiency: average inverse shortest path length, reflecting node’s ability to propagate information with other nodes in the graph.
- Local Efficiency: measures the global efficiency for the neighborhood of a node, averaged over nodes, indicating efficiency of transmitting information by nodes with their neighborhood in the graph.

The results of the graph measures for different stimuli are summarized by boxplots. The boxplot of a graph measure (e.g. betweenness centrality) for a stimulus (e.g. natural scenes) shows the distribution of the values of the graph measure over trials for that stimulus, with the top and bottom of the box indicating the upper and lower quartiles of the distribution, the middle of the box indicating the median, while the whiskers extend to show the rest of the distribution excluding outliers, which are marked by points (See Figure 3.7).

In Figure 3.7, in terms of *betweenness centrality*, natural scenes have higher score compared to other stimuli, followed by static gratings, then flashes and lastly Gabor patches. This shows that during natural scenes, the active neurons have a more remarkable effect on the neural information flow, compared to other stimuli, while Gabor patches have the least remarkable effect. In terms of *transitivity*, the scores between the stimuli are close, though static gratings have a relatively higher transitivity, followed by natural scenes, flashes and Gabors. This indicates that static gratings evoked a relatively higher prevalence of clustered connectivity, followed by natural scenes, flashes and lastly Gabor patches. In terms of *assortativity*, the stimuli did not evoke a comparatively distinct score as well. In terms of *clustering coefficient*, natural scenes and static gratings have higher scores compared to flashes and Gabors. This shows that, natural scenes and static gratings have the most tightness of connections between nodes in the graph, flashes have a comparatively lower tightness of connections, while Gabor patches have the least score. In terms of *global efficiency*, natural scenes have the highest score, while, static gratings have comparatively lower score, followed by flashes and Gabor patches. This shows that natural scenes evoked comparatively highest efficiency of information propagation in the CFC globally, followed by static gratings, flashes and Gabors. In terms of *local efficiency*, the trend across stimuli is similar to global efficiency, while natural scenes and static gratings evoked a more similar local efficiency compared to global efficiency. This shows that the efficiency in information propagation in local neighborhoods of neurons is higher for natural scenes and static gratings to a similar extent, but with higher efficiency compared to flashes and Gabor patches. In conclusion, betweenness centrality, clustering coefficient, global and local efficiency indicate a differential involvement of brain regions for different stimulus types.

It is noteworthy that while transitivity and assortativity are measures found on the graph as a whole, the remaining measures are found for each node and then averaged upon the nodes of the graph. So, transitivity and assortativity gives a high weight upon nodes with a high degree while the other measures give comparatively greater weight upon nodes with a low degree. This can explain the different pattern seen for transitivity and assortativity compared to the other measures in Figure 3.7. It is also implied that the greater weight for low degree nodes leads to a different distribution of the measures of betweenness centrality, clustering coefficient, global and local efficiency across different stimulus types, in contrast to similar measures for transitivity and assortativity across different stimulus types.

3.3 Discussion

In this chapter, we propose a novel methodology, the TPC Algorithm, for finding causal functional connectivity between neurons from neural time series using Directed Probabilistic Graphical Models (DPGM). In particular, we extend the applicability of DPGM to CFC inference from time series by unrolling and implementing the Directed Markov Property (DMP) to obtain the unrolled DAG reflecting causal spatial and temporal interactions. We then roll the DAG back to obtain the CFC graph. The methodology exhibits interpretability of causal interactions over time between neural entities. It also incorporates time delays in interactions between neurons as well as the presence of feedback-loops. The model and the approach are non-parametric, meaning that it does not require the specification of a parametric dynamical equation for neural activity. We show that if the neural activity obeys an arbitrary dynamical process, the Rolled CFC-DPGM agrees with respect to the causal relationships implied by the dynamical process. We determine that the Rolled CFC-DPGM is predictive of counterfactual queries such as ablation or modulation. We show that the answers can be provided by using simple causal reasoning with the edges of the rolled CFC-DPGM. We demonstrate the utilization of the methodology in obtaining CFC from simulations and compare the performance of TPC with other methods such as Granger Causality (GC) and common DPGM. Furthermore, we apply the methods to benchmarks of time-series causal inference and neurobiological dataset from mice brain presented with various visual stimuli. The results provide insights into the CFC between neurons in the mouse brain in a variety of stimuli scenario. We also compare the topological patterns in the estimated CFC between different stimuli using graph-theoretic measures.

The distinctive and useful aspect of TPC is that it takes into consideration neural interactions over time. In neuroscience literature, causality is typically referred to as “a cause of an observed neural event (the ‘effect’) as a preceding neural event whose occurrence is necessary to observe the effect” [109]. The approach of unrolling causal graphs over time and considering time-delays in TPC incorporates this definition by essentially finding whether the previous time values of the neurons impact the present value of a particular neuron by the obtained directed graph. By virtue of the Directed Markov Property, TPC incorporates causality of neural interactions in a non-parameteric, model-free manner and incorporates interventional properties. The subsampling step filters spurious connections by repeating the inference of CFC over several subsampled blocks of the time series and discarding those connections that are absent in multiple repetitions.

Pruning further selects edges by discarding edges with exceedingly low edge weight. In contrast, parametric approaches, such as GC, do not satisfy the Directed Markov Property, thereby not incorporating interventional properties. Furthermore, in comparative studies with simulated, benchmark and real neurobiological datasets, we found the performance of TPC to be better compared to other approaches. We conclude by summarizing, in Table 3.2, the differences and benefits of TPC in a comparison with other approaches including variants of the PC algorithm (DPGM) which satisfy DMP in a static data setting, and other approaches that do not obey the DMP, such as Granger Causality (GC) and Dynamic Causal Model (DCM), outlining their strengths and weaknesses with respect to several criteria of causality in functional connectomics. Indeed, capturing as many causal criteria is fundamental to any approach from statistical and application points of view.

Our exposition of properties of each approach and the comparative study show that each of the methods address different aspects of modeling causality of neural interaction and mapping them in the form of a graph [1]. The comparative table demonstrates that with respect to the model that each approach is assuming, GC requires a linear model in its common use, though recent non-linear and non-parametric extensions, have been applied. DCM requires a strict well defined mechanistic biological model and thus can only compare different models based on evidence from data. In comparison, DPGM and TPC have the advantage of not requiring modeling of the neural dynamics using a parametric equation or assumption of a linear model. While DPGM is developed for static variables, and as such cannot address temporal and contemporaneous relationships and must obey DAG architecture, TPC is suited for time series setting and extends DPGM for spatiotemporal data. TPC obtains the CFC that follows Directed Markov Property extended to include inter-temporal relationships in a time-series setting such that parent-child relations between neurons are equivalent to arbitrary functional relationships between their neural activity over time. In terms of incorporating contemporaneous interactions arising when causal interactions happen faster than sampling rate, while GC and DPGM do not, TPC is able to incorporate contemporaneous interactions by design. In terms of incorporating self-loops in neural activity, while DPGM typically produces a DAG, TPC incorporates self-loops in neural activity. In regards to guarantee of causality, GC can provide useful insights into a system's dynamical interactions in different conditions, however its causal interpretation is not guaranteed as it focuses on the predictability of future based on past observations of variables. DCM uses the parameters for coupling between hidden neural states in competing biological models to indicate CFC, however it compares hypothetical models based on evidence

from data which relevance to causality is not guaranteed (Friston et al., 2003). In summary, TPC extends DPGM to the time-series setting and provides a probabilistic foundation for causality which is predictive of the consequence of possible intervention like neuron ablation and neuromodulation.

While the unrolled graphical modeling can capture contemporaneous interactions over distinct neurons, it does not capture contemporaneous self-interactions for the same neuron. This is due to the requirement of DPGM for the unrolled graphical model to be a directed acyclic graph. Although this is a limitation, it is still an improvement from Granger Causality for which the VAR model does not allow any contemporaneous relations [199].

While TPC provides a powerful causal framework for time series, its current version relies on the PC algorithm which in turn assumes *causal sufficiency*, that is, all the causes of the input variables are present within the input variables. In terms of neurons, causal sufficiency means that neurons that have a causal influence on the observed neurons are also observed. There is no other requirement on the number of neurons. However, causal sufficiency in real neural data is challenging to achieve since it is often the case that there are uncaptured neurons which may causally influence the observed neurons [181; 200]. In the absence of causal sufficiency, Step 3 of TPC involving PC algorithm would estimate the true causal edges between the observed neurons but also a spurious edge between a pair of observed neurons whenever an unobserved neuron is a common cause to both of the observed neurons in the pair [201; 202]. These are partly remediated by the TPC algorithm by discarding spurious edges in the Subsampling and Pruning steps, but some of spurious edges may remain. To address this, an alternative strategy that could be considered is replacing the PC algorithm with the FCI algorithm in Step 3 of TPC (1) since the FCI algorithm exhibits statistical consistency, that converges in probability to the true causal relationships, given i.i.d. samples in the presence of latent confounders [163]. However, the FCI algorithm identifies indirect connections only and not direct connections, thereby would result in a CFC with indirect causal connections.

Under the assumptions of 1) causal sufficiency, 2) faithfulness, and 3) given i.i.d. samples, the PC algorithm's estimated causal DAG is statistically consistent. In light of these assumptions, the step 2 of the TPC algorithm constructs samples with a time-advance of $2(\tau + 1)$ between samples from a subsampling window, for max time-delay of interaction τ , which are used as an input to the PC algorithm in the next step. The time-gap between samples of $2(\tau + 1)$ reduces between-sample dependence. And, the samples being

constructed from a short subsampling window instead of the entire time series aids to make their distribution fairly identical. Yet, specifications of τ of lower value could lead to between-sample dependence and rapid perturbations to the time series in a subsampling window can lead to distribution changes between samples in a subsampling window, leading to potential reduction in efficacy of edge detection by the PC algorithm. To improve overall efficacy in such scenarios and curb spurious edge detections, the subsampling step of the TPC algorithm outputs many CFCs over random time windows and preserves only stable edges over the set of CFCs. In the pruning step, the edges in the CFC are pruned if having an exceedingly low connectivity weight. The mathematical guarantees of the TPC algorithm has been treated in Biswas and Mukherjee [3].

The primary tuning parameters or hyperparameters in TPC are the significance level of the conditional dependence tests (α) and the maximum time-delay of interaction (τ). Due to subsampling, it has other tuning parameters: window width (L), number of iterations (m) and subsampling stability cutoff (γ). It is noteworthy that the subsampling procedure implemented is the stability selection procedure as discussed in Maathuis et al. [184] for tuning the α parameter in the PC algorithm, and introduced in Meinshausen and Bühlmann [203]. It is shown in [203] that stability selection ensures that a certain familywise type I error rate in multiple testing is conservatively controlled for finite sample size, and provides a formula for determining the stability cut-off γ via the multiple testing error rate. Among the stability selection parameters, it has been studied in Meinshausen and Bühlmann [203] that the influence of the cut-off γ is very small and for values in the range of $\gamma \in (0.6, 0.9)$, results tend to be very similar. The parameter τ , representing the maximum time delay of direct interaction between variables over time, depends on the application and should be chosen according to the maximum time delay of direct interactions expected in the complex system. In practice, τ should be larger or equal to the true value of the maximum time delay. One also needs to choose the kind of conditional dependence test, such as test based on Fisher's z-transform of partial correlation that is applicable to data having a Gaussian distribution [57], or test based on Hilber-Schmidt criterion that is applicable to Gaussian as well as non-Gaussian distributions [3].

For completeness of exposition of the comparative performance of TPC, we also compare its performance with a graph neural network (GNN) based approach developed by Yu et al. [204] (coined DAG-GNN; DGNN in short). DGNN infers the DAG from a generalized version of the linear structural equation model, leveraging a continuous characterization to the combinatorial DAG constraint in linear structural equation model, making

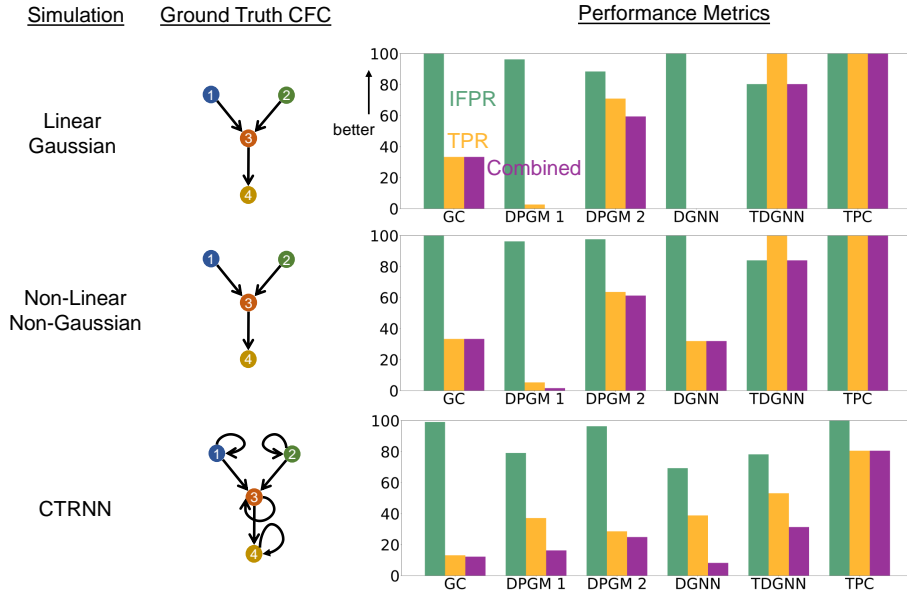


Figure 3.8 Comparison of performance with Graph Neural Networks. This figure compares the performance of GC, DPGM 1, DPGM 2, DGNN, TDGNN and TPC in terms of IFPR (green), TPR (orange) and Combined Score (purple) on three simulation settings; from top to bottom: Linear Gaussian, Non-linear Non-Gaussian and CTRNN.

it amenable to gradient descent-based parameter estimation [205]. However, DGNN is applicable for independent data instead of time series akin to PC algorithm, and does not incorporate across-time causal interactions akin to naive application of the PC algorithm to time series data (DPGM 1; See section 2.2.5). Thereby, DGNN suffers in performance in the context of time series data. We considered an extension of DGNN for time series data (coined TDGNN) by applying DGNN on time advanced samples as defined in Step 1 of TPC to obtain the unrolled DAG and then rolling back to obtain the rolled graph (see Figure 3.2 and Section 3.1.3). By virtue of this extension, TDGNN incorporates across-time interactions of variables within each time-advanced sample while still assuming independence between different time-advanced samples. In Figure 3.8, we compared the performance of TPC and other approaches with DGNN and TGNN in the simulated datasets of Figure 3.3 in Section 3.2.1. In terms of Combined Score, TPC has a better performance than both DGNN and TGNN across the simulation settings. TDGNN has relatively better performance than DGNN since TDGNN incorporates across-time causal interactions while DGNN does not. TDGNN's lower performance than TPC on the simulated benchmarks can be because TDGNN assumes a generalized version of linear structural equation model and independence in sampling of the variables, whereas TPC neither assumes any specific functional form for the dynamical system nor does it assume the independence

of samples. A detailed analysis of TDGNN for causal inference in time series is kept for future study.

In conclusion, TPC provides a probabilistic and interpretable formulation for CFC modeling and inference in the context of neural time series. We have established the statistical properties of the model as well as demonstrated its performance in estimation of CFC. We have demonstrated TPC application in continuous time series datasets, however TPC is similarly applicable to discrete time series datasets by simply using a statistical conditional independence test for discrete data in the algorithm. This can be especially relevant for count datasets such as spiking neuron datasets.

3.4 Materials and Methods

3.4.1 Causal Functional Connectivity for Static Variables - Review

In this section, we provide a concise summary of DPGM for finding CFC for static variables, i.e. variables with **i.i.d. samples**, and extend in later sections to incorporate temporal dependence in time series. Let us consider a brain network $V = \{v_1, \dots, v_N\}$ with N neurons labeled as v_1, \dots, v_N and $X_v(t) \in \mathbb{R}$ denote a random variable measuring the activity of neuron v at time t . Examples for such variables are instantaneous membrane potential, instantaneous firing rate, etc. Let Y_v denote a scalar-valued random variable corresponding to $v \in V$, e.g., the neural recording at time t : $Y_v = X_v(t)$, average of recordings over time $Y_v = \bar{X}(v)$, and for a set of neurons $A \subset V$, \mathbf{Y}_A denotes the random vector $(Y_v, v \in A)$. Let $G = (V, E)$ denote a *directed acyclic graph* (DAG), i.e., a directed graph without directed cycles, over the neurons in V and with directed edges E . Nodes u and $v \in V$ are said to be *adjacent* if $v \rightarrow u \in E$ or $u \rightarrow v \in E$. A *path* is a sequence of distinct nodes in which successive nodes are adjacent. For a path $\pi = (v_0, \dots, v_k)$, if every edge of π is of the form $v_{i-1} \rightarrow v_i$ then v_0 is an *ancestor* of v_k and v_k is a *descendant* of v_0 . The set of *non-descendants* of v , denoted $nd_G(v)$, contains nodes $u \in V \setminus \{v\}$ that are not descendants of v . The set of *parents* of $v \in V$ is denoted as $pa_G(v) = \{u \in V : u \rightarrow v \in E\}$. We mark the set $nd_G(v) \setminus pa_G(v)$ as the set that contains all nodes which are older ancestors of v before its parents [27; 28].

With these notations, we highlight the *Directed Markov Property* (DMP) and its functional equivalence for DPGM. The DMP connects probabilistic conditional independencies between nodes of a directed graph

with relationships of causal influence by ensuring that the influence of each node's ancestors beyond parents reaches to the node exclusively via its parents. Furthermore, the functional equivalence for DPGM shows that, the edges in a DPGM satisfying the DMP are in agreement with causal functional interactions among the nodes.

Directed Markov Property (DMP) $(Y_v, v \in V)$ is said to satisfy the *Directed Markov Property* with respect to the DAG G if and only if,

$$Y_v \perp\!\!\!\perp \mathbf{Y}_{nd_G(v) \setminus pa_G(v)} \mid \mathbf{Y}_{pa_G(v)} \quad (3.1)$$

The DMP translates the edges in the DAG into conditional independencies, such that for each node v and its older ancestors $nd_G(v) \setminus pa_G(v)$, Y_v is conditionally independent of $\mathbf{Y}_{nd_G(v) \setminus pa_G(v)}$ given the variables corresponding to the parents of v , $\mathbf{Y}_{pa_G(v)}$. The DMP can be equivalently represented with functional relationships between parents and child instead of conditional independencies, which is described in the following theorem [162].

Functional Equivalence of DMP If Y_v satisfies

$$Y_v = g_v(\mathbf{Y}_{pa_G(v)}, \epsilon_v), v \in V \quad (3.2)$$

where ϵ_v are independent random variables, and g_v are measurable functions for nodes $v \in V$, then $Y_v, v \in V$ satisfies the Directed Markov Property with respect to the DAG G . Conversely, if $Y_v, v \in V$ satisfies the Directed Markov Property with respect to G , then there are independent random variables ϵ_v and measurable functions g_v for which Eq. (3.2) holds. This shows that if $Y_v, v \in V$ satisfies the DMP with respect to the DAG G , then G admits a natural causal interpretation, due to its functional equivalence: parent nodes of v in G causally influence the child node v [87].

Peter & Clark (PC) algorithm Let $Y_v, v \in V$ satisfy the DMP with respect to the DAG G . The PC algorithm is a popular method to infer G from observed data [52]. The PC algorithm uses a consistent statistical test, such as Fisher's Z-transform when $Y_v, v \in V$ are Gaussian variables, and kernel and distance

based tests for non-Gaussian variables [57; 166]. The algorithm first represents the observed variables by nodes of a graph and starts with an empty set of edges and puts an undirected edge between each pair of nodes if they are not independent or conditionally independent given any other variable(s) determined by the statistical test at level α . This results in the undirected skeleton graph, which is then converted into a DAG by directing the undirected edges using rules for orientation. The PC algorithm estimates several DAGs \hat{G}_i based on i.i.d. samples of $Y_v, v \in V$, and outputs a single completed partially directed acyclic graph (CPDAG) \hat{G} defined as follows: \hat{G} has a directed edge from node $v \rightarrow w$ if $v \rightarrow w$ is present in all the DAGs \hat{G}_i . It has an undirected edge between v and w if either directions between them are present among the DAGs \hat{G}_i . It has no edge between v and w if no edge is present between them in any of the DAGs \hat{G}_i . The CPDAG \hat{G} is uniquely identifiable from observed data.

The PC algorithm assumes *causal sufficiency* of the input variables: A set V of variables is causally sufficient for a population if and only if in the population every common cause of any two or more variables in V is in V , or has the same value for all units in the population. Another method, the FCI algorithm, is applicable when causal sufficiency does not hold [163]. However, while the PC algorithm can identify direct causes, FCI algorithm cannot distinguish between direct and indirect causes.

Let P denote the probability distribution of $Y_v, v \in V$. The PC algorithm also assumes *faithfulness* of the DAG G to $Y_v, v \in V$: if the DMP with respect to G encompasses all the conditional independence relations due to P and does not have any conditional independence relation that is not given by P , G is said to be *faithful* to $Y_v, v \in V$. Using a consistent statistical test for conditional independence, and assuming causal sufficiency and faithfulness, the PC algorithm estimate, \hat{G} , is consistent for G ; that is, \hat{G} converges in probability to G with increasing number of samples in data [206; 52].

3.4.2 Choice of gap between time-advanced samples in TPC algorithm

The TPC algorithm considers a choice of $2(\tau + 1)$ for gap between time-advanced samples to help reduce between-sample dependence for the time-advanced samples. To justify the choice of $2(\tau + 1)$, let the first two time advanced samples with a gap of r be $\chi_0 = (\mathbf{X}_0, \dots, \mathbf{X}_\tau)$ and $\chi_1 = (\mathbf{X}_r, \dots, \mathbf{X}_{\tau+r})$. If $r < 2(\tau + 1)$ then $r - \tau \leq \tau$, and so time $t = \tau$ and time $t = r$ are within the maximum time delay of interaction, thereby \mathbf{X}_τ and \mathbf{X}_r can have a dependence between them, which implies a dependence between χ_0 and χ_1 . If

$r \geq 2(\tau + 1)$ then $r - \tau > \tau$ so time $t = \tau$ and time $t = r$ have a gap greater than the maximum delay of interaction τ , therefore, \mathbf{X}_τ and \mathbf{X}_r will be independent. This justifies the choice for the gap r to be $2(\tau + 1)$.

3.4.3 Connectivity weights in the CFC for time series

In this section, we define the connectivity weights obtained by TPC. Connectivity weights refer to a weight w_{uv} for connections $u \rightarrow v$ in the CFC graph. We consider *interventional causal effects* to define connectivity weights. Interventional causal effects quantify how much effect an intervention applied to neuron u will have on neuron v .

Definition 2: Interventional causal effects in Unrolled DAG. Let \mathbf{X} satisfy the Directed Markov Property with respect to $\mathbf{G} = (\mathbf{V}, \mathbf{E})$. The *interventional causal effect* of $X_u(t) = x_{u,t}$ on $X_v(t')$, where $x_{u,t}$ are fixed values, for $u, v \in V$, and $t, t' \in \{0, 1, \dots, T\}$ with $(u, t) \neq (v, t')$, is defined in interventional calculus by Pearl et. al as follows [48; 207; 184]

$$\frac{\partial}{\partial x} E(X_v(t') | X_u(t) = x) |_{x=x_{u,t}} \quad (3.3)$$

where, for a random variable Y , $E(Y)$ denotes the expectation of Y .

Assuming $X_u(t)$, $u \in V, t \in \{0, 1, \dots, T\}$ are jointly Gaussian, the causal effect does not depend on the value of $x_{u,t}$, and the causal effect of $X_u(t)$ on $X_v(t')$ from Eq. (3.3) takes the following form,

$$w_{u,t}^{v,t'} = \begin{cases} 0, & \text{if } (v, t') \in pa_{\mathbf{G}}((u, t)), \\ \text{coefficient of } X_u(t) \text{ in } X_v(t') \sim X_u(t) + X_{pa_{\mathbf{G}}((u,t))} & \text{if } (v, t') \notin pa_{\mathbf{G}}((u, t)) \end{cases} \quad (3.4)$$

where $X_v(t') \sim X_u(t) + X_{pa_{\mathbf{G}}((u,t))}$ is shorthand for linear regression of $X_v(t')$ on $X_u(t)$ and $X_{pa_{\mathbf{G}}((u,t))} = \{X_a(b) : (a, b) \in pa_{\mathbf{G}}((u, t))\}$. Note that (u, t) is a node of \mathbf{G} so $pa_{\mathbf{G}}((u, t))$ denotes the set of parents of (u, t) in \mathbf{G} as defined in Section 3.4.1.

Note that when \mathbf{X} satisfy the Directed Markov Property with respect to $\mathbf{G} = (\mathbf{V}, \mathbf{E})$, the causal effects are defined in interventional calculus literature for all pairs of nodes in \mathbf{V} and not that only for those pairs which are adjacent. And, if $u \rightarrow v \in \mathbf{E}$ then the interventional causal effect from v to u is 0 [184].

Under the Gaussian assumption, we define the *interventional causal effect* of the activity of neuron u at

time t on the activity of neuron v at time t' to be $w_{u,t}^{v,t'}$ for $t < t'$ where $t, t' \in \{0, 1, \dots, T\}$, $u, v \in V$. Using this, we define weights for the connection from u to v for u, v in the rolled CFC-DPGM F_τ , following the way F_τ is defined from G .

Definition 3: Interventional connectivity weights in Rolled CFC-DPGM. Let X satisfy DMP with respect to the DAG $G = (V, E)$ and F_τ is the Rolled CFC-DPGM with max delay τ . If neurons u, v are connected as $u \rightarrow v$ in F_τ , then, the *weight of connection* from neuron u to v with max delay τ , denoted by $w_\tau(u, v)$, is defined as the average of the causal effects: $w_{u,t}^{v,t'}$ for $(u, t) \rightarrow (v, t') \in E$, $t \leq t' \leq t + \tau$, $t, t' \in \{0, 1, \dots, T\}$.

Connectivity weights in TPC Algorithm. After the CFC graph F_τ^b is obtained in Step 3-5a in TPC algorithm, the interventional connectivity weights for connections in F_τ^b are obtained in Step 5b to define the connectivity weights $w_\tau^b(u, v)$ for connections $u \rightarrow v \in F_\tau^b$. Then subsampling in Step 6 ensures greater stability of the estimated connectivity weights. Step 6 outputs a single connectivity weight $w_\tau(u, v)$ for connections $u \rightarrow v$ in F_τ , as the average of $\{w_\tau^b(u, v) : u \rightarrow v \in F_\tau, u \rightarrow v \in F_\tau^b, b \in 1, \dots, m\}$ when the set is non-empty and 0 otherwise. Therefore, this finds a connectivity weight for the edge $u \rightarrow v \in F_\tau$ by taking the average of connectivity weight of the edge $u \rightarrow v \in F_\tau^b$ whenever it exists over b .

Pruning by Connectivity Weights. After the rolled CFC-DPGM and Connectivity Weights have been inferred by the TPC Algorithm, spurious connections can be pruned further in Step 7 of TPC (1), by discarding those connections whose connectivity weight is less than a threshold. For this threshold, we use a factor of 10 of the maximum Connectivity Weight in the rolled CFC-DPGM.

From the interpretation of regression coefficients in Eq. (3.4), a negative connectivity weight (3.4.3) from neuron $u \rightarrow v$ in F_τ indicates an inhibitory connection, in which, increased activity $X_u(t)$ of the pre-synaptic neuron u at time t causes subjugation of activity $X_v(t')$ of the post-synaptic neuron v at a following time t' , when activity of the neurons that are causally connected to neuron u at time t , are kept fixed. In a similar manner, a positive FC weight from neuron $u \rightarrow v$ in F_τ indicates an excitatory connection. In this way the strength of the functional connection, which also indicates it's excitatory and inhibitory nature is learnt from the data.

3.4.4 Properties of Rolled CFC-DPGM

We highlight properties of the Rolled CFC-DPGM, obtained by TPC, in capturing causal relationships in neural dynamics in a non-parametric manner and being predictive of the impact of counterfactual interventions to the neurons in the CFC.

Well-defined. The transformation from unrolled DAG G to Rolled CFC, F_τ , is a well-defined function, meaning that starting from the same unrolled DAG G we will not have multiple possible CFC and there will be a unique CFC F_τ .

Proof: By contradiction, consider two distinct CFCs $F_{\tau,1}$ and $F_{\tau,2}$ with $F_{\tau,1} \neq F_{\tau,2}$ obtained from the unrolled DAG $G = (\mathbf{V}, \mathbf{E})$. Since $F_{\tau,1} \neq F_{\tau,2}$, so $\exists i, j \in V$ such that $i \rightarrow j \in F_{\tau,1}$ but $i \rightarrow j \notin F_{\tau,2}$, where i, j could be either the same or distinct neurons. Using the definition 2 of Rolled CFC-DPGM, $i \rightarrow j \in F_{\tau,1}$ implies that for some $0 \leq t_1 \leq t_2 \leq t_1 + \tau$, $(i, t_1) \rightarrow (j, t_2) \in \mathbf{E}$. But $i \rightarrow j \notin F_{\tau,2}$ contradicts this as it implies that $(i, t_1) \not\rightarrow (j, t_2) \in \mathbf{E}$ for any $0 \leq t_1 \leq t_2 \leq t_1 + \tau$.

Non-parametric Causal Relations. The following theorem shows that the CFC given by the model agrees with the ground truth causal relationships between neural activity at different time without requiring any assumptions on the functional form of the relationships. That is, we show that if past time points of neurons in $A_v \subset V$ influence the present time point of $v \in V$ by an arbitrary function with independent random noise, then neurons in A_v are connected to the neuron v in their Rolled CFC-DPGM. This means that causal relationships among the neurons, in their unknown arbitrary dynamical equation, are accurately represented by the CFC without prior knowledge of the functional form of the relationships. The benefit of Rolled CFC-DPGM is that it uses a non-parametric graphical model for the temporal relationships between neurons and does not assume a parametric equation for the temporal relationships. Furthermore, the Rolled CFC-DPGM provides a framework to answer causal questions related to the consequence of interventions and counterfactuals.

Theorem 3.4.1 (Agreement with Time Series Causal Relations). For neurons $v \in V$ with activity $X_v(t)$ at time $t \in \{0, 1, \dots, T\}$, if $X_v(t)$ satisfies

$$X_v(t) = g_{v,t}(X_{u_{v,1}}(t_{v,1}), \dots, X_{u_{v,p}}(t_{v,p}), \epsilon_v(t)), \quad (3.5)$$

for some $p \geq 1$, nodes $u_{v,i} \in V$, times $t_{v,i} \in \{t - \tau, t - \tau + 1, \dots, t\}$ with either $u_{v,i} \neq v$ or $t_{v,i} \neq t$, $i \in \{1, \dots, p\}$, where τ is the maximum time-delay of interaction, and $g_{v,t}$ is a measurable function and $\epsilon_v(t)$ are independent random variables, then the graph F_τ with nodes V and parents of v , $pa_{F_\tau}(v)$, given by

$$pa_{F_\tau}(v) = \{u_{v,1}, \dots, u_{v,p}\}$$

is the Rolled CFC-DPGM between the neurons in V .

Proof. See Appendix A.1. □

According to Theorem 1, the edges in the Rolled CFC-DPGM corresponds to actual causal relations in the neural time series. TPC takes steps to estimate the Rolled CFC-DPGM, and reports good performance in simulations in recovering the true CFC in comparison to GC and DPGM.

TPC is a non-parametric approach assuming DMP with respect to the unrolled DAG which is satisfied for arbitrary functional relationships (Theorem 1). Estimation using PC algorithm further assumes the converse of DMP, i.e. faithfulness, which is a generic assumption facilitating causal inference. These are broad assumptions which include arbitrary functional relationships and noise distributions, under which TPC is expected to perform well. Estimation using PC algorithm also assumes causal sufficiency, therefore there is a possibility for unobserved causal influence to lead to erroneous edges being outputted by TPC. This is remediated (to some extent) by excluding spurious edges in the subsampling step of TPC. Even though a non-parametric approach like TPC is expected to perform well irrespective of parametric model assumptions, parametric approaches can outperform it when the parametric model assumptions are precisely satisfied. For example, selVAR is seen to outperform TPC in the Logistic Map synthetic dataset. Also GC, which is a parametric model-based approach, assuming a stationary Gaussian vector autoregressive (VAR) model, expected to perform well when its model assumptions are satisfied. However, irrespective of whether the Gaussian VAR model assumption is satisfied or not, TPC is seen to perform at par with or better than GC in simulations. DPGM is expected to excel in performance in i.i.d. scenarios but would not be generally applicable in capturing functional relationships across-time which require setting τ in TPC to be greater than zero.

Interventional Properties The Rolled CFC-DPGM can answer questions concerning counterfactual interventions on neurons without experimentally performing the interventions, i) Ablation of a neuron A , ii) Activity of neuron B is externally modulated. In the following corollary we show how conclusions can be drawn for such queries.

Corollary 3.4.1 (Intervention). For neurons $v \in V$ following the dynamics in equation Eq. (3.5), let us consider there is an experimental or counterfactual intervention on neurons $v_1, \dots, v_k \in V$ during the set of times $I \subseteq \{0, 1, \dots, T\}$, such as ablation or external control. 1) For ablation of v_1, \dots, v_k the connections incident as well as outgoing from them are removed. 2) During $t \in I$, for external control, all connections incident on v_1, \dots, v_k are removed in the Rolled CFC-DPGM F_τ and other connections remain intact.

Proof. See Appendix A.2. □

This corollary justifies the usage of causal reasoning with the edges of the CFC alone to answer the interventional queries. Answers to the questions by causally reasoning are as follows:

i) When neuron A is ablated, one just deletes all the edges incident and originating from neuron A since A has a fixed value after ablation and neither do other neurons influence the activity of neuron A nor does A influence the activity of any other neuron.

ii) When activity of neuron B is externally controlled, one simply removes the edges incident on neuron B , because activity of neuron B no longer depends on its parent neurons in the CFC obtained before intervention rather the activity of neuron B depends on the external control. Edges originating from neuron B in the CFC from before the intervention should remain intact during the intervention since the functional pathways from B to its descendant neurons in the CFC remain intact during external control. To illustrate the properties of non-parametric causal relations (Theorem 3.4.1) and interventions (Corollary 3.4.1), we consider the following example.

Example 3.4.1. Let V denote a network of 4 neurons, labeled $\{1, 2, 3, 4\}$ with neural activity $X_v(t), v \in V$




Figure 3.9 Impact of intervention. Rolled CFC-DPGM (left) for neurons 1-4 with dynamics as in Example 3.4.1, and consequence of intervention on neurons labelled A and B by (i) Ablation of A and (ii) External modulation of B.

related as,

$$X_1(t) = g_{1,t}(\epsilon_1(t))$$

$$X_2(t) = g_{2,t}(X_2(t-1), X_3(t-1), X_3(t-2), \epsilon_2(t))$$

$$X_3(t) = g_{3,t}(X_3(t-1), X_1(t-1), \epsilon_3(t))$$

$$X_4(t) = g_{4,t}(X_2(t-1), \epsilon_4(t))$$

for independent random variables $\epsilon_v(t)$ and measurable functions $g_{v,t}$, $v = 1, 2, 3, 4; 0 \leq t \leq 1000$ msec. By Theorem 3.4.1 it follows that the graph: $2 \rightarrow 2, 3 \rightarrow 2, 3 \rightarrow 3, 1 \rightarrow 3, 2 \rightarrow 4$ as in Figure 3.9-left is the CFC among the neurons considering maximum time delay of interaction to be 1 msec or higher. Suppose one asks the question of type (i), how would the functional connectome change if neuron (A) were ablated? According to Corollary 3.4.1, the resulting CFC would be $1 \rightarrow 3, 3 \rightarrow 3$ as in Figure 3.9-middle by removing the connections to and from neuron 2 according to Corollary 3.4.1. Suppose one asks the question of type (ii), how would the functional connectome change if activity of neuron (B) were to be externally controlled by optogenetics? According to Corollary 3.4.1, the resulting CFC would be $3 \rightarrow 2, 2 \rightarrow 2, 2 \rightarrow 4$ as in Figure 3.9-right by removing the parent connections of neuron 3 according to Corollary 3.4.1.

Data Availability

The software package for the methods in this chapter and example code are available in <https://github.com/shlizee/TimeAwarePC>.

Table 3.2 Comparative summary of different approaches for causal modeling.

	GC	DCM	DPGM	TPC
Form of Causality	Non-zero parameters in VAR model	Coupling parameters in biological model	Directed Markov Graph	Directed Markov Graph over time-advanced variables
Inclusion of temporal relationships	Yes	Yes	No , formulation for static variables	Yes , adapts DPGM for inter-temporal relationships
Inclusion of contemporaneous relationships	No	Yes , by a differential equation	No	Yes , if $(i, t) \rightarrow (j, t)$ then $i \rightarrow j$.
Generalizable Statistical Model	Yes	No	Yes	Yes
Non-parametric Model	Yes , parametric and non-parametric approaches exist.	No , biologically mechanistic non-linear model.	Yes , equivalent to an arbitrary functional relationship between nodes.	Yes , equivalent to an arbitrary functional relationship between neural activity at different times.
Supports CFC Inference	Yes	No , suitable for comparing model hypotheses	Yes	Yes
Cycles (including self-loops) Occuring In The Model	Yes for VAR model (neuron $i \rightarrow i$ when $A_{ii}(k) \neq 0$ for some k).	Yes ($i \rightarrow i$ when $\theta_{ii} \neq 0$)	No , it is a DAG	Yes ($i \rightarrow i$ when $(i, t) \rightarrow (i, t')$ for some $t < t'$)
Incorporation of Interventional and Counterfactual queries	No	No	Yes but for static variables.	Yes , adapts for temporal scenario, can predict the consequence on CFC of counterfactual intervention to neural activity.

Chapter 4

Mathematical Guarantees

This chapter is a reproduction of published work in Biswas and Mukherjee [3], with minor adaptations to fit the thesis structure and formatting.

Directed probabilistic graphical models are a popular tool to find causal relations between variables from their observational data. The causal relations are often represented by directed acyclic graphs (DAGs), which have its nodes as the variables and edges encoding information on conditional dependence between the variables. The conditional dependencies are specified by the directed Markov property [208].

Estimation of the DAG from independent and identically distributed (i.i.d.) data can be conducted by score-based, constraint-based or hybrid methods. Score-based methods search over the space of all possible DAGs and maximize a goodness-of-fit score such as GES (greedy equivalent search) [164] and GIES (greedy interventional equivalent search, for data with experimental interventions) [165]. Constraint-based methods first use conditional dependence tests to find an undirected skeleton graph and then use the conditional dependence information in the skeleton to partially direct the edges, such as in the PC algorithm (acyclic graph, no latent confounders, no selection bias) [52] and FCI algorithm (acyclic graph, latent confounders, selection bias) [163]. Finally, hybrid methods combine score and constraint-based approaches for example the max-min hill climbing method [209].

Among the different approaches, PC algorithm remains one of the most popular methods for causal inference [52; 210]. The PC algorithm estimates the completed partially directed acyclic graph (CPDAG) which represents the class of all DAGs which are Markov Equivalent and thereby indistinguishable from

observational data. In terms of computational time complexity, the PC algorithm is exponential (as a function of the number of nodes) in the worst case, but if the true DAG is sparse, which is often a reasonable assumption, it reduces to polynomial complexity. The question of asymptotic guarantees in estimation of the CPDAG by the PC algorithm has been studied in [52; 56; 57]. In particular, assuming only faithfulness, it has been shown that pointwise consistency of the PC algorithm can be achieved based on i.i.d. observations [52]. Furthermore, high dimensional uniform consistency of the PC algorithm based on i.i.d. samples has been shown to be achieved with a further Gaussian distribution assumption [57], which popularized the use of PC algorithm in high dimensional settings.

In the context of time series causality, it is of interest to find whether a variable at a certain time influences another variable at a later time [211; 212; 213; 212; 214]. For example, in neuroscience, a question of interest is whether the activity of neuron u at time t influences the activity of neuron v at a later time t' where $t' \geq t$ [109; 1; 2]. In climate studies, it is of interest to know whether change in certain climatic variables (e.g. river run-off at a station on a day) has a causal effect on other climatic variables at a later time (e.g. river run-off at a different station on a later day) [215; 2; 216]. In econometrics, it is of interest to know whether change in one macro-economic variable (e.g. real gross domestic product) at a certain month influences another macro-economic variable (e.g. borrowed reserves) at a later month [217; 132; 169]. We underscore that such across-time interactions are not a nuisance parameter interfering with causal inference, instead are the objects of inference.

The across-time causal relationships between the variables over time are typically represented by an unrolled graph with nodes (v, t) corresponding to variable v at time t and edges $(u, t_1) \rightarrow (v, t_2)$ representing the causal influence of variable u at time t_1 to variable v at time t_2 . The unrolled graph is summarized by a *rolled graph* with nodes as the variables and an edge $u \rightarrow v$ for variables u, v if variable u at some time t has a causal effect on variable v at a later time t' , where $t' \geq t$ (see Figure 4.1) [214; 2]. The notion of the rolled graph summarizes the interaction of the variables over time. Therefore, an alternative name used for rolled graph in literature is *summary causal graph* for time series [214; 218]. In neurosciences, the rolled graph of across-time causal relationships of activity of multiple neurons over time is termed as the *causal functional connectivity* between the neurons [109; 63; 2]. The causal functional connectivity is considered to provide important information for brain function, dysfunction and cognition [16; 15].

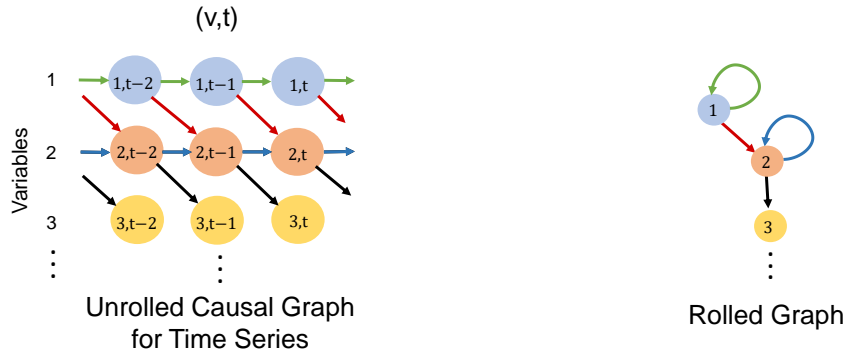


Figure 4.1 Schematic depiction of causal modeling of time series by the Unrolled Causal Graph (left), and then rolling back its edges (colored) to define the Rolled Graph (right).

As noted earlier, for causal inference in time series, across-time causal relationships are of interest, i.e. whether variable u at time t has a causal influence on a variable v at a later time t' where $t' \geq t$. Traditionally, Granger Causality based on vector-autoregressive models have been used for answering causal queries in time series [117; 219], whose outcome have been represented by unrolled and rolled graphs based on Granger causality [134; 199]. In recent years, Chu et al. [213] proposed a causal inference algorithm designed for additive, nonlinear and stationary time series, using conditional dependence tests based on additive regression model. Using the fact that one can incrementally construct the mutual information between a cause and an effect based on the mutual information values between the effect and the previously found causes, Jangyodsuk et al. [220] proposed to obtain a causal graph with each time series as a node and the difference in time steps between the occurrence of cause and the occurrence of the effect as the edge weights. Recently, methods based on manipulations of the PC algorithm for inferring causal graphs in time series has also drawn attention [1; 2; 132; 31; 188].

It is noteworthy that the naive application of the PC algorithm to time series assumes faithfulness of the random vector at each time with respect to a (contemporaneous) DAG that is same for each time, and infers that contemporaneous DAG at each time (see Section 4.2 for details). Only inferring the contemporaneous DAG at a fixed time would provide no information about across-time causal relationships of interest. Recent approaches such as the Time-Aware PC (TPC) algorithm uses the PC algorithm after transformations on the time series and has been shown to have effective performance in inferring across-time causal relationships in simulated and benchmarking real datasets [2]. In contrast to the naive PC, the TPC algorithm infers the DAG

across time points under faithfulness with respect to that DAG. This results in inference of across-time causal relationships.

The goal of the current chapter is to establish consistency of the TPC algorithm in inferring the across-time DAG and its rolled graph in a time series scenario. However, for clarity on the consistency of the TPC algorithm which builds upon the PC algorithm in time series, it is first essential to have clarity on the consistency of the PC algorithm under temporal dependence.

In this chapter we prove consistency of the PC algorithm under temporal dependence, where the number of variables p is fixed, and there are n dependent identically distributed samples over time. A common example is a stationary multivariate time series. In particular, we show that under a standard set of assumptions, the PC algorithm is consistent in time series, for two well-known tests for conditional dependence. Specifically, we show that the consistency of the PC algorithm holds for a multivariate time series, under conditions of ρ -mixing (see 4.7) with ‘sufficiently fast’ convergence of the maximal correlation coefficients and faithfulness of the random vector at each time. This enables us to establish the consistency of PC algorithm in common time series models such as vector auto-regressive moving average models and linear processes. We also show that the Time-Aware PC algorithm achieves consistency under similar assumptions but with faithfulness with respect to a DAG across time points. The latter is more reasonable in a time series context with interactions between variables across different times.

The rest of the chapter is organized as follows. In section 4.1 we provide a succinct background on causal structure learning. In section 4.2, we establish consistency of PC algorithm to infer the contemporaneous DAG from ρ -mixing time series. In section 4.3, we present analogous results for strongly mixing processes. In section 4.4, we exemplify the results in common time series models such as vector auto-regressive moving average and linear processes. In section 4.5, we show consistency of the TPC algorithm in inferring the across-time and rolled DAGs from time series under mixing. In section 4.6, we demonstrate the comparative performance of the methods in simulation studies, and then in a real benchmark dataset from the public benchmarking platform *CauseMe* in section 4.7. This is followed by a Conclusion.

4.1 Causal Structure Learning

We begin by describing the general framework for causal structure learning by directed graphical models. A graph $G = (V, E)$ consists of a set of nodes $V = \{1, \dots, p\}$ for fixed p , and a set of edges $E \subseteq V \times V$. In our setting the set of nodes corresponds to the components of a p -variate random vector $\mathbf{X} \in \mathbb{R}^p$. An edge $(u, v) \in E$ is called directed and denoted by $u \rightarrow v$ if $(u, v) \in E$ but $(v, u) \notin E$. An edge between u and v is called undirected and denoted by $u - v$ if both $(u, v) \in E$ and $(v, u) \in E$. A Directed Acyclic Graph (DAG) is a graph with all edges directed and devoid of directed cycles.

We will now define some graphical preliminaries which will be used in this section (see [221]). In a DAG G , we call a pair of nodes v and w adjacent if $v \rightarrow w \in E$ or $w \rightarrow v \in E$. A path is a sequence of distinct nodes in which successive nodes are adjacent. v_0 and v_k are called endpoints of the path $\pi = (v_0, v_1, \dots, v_k)$. For a non-endpoint v_i , if $v_{i-1} \rightarrow v_i \leftarrow v_{i+1}$ is a subpath of π then v_i is called a collider on π . Else, v_i is called a non-collider on π . v_0 is called an ancestor of v_k (equivalently, v_k is called a descendant of v_0) if every edge on π is of the form $v_{i-1} \rightarrow v_i$. A common convention is to refer to v as an ancestor and descendant of itself. We denote the set of ancestors and descendants of v in G by $an_G(v)$ and $de_G(v)$ respectively, and define $an_G(C) := \cup_{v \in C} an_G(v)$, $de_G(C) := \cup_{v \in C} de_G(v)$. We refer to the skeleton of a DAG G as the undirected graph obtained by replacing directed edges in G with undirected edges. An ordered triplet of nodes (u, v, w) with u and w not adjacent in G and G containing the directed edges $u \rightarrow v$ and $w \rightarrow v$ is called a v-structure in G . The following definition of d-separation is taken from [221].

Definition 3 (d-separation). Two nodes v and w in a DAG G are *d-connected* given $C \subset V \setminus \{v, w\}$ if G contains a path π with endpoints v and w such that (i) the colliders on π are in $an_G(C)$, and (ii) no non-collider on π is in C . Generalizing to sets, two disjoint sets $A, B \subset V$ are *d-connected* given $C \subset V \setminus (A \cup B)$ if there are two nodes $v \in A$ and $w \in B$ that are d-connected given C . If this is not the case, then C *d-separates* A and B .

4.1.1 Conditions for Causal Inference

Spirtes and Pearl pioneered the use of DAGs in causal inference using the condition of *causal sufficiency*, i.e. absence of hidden (or latent) variables [52], and two conditions that relate the DAG and probability distributions, namely *directed Markov property* and *faithfulness*.

1. **Directed Markov Property.** Let $(X_v : v \in V) \sim P$. P is said to satisfy the Directed Markov Property with respect to G if, for all $A, B, C \subset V$,

$$C \text{ d-separates } A \text{ and } B \Rightarrow X_A \perp\!\!\!\perp X_B | X_C \quad (4.1)$$

where $\perp\!\!\!\perp$ represents independence of two random variables.

2. **Faithfulness.** Let P satisfy the Directed Markov Property. P is said to be faithful with respect to G if the converse of (4.1) holds, consequently i.e., for all $A, B, C \subset V$,

$$X_A \perp\!\!\!\perp X_B | X_C \Leftrightarrow C \text{ d-separates } A \text{ and } B \quad (4.2)$$

It is noteworthy that faithfulness constrains the class of probability distributions. See Chapter 3.5.2 in [52] for an example of a non-faithful distribution. In this chapter, we mostly limit ourselves to the multivariate Gaussian family, where non-faithful distributions form a Lebesgue null set in the space of distributions associated with a DAG G (see [222]).

Most algorithms for estimating a DAG G under the Conditions 1-3 cannot distinguish between two DAGs which are *Markov equivalent* (two DAGs G_1 and G_2 are said to be Markov equivalent when the set of distributions that are faithful with respect to G_1 is the same as the set of distributions that are faithful with respect to G_2 [223]). Instead, one can identify the Markov equivalence class of DAGs. The Markov equivalence classes of DAGs can also be characterized by the following criterion: Two DAGs G_1 and G_2 are Markov equivalent if and only if they have the same skeleton and v-structures [224]. It is commonplace for a Markov equivalence class to be conveniently represented by a Completed Partially Directed Acyclic Graph (CPDAG) which has a directed edge $v \rightarrow w$ if $v \rightarrow w$ is present in all the DAGs belonging to the Markov equivalence class of the DAG and has an undirected edge between v and w if both $v \rightarrow w$ and $w \rightarrow v$ are present among the DAGs. A CPDAG represents a Markov equivalence class uniquely, since two CPDAGs are identical if and only if they represent the same Markov equivalence class [223].

Therefore, instead of estimating the DAG satisfying the conditions of causal inference, the main goal is to estimate the CPDAG. We focus on the PC algorithm in this chapter which is a popular method for this estimation task [52; 48]. Although the main goal is to estimate the CPDAG, PC algorithm divides the task

into two parts 1) Estimation of the skeleton of the true DAG, and 2) Orientation of edges of the skeleton to obtain the CPDAG estimate. All inference from data takes place in the first part. In the following section we describe the steps of the PC algorithm which will be used later to establish conditions for its consistency in temporal dependence scenario.

4.1.2 Estimation of the CPDAG: The PC algorithm

The PC algorithm is a popular method to estimate the CPDAG from observed data. It has been shown to be pointwise consistent based on independent and identically distributed data using a consistent test for conditional dependence, and uniformly consistent¹ when additionally the data distribution is Gaussian [57].

The population version of the algorithm has two parts:

1. To find the skeleton G_{skel} of the true DAG G by representing the variables as nodes of an empty DAG and putting an undirected edge between each pair of nodes if they are independent or conditionally independent given any other variable(s).
2. The skeleton is then converted into the CPDAG G_{CPDAG} by orienting the undirected edges using rules for orientation.

In the sample version the independence and conditional independence statements are replaced by a statistical test based on the sample. The parts of the population version of the PC algorithm are outlined in Algorithms 2 and 3. For a DAG G such that the distribution P is faithful to G , it is proved that the PC algorithm 2 constructs the true skeleton of the DAG [57], and algorithm 3 outputs the true CPDAG [225]. In the sample version only the conditional independence statements in finding the skeleton in Algorithm 2 are replaced with a statistical test for conditional dependence based on data.

4.1.3 Causal recovery in time series: The Time-aware PC algorithm

The Time-Aware PC algorithm is an adaptation of the PC algorithm for causal inference in a time series scenario by considering DAGs between variables across different times instead of at a fixed time. This is especially relevant in practice where past observations in a time series have a causal influence on future

¹See [56] for definition of uniform consistency in this setting.

Algorithm 2: Population PC - Finding the skeleton graph

Input : Node set V , Conditional Independence Information
Output : Skeleton G_{skel} , separation sets S

- 1 Start with a complete undirected graph C on the vertex set V .
- 2 Set $\ell = -1$; $G_{skel} = C$
- 3 **repeat**
- 4 $\ell = \ell + 1$
- 5 **repeat**
- 6 Select a new ordered pair i, j adjacent in G_{skel} with $|adj(G_{skel}, i) \setminus \{j\}| \geq \ell$
- 7 **repeat**
- 8 Choose $k \subseteq adj(G_{skel}, i) \setminus \{j\}$ with $|k| = \ell$
- 9 **if** $i \perp\!\!\!\perp j \mid k$ **then**
- 10 Delete edge $i - j$
- 11 Denote this new graph by G_{skel}
- 12 Save k in $S(i, j)$ and $S(j, i)$
- 13 **until** edge i, j is deleted or all $k \subset adj(G_{skel}, i) \setminus \{j\}$ with $|k| = \ell$ are selected;
- 14 **until** all i, j adjacent such that $|adj(G_{skel}, i) \setminus \{j\}| \geq \ell$ and $k \subseteq adj(G_{skel}, i) \setminus \{j\}$ with $|k| = \ell$ have been tested for conditional independence.;
- 15 **until** for each i, j adjacent: $|adj(G_{skel}, i) \setminus \{j\}| \leq \ell$;

Algorithm 3: Population PC - Directing the skeleton into a CPDAG

Input : Skeleton G_{skel} , separation sets S
Output : G_{CPDAG}

- 1 **for all** i, j non-adjacent
- 2 Orient $i - c - j$ whenever i, j are not conditionally independent given c into $i \rightarrow c \leftarrow j$.
- 3 Orient $i \rightarrow j - c$ with i, j non-adjacent into $i \rightarrow j \rightarrow c$
- 4 **end for**
- 5 **for all** i, j adjacent
- 6 Orient $i - j$ whenever $i \rightarrow c \rightarrow j$ into $i \rightarrow j$
- 7 Orient $i - j$ whenever $i - c \rightarrow j$ and $i - d \rightarrow j$ for some c, d nonadjacent into $i \rightarrow j$.
- 8 Orient $i - j$ whenever $i - c \rightarrow d$ and $c \rightarrow d \rightarrow j$ for some c, d nonadjacent into $i \rightarrow j$.
- 9 **end for**

observations, that is, there exist across-time causal relationships. Let us denote a time series of variables $V = \{1, \dots, p\}$ by $\mathbf{X}_t = (X_{t1}, \dots, X_{tp}), t \geq 1$. An object of general interest for causal inference in the time series scenario is a graph G_R with nodes as the variables $V = \{1, \dots, p\}$ and edge $u \rightarrow v$ if $X_{t_1u} \rightarrow X_{t_2v}$ for some $t_1 \leq t_2$, where $X_{t_1u} \rightarrow X_{t_2v}$ are defined by a causal model with respect to a graph G with nodes X_{tu} . Some common examples are in interventional, structural and Granger Causality [134; 226], and applications in neurosciences [1; 23], and econometrics [211]. We call G_R as the *Rolled Graph* of G (see Figure 4.1).

Algorithm 4: Time-Aware PC Algorithm

Input : Vertex Set \mathbf{V} , $\{\chi_t\}_{t=1}^N$

Output : $\hat{\mathbf{G}}_{CPDAG,R}$ (estimate of $\mathbf{G}_{CPDAG,R}$)

- 1 Use PC Algorithm to estimate \mathbf{G}_{CPDAG} based on χ_1, \dots, χ_N by $\hat{\mathbf{G}}_{CPDAG}$.
 - 2 Next, $\hat{\mathbf{G}}_{CPDAG}$ is transformed to obtain $\hat{\mathbf{G}}_{CPDAG,R}$.
-

Let $\mathbf{X}_1, \dots, \mathbf{X}_n \in \mathbb{R}^p$ be a strictly stationary time series and $\mathbf{X}_t = (X_{t1}, \dots, X_{tp})$. Let the graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ consist of the set of nodes $\mathbf{V} = \{(v, t) : v \in \{1, \dots, p\}, t \in \{1, \dots, \tau\}\}$ for $\tau \geq 1$, and directed edges $\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$. Assume that \mathbf{G} is a DAG. In \mathbf{G} , an edge $(u, t_1) \rightarrow (v, t_2)$ represents a connection from variable u at time t_1 to variable v at time t_2 .

Let $\chi_t = (\mathbf{X}_{1+(t-1)r}, \dots, \mathbf{X}_{\tau+(t-1)r}) \in \mathbb{R}^{p\tau}$, $r \geq 1, t = 1, 2, \dots, N := \lfloor \frac{n-\tau}{r} \rfloor + 1$. Note that χ_t is a concatenation of the vectors $\mathbf{X}_{1+(t-1)r}, \dots, \mathbf{X}_{\tau+(t-1)r}$, each of dimension p , into a resulting vector of dimension $p(\tau + 1)$. For our setting, the components of χ_t have the following correspondence with the nodes in the unrolled graph \mathbf{G} : the $(p(t' - 1) + v)$ -th component of χ_t correspond to node (v, t') in \mathbf{G} . It follows that $\{\chi_t\}_{t=1}^\infty$ is also strictly stationary.

The Time-Aware PC Algorithm uses PC Algorithm to estimate \mathbf{G}_{CPDAG} based on χ_1, \dots, χ_n by $\hat{\mathbf{G}}_{CPDAG}$. Next, the estimate $\hat{\mathbf{G}}_{CPDAG}$ is transformed to obtain $\hat{\mathbf{G}}_{CPDAG,R}$. See Algorithm 4 for an outline of the algorithm and [2] for more details.

4.1.4 Conditional Dependence Test in the Gaussian Regime: Pearson's Partial Correlations

In the Gaussian scenario conditional dependence can be inferred from partial correlations. Assume that $\mathbf{X} = (X_1, \dots, X_p)$ is a p -dimensional Gaussian random vector, for some fixed integer p . For $i \neq j \in \{1, \dots, p\}$, $\mathbf{k} \subseteq \{1, \dots, p\} \setminus \{i, j\}$, denote by $\rho_{i,j|\mathbf{k}}$ the partial correlation between X_i and X_j given $\{X_r : r \in \mathbf{k}\}$. We can infer conditional dependence by inferring partial correlations due to the following elementary property of the multivariate Gaussian distribution (see Prop. 5.2 in [208]) that,

$$\rho_{i,j|\mathbf{k}} = 0 \text{ if and only if } X_i \perp\!\!\!\perp X_j \mid \{X_r : r \in \mathbf{k}\}.$$

The sample partial correlation $\hat{\rho}_{i,j|\mathbf{k}}$ can be calculated via regression or by using the following identity. Denote $k = |\mathbf{k}|$ and let without loss of generality $\{X_r; r \in \mathbf{k}\}$ be the last k entries in \mathbf{X} . Let $\Sigma := \text{cov}(\mathbf{X})$

with $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$ where Σ_{11} is of dimension $(p - k) \times (p - k)$, Σ_{22} is of dimension $k \times k$, $\Sigma_{11.2} = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}$, and $\hat{\Sigma}$ and $\hat{\Sigma}_{11.2}$ be the sample versions of Σ and $\Sigma_{11.2}$ based on the sample covariance matrix. Let $\mathbf{e}_1, \dots, \mathbf{e}_p$ be the canonical basis vectors of \mathbb{R}^p . It follows from [227] that,

$$\rho_{i,j|\mathbf{k}} = \frac{\mathbf{e}_i^\top \Sigma_{11.2} \mathbf{e}_j}{\sqrt{(\mathbf{e}_i^\top \Sigma_{11.2} \mathbf{e}_i)(\mathbf{e}_j^\top \Sigma_{11.2} \mathbf{e}_j)}},$$

$$\hat{\rho}_{i,j|\mathbf{k}} = \frac{\mathbf{e}_i^\top \hat{\Sigma}_{11.2} \mathbf{e}_j}{\sqrt{(\mathbf{e}_i^\top \hat{\Sigma}_{11.2} \mathbf{e}_i)(\mathbf{e}_j^\top \hat{\Sigma}_{11.2} \mathbf{e}_j)}}.$$

For testing whether a partial correlation is zero or not, we first apply Fisher's z-transform

$$Z(i, j|\mathbf{k}) = g(\hat{\rho}_{i,j|\mathbf{k}}) = \frac{1}{2} \log \left(\frac{1 + \hat{\rho}_{i,j|\mathbf{k}}}{1 - \hat{\rho}_{i,j|\mathbf{k}}} \right). \quad (4.3)$$

Let also, $z(i, j|\mathbf{k}) = g(\rho_{i,j|\mathbf{k}})$. Note that $z(i, j|\mathbf{k}) = 0 \Leftrightarrow \rho(i, j|\mathbf{k}) = 0$, and hence, $z(i, j|\mathbf{k}) = 0 \Leftrightarrow X_i \perp\!\!\!\perp X_j \mid \{X_r : r \in \mathbf{k}\}$. We will show that $Z(i, j|\mathbf{k})$ is a consistent estimator of the population parameter $z(i, j|\mathbf{k})$, and hence, a consistent conditional dependence test can be constructed based on the statistic $Z(i, j|\mathbf{k})$. Using such a test for conditional dependence, the sample version of the PC algorithm (see 5) only replaces line 9 in the population PC in 2 about conditional independence with the statistical test.

4.1.5 Conditional Dependence Test in the Non-Gaussian Regime: The Hilbert Schmidt Criterion

When the data is non-Gaussian, zero partial correlations do not necessarily imply conditional independence. In such a situation, the Hilbert Schmidt criterion acts as a convenient test for conditional dependence. We describe this concept below.

Given \mathbb{R} -valued random variables X, Y and the random vector \mathbf{Z} we propose the use of the following statistic for testing the conditional dependence of $X, Y|\mathbf{Z}$ (see [228]):

$$\hat{H}_n(X, Y|\mathbf{Z}) = Tr[R_{\dot{Y}} R_{\dot{X}} - 2R_{\dot{Y}} R_{\dot{X}} R_{\mathbf{Z}} + R_{\dot{Y}} R_{\mathbf{Z}} R_{\dot{X}} R_{\mathbf{Z}}]$$

Algorithm 5: Sample PC algorithm

Input : Dataset $\{\mathbf{X}_i\}_{i=1}^n$ **Output** : CPDAG estimate \hat{G}_{CPDAG}

- 1 Run the Population PC Algorithm 2 to find the skeleton but replace in line 9 the statement about conditional independence of i, j given \mathbf{k} by the event of acceptance of conditional independence by a consistent test for conditional dependence.
 - 2 Extend the skeleton to a CPDAG using Algorithm 3.
-

where G_X, G_Y, G_Z are the centered Gram matrices with respect to a positive definite and integrable kernel k , that is, $G_{X,ij} = \langle k(\cdot, X_i) - \hat{m}_X^{(n)}, k(\cdot, X_j) - \hat{m}_X^{(n)} \rangle$ with $\hat{m}_X^{(n)} = \frac{1}{n} \sum_{i=1}^n k(\cdot, X_i)$, and $R_A = G_A(G_A + n\epsilon_n I_n)^{-1}$ for $A \in \{X, Y, Z\}$, and $\ddot{X} = (X, Z), \ddot{Y} = (Y, Z)$. Under some regularity assumptions mentioned below, it follows from the proof of Theorem 5 in [228] that $\hat{H}_n(X, Y|Z)$ is a consistent estimator of $H(X, Y|Z) := \|V_{\ddot{X}\ddot{Y}|Z}\|^2$, where

$$V_{\ddot{X}\ddot{Y}|Z} := \Sigma_{\ddot{X}\ddot{X}}^{-1/2} (\Sigma_{\ddot{X}\ddot{Y}} - \Sigma_{\ddot{X}Z} \Sigma_{ZZ}^{-1} \Sigma_{Z\ddot{Y}}) \Sigma_{\ddot{Y}\ddot{Y}}^{-1/2}$$

and Σ_{UV} denotes the covariance matrix of U and V . It follows from [228] that $X \perp\!\!\!\perp Y|Z \Leftrightarrow H(X, Y|Z) = 0$. Hence, a consistent conditional dependence test of $X, Y|Z$ can be constructed based on the statistic $H_n(X, Y|Z)$. The advantage of this method is that unlike the Pearson partial correlation, it does not require Gaussianity of the data to decide conditional independence, and hence can be used in the PC algorithm if the underlying time series is non-Gaussian.

4.2 Consistency of the PC Algorithm under temporal dependence

While the sample PC algorithm as described in Algorithm 5 is known to be consistent in estimating the true CPDAG from independent and identically distributed data provided one uses a consistent conditional dependence test [57; 228], its consistency, to the best of our knowledge, is unknown when the data are not independent such as in a time series scenario. In the simple case when the data is Gaussian, the analysis for i.i.d. data in [57; 229] involves calculation of the density function of sample partial correlations and thus the conclusions do not trivially follow in a setting with temporally dependent samples. In this section, we establish some structure on the dependent time series data that ensures consistency of the PC algorithm. An

important structure that we assume for the time series is strongly mixing, which assumes a form of weak dependence among the time series variables with increasing time gap between the variables. As a preparation we introduce a general consistency framework below.

4.2.1 The General Consistency Framework

The aim of this section is to show that under certain assumptions on the time series, the CPDAG estimate of the PC algorithm is consistent, that is, the probability of an error in estimation of the CPDAG converges to 0. An error happens in estimating an edge in the CPDAG if and only if the underlying conditional dependence tests performs false rejection or acceptance. Therefore, for showing that the probability of an error in estimation of the CPDAG converges to 0, it suffices to show that the probability of an error in testing conditional dependence converges to 0, a fact that is guaranteed if one works with a consistent test for conditional dependence in the PC algorithm. To be precise, suppose that we start with a measure $\mu(i, j|\mathbf{k})$ of conditional dependence, which satisfies the property that

$$\mu(i, j|\mathbf{k}) = 0 \text{ if and only if } X_i \perp\!\!\!\perp X_j \mid \{X_r : r \in \mathbf{k}\},$$

and suppose that for each i, j, \mathbf{k} , we have a sequence $\hat{\mu}_n(i, j|\mathbf{k})$ of consistent estimators of $\mu(i, j|\mathbf{k})$. Then, the test of conditional dependence based on $\hat{\mu}_n(i, j|\mathbf{k})$, which will reject conditional independence of X_i and X_j given $\{X_r\}_{r \in \mathbf{k}}$ if and only if $|\hat{\mu}_n(i, j|\mathbf{k})| > \gamma$ for some thresholding parameter γ , will be consistent. Consequently, the error probability in estimating the CPDAG converges to 0. The threshold γ can be obtained using a bootstrap method, e.g. the stationary bootstrap for dependent samples constituting a stationary time series [230].

Below, we state the result that guarantees consistent estimation of the DAG skeleton.

Theorem 4.2.1. Denote by $\hat{G}_{skel,n}$ the estimate of the skeleton graph from the PC algorithm with $\hat{\mu}_n$ as the test statistic for conditional dependence, and by G_{skel} the true skeleton of the DAG G . Then,

$$P(\hat{G}_{skel,n} = G_{skel}) \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

Proof. An error occurs in the sample PC algorithm if there is a pair of nodes i, j and a conditioning set

$\mathbf{k} \in K_{i,j}$ where an error event $E_{i,j|\mathbf{k}}$ occurs, where $E_{i,j|\mathbf{k}}$ denotes that “an error occurred when testing partial correlation for zero at nodes i, j with conditioning set \mathbf{k} ”, i.e.,

$$E_{i,j|\mathbf{k}} = E_{i,j|\mathbf{k}}^I \cup E_{i,j|\mathbf{k}}^{II},$$

where

$$E_{i,j|\mathbf{k}}^I := \{|\hat{\mu}_n(i, j|\mathbf{k})| > \gamma \text{ and } \mu(i, j|\mathbf{k}) = 0\}$$

$$E_{i,j|\mathbf{k}}^{II} := \{|\hat{\mu}_n(i, j|\mathbf{k})| \leq \gamma \text{ and } \mu(i, j|\mathbf{k}) \neq 0\}$$

denote the events of Type I error and Type II error respectively. Thus,

$$\begin{aligned} P(\text{an error occurs in the PC algorithm}) &\leq P\left(\bigcup_{i,j,\mathbf{k} \in K_{i,j}} E_{i,j|\mathbf{k}}\right) \\ &\leq O(1) \sup_{i,j,\mathbf{k} \in K_{i,j}} P(E_{i,j|\mathbf{k}}) \end{aligned} \quad (4.4)$$

using that the cardinality of the set $|\{i, j, \mathbf{k} \in K_{i,j}\}| = 2^{p-2}p^2$. Then, for any $\gamma > 0$, we have:

$$\sup_{i,j,\mathbf{k} \in K_{i,j}} P(E_{i,j|\mathbf{k}}^I) = \sup_{i,j,\mathbf{k} \in K_{i,j}} P(|\hat{\mu}_n(i, j|\mathbf{k}) - \mu(i, j|\mathbf{k})| > \gamma)$$

It now follows from the consistency of $\hat{\mu}_n$ that

$$\sup_{i,j,\mathbf{k} \in K_{i,j}} P(E_{i,j|\mathbf{k}}^I) \rightarrow 0 \quad (4.5)$$

Next, we control the type II error probability. Towards this, let $c = \inf\{|\mu(i, j|\mathbf{k})| : \mu(i, j|\mathbf{k}) \neq 0\} > 0$, and choose $\gamma = c/2$. Then,

$$\begin{aligned} \sup_{i,j,\mathbf{k} \in K_{i,j}} P(E_{i,j|\mathbf{k}}^{II}) &= \sup_{i,j,\mathbf{k} \in K_{i,j}} P(|\hat{\mu}_n(i,j|\mathbf{k})| \leq \gamma, \mu(i,j|\mathbf{k}) \neq 0) \\ &\leq \sup_{i,j,\mathbf{k} \in K_{i,j}} P(|\hat{\mu}_n(i,j|\mathbf{k}) - \mu(i,j|\mathbf{k})| > c/2) \end{aligned}$$

Since the right hand side $\rightarrow 0$ as $n \rightarrow \infty$, it follows that

$$\sup_{i,j,\mathbf{k} \in K_{i,j}} P(E_{i,j|\mathbf{k}}^{II}) \rightarrow 0 \quad (4.6)$$

Recall that from (4.4), $P(\text{an error occurs in the PC algorithm}) \leq O(1) \sup_{i,j,\mathbf{k} \in K_{i,j}} P(E_{i,j|\mathbf{k}})$. Now, by union bound, we have that $P(E_{i,j|\mathbf{k}}) \leq P(E_{i,j|\mathbf{k}}^I) + P(E_{i,j|\mathbf{k}}^{II})$. Hence, by (4.5) and (4.6) we get

$$P(\text{an error occurs in the PC algorithm}) \rightarrow 0$$

This completes the proof. \square

Since all inference is done while finding the skeleton, if this part is completed perfectly, that is, if there was no error while testing for conditional dependence, then G_{CPDAG} will be estimated without error (See [225]). Therefore, we easily obtain:

Theorem 4.2.2. Denote by $\hat{G}_{CPDAG,n}$ the estimate from the entire PC algorithm and by G_{CPDAG} the true $CPDAG$ from the DAG G . Then,

$$P(\hat{G}_{CPDAG,n} = G_{CPDAG}) \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

4.2.2 Consistency of Tests for Conditional Dependence

In the Gaussian regime (Section 4.1.4), one can take $\mu(i,j|\mathbf{k}) = z(i,j|\mathbf{k})$ and $\hat{\mu}_n(i,j|\mathbf{k}) := Z(i,j|\mathbf{k})$, and in the non-Gaussian regime (Section 4.1.5), one can take $\mu(i,j|\mathbf{k}) = H(X_{1i}, X_{1j}|X_{1r} : r \in \mathbf{k})$ and $\hat{\mu}_n(i,j|\mathbf{k}) := H_n(X_{1i}, X_{1j}|X_{1r} : r \in \mathbf{k})$, where $\mathbf{X}_t = (X_{t1}, X_{t2}, \dots, X_{tp})$. All that we need to show in order to establish consistency of the PC algorithm in these two settings, is the consistency of $\hat{\mu}_n(i,j|\mathbf{k})$ as an

estimator of $\mu(i, j|\mathbf{k})$. With this in mind, we will delineate the conditions on the time series under which the consistency of the sample versions of the conditional dependence tests holds. We need a few notations first.

Let $\{\mathbf{X}_t\}_{t=1}^{\infty}$ be a strictly stationary \mathbb{R}^p -valued time series, where $\mathbf{X}_t \sim P$ for all $t \geq 1$. Let $\mathbf{Y}_t = \mathbf{X}_t - \mathbb{E}\mathbf{X}_t$ and denote by ρ_{ij} the population (Pearson's) correlation between X_{1i} and X_{1j} (equivalently between Y_{1i} and Y_{1j}), and by $\hat{\rho}_{n;i,j}$ the sample correlation between $\mathbf{X}^{(i)} = (X_{1i}, X_{2i}, \dots, X_{ni})$ and $\mathbf{X}^{(j)}$ (equivalently between $\mathbf{Y}^{(i)}$ and $\mathbf{Y}^{(j)}$) for n samples.

We now introduce the concept of ρ -mixing. For fixed $i, j \in 1, \dots, p$, let \mathcal{F}_a^b be the σ -field of events generated by the random variables $\{X_{ti}, X_{tj} : a \leq t \leq b\}$, and $L_2(\mathcal{F}_a^b)$ be the collection of all second-order random variables which are \mathcal{F}_a^b -measurable. The stationary process $\{X_{ti}, X_{tj} : t = 1, 2, \dots\}$ is called ρ -mixing [231; 232] if the maximal correlation coefficient:

$$\xi_{ij}(k) := \sup_{l \geq 1} \sup_{\substack{U \in L_2(\mathcal{F}_1^l) \\ V \in L_2(\mathcal{F}_{l+k}^{\infty})}} \frac{|\text{cov}(U, V)|}{\text{var}^{1/2}(U)\text{var}^{1/2}(V)} \rightarrow 0 \text{ as } k \rightarrow \infty. \quad (4.7)$$

Let us assume the following conditions.

- (A.1) $\{X_{ti}, X_{tj} : t = 1, 2, \dots\}$ is ρ -mixing for all i, j , with maximal correlation coefficients $\xi_{ij}(k), k \geq 1$.
- (A.2) $\mathbb{E}X_{ti}^4 < \infty$ for all t, i and $\sum_{k=1}^{\infty} \xi_{ij}(k) < \infty$ for $i, j \in 1, \dots, p$.
- (A.3) There exists a sequence of positive integers $s_n \rightarrow \infty$ and $s_n = o(n^{1/2})$ such that $n^{1/2}\xi_{ij}(s_n) \rightarrow 0$ as $n \rightarrow \infty$ for $i, j \in 1, \dots, p$.
- (A.4) P is faithful to a DAG G .

The Gaussian Regime

The following lemma shows that under assumptions A.1-A.3, the sample correlation between every pair of variables converges to the corresponding population correlation.

Lemma 4.2.1. Under A.1-A.3,

$\hat{\rho}_{n;i,j}$ converges to $\rho_{i,j}$ in probability.

Proof. Note that $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ are centered random variables, and under A.1-A.3 together with the Gaussianity assumption on the data, each of their entries have finite fourth moment. Furthermore, the maximal correlation coefficients are invariant to centering and scaling of the random variables. Therefore, assumptions A.1-A.3 hold for $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ which are also Gaussian. Hence, by Theorem 3 in [233], the lemma follows. \square

Denote the population partial correlation between X_{1i} and $X_{1j} \mid \{X_{1r} : r \in \mathbf{k}\}$ for some $\mathbf{k} \subset \{1, \dots, p\} \setminus \{i, j\}$ by $\rho_{i,j|\mathbf{k}}$, and let $k := |\mathbf{k}|$. Similarly denote by $\hat{\rho}_{n;i,j|\mathbf{k}}$, the sample Partial Correlation between $\mathbf{X}^{(i)}$ and $\mathbf{X}^{(j)} \mid \{\mathbf{X}^{(r)} : r \in \mathbf{k}\}$ which is also the partial correlation between $\mathbf{Y}^{(i)}$ and $\mathbf{Y}^{(j)} \mid \{\mathbf{Y}^{(r)} : r \in \mathbf{k}\}$, for $\mathbf{k} \subset \{1, \dots, p \setminus \{i, j\}\}$.

We will now show that Lemma 4.2.1 can in fact be used to prove convergence of the pairwise sample partial correlations to the corresponding population correlations.

Lemma 4.2.2. Assume (A.1)-(A.3). Then,

$\hat{\rho}_{n;i,j|\mathbf{k}}$ converges to $\rho_{i,j|\mathbf{k}}$ in probability.

Proof. Without loss of generality let $\{X_{1r} : r \in \mathbf{k}\}$ be the last k entries of \mathbf{X}_1 . We will define a function f on the set of all non-singular $M \in \mathbb{R}^{p \times p}$. Let $M = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix}$ where M_{11} is of dimension $(p-k) \times (p-k)$ and M_{22} is of dimension $k \times k$, $M_{11,2} = M_{11} - M_{12}M_{22}^{-1}M_{21}$ and e_1, \dots, e_p denote the canonical basis vectors of \mathbb{R}^p . Define $f(M) = \frac{e_i^\top M_{11,2} e_j}{\sqrt{(e_i^\top M_{11,2} e_i)(e_j^\top M_{11,2} e_j)}}$. Clearly f is a continuous function. Let $\hat{\Sigma}$ denote the sample covariance matrix of $\mathbf{X}_1, \dots, \mathbf{X}_n$. Therefore as seen in Section 5.3 of [227],

$$\rho_{i,j|\mathbf{k}} = f(\Sigma)$$

and,

$$\hat{\rho}_{i,j|\mathbf{k}} = f(\hat{\Sigma})$$

Also note that $\Sigma = ((\rho_{ij}\sigma_i\sigma_j))$ and $\hat{\Sigma} = ((\hat{\rho}_{ij}\hat{\sigma}_i\hat{\sigma}_j))$ where ρ_{ij} is the population correlation of X_{1i} and X_{1j} and $\hat{\rho}_{ij}$ is the sample correlation of $\mathbf{X}^{(i)}$ and $\mathbf{X}^{(j)}$, σ_i^2 is the population variance of X_{1i} , and $\hat{\sigma}_i^2$ sample variance of $\mathbf{X}^{(i)}$, where $\mathbf{X}^{(i)} = (X_{1i}, X_{2i}, \dots, X_{ni})$. Now, from Corollary 1 in [233], it

follows that $\hat{\sigma}_i \xrightarrow{P} \sigma_i$ for all i under (A.0)-(A.3). Furthermore, $\hat{\rho}_{ij} \xrightarrow{P} \rho_{ij}$ according to Lemma 4.2.1 under (A.0)-(A.3). These together imply that $\hat{\Sigma} \xrightarrow{P} \Sigma$. Since, f is a continuous function, we also have $\hat{\rho}_{i,j|\mathbf{k}} = f(\hat{\Sigma}) \xrightarrow{P} f(\Sigma) = \rho_{i,j|\mathbf{k}}$, which completes the proof. \square

The PC algorithm tests partial correlations after the Z-transform $g(\rho) = 0.5 \log((1 + \rho)/(1 - \rho))$ (see Section 4.1.4 for details). Denote $Z_{n;i,j|\mathbf{k}} = g(\hat{\rho}_{n;i,j|\mathbf{k}})$ and $z_{i,j|\mathbf{k}} = g(\rho_{i,j|\mathbf{k}})$.

Lemma 4.2.3. Under Assumptions (A.1)-(A.3), we have, for all $i, j, \mathbf{k} \subseteq K_{i,j}$:

$Z_{n;i,j|\mathbf{k}}$ converges to $z_{i,j|\mathbf{k}}$ in probability.

Proof. Lemma 4.2.3 follows trivially from Lemma 4.2.2, on observing that the function g is continuous. \square

Since $Z_{n;i,j|\mathbf{k}}$ is a consistent sequence of estimators of $z_{i,j|\mathbf{k}}$, consistency of the PC algorithm using $Z_{n;i,j|\mathbf{k}}$ as a statistic for testing conditional dependence in the Gaussian regime now follows from Theorems 4.2.1 and 4.2.2.

The Non-Gaussian Regime

In this section, we will show that under assumptions A.1 - A.3, $H_n(X_{1i}, X_{1j}|X_{1r} : r \in \mathbf{k})$ is a consistent estimator of $H(X_{1i}, X_{1j}|X_{1r} : r \in \mathbf{k})$. For notational convenience, let us abbreviate $H(X_{1i}, X_{1j}|X_{1r} : r \in \mathbf{k})$ by $H_{i,j|\mathbf{k}}$ and $H_n(X_{1i}, X_{1j}|X_{1r} : r \in \mathbf{k})$ by $H_{n;i,j|\mathbf{k}}$.

Lemma 4.2.4. Under A.1-A.3, if the regularization constant ϵ_n satisfies $n^{-1/3} \ll \epsilon_n \ll 1$, then

$H_{n;i,j|\mathbf{k}}$ converges to $H_{i,j|\mathbf{k}}$ in probability.

Proof. The proof of Lemma 4.2.4 essentially follows from the proof of Theorem 5 in [228], modulo the fact that the samples are no longer independent. To begin with, it follows from the proof of Theorem 5 in [228] that it suffices to establish equations (14) and (15) in the supplement of [228] in our time-series setting, in order to prove Lemma 4.2.4. Equation (15) in the supplement of [228] is purely a population version which does not depend on the samples, so will go through in our case, too. Hence, we only need to show the validity

of Equation (14) in the supplement of [228] for our setting, in order to complete the proof. This, in turn, follows from Corollary 1 and Theorem 5 in [233], which gives the following for all $i \neq j$:

1. $\hat{\sigma}_i^2 := \frac{1}{n} \sum_{t=1}^n (X_{ti} - \bar{X}^{(i)})^2 \rightarrow \sigma_i^2 := \text{Var}(X_{1i})$, where $\bar{X}^{(i)} := \frac{1}{n} \sum_{t=1}^n X_{ti}$.
2. $\hat{\rho}_{n;i,j} = \rho_{i,j} + O_P(n^{-1/2})$.

The proof of Lemma 4.2.4 is now complete. □

Once again, since $H_{n;i,j|\mathbf{k}}$ is a consistent estimator of $H_{i,j|\mathbf{k}}$, consistency of the PC algorithm using $H_{n;i,j|\mathbf{k}}$ as a statistic for testing conditional dependence in the non-Gaussian regime now follows from Theorems 4.2.1 and 4.2.2.

4.3 Analogous Results for Strongly Mixing Processes

The class of ρ -mixing processes is contained in the class of so called strongly mixing processes, and it turns out that all our results can easily be adopted in this more general framework, too. Let us first introduce the concept of strong mixing. For fixed $i, j \in 1, \dots, p$, let \mathcal{F}_a^b be the σ -field of events generated by the random variables $\{X_{ti}, X_{tj} : a \leq t \leq b\}$, and $L_2(\mathcal{F}_a^b)$ be the collection of all second-order random variables which are \mathcal{F}_a^b -measurable. The stationary process $\{X_{ti}, X_{tj} : t = 1, 2, \dots\}$ is called strongly mixing [231; 232] if:

$$\alpha_{ij}(k) := \sup_{l \geq 1} \sup_{\substack{A \in L_2(\mathcal{F}_1^l) \\ B \in L_2(\mathcal{F}_{l+k}^\infty)}} |P(A \cap B) - P(A)P(B)| \rightarrow 0 \text{ as } k \rightarrow \infty.$$

In this case, $\alpha_{ij}(k)$ are called the strongly mixing coefficients.

All our results will go through for a strongly mixing process too, under the following slightly different set of assumptions:

(A.1)* $\{X_{ti}, X_{tj} : t = 1, 2, \dots\}$ is strongly mixing for all i, j , with coefficients $\alpha_{ij}(k), k \geq 1$.

(A.2)* $\mathbb{E}|X_{ti}|^{2\delta} < \infty$ for some $\delta > 2$ and all t, i , and the strongly mixing coefficients satisfy: $\sum_{k=1}^{\infty} \alpha_{ij}(k)^{1-2/\delta} < \infty$ for $i, j \in 1, \dots, p$.

(A.3)* There exists a sequence of positive integers $s_n \rightarrow \infty$ and $s_n = o(n^{1/2})$ such that $n^{1/2} \alpha_{ij}(s_n) \rightarrow 0$ as $n \rightarrow \infty$ for $i, j \in 1, \dots, p$.

(A.4)* P is faithful to a DAG G .

The proof follows directly from [233], in a manner exactly similar to the ρ -mixing case, so we skip it.

Remark 4.3.1. This section shows that one can derive consistency of the PC algorithm for the more general class of strongly mixing processes, if one agrees to assume conditions (A.1)* - (A.4)*. The results under the ρ -mixing assumption are still relevant, because if a ρ -mixing time series has finite fourth moment but all higher moments are infinite, then (A.2)* will not hold. In that case, one must appeal to the assumptions (A.1) - (A.4) in order to conclude consistency of the PC algorithm.

4.4 Common Time Series Models

We will consider two classes of examples to demonstrate the consistency of the PC algorithm in two commonly used time series models, namely VARMA and Linear Processes. Throughout this section, we are going to assume that the time series \mathbf{X}_t satisfies $\mathbb{E}|X_{ti}|^{2\delta} < \infty$ for some $\delta > 2$.

4.4.1 VARMA Process

Assume now that \mathbf{X}_t is a stationary vector autoregressive moving average (VARMA) process with values in \mathbb{R}^p . Then it admits a Markovian representation (See [234]),

$$\mathbf{X}_t = HZ_t, \quad Z_t = FZ_{t-1} + G\epsilon_t \quad (4.8)$$

where Z_t are random vectors, H, F, G are appropriate matrices with all eigenvalues of F being strictly less than 1, and ϵ_t are i.i.d. error random vectors with density g .

Theorem 4.4.1. Let \mathbf{X}_t be a stationary VARMA process which is faithful with respect to a DAG G . Also, suppose that the density g satisfies $\int \|x\|^{\beta_1} g(x) dx < \infty$ and $\int |g(x) - g(x - \theta)| dx = O(\|\theta\|^{\beta_2})$ for some $\beta_1, \beta_2 > 0$. Then, we have:

$$P(\hat{G}_{CPDAG,n} = G_{CPDAG}) \rightarrow 1 \quad \text{as } n \rightarrow \infty$$

In particular, the conclusion holds for a stationary Gaussian VARMA process.

Proof. Note that Lemma B.1.1 implies (A.1)* - (A.3)*. Hence, the consistency of the PC algorithm follows from the arguments presented in Section 4.3. Theorem 4.4.1 now follows from Lemma B.1.1, on observing that the Gaussian density satisfies its hypothesis with $\beta_1 = \beta_2 = 1$. \square

4.4.2 Linear Process

Assume now that \mathbf{X}_t is a stationary linear process with values in \mathbb{R}^p . That is, there exist i.i.d. random vectors ϵ_t having density g (see [234]), such that

$$\mathbf{X}_t = \sum_{i=0}^{\infty} A_i \epsilon_{t-i}, \quad A_0 = I,$$

for matrices A_i , where I is the identity matrix. We will assume that the density function g of the errors ϵ_t satisfies the condition $\int |g(x) - g(x - \theta)| dx = O(\|\theta\|)$, which is proved in (B.1) for a Gaussian density. Let the generating function of A_i be denoted as $A(z) = \sum_{k=0}^{\infty} A_k z^k$. Let us also assume that $E\|\epsilon_i\| < \infty$. Then, the conditions of Theorem 2.1 in [234] are satisfied. Denote $S_i = \sum_{j=i}^{\infty} \|A_j\|$, $\beta_\lambda(k) = \sum_{i=k}^{\infty} (S_i)^{\lambda/(1+\lambda)}$.

Theorem 4.4.2. If \mathbf{X}_t is a stationary linear process, which is faithful with respect to a DAG G , and satisfies the following conditions:

1. The density g satisfies $\int |g(x) - g(x - \theta)| dx = O(\|\theta\|)$,
2. $A(z) \neq 0$ for $|z| \leq 1$,
3. There exists $\lambda > 0$ such that $\sup_t \mathbb{E}\|\epsilon_t\|^\lambda < \infty$,
4. $\sum_{k=0}^{\infty} \beta_\lambda(k)^{1-2/\delta} < \infty$,
5. $n^{1/2} \beta_\lambda(s_n) \rightarrow 0$ for some $s_n = o(n^{1/2})$.

Then we have,

$$P(\hat{G}_{CPDAG,n} = G_{CPDAG}) \rightarrow 1 \quad \text{as } n \rightarrow \infty$$

For proving Theorem 4.4.2, we need the following lemma:

Lemma 4.4.1. Under the assumptions of Theorem 4.4.2, $\{X_{ti}, X_{tj}, t \geq 1\}$ is strongly mixing with coefficient $\alpha_{ij}(k) \leq C\beta_\lambda(k)$, where C is a constant. Also, $\mathbf{X}_t, t \geq 1$ satisfies (A.2)* and (A.3)*.

Proof. By Theorem 2.1 in [234], the process is strongly mixing. Note that $\alpha_{ij}(k) \leq 4\|\Delta_k\|_1$ and hence, $\alpha_{ij}(k) \leq C\beta_\lambda(k)$ (by Theorem 2.1 in [234]). This proves Lemma 4.4.1. \square

Proof of Theorem 4.4.2. Theorem 4.4.2 follows from its assumptions, together with Lemma 4.4.1 and Section 4.3. \square

4.5 Consistency of the Time-Aware PC algorithm

In the traditional PC algorithm, (A.4) is assumed, that is, \mathbf{X}_t is faithful with respect to a DAG G , which models causal relations at a fixed t by the DAG G . In practice, this assumption can be unreasonable in a time series scenario which typically includes relationships between variables over time such as, $X_{1u} \rightarrow X_{2v}$. To better accommodate inter-temporal causal relations in a time-series scenario, we assume faithfulness with respect to \mathbf{G} which has edges from variable u at time t_1 to variable v at time t_2 . That is, instead of (A.4), we assume (B.1) as follows.

(B.1): χ_t is faithful with respect to such a DAG \mathbf{G} as above.

Moreover, if one is interested in using partial correlations as the tests for dependence, then the following condition on the time-series may be useful:

(B.0): $\{\mathbf{X}_t\}_{t=1}^n$ is a strictly stationary Gaussian process.

Under (B.1), let $\mathbf{G}_{CPDAG} = (\mathbf{V}, \mathbf{E}(\mathbf{G}_{CPDAG}))$ be the *CPDAG* from the DAG \mathbf{G} . Next, \mathbf{G}_{CPDAG} is transformed to obtain $\mathbf{G}_{CPDAG,R}$ with nodes $1, \dots, p$, and directed edges $\mathbf{E}(\mathbf{G}_{CPDAG,R})$ such that $u \rightarrow v \in \mathbf{E}(\mathbf{G}_{CPDAG,R})$ if and only if $(u, t_1) \rightarrow (v, t_2) \in \mathbf{E}(\mathbf{G}_{CPDAG})$ for some $1 \leq t_1 \leq t_2 \leq \tau$. $\mathbf{G}_{CPDAG,R}$ is the Rolled Graph of \mathbf{G}_{CPDAG} and referred to as the Rolled Markov Graph with respect to χ_t .

For any DAG $\tilde{\mathbf{G}}$ with nodes \mathbf{V} , recall that $pa_{R,\tilde{\mathbf{G}}}(v)$ denotes the parents of v in the Rolled Graph of $\tilde{\mathbf{G}}$. We show that the parents of a node v in a Rolled Markov Graph with respect to χ_t is the union of parents of v in the Rolled Graphs from a Markov equivalence class. For any graph $\tilde{\mathbf{G}} = (\mathbf{V}, \mathbf{E}(\tilde{\mathbf{G}}))$ in the Markov

equivalence class $\mathcal{M}(\mathbf{G})$. Finally, let $pa_R(v)$ denote the parent set of v in $\mathbf{G}_{CPDAG,R}$. Below, we show that $pa_R(v)$ can be expressed as the union of $pa_{R,\tilde{\mathbf{G}}}(v)$ over all rolled graphs $\tilde{\mathbf{G}}$ in the Markov equivalence class of \mathbf{G} .

Lemma 4.5.1. If χ_t is faithful with respect to \mathbf{G} , then

$$pa_R(v) = \bigcup_{\tilde{\mathbf{G}} \in \mathcal{M}(\mathbf{G})} pa_{R,\tilde{\mathbf{G}}}(v)$$

Proof. Choose $v' \in pa_{R,\tilde{\mathbf{G}}}(v)$ for some $\tilde{\mathbf{G}} \in \mathcal{M}(\mathbf{G})$. Then, by definition, $(v', t') \rightarrow (v, t) \in \mathbf{E}(\tilde{\mathbf{G}})$ for some $t' \in t - \tau + 1, \dots, t$. Therefore, $(v', t') \rightarrow (v, t) \in \mathbf{E}(\mathbf{G}_{CPDAG})$. So, $v' \rightarrow v \in \mathbf{G}_{CPDAG,R}$. Therefore, $v' \in pa_R(v)$. This shows that $\bigcup_{\tilde{\mathbf{G}} \in \mathcal{M}(\mathbf{G})} pa_{R,\tilde{\mathbf{G}}}(v) \subseteq pa_R(v)$.

For the other direction, choose $v' \in pa_R(v)$, whence $v' \rightarrow v \in \mathbf{E}(\mathbf{G}_{CPDAG,R})$. Therefore, there exists $t' \leq t$ such that $(v', t') \rightarrow (v, t) \in \mathbf{E}(\mathbf{G}_{CPDAG})$. Hence there exists $\tilde{\mathbf{G}} \in \mathcal{M}(\mathbf{G})$ such that $(v', t') \rightarrow (v, t) \in \mathbf{E}(\tilde{\mathbf{G}})$. By definition $v' \in pa_{R,\tilde{\mathbf{G}}}(v)$. Therefore, $pa_R(v) \subseteq \bigcup_{\tilde{\mathbf{G}} \in \mathcal{M}(\mathbf{G})} pa_{R,\tilde{\mathbf{G}}}(v)$. \square

4.5.1 Main Result

In this section, we show consistency of the Time-Aware PC algorithm.

Theorem 4.5.1 (Consistency of Time-Aware PC algorithm). Assume (B.1), (A.1)-(A.3) (or (A.1)*-(A.3)*).

Denote by $\hat{\mathbf{G}}_{CPDAG,n}$ the estimate of \mathbf{G}_{CPDAG} by the PC algorithm based on $\chi_1, \dots, \chi_{n-\tau+1}$. Then,

$$P(\hat{\mathbf{G}}_{CPDAG,n} = \mathbf{G}_{CPDAG}) \rightarrow 1 \text{ as } n \rightarrow \infty$$

Proof. By Lemma B.1.2, if \mathbf{X}_t satisfies (A.1)-(A.3) (or (A.1)*-(A.3)*), then so does χ_t . Also, note that (B.1) is same as (A.4) with respect to χ_t being faithful with \mathbf{G} . The statement then follows from Theorem 4.2.2. \square

Corollary 4.5.1. Since $\hat{\mathbf{G}}_{CPDAG,n} = \mathbf{G}_{CPDAG} \implies \hat{\mathbf{G}}_{CPDAG,R,n} = \mathbf{G}_{CPDAG,R}$, therefore, also, $P(\hat{\mathbf{G}}_{CPDAG,R,n} = \mathbf{G}_{CPDAG,R}) \rightarrow 1$ as $n \rightarrow \infty$.

Remark 4.5.1 (Subsampled Time-Aware PC for non-stationary processes (TPCNS)). In practice, the entire time series may not be strictly stationary, in such a scenario, the Time-Aware PC algorithm is conducted in

a shorter time-window of length L selected at random, i.e. $[n_i, n_i + L] \subseteq \{1, \dots, \lfloor \frac{n-\tau}{r} \rfloor + 1\}$, and a set of graphs $\hat{\mathbf{G}}_{CPDAG}^{(i)}$ is estimated with $\chi_t, t \in [n_i, n_i + L]$. If there is an edge $(u, t_1) \rightarrow (v, t_2)$ for $t_1 > t_2$ in $\hat{\mathbf{G}}_{CPDAG}^{(i)}$ then the edge is re-oriented to $(u, t_1) \rightarrow (v, t_2)$ in $\hat{\mathbf{G}}_{CPDAG}^{(i)}$, which is then transformed to $\hat{\mathbf{G}}_{CPDAG,R}^{(i)}$. Next, a single graph $\hat{\mathbf{G}}_{CPDAG,R}^*$ is obtained by union of $\hat{\mathbf{G}}_{CPDAG,R}^{(i)}$ while excluding those edges with exceedingly low frequency of occurrence. At the last step, $\hat{\mathbf{G}}_{CPDAG,R}^{(i)}$ is pruned by excluding those edges with exceedingly low edge weights which can be obtained by interventional causal effects.

Example 4.5.1 (VARMA process). Assume now that \mathbf{X}_t is the stationary VARMA process (4.8) satisfying the assumptions of Theorem 4.4.1. In this example, consistency of the Time-Aware PC algorithm holds under B.1, by Theorem 4.5.1.

Example 4.5.2 (Linear Process). Similarly, if \mathbf{X}_t is the stationary linear process (4.4.2) satisfying the assumptions of Theorem 4.4.2, then once again, consistency of the Time-Aware PC algorithm holds under B.1, by Theorem 4.5.1.

4.6 Simulation Studies

We compare the performance of Time-Aware PC for stationary data (TPCS), subsampled Time-Aware PC for non-stationary data (TPCNS), the usual PC algorithm using partial correlation test as well as Hilbert-Schmidt test for conditional independence (denoted TPCSHS, TPCNSHS and PCHS for the latter), and Granger Causality (GC) to recover the ground truth causal relations from four simulation paradigms. The simulation paradigms correspond to specific model assumptions to assess the impact of model specifications on the performance of the approaches (See Appendix B.2).

We generated 25 simulated time series from each simulation paradigm and estimated the causal relationships from each time series. We summarized the performance of the methods in recovering the ground truth causal relationships using the following three metrics: (1) Combined Score (CS), (2) True Positive Rate (TPR), (3) 1 - False Positive Rate (IFPR). Let True Positive (TP) represent the number of correctly identified edges, True Negative (TN) represent the number of correctly identified missing edges, False Positive (FP) represent the number of incorrectly identified edges, and False Negative (FN) represent the number of

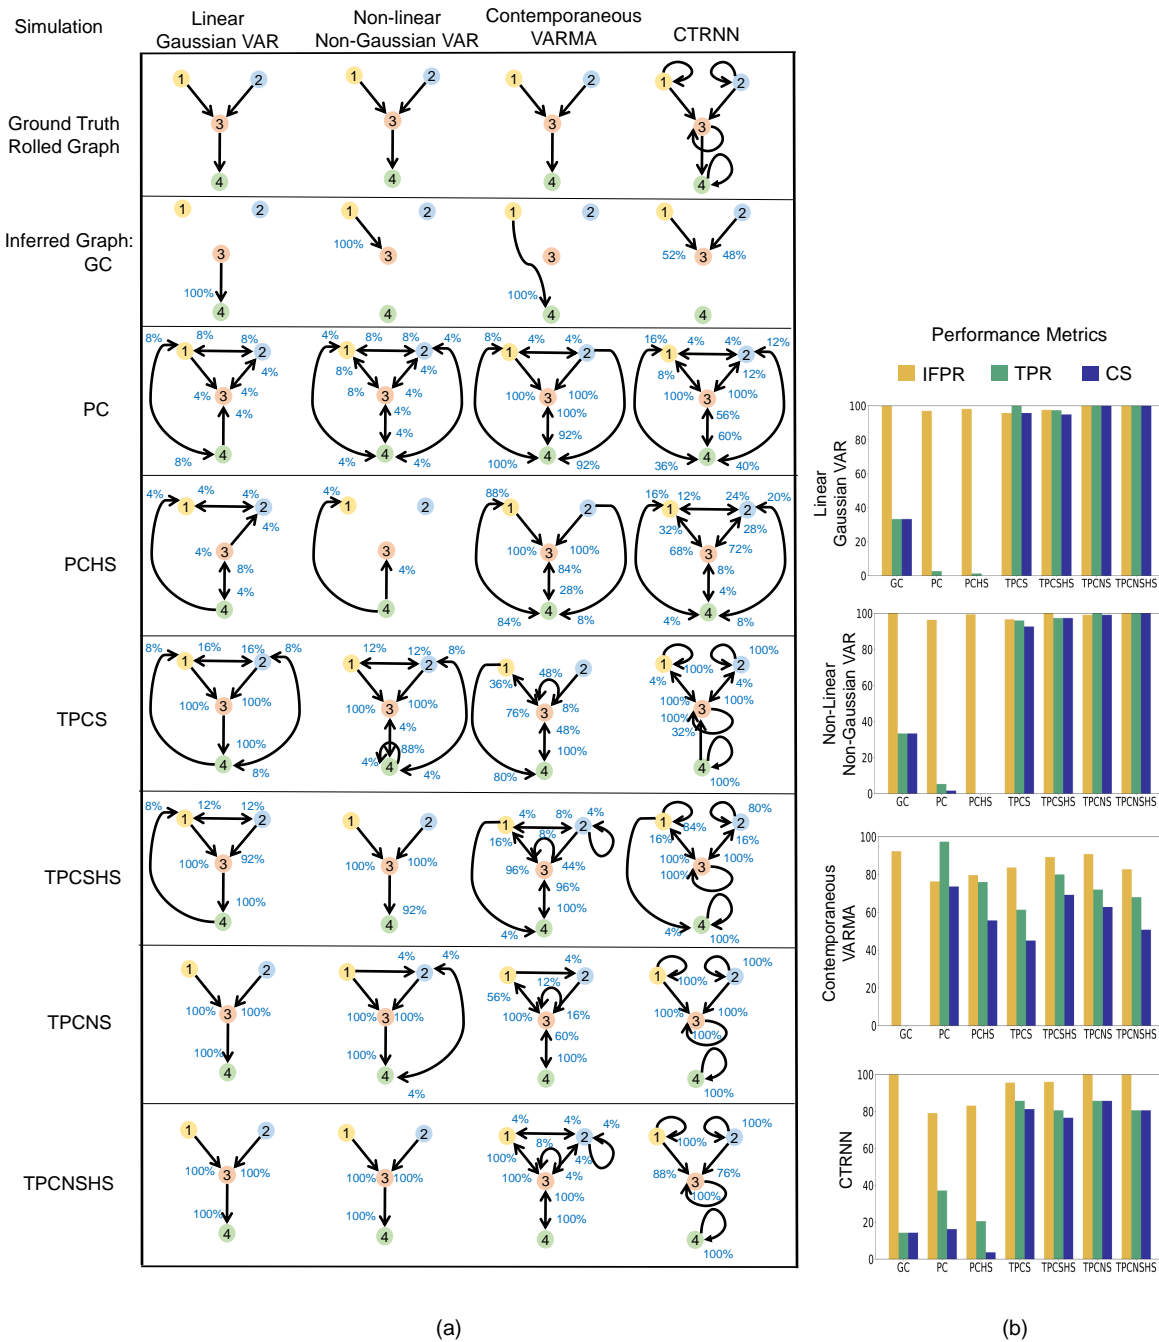


Figure 4.2 (a) The outcomes of GC, PC, PCHS, TPCS, TPCHS, TPCNS and TPCNSHS are compared on four examples of motifs and simulation paradigms; from left to right: Linear Gaussian VAR, Non-linear Non-Gaussian VAR, Contemporaneous VARMA and Continuous-Time Recurrent Neural Network (CTRNN). Table: 4-neurons motifs that define the Ground Truth (row 1) are depicted along with inferred graphs over simulation instances by the different methods (row 2-8). Each inferred graph has an edge $v \rightarrow w$ that corresponds to an edge detected in any of the inference instances. The percentage (blue) next to each edge indicates the number of times the edge was detected out of all instances. (b) IFPR (gold), TP rate (green) and Combined Score (blue) of each method are shown for each motif.

incorrectly identified missing edges across simulations. IFPR is defined as:

$$\text{IFPR} = \left(1 - \frac{\text{FP}}{\text{FP} + \text{TN}}\right) \cdot 100,$$

which measures the ratio of the number of correctly identified missing edges by the algorithm to the total number of true missing edges. Note that the rate is reported such that 100% corresponds to no falsely detected edges. TPR is defined as $\text{TPR} = \left(\frac{\text{TP}}{\text{TP} + \text{FP}}\right) \cdot 100$ as i.e. the ratio of the number of correctly identified edges by the algorithm to the total number of true edges in percent. The Combined Score (CS) is given by Youden’s Index [185; 186], as follows, $\text{CS} = \text{TPR} - \text{FPR}$.

In implementations of the PC algorithm in the *pcalg* package in *R* and other software such as *TETRAD^{IV}*, the Gaussian conditional dependence tests use a fixed level α and $\sqrt{n - k - 3} g(\hat{\rho}) \leq \Phi^{-1}(1 - \alpha)$, thereby we use $\gamma = \Phi^{-1}(1 - \alpha) / \sqrt{n - k - 3}$ which also gives rise to a consistent test. The PC algorithm with Hilbert-Schmidt conditional independence criterion is implemented using the *kpcalg* package in *R*.

The true graphs in each of the simulations consist of 4 nodes and 16 possible edges (including self-loops) between the nodes. Therefore for the 25 simulations in a simulation setting, there are a total of 25 graphs to infer and 400 possible edges. Figure 4.2 shows the ground truth graph in each simulation setting for noise standard deviation $\eta = 1$ and performance of the estimation for $\alpha = 0.05$. We report the percentage of the simulations in which an edge is estimated to be present. The percentages of occurrence among simulations are also shown in Table 4.1. We also compare the approaches in estimating the true edges by TPR, IFPR and CS scores of performance. A higher percentage indicates higher confidence in the detection of that edge. Figure 4.4 compares the Combined Score of performance of the approaches over different values of α and noise standard deviation η in each simulation setting.

Linear Gaussian VAR								
Edge	Ground Truth	GC	PC	PCHS	TPCS	TPCSHS	TPCNS	TPCNSHS
1 → 1	0	0	0	0	0	0	0	0
1 → 2	0	0	8	4	16	12	0	0
1 → 3	100	0	4	0	100	100	100	100
1 → 4	0	0	8	0	8	0	0	0

2 → 1	0	0	8	4	16	12	0	0
2 → 2	0	0	0	0	0	0	0	0
2 → 3	100	0	4	0	100	92	100	100
2 → 4	0	0	0	0	8	0	0	0
3 → 1	0	0	0	0	0	0	0	0
3 → 2	0	0	4	4	0	0	0	0
3 → 3	0	0	0	0	0	0	0	0
3 → 4	100	100	0	4	100	100	100	100
4 → 1	0	0	8	0	8	8	0	0
4 → 2	0	0	4	0	8	0	0	0
4 → 3	0	0	4	8	0	0	0	0
4 → 4	0	0	0	0	0	0	0	0

Non-linear Non-Gaussian VAR

Edge	Ground Truth	GC	PC	PCHS	TPCS	TPCSHS	TPCNS	TPCNSHS
1 → 1	0	0	0	0	0	0	0	0
1 → 2	0	0	8	0	12	0	4	0
1 → 3	100	100	8	0	100	100	100	100
1 → 4	0	0	4	0	0	0	0	0
2 → 1	0	0	8	0	12	0	0	0
2 → 2	0	0	0	0	0	0	0	0
2 → 3	100	0	4	0	100	100	100	100
2 → 4	0	0	4	0	4	0	4	0
3 → 1	0	0	8	0	0	0	0	0
3 → 2	0	0	4	0	0	0	0	0
3 → 3	0	0	0	0	0	0	0	0
3 → 4	100	0	4	0	88	100	100	100
4 → 1	0	0	4	4	0	0	0	0
4 → 2	0	0	4	0	8	0	4	0

4 → 3	0	0	4	4	4	0	0	0
4 → 4	0	0	0	0	4	0	0	0

Contemporaneous VARMA

Edge	Ground Truth	GC	PC	PCHS	TPCS	TPCSHS	TPCNS	TPCNSHS
1 → 1	0	0	0	0	0	0	0	0
1 → 2	0	0	4	0	0	8	4	4
1 → 3	100	0	100	100	76	96	100	100
1 → 4	0	100	100	84	80	4	0	0
2 → 1	0	0	4	0	0	4	0	4
2 → 2	0	0	0	0	0	4	0	4
2 → 3	100	0	100	100	8	44	16	4
2 → 4	0	0	92	8	0	0	0	0
3 → 1	0	0	0	0	36	16	56	100
3 → 2	0	0	0	0	0	0	0	4
3 → 3	0	0	0	0	48	8	12	8
3 → 4	100	0	92	28	100	100	100	100
4 → 1	0	0	8	88	0	0	0	0
4 → 2	0	0	0	0	0	0	0	0
4 → 3	0	0	100	84	48	96	60	100
4 → 4	0	0	0	0	0	0	12	0

CTRNN

Edge	Ground Truth	GC	PC	PCHS	TPCS	TPCSHS	TPCNS	TPCNSHS
1 → 1	100	0	0	0	100	64	100	100
1 → 2	0	0	4	24	0	0	0	0
1 → 3	100	52	100	68	100	100	100	88
1 → 4	0	0	36	4	0	4	0	0
2 → 1	0	0	4	12	0	0	0	0
2 → 2	100	0	0	0	100	80	100	100

$2 \rightarrow 3$	100	48	100	72	100	100	100	76
$2 \rightarrow 4$	0	0	40	8	0	0	0	0
$3 \rightarrow 1$	0	0	8	32	4	16	0	0
$3 \rightarrow 2$	0	0	12	28	4	16	0	0
$3 \rightarrow 3$	100	0	0	0	100	100	100	100
$3 \rightarrow 4$	100	0	60	4	0	0	0	0
$4 \rightarrow 1$	0	0	16	16	0	0	0	0
$4 \rightarrow 2$	0	0	12	20	0	0	0	0
$4 \rightarrow 3$	0	0	56	8	32	0	0	0
$4 \rightarrow 4$	100	0	0	0	100	100	100	100

Table 4.1 Percentage of occurrence of edges among simulation trials.

In *Linear Gaussian VAR* (left column in Figure 4.2), GC estimates a single edge $3 \rightarrow 4$ among the three edges of the true graph in every iteration leading to IFPR 100%, TPR 33.3% and CS 33.3%. In comparison, PC and PCHS do not estimate any edge more than 10% of the iterations. Overall, PC has IFPR 96.9%, TPR 2.7% and CS -0.4% and closely followed by PCHS. TPCS and TPCHS estimates the true edges for $1 \rightarrow 3$ and $3 \rightarrow 4$ for 100% of the iterations, and the true edge $2 \rightarrow 3$ in 100%, 92% of the iterations for TPCS and TPCHS respectively. There are other spurious edges also obtained in 8%-16% of the iterations. Overall, TPCS has IFPR 95.7%, TPR 100%, CS 95.7% and TPCHS has IFPR 97.5%, TPR 97.3% and CS 94.9%. TPCNS and TPCNSHS obtains the Ground Truth, with no spurious edges and obtains the true edges in all of the trials ($1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$ with 100%, 100% and 100% respectively), overall both having IFPR, TPR and CS of 100%. Thereby, among the three methods, we conclude that TPCNS, TPCNSHS, TPCS and TPCSHS have improvement in performance by a great margin (+66.7%) in CS while TPCNS and TPCNSHS detects the true edges perfectly. The non-detection of the edges by PC algorithm can be explained by Theorem 4.4.1, since at a fixed time the simulated process is faithful with respect to the empty graph, even though there are across-time causal relations leading to the ground truth rolled graph as in Figure 4.2.

In the *Non-linear Non-Gaussian VAR* (second column), as previously, GC always detects a single edge $1 \rightarrow 3$ with 100% out of the three true edges. PC and PCHS does not detect any edge with greater than

8% occurrence in iterations. In contrast, TPCS outputs the true edges with a high percentage ($1 \rightarrow 3$, $2 \rightarrow 3$, $3 \rightarrow 4$ with 100%, 100%, 88%) and some spurious edges with a low percentage of less than 12%, while TPCSHS outputs the true edges with a high percentage ($1 \rightarrow 3$, $2 \rightarrow 3$, $3 \rightarrow 4$ with 100%, 100%, 92%) and without spurious edges. TPCNS detects the true edges in all the trials ($1 \rightarrow 3$, $2 \rightarrow 3$, $3 \rightarrow 4$ with 100%, 100%, 100%) and relatively lower spurious edges in 4% of the trials. TPCNSHS detects the true edges in all the trials ($1 \rightarrow 3$, $2 \rightarrow 3$, $3 \rightarrow 4$ with 100%, 100%, 100%) and no spurious edges. In summary, all approaches yielded IFPR close to 100%, while for TPR TPCNSHS and TPCNS outperformed the other approaches with TPR 100%, followed by TPCSHS and TPCS with 97%, 96% respectively, and GC with 33.3%, and for CS also, TPCNSHS and TPCNS outperformed at 100%, 99.1%, closely followed by TPCSHS and TPCS with 97.3%, 92.6% and GC at 33.3%. For this scenario, TPCNSHS has the best performance followed closely by other variants of TPC such as TPCSNS, TPCSHS and TPCS among the methods.

In *Contemporaneous VARMA (third column)*, the ground truth causal graph encode the linear dependence between variables at each fixed time as well as weaker relationship following the same rolled graph across time. GC obtains a spurious edge $1 \rightarrow 4$ for 100% of the iterations. PC estimates the true edges $1 \rightarrow 3$, $2 \rightarrow 3$, $3 \rightarrow 4$ in 100%, 100%, 92% of the trials, TPCS in 76%, 8%, 100%, TPCSHS in 96%, 44%, 100%, and TPCNS in 100%, 16%, 100% trials with other spurious edges also detected. In summary, GC, PC, TPCS and TPCNS yielded IFPR = 92.3%, 76.3%, 83.7%, 90.7%, TPR = 0%, 97.3%, 61.3%, 72.0% and CS = -7.7%, 73.5%, 45.0%, 62.8%. TPCNS performed better than TPCS with IFPR 76%, TPR 80%, CS 80%. For this scenario, PC has the highest performance, closely followed by TPCSHS, PCHS, TPCNS, TPCNSHS, TPCS and then followed by GC. This is an impractical scenario when the DAG encoding causal relations at each fixed time is also the same as the Ground Truth Rolled Graph, and the PC algorithm is shown to be consistent in Theorem 4.4.1 for the causal relations at each fixed time.

In the *CTRNN scenario (fourth column)*, self-loops are present for each neuron. GC obtains two of the three true non-self edges $1 \rightarrow 3$, $2 \rightarrow 3$ for 52%, 48% of the trials. PC detects spurious edges in up to 12% of the trials, but also infers the non-self true edges $1 \rightarrow 3$, $2 \rightarrow 3$ for 100% of the trials. TPCS outputs all the self true edges and non-self true edges $1 \rightarrow 3$, $2 \rightarrow 3$ for 100% of the trials while also detecting false edges in relatively few 4% of the trials. In comparison, TPCNS infers no false edges and all the self true edges for 100% of the trials and non-self true edges $1 \rightarrow 3$ and $2 \rightarrow 3$ for 100%, 100% of the trials. TPCNSHS also

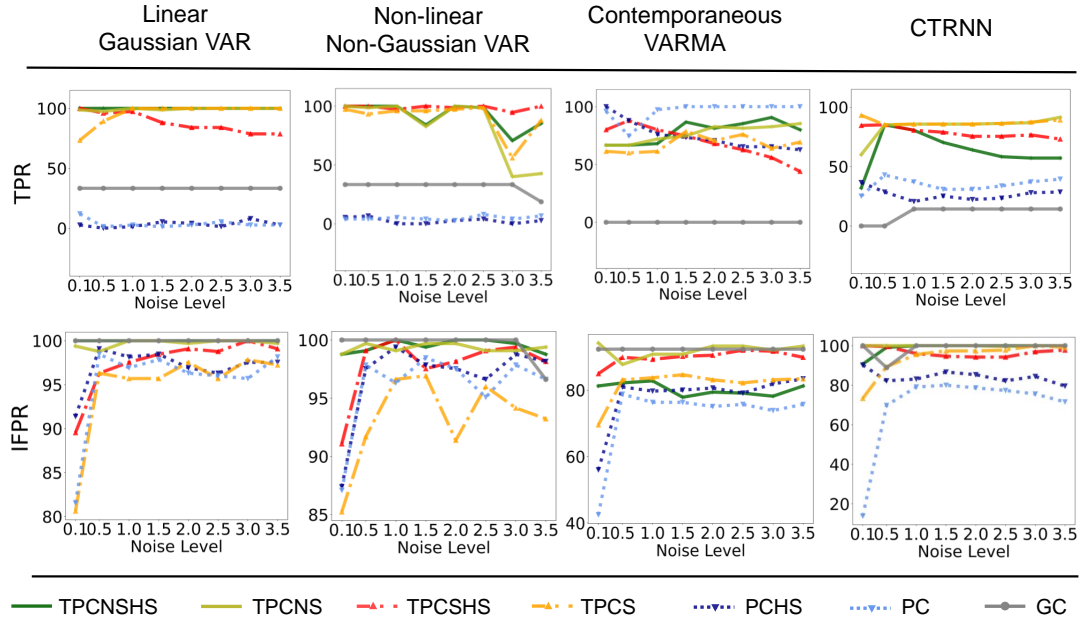


Figure 4.3 Performance of the algorithms - TPCNSHS (dark green), TPCNS (light green), TPCSHS (red), TPCS (orange), PCHS (dark blue), PC (light blue), GC(gray), as measured by TPR and IFPR in different simulation paradigms: Linear Gaussian VAR, Non-linear Non-Gaussian VAR, Contemporaneous VARMA and CTRNN, over increasing levels of noise variance and $\alpha = 0.05$. Note that the y axis range is not common between the rows.

infers no false edges and all the self true edges for 100% of the trials and non-self true edges $1 \rightarrow 3$ and $2 \rightarrow 3$ for 88%, 78% of the trials. In summary, IFPR of GC, PC, TPCS and TPCNS are 100%, 79.1%, 95.6%, 100%, that for TPR are 14.3%, 37.1%, 85.7%, 85.7%, and that for CS are 14.3%, 16.3%, 81.3%, 85.7% respectively and the HS versions closely follow in performance. Among all methods, TPCS and TPCNS have the highest TPR together, followed by PC and lastly GC. TPCNS and GC have the highest IFPR having not detected any false edges, followed by TPCS and then PC. In terms of CS, TPCNS and TPCNSHS have the highest performance, followed by TPCS and TPCSHS, compared to other methods.

We compare the TPR and IFPR of the approaches across varying levels of noise standard deviation η from 0.1 to 3.5 and $\alpha = 0.05$ in Figure 4.3. In the Linear Gaussian scenario, we note that TPCNS and TPCNSHS has a TPR of $\approx 100\%$ across all levels of simulation noise, and is followed by TPCS and TPCNSHS, and then by GC in performance and lastly PC. In terms of IFPR also, TPCNSHS and TPCNS have the best performance overall. In the Non-linear Non-Gaussian scenario, TPCSHS has a TPR of $\approx 100\%$, while TPCNSHS has some fluctuations below 100% at a few noise levels, and all the TPC variants have a TPR

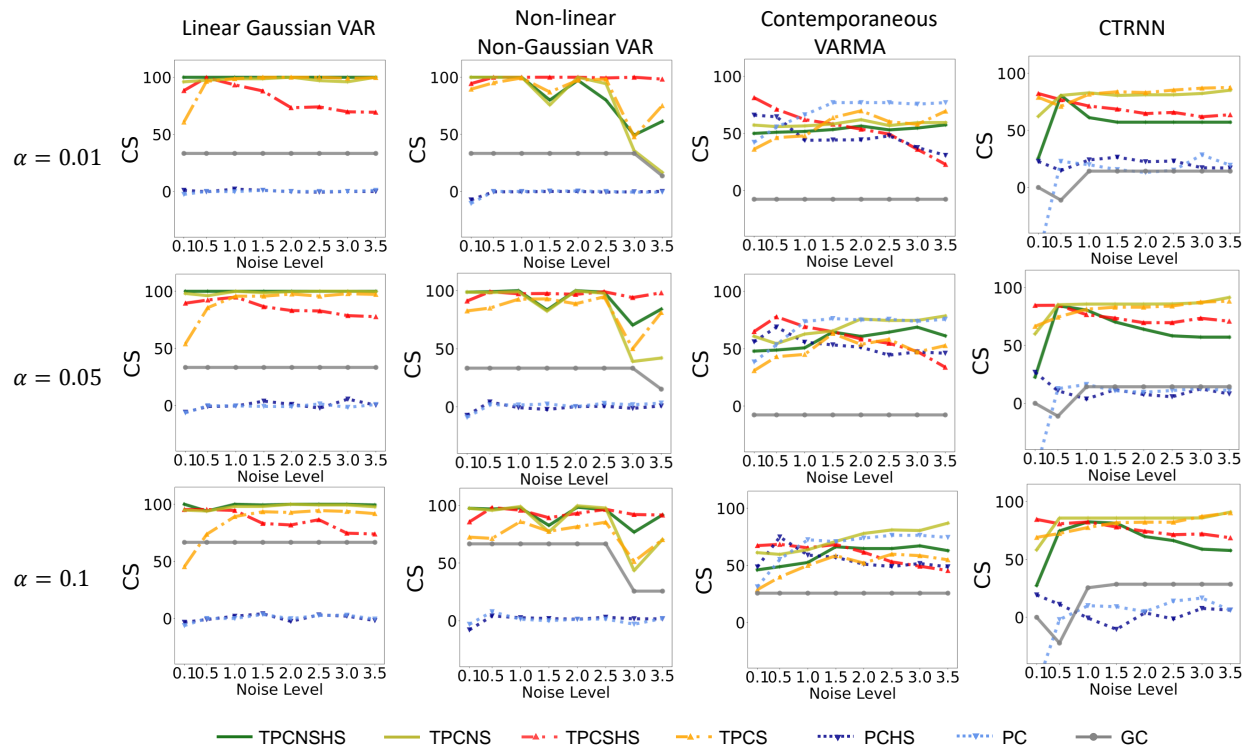


Figure 4.4 Combined Score of the four methods - TPCNSHS (dark green), TPCNS (light green), TPCS (orange), PCHS (dark blue), PC (light blue), GC (gray), over varying noise levels in simulation $\eta = 0.1, 0.5, 1.0, \dots, 3.5$, for simulated motifs from Linear Gaussian VAR, Non-linear Non-Gaussian VAR, Contemporaneous VARMA and CTRNN paradigms (left to right), with $\alpha = 0.01, 0.05, 0.1$ for conditional dependence tests (top to bottom).

exceeding that of PC and GC. In terms of IFPR, TPCNSHS and TPCNS have an IFPR of $\approx 100\%$ and exceeds the other approaches. In the Contemporaneous scenario, PC and TPCNS have comparatively higher TPR while TPCNSHS, TPCNS and TPCSHS have comparatively higher IFPR than other approaches. In the CTRNN scenario, both the TPR and IFPR of TPCS and TPCNS are higher compared to the other methods over the different noise levels. We also compare the Combined Score of the approaches across varying levels of simulation noise η from 0.1 to 3.5 and thresholding parameter $\alpha = 0.01, 0.05, 0.1$ in Figure 4.4. In the Linear Gaussian scenario, we note that TPCNSHS and TPCNS have a CS of $\approx 100\%$ across all levels of simulation noise and thresholding parameter α , and is closely followed by TPCS, then TPCSHS, GC, and lastly PC. In the Non-linear Non-Gaussian scenario, TPCSHS has the highest CS, closely followed by TPCS, TPCNS, and TPCNSHS, compared to other methods across levels of noise and α . In the Contemporaneous scenario, PC has better CS for $\alpha = 0.01$ and for $\alpha = 0.05, 0.1$ TPCNS exceeds in performance for higher levels of noise, closely followed by PC and other variants of TPC and then GC. In the CTRNN scenario, TPCNS and TPCS have the highest CS overall followed by TPCNSHS and TPCSHS, compared to the other methods.

4.7 Application to River Runoff Benchmark Data

We also conducted analysis of a real benchmarking dataset in the public benchmarking platform - *CauseMe* [187]. We used the *River Runoff* dataset (See Appendix B.3). The River Runoff dataset is expected to have contemporaneous as well as across-time interactions with noise and interaction strengths not controlled for, thereby would demonstrate the performance of the methods in an empirical setting. We compare the approaches selVAR [189], SLARAC [235], PCMCI-GPDC [188] - which are among the top of the leaderboard for performance on the benchmarking dataset, GC, PC, TPC (Our) and TPCNS (Our). We used 1 - False Positive Rate (IFPR), True Positive Rate (TPR) and Combined Score given by Youden's Index ($CS = TPR - FPR$) (See Table 4.2) to measure the performance of the algorithms.

In terms of performance with respect to CS, TPCNS yielded the best performance with a CS of 72%, closely followed by TPCS at 68% and selVAR, SLARAC, PCMCI-GPDC, PC and GC at 54%, 50%, 45%, 38%, 37% respectively. TPCNS and TPCS exceeds the best among the existing approaches **by 18%** and 14% respectively. With respect to TPR, both TPCNS and PCMCI achieve the highest TPR at 100%, followed by selVAR,

Combined [True, 1-False] Rates (%)	
Algorithm	River-Runoff (Real)
GC	37 [45, 92]
PC	38 [55, 83]
PCMCI	45 [100, 45]
SLARAC	50 [86, 64]
PCHS	50 [64, 86]
TPCSHS	50 [64, 87]
seIVAR	54 [91, 63]
TPCS	68(+14%) [82, 86]
TPCNS	72(+18%) [82, 90]
TPCNSHS	72(+18%) [100, 72]

Table 4.2 Comparison of performance on benchmarking datasets. For each dataset, each method’s Combined Score, True Positive Rate, and 1-False Positive Rate are reported (Higher value is better).

SLARAC, TPCNS, TPCS, PC and GC at 91%, 86%, 82%, 82%, 55%, 45% respectively. With respect to IFPR, GC has the best performance with a score of 92% closely followed by TPCNS and TPCS at 90%, 86%, PC at 83%, and SLARAC, seIVAR and PCMCI with at 64%, 63%, 45% respectively. PC and GC have a higher IFPR than one of TPCS or TPCNS because they detect less number of false edges, but that is achieved at the expense of detecting a less number of edges altogether including a less number of true edges which in turn leads to their TPR among the lowest in Table 4.2. In comparison, TPCNS and TPCS detects true edges with greater sensitivity, thereby achieving a higher TPR. TPCS and TPCNS better maintains the trade-off between TPR and FPR, thereby leading to a better CS in comparison to the other methods.

Since the coupling between variables as well as noise are not controlled and contemporaneous interactions are expected to be present as the sampling resolution is greater than the time taken for interactions between the variables, the real dataset provides a challenge for the methods. TPCS and TPCNS outperform other approaches by a CS of 14% and 18% respectively and shows significant improvement in performance than the other methods. This performance in a benchmarked real data setting demonstrates the applicability of TPC to real data scenarios.

4.8 Discussion

We show that the PC algorithm is consistent for estimating the equivalence class of the DAG between variables from a stationary Gaussian time series with ρ -mixing properties, thereby demonstrating the asymptotic correctness of the PC algorithm in a temporal dependence setting. This enables us to show consistency of the PC algorithm in popular time series such as vector auto-regressive models and linear processes.

We also show that the Time-Aware PC algorithm consistently estimates the equivalence class of the DAG whose nodes are variables at different times using the PC algorithm and then transforms it into a rolled Markov graph between the variables. In contrast to PC, the consistency of the Time-Aware PC algorithm assumes faithfulness with respect to a DAG whose nodes are tuples of variable and time. The latter indicates causal relations between variables at different times. Therefore, in the time series scenario, while the PC algorithm supports finding of contemporaneous causal relations between the variables, Time-Aware PC provides a more general setting in which it first finds causal relations between the variables across different times and then transforms the temporal causal relations into the rolled Markov graph between the variables.

We compare the performance in recovery of ground truth causal relations for PC, GC and Time-Aware PC for stationary and non-stationary time series in simulation studies with different model assumptions and benchmark real data. The results indicate the greater performance of Time-Aware PC in general when interactions in the time series are across different times and greater performance of PC in the specific scenario when the interactions are restricted to be contemporaneous.

Chapter 5

Causal Functional Connectivity in Alzheimer's Disease Computed from Time Series fMRI data

This chapter is a reproduction of published work in Biswas and Sripada [4], with minor adaptations to fit the thesis structure and formatting.

Alzheimer's disease (AD) is the most common age-related progressive neurodegenerative disorder. It typically begins with a preclinical phase and advances through mild cognitive impairment (MCI) to clinically significant AD, a form of dementia [236]. Despite substantial efforts to identify biomarkers for AD, it still relies on clinical diagnosis, and early and accurate disease prediction remains limited [237; 238]. Abnormal resting-state functional connectivity (FC) between brain regions has been observed as early as two decades before brain atrophy and the emergence of AD symptoms [239; 240]. Therefore, resting-state FC can potentially determine the relative risk of developing AD [241; 242].

Resting-state functional magnetic resonance imaging (rs-fMRI) records the blood-oxygen-level-dependent (BOLD) signals from different brain regions while individuals are awake and not engaged in any specific task. The BOLD signal is popularly used to infer FC between brain regions partly due to the advantage that BOLD signal provides high spatial resolution [243; 244; 245; 246].

FC refers to the stochastic relationship between brain regions with respect to their activity over time.

Popularly, FC involves measuring the statistical association between signals from different brain regions. The statistical association measures are either pairwise associations between pairs of brain regions, such as Pearson's correlation, or multivariate i.e. incorporating multi-regional interactions such as undirected graphical models [1]. Detailed technical explanations of FC in fMRI can be found in Chen et al. [247]; Keilholz et al. [248]; Scarapicchia et al. [249]. The findings from studies using FC [250; 251], and meta-analyses [252; 253; 254] indicate a decrease in connectivity in several brain regions with AD, such as the posterior cingulate cortex and hippocampus. These regions play a role in attentional processing and memory. On the other hand, some studies have found an increase in connectivity within brain regions in the early stages of AD and MCI [255; 256; 257]. Such an increase in connectivity is a well known phenomenon that occurs when the communication between other brain regions is impaired. Such hyperconnectivity has been interpreted as a compensatory mechanism where alternative paths within the brain's network are recruited [257; 258; 259].

In contrast to Associative FC (AFC), Causal FC (CFC) represents functional connectivity between brain regions more informatively by a directed graph, with nodes as the brain regions, directed edges between nodes indicating causal relationships between the brain regions, and weights of the directed edges quantifying the strength of the corresponding causal relationship [52]. However, functional connectomics studies in general, and those concerning fMRI from AD in particular, have predominantly used associative measures of FC [109]. There are a few studies that deal with comparing broad hypotheses of alteration within the CFC in AD [260; 261]. However, this area is largely unexplored, partly due to the lack of methods that can infer CFC in a desirable manner, as explained next.

Several properties are desirable in the context of causal modeling of FC [1; 23]. Specifically, the CFC should represent causality while free of limiting assumptions such as linearity of interactions. In addition, since the activity of brain regions are related over time, such temporal relationships should be incorporated in defining causal relationships in neural activity. The estimation of CFC should be computationally feasible for the whole brain FC instead of limiting it to a smaller brain network. It is also desirable to capture beyond-pairwise multi-regional cause-and-effect interactions between brain regions. Furthermore, since the BOLD signal occurs and is sampled at a temporal resolution that is far slower than the neuronal activity, thereby causal effects often appear as contemporaneous [112; 23]. Therefore, the causal model in fMRI data

should support contemporaneous interactions between brain regions.

Among the methods for finding CFC, *Dynamic Causal Model* (DCM) requires a mechanistic biological model and compares different model hypotheses based on evidence from data, and is unsuitable for estimating the CFC of the whole brain [141; 23]. On the other hand, Granger Causality (GC) typically assumes a vector auto-regressive linear model for the activity of brain regions over time, and it tells whether a regions's past is predictive of another's future [211]. Furthermore, GC does not include contemporaneous interactions. This is a drawback since fMRI data often consists of contemporaneous interactions [23]. In contrast, *Directed Graphical Modeling* (DGM) has the advantage that it does not require the specification of a parametric equation of the neural activity over time, it is predictive of the consequence of interventions, and supports estimation of whole brain CFC. Furthermore, the approach inherently goes beyond pairwise interactions to include multi-regional interactions between brain regions and estimating the cause and effect of such interactions. The *Time-aware PC* (TPC) algorithm is a recent method for computing the CFC based on DGM in a time series setting [2]. In addition, TPC also accommodates contemporaneous interactions among brain regions. A detailed comparative analysis of approaches to find CFC is provided in Biswas and Shlizerman [1, 2]. With the development of methodologies such as TPC, it would be possible to infer the whole brain CFC with the aforementioned desirable properties.

In this chapter, we apply the TPC algorithm to infer the CFC between brain regions from resting-state fMRI data. The TPC algorithm estimates the subject-specific CFC for each subject from their fMRI data. We compare the CFC outcome of TPC with GC and Sparse Partial Correlation (SPC), which are approaches to find the CFC and AFC, respectively. We then use the CFC outcome of TPC to investigate the alteration of CFC in AD. In this regard, we conducted an exploratory analysis for the difference in strength of causal connections in AD compared to CN subjects (and MCI compared to CN subjects), based on their edge-wise p-values given by Welch's t-test. We reported the resulting CFC edges with lowest edge-wise p-values for altered connectivity in AD compared to CN subjects and their corresponding brain regions. The brain regions identified in those analyses are consistent with published literature on regions impacted by AD, with each such publication being a report from a team involving a clinical setting and at least one medical expert, thereby validating the approach.

Characteristic	CN	MCI	AD	<i>P</i>
Number of subjects	41	54	34	-
Sex (M/F)	19/22	29/26	16/18	0.16
Age (years)	74.9 ± 6.4	74.2 ± 7.1	74.4 ± 7.4	0.86
Education (years)	16.5 ± 2.3	15.7 ± 2.6	15.4 ± 2.5	0.22
MMSE	29.1 ± 1.4	27.8 ± 1.9	21.9 ± 4.2	< 10 ⁻¹⁴

Table 5.1 Summary of demographic information and Mini Mental State Examination (MMSE) for CN, MCI and AD subjects. The second to fourth columns present group characteristics, mean ± SD. The fifth column presents *P* values for the statistical significance of the inter-group differences. Differences in Sex was assessed using a Chi-Squared test and differences in Age, Education and MMSE using non-parametric analysis of variance by Kruskal-Wallis test.

5.1 Materials and Methods

5.1.1 Participants

The resting fMRI and demographic data were downloaded from the Alzheimer’s Disease Neuroimaging Initiative (ADNI; <http://adni.loni.usc.edu/>). A total of 129 subjects were included in the study: 41 subjects who are CN, 54 subjects with MCI, and 34 subjects with AD.

Table 5.1 includes a summary of the participants’ demographic and medical information. In the experiments, the subjects with AD presented significantly lower scores in the screening assessment cognitive test Mini-Mental State Examination (MMSE) in comparison with the other groups. The subjects were age-matched (Kruskal-Wallis test: $p > 0.8$), gender-matched (Chi-Squared test: $p > 0.1$), and matching number of years of education (Kruskal-Wallis test: $p > 0.2$). As expected, MMSE scores had a significant difference between all pairs of groups (Kruskal-Wallis test: $p < 10^{-14}$).

5.1.2 Image Acquisition

The acquisition of fMRI images was performed using Philips Medical Systems scanner. The fMRI images were obtained using an echo planar imaging sequence at a field strength of 3.0 Tesla, with a repetition time (TR) of 3 seconds, an echo time (TE) of 30 milliseconds, and a flip angle of 80 degrees. The matrix size was 64 × 64 pixels, 140 volumes, 48 slices per volume, slice thickness of 3.3 mm, and voxel size of 3.3 × 3.3 × 3.3 mm³.

5.1.3 fMRI Preprocessing

The fMRI pre-processing steps were carried out using the CONN toolbox version 21a, which utilizes the Statistical Parametric Mapping (SPM12), both of which are MATLAB-based cross-platform software [262; 263]. We used the default pre-processing pipeline in CONN, consisting of the following steps in order: functional realignment and unwarp (subject motion estimation and correction), functional centering to (0,0,0) coordinates (translation), slice-time correction with interleaved slice order, outlier identification using Artifact Detection and Removal Tool, segmentation into gray matter, white matter and cerebrospinal fluid tissue, and direct normalization into standard Montreal Neurological Institute (MNI) brain space, and lastly, smoothing using spatial convolution with a Gaussian kernel of 8mm full-width half maximum. This pipeline was followed by detrending and bandpass filtering (0.001-0.1 Hz) to remove low-frequency scanner drift and physiological noise in the fMRI images. The first four time points have been filtered out to remove any artifacts.

For the extraction of Regions-Of-Interest (ROIs), the automated anatomical labeling (AAL) atlas was utilized on the pre-processed rs-fMRI dataset [264]. The list of all regions in the AAL atlas is provided in Appendix C.1 along with their abbreviated, short, and full region names. This parcellation method has been demonstrated to be optimal for studying the FC between brain regions [265]. The voxels within each ROI were averaged, resulting in a time series for each ROI.

5.1.4 Inference of causal functional connectivity: Time-aware PC algorithm

The TPC Algorithm finds CFC between brain regions from time series based on DGM [52; 48; 1; 2; 3]. While traditional DGM applies to static data, TPC extends the applicability of DGM to CFC inference in time series by first implementing the Directed Markov Property to model causal spatial and temporal interactions in the time series by an unrolled Directed Acyclic Graph (DAG) of the time series. The unrolled DAG consists of nodes (v, t) , for region of interest v and time t , and edge $(v_1, t_1) \rightarrow (v_2, t_2)$ reflecting causal interaction from the BOLD signal in region v_1 at time t_1 to the BOLD signal in region v_2 at time t_2 . The estimation of the unrolled DAG is carried out by first transforming the time series into sequential variables with a maximum time delay of interaction τ and then applying the Peter-Clark (PC) algorithm to infer the unrolled DAG based on the sequential variables [57]. TPC then rolls the DAG back to obtain the CFC graph between the regions




Figure 5.1 Steps conveying the concept of the TPC algorithm to infer CFC from observed neural time series data: First the neural time series is transformed to form sequential samples with a maximum time delay of interaction, τ (here $\tau = 1$). Then, Peter-Clark (PC) algorithm is applied on the sequential samples to obtain the unrolled DAG satisfying the Directed Markov Property. Finally the unrolled DAG is transformed to obtain the Rolled CFC between regions. This figure focuses on the key steps of the TPC algorithm from Figure 3.2: firstly Step 1 to construct Time Advanced Samples referred here as Sequential Samples, then Step 3 that applies the PC algorithm, and lastly Step 4 that transforms the Unrolled DAG into the Rolled graph, while excluding the other steps of Subsampling, Robust edges and Pruning.

of interest (see Figure 5.1) [2]. We consider $\tau = 1$ for our analyses, which would include interactions of the BOLD signal between regions of interest with a maximum time delay of 3 s, the TR of the fMRI acquisition. The Python package *TimeAwarePC* is used for implementation [2].

The CFC outcome of this methodology is interpretable in the following manner: An edge from region $i \rightarrow j$ in the CFC estimate represents significant causal interaction from brain region i at preceding times to region j at following times. The model and the approach are non-parametric, meaning that it does not require the specification of a parametric dynamical equation for neural activity. The method captures beyond-pairwise multivariate interactions between brain regions. It also supports the estimation of the CFC for the whole brain in a computationally feasible manner. It also allows for time delays in interactions between the brain units and the presence of feedback loops. Furthermore, it has been shown that if the neural activity obeys an arbitrary dynamical process, the model outcome of TPC is consistent with respect to the causal relationships implied by the dynamical process and is predictive of counterfactual queries such as ablation or modulation [2].

It is noteworthy that implementing the Directed Markov Property on the unrolled DAG to model causal relationships over time enables contemporaneous interactions e.g., from region u to region v at time t [2]. Such contemporaneous interactions are represented by the edge $(u, t) \rightarrow (v, t)$ in the unrolled DAG, and the presence of such an edge in the unrolled DAG would be reflected as an edge $u \rightarrow v$ in the Rolled CFC

outcome. Such contemporaneous interactions are especially relevant in fMRI due to the relatively slow temporal resolution of the BOLD signal compared to the underlying neural activity [23].

5.1.5 Comparison with functional connectivity using other approaches

In Biswas and Shlizerman [2], the authors have demonstrated that TPC performs better in computing CFC compared to other methods such as GC on simulated and public benchmarking datasets as well as on a real neurobiological dataset of single neuron signals obtained using Neuropixels. Additionally, the authors have drawn contrast (in Biswas and Shlizerman, 2022) with SPC, which is a popular method for inferring AFC. In this chapter, we computed AFC using SPC and CFC using GC from fMRI data [266; 267]. We compared these two outputs with the CFC obtained by TPC from fMRI data. The GC graph is computed using the Nitime Python library, which fits a Multi-variate Auto-Regressive (MVAR) model followed by the use of GrangerAnalyzer to compute the GC [167]. We consider MVAR model of order 1, and GC likelihood ratio statistic of greater than 95 percentile as indicating edges [108]. The SPC was estimated by Graphical Lasso penalized Maximum Likelihood Estimation, whose optimal penalization was obtained by a 5-fold cross-validation [24].

5.1.6 Alterations of CFC edges in Alzheimer’s disease

We perform an exploratory analysis of statistical trends for edge-wise inter-group differences. Using the subject-specific CFC computed by TPC algorithm, for each detected CFC edge, we reported the p-value in the Welch’s t-test for greater average edge weight in one clinical group compared to another clinical group [268]. Specifically, we listed the CFC edges with ten lowest p-values for greater average weight in CN compared to AD group (and for greater average weight in AD compared to CN). For a CFC edge from region u to region v , we refer to u as the *source* brain region and v as the *destination* brain region. The source brain regions of the CFC edges with lowest p-values are found to be in agreement with literature for regions impacted by AD.

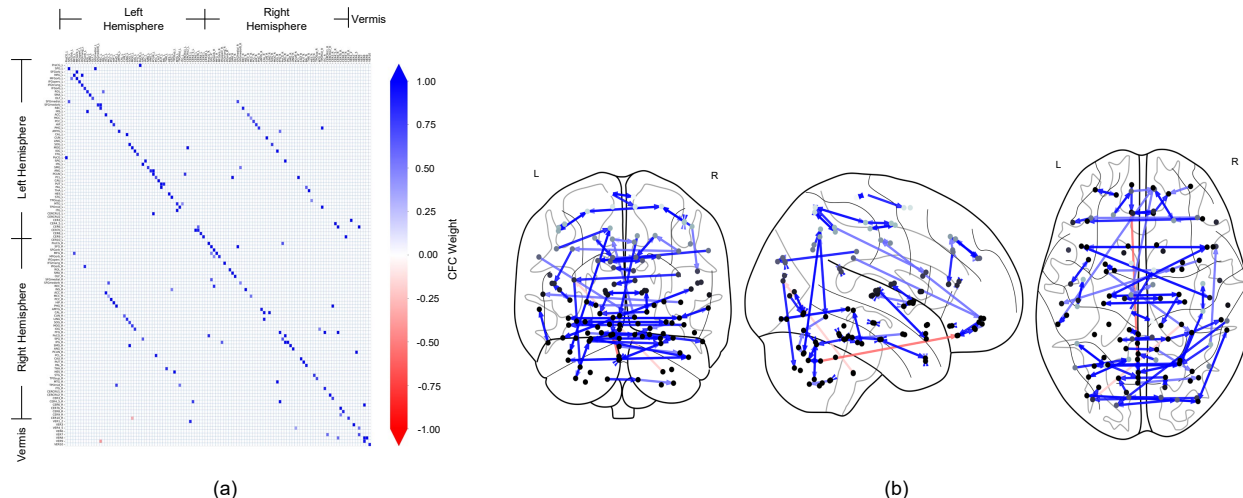


Figure 5.2 CFC for an example subject who is CN, estimated by TPC algorithm. (a) The estimated CFC is represented by its adjacency matrix, whose non-zero entry (i, j) represents the connection of region $i \rightarrow j$. (b) The CFC is visualized with directed graph edges on the Frontal, Axial and Lateral brain maps (left to right). The nodes correspond to brain region centers, ranging from superficial (light gray) to deeper (darker gray) regions, in the AAL brain atlas.

5.2 Results

5.2.1 Subject-specific Causal Functional Connectivity

Figure 5.2 shows the CFC estimated using the TPC algorithm for an example subject (ID: 129_S_4396) in the CN group. In Figure 5.2(a), the CFC is represented in the form of a matrix, whose entry (i, j) indicates the presence of connectivity from region index $i \rightarrow j$, and the value at entry (i, j) represents the weight of that causal connection. A positive value (blue) of the weight indicates excitatory influence, whereas a negative value (red) indicates inhibitory influence. The diagonal of the matrix representing self-connections for regions has been filtered out. In Figure 5.2(b), the CFC is represented by a directed graph overlaid on schematics of the brain. The schematics of the brain comprise 2-dimensional brain projections in the Frontal, Axial, and Lateral planes. The nodes of the CFC graph correspond to the centers of brain regions in the AAL atlas. The nodes are colored light to dark gray according to their depth in the brain, with light gray representing superficial and dark gray representing deeper brain regions. The CFC graph provides a highly informative map of causal interactions between brain regions.

It is noteworthy that the CFC computed by TPC is sparse since the edges are filtered by conditional dependence tests. We quantified the sparsity of a CFC graph by its edge density. Edge density of a directed

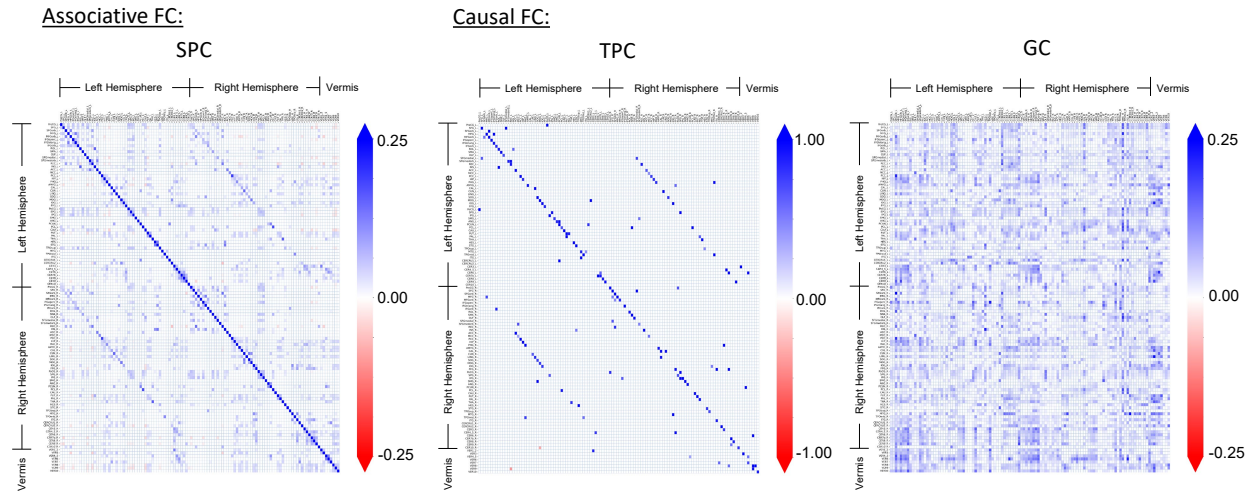


Figure 5.3 Comparison and demonstration of FC inferred by three methods: Associative FC using SPC, and Causal FC using TPC and GC. The estimated FC is represented by its adjacency matrix with edge weights, which is symmetric for Associative FC and asymmetric for Causal FC. In the adjacency matrices, a non-zero entry in (i, j) represents the connection of region $i \rightarrow j$.

graph is the proportion of the number of edges in the directed graph over the total number of edges in the corresponding fully connected graph. Therefore, the edge density of an empty graph is 0 and that of a fully connected graph is 1. For the CFC graphs computed by TPC, the edge density for subjects in the CN group is (mean \pm standard deviation) 0.0117 ± 0.0008 , MCI group is 0.0118 ± 0.0009 , and AD group is 0.0118 ± 0.0008 , indicating a sparse CFC outcome of TPC for subjects in each of the groups.

5.2.2 Comparison with functional connectivity using other approaches

Figure 5.3 shows the adjacency matrices for the FC obtained by different methods for an example subject (ID: 129_S_4396) in the CN group. The AFC constitutes a distinct pattern of associative connectivity among the regions. It is expected that the CFC will be a directed subgraph of the AFC and be consistent with the overall patterns present in the AFC [193; 194]. However, the patterns present in the CFC obtained by GC do not match with the AFC upon visual inspection. In comparison, the overall patterns present in the CFC obtained by TPC indeed match with the AFC obtained by SPC. On a detailed level, there are differences between TPC-CFC and AFC: TPC results in a directed graph thereby its adjacency matrix is asymmetric while AFC is an undirected graph with symmetric adjacency matrix. Furthermore, the CFC obtained by TPC includes self-loops represented by the diagonals of the adjacency matrix in contrast to GC, and results in a

(a) CN>AD			
Edge	P-value	Region Name	Reported by
HES_L → ROL_L	0.0008	Heschl's gyrus	Hänggi et al. [269]; Dhanjal et al. [270]
ITG_R → ITG_R	0.001	Inferior temporal gyrus	[271; 272]
SOG_L → SOG_R	0.001	Superior occipital gyrus	Beyer et al. [273]; Mao et al. [274]
SFG_L → SFG_R	0.001	Superior frontal gyrus	Brachova et al. [275]; Lue et al. [276]
MFGorb_R → IFGorb_R	0.002	Middle frontal gyrus	Neufang et al. [277]; Zhou et al. [278]
THA_R → THA_L	0.002	Thalamus	Braak and Braak [279]; de Jong et al. [280]
SMG_L → SMG_R	0.002	SupraMarginal gyrus	Grignon et al. [281]; Desikan et al. [282]
IOG_L → IOG_R	0.005	Inferior occipital gyrus	Wu et al. [283]; Johnen et al. [284]
SMG_R → SMG_L	0.005	SupraMarginal gyrus	Grignon et al. [281]; Desikan et al. [282]
PCC_L → ANG_L	0.006	Posterior cingulate gyrus	Villain et al. [285]; Caminiti et al. [286]; Mascali et al. [287]

(b) CN<AD			
Edge	P-value	Region Name	Reported by
PHG_R → PHG_R	0.0008	Parahippocampal gyrus	Van Hoesen et al. [288]; Thangavel et al. [289]
REC_R → REC_R	0.004	Gyrus rectus	Mölsä et al. [290]; Noehlin et al. [291]; Sheline et al. [292]
CER6_R → CER4_5_R	0.008	Cerebellum	Joachim et al. [293]; Jacobs et al. [294]
IFGtriang_R → MFG_R	0.008	Inferior frontal gyrus	Eliasova et al. [295]; Cajanus et al. [296]
CER7b_L → ITG_R	0.008	Cerebellum	Joachim et al. [293]; Jacobs et al. [294]
FFG_R → ITG_R	0.008	Fusiform gyrus	Ma et al. [297]; Whitwell [298]
CUN_L → SOG_L	0.014	Cuneus	He et al. [299]; Niskanen et al. [300]
VER1_2 → VER3	0.014	Vermis	Sjöbeck and Englund [301]; A Mavroudis et al. [302]
CAL_L → SOG_L	0.014	Calcarine fissure	[303; 304]
PCL_L → PCL_L	0.017	Paracentral lobule	Garcia Martin et al. [305]; Yang et al. [306]

Table 5.2 CFC edges with lowest edge-wise p-values for (a) greater weight in CN compared to AD group and (b) greater weight in AD compared to CN group. The corresponding source brain regions are in agreement with regions reported in literature (right column) as impacted by AD.

sparse matrix devoid of noise since the connections are filtered by conditional dependence tests.

5.2.3 Alterations of CFC edges in Alzheimer's disease

Figure 5.4 shows the edge-wise p-values for greater average edge weight in one clinical group compared to another, based on Welch's t-test. This provides insights into statistical trends for CFC edges that have an increase (or decrease) in strength in CN compared to MCI, CN compared to MCI, and MCI compared to AD subjects.

In Table 5.2, we report ten CFC edges that show the lowest p-values for greater average strength in CN subjects compared to AD, and their source brain regions. Similarly, we report another list of ten CFC edges corresponding to lowest average strength in subjects with AD compared to CN, and their source brain regions.

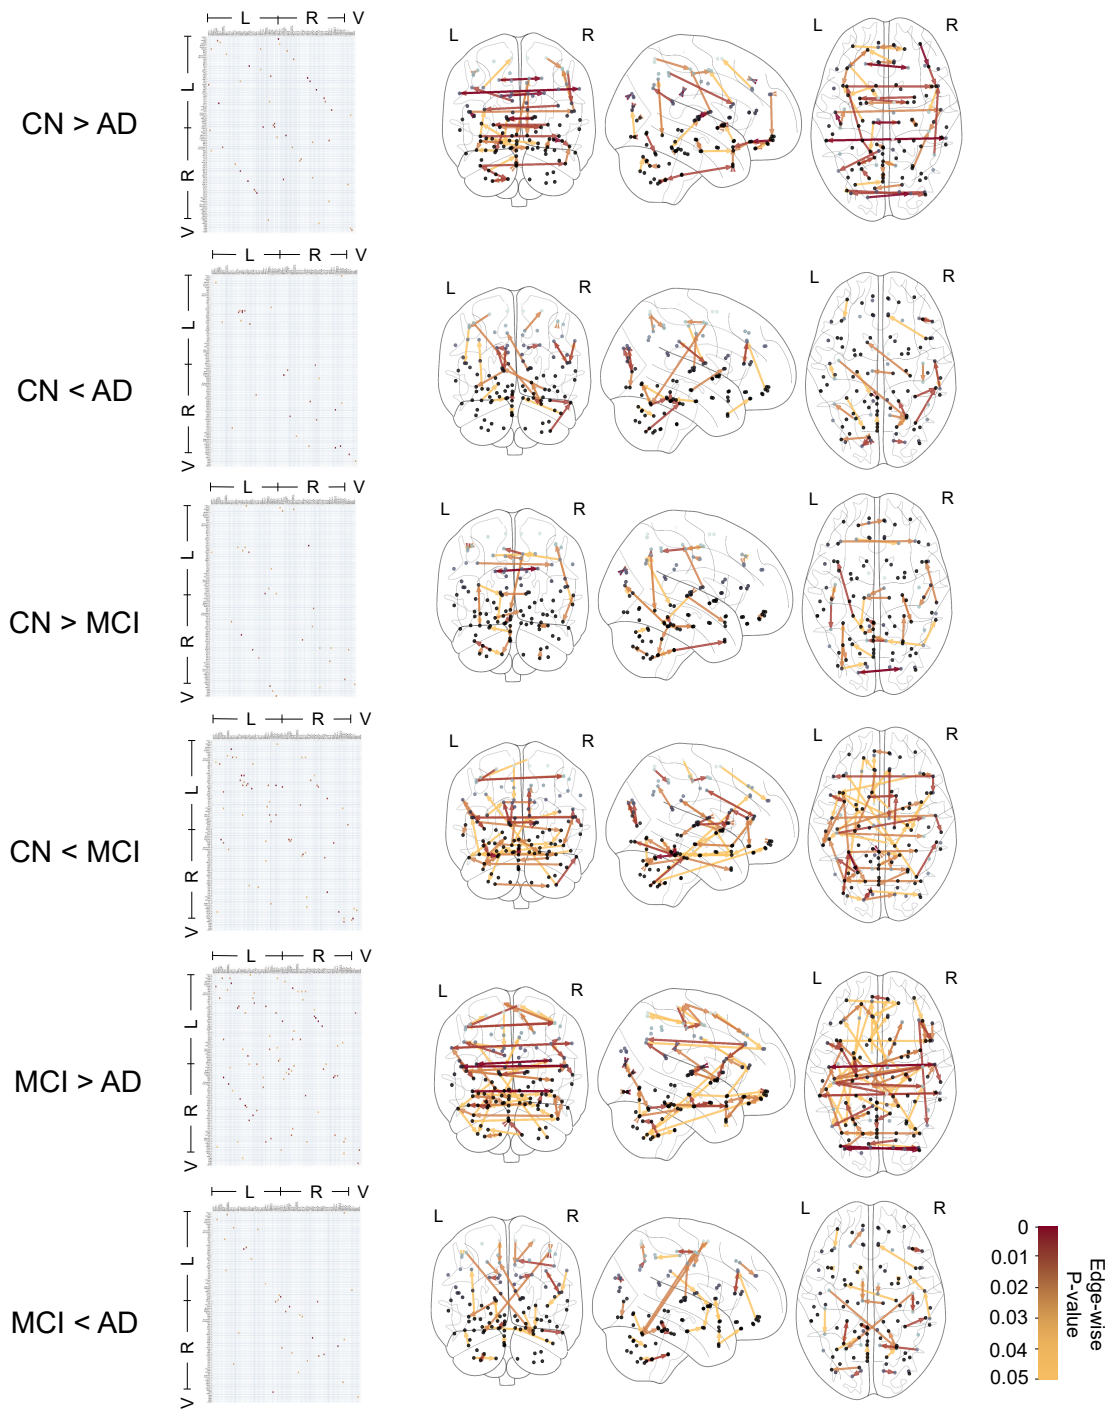


Figure 5.4 Causal functional connections with edge-weights differing between clinical groups with edge-wise p-values ranging in $0 - 0.05$ based on t-test. The edge-wise p-values are represented by a matrix whose entry in (i, j) corresponds to the edge $i \rightarrow j$ and also represented by graph edges on brain schematics. The brain regions are annotated by Left (L) and Right (R) hemispheres of the brain and Vermis (V).

The reported brain regions are in agreement with published medical literature cited in Table 5.2.

5.3 Discussion

In this study, we have obtained the CFC of the whole brain from its resting state fMRI time series. We used the recently developed TPC algorithm based on directed graphical modeling in time series, to compute the CFC. In the dataset, the subjects belonged to three clinical categories: CN, MCI and AD. We computed the subject-specific CFC using TPC and compared it with those obtained by other approaches, such as GC. We then used the CFC outcomes of TPC for further investigation into the alteration of CFC in AD. In this regard, we explored statistical trends for edges that have a difference in strength between clinical categories, based on their edge-wise p-values obtained by Welch's t-test. We reported the causal connections with lowest p-values for greater strength in CN compared to AD (and greater strength in AD compared to CN) and their corresponding brain regions. The brain regions identified in the above analyses were found to be in agreement with medical literature for regions impacted by AD.

In Figure 5.4 and Table 5.2, the presence of CFC edges with weight in AD greater than that in CN (in addition to edges with weight in AD less than that in CN) is consistent with published studies in the literature. While several studies have concluded decreased connectivity in MCI and AD compared to CN [252; 253; 254], others have highlighted that MCI and early stages of AD can involve an increase in FC between brain regions [307; 308]. This increase occurs when the communication between specific brain regions is impaired and has been interpreted as a compensatory mechanism where alternative paths within the brain's network are recruited [257; 258; 259]. In the short term, the augmentation of FC along alternative pathways exhibits efficiency and adaptability of the brain. However, it is imperative to acknowledge the susceptibility of these densely interconnected hubs to beta-amyloid deposition, which can elicit secondary damage through metabolic stress, ultimately culminating in system breakdown [257]. Consequently, the initial state of hyperconnectivity observed in neurodegenerative disorders may gradually transition into hypoconnectivity among the engaged pathways, thereby contributing to cognitive decline as the disease advances [258].

In Table 5.2(a), the Heschl's gyrus (Heschl's gyrus Left \rightarrow Rolandic operculum Left with edge-wise p-value 0.0008) is prominent for lower CFC weight in AD compared to CN subjects. The Heschl's gyrus is

not only important for language comprehension, but it also has a crucial role in speech production, phonologic retrieval, and semantic processing [309; 310], and has been reported in the literature to be impacted by AD [269; 270]. The Thalamus is also present among the list of regions in Table 5.2(a) (Thalamus Right → Thalamus Left with edge-wise p-value 0.002). The Thalamus functions as a relay station between different sub-cortical areas and the cerebral cortex and also plays a role in sleep, wakefulness, consciousness, and memory [311; 312; 313; 314], and is also known to be impacted by AD [279; 280]. Also present in the table is the Posterior cingulate gyrus (Posterior cingulate Left → Angular gyrus Left with edge-wise p-value of 0.006), which plays an essential role in memory integration and attentional processing, and is widely considered to be impacted by AD [285; 252; 253; 254]. The Hippocampus, which is involved in long-term memory formation and memory retrieval, is not in the list of regions, yet exhibits a trend of reduction in CFC weight in AD compared to CN (Hippocampus Right → Parahippocampal gyrus Right, edge-wise p-value 0.033) [315; 316]. Self-connections in Hippocampus have been reported to be often involved in compensatory mechanisms leading to increased strength in AD [317]. In Table 5.2(b), the self-connection Parahippocampal gyrus Right → Parahippocampal gyrus Right (edge-wise p-value 0.0008) is prominent for greater weight in AD compared to CN. It is known that the Parahippocampal gyrus is highly impacted by AD and is the focus of damage during disease onset, in a manner such that its connectivity to other regions of the brain decreases with AD, while its activity and intrinsic connectivity within the region increases with AD [288; 318; 317; 319; 320].

TPC identified 1475 edges in the CFCs across subjects with CN and AD. To obtain the subset of edges which have significant inter-group difference at Bonferonni family-wise error rate of 0.05 requires a total of 6352 unique subjects across three groups (2117 per group) to ensure a family-wise power of 0.95 in detecting mean differences of a quarter of the standard deviation (Cohen's $D = 0.25$), computed by `power_t_test` function in MESS package in R [321]. None of the databases that are available publicly have so many subjects. For example, ADNI has under 2000 subjects [322, <https://adni.loni.usc.edu/adni-3/>], and the Australian database has 2359 subjects [323, <https://aibl.org.au/about/>]. Therefore we took a subset of the ADNI dataset that is captured using 3T fMRI scanner while matching education and age levels for exploratory analysis.

Based on the whole-brain CFC outcome alone, this study obtained brain regions that have been reported

across more than 30 different studies of altered connectivity in AD, using different feature extraction methods and advanced imaging technologies (see Table 5.2). This demonstrates the promise of CFC computed by the TPC algorithm based on directed graphical models in a time series setting. Given the nature of AD, progressively more and more regions of the brain get impacted. Therefore, we make the case for the collection of larger datasets to enable the identification, at desirable levels of significance, of various subnetworks that alter with AD. This would promote the maturation and the use of the TPC-CFC (and other approaches) for prognostic and diagnostic purposes for AD.

It is noteworthy that machine-learning-based classifiers can help predict the clinical category of subjects and diagnose AD [324; 325; 326; 327; 328]. Recently, researchers have proposed robust multi-class classification methods in the presence of incorrect labeling of classes using the broad learning system [329; 330]. Such classifiers would be able to classify a subject as belonging to one of the clinical categories, given a subject's fMRI time-series data as input. However, such classifiers do not compute the CFC between brain regions. Computing the CFC can nicely complement a classifier by providing insights into specific causal functional connections and subnetworks that are altered by AD [331; 332]. Abnormal resting-state FC between brain regions is known to predate brain atrophy and the emergence of AD symptoms by upto two decades or more [239; 240; 241; 242]. Therefore, a subject's computed CFC can shed light on such abnormalities and promises to be a biomarker for early diagnosis and prognosis of the disease.

In this chapter, we have demonstrated the following: (a) Application of the TPC algorithm to compute whole-brain CFC for each subject (b) Comparison of CFCs computed using other approaches (c) Interpretation of CFC in the context of AD using domain (neuropathological) knowledge, and (d) Exploratory analysis for edge-wise differences and corresponding brain regions with altered connectivity in subjects with AD compared to CN. The findings are consistent with published medical literature. In summary, our results show the promise of computing the whole-brain CFC from fMRI data using the TPC algorithm to gain prognostic and diagnostic insights.

Chapter 6

Conclusion

In this thesis, we first established a statistical framework for causal functional connectomics particularly using directed probabilistic graphical models, and studied the applicability and limitations of existing approaches in a time series scenario while comparing their performance in simulated time series data. We then developed a novel methodology that improve upon limitations of existing approaches of directed graphical modeling in the time series scenario by better incorporating across-time as well as contemporaneous relationships prevalent in time series. We benchmarked the approaches in simulated and public benchmark datasets and also showed its consistency under standard assumptions on the time series involving mixing. We then applied the methods and obtained insights on Neuropixels and resting-state fMRI datasets. We also made available a Python package with example codes.

In Chapter 2, we established a statistical guide for neural connectomics. We discussed the possibility of using a graph to map neural anatomical connections, or a “anatomical connectome”, and we then showed how this differs from a graph that captures functional interactions, or a associative functional connectome (AFC), which captures interactions between different brain regions. We then defined the graphical tools and properties required to convert AFC into a directional graph, i.e., causal functional connectome (CFC), which represents flows of cause and effect in the interaction between neurons in recognition of the fact that the ultimate goal of functional connectomics is to infer causal interactions between neurons. Then, in the context of functional connectomics, we contrasted several popular methodologies with the ultimate goal of discovering “causation”, such as Granger Causality (GC), Dynamic Causal Model (DCM), and Directed

Probabilistic Graphical Model (DPGM). Specifically, we related the statistical developments of DPGM using the Directed Markov Property to the topic of CFC inference. We demonstrated the importance of this property in probabilistic graphical models that could make up neural CFC. The PC algorithm, a popular statistical method for inference of such graphs, is next described. Based on these notions and outcomes of the Directed Markov Property, we formulated a list of criteria for comparing CFC models. We performed a holistic comparison of GC, DCM, and PC algorithm (DPGM), comparing them with respect to the criteria that we have outlined. We demonstrated the applicability and the challenges for inference of CFC from measured neural activity for each of the approaches on simulated motifs.

While DPGM is a powerful causal framework, existing DPGM algorithms do not reflect the across-time causal dependencies within and between the neural time series and are suitable for variables with i.i.d. replicates. Yet in the neural time series setting, nodes of the connectivity graph are neurons that correspond to an entire time series of neural activity and comprise across-time dependence. Thus the remaining challenge is the adaptation of the DPGM based formulation of CFC to incorporate across-time dependencies in the neural time series. Such an adaptation will further increase the strength of using DPGM for CFC inference from neural dynamics.

In Chapter 3, we presented a novel method, the TPC algorithm, for finding the causal functional connectivity between neurons from neural time series. TPC extends directed graphical modeling to time series scenario by modeling across-time relationships used an unrolled CPDAG and then roll it back to obtain the CFC graph. It also incorporates time delays in interactions between neurons and the presence of feedback-loops. We showed that if the neural activity obeys an arbitrary dynamical process, i.e. with arbitrary functional relationships over time, then the Rolled CPDAG for the process agrees with the causal relationships implied by the dynamical process. We described that the Rolled CPDAG is predictive of counterfactual queries of ablation or modulation. We showed that the answers to such queries can be provided by reasoning with the edges of the Rolled CPDAG. We used the methodology in obtaining CFC from simulated datasets and demonstrated comparatively better performance than Granger Causality (GC) and common DPGM. Furthermore, we applied the methods to benchmarks of time series causal inference and neurobiological dataset from mice brain presented with various stimuli. The results provided insights into the CFC between neurons in the mouse brain in a variety of stimuli scenario. We also compared the topological patterns in the

estimated CFC between different stimuli using graph-theoretic measures. We concluded by summarizing the differences and benefits of TPC in a comparison with other approaches including the PC algorithm (DPGM) which satisfy Directed Markov Property (DMP) in a static data setting, and other approaches such as GC and DCM, outlining their strengths and weaknesses with respect to several criteria of causality in functional connectomics.

In Chapter 4, we established the mathematical guarantees of the TPC algorithm under standard conditions on the underlying time series that encompasses popular time series models such as vector autoregressive moving average and linear models. The findings are supported by simulated and benchmark real data analyses involving different variations of TPC that uses different conditional dependence tests and with/without subsampling.

In Chapter 5, we computed the CFC in Alzheimer's disease (AD) from time series rs-fMRI data. In this regard, we applied the TPC algorithm to compute whole-brain CFC for each subject in three clinical categories: cognitively normal, mild cognitive impairment, Alzheimer's disease. We then compared the CFC outcome of TPC with other related approaches. We interpreted the CFC in the context of AD using domain (neuropathological) knowledge. We then performed exploratory analysis for edge-wise differences and corresponding brain regions with altered connectivity in subjects with AD compared to CN. The findings are consistent with published medical literature. In summary, our results show the promise of inferring the whole-brain CFC from fMRI data using the TPC algorithm to gain prognostic and diagnostic insights.

In summary, in this thesis, we have established the following: 1) A comparative study of different approaches for inferring causal functional connectivity from neural activity, 2) Development of a novel method - the TPC algorithm, based on directed graphical models in time series, 3) Establishment of mathematical guarantees of the TPC algorithm in capturing causal relationships from time series data, 4) A comparative study of performance in simulated and benchmark real datasets, 5) A comparative study using different variations of TPC algorithm that uses different conditional dependence tests and with/without subsampling, 6) Application to analysis of the CFC for mice brain from Neuropixels dataset, 7) Computation of CFC for human brain in Alzheimer's disease from resting-state fMRI data and 8) Exploratory analysis for CFC edges and corresponding brain regions with altered connectivity in subjects with Alzheimer's compared to cognitively normal subjects.

6.1 Code Availability

The software package for the Time-aware PC algorithm and example code are available in <https://github.com/shlizee/TimeAwarePC> [2].

Rahul Biswas's Publications

- [RB1] Rahul Biswas and Somabha Mukherjee. Consistent causal inference from time series with pc algorithm and its time-aware extension. *Statistics and Computing*, 34(1):14, 2024.
- [RB2] Tianyu Liu, Somabha Mukherjee, and Rahul Biswas. Tensor recovery in high-dimensional ising models. *Journal of Multivariate Analysis*, page 105335, 2024.
- [RB3] Rahul Biswas and SuryaNarayana Sripada. Alterations in causal functional brain networks in alzheimer's disease: A resting-state fmri study. *BioRxiv*, 2024. doi: 10.1101/2024.05.12.593795.
- [RB4] Rahul Biswas and SuryaNarayana Sripada. Causal functional connectivity in alzheimer's disease computed from time series fmri data. *Frontiers in Computational Neuroscience*, 17:1251301, 2023.
- [RB5] Rahul Biswas, SuryaNarayana Sripada, and Somabha Mukherjee. Inferring causality from time series data based on structural causal model and its application to neural connectomics. *arXiv preprint arXiv:2312.09604*, 2023.
- [RB6] Rahul Biswas, Marie Thoma, and Xiangrong Kong. Functional data analysis to characterize disease patterns in frequent longitudinal data: application to bacterial vaginal microbiota patterns using weekly nugent scores and identification of pattern-specific risk factors. *BMC medical research methodology*, 23(1):251, 2023.
- [RB7] Rahul Biswas and Eli Shlizerman. Statistical perspective on functional and causal neural connectomics: The time-aware pc algorithm. *PLOS Computational Biology*, 18(11):e1010653, 2022.
- [RB8] Rahul Biswas and Eli Shlizerman. Statistical perspective on functional and causal neural connectomics: A comparative study. *Frontiers in Systems Neuroscience*, 16:817962, 2022.

- [RB9] Rahul Biswas and Eli Shlizerman. Neuro-pc: Causal functional connectivity from neural dynamics. *arXiv preprint arXiv:2011.03913*, 2020.
- [RB10] Soham Sarkar, Rahul Biswas, and Anil K Ghosh. On some graph-based two-sample tests for high dimension, low sample size data. *Machine Learning*, 109(2):279–306, 2020.

Bibliography

- [1] Rahul Biswas and Eli Shlizerman. Statistical perspective on functional and causal neural connectomics: A comparative study. *Frontiers in Systems Neuroscience*, 16, 2022. ISSN 1662-5137. doi: 10.3389/fnsys.2022.817962.
- [2] Rahul Biswas and Eli Shlizerman. Statistical perspective on functional and causal neural connectomics: The time-aware pc algorithm. *PLOS Computational Biology*, 18(11):e1010653, 2022.
- [3] Rahul Biswas and Somabha Mukherjee. Consistent causal inference from time series with pc algorithm and its time-aware extension. *Statistics and Computing*, 34(1):14, 2024.
- [4] Rahul Biswas and SuryaNarayana Sripada. Causal functional connectivity in alzheimer’s disease computed from time series fmri data. *Frontiers in Computational Neuroscience*, 2023.
- [5] Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: a structural description of the human brain. *PLOS Computational Biology*, 1(4):e42, 2005.
- [6] Yonggang Shi and Arthur W Toga. Connectome imaging for mapping human brain pathways. *Molecular psychiatry*, 22(9):1230–1240, 2017.
- [7] Tabinda Sarwar, Caio Seguin, Kotagiri Ramamohanarao, and Andrew Zalesky. Towards deep learning for connectome mapping: A block decomposition framework. *NeuroImage*, 212:116654, 2020.
- [8] C Shan Xu, Michal Januszewski, Zhiyuan Lu, Shin-ya Takemura, Kenneth Hayworth, Gary Huang, Kazunori Shinomiya, Jeremy Maitin-Shepard, David Ackerman, and Stuart Berg. A connectome of the adult drosophila central brain. *BioRxiv*, 2020.

- [9] Wei-Chung Allen Lee and R Clay Reid. Specificity and randomness: structure–function relationships in neural circuits. *Current opinion in neurobiology*, 21(5):801–807, 2011.
- [10] Nancy J Kopell, Howard J Gritton, Miles A Whittington, and Mark A Kramer. Beyond the connectome: the dynamome. *Neuron*, 83(6):1319–1328, 2014.
- [11] Jimin Kim, William Leahy, and Eli Shlizerman. Neural interactome: Interactive simulation of a neuronal system. *Frontiers in Computational Neuroscience*, 13:8, 2019.
- [12] Jimin Kim, Julia A Santos, Mark J Alkema, and Eli Shlizerman. Whole integration of neural connectomics, dynamics and bio-mechanics for identification of behavioral sensorimotor pathways in *Caenorhabditis elegans*. *BioRxiv*, page 724328, 2019.
- [13] Cornelia I Bargmann and Eve Marder. From the connectome to brain function. *Nature methods*, 10(6):483, 2013.
- [14] R Clay Reid. From functional architecture to functional connectomics. *Neuron*, 75(2):209–217, 2012.
- [15] Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258, 2017.
- [16] Emily S Finn, Xilin Shen, Dustin Scheinost, Monica D Rosenberg, Jessica Huang, Marvin M Chun, Xenophon Papademetris, and R Todd Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature neuroscience*, 18(11):1664–1671, 2015.
- [17] Eli Shlizerman, Konrad Schroder, and J Nathan Kutz. Neural activity measures and their dynamics. *SIAM Journal on Applied Mathematics*, 72(4):1260–1291, 2012.
- [18] 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.
- [19] Maria Giulia Preti, Thomas AW Bolton, and Dimitri Van De Ville. The dynamic functional connectome: State-of-the-art and perspectives. *Neuroimage*, 160:41–54, 2017.
- [20] Yuting Xu and Martin A Lindquist. Dynamic connectivity detection: an algorithm for determining functional connectivity change points in fmri data. *Frontiers in neuroscience*, 9:285, 2015.

- [21] Chong-Yaw Wee, Pew-Thian Yap, and Dinggang Shen. Diagnosis of autism spectrum disorders using temporally distinct resting-state functional connectivity networks. *CNS neuroscience & therapeutics*, 22(3):212–219, 2016.
- [22] Gael Varoquaux, Alexandre Gramfort, Jean-Baptiste Poline, and Bertrand Thirion. Brain covariance selection: better individual functional connectivity models using population prior. In *Advances in neural information processing systems*, pages 2334–2342, 2010.
- [23] Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for fmri. *Neuroimage*, 54(2):875–891, 2011.
- [24] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- [25] 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.
- [26] 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.
- [27] Steffen L Lauritzen. Causal inference from graphical models. *Complex stochastic systems*, pages 63–107, 2001.
- [28] Marloes Maathuis, Mathias Drton, Steffen Lauritzen, and Martin Wainwright. *Handbook of graphical models*. CRC Press, 2018.
- [29] Ana María Estrada Gómez, Kamran Paynabar, and Massimo Pacella. Functional directed graphical models and applications in root-cause analysis and diagnosis. *Journal of Quality Technology*, pages 1–17, 2020.
- [30] Daniel Felix Ahelegbey. The econometrics of bayesian graphical models: a review with financial application. *Journal of Network Theory in Finance*, 2(2):1–33, 2016.

- [31] Imme Ebert-Uphoff and Yi Deng. Causal discovery for climate research using graphical models. *Journal of Climate*, 25(17):5648–5665, 2012.
- [32] Markus Kalisch, Bernd AG Fellinghauer, Eva Grill, Marloes H Maathuis, Ulrich Mansmann, Peter Bühlmann, and Gerold Stucki. Understanding human functioning using graphical models. *BMC Medical Research Methodology*, 10(1):1–10, 2010.
- [33] Ke Deng, Delin Liu, Shan Gao, and Zhi Geng. Structural learning of graphical models and its applications to traditional chinese medicine. In *International Conference on Fuzzy Systems and Knowledge Discovery*, pages 362–367. Springer, 2005.
- [34] Michael S Haigh and David A Bessler. Causality and price discovery: An application of directed acyclic graphs. *The Journal of Business*, 77(4):1099–1121, 2004.
- [35] Huange Wang, Fred A van Eeuwijk, and Johannes Jansen. The potential of probabilistic graphical models in linkage map construction. *Theoretical and Applied Genetics*, 130(2):433–444, 2017.
- [36] Christine Sinoquet. *Probabilistic graphical models for genetics, genomics, and postgenomics*. OUP Oxford, 2014.
- [37] Raphaël Mourad, Christine Sinoquet, and Philippe Leray. Probabilistic graphical models for genetic association studies. *Briefings in bioinformatics*, 13(1):20–33, 2012.
- [38] Junbai Wang, Leo Wang-Kit Cheung, and Jan Delabie. New probabilistic graphical models for genetic regulatory networks studies. *Journal of biomedical informatics*, 38(6):443–455, 2005.
- [39] Hexuan Liu, Jimin Kim, and Eli Shlizerman. Functional connectomics from neural dynamics: probabilistic graphical models for neuronal network of caenorhabditis elegans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1758):20170377, 2018.
- [40] Nir Friedman. Inferring cellular networks using probabilistic graphical models. *Science*, 303(5659):799–805, 2004.
- [41] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

- [42] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [43] KR1442 Chowdhary and KR Chowdhary. Natural language processing. *Fundamentals of artificial intelligence*, pages 603–649, 2020.
- [44] Yinglong Li. Research and application of deep learning in image recognition. In *2022 IEEE 2nd International Conference on Power, Electronics and Computer Applications (ICPECA)*, pages 994–999. IEEE, 2022.
- [45] Mattia Prosperi, Yi Guo, Matt Sperrin, James S Koopman, Jae S Min, Xing He, Shannan Rich, Mo Wang, Iain E Buchan, and Jiang Bian. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence*, 2(7):369–375, 2020.
- [46] Kevin Xia, Kai-Zhan Lee, Yoshua Bengio, and Elias Bareinboim. The causal-neural connection: Expressiveness, learnability, and inference. *Advances in Neural Information Processing Systems*, 34: 10823–10836, 2021.
- [47] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- [48] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [49] Judea Pearl. The foundations of causal inference. *Sociological Methodology*, 40(1):75–149, 2010.
- [50] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- [51] Bernhard Schölkopf. Causality for machine learning. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 765–804. 2022.
- [52] Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. *Causation, prediction, and search*. MIT press, 2000.
- [53] Thomas S Richardson and James M Robins. Single world intervention graphs (swigs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper*, 128(30):2013, 2013.

- [54] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *Advances in neural information processing systems*, 30, 2017.
- [55] Peng Cui and Susan Athey. Stable learning establishes some common ground between causal inference and machine learning. *Nature Machine Intelligence*, 4(2):110–115, 2022.
- [56] James M Robins, Richard Scheines, Peter Spirtes, and Larry Wasserman. Uniform consistency in causal inference. *Biometrika*, 90(3):491–515, 2003.
- [57] Markus Kalisch and Peter Bühlman. Estimating high-dimensional directed acyclic graphs with the pc-algorithm. *Journal of Machine Learning Research*, 8(3), 2007.
- [58] Ila R Fiete, Walter Senn, Claude ZH Wang, and Richard HR Hahnloser. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron*, 65(4):563–576, 2010.
- [59] Mohammad R Arbabshirani, Adrian Preda, Jatin G Vaidya, Steven G Potkin, Godfrey Pearlson, James Voyvodic, Daniel Mathalon, Theo van Erp, Andrew Michael, Kent A Kiehl, et al. Autoconnectivity: A new perspective on human brain function. *Journal of neuroscience methods*, 323:68–76, 2019.
- [60] Roman Borisyuk and Frank Hoppensteadt. Oscillatory models of the hippocampus: a study of spatio-temporal patterns of neural activity. *Biological cybernetics*, 81(4):359–371, 1999.
- [61] Michael J Jutras and Elizabeth A Buffalo. Synchronous neural activity and memory formation. *Current opinion in neurobiology*, 20(2):150–155, 2010.
- [62] Thomas S Richardson, Peter Spirtes, et al. *Automated discovery of linear feedback models*. Carnegie Mellon [Department of Philosophy], 1996.
- [63] Pedro A Valdes-Sosa, Alard Roebroeck, Jean Daunizeau, and Karl Friston. Effective connectivity: influence, causality and biophysical modeling. *Neuroimage*, 58(2):339–361, 2011.
- [64] Joseph D Ramsey, Stephen José Hanson, Catherine Hanson, Yaroslav O Halchenko, Russell A Poldrack, and Clark Glymour. Six problems for causal inference from fmri. *neuroimage*, 49(2):1545–1558, 2010.

- [65] Rahul Biswas and Eli Shlizerman. Neuro-PC: Causal functional connectivity from neural dynamics. *arXiv preprint arXiv:2011.03913*, 2020.
- [66] John G White, Eileen Southgate, J Nichol Thomson, and Sydney Brenner. The structure of the nervous system of the nematode *caenorhabditis elegans*. *Philos Trans R Soc Lond B Biol Sci*, 314(1165):1–340, 1986.
- [67] Thomas E Conturo, Nicolas F Lori, Thomas S Cull, Erbil Akbudak, Abraham Z Snyder, Joshua S Shimony, Robert C McKinstry, Harold Burton, and Marcus E Raichle. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences*, 96(18):10422–10427, 1999.
- [68] Denis Le Bihan. Looking into the functional architecture of the brain with diffusion mri. *Nature reviews neuroscience*, 4(6):469–480, 2003.
- [69] Marco Catani, Derek K Jones, Rosario Donato, and Dominic H Ffytche. Occipito-temporal connections in the human brain. *Brain*, 126(9):2093–2107, 2003.
- [70] Timothy Ragan, Lolahon R Kadiri, Kannan Umadevi Venkataraju, Karsten Bahlmann, Jason Sutin, Julian Taranda, Ignacio Arganda-Carreras, Yongsoo Kim, H Sebastian Seung, and Pavel Osten. Serial two-photon tomography for automated ex vivo mouse brain imaging. *Nature methods*, 9(3):255–258, 2012.
- [71] Flaviano Morone and Hernán A Makse. Symmetry group factorization reveals the structure-function relation in the neural connectome of *caenorhabditis elegans*. *Nature communications*, 10(1):1–13, 2019.
- [72] Saul Kato, Harris S Kaplan, Tina Schrödel, Susanne Skora, Theodore H Lindsay, Eviatar Yemini, Shawn Lockery, and Manuel Zimmer. Global brain dynamics embed the motor command sequence of *caenorhabditis elegans*. *Cell*, 163(3):656–669, 2015.
- [73] Peter Zátka-Haas, Nicholas A Steinmetz, Matteo Carandini, and Kenneth D Harris. A perceptual decision requires sensory but not action coding in mouse cortex. *BioRxiv*, page 501627, 2020.

- [74] Vincent Villette, Mariya Chavarha, Ivan K Dimov, Jonathan Bradley, Lagnajeet Pradhan, Benjamin Mathieu, Stephen W Evans, Simon Chamberland, Dongqing Shi, Renzhi Yang, et al. Ultrafast two-photon imaging of a high-gain voltage indicator in awake behaving mice. *Cell*, 179(7):1590–1608, 2019.
- [75] Nicholas A Steinmetz, Peter Zátka-Haas, Matteo Carandini, and Kenneth D Harris. Distributed coding of choice, action and engagement across the mouse brain. *Nature*, 576(7786):266–273, 2019.
- [76] Nicholas A Steinmetz, Christof Koch, Kenneth D Harris, and Matteo Carandini. Challenges and opportunities for large-scale electrophysiology with neuropixels probes. *Current opinion in neurobiology*, 50:92–100, 2018.
- [77] Andrea Stocco, Zoe Steine-Hanson, Natalie Koh, John E Laird, Christian J Lebiere, and Paul Rosenbloom. Analysis of the human connectome data supports the notion of a “common model of cognition” for human and human-like intelligence. *BioRxiv*, page 703777, 2019.
- [78] David C Van Essen, Kamil Ugurbil, Edward Auerbach, Deanna Barch, Timothy EJ Behrens, Richard Bucholz, Acer Chang, Liyong Chen, Maurizio Corbetta, Sandra W Curtiss, et al. The human connectome project: a data acquisition perspective. *Neuroimage*, 62(4):2222–2231, 2012.
- [79] Sean L Hill, Yun Wang, Imad Riachi, Felix Schürmann, and Henry Markram. Statistical connectivity provides a sufficient foundation for specific functional connectivity in neocortical neural microcircuits. *Proceedings of the National Academy of Sciences*, 109(42):E2885–E2894, 2012.
- [80] Luca Passamonti, A Terracciano, R Riccelli, G Donzuso, Antonio Cerasa, MG Vaccaro, F Novellino, Francesco Fera, and Aldo Quattrone. Increased functional connectivity within mesocortical networks in open people. *Neuroimage*, 104:301–309, 2015.
- [81] Jeanette A Mumford and Joseph D Ramsey. Bayesian networks for fmri: a primer. *Neuroimage*, 86:573–582, 2014.
- [82] Xiang Li, Chulwoo Lim, Kaiming Li, Lei Guo, and Tianming Liu. Detecting brain state changes via fiber-centered functional connectivity analysis. *Neuroinformatics*, 11(2):193–210, 2013.

- [83] F DuBois Bowman, Lijun Zhang, Gordana Derado, and Shuo Chen. Determining functional connectivity using fmri data with diffusion-based anatomical weighting. *NeuroImage*, 62(3):1769–1779, 2012.
- [84] Travis A Jarrell, Yi Wang, Adam E Bloniarz, Christopher A Brittin, Meng Xu, J Nichol Thomson, Donna G Albertson, David H Hall, and Scott W Emmons. The connectome of a decision-making neural network. *science*, 337(6093):437–444, 2012.
- [85] Lav R Varshney, Beth L Chen, Eric Paniagua, David H Hall, and Dmitri B Chklovskii. Structural properties of the caenorhabditis elegans neuronal network. *PLoS computational biology*, 7(2):e1001066, 2011.
- [86] Martin J Wainwright and Michael Irwin Jordan. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- [87] Mathias Drton and Marloes H Maathuis. Structure learning in graphical modeling. *Annual Review of Statistics and Its Application*, 4:365–393, 2017.
- [88] Sacha Epskamp, Lourens J Waldorp, René Möttöus, and Denny Borsboom. The gaussian graphical model in cross-sectional and time-series data. *Multivariate Behavioral Research*, 53(4):453–480, 2018.
- [89] Martin Dyrba, Reza Mohammadi, Michel J Grothe, Thomas Kirste, and Stefan J Teipel. Gaussian graphical models reveal inter-modal and inter-regional conditional dependencies of brain alterations in alzheimer’s disease. *Frontiers in aging neuroscience*, 12:99, 2020.
- [90] T. P. Speed and H. T. Kiiveri. Gaussian Markov Distributions over Finite Graphs. *The Annals of Statistics*, 14(1):138 – 150, 1986. doi: 10.1214/aos/1176349846.
- [91] Ali Shojaie and George Michailidis. Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs. *Biometrika*, 97(3):519–538, 2010.
- [92] Onureena Banerjee, Laurent El Ghaoui, and Alexandre d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine learning research*, 9(Mar):485–516, 2008.

- [93] Adam J Rothman, Peter J Bickel, Elizaveta Levina, Ji Zhu, et al. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.
- [94] Nicolai Meinshausen, Peter Bühlmann, et al. High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, 34(3):1436–1462, 2006.
- [95] Karl J Friston. Functional and effective connectivity: a review. *Brain connectivity*, 1(1):13–36, 2011.
- [96] Barry Horwitz. The elusive concept of brain connectivity. *Neuroimage*, 19(2):466–470, 2003.
- [97] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 106(3):1125, 2011.
- [98] Jonathan D Power, Alexander L Cohen, Steven M Nelson, Gagan S Wig, Kelly Anne Barnes, Jessica A Church, Alecia C Vogel, Timothy O Laumann, Fran M Miezin, Bradley L Schlaggar, et al. Functional network organization of the human brain. *Neuron*, 72(4):665–678, 2011.
- [99] Jonathan D Power and Steven E Petersen. Control-related systems in the human brain. *Current opinion in neurobiology*, 23(2):223–228, 2013.
- [100] Stephen M Smith, Peter T Fox, Karla L Miller, David C Glahn, P Mickle Fox, Clare E Mackay, Nicola Filippini, Kate E Watkins, Roberto Toro, Angela R Laird, et al. Correspondence of the brain’s functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106(31):13040–13045, 2009.
- [101] James A Carr. I’ll take the low road: the evolutionary underpinnings of visually triggered fear. *Frontiers in neuroscience*, 9:414, 2015.
- [102] Luiz Pessoa and Ralph Adolphs. Emotion processing and the amygdala: from a ‘low road’ to ‘many roads’ of evaluating biological significance. *Nature reviews neuroscience*, 11(11):773–782, 2010.
- [103] Caterina Bertini, Roberto Cecere, and Elisabetta Làdavas. I am blind, but i “see” fear. *Cortex*, 49(4):985–993, 2013.

- [104] John S Morris, Arne Öhman, and Raymond J Dolan. A subcortical pathway to the right amygdala mediating “unseen” fear. *Proceedings of the National Academy of Sciences*, 96(4):1680–1685, 1999.
- [105] Andreas Horn and Michael D Fox. Opportunities of connectomic neuromodulation. *Neuroimage*, 221: 117180, 2020.
- [106] George Casella and Roger L Berger. *Statistical inference*. Cengage Learning, 2021.
- [107] John H Byrne, Ruth Heidelberger, and M Neal Waxham. *From molecules to networks: an introduction to cellular and molecular neuroscience*. Academic Press, 2014.
- [108] Christoph Schmidt, Britta Pester, Nicole Schmid-Hertel, Herbert Witte, Axel Wismüller, and Lutz Leistritz. A multivariate granger causality concept towards full brain functional connectivity. *PloS one*, 11(4):e0153105, 2016.
- [109] Andrew T Reid, Drew B Headley, Ravi D Mill, Ruben Sanchez-Romero, Lucina Q Uddin, Daniele Marinazzo, Daniel J Lurie, Pedro A Valdés-Sosa, Stephen José Hanson, Bharat B Biswal, et al. Advancing functional connectivity research from association to causation. *Nature neuroscience*, 1(10), 2019.
- [110] Judea Pearl et al. Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*, 2000.
- [111] Judea Pearl et al. Causal inference in statistics: An overview. *Statistics surveys*, 3:96–146, 2009.
- [112] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.
- [113] Sumanta Basu, Ali Shojaie, and George Michailidis. Network granger causality with inherent grouping structure. *The Journal of Machine Learning Research*, 16(1):417–453, 2015.
- [114] Karl Friston, Rosalyn Moran, and Anil K Seth. Analysing connectivity with granger causality and dynamic causal modelling. *Current opinion in neurobiology*, 23(2):172–178, 2013.
- [115] Helmut Lütkepohl. *New introduction to multiple time series analysis*. Springer Science & Business Media, 2005.

- [116] Alex Tank, Ian Covert, Nicholas Foti, Ali Shojaie, and Emily Fox. Neural granger causality for nonlinear time series. *arXiv preprint arXiv:1802.05842*, 2018.
- [117] Mukeshwar Dhamala, Govindan Rangarajan, and Mingzhou Ding. Analyzing information flow in brain networks with nonparametric granger causality. *Neuroimage*, 41(2):354–362, 2008.
- [118] Luca Pollonini, Udit Patidar, Ning Situ, Roozbeh Rezaie, Andrew C Papanicolaou, and George Zouridakis. Functional connectivity networks in the autistic and healthy brain assessed using granger causality. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 1730–1733. IEEE, 2010.
- [119] Xinling Guo, Qiaosheng Zhang, Amrita Singh, Jing Wang, and Zhe Sage Chen. Granger causality analysis of rat cortical functional connectivity in pain. *Journal of Neural Engineering*, 17(1):016050, 2020.
- [120] Daniele Marinazzo, Mario Pellicoro, and Sebastiano Stramaglia. Kernel method for nonlinear granger causality. *Physical review letters*, 100(14), 2008.
- [121] Alex Tank, Emily B Fox, and Ali Shojaie. Granger causality networks for categorical time series. *arXiv preprint arXiv:1706.02781*, 2017.
- [122] John F Geweke. Measures of conditional linear dependence and feedback between time series. *Journal of the American Statistical Association*, 79(388):907–915, 1984.
- [123] Corrado Bernasconi and Peter KoÈnig. On the directionality of cortical interactions studied by structural analysis of electrophysiological recordings. *Biological cybernetics*, 81(3):199–210, 1999.
- [124] Mingzhou Ding, Steven L Bressler, Weiming Yang, and Hualou Liang. Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. *Biological cybernetics*, 83(1):35–45, 2000.
- [125] Andrea Brovelli, Mingzhou Ding, Anders Ledberg, Yonghong Chen, Richard Nakamura, and Steven L Bressler. Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by granger causality. *Proceedings of the National Academy of Sciences*, 101(26):9849–9854, 2004.

- [126] Adam B Barrett, Michael Murphy, Marie-Aur lie Bruno, Quentin Noirhomme, M lanie Boly, Steven Laureys, and Anil K Seth. Granger causality analysis of steady-state electroencephalographic signals during propofol-induced anaesthesia. *PloS one*, 7(1):e29072, 2012.
- [127] Olivier David, Isabelle Guillemain, Sandrine Sallet, Sebastien Reyt, Colin Deransart, Christoph Segebarth, and Antoine Depaulis. Identifying neural drivers with functional mri: an electrophysiological validation. *PLoS Biol*, 6(12):e315, 2008.
- [128] Alard Roebroeck, Elia Formisano, and Rainer Goebel. Mapping directed influence over the brain using granger causality and fmri. *Neuroimage*, 25(1):230–242, 2005.
- [129] Xiaotong Wen, Li Yao, Yijun Liu, and Mingzhou Ding. Causal interactions in attention networks predict behavioral performance. *Journal of Neuroscience*, 32(4):1284–1292, 2012.
- [130] Steven L Bressler, Wei Tang, Chad M Sylvester, Gordon L Shulman, and Maurizio Corbetta. Top-down control of human visual cortex by frontal and parietal cortex in anticipatory visual spatial attention. *Journal of Neuroscience*, 28(40):10056–10061, 2008.
- [131] 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.
- [132] Rainer Dahlhaus and Michael Eichler. Causality and graphical models in time series analysis. *Oxford Statistical Science Series*, pages 115–137, 2003.
- [133] Ruocheng Guo, Lu Cheng, Jundong Li, P Richard Hahn, and Huan Liu. A survey of learning causality with data: Problems and methods. *arXiv preprint arXiv:1809.09337*, 2018.
- [134] Michael Eichler. Causal inference with multiple time series: principles and problems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1997):20110613, 2013.
- [135] Karl Friston. Causal modelling and brain connectivity in functional magnetic resonance imaging. *PLoS biology*, 7(2), 2009.

- [136] Greta Grassmann. New considerations on the validity of the wiener-granger causality test. *Heliyon*, 6(10):e05208, 2020.
- [137] Jianping Qiao, Zhishun Wang, Guihu Zhao, Yuankai Huo, Carl L Herder, Chamonix O Sikora, and Bradley S Peterson. Functional neural circuits that underlie developmental stuttering. *PLoS One*, 12(7):e0179255, 2017.
- [138] Lionel Barnett, Adam B Barrett, and Anil K Seth. Granger causality and transfer entropy are equivalent for gaussian variables. *Physical review letters*, 103(23):238701, 2009.
- [139] Mauro Ursino, Giulia Ricci, and Elisa Magosso. Transfer entropy as a measure of brain connectivity: a critical analysis with the help of neural mass models. *Frontiers in computational neuroscience*, 14:45, 2020.
- [140] Michael Wibral, Nicolae Pampu, Viola Priesemann, Felix Siebenhühner, Hannes Seiwert, Michael Lindner, Joseph T Lizier, and Raul Vicente. Measuring information-transfer delays. *PloS one*, 8(2):e55809, 2013.
- [141] Karl J Friston, Lee Harrison, and Will Penny. Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302, 2003.
- [142] Klaas Enno Stephan, Nikolaus Weiskopf, Peter M Drysdale, Peter A Robinson, and Karl J Friston. Comparing hemodynamic models with dcm. *Neuroimage*, 38(3):387–401, 2007.
- [143] Karl J Friston, Andrea Mechelli, Robert Turner, and Cathy J Price. Nonlinear responses in fmri: the balloon model, volterra kernels, and other hemodynamics. *NeuroImage*, 12(4):466–477, 2000.
- [144] Klaas E Stephan, Lee M Harrison, Stefan J Kiebel, Olivier David, Will D Penny, and Karl J Friston. Dynamic causal models of neural system dynamics: current state and future extensions. *Journal of biosciences*, 32(1):129–144, 2007.
- [145] William D Penny. Comparing dynamic causal models using aic, bic and free energy. *Neuroimage*, 59(1):319–330, 2012.

- [146] André C Marreiros, Stefan J Kiebel, and Karl J Friston. A dynamic causal model study of neuronal population dynamics. *Neuroimage*, 51(1):91–101, 2010.
- [147] Klaas Enno Stephan, Will D Penny, Rosalyn J Moran, Hanneke EM den Ouden, Jean Daunizeau, and Karl J Friston. Ten simple rules for dynamic causal modeling. *Neuroimage*, 49(4):3099–3109, 2010.
- [148] Baojuan Li, Jean Daunizeau, Klaas E Stephan, Will Penny, Dewen Hu, and Karl Friston. Generalised filtering and stochastic dcm for fmri. *neuroimage*, 58(2):442–457, 2011.
- [149] Klaas Enno Stephan, Lars Kasper, Lee M Harrison, Jean Daunizeau, Hanneke EM den Ouden, Michael Breakspear, and Karl J Friston. Nonlinear dynamic causal models for fmri. *Neuroimage*, 42(2):649–662, 2008.
- [150] Stefan J Kiebel, Marta I Garrido, Rosalyn J Moran, and Karl J Friston. Dynamic causal modelling for eeg and meg. *Cognitive neurodynamics*, 2(2):121–136, 2008.
- [151] Rosalyn J Moran, Dimitris A Pinotsis, and Karl J Friston. Neural masses and fields in dynamic causal modeling. *Frontiers in computational neuroscience*, 7:57, 2013.
- [152] Richard B Stein, E Roderich Gossen, and Kelvin E Jones. Neuronal variability: noise or part of the signal? *Nature Reviews Neuroscience*, 6(5):389–397, 2005.
- [153] Amit Manwani and Christof Koch. Signal detection in noisy weakly-active dendrites. *Advances in neural information processing systems*, pages 132–138, 1999.
- [154] Lasitha S Vidyaratne and Khan M Iftekharuddin. Real-time epileptic seizure detection using eeg. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(11):2146–2156, 2017.
- [155] Rahul Biswas, Koulik Khamaru, and Kaushik K Majumdar. A peak synchronization measure for multiple signals. *IEEE Transactions on Signal Processing*, 62(17):4390–4398, 2014.
- [156] Sven Kroener, L Judson Chandler, Paul EM Phillips, and Jeremy K Seamans. Dopamine modulates persistent synaptic activity and enhances the signal-to-noise ratio in the prefrontal cortex. *PloS one*, 4(8):e6507, 2009.

- [157] Zafer Soybaş, Sefa Şimşek, FM Betül Erol, U Çiya Erdoğan, Esra N Şimşek, Büşra Şahin, Merve Marçalı, Bahattin Aydoğdu, Çağlar Elbüken, and Rohat Melik. Real-time in vivo control of neural membrane potential by electro-ionic modulation. *Iscience*, 17:347–358, 2019.
- [158] John D Fernstrom. Effects of the diet on brain neurotransmitters. *Metabolism*, 26(2):207–223, 1977.
- [159] Francisco Mora, Gregorio Segovia, Alberto Del Arco, Marta de Blas, and Pedro Garrido. Stress, neurotransmitters, corticosterone and body–brain integration. *Brain research*, 1476:71–85, 2012.
- [160] Pavitra Krishnaswamy, Gabriel Obregon-Henao, Jyrki Ahveninen, Sheraz Khan, Behtash Babadi, Juan Eugenio Iglesias, Matti S Hämäläinen, and Patrick L Purdon. Sparsity enables estimation of both subcortical and cortical activity from meg and eeg. *Proceedings of the National Academy of Sciences*, 114(48):E10465–E10474, 2017.
- [161] T Verma and J Pearl. Causal networks: Semantics and expressiveness, in proceedings. In *4th Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN, Mountain View, CA*, page 352, 1988.
- [162] Kenneth A Bollen. Structural equations with latent variables wiley. *New York*, 1989.
- [163] Peter Spirtes, Christopher Meek, and Thomas Richardson. An algorithm for causal inference in the presence of latent variables and selection bias. *Computation, causation, and discovery*, 21:1–252, 1999.
- [164] David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.
- [165] Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *The Journal of Machine Learning Research*, 13(1): 2409–2464, 2012.
- [166] Robert E Tillman, Arthur Gretton, and Peter Spirtes. Nonlinear directed acyclic structure learning with weakly additive noise models. *Advances in Neural Information Processing Systems*, pages 1847–1855, 2009.
- [167] Ariel Rokem, M Trumpis, and F Perez. Nitime: time-series analysis for neuroimaging data. In

- Proceedings of the 8th Python in Science Conference*, pages 68–75, 2009. URL <http://nipy.org/nitime/>.
- [168] K.E. Stephan and K.J. Friston. Functional connectivity. In Larry R. Squire, editor, *Encyclopedia of Neuroscience*, pages 391–397. Academic Press, Oxford, 2009. ISBN 978-0-08-045046-9. doi: 10.1016/B978-008045046-9.00308-9. URL <https://www.sciencedirect.com/science/article/pii/B9780080450469003089>.
- [169] Michael Eichler and Vanessa Didelez. On granger causality and the effect of interventions in time series. *Lifetime data analysis*, 16(1):3–32, 2010.
- [170] Daniel Chicharro and Anders Ledberg. When two become one: the limits of causality analysis of brain dynamics. *PLoS One*, 7(3):e32466, 2012.
- [171] Anil K Seth, Paul Chorley, and Lionel C Barnett. Granger causality analysis of fmri bold signals is invariant to hemodynamic convolution but not downsampling. *Neuroimage*, 65:540–555, 2013.
- [172] Douglas Zhou, Yaoyu Zhang, Yanyang Xiao, and David Cai. Analysis of sampling artifacts on the granger causality analysis for topology extraction of neuronal dynamics. *Frontiers in computational neuroscience*, 8:75, 2014.
- [173] Victor Solo. State-space analysis of granger-geweke causality measures with application to fmri. *Neural computation*, 28(5):914–949, 2016.
- [174] Lionel Barnett and Anil K Seth. Detectability of granger causality for subsampled continuous-time neurophysiological processes. *Journal of neuroscience methods*, 275:93–121, 2017.
- [175] Luca Faes, Giandomenico Nollo, Sebastiano Stramaglia, and Daniele Marinazzo. Multiscale granger causality. *Physical Review E*, 96(4):042150, 2017.
- [176] Irene Winkler, Danny Panknin, Daniel Bartz, Klaus-Robert Müller, and Stefan Haufe. Validity of time reversal for testing granger causality. *IEEE Transactions on Signal Processing*, 64(11):2746–2760, 2016.

- [177] Ruben Sanchez-Romero, Joseph D Ramsey, Kun Zhang, Madelyn RK Glymour, Biwei Huang, and Clark Glymour. Estimating feedforward and feedback effective connections from fmri time series: Assessments of statistical methods. *Network Neuroscience*, 3(2):274–306, 2019.
- [178] Norman R Swanson and Clive WJ Granger. Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions. *Journal of the American Statistical Association*, 92(437):357–367, 1997.
- [179] Jakob Runge, Sebastian Bathiany, Erik Bollt, Gustau Camps-Valls, Dim Coumou, Ethan Deyle, Clark Glymour, Marlene Kretschmer, Miguel D Mahecha, Jordi Muñoz-Marí, et al. Inferring causation from time series in earth system sciences. *Nature communications*, 10(1):1–13, 2019.
- [180] Shigeru Shinomoto. Estimating the firing rate. In *Analysis of Parallel Spike Trains*, pages 21–35. Springer, 2010.
- [181] 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.
- [182] Caitlin Mullins, Gord Fishell, and Richard W Tsien. Unifying views of autism spectrum disorders: a consideration of autoregulatory feedback loops. *Neuron*, 89(6):1131–1156, 2016.
- [183] Julia M Sheffield and Deanna M Barch. Cognition and resting-state functional connectivity in schizophrenia. *Neuroscience & Biobehavioral Reviews*, 61:108–120, 2016.
- [184] Marloes H Maathuis, Markus Kalisch, Peter Bühlmann, et al. Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics*, 37(6A):3133–3164, 2009.
- [185] Ana-Maria Šimundić. Measures of diagnostic accuracy: basic definitions. *Ejifcc*, 19(4):203, 2009.
- [186] Jørgen Hilden and Paul Glasziou. Regret graphs, diagnostic uncertainty and youden’s index. *Statistics in medicine*, 15(10):969–986, 1996.
- [187] Bart Bussmann, Jannes Nys, and Steven Latré. Neural additive vector autoregression models for causal discovery in time series. In *International Conference on Discovery Science*, pages 446–460. Springer, 2021.

- [188] Jakob Runge, Peer Nowack, Marlene Kretschmer, Seth Flaxman, and Dino Sejdinovic. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5(11): eaau4996, 2019.
- [189] Sebastian Weichwald, Martin E. Jakobsen, Phillip B. Mogensen, Lasse Petersen, Nikolaj Thams, and Gherardo Varando. Causal structure learning from time series: Large regression coefficients may predict causal links better in practice than small p-values. In Hugo Jair Escalante and Raia Hadsell, editors, *Proceedings of the NeurIPS 2019 Competition and Demonstration Track*, volume 123 of *Proceedings of Machine Learning Research*, pages 27–36. PMLR, 08–14 Dec 2020. URL <https://proceedings.mlr.press/v123/weichwald20a.html>.
- [190] Saskia EJ de Vries, Jerome A Lecoq, Michael A Buice, Peter A Groblewski, Gabriel K Ocker, Michael Oliver, David Feng, Nicholas Cain, Peter Ledochowitsch, Daniel Millman, et al. A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience*, 23(1):138–151, 2020.
- [191] Allen-Brain-Observatory. Allen institute for brain science. Available from: <https://portal.brain-map.org/explore/circuits/visual-coding-neuropixels>, October, 2019.
- [192] James J Jun, Nicholas A Steinmetz, Joshua H Siegle, Daniel J Denman, Marius Bauza, Brian Barbarits, Albert K Lee, Costas A Anastassiou, Alexandru Andrei, Çağatay Aydın, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, 2017.
- [193] Mehrdad Dadgostar, Seyed Kamaledin Setarehdan, Sohrab Shahzadi, and Ata Akin. Functional connectivity of the pfc via partial correlation. *Optik*, 127(11):4748–4754, 2016.
- [194] Yikai Wang, Jian Kang, Phebe B Kemmer, and Ying Guo. An efficient and reliable statistical method for estimating functional connectivity in large scale brain networks using partial correlation. *Frontiers in neuroscience*, 10:123, 2016.
- [195] Olaf Sporns, Dante R Chialvo, Marcus Kaiser, and Claus C Hilgetag. Organization, development and function of complex brain networks. *Trends in cognitive sciences*, 8(9):418–425, 2004.

- [196] Martijn P van den Heuvel, Cornelis J Stam, Maria Boersma, and HE Hulshoff Pol. Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *Neuroimage*, 43(3):528–539, 2008.
- [197] Issei Ueda, Shingo Kakeda, Keita Watanabe, Koichiro Sugimoto, Natsuki Igata, Junji Moriya, Kazuhiro Takemoto, Asuka Katsuki, Reiji Yoshimura, Osamu Abe, et al. Brain structural connectivity and neuroticism in healthy adults. *Scientific reports*, 8(1):1–8, 2018.
- [198] Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
- [199] Ali Shojaie and Emily B Fox. Granger causality: A review and recent advances. *Annual Review of Statistics and Its Application*, 9:289–319, 2022.
- [200] Jonathan Pillow and Peter Latham. Neural characterization in partially observed populations of spiking neurons. *Advances in Neural Information Processing Systems*, 20, 2007.
- [201] Richard Scheines. An introduction to causal inference. Technical report, Carnegie Mellon University, 1997.
- [202] Frederick Eberhardt and Richard Scheines. Interventions and causal inference. *Philosophy of science*, 74(5):981–995, 2007.
- [203] Nicolai Meinshausen and Peter Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- [204] Yue Yu, Jie Chen, Tian Gao, and Mo Yu. Dag-gnn: Dag structure learning with graph neural networks. In *International Conference on Machine Learning*, pages 7154–7163. PMLR, 2019.
- [205] Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems*, 31, 2018.
- [206] Jiji Zhang and Peter L Spirtes. Strong faithfulness and uniform consistency in causal inference. *arXiv preprint arXiv:1212.2506*, 2012.

- [207] Yimin Huang and Marco Valtorta. Pearl’s calculus of intervention is complete. *arXiv preprint arXiv:1206.6831*, 2012.
- [208] Steffen L Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.
- [209] Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78, 2006.
- [210] Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.
- [211] Clive WJ Granger. *Essays in econometrics: collected papers of Clive WJ Granger*, volume 32. Cambridge University Press, 2001.
- [212] Michael Eichler. Graphical modelling of multivariate time series. *Probability Theory and Related Fields*, 153:233–268, 2012.
- [213] Tianjiao Chu, Clark Glymour, and Greg Ridgeway. Search for additive nonlinear time series causal models. *Journal of Machine Learning Research*, 9(5), 2008.
- [214] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Causal inference on time series using restricted structural equation models. *Advances in neural information processing systems*, 26, 2013.
- [215] José-Luis Molina and Santiago Zazo. Causal reasoning for the analysis of rivers runoff temporal behavior. *Water Resources Management*, 31:4669–4681, 2017.
- [216] Peter Miersch, Shijie Jiang, Oldrich Rakovec, and Jakob Zscheischler. Identifying drivers of river floods using causal inference. Technical report, Copernicus Meetings, 2023.
- [217] John Geweke. Inference and causality in economic time series models. *Handbook of econometrics*, 2: 1101–1144, 1984.
- [218] Charles K Assaad, Emilie Devijver, and Eric Gaussier. Entropy-based discovery of summary causal graphs in time series. *Entropy*, 24(8):1156, 2022.

- [219] Fatimah Abdul Razak and Henrik Jeldtoft Jensen. Quantifying ‘causality’ in complex systems: understanding transfer entropy. *PLoS One*, 9(6):e99462, 2014.
- [220] Piraporn Jangyodsuk, Dong-Jun Seo, and Jean Gao. Causal graph discovery for hydrological time series knowledge discovery. 2014.
- [221] Mathias Drton and Marloes H Maathuis. Structure learning in graphical modeling. *arXiv preprint arXiv:1606.02359*, 2016.
- [222] Christopher Meek. Strong-completeness and faithfulness in belief networks. Technical report, Carnegie Mellon University, 1995.
- [223] David Maxwell Chickering. Learning equivalence classes of bayesian-network structures. *The Journal of Machine Learning Research*, 2:445–498, 2002.
- [224] Thomas S Verma and Judea Pearl. Equivalence and synthesis of causal models. In *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence, 1991*. Elsevier, 1991.
- [225] C Meek. Causal inference and causal explanation with background knowledge in uncertainty in artificial intelligence 11, 1995.
- [226] Raha Moraffah, Paras Sheth, Mansooreh Karami, Anchit Bhattacharya, Qianru Wang, Anique Tahir, Adrienne Raglin, and Huan Liu. Causal inference for time series analysis: Problems, methods and evaluation. *Knowledge and Information Systems*, pages 1–45, 2021.
- [227] Robb J Muirhead. *Aspects of multivariate statistical theory*. John Wiley & Sons, 2009.
- [228] Kenji Fukumizu, Arthur Gretton, Xiaohai Sun, and Bernhard Schölkopf. Kernel measures of conditional dependence. *Advances in neural information processing systems*, 20, 2007.
- [229] Harold Hotelling. New light on the correlation coefficient and its transforms. *Journal of the Royal Statistical Society. Series B (Methodological)*, 15(2):193–232, 1953.
- [230] Dimitris N Politis and Joseph P Romano. The stationary bootstrap. *Journal of the American Statistical association*, 89(428):1303–1313, 1994.

- [231] Andrei Nikolaevich Kolmogorov and Yu A Rozanov. On strong mixing conditions for stationary gaussian processes. *Theory of Probability & Its Applications*, 5(2):204–208, 1960.
- [232] Richard C Bradley. Basic properties of strong mixing conditions. a survey and some open questions. *Probability surveys*, 2:107–144, 2005.
- [233] Elias Masry. The estimation of the correlation coefficient of bivariate data under dependence: Convergence analysis. *Statistics & Probability Letters*, 81(8):1039–1045, 2011.
- [234] Tuan D Pham and Lanh T Tran. Some mixing properties of time series models. *Stochastic processes and their applications*, 19(2):297–303, 1985.
- [235] Sebastian Weichwald, Martin E Jakobsen, Phillip B Mogensen, Lasse Petersen, Nikolaj Thams, and Gherardo Varando. Causal structure learning from time series: Large regression coefficients may predict causal links better in practice than small p-values. In *NeurIPS 2019 Competition and Demonstration Track*, pages 27–36. PMLR, 2020.
- [236] Henry W Querfurth and Frank M LaFerla. Alzheimer’s disease. *New England Journal of Medicine*, 362(4):329–344, 2010.
- [237] Christoph Laske, Hamid R Sohrabi, Shaun M Frost, Karmele López-de Ipiña, Peter Garrard, Massimo Buscema, Justin Dauwels, Surjo R Soekadar, Stephan Mueller, Christoph Linnemann, et al. Innovative diagnostic tools for early detection of alzheimer’s disease. *Alzheimer’s & Dementia*, 11(5):561–578, 2015.
- [238] Hongming Li, Mohamad Habes, David A Wolk, Yong Fan, Alzheimer’s Disease Neuroimaging Initiative, et al. A deep learning model for early prediction of alzheimer’s disease dementia based on hippocampal magnetic resonance imaging data. *Alzheimer’s & Dementia*, 15(8):1059–1070, 2019.
- [239] A Ashraf, Z Fan, DJ Brooks, and P Edison. Cortical hypermetabolism in mci subjects: a compensatory mechanism? *European journal of nuclear medicine and molecular imaging*, 42:447–458, 2015.
- [240] Akinori Nakamura, Pablo Cuesta, Takashi Kato, Yutaka Arahata, Kaori Iwata, Misako Yamagishi,

- Izumi Kuratsubo, Kimiko Kato, Masahiko Bundo, Kersten Diers, et al. Early functional network alterations in asymptomatic elders at risk for alzheimer's disease. *Scientific Reports*, 7(1):1–11, 2017.
- [241] Matthew R Brier, Jewell B Thomas, Anne M Fagan, Jason Hassenstab, David M Holtzman, Tammie L Benzinger, John C Morris, and Beau M Ances. Functional connectivity and graph theory in preclinical alzheimer's disease. *Neurobiology of aging*, 35(4):757–768, 2014.
- [242] Yvette I Sheline and Marcus E Raichle. Resting state functional connectivity in preclinical alzheimer's disease. *Biological psychiatry*, 74(5):340–347, 2013.
- [243] Takao Yamasaki, Hiroyuki Muranaka, Yumiko Kaseda, Yasuyo Mimori, and Shozo Tobimatsu. Understanding the pathophysiology of alzheimer's disease and mild cognitive impairment: A mini review on fmri and erp studies. *Neurology research international*, 2012, 2012.
- [244] Olaf Sporns. The human connectome: origins and challenges. *Neuroimage*, 80:53–61, 2013.
- [245] Sidong Liu, Weidong Cai, Siqi Liu, Fan Zhang, Michael Fulham, Dagan Feng, Sonia Pujol, and Ron Kikinis. Multimodal neuroimaging computing: the workflows, methods, and platforms. *Brain informatics*, 2(3):181–195, 2015.
- [246] Chen Xue, Baoyu Yuan, Yingying Yue, Jiani Xu, Siyu Wang, Meilin Wu, Nanxi Ji, Xingzhi Zhou, Yilin Zhao, Jiang Rao, et al. Distinct disruptive patterns of default mode subnetwork connectivity across the spectrum of preclinical alzheimer's disease. *Frontiers in Aging Neuroscience*, 11:307, 2019.
- [247] Jingyuan E Chen, Mikail Rubinov, and Catie Chang. Methods and considerations for dynamic analysis of functional mr imaging data. *Neuroimaging Clinics*, 27(4):547–560, 2017.
- [248] Shella Keilholz, Cesar Caballero-Gaudes, Peter Bandettini, Gustavo Deco, and Vince Calhoun. Time-resolved resting-state functional magnetic resonance imaging analysis: current status, challenges, and new directions. *Brain connectivity*, 7(8):465–481, 2017.
- [249] Vanessa Scarapicchia, Erin L Mazerolle, John D Fisk, Lesley J Ritchie, and Jodie R Gawryluk. Resting state bold variability in alzheimer's disease: a marker of cognitive decline or cerebrovascular status? *Frontiers in aging neuroscience*, 10:39, 2018.

- [250] Kun Wang, Meng Liang, Liang Wang, Lixia Tian, Xinqing Zhang, Kuncheng Li, and Tianzi Jiang. Altered functional connectivity in early alzheimer's disease: A resting-state fmri study. *Human brain mapping*, 28(10):967–978, 2007.
- [251] Hee Jin Kim, Jungho Cha, Jong-Min Lee, Ji Soo Shin, Na-Yeon Jung, Yeo Jin Kim, Yearn Seong Choe, Kyung Han Lee, Sung Tae Kim, Jae Seung Kim, et al. Distinctive resting state network disruptions among alzheimer's disease, subcortical vascular dementia, and mixed dementia patients. *Journal of Alzheimer's Disease*, 50(3):709–718, 2016.
- [252] Heidi IL Jacobs, Joaquim Radua, Helen C Lückmann, and Alexander T Sack. Meta-analysis of functional network alterations in alzheimer's disease: toward a network biomarker. *Neuroscience & Biobehavioral Reviews*, 37(5):753–765, 2013.
- [253] Hui-Jie Li, Xiao-Hui Hou, Han-Hui Liu, Chun-Lin Yue, Yong He, and Xi-Nian Zuo. Toward systems neuroscience in mild cognitive impairment and alzheimer's disease: A meta-analysis of 75 fmri studies. *Human brain mapping*, 36(3):1217–1232, 2015.
- [254] AmanPreet Badhwar, Angela Tam, Christian Dansereau, Pierre Orban, Felix Hoffstaedter, and Pierre Bellec. Resting-state network dysfunction in alzheimer's disease: a systematic review and meta-analysis. *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring*, 8:73–85, 2017.
- [255] Natalina Gour, Olivier Felician, Mira Didic, Lejla Koric, Claude Gueriot, Valérie Chanoine, Sylviane Confort-Gouny, Maxime Guye, Mathieu Ceccaldi, and Jean Philippe Ranjeva. Functional connectivity changes differ in early and late-onset alzheimer's disease. *Human brain mapping*, 35(7):2978–2994, 2014.
- [256] Marco Bozzali, Claire Dowling, Laura Serra, Barbara Spano, Mario Torso, Camillo Marra, Diana Castelli, Nicholas G Dowell, Giacomo Koch, Carlo Caltagirone, et al. The impact of cognitive reserve on brain functional connectivity in alzheimer's disease. *Journal of Alzheimer's Disease*, 44(1):243–250, 2015.
- [257] Frank G Hillary and Jordan H Grafman. Injured brains and adaptive networks: the benefits and costs of hyperconnectivity. *Trends in cognitive sciences*, 21(5):385–401, 2017.

- [258] Scott Marek and Nico UF Dosenbach. The frontoparietal network: function, electrophysiology, and importance of individual precision mapping. *Dialogues in clinical neuroscience*, 2022.
- [259] Stuart Oldham and Alex Fornito. The development of brain network hubs. *Developmental cognitive neuroscience*, 36:100607, 2019.
- [260] Romana Rytsar, Eleonora Fornari, Richard S Frackowiak, Joseph A Ghika, and Maria G Knyazeva. Inhibition in early alzheimer’s disease: an fmri-based study of effective connectivity. *Neuroimage*, 57(3):1131–1139, 2011.
- [261] Uttam Khatri, Ramesh Kumar Lama, and Goo-Rak Kwon. Diagnosis of alzheimer’s disease using effective connectivity of rs-fmri. In *2021 36th International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC)*, pages 1–4. IEEE, 2021.
- [262] Alfonso Nieto-Castanon. Conn functional connectivity toolbox (rrid: Scr_009550), version 21, 2021.
- [263] Karl J Friston, Andrew P Holmes, Keith J Worsley, J-P Poline, Chris D Frith, and Richard SJ Frackowiak. Statistical parametric maps in functional imaging: a general linear approach. *Human brain mapping*, 2(4):189–210, 1994.
- [264] Nathalie Tzourio-Mazoyer, Brigitte Landeau, Dimitri Papathanassiou, Fabrice Crivello, Octave Etard, Nicolas Delcroix, Bernard Mazoyer, and Marc Joliot. Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. *Neuroimage*, 15(1):273–289, 2002.
- [265] Salim Arslan, Sofia Ira Ktena, Antonios Makropoulos, Emma C Robinson, Daniel Rueckert, and Sarah Parisot. Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *Neuroimage*, 170:5–30, 2018.
- [266] Gopikrishna Deshpande, Stephan LaConte, George Andrew James, Scott Peltier, and Xiaoping Hu. Multivariate granger causality analysis of fmri data. *Human brain mapping*, 30(4):1361–1373, 2009.
- [267] Tijn M Schouten, Marisa Koini, Frank De Vos, Stephan Seiler, Jeroen Van Der Grond, Anita Lechner, Anne Hafkemeijer, Christiane Möller, Reinhold Schmidt, Mark De Rooij, et al. Combining anatomical,

- diffusion, and resting state functional magnetic resonance imaging for individual classification of mild and moderate alzheimer's disease. *NeuroImage: Clinical*, 11:46–51, 2016.
- [268] Karen K Yuen. The two-sample trimmed t for unequal population variances. *Biometrika*, 61(1): 165–170, 1974.
- [269] Jürgen Hänggi, Johannes Streffer, Lutz Jäncke, and Christoph Hock. Volumes of lateral temporal and parietal structures distinguish between healthy aging, mild cognitive impairment, and alzheimer's disease. *Journal of Alzheimer's Disease*, 26(4):719–734, 2011.
- [270] Novraj S Dhanjal, Jane E Warren, Maneesh C Patel, and Richard JS Wise. Auditory cortical function during verbal episodic memory encoding in alzheimer's disease. *Annals of neurology*, 73(2):294–302, 2013.
- [271] Alan M Palmer and Mark A Burns. Selective increase in lipid peroxidation in the inferior temporal cortex in alzheimer's disease. *Brain research*, 645(1-2):338–342, 1994.
- [272] Stephen W Scheff, Douglas A Price, Frederick A Schmitt, Melissa A Scheff, and Elliott J Mufson. Synaptic loss in the inferior temporal gyrus in mild cognitive impairment and alzheimer's disease. *Journal of Alzheimer's Disease*, 24(3):547–557, 2011.
- [273] Nancy Beyer, David TR Coulson, Shirley Heggarty, Rivka Ravid, G Brent Irvine, Jan Hellemans, and Janet A Johnston. Znt3 mrna levels are reduced in alzheimer's disease post-mortem brain. *Molecular neurodegeneration*, 4:1–10, 2009.
- [274] Yanping Mao, Zhengluan Liao, Xiaozheng Liu, Ting Li, Jiaojiao Hu, Dansheng Le, Yangliu Pei, Wangdi Sun, Jixin Lin, Yaju Qiu, et al. Disrupted balance of long and short-range functional connectivity density in alzheimer's disease (ad) and mild cognitive impairment (mci) patients: a resting-state fmri study. *Annals of Translational Medicine*, 9(1), 2021.
- [275] Libuse Brachova, Lih-Fen Lue, James Schultz, Tarek El Rashidy, and Joseph Rogers. Association cortex, cerebellum, and serum concentrations of c1q and factor b in alzheimer's disease. *Molecular brain research*, 18(4):329–334, 1993.

- [276] Lih-Fen Lue, Libuse Brachova, W Harold Civin, and Joseph Rogers. Inflammation, $a\beta$ deposition, and neurofibrillary tangle formation as correlates of alzheimer's disease neurodegeneration. *Journal of neuropathology & experimental neurology*, 55(10):1083–1088, 1996.
- [277] Susanne Neufang, Atae Akhrif, Valentin Riedl, Hans Förstl, Alexander Kurz, Claus Zimmer, Christian Sorg, and Afra M Wohlschläger. Disconnection of frontal and parietal areas contributes to impaired attention in very early alzheimer's disease. *Journal of Alzheimer's Disease*, 25(2):309–321, 2011.
- [278] Bo Zhou, Yong Liu, Zengqiang Zhang, Ningyu An, Hongxiang Yao, Pan Wang, Luning Wang, Xi Zhang, and Tianzi Jiang. Impaired functional connectivity of the thalamus in alzheimer's disease and mild cognitive impairment: a resting-state fmri study. *Current Alzheimer Research*, 10(7):754–766, 2013.
- [279] Heiko Braak and E Braak. Alzheimer's disease affects limbic nuclei of the thalamus. *Acta neuropathologica*, 81(3):261–268, 1991.
- [280] Laura W de Jong, Karin van der Hiele, Ilya M Veer, JJ Houwing, RGJ Westendorp, ELEM Bollen, Paul W de Bruin, HAM Middelkoop, Mark A van Buchem, and Jeroen van der Grond. Strongly reduced volumes of putamen and thalamus in alzheimer's disease: an mri study. *Brain*, 131(12):3277–3285, 2008.
- [281] Y Grignon, C Duyckaerts, M Benneceb, and J-J Hauw. Cytoarchitectonic alterations in the supra-marginal gyrus of late onset alzheimer's disease. *Acta neuropathologica*, 95:395–406, 1998.
- [282] Rahul S Desikan, Howard J Cabral, Christopher P Hess, William P Dillon, Christine M Glastonbury, Michael W Weiner, Nicholas J Schmansky, Douglas N Greve, David H Salat, Randy L Buckner, et al. Automated mri measures identify individuals with mild cognitive impairment and alzheimer's disease. *Brain*, 132(8):2048–2057, 2009.
- [283] Huimin Wu, Yu Song, Xinyi Yang, Shanshan Chen, Honglin Ge, Zheng Yan, Wenzhang Qi, Qianqian Yuan, Xuhong Liang, Xingjian Lin, et al. Functional and structural alterations of dorsal attention network in preclinical and early-stage alzheimer's disease. *CNS Neuroscience & Therapeutics*, 29(6): 1512–1524, 2023.

- [284] A Johnen, L Brandstetter, H Lohmann, and T Duning. P12. neural correlates of apraxia in mild dementia of alzheimer's disease—a voxel-based morphometry study. *Clinical Neurophysiology*, 126(8): e93, 2015.
- [285] Nicolas Villain, Béatrice Desgranges, Fausto Viader, Vincent De La Sayette, Florence Mézenge, Brigitte Landeau, Jean-Claude Baron, Francis Eustache, and Gaël Chételat. Relationships between hippocampal atrophy, white matter disruption, and gray matter hypometabolism in alzheimer's disease. *Journal of Neuroscience*, 28(24):6174–6181, 2008.
- [286] Silvia Caminiti, Arianna Sala, Andrea Pilotto, Luca Presotto, Valentina Garibotto, Marcello D'Amelio, Claudio Liguori, Nicola Mercuri, Alessandro Padovani, and Daniela Perani. Imaging dopamine system transporter activity and connectivity in alzheimer's dementia: Neuroimaging/new imaging methods. *Alzheimer's & Dementia*, 16:e043304, 2020.
- [287] Daniele Mascali, Mauro DiNuzzo, Tommaso Gili, Marta Moraschi, Michela Fratini, Bruno Maraviglia, Laura Serra, Marco Bozzali, and Federico Giove. Intrinsic patterns of coupling between correlation and amplitude of low-frequency fmri fluctuations are disrupted in degenerative dementia mainly due to functional disconnection. *PLOS ONE*, 10(4):1–18, 04 2015. doi: 10.1371/journal.pone.0120988. URL <https://doi.org/10.1371/journal.pone.0120988>.
- [288] Gary W Van Hoesen, Jean C Augustinack, Jason Dierking, Sarah J Redman, and Ramasamy Thangavel. The parahippocampal gyrus in alzheimer's disease: clinical and preclinical neuroanatomical correlates. *Annals of the New York Academy of Sciences*, 911(1):254–274, 2000.
- [289] Ramasamy Thangavel, Gary W Van Hoesen, and Asgar Zaheer. Posterior parahippocampal gyrus pathology in alzheimer's disease. *Neuroscience*, 154(2):667–676, 2008.
- [290] PK Mölsä, E Säkö, L Paljärvi, JO Rinne, and UK Rinne. Alzheimer's disease: neuropathological correlates of cognitive and motor disorders. *Acta neurologica scandinavica*, 75(6):376–384, 1987.
- [291] D Nochlin, G Van Belle, TD Bird, and SM Sumi. Comparison of the severity of neuropathologic changes in familial and sporadic alzheimer's disease. *Alzheimer disease and associated disorders*, 7(4): 212–222, 1993.

- [292] Yvette I Sheline, Marcus E Raichle, Abraham Z Snyder, John C Morris, Denise Head, Suzhi Wang, and Mark A Mintun. Amyloid plaques disrupt resting state default mode network connectivity in cognitively normal elderly. *Biological psychiatry*, 67(6):584–587, 2010.
- [293] Catharine L Joachim, James H Morris, and Dennis J Selkoe. Diffuse senile plaques occur commonly in the cerebellum in alzheimer’s disease. *The American journal of pathology*, 135(2):309, 1989.
- [294] Heidi IL Jacobs, David A Hopkins, Helen C Mayrhofer, Emiliano Bruner, Fred W van Leeuwen, Wijnand Raaijmakers, and Jeremy D Schmahmann. The cerebellum in alzheimer’s disease: evaluating its role in cognitive decline. *Brain*, 141(1):37–47, 2018.
- [295] Ilona Eliasova, Lubomira Anderkova, Radek Marecek, and Irena Rektorova. Non-invasive brain stimulation of the right inferior frontal gyrus may improve attention in early alzheimer’s disease: a pilot study. *Journal of the neurological sciences*, 346(1-2):318–322, 2014.
- [296] Antti Cajanus, Eino Solje, Juha Koikkalainen, Jyrki Lötjönen, Noora-Maria Suhonen, Ilona Hallikainen, Ritva Vanninen, Päivi Hartikainen, Matteo de Marco, Annalena Venneri, et al. The association between distinct frontal brain volumes and behavioral symptoms in mild cognitive impairment, alzheimer’s disease, and frontotemporal dementia. *Frontiers in neurology*, 10:1059, 2019.
- [297] Dingailu Ma, Irfete S Fetahu, Mei Wang, Rui Fang, Jiahui Li, Hang Liu, Tobin Gramyk, Isabella Iwanicki, Sophie Gu, Winnie Xu, et al. The fusiform gyrus exhibits an epigenetic signature for alzheimer’s disease. *Clinical epigenetics*, 12(1):1–16, 2020.
- [298] Jennifer L Whitwell. Progression of atrophy in alzheimer’s disease and related disorders. *Neurotoxicity research*, 18(3-4):339–346, 2010.
- [299] Yong He, Liang Wang, Yufeng Zang, Lixia Tian, Xinqing Zhang, Kuncheng Li, and Tianzi Jiang. Regional coherence changes in the early stages of alzheimer’s disease: a combined structural and resting-state functional mri study. *Neuroimage*, 35(2):488–500, 2007.
- [300] Eini Niskanen, Mervi Könönen, Sara Määttä, Merja Hallikainen, Miia Kivipelto, Silvia Casarotto, Marcello Massimini, Ritva Vanninen, Esa Mervaala, Jari Karhu, et al. New insights into alzheimer’s disease progression: a combined tms and structural mri study. *PLoS One*, 6(10):e26113, 2011.

- [301] Martin Sjöbeck and Elisabet Englund. Alzheimer's disease and the cerebellum: a morphologic study on neuronal and glial changes. *Dementia and geriatric cognitive disorders*, 12(3):211–218, 2001.
- [302] Ioannis A Mavroudis, Marina G Manani, Foivos Petrides, Konstantina Petsoglou, Samuel D Njau, Vasiliki G Costa, and Stavros J Baloyannis. Dendritic and spinal pathology of the purkinje cells from the human cerebellar vermis in alzheimer's disease. *Psychiatria Danubina*, 25(3):0–226, 2013.
- [303] Shuhua Ren, Qi Huang, Donglang Jiang, Lin Huang, Ying Wang, Yihui Guan, Qihao Guo, and Fang Xie. Brain amyloid accumulation and glucose hypometabolism in chinese alzheimer's disease population: Neuroimaging/multi-modal comparisons. *Alzheimer's & Dementia*, 16:e043567, 2020.
- [304] Liu Yang, Yan Yan, Yuxia Li, Xiaochen Hu, Jie Lu, Piu Chan, Tianyi Yan, and Ying Han. Frequency-dependent changes in fractional amplitude of low-frequency oscillations in alzheimer's disease: a resting-state fmri study. *Brain imaging and behavior*, 14:2187–2201, 2020.
- [305] ES Garcia Martin, R De Hoz, B Rojas, P Gil, R Yubero, and JM Ramirez. Macular nerve-fiber-layer measurement in early stage alzheimer's disease using optical coherence tomography. *Acta Ophthalmologica*, 91, 2013.
- [306] Huanqing Yang, Hua Xu, Qingfeng Li, Yan Jin, Weixiong Jiang, Jinghua Wang, Yina Wu, Wei Li, Cece Yang, Xia Li, et al. Study of brain morphology change in alzheimer's disease and amnesic mild cognitive impairment compared with normal controls. *General psychiatry*, 32(2), 2019.
- [307] Carolyn A Fredericks, Virginia E Sturm, Jesse A Brown, Alice Y Hua, Murat Bilgel, Dean F Wong, Susan M Resnick, and William W Seeley. Early affective changes and increased connectivity in preclinical alzheimer's disease. *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring*, 10:471–479, 2018.
- [308] Lucía Penalba-Sánchez, Patrícia Oliveira-Silva, Alexander Luke Sumich, and Ignacio Cifre. Increased functional connectivity patterns in mild alzheimer's disease: A rsfmri study. *Frontiers in Aging Neuroscience*, 14:1037347, 2023.
- [309] Catherine Warriier, Patrick Wong, Virginia Penhune, Robert Zatorre, Todd Parrish, Daniel Abrams,

- and Nina Kraus. Relating structure to function: Heschl's gyrus and acoustic processing. *Journal of Neuroscience*, 29(1):61–69, 2009.
- [310] Leticia Fernández, Carlos Velásquez, Juan A García Porrero, Enrique Marco de Lucas, and Juan Martino. Heschl's gyrus fiber intersection area: a new insight on the connectivity of the auditory-language hub. *Neurosurgical focus*, 48(2):E7, 2020.
- [311] Michel S Gazzaniga, Richard B Ivry, and George R Mangun. Cognitive neuroscience: The biology of the mind . new york: W. w. Norton & Company, 199:166, 2002.
- [312] Mircea Steriade and Rodolfo R Llinás. The functional states of the thalamus and the associated neuronal interplay. *Physiological reviews*, 68(3):649–742, 1988.
- [313] Marie-Aurélie Bruno, Steven Laureys, and Athena Demertzi. Coma and disorders of consciousness. *Handbook of clinical neurology*, 118:205–213, 2013.
- [314] John P Aggleton, Shane M O'Mara, Seralynne D Vann, Nick F Wright, Marian Tsanov, and Jonathan T Erichsen. Hippocampal–anterior thalamic pathways for memory: uncovering a network of direct and indirect actions. *European Journal of Neuroscience*, 31(12):2292–2307, 2010.
- [315] Claire Boutet, Marie Chupin, Stéphane Lehericy, Linda Marrakchi-Kacem, Stéphane Epelbaum, Cyril Poupon, Christopher Wiggins, Alexandre Vignaud, Dominique Hasboun, Bénédicte Defontaine, et al. Detection of volume loss in hippocampal layers in alzheimer's disease using 7 t mri: a feasibility study. *NeuroImage: Clinical*, 5:341–348, 2014.
- [316] Y Lakshmisha Rao, B Ganaraja, BV Murlimanju, Teresa Joy, Ashwin Krishnamurthy, and Amit Agrawal. Hippocampus and its involvement in alzheimer's disease: a review. *3 Biotech*, 12(2):55, 2022.
- [317] Lorenzo Pasquini, Martin Scherr, Masoud Tahmasian, Chun Meng, Nicholas E Myers, Marion Ortner, Mark Mühlau, Alexander Kurz, Hans Förstl, Claus Zimmer, et al. Link between hippocampus' raised local and eased global intrinsic connectivity in ad. *Alzheimer's & Dementia*, 11(5):475–484, 2015.
- [318] Guangyu Chen, B Douglas Ward, Gang Chen, and Shi-Jiang Li. Decreased effective connectivity from

- cortices to the right parahippocampal gyrus in alzheimer's disease subjects. *Brain connectivity*, 4(9): 702–708, 2014.
- [319] Masoud Tahmasian, Lorenzo Pasquini, Martin Scherr, Chun Meng, Stefan Förster, Satja Mulej Bratec, Kuangyu Shi, Igor Yakushev, Markus Schwaiger, Timo Grimmer, et al. The lower hippocampus global connectivity, the higher its local metabolism in alzheimer disease. *Neurology*, 84(19):1956–1963, 2015.
- [320] Lorenzo Pasquini, Martin Scherr, Masoud Tahmasian, Nicholas E Myers, Marion Ortner, Alexander Kurz, Hans Förstl, Claus Zimmer, Timo Grimmer, Atae Akhrif, et al. Increased intrinsic activity of medial-temporal lobe subregions is associated with decreased cortical thickness of medial-parietal areas in patients with alzheimer's disease dementia. *Journal of Alzheimer's Disease*, 51(1):313–326, 2016.
- [321] Claus Thorn Ekstrøm. *MESS: Miscellaneous Esoteric Statistical Scripts*, 2023. URL <https://CRAN.R-project.org/package=MESS>. R package version 0.5.12.
- [322] Michael W Weiner, Dallas P Veitch, Paul S Aisen, Laurel A Beckett, Nigel J Cairns, Robert C Green, Danielle Harvey, Clifford R Jack Jr, William Jagust, John C Morris, et al. The alzheimer's disease neuroimaging initiative 3: Continued innovation for clinical trial improvement. *Alzheimer's & Dementia*, 13(5):561–571, 2017.
- [323] Christopher Fowler, Stephanie R Rainey-Smith, Sabine Bird, Julia Bomke, Pierrick Bourgeat, Belinda M Brown, Samantha C Burnham, Ashley I Bush, Carolyn Chadunow, Steven Collins, et al. Fifteen years of the australian imaging, biomarkers and lifestyle (aibl) study: progress and observations from 2,359 older adults spanning the spectrum from cognitive normality to alzheimer's disease. *Journal of Alzheimer's disease reports*, 5(1):443–468, 2021.
- [324] Daoqiang Zhang, Yaping Wang, Luping Zhou, Hong Yuan, Dinggang Shen, Alzheimer's Disease Neuroimaging Initiative, et al. Multimodal classification of alzheimer's disease and mild cognitive impairment. *Neuroimage*, 55(3):856–867, 2011.
- [325] Yu-Dong Zhang, Shuihua Wang, and Zhengchao Dong. Classification of alzheimer disease based on structural magnetic resonance imaging by kernel support vector machine decision tree. *Progress In Electromagnetics Research*, 144:171–184, 2014.

- [326] Katherine R Gray, Paul Aljabar, Rolf A Heckemann, Alexander Hammers, Daniel Rueckert, Alzheimer's Disease Neuroimaging Initiative, et al. Random forest-based similarity measures for multi-modal classification of alzheimer's disease. *NeuroImage*, 65:167–175, 2013.
- [327] Junhao Wen, Elina Thibeau-Sutre, Mauricio Diaz-Melo, Jorge Samper-González, Alexandre Routier, Simona Bottani, Didier Dormont, Stanley Durrleman, Ninon Burgos, Olivier Colliot, et al. Convolutional neural networks for classification of alzheimer's disease: Overview and reproducible evaluation. *Medical image analysis*, 63:101694, 2020.
- [328] Ahmad Waleed Salehi, Preety Baglat, Brij Bhushan Sharma, Gaurav Gupta, and Ankita Upadhyia. A cnn model: earlier diagnosis and classification of alzheimer disease using mri. In *2020 International Conference on Smart Electronics and Communication (ICOSEC)*, pages 156–161. IEEE, 2020.
- [329] Junwei Jin, Yanting Li, and CL Philip Chen. Pattern classification with corrupted labeling via robust broad learning system. *IEEE Transactions on Knowledge and Data Engineering*, 34(10):4959–4971, 2021.
- [330] Junwei Jin, Biao Geng, Yanting Li, Jing Liang, Yang Xiao, and CL Philip Chen. Flexible label-induced manifold broad learning system for multiclass recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [331] Gang Chen, B Douglas Ward, Chunming Xie, Wenjun Li, Zhilin Wu, Jennifer L Jones, Malgorzata Franczak, Piero Antuono, and Shi-Jiang Li. Classification of alzheimer disease, mild cognitive impairment, and normal cognitive status with large-scale network analysis based on resting-state functional mr imaging. *Radiology*, 259(1):213–221, 2011.
- [332] Yuhui Du, Zening Fu, and Vince D Calhoun. Classification and prediction of brain disorders using functional connectivity: promising but challenging. *Frontiers in neuroscience*, 12:525, 2018.

Appendix A

Appendix to the Time-Aware PC Algorithm

A.1 Proof of Theorem 1

Let $V = \{(v, t) : v \in V, t \in T\}$, $E = \{(u_{v,i}, t_{v,i}) \rightarrow (v, t) : 1 \leq i \leq p, v \in V, t \in \{0, 1, \dots, T\}\}$, and $G = (V, E)$. Rewriting Eq. (3.5) we get,

$$\mathbf{X}(v, t) = g_{v,t}(\{\mathbf{X}(u, k) : (u, k) \in pa_G((v, t))\}, \epsilon_v(t))$$

By Theorem 1.4.1 in [48], the above implies that \mathbf{X} satisfies the DMP with respect to G .

Therefore by Definition 2, the rolled CFC-DPGM, F_τ , has nodes V and edges given by $u_{v,i} \rightarrow v$, since $(u_{v,i}, t_{v,i}) \rightarrow (v, t) \in E$, for $v \in V$ and $1 \leq i \leq K$. That is, $pa_{F_\tau}(v) = \{u_{v,1}, \dots, u_{v,K}\}$.

A.2 Proof of Corollary 7.1

For $t \in I$, i.e. during the experimental/counterfactual intervention, such as controlling the activity of neuron or neuron ablation, the activity $X_{v_1}(t), \dots, X_{v_k}(t)$ would no longer be a function of the activities of the neurons at preceding time points, while the activity of neurons which are not intervened, would still be a function of activity at preceding time points. That is, $X_v(t) = g_{v,t}(X_{pa_{F_\tau}(v)}(t-), \epsilon_v(t))$, for $v \notin \{v_1, \dots, v_k\}$ and $X_v(t) = g_{v_i,t}(\epsilon_{v_i}(t))$ for $v \in \{v_1, \dots, v_k\}$ for $t \in I$, where $g_{v_i,t}(\epsilon_{v_i}(t))$ represents the distribution of neural activity due to the experimental intervention. For example, $g_{v_i,t}$ is identically 0 for neuron ablation, and can

be an oscillating function with a high amplitude and random noise $\epsilon_{v_i}(t)$ for stimulation through external control. The dynamics can be written as

$$X_v(t) = g_{v,t}(X_{pa_{F'_\tau}(v)}(t-), \epsilon_v(t)) \quad (\text{A.1})$$

for $v \in V, t \in I$, where $pa_{F'_\tau}(v) = pa_{F_\tau}(v)$ for $v \notin \{v_1, \dots, v_k\}$ and $pa_{F'_\tau}(v_1) = \dots = pa_{F'_\tau}(v_k) = \Phi$, where Φ denotes the null set. In other words, F'_τ has all connections same as F_τ except that all connections directing to the intervened neurons v_1, \dots, v_k are removed. It follows from Eq. (A.1) and Theorem 3.4.1 that F'_τ is the causal functional connectivity between the neurons in V during time I when v_1, \dots, v_k are subject to experimental intervention.

For ablations, the edges originating from v_1, \dots, v_k , can be removed since the activity of these neurons would be fixed (at zero) and would be trivial variables that can be excluded from the argument of the function $g_{v,t}$ in Eq. (A.1).

A.3 Simulation Study Details

We study the following simulation paradigms.

1. Linear Gaussian Time Series (Figure 3.3a left-column). Let $N(0, \eta)$ denote a Normal random variable with mean 0 and standard deviation η . We define $X_v(t)$ as a linear Gaussian time series for $v = 1, \dots, 4$ whose true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Let $X_v(0) = N(0, \eta)$ for $v = 1, \dots, 4$, and for $t = 1, 2, \dots, 1000$,

$$\begin{aligned} X_1(t) &= 1 + N(0, \eta), & X_2(t) &= -1 + N(0, \eta), \\ X_3(t+1) &= 2X_1(t) + X_2(t) + N(0, \eta), & X_4(t+1) &= 2X_3(t) + N(0, \eta). \end{aligned}$$

We obtain 25 simulations of the entire time series each for different noise levels $\eta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5\}$.

2. Non-linear Non-Gaussian Time Series (Figure 3.3a middle-column). Let $U(0, \eta)$ denote a *Uniformly* distributed random variable on the interval $(0, \eta)$. We define $X_v(t)$ as a non-linear non-Gaussian time series for $v = 1, \dots, 4$ whose true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Let $X_v(0) = U(0, \eta)$ for $v = 1, \dots, 4$

and for $t = 1, 2, \dots, 1000$,

$$\begin{aligned} X_1(t) &= U(0, \eta), & X_2(t) &= U(0, \eta), \\ X_3(t+1) &= 4 \sin(X_1(t)) + 3 \cos(X_2(t)) + U(0, \eta), & X_4(t+1) &= 2 \sin(X_3(t)) + U(0, \eta). \end{aligned}$$

We obtain 25 simulations of the entire time series each for different noise levels $\eta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5\}$.

3. Continuous Time Recurrent Neural Network (CTRNN) (Figure 3.3a right-column). We simulate neural dynamics by Continuous Time Recurrent Neural Networks, Eq. (A.2). $u_j(t)$ is the instantaneous firing rate at time t for a post-synaptic neuron j , w_{ij} is the linear coefficient to pre-synaptic neuron i 's input on the post-synaptic neuron j , $I_j(t)$ is the input current on neuron j at time t , τ_j is the time constant of the post-synaptic neuron j , with i, j being indices for neurons with m being the total number of neurons. Such a model is typically used to simulate neurons as firing rate units,

$$\tau_j \frac{du_j(t)}{dt} = -u_j(t) + \sum_{i=1}^m w_{ij} \sigma(u_i(t)) + I_j(t), j = 1, \dots, m. \quad (\text{A.2})$$

We consider a motif consisting of 4 neurons with $w_{13} = w_{23} = w_{34} = 10$ and $w_{ij} = 0$ otherwise. We also note that in Eq. A.2, activity of each neuron $u_j(t)$ depends on its own past. Therefore, the true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4, 1 \rightarrow 1, 2 \rightarrow 2, 3 \rightarrow 3, 4 \rightarrow 4$. The time constant τ_i is set to 10 msec for each neuron i . We consider $I_i(t)$ to be distributed as independent Gaussian process with the mean of 1 and the standard deviation of η . The signals are sampled at a time gap of $e \approx 2.72$ msec for a total duration of 1000 msec. We obtain 25 simulations of the entire time series each for different noise levels $\eta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5\}$.

The GC graph is computed using the *Nitime* Python library, which fits an MVAR model followed by using the *GrangerAnalyzer* to compute the Granger Causality [167]. The PC algorithm, which requires several samples of a scalar-valued random variable Y_v (measured activity) for neurons $v \in V$, is used to compute DPGM. We define Y_v as a windowed average of recordings over a duration of 50 msec: $Y_v = X_v, v \in V$, and averaging over different 50 msec windows with a gap of 50 msec between consecutive windows yields different Y_v samples. This choice of Y_v performs better than considering Y_v to be neural recordings at time t :

$Y_v = X_v(t), v \in V$, with different t giving different samples of Y_v in previous work [1]. The TPC algorithm computes the rolled CFC-DPGM directly from the signals and, we use a maximum time-delay of interaction of 1 msec.

The choice of thresholds tunes the decision whether a connection exists in the CFC. For DPGM and TPC, increasing the significance level α for conditional independence tests increases the rate of detecting edges, but also increasing the rate of detecting false positives. We consider $\alpha = 0.01, 0.05, 0.1$ for DPGM and TPC. For GC, a likelihood ratio statistic L_{uv} is obtained for testing $A_{uv}(k) = 0$ for $k = 1, \dots, K$. An edge $u \rightarrow v$ is outputted if L_{uv} has a value greater than a threshold. We use a percentile-based threshold, and output an edge $u \rightarrow v$ if L_{uv} is greater than $100(1 - \alpha)$ percentile of L_{ij} 's over all pairs of neurons (i, j) in the graph [108]. We consider $\alpha = 0.01, 0.05, 0.1$ which corresponds to percentile thresholds of 99%, 95%, 90%. For the bootstrap procedure in TPC, we consider 50 bootstrap iterations with bootstrap window length of 50 msec and bootstrap stability cutoff $\gamma = 25\%$.

A.4 Benchmark Datasets

We use the following benchmark datasets from *Causeme* [179; 187].

1. River Runoff. This is a real dataset that consists of time series of river runoff at different stations. The time series have a daily time resolution and only include summer months (June-August). The physical time delay of interaction (as inferred from the river velocity) are roughly below one day, hence the dataset has contemporaneous time interactions. This dataset has 12 variables and 4600 time recordings for each variable.
2. Logistic Map. This is a synthetic dataset generated from logistic map with maximum time delay of 3 and a low dynamical noise and moderate strength of coupling between the variables. This dataset has 5 variables and 300 time recordings per variable.

The PC algorithm is used to compute the DPGM from the scalar-valued random variables Y_v , defined by the average of recordings of a time window of length Δt , and averaging over alternate time windows of length Δt and a gap of Δt between consecutive windows results in samples of Y_v . Considering river runoff has large number of 4600 time recordings while logistic map has only 300 recordings, we used $\Delta t = 50$

and 3 for river runoff and logistic map data respectively. The PC algorithm was implemented with p-value of 0.1 for kernel-based non-linear conditional independence tests. In river-runoff and logistic map data, the TPC algorithm was implemented with maximum time-delay of interaction to be 1 and 3 recordings respectively, as per specification in the datasets, and a significance level $\alpha = 0.1$ for kernel-based conditional independence tests. For the bootstrap procedure in TPC, 50 bootstrap iterations with bootstrap window length of 50 recordings and bootstrap stability threshold $\gamma = 0.01, 0.15$ for river-runoff and logistic map datasets respectively.

A.5 Visual Coding Neuropixels Dataset

For the purpose of application and comparison of the results of the methods discussed in this paper, we restrict our analysis to a 116 days old male mouse (Session ID 791319847) with 555 neurons whose spike trains are recorded simultaneously by six Neuropixel probes. The spike trains during the entire experiment were recorded at a frequency of 1 KHz. We analyze the spike trains for four stimuli categories:

1. Natural scenes, consisting of 118 natural scenes selected from three databases (Berkeley Segmentation Dataset, van Hateren Natural Image Dataset and McGill Calibrated Colour Image Database), with each scene presented briefly for 250ms and then replaced with the next scene image. Each scene is repeated 50 times in random order with intermittent blank intervals.
2. Static gratings consisting of full-field sinusoidal gratings with 6 orientations (the angle of the grating), 5 spatial frequencies (the width of the grating), and 4 phases (the position of the grating) resulting in 120 stimulus conditions. Each grating is presented briefly (250 ms) before being replaced with a different orientation, spatial frequency and phase condition and each condition is repeated 50 times, in random order with intermittent blank intervals.
3. Gabor patches with 3 orientations where the patch center is lying at one of the points in a 9×9 visual field. Each Gabor patch is being presented for 250ms and then replaced by a different patch, and each condition is repeated 50 times in random order with intermittent blank intervals.
4. Full-field flashes, lasting for 250 ms followed by a blank interval of 1.75 s, and then the next flash, totaling 150 repetitions.

This variety of stimuli is ranging from relatively *natural stimuli* invoking mice’s natural habitats (natural scenes) to *artificial stimuli* (static gratings, Gabors and flashes). Among the artificial stimuli, static gratings incorporate sinusoidal patches, while full-field flashes incorporate sharp changes in luminosity in the whole visual field in short period of time, and Gabor patches incorporate sinusoidal patches with declining luminosity with distance from the center of the patch. With this choice of four stimuli we investigate how the variety of stimuli possibly invokes distinct patterns of neuronal interactions and connectivity. We exclude dynamic stimuli like natural movies, and drifting gratings, from this analysis because their results would require more nuanced study and interpretation, which we defer for future analysis.

Preprocessing We convert the spike trains recorded at 1 KHz to bin size of 10 ms by aggregating and then separating by start and end times of each stimuli presentation and obtain the Peri-Stimulus Time Histograms (PSTH) with bin-size 10 ms. We smooth the PSTHs by a Gaussian smoothing kernel of bandwidth 16ms which provides a smoothed version of the PSTH for each neuron and each stimulus presentation. Some examples of the smoothed PSTH are displayed in Figure 3.6. We use the smoothed PSTHs for neurons over each stimuli type as input for inference of the FC between the neurons for each stimuli presentation. For each stimulus presentation, we first selected the set of neurons that were active in at least 25% of the bins in the PSTH, and then collected the set of unique neurons over all stimuli, which resulted in 54, 43, 19 and 36 active neurons for natural scenes, static gratings, Gabor patches and flashes respectively, and 68 unique active neurons overall. We considered separated the entire duration of stimulus presentation to yield 58 trials of natural scenes, 60 trials of static gratings, 58 trials of Gabor patches, and, 3 trials of flashes, where each trial is of duration 7.5 s.

We compare TPC with two popular methods for inferring the FC from neural signals: Granger Causality (GC) and Sparse Partial correlation via Graphical Lasso penalized Maximum Likelihood Estimation (Sparse Partial Corr). The TPC algorithm was implemented with maximum time-delay of interaction 10 ms, significance level $\alpha = 0.3$, 50 bootstrap iterations, 250 ms bootstrap window width, and stability threshold $\gamma = 0.01$. For GC, we consider VAR model of order 1, and GC likelihood ratio statistic of greater than 90 percentile as indicating edges [108]. For Sparse Partial Corr, the optimal penalization was determined by 5-fold cross-validation. A summary of the results is provided in Figure 3.6.

Appendix B

Appendix to Mathematical Guarantees

B.1 Facts about the VARMA Process (4.8)

Lemma B.1.1. Suppose that the eigenvalues of F in (4.8) are of modulus strictly less than 1. Then $\{X_{ti}, X_{tj}, t \geq 1\}$ is strongly mixing at an exponential rate for all i, j , that is, $\xi_{ij}(t) = \exp(-a_{ij}t)$ for some $a_{ij} > 0$. Furthermore, the conditions $\int \|x\|g(x)dx < \infty$ and $\int |g(x + \theta) - g(x)|dx = O(\|\theta\|)$ hold for the Gaussian density g .

Proof. It follows from Theorem 3.1 in [234], that $\|\Delta_t\|_1 \rightarrow 0$ at exponential rate, where $\|\Delta_t\|_1$ is a Gastwirth and Rubin mixing coefficient of $\mathbf{X}_t, t \geq 1$. Furthermore, $\alpha_{ij}(t) \leq 4\|\Delta_t\|_1$ (See [234]), where $\alpha_{ij}(t)$ are the strong mixing coefficients of $\mathbf{X}_t, t \geq 1$. Hence, $\alpha_{ij}(t) \rightarrow 0$ as $t \rightarrow \infty$, at an exponential rate.

Now, $\int \|x\|g(x)dx < \infty$ holds trivially for the Gaussian density g . Therefore, it only remains to show that $\int |g(x) - g(x - \theta)|dx = O(\|\theta\|)$ for the Gaussian density. Towards this, note that:

$$\begin{aligned}
\int |g(x + \theta) - g(x)| dx &= \int \left| \int_0^1 \nabla g(x + z\theta) dz \cdot \theta \right| dx \\
&\leq \|\theta\| \int \int_0^1 \|\nabla g(x + z\theta)\| dz dx \\
&= \|\theta\| \int_0^1 \int \|\nabla g(x + z\theta)\| dx dz \\
&= \|\theta\| \int_0^1 \int \|\nabla g(x)\| dx dz \\
&\text{by change of variable from } x + z\theta \text{ to } x \\
&\lesssim \|\theta\| \int_0^1 \sup_i \mathbb{E}_{\mathbf{Z} \sim N(\boldsymbol{\mu}, \Sigma)} |Z_i| dz = O(\|\theta\|). \tag{B.1}
\end{aligned}$$

where $a \lesssim b$ denotes that $a \leq Cb$ for some universal constant $C > 0$.

□

Lemma B.1.2. If \mathbf{X}_t satisfies (A.1)-(A.3) (or (A.1)* - (A.3)*), then so does χ_t .

Proof. We only address the case when \mathbf{X}_t satisfies (A.1)-(A.3), as the case when \mathbf{X}_t satisfies (A.1)*-(A.3)* can be proved exactly similarly. For $u \in \{i, j\}$, let q_u and r_u denote the quotient and remainder (respectively) on dividing u by p , with the slightly different convention of redefining r_u to be p if $r_u = 0$. Note that $\chi_{tu} = X_{(t-1)r+q_u+1(r_u \neq p), r_u}$. Hence, if we define \mathcal{G}_a^b as the σ -field generated by the random variables $\{\chi_{si}, \chi_{sj} : a \leq s \leq b\}$, then assuming $i \leq j$ without loss of generality, we have:

$$\mathcal{G}_1^l \subseteq \mathcal{F}_1^{(l-1)r+q_j+1}(r_i, r_j) \quad \text{and} \quad \mathcal{G}_{l+k}^\infty \subseteq \mathcal{F}_{(l+k-1)r+q_i}^\infty(r_i, r_j),$$

where $\mathcal{F}_a^b(u, v)$ denotes the σ -field generated by the random variables $\{X_{tu}, X_{tv} : a \leq t \leq b\}$. Hence, denoting $\tilde{\xi}_{ij}(k)$ to be the maximal correlation coefficients for the process $\{\chi_{ti}, \chi_{tj} : t = 1, 2, \dots\}$, we have:

$$\tilde{\xi}_{ij}(k) \leq \xi_{r_i, r_j}(kr - (q_j - q_i) - 1)$$

for all $k > (q_j - q_i + 1)/r$. Assumptions (A.1) and (A.2) for the process $\{\chi_t\}$ now follow immediately. For (A.3), note that if $n^{1/2} \xi_{r_i, r_j}(s_n) \rightarrow 0$ for some $s_n = o(n^{1/2})$, then $t_n := r^{-1}(s_n + (q_j - q_i) + 1)$ satisfies

$t_n = o(n^{1/2})$, and $n^{1/2}\tilde{\xi}_{ij}(t_n) \rightarrow 0$, thereby verifying assumption (A.3) for the process $\{\chi_t\}$. \square

B.2 Simulation Study Details

We study the following simulation paradigms.

1. Linear Gaussian Vector Auto-Regressive (VAR) Model (Figure 4.2a left-column). Let $N(0, \eta^2)$ denote a normal random variable with mean 0 and standard deviation η . We define X_{tv} as a linear Gaussian VAR for $v = 1, \dots, 4$ and $t = 1, 2, \dots, 1000$, whose true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Let $X_{0v} \sim N(0, \eta^2)$ for $v = 1, \dots, 4$ and for $t \geq 1$,

$$\begin{aligned} X_{t1} &= 1 + \epsilon_{t1}, & X_{t2} &= -1 + \epsilon_{t2}, \\ X_{t3} &= 2X_{(t-1)1} + X_{(t-1)2} + \epsilon_{t3}, & X_{t4} &= 2X_{(t-1)3} + \epsilon_{t4}. \end{aligned}$$

where $\epsilon_t \sim N(0, \eta^2)$. It follows that the Rolled Markov Graph with respect to χ_t has edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. We obtain 25 simulations of the entire time series each for different noise levels $\eta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5\}$.

2. Non-linear Non-Gaussian VAR Model (Figure (4.2a) 2nd left-column). Let $U(0, \eta)$ denote a *Uniformly* distributed random variable on the interval $(0, \eta)$. We define X_{tv} as a non-linear non-Gaussian VAR for $v = 1, \dots, 4$ and for $t = 1, 2, \dots, 1000$, whose true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Let $X_{0v} \sim U(0, \eta)$ for $v = 1, \dots, 4$ and for $t \geq 1$,

$$\begin{aligned} X_{t1} &\sim U(0, \eta), & X_{t2} &\sim U(0, \eta), \\ X_{t3} &= 4 \sin(X_{(t-1)1}) + 3 \cos(X_{(t-1)2}) + U(0, \eta), & X_{t4} &= 2 \sin(X_{(t-1)3}) + U(0, \eta). \end{aligned}$$

The Rolled Markov Graph with respect to χ_t has edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. We obtain 25 simulations of the entire time series each for different noise levels $\eta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5\}$.

3. Contemporaneous Vector Auto-Regressive Moving Average (VARMA) Model (Figure (4.2a) 3rd left-column)

Let $N(0, \eta)$ denote a normal random variable with mean 0 and standard deviation η . We define $X_v(t)$ as

a linear Gaussian VAR for $v = 1, \dots, 4$ whose true CFC has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Let $X_v(0) = N(0, \eta)$ for $v = 1, \dots, 4$, and $t = 1, 2, \dots, 1000$,

$$\mathbf{X}_t = \begin{pmatrix} 1 \\ -1 \\ 1 \\ 2 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \end{pmatrix} \mathbf{X}_{t-1} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \end{pmatrix} \boldsymbol{\epsilon}_{t-1} + \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 2 & 1 & 1 & 0 \\ 2 & 1 & 1 & 1 \end{pmatrix} \boldsymbol{\epsilon}_t$$

It follows that the Rolled Markov Graph with respect to $\boldsymbol{\chi}_t$ has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4$. Furthermore, \mathbf{X}_t is faithful with respect to the same graph. We obtain 25 simulations of the entire time series each for different noise levels $\eta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5\}$.

4. Continuous Time Recurrent Neural Network (CTRNN) (Figure (4.2a) right-column). We simulate neural dynamics by Continuous Time Recurrent Neural Networks (B.2). $u_j(t)$ is the instantaneous firing rate at time t for a post-synaptic neuron j , w_{ij} is the linear coefficient to pre-synaptic neuron i 's input on the post-synaptic neuron j , $I_j(t)$ is the input current on neuron j at time t , τ_j is the time constant of the post-synaptic neuron j , with i, j being indices for neurons with m being the total number of neurons. Such a model is typically used to simulate neurons as firing rate units,

$$\tau_j \frac{du_j(t)}{dt} = -u_j(t) + \sum_{i=1}^m w_{ij} \sigma(u_i(t)) + I_j(t), \quad j = 1, \dots, m. \quad (\text{B.2})$$

We consider a motif consisting of 4 neurons with $w_{13} = w_{23} = w_{34} = 10$ and $w_{ij} = 0$ otherwise. We also note that in Eq. B.2, activity of each neuron $u_j(t)$ depends on its own past. Therefore, the Rolled Markov Graph with respect to $\boldsymbol{\chi}_t$ has the edges $1 \rightarrow 3, 2 \rightarrow 3, 3 \rightarrow 4, 1 \rightarrow 1, 2 \rightarrow 2, 3 \rightarrow 3, 4 \rightarrow 4$. The time constant τ_i is set to 10 msec for each neuron i . We consider $I_i(t)$ to be distributed as an independent Gaussian process with mean 1 and the standard deviation η . The signals are sampled at a time gap of $e \approx 2.72$ msec for a total duration of 1000 msec. We obtain 25 simulations of the entire time series each for different noise levels $\eta \in \{0.1, 0.5, 1, 1.5, 2, 2.5, 3, 3.5\}$.

The GC graph is computed using the *Nitime* Python library, which fits an MVAR model followed by using the *GrangerAnalyzer* to compute the Granger Causality [167]. For PC, TPCS and TPCNS, the computation

is done using the *TimeAwarePC* Python library [2]. The TPCS and TPCNS algorithms estimate the Rolled Markov Graph from the signals with $\tau = 1, r = 2\tau$.

The choice of thresholds tunes the decision whether a connection exists in the estimate. For PC, TPCS and TPCNS, increasing α in conditional dependence tests increases the rate of detecting edges, but also increases the rate of detecting false positives. We consider $\alpha = 0.01, 0.05, 0.1$ for PC, TPCS and TPCNS. For GC, a likelihood ratio statistic L_{uv} is obtained for testing $A_{uv}(k) = 0$ for $k = 1, \dots, K$. An edge $u \rightarrow v$ is outputted if L_{uv} has a value greater than a threshold. We use a percentile-based threshold, and output an edge $u \rightarrow v$ if L_{uv} is greater than the $100(1 - \alpha)^{\text{th}}$ percentile of L_{ij} 's over all pairs of neurons (i, j) in the graph [108]. We consider $\alpha = 0.01, 0.05, 0.1$ which corresponds to percentile thresholds of 99%, 95%, 90%. TPCNS is conducted with 50 subsamples with window length of 50 msec and frequency cutoff for edges to be equal to 40%.

B.3 Benchmark Datasets

We use the River Runoff benchmark real dataset from *Causeme* [179; 187]. This is a real dataset that consists of time series of river runoff at different stations. This time series has a daily time resolution and only includes summer months (June-August). The physical time delay of interaction (as inferred from the river velocity) are roughly below one day, hence the dataset has contemporaneous time interactions. This dataset has 12 variables and 4600 time recordings for each variable.

The PC algorithm was implemented with p-value 0.1 for kernel-based non-linear conditional dependence tests. In river-runoff data, the TPCS algorithm was implemented with $\alpha = 0.05$, and the TPCNS algorithm was implemented with $\tau = 2$ and 4 recordings respectively, as per specification in the datasets, and $\alpha = 0.05$ for the conditional dependence tests. TPCNS was conducted with 50 subsamples with window length of 50 recordings and frequency cut-off for edges to be equal to 0.1.

Appendix C

Appendix to Causal Functional Connectivity in Alzheimer's Disease

C.1 Automated Anatomical Labeling (AAL) Atlas

The regions in the AAL atlas along with their abbreviated, short and full names are listed in Table C.1.

Table C.1 Names of regions in the AAL Atlas

No	Abbr. Name	Short Name	Full Region Name
1	PreCG_L	Precentral_L	Precentral gyrus Left
2	PreCG_R	Precentral_R	Precentral gyrus Right
3	SFG_L	Frontal_Sup_L	Superior frontal gyrus, dorsolateral Left
4	SFG_R	Frontal_Sup_R	Superior frontal gyrus, dorsolateral Right
5	SFGorb_L	Frontal_Sup_Orb_L	Superior frontal gyrus, pars orbitalis Left
6	SFGorb_R	Frontal_Sup_Orb_R	Superior frontal gyrus, pars orbitalis Right
7	MFG_L	Frontal_Mid_L	Middle frontal gyrus Left
8	MFG_R	Frontal_Mid_R	Middle frontal gyrus Right
9	MFGorb_L	Frontal_Mid_Orb_L	Middle frontal gyrus, pars orbitalis Left
10	MFGorb_R	Frontal_Mid_Orb_R	Middle frontal gyrus, pars orbitalis Right

11	IFGoperc_L	Frontal_Inf_Oper_L	Inferior frontal gyrus, opercular part Left
12	IFGoperc_R	Frontal_Inf_Oper_R	Inferior frontal gyrus, opercular part Right
13	IFGtriang_L	Frontal_Inf_Tri_L	Inferior frontal gyrus, triangular part Left
14	IFGtriang_R	Frontal_Inf_Tri_R	Inferior frontal gyrus, triangular part Right
15	IFGorb_L	Frontal_Inf_Orb_L	Inferior frontal gyrus, pars orbitalis, Left
16	IFGorb_R	Frontal_Inf_Orb_R	Inferior frontal gyrus, pars orbitalis, Right
17	ROL_L	Rolandic_Oper_L	Rolandic operculum Left
18	ROL_R	Rolandic_Oper_R	Rolandic operculum Right
19	SMA_L	Supp_Motor_Area_L	Supplementary motor area Left
20	SMA_R	Supp_Motor_Area_R	Supplementary motor area Right
21	OLF_L	Olfactory_L	Olfactory cortex Left
22	OLF_R	Olfactory_R	Olfactory cortex Right
23	SFGmedial_L	Frontal_Sup_Medial_L	Superior frontal gyrus, medial Left
24	SFGmedial_R	Frontal_Sup_Medial_R	Superior frontal gyrus, medial Right
25	SFGmedorb_L	Frontal_Med_Orb_L	Superior frontal gyrus, medial orbital Left
26	SFGmedorb_R	Frontal_Med_Orb_R	Superior frontal gyrus, medial orbital Right
27	REC_L	Rectus_L	Gyrus rectus Left
28	REC_R	Rectus_R	Gyrus rectus Right
29	INS_L	Insula_L	Insula Left
30	INS_R	Insula_R	Insula Right
31	ACC_L	Cingulum_Ant_L	Anterior cingulate & paracingulate gyri Left
32	ACC_R	Cingulum_Ant_R	Anterior cingulate & paracingulate gyri Right
33	MCC_L	Cingulum_Mid_L	Middle cingulate & paracingulate gyri Left
34	MCC_R	Cingulum_Mid_R	Middle cingulate & paracingulate gyri Right
35	PCC_L	Cingulum_Post_L	Posterior cingulate gyrus Left
36	PCC_R	Cingulum_Post_R	Posterior cingulate gyrus Right
37	HIP_L	Hippocampus_L	Hippocampus Left
38	HIP_R	Hippocampus_R	Hippocampus Right

39	PHG_L	ParaHippocampal_L	Parahippocampal gyrus Left
40	PHG_R	ParaHippocampal_R	Parahippocampal gyrus Right
41	AMYG_L	Amygdala_L	Amygdala Left
42	AMYG_R	Amygdala_R	Amygdala Right
43	CAL_L	Calcarine_L	Calcarine fissure and surrounding cortex Left
44	CAL_R	Calcarine_R	Calcarine fissure and surrounding cortex Right
45	CUN_L	Cuneus_L	Cuneus Left
46	CUN_R	Cuneus_R	Cuneus Right
47	LING_L	Lingual_L	Lingual gyrus Left
48	LING_R	Lingual_R	Lingual gyrus Right
49	SOG_L	Occipital_Sup_L	Superior occipital gyrus Left
50	SOG_R	Occipital_Sup_R	Superior occipital gyrus Right
51	MOG_L	Occipital_Mid_L	Middle occipital gyrus Left
52	MOG_R	Occipital_Mid_R	Middle occipital gyrus Right
53	IOG_L	Occipital_Inf_L	Inferior occipital gyrus Left
54	IOG_R	Occipital_Inf_R	Inferior occipital gyrus Right
55	FFG_L	Fusiform_L	Fusiform gyrus Left
56	FFG_R	Fusiform_R	Fusiform gyrus Right
57	PoCG_L	Postcentral_L	Postcentral gyrus Left
58	PoCG_R	Postcentral_R	Postcentral gyrus Right
59	SPG_L	Parietal_Sup_L	Superior parietal gyrus Left
60	SPG_R	Parietal_Sup_R	Superior parietal gyrus Right
61	IPG_L	Parietal_Inf_L	Inferior parietal gyrus, excluding supramargina...
62	IPG_R	Parietal_Inf_R	Inferior parietal gyrus, excluding supramargina...
63	SMG_L	SupraMarginal_L	SupraMarginal gyrus Left
64	SMG_R	SupraMarginal_R	SupraMarginal gyrus Right
65	ANG_L	Angular_L	Angular gyrus Left
66	ANG_R	Angular_R	Angular gyrus Right

67	PCUN_L	Precuneus_L	Precuneus Left
68	PCUN_R	Precuneus_R	Precuneus Right
69	PCL_L	Paracentral_Lobule_L	Paracentral lobule Left
70	PCL_R	Paracentral_Lobule_R	Paracentral lobule Right
71	CAU_L	Caudate_L	Caudate nucleus Left
72	CAU_R	Caudate_R	Caudate nucleus Right
73	PUT_L	Putamen_L	Lenticular nucleus, Putamen Left
74	PUT_R	Putamen_R	Lenticular nucleus, Putamen Right
75	PAL_L	Pallidum_L	Lenticular nucleus, Pallidum Left
76	PAL_R	Pallidum_R	Lenticular nucleus, Pallidum Right
77	THA_L	Thalamus_L	Thalamus Left
78	THA_R	Thalamus_R	Thalamus Right
79	HES_L	Heschl_L	Heschl's gyrus Left
80	HES_R	Heschl_R	Heschl's gyrus Right
81	STG_L	Temporal_Sup_L	Superior temporal gyrus Left
82	STG_R	Temporal_Sup_R	Superior temporal gyrus Right
83	TPOsup_L	Temporal_Pole_Sup_L	Temporal pole: superior temporal gyrus Left
84	TPOsup_R	Temporal_Pole_Sup_R	Temporal pole: superior temporal gyrus Right
85	MTG_L	Temporal_Mid_L	Middle temporal gyrus Left
86	MTG_R	Temporal_Mid_R	Middle temporal gyrus Right
87	TPOmid_L	Temporal_Pole_Mid_L	Temporal pole: middle temporal gyrus Left
88	TPOmid_R	Temporal_Pole_Mid_R	Temporal pole: middle temporal gyrus Right
89	ITG_L	Temporal_Inf_L	Inferior temporal gyrus Left
90	ITG_R	Temporal_Inf_R	Inferior temporal gyrus Right
91	CERCRU1_L	Cerebellum_Crus1_L	Crus I of cerebellar hemisphere Left
92	CERCRU1_R	Cerebellum_Crus1_R	Crus I of cerebellar hemisphere Right
93	CERCRU2_L	Cerebellum_Crus2_L	Crus II of cerebellar hemisphere Left
94	CERCRU2_R	Cerebellum_Crus2_R	Crus II of cerebellar hemisphere Right

95	CER3_L	Cerebellum_3_L	Lobule III of cerebellar hemisphere Left
96	CER3_R	Cerebellum_3_R	Lobule III of cerebellar hemisphere Right
97	CER4_5_L	Cerebellum_4_5_L	Lobule IV, V of cerebellar hemisphere Left
98	CER4_5_R	Cerebellum_4_5_R	Lobule IV, V of cerebellar hemisphere Right
99	CER6_L	Cerebellum_6_L	Lobule VI of cerebellar hemisphere Left
100	CER6_R	Cerebellum_6_R	Lobule VI of cerebellar hemisphere Right
101	CER7b_L	Cerebellum_7b_L	Lobule VIIB of cerebellar hemisphere Left
102	CER7b_R	Cerebellum_7b_R	Lobule VIIB of cerebellar hemisphere Right
103	CER8_L	Cerebellum_8_L	Lobule VIII of cerebellar hemisphere Left
104	CER8_R	Cerebellum_8_R	Lobule VIII of cerebellar hemisphere Right
105	CER9_L	Cerebellum_9_L	Lobule IX of cerebellar hemisphere Left
106	CER9_R	Cerebellum_9_R	Lobule IX of cerebellar hemisphere Right
107	CER10_L	Cerebellum_10_L	Lobule X of cerebellar hemisphere Left
108	CER10_R	Cerebellum_10_R	Lobule X of cerebellar hemisphere Right
109	VER1_2	Vermis_1_2	Lobule I, II of vermis
110	VER3	Vermis_3	Lobule III of vermis
111	VER4_5	Vermis_4_5	Lobule IV, V of vermis
112	VER6	Vermis_6	Lobule VI of vermis
113	VER7	Vermis_7	Lobule VII of vermis
114	VER8	Vermis_8	Lobule VIII of vermis
115	VER9	Vermis_9	Lobule IX of vermis
116	VER10	Vermis_10	Lobule X of vermis
