Causal Discovery of Non-Stationary Causal Graphical Models: Applications in Brain Effective Connectivity

PhD Dissertation Proposal

by

Mario Andrés De Los Santos Hernández

Doctoral Advisors:

Dr. Luis Enrique Sucar Succar, INAOE
Dr. Felipe Orihuela-Espina, University of Birmingham

Instituto Nacional de Astrofísica, Óptica y Electrónica

August, 2024
Santa María de Tonantzintla, Puebla, CP 72840

Abstract

Non-stationary processes, where statistical properties evolve, pose a significant challenge in causal discovery. Traditional models often assume stationarity, limiting their ability to uncover causal structures amid unpredictable changes. This research addresses the representation and learning of non-stationary Dynamic Causal Bayesian Networks (nsDCBNs) using statistical tests to detect changes in the properties of the signals, thereby indicating structural and parametric shifts; and updating the structure and parameters of the causal model. This research aims to be applied in modeling brain plasticity behavior, specifically effective connectivity during neurorehabilitation, using functional near-infrared spectroscopy (fNIRS) neuroimages. The nsDCBN framework integrates a dynamic representation of evolving causal relationships, accommodating both temporal and structural non-stationarity. This allows for accurate modeling processes where the parameters and the causal structure can change over time. The main contributions of this research include: i) a novel non-stationarity detection algorithm based on statistical tests, ii) an innovative model to learn nsDCBNs through causal discovery, iii) a synthetic causal data generator that emulates fNIRS signals, and iv) a model to recover and represent brain effective connectivity from fNIRS data. These advancements aim to enhance the understanding and modeling of non-stationary processes, with potential applications in improving neurorehabilitation strategies.

Keywords: Causal Discovery, non-stationary Dynamic Causal Bayesian Networks, Causal Graphical Models, fNIRS, Effective Connectivity

Contents

1 Introduction

Given a collection of random variables defined as processes, non-stationarity is a property of those variables that indicates the fluctuation of statistical moments across the sampling topology (often time or space in physical manifestations). Meanwhile, stationarity is a deceivingly simple concept to grasp as the same state of observations is, more often than not, both stationary or not at once, with the scale of consideration being the discriminant factor.

These non-stationary processes, characterized by their continuously evolving statistical properties, are valuable tools for studying various phenomena, from environmental dynamics and the spread of diseases to the intricacies of cognitive brain functions. The dynamic nature of brain processes poses challenges to our ability to understand, model, or even predict changes within these complex systems [1]. These characteristics represent an obstacle in the path of causal discovery, particularly under traditional assumptions in the state-of-art: a static structure and parameters in the causal model.

Traditional approaches facing these non-stationary scenarios often fall short, constrained by the reliance on static causal frameworks. This research proposes to face the application of neuroimage as a non-stationary system, focusing on effective connectivity that expresses the causal influences that neuronal regions exert over one another, let us say, the causal effect observed between nodes in a causal graph. Effective connectivity offers insights into the interactions and influences between brain regions. Our focus will be on neurorehabilitation, an application where injury disrupts regular connectivity patterns, leading to plastic reorganization and modulation (i.e., non-stationary) during therapy administration to patients.

Techniques exist that allow the estimation and visualization of these causal connections based on various neuroimaging modalities, such as functional magnetic resonance imaging (fMRI), electroencephalography (EEG), magnetoencephalography (MEG), positron emission tomography (PET), single-photon emission computed tomography (SPECT), and computed tomography (CT), among others [2, 3, 4, 5, 6]. This also includes optical neuroimaging modalities such as functional near-infrared

spectroscopy (fNIRS), diffuse optical tomography (DOT), and diffuse correlation spectroscopy (DCS).

Although brain networks are more complex than simple graphs, they are perhaps the most widespread mathematical objects used to represent such networks [7]. Visualizing brain connectivity as a causal graph helps understand how information flows and is processed through causal connections, shedding light on the networks of interactions underlying brain function. This thesis aims to develop a model capable of identifying the non-stationarity in the observational data and representing it through a non-stationary Dynamic Causal Bayesian Network (nsDCBN), learning the structural and parametric changes of the observed phenomena (i.e., brain-effective connectivity in the reorganization process during rehabilitation).

However, the classical Causal Bayesian Networks (CBN) framework relies on multiple assumptions, including a static causal structure [8]. On the other hand, introducing Dynamic Causal Bayesian Networks (DCBNs) allows for the analysis of temporal signals but remains a fixed structure. These principles pose difficulties in scenarios where structure and parameters can change, limiting the capacity to capture the full extent of non-stationary phenomena or, in our application, the brain plasticity adaptions noticed in neurological rehabilitation [**?**].

Problems such as the plastic aspects of the brain's effective connectivity motivate the exploration of nsDCBN. This approach captures the statistical fluctuation in time series, representing it as a probabilistic graphical model (PGM). Unlike DCBNs, which handle temporal changes within a stationary structural context, nsDCBNs are uniquely equipped to represent significant evolutions in causal relations over time from either structural changes or parametric [9]. The approach of nsDCBNs can be a convenient, dynamic representation and modeling of a process with constant stochastic changes where understanding the changes can underline aspects such as brain neurorehabilitation.

Some approaches address the problem of temporal changes within the Granger causality framework [10]. Certain methods tackle this issue by dividing a non-stationary process into temporal windows, allowing the analysis to be performed under the assumption of stationarity, observing relational changes between time points

t and $t-1$ [4]. However, other methodologies propose modifications to dominant causal discovery algorithms, focusing on temporal dynamic changes and non-linearity rather than non-stationarity, assuming that non-linearity implies non-stationarity. Techniques such as variational inference, state-space models, local linear modeling, and modified transfer entropy have been explored in this context [11, 12, 13, 14]. These methods can extract causal structures when non-linearity is detected but do not fully represent the process parametrically; on the other hand, there are methods designed explicitly for nsDCBN that cover some of the mentioned challenges but relate under the assumption of predefined change points in the time series, which in observational scenarios is not possible to know when the changes will occur [15].

Based on these observations, we propose a causal discovery approach that detects non-stationarity in observational data. This approach iteratively learns not only the causal structure but also a representation of the phenomena through probabilistic graphical models (PGMs), such as non-stationary causal Bayesian networks (nsDCBNs). Our method focuses on the online analysis of non-stationary signals, identifying significant changes that trigger the causal discovery algorithm. This algorithm then updates the causal structure to iteratively learn nsDCBNs that characterize the process under analysis, such as the brain's effective connectivity.

Although we aim to develop general algorithms, we will apply the proposed approach to data from functional Near-Infrared Spectroscopy (fNIRS), a noninvasive optical neuroimaging technique that indirectly interrogates brain function through observed hemodynamic changes. However, like other neuroimaging modalities, fNIRS is susceptible to physiological noise and artifacts that compromise signal quality [16]. We propose a preprocessing step to address this issue to enhance signal quality for more reliable causal discovery.

This research aims to develop a novel model representing causal relationships in a noisy, non-stationary environment through a nsDCBN. Our approach seeks to improve brain connectivity understanding, particularly in contexts where detecting dynamic changes is crucial.

1.1 Motivation

Non-stationary scenarios are prevalent in various fields, including economics, environmental studies, social sciences, and medicine. Despite the importance of modeling non-stationarity, current causal approaches often rely on strong assumptions such as stationarity (constant statistical properties), linearity (relationships between variables are linear and additive), and the Markov property (future states depend only on the present state, not on the sequence of preceding events) [8]. These assumptions can hinder the accurate modeling of non-stationary processes, which may exhibit changing structures, non-linear interactions, dependencies on past states, and cyclic or self-loop behaviors [17]. Such limitations are particularly evident in brain-effective connectivity, where some behaviors are simultaneously expected [18].

The application of causal discovery in neuroimaging can significantly enhance the understanding, diagnosis, and development of rehabilitation strategies for neurological disorders such as Alzheimer's disease [15] and cerebrovascular accidents, particularly strokes. The aftermath of strokes often includes motor impairments, with an incidence rate ranging from 33% to 78% [19]. Depending on the prognosis, neurorehabilitation aims to restore patient independence through motor retraining. This retraining follows a hierarchical approach: by attempting to recover pre-injury motor behavior, by compensating through teaching new motor strategies, and by substituting with habilitating devices, such as wheelchairs or walking sticks, when recovery is not feasible. However, accurately representing the brain's plasticity can be crucial for effective and efficient rehabilitation during the brain network's rewiring process.

1.2 Justification

The challenges of non-stationary data arise from the difficulty of understanding the evolving nature of data over time. Traditional strategies often model the problem within a Dynamic Causal Bayesian network (DCBN) confines, assuming changes are temporally bound rather than structurally significant [20]. This simplifies the dynamic nature of causal relationships into a more manageable form. However,

the introduction of non-stationary Dynamic Causal Bayesian Networks (nsDCBNs) marks a significant advancement in adapting to alterations in causal structures, enabling the learning process to evolve with new data insights.

Existing methodologies employing nsDCBNs often rely on assumptions restricting their applicability across a broader spectrum of disciplines [20, 9]. A common assumption is the supposition of pre-labeled temporal changes, simplifying the representation of structural evolution within the data [15]. However, this approach falls short in scenarios where the exact nature and timing of changes are unknown. This highlights the necessity for a more adaptable model that iteratively searches for and detects changes, learns their implications on the causal structure, and swiftly updates the nsDBNs using previous causal observations to accurately reflect the current state of knowledge.

1.3 Problem Statement

The research problem consists on learning and modeling non-stationary processes through non-stationary Dynamic Causal Bayesian Networks (nsDCBN). Non-stationary processes exhibit constant, stochastic, and unpredictable changes in their time series, posing significant challenges in detecting when the statistical properties of the signal change substantially. Furthermore, causal discovery must ascertain an initial causal structure that iteratively updates in response to the detected changes in the signal properties. This also involves updating the nsDCBN iteratively, using the discovered structure as prior knowledge, and modeling the joint probabilities behavior observed across the entire time series. Additionally, applying this approach to neuroimaging data, such as fNIRS signals, requires effectively removing dominant noise while maintaining the statistical power necessary to prevent noise from directly impacting the causal structure discovery.

Formally:

Let $\bar{X}_t = \{x_1(t), x_2(t), \ldots, x_n(t)\}$ be a collection of non-stationary random variables observed over time t, and let $G_t\Theta_t$ be a non-stationary Dynamic Causal Bayesian Network (nsDCBN) at time t, where Θ_t represents the Conditional Prob-

ability Tables (CPTs) and G_t the causal structure graph. The problem is to define a learning algorithm f capable of (1) establishing an initial causal structure G_1 based on the observed data $\bar{X}_1$ at time $t = 1$, $G_1 = f(\bar{X}_1)$; (2) identifying time points $T = \{t_1, t_2, \ldots, t_k\}$ where significant changes in the statistical properties of $\bar{X}_t$ occur, with k being the total number of detected change points, $T = \{t \mid \text{statistical properties of } \bar{X}_t \text{ change significantly}\}$; then, (3) for each detected change point $t_i \in T$, updating the causal structure in $G_{t_i}\Theta_{t_i}$ using the algorithm f, incorporating prior knowledge from previous structures $G_{t_{i-1}}$, $G_{t_i} = f(\bar{X}_{t_i}, G_{t_{i-1}})$; and (4) modeling the joint probabilities $P(\bar{X}_t \mid G_t, \Theta_t)$ across the entire time series, ensuring an accurate representation of the non-stationary process, $P(\bar{X}_t \mid G_t, \Theta_t) = \prod_{i=1}^{n} P(x_i(t) \mid \text{Parents}(x_i(t)), \Theta_t)$. Here, $\text{Parents}(x_i(t))$ refers to the set of variables that have a direct causal influence on $x_i(t)$ according to the structure G_t.

Moreover, applying this approach to neuroimaging data, such as fNIRS signals, involves the removal of dominant noise. Formally, this requires: (1) Noise removal, $\bar{X}_t^{clean} = h(\bar{X}_t)$, where h is a noise removal function and $\bar{X}_t^{clean}$ represents the noise-free signals obtained from $\bar{X}_t$; (2) Robustness to changing statistical properties, where the algorithm f should adapt to $\bar{X}_t^{clean}$ as t varies; (3) Accurate change point detection, $T = \{t \mid \text{statistical properties of } \bar{X}_t^{clean} \text{ change significantly}\}$; and (4) Iterative nsDCBN update, $G_{t_i} = f(\bar{X}_{t_i}^{clean}, G_{t_{i-1}})$. The algorithm must be robust enough to adapt to the changing statistical properties of the non-stationary process, accurately detect change points, and update the nsDCBN accordingly. The algorithm should maintain the statistical power to distinguish between true causal relationships and noise, particularly in the neuroimaging data. The goal is to enhance the understanding of dynamic causal relationships in non-stationary systems, focusing on applications in brain-effective connectivity.

1.4 Research Questions

- *Non-Stationarity Detection*
  - Can the proposed algorithm, based on observing statistical moments and applying statistical tests, detect changes in time-series data that indicate significant structural or parameter changes in the causal model?

- *Causal Discovery*
  - Can a non-stationary Dynamics Causal Bayesian Network extract and interpret the causal relationships within time-series data iteratively adapting to non-stationary changes detected in the signal?
- *Causal Discovery in the domain Neuroimaging*
  - How can a non-stationary Dynamics Causal Bayesian Network concurrently represent the relationships in online neuroimaging samples where the causal non-stationary time series indirectly represents brain-effective connectivity?

1.5 Hypothesis

- *Non-Stationarity Detection*
  - The proposed algorithm, based on statistical moments and statistical tests, can effectively detect non-stationary changes in time-series data that indicate significant structural or parametric causal changes.
- *Causal Discovery of Non-Stationary time-series*
  - A nsDCBN can effectively extract and interpret causal relationships within time-series data by iteratively detecting non-stationary changes. This is due to its capacity to identify and adapt to structural and parametric changes in synthetic and semisynthetic data, where the ground truth is controlled and validated using the Structural Hamming Distance (SHD) as a metric.
- *Causal Discovery in the Neuroimaging domain*
  - The nsDCBN represents a causal non-stationary time series in neuroimaging that indirectly represents brain-effective connectivity. This is due to the capacity to identify and iteratively detect the non-stationary changes in the signal, where the validation is based on concurrence and expert validation of the regions of interest in the brain.

1.6 Objectives

To develop and validate a novel causal discovery model that extracts and learns a non-stationary Dynamic Causal Bayesian Network (nsDCBN) representing causal relationships in non-stationary processes. The model will be validated using synthetic, semisynthetic, and observational online neuroimaging data. The evaluation will be based on ground truth comparison using the Structural Hamming Distance (SHD) and expert concurrence validation.

1.7 Specific Objectives

1. To develop a synthetic data generator based on a causal graph representing effective connectivity, producing synthetic neuroimaging data with a controllable ground truth. This generator will support the validation of the nsDCBN learning model by simulating its structural and parametric changes over time.

2. To conduct denoise processing based on a deep learning approach on domain-specific data, such as fNIRS signals, to enhance data quality and improve causal discovery. The effectiveness of denoising will be evaluated through statistical analysis of signal-to-noise ratios before and after processing and by comparing the results with traditional methods.

3. To design and implement an algorithm that detects structural changes in time-series data, showcasing brain plasticity. The algorithm's responsiveness will be evaluated using synthetic data with predefined structural modifications, aiming for application in real-world scenarios.

4. To create a non-stationary Dynamic Causal Bayesian Network (nsDCBN) learning model capable of capturing and adapting to causal relationships detected in time-series data. The model's adaptability will be evaluated by assessing its response to structural changes uncovered by the causal discovery algorithm triggered by statistical fluctuations in the data.

5. To evaluate the performance of the proposed nsDCBN model, enhanced with the proposed denoise processing, in extracting and interpreting dynamic causal rela-

tionships in time-series data. The accuracy and reliability will be assessed using the Structural Hamming Distance (SHD) in-ground truth-controlled scenarios. The model's performance with and without denoise processing will be compared to uncover causal relationships in synthetic and semisynthetic noise scenarios.

6. To evaluate the performance of the proposed nsDCBN model using observational data enhanced with the domain-specific denoise processing to extract and model brain-effective connectivity. A concurrent expert validation and statistical analysis will assess the accuracy and reliability.

1.8 Scope and Limitations

This work addresses certain limitations and assumptions identified as of the current date with a forward-looking approach:

- The proposed Causal Discovery model adopts Pearl's Causality [8], directly integrating its axioms and specific assumptions into our approach.
- This research application domain is focused on neuroimaging, particularly using functional Near-Infrared Spectroscopy (fNIRS) to study brain effective connectivity.
- The iterative nature of the nsDCBN model and the preprocessing steps may result in high computational demands, potentially limiting real-time application or scalability to large datasets.

1.9 Expected Contributions

This research proposes developing a novel approach for Causal Discovery to learn a non-stationary Dynamic Causal Bayesian Network to represent non-stationary time-series scenarios applied to recover the brain's effective connectivity equipped with domain-specific denoise processing. Then, the expected contributions are listed:

- *A novel learning algorithm for nsDCBN*, capable of iteratively using the uncovered causal structure from the causal discovery algorithm as prior knowledge. This

model aims to characterize non-stationary processes, particularly in modeling the brain's effective connectivity.

- *An innovative non-stationarity detection algorithm* based on observing statistical moments and applying statistical tests. This algorithm can detect significant changes due to stochastic processes in non-stationary signals and indicates the trigger for causal discovery.
- *A pioneer generative deep learning architecture* capable of generating realistic fNIRS signals, where its inverse analog architecture allows the online denoising process to preserve the signal's statistical power.
- *A framework to recover and model* the brain's effective connectivity from fNIRS samples during a neuro-rehabilitation process.

1.10 Summary

This chapter introduces the proposal's foundational concepts and outlines the research questions, hypothesis and objective; aiming to establish a comprehensive background and methodological framework for the research.

Chapter 2 presents the theoretical foundations of the proposed research, followed by a summary of the relevant literature in Chapter 3. Chapter 4 details the formal structure of the research, including the methodology, structured work plan, and an outline of the expected publications plan. Chapters 5 covers the preliminary work aligned with our specific objectives, and Chapter 6 provides the final remarks.

2 Background

The following chapter overviews the research proposal's fundamental theoretical framework.

2.1 Non-Stationary Signal

A non-stationary signal is one whose statistical properties change over time. Unlike stationary signals, where statistical properties such as mean, variance, and auto-correlation are invariant over time, non-stationary signals exhibit time-dependent behavior.

Let $X(t)$ be a time series or signal. For $X(t)$ to be considered non-stationary, at least one of its statistical properties must be a function of time; this can be formally expressed as follows:

1. **Mean**: The mean of $X(t)$, denoted by $\mu_X(t)$, is time-dependent:
$$\mu_X(t) = \mathbb{E}[X(t)] \quad \text{such that} \quad \frac{d\mu_X(t)}{dt} \neq 0. \tag{1}$$
2. **Variance**: The variance of $X(t)$, denoted by $\sigma_X^2(t)$, is time-dependent:
$$\sigma_X^2(t) = \mathbb{V}[X(t)] = \mathbb{E}[(X(t) - \mu_X(t))^2] \quad \text{such that} \quad \frac{d\sigma_X^2(t)}{dt} \neq 0. \tag{2}$$
3. **Skewness**: The skewness of $X(t)$, denoted by $\gamma_X(t)$, is time-dependent:
$$\gamma_X(t) = \frac{\mathbb{E}[(X(t) - \mu_X(t))^3]}{\sigma_X^3(t)} \quad \text{such that} \quad \frac{d\gamma_X(t)}{dt} \neq 0. \tag{3}$$
4. **Kurtosis**: The kurtosis of $X(t)$, denoted by $\kappa_X(t)$, is time-dependent:
$$\kappa_X(t) = \frac{\mathbb{E}[(X(t) - \mu_X(t))^4]}{\sigma_X^4(t)} \quad \text{such that} \quad \frac{d\kappa_X(t)}{dt} \neq 0. \tag{4}$$

Consider a non-stationary signal $X(t)$ where the mean evolves. For example, a sine wave with a frequency that changes linearly with time:

$$X(t) = A\sin(2\pi f(t)t + \phi), \tag{5}$$

Where A is the amplitude, ϕ is the phase, and $f(t)$ is a time-varying frequency, such as $f(t) = f_0 + \alpha t$ with f_0 being the initial frequency and α a constant rate of change of the frequency. Here, the frequency component $f(t)$ introduces non-stationarity into the signal.

Non-stationary signals are common in real-world applications, such as financial time series, physiological signals (i.e., neuroimaging), and environmental data, where underlying processes and external influences cause time-dependent changes in statistical properties.

2.2 Probabilistic Graphical Models

A Probabilistic Graphical Model (PGM) is defined as a pair (G, Ω), where $G = (W, E)$ represents the graphical structure of the model, and $\Omega = \omega_v$ is a set of local functions ω_v that define the distribution parameters [21]. The joint probability distribution implied by the PGM is calculated as the product of the local functions:

$$P(V_1, V_2, \ldots, V_n) = \prod_{i=1}^{n} \omega_{V_i} \tag{6}$$

PGMs can be classified according to three dimensions:

- Directed or Undirected:
  - Directed Graphs: Represent significant directional relationships between variables, indicating causal or conditional dependencies.
  - Undirected Graphs: Represent symmetric relationships where the direction of influence is not specified, focusing on mutual dependencies.
- Static or Dynamic:
  - Static Models: Capture the relationships between variables at a specific time, assuming no change in variables or their relationships over time.
  - Dynamic Models: Represent variables and their dependencies over time, capturing relationship changes and evolutions.
- Probabilistic or Decisional:
  - Probabilistic Models: Include only random variables, representing uncertainties in relationships between variables.
  - Decisional Models: Incorporate decision, utility, and random variables, focusing on decision-making under uncertainty.

Learning and inference are the primary aims when using probabilistic graphical models (PGMs). Learning involves estimating the structure, E, and parameters, ω, of the model from an observational dataset over a set of variables, V. Inference involves answering probabilistic queries by obtaining conditional or marginal

probability distributions for a subset of variables. Bayesian networks are a type of PGMs.

2.2.1 Bayesian Networks

A Bayesian network (BN) is a directed acyclic graph that represents conditional dependencies and independencies between random variables [21]. The set $V = \{V_1, V_2, \ldots, V_n\}$ contains all the variables in the joint probability distribution, and each directed edge $(V_i, V_j) \in E$ indicates a conditional dependency between V_i and V_j. The joint probability distribution of a BN is:

$$P(V_1, V_2, \ldots, V_n) = \prod_{i=1}^{n} P(V_i|\text{pa}(V_i)) \tag{7}$$

Here:

- V_i is the i-th variable.
- $\text{pa}(V_i)$ denotes the parent variables of V_i.
- $P(V_i|\text{pa}(V_i))$ is the conditional probability of V_i given its parents.

Structure learning determines the BN's topology from data (see Figure 1), with algorithm complexity growing exponentially with the number of variables [21]. Methods are grouped into:

1. **Global Methods**: Heuristic searches over network structures, using scores like BIC [22] and MDL [23] to find the best fit.
2. **Local Methods**: Sequential evaluation of independent relationships between subsets of variables.

2.3 Causality

A causal system is defined by its dynamics and governed by causal relationships. The study of causality involves several interconnected tasks: uncovering the underlying

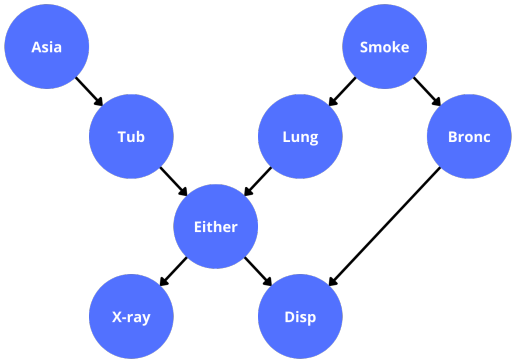

Figure 1: A Bayesian Network representing probabilistic relationships between respiratory conditions. For example, this network answers queries like the probability of lung disease given a history of tuberculosis and smoking: P(Lung | Tub, Smoke).

structure of a causal system to identify how different components influence one another, making inferences, and testing counterfactuals based on the established causal structure. These steps allow for predictions and understanding of potential outcomes under varying conditions.

Causality has been approached from different perspectives across various domains, starting with philosophical interpretations and evolving into more mathematical frameworks by the mid-twentieth century. The concept of causality lacks a single operational definition; instead, each approach uses specific mathematical properties to develop theories and derive causal relationships.

In time series analysis, causality is initially considered concerning time (or, more precisely, order): "A time series X is called causal to a second-time series Y if knowledge about the past of X and Y together allows one to predict the future of Y better than knowledge about the past of Y alone." This introduces the concept of temporal precedence as a key element in causality, though it did not account for the third variable context, whether observed or hidden. Later, Granger formalizes by incorporating context and stating that "Y is causing $X(Y \rightarrow X)$ if we are better able to predict X using all available information (Z) than if the information apart from Y had been used," which is known as Wiener-Granger causality [24].

From probability and statistics, Hume stated that causation concerns experiences rather than facts. He posited that we cannot empirically demonstrate that a cause produces an effect, only that events referred to as causes are followed by events called effects [25]. Suppes extended Hume's approach, suggesting that one event causes another if the first event's appearance is followed by the second with high probability, and no third event can factor out this relationship. Specifically, C is a genuine cause of E if $P(E|C) > P(E)$ and not $(P(E|C, D) = P(E|D)$ and $P(E|C, D) \geq P(E|C))$, indicating a spurious cause [26].

Holland framed the causal inference problem as the difference between the response of a unit u when exposed to treatment t versus control c, represented by $Y_t(u) - Y_c(u)$. He noted the Fundamental Problem of Causal Inference: it is impossible to observe $Y_t(u)$ and $Y_c(u)$ on the same unit, thus making it impossible to observe the effect of t on u directly [27].

Spirtes, Glymour, and Scheines, described causation as a relationship between particular events: an event A causes an event B. While A can have multiple causes, none alone suffices to produce A. Causation is usually seen as transitive, irreflexive, and antisymmetric: if A causes B and B causes C, then A causes C; an event cannot cause itself; and if A causes B, then B cannot cause A [8].

Pearl introduced the causal discovery problem as an induction game where scientists play against Nature, which possesses stable causal mechanisms that are deterministic functional relationships between variables, some of which are unobservable [8].

2.4 Probabilistic Causality

Two major schemes emerge when studying causality: learning causal relationships (structure discovery) and using these relationships for inference. Identifying causal relationships typically involves controlling potential external factors or assuming they don't have a direct effect, then perturbing the suspected cause variable and measuring the effect on the dependent variable, a process known as intervention. This method requires a complete causal model, which physical or ethical considera-

tions may constrain. Once established, this model can predict outcomes (inference) or answer hypothetical questions (counterfactuals) [8].

This research relies on the probabilistic definition of causality under Pearl's framework [8]. Probabilistic causality uses probability theory to obtain causal information and make causal inferences, assessing whether knowledge about variable A increases the likelihood of predicting the outcome of variable B [28].

2.5 Causal Graphical Models

Probabilistic Graphical Models (PGMs), such as Bayesian networks, are limited to encoding statistical associations and cannot express causality [8]. To be able to represent a causal interpretation, the relation implied by $X \rightarrow Y$ must extended to signify that X is a cause of Y. Causal graphical models provide a mathematical framework that supports this causal semantics.

A causal model consists of a triplet $M = \langle U, V, F \rangle$. Here, U represents a set of background variables, known as exogenous variables, determined by factors outside the model. V is a set $\{V_1, V_2, \ldots, V_n\}$ of endogenous variables, determined by the variables within the model, including those in $U \cup V$. The set F comprises functions $\{f_1, f_2, \ldots, f_n\}$, where each function f_i maps from a subset of U and the parents of V_i (PA_i) to V_i. Essentially, each f_i in $V_i = f_i(PA_i, U_i)$, for $i = 1, 2, \ldots, n$, assigns a value to V_i based on selected variables in $V \cup U$, and the entire set F has a unique solution $V(u)$ [8].

A probabilistic causal model extends this concept and is defined as a pair $\langle M, P(u) \rangle$, where M is a causal model and $P(u)$ is a probability function over the domain of U. These models must represent sudden changes in the system, indicating disturbances in variable behavior due to external factors, regardless of prior probability distributions. An appropriate semantic can represent these changes, addressing various causal questions. The most common way to represent these changes is through interventions.

Interventions in a graphical model are performed using the do() operator [8]. An intervention on X, denoted by $do(X = x)$ or $\hat{x}$, involves removing the equation

$X_i = f_i(PA_i, U_i)$ from the model and substituting $X_i = x_i$ in the remaining equations. The post-intervention joint distribution P' over $X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n$ given $\hat{x}_i$ is given by:

$$P'(x_1, x_2, \ldots, x_n|\hat{x}_i) = \begin{cases} \prod_{j \neq i} P(x_j|PA_j) & \text{if } x_i = x_i', \\ 0 & \text{if } x_i \neq x_i'. \end{cases} \tag{8}$$

Given the intervention on a variable X, denoted by $do(X = x)$ or $\hat{x}$, the probability of a variable Y is expressed as $P(Y|do(X = x))$, or $P(Y|\hat{x})$. Intervening on X_i means isolating X_i from its direct causes (parents), moving X_i to take a specific value, e.g., $X_i = x_i$, and setting the probability of X_i having that value to 1.0. In other words, when $do(X_i = x_i)$ is performed, it (i) removes the links from all of its parents, (ii) changes X_i to a new value, and (iii) assigns a probability of 1 to this event.

An intervention on any specific variable leads to a causal effect as a result [8, 29]. A causal effect is the observed response of one system element given the intervention on another. The causal effect of X on Y, denoted as $P(y|do(x))$, is a function from X to the space of the probability distribution of Y. For each intervention of X, $P(y|do(x))$ provides the probability of $Y = y$ induced by removing from every function in F all other influences on the variables in X and substituting $X = x$ in the remaining equations.

2.5.1 Causal Bayesian Networks (CBN)

A Causal Bayesian Network (CBN) is the extension of a Bayesian network with a causal interpretation. A CBN denoted as C, is a pair (G, ω), where G is a Directed Acyclic Graph (DAG) and ω represents the parameters, as in a classical Bayesian network. In this context, each edge $(V_i, V_j) \in E$ signifies that V_i is a cause of V_j, aligning with the semantics of a causal model.

Let us consider a probability distribution $P(V = v)$ over a set of variables V and an interventional distribution $P(V = v \mid do(X = x))$, which sets a subset X of variables to a constant x. Denote by $P^\star$ the set of all interventional distributions

$P(V = v \mid do(X = x))$ for $X \subseteq V$, including $P(V = v)$ representing no intervention (i.e., $X = \emptyset$). A DAG G represents a causal Bayesian network compatible with $P^\star$ if the following conditions hold for every distribution in $P^\star$:

1. $P(V = v \mid do(X = x))$ is Markov compatible with $X = x$
2. $P(V_i = v_i \mid do(X = x)) = 1$ for all $V_i \in X$ whenever v_i is consistent with $X = x$
3. $P(V_i = v_i \mid do(X = x), \mathrm{pa}_i) = P(V_i = v_i \mid \mathrm{pa}_i)$ for all $V_i \notin X$ whenever pa_i is consistent with $X = x$, meaning each $P(V_i = v_i \mid \mathrm{pa}_i)$ remains invariant to interventions not involving V_i.

Beyond the interpretation of directed edges, the main distinction between a Bayesian network (BN) and a CBN is that a CBN must be compatible with all probability distributions resulting from interventions on a subset of variables [8]. In the literature, the notation can be ambiguous, which often does not directly differentiate between a Bayesian network $B = (G, \omega)$ and a Causal Bayesian Network $C = (G, \omega)$ unless explicitly stated.

2.5.2 Dynamic Causal Bayesian Networks (DCBN)

Dynamic Causal Bayesian Networks (DCBNs) extend Causal Bayesian Networks (BNs) to model time-series data [8]. A DCBN combines Dynamic Bayesian Networks (DBNs) and Causal Bayesian Networks (CBNs) to represent temporal dynamics and causal relationships [30].

In a DCBN, let $V = \{V_1, V_2, \ldots, V_n\}$ be a set of random variables observed over discrete time points $t = 1, 2, \ldots, T$. The goal is to model the joint probability distribution of these variables while capturing their temporal and causal dependencies [30]. Formally, a DCBN consists of an initial network structure $G_0 = (V_0, E_0)$ and a transition network structure $G_t = (V_t, E_t)$ that repeats for all $t \geq 1$. The network structure and parameters are assumed to be stationary, meaning they do not change over time [21].

The initial network G_0 captures the dependencies among variables at the starting time point, while the transition network G_t describes how variables at time t depend on variables at time $t-1$. This is typically represented as a first-order Markov

model, where each variable at time t depends only on the variables at time $t-1$ and potentially on other variables at time t.

Each $G_t = (V_t, E_t)$ represents the network structure at time t, where V_t is the set of nodes (variables) and E_t is the set of directed edges indicating causal relationships. The CPD θ_t specifies the conditional probabilities for each variable given its parents in G_t:

$$P(V_t \mid \mathrm{Pa}(V_t)) = \prod_{i=1}^{n} P(V_{i,t} \mid \mathrm{Pa}(V_{i,t})) \tag{9}$$

Where $\mathrm{Pa}(V_{i,t})$ denotes the parent set of $V_{i,t}$ in G_t. The joint probability distribution of the entire time-series data is:

$$P(V_1, V_2, \ldots, V_T) = P(V_1) \prod_{t=2}^{T} \prod_{i=1}^{n} P(V_{i,t} \mid \mathrm{Pa}(V_{i,t-1})) \tag{10}$$

The causal interpretation in DCBNs extends static causal relationships in CBNs to the temporal domain [31], allowing inference of both contemporaneous and temporal causal effects [8].

Learning the structure and parameters of a DCBN involves estimating G_0, G_t, and θ_t using algorithms like Dynamic Hill Climbing (DHC) or Dynamic Programming (DP), optimizing a scoring function such as BIC or MDL [31].

2.5.3 Non-Stationary Dynamic Bayesian Networks (nsDBN)

Traditional DCBNs assume the underlying data-generation process is stationary, meaning the conditional dependencies and structure between variables do not change over time. However, this assumption often only applies in real-world scenarios where the network structure and parameters do not change. To address this, Non-Stationary Dynamic Bayesian Networks (nsDBNs) accommodate changes in the network structure and parameters over time [9].

An nsDBN is a Dynamic Bayesian network where the conditional dependence

changes at different time points. Formally, let D be a multivariate time-series dataset consisting of n variables observed at N discrete time points. The goal is to identify a sequence of network structures $\{G_1, G_2, \ldots, G_m\}$ that describe the conditional dependencies in the data across m epochs, with transitions occurring at times $T = \{t_1, t_2, \ldots, t_{m-1}\}$.

Each network G_i represents the structure during the i-th epoch, and the changes between successive networks are captured by the set of edge modifications Δg_i. The number of edge changes between epochs i and $i+1$ is denoted by S_i. The marginal likelihood of the observed data given the sequence of network structures and transition times is:

$$P(D \mid G_1, \Delta g_1, \ldots, \Delta g_{m-1}, T) = \prod_{i=1}^{n} \prod_{h=1}^{p_i} \prod_{j=1}^{q_{ih}} \frac{\Gamma(\alpha_{ij}(I_h))}{\Gamma(\alpha_{ij}(I_h) + N_{ij}(I_h))} \prod_{k=1}^{r_i} \frac{\Gamma(\alpha_{ijk}(I_h) + N_{ijk}(I_h))}{\Gamma(\alpha_{ijk}(I_h))} \tag{11}$$

Where I_h denotes the interval during which the parent set π_{ih} is operative, N_{ij} and N_{ijk} are the counts of occurrences, and α_{ij} and α_{ijk} are Dirichlet hyperparameters adjusted for the interval lengths.

The posterior probability of the nsDBN structure, given the data and transition times, incorporates prior knowledge about the network's evolution. Then, use exponential priors on the total number of edge changes $s = \sum_i S_i$ and the number of epochs m with rates λ_s and λ_m, respectively. The posterior distribution is:

$$P(G_1, \Delta g_1, \ldots, \Delta g_{m-1}, T \mid D) \propto P(D \mid G_1, \Delta g_1, \ldots, \Delta g_{m-1}, T)\, e^{-\lambda_s s}\, e^{-\lambda_m m} \tag{12}$$

To efficiently sample from this posterior, the Markov Chain Monte Carlo (MCMC) approach allows us to handle the complexity of the large state space [20].

2.5.4 Non-Stationary Dynamic Causal Bayesian Networks (nsDCBN)

Non-stationary dynamic Causal Bayesian Networks extend nsDBNs by explicitly modeling the causal relationships among variables and their evolution over time [8, 9].

An nsDCBN is a dynamic causal Bayesian network where the structure and conditional dependencies can change at different time points. Let D be a multivariate time-series dataset consisting of n variables observed at N discrete time points. The goal is to identify a sequence of causal network structures $\{G_1, G_2, \ldots, G_m\}$ that describe the causal dependencies in the data across m epochs, with transitions occurring at times $T = \{t_1, t_2, \ldots, t_{m-1}\}$.

Each network G_i in nsDCBN represents the causal structure during the i-th epoch, and the changes between successive networks are captured by the set of edge modifications Δg_i. The difference from nsDBNs is the explicit modeling of causal relationships, where each edge $(V_i \to V_j) \in G_i$ signifies a causal influence of V_i on V_j.

The posterior probability of the nsDCBN structure, given the data and transition times, incorporates prior knowledge about the network's evolution. The priors on the total number of edge changes $s = \sum_i S_i$ and the number of epochs m with rates λ_s and λ_m, respectively, remain the same as in nsDBNs. The posterior distribution is:

$$P(G_1, \Delta g_1, \ldots, \Delta g_{m-1}, T \mid D) \propto P(D \mid G_1, \Delta g_1, \ldots, \Delta g_{m-1}, T)\, e^{-\lambda_s s}\, e^{-\lambda_m m} \quad (13)$$

By capturing both temporal and causal dependencies while accounting for non-stationarity, nsDCBNs provide a framework for understanding complex, time-varying systems [20]. This framework is suited for applications with unknown underlying evolving processes, such as gene regulatory networks, neural connectivity studies, and financial market analysis. By explicitly modeling causal relationships and their evolution, nsDCBNs enable a more accurate and insightful analysis of dynamic systems.

2.6 Causal Discovery

Without direct interventions, learning causal models from observational data presents significant challenges. Identifying a dependency between two variables, X and Y, does not reveal whether X causes Y, Y causes X, or if an unobserved variable

Z influences both [21]. For instance, observing that individuals who drink wine have fewer heart attacks might suggest a causal effect of wine consumption on heart attack reduction. However, income level could confound this relationship, where higher income leads to increased wine consumption and better healthcare access, thus reducing heart attack incidence [21].

Reichenbach's common cause principle formalizes this by stating that if two variables, X and Y, are statistically dependent, a third variable, Z, causally influences both. This principle necessitates including all relevant factors to model causal relationships accurately.

Determining a unique causal structure from observational data alone is generally impossible. Instead, we derive a Markov equivalence class (MEC), which consists of all directed acyclic graphs (DAGs) that encode the same set of conditional independence relationships. For example, given three variables X, Y, and Z with the structure $X - Y - Z$, multiple DAGs can represent this structure, forming different MECs [21].

A Markov equivalence class includes all graphs that share the following properties [8]: (1) Identical Skeleton: The underlying undirected graph is the same across all graphs in the MEC; and (2) Consistent V-Structures: Subgraphs of the form $X \rightarrow Y \leftarrow Z$ are the same, with no direct edge between X and Z.

When causal sufficiency is not assumed, the models are represented as maximal ancestral graphs (MAGs) and their equivalence classes as partial ancestral graphs.

To infer the structure of a causal model from observational data, certain assumptions are required [8]:

- **Causal Markov Condition**: A variable is independent of its non-descendants given its direct causes.
- **Faithfulness**: No additional independencies exist between variables other than those implied by the Causal Markov Condition.
- **Causal Sufficiency**: There are no unobserved common confounders of the observed variables.
- **Acyclic Structure**: The model is free from cycles.

- **Causal Stationarity**: The causal relationships between variables remain constant over time.

These are strong assumptions for constraint-based algorithms, which rely on statistical tests and estimators influenced by sample size and data quality to construct a graphical model that encodes the joint probability distribution. Prominent algorithms for causal discovery that extend Bayesian network learning include the PC algorithm, Greedy Equivalence Search (GES), and Linear Non-Gaussian Acyclic Model (LiNGAM).

The PC algorithm systematically tests conditional independencies to construct the graph, starting with a fully connected graph and removing edges based on these tests [32]. GES uses a score-based approach, starting with an empty graph, adding edges to maximize a scoring criterion, and then removing edges to refine the graph [33]. LiNGAM identifies causal directions in non-Gaussian data by leveraging independent component analysis [34]. Each of these algorithms adapts to different assumptions and data characteristics, showcasing the diversity of approaches in causal discovery.

2.7 Brain Connectivity

Brain connectivity is the intricate network of anatomical, functional, and effective connections within the brain. It encompasses the structural links between neurons at a microscopic scale (neuroanatomical connectivity), the statistical dependencies between different brain regions (functional connectivity), and the causal interactions indicating the influence of one neural region over another (effective connectivity) [35].

Brain connectivity helps us understand how information is processed in the brain and plays a fundamental role in various aspects of neuroscience research, including neuroimaging studies, investigations into brain disorders, and exploration of brain architecture.

2.7.1 Effective Connectivity

Effective connectivity refers to the influence one neural system exerts over another at a synaptic or population level. It is dynamic and context-dependent, meaning it can change based on the state of the neural systems involved and the specific tasks or conditions being performed [35]. Effective connectivity is typically inferred using models that describe the causal relationships between neural elements, such as Dynamic Causal Modeling (DCM) [2] or Granger causality [10]. Usually, models like Dynamic Causal Modeling (DCM) investigate the causal interactions within the regions of interest (ROI).

2.7.2 Stroke

A stroke, or cerebrovascular accident, significantly impacts human health by affecting the brain's arteries. It occurs when a blood clot blocks a vessel or when a vessel ruptures, causing internal bleeding. This deprives specific brain regions of blood and oxygen, leading to localized neuron death. The extent and location of the damage are crucial in determining the stroke's effects [19]. Strokes are classified into two main types: ischemic (blocked arteries) and hemorrhagic (ruptured vessels). Post-stroke damage often results in motor neuron impairment, causing conditions like hemiparesis or hemiplegia, where patients lose voluntary movement in one side of their body [36].

2.7.3 Functional Near-Infrared Spectroscopy

fNIRS is a non-invasive optical technique that estimates brain activity by detecting changes in optical absorption through the intact scalp and skull. It examines light attenuation between 5 and 10 mW at red and near-infrared wavelengths of 650–850 nm, the "optical window" where biological tissues are nearly transparent [37]. During brain activity, increased local oxygen demand changes the concentrations of oxyhemoglobin (HbO2) and deoxyhemoglobin (HHb), altering light absorption [38].

fNIRS uses optodes placed on the scalp (Figure 2A). Emitted light traverses

tissues and differences in returning light that crossed the sample are detected (Figure 2B). Due to low absorption at the optical window, infrared light penetrates several centimeters into the tissue, reaching the cerebral cortex and returning to the surface. This allows light to enter 1.5 to 2 centimeters into the head and reach 5 to 10 millimeters into the brain tissue [39, 37]. The light penetration depth depends on the source-detector distance, generally between one-third and one-half of the gap, making the technique sensitive to changes in hemoglobin in the outer cortex.

Biological tissues are mostly forward scatterers, but after numerous scattering events, photon paths become random (diffusion regime). Only a small proportion of photons return to the scalp, completing the optical path between the emitter and detector in a "banana" shape [40] (Figure 2B).

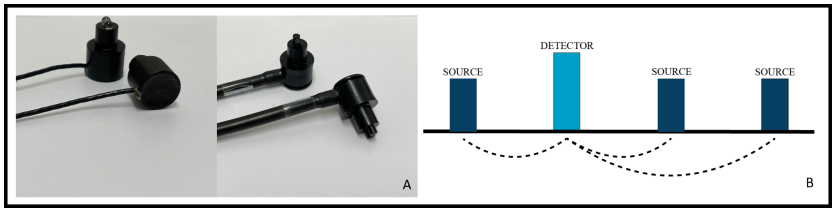

Figure 2: Left: Source and detector optodes. Right: Schematic depiction of the banana shape principle in continuous wave fNIRS.

Light and tissue interact through absorption and scattering, which are often modeled separately despite being aspects of the same extinction phenomenon. Diffraction, though physically different, is treated as a type of scattering.

2.8 Summary

This chapter establishes the theoretical foundation for the research proposal by exploring causality from philosophical, temporal, and probabilistic perspectives. We examine probabilistic graphical models (PGMs), focusing on Bayesian networks (BNs) and causal Bayesian networks (CBNs) for representing and inferring variable relationships.

The chapter highlights causal graphical models (CGMs) structure and learning processes, emphasizing distinctions between different model types and the impor-

tance of identifying and validating causal structures.

Effective connectivity in brain research is discussed to demonstrate the practical use of causal models in neuroscience. This overview provides a foundation for applying causality and causal models in the research proposal, particularly in brain connectivity studies.

3 Related work

This chapter recompiles the related work for this research; the sections presented are aligned with each component that composes the general objective of the research, starting from the non-stationary process in causal discovery, passing from the learning of nsDCBN, detection of non-stationarity to the particularity of denoising deep learning for our domain.

3.1 Non-Stationary Processes and Causal Discovery

The study of causal discovery from time series has significantly evolved. This approach identifies causal relationships where underlying distributions change over time. This approach has shown promising results in various fields, such as stock price prediction and social analysis. In these areas, dynamic changes are often due to external interventions or intrinsic process changes, rendering the time series data non-stationary.

Several frameworks have been adapted for non-stationary scenarios. State-space models, for example, adjust to changes in structure and noise variance, enhancing forecasting accuracy and incorporating previous causal knowledge [12]. However, they can become overly complex with multiple time-varying parameters and may struggle with rapid changes or limited observations [12, 41]. Additionally, these models often usually assume distributed processes, which may not always be the case [42].

The Just-in-Time (JIT) modeling framework addresses non-stationary and

non-linear time-series data using a local linear structural causal model that adapts to new data inputs [13]. This approach enhances interpretability and reduces computational complexity but may be inaccurate for highly non-linear observations and dependent on data availability [42].

The Constraint-based Causal Discovery from Non-stationary/Heterogeneous Data (CD-NOD) method detects changes in local causal mechanisms and utilizes shifts in data distribution to determine causal orientations [17]. While effective, it is sensitive to noise and assumes certain variables are conditionally independent, which may not hold in complex systems. It also does not address changes in causal direction over time [43].

Despite the advances and diverse methodologies for causal discovery in non-stationary scenarios, challenges remain. These include handling rapid shifts in causal structures, managing noise, and ensuring robustness across different domains. Addressing these issues will improve forecasting accuracy and enhance our understanding of complex systems and their underlying causal relationships, providing valuable insights across various practical applications.

3.2 Learning Non-Stationary Dynamic Causal Bayesian Networks

The learning process of Dynamic Bayesian Network (DBN) structures provides mechanisms for identifying conditional dependencies in time-series data, assuming stationarity. This assumption often limits the models' effectiveness [44]. The introduction of non-stationary Dynamic Bayesian Networks (nsDBN) added a new class of graphical models that allow the process to change over time [9]. This approach identifies a discrete Bayesian network evolving through a piecewise stationary process where edges can be gained or lost over time. The innovation lies in allowing the network structure and parameters to evolve dynamically and using Markov Chain Monte Carlo sampling to learn the model structure from time-series data [9, 20]. This algorithm identifies linear, non-linear, and combinatorial interactions between variables in a non-stationary process [20].

However, Robinson's methods have limitations. The assumption of known

transition points is often unrealistic in practical scenarios. Furthermore, the piecewise stationarity and smooth evolution assumptions may not adequately capture abrupt structural changes. The model's sensitivity to hyperparameter selection necessitates precise tuning to avoid suboptimal inferences. Additionally, the framework restricts network evolution to single-edge changes at a time, limiting the capacity to model more complex, simultaneous changes. Validation on real biological data presents further challenges, requiring additional empirical studies to confirm the model's effectiveness in practical applications [20, 9].

The introduction of nsDCBN supported other algorithms discussed in this section, such as Flexible Lag nsDCBN (FLnsDCBNs) [45]. This approach models non-stationary time series, assuming they can be divided into multiple sections, each governed by its graph structure. Thus, different causal relationships are identified in each segment [45, 17]. The flexible lag in FLnsDCBNs helps determine the optimal time lag for each segment, enhancing the model's adaptability to changes in the data; this particular method is limited by the detection accuracy of the changes, resulting in some scenarios requiring extra detection methods. On the other hand, as the number of changes increases, the computational complexity of the model increases too significantly.

Then, similarly to FLnsDCBNs, Reversible Jump Markov Chain Monte Carlo Non-Stationary DCBNs [46] are based on the proposed work by Robinson [20]. Here, the model allows us to adapt to changes in the processing, permitting transitions between different network structures over time. This capability facilitates the exploration of the models with varying dimensions, making it particularly useful to identify the number and timing of causal structural changes in the network. However, in this case, it only captures the structural changes, not the parametric ones.

Allocation Sampler Non-Stationary DBNs (ASnsDBNs) [47] utilize an allocation sampler to learn the structure of non-stationary Bayesian networks dynamically. This method allows for identifying the number and timing of transitions between different network structures without prior knowledge of these parameters. The allocation sampler effectively allocates data segments to different models, capturing the evolving nature of the underlying processes. While ASnsDBNs perform competitively with other non-stationary modeling approaches, they may face challenges

in high-variability scenarios where the data patterns change rapidly. However, the reversible jump MCMC method can be computationally intensive, especially with larger datasets, and may require careful tuning of parameters to achieve optimal performance [47, 45].

Non-Stationary Continuous Dynamic Bayesian Networks extend the traditional DCBN framework to accommodate continuous data, allowing for the modeling of systems where both structure and parameters can change over time [48]. This approach utilizes a Bayesian multiple change-point process to identify when changes occur in the data, enabling the model to adapt to evolving relationships. This method facilitates information sharing across different periods, allowing parameters to vary among segments while maintaining a common network structure. However, the complexity of continuous data modeling can lead to challenges in the convergence and stability of the inference algorithms, and the need for continuous data may limit its applicability in discrete time series contexts.

These methods have served as references for more modern approaches with concrete applications in neuroimaging and environmental or financial analysis scenarios. However, significant limitations remain, particularly in handling abrupt structural changes and adapting to unknown transition points [44, 45].

In contrast, our proposed approach aims to overcome these limitations by developing a model capable of dynamically detecting and adapting to non-stationary changes in the data without relying on predefined change points. By iteratively updating the causal structure and parameters, our method provides a more flexible and accurate representation of the underlying processes, particularly in brain-effective connectivity during neurorehabilitation.

3.3 Detection of Non-Stationarity

Developing a non-stationarity detection algorithm is an innovative field without a dominant reference method for time series analysis. Various approaches offer unique insights and benefits, including statistical tests and non-parametric methods. Among statistical tests, several are particularly noteworthy. The Augmented Dickey-Fuller

(ADF) test is widely used to test for a unit root in a univariate time series [49]. It involves estimating a regression model and examining the significance of the lagged level of the series. If the test statistic is less than the critical value, the null hypothesis of non-stationarity is rejected, indicating that the series is stationary.

Similarly, the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test examines stationarity around a deterministic trend [50]. In this test, the null hypothesis asserts that the series is stationary. If the computed test statistic is significant, the null hypothesis is rejected, suggesting the presence of a unit root and, therefore, non-stationarity. The Phillips-Perron (PP) test, another critical method, extends the ADF test by accounting for serial correlation and heteroskedasticity in the error terms, providing a more robust measure for non-stationarity [51]. Additionally, the Variance Ratio test examines random walk behavior in time series data by comparing the variance of the series over different time intervals [52]. Significant deviations from the expected variance under the random walk hypothesis indicate non-stationarity.

On the other hand, some non-parametric tests have emerged as an alternative for detecting non-stationarity in time series data. These tests do not make assumptions about the underlying distribution of the data, making them more flexible in handling complex real-world scenarios. However, they are limited in their applicability to specific types of data or processes. For instance, Kanaya (2011) proposed a non-parametric test for stationarity in continuous-time Markov processes [53]. This test constructs a kernel-based test statistic and conducts Monte-Carlo simulations to study the finite-sample size and power properties applicable only to univariate time-homogeneous Markov processes. Another example is a non-parametric test for stationarity in functional time series, suggested by van Delft et al. (2017), which is limited to functional time series data obtained by separating a continuous time record into natural consecutive intervals (e.g., days) [54]. van Delft and Eichler (2018) also proposed a test for local stationarity in functional time series.

Additionally, Basu et al. (2009) proposed a non-parametric test for stationarity based on local Fourier analysis, applicable to any zero-mean discrete-time random process by transforming any finite sample of a discrete process to have zero mean [55]. Other notable non-parametric tests include those proposed by Breitung (2002) for unit roots and cointegration [56] and wavelet packet tests for second-order sta-

tionarity developed by Cardinali and Nason (2018) [57].

In summary, statistical tests like ADF, KPSS, PP, and Variance Ratio are well-established but rely on assumptions (e.g., ADF assumes no structural breaks, KPSS assumes homoscedasticity) that may not hold in all scenarios. Non-parametric tests offer greater flexibility and can handle complex data structures without predefined parameters. However, they may be limited in their applicability to specific data or processes and might not be as broadly applicable as statistical tests.

3.4 Causal Discovery Applications in Neuroimaging

Understanding the brain's intricate connectivity patterns is a key objective in neuroscience. The advent of causal discovery methods in neuroimaging has provided tools to explore these complexities. Various approaches have been developed, each offering specific insights and facing distinct challenges.

Montero-Hernandez et al. [58] introduced a method that estimates intervals of causal effects to infer graphical models from neuroimaging data. This technique addresses the uncertainty in causal effect estimation, a common issue in neuroimaging due to the noisy and variable nature of brain signals. The method involves estimating causal effects using a framework that computes confidence intervals for these estimates, incorporated into the graphical model learning process. This approach assumes linear relationships between variables, which may limit its ability to capture the non-linear dynamics often present in brain connectivity. Furthermore, the method assumes that the data is stationary, meaning that the statistical properties of the brain signals do not change over time. This is a strong assumption, considering that the brain's connectivity patterns may evolve, potentially leading to inaccurate causal inferences. Additionally, the method's effectiveness can be hindered by unmeasured confounders, which are common in neuroimaging data and can introduce bias into the causal effect estimates.

Sanchez-Romero et al. [5] developed causally informed activity flow models that integrate causal discovery techniques to enhance the interpretation of brain activity. This method identifies directional influences between neural activities and

incorporates these causal directions into activity flow models, improving their explanatory power. This approach has shown improved accuracy in modeling brain activity and provides a mechanistic understanding of brain function. However, it relies on several assumptions that may limit its applicability to non-stationary processes. For instance, the method assumes acyclicity in the underlying network and has no unmeasured confounders; additionally, it assumes the stationarity of the BOLD signals, meaning the statistical properties of these signals do not change over time. While they used strategies like linear detrending to address temporal trends, the method is primarily designed for stationary processes. Furthermore, the focus on resting-state data and the limited temporal resolution of fMRI constrain its ability to capture dynamic changes over shorter timescales, which is characteristic of non-stationary processes.

Liu et al. [15] employed dynamic causal modeling (DCM) to estimate effective connectivity from fMRI data. This method uses Bayesian inference to estimate model parameters and their uncertainties and model comparison techniques to identify the best-fitting connectivity model. DCM captures the dynamic interactions between brain regions and provides detailed insights into their causal relationships. However, it is computationally intensive and requires strong priors about the network structure, which may limit its applicability in exploratory analyses. Additionally, although Liu et al. utilize non-stationary dynamic Bayesian networks (nsDBN), they operate under the strong assumption that the points in time where changes occur in the connectivity network are known *a priori*. This assumption may limit the model's flexibility and applicability in observational scenarios without predefined change points.

Saetia et al. [4] applied causal discovery algorithms tailored for neuroimaging data to construct brain connectivity models. This method uses constraint-based and score-based causal discovery algorithms and validates the inferred models against known anatomical connectivity. The study demonstrates that causal discovery can effectively identify meaningful connectivity patterns in the brain. The main limitation is the sensitivity to noise and the need for large, high-quality datasets to achieve reliable results.

Cai et al. [41] introduced a method known as the Gaussian-based Variational

Temporal Abstraction (GVTA) model, which detects stationary states within non-stationary time series and estimates causal mechanisms for each state using Gaussian processes. This method addresses the challenges of non-stationary data, providing a framework for capturing causal coefficients and direction changes. The GVTA model assumes that the data can be segmented into stationary intervals, which might not always be accurate in real-world scenarios. The GVTA model assumes that the time series data can be divided into distinct segments where each segment is stationary. Furthermore, the Gaussian process framework for estimating causal mechanisms requires that the data within each segment follow a Gaussian distribution, an assumption that might not hold in all neuroimaging datasets. While necessary for the model's theoretical framework, these assumptions can limit its flexibility and accuracy in practical applications where data may exhibit more complex, non-Gaussian, and continuously evolving characteristics.

In summary, these methodologies illustrate the varied approaches to causal discovery in neuroimaging. Montero-Hernandez et al. [58] focus on robustness through confidence intervals but may need help with non-linear dynamics and stationarity assumptions. Sanchez-Romero et al. [5] enhance activity flow models with causal directions but depend on accurate causal discovery and assume stationarity. Liu et al. [15] offer detailed dynamic modeling with DCM but at a high computational cost and rely on predefined change points. Saetia et al. [4] validate causal models against anatomical data but are noise-sensitive and rely on solid assumptions. Cai et al. [41] provide a framework for non-stationary data but face challenges in model complexity, resource requirements, and data segmentation and distribution assumptions.

3.5 Deep Learning for Denoising in Neuroimaging

Deep learning (DL) techniques for denoising neuroimaging data, especially functional near-infrared spectroscopy (fNIRS), have shown great promise in mitigating noise and artifacts. These DL models mainly address physiological noise, such as heart rate, blood pressure variations, and motion artifacts [59], which can degrade fNIRS signal quality and lead to incorrect interpretations. Common DL architectures used

for these tasks include convolutional neural networks (CNNs), autoencoders (AEs), and long short-term memory networks (LSTMs).

Several studies have shown DL's effectiveness in enhancing fNIRS data quality by removing unwanted noise. Gao et al. (2020) used a denoising autoencoder (DAE) with nine convolutional layers to tackle motion artifacts during a precision cutting surgical task [60]. Their model, trained on both simulated and real data, achieved 93% artifact removal in simulated data and 100% in real data, outperforming traditional methods like wavelet filtering and principal component analysis. Kim et al. (2022) also focused on motion artifact removal using a CNN, comparing its performance to wavelet denoising and autoregressive methods. The CNN achieved a mean square error (MSE) of approximately 0.004 to 0.005, better than the combined wavelet and autoregressive method's MSE of about 0.009 [61].

Lee et al. (2018) explored using CNNs for motion artifact removal, training their model on raw fNIRS time series and estimated canonical responses. Their model achieved a contrast-to-noise ratio (CNR) of 0.63, significantly better than wavelet denoising's CNR of 0.36, highlighting CNNs' potential in enhancing fNIRS signal quality [62]. Then, Liu et al. (2021) used an echo state network autoencoder (ESN AE) to extract features from fNIRS data, which were then classified using a multilayer perceptron (MLP). The ESN AE outperformed convolutional autoencoders (CAE) and manual feature extraction, achieving a four-class classification accuracy of 52.45%. This study underscored the importance of combining autoencoders and classifiers for improved feature extraction and denoising performance [63].

Woo et al. (2020) used a deep convolutional generative adversarial network (DCGAN) to generate clean fNIRS activation t-maps. These generated t-maps were then used to augment the CNN training dataset, significantly increasing the classification accuracy of a finger-tapping task from 92% to 97%. The DCGAN was trained on noisy data to produce clean data, showcasing the potential of generative models in denoising applications [64].

DL techniques for denoising in neuroimaging, particularly fNIRS, have significantly improved the removal of motion artifacts and physiological noise. These meth-

ods use architectures like CNNs, AEs, and GANs to enhance data quality, leading to better performance in classification and connectivity analysis tasks. Advancements in DL-based denoising are promising for real-time neuroimaging, reducing extensive preprocessing and enabling more accurate analyses [59].

3.6 Summary

This chapter explored advancements in causal discovery for non-stationary time series, underscoring their relevance in finance, social analysis, and neuroimaging. Various frameworks have been developed to manage dynamic data changes, each with strengths, assumptions, and challenges. We examined methods for learning dynamic causal relationships, highlighting the complexity of modeling non-stationary systems.

We also reviewed approaches for detecting non-stationarity in time series, comparing statistical and non-parametric methods. We discussed the application of causal discovery in neuroimaging, emphasizing the need to understand brain connectivity patterns. The chapter concluded by demonstrating how deep learning techniques have enhanced neuroimaging data quality by effectively removing noise and artifacts and improving the accuracy of subsequent analyses.

However, many existing methods face significant limitations. They often assume known transition points, rely on smooth evolutionary changes, or treat processes as stationary, which can miss abrupt structural changes and necessitate precise hyperparameter tuning. These approaches frequently limit network evolution to single-edge changes, hindering their ability to model simultaneous changes.

In contrast, our proposed approach addresses these limitations by dynamically detecting and adapting to non-stationary changes in the data without relying on predefined change points. By iteratively updating the causal structure and parameters, our method offers a more flexible and accurate representation of underlying processes, particularly for modeling brain-effective connectivity during neurorehabilitation. This contributes to a more robust understanding and analysis of dynamic causal relationships in non-stationary environments.

4 Research Proposal

This chapter outlines the methodology for the research aims. It details the methodology, activities plan, and publications plan. Figure 3 shows an overview of the research methodology.

4.1 Methodology

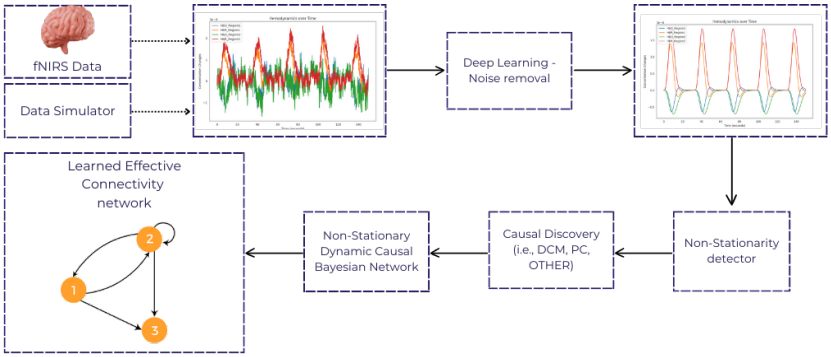

Figure 3: Methodology Block Diagram of the Research Proposal. The process begins with signal acquisition, initially from synthetic and semisynthetic data [3], then observational fNIRS data. The data undergoes denoising to remove physiological and common artifacts, such as optode movements. An iterative method detects significant parametric and structural changes, enabling causal discovery based on these detections. This iterative process allows learning a non-stationary Dynamic Causal Bayesian Network only when a change is present, using the causal structure as a skeleton.

4.1.1 Synthetic Data Generator

Synthetic data is essential for validating the proposed model, allowing us to understand model behavior before using observational data. It provides controlled structural and parametric changes, offering a ground truth. Our synthetic data generator extends previous work [2] by adding realistic noises categorized into three main groups [65]:

- Motion Artifacts: Caused by optode movement on the participant's head, lead-

ing to abrupt changes in signal intensity.

- Physiological Noise: Originating from the participant's internal state, including heart rate, breathing, vasomotion, and other systemic fluctuations [16].
- Instrumental Noise: Associated with the instrument's performance, such as fluctuations in light source intensity and detector sensitivity.

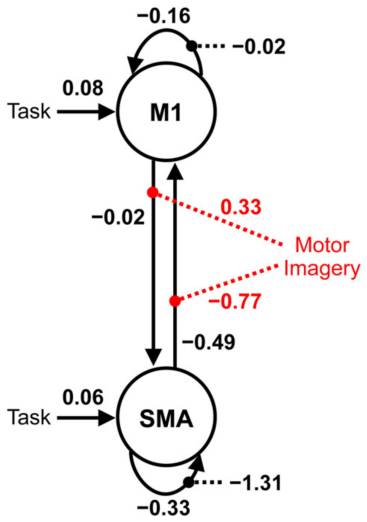

Figure 4: Causal Structure Example. A causal model was used to generate a synthetic fNIRS sample.

The generator consists of three stages: Neurodynamics, Hemodynamics, and Optics. The first two stages evoke neural responses and subsequent hemodynamics in the targeted regions, simulating blood flow changes in active regions. An extension of the original model emulates optical density changes representing fNIRS signals [3], as illustrated in Figure 5.

Our implementation is based on a causal model (Figure 4), allowing the emulation of multiple regions or nodes using a graphical causal model as a reference, which can be considered ground truth to evaluate the causal models learned from simulated data.

To evaluate the synthetic data generator, we will employ a multifaceted approach. First, we will conduct a qualitative evaluation by visually comparing the

generated data with real data and seeking expert reviews to ensure the synthetic data accurately captures essential characteristics. We will also validate the noise and artifact characteristics by comparing them to those found in real data. Additionally, we will conduct controlled experiments to compare synthetic data outputs with known ground truths and evaluate the generator's sensitivity to parameter changes.

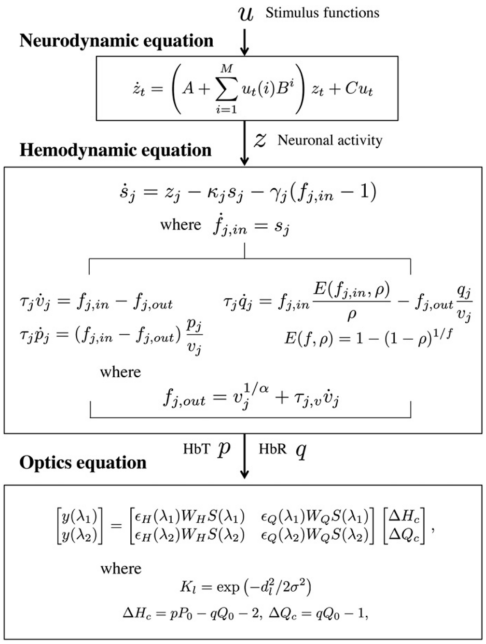

Figure 5: Schematic of the Generative Model of fNIRS Data [3]. **Neurodynamic equation**: Models neural activity $\mathbf{z_t}$ with connectivity matrix $\mathbf{A}$, experimental modulation $\mathbf{u_t(i)}$ via matrices $\mathbf{B_i}$, and input influences $\mathbf{C}$. **Hemodynamic equation**: Describes how neural activity affects blood flow ($\mathbf{s_j}$), volume ($\mathbf{v_j}$), and deoxyhemoglobin levels ($\mathbf{q_j}$) using the Balloon model, with variables $\mathbf{f_j}$ for inflow and $\tau_\mathbf{j}$ as time constants. **Optics equation**: Links hemodynamic changes to optical measurements $\mathbf{y}(\lambda)$ via sensitivity matrix $\mathbf{S}(\lambda)$, correcting for pial vein contamination with matrices $\mathbf{W_H}(\lambda)$ and $\mathbf{W_Q}(\lambda)$. Gaussian spatial smoothing kernel $\mathbf{K_I}$ generates the spatially distributed hemodynamic response.

4.1.2 Deep learning denoise process for fNIRS

fNIRS signal samples are often contaminated by various types of noise, significantly affecting the quality and interpretation of the data. Noise in these signals usually requires a preprocessing stage to mitigate its presence while minimizing the loss

of statistical power and allowing an indirect analysis of neuronal activity based on hemodynamics. Given the application's domain, this preprocessing is an initial phase of the process where we propose to develop a generative deep-learning model trained on semisynthetic and synthetic physiological noise.

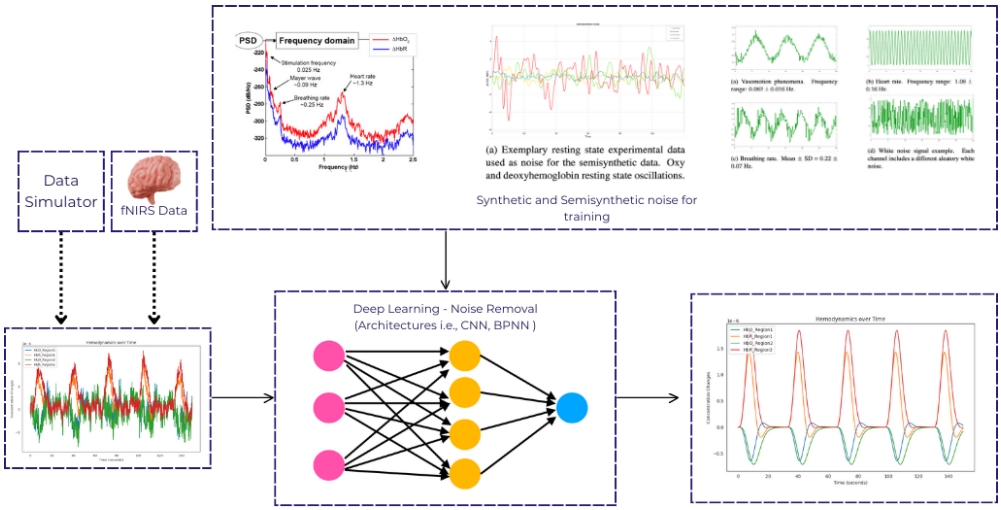

Figure 6: Deep Learning Methodology Approach: This diagram illustrates a deep learning architecture trained on synthetic and semisynthetic noise in fNIRS neuroimaging. The model denoises signals by mitigating motion artifacts and physiological and instrumental noise.

The proposed architecture aims to be a generative model of realistic fNIRS time courses and, inversely, as a denoising tool. Unlike current methods, our model, once trained, can apply an online denoising process, figure 6, whereas traditional methodologies typically rely on offline analysis.

We plan to use synthetic physiological noise generated by our data generator and previously collected semisynthetic data in a controlled group reported by [66].

The deep learning architecture in this research will serve as a denoiser and provide validation scenarios for generating synthetic data. We will validate our model's denoising performance against state-of-the-art algorithms [62, 65, 61, 60] using controlled ground truth from our synthetic data generator and observational data from healthy volunteers. This validation will involve statistical analysis, normality tests, and comparisons to reference values, as systematically reported by [65]

4.1.3 Non-Stationarity Detector

Applying causal discovery to non-stationary time series is challenging; most state-of-the-art approaches assume stationary data or treat non-stationary cases as an extension of non-temporal methods by expanding the causal graph in time with a static time lag. Directly applying traditional methods to time series can lead to misleading causal relationships.

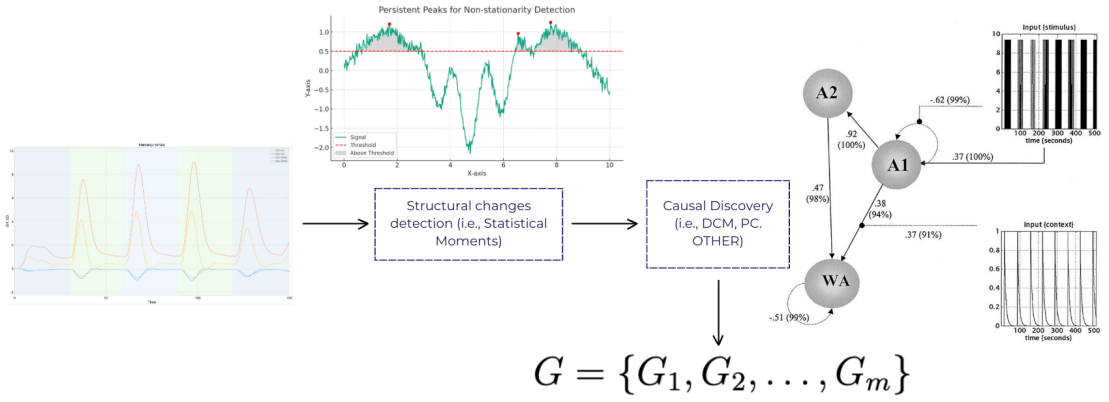

Figure 7: Non-Stationarity Detector: The process iteratively applies two modalities of timestamp windows: one with 50% overlap and another with independent, non-overlapping windows. When a change from a stationary to a non-stationary trend is detected, it triggers the causal discovery process, which generates a new causal structure using the current window and a defined number of past windows.

Detecting non-stationary changes is crucial for accurate causal discovery and iterative learning of a non-stationary Dynamic Causal Bayesian Network (nsDCBN). The detection process identifies changes in the signal's statistical moments, determines their persistence, and establishes criteria for significant changes to trigger the causal discovery and learning process.

We propose developing a non-stationarity detector algorithm, figure 7, based on the iterative time stamps windows analysis using the first four statistical moments: mean (shifts in the signal's central tendency), variance (changes in signal volatility), skewness (asymmetry variations), and kurtosis (frequency of extreme values). However, these statistical moments alone may not be sufficient for detection, so we plan

to employ several statistical methods to enhance the reliability of change detection from stationary to non-stationary trends.

We will use synthetic datasets to validate our non-stationarity detection algorithm where the ground truth is controlled. This approach allows us to verify the algorithm's accuracy in detecting non-stationary changes with labeled shifts and uncover causal structures using structural hamming distance (SHD). Additionally, cross-validation techniques and robustness tests against different changes and noise levels (i.e., synthetic and semisynthetic) and signal perturbations will be employed to ensure the algorithm's reliability further. An initial implementation of the algorithm is presented in the preliminary work, section 5.2.

4.1.4 Causal Discovery framework

Given the discovered causal structure, the learning nsDCBM proposed model aims to be capable of taking the structure as a priori structural knowledge over time to represent the non-stationary process in a probabilistic graphical model, Figure 8. The learning model will provide a mechanism to identify conditional dependencies in the non-stationary time-series process. We can observe this as an iterative process where we evolve according to the piecewise extracted from a change between a stationary piece and a non-stationary that occurs in the process over time.

Hence, while the previous algorithm detects changes, uncovering the causal structure is crucial for iterative learning of the nsDCBN. This discovery stage identifies a unique causal structure for each timeslice where non-stationary detection occurs. Each unique structure is then compared with the previous state of the causal model to update the model structure and parameters.

This proposal aims to analyze and modify some of the state-of-the-art algorithms for causal discovery in scenarios with changing causal mechanisms (structure and parametric changes), such as Constraint-based Causal Discovery from Non-stationary Data algorithm (CD-NOD) [67] or the Gaussian-based Variational Temporal Abstraction model (GVTA) [41]. We propose to analyze and, if needed, modify classical causal discovery algorithms under the assumption that we can trait a tem-

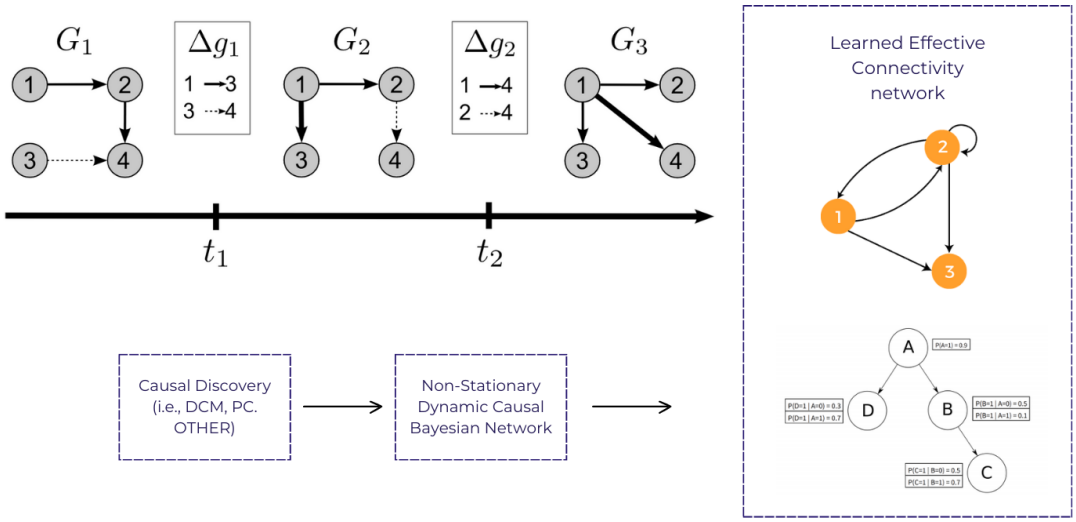

Figure 8: Learning nsDCBN. The learning process follows the causal discovery algorithm, which, once it uncovers the causal relationships within the data, allows the use of that structure as apriori knowledge given a difference between the current observation and the pass.

poral window as stationary, allowing us to iteratively treat it as non-stationary (e.g., PC [68], LiNGAM [13]).

The choice of the algorithm will be based on the specific assumptions and behavior of the algorithms in the synthetic and semisynthetic conditions we propose from our simulator, where we can control the ground truth. Given the application domain, we will analyze the implementation of Dynamic Causal Modelling (DCM) [2], a widely used model to model brain activity's inherent time-dependent nature in the application's neuroimaging domain. However, using DCM, based on Granger's Causality framework, requires a modification to rely on Pearl's Causality [35].

The proposed framework will be validated and evaluated in two stages. Initially, we will use the proposed synthetic and semisynthetic data generator to control the ground truth and directly evaluate using the structural hamming distance (SHD). In this way, we can evaluate multiple scenarios, varying the causal parameters and structure, the noise levels, and changes in sensitivity.

The second stage implies validating and evaluating observational data samples from fNIRS neuroimages, involving expected knowledge and statistical co-occurrence

tests; this test will ensure the accuracy and relevance of the discovered causal relationships based on the expert's expectations.

4.2 Performance Evaluation

The performance evaluation will be conducted in three scenarios:

- Synthetic data where the ground truth is initially a parametric non-stationary model, then structural and parametric non-stationary. The primary metric will be the normalized Structural Hamming Distance (SHD), allowing direct comparison between ground truth and inferred causal structures.
- Synthetic and semisynthetic fNIRS data with structural and parametric non-stationary conditions at different noise levels, including our proposed preprocessing model. Validation will involve improvements in signal-to-noise ratios and statistical measures to assess the denoising model against state-of-the-art algorithms. Using SHD, the causal discovery framework will be evaluated by comparing inferred structures to known ground truth.
- fNIRS trials with healthy volunteers, over observational data (resting state) and experimental data (with external stimulation, e.g., finger tapping tasks). Then, given the results in this phase and the approval of an ethical protocol, samples in clinical patients post-stroke in the same scenarios, resting state, and rehabilitation process. Observational data from fNIRS samplings will be validated using a co-occurrence statistical test, leveraging expert knowledge on the regions of interest analyzed. Results will be compared offline using Dynamic Causal Modeling (DCM) and expert validation to determine the relevance and correctness of the revealed causal correlations.

Cross-validation and robustness testing will ensure the models' generalizability and stability. K-fold cross-validation will be applied across all scenarios to evaluate performance comprehensively. Robustness tests will determine the models' response to noise levels and input parameter changes. The methodology will be applied to clinical fNIRS data from post-stroke patients, with performance evaluated similarly

to observational data and additional considerations for clinical relevance in rehabilitation outcomes. This comprehensive evaluation framework establishes the proposed methodology's robustness, reliability, and applicability in synthetic and real-world scenarios, ensuring accurate non-stationary causal discovery in fNIRS neuroimaging data.

4.3 Work Plan

Figure 9 presents the comprehensive work plan for this research project, which spans four years from August 2023 to August 2027, aligning with the duration of the doctoral program. We are currently in the initial phase of the plan. Due to the timing of the proposal presentation, there is a slight overlap in the first month of the second phase.

4.4 Publications Plan

The expected publications and their principal contributions are outlined as follows:

4.4.1 Conference Articles

- **fNIRS Conference**
  1. **Objective:** Present the synthetic data generation framework based on bilinear model [3], the model emulates typical fNIRS samples. Section 5.1 presents the implementation.

     Estimated Submission Dates: April 2024

     Status: Accepted
  2. **Objective:** Present the advances in fNIRS denoise deep learning architecture developed during the research.

     Estimated Submission Dates: April 2026

- **Causal Learning and Reasoning (CLeaR) Conference**

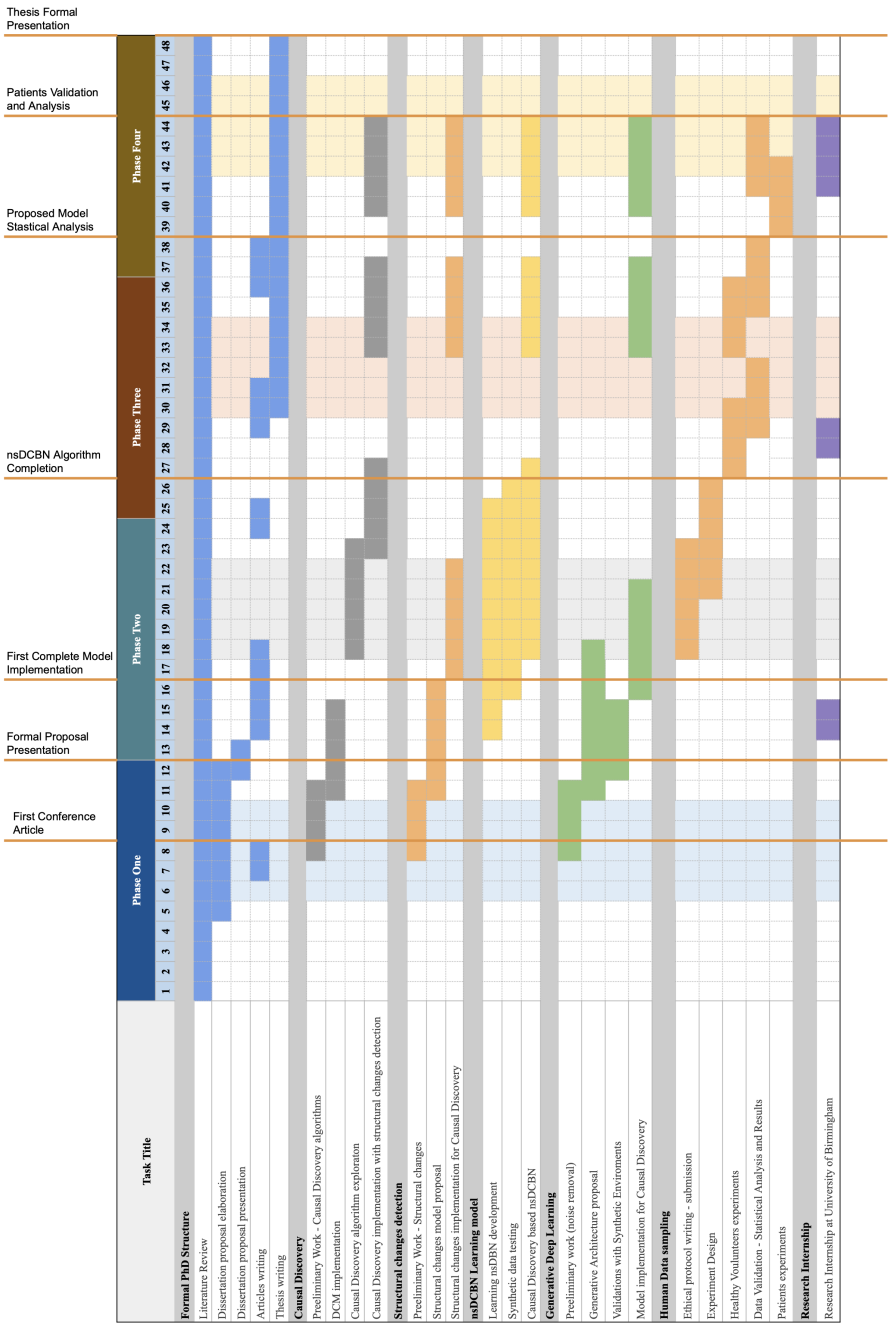

Figure 9: This timeline outlines the major milestones and tasks for achieving PhD accreditation. The left side lists key activities, while the top headers indicate the phases and specific milestones over 48 months. Each colored block represents the duration and overlap of various tasks throughout the timeline.

Objective: Present advances in the proposed causal discovery methods for learning nsDCBN in synthetic non-stationary Bayesian networks based on non-stationarity detection.
Estimated Submission Date: October 2025

- **International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)**
  - **Objective:** Present the advances in brain dynamic effective connectivity representation through nsDCBN.
  - **Estimated Submission Dates:** February 2026

4.4.2 Journal Articles

- **Journal of Causal Inference**
  - **Objective:** Publish a Q1 article on our advances in learning nsDCBN through causal discovery, using the proposed algorithm to detect structural changes in the data.
  - **Estimated Submission Date:** November 2025
- **International Journal of Approximate Reasoning**
  - **Objective:** Publish a Q1 article presenting our Causal Discovery in a non-stationary time-series model with applications in Brain Effective Connectivity.
  - **Estimated Submission Date:** November 2026
- **SPIE Neurophotonics**
  - **Objective:** Present the advances in the domain of effective brain connectivity reached in this research, mainly the capability of our model to characterize brain behavior during a neurorehabilitation process.
  - **Estimated Submission Date:** February 2026

4.5 Summary

In this chapter, we outline our methodology for achieving the research aims, detailing the processes of signal acquisition, non-stationarity detection, causal discovery, and developing a learning model for non-stationary Dynamic Causal Bayesian Networks (nsDCBN). Signals are acquired from synthetic and semisynthetic data before progressing to observational fNIRS data. We propose an algorithm for non-stationarity detection that uses changes in statistical moments to trigger the causal discovery process accurately.

Our causal discovery framework aims to identify unique causal structures within each time slice, modifying and adapting state-of-the-art algorithms and domain-specific alternatives. This iterative learning process will evolve the nsDCBN model over time, capturing both parametric and structural changes. We will validate the framework using synthetic, semisynthetic, and observational data, ensuring accuracy and reliability through expert knowledge and statistical tests.

The chapter also outlines a four-year work plan with specific milestones and phases and describes our publications plan, highlighting anticipated conference and journal articles to disseminate our findings.

5 Preliminary Results

This chapter details the progress made in the program's first year, which is aligned with our specific objectives and demonstrates the feasibility of the proposed research.

5.1 Synthetic Data Generator

We developed a synthetic data generator grounded in the bilinear model, figure 5, as outlined in [3]. The model effectively emulates neurodynamics, which can be extrapolated to model the hemodynamic response with control over the ground truth through a causal model. Subsequently, the hemodynamic response simulates the optical density changes emulating fNIRS samples, thus enhancing the realism of

the synthetic data used for system validation.

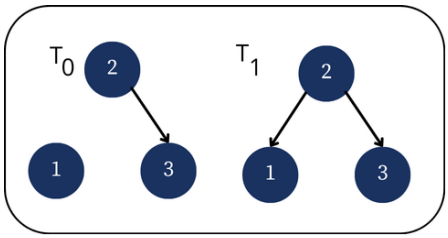

Figure 10: Ground truth causal structure showing causal relationships between three regions at two different structures. These structures generate the synthetic fNIRS dataset for the non-stationarity detection experiment (see Figure 13). Each node represents a region producing a channel, and the directed edges indicate causal influence.

The generator produces clean synthetic fNIRS signals given the causal model, figure 4. Then, we incorporated synthetic and semisynthetic noise to enhance realism and represent more complex scenarios. Synthetic noise includes typical physiological disturbances such as heart rate variability, breathing patterns, and vasomotion, along with a generic colored noise model to simulate additional external perturbations [69]. On the other hand, semisynthetic noise utilizes a predefined set of resting state data [66], which embodies the natural noises typically found in fNIRS resting state recordings.

Implementing this model enables us to propose a causal structure, figure 10, to generate models per the guidelines in [3]. Initially, the design of the structures is for experimental and validation purposes, figure 11. Nevertheless, more formal methodologies are available for crafting realistic causal models that accurately reflect brain responses [2].

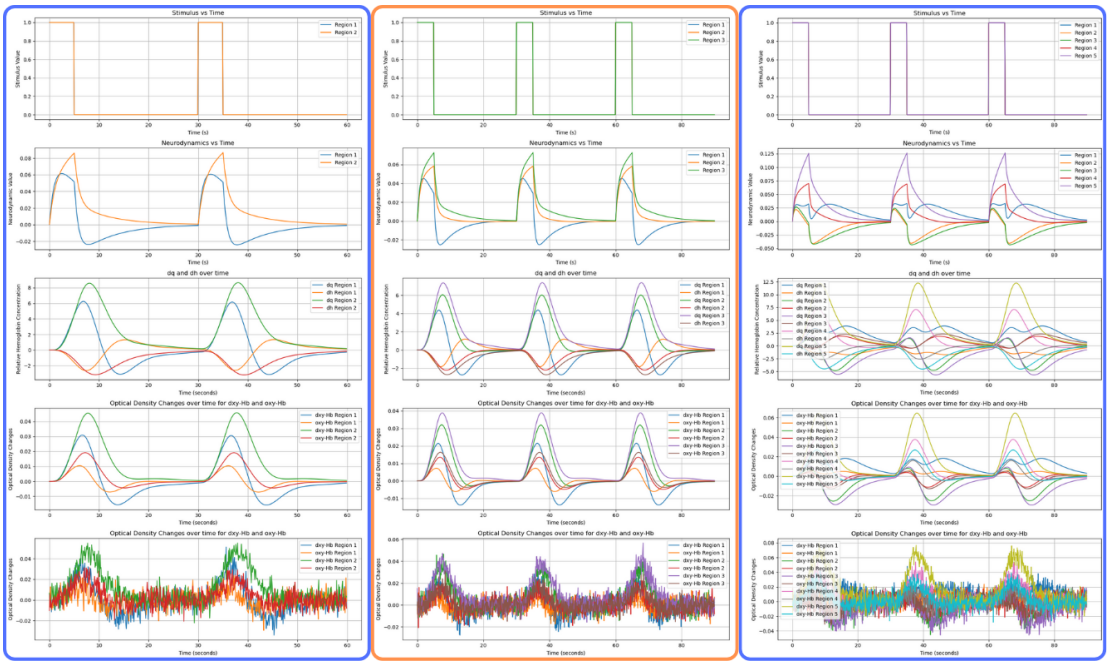

Figure 11: Examples of synthetic data generation for fNIRS signal analysis demonstrating the bilinear model's application across varying regions of interest (from left to right: two, three, and five regions, denoted by colored borders). Each scenario shows the designed stimulus pattern, followed by plots of neurodynamics activation, hemodynamics response, pristine fNIRS signal output, and the signal with synthetic and semisynthetic noise. These stages illustrate the model's ability to emulate physiological and environmental noise influences on neuroimaging data.

5.2 Non-Stationarity Detector

We propose a non-stationarity detector algorithm that uses a signal's first four statistical moments: Mean, Variance, Skewness, and Kurtosis (Figure 14). Each moment provides insight into the signal's specific properties. The algorithm analyzes two types of timestamp windows for each slice k: one with a 50% overlap between windows k and $k-1$ and the other as an independent window in slice k.

The algorithm also applies the KPSS test, section 3.3, to determine the signal's stationarity in both window strategies. The KPSS test categorizes a series as stationary based on its stationarity around a mean or linear trend. A significant test

statistic indicates the presence of a unit root, suggesting non-stationarity.

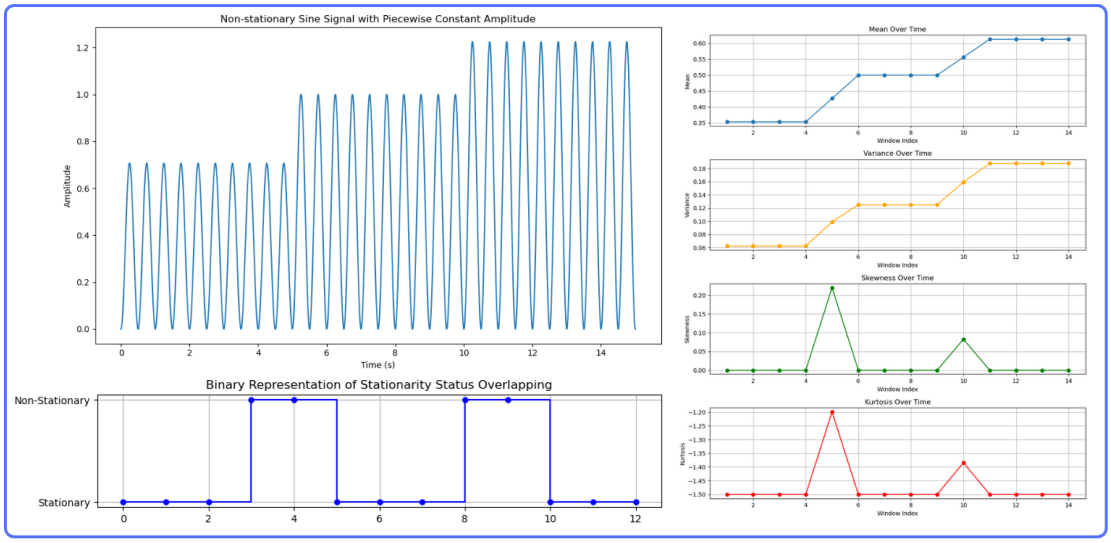

Figure 12: Non-stationarity detection in a sine signal. This experiment illustrates a sine signal with piecewise constant amplitude changes. The signal's amplitude varies following three different values [0.5, 1, 1.5] Hz, each active for a specific duration to allow controlled observation of changes. Below the signal plot, the binary representation indicates that non-stationarity was detected over time. The mean, variance, skewness, and kurtosis plots on the right show clear and consistent changes in these statistical moments. Each point in the plots represents a window in a slice k with overlapping segments.

To validate this approach, we implemented various controlled scenarios, starting with a sine signal exhibiting controlled non-stationary amplitude (Figure 12), followed by scenarios using data from the synthetic data generator (Figure 13) where we control the variation of the causal structure of the analyzed sample using the models described in figure 10.

The algorithm determines the stationarity status by evaluating the results of the statistical tests and comparing the statistical moments of the current and previous windows. A threshold, defined as twice the standard deviation of the accumulated moments' mean, is used to detect significant changes. If both windows indicate non-stationarity, the algorithm infers a transition from a stationary trend to a non-stationary one. The experiment results show that the model can detect a change in signal stationarity. Although this initial approach shows promising results, we aim

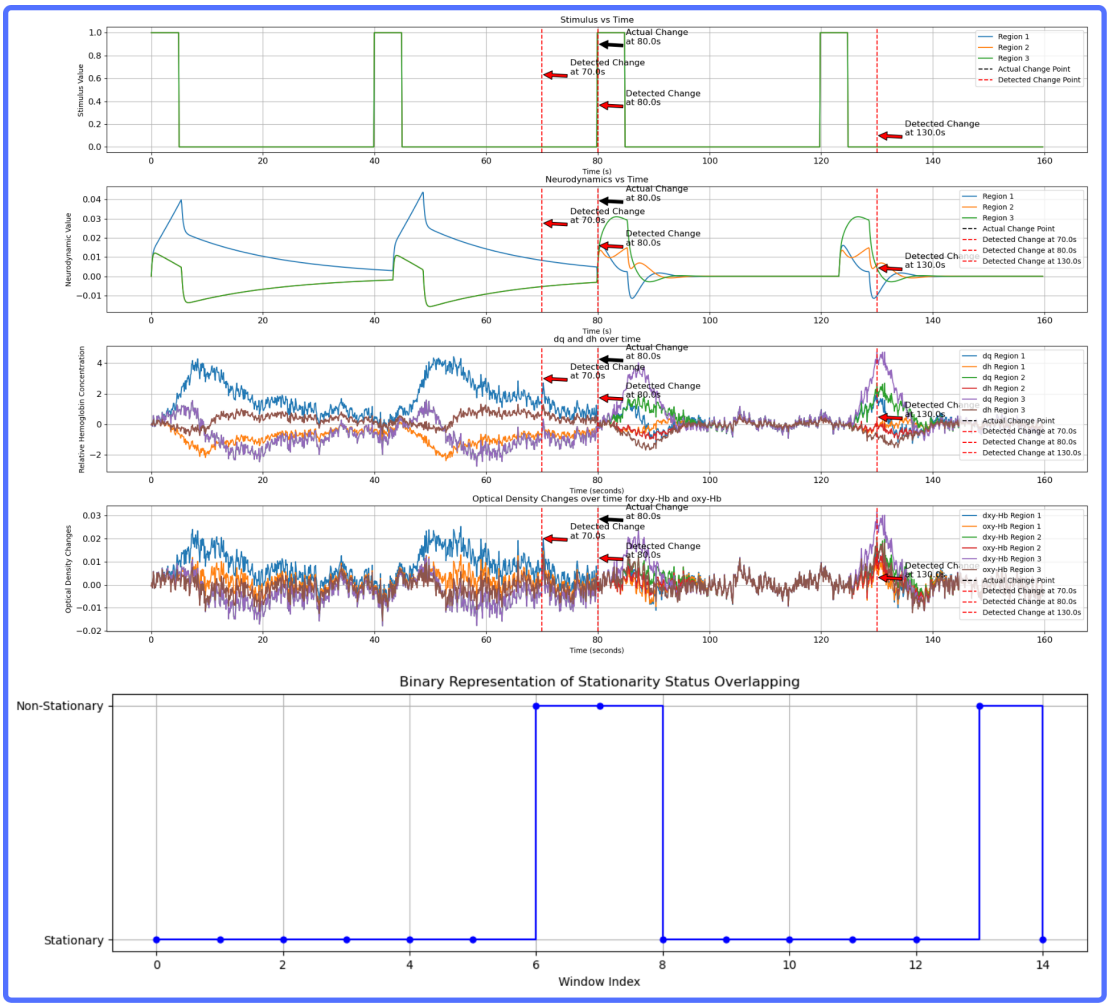

Figure 13: Non-stationarity detection in synthetic fNIRS signals based on the ground truth (Figure 10). During the first 80 seconds of the sample, the first scenario is observed, followed by a change to the second causal structure. Detection is performed individually on each channel, affecting the entire signal set if changes occur. The top plots show the synthetic fNIRS data, including stimulus and neurodynamic values. The bottom plot indicating binary non-stationarity detection using overlapping windows. The actual change points and detected change points are clearly annotated on the plots (actual change in black, while detected change in red).

to include other statistical tests and a decision based on a voting mechanism.

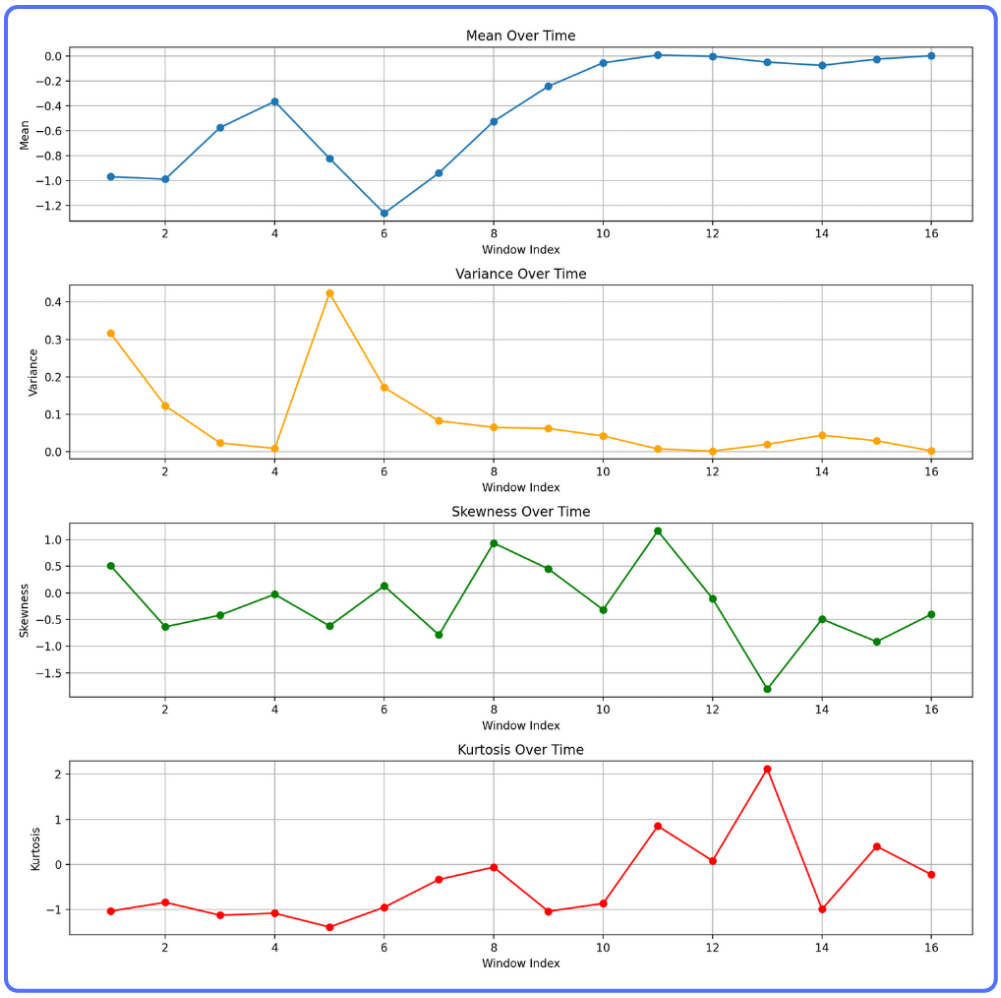

Figure 14: Non-stationarity detection in synthetic fNIRS signals based on the ground truth (Figure 10). Plots display the first four statistical moments evolution over time.

We evaluated the performance of the non-stationarity detector by measuring the time difference between the actual change points and the detected change points. The results indicate that the algorithm reliably identifies changes within a short time lag, with an average time difference of ± 10 seconds between actual and detected change points (in a sequence of 160 seconds); this is achieved given the overlapping strategy used in the observational windows, demonstrating the algorithm's capability to detect changes in the signal. In this example, we observed a false positive at window 13, likely due to high noise levels in the signal. Despite this, the overall false positive rate is low. To further improve performance, we aim to reduce these errors

through the proposed denoising stage (Section 3.5).

5.3 Summary

This chapter presents the preliminary work developed in line with the objectives outlined earlier. We introduce an innovative algorithm for detecting non-stationarity, a critical research component planted in the specific objective 3. This contribution advances the state of the art in this area.

Additionally, developing the synthetic model generator enables the creation of realistic scenarios for validating our research under various dynamic conditions, noise levels, and parameters following our specific objective 1. This chapter highlights the initial efforts and progress in developing fundamental elements for subsequent research stages.

6 Final Remarks

Non-stationary processes remain a significant challenge in causal discovery. Inspired by their potential applications in neurorehabilitation, this research proposal aims to develop a framework for learning non-stationary Dynamic Causal Bayesian Networks (nsDCBNs). This framework is envisioned to model brain connectivity dynamics, particularly in neurorehabilitation, effectively.

The proposed research includes several key stages: the development of a denoising deep learning architecture 4.1.2, the creation of a novel non-stationarity detector 4.1.3, and the iterative modification of a causal discovery model to accommodate structural changes 4.1.4. We are pleased to report that we have successfully developed a synthetic data generator that emulates realistic fNIRS time series, providing a robust foundation for the initial design and validation of our models 5.1. Additionally, the first version of a non-stationarity detector has been introduced 5.2, leveraging statistical moments and tests to identify changes. Preliminary results from this detector are promising, indicating its potential effectiveness.

These preliminary achievements validate two specific objectives (1 and 3) of our research and demonstrate the feasibility and potential of our approach. Moving forward, we will refine and expand upon these initial developments, maintaining a rigorous and iterative approach to enhance the models' accuracy and applicability.

In summary, while the journey of this research is still ongoing, the foundational work completed thus far offers an optimistic outlook. The proposed framework, with its promising potential, is poised to address the complexities of non-stationary processes in brain connectivity, paving the way for exciting advancements in neurorehabilitation strategies.

References

[1] R. e. a. Hutchison, "Dynamic functional connectivity: Promise, issues, and interpretations," *NeuroImage*, vol. 80, pp. 360–378, 2013.

[2] K. J. Friston, L. Harrison, and W. Penny, "Dynamic causal modelling," *Neuroimage*, vol. 19, no. 4, pp. 1273–1302, 2003.

[3] S. Tak, A. Kempny, K. J. Friston, A. P. Leff, and W. D. Penny, "Dynamic causal modelling for functional near-infrared spectroscopy," *Neuroimage*, vol. 111, pp. 338–349, 2015.

[4] S. Saetia, N. Yoshimura, and Y. Koike, "Constructing brain connectivity model using causal network reconstruction approach," *Frontiers in Neuroinformatics*, vol. 15, p. 619557, 2021.

[5] R. e. a. Sanchez-Romero, "Causally informed activity flow models provide mechanistic insight into network-generated cognitive activations," *NeuroImage*, vol. 278, p. 120300, 2023.

[6] A. R. Anwar, M. Muthalib, S. Perrey, A. Galka, O. Granert, S. Wolff, and M. Muthuraman, "Effective connectivity of cortical sensorimotor networks during finger movement tasks: a simultaneous fnirs, fmri, eeg study," *Brain Topography*, vol. 29, pp. 645–660, 2016.

[7] O. Sporns, *Networks of the Brain*. Cambridge, MA: MIT Press, 2010.

[8] J. Pearl, *Causality*. Cambridge University Press, 2009.

[9] J. Robinson and A. Hartemink, "Non-stationary dynamic bayesian networks," *Preprint*. No publication date.

[10] K. J. Friston, A. M. Bastos, A. Oswal, B. Van Wijk, C. Richter, and V. Litvak, "Granger causality revisited," *Neuroimage*, vol. 101, pp. 796–808, 2014.

[11] B. Huang, K. Zhang, and B. Schölkopf, "Identification of time-dependent causal model: a gaussian process treatment," *Twenty-fourth international joint conference on artificial intelligence*, 2015.

[12] B. Huang, K. Zhang, M. Gong, and C. Glymour, "Causal discovery and forecasting in nonstationary environments with state-space models," *arXiv preprint arXiv:1905.10857*, 2019.

[13] D. Fujiwara, K. Koyama, K. Kiritoshi, T. Okawachi, T. Izumitani, and S. Shimizu, "Causal discovery for non-stationary non-linear time series data using just-in-time modeling," in *Proceedings of the Second Conference on Causal Learning and Reasoning*, vol. 213 of *Proceedings of Machine Learning Research*, pp. 880–894, PMLR, 2023.

[14] J. Runge, "Tigramite: Causal discovery for time series analysis." `https://jakobrunge.github.io/tigramite/`, 2020.

[15] J. e. a. Liu, "Estimating brain effective connectivity in fmri data by non-stationary dynamic bayesian networks," in *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, (San Diego, CA, USA), pp. 834–839, IEEE, 2019.

[16] E. Kirilina, N. Yu, A. Jelzow, H. Wabnitz, A. M. Jacobs, and I. Tachtsidis, "Identifying and quantifying main components of physiological noise in functional near infrared spectroscopy on the prefrontal cortex," *Frontiers in Human Neuroscience*, vol. 7, p. 864, 2013.

[17] B. Huang, K. Zhang, J. Zhang, J. Ramsey, R. Sanchez-Romero, C. Glymour, and B. Schölkopf, "Causal discovery from heterogeneous/nonstationary data," *Journal of Machine Learning Research*, vol. 21, no. 89, pp. 1–53, 2020.

[18] Y.-t. Fan, C.-y. Wu, H.-l. Liu, K.-c. Lin, Y.-y. Wai, and Y.-l. Chen, "Neuroplastic changes in resting-state functional connectivity after stroke rehabilitation," *Frontiers in Human Neuroscience*, vol. 9, 2015.

[19] R. Allison, L. Shenton, K. Bamforth, C. Kilbride, and D. Richards, "Incidence, time course and predictors of impairments relating to caring for the profoundly affected arm after stroke: a systematic review," *Physiotherapy Research International*, vol. 21, no. 4, pp. 210–227, 2016.

[20] J. Robinson and A. Hartemink, "Learning non-stationary dynamic bayesian networks," *Preprint*. No publication date.

[21] L. E. Sucar, *Probabilistic graphical models*. Advances in Computer Vision and Pattern Recognition, London: Springer London, 2021.

[22] G. Cooper and E. Herskovitz, "A bayesian method for the induction of probabilistic networks from data," *Machine Learning*, vol. 9, no. 4, pp. 309–348, 1992.

[23] D. Heckerman, "A tutorial on learning with bayesian networks," in *Innovations in Bayesian Networks*, pp. 33–82, Netherlands: Springer, 2008.

[24] C. W. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica: Journal of the Econometric Society*, pp. 424–438, 1969.

[25] D. Hume, *A Treatise of Human Nature*. Oxford University Press, 1740.

[26] P. Suppes, *A Probabilistic Theory of Causality*. Amsterdam: North Holland Publishing Company, 1970.

[27] P. W. Holland, "Statistics and causal inference," *Journal of the American Statistical Association*, vol. 81, no. 396, pp. 945–960, 1986.

[28] J. Pearl, "Causal inference in statistics: An overview," *Statistics Surveys*, pp. 96–146, 2009.

[29] A. Darwiche, *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press, 2009.

[30] V. Mihajlovic and M. Petkovic, "Dynamic bayesian networks: A state of the art," *Pattern Recognition*, vol. 35, no. 12, pp. 2431–2452, 2001.

[31] S. F. Galán and F. J. Díez, "Modeling dynamic causal interaction with bayesian networks: Temporal noisy gates," in *Proceedings of the 2nd International Workshop on Causal Networks*, pp. 1–5, August 2000.

[32] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search*. MIT Press, 2000.

[33] D. M. Chickering, "Optimal structure identification with greedy search," *Journal of Machine Learning Research*, vol. 3, pp. 507–554, 2002.

[34] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen, "A linear non-gaussian acyclic model for causal discovery," *Journal of Machine Learning Research*, vol. 7, pp. 2003–2030, 2006.

[35] K. J. Friston, "Functional and effective connectivity: a review," *Brain Connectivity*, vol. 1, no. 1, pp. 13–36, 2011.

[36] P. Langhorne, J. Bernhardt, and K. G, "Stoke rehabilitation," *The Lancet*, vol. 377, no. 9778, pp. 1693–1702, 2011.

[37] P. Pinti, I. Tachtsidis, A. Hamilton, J. Hirsch, C. Aichelburg, S. Gilbert, and P. W. Burgess, "The present and future use of functional near-infrared spectroscopy (fnirs) for cognitive neuroscience," *Annals of the New York Academy of Sciences*, vol. 1464, no. 1, pp. 5–29, 2020.

[38] T. Huppert, "Introduction to the nirs-toolbox," in *4th Biennial Meeting of the Society for Functional Near Infrared Spectroscopy (fNIRS)*, p. 215, 2016.

[39] M. Ferrari and V. Quaresima, "A brief review on the history of human functional near-infrared spectroscopy (fnirs) development and fields of application," *Neuroimage*, vol. 63, no. 2, pp. 921–935, 2012.

[40] A. Gallagher, M. Thériault, E. Maclin, K. Low, G. Gratton, M. Fabiani, and M. Lassonde, "Near-infrared spectroscopy as an alternative to the wada test for language mapping in children, adults and special populations," *Epileptic Disorders*, vol. 9, no. 3, pp. 241–255, 2007.

[41] R. Cai, L. Huang, W. Chen, J. Qiao, and Z. Hao, "Learning dynamic causal mechanisms from non-stationary data," *Applied Intelligence*, vol. 53, no. 5, pp. 5437–5448, 2023.

[42] P. de Jong and S. Chu-Chun-Lin, "Stationary and non-stationary state space models," *Journal of Time Series Analysis*, vol. 15, no. 2, pp. 151–166, 1994.

[43] K. Zhang, B. Huang, J. Zhang, C. Glymour, and B. Schölkopf, "Causal discovery from nonstationary/heterogeneous data: Skeleton estimation and orientation determination," *IJCAI: Proceedings of the Conference*, vol. 2017, p. 1347, 2017.

[44] P. Shiguihara, A. D. A. Lopes, and D. Mauricio, "Dynamic bayesian network modeling, learning, and inference: A survey," *IEEE Access*, vol. 9, pp. 117639–117648, 2021.

[45] Y. Jia and J. Huan, "Constructing non-stationary dynamic bayesian networks with a flexible lag choosing mechanism," *BMC Bioinformatics*, vol. 11 Suppl 6, no. Suppl 6, p. S27, 2010.

[46] P. J. Green, "Reversible jump markov chain monte carlo computation and bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.

[47] M. Grzegorczy, D. Husmeier, K. Edwards, P. Ghazal, and A. Millar, "Modelling non-stationary gene regulatory processes with a non-homogeneous bayesian network and the allocation sampler," *Bioinformatics*, vol. 24, pp. 2071–2078, 2008.

[48] M. Grzegorczyk and D. Husmeier, "Non-stationary continuous dynamic bayesian networks," in *Advances in Neural Information Processing Systems 22*, 2009.

[49] D. A. Dickey and W. A. Fuller, "Distribution of the estimators for autoregressive time series with a unit root," *Journal of the American Statistical Association*, vol. 74, no. 366a, pp. 427–431, 1979.

[50] D. Kwiatkowski, P. C. Phillips, P. Schmidt, and Y. Shin, "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?," *Journal of Econometrics*, vol. 54, no. 1-3, pp. 159–178, 1992.

[51] P. C. Phillips and P. Perron, "Testing for a unit root in time series regression," *Biometrika*, vol. 75, no. 2, pp. 335–346, 1988.

[52] A. W. Lo and A. C. MacKinlay, "Stock market prices do not follow random walks: Evidence from a simple specification test," *The Review of Financial Studies*, vol. 1, no. 1, pp. 41–66, 1988.

[53] S. Kanaya, "A nonparametric test for stationarity in continuous-time markov processes," *Journal of Econometrics*, vol. 164, no. 2, pp. 329–353, 2011.

[54] A. van Delft and M. Eichler, "A nonparametric test for stationarity in functional time series," *Journal of Time Series Analysis*, vol. 39, no. 6, pp. 883–904, 2018.

[55] S. Basu, R. A. Davis, and T. Mikosch, "A nonparametric test for stationarity in the mean based on local fourier analysis," *Journal of Time Series Analysis*, vol. 30, no. 2, pp. 169–190, 2009.

[56] J. Breitung, "Nonparametric tests for unit roots and cointegration," *Journal of Econometrics*, vol. 108, no. 2, pp. 343–363, 2002.

[57] A. Cardinali and G. P. Nason, "Practical powerful wavelet packet tests for second-order stationarity," *Journal of Time Series Analysis*, vol. 39, no. 6, pp. 764–787, 2018.

[58] S. e. a. Montero-Hernandez, "Causal probabilistic graphical models for decoding effective connectivity in functional near infrared spectroscopy," *Preprint*. No publication date.

[59] C. e. a. Eastmond, "Deep learning in fnirs: a review," *Neurophotonics*, vol. 9, no. 04, 2022.

[60] Y. Gao *et al.*, "Deep learning-based motion artifact removal in functional near-infrared spectroscopy," *Neurophotonics*, vol. 9, no. 4, p. 041406, 2022.

[61] M. Kim, S. Lee, I. Dan, and S. Tak, "A deep convolutional neural network for estimating hemodynamic response function with reduction of motion artifacts in fnirs," *Journal of Neural Engineering*, vol. 19, no. 1, p. 016017, 2022.

[62] G. Lee, S. H. Jin, and J. An, "Motion artifact correction of multi-measured functional near-infrared spectroscopy signals based on signal reconstruction using an artificial neural network," *Sensors*, vol. 18, no. 9, p. 2957, 2018.

[63] R. Liu *et al.*, "Unsupervised fnirs feature extraction with cae and esn autoencoder for driver cognitive load classification," *Journal of Neural Engineering*, vol. 18, no. 3, p. 036002, 2021.

[64] S.-W. Woo, M.-K. Kang, and K.-S. Hong, "Classification of finger tapping tasks using convolutional neural network based on augmented data with deep convolutional generative adversarial network," in *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, pp. 559–564, IEEE, 2020.

[65] P. Pinti, F. Scholkmann, A. Hamilton, P. Burgess, and I. Tachtsidis, "Current status and issues regarding pre-processing of fnirs neuroimaging data: An investigation of diverse signal filtering methods within a general linear model framework," *Frontiers in Human Neuroscience*, vol. 12, p. 505, 2019.

[66] S. Montero-Hernandez, F. Orihuela-Espina, and L. E. Sucar, "Intervals of causal effects for learning causal graphical models," in *International Conference on Probabilistic Graphical Models*, pp. 296–307, PMLR, August 2018.