UNIVERSITY
OF COLOGNE

DISSERTATION for
the award of the doctoral degree
of the Faculty of Mathematics and Natural Sciences

Learning the Topology of Latent Signaling Networks from High Dimensional Transcriptional Intervention Effects

Zahra Sadat Hajseyed Nasrollah

March 25, 2025

Supervisors: Prof. Dr. Achim Tresch, Prof. Dr. Holger Froehlich
1st Examiner: Prof. Dr. Achim Tresch
2nd Examiner: Prof. Dr. Kay Hofmann

Prof. Froehlich is affiliated with the *Bonn-Aachen International Center for IT (b-it)* and the *Fraunhofer Institute for Algorithms and Scientific Computing (SCAI)*.

To the memory of the 176 innocent lives lost on board
Flight PS752 on January 8, 2020.
May we always strive for truth, justice, and humanity
in your honor.

Acknowledgements

I would like to express my deepest gratitude to all those who have guided and supported me throughout the journey of completing this thesis. Their encouragement, insights, and belief in my abilities have been invaluable. This work would not have been possible without the mentorship and kindness of so many individuals who shaped both my research and personal growth.

I am profoundly thankful to my supervisors, Prof. Dr. Achim Tresch and Prof. Dr. Holger Froehlich, for their exceptional guidance and patience. From Achim, I rediscovered the joy of mathematics and the fact that a solid work needs a solid proofs on paper first. Holger always approached our long discussions with care and thoughtfulness. His offering and solutions always inspired me throughout my Ph.D study. I am truly grateful to both for creating the opportunity for me to continue this work, even when the path was neither easy nor straightforward. Their encouragement has been a constant source of strength, and I feel privileged to have had their mentorship.

My heartfelt thanks go to my husband, my best friend, Lutz Werner, who stood by me in every moment of this study. His belief in me and my abilities gave me the strength to keep going. Without his presence, this path would have been far more challenging.

I am also deeply thankful to my friend, Ali Darijani, who was always there to answer my endless questions, especially in mathematics. He was the one who pointed out that placing parameters inside a tanh function could be risky.

I feel incredibly fortunate to have shared my Ph.D. time with Dr. Katharina Moos and Dr. Rafael Campos Martin. Thank you for filling this time with laugh, joy, and countless unforgettable moments. Your friendship is a true gift to me.

I warmly thank my Ph.D. group, the AG Tresch. To Dr. Niklas Kleinenkuhnen and Dr. Jason Müller, your genuine kindness and encouragement could inspire anyone to move

mountains. To Dr. Till Baar and Karin Sablonsky, I appreciate your patience in answering my many boring questions about forms and bureaucracies. To Dr. Vlada Milchevskaya and our balcony meetings time. To Kiarash, Thanina, Abdul and Mohammad the new Ph.D. students in our group. Your energy and laughter reminded me of the enthusiasm I felt in my early days. It has been a privilege to have such great people next to me.

I would also like to express my gratitude to the wonderful colleagues I worked with at company Operaize. In particular, I am deeply thankful to my senior, Fabian Killus. His mentorship have been a great honor and an invaluable part of my growth. Through his guidance, I learned the importance of persistency, even in tasks that may seem less significant, and the value of being truly team oriented. These lessons have had a lasting impact on both my personal and professional path.

I am deeply thankful to my wonderful friends who have always brought joy and stability, what I need the most :
Mahsa Barzegar-Keshteli, Paniz Rasooli, Mahdi Saffar, Narges Jafari, Hossein Khazaei, Dr. Amir Sabzalipour, Dr. Maryam Habibi, Donya Saberi, Dr. Mina Morshed Behbahani, Dr. Marzie Karimi, Dr. Sarah Rahmati, Nina Hopen, Patricia Moutinho, Elis Martineli, Lukas Werner, Paul Weyer and many more lovely people.

Finally, my special thanks go to my parents and my dearest sister and brothers. Despite the distance, they have stood by me with their unconditioned love, accepting our separation with understanding and faith in me. Their belief in my journey has meant the world to me.

Abstract

Biological systems rely on complex networks of interacting molecules, genes, proteins, and other regulators to govern cellular processes. However, inferring the causal structure of these networks is challenging. Particularly when important regulatory factors are only indirectly measured or entirely hidden. This difficulty is combined with the need to distinguish direct from indirect interactions and to capture dynamic responses under different perturbations such as gene knockouts or drug treatments. Consequently, developing a method that can incorporate both time resolved data and partial observability is critical for inference of the intricate architectures underlying signal transduction and gene regulation.

To address these challenges, we propose an Ordinary Differential Equation based Nested Effects Model (odeNEM). Our approach combines the established Nested Effects Model framework,which maps hidden regulatory nodes to observable downstream effects, with a continuous time ODE formulation. This combination can capture saturable, nonlinear kinetics. By explicitly modeling the propagation of perturbation signals over time, odeNEM infers which hidden regulators are causally upstream of others. Additionally, we developed mixture models for the downstream expression data which is considering the noisy high throughput nature of biological experiments.

We validate our method on two key applications. First, we reconstruct signaling pathways in breast cancer cells using phosphoprotein time course data from the HPN-DREAM challenge. This shows how odeNEM captures context specific interactions among kinases such as AKT, mTOR, and MEK. Second, we infer a gene regulatory network in pluripotent stem cells from CRISPR single cell transcriptomics (RENGE), highlighting both canonical regulators like POU5F1 and SOX2 and novel interactions that warrant further biological investigation. Across both use cases, our results align with known biology and external data sources (e.g., ChIP-seq). This also confirms that a continuous time hidden node model can robustly uncover causal relationships.

Overall, odeNEM expands the applicability of NEM approache by explicitly modeling. The model's synergy with perturbation data, prior knowledge, and advanced inference strategies (e.g., MCMC) enables a more robust and biologically realistic reconstruction

of latent signaling networks. This contributes a promising picture for unraveling the complex interplay of molecular interactions in systems biology, with broad implications for identifying therapeutic targets and understanding disease mechanisms.

Zusammenfassung

Biologische Systeme basieren auf komplexen Netzwerken, in denen Gene, Proteine und andere regulatorische Faktoren zusammenwirken, um zelluläre Prozesse zu steuern. Die Rekonstruktion der zugrunde liegenden kausalen Strukturen ist eine große Herausforderung, insbesondere wenn wichtige Regulatoren nur indirekt oder gar nicht direkt messbar sind. Hinzu kommt die Schwierigkeit, zwischen direkten und indirekten Interaktionen zu unterscheiden und zeitabhängige Reaktionen auf verschiedene Störungen,etwa Gen Knockouts oder medikamentöse Inhibitoren, korrekt abzubilden. Ein entscheidender Schritt ist daher die Entwicklung von Methoden, die sowohl zeitaufgelöste Daten als auch teilweise verborgene Regulationsmechanismen integrieren können.
Um diesen Anforderungen zu begegnen, stellen wir ein „Ordinary Differential Equation basiertes Nested Effects Model“ (odeNEM) vor. Dieser Ansatz kombiniert das Nested Effects Model Konzept, bei dem versteckte regulatorische Knoten über beobachtete Downstream Effekte erschlossen werden, mit einer kontinuierlichen, nichtlinearen ODE Modellierung. Durch die explizite Abbildung der Signalweiterleitung über die Zeit ermöglicht odeNEM die Identifikation kausaler Beziehungen zwischen verborgenen Regulatoren, selbst wenn deren Aktivitäten nicht direkt gemessen werden können. Ergänzend verwenden wir Mischmodelle für die Rausch und Heterogenitätseffekte in den Experimentaldaten, sodass unser Verfahren robust gegenüber typischen Messungenauigkeiten in Hochdurchsatz Experimenten ist.
Zur Validierung haben wir odeNEM auf zwei zentrale Anwendungen übertragen. Erstens rekonstruieren wir Signalwege in Brustkrebszellen mithilfe von Phosphoprotein Zeitreihen aus der HPN-DREAM-Challenge. Dabei erfasst odeNEM kontextspezifische Interaktionen zwischen Kinase Netzwerken wie AKT, mTOR und MEK. Zweitens inferieren wir in pluripotenten Stammzellen (hiPSCs) ein Genregulationsnetzwerk auf Basis von CRISPR Experimenten mit Einzelzell RNA Sequenzierung (RENGE). Hier identifizieren wir sowohl etablierte Regulatoren wie POU5F1 und SOX2 als auch neue potenzielle Interaktionspartner, die weiterführende biologische Untersuchungen nahelegen. In beiden Fällen stimmen unsere Ergebnisse gut mit externen Datensätzen (z.B. ChIP-seq) und etablierter Literatur überein, was die Verlässlichkeit des Modells

unterstreicht.
Insgesamt liefert odeNEM einen einheitlichen Ansatz, um zeitabhängige Daten, verdeckte Regulatoren und nichtlineare Kinetik in einem Modell zu vereinen. Durch die Verknüpfung von ODE Dynamik und Nested Effects Logik stellt dieses Verfahren ein flexibles Werkzeug zur Analyse komplexer molekularer Netzwerke dar sei es im Bereich der Proteinsignalgebung oder der Genregulation. Dies trägt nicht nur zu einem vertieften Verständnis grundlegender biologischer Prozesse bei, sondern bildet auch eine Grundlage für gezieltere therapeutische Eingriffe, indem es die Ausbreitung von Störungen über bisher unsichtbare Schichten zellulärer Netzwerke aufdeckt.

Contents

1 Introduction

This chapter introduces the foundational concepts and challenges addressed in this thesis, focusing on signaling pathways and their interplays. A central objective of this work is to develop methodologies for reconstructing biological networks from time-series interventional data, particularly under conditions where hidden variables and incomplete observations complicate analysis. The chapter begins by discussing the basic building blocks of cellular systems, emphasizing gene regulation, signal transduction, and their dynamic and interconnected roles in cellular processes. Experimental techniques like RNA interference (RNAi), which provide key interventional data, are highlighted for their transformative role in functional genomics. The discussion then transitions to the challenges posed by non-transcriptional regulation and hidden variables, both of which necessitate advanced computational approaches for network inference. Through this background, the chapter lays the groundwork for exploring reverse engineering of biological networks and introduces the motivations behind the methodologies developed in subsequent chapters.

1.1 Biological Systems and Cellular Processes

Cells are basic building units of living organisms, and they function independently to carry out the processes needed for their survival. Surrounded by a protective membrane, the cell maintains its internal balance while interacting with its environment. Inside, molecules like proteins, enzymes, and nucleotides work together to perform essential tasks, allow- ing the cell to adapt and function effectively.
Each cell holds deoxyribonucleic acid (DNA), which contains the instructions needed for an organism's development, function, and reproduction. These instructions are organized into units called genes. Genes will lead the production of RNA and proteins. Proteins take on critical roles in shaping the cell's structure, facilitating its activities, and managing interactions within the cell. This flow of genetic information ensures that cells can adapt and respond to their surroundings [1]. The central dogma, Figure 1.1, of molecular biology explains how genetic information flows within a cell to create

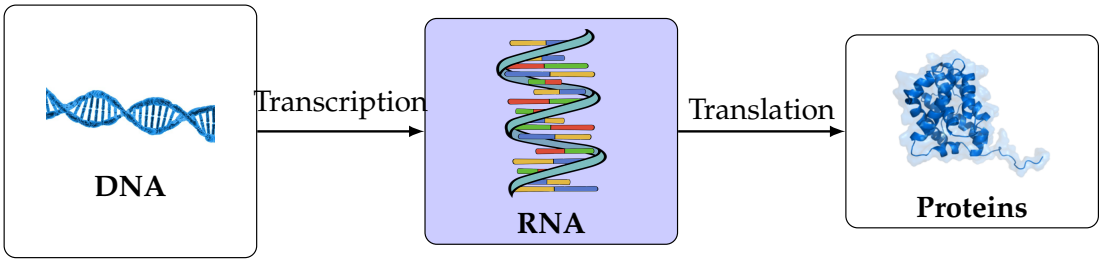

Figure 1.1: Figure: A simplified representation of the central dogma of molecular biology, illustrating the flow of genetic information from DNA to RNA (Transcription) and subsequently to Proteins (Translation). Icons sourced from Pixabay.

proteins [2]. DNA, the molecule that stores genetic instructions, undergoes transcription to produce RNA. This RNA, specifically messenger RNA (mRNA), serves as an intermediary to transfers the genetic code that is needed for proteins to be synthesized. During transcription, the mRNA undergoes essential modifications, such as splicing, where non-coding sequences (introns) are removed, and a 5′ cap and a poly-A tail are added. These modifications stabilize the mRNA and prepare it for translation.
Translation is the cellular process in which ribosomes decode an mRNA sequence to assemble amino acids into a growing polypeptide chain, ultimately forming a functional protein. This mechanism is carefully regulated to maintain accuracy and efficiency. During translation, ribosomes assemble amino acids into proteins by reading the instructions encoded in mRNA. The process is controlled by specialized factors that help ribosomes bind to mRNA and initiate synthesis [3]. Additionally, small RNA molecules can target specific mRNA for degradation or prevent its translation [4]. This precise regulation ensures proteins are made only when and where they are needed, maintaining cellular efficiency and balance.

Regulation of gene expression happens at both transcriptional and non-transcriptional levels, each contributing uniquely to cellular processes. Transcriptional regulation, where transcription factors influence gene expression, leaves measurable traces such as mRNA levels, making it more accessible through high-throughput techniques like RNA sequencing [5]. In contrast, non-transcriptional regulation (post-translational modifications, protein interactions, and enzymatic effects) lacks direct molecular signatures [6]. These layers of regulation ensure that cellular resources are used efficiently and that specific proteins are expressed at the right time to contribute to cellular functions [7].

1.1.1 Modeling Biological Systems with Networks

Gene regulation occurs at the level of individual genes and their products but is part of broader, interconnected cellular networks that integrate and coordinate molecular activities [7].

To represent cellular systems, networks are used broadly with nodes corresponding to molecular entities (genes, proteins, metabolites) and edges representing interactions (e.g., transcriptional regulation, protein-protein interactions, or metabolic flux) [1]. These networks are typically categorized into signaling, metabolic, and transcriptional networks, which interconnect to form a complex web of cellular activity.

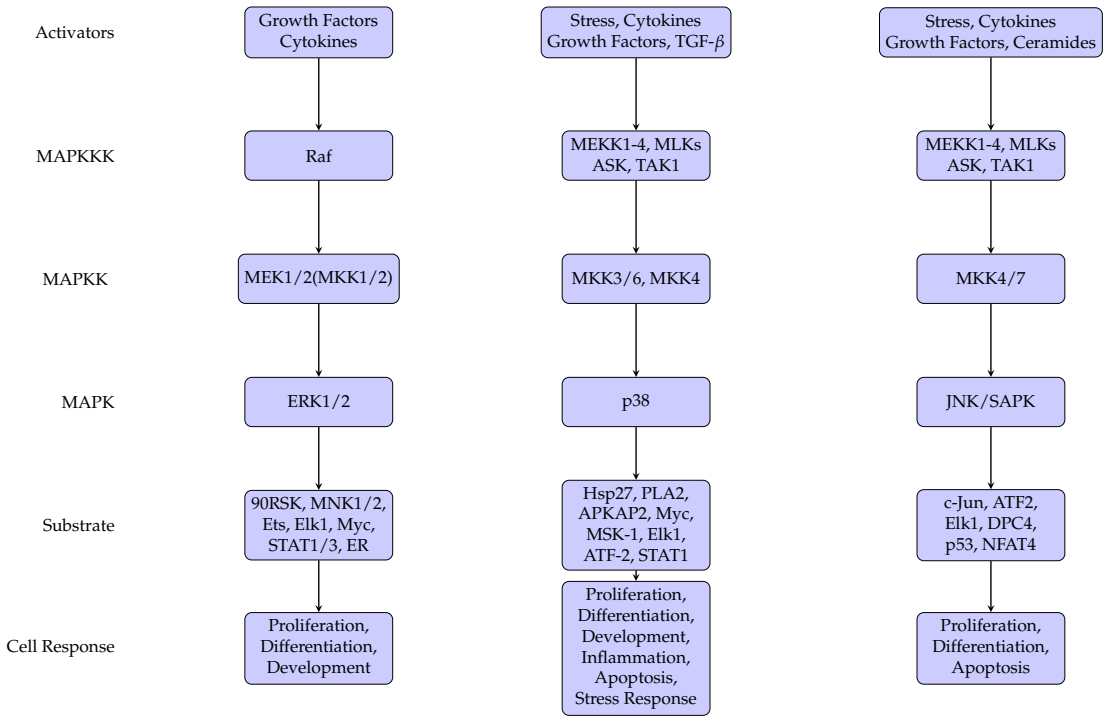

Figure 1.2: Simplified representation of MAPK signaling pathways in regulation of cell proliferation. This figure is adapted and redrawn from Figure 1 in [8]

Signaling Networks: Cells rely on signal transduction pathways to interpret and respond to their environment. These pathways transmit signals from receptors on the cell surface to downstream effectors through cascades of molecular interactions [9]. For instance, mitogen-activated protein kinase (MAPK) pathways help cells adapt to growth signals or stress by regulating key processes such as proliferation and differentiation, as it is shown in Figure 1.2 from [8]. These pathways are dynamic, featuring feedback

mechanisms and interactions with other signaling networks. Also, they may involve mechanisms that amplify the signal. Understanding the function of signaling networks is essential, as it uncovers how cells make decisions under normal and disease conditions.

Metabolic Networks: These networks map the biochemical reactions within cells that produce energy and synthesize cellular building blocks. Glycolysis and the citric acid cycle are fundamental metabolic pathways crucial for cell survival and growth ([10]). These networks are not isolated; they interact closely with signaling and transcriptional networks, as metabolic states can influence gene expression and protein activity [11].

Transcriptional Networks: These networks capture the regulatory relationships between transcription factors and their target genes, forming an intricate system of control over gene expression. Cells respond to environmental changes, developmental signals, and cellular stress via transcriptional regulation and modulating the production of RNA and proteins.
In transcriptional networks, nodes represent genes or regulatory elements, while edges denote interactions between transcription factors and their targets [12]. These interactions can be activation or repression, direct (such as physical binding to a promoter region), or inferred (through correlations in gene expression profiles). The integration of multiple regulatory signals allows transcriptional networks to act as control systems, determining spatial and temporal patterns of gene expression.

These networks, while distinct in their functions, are interconnected, forming a dynamic web that governs cellular behavior [13]. However, despite advancements in experimental techniques, many interactions and regulatory mechanisms remain difficult to catch. This is where defining a method to infer a network of relations from data becomes necessary.

1.2 Reverse Engineering Biological Networks

Reverse engineering is reconstructing the structure of biological networks by analyzing experimental data and reasoning backward to identify the underlying interactions among genes, proteins, and metabolites [14]. This approach is vital for understanding the regulatory and functional architecture of cellular systems, allowing to uncover pathways that drive complex biological processes.

High-throughput Observational Data: Data for reverse engineering largely comes from high-throughput experimental techniques, which capture the molecular and cellular processes under unperturbed conditions. RNA sequencing provides critical insights into transcriptional regulation by profiling gene expression under various conditions [15]. Proteomics reveals protein-protein interactions, post-translational modifications, and enzymatic activities [16]. ChIP-seq identifies binding sites of DNA-associated proteins, such as transcription factors [17]. This maps their roles into gene regulation. Additional techniques like mass spectrometry and microarrays provide further insights into protein abundance, transcriptomic changes, and other cellular dynamics.
While observational data is foundational, it often lacks the resolution to disentangle the complexities of network interactions. In contrast to observational data, perturbation-based data (which falls under interventional data) directly tests hypotheses about the relationships between components. Such interventional data enables a deeper understanding of cellular processes [14].

Perturbation-based Techniques: Perturbation-based approaches can reveal causal relationships within biological networks if the provided intervention is targeted to a known component. If it is unclear which part of the molecular network is perturbed (e.g., by an uncharacterized chemical), causal insights may be no greater than those from purely observational data. In practice, many perturbation strategies exist, such as environmental (e.g., heat or salt stress), pharmacological (e.g., small molecule drugs), and genetic approaches (e.g., gene knockdowns and knockouts). A gene knockdown (KD) partially reduces gene expression (often by degrading messenger RNA (RNAi)). KDs particularly are useful for studying essential genes without completely silencing targets [18]. A gene knockout (KO), by contrast, completely eliminates gene function. It is often implemented using CRISPR (Clustered Regularly Interspaced Short Palindromic Repeats) technology [19]. By combining different perturbation methods (knockdowns, knockouts, or environmental/drug treatments), researchers can examine both direct and adaptive responses, thereby refining causal inferences about molecular interactions within the network.

Data Driven Reverse Engineering of Biological Networks: Biological networks can be reconstructed using various data types, including purely observational data (e.g., transcriptome or proteome measurements under normal conditions) and perturbation-based data (which actively manipulate systems to reveal causal relationships). While observational data are widely used in the literature and can yield correlations among genes or

proteins, correlations alone do not imply causation. For instance, RNA sequencing might reveal that certain genes are coexpressed across conditions, yet it cannot distinguish direct regulatory links from indirect influences through intermediate molecules. To overcome this limitation, perturbation-based approaches (e.g., targeted knockdowns, knockouts, or chemical inhibitors) actively alter specific genes or proteins. Observing how these interventions affect the rest of the system provides stronger evidence of causal interactions [20].

The integration of these data sources calls for robust computational techniques to model and analyze the complexity of biological systems. By network inference, we mean the process of reconstructing or reverse engineering a biological network (e.g., a gene regulatory or protein interaction network) from the patterns observed in the data. Such methods typically draw on statistical models and machine learning to detect relationships and predict functional links. For instance, clustering algorithms group genes with similar expression profiles, while Bayesian networks and dynamic models capture temporal dependencies and probabilistic interactions [14, 20]. Beyond these correlation and probability based approaches, more formally defined causal frameworks exist. For instance, structural equation models (SEMs) provide a mathematical means for specifying causal dependencies among exogenous and endogenous variables, although they are not themselves a direct inference procedure. In contrast, do-calculus describes how targeted interventions can be represented in a graphical model and has proven valuable in distinguishing direct from indirect effects [21].

Experimental techniques we have mentioned, such as RNAi and CRISPR-based editing, provide the necessary perturbation data, systematically silencing or modifying genes to observe downstream effects. These methods complement observational data by revealing causal dependencies and highlighting critical nodes in the network. Computational tools further enhance this integration by mapping molecular information and reconstructing network topology, bridging the gap between raw experimental data and functional understanding.

1.3 Hidden Variables in Gene regulation

Hidden layers of regulation, such as phosphorylation events or secondary signaling pathways, can significantly influence cellular decisions. However, many large-scale studies (e.g., standard transcriptomic or proteomic analyses) primarily measure steady-state mRNA or protein abundance, making transient post-translational modifications (like phosphorylation) difficult to capture. Phosphorylation often occurs rapidly and

in context-specific bursts that can be missed by snapshot high throughput proteomic datasets [22] unless specialized phosphoproteomic methods are employed.
Unmeasured proteins, metabolites, or modifications in nontranscriptional regulation can complicate network reconstruction. Effects of such hidden variables mix with those known components, resulting in confounding and biased interpretations [23]. For example, in signaling pathways, secondary effects may obscure direct causal interactions, making it difficult to infer upstream downstream relationships accurately.
Addressing these complexities requires advanced computational approaches to mitigate confounding effects introduced by hidden variables, such as unmeasured genes or signaling intermediates. Hidden variables can distort the observed interplay between known components, leading to biased or incomplete network models. Although methods exist for learning networks with latent variables (using Bayesian Networks [24]), such approaches may still struggle when data sparsity or noise levels are high.
Nested Effects Models (NEMs) provide a promising framework for dealing with hidden variables and noise in observations. Conceptually, NEMs can be viewed as a specialized form of Bayesian Networks where subset relationships are emphasized [25], [26]. This makes NEMs well suited for inferring upstream downstream interactions in signaling genes, even when significant data gaps exist. By addressing the confounding influence of hidden variables, NEMs enable a more accurate reconstruction of biological networks.

1.4 Objectives and Thesis Organization

The goal of this thesis is to develop a new class of probabilistic graphical models within the DBN framework, which overcomes some major drawbacks of NEMs:

- Perturbation effects are propagated in a strictly deterministic sense through the hidden variable network: If there is a path from node A to node B, it is assumed that perturbation of A always yields a perturbation of B. This idea is rather simplistic and neglects possible compensatory effects, which occur frequently in biological systems. These compensatory effects may be due to nonlinearities in the system or due to the activation of unknown system components, which itself influence B.

- In the present NEM formulation, states correspond to binary hidden variables (protein ac- tivities), for which downstream perturbation effects (gene expression data) are observed. A gene knock-down presently corresponds to decoupling a specific hidden variable from its inputs and setting its activity constantly to 0.

Therefore, NEMs perform ideal interven- tions (from [27]) in the system of hidden variables. Experimental perturbations rarely achieve such an efficiency.

- NEMs require all nodes in the hidden variable network to be perturbed individually. This restriction to isolated perturbations is neither realistic nor desirable. Targeting several network nodes simultaneously (combinational interventions) may help to reveal the true network structure much more efficiently. Moreover, many experimental perturbations tar- get multiple network nodes simultaneously because the perturbing substance (e.g., a drug) is not specific. Other perturbations, like somatic mutations in cancer, accumu- late over time and characterize a set of genetic alterations, which represent combinational interventions into a patient's signaling network.

In Chapter 2, we will provide a comprehensive overview of Nested Effects Models (NEMs), including their principles and formulas, with a particular focus on the basic and general framework. Additionally, we will discuss further advancements in NEMs, leading to conclusions that set the stage for our own developments.
In Chapter 3, we introduce our approach called odeNEM (Ordinary Differential Equational NEMs). This method provides a more realistic representation of time compared to previous models. Furthermore, this chapter includes results from simulation studies, showcasing the performance of our approach.
The next two chapters are dedicated to demonstrating the application of our model on real datasets. Chapter 4 describes how our model was applied to infer protein-protein interaction networks. The final chapter, Chapter 5, extends the application of odeNEM to reconstruct regulatory networks, emphasizing its effectiveness in analyzing time-resolved experimental datasets and comparing them with existing methodologies. Both chapters include a detailed discussion of the datasets used, preprocessing steps, and comprehensive results.

2 Nested Effect Models (NEMs)

In the last chapter, we opened the discussion about reconstruction of biological network in the absence of observation for some elements. In this chapter, we will introduce Nested Effect Model as a model for such conditions. Since this model will be the foundation of our developed version, we will explain it thoroughly in this chapter.

2.1 Introduction

By external signals, cells respond through a complex and highly regulated sequence of events, which can be viewed as two distinct transduction layers. The first layer, called the transcriptional level, directly influences gene expression. Signaling transduction at this level activates transcription factors, which then target gene promoters, resulting in either activating or suppressing specific genes [28].
Then the non-transcriptional level comes, which is not easily detectable but plays a prominent role in the regulation of gene expression profiles via post-translational modifications (PTMs). This way of transduction, which consists of PTMs in conjunction with other factors like epigenetic alterations, adds a more complex layer to gene regulations and dynamic of cellular responses to external stimuli [29, 30]. This type of modification can not be directly measurable via profiling gene expressions. The occurrence of such modifications will affect the expression of related genes.
Understanding the downstream effects of PTMs is essential for a comprehensive view of cellular regulation. PTMs can modulate transcription factor activity, chromatin structure, and signaling cascades, leading to secondary effects on gene expression and cellular responses. By integrating both transcriptional and post-translational mechanisms, cells achieve a finely tuned response to external signals govern differentiation, adaptation, and stress response [13]. In the importance of exploring such indirect effects, [31] brought a hypothetical pathway example, which you can see the related schema in Figure 2.1. It contains three different levels of abstractions in biological concepts: DNA level (genome), mRNA level (transcriptome), and proteome level. When gene 1 is expressed, it produces a transcription factor that binds to the promoter of gene 2, regulating its expression. This

will encode a protein kinase, which yields the protein phosphatase expressed from gene 3 to be phosphorylated. In the next step, the protein phosphates will dephosphorylate the transcription factor of gene 4. By that event, the last mentioned transcription factor will be activated, bind to the gene 5 promoter, and change its expression.
To understand this concept, [31] expands the example by comparing two network models inferred from different data types. The first model (Figure 2.1a), which is inferred purely from observational gene expression data (mRNA level), fails to capture the effects of gene 3 and gene 4 on the expression of gene 5. This limitation arises because the influence of gene 3 and gene 4 occurs at the protein level, involving post-translational modifications such as phosphorylation and dephosphorylation do not directly appear in gene expression profiles. Consequently, methods relying only on statistical dependencies in expression data misinterpret the regulatory structure and incorrectly infer a direct connection between gene 2 and gene 5 while completely missing the intermediate molecular events involving gene 3 and gene 4.
This demonstrates a fundamental challenge in network reverse engineering from observational gene expression data. Although gene expression data can reveal statistical correlations between genes, it cannot establish causal relationships without additional structural assumptions or interventional experiments. Correlation-based methods may capture co-expression patterns but do not differentiate direct from indirect influences. A high correlation between genes does not imply a direct regulatory interaction—correlation may arise due to shared upstream regulators, hidden confounders, or indirect pathways that do not correspond to actual molecular interactions.
To overcome this limitation, Figure 2.1b incorporates interventional data, such as RNAi silencing of gene 3. By intervening in gene 3 and disrupting its activity, we block the normal flow of regulatory influence that would otherwise propagate through gene 3 to gene 5. If this intervention causes a measurable change in the expression of gene 5, it indicates that gene 3 plays an indirect but essential role in regulating gene 5.

Given that both transcriptional processes and post-translational modifications influence cellular responses, a computational model must accurately reflect these two types of transduction. As we mentioned in the above example, interventional studies like RNA interference (RNAi) can capture the profile of such hidden processes regulating pathways and their downstream effects. RNAi functions by silencing specific genes at the mRNA level, but it can indirectly cause hidden post-translational mechanisms, such as phosphorylations or other regulatory processes. This is achieved by observing the resulting changes in downstream gene expressions and protein activities, which may

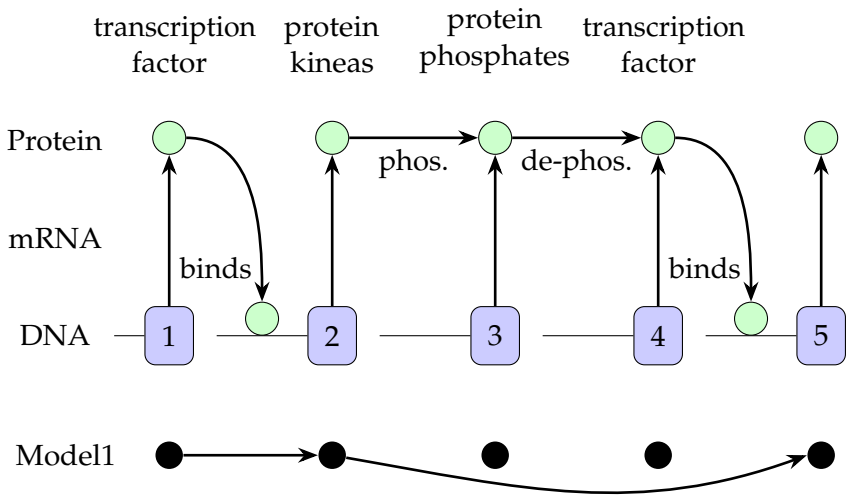

(a)

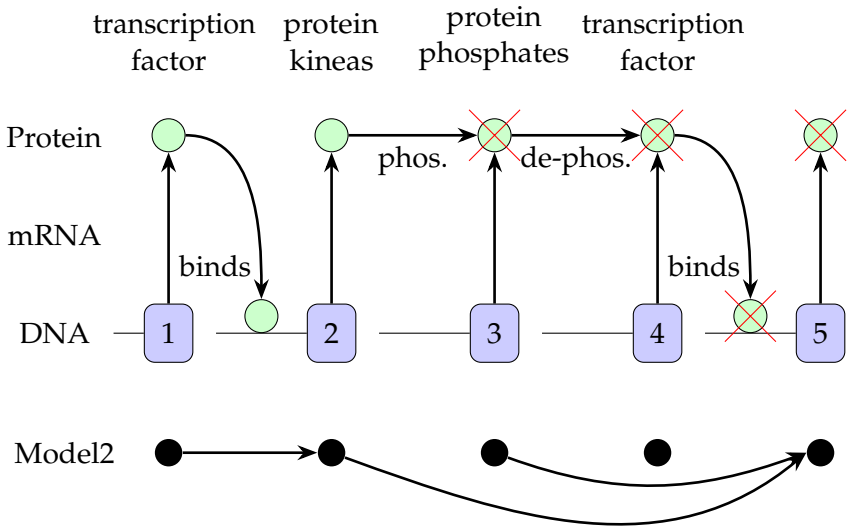

(b)

Figure 2.1: A hypothetical biochemical pathway- (a) illustrates a simplified biochemical pathway, showcasing interactions across biological levels: proteins, mRNA, and DNA. This model, derived from [31], highlights the limitations of inference based solely on gene expression data, where key contributions from genes 3 and 4 are missed. (b) demonstrates how interventional data, such as RNAi silencing of gene 3, disrupts pathway information flow and reveals an expression change at gene 5. This inclusion refines the model by uncovering a direct link between gene 3 and gene 5, thereby expanding its accuracy and predictive power. Adapted from [31]

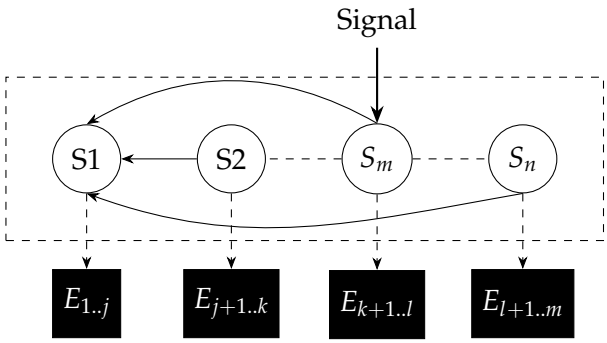

Figure 2.2: General Structure of Nested Effect Model (NEM)

provide insights into the underlying regulatory networks [32].
This mechanism includes two key steps. First, double-stranded RNA (dsRNA) has been split into small interfering RNAs (siRNAs) using RNA nucleases. SiRNAs join the RAN-induced silencing complex (RISC), which targets and degrades single-stranded mRNAs. The described process allows for the selective gene silencing without harming unrelated factors [33].

General class of probabilistic graphical models (PGMs) can address the challenge of modeling uncertain and complex relations, particularly those involving hidden variables with random characteristics [24]. The answer to this question remains: How can a probabilistic model infer such a hidden network from the downstream effects of perturbations?
Nested Effects Models (NEMs), as a specific class of PGMs, have been developed to leverage the alterations seen in transcriptional downstream response profiles due to a gene knockdown. This model can map these observed effects to a hidden regulatory network that causes the effect. However, the accuracy of NEMs depends on the informativeness of the perturbations and the complexity of the network.

2.2 NEMs

NEMs emerged from the pioneering work [25], builds upon preliminary work of [34], offering a unique approach to mapping secondary profiles (particularly in the context of RNAi screens) to the hidden network caused it. These nested downstream effects are not part of the model directly but rather serve as observable outcomes to inform NEMs from underlying hidden dynamics.

2.2.1 Basics

This graphical model uses two types of nodes, which can also mimic "genes" and categorized into two primary groups:

- $S-$genes (signaling or hidden nodes): These are the (nontranscriptional) nodes that are subjected to perturbation (silenced in RNAi experiments) in the pathway of a study. We denote these as $S = \{S_1, S_2, ..., S_n\}$. The received effects of perturbations on these nodes are not directly measurable, and we are interested in finding the interplay of such elements with other key elements in the regulatory network.
- E-genes (Effects): These are the (transcriptional) genes represented as $E = \{E_1, E_2, ..., E_m\}$ that show changes in their expression profiles due to an intervention on an $S-$gene.

The conceptual architecture of NEMs, shown by Figure 2.2, is a hypothesized directed graph between signaling $S-$genes against the observed cascading structure of downstream $E-$genes.
The main idea of NEM benefits the mentioned structure: the main goal is to capture the changes in information flow within the pathway ($S-$genes) between two states: before and after the interventions. Before any perturbation, connections between elements of such a regulatory network maintain a dynamic equilibrium where the rates of changes are balanced. When a perturbation, such as a gene knockdown, occurs, it disrupts the balance of a few genes. As a result, alters the dynamic of a (biological) system from a steady-state to some new activity level. This causes the expression of some $E-$genes to be impacted by perturbation while the others stay in their basic expressions [35, 36].
The $S-$genes can assume binary states '1' or '0'. While '1' represents an interruption in signaling, '0' exemplifies a node that is an active participant in the pathway's information flow. [25] assumed data of phenotypic profiles has a form of a binary matrix $\mathcal{D}$ which columns refer to experiments (on each $S-$genes),and rows are observed effects as $E-$genes, like what we see in Figure 2.3. Disregard of strength or direction of change, $\mathcal{D}_{ij}$ will indicate whether the E_i notices the interruption of S_j signaling node or not.
To expand the topic further, let us assume an example summarized in Figure 2.3. Since S_1 is a direct or undirect upstream of all other signaling genes S_2, S_3 and S_4, all receive silencing information on node S_1. In another example, the effect of perturbing S_2 will not affect S_3 and vice versa but will overlap as S_4 perceives both these two perturbations.
Consequently, the effect of genes ($E-$genes) receive perturbation effects from their direct signaling gene ($S-$genes) parent. This can be quantified on the level of expression of associated $E-$genes.

In NEMs, this simple concept had been extended to form a subset of relations on whole observed effects regarding perturbations which caused those. The set of $E-$genes that show changes following the perturbation of S_i are represented as $\mathcal{D}_i$ then:

$$S_i \rightarrow S_j \Leftrightarrow \mathcal{D}_j \subset \mathcal{D}_i \tag{2.1}$$

In essence, within the NEM framework, the assertion that S_j functions downstream of S_i can be inferred if the observable changes under j-th perturbation is subset of observable changes under i-th perturbation. This indicates that any downstream effects incited by a perturbation in S_j could also be initiated by intervening S_i.
The collective of $S-$gene, which transitions to state '1' when an $S-$gene S is silenced, defining the 'influence region' of S. The compilation of these influence regions is termed a silencing scheme. This scheme captures the predicted intervention effects based on the presumed pathway structure. In mathematical terms, it is conceptualized as a transitively closed graph that describes a partial order of $S-$genes based on the anticipated nested framework of downstream impacts. Adopting the assumed detailed by [25], $E-$gene's positions are taken as intrinsic model parameters. The underlying assumption is that an $E-$gene aligns exclusively with a singular $S-$gene.
Each perturbation is considered an action A_i that directly influences a specific $S-$node in the system. The perturbed $S-$node can, in turn, exert a causal influence on other $S-$nodes, forming a directed graph (or adjacency matrix) Φ of size $S \times S$, where an edge from node i to node j indicates that S_i has a causal effect on S_j. As we mentioned before, there is a second type of connections, called Θ (matrix $E \times S$), depicts how observations are linked to actions. [37] defines NEM as an effect model concluded as a by-product of these two parts:

$$F = \Phi\Theta \tag{2.2}$$

The main goal of NEM is to discover a causality structure Φ among $S-$genes with a subset relations between phenotypic profile best matching to observed $E-$genes [25, 37]. Leveraging Bayes' theorem, a particular network Φ can be evaluated and scored accordingly:

$$P(\Phi|\mathcal{D}) \propto P(\mathcal{D}|\Phi)P(\Phi) \tag{2.3}$$

$\Theta \in \{0,1\}^{m\times n}$ associates m $E-$genes to n $S-$genes. The fact that observable genes are independent from each other, given Φ and Θ, likelihood function can be rewritten as

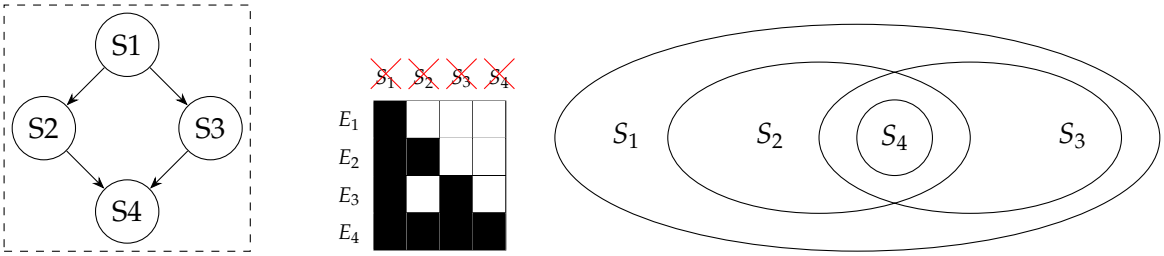

(a) (b) (c)

Figure 2.3: NEMs and the representation of perturbation effects - (a) The directed acyclic graph (DAG) illustrates causal relationships among S-genes (S1, S2, S3, S4), where perturbations propagate effects through the network. The binary effect matrix shows the effects of each experiment (column) in E-genes, with black cells representing observed effects. S1, the upstream regulator influences all E-genes, while S4 has the smallest effects. (b) The Venn diagram demonstrates the nested structure of effects, where S1 influences all other S-genes' effects, and S4's effects are a subset of the others. This hierarchy of nested effects forms the basis of Nested Effects Models (NEMs).

products of likelihood of each observation E_i:

$$P(\mathcal{D}|\Phi,\Theta) = \prod_{i=1}^{m} P(\mathcal{D}_i|\Phi,\Theta_i) \tag{2.4}$$

This independence assumption is a key simplification that allows the likelihood to factorize, making the computation more tractable. However, in real biological systems, dependencies between observable genes may exist, and deviations from this assumption should be considered when interpreting results. As all Θ_is are also independent,

$$P(\Theta|\Phi) = \prod_{i=1}^{m} P(\Theta_i|\Phi) \tag{2.5}$$

This independence assumption simplifies the inference process by allowing the posterior over Θ to be computed separately for each $E-$gene. It is based on the assumption that the connections between $S-$genes and $E-$genes $(S \rightarrow E)$ are independent for each $E-$gene given Φ. In other words, given the causal structure Φ, assigning of an $E-$gene to an $S-$gene does not influence the assignment of other $E-$genes. While this assumption reduces computational complexity, it is important to recognize that dependencies between different $E-$genes might exist in real biological systems due to shared regula-

tory mechanisms. Ultimately for all experiments k on $S-$nodes, marginalized likelihood over all parameters Θ is,

$$\begin{aligned} P(\mathcal{D}|\Phi) &= \int_{\Theta} P(\mathcal{D}|\Phi,\Theta)P(\Theta|\Phi)d\Theta \\ &= \prod_{i=1}^{m}\sum_{j=1}^{n}\prod_{k=1}^{n} P(\mathcal{D}_{ik}|\Phi,\Theta_{ij}=1)P(\Theta_{ij}=1) \end{aligned} \tag{2.6}$$

[38] allows the model to reflect the degree of belief in the existence of an edge. Generally, including proper priors helps the Bayesian models to avoid overfitting by providing regularization [39]. The prior matrix of all edge probabilities is an $n \times n$ matrix $\hat{\Phi}$,

$$P(\hat{\Phi}) = \prod_{i,j} P(\hat{\Phi}_{ij}) \tag{2.7}$$

2.2.2 Likelihood function

The discussion around changes in likelihood models of NEMs is crucial because it reflects the ongoing efforts to adapt this powerful analytical tool to various experimental setups and data characteristics. This evolution in likelihood models allows NEMs to remain relevant and practical in analyzing regulatory networks across a broader range of experimental contexts. By examining these changes, we gain insight into how the model has addressed practical challenges and expanded its applicability, ultimately enhancing our ability to interpret and understand complex biological systems.

NEMs were originally formulated based on specific experimental design assumptions. The initial work [34] relied on the presence of both positive and negative controls alongside RNAi perturbation data. This setup allowed for a straightforward counting approach to measure observables, where effects could be quantified by comparing against these controls. They have modeled the expression levels as binary random variable E_{ik} for observation E_i under experiment k.

$$p\left(e_{ik} \mid \Phi, \theta_i = j\right) = \begin{cases} \begin{array}{cc} e_{ik}=1 & e_{ik}=0 \\ \hline \alpha & 1-\alpha \end{array} & \text{if } S_j = 0 \\ \begin{array}{cc} 1-\beta & \beta \end{array} & \text{if } S_j = 1 \end{cases} \tag{2.8}$$

So the probability of $e_{ik} = 1$ when its parent S_j is not in the silence region Φ region (of silenced hidden gene) is the type-I error (α). Similarly, the probability to observe $e_iik = 0$ even though there is an effect on parent node S_j is equal to type-II error (β). For the

binary data matrix D and uniform prior $\frac{1}{p}$ of E to S attachments, likelihood has been defined as:

$$P(D|\Phi) = \frac{1}{p^m} \prod_{i=1}^{m} \sum_{j=1}^{n} \prod_{k=1}^{l} p_{\alpha,\beta} \left(e_{ik} \mid \Phi, \theta_i = j \right) \tag{2.9}$$

However, this requirement poses significant limitations, as most real-world experiments typically feature only a single control condition relative to the treatment (RNAi effects). This discrepancy between ideal experimental design and common practice has led to challenges in applying the original NEM formulation to many datasets.

Beta Uniform Mixture (BUM) models represent an advancement in statistical approaches for analyzing gene expression data. These models address limitations inherent in previous versions of NEM. It operates on p-values derived from differential gene expression analyses, typically conducted against a single control group. Under the null hypothesis, p-values exhibit a uniform distribution, characterized by a flat probability density function (PDF). Conversely, under the alternative hypothesis, p-values demonstrate a skewed distribution with higher density near zero, decreasing monotonically as p-values increase. [40] formulated the overall probability density function as a mixture of these two distributions:

$$P(D_{ik}) = \gamma_k + (1 - \gamma_k) \cdot f_1(D_{ik}),\ \gamma \in (0, 1) \tag{2.10}$$

So $P(D_{ik}|\Phi, \theta_i)$ is a conditional value:

$$P(D_{ik} \mid \Phi, \theta_i) = \begin{cases} f_1(D_{ik}) & \text{if } \Phi \text{ predicts an effect} \\ 1 & \text{otherwise} \end{cases} \tag{2.11}$$

definition of f_1, as knockdown effect, remains yet ambigious. A subsequent refinement introduced by [38] is a three-component mixture:

$$f(D_{ik}) = \pi_{1k} + \pi_{2k}\text{Beta}(D_{ik}, \alpha_k, 1) + \pi_{3k}\text{Beta}(D_{ik}, 1, \beta_k) \tag{2.12}$$

where the summation of three mixture coefficients (π_{ck}) is equal to one. Since f_1 indicates knockdown effect, they have checked how the estimated f (obtained via EM algorithm) is far from the maximum uniform part of the BUM model ($\hat{\pi}$).

$$f(D_{ik}) = \frac{f(D_{ik}) - \hat{\pi}}{1 - \hat{\pi}} \tag{2.13}$$

2.2.3 Formulation of NEM as Bayesian Networks

NEMs are a crucial tool for understanding how perturbations in biological systems, such as signaling pathways, propagate across genes and proteins. Traditionally, NEMs provided a simplified framework that imposed certain restrictions to ensure computational efficiency. However, a more flexible and generalizable approach models NEMs within the framework of Bayesian networks, which not only justifies prior assumptions but also significantly expands the model's capacity for probabilistic inference, as explored in [26].

By framing NEMs as Bayesian networks, we can more effectively manage complex signaling hierarchies while accounting for both discrete and continuous data. This formulation facilitates efficient inference, significantly reducing the search space for determining the optimal network topology. The Bayesian approach not only provides flexibility in modeling the probabilistic relationships between signaling and effect genes, but also allows the integration of prior knowledge and real-world experimental data. As a result, this reformulation enhances both the interpretability of the model's outcomes and the computational feasibility of uncovering intricate gene interactions.

In a Bayesian network (BN), we describe the probabilistic relationships between random variables using a directed acyclic graph (DAG). As we point before, in the context of NEMs, the random variables represent two main categories: signaling genes ($S-$genes) and effect genes ($E-$genes). The $S-$genes are the direct target of perturbations, but they are typically latent (not directly observed). The $E-$genes are the observed effects of perturbations, making the system hierarchical, with $S-$genes influencing $E-$genes through intermediate probabilistic pathways. The overall system can be modeled as a Directed Acyclic Graph (DAG) where the vertices $V(G)$ consist of both the signaling genes S and the effect genes E, such that:

$$V(G) = S \cup E \tag{2.14}$$

This forms a two-layer structure where the $S-$genes influence the $E-$genes. The perturbation of a signal $s \in S$ leads to the observation of an effect $e \in E$. The model is designed so that an effect is observable if there exists a directed path from an active $S-$gene to that effect gene. The conditional probability of observing an effect given the state of its

parent $S-$genes is defined as:

$$P(x = 1 \mid \text{pa}(x)) = \begin{cases} 1, & \text{if any parent pa}(x) = 1 \\ 0, & \text{otherwise} \end{cases} \tag{2.15}$$

To compute the likelihood of the data D given a graph structure G, we need to assume that the data is generated independently for each gene. The likelihood under the NEM framework is given by the product of local probabilities across all $S-$genes and $E-$genes:

$$P(D \mid G, L) = \prod_{s \in S} \prod_{e \in E} P(D_{se} \mid e_{se}) \tag{2.16}$$

where D_{se} represents the observation of effect e caused by signal s. When maximizing this likelihood, we aim to infer the graph structure G that best explains the observed data. This task is typically framed within the context of Maximum A Posteriori (MAP) estimation, where we seek to maximize the posterior probability:

$$P(G, \Theta \mid D) = \frac{P(D \mid G, \Theta) P(G) P(\Theta)}{P(D)} \tag{2.17}$$

A simultaneous maximum a posteriori (MAP) for G and Θ can be formulated as [37] suggests,

$$\begin{aligned} (\hat{G}, \hat{\Theta}) &= \arg\max_{G,\Theta} P(G, \Theta \mid D) \\ &= \arg\max_{G} \left(\arg\max_{\Theta} P(D \mid G, \Theta) P(\Theta) \right) P(G) \end{aligned} \tag{2.18}$$

Given prior knowledge about the network topology and parameters Θ, the goal is to identify the most likely graph structure G and parameters Θ based on the data. The posterior probability of a graph G, marginalizing over Θ, is given by:

$$\begin{aligned} P(G \mid D) &= \int P(G, \Theta \mid D) d\Theta \\ &\propto P(G) \int P(D|G, \Theta) P(\Theta) d\Theta \end{aligned} \tag{2.19}$$

which, under certain assumptions, can be approximated by maximizing the product of the likelihood $P(D \mid G, \Theta)$ and the prior terms $P(G)$ and $P(\Theta)$.

In the Bayesian framework, we often assume independent priors for the graph topology and the parameters Θ. When observations are binary (active/inactive), beta priors can be used to model the local probabilities [41]:

$$P(D \mid G) \propto \prod_{j=1}^{N} \prod_{e \in E} \prod_{i \in \{0,1\}} \frac{\Gamma(N_{e,i} + \alpha_i)\Gamma(N_{e,i} + \beta_i)}{\Gamma(N_{e,i} + \alpha_i + \beta_i)} \tag{2.20}$$

where $N_{e,i}$ represents the number of observations for each binary state of the effect gene, and α_i, β_i are the parameters of the beta distribution. This formula encapsulates the core likelihood in the presence of prior information about the graph. For continuous observations, the Cooper-Herskovits formula, Equation (2.20), can be used, assuming a normal distribution for the observed data. This is particularly relevant when dealing with real-valued measurements from experiments:

$$P(D \mid G) \propto \prod_{e \in \epsilon} \prod_{i \in \{0,1\}} \prod_{e \in E} \sqrt{\left(\frac{v}{v + N_{e,k}}\right)} \frac{\Gamma\left(\frac{\alpha + N_{e,k}}{2}\right)}{\left(\beta + s_{e,k} + \left(\frac{v_{nek}}{v + N_{e,k}}\right)(\bar{x}_{e,k} - \mu)^2\right)^{(\alpha + N_{e,k})/2}} \tag{2.21}$$

This continuous adaptation extends the application of NEMs to a broader set of experimental data, making it more flexible.

2.2.4 Factor Graph

A factor graph is a bigraph used to represent the decomposition of a more complex global function into simpler and smaller local functions, making it a generalization of a Bayesian Network [42]. While Bayesian Networks represent conditional dependencies among variables, factor graphs allow for a more flexible representation of any factorizable function, including but not limited to probability distributions.
The original NEMs by [34] modeled relationships in a simple binary manner, where pairwise interactions were either present or absent. [43] extend this idea by incorporating a more detailed approach to model gene interactions, particularly distinguishing between activating and inhibiting regulations in biological pathways. This formulation uses a factor graph framework, which offers a more flexible representation of relationships between signaling genes ($S-$genes) and effect genes ($E-$genes). In Factor Graph Nested Effects Models (FG-NEMs), the factor graph structure introduces two types of factor nodes:

- Transitive factors (ω): These factors model the interactions between $S-$genes and

capture whether the relationships are transitive (activating or inhibiting).

- Interaction factors (ρ): These factors represent the connections between $S-$genes and $E-$genes.

It allow for six possible interaction modes between $S-$genes:

(1) $A \rightarrow B$ (A activates B)

(2) $A \dashv B$ (A inhibits B)

(3) $A = B$ (A is equivalent to B)

(4) $A \neq B$ (A and B do not interact)

(5) $B \rightarrow A$ (B activates A)

(6) $B \dashv A$ (B inhibits A)

The factor graph representation of NEMs expresses the likelihood of the observed data D, given a structure G as:

$$P(D \mid G) = \sum_{S} \prod_{\gamma_{eS_1}, \gamma_{eS_2}} P(\gamma_{eS_1}, \gamma_{eS_2} \mid G_{S_1,S_2}, \theta_{eS_1,S_2}) P(S_e, S_1 \mid \gamma_{eS1}) P(S_e, S_2 \mid \gamma_{eS_2}) \quad (2.22)$$

Here, the likelihood is defined in terms of the transitive factors ω and interaction factors ρ, which regulate how the data D relate to the hidden states and interactions between the genes. The prior probability of the S-gene network $P(G)$ is defined through the factor graph as:

$$P(G) = \prod_{\omega_S} \prod_{\rho_{S_1,S_2}} \omega(G_{S_1,S_2}) \rho(G_{S_1,S_2}) \quad (2.23)$$

This formulation indicates that each pairwise relationship between $S-$genes (S_1, S_2) is associated with a transitive factor (ω) that ensures the consistency of the interactions and an interaction factor (ρ) that defines whether the interaction is activating, inhibiting, or neutral.

To infer the optimal $S-$gene network that maximizes the posterior probability $P(G \mid D)$, NEM uses the maximum a posteriori (MAP) estimation:

$$\hat{G} = \arg\max_{G} P(G \mid D) = \arg\max_{G} \sum_{\Theta, H} P(G, \Theta, H \mid D) \quad (2.24)$$

where Θ represents how $E-$genes are attached to the network of $S-$genes and H represents the hidden states of $E-$genes (e.g., upregulated, downregulated, or no change). By applying the same assumptions as in the original NEM framework, we can simplify the MAP expression:

$$\hat{G} = \arg\max_{G} P(G) \sum_{\Theta,H} P(H \mid G, \Theta) P(D \mid H) \tag{2.25}$$

Since the data D consists of observations for many $E-$genes, each with a hidden state H_e, we can decompose this sum across all $E-$genes:

$$\hat{\Phi} = \arg\max_{\Phi} P(\Phi) \prod_{e \in E} L_e(\Phi) \tag{2.26}$$

Where $L_e(\Phi)$ is the marginal likelihood for each $E-$gene, conditioned on the network structure Φ. Each L_e is further factored into pairwise $S-$gene terms, making it computationally efficient to calculate.

2.2.5 Learning Networks in NEMs

To comprehensively evaluate a network hypothesis within the NEMs framework, one often resorts to the Bayes formula. As we declared in Equation (2.3), bayesian approach allows the integration of prior knowledge with observed data, accompaining a posterior probability that can then be utilized to score and rank different network models.
The key challenge in learning the structure of $S-$gene networks is that the number of possible network topologies grows exponentially with the number of genes, making exhaustive enumeration feasible only for very small networks. This combinatorial explosion necessitates the use of propper methods for exploring larger networks.
Another issue is *likelihood equivalence* that arises because NEMs represent subset relationships that are inherently transitive [34]. In other words, different networks that only differ in their transitive edges can not be distinguished by the standard Bayesian scoring scheme. Several approaches have been proposed to ease this limitation, but they often focus on estimating a single high-scoring network rather than differentiating between transitively equivalent structures.
To address the scalability issues, pairwise and triple search methods were introduced to restrict the search space and make network inference more tractable[44]. The pairwise search method focuses on inferring the relationships between pairs of $S-$genes. For each pair of genes S_1 and S_2, four possible interaction models are considered:

(1) $S_1 \rightarrow S_2$ (effects of S_1 are a superset of the effects of S_2),

(2) $S_1 \leftarrow S_2$ (effects of S_1 are a subset of the effects of S_2),

(3) $S_1 \leftrightarrow S_2$ (effects of S_1 and S_2 are indistinguishable),

(4) $S_1 \cdot S_2$ (the genes are unrelated).

For each pair, a Bayesian score is computed, and the maximum a posteriori (MAP) model M_{S_1,S_2} is selected. This approach enhances efficiency, allowing for fast inference even with a large number of genes. But it treats each pairwise edge independently, which limits the accuracy of the inferred network due to the potential introduction of false edges or missed connections in the presence of noise.

To address the limitations of pairwise search, [44] extends the analysis to triplets of $S-$genes in a method called the triple search. This method evaluates all possible interactions between three genes, considering 29 possible quasi-order structures (or transitively closed graphs) for each triplet. The goal is to compute the MAP model M_{S_1,S_2,S_3} for each triplet. Once the highest-scoring models for each triplet are identified, they are combined into a final graph using edge-wise model averaging. For each edge, a confidence score is computed based on how frequently it appears in the triplet models. The final graph contains edges with confidence scores exceeding a predefined threshold.

One of the key challenges in network inference is the presence of noise in the data, which can lead to false positives (FP) and false negatives (FN). For instance, the true relationship $S_1 \rightarrow S_3$ may be missed due to incomplete overlap in the perturbation effects of S_1 and S_3. Although triple search does not always guarantee a transitively closed graph, it is better equipped to handle such issues. It considers a more comprehensive set of possible interactions by evaluating triplets rather than pairs, resulting in structures that are closer to a quasi-order [44].

In other approach, [38, 37], had applied greedy hill climbing search in NEM context. Greedy hillclimbing search is an optimization algorithm that iteratively improves a solution by changing a single element of the current solution that results in the greatest increase in the objective function. It continues this process until no better neighboring solutions are found. This can potentially result in a local maximum rather than the global optimum.

In NEMs, this search begins with an empty initial graph and incrementally adds edges as the observed data better fits the updated graph.

The module network approach, introduced by [38], was also developed to address the challenges associated with inferring large networks in NEMs. The module network

method breaks down this complexity by dividing the network into smaller, manageable subgraphs, or modules. This approach begins by applying hierarchical clustering to the $S-$genes, meaning grouping $S-$genes that exhibit similar $E-$gene response patterns into one cluster. The rationale is that genes with similar effects on $E-$genes are likely to be positioned close to each other in the signaling pathway. Each module consists of no more than four $S-$genes. By limiting each module to four nodes, the network structure within each module can be thoroughly explored, identifying the optimal subnetwork for that module.
Once the optimal structure for each module is identified, the modules must be reconnected to form the complete network. [38, 45] proposed two strategies for this task. The first approach uses pairwise model testing, where connections between nodes from different modules are established by evaluating the likelihood of pairwise interactions. The second approach, utilizes the log-likelihood of the entire network, reconnecting modules by selecting links that result in the highest overall network likelihood.
By dividing the network into smaller, modular subgraphs, this method significantly reduces the computational burden of large-scale NEMs. Additionally, it aligns well with the biological reality of modularity in signaling pathways, where certain sub-networks function independently but are connected to form larger regulatory systems.

2.3 Temporal Dynamics in Nested Effects Models with DynoNEM

The motivation for including time dimension to NEMs stems from the fact that traditional versions of NEMs were analyzing biological networks from static perturbation data. So they fall short in capturing the temporal nature of biological processes, especially in systems with distinct time scales, such as transcriptional networks. These networks involve sequential regulatory effects that occur over varying time intervals, making static approaches insufficient for distinguishing between direct and indirect interactions or capturing feedback mechanisms.
Dynamic NEMs (D-NEMs) [46] partially addressed this gap by incorporating time delays into their models. However, they relied on computationally intensive methods, such as Gibbs sampling, and their approach to inferring upstream signaling times was less biologically grounded. DynoNEMs were introduced in [47], advances these efforts by integrating time-series perturbation data into NEMs through a probabilistic framework. This allows for a more biologically realistic and computationally efficient

analysis, enabling the modeling of temporal dynamics and providing deeper insights into regulatory networks.
If an experiment affects S_j at time t, the changes may appear in downstream E_i with a delay of Δt. This delay explicitly allows the model to capture chain of effects unfolded over time. The model operates on the principle of unrolling the hidden network Ψ over time, Influence from nodes i to node j transmited by delay Ψ_{ij}:

$$S_i(t) \rightarrow S_j(t + \Psi_{ij}),\ t = 1, \ldots, T - \Psi_{ij} \tag{2.27}$$

Unlike traditional NEMs, DynoNEM assigns weights to connections Ψ_{ij}, reflecting the time-dependent strength of influence. Also this model operates on continouse time-series data (p-values from differentialy expressed effects). Let $\mathcal{D}_i(t)$ denote the observed effect profile for E_i at time t under experiment k. Likelihood function was described as follow:

$$p(D|\Psi) = \prod_{i \in E} \sum_{s \in K} \prod_{k \in K} \prod_{t=1}^{T} p(D_{ik}(t)|\Psi, \Theta_{iS} = 1) \Pr(\Theta_{iS} = 1) \tag{2.28}$$

which θ describes the connections between S and E genes. To model perturbation response under alternativehypothesis, for an expected effect, dynoNEM use the BUM (Beta-Uniform Mixture) model from [40]. Additionally, DynoNEM incorporates a prior on the network structure Ψ to penalize larger time delays and promote sparse, biologically plausible networks.

$$P(\Psi|\nu) \propto \prod_{i,j} \exp\left(-\frac{|\Psi_{ij} - \hat{\Psi}_{ij}|}{\nu}\right)$$

DynoNEM differs from traditional NEM models by moving beyond the first-order Markov assumption, which only considers immediate interactions ($k = 1$). Instead, by having weighted edges in mimicing delay behavior, it allows for connections that span multiple time steps ($k > 1$), capturing longer-term dependencies. Such assumption is often missed in simpler models like Dynamic Bayesian Networks. Also, by integrating time-series data, DynoNEM overcomes issues with static NEMs, such as indistinguishable network structures and the inability to model feedback loops. This makes it both biologically more realistic and computationally practical.

2.4 Objectives of this Thesis

NEMs have been widely successful in learning latent, causal network structures in biological systems. However, they exhibit key limitations that need to be addressed to enhance their applicability and accuracy. The primary goal of this work is to develop a new class of probabilistic graphical models within the framework of Dynamic Bayesian Networks (DBNs), overcoming challenges inherent in NEMs.

First, NEMs assume a strictly deterministic propagation of perturbation effects through the hidden variable network. If a path exists from node A to node B, the model assumes that perturbing A will always perturb B. This simplistic view neglects compensatory effects, which are common in biological systems. These effects may arise due to nonlinearities or the activation of unknown system components that influence downstream nodes, such as B, making the current formulation biologically unrealistic in such cases.

Second, the current NEM approach represents hidden states (e.g., protein activities) as binary variables, where a gene knockdown corresponds to decoupling a specific hidden variable from its inputs and setting its activity to 0. This approach performs what are referred to as "ideal interventions" [27], assuming perfect and complete disruption. However, experimental perturbations in biology rarely achieve such precision, which limits the real-world applicability of the model.

Additionally, NEMs require each hidden variable (node) to be perturbed individually, restricting their utility in scenarios involving combinatorial perturbations. Real-world biological systems often involve simultaneous perturbations of multiple network components, either due to non-specific effects of drugs or the cumulative effects of genetic mutations, such as those in cancer. These combinatorial perturbations often provide richer insights into the true network structure, which NEMs fail to exploit effectively.

One of the key limitations of DynoNEM is its reliance on discrete time steps to represent delays in signal propagation. While this approach provides a simplified view of time-dependent interactions, it imposes artificial constraints on the temporal dynamics of biological systems. Biological processes, such as gene regulation and protein signaling, occur in continuous time, often governed by nonlinear and dynamic changes that cannot be fully captured by a discrete time-step framework.

2.5 Toward a Unified Framework for Dynamic and Causal Network Inference

To address the limitations of existing Nested Effects Models (NEMs) and their extensions like DynoNEM, we propose a new modeling framework that combines the strengths of State Space Models (SSMs) and Structural Causal Models (SCMs). SSMs provide a robust foundation for representing dynamic systems, particularly in capturing the temporal evolution of hidden states and their relationship to observable variables. However, SSMs lack an explicit causal framework, which is essential for understanding the underlying mechanisms driving system dynamics. SCMs complement this by explicitly modeling causal relationships and addressing questions of "what if" scenarios. Together, these approaches offer a comprehensive solution for reconstructing dynamic and causal networks in complex biological systems.
The following sections explore the foundational principles of SSMs and SCMs, highlighting their unique strengths and their integration into a unified framework that bridges temporal and causal modeling.

2.5.1 State-space models (SSMs)

State Space Models (SSMs) are a natural choice for modeling the evolution of dynamic systems because they capture how hidden states change over time and relate to observable data. These models rely on two fundamental equations: the state transition equation, which describes how the system's hidden states evolve over time, and the observation equation, which links these hidden states to observable outputs. The continuous-time dynamics are described by an ordinary differential equation (ODE):

$$\frac{dS}{dt} = f(S_t, U_t) \tag{2.29}$$

where S_t represents the hidden states at time t, and f defines how these states evolve dynamically. For systems modeled in discrete time, the state transition can be approximated by this first-order Markov process:

$$S_t = F_t S_{t-1} + B_t U_t + w_t \tag{2.30}$$

where F_t describes the state transition matrix, B_t incorporates external influences (U_t), and w_t captures process noise.
The relationship between hidden states and observed data can also be framed proba-

bilistically:

$$D(t) \sim g(S(t)) + v_t \tag{2.31}$$

where $D(t)$ are observed variables dependent on the hidden states $S(t)$, with g describing the probabilistic mapping. A linear case of this mapping, often used in SSMs, is given by:

$$D_t = H_t S_t + v_t \tag{2.32}$$

H_t is the observation matrix, and v_t represents observation noise.
SSMs also share a strong relationship with Probabilistic Graphical Models (PGMs) and in particular with Dynamic Bayesian Networks (DBNs). SSMs can be seen as a specific instantiation of discrete time DBNs, where the nodes represent state variables over time, and the edges encode temporal dependencies. The state transition equations in SSMs parallel the transition models in DBNs, while the observation equations link hidden states to observable variables, much like observation models in DBNs. However, not all DBNs are state space models. In particular, continuous time DBNs (CT-DBNs) [48] generalize DBNs to continuous temporal dynamics but do not necessarily conform to the structured latent variable framework of classical SSMs. This connection enables the use of inference techniques developed for DBNs, such as Kalman Filtering [49] for linear models and Particle Filtering for non-linear systems, to estimate hidden states in SSMs.

2.5.2 Structural Causal Models (SCMs)

While SSMs excel at modeling temporal dynamics and hidden state evolution, they do not inherently capture causal relationships between variables. In biological systems, understanding causality is crucial for deciphering the mechanisms driving dynamic processes. Structural Causal Models (SCMs) fill this gap by providing an explicit framework for representing and reasoning about cause effect relationships. Incorporating SCMs into the modeling framework enables researchers to move beyond correlations and temporal dependencies to uncover the causal structure of biological networks.
This integration not only complements our previous discussion on SSMs and DBNs but also extends our analytical capabilities to more nuanced and realistic representations of how systems evolve and respond to various influences over time.

2.5.3 Integrating SCMs and SSMs

Taking the basic idea from Structural Causal Models (SCMs), we can show our hidden system as a graph with S nodes and directed edges among the nodes as causal relations and a function which translates causality relations to outputs usable by observations [21].

However, real-world biological systems are dynamic. So, causal effects unfold over time rather than occurring instantaneously. SSMs extend SCMs by incorporating temporal dynamics, modeling how the states of the system evolve continuously and how causal effects propagate across time. To model this, the tempotal dynamics of the hidden layer can be defined by ODE in Equation (2.29). This ODE describes how changes in one hidden state causally propagate to others over time, reflecting the dynamic nature of biological systems.

Biological systems are often partially observable, meaning that we cannot directly measure the hidden states (S). Instead, observe downstream effects (D), which are probabilistically linked to the hidden states S, can be described by Equation (2.31). The causal structure provided by SCMs ensures that we can interpret g in terms of cause-effect relationships, while SSMs model how these effects are dynamically propagated.

Interventions play a key role in revealing the causal structure. When an external force (e.g., gene knockout, drug treatment) is applied, it modifies the dynamics: $\frac{dS}{dt}$. By comparing the observed downstream effects (D) before and after the intervention, we can infer the causal relationships and refine the dynamic model.

This unified approach lays the foundation for refining network inference techniques, as explored in the subsequent chapter.

3 Ordinary Differential Equation based NEMs

While it is challenging to determine causal dependencies from observational data in general, recent developments in structural causal models [21] show that causal inference is sometimes possible under specific assumptions (e.g., causal Markov condition, faithfulness, no hidden confounders). Nonetheless, interventional or time-resolved data can substantially reduce the reliance on such strong assumptions, making causal relationships easier to detect and validate. This thesis focuses on the benefits of time-resolved data—often the most direct and robust way to uncover causal structure.

It is well known that causal dependence of variables in general cannot be inferred from observational data alone. It first needs some additional assumptions that define causality in the probabilistic context [21]. Ideally, one would like to have interventional data, and even better, one would like to have time-resolved data. My thesis deals with the latter situation, which is best suited for revealing causal relationships between variables.

The main objective is an identification of the causal relations between a set of variables with continuous dynamics. $S = (S_i(t))_{i=1,\dots,n}$ which are not observable directly (hidden nodes).

3.1 General Structure

We distinguish between observables E and principally unobservable nodes S. The weighted graph between all nodes S can be represented via an adjacency matrix $W \in \mathbb{R}^{n\times n}$ with $n := |S|$. Connections between S and E nodes are represented by another rectangular adjacency matrix $\theta \in \{0,1\}^{n\times m}$ with $m := |E|$.

3.2 Modeling the dynamics of hidden nodes

The matrix $W = (w_{ij})$ is meant to determine the dynamics of $S(t)$ such that w_{ij} describes the effect of node S_i on S_j. Using an ordinary differential equation (ODE), an obvious choice would be:

$$\frac{dS}{dt}(t) = W^T S(t) \tag{3.1}$$

However, such a system has an exponential solution:

$$S(t) = \exp(Dt) * S(0) \tag{3.2}$$

for some matrix D. The (complex) eigenvalues $d_1, ..., d_n$ of D characterize the qualitative behavior of the system: If $|d_i| > 1$ for some i, the system will diverge (almost surely), which is not compatible with the behavior of a biological system. If $|d_i| < 1$ for all i, the system will converge to the zero solution, $S(t) = 0$, also not observed in biological systems. Even if $|d_i| = 1$ for all i (which, by the way, would require a an extensive tuning of W) the imaginary parts of the eigenvalues would lead to oscillatory behavior. Achieving $|d_i| = 1$ for each eigenvalue is a highly restrictive condition, requiring very careful or 'fine-tuned' choices of W. In practice, any slight change in W would typically push some eigenvalues inside or outside the unit circle, causing convergence or divergence of $S(t)$ rather than purely sustained oscillations.
While this can occur in linear systems (think of the cell cycle as the main pseudo-oscillatory process), the regulatory circuits we will investigate aim to maintain cellular homeostasis, i.e., they are designed to reach a steady-state. Therefore a standard linear ODE approach does not satisfy our needs.
As linear ODEs have proven useful in biological network reconstruction ([50, 51]), our goal is to design a differential equation system which is locally close to a linear ODE. To avoid exponential dynamics, we limit the influence of individual nodes $S_i(t)$ on other nodes by a sigmoid function $\sigma : \mathbb{R} \to [-1, 1]$,

$$\frac{dS}{dt}(t) = W^T \sigma(S(t)) \tag{3.3}$$

Specifically, we chose $\sigma(x) = \tanh(x)$. Based on our experience the actual choice of σ has little effect on the qualitative outcome. While this step limits the values of $\frac{dS}{dt}$, it is not enough to avoid divergence. We therefore add a dampening term $\alpha S(t)$, $\alpha > 0$, which

limits the values of S to a finite range and also dampens oscillatory behavior.

$$\frac{d\mathbf{S}(t)}{dt} = W^T * tanh(\boldsymbol{S}(t)) - \alpha S(t) \tag{3.4}$$

This equation that is at the heart of our model, mimics the rate of enzymatic reactions in Michaelis–Menten kinetics [52].
The tanh function (or sigmoidal/logistic functions) saturates at large positive or negative arguments, much like the Michaelis–Menten equation saturates at high substrate concentration. It is a convenient smooth approximation for "bounded" enzymatic or regulatory rates, ensuring that reaction velocities do not grow unboundedly with increasing substrate levels.

3.3 steady-state and attractors of hidden nodes

To begin with, it is important to highlight that Equation (3.4) implies that irrespective of the starting condition $S(0)$, the state vector $S(t)$ can take only values in a bounded range of $\mathbb{R}^n$. We show that $\|S(t)\|$ remains in a compact domain because if any $S_i(t)$ grows too large, the ODE's negative feedback pulls it back down (and vice versa). This ensures solutions are globally bounded.
To see this, let $|W| = \sum_{ij} |w_{ij}|$ and realizing that for $S_i(t) > |W|/\alpha$

$$\frac{dS_i(t)}{dt} \leq |W| - \alpha S_i(t) < 0 \tag{3.5}$$

and for $S_i(t) < -|W|/\alpha$

$$\frac{dS_i(t)}{dt} \geq |W| - \alpha S_i(t) > 0 \tag{3.6}$$

Thus, irrespective of $S(0)$, the trajectory $S(t)$ will approach the compact cube $[-|W|/\alpha, |W|/\alpha]^n$ monotonically in every coordinate and never leave it again. As a consequence, $S(t)$ has an attractor ([53]) (this implies that if the attractor is a point, it converges. if the attractor is a curve, it may run exactly on that curve, which means it oscillates, or it might approach the attractor, which means this is a genuine attractor.). This in turn implies that the ergodic limit

$$S_i(\infty) = \lim_{t\to\infty} \frac{1}{t} \int_0^t S_i(t')dt' \tag{3.7}$$

exists. We use this value to describe the (average) long term behavior of our system. Equation Equation (3.4) has one obvious steady-state solution, $S(t) = 0$. To find non-zero steady-state solutions S^*, setting the derivatives to zero leads to the condition

$$S^* = \alpha^{-1} W^T \sigma(S^*) \tag{3.8}$$

We search for steady-state solutions by minimizing

$$\left\| S^* - \alpha^{-1} W^T \sigma(S^*) \right\|_2^2 \tag{3.9}$$

by starting at some random state $S(0)$ and running an ODE solver long enough.
Comment: There might be several solutions, but rare. We take the steady-state solution with the largest value of S^* and call it as $\|S_\infty\|_2$. We merely restart as long as we do not find a solution different from 0.

3.3.1 Notes on steady-states of a model

In system theory, the steady-state of a dynamical system is considered as a state where the behavior of variables (state variables) without any external force or perturbation does not change over time. This aligns with the fact that the system tends to stay in configurations of a local or global optima of the energy landscape (minimal energy or maximal stability). Such optimas are characterized by the minimization of the system's gradient.
The gradient of function f at point p is denoted by:

$$\nabla f(p) = \begin{bmatrix} \frac{\partial f}{\partial s_1}(p) \\ \frac{\partial f}{\partial s_2}(p) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \tag{3.10}$$

While the gradient could be derived analytically, in this case, we relied on numerical solvers to find the roots of Equation (3.8) directly.
In the two dimensional examples of our dynamical system Equation (3.4), we have $S = (s_1, s_2)$. The function f is then the *"energy-like" (or Lyapunov-like) potential* whose gradient matches (up to sign) the right-hand side of the ODE. Concretely:

- s_1, s_2 are simply the two coordinates of the state vector $S \in \mathbb{R}^2$.
- $f(s_1, s_2)$ is the scalar function whose gradient satisfies :

$$\nabla f(S) \;=\; \alpha\, S \;-\; W^T \tanh(S).$$

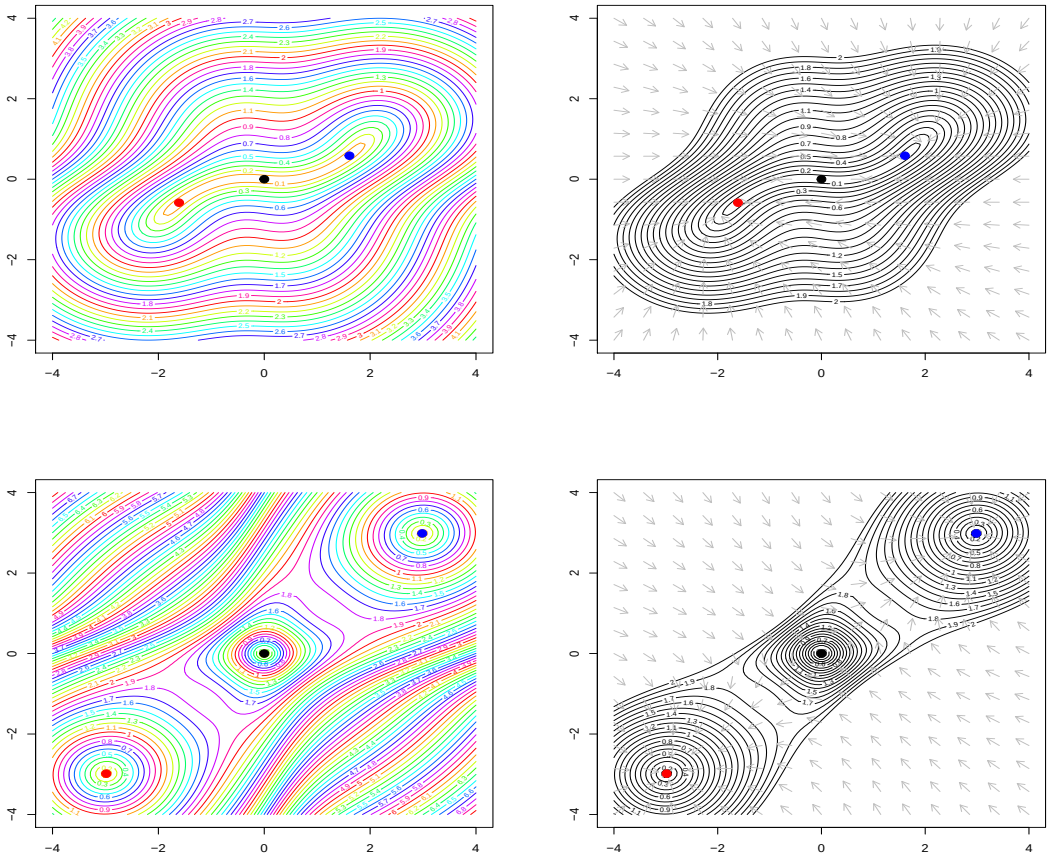

Figure 3.1: steady-states and Energy Landscape in Dynamical Systems described in Equation (3.4) for two different W matrices. The steady-states, represented as dots.

Setting $\nabla f(S) = \mathbf{0}$ recovers the steady-state condition

$$\dot{S} \;=\; 0 \;\iff\; W^T \tanh(S) \;-\; \alpha\, S \;=\; 0.$$

For two different W matrices, the minimum of gradient of our dynamical system Equation (3.4), indicative of steady-states, are depicted as dots in the contour plot of Figure 3.1. The contours plot depicts *level sets* of f. The dots (steady-states) are precisely those points where the system's gradient vanishes, i.e. where

$$W^T \tanh(S) \;-\; \alpha\, S \;=\; 0.$$

As you can see, under specific assumptions, different initial conditions subjected to the same constraints (function) can result in divergent dynamics, leading to various steady-states, marked by the dots in the figure. However, the symmetric behavior of *tanh* function causes our dynamics to have two attractors with the same magnitude but opposite signs.
It is necessary to mention that, in biological systems modeling we assume the system is initially at a stable steady-state; for computational stability, this initial state should not be in the vicinity of zero. Additionally, having negative state is not is not biologically meaningful for many modeling scenarios, as variables such as concentrations of proteins, metabolites, or gene expression levels cannot take negative values in real systems. So, when the ODE of Equation (3.4) starts from any sufficiently large positive value, it will settle into one positive steady-state also.

3.4 Modeling interventions of the system

Our main objective is to find causal relations between S nodes as it describes our observations in E. To bridge the gap between our ODE model described in Section 3.2 and our overarching goal of learning the causalities, we draw upon the concept of a Structural Causal Model (SCM). As we discussed in Section 2.5.2, an SCM is a framework that enables us to infer causal relationships by introducing interventions on variables within a system. In this context, we can view a causal model as equivalent to an ODE, where any causal connections between variables can be interpreted as non-zero coefficients within the ODE and vice versa. This connection serves as the foundation for our approach.

In designing our interventions, we introduced a simplifying assumption for the sake

of clarity. We assumed that the flow of information from direct targets of interventions was consistently interrupted by perfect inhibition. This implies that the state transitions of such target nodes become constantly zero, irrespective of the dynamics of any other related S nodes. However, the state of child nodes (direct or indirect) of the inhibited targets in our causality structure is influenced by this inhibition. By this description, we now define our objective as observing the deviations in the state of a variable S compared to its steady-state due to the excitation of another hidden variable to gain insights into causalities.

In the language of an SCM, these 'perfect inhibitions' correspond to do-interventions on specific nodes. Concretely, if we denote a node $\tilde{S}$ as a direct target for intervention, then we can write do$(\tilde{S} = 0)$ to indicate that $\tilde{S}(t)$ remains fixed at zero regardless of the usual ODE dynamics. Equivalently, from the ODE perspective, we remove $\tilde{S}$'s usual update equation and force $\frac{d\tilde{S}}{dt} = 0$ and $\tilde{S}(t) = 0$ for all t, reflecting a perfect inhibition. This manipulation severs any causal influence of $\tilde{S}$ on its children, allowing us to observe how downstream states in the network respond to (or no longer receive) inputs from $\tilde{S}$. For nodes not in the path from the inhibited target, their activity levels remain unchanged, as they would in a steady-state scenario. Although depletion may cause a slight decrease in activity levels, this effect is considered negligible. This approach allows us to distinguish between changes in activity levels resulting from depletion and those caused by inhibitions within our model.

When imposing an experimental perturbation j to the system, we model this by replacing W in Equation (3.4) by a matrix W^j. Here, W^j is obtained from W in a canonical way: When $S-$node i is perturbed in experiment j, the i-the column of W, corresponding to the inputs of node i, are set to 0. This obviously leads to the silencing of node i ($S_i(t) \to 0$ as $t \to \infty$), and $S_i^{j*} = 0$ for every (pseudo) steady-state of in experiment j. Denote by $S^j(t; W^j, \alpha)$ the dynamics of the system under perturbation/experiment j, given by equation Equation (3.4):

$$\frac{dS^j}{dt}(t; W^j, \alpha) = (W^j)^T \sigma(S^j(t)) - \alpha S^j(t) \tag{3.11}$$

By $\mathcal{W} = (W^1, ..., W^V)$, we denote the set of matrices encoding all perturbation experiments. By this definition, $s_k^j(t)$ will be state of hidden variable k under j experiment.

For W^j, Equation (3.9) can be rewritten as:

$$S^{j*} = \alpha^{-1} {W^j}^T \sigma(S^{j*}) \tag{3.12}$$

Regarding $S^0 = S(0)$ as the basic steady-state of hidden variables in control situation (no intervention), we run ODE solver long enough to find roots $S^j(\infty)$ for Equation (3.12). So $s_k^j(\infty)$ is the ergodic limit for variable k under experiment j as $t \to \infty$.

3.5 Linking hidden nodes to observables

Recall that $\theta_{ki} \in \{0, 1\}$ indicates whether hidden node k is a parent of the observabe i, meaning that changes in k will affect i. We assume that each observable can have exactly one parent. The parent of $i \in E$ is denoted by $\mathrm{pa}(i) \in S$. The matrix $\theta = (\theta_{ij})$ is used to propagate the dynamic changes in the hidden network to the observables.
Let $o_i^j(t) \in \mathbb{R}$ denote the observation made at node i, $i \in E$, under some experimental condition j, $j \in \{1, ..., V\}$, at time t. The full data therefore is $D = \{o_1^1, ..., o_i^j ...\}$. We assume that there is a "control" or "null" condition in which the system is in its basic steady-state. Upon perturbation, we are interested in modeling the relative deviation from this basic state. Therefore, we take $o_i^j(t)$ as the absolute log fold change of the value $o_i^j(t)$ over its basic state value. When doing so, small positive values will indicate small effects, whereas strong effects will lead to large positive values. As there will be random fluctuations of our measurements, be it due to technical limitations or to biological variation, we assume that the observations $o_i^j(t)$ are drawn from a mixture distribution

$$o_i^j(t) \sim (1 - \gamma_i^j(t)) \cdot f_0^j + \gamma_i^j(t) \cdot f_1^j \tag{3.13}$$

where f_0^j is the distribution modeling random fluctuations and f_1^j models the distribution of relevant changes. These distributions can depend on the experimental condition j. For simulation study, we use normal distributions for both distributions that were truncated at 0,

$$f_0^j \equiv \mathcal{N}_+(0, \sigma_0^j) \tag{3.14}$$

$$f_1^j \equiv \mathcal{N}_+(\mu_1^j, \sigma_1^j) \tag{3.15}$$

The mean and variance parameters are estimated from the data (see in Section 3.6). Let $\Gamma = \{\sigma_0^j, \sigma_1^j, \mu_1^j; j \in \mathcal{E}\}$.
The mixture coefficient $\gamma_i^j(t) \in [0, 1]$ is a function of the activity of the (hidden) parent

$k = \text{pa}(i)$ of i:

$$\gamma_i^j(t) = \tanh \frac{\left| s_k^j(t) - s_k^0 \right|}{\left| s_k^j(\infty) - s_k^0 \right|} \tag{3.16}$$

Notable, the magnitude of perturbation probability $\gamma_i^j(t)$ corresponds to a higher likelihood of perturbation of S node at a time. In other words, we want to see how dynamical state of node $k = \text{pa}(i)$ under experiment j is far from the wild-type steady-state s_k^0, and normalize it with the distance of two steady-states (under the corresponding experiment vs wild-type). Recall W, Θ, Γ , where Γ is the set of all parameters of the mixture model in all experimental conditions $\Gamma = \{\sigma_0^j, \sigma_1^j, \mu_1^j; j \in \mathcal{E}\}$. Recall that θ_{ik} is the probability that $S-$node k is the parent of $E-$node i. Then

$$P(D \mid \mathcal{W}, \alpha, \Theta, \Gamma) = \prod_{j=1}^{V} \prod_{i \in \mathcal{E}} \sum_{k \in \mathcal{S}} \theta_{ik} \prod_{t=1}^{T} p(o_i^j(t) \mid S_k^j(t; W^j, \alpha), \Gamma) \tag{3.17}$$

The concerns of numerical underflow of Equation (3.17) is described in Section .0.

3.6 Parameter Learning

The assignment probabilities : Θ need to be specified beforehand, by default we choose a uniform prior $\theta_{ik} = 1/|\mathcal{S}|$. In the presence of prior knowledge, this prior can be exchanged by an informative prior.

Estimating Parameters of the mixture model : The mean and the variance parameters Γ are estimated from the data using an EM algorithm. The Expectation-Maximization (EM) algorithm is a powerful iterative method that is well-suited for scenarios involving latent variables or incomplete data to estimate the parameters of the distributions. In our modeling, it allows us to iteratively estimate the parameters of the truncated normal distributions that model the random fluctuations (f_0) and the relevant changes (f_1) in the data.

EM Algorithm Steps:

- Initialization: Start with initial guesses for the parameters σ_{0j}, μ_{1j} and σ_{1j}. These can be derived from simple statistics of the data, such as the overall mean and

variance, or by using prior knowledge.

- E-step (Expectation step): In this step, the posterior probability that i-th data point belongs to the f_1 by estimated parameters j will be updated.

$$\tau_{ij} = \frac{\pi_j f_{1j}(x_i)}{\pi_j f_{1j}(x_i) + (1-\pi_j) f_{0j}(x_i)} \tag{3.18}$$

  where π_j is the mixing factor.

- M-step (Maximization step): In this step, we update the parameter estimates by maximizing the expected log-likelihood computed in the E-step. The updates for the parameters, for f_0 and f_1 (Equation (3.14) and Equation (3.15)), are given by:

$$\hat{\mu}_{1j}^{(t+1)} = \frac{\sum_{i=1}^{N} \tau_{ij} x_i}{\sum_{i=1}^{N} \tau_{ij}} \tag{3.19}$$

$$\hat{\sigma}_{1j}^{(t+1)} = \sqrt{\frac{\sum_{i=1}^{N} \tau_{ij} \left(x_i - \hat{\mu}_{1j}^{(t+1)}\right)^2}{\sum_{i=1}^{N} \tau_{ij}}} \tag{3.20}$$

$$\hat{\sigma}_{0j}^{(t+1)} = \sqrt{\frac{\sum_{i=1}^{N}(1-\tau_{ij}) x_i^2}{\sum_{i=1}^{N}(1-\tau_{ij})}} \cdot \sqrt{1 + \frac{\phi\left(\frac{a}{\sigma_{0j}^{(t)}}\right)}{1-\Phi\left(\frac{a}{\sigma_{0j}^{(t)}}\right)} \left(\frac{\phi\left(\frac{a}{\sigma_{0j}^{(t)}}\right)}{1-\Phi\left(\frac{a}{\sigma_{0j}^{(t)}}\right)} - \frac{a}{\sigma_{0j}^{(t)}} \right)} \tag{3.21}$$

- Iteration: Repeat the E-step and M-step until convergence, which is typically determined when the changes in the parameter estimates between iterations fall below a predefined threshold.

Derivation of the $\hat{\sigma}_{0j}^{(t+1)}$ is included in Section .0.

Learning $\mathcal{W}$ and α via MCMC: The parameters of the dynamic system, $\mathcal{W}$ and α , are learned by a Markov Chain Monte Carlo (MCMC) algorithm. Note that for learning $\mathcal{W}$, only W needs to be learned, as the matrices W^j in $\mathcal{W}$ are constructed from W using a fixed scheme.
One possibility to learn W is to sample from the posterior distribution $P(W \mid D, \mathbf{q})$.

This can be achieved via Metropolis-Hastings Markov-Chain Monte Carlo (MCMC) sampling. More specifically, we define the following possible MCMC moves for any uniform randomly selected edge $i \to j (i \neq j)$:

(i) insert new edge $i \to j$, i.e. sample weight w_{ij} from $N(\nu, \tau)$

(ii) delete edge $i \to j$, i.e. set $w_{ij} = 0$

(iii) swap weights of two edges $i \to j$ and $j \to i$

(iv) modify edge weight, i.e. $\log w_{ij} \mapsto \log w_{ij} + N(0, \zeta^2)$

Each random move from these four moves is chosen uniformly. Each MCMC move (resulting in a modified weight matrix W') is accepted with probability

$$\min\left(1, \frac{p(D \mid W', Q)p(W' \mid \boldsymbol{\rho})}{p(D \mid W, Q)p(W \mid \boldsymbol{\rho})}\right) \tag{3.22}$$

as Metropolis Hasting ratio.
Since the true $W = (w_{ij})$ will typically be non-identifiable from limited data, it is relevant to define a prior distribution. Here we use an edge-wise prior motivated by altered defenition from the spike and slab [54], which allows for sparse models and at the same time allows for incorporating prior knowledge, if available:

$$w_{ij} \mid \rho_{ij} \sim (1 - \rho_{ij})\delta_0 + \rho_{ij} \tag{3.23}$$

where δ_0 denotes a point mass at zero with $Lap(0, \beta)$ distribution. The prior can be understood as follows:
ρ_{ij} defines the degree of certainty about the existence of edge $i \to j$. A larger value increases the chance for $w_{ij} \sim Uniform(0, 1)$.
That means the edge weight w_{ij} is given a higher chance to be different from zero. On the other hand, a small ρ_{ij} increases the chance for $w_{ij} = 0$. Hence, our prior enforces sparsity for edges with low a priori confidence and increases the probability of non-zero weights for high confidence edges.
Note that due to edge-wise independence $p(W \mid \boldsymbol{\rho}) = \prod_{ij} p(w_{ij} \mid \rho_{ij})$.

3.7 Simulation Studies

We evaluated the performance of our proposed model in Equation 3.4 and Equation 3.17 by reconstructing networks of different structures and complexity by various parameters.

This evaluation process was methodically structured to assess the model's robustness and accuracy across a spectrum of network configurations. In the first part of our eval-

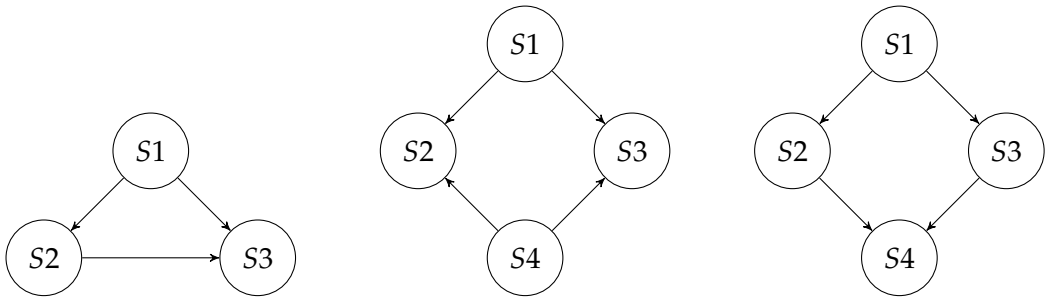

(a) Feed forward (b) Bifan (c) Diamond

Figure 3.2: Motif Structures

uation, we began with simple network patterns known as motifs, which are common in real-world networks. These motifs which are shown in Figure 3.2, are like building blocks in the study of complex networks. Starting with them helped us understand how well our model could reproduce basic interaction patterns, providing valuable insights into network principles.
To further evaluate our model on larger structures, we extended our analysis to subnetworks of KEGG pathways [55]. Specifically, we used 5 and 10-node graphs, as defined in the supplementary material of [47], to allow a one-to-one comparison of odeNEM results with dynoNEM. This transition from small motifs to larger, biologically relevant networks allowed us to carefully assess how effectively odeNEM handled more intricate network structures and parameter settings.

Importantly, while dynoNEM results were referenced from the original publication, the method was not rerunnable, which restricted direct re-validation of its performance against our approach.

Having graph structures, related data has been generated also. As we described in Section 3.4, perfect inhibition of a direct target node can simulate an RNAi experiment. This idea were also applied to simulate knockdown of a hidden node ($S-$gene) in our sampled networks.
For a W matrix matching to a motif structure and starting from steady-states as initial values for the hidden variables, which are set sufficiently away from zero to avoid computational instabilities, we utilize Equation (3.11) to capture the dynamics of a hidden

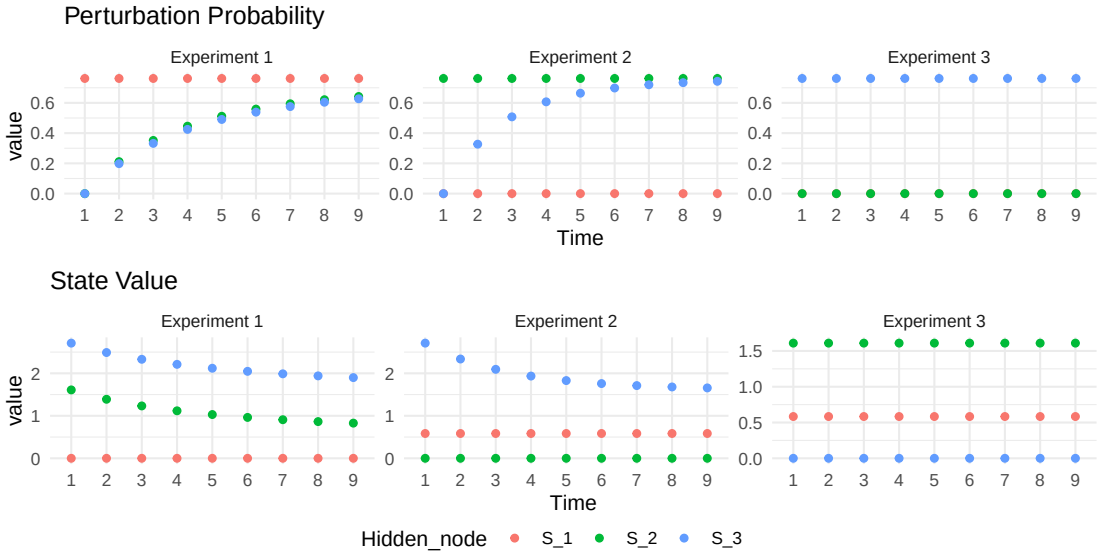

Figure 3.3: Dynamics of hidden layer in feed forward network

layer under inhibition of each $S-$nodes. These dynamics levels have been translated to perturbation levels γ, by Equation (3.16), to generate data with our mixture model Equation (3.13). An extensive example of state levels and perturbation levels of feed forward motif network is provided in Figure 3.3. This captures how odeNEM capture expected hidden dynamic behavior upon each experiments correctly. Experiments in our simulation studies are single complete inhibition of a node. Figure 1 and Figure 2 are also showing the same in supplementary.

To assess the convergence of the MCMC chains, we performed multiple diagnostic checks, including trace plots, likelihood convergence, and autocorrelation analysis. The trace plots were examined to ensure the stability and proper mixing of the chains over iterations, indicating stationarity. Likelihood convergence was analyzed to confirm that the overall posterior likelihood stabilized, demonstrating the chains had reached the target distribution. Additionally, autocorrelation was evaluated to verify the independence of samples and the effective exploration of the parameter space. Together, these checks provide a validation of the convergence and reliability of the MCMC results.

To mitigate the inherent randomness in the MCMC process, we have averaged the simulation and inference results over five iterations for each network configuration. All computational runs described in this chapter, as well as the subsequent chapters, were performed on the CHEOPS cluster using the batch job scheduler, supported by the Regional Computing Center (RRZK) at the University of Cologne.

3.7.1 Evaluation Metrics

The performance of our proposed propagation model in capturing the structural properties of true networks is first validated by using two statistical measures:
Area under the Receiver Operating Characteristic curve (AUC-ROC) : The ROC curve illustrates the relationship between the True Positive Rate (TPR) and the False Positive Rate (FPR) across different threshold values. The TPR and FPR are calculated as follows:

$$TPR = \frac{TP}{TP + FN} \tag{3.24}$$

$$FPR = \frac{FP}{FP + TN} \tag{3.25}$$

where TP is the number of true positives, FN is the number of false negatives, FP is the number of false positives and TN is the number of true negatives.
The AUC-ROC is the area under this curve, measuring model's ability to distinguish between classes and is particularly useful in situations where there is an imbalance between the classes.
Area under the Precision-Recall curve (AUC-PR) : This curve plots Precision against Recall for different threshold values. Precision is defined as follows:

$$Precision = \frac{TP}{TP + FP} \tag{3.26}$$

$$Recall = \frac{TP}{TP + FN} \tag{3.27}$$

and Recall is the same as TPR in the ROC curve.
AUC-PR provides insight into the model's performance in terms of both precision (the proportion of true positive results among all positive predictions) and recall (the proportion of true positive results detected among all actual positives) specifically in cases of imbalanced datasets where one class is much more prevalent than the other.
Negative Predictive Value (NPV) and Positive Predictive Value (PPV) : NPV and PPV can be interpreted as accuracy measures for a diagnostic test which are formulated as:

$$NPV = \frac{TN}{TN + FN} \tag{3.28}$$

the probability that subjects with a negative screening test truly don't have the edge in our gold standard network.

$$PPV = \frac{TP}{TP + FP} \tag{3.29}$$

the probability that subjects with a positive screening test truly have the edge in our gold standard network. Furthermore, the analysis of the quantiles within the sampled distributions assists in identifying which elements of the weight matrices are likely to represent an edge in the inferred causal model. Quantiles are values that partition a probability distribution into equal-sized, continuous intervals. The q_{th} quantile of a distribution is typically based on the cumulative distribution function (CDF).
CDF Calculation : For a given distribution with $F(x)$ as CDF, the q_{th} quantile is defined such that

$$F(x_q) = q \tag{3.30}$$

For instance, the median is the $50 - th$ percentile, where $q = 0.5$. The median is a resilient measure of central tendency, offering greater resistance to outliers and skewed distributions than the mean. So in our evaluations, a variable whose $50 - th$ percentile value is zero in the sampled distribution may suggest a weaker causal relationship, providing a pivotal clue in our analysis of causal models.

The most robust conclusions are drawn from analyses using High Posterior Density (HPD) intervals. HPD offers a quantifiable measure of the uncertainty or confidence in our parameter estimation based on the posterior distribution. This way offers a Bayesian counterpart to frequentist confidence intervals.
For example, an HPD interval with credible level of 95%($\alpha = 0.05$) covers range of values for the parameter that contains 95% of the probability mass. In more general definition, HPD interval for a given probability p is calculated by finding the smallest interval where the probability of the parameter being in this interval is at least p.
HPD Calculation : Mathematically, for a continuous parameter θ and its posterior distribution $P(\theta \mid D)$, an HPD interval $[a, b]$ satisfies:

$$P(a \leq \theta \leq b \mid D) = p \tag{3.31}$$

For all intervals $[c, d]$ where $P(c \leq \theta \leq d \mid D) = p$, it holds that $b - a \leq d - c$. In simpler terms, $[a, b]$ is the narrowest interval containing the specified proportion p of

the posterior distribution.

3.8 Evaluations

To have an initial validations of our model, we examine whether the temporal trajectories of the hidden states align with the anticipated behavior when subjected to inhibition of a single node. Figure 3.3 and subsequent Figures in supplementary illustrates the model's response over time to specific modifications (here is the complete inhibition of a node in each experiment). These trajectory analyses offer an initial confirmation that our model accurately captures the dynamics we expect to observe.
Notably, in Figure 3.3, second experiment which inhibits node S_2, leads to its inactivity. This is evidenced by a zero-state value. In such one-directional information flow system, the inhibition of S_2 results in a decreased activity level of S_3, while S_1 remains stable maintaining its steady-state value. It is also relevant here to mention that simulations are extended over prolonged periods to capture the true effects as some effects can develop and stabilize more slowly.

HPD Validations

Validation of MCMC results using HPD offers a more generalized perspective on estimated parameters by providing confidence intervals. This approach primarily emphasizes the range of the most likely values, rather than the detailed shape of the posterior distribution. In our evaluations, we consider an HPD interval "narrowly centered around zero" if it (i) has a small width (i.e., low posterior uncertainty) and (ii) includes zero as its midpoint or close to it. This signals that the data strongly support a parameter value near zero, implying no substantial effect. Conversely, an HPD interval that is "sufficiently broad and far from zero" reflects two aspects: the interval does not contain zero (indicating a likely nonzero effect) and its greater width denotes higher uncertainty regarding the effect's exact magnitude. Thus, both the position of the HPD relative to zero and its width inform us whether a parameter is credibly nonzero—and how certain we can be of its estimated value. Figure 3.4 demonstrates how increasing the number of time points improves our model's ability to reconstruct the underlying feed forward network. As t grows from 2 to 5, 10, and finally 20, the posterior distributions for the true edges (in the upper triangle) move further away from zero, indicating stronger confidence in the presence of these connections. Simultaneously, the distributions for the absent edges (in the lower triangle) tightly concentrate around zero, reflecting high certainty that no edge

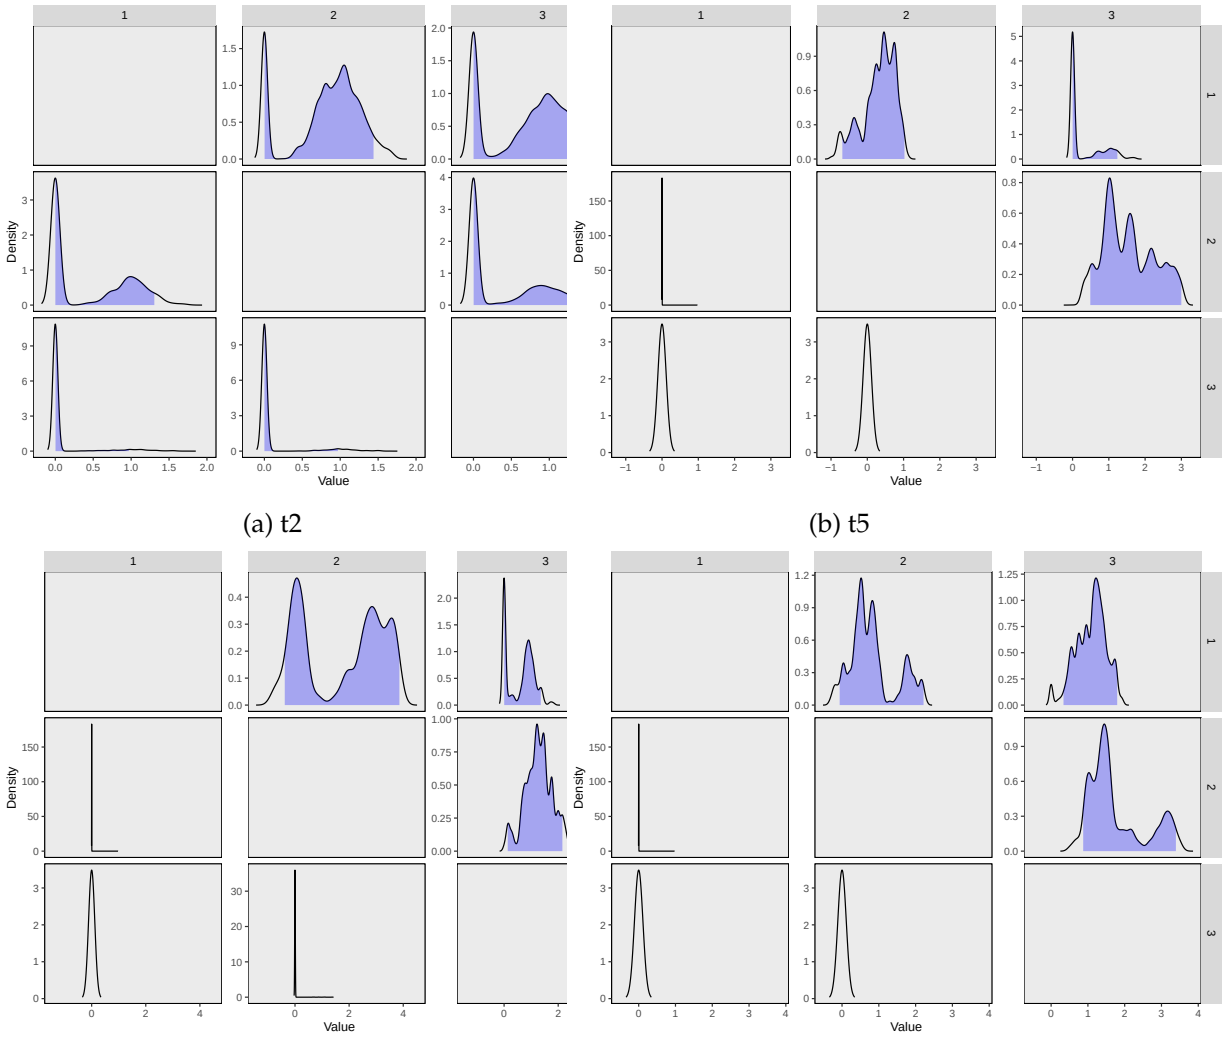

(a) t2

(b) t5

(c) t10

(d) t20

Figure 3.4: Effect of increasing the number of time points ($t \in \{2, 5, 10, 20\}$) on reconstructing feed forward networks. Each panel (a)–(d) corresponds to a specific $t-$value and is arranged in a matrix format where the rows represent the source nodes and the columns represent the target nodes. Consequently, subfigure (i, j) in each panel shows the posterior density of the estimated weight for the edge $i \to j$.
Within each cell, the shaded area corresponds to the highest posterior density (HPD) region, indicating the most credible values of that edge's weight. In the lower triangle cells (e.g., $(3, 1)$), the true network specifies no direct edge, so we expect the posterior distribution to center around zero. In the upper triangle cells (e.g., $(1, 3)$), the network has a true edge, so we expect the posterior to deviate from zero.

should be inferred there. By $t = 20$, most true edges show posterior densities whose HPD intervals clearly exclude zero, signifying robust detection. Conversely, edges that are not in the true network remain near zero with narrow HPDs, thus reinforcing the model's capacity for sparse, accurate reconstruction. Overall, this pattern confirms that collecting more time points significantly boosts the reliability of inferred edge weights and supports the feasibility of our ODE based approach for network inference.
While the HPD figure effectively illustrates the ability to detect both present and absent edges, it does not fully capture overall model accuracy or predictive power. Metrics such as AUC, PPV, and sensitivity (as discussed in Section 3.7.1) complement HPD by offering a broader evaluation of performance. Together, these insights guide model refinement and improvement in predictive accuracy.

In following, we did a comprehensive performance analysis across various metrics. To cover the vast effects of different parameters on performance, simulations have been done under various settings as follows.

Hidden nodes

Number of $S-$genes play a prominent role in describing how our model can perform. We examined performance of odeNEM in the reconstruction of a sample network with $n = \{3, 4, 5, 10\}$ hidden nodes. The ground truth networks corresponding to these cases are described in Section 3.7. To observe the true effect, other parameters were set to constant values (timepoints=20, number of $E-$genes=20). The chain of networks sampled by MCMC after burn-in time and thining filters, have been evaluated by metrics described in Section 3.7.1. As it is obvious in Figure 3.5, increasing number of hidden nodes comes with the cost of sensitivity reduction. Specifically, adding more nodes seems to reduce the precision (as seen in PPV) and increase the variability in sensitivity and specificity, indicating that larger networks may introduce more complexity, potentially leading to more varied and less predictable performance outcomes.
When examining the effect of increasing the number of hidden nodes on AUC and AUPR metrics, it becomes clear that the model's performance decreases as the network grows larger. For smaller networks with 3 or 4 nodes, the AUC is nearly perfect at 1.0, indicating the method's strong capacity to distinguish between positive and negative cases. However, with 10 hidden nodes, the AUC drops below 0.75. This is a marked decrease in the model's overall classification ability.
A similar pattern emerges for AUPR. It remains relatively high in smaller networks (3

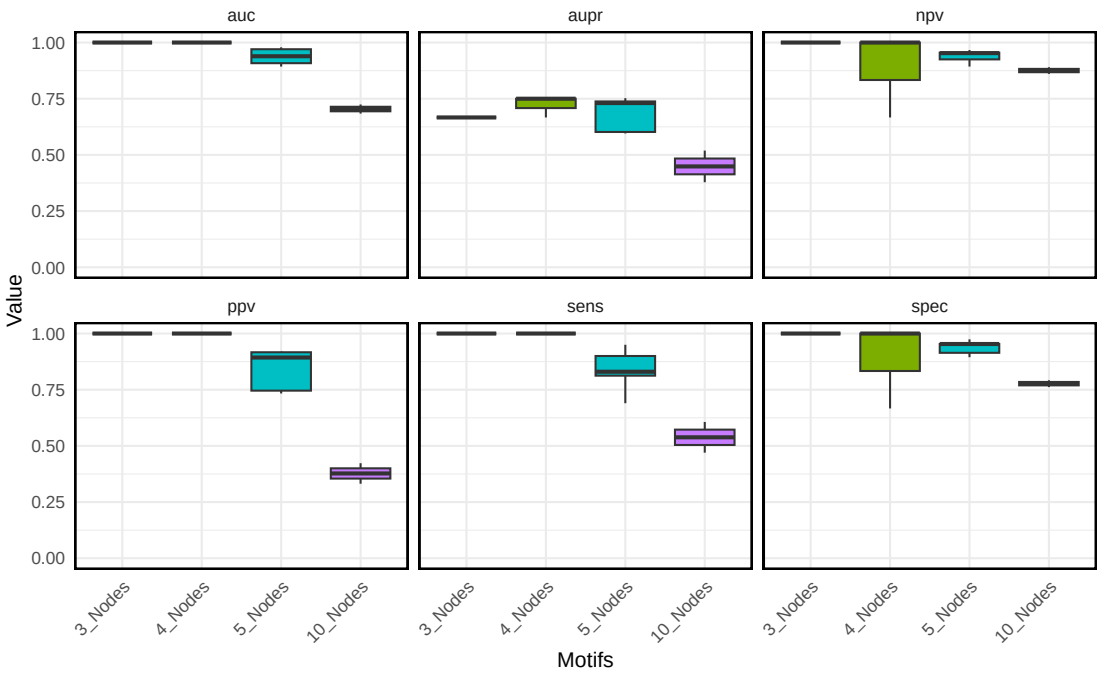

Figure 3.5: Motifs' Reconstruction Evaluations with Different Hidden Nodes

and 4 nodes), but declines substantially as the number of hidden nodes increases. The limited data set (20 observations) cannot sufficiently support the complexity introduced by a larger network, making it harder to maintain high precision recall performance. One major factor is the high ratio of free parameters to available data, which can degrade inference. In a more predictive context, such a scenario might be referred to as overfitting where the model appears to capture patterns that do not generalize. However, since we are inferring network structure (rather than a predictive function), it is more accurate to say that the increased model flexibility leads to identifiability challenges. Many plausible networks may fit the limited data equally well, diluting true connections and introducing false edges. This mechanism can lower precision (due to false positives) and reduce sensitivity (due to increased false negatives).

These observations underscore the bias variance trade off in statistical modeling. A larger network can capture more complex relationships (thus lowering bias). However, with insufficient data or regularization, this complexity also increases variance. As a consequence, parameter estimates (edges and weights) can become unstable and performance metrics such as sensitivity, PPV, or AUPR will decrease. Finding an optimal

network complexity thus depends critically on data quantity and quality.

Effect nodes

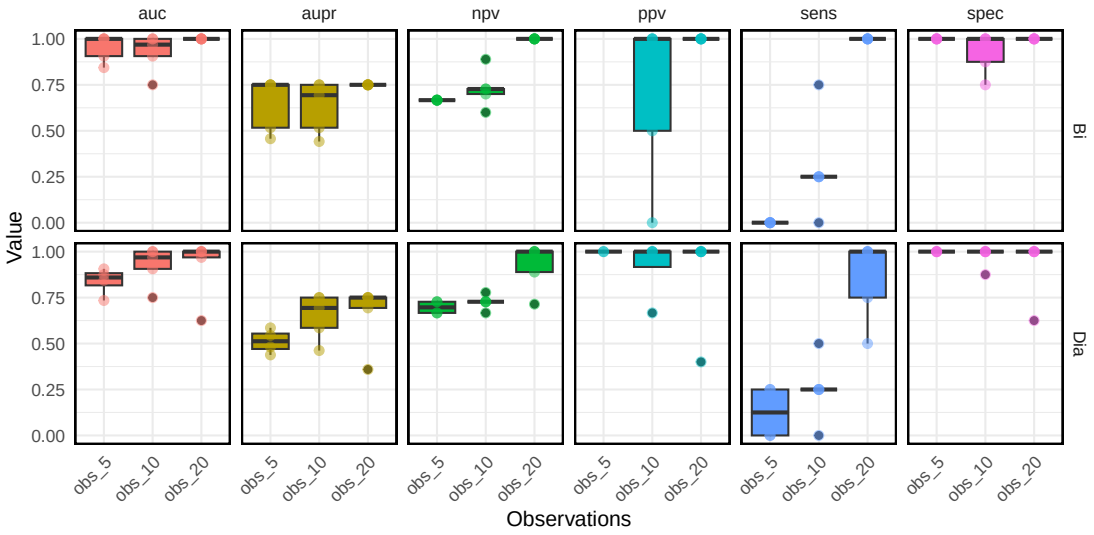

Figure 3.6: Motifs' Reconstruction Evaluations with Different Observation Numbers: Performance metrics (AUC, AUPR, NPV, PPV, sensitivity, and specificity) across two motifs ('Bifan,' 'Diamond) with varying observation numbers (E = 5, 10, 20).

As a next step, we show how number of $E-$genes can affect performance. We observed performance of the reconstruction of "Diamond" and "Bifan" motifs with a varying number of $E-$genes as $m = \{5, 10, 20\}$. Once more, we simulated data for constant number of timepoint with $T = 10$. In both cases, adding number of observations enhances all metrics.

We observed that as the number of $E-$genes increases, there is a consistent improvement in reconstruction performance across all metrics, including AUC, AUPR, NPV, PPV, sensitivity, and specificity. This pattern is apparent in both the "Diamond" and "Bifan" motifs, indicating that a higher density of observational data enhances motif recovery. The improvements are particularly pronounced in metrics such as AUPR and PPV, which are indicative of the model's ability to correctly identify true relationships and minimize false positives.

For both motifs, increasing the number of $E-$genes consistently improves reconstruction performance across all metrics. While the improvements are most notable when

transitioning from smaller to moderate values (e.g., from 5 to 10), the effect of additional observations remains substantial even at higher values (e.g., from 10 to 20). Notably, Sensitivity and NPV exhibit significant enhancements with higher observation counts, indicating that the model becomes more adept at identifying true connections and correctly ruling out false ones as the dataset becomes more comprehensive. This underscores the value of larger datasets for robust motif reconstruction, even as the rate of improvement for some metrics may taper off.

Additionally, the results indicate that "Diamond" motifs tend to achieve higher sensitivity and specificity compared to "Bifan" motifs under similar conditions, pointing to potential differences in the structural complexity of these motifs and their reconstruction challenges. This highlights the importance of motif type and network structure in determining the reconstruction success and provides a basis for tailoring observational strategies based on the underlying motif complexity.

Overall, these findings emphasize the crucial role of observational density in enhancing reconstruction accuracy, especially for more complex motifs. This has implications for experimental designs, where resource allocation to increase observational data can significantly benefit network reconstruction efforts.

Timepoints

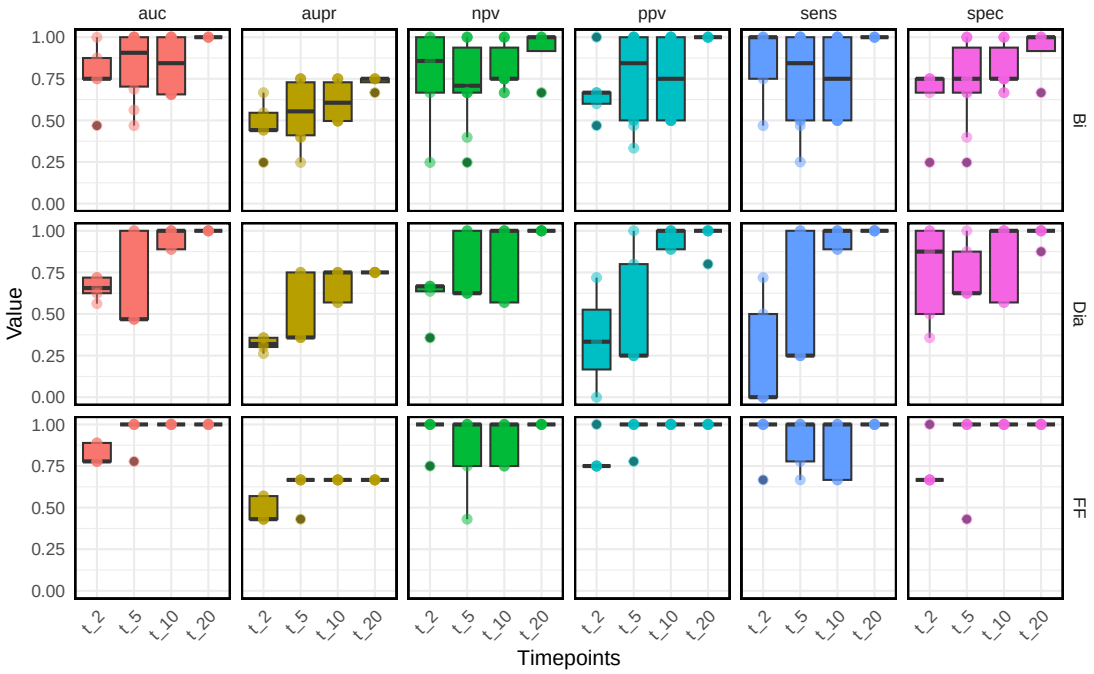

Figure 3.7: Motifs' Reconstruction Evaluations with Different Timepoints : Performance metrics (AUC, AUPR, NPV, PPV, sensitivity, and specificity) across three motifs ('Bifan,' 'Diamond' and 'Feed forward') with varying timepoints (t = 2, 5, 10, 20).

One crucial factor in dynamic network reconstruction is the length of the time series. Insufficient timepoints may fail to capture slower-developing effects, resulting in incomplete or inaccurate reconstruction. To address this, we investigated the impact of varying timepoints ($T = \{2, 5, 10, 20\}$) on the reconstruction performance of three motifs ("Feed forward," "Diamond," and "Bifan") and five randomly selected networks from KEGG pathways with five hidden nodes (from [47]). For this analysis, the number of observations was held constant at 20.

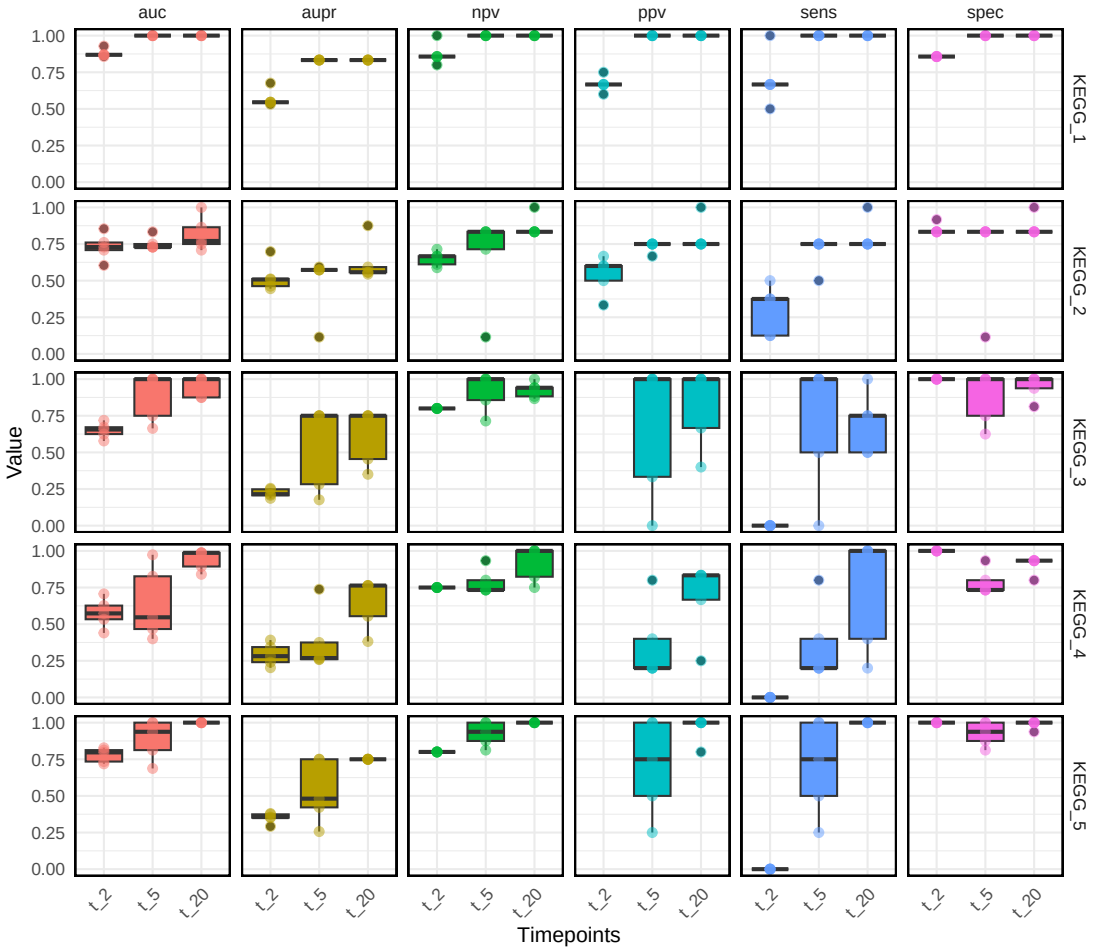

Figure 3.8: Reconstruction performance of examples of 5-Node KEGG networks across different timepoints (t = 2, 5, 10, 20) using six metrics: AUC, AUPR, NPV, PPV, sensitivity, and specificity.

Figure 3.7 and Figure 3.8 illustrate the impact of increasing timepoints on the reconstruction of motifs and KEGG pathway networks, respectively. Across all cases, extending the length of the time series improves performance metrics, including AUC, AUPR, NPV, PPV, sensitivity, and specificity.

For motifs (Figure 3.7), AUC consistently reaches perfection (1.0) at just 5 timepoints, underscoring the model's robustness in distinguishing positive and negative cases, even with relatively short time series. AUPR, however, shows greater variability, particularly for shorter timepoints ($T = 2, 5$), as the limited temporal resolution can miss key relationships within the network. Metrics such as PPV and sensitivity are initially variable

but show improvements and stabilization as the number of timepoints increases.
For KEGG pathway networks (Figure 3.8), the trends are similar but more gradual. AUC and NPV steadily improve, reaching near-perfect values at 20 timepoints. Sensitivity and PPV demonstrate higher variability for shorter timepoints, highlighting their sensitivity to limited data. Specificity remains consistent across all timepoints, reflecting the model's reliability in identifying negative cases.
These observations highlight the critical role of adequate timepoints in improving model performance. A shorter time series may capture less information, leading to variability in metrics such as AUPR and sensitivity. As the timepoints increase, the temporal resolution becomes sufficient to detect network relationships accurately, leading to stabilization in performance metrics. The consistency of AUC and NPV across both motifs and KEGG networks further underscores the model's robustness, while the variability in PPV and sensitivity emphasizes the need for sufficient temporal data to balance precision and recall.
This analysis helps to reveal how the model's predictive accuracy evolves, ensuring that it remains robust across different scenarios and stages.
The metrics, as described in Section 3.7.1 include the area under the receiver operating characteristic curve (AUC), precision-recall curve (AUPR), negative predictive value (NPV), positive predictive value (PPV), sensitivity (sens), specificity (spec) and high posterior density (HPD). These indicators give us a rounded understanding of the model's strengths and potential areas for improvement, encompassing both its ability to correctly identify true positives/negatives as well as its precision in avoiding false positives/negative.

3.8.1 Performance Comparison with dynoNEM

We compare the performance of odeNEM with dynoNEM for KEGG networks with 5 nodes to evaluate their respective strengths and limitations. Our method demonstrates strong specificity, maintaining high reliability in rejecting false positives, comparable to the dynoNEM results (Figure 3 and 4 from [47]). AUC achieves near-perfect levels for ($t > 5$), showcasing excellent classification performance. However, sensitivity falls behind dynoNEM, particularly at lower timepoints ($t = 2, 5$), indicating limitations in detecting true positives with sparse data. Trade-offs between sensitivity, AUPR, and PPV are apparent, in some examples of Figure 3.8, underscoring the need for improved balance between precision and recall.
For 10-node networks, the complexity of reconstruction increases significantly. In odeNEM,

Figure 3.5, sensitivity shows a pronounced decline, indicating that identifying true positive connections becomes more challenging as the network size grows. Specificity, while generally stable, also exhibits slight declines, suggesting that distinguishing true negatives may be affected by increased network complexity. In contrast, dynoNEM maintains robust sensitivity as the number of $E-$genes increases, demonstrating its capacity to identify true positives effectively in larger networks (Figure 3 supplementry of [47]).
We mention once more, while dynoNEM results were referenced from the original publication, the method was not rerunnable. This restricted direct re-validation of its performance against our approach.

3.9 Discussion

In this chapter, we developed a novel method for reconstructing hidden networks. This integrated approach increases the strength of NEMs in capturing hierarchical relationships while incorporating SSMs and SCMs to account for temporal dynamics and causal interactions within complex systems.
Unlike traditional static methods, our framework accounts for the nonlinear and dynamic nature of real-world systems, offering greater flexibility and accuracy in reconstructing complex networks. A major contribution of this work lies in the integration of ODE-based modeling with probabilistic tools. This combination allows us to quantify uncertainty in parameter estimates while ensuring biologically meaningful results.

Various validations were used to identify significant connections, distinguishing true signals from noise. This demonstrated the effectiveness of this approach across various network motifs and KEGG pathways. For simpler motifs such as feed forward, Bifan, and Diamond, the method achieved high accuracy even with a limited number of timepoints, as seen in the perfect AUC scores for as few as five timepoints. This efficiency highlights the method's potential for scenarios where experimental data is scarce. However, the performance for larger, more complex networks, such as those in KEGG pathways, revealed the need for sufficient observational data to maintain sensitivity and precision. These findings underline the trade-offs between network complexity and data availability, emphasizing the importance of tailoring the model to the data at hand.

Beyond reconstruction performance, the method also addressed key aspects of biological system. For example, modeling steady-states through gradient minimization and identifying attractors, captures essential biological features. This focus is particularly important for systems where variables, such as protein concentrations or gene expression

levels, must remain within physiological bounds. Additionally, the incorporation of a depletion factor allowed us to model the decay and turnover of biological entities over time, reflecting the natural dynamics of cellular processes. The method also accounted for noise in experimental data, a common challenge in biological systems. By integrating probabilistic tools with deterministic ODE models, we ensured robust performance even under uncertain or noisy conditions, enabling the identification of true network connections while minimizing false positives.

Despite these achievements, challenges remain. The decline in sensitivity and precision with increasing network complexity suggests that future work could explore improved regularization techniques or adaptive priors.

4 Inference of Protein-Protein Interaction Networks

In the last chapter, we presented the formal description of the odeNEM, a method designed to reconstruct the underlying hidden network of $S-$nodes that cause the observational effects in $E-$nodes. The next two chapters explain the applications of odeNEMs: the first focuses on protein phosphorylation measurements, and the second one focuses on RNA-seq data. However, before going through these applications, we introduce an abstract description of the main questions and further expand on the details of each dataset and corresponding experiments. Also, the semantics of $S-$nodes and $E-$nodes must be specified for each experiment, as the meaning of these nodes varies entirely across different experimental contexts. By applying odeNEM to protein phosphorylation time series data, we confirm known and reveal novel interactions between pathways that are active in cancer.

4.1 Breast Cancer

Breast cancer, the most prevalent cancer in women worldwide, is a complex and multifaceted cancer that requires a comprehensive approach to diagnosis and treatment. Global health reports show 11.7% of all yearly new cases are breast cancers [56], and global numbers show breast cancer accounts for cancer related death in high ranks [56]. Despite the fact that many early-stage cases of this cancer show no noticeable symptoms, common signs in more advanced stages can be a detectable change in the structure of breast(size, shape, thickening, lump or dimple) or nipple discharge of blood. Within the main two categories of breast cancer, ductal or lobular carcinoma, further subtypes are defined based on activity of hormone receptors or epidermial growth factor receptor [57]. Each breast cancer subtype can activate or suppress some pathways, depicted in Figure 4.1. That leads to variations in the tumor's susceptibility to different therapies. We introduce the main subtypes and receptor statuses, which are crucial for accurate modeling of complex dynamics of cancer.

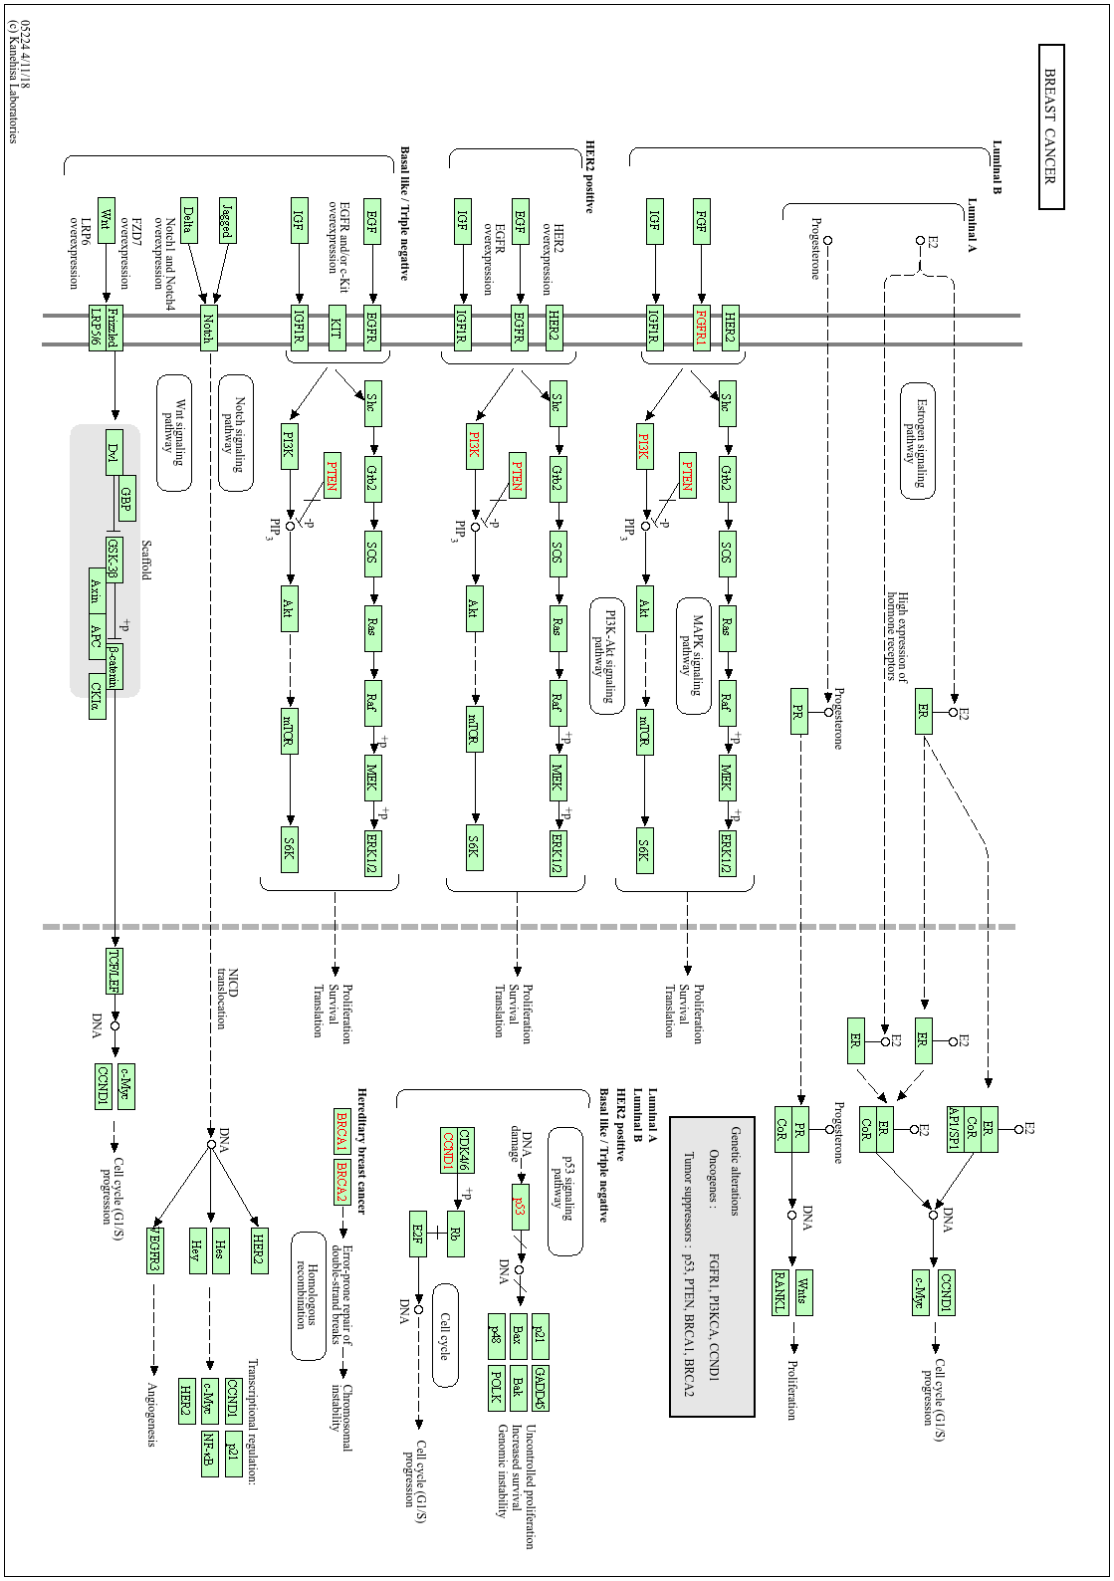

Figure 4.1: Breast Cancer Pathway (KEGG Pathway ID: hsa05224) [55]

HER2 Positive

HER2 (Human Epidermal Growth Factor Receptor 2) is a cell surface receptor that can be overexpressed in some subtypes of breast cancers. This overexpression increases cell proliferation and survival, leading this subtype to be more aggressive and historically more challenging to treat. However, the development of targeted therapies against HER2 expression has improved outcomes for patients with HER2-positive breast cancer [58], [59], [60].

Luminal

The luminal subtype is characterized by high levels of expression of the estrogen receptor (ER) and progesterone receptor (PR) [61].
The ER is a receptor that, upon binding to estrogen, regulates the transcription of genes involved in cell growth, survival, and differentiation. High levels of ER are commonly found in breast cancer and are often associated with a better response to hormonal therapy. However, resistance to hormonal therapy can develop, especially in advanced stages of the disease [62].
The PR is another nuclear receptor that interacts with progesterone. It plays a crucial role in the development and progression of certain types of breast cancer. Breast cancers that are PR positive are often sensitive to hormonal fluctuations and may respond to special therapies modulating progesterone[63].
There are two main luminal categories: luminal A and luminal B. Luminal A are typically ER and PR positive and HER2 negative with low proliferation rates. On the other hand, luminal B tumors are ER positive, PR can be positive or negative and HER2 can be positive. Luminal B often has higher proliferation rates.

Basal like/Triple Negative

The basal subtype has lower expressions of the three hormone receptors (ER, PR, and HER2). This pattern of receptor expression is referred to as triple-negative breast cancer (TNBC). Due to low expression of these three receptors, that are typically used to guide hormone therapy, TNBC is resistant to hormone therapy [64]. However, recent advances in cancer research have led to the discovery of other markers and biomarkers that can be used to guide treatment decisions in TNBC. For instance some types of cell lines in this subtype are sensitive to englerin A and digoxin, drugs that alter intracellular concentrations [65]. This suggests that these cell lines could be particularly responsive to

treatments that target these mechanisms.
The data that will be presented in this chapter, includes cell lines which are from all the mentioned subtypes above. We will discuss these cell lines and more features from the data in rest of this chapter. But before, we highlight a few pathways which are prominent in cell signal transduction of cancer. While there is abundant literature on cancer pathways, we focus on [66].

4.2 Pathway Interactions in Cancer

The heterogeneity of breast cancer is reflected in the variety of pathways that can be active in a tumor cell. To understand the systemic consequences of their activity, it is necessary to know their interactions. The KEGG pathway for breast cancer signaling (hsa05224), in Figure 4.1, provides an overview of the molecules involved in the most important pathways, and it shows molecules at which different pathways converge:

- PI3K/AKT/mTOR pathway: This pathway is crucial for keeping cells alive, helping them grow and managing their energy. Mutations and alterations in this pathway are common in breast cancer and can make the cancer grow faster and resist treatments [67]. Also it can interact with other pathways like MAPK/ERK [68].

- MAPK/ERK pathway: The MAPK pathway helps cells grow and become different types of cells. Often cross-talk between this pathway and other pathways like ER or HER2 can cause breast cancer cells to grow uncontrollably [69]. Mutation in this pathway leads to start of activation of many cancerous functions including uncontrolled cell proliferation, resistance to apoptosis, enhanced angiogenesis, and increased metastatic potential [70].

- WNT pathway: Abnormal activation of this pathway contributes to cancer initiation and progression, especially because the pathway is responsible for stem cell maintenance and tissue homeostasis. Its interaction with Notch and Hedgehog causes a transition of epithelial to mesenchymal (EMT) [71].

- NF-kβ pathway: It plays a central role in inflammation and immune responses. Interaction of this pathway with PI3K/AKT enhances oncogenic signals [72].

- p53 pathway: Normally it is crucial in DNA damage repair and cell cycle regulation. In various cancers, this pathway is frequently mutated which leads to

genomic instability and proliferation [73].

4.2.1 Importance of Studying Signaling Pathways in Specific Context

The behavior of signaling pathways can vary depending on genetic and epigenetic factors. This means that the same external intervention (stimulus or inhibitor) can lead to different cellular responses under distinct conditions. Due to the unique genetic or epigenetic composition of each lineage, these differences become more clear when responses to identical perturbations have been compared across different lineages [74].
Such variability underscores the need for computational approaches that can adapt to the unique characteristics of diverse biological contexts. Additionally, it is crucial to predict how structure of such networks will change in response to new stimuli or treatments. Therefore, the proposed model should simulate the dynamic behavior of the network under various conditions.
The importance of understanding signaling pathways in context is the topic of HPN-DREAM (Heritage Provider Network) challenge [75], [76]. This challenge emphasizes the need for models that can adapt to various biological conditions and predict how signaling networks behave in different contexts. The set of experiments in the large-scale data of [75] can be effectively described by our model, Section 3.4.
Also our approach, outlined in Chapter 3, aims to address these challenges by considering the unique characteristics of different cell types and their specific biological contexts as feeded observations under variouse experiments. In the end, by integrating all inffered context-specific networks, we can better model a holistic picture of the signaling pathways which are central in cells responses to various perturbations.

4.3 HPN-DREAM Challenge

The HPN-DREAM holds a challenge in the field of computational biology that focuses on learning causal influences in signaling networks using both experimental and in silico data. The 8th version, [76], consists of three sub-challenges:
causal network inference : It involved inferring causal signaling networks using protein time course data from various biological contexts, each defined by a combination of cell line and stimulus. Using inferred networks, the causal validity was evaluated using unseen interventional data.
time-course prediction : The inferred network from previous sub-challenge had been

applied to predict trajectory of effects under new perturbation (not included in data). It basically evaluates suggested model in generalization aspect.
Among these three main concepts, we will mainly deal with first one, applying our model to infer causal relations that may vary across different contexts.

4.3.1 DREAM Data

The challenge involved 32 different biological contexts, each defined by a unique combination of a cell line and a stimulus to capture the variability in signaling responses due to differences in genetic and epigenetic backgrounds. For each context, they provided the time course data of 45 phosphoproteins using kinase inhibitors as interventions, Figure 4.2a.
Each cell line was firstly deprived of Serum to minimize effects from nonplanned factors in the medium. This is normally a standard practice in cell culture to ensure that the experimental conditions are controlled and that the results are just obtained by the experiments and not a random effect. Then the pre-treated cells were subjected to kinase inhibitors or DMSO as a control. As we will describe in Section 4.3.1, kinase inhibitors are commonly used to block the activity of specific enzymes while DMSO serves as a control. After stimulation with each stimulus, Section 4.3.1, Reverse-Phase Protein Array (RPPA) was conducted at 10 time points targeting 40 phosphoproteins. 150 high quality antibodies, aimed at detecting total and phosphorylated proteins. The RPPA methodology in [76] involves serial dilutions of cell lysates, printed on slides, and probed with validated antibodies. After capturing signal intensity, the data is processed using a B-spline model to fit a "supercurve" for each slide, relating signal intensity to protein concentration. Quality control checks ensure slide consistency, excluding those below a certain threshold. For normalization, values are adjusted by median centering across antibodies and correcting for each sample to account for protein loading differences. To isolate the immediate and direct effects of the inhibitors on the phosphoproteins, the paper in which the experiment was carried out [76] focuses on short changes up to 4 hours post-stimulation.
So, this data can be understood through these variables: *cell Lines*, *kinase Inhibitions*, *stimuli*, *timepoints*, *proteins*. We provide below an overview of the values these variables can assume.

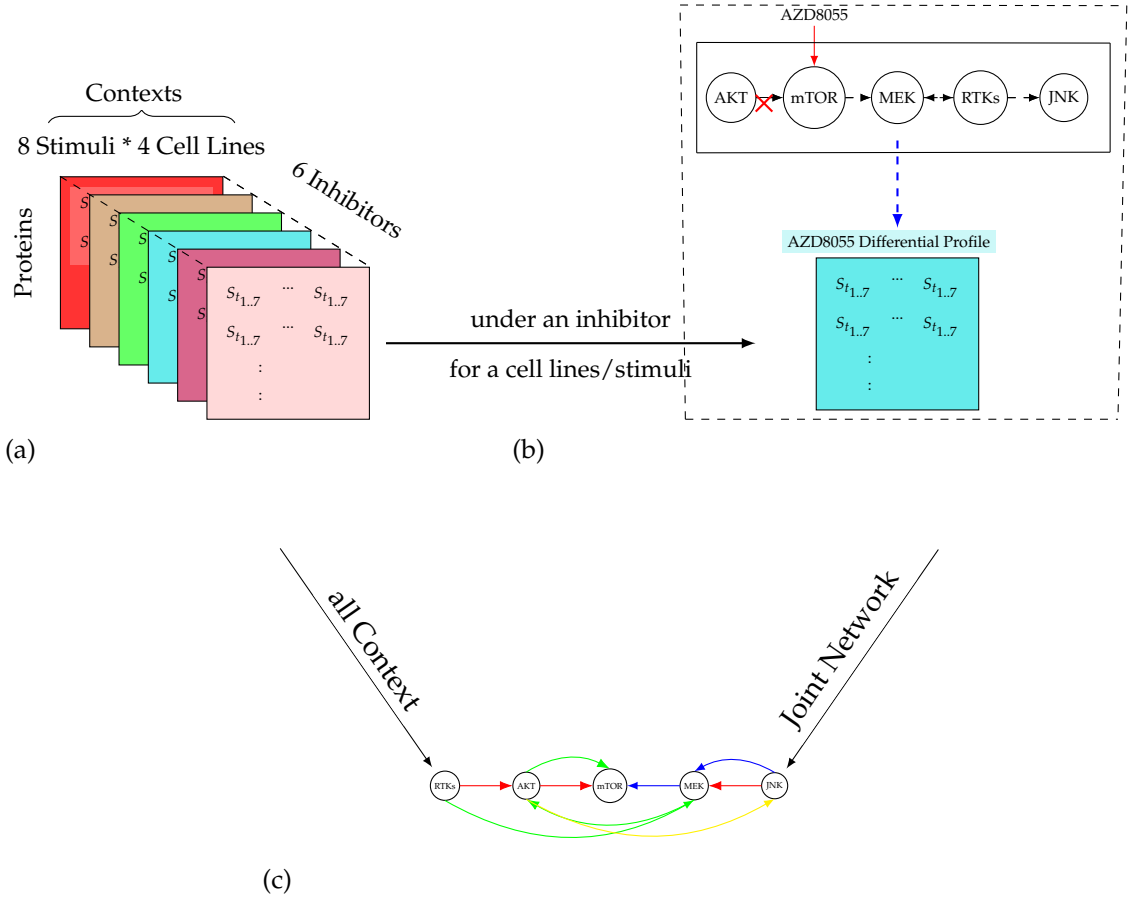

Figure 4.2: **(a) Overview of Biological Contexts and Data Dimensions:** Each context is defined by the combination of 8 stimuli, 4 cell lines, and 6 kinase inhibitors. The time course data captures phosphoprotein responses across these contexts, reflecting variability in signaling due to genetic and epigenetic differences. **(b) Differential Profiles Under Kinase Inhibition:** Using kinase inhibitors (e.g., AZD8055), specific signaling pathways are perturbed. Reverse-Phase Protein Array (RPPA) data measures changes in 45 phosphoproteins over time to characterize the inhibitor's effects. **(c) Joint Network Inference Across All Contexts:** Integrating data from all biological contexts to infer a joint signaling network. Nodes represent proteins, and edges represent interactions, color-coded by the pathway or context specificity.

Cell Lines

Various breast cancer cell lines can contain different genomic/epigenomic variants. In addition to that, the subtype which a cell line belongs to can highlight some signaling pathways while supressing the rests, Section 4.1. [76] included four cell lines: MCF7, UACC812, BT20 and BT549. Each cell line represents a different subtype of breast cancer, based on the expression of certain receptors and the presence of specific mutations.
For instance, the MCF7 is a luminal subtype known to express high levels of the estrogen receptor (ER) and progesterone receptor (PR). UACC812 is another luminal subtype whith overexpression of human epidermal growth factor receptor-2 (HER2) [61].
On the other hand, BT20 and BT549 are believed to belong to the basal subtype, given their lower expression of these three receptors. This pattern of receptor expression is referred to as "triple-negative".

While these cell lines serve as widely used models for studying breast cancer subtypes, it remains debated how accurately they reflect the complexity of tumor heterogeneity in real patients. Differences in the tumor microenvironment, genetic diversity, and epigenetic modifications may limit their direct translational relevance [77].

Stimuli

DREAM 8 challenge includes eight stimuli: Epidermal Growth Factor (EGF), Fibroblast Growth Factor (FGF), Insulin-like Growth Factor 1 (IGF-1), Neuregulin 1 (NRG1), Hepatocyte Growth Factor (HGF), Insulin, PBS, Serum. As we see previously in Figure 4.1, each of the mentioned stimuli can interact with specific receptors on cell surface (under specific conditions). This triggers different signaling pathways and leads to an alternation in regulation of cell functions such as proliferation, survival, and many other cellular processes. For example, EGF as a potent mitogen binds to receptors like EGFR (ErbB1, HER1) or HER2 (ErbB2) to trigger intracellular signal transduction pathways in cell proliferation, differentiation, and survival [78]. So stimulating each cell line with these stimuli mimics both physiological and pathological conditions that breast cancer cells might encounter in the body.

Kinase Inhibitions

Before stimulating a cell line with a stimuli, it has been treated with one of the five inhibitors: PD173074, BEZ235, AZD8055, GSK690693, GSK690693-GSK1120212. Also the

control case (DMSO) has been considered. Each of the mentioned inhibitors restrains activity of one or more direct targets. Such experiments aimed to uncover specific causal relationships between targets (directs or undirects) that might not otherwise be evident. This has been done by isolating effect of each experiment (in our case inhibition) to see how dynamics change between nodes [79], [80] and [81].
Inhibitors in the DREAM 8 challenge target various key roled proteins, each playing a crucial role in cellular signaling pathways: PD173074 $\rightarrow (FGFR1, FGFR2)$, BEZ235 $\rightarrow (p110, mTOR1, mTOR2)$, AZD8055 $\rightarrow (mTOR1, mTOR2)$, GSK690693 $\rightarrow (panAKT)$, GSK690693-GSK1120212 $\rightarrow (panAKT, MEK1, MEK2)$

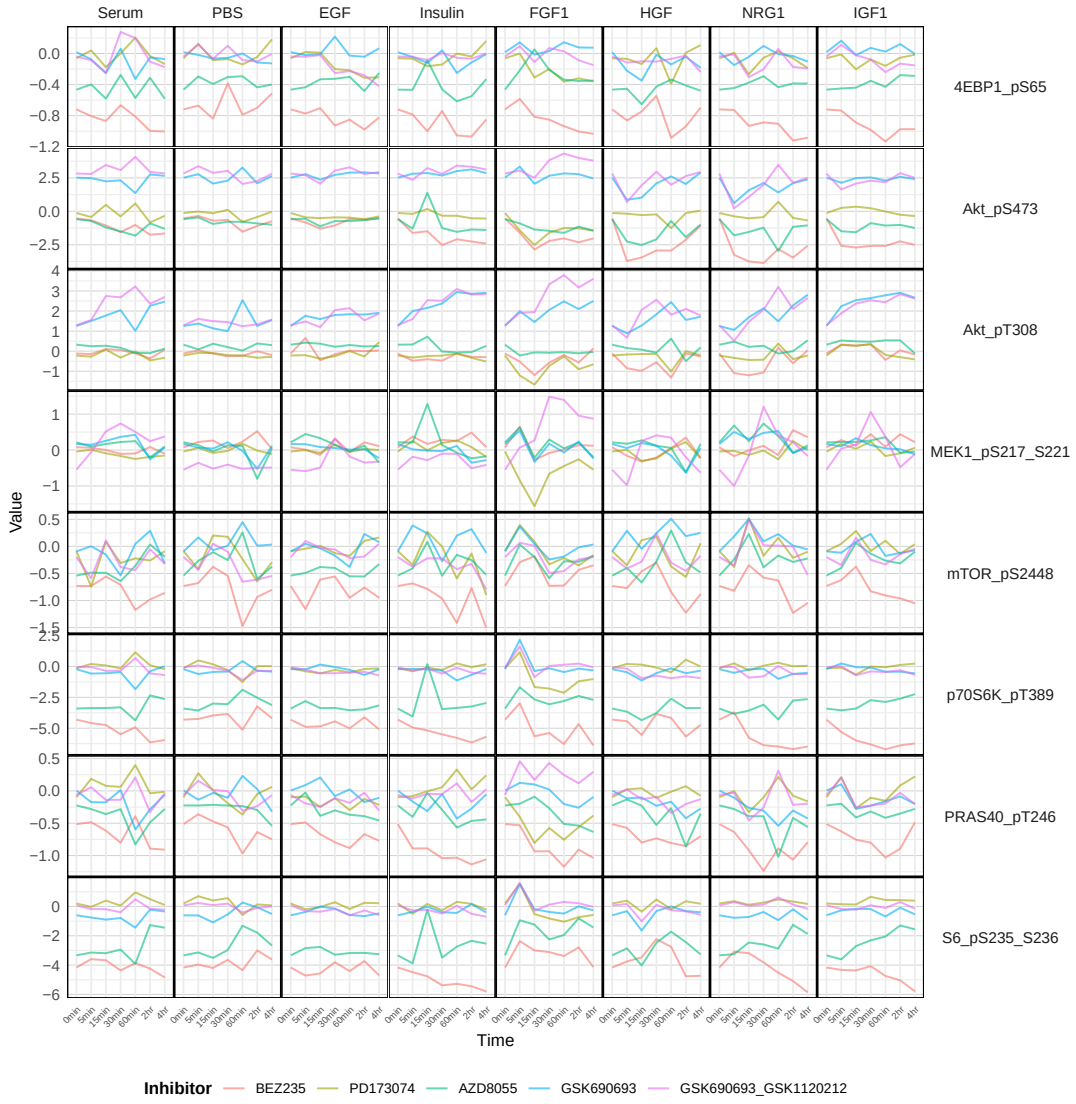

Figure 4.3: Stimulus Dependent Changes in Phosphoprotein Abundance Under Kinase Inhibition

Timepoints

To capture the temporal dynamics of signaling pathways, the experiments include measurements at multiple timepoints following stimulation. These timepoints allow for observing both the immediate and delayed effects of stimuli and inhibitors on protein activity. Such temporal data are critical for understanding transient (short response) versus sustained (remain active longer) signaling events and for inferring causal interactions

among network components.

Proteins

The dataset encompasses measurements of key signaling proteins across various conditions, including total and phosphorylated forms. These proteins serve as nodes in the signaling network and include those involved in critical pathways such as the MAPK, PI3K-AKT, and mTOR pathways. Like the main paper of DREAM, we have used phosphoproteins because they directly represent the active state of signaling pathways, providing functional and dynamic insights into cellular processes, unlike whole protein levels.

4.3.2 Data Analysis

The time series data is used to determine which phosphoproteins show significant changes in abundance under an experiment (kinase inhibition effect) versus control state (DMSO). This set of proteins will be regarded as our $E-$nodes later.

Preprocessing

In [76], the paired t-tests are used to compare the mean phosphoprotein abundance under DMSO control and the inhibitor regime. The results are then corrected for multiple testing using the adaptive linear step up procedure for controlling the False Discovery Rate (FDR) and evaluated against specific criteria to identify significant changes. This is necessary in multiple hypotheses testing, to keep the probability of a false positive result (Type I error) below a predetermined threshold. 35 out of 48 phosphoproteins significantly changed under a given inhibitor regarding these two conditions:

- FDR value less than 5 percent (from paired t-tests comparing DMSO and inhibitor time courses)
- The inhibition effect size (measured as the log2 ratio of abundance under inhibition versus DMSO control) exceeds the replicate standard deviation:

$$A(t) = \log_2\left(\frac{\text{abundance under a condition}}{\text{abundance under DMSO}}\right) \tag{4.1}$$

Detailed results are accessible in [76]. These 35 proteins will be our set of $E-$nodes in setting up of our model. For the selected proteins, we also calculated the log ratio of changes $A(t)$, which will be the main profile of observed measurements.

Context Dependent Effects

As it has been shown in Figure 4.3, there are few specific changes that can be described exclusively due to stimulus addition. For example, relative changes of some phosphoproteins, like Akt or MEK1 phosphoproteins, under FGFR inhibition (PD173074) are more prominent when cell line is stimulated with FGF1. Also, large variations in relative abundance of AKT phosphorylations seems to be more stimulus dependent under inhibition with BEZ235.

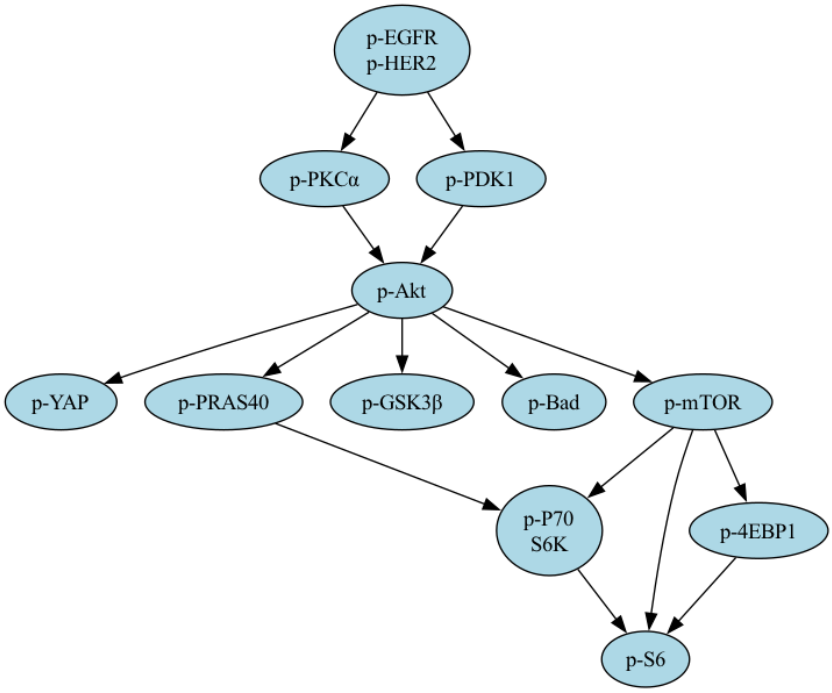

Figure 4.4: Canonical Signaling Pathway Highlighting Context Independent Interactions from [76]

Before we expand on this topic, let's have a look to canonical signaling pathway which [76] introduces. Such networks are applied to make a more general overview about pathways regardless of the conditions they have been concluded in. Most of the prominent known connections in such networks are also valid under different circumstances.
While such canonical signaling pathways provide a general framework for regulatory interactions, the actual edges (connections) between proteins can change depending on the specific biological context. This means that under different conditions (e.g., different cell lines, stimuli, or inhibitors), some interactions may become more or less prominent, reflecting context-specific regulation rather than universally fixed connections. For exam-

ple, mTOR and its more direct descendants in canonical networks, in Figure 4.4, are being affected mainly under inhibition of mTOR (via AZD8055 or BEZ235). This is regardless of cell lines or stimuli. For example, Figure 4.3 shows a noticeable phosphorylation level reduction for p70S6K just under inhibition of mTOR (AZD8055 and BEZ235) regardless of presence of any stimulus.
So, screening phosphoprotein changes in specific context c (combination of cell lines, Stimuli) shows notable changes mainly when the inhibitor targets the parent directly. However there are yet influences that suggest effects are not just due to known connections in canonical pathways.
Complete plots for all phosphoproteins in all four cell lines are provided at the end of this thesis for comparison.

4.3.3 Inferred Network by Challenge Contest

Stacking all data together and utilizing techniques which infer causalities under different conditions, [76] reconstructed networks for each context. The STAR (STructure learning with Adaptive Regulators) method is a Bayesian network inference approach designed to reconstruct causal signaling networks from perturbation data. It integrates prior knowledge with experimental data to identify significant regulatory interactions. The method employs an adaptive strategy to update the network structure iteratively, refining the causal relationships as more data becomes available. This approach is particularly useful in capturing context-specific interactions that may not be evident from prior knowledge alone. For a detailed description of the STAR method, including its algorithmic implementation, please refer to the supplementary materials of [76].
The STAR method reconstructs causal signaling networks where each variable (node) is directly observable. In contrast, our approach (odeNEM) is designed for networks with hidden variables (directly observable). Due to this fundamental difference in network structure and inference methodology, we do not perform a head to head comparison between the two methods. Instead, we leverage the STAR inferred network as a reference to contextualize our findings.

The comprehensive network in Figure 4.5 is a summarization of causal networks of all contexts in [76]. By this unification, we can see many edges appeared as significant cell specific relations that were not present in the applied prior networks of the corresponding paper. This can strongly suggest the necessity of studying causal networks regarding specific setups and conditions.

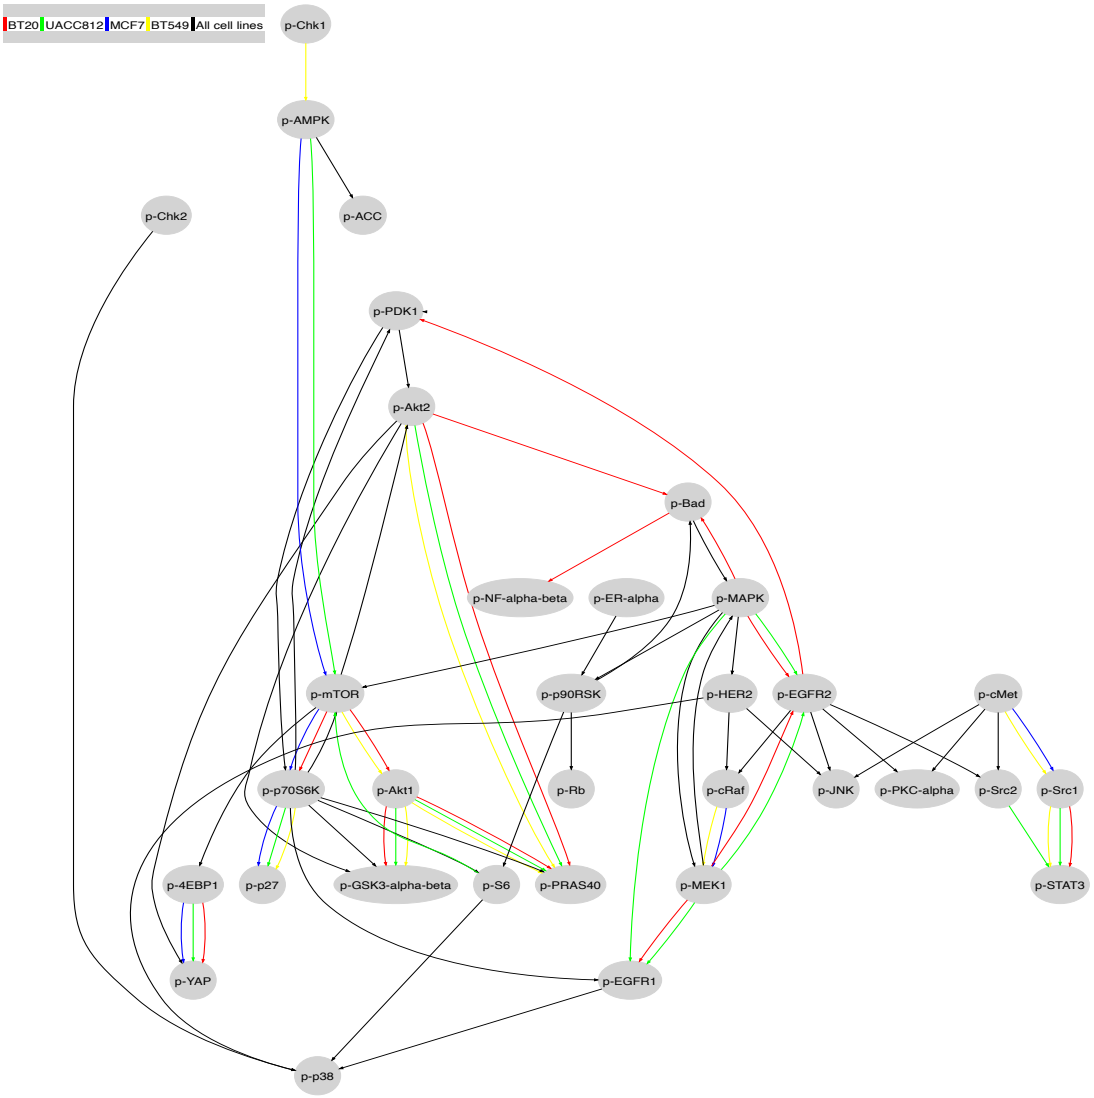

Figure 4.5: Comprehensive Causal Network Summarizing All Contexts using the STAR Method.
This network, recreated from [76], represents inferred causal relationships across multiple cell lines. Nodes correspond to phosphoproteins, and edges indicate significant interactions inferred across experimental conditions. The network includes context dependent interactions, meaning relationships may vary across cell lines and experimental conditions.

4.4 Applying odeNEM

A computational model, in particular Nested Effect Model (NEM), can be used to study dynamics of interactions between pathways under perturbation effects. Especially the complete inhibition experiments, like how DREAM 8 data is produced based on, can be modeled easily by our odeNEM model (Section 3.4).
The primary motivation for applying odeNEM to these data is benchmarking rather than discovering new interactions. By evaluating its performance on this well characterized dataset, we assess the model's ability to recover interactions between pathways which are known.

4.4.1 Selecting hidden nodes and effect nodes

The $E-$nodes in the odeNEM model represent the observables. In our data, as it was described in Section 4.3.2, phosphorylation levels of 35 proteins had been considered as the observed effect nodes E.
The $S-$nodes are variables that transmit signals to the $E-$nodes. Ideally, one would like to include all molecules involved in the signaling pathways (Figure 4.5), but the network of such a large number of $S-$nodes can not be learned reliably from our noisy observations. We therefore decided to reduce the complexity by merely modeling 3 main signaling pathways (mTOR, AKT, MEK) as $S-$nodes. In addition to that, there are quite a few phosphoproteins in the data which play role as receptors of external stimuli in the big picture of breast cancer pathways, Figure 4.1. So, we have added a "Receptor" node to the set of $S-$nodes to indicate this class of proteins.
It is worth mentioning back to the definition of interventions in Section 3.4. Under each experiment, one (or several) $S-$node becomes the target of a complete inhibition and passing information through this node becomes fully blocked. So, the set of direct targets of an inhibitor can be potentially regarded as hidden nodes. Also, it is necessary that for each of those hidden nodes there should be relevant observations in set of effects, $E-$nodes, to support what we observe from experiments. As in previous parts have been described, inhobitors used in DREAM challenge inhibit these targets: FGFR1, FGFR2, p110, mTOR1, mTOR2, panAKT, MEK1 and MEK2. From these direct targets, we do not have measured phosphoproteins of $FGFR_i$ and p110. So, we regard mTOR, AKT, MEK as the hidden nodes.

4.4.2 Prior on S-E Connections

We considered three different priors for the connections between hidden nodes S and observations E: uniform, strict, and DREAM-based priors. In all three cases, few phosphoproteins ($E-$genes) were directly and exclusively connected to their corresponding parent S-nodes :
(Akt-pS473, Akt-pT308) $->$ AKT, mTOR-pS2448 $->$ mTOR, MEK1-pS217-S221 $->$ MEK
For certain other proteins which are well documented in the literature as downstream markers of specific pathways, our strict prior considered these proteins exclusively as downstream effects of their relevant parent. For all other proteins, prior distribute the connections evenly across all possible hidden nodes. For example, PRAS40, which is known to be regulated by the PI3K-AKT pathway [82], was connected to the parent node with the highest prior probability.
In addition to using biological knowledge, we incorporated canonical pathway information and the prior network described in [76] to refine the S-E connections. This approach formed the DREAM-based prior, as illustrated in Figure 4.6.
The uniform prior distributed observations evenly among all hidden nodes, except for the specifically assigned phosphoproteins (Akt-pS473, Akt-pT308, mTOR-pS2448, MEK1-pS217-S221), which were directly linked to their corresponding $S-$nodes.
It is important to note that while our priors specifically influence S-E connections, the S-S connections themselves are not determined by these priors but are inferred through the model.

4.4.3 Emission Probability Modeling

After calculating log ratio of changes for each condition versus the wild type by Equation (4.1), a gamma mixture model has been fitted to the observation profile for each phosphoprotein under a perturbation. To determine the most appropriate model, we used both Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). AIC assesses model fit while balancing the number of parameters, whereas BIC provides a stricter penalty for model complexity. Using these two criteria, mixture of two gamma models showed best trade off between goodness of fit and model simplicity (lowest AIC and BIC) for most of the proteins and conditions.
An example of mixture model fit for protein p70S6K under two expreiments has been shown in Figure 4.7. More details on the EM algorithm and parameters fitting can be found in Section 3.6.

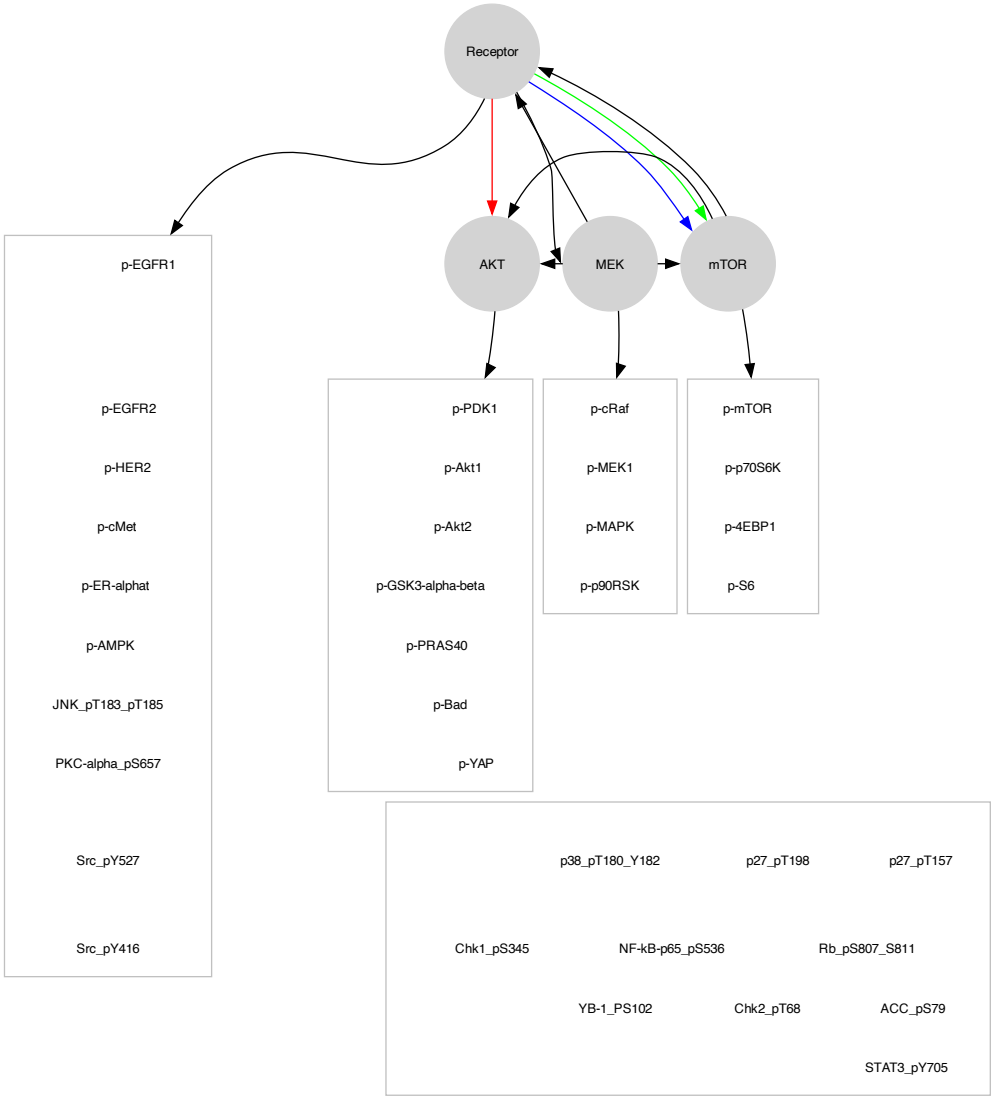

Figure 4.6: Pathway Level Network Derived from the STAR Inference.
This network restructures Figure 4.5 by grouping phosphoproteins into hidden nodes, each representing a signaling pathway. The grouping shown here follows the DREAM-based prior, described Section 4.4.2. This structure serves as a ground truth for evaluating odeNEM's performance in using DREAM-based prior in Figure 4.8.

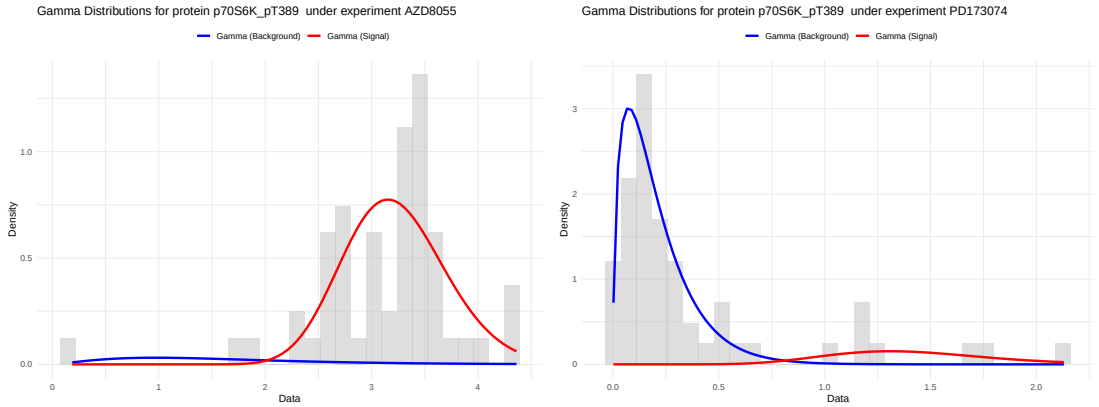

Figure 4.7: Gamma Mixture Modeling of Phosphoprotein Profiles Across Experimental Conditions

4.5 Results

Our model focuses on the interaction between various signaling pathways that play key roles in breast cancer. To highlight pathway level interactions, we derived Figure 4.6 by restructuring the network inferred by the STAR method (Figure 4.5). Specifically, we grouped phosphoproteins into hidden nodes representing their associated signaling pathways, reducing node complexity. Unlike Figure 4.5, which includes protein level interactions across different contexts, Figure 4.6 retains only edges that reflect regulatory relationships between pathways. By this transformation, we ensure structural comparability with the pathway level network inferred by odeNEM, allowing a direct evaluation of its predictive performance.

The grouping shown in Figure 4.6 was determined by the DREAM-based prior, which integrates biological knowledge and canonical pathway structures. Since each prior (uniform, strict, and DREAM-based described in Section 4.4.2) assigns phosphoproteins to hidden nodes differently, each prior defines a different possible ground truth. For clarity, Figure 4.6 presents only the network derived from the DREAM-based prior, while Figure 4.8 evaluates odeNEM's performance across all priors by comparing inferred networks to their respective ground truths.
From the odeNEM model, we obtain a posterior probability for each potential edge $S_i \rightarrow S_j$. An edge is considered as predicted (positive) if its posterior probability exceeds a chosen threshold, and not predicted (negative) otherwise. In these experiments, the typical threshold was set to 0.7 on the posterior probability. This keeps our results far

beyond random inferred effects. Once both the inferred adjacency matrix (based on the thresholded posterior probabilities) and the ground truth adjacency matrix (as per the uniform, strict, or DREAM-based prior) are in binary form, each potential edge falls into one of four categories: TP, FP, TN and FN. Using these counts, the standard metrics, described Section 3.7.1, were computed.

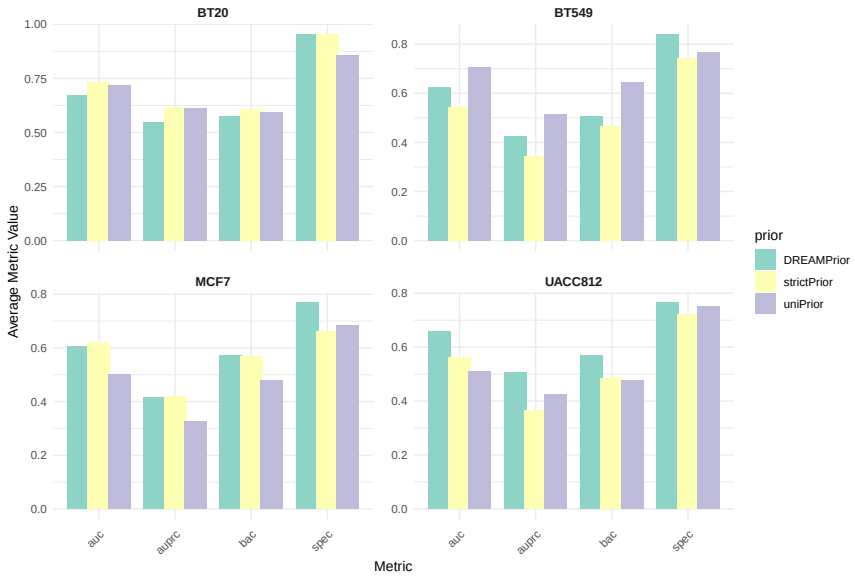

Figure 4.8: Inference Performance of odeNEM Under Different Priors (uniform, strict, and DREAM-based)
The inferred network for each prior were compared to their respective ground truths

To ensure that our MCMC sampling adequately explored the posterior distribution and reached a stationary regime, we conducted extensive diagnostics on the convergence behavior. We ran multiple independent MCMC chains, monitoring their evolution over 200,000 iterations. Convergence was assessed through visual inspection of log-likelihood trace plots and the stabilization of posterior samples.

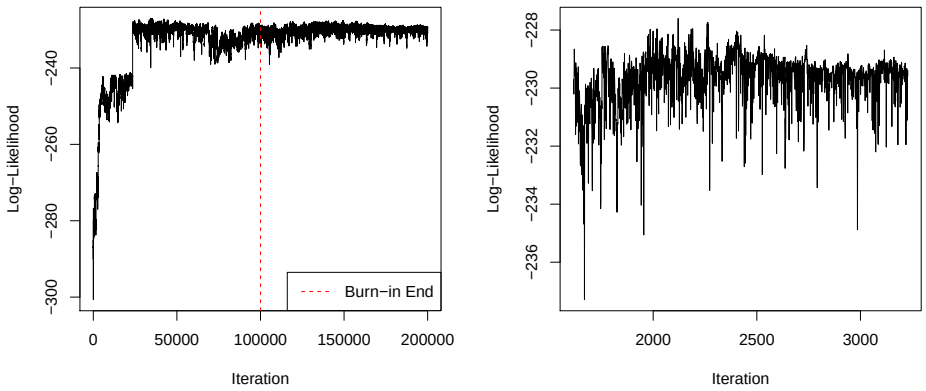

(a) Chain 1

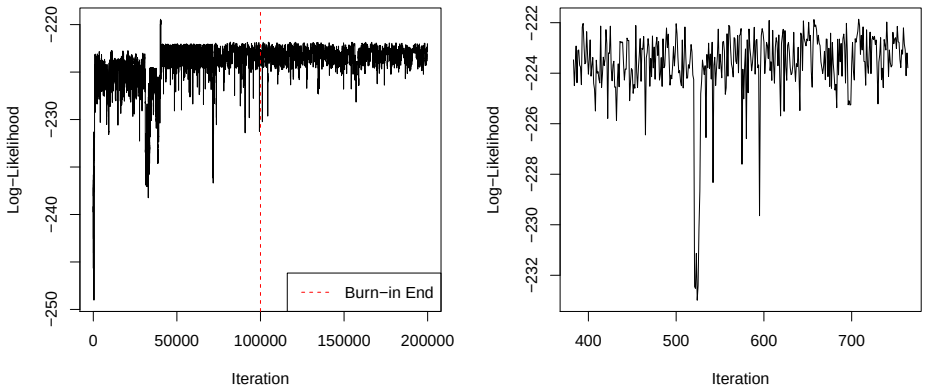

(b) Chain 2

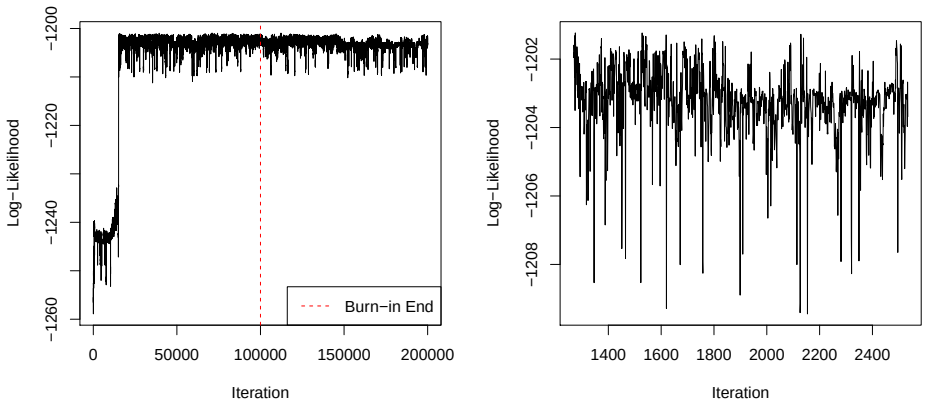

(c) Chain 3

Figure 4.9: Convergence diagnostics of the MCMC sampler based on log-likelihood trace plots. Each panel shows the log-likelihood evolution over iterations for different chains on left side. On right side, the accepted log-likelihoods after burn-in and thining are depicted.
Chains are from different contexts (Cell line, Stimuli). So, the likelihood values are not comparable. The burn-in phase is visually evident as an initial drift, after which the likelihood stabilizes, indicating proper convergence.

Figure 4.9 presents representative log-likelihood trace plots for different chains. The left panel of each plot illustrates the complete trajectory of log-likelihood values over the full set of iterations, highlighting the initial burn-in period where the likelihood exhibits a transient phase before stabilizing. A vertical red dashed line at 100,000 iterations marks the empirically determined burn-in phase, beyond which the samples appear to have reached stationarity. The right panel zooms in on the post burn-in region, showing the fluctuations of log-likelihood values in a stable regime.
To further confirm the reliability of our sampling strategy, we applied a thinning interval of 50, reducing autocorrelation in the posterior samples while ensuring sufficient effective sample size (ESS). The trace plots confirm that, after the burn-in phase, the likelihood stabilizes, and the sampling process maintains sufficient variance to capture the posterior distribution adequately.

4.5.1 Evaluation of Priors

We evaluated our model using the three described priors in Section 4.4.2: uniform, strict, and DREAM-based priors. Figure 4.8 illustrates the reconstruction performance across these priors, measured using metrics such as AUC, AUPRC, BAC, and specificity. The DREAM-based prior consistently outperformed the other priors across most metrics and cell lines, demonstrating its robustness. It likely performed better because it incorporates biological knowledge and canonical pathway information, making it more specific while still allowing flexibility in modeling noisy data.
AUC is generally high for all priors and cell lines, particularly for the DREAMPrior, showing its robustness in identifying true signal from noise. Also it is observable that DREAMPrior slightly outperforms other priors in BT20 and BT549, which indicates it captures the signal to noise separation better for these cell lines. Specificity is high for all priors, particularly for BT20 and UACC812, where all priors perform almost equally. This shows the model has a low false positive rate and performs well in identifying true negatives. While AUC and specificity are generally high, AUPRC (and to some extent BAC) highlight the challenges of balancing true positive identification with false positive minimization. AUPRC is consistently lower than AUC, which is expected for imbalanced datasets (sparse networks). BAC (Balanced Accuracy) is moderate to high for all priors. DREAMPrior performs slightly better, again showing balanced performance in both positives and negatives (high BAC shows balanced sensitivity (true positive rate) and specificity (true negative rate)).
In conclusion, the DREAM-based prior demonstrates consistent superiority across most

metrics, particularly in AUC, AUPRC, and BAC, making it a reliable framework for capturing pathway dynamics and modeling sparse, noisy data. Based on these results, we selected the DREAM-based prior as the foundation for further analyses and assessments. The differences observed across cell lines (e.g., BT20 and BT549 are performing better) may reflect underlying biological variability or differences in pathway activation patterns as well as the level of noise in the data.

4.5.2 Heatmap Analysis

To assess how well the model predictions align with the observed data across various conditions, we divided the data into two parts: a training set and a test set. We split the data by inhibitor: for each chosen test inhibitor, we excluded all time course measurements under that inhibitor from the training set and reserved them as the test set. Concretely, if AZD8055 was designated as the test inhibitor, then all AZD8055 data were excluded from training, and the remaining inhibitors' data (which do not target same node as AZD8055 does) formed the training set.
Our model was first trained on the training data yielding a collection of candidate networks ranked by their posterior likelihood. From these, we selected the first few hundreds best inferred networks. In practice, this often means taking a few hundred top scoring networks from the MCMC sampling. Using these top networks, we predicted the activity trajectory of the hidden node that the test inhibitor targets (e.g., mTOR for AZD8055). The activity level was computed directly from Equation (3.4). Since the ODE output (predicted activity) and the experimental data (log fold changes from Equation (4.1)) are on different scales, we could not compare them numerically. Instead, we compared their direction (increase or decrease) over time. Figure 4.10 summarizes these comparisons by displaying the match proportion for each phosphoprotein, stimulus, and cell line under the held-out inhibitor regime. Concretely, the match proportion is the fraction of time points at which the sign (positive or negative) of the predicted activity agrees with the sign of the observed log fold change. A value of 1.0 (dark red) indicates perfect sign alignment over all measured time points, whereas 0.0 (white) indicates no agreement in direction. Our results show that while the model generally captures pathway dynamics for many proteins and stimuli, there is variability in alignment, suggesting potential areas for refinement.

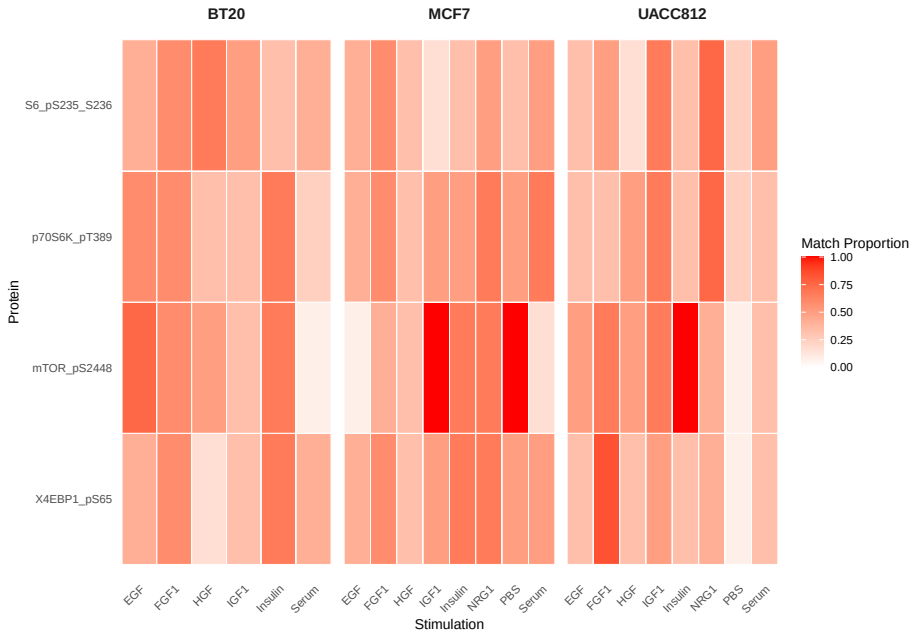

Figure 4.10: Heatmap of match proportion between observed and predicted trajectories across proteins, stimuli, and cell lines.
The match proportion represents the fraction of time points where predicted changes in pathway activity match the observed phosphoprotein response. Each cell shows the fraction of time points (0 to 1) at which the direction of the predicted activity matches the sign of the observed log fold change for that phosphoprotein, stimulus, and cell line. Higher match proportions (dark red) indicate stronger agreement between model predictions and experimental data.

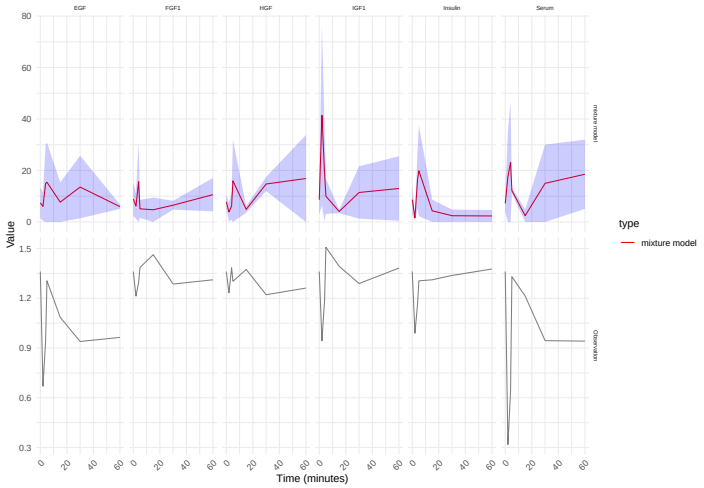

(a) BT20

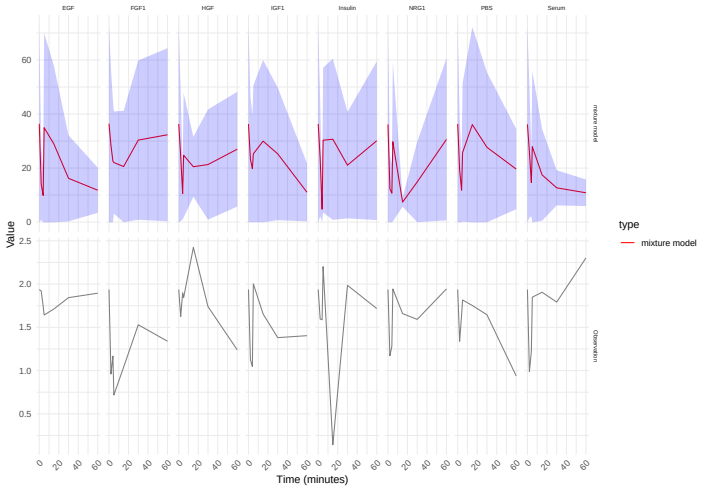

(b) MCF7

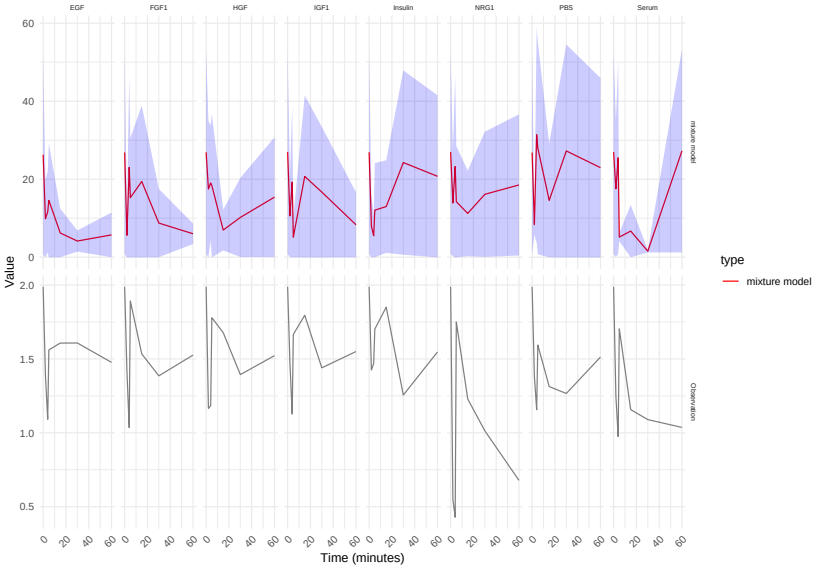

(c) UACC812

Figure 4.11: Assessment of causal learning by comparing true observation and our mixture model. Each subfigure illustrates the comparison between observed (solid black) and mixture model (red) for the specified cell line.

4.5.3 Trajectory Comparisons

To visually assess the model's performance, we generated plots in Figure 4.11 that overlay the observed trajectories (black) with the predicted trajectories (red). The observed trajectories are average log fold changes (Equation (4.1)) for the proteins specifically connected to the hidden node inhibited in test experiment, while the predicted trajectories come from the a mixture distribution (Equation (3.13)). In other words, we want to see how our mixture model can mimic the true observation for the $E-$nodes which the parent is the target in test set. Although these two curves appear on different numerical scales, their shapes can be compared to judge whether the model captures the qualitative trends of activation or deactivation under each stimulus.
Because of this mismatch, we do not directly compare the black and red lines' absolute values. Instead, we look at whether increases or decreases occur at similar time intervals and whether the predicted curve stabilizes or fluctuates in a manner consistent with the observed protein-level data.

For instance, in BT20 with Insulin, EGF or FGF1, the observed log fold changes and the mixture model increases early on and then stabilizes, indicating the model is capturing the timing of the major transitions.
In MCF7, mixture model for most stimuli, except HGF, captures general trends effectively.
For UACC812, predictions for PBS, FGF1 and HGF align closely with observed data.
Overall, the model captures trends well and more prominently it captures the jump in the time points of the occurence.

4.6 Discussion

In this chapter, we applied the odeNEM framework to infer and evaluate causal networks underlying protein signaling in breast cancer cell lines. By modeling the interplay between key pathways, we successfully captured both known and novel interactions, demonstrating the model's ability to integrate noisy experimental data into meaningful biological insights. Below, we summarize the main findings and discuss their implications.
Our comparison of prior configurations (uniform, strict, and DREAM-based priors) revealed the superior performance of the DREAM-based prior in reconstructing networks across all cell lines and metrics. This finding highlights the importance of incorporating

prior biological knowledge to address the inherent complexity and variability in cancer signaling. The analysis highlighted significant cell line specific differences in signaling pathway interactions, underscoring the heterogeneity of breast cancer. For instance, BT20, a basal like cell line, showed strong alignment between observed and predicted trajectories, particularly under stimuli. This is also valid for most stimuli in MCF7 cell line.
In contrast, UACC812 exhibited greater variability in model performance. This variability might reflect biological complexities such as pathway cross-talk, feedback mechanisms, or context dependent signaling that are not fully captured by the current model.

While the model demonstrates promising results, it has limitations that should be addressed in future work. It struggles to fully capture complex cross-talk between pathways, particularly in challenging contexts, which suggests a need for additional regulatory mechanisms or expanded hidden node networks. The reliance on limited experimental data and noisy observations limits its ability to generalize, highlighting the potential value of integrating complementary datasets, such as RNA-seq or proteomics, for greater robustness. Additionally, while simplifying the hidden node network improves reliability, it may overlook subtle dynamics, which could be addressed by developing computational strategies for handling larger and more complex networks.
Also, the characteristic activation and stabilization patterns of certain stimuli, were well predicted but underestimated in late phases. This indicates that while the model can effectively capture short term dynamics, it may require further refinement to account for delayed feedback or long term stabilization effects.

In conclusion, the odeNEM framework demonstrates significant potential for uncovering causal relationships in complex signaling networks. The application of the DREAM-based prior proved to be a reliable foundation for capturing pathway dynamics, particularly in basal like cell lines such as BT20. However, the observed variability in performance across cell lines and conditions highlights areas for refinement, including the need to account for cross-talk, feedback mechanisms, and experimental variability. By integrating additional data types and refining computational strategies, future iterations of the model can achieve greater robustness and accuracy, advancing our understanding of signaling dynamics in breast cancer.

5 Inference of Gene Regulatory Networks (GRNs)

In this chapter, we apply the odeNEM model to generate gene regulatory networks (GRNs) using data obtained from time-series CRISPR perturbation experiments. We introduced odeNEM, in Chapter 3, as a method designed to reconstruct the hidden networks underlying the interactions between $S-$nodes using observations of $E-$nodes. Here, we leverage RENGE's dataset, from [83], to infer the GRN using odeNEM.
RENGE contains both a regression-based computational framework and a dataset. It infers GRNs from time series observational gene expression data to identify interactions between genes. The resulting network is then systematically compared with the GRN inferred directly by RENGE, enabling us to assess consistency, identify discrepancies, and explore unique features detected by each method.

Before diving into these comparisons, we provide an overview of the experimental data and discuss the integration process with odeNEM, emphasizing how our model contributes to validating and expanding upon the inferred gene regulatory structures.

5.1 Importance of GRNs in Understanding Cellular Functions

GRNs are vital for understanding how genes regulate cellular functions. They map the transcriptional and post-transcriptional relationships between genes, proteins, and other regulatory components, revealing mechanisms of growth, differentiation, and cellular responses to stimuli. Beyond advancing our knowledge of cellular biology, GRNs have clinical significance. Dysregulation of these networks is associated with diseases such as cancer, autoimmune disorders, and metabolic syndromes. By identifying key regulatory genes, or "master regulators," GRN analysis provides a basis for therapeutic intervention, offering new opportunities to target disease-driving pathways.

5.1.1 Challenges and Advances in GRN Inference

Inferring GRNs is inherently challenging due to the complexity and dynamic nature of biological systems. Traditional approaches, such as co-expression analysis, often fail to distinguish between direct and indirect regulatory interactions, leading to ambiguous or incomplete network reconstructions. Snapshot-based methods, which capture gene expression at a single time point, further intensify this issue by missing the temporal sequence of regulatory events and transient interactions.
To overcome these challenges, temporal data and perturbation-based methods have become necessary. Temporal data allow for the differentiation between early, direct regulatory responses and later, downstream effects, providing a dynamic view of gene interactions. Meanwhile, perturbation experiments, such as CRISPR-based knockouts, introduce targeted disruptions that help isolate the functional impact of specific genes, improving the accuracy of causal network inference. This combination of temporal and perturbation data not only reduces ambiguity but also improves the accuracy of causal network inference.
The RENGE framework exemplifies this integration by combining time-series scRNA-seq with systematic CRISPR perturbations. Time-series data captures the progression of gene expression changes, while CRISPR enables controlled interventions to identify causal regulatory relationships. This approach addresses key questions in gene regulation, such as identifying causal relationships and distinguishing between direct and indirect effects.

5.2 RENGE's Framework

RENGE integrates time-series scRNA-seq with systematic CRISPR perturbations to reconstruct GRNs. The use of scRNA-seq captures the temporal progression of gene expression changes following perturbation, while CRISPR enables direct identification of causal regulatory relationships by systematically targeting specific genes.
A major strength of RENGE lies in its capacity to differentiate between early (direct) and late (indirect) regulatory responses. By capturing the progression of gene interactions over time, RENGE enables the reconstruction of regulatory networks that account for both transient and sustained interactions, including those involving genes not directly knocked out. This comprehensive approach addresses limitations in static and correlation-based methods, offering a more detailed and dynamic view of GRNs.
The experimental design, depicted in Figure 5.1, highlights the targeted perturbations of

23 key genes using CRISPR-based knockouts, followed by single-cell RNA sequencing across days 2 to 5 post-transduction days, capturing temporal dynamics crucial for GRN reconstruction.

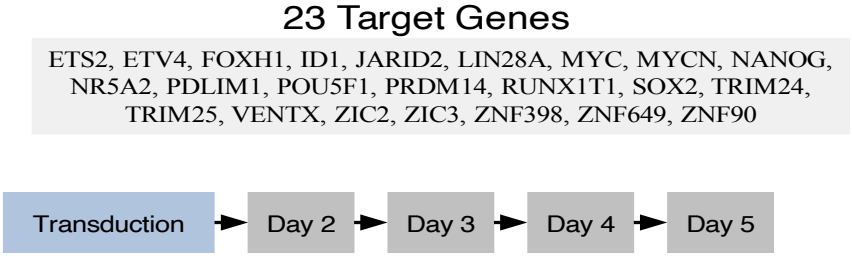

Figure 5.1: Experimental Design of RENGE: scRNA-seq Profiling Across Days 2-5 Post-Transduction with gRNA Vectors

5.2.1 RENGE Dataset

Below, we outline the dataset's composition and the experimental framework used to generate it:

Cell Types

The dataset includes data from human induced pluripotent stem cells (hiPSCs), which were chosen due to their ability to differentiate into various cell types, making them ideal for studying complex regulatory mechanisms. These characteristics make hiPSCs an excellent model for understanding regulatory dynamics at the early stages of cell differentiation and for identifying key factors that maintain pluripotency .

Temporal Data Collection and Perturbation

To capture the temporal dynamics of CRISPR-mediated gene knockouts, samples were collected on days 2, 3, 4, and 5 post-transduction from hiPSCs. The experiment targeted 23 genes (along with 4 control gRNAs), with two gRNAs per gene, ensuring coverage of the pluripotency network. Cells were transduced with CRISPR gRNAs and sorted for expression of gRNAs using an MA900 cell sorter. Approximately 5,000 single cells were captured per sample using the 10x Genomics Chromium platform, aiming for 100 cells

per gRNA. Single-cell RNA sequencing (scRNA-seq) was performed using the Illumina NovaSeq system, with 26 cycles for read 1, 91 cycles for read 2, and 8 cycles for the sample index, ensuring high sequencing depth.

Cell Culture and Experimental Conditions

The hiPSCs (OILG-3 line) were cultured under controlled conditions that supported their growth and maintenance of pluripotency. Cells were grown on specific substrates and treated with factors that ensured their viability and ability to respond predictably to gene perturbations. By using consistent culture conditions, the dataset captures genuine biological responses to perturbations rather than variability introduced by differing growth environments.

5.2.2 Data Processing and Normalization

Once the sequencing data was obtained, it underwent several processing steps to prepare it for GRN inference. First, the raw sequencing reads were aligned to the reference genome to determine which genes were expressed in each cell. Cells that did not meet certain quality criteria were filtered out to ensure that only high-quality data was used. For example, cells expressing >200 and <10,000 genes were preserved and cells with more than 20% mitochondrial reads were excluded.
Next, the expression data was normalized to account for differences in sequencing depth across cells. This normalization step is crucial, as it allows for meaningful comparisons of gene expression levels between cells, ensuring that the inferred regulatory relationships are reliable. Additionally, each cell was assigned a gRNA label, indicating which gene had been targeted, which helped link the observed changes in gene expression to specific gene knockouts. In all these steps, the Seurat software package was used for data processing, including quality control, normalization, and scaling of gene expression data [84].
The variance vs. mean plot, Figure 5.2, illustrates the relationship between gene expression variability and mean expression levels. We expect the normalization could to some extent stabilize the variance of the observations. The plot highlights how variance initially increases with mean expression, reaches a peak, and then decreases at higher expression levels. This pattern reflects biological and technical sources of variability. The observed reduction in variance at higher expression levels suggests effective stabilization of the mean-variance relationship by normalization, ensuring comparability of gene expression across cells.

The analysis of expression values between experimental and control groups provided

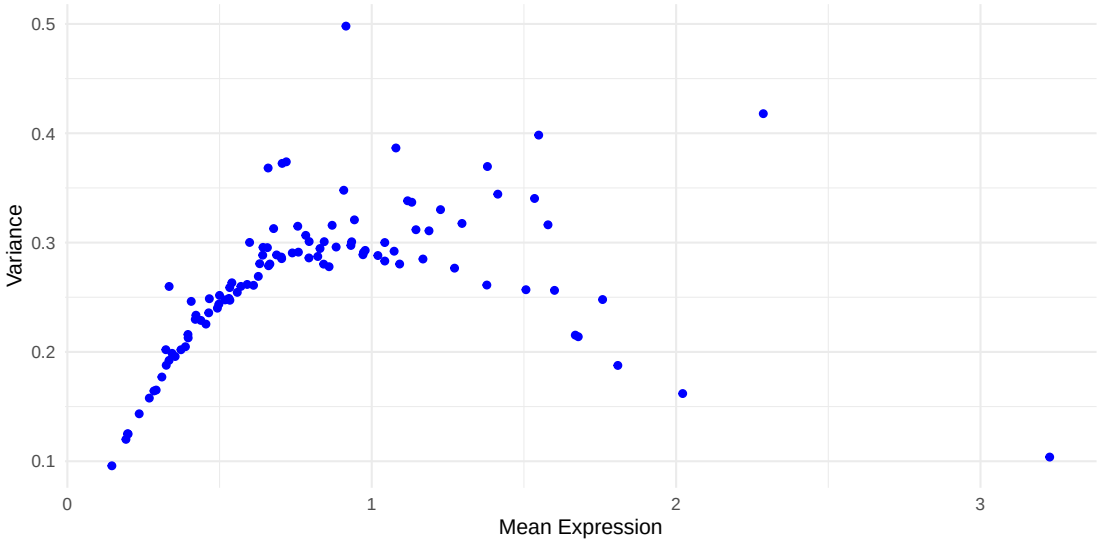

Figure 5.2: Variance and mean expression, post-normalization

critical insights into the impact of gene perturbations. The density plot in Figure 5.3 revealed a higher density of low-expression values in the experimental group compared to the control, suggesting transcriptional changes triggered by CRISPR-based perturbations. Violin plots showed variability in expression values, particularly within the experimental group. This variability reflects heterogeneity in gene expression responses to knockouts. Subtle differences between the groups, as seen in both density and violin plots, indicate potential shifts in variability at specific expression levels. Additionally, the presence of outliers in both groups underscores the inherent variability of single-cell RNA-seq data.

Figure 5.3, directly compares mean expression levels between experimental conditions and controls across all genes and timpoints. It highlights deviations from the diagonal (i.e., experimental vs. control parity). These deviations indicate experiments which affect genes more responsive. For instance, POU5F1, PRDM14, RUNX1T1 and SOX2 exhibit strong deviations, suggesting that these knockdowns are highly more effective. In contrast, genes such as ID1 and ZNF649 show minimal deviations.

This analysis demonstrates that the number and magnitude of reacting genes vary widely across different perturbations, reflecting the differential sensitivity of the regulatory network components to targeted gene knockouts.

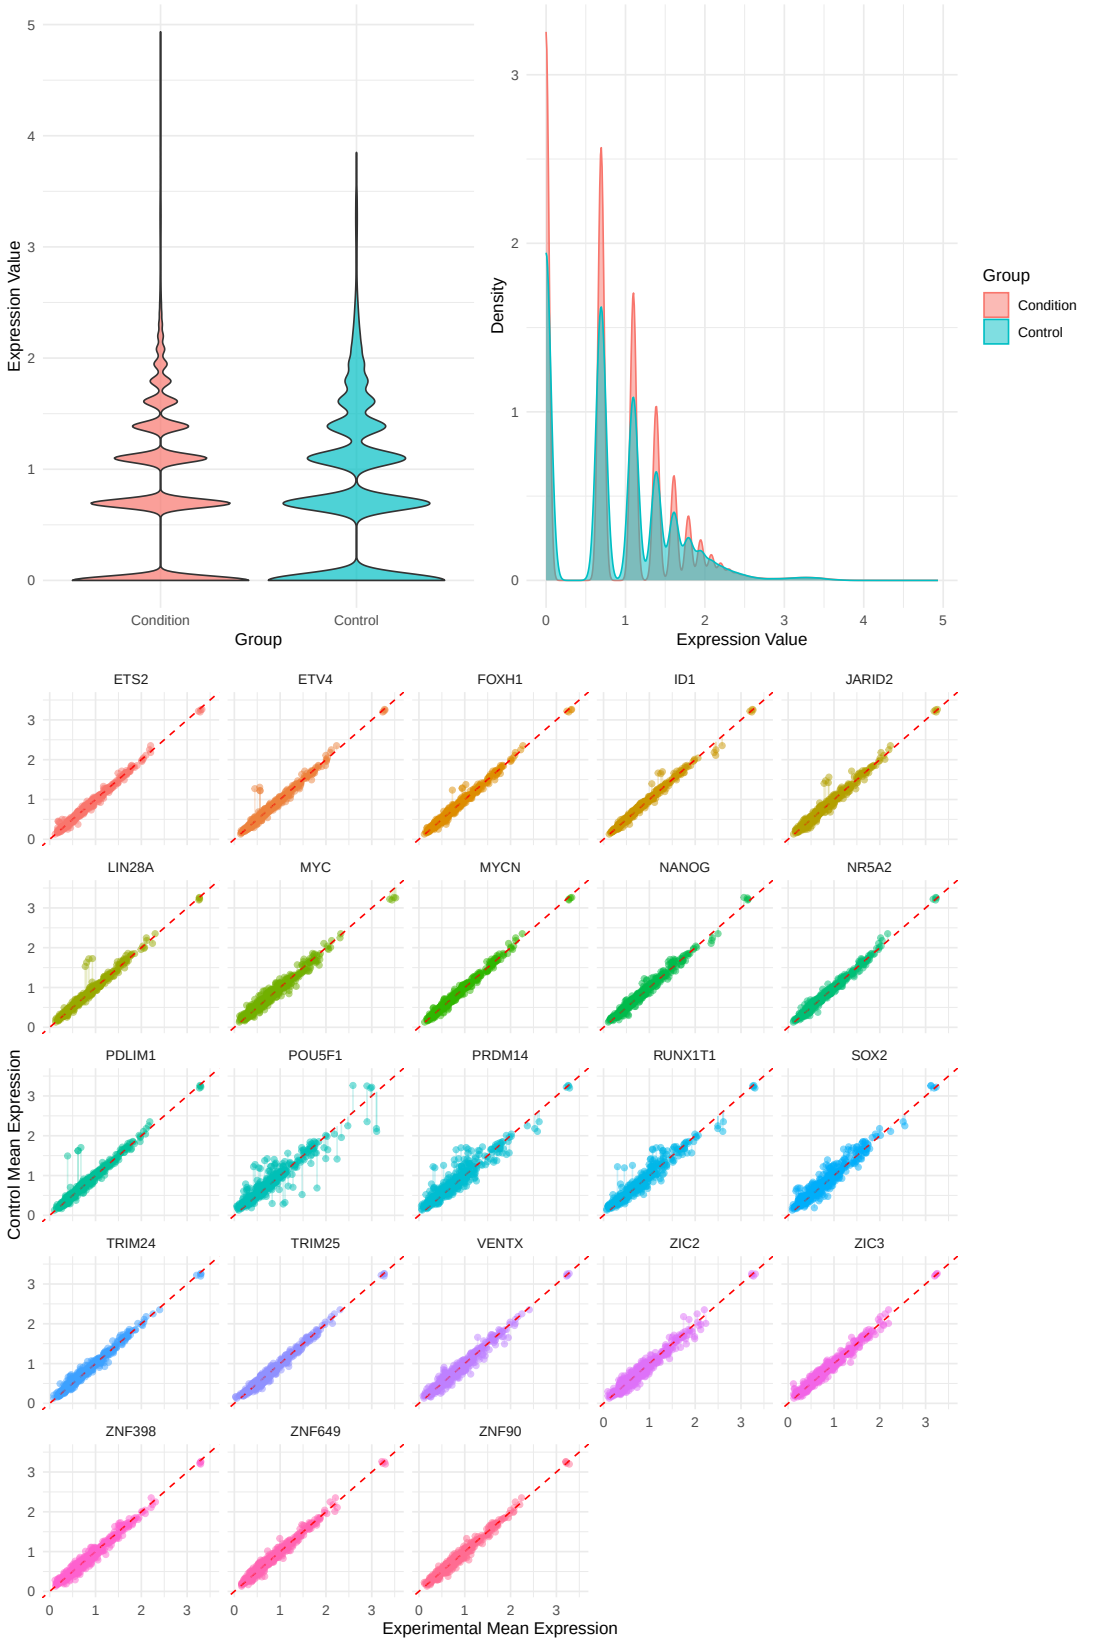

Figure 5.3: Visualizations of Gene Expression Distributions and Comparisons Across Conditions : Violin plots of the aggregated gene expression distributions in the control condition (turquoise) versus all perturbations (pink), representing expression levels across cells in the control group and cells in the experimental group. Density plot of the same data as in the violin plot, showing the distribution of expression values across all cells in the control and experimental groups. Scatterplots comparing the mean expression levels of individual genes between experimental (x-axis) and control (y-axis) conditions for each of the 23 perturbations. Each panel represents one perturbation, with points indicating individual genes in each timepoint, and the diagonal (red line) representing parity between experimental and control mean expression values.

5.2.3 RENGE Algorithm Overview

The core model can be expressed as:

$$E_{g,t} = \sum_{k=1}^{K} w(t,k,g) A^k X_g + b_t \tag{5.1}$$

where:

- E_g: is the expected gene expression vector when gene g is knocked out.
- $w(t,k,g)$: represents the effect strength at time t for regulation order k.
- A: is the gene regulatory matrix.
- X_g: is the decrease in expression of the knocked-out gene g.
- b_t: is the baseline expression level.

5.2.4 Inferring Pluripotency Network in hiPSCs

By focusing on transcription factors central to maintaining pluripotency, RENGE predicted key regulatory relationships vital to this state. Different methods, including RENGE, were used to infer gene regulatory networks for 103 transcription factors, encompassing both the 23 knockout genes and additional transcription factors showing significant changes in expression. Compared to other methods, such as GENIE3 and MIMOSCA, RENGE demonstrated superior accuracy, especially in identifying interactions involving non-knockout genes (see Figure 2 and 4 in [83]). To validate the inferred networks, a subnetwork of 19 genes was constructed using ChIP-seq data from human pluripotent stem cells as a reference (Figure 5d in [83]).
This analysis revealed both known regulators of pluripotency, such as POU5F1, NANOG, and SOX2, as well as new potential regulators that may contribute to the maintenance of pluripotency. For instance, the model highlighted novel regulatory relationships involving a complex between PRDM14 and RUNX1T1, suggesting new avenues for exploring the intricate network governing stem cell fate (see Figure 6 in [83]).

5.3 Application of odeNEM

The successful application of odeNEM relies on a meaningful selection of the hidden nodes for which to infer regulatory interactions (S) and the model for which we obtain

informative measurements (E). Before looking into the data preparation for odeNEM, it is important to outline the specific requirements for setting up our model.

5.3.1 Hidden Nodes and Observations

In odeNEM, the regulatory genes targeted by CRISPR (the 23 knockout genes in Figure 5.1) are referred to as the hidden nodes. These hidden nodes will represent a regulatory network whose edges are captured indirectly by observing changes in downstream genes. In this study, observation profile contains the log fold change of mean expression of cells receiving an inhibition and the control cells for each gene and time point. This differential expression, quantified as log fold change of the means, helps in capturing the true effect of each knockout while accounting for baseline variation in gene expression. A simple diagram of data array preparation is presented in Figure 5.4.

Since the transformed expression profile is reported in a natural log, we took it to the

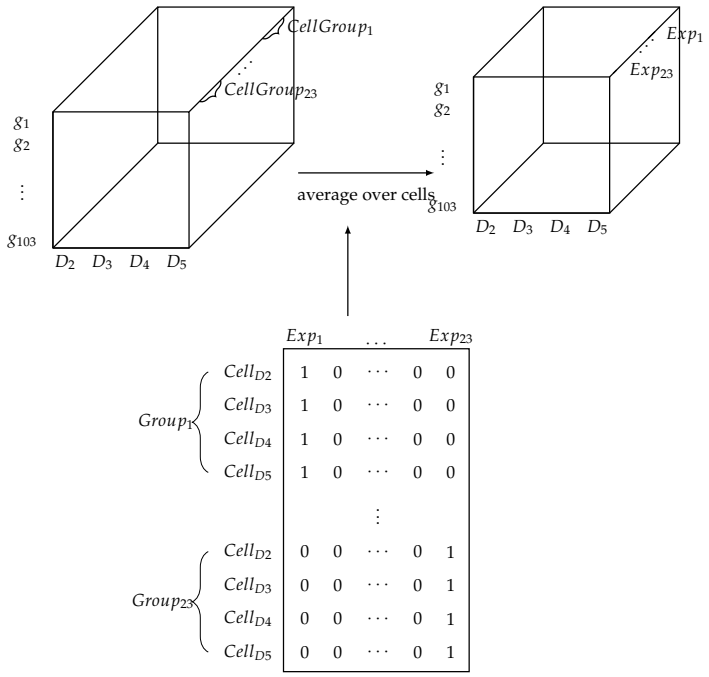

Figure 5.4: Representation of Cell Groups and Gene Expression Averaging Across Experimental Conditions.
This figure illustrates the organization of data across genes (g), Days (D), and cell groups. Each cell group corresponds to a subset of cells measured under specific experimental conditions(knock out of one of 23 target genes). The computation of the mean gene expression values across all cells in each cell group, resulted in a reduced cube.

log2 fold-change format, which is commonly used in transcriptomics studies for easier interpretation [84].

$$\text{LogFC} = \frac{mean(\text{Expression}_{g_i,D_j,ExperimentCells_k}) - mean(\text{Expression}_{g_i,D_j,ControlCells})}{log2} \quad (5.2)$$

As shown in Figure 5.5, the distribution of LogFC values was generally centered around zero, with positive and negative shifts indicating upregulation and downregulation, respectively. These observations underscore the biological variability in response to perturbations and validate the experimental design's ability to capture dynamic regulatory changes.

AIC and BIC were used to evaluate and compare the different types of mixture models

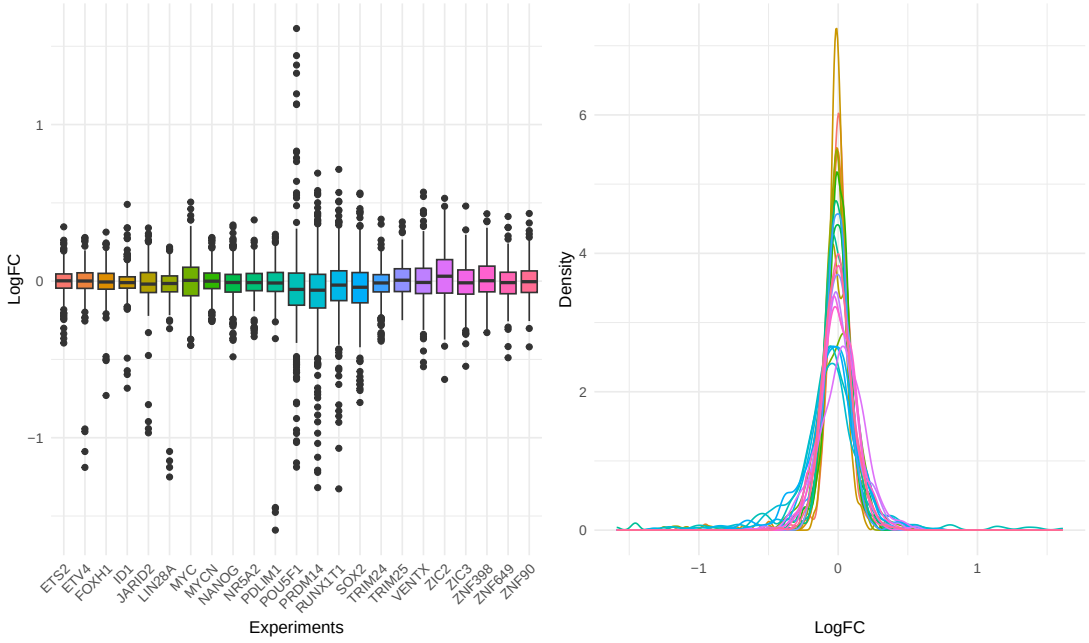

Figure 5.5: Distribution of Log Fold Change (LogFC) across experiments
Boxplots represent the distribution of LogFC values for each experiment, with the boxes showing the interquartile range (IQR). Outliers beyond this range are plotted individually as points, highlighting genes with extreme expression changes. The right panel presents density plots of LogFC values for each experiment, allowing for a visual comparison of the overall distribution and variance of LogFC values across different experiments.

applied to the gene expression data. These metrics identify the best fitting model for each gene by balancing model complexity and fit quality [85, 86]. AIC and BIC values were

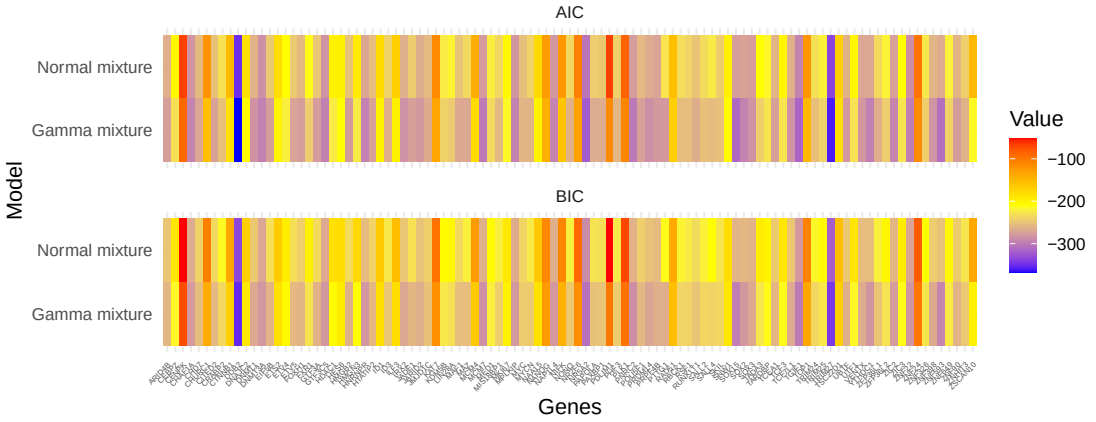

Figure 5.6: Heatmap of AIC and BIC values across genes for the tested models. Each row corresponds to a mixture model (Normal or Gamma), and each column represents a gene. The color intensity indicates the model fit, with lower AIC/BIC values (blue) representing better fits.

calculated for the Normal mixture and Gamma mixture. The heatmaps below (Figure 5.6) illustrate the AIC and BIC values across all genes, with lower values indicating better fits.

From the heatmap (Figure 5.6), it is evident that the Gamma mixture model consistently achieves lower AIC and BIC values across a majority of the genes compared to the Normal mixture model. This suggests that the Gamma mixture model better captures the variability and skewness in the gene expression data. Consequently, we employed the Gamma mixture model in Equation (3.13) to further investigate the distribution of log-transformed fold-change values:

$$o^g \sim (1 - \gamma^g) \cdot f_0^g + \gamma^g \cdot f_1^g \tag{5.3}$$

where $f_i(o^g | \alpha_i^g, \beta_i^g)$ is Gamma density function for the responsive component with parameters α_i^g (shape) and β_i^g (rate). The input to the model consists of the log-transformed fold change values for each gene. This gamma mixture model aims to capture heterogeneity in the response by fitting distinct gamma distributions to a gene expression, each representing a different regulatory effect (strongly responsive conditions f_1, weakly responsive conditions f_0).

In the E-step, the responsibilities τ^g from Equation (3.18), representing the posterior probabilities that a log-transformed fold change o^g belongs to the responsive component

(f_1), are calculated as:

$$\tau^g = \frac{\pi^g \cdot f_1^g(o^g \mid \alpha_1^g, \beta_1^g)}{\pi^g \cdot f_1^g(o^g \mid \alpha_1^g, \beta_1^g) + (1 - \pi^g) \cdot f_0^g(o^g \mid \alpha_0^g, \beta_0^g)}, \tag{5.4}$$

where π^g is the mixing proportion. The responsibilities are computed for each observation o^g, allowing the model to assign weights to the contributions of f_0 and f_1.
In the maximization step, the model parameters are updated to maximize the expected log-likelihood. The mixing proportion π^g is updated based on the average responsibility:

$$\pi^g = \frac{1}{N} \sum_{i=1}^{N} \gamma_i^g, \tag{5.5}$$

where N is the number of observations.
The shape (α_0^g) and rate (β_0^g) parameters for the weakly responsive component are updated using the method of moments:

$$\mu_0^g = \frac{\sum_{i=1}^N (1 - \gamma_i^g) \cdot o_i^g}{\sum_{i=1}^N (1 - \gamma_i^g)}, \quad \sigma_0^{g2} = \frac{\sum_{i=1}^N (1 - \gamma_i^g) \cdot (o_i^g - \mu_0^g)^2}{\sum_{i=1}^N (1 - \gamma_i^g)}. \tag{5.6}$$

$$\alpha_0^g = \frac{\mu_0^{g2}}{\sigma_0^{g2}}, \quad \beta_0^g = \frac{\sigma_0^{g2}}{\mu_0^g}. \tag{5.7}$$

The parameters for the strongly responsive component are similarly updated:

$$\mu_1^g = \frac{\sum_{i=1}^N \gamma_i^g \cdot o_i^g}{\sum_{i=1}^N \gamma_i^g}, \quad \sigma_1^{g2} = \frac{\sum_{i=1}^N \gamma_i^g \cdot (o_i^g - \mu_1^g)^2}{\sum_{i=1}^N \gamma_i^g}. \tag{5.8}$$

$$\alpha_1^g = \frac{\mu_1^{g2}}{\sigma_1^{g2}}, \quad \beta_1^g = \frac{\sigma_1^{g2}}{\mu_1^g}. \tag{5.9}$$

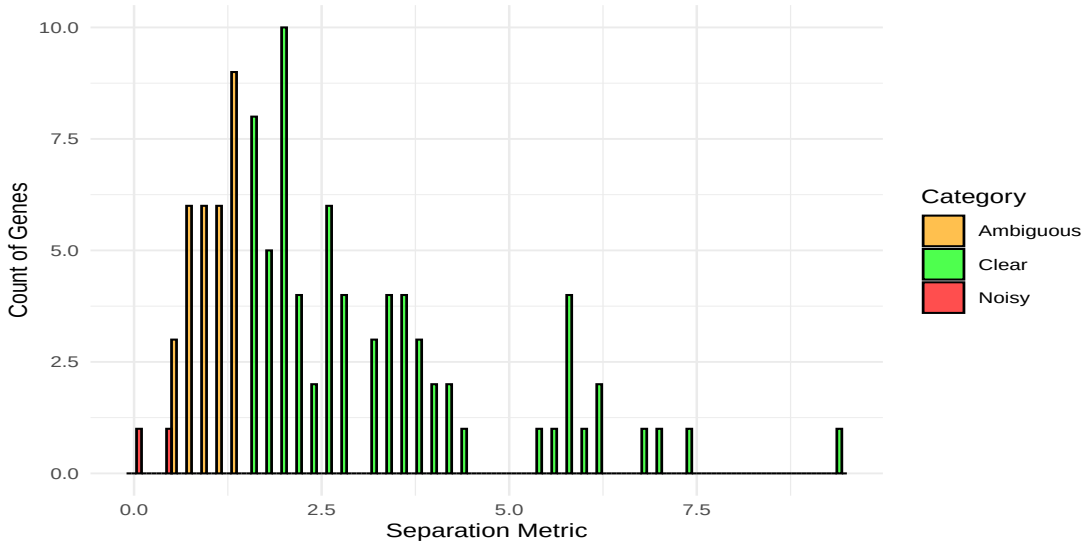

Figure 5.7: Distribution of genes by model fit category Distribution of genes across categories ("clear," "ambiguous," and "noisy") based on the separation metric, Sparation metric (difference in means) quantifies the difference between f_0 and f_1. Clear genes exhibit strong separation, ambiguous genes show moderate separation, and noisy genes demonstrate high overlap or minimal separation between the components.

Finally, genes were clustered based on the separation between f_0 and f_1 (the difference in means or overlap of the gamma distributions). Genes were categorized as "clear," "ambiguous," or "noisy" based on the separation metric. The distribution of genes across these categories is shown in Figure 5.7.
Together, these analyses demonstrate the ability of the gamma mixture model and its comparisons with other models to capture the heterogeneity in gene expression and regulatory responses.

5.3.2 Model Assumptions and Parameters

The connection between the hidden nodes and the observation nodes was established without any specific prior knowledge of the regulatory network except for 23 genes which are direct targets of inhibitions. We connected those exclusively with the corresponding hidden node. The remaining 80 informative genes were uniformly connected to all hidden nodes, ensuring that no prior bias or specific regulatory information was imposed on their relationships.

5.4 Results

Using Markov Chain Monte Carlo (MCMC), we ran multiple independent chains to better diagnose convergence and to ensure the robustness of our posterior estimates. Preliminary runs suggested potential slow mixing for certain parameters. This motivated the choice of 200,000 iterations. The selection of the burn-in period and thinning interval was carefully tailored to ensure convergence and reduce autocorrelation in the sampled chains. The burn-in length was determined empirically by monitoring trace plots of parameters for signs of stationarity. Figure 5.8 displays a representative trace plot for two parameters. It shows initial transients (burn-in phase), followed by a region of relative stationarity. Specifically, the first 100,000 samples exhibit a noticeable drift, and convergence appears to stabilize around iteration 100,000. Based on these observations, we discarded the first 100,000 iterations from subsequent analyses. To further confirm convergence, we examined the log-likelihood trace plots across the three independent MCMC chains (Figure 5.9). Each panel shows the full trajectory of log-likelihood values over iterations, with a vertical dashed line indicating the burn-in phase. Convergence was assessed based on the stabilization of log-likelihood values post burn-in. The final segments of the traces exhibit relatively minor fluctuations, supporting the assumption that the Markov chains reached a stationary distribution.

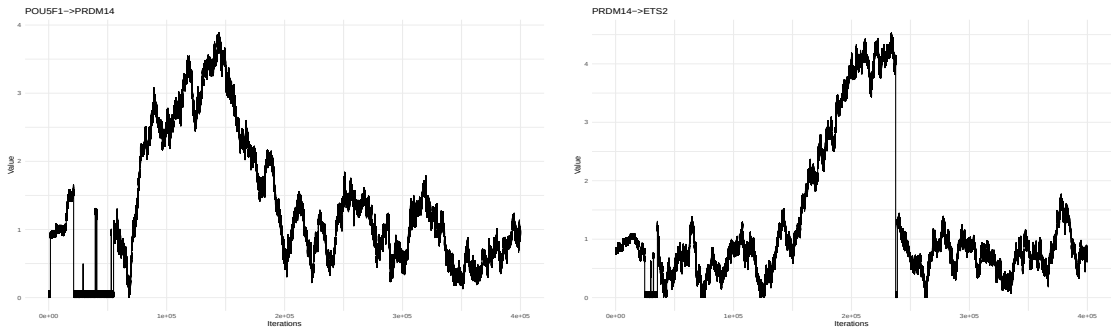

(a) POU5F1 to PRDM14 (b) PRDM14 to ETS2

Figure 5.8: Representative trace plots for two inferred gene regulatory parameters from the odeNEM model, highlighting the MCMC sampling process and convergence characteristics. Y-axis is the value sampled for that parameter (element of W matrix). Each plot spans 200,000 iterations. The first half of iterations represent the burn-in phase.

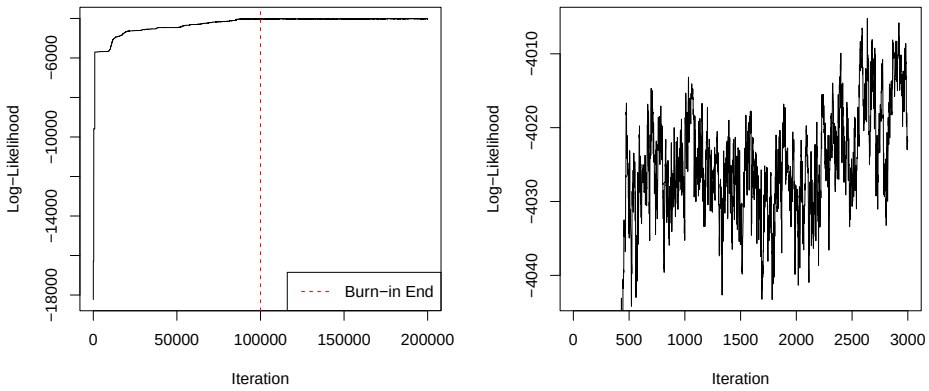

(a) Chain 1

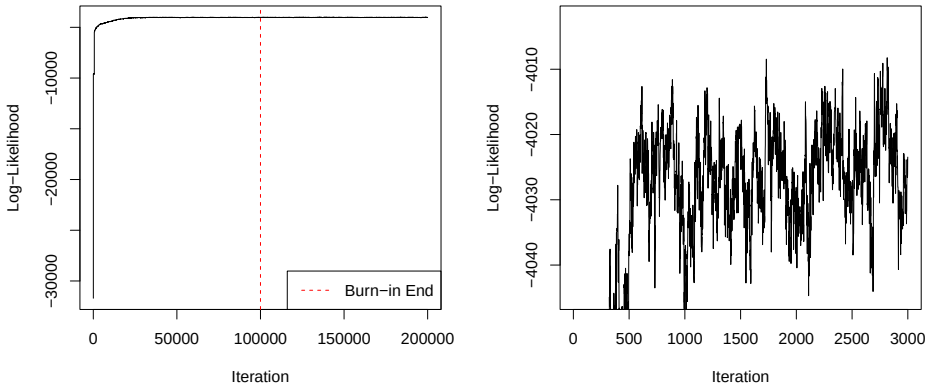

(b) Chain 2

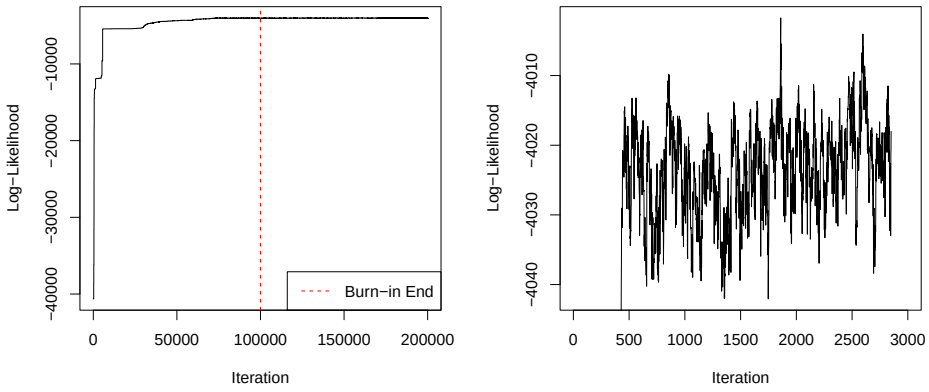

(c) Chain 3

Figure 5.9: Log-likelihood trace plots for the three independent MCMC chains. Each plot shows the progression of log-likelihood values across iterations. The red dashed line indicates the burn-in phase, after which the chains exhibit relatively stable behavior. The left plots are the whole chain, and the right plots are the window after burn-in and after thinning.

To address autocorrelation in the MCMC samples, we applied thinning by selecting every 50th sample. We determined this thinning interval by inspecting the autocorrelation function and ensuring an effective sample size (ESS) above 200 for the majority of parameters. This threshold of 200 is a commonly accepted heuristic in Bayesian inference, roughly ensuring that the Monte Carlo error remains small relative to the posterior uncertainty [87].

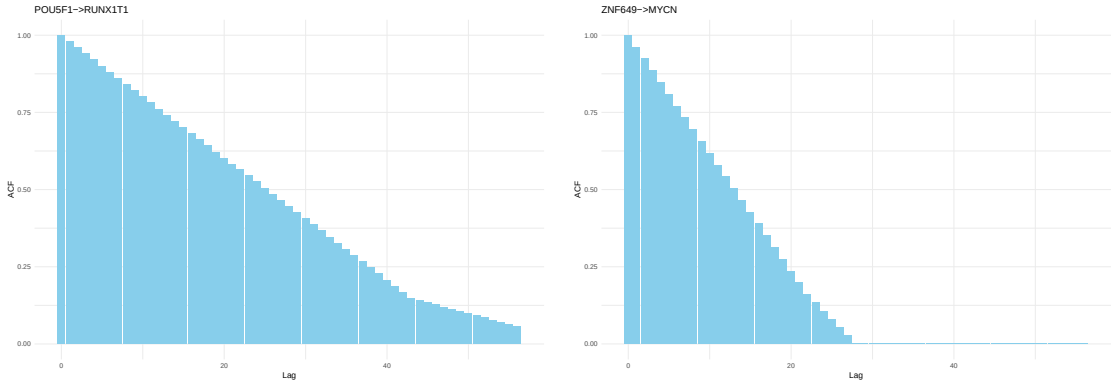

(a) POU5F1 to RUNX1T1 (b) ZNF649 to MYCN

Figure 5.10: Autocorrelation functions (ACFs) for two representative regulatory parameters taken from the MCMC posterior samples. The y-axis indicates the autocorrelation at each lag (x axis).

Figure 5.10 shows the autocorrelation function (ACF) for two representative examples of parameters. It illustrates that substantial autocorrelation persists up to lag 40, motivating our choice of a thinning interval of 50.
Figure 5.11 indicates that, despite the initial slow mixing, our final sampling design provides enough independent information to support robust downstream analyses. After thinning, each chain yielded around 2,000 effectively independent samples.
Finally, we summarized HPD for each regulatory parameter by computing the mean and 95% credible intervals (CIs) across the thinned samples, Section .0. Edges in the GRN can be classified as significant if their 95% CIs excluded zero, thereby indicating robust posterior support for the corresponding interaction. For completeness, we also reported Bayes Factors (BF) and posterior probabilities of edge presence, which reinforce the CI-based findings and offer an alternative perspective on the strength of evidence.

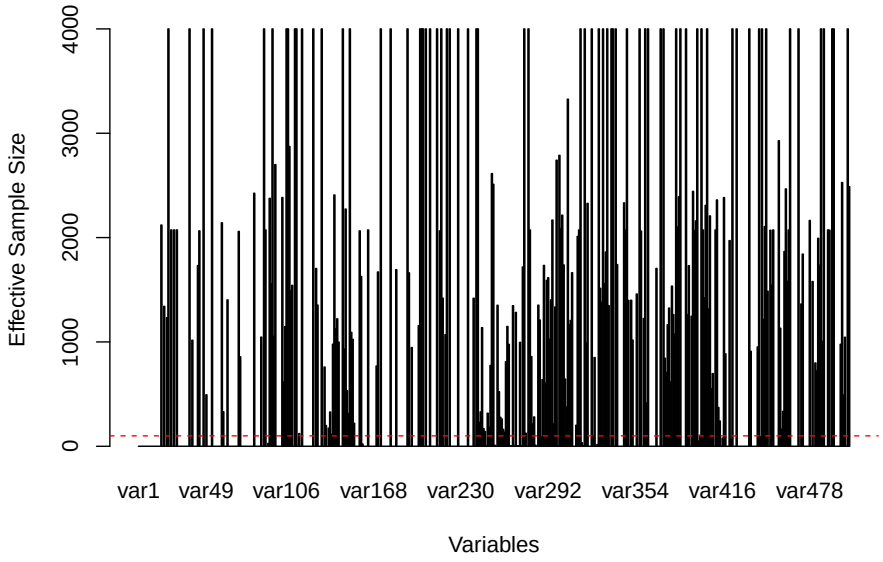

Figure 5.11: Effective sample sizes (ESS) for all parameters after burn-in and thinning. Each bar represents a parameter, and the dashed red line marks an ESS of 200, a commonly recommended threshold for stable posterior estimation.

5.4.1 Network Construction and Key Findings

Our network analysis, presented in Figure 5.12, shows edges between 23 hidden nodes. The main focus of us will be on key pluripotency factors like POU5F1, NANOG, SOX2, PRDM14.

At this point, it worths to mention shortly overview of these factors [88]. POU5F1 also known as OCT4 is a transcription factor essential for maintaining the undifferentiated state of pluripotency of stem cells. SOX2 often working in conjunction with POU5F1, is necessary in self-renewal of embryonic stem cells. NANOG reinforces the pluripotent state by inhibiting differentiation. PRDM14, as a cofactor, is a transcriptional regulator essential for maintaining pluripotency in embryonic stem cells (ESCs) and specifying primordial germ cells (PGCs) by repressing differentiation-associated genes. It also facilitates DNA demethylation, ensuring the epigenetic landscape required for pluripotency and germ cell development [89].

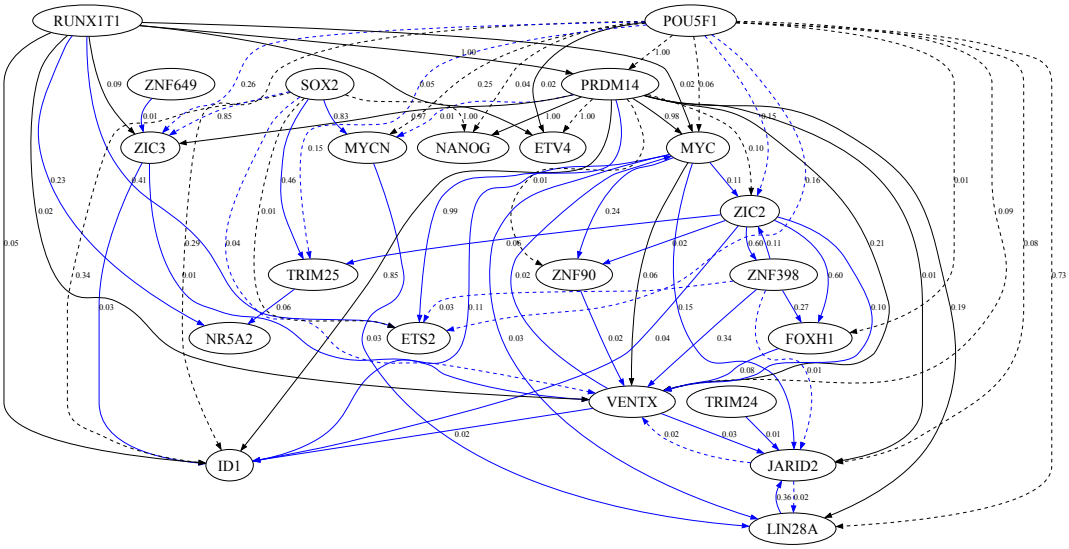

Figure 5.12: Inferred Gene Regulatory Network Capturing Pluripotency Dynamics, Edge labels represent posterior probabilities from the MCMC inference, indicating the likelihood that a regulatory interaction exists.
Black lines represent interactions supported by RENGE regulatory coefficients, and blue are edges reported only by odeNEM. Dashed lines indicating predicted edges supported by ChIP-seq data ($-10 \cdot \log_{10}(\text{MACS2 q-value}) > 50$).

Findings of odeNEM model on RENGE data sets show POU5F1, PRDM14, RUNX1T1 and SOX2 emerged as hub genes, regulating multiple targets in a scale-free structure. MYC, ZIC2 and ZNF398, less known as main elements of pluripotency core, regulate several other nodes in our inferred network.

The regulatory coefficient in the RENGE framework, [83], quantifies the influence of a transcription factor (TF) on its target genes. This coefficient reflects the strength and direction of regulatory relationships, enabling the construction of the GRN that maps how TFs control gene expression over time.

For a detailed explanation of the regulatory coefficient and its computation, refer to the Supplementary Information of the RENGE study in [83].

A summary of edge evidence between nodes, showing pairwise relationships with Bayesian statistics is included in Section .0. The table includes lower, median, and upper quantiles of the edge weights, the mean and variance of the edge distributions, the probabilities of edge presence and absence, and BF, which quantifies the strength of evidence for each edge (probabilities of edge presence/probabilities of absence). High BF

values indicate strong evidence for the presence of an edge, with values calculated based on posterior probabilities after appropriate burn-in and thinning of MCMC samples [90].

5.4.2 Validation of Regulatory Interactions Using ChIP-seq Data

To validate the inferred GRN, we constructed a ground-truth network using publicly available ChIP-seq data from the ChIP-Atlas database from [91], following the approach described in the RENGE paper. Regulatory interactions were defined based on transcription factor (TF) binding scores. Only interactions exceeding a threshold score were considered as high-confidence edges in the ground-truth network.
To systematically assess the alignment of our inferred network with experimental evidence, we varied the ChIP-seq threshold across a range of values. This allowed us to evaluate the precision and recall of our network predictions at different levels of thresholds. The Area Under the Precision-Recall Curve (AUPRC) was calculated for each threshold to measure the consistency between our inferred edges and those supported by ChIP-seq data.

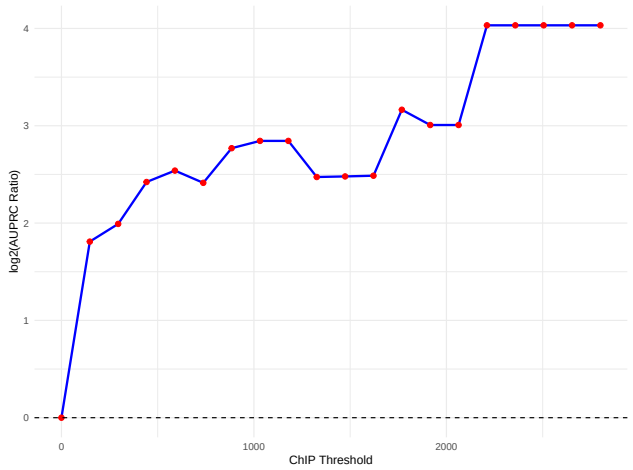

Figure 5.13: Validation of Network Predictions Using ChIP-seq Data Thresholds: AUPRC ratio of comparison between odeNEM network and ground truth from Chip-seq data under different thresholds of confidence.
The red dots represent the specific thresholds at which we evaluated the AUPRC ratio.

To calculate the AUPRC ratio, the mean AUPRC across all transcription factors was divided by the random baseline, which represents the expected performance of a random predictor (proportional to the fraction of positive interactions in the binarized ChIP-seq

data at each threshold). The log2 transformation of this ratio was then used for evaluation. This step not only provided a quantitative metric for validation but also ensured that the inferred network captured both direct and indirect regulatory interactions with high accuracy.

Figure 5.12 shows the inferred GRN with dashed edges representing regulatory interactions supported by ChIP-seq data with confidence exceeding the defined threshold. High-confidence edges are visually highlighted while retaining the original structure and color coding of the network.

Figure 5.13 presents the log2(AUPRC ratio) as a function of the ChIP-seq threshold, illustrating how the model's alignment with experimental data improves as the threshold increases. The trend reflects the model's strong ability to predict high-confidence regulatory interactions, with performance stabilizing at higher thresholds. This dual approach of visual and quantitative validation highlights the robustness of the inferred network and its consistency with experimentally validated regulatory interactions.

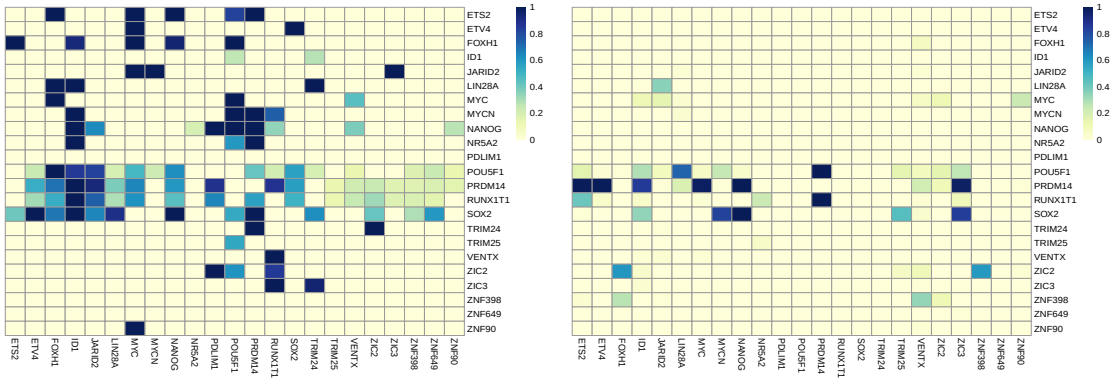

(a) RENGE regulatory effects

(b) odeNEM posterior probabilities

Figure 5.14: Comparison of Inferred Regulatory Effects between RENGE and odeNEM (a) presents RENGE regulatory effects and (b) posterior probabilities for genes from rows to genes in the columns.

5.4.3 Basic Comparison with RENGE Results

Our network shares 29 key regulatory edges with RENGE's results. To further highlight the network similarities and differences, Figure 5.14 compares the regulatory effects inferred by RENGE (left panel) with those derived from our posterior probabilities (right panel). The heatmaps illustrate shared regulatory patterns, such as the strong interactions for POU5F1 -> PRDM14, Sox2 -> NANOG, PRDM14 -> NANOG and RUNX1T1 ->PRDM14, as well as distinct differences, such as weaker connectivity for certain nodes

in the posterior results. One prominent example is the NANOG outgoing interactions from RENGE which our model did not infer those. This emphasizes the overall alignment while identifying unique edges specific to our approach. The combination of network topology and heatmap comparison underscores the robustness of our findings while also revealing potential novel regulatory relationships.

When comparing the out-degree of nodes between our network and RENGE's, key similarities and differences were observed which is summed up in Figure 5.15. Both networks identified POU5F1 and PRDM14 as the most prominent hub genes with high out-degrees. Specifically, POU5F1 has an out-degree of 15 in our network, aligning closely with its out-degree of 16 in RENGE. PRDM14 demonstrated similarly high connectivity in our network compared to RENGE's results. Nodes such as RUNX1T1 and SOX2 show comparable out-degrees in both networks, confirming their established roles in pluripotency regulation.

One key difference is NANOG node. This node exhibits a higher out-degree in RENGE (9) compared to our network (1), suggesting additional downstream regulatory targets identified by RENGE. Nodes like ZIC2, ZNF398 and MYC exhibit also variations, with our network highlighting more connections compared to RENGE.

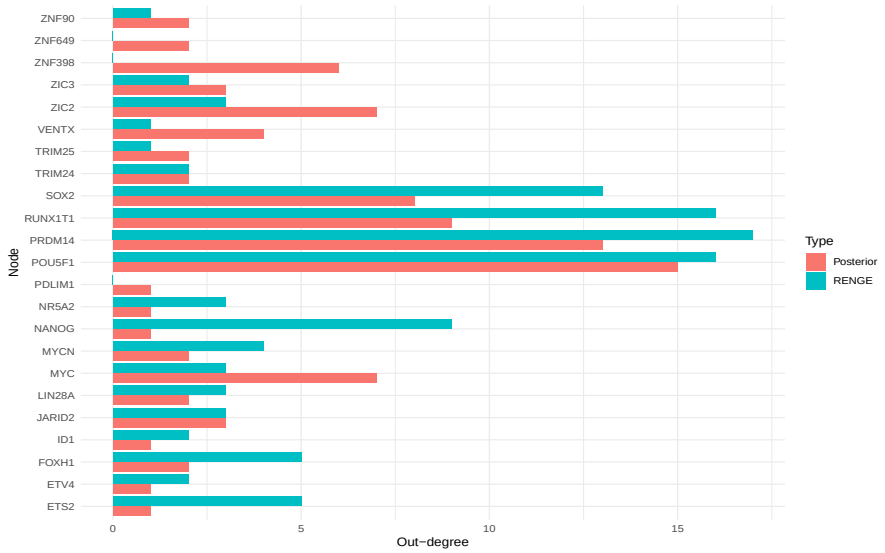

Figure 5.15: Node Out-Degree Comparison Between RENGE and odeNEM inferred Networks

5.4.4 Predictive Accuracy and Temporal Correlation Analysis

To further evaluate the predictive power of our model, we compared the predicted activity (Equation (3.16)) of hidden nodes under inhibitions versus the experimental observations. Activity levels highlight probability of change in state of a hidden variable. All following evaluations are done for first 1000 inferred networks with best likelihood values.
Figure 5.16 depicts predicted (blue) vs. observed (red) gene activity over time, demonstrating the model's performance under two different knockouts: (a) POU5F1 and (b) NANOG. For knockdown of POU5F1, Figure 5.16a, the model accurately mirrored the expression patterns. Under NANOG inhibition in Figure 5.16b, while the model captured few trends, it's deviations in certain gene activities. These results affirm the model's reliability while identifying some dynamics more challenging to infer.
To evaluate the performance of the model with more detail, we analyzed the correlation between state levels (Equation (3.11)) and observed data (Equation (5.2)) across multiple time points for all inhibitions. Figure 5.17 shows the density distribution of correlation coefficients for each time point. The density curves summarize how well the model's predictions align with experimental observations under different KOs.
Our finding for some inhibitions like POU5F1 KO is highlighting correlation near zero at early time points. As we can see, for POU5F1 KO, correlation shifts more to right as the perturbation's effect accumulates during time. This indicates that the knockout's impact expands over time, with the transcriptome diverging more significantly from its original state. This finding is consistent with the PCA findings in the RENGE paper (Figure 3.d in [83]), where the impact of knockout on the transcriptome became more pronounced over time. The progressive divergence in correlation values reflects how each experiment disrupts the regulatory network, causing cascading effects that accumulate over time. And our model could capture this progress in effects during time.
A negative correlation in this context indicates that the model failed to capture the correct direction of change in gene expression. This means that while the true expression levels increased, the predicted state values decreased, or vice versa. However, in our model, the primary focus is on the probability of change rather than the precise direction (upward or downward). As long as the model accurately captures whether a gene's expression is affected by the perturbation, the inference remains meaningful.

Figure 5.18 shows the distribution of correlation coefficients per experiment, allowing us to see under which KO the model predicts accurately (with peaks toward higher

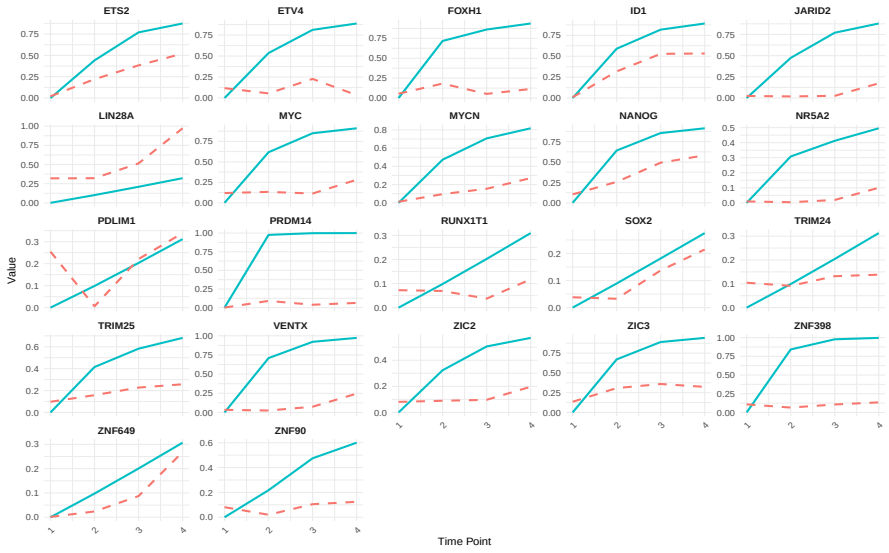

(a) POU5F1 KO

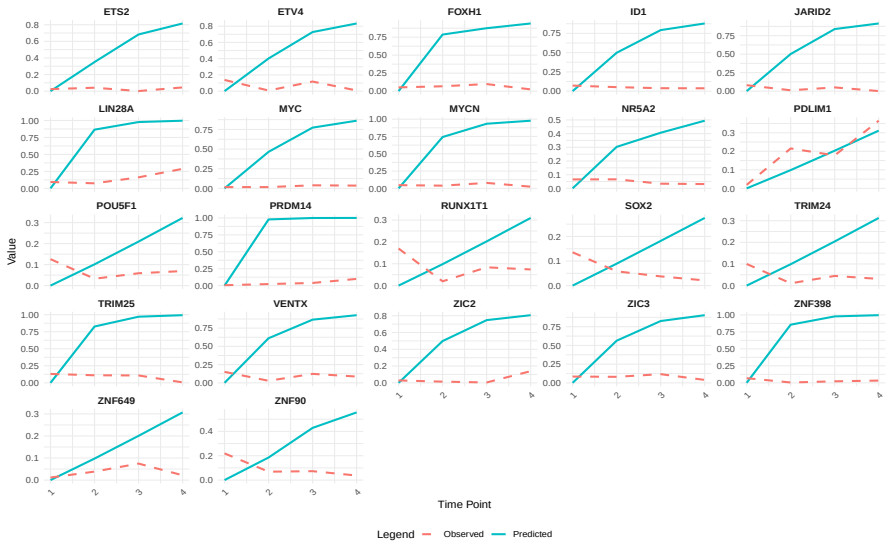

(b) NANOG KO

Figure 5.16: Model Performance in Predicting Gene Activity Under Knockouts:
Line plots showing predicted (blue) vs. observed (red) gene activity over time for key pluripotency factors (a: POU5F1, b: NANOG)
value in y-axis is predicted probability of change for blue line and gene expression change for red line.

correlations) versus those where it underperforms (peaks near zero or negative correlations). This view reveals that while several KOs (like SOX2, POU5F1, ZIC3, JARID2) exhibit stronger alignment between predicted and observed expressions, others (like NANOG) display multiple or shifted peaks, suggesting more complex dynamics.

Figure 5.19 then aggregates the correlation data across knockout experiments, but splits them by time point. Notably, later time points often shift right (toward higher correlations), indicating that the model's predictions align better once the perturbation has had time to manifest.

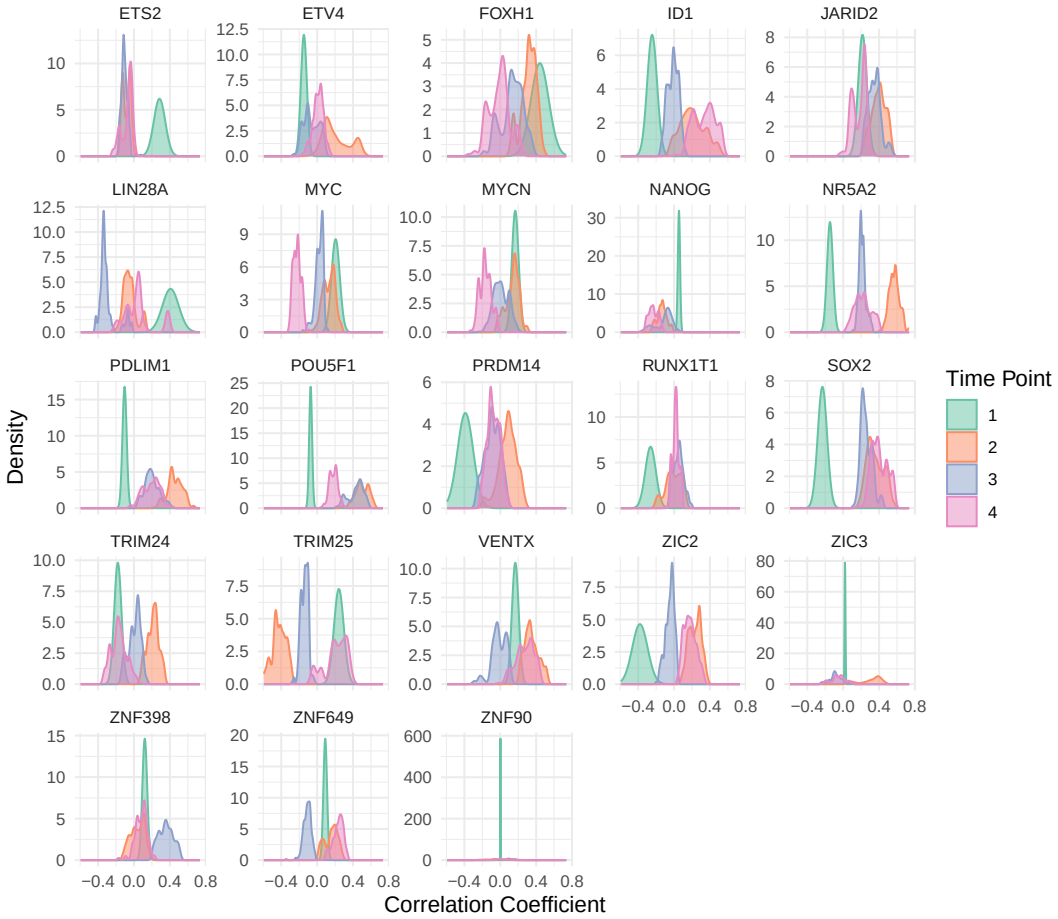

Figure 5.17: Temporal Dynamics of Model Predictions and Observations:
Density plots showing the correlation coefficients between model-predicted states and observed gene expressions over time under knockouts of key pluripotency factors.

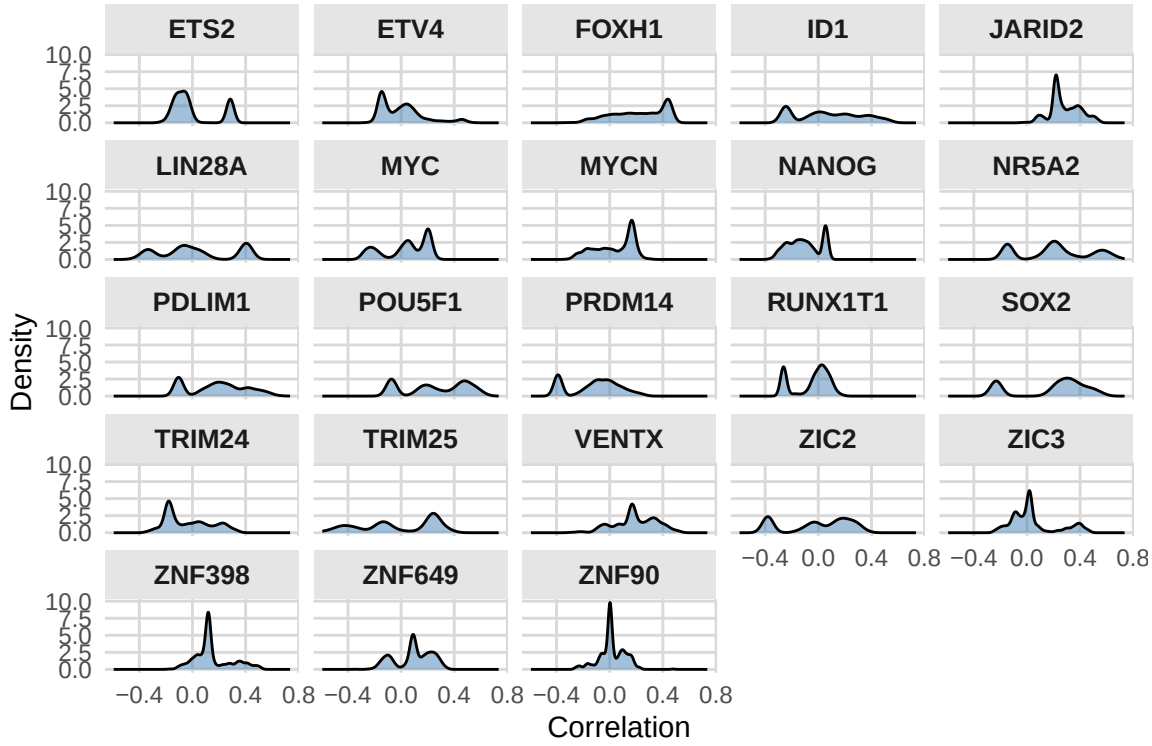

Figure 5.18: Distribution of correlation coefficients between predicted and observed gene expression across knockout conditions, aggregated over all time points.

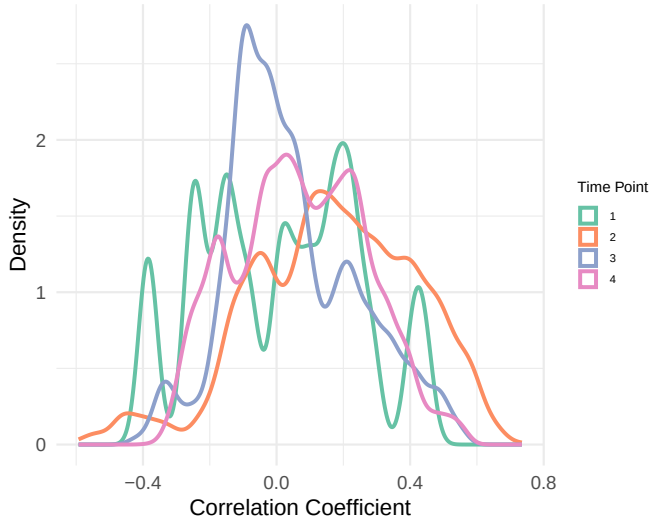

Figure 5.19: Density plot of correlation coefficients (predicted vs. observed expression) across all knockouts, stratified by time point.

5.4.5 Integration of RENGE Insights

The findings of the RENGE paper provide several key insights that align with and complement our results. By comparing these insights with our observations, we identify notable areas of agreement and divergence.

RENGE emphasizes the importance of protein-protein interactions (PPIs) between key pluripotency factors, such as POU5F1, SOX2, and NANOG, as also noted in [88]. These interactions are essential for coordinating transcriptional regulation and maintaining the undifferentiated state of stem cells. Supporting this, STRING scores ([92]) and colocalization data ([91]) further highlight the robustness of these PPIs. Specifically, in Figure 6.d [83] identifies gene pairs, such as (POU5F1, NANOG) and (SOX2, NANOG), with high regulatory correlation. These pairs likely reflect cooperative regulation, further supported by evidence from STRING scores(>900) and significant colocalization, emphasizing their roles in sustaining pluripotency through self-renewal.

In our inferred network, the edge (POU5F1, SOX2) $\rightarrow$ NANOG consistently emerges as a strong interaction. The pair (POU5F1, SOX2) exhibits a relatively low correlation of posterior probabilities (0.15) in our results, combined with a high STRING score (924) and significant affinity (<2). These findings support the hypothesis that POU5F1 and SOX2 often act in conjunction with co-regulators to maintain the pluripotent state.

Both networks highlight novel regulatory interactions that extend beyond the canonical pluripotency core. RENGE identifies PRDM14 and RUNX1T1 as a key regulatory pair, with a strong regulatory correlation (0.92) between them, suggesting a potential complex formation. In mice, such a complex has been supported by experimental evidence using liquid chromatography-tandem mass spectrometry (LC-MS/MS), indicating its potential importance in maintaining pluripotency [93].

In our inferred network, PRDM14 and RUNX1T1 also emerge as significant regulators, both targeting several downstream nodes. Although the correlation between their posterior probabilities in our network is relatively modest (0.34), their shared targets and connectivity patterns align with the idea of cooperative regulation. These findings suggest that PRDM14 and RUNX1T1 may work in conjunction to maintain pluripotency, consistent with RENGE's hypothesis of their regulatory partnership.

In addition, our inferred network highlights novel connections which were not explicitly identified in the STRING, Chip-seq data or RENGE model. Notably, edges such as PRDM14 $\rightarrow$ ETS2, ZIC2 $\rightarrow$ FOXH1, SOX2 $\rightarrow$ MYCN, SOX2 $\rightarrow$ ZIC3, and ZIC2 $\rightarrow$ ZNF398 emerged with high posterior probabilities (>0.5) in our analysis. These novel regulatory interactions may represent previously uncharacterized relationships in the

pluripotency network, suggesting the potential for unique regulatory mechanisms that were not captured in earlier studies. Such findings need further experimental validation to assess their biological significance and to potentially extend insights derived from RENGE.

5.5 Discussion

This chapter was built upon the data of the RENGE paper for inferring a GRN of pluripotency dynamics using the odeNEM model. Our inferred network not only aligns with many of RENGE's key findings but also highlights novel regulatory interactions, emphasizing the robustness and potential of our approach.
One notable finding is the identification of strong regulatory edges such as (POU5F1, SOX2) $\rightarrow$ NANOG, which were also highlighted in RENGE as central to maintaining pluripotency. RENGE's identification of PRDM14 and RUNX1T1 as a regulatory pair with a high regulatory correlation is also supported by our network, where these genes share several downstream targets.
Interestingly, our network highlights novel regulators like ZIC2, ZNF398, and MYC, which exhibit higher connectivity compared to RENGE's network. These findings suggest additional layers of regulatory complexity in the pluripotency network that warrant further investigation.
Discrepancies between our findings and RENGE's, such as the lower out-degree of NANOG in our network, may reflect differences in model assumptions or methodological approaches. NANOG's role as a pluripotency factor is well-established, and its limited connectivity in our network suggests potential areas for refinement or experimental validation.
Validation of our network using ChIP-seq data further supports the inferred regulatory interactions. The observed improvement in AUPRC at higher ChIP-seq thresholds demonstrates the consistency of our network with experimental evidence. Also, our model demonstrated considerable predictive accuracy in capturing gene activity dynamics under perturbations of key pluripotency factors such as POU5F1. The temporal correlation analysis revealed that, for POU5F1 knockouts, the impact on the transcriptome increases progressively over time, with correlation distributions broadening and showing bimodal patterns at later stages. These findings align with observations from the PCA plot of POU5F1 KO in the RENGE study, further supporting the robustness of our model in progress of knockout effects in time.

In conclusion, our findings align with and extend RENGE's insights into the regulatory landscape of pluripotency. By identifying novel interactions and validating them against experimental data, this study contributes to a deeper understanding of the transcriptional and epigenetic networks governing stem cell states. Future work should focus on experimentally validating novel edges and exploring the biological implications of regulatory interactions specific to our network.

6 Conclusion

Biological systems exhibit multiple layers of regulation, from DNA to mRNA to proteins and post translational modifications. The overarching goal of this thesis was to infer causal regulatory structures by bridging partial observables (gene expression data) and hidden signaling nodes whose activity is reflected only indirectly. Nested Effects Models (NEMs) provide a principled way to analyze perturbations by linking each hidden S-node to sets of observable E-nodes (genes). Our work uses ordinary differential equations to capture the time resolved propagation of signals among S-nodes, with the chosen form $\dot{S}(t) = W^T \tanh(S(t)) - \alpha S(t)$. This highlights saturable, bounded dynamics. We also emphasize the synergy of data (time series expression profiles) and domain knowledge (priors about edges or functional modules).
This integrated approach allows the model to exploit multiple time points and combinatorial perturbations, leading to more confident inferences about who regulates whom in the hidden signaling layer. The resulting graph Φ (or adjacency matrix W) can be interpreted as a causal structure from which we can study downstream effects and hypothesize novel interventions.

6.1 Key Strengths of the Model

One key strength is the ability to perform causal inference on time series data. Our model captures how the system evolves under perturbations across multiple time points. As it was shown, Figure 3.3, our model handle simplest motif as transient or feed forward where static snapshots are insufficient. Another important advantage is the capacity to handle combinatorial interventions across many S-nodes, (Section 4.3.1). This feature extends classical NEMs that usually focus on simpler or single gene knockdowns.
Our model is built on established NEM ideas and preserves the nested effects principle, such that the presence of an edge $S_i \rightarrow S_j$ implies partially overlapping sets of E-genes for each perturbed node. Consequently, odeNEM extends the original NEM framework as it keeps the familiar Bayesian interpretation. A further strength lies in the possibility of exploiting hierarchical by allowing priors on edges. For instance from

known pathways, the model can filter out implausible structures, thereby improving the speed and correctness of convergence. Lastly, the continuous saturable dynamics based on tanh resemble many biochemical reaction forms (such as Michaelis–Menten). This will cover broad applicability to transcriptional, signaling, or metabolic contexts.

6.2 Limitations and Challenges

Despite these strengths, certain limitations remain. First, a reliance on strong priors can bias the model if those priors are incomplete or inaccurate. This makes it difficult to unravel biological relationships from prior assumptions without careful calibration or hierarchical methods. Second, fitting is slow on large networks: MCMC or other search strategies become computationally heavy as the number of S-nodes increases, especially since each candidate network requires solving ODE trajectories. Third, the biological data may exhibit abrupt or switch-like changes that a smooth tanh-based ODE cannot always capture, resulting in potential mismatches if regulation is more stepwise or bursty. Fourth, our mixture model relies on a "baseline" distribution f_0 and an "altered" distribution f_1. Misspecification of f_1 can weaken the model's ability to distinguish affected from unaffected E-genes. Finally, some cellular processes are non-monotonic or even oscillatory, whereas the model's ODE tends to settle in a single stable point, complicating MCMC sampling if the data contain repeated up/down fluctuations.

6.3 Further Enhancements

Real signaling pathways can exhibit multi stability, oscillations, or bistable switches. Simple *tanh* saturations might be replaced by piecewise or delay differential equations for sharper transitions. Alternatively, the model might incorporate stochastic elements to handle single cell variability or in vivo noise. Another potential improvements to odeNEM is incorporating direction or sign of regulation. So, the adjacency matrix reflects activation or inhibition (rather than just presence and absence). This enhacement is enabling more precise interpretations of downstream effects. On the prior side, adopting more flexible or hierarchical Bayesian priors would allow the model to learn the reliability of each knowledge source. It will decrease the risk of using a single, rigid assumption base. Additionally, faster or more scalable inference methods, such as parallelized MCMC or variational techniques, could help tackle large networks efficiently. Alternatively, block wise MCMC moves instead of single altration can exploit the hierarchical structure of biology faster. Finally, accommodating non monotonic or more

complex altered distributions by replacing f_1 with either a mixture of Gaussians or a nonparametric approach could help the model adapt to more irregular expression patterns.

6.4 Final Remarks

In conclusion, this thesis connects the original NEM framework with the richness of time dependent hidden node dynamics governed by an ODE. These extensions introduce new complexities such as computational overhead, distributional assumptions, and potential mismatches with abrupt biological changes. But at the same time, they significantly expand the kinds of causal questions we can address about multi layered and time resolved cellular processes. By combining combinatorial perturbation designs, saturated ODE dynamics, and strong domain priors, we have shown that odeNEM can recover intricate regulatory relationships in silico. With further enhancements, such as sign inference, multi stability modeling, and more flexible altered distributions, we expect this approach to become increasingly robust and scalable for large-scale systems biology.

Supplementary Material

Numerically Stable Log-Space Summation for Likelihood Calculations

The numerical evaluation of Equation (3.17) needs to be performed in log space, because of numerical underflow for large products of small numbers. To calculate the sum of values represented in log space without having to transform back to absolute values, we use the following method:
Suppose we want to calculate the sum $p = p_1 + ... + p_n$ for very small values p_i. Let the log values of the p_i be given, $x_i = \log p_i, i = 1, ..., n$, and

$$m = \text{argmax}_i \left(x_i; i = 1, ..., n \right) \tag{.1}$$

Then, $\log p$ can be calculated in a numerically stable fashion by

$$\log p = \log \left(\sum_{i=1}^{n} p_i \right) = x_m + \text{log1p} \left(\sum_{i=1, i \neq m}^{n} \exp(x_i - x_m) \right) \tag{.2}$$

where log1p(x) is a faster and numerically precise implementation of the expression $\log(1 + x)$. Note that in our calculations, p_i might be zero, hence $x_i = -\infty$. One therefore has to care about cases in which all p_i are zero to avoid undefined $-\infty - (-\infty)$ expressions.

Derivation of the Corrected Formula for $\hat{\sigma}_{0j}^{(t+1)}$

- **Understanding the truncated normal distribution:** Given the truncated normal distribution $f_0(x)$ with mean $\mu_0 = 0$ and variance σ_0^2, truncated at $x = a$:

$$f_0(x) = \frac{\phi\left(\frac{x}{\sigma}\right)}{1 - \Phi\left(\frac{a}{\sigma}\right)}, \quad x > a \tag{.3}$$

  where $\phi(\cdot)$ is the PDF of the standard normal distribution and $\Phi(\cdot)$ is its cumulative distribution function (CDF).

- **Expected value and variance of truncated normal distribution:** The expected value and variance of a truncated normal distribution with truncation at $x = a$ are given by:

$$\mathbb{E}[X \mid X > a] = \mu + \sigma \cdot \frac{\phi\left(\frac{a-\mu}{\sigma}\right)}{1 - \Phi\left(\frac{a-\mu}{\sigma}\right)} \tag{.4}$$

  The variance X willl be,

$$\text{Var}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2 \tag{.5}$$

$$\text{Var}(X \mid X > a) = \sigma^2 \left[1 + \frac{\phi\left(\frac{a-\mu}{\sigma}\right)}{1 - \Phi\left(\frac{a-\mu}{\sigma}\right)} \left(\frac{\phi\left(\frac{a-\mu}{\sigma}\right)}{1 - \Phi\left(\frac{a-\mu}{\sigma}\right)} - \frac{a-\mu}{\sigma} \right)\right] \tag{.6}$$

  For $\mu_0 = 0$:

$$\text{Var}(X \mid X > a) = \sigma^2 \left[1 + \frac{\phi\left(\frac{a}{\sigma}\right)}{1 - \Phi\left(\frac{a}{\sigma}\right)} \left(\frac{\phi\left(\frac{a}{\sigma}\right)}{1 - \Phi\left(\frac{a}{\sigma}\right)} - \frac{a}{\sigma} \right)\right] \tag{.7}$$

- **Formulation in the EM algorithm:** In the EM algorithm, the update for σ_0^2 considering the truncated distribution should account for this additional variance term. Thus, the correct update for $\hat{\sigma}_{0j}^{(t+1)}$ becomes as Equation (3.21).

When $a = 0$, the PDF and CDF of the standard normal distribution at 0 are given by:

$$\phi(0) = \frac{1}{\sqrt{2\pi}} \tag{.8}$$

$$\Phi(0) = 0.5 \tag{.9}$$

Substituting $a = 0$ in Equation (3.21):

$$\text{Var}(X \mid X > 0) = \sigma^2 \left[1 + \frac{2}{\pi}\right] \tag{.10}$$

Dynamics of hidden layer in motifs

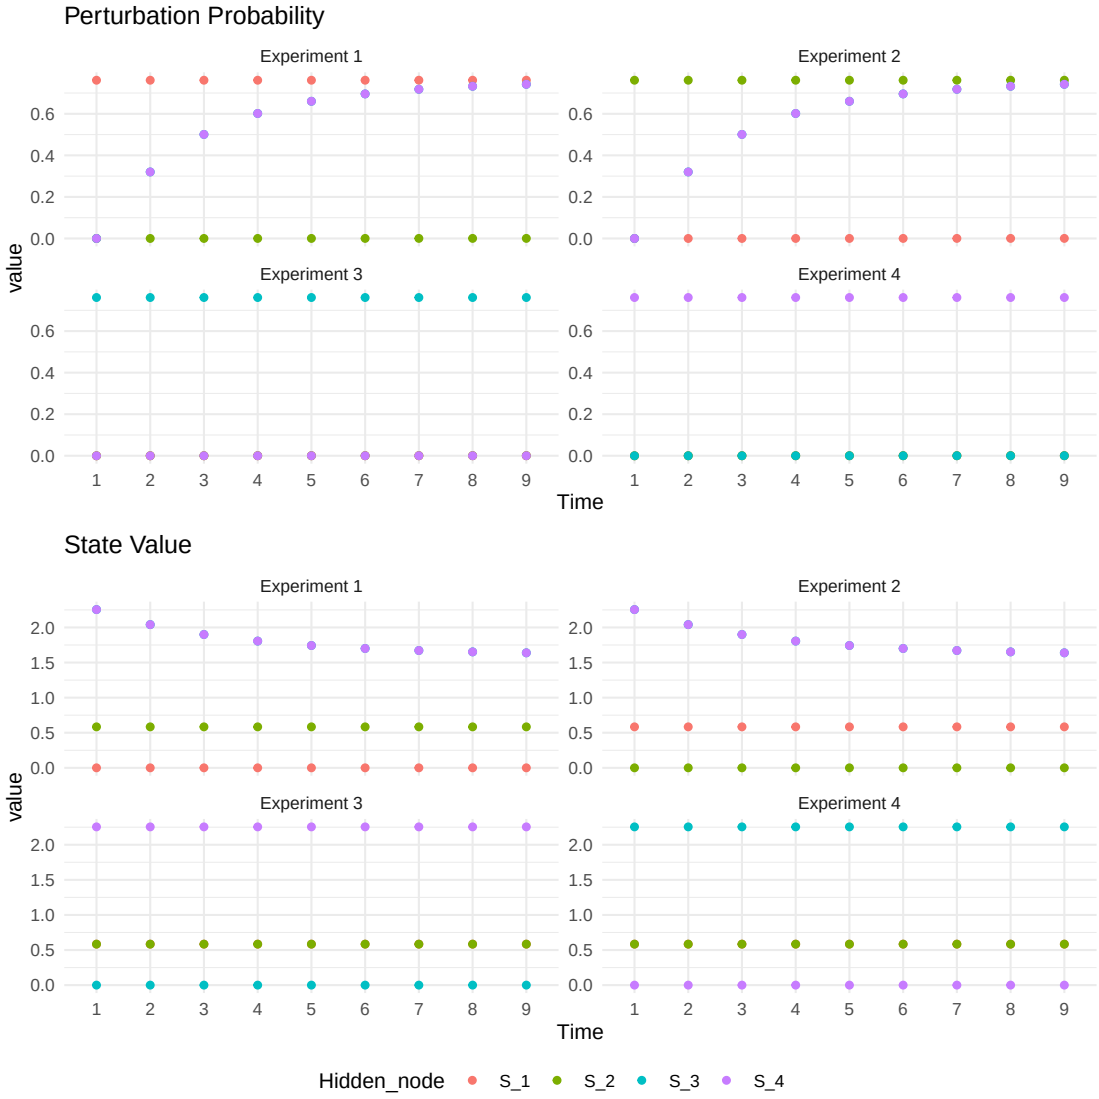

Figure 1: Dynamics of hidden layer in Bifan network

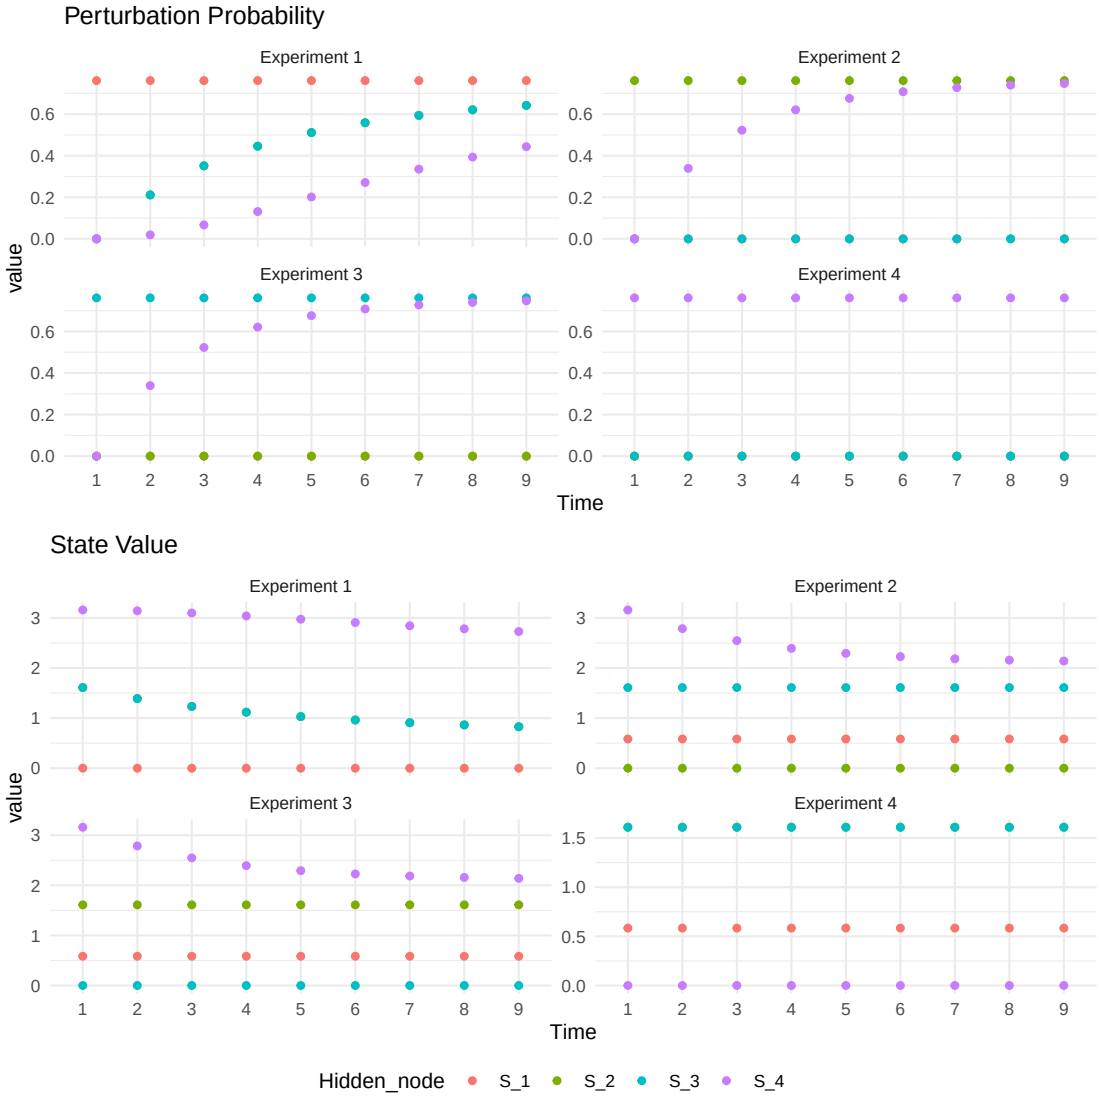

Figure 2: Dynamics of hidden layer in Diamond network

Table 1: Summary of Results

pair	median_q	upper_q	mean	var	p_present	p_absent	bf
RUNX1T1 >PRDM14	6.4746890	7.3150030	6.4480510	0.3476173	1.00000	0.0011081	902.4220000
POU5F1- >PRDM14	0.7950008	1.4914710	0.8374263	0.1365132	1.00000	0.0011909	839.7049000
SOX2- >NANOG	1.8034490	2.6281650	1.5082150	0.6468475	1.00000	0.0032024	312.2703000
PRDM14- >ETV4	1.3403420	2.8285380	1.4299390	0.5178705	0.99505	0.0049500	201.0202000
PRDM14- >NANOG	1.9560020	2.7240500	1.8834930	0.2909985	1.00000	0.0068921	145.0929000
PRDM14- >ETS2	0.8015941	4.2194960	1.3306160	1.7404720	0.99260	0.0074000	134.1351000
PRDM14- >MYC	1.9027220	4.5241130	2.1622090	1.2489650	0.97860	0.0214000	45.7289700
PRDM14- >ZIC3	0.6587175	1.8334060	0.7617404	0.2944873	0.97320	0.0268000	36.3134300
PRDM14- >ID1	1.0375690	4.1911660	1.3605710	1.4308950	0.85180	0.1482000	5.7476380
SOX2- >ZIC3	0.8757106	2.0030330	0.8912956	0.3755631	0.84590	0.1541000	5.4892930
SOX2- >MYCN	0.7880227	2.2818660	0.9316865	0.5593425	0.82510	0.1749000	4.7175530
POU5F1- >LIN28A	1.1327970	2.6074410	1.1063640	0.7953053	0.73405	0.2659500	2.7601050
ZIC2- >FOXH1	0.5652532	2.1826640	0.6398365	0.4055791	0.60310	0.3969000	1.5195260
ZIC2- >ZNF398	0.4269766	1.3946520	0.4578436	0.2112489	0.59570	0.4043000	1.4734110
SOX2- >TRIM25	0.0000000	1.5526440	0.4157347	0.2575999	0.45600	0.5440000	0.8382353
RUNX1T1- >ETS2	0.0000000	1.5786510	0.4224065	0.2898703	0.41115	0.5888500	0.6982254

LIN28A->JARID2	0.0000000	1.3574040	0.3235254	0.2202219	0.35815	0.6418500	0.5579964
SOX2->ID1	0.0000000	1.4128070	0.3141769	0.2292688	0.34005	0.6599500	0.5152663
ZNF398->VENTX	0.0000000	1.3690690	0.3334506	0.2445107	0.33830	0.6617000	0.5112589
POU5F1->ID1	0.0000000	1.3636930	0.2919352	0.2365154	0.28730	0.7127000	0.4031149
ZNF398->FOXH1	0.0000000	1.4023220	0.2681359	0.2143368	0.27335	0.7266500	0.3761784
POU5F1->ZIC3	0.0000000	1.2985350	0.2644823	0.2088777	0.26295	0.7370500	0.3567601
POU5F1->MYCN	0.0000000	1.2657650	0.2384594	0.1886451	0.25240	0.7476000	0.3376137
MYC->ZNF90	0.0000000	1.1748460	0.2017385	0.1504779	0.23715	0.7628500	0.3108737
RUNX1T1->NR5A2	0.0000000	1.2964690	0.2161474	0.1695062	0.23450	0.7655000	0.3063357
PRDM14->VENTX	0.0000000	1.2324570	0.2087356	0.1718841	0.21245	0.7875500	0.2697607
PRDM14->LIN28A	0.0000000	2.6686590	0.3463562	0.6135842	0.18745	0.8125500	0.2306935
POU5F1->ETS2	0.0000000	1.2669400	0.1583764	0.1450582	0.15620	0.8438000	0.1851150
POU5F1->ZIC2	0.0000000	1.3195990	0.1430067	0.1335299	0.14885	0.8511500	0.1748810
POU5F1->TRIM25	0.0000000	1.3119600	0.1498684	0.1427644	0.14750	0.8525000	0.1730205
MYC->JARID2	0.0000000	1.2671750	0.1433328	0.1307450	0.14555	0.8544500	0.1703435
ZNF398->ZIC2	0.0000000	1.0551990	0.0941570	0.0801639	0.11270	0.8873000	0.1270145
MYC->ZIC2	0.0000000	1.1038130	0.1071574	0.0959141	0.11200	0.8880000	0.1261261

MYC->ID1	0.0000000	1.0901400	0.1037899	0.0911164	0.11070	0.8893000	0.1244799
PRDM14->ZIC2	0.0000000	1.1254380	0.0992230	0.0963751	0.10210	0.8979000	0.1137098
ZIC2->VENTX	0.0000000	1.3045530	0.1107054	0.1208383	0.09840	0.9016000	0.1091393
RUNX1T1->ZIC3	0.0000000	1.2102070	0.0947746	0.0923294	0.09415	0.9058500	0.1039355
POU5F1->VENTX	0.0000000	1.1341540	0.0922256	0.0909104	0.08945	0.9105500	0.0982373
POU5F1->JARID2	0.0000000	1.3105490	0.0952872	0.1056370	0.08380	0.9162000	0.0914648
FOXH1->VENTX	0.0000000	1.1621010	0.0792190	0.0799257	0.07560	0.9244000	0.0817828
ZIC2->TRIM25	0.0000000	1.0123980	0.0602250	0.0587652	0.06150	0.9385000	0.0655301
POU5F1->MYC	0.0000000	1.0373670	0.0570428	0.0545404	0.05860	0.9414000	0.0622477
MYC->VENTX	0.0000000	1.0086030	0.0545427	0.0539499	0.05845	0.9415500	0.0620785
TRIM25->NR5A2	0.0000000	1.0394730	0.0567033	0.0568876	0.05650	0.9435000	0.0598834
RUNX1T1->ID1	0.0000000	0.9783141	0.0533047	0.0568467	0.05370	0.9463000	0.0567473
RUNX1T1->ETV4	0.0000000	1.0826460	0.0528603	0.0555182	0.05005	0.9499500	0.0526870
POU5F1->NANOG	0.0000000	0.9481680	0.0446980	0.0461321	0.04405	0.9559500	0.0460798
ZIC2->ID1	0.0000000	0.9884844	0.0464124	0.0520004	0.04345	0.9565500	0.0454237
SOX2->VENTX	0.0000000	0.9196117	0.0398374	0.0417424	0.03875	0.9612500	0.0403121
VENTX->JARID2	0.0000000	0.8092178	0.0315469	0.0332457	0.03030	0.9697000	0.0312468

MYC->LIN28A	0.0000000	0.9286524	0.0318945	0.0335746	0.03030	0.9697000	0.0312468
MYCN->LIN28A	0.0000000	0.8091429	0.0305633	0.0315134	0.02990	0.9701000	0.0308216
ZIC3->ID1	0.0000000	0.7885552	0.0294079	0.0309005	0.02795	0.9720500	0.0287537
ZNF398->ETS2	0.0000000	0.3485535	0.0241374	0.0235879	0.02530	0.9747000	0.0259567
JARID2->VENTX	0.0000000	0.0000000	0.0232442	0.0245158	0.02285	0.9771500	0.0233843
ZIC2->ZNF90	0.0000000	0.0000000	0.0224998	0.0230802	0.02215	0.9778500	0.0226517
RUNX1T1>VENTX	0.0000000	0.0000000	0.0196837	0.0183906	0.02185	0.9781500	0.0223381
JARID2->LIN28A	0.0000000	0.0000000	0.0500541	0.1189002	0.02070	0.9793000	0.0211376
POU5F1->ETV4	0.0000000	0.0000000	0.0198125	0.0203147	0.01965	0.9803500	0.0200439
VENTX->ID1	0.0000000	0.0000000	0.0177111	0.0172241	0.01925	0.9807500	0.0196278
ZNF90->VENTX	0.0000000	0.0000000	0.0190213	0.0205706	0.01845	0.9815500	0.0187968
VENTX->MYC	0.0000000	0.0000000	0.0171411	0.0169971	0.01800	0.9820000	0.0183299
RUNX1T1>MYC	0.0000000	0.0000000	0.0163069	0.0166596	0.01650	0.9835000	0.0167768
ZNF649->ZIC3	0.0000000	0.0000000	0.0159051	0.0170374	0.01520	0.9848000	0.0154346
TRIM24->JARID2	0.0000000	0.0000000	0.0135658	0.0131592	0.01460	0.9854000	0.0148163
SOX2->ETS2	0.0000000	0.0000000	0.0133933	0.0133330	0.01420	0.9858000	0.0144045
POU5F1->FOXH1	0.0000000	0.0000000	0.0133988	0.0140880	0.01330	0.9867000	0.0134793

ZNF398->JARID2	0.0000000	0.0000000	0.0132537	0.0139089	0.01310	0.9869000	0.0132739
PRDM14->MYCN	0.0000000	0.0000000	0.0109534	0.0102865	0.01200	0.9880000	0.0121458
PRDM14->JARID2	0.0000000	0.0000000	0.0115073	0.0117758	0.01195	0.9880500	0.0120945
PRDM14->ZNF90	0.0000000	0.0000000	0.0115591	0.0127688	0.01125	0.9887500	0.0113780
ZIC3->VENTX	0.0000000	0.0000000	0.0095233	0.0093299	0.01015	0.9898500	0.0102541

Bibliography

[1] D. Bray. "Protein molecules as computational elements in living cells." In: *Nature* 376.6538 (July 1995), pp. 307–312.

[2] F. CRICK. "Central Dogma of Molecular Biology." In: *Nature* 227.5258 (Aug. 1970), pp. 561–563.

[3] A. G. Hinnebusch. "Molecular Mechanism of Scanning and Start Codon Selection in Eukaryotes." In: *Microbiology and Molecular Biology Reviews* 75.3 (Sept. 2011), pp. 434–467.

[4] D. P. Bartel. "MicroRNAs." In: *Cell* 116.2 (Jan. 2004), pp. 281–297.

[5] K. M. Lelli, M. Slattery, and R. S. Mann. "Disentangling the Many Layers of Eukaryotic Transcriptional Regulation." In: *Annual Review of Genetics* 46.1 (Dec. 2012), pp. 43–68.

[6] C. T. Walsh, S. Garneau-Tsodikova, and G. J. Gatto. "Protein Posttranslational Modifications: The Chemistry of Proteome Diversifications." In: *Angewandte Chemie International Edition* 44.45 (Nov. 2005), pp. 7342–7372.

[7] C. Vogel and E. M. Marcotte. "Insights into the regulation of protein abundance from proteomic and transcriptomic analyses." In: *Nature Reviews Genetics* 13.4 (Mar. 2012), pp. 227–232.

[8] W. ZHANG and H. T. LIU. "MAPK signal pathways in the regulation of cell proliferation in mammalian cells." In: *Cell Research* 12.1 (Mar. 2002), pp. 9–18.

[9] W. A. Lim and T. Pawson. "Phosphotyrosine Signaling: Evolving a New Cellular Communication System." In: *Cell* 142.5 (Sept. 2010), pp. 661–667.

[10] H. A. Krebs and W. A. Johnson. "The Role of Citric Acid in Intermediate Metabolism in Animal Tissues." In: *Source Book in Chemistry, 1900–1950*. Harvard University Press, Dec. 1968, pp. 383–390.

[11] A. R. Fernie and M. Stitt. "On the Discordance of Metabolomics with Proteomics and Transcriptomics: Coping with Increasing Complexity in Logic, Chemistry, and Network Interactions Scientific Correspondence." In: *Plant Physiology* 158.3 (Jan. 2012), pp. 1139–1145.

[12] U. Alon. "Network motifs: theory and experimental approaches." In: *Nature Reviews Genetics* 8.6 (June 2007), pp. 450–461.

[13] A.-L. Barabási and Z. N. Oltvai. "Network biology: understanding the cell's functional organization." In: *Nature Reviews Genetics* 5.2 (Feb. 2004), pp. 101–113.

[14] M. Hecker, S. Lambeck, S. Toepfer, E. van Someren, and R. Guthke. "Gene regulatory network inference: Data integration in dynamic models—A review." In: *Biosystems* 96.1 (Apr. 2009), pp. 86–103.

[15] A. Mortazavi, B. A. Williams, K. McCue, L. Schaeffer, and B. Wold. "Mapping and quantifying mammalian transcriptomes by RNA-Seq." In: *Nature Methods* 5.7 (May 2008), pp. 621–628.

[16] R. Aebersold and M. Mann. "Mass spectrometry-based proteomics." In: *Nature* 422.6928 (Mar. 2003), pp. 198–207.

[17] P. J. Park. "ChIP–seq: advantages and challenges of a maturing technology." In: *Nature Reviews Genetics* 10.10 (Sept. 2009), pp. 669–680.

[18] C. J. Echeverri and N. Perrimon. "High-throughput RNAi screening in cultured cells: a user's guide." In: *Nature Reviews Genetics* 7.5 (Apr. 2006), pp. 373–384.

[19] P. D. Hsu, E. S. Lander, and F. Zhang. "Development and Applications of CRISPR-Cas9 for Genome Engineering." In: *Cell* 157.6 (June 2014), pp. 1262–1278.

[20] N. Friedman, M. Linial, I. Nachman, and D. Pe'er. "Using Bayesian Networks to Analyze Expression Data." In: *Journal of Computational Biology* 7.3–4 (Aug. 2000), pp. 601–620.

[21] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Sept. 2009.

[22] R. Linding, L. J. Jensen, G. J. Ostheimer, M. A. van Vugt, C. Jørgensen, I. M. Miron, F. Diella, K. Colwill, L. Taylor, K. Elder, P. Metalnikov, V. Nguyen, A. Pasculescu, J. Jin, J. G. Park, L. D. Samson, J. R. Woodgett, R. B. Russell, P. Bork, M. B. Yaffe, and T. Pawson. "Systematic Discovery of In Vivo Phosphorylation Networks." In: *Cell* 129.7 (June 2007), pp. 1415–1426.

[23] T. Pawson and J. D. Scott. "Protein phosphorylation in signaling – 50 years and counting." In: *Trends in Biochemical Sciences* 30.6 (June 2005), pp. 286–290.

[24] D. Koller and N. Friedman. *Probabilistic graphical models. Principles and techniques*. [Nachdr.] Adaptive computation and machine learning. Includes bibliographical references and index. Cambridge, Mass. [u.a.]: MIT Press, 2010. 1231 pp.

[25] F. Markowetz, D. Kostka, O. G. Troyanskaya, and R. Spang. "Nested effects models for high-dimensional phenotyping screens." In: *Bioinformatics* 23.13 (July 2007), pp. i305–i312.

[26] C. Zeller, H. Fröhlich, and A. Tresch. "A Bayesian Network View on Nested Effects Models." In: *EURASIP Journal on Bioinformatics and Systems Biology* 2009.1 (2009), p. 195272.

[27] J. PEARL. "Causal diagrams for empirical research." In: *Biometrika* 82.4 (1995), pp. 669–688.

[28] P. Weidemüller, M. Kholmatov, E. Petsalaki, and J. B. Zaugg. "Transcription factors: Bridge between cell signaling and gene regulation." In: *PROTEOMICS* 21.23–24 (Aug. 2021).

[29] S. Ramazi and J. Zahiri. "Post-translational modifications in proteins: resources, tools and prediction methods." In: *Database* 2021 (Jan. 2021).

[30] C. Agatemor, S. A. D. Middleton, and D. Toledo. "How pervasive are post-translational and -transcriptional modifications?" In: *Trends in Cell Biology* 32.6 (June 2022), pp. 475–478.

[31] A. Wagner. "How to reconstruct a large genetic network from n gene perturbations in fewer than n2 easy steps." In: *Bioinformatics* 17.12 (Dec. 2001), pp. 1183–1197.

[32] R. C. Wilson and J. A. Doudna. "Molecular Mechanisms of RNA Interference." In: *Annual Review of Biophysics* 42.1 (May 2013), pp. 217–239.

[33] S. Mocellin and M. Provenzano. In: *Journal of Translational Medicine* 2.1 (2004), p. 39.

[34] F. Markowetz, J. Bloch, and R. Spang. "Non-transcriptional pathway features reconstructed from secondary effects of RNA interference." In: *Bioinformatics* 21.21 (Sept. 2005), pp. 4026–4032.

[35] J. Tegnér and J. Björkegren. "Perturbations to uncover gene networks." In: *Trends in Genetics* 23.1 (Jan. 2007), pp. 34–41.

[36] S. Nelander, W. Wang, B. Nilsson, Q.-B. She, C. Pratilas, N. Rosen, P. Gennemark, and C. Sander. "Models from experiments: combinatorial drug perturbations of cancer cells." In: *Molecular Systems Biology* 4.1 (Jan. 2008).

[37] A. Tresch and F. Markowetz. "Structure Learning in Nested Effects Models." In: *Statistical Applications in Genetics and Molecular Biology* 7.1 (Jan. 2008).

[38] H. Froehlich, M. Fellmann, H. Sueltmann, A. Poustka, and T. Beissbarth. "Large scale statistical inference of signaling pathways from RNAi and microarray data." In: *BMC Bioinformatics* 8.1 (Oct. 2007).

[39] Y. Feng and J. Pan. "How Does Prior Distribution Affect Model Fit Indices of Bayesian Structural Equation Model?" In: *Fudan Journal of the Humanities and Social Sciences* (May 2024).

[40] S. Pounds and S. W. Morris. "Estimating the occurrence of false positives and false negatives in microarray studies by approximating and partitioning the empirical distribution of p-values." In: *Bioinformatics* 19.10 (July 2003), pp. 1236–1242.

[41] G. F. Cooper and E. Herskovits. "A Bayesian method for the induction of probabilistic networks from data." In: *Machine Learning* 9.4 (Oct. 1992), pp. 309–347.

[42] F. Kschischang, B. Frey, and H.-A. Loeliger. "Factor graphs and the sum-product algorithm." In: *IEEE Transactions on Information Theory* 47.2 (2001), pp. 498–519.

[43] C. J. Vaske, C. House, T. Luu, B. Frank, C.-H. Yeang, N. H. Lee, and J. M. Stuart. "A Factor Graph Nested Effects Model To Identify Networks from Genetic Perturbations." In: *PLoS Computational Biology* 5.1 (Jan. 2009). Ed. by A. Asthagiri, e1000274.

[44] F. Markowetz and R. Spang. "Inferring cellular networks – a review." In: *BMC Bioinformatics* 8.S6 (Sept. 2007).

[45] H. Fröhlich, M. Fellmann, H. Sültmann, A. Poustka, and T. Beissbarth. "Estimating large-scale signaling networks through nested effect models with intervention effects from microarray data." In: *Bioinformatics* 24.22 (Jan. 2008), pp. 2650–2656.

[46] B. Anchang, M. J. Sadeh, J. Jacob, A. Tresch, M. O. Vlad, P. J. Oefner, and R. Spang. "Modeling the temporal interplay of molecular signaling and gene expression by using dynamic nested effects models." In: *Proceedings of the National Academy of Sciences* 106.16 (Apr. 2009), pp. 6447–6452.

[47] H. Fröhlich, P. Praveen, and A. Tresch. "Fast and efficient dynamic nested effects models." In: *Bioinformatics* 27.2 (Nov. 2010), pp. 238–244.

[48] U. Nodelman, C. R. Shelton, and D. Koller. "Continuous Time Bayesian Networks." In: (Dec. 2013). arXiv: 1301.0591 [cs.AI].

[49] R. E. Kalman. "A New Approach to Linear Filtering and Prediction Problems." In: *Journal of Basic Engineering* 82.1 (Mar. 1960), pp. 35–45.

[50] M. K. S. Yeung, J. Tegnér, and J. J. Collins. "Reverse engineering gene networks using singular value decomposition and robust regression." In: *Proceedings of the National Academy of Sciences* 99.9 (Apr. 2002), pp. 6163–6168.

[51] T. S. Gardner, D. di Bernardo, D. Lorenz, and J. J. Collins. "Inferring Genetic Networks and Identifying Compound Mode of Action via Expression Profiling." In: *Science* 301.5629 (July 2003), pp. 102–105.

[52] K. A. Johnson and R. S. Goody. "The Original Michaelis Constant: Translation of the 1913 Michaelis–Menten Paper." In: *Biochemistry* 50.39 (Sept. 2011), pp. 8264–8269.

[53] J. Hale. *Asymptotic Behavior of Dissipative Systems*. American Mathematical Society, Jan. 2010.

[54] E. I. George and R. E. McCulloch. "Approaches for Bayesian Variable Selection." In: *Statistica Sinica* 7.2 (1997), pp. 339–373.

[55] M. Kanehisa. "KEGG: Kyoto Encyclopedia of Genes and Genomes." In: *Nucleic Acids Research* 28.1 (Jan. 2000), pp. 27–30.

[56] J. Ferlay, M. Colombet, I. Soerjomataram, D. M. Parkin, M. Piñeros, A. Znaor, and F. Bray. "Cancer statistics for the year 2020: An overview." In: *International Journal of Cancer* 149.4 (Apr. 2021), pp. 778–789.

[57] E. Orrantia-Borunda, P. Anchondo-Nuñez, L. E. Acuña-Aguilar, F. O. Gómez-Valles, and C. A. Ramírez-Valdespino. "Subtypes of Breast Cancer." In: *Breast Cancer*. Exon Publications, Aug. 2022, pp. 31–42.

[58] F. Schettini, G. Buono, C. Cardalesi, I. Desideri, S. De Placido, and L. Del Mastro. "Hormone Receptor/Human Epidermal Growth Factor Receptor 2-positive breast cancer: Where we are now and where we are going." In: *Cancer Treatment Reviews* 46 (May 2016), pp. 20–26.

[59] F. Montemurro, S. Di Cosimo, and G. Arpino. "Human epidermal growth factor receptor 2 (HER2)-positive and hormone receptor-positive breast cancer: new insights into molecular interactions and clinical implications." In: *Annals of Oncology* 24.11 (Nov. 2013), pp. 2715–2724.

[60] R. S. Finn, M. F. Press, J. Dering, M. Arbushites, M. Koehler, C. Oliva, L. S. Williams, and A. Di Leo. "Estrogen Receptor, Progesterone Receptor, Human Epidermal Growth Factor Receptor 2 (HER2), and Epidermal Growth Factor Receptor Expression and Benefit From Lapatinib in a Randomized Trial of Paclitaxel With Lapatinib or Placebo As First-Line Treatment in HER2-Negative or Unknown Metastatic Breast Cancer." In: *Journal of Clinical Oncology* 27.24 (Aug. 2009), pp. 3908–3915.

[61] X. Dai, H. Cheng, Z. Bai, and J. Li. "Breast Cancer Cell Line Classification and Its Relevance with Breast Tumor Subtyping." In: *Journal of Cancer* 8.16 (2017), pp. 3131–3141.

[62] W. Fan, J. Chang, and P. Fu. "Endocrine therapy resistance in breast cancer: current status, possible mechanisms and overcoming strategies." In: *Future Medicinal Chemistry* 7.12 (Aug. 2015), pp. 1511–1519.

[63] Z. Li, H. Wei, S. Li, P. Wu, and X. Mao. "The Role of Progesterone Receptors in Breast Cancer." In: *Drug Design, Development and Therapy* Volume 16 (Jan. 2022), pp. 305–314.

[64] A. Grigoriadis, A. Mackay, E. Noel, P. Wu, R. Natrajan, J. Frankum, J. S. Reis-Filho, and A. Tutt. "Molecular characterisation of cell line models for triple-negative breast cancers." In: *BMC Genomics* 13.1 (2012), p. 619.

[65] C. V. Grant, C. M. Carver, S. D. Hastings, K. Ramachandran, M. Muniswamy, A. L. Risinger, J. A. Beutler, and S. L. Mooberry. "Triple-negative breast cancer cell line sensitivity to englerin A identifies a new, targetable subtype." In: *Breast Cancer Research and Treatment* 177.2 (June 2019), pp. 345–355.

[66] D. Fu, Z. Hu, X. Xu, X. Dai, and Z. Liu. "Key signal transduction pathways and crosstalk in cancer: Biological and therapeutic opportunities." In: *Translational Oncology* 26 (Dec. 2022), p. 101510.

[67] D. Miricescu, A. Totan, I.-I. Stanescu-Spinu, S. C. Badoiu, C. Stefani, and M. Greabu. "PI3K/AKT/mTOR Signaling Pathway in Breast Cancer: From Molecular Landscape to Clinical Aspects." In: *International Journal of Molecular Sciences* 22.1 (Dec. 2020), p. 173.

[68] Q. Li, Z. Li, T. Luo, and H. Shi. "Targeting the PI3K/AKT/mTOR and RAF/MEK/ERK pathways for cancer therapy." In: *Molecular Biomedicine* 3.1 (Dec. 2022).

[69] H. Lu, Y. Guo, G. Gupta, and X. Tian. "Mitogen-Activated Protein Kinase (MAPK): New Insights in Breast Cancer." In: *Journal of Environmental Pathology, Toxicology and Oncology* 38.1 (2019), pp. 51–59.

[70] M. Burotto, V. L. Chiou, J.-M. Lee, and E. C. Kohn. "The MAPK pathway across different malignancies: A new perspective." In: *Cancer* 120.22 (June 2014), pp. 3446–3456.

[71] N. Takebe, L. Miele, P. J. Harris, W. Jeong, H. Bando, M. Kahn, S. X. Yang, and S. P. Ivy. "Targeting Notch, Hedgehog, and Wnt pathways in cancer stem cells: clinical update." In: *Nature Reviews Clinical Oncology* 12.8 (Apr. 2015), pp. 445–464.

[72] Y. Xia, S. Shen, and I. M. Verma. "NF-kBeta, an Active Player in Human Cancers." In: *Cancer Immunology Research* 2.9 (Sept. 2014), pp. 823–830.

[73] M. Olivier, M. Hollstein, and P. Hainaut. "TP53 Mutations in Human Cancers: Origins, Consequences, and Clinical Use." In: *Cold Spring Harbor Perspectives in Biology* 2.1 (Nov. 2009), a001008–a001008.

[74] M. Guo, Y. Peng, A. Gao, C. Du, and J. G. Herman. "Epigenetic heterogeneity in cancer." In: *Biomarker Research* 7.1 (Oct. 2019).

[75] S. M. Hill, L. M. Heiser, T. Cokelaer, M. Unger, N. K. Nesser, D. E. Carlin, Y. Zhang, A. Sokolov, E. O. Paull, C. K. Wong, K. Graim, A. Bivol, H. Wang, F. Zhu, B. Afsari, L. V. Danilova, A. V. Favorov, W. S. Lee, D. Taylor, C. W. Hu, B. L. Long, D. P. Noren, A. J. Bisberg, G. B. Mills, J. W. Gray, M. Kellen, T. Norman, S. Friend, A. A. Qutub, E. J. Fertig, Y. Guan, M. Song, J. M. Stuart, P. T. Spellman, H. Koeppl, G. Stolovitzky, J. Saez-Rodriguez, and S. Mukherjee. "Inferring causal molecular networks: empirical assessment through a community-based effort." In: *Nature Methods* 13.4 (Feb. 2016), pp. 310–318.

[76] S. M. Hill, N. K. Nesser, K. Johnson-Camacho, M. Jeffress, A. Johnson, C. Boniface, S. E. Spencer, Y. Lu, L. M. Heiser, Y. Lawrence, N. T. Pande, J. E. Korkola, J. W. Gray, G. B. Mills, S. Mukherjee, and P. T. Spellman. "Context Specificity in Causal Signaling Networks Revealed by Phosphoprotein Profiling." In: *Cell Systems* 4.1 (Jan. 2017), 73–83.e10.

[77] J. Kao, K. Salari, M. Bocanegra, Y.-L. Choi, L. Girard, J. Gandhi, K. A. Kwei, T. Hernandez-Boussard, P. Wang, A. F. Gazdar, J. D. Minna, and J. R. Pollack. "Molecular Profiling of Breast Cancer Cell Lines Defines Relevant Tumor Models and Provides a Resource for Cancer Gene Discovery." In: *PLoS ONE* 4.7 (July 2009). Ed. by M. V. Blagosklonny, e6146.

[78] K. Oda, Y. Matsuoka, A. Funahashi, and H. Kitano. "A comprehensive pathway map of epidermal growth factor receptor signaling." In: *Molecular Systems Biology* 1.1 (Jan. 2005).

[79] K. Imai, D. Tingley, and T. Yamamoto. "Experimental Designs for Identifying Causal Mechanisms." In: *Journal of the Royal Statistical Society Series A: Statistics in Society* 176.1 (Nov. 2012), pp. 5–51.

[80] J. Lewis. "Experimental Design: Ethics, Integrity, and the Scientific Method." In: *Handbook of Research Ethics and Scientific Integrity*. Springer International Publishing, 2020, pp. 459–474.

[81] P. Tigas, Y. Annadani, A. Jesson, B. Schölkopf, Y. Gal, and S. Bauer. *Interventions, Where and How? Experimental Design for Causal Models at Scale*. 2022.

[82] O. Raglan, N. Assi, J. Nautiyal, H. Lu, H. Gabra, M. J. Gunter, and M. Kyrgiou. "Proteomic analysis of malignant and benign endometrium according to obesity and insulin-resistance status using Reverse Phase Protein Array." In: *Translational Research* 218 (Apr. 2020), pp. 57–72.

[83] M. Ishikawa, S. Sugino, Y. Masuda, Y. Tarumoto, Y. Seto, N. Taniyama, F. Wagai, Y. Yamauchi, Y. Kojima, H. Kiryu, K. Yusa, M. Eiraku, and A. Mochizuki. "RENGE infers gene regulatory networks using time-series single-cell RNA-seq data with CRISPR perturbations." In: *Communications Biology* 6.1 (Dec. 2023).

[84] Y. Hao, S. Hao, E. Andersen-Nissen, W. M. Mauck, S. Zheng, A. Butler, M. J. Lee, A. J. Wilk, C. Darby, M. Zager, P. Hoffman, M. Stoeckius, E. Papalexi, E. P. Mimitou, J. Jain, A. Srivastava, T. Stuart, L. M. Fleming, B. Yeung, A. J. Rogers, J. M. McElrath, C. A. Blish, R. Gottardo, P. Smibert, and R. Satija. "Integrated analysis of multimodal single-cell data." In: *Cell* 184.13 (June 2021), 3573–3587.e29.

[85] H. Akaike. "A New Look at the Statistical Model Identification." In: *Selected Papers of Hirotugu Akaike*. Springer New York, 1974, pp. 215–222.

[86] G. Schwarz. "Estimating the Dimension of a Model." In: *The Annals of Statistics* 6.2 (Mar. 1978).

[87] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, Nov. 2013.

[88] M. Li and J. C. I. Belmonte. "Ground rules of the pluripotency gene regulatory network." In: *Nature Reviews Genetics* 18.3 (Jan. 2017), pp. 180–191.

[89] Y. Seki. "PRDM14 Is a Unique Epigenetic Regulator Stabilizing Transcriptional Networks for Pluripotency." In: *Frontiers in Cell and Developmental Biology* 6 (Feb. 2018).

[90] R. E. Kass and A. E. Raftery. "Bayes Factors." In: *Journal of the American Statistical Association* 90.430 (June 1995), pp. 773–795.

[91] S. Oki, T. Ohta, G. Shioi, H. Hatanaka, O. Ogasawara, Y. Okuda, H. Kawaji, R. Nakaki, J. Sese, and C. Meno. "Ch IP -Atlas: a data-mining suite powered by full integration of public Ch IP -seq data." In: *EMBO reports* 19.12 (Nov. 2018).

[92] D. Szklarczyk, A. L. Gable, K. C. Nastou, D. Lyon, R. Kirsch, S. Pyysalo, N. T. Doncheva, M. Legeay, T. Fang, P. Bork, L. J. Jensen, and C. von Mering. "The STRING database in 2021: customizable protein–protein networks, and functional characterization of user-uploaded gene/measurement sets." In: *Nucleic Acids Research* 49.D1 (Nov. 2020), pp. D605–D612.

[93] M. Yamamoto, Y. Suwa, K. Sugiyama, N. Okashita, M. Kawaguchi, N. Tani, K. Matsubara, A. Nakamura, and Y. Seki. "The PRDM14–CtBP1/2–PRC2 complex regulates transcriptional repression during the transition from primed to naïve pluripotency." In: *Journal of Cell Science* 133.15 (Aug. 2020).

Erklärung zur Dissertation
gemäß der Promotionsordnung vom 12. März 2020

Diese Erklärung muss in der Dissertation enthalten sein.
(*This version must be included in the doctoral thesis*)

„Hiermit versichere ich an Eides statt, dass ich die vorliegende Dissertation selbstständig und ohne die Benutzung anderer als der angegebenen Hilfsmittel und Literatur angefertigt habe. Alle Stellen, die wörtlich oder sinngemäß aus veröffentlichten und nicht veröffentlichten Werken dem Wortlaut oder dem Sinn nach entnommen wurden, sind als solche kenntlich gemacht. Ich versichere an Eides statt, dass diese Dissertation noch keiner anderen Fakultät oder Universität zur Prüfung vorgelegen hat; dass sie - abgesehen von unten angegebenen Teilpublikationen und eingebundenen Artikeln und Manuskripten - noch nicht veröffentlicht worden ist sowie, dass ich eine Veröffentlichung der Dissertation vor Abschluss der Promotion nicht ohne Genehmigung des Promotionsausschusses vornehmen werde. Die Bestimmungen dieser Ordnung sind mir bekannt. Darüber hinaus erkläre ich hiermit, dass ich die Ordnung zur Sicherung guter wissenschaftlicher Praxis und zum Umgang mit wissenschaftlichem Fehlverhalten der Universität zu Köln gelesen und sie bei der Durchführung der Dissertation zugrundeliegenden Arbeiten und der schriftlich verfassten Dissertation beachtet habe und verpflichte mich hiermit, die dort genannten Vorgaben bei allen wissenschaftlichen Tätigkeiten zu beachten und umzusetzen. Ich versichere, dass die eingereichte elektronische Fassung der eingereichten Druckfassung vollständig entspricht."

Teilpublikationen:

13.12.2024

Zahra Sadat Hajseyed Nasrollah 

Datum, Name und Unterschrift