

Microbial Communities: Network Reconstruction and Control

Ao Fu

St Peter's College

University of Oxford

A thesis presented for the degree of

Doctor of Philosophy

January 2024

Acknowledgements

I am profoundly grateful to my supervisor, Professor Aidong Yang, for his invaluable guidance, patience, and expertise throughout my PhD journey. His intellectual rigor and insightful feedback have been instrumental in shaping both this thesis and my growth as a researcher.

I must also extend my heartfelt thanks to my parents, whose unwavering love and encouragement have been my constant source of strength and motivation. Their sacrifices and belief in my abilities have been the cornerstone of my success.

To my wife, Luyu Yang, I owe a debt of gratitude that words cannot fully express. Her endless support, understanding, and companionship have been my sanctuary throughout the challenges and triumphs of this academic pursuit. Her presence in my life is a constant reminder of what truly matters.

To my dearest puppy, Pupu, the joyous spark in our family and my cherished companion. Your unwavering presence has brought warmth and affection into the long days and nights of my doctoral journey. Thank you for being my furry confidant and for the countless moments of solace and cheer.

Lastly, I would like to acknowledge my colleague, Zhongyi Hu, for his valuable contributions and collaborative spirit. Our discussions and shared moments of discovery have greatly enriched my research experience and added depth to my work.

This thesis is not just a reflection of my efforts but a testament to the collective support and guidance of these remarkable individuals. I am eternally grateful for their presence in my life and their contributions to my academic journey.

Abstract

Microbial communities, prevalent in natural and artificial settings, have garnered increasing attention for their applications in environmental enhancement, bio-engineering, and biomedicine. This thesis embarks on an exploration of microbial networks, focusing on the evaluation of existing reconstruction methods, the integration of network inference with causal inference, and the use of reinforcement learning for network control.

A critical component of this research is the comparison of existing microbial network reconstruction methods. These methods are evaluated based on various criteria, such as data availability and network sparsity, highlighting their assumptions, advantages, and limitations. This evaluation not only provides a guide for method selection but also discusses the challenges in microbial interaction studies and suggests future research directions.

The thesis then introduces the Microbial Causal Inference (MCI) framework, designed to infer microbial species interactions, particularly in scenarios with feedback loops. MCI employs structural causal models to accurately represent microbial community dynamics under equilibrium conditions. Its efficacy is demonstrated in both acyclic and cyclic cases, including feedback loops, using algorithms such as Fast Causal Inference (FCI) and Greedy Fast Causal Inference (GFCI). The framework is validated through systematic analysis using simulated data, showcasing its applicability to different community parameters.

Addressing the limitations of observational equilibrium data, the thesis proposes a set of rules and algorithms for designing experiments to obtain interventional data. These experiments are crucial for transforming Partial Ancestral Graph (PAG)-inferred structures into fully identified causal models. This approach not only enhances the understanding of microbial networks but also guides the experimental process in confirming or refining inferred causal relationships.

Further, the thesis explores network control within microbial communities, introducing a novel reinforcement learning-based method for developing control algo-

rithms. This venture into the application of reinforcement learning for network control opens new avenues for research and potential breakthroughs in this domain.

The novelty of this research lies in the development and integration of the MCI framework with existing methods, providing a robust approach to infer microbial interactions under complex conditions such as feedback loops. This framework uniquely bridges network reconstruction and causal inference, validated with innovative algorithms (FCI and GFCI). This thesis extends this novelty by addressing the challenge of inferring complete causal structures from observational equilibrium data in the presence of feedback loops. It proposes adaptive learning methods and experimental design rules to transform PAG-inferred structures into fully identified causal models, thus enhancing our understanding of microbial dynamics and providing a systematic approach for future research in causal inference within complex biological systems. Additionally, the introduction of reinforcement learning for network control is a pioneering step, offering new strategies for manipulating microbial networks.

The thesis concludes with a summary of the research findings, discussing the limitations and potential future directions. It emphasizes the need for ongoing development and refinement of methodologies to better understand and manipulate microbial community networks. The thesis posits that the integration of network reconstruction, causal inference, and reinforcement learning provides a comprehensive approach to studying and controlling microbial communities.

Table of Contents

1 Introduction **10**

1.1 Background 10

1.2 Project Aim and Thesis Structure 13

References **16**

2 Literature Review **18**

2.1 Structured Overview 18

2.2 Microbial Communities Interactions and Network Model 19

2.3 Microbial Network Reconstruction 22

2.3.1 Correlational Analysis Approach 22

2.3.2 Graphical Model Based Approach 24

2.4 Causal Inference 29

2.4.1 Constraint Based Approach 31

2.4.2 Score Based Approach 36

2.4.3 Cyclic Causal Inference 37

2.4.4 Intervention Calculus 40

2.5 Network Control: Model-Based and Data-Driven Approaches 42

2.6 Research Gaps 46

References **48**

3 Comparative Analysis of Microbial Network Reconstruction Methods **52**

3.1 Introduction 53

3.2 Network Reconstruction in Microbial Ecosystems: Methodologies and Challenges 54

3.2.1 Varieties of Microbial Interaction Networks 54

3.2.2 Challenges 56

3.3 Methods Overview 59

3.3.1 Undirected Network Inference Approaches 60

3.3.2 Directed Network Approaches . . . 65
3.3.3 Causality in Microbial Communities . . . 70
3.4 Future Directions . . . 75

References **79**

4 Microbial Causal Inference Framework for Community Network Reconstruction **82**
4.1 Introduction . . . 82
4.2 Methods . . . 87
4.2.1 Preliminaries . . . 87
4.2.2 Assumed Underlying Dynamics of Use . . . 89
4.2.3 The SCM . . . 90
4.2.4 Fast Causal Inference (FCI) and Greedy Fast Causal Inference (RFCI) . . . 93
4.2.5 Generating the Ground Truth . . . 96
4.2.6 Partially Observed System and Latent Confounding . . . 100
4.2.7 Inference of Direct Causes and Possible Cycles . . . 103
4.3 Results . . . 104
4.3.1 Metrics . . . 104
4.3.2 Performance at Different Scales . . . 106
4.3.3 Influence of Sample Sizes . . . 109
4.3.4 Comparison of Methods . . . 112
4.4 Conclusion . . . 115

References **117**

5 Adaptive Learning and Experimental Design with Interventional Data **121**
5.1 Introduction . . . 121
5.2 Problem Setting . . . 124
5.2.1 Preliminaries in Graphs . . . 124
5.2.2 From ODEs to SCMs . . . 125
5.2.3 Interventions . . . 128

5.2.4 Possible Cycles . . . 130
5.2.5 Problem Description . . . 133
5.3 Experiment Design Strategy and Illustrative Examples . . . 133
5.3.1 Breaking the Cycles . . . 133
5.3.1.1 Cycle-Breaking in DPAG Using Greedy Algorithm for Minimal Set Cover . . . 134
5.3.2 Determining Undetermined Edgemarks in DPAGs . . . 139
5.3.3 Recovering the Cycles . . . 145
5.4 Conclusion and Future Directions . . . 148

References **150**

6 Reinforcement Learning in Network Control **152**
6.1 Introduction . . . 152
6.2 Reinforcement Learning . . . 157
6.2.1 General Philosophy . . . 157
6.2.2 Reinforcement Learning Algorithms . . . 160
6.3 Case Example: A Model-Free RL Approach to Network Control . . . 163
6.4 Conclusion . . . 169

References **171**

7 Conclusion and Future Work **173**
7.1 Summary of Work and Key Contributions . . . 173
7.2 Limitations and Future Work . . . 174

List of Figures

2.1 Summary of ecological interactions. 20

3.1 Illustration of undirected and directed network 55
3.2 Illustration of variable-mediated network such as metabolite-mediated, host-mediated, and environment-mediated interactions 56

4.1 MCI pipeline 96
4.2 Example of acyclification and generating ground truth. 99
4.3 Confounder toy example: Microbial Causal Inference vs correlation analysis. 102
4.4 Scalability analysis of our pipeline MCI. 107
4.5 Sample size analysis of our pipeline MCI. 110
4.6 Comparison between methods: precision and recall for number of species 5,10,15 114

5.1 Graph of microbial community Example 5.1 and corresponding perfect intervention on X_1 130
5.2 Example in the case of cycle-breaking 140
5.3 Example in the case of edgemarks determination 144

6.1 Agent-environment interaction 158
6.2 Policy Iteration 161
6.3 Controlled ecological network of test example 164
6.4 Discretization of state space: This figure demonstrates the discretization of a continuous 3D state space into a grid of smaller, discrete cubes. The original 3D cube is partitioned into multiple smaller cubes. Each small cube represents a distinct state that the reinforcement learning agent can occupy within the 3D space. This discretization reduces the complexity of the continuous state space, enabling the agent to efficiently learn and make decisions by navigating through a finite set of states. 166

6.5 RL control scheme 167
6.6 Experiment results of test example. In each row, the first number denotes its episode number and $[x_1, x_2, x_3]$ denotes the states that agent has reached in that episode 168
6.7 Simulated system dynamics to test the control policy learned from RL: the x-axis represents the time and the y-axis represents the absolute abundances of species 169

List of Tables

2.1 Examples of biological networks 21

3.1 Summary of microbial network reconstruction methods of undirected network 77

3.2 Summary of microbial network reconstruction methods of directed network 78

4.1 Parameters selection for scalability analysis. 108

4.2 Parameters selection for sample size analysis. 111

1 | Introduction

1.1 Background

The microbiota encompasses a diverse array of microorganisms, including bacteria, archaea, viruses, and unicellular fungi, that inhabit various environments and ecosystems. These environments range from specific human sites like the gut and saliva (Sender, Fuchs, & Milo, 2016) (Qin et al., 2010) to broader ecological niches found in animals, plants, water bodies, and soil (Fierer & Jackson, 2006) (Gilbert, Jansson, & Knight, 2014). This intricate micro-world is defined by a web of interactions that shape its overall character and functionality (Faust & Raes, 2012) (Costello, Stagaman, Dethlefsen, Bohannan, & Relman, 2012). In a microbial community, interactions can be broadly categorized into two types: microbial and ecological, and in this thesis we are primarily interested in microbe-microbe interactions. Microbial interactions encompass a range of relationships between different taxa, each producing varying effects on the involved parties (Faust & Raes, 2012):

1. **Neutralism (0, 0)**: A state where interactions are either absent or insignificant.

2. **Mutualism (+, +)**: A mutually beneficial relationship where organisms cooperate or exchange metabolic products, like syntrophy or cross-feeding.

3. **Competition (-, -)**: Two species compete for limited resources, potentially leading to extinction or niche adaptation, as explained by Gause's competitive exclusion principle.

4. **Parasitism and Predation (+, -)**: One organism (parasite or predator) relies on another (host or prey) for survival, such as bacteriophages infecting bacteria or predators like *Bdellovibrio* attacking other bacteria.

5. **Commensalism (+, 0)**: One organism benefits from others without affecting them, often seen in biodegradation processes.

6. **Amensalism (-, 0)**: An organism negatively impacts another without benefit

or harm to itself, which can occur when metabolic by-products change the environment detrimentally for other species.

In exploring the intricate dynamics of bacterial communities, network analysis emerges as a potent tool. In molecular biology, graphs are routinely employed to map out the interconnections among various elements, represented as nodes in a network. The connections between these nodes, or edges, signify the interactions occurring within the network. These edges can be either directed, indicating a one-way interaction from one node to another, or undirected, denoting a bidirectional or mutual interaction. Additionally, these edges can be weighted, providing a quantitative measure of the interaction's strength. Biological networks, thus, effectively encapsulate the web of relationships among key molecular players like proteins, genes, or biomolecules, crucial in understanding the underlying physiological processes.

In the past few years, advancements in sequencing technologies have significantly enhanced our ability to obtain genomic data from microbial communities in natural settings. This technological progression has led to the discovery of numerous previously unidentified bacterial species, resulting in an exponential increase in the volume of microbial sequencing data (Proal, Lindseth, & Marshall, 2017). The advance in the technology of sequencing data enables us to characterize and infer the interactions between various microbes in a microbial community. In the realm of research, two primary types of experimental designs are prevalent: cross-sectional and longitudinal studies. Cross-sectional studies capture data from various samples at a single moment in time. This approach offers a 'snapshot' view of the microbial communities across different subjects, but lacks the ability to track dynamic changes over time. The networks derived from such studies typically take the form of undirected graphs, illustrating the pairwise associations between the abundances of different taxa. The underlying principle is that taxa with similar abundance patterns, as per a specified metric, are considered co-occurring, leading to the formation of a link in the network. Consequently, cross-sectional studies primarily yield static co-occurrence networks. On the other hand, longitudinal studies involve repeated observations over multiple time points, allowing for the analysis of how bacterial

interactions evolve. Such studies can reveal potential cause-effect relationships and their association with external factors impacting the microbial ecosystem. The networks inferred from longitudinal data are often directed, reflecting cause-effect dynamics within the observed time frame.

Another area that draws much attention in the field of microbial communities is the corresponding controlling of the system. Microbial communities, forming elaborate networks, are essential in sustaining the health of their hosts or environments (Pepper & Rosenfeld, 2012) (DeLeon-Rodriguez et al., 2013). Disruptions to these communities can lead to significant adverse outcomes. In human health, for example, an imbalance in the gut microbiota, the collective microorganisms in the intestines, has been linked to various health issues, such as irritable bowel syndrome, Clostridium difficile infection, autism spectrum disorders, obesity, and certain neurological conditions (Karczewski, Poniedziałek, Adamski, & Rzymski, 2014). Similarly, in agriculture, alterations in the rhizosphere microbiota can compromise disease resistance in crops, potentially reducing yields(East, 2013). In marine environments, changes in microbial populations can affect global climate dynamics by influencing carbon capture mechanisms (Buchan, LeCleir, Gulvik, & Gonz'alez, 2014). Restoring imbalanced microbial communities to their natural, healthy states presents an opportunity for innovative approaches in several critical areas. It holds potential not only for preventing and treating multifaceted human diseases but also for bolstering sustainable agricultural practices. Furthermore, such restorative measures could play a significant role in mitigating global warming, showcasing the far-reaching impact of maintaining microbial equilibrium.

Previous research on the control of complex networks has primarily centered on model-based approaches (Liu, Slotine, & Barab'asi, 2011) (Pasqualetti, Zampieri, & Bullo, 2014). However, developing precise models for large-scale networks is frequently a daunting and sometimes impractical task, as indicated by various studies in the field. This challenge is particularly pronounced due to the intricate and expansive nature of such networks (Bof, Baggio, & Zampieri, 2016) (Shandilya & Timme, 2011). A pertinent question arising in this context is the feasibility of de-

riving network control mechanisms solely from data, leading to the development of data-driven control policies. This inquiry delves into the potential of harnessing raw data to inform and dictate control strategies without relying on pre-established models.

The field of data-driven control for dynamical systems has seen a surge in interest recently, fueled by breakthroughs in machine learning and artificial intelligence (Levine, Finn, Darrell, & Abbeel, 2016). Traditionally, control learning from data involves a two-step process: initially, system identification is performed to create a model from existing data, followed by the computation of controls based on this model (Gevers, 2005). However, this approach can sometimes lead to inaccuracies and be time-intensive. Consequently, various direct data-driven methodologies have emerged, circumventing the need for initial system identification. These methods include (model-free) reinforcement learning (Lewis, Vrabie, & Vamvoudakis, 2012), iterative learning control, adaptive and self-tuning control mechanisms, and behavior-based approaches, offering more streamlined and efficient ways to derive control strategies from data (Markovsky & Rapisarda, 2008).

1.2 Project Aim and Thesis Structure

The primary aim of this thesis is to delve into the complex world of microbial communities, examining them through the lens of network theory and control. It seeks to:

- Analyze and compare current methodologies in microbial network reconstruction, elucidating their strengths, limitations, and applicability.
- Integrate the principles of network inference and statistical causal inference, particularly in the context of microbial communities with feedback loops.
- Develop and validate the Microbial Causal Inference (MCI) framework for understanding the dynamic interactions within microbial networks.
- Propose innovative experimental design strategies for obtaining interventional

data, thereby enhancing the accuracy of network models.

- Explore the application of reinforcement learning for effective control and manipulation of microbial networks, potentially revolutionizing approaches in biomedicine, agriculture, and environmental management.
- Synthesize these elements into a comprehensive framework, offering new insights and methodologies for the study and control of microbial communities.

The remainder of this thesis is organised as follows:

Chapter 1: Introduction — provides an overview of the background and motivation for this thesis. It sets the stage by elucidating the significance of microbial communities in various environments and their impact on human health, agriculture, and ecological systems. This chapter also outlines the research objectives and the scope of the thesis.

Chapter 2: Literature Review — delves into an extensive review of existing literature, encompassing the methodologies employed in microbial network reconstruction, advancements in statistical causal inference with a focus on complex microbial dynamics, and the evolution of reinforcement learning techniques in biological systems.

Chapter 3: Comparative Analysis of Microbial Network Reconstruction Methods — presents a critical comparison of various methodologies used in microbial network reconstruction. This chapter evaluates these methods against specific criteria, highlighting their strengths and limitations in the context of microbial interaction studies.

Chapter 4: Microbial Causal Inference Framework — introduces the Microbial Causal Inference (MCI) framework. This chapter discusses how the framework facilitates the understanding of microbial species interactions, especially under scenarios involving feedback loops, and demonstrates its application through systematic analysis.

Chapter 5: Adaptive Learning and Experimental Design with Interven-

tional Data — explores strategies for designing experimental interventions for adaptive learning in microbial networks. It elaborates on how these methodologies can transform theoretical models into practical, applicable solutions for microbial network control.

Chapter 6: Reinforcement Learning in Network Control — investigates the application of reinforcement learning for controlling microbial networks. This chapter assesses the potential of reinforcement learning algorithms in manipulating complex microbial ecosystems, paving the way for novel interventions in various fields.

Chapter 7: Conclusion and Future Directions —synthesizes the findings of the thesis, discussing the broader implications of the research. It addresses the limitations encountered during the study and proposes future research directions. Emphasis is placed on the need for continued exploration and innovation in the study of microbial communities, highlighting the potential for further advancements and applications in this field.

References

Bof, N., Baggio, G., & Zampieri, S. (2016). On the Role of Network Centrality in the Controllability of Complex Networks. *IEEE Transactions on Control of Network Systems*, *4*(3), 643–653.

Buchan, A., LeCleir, G. R., Gulvik, C. A., & Gonz'alez, J. M. (2014). Master Recyclers: Features and Functions of Bacteria Associated With Phytoplankton Blooms. *Nature Reviews Microbiology*, *12*(10), 686–698.

Costello, E. K., Stagaman, K., Dethlefsen, L., Bohannan, B. J., & Relman, D. A. (2012). The Application of Ecological Theory Toward an Understanding of the Human Microbiome. *Science*, *336*(6086), 1255–1262.

DeLeon-Rodriguez, N., Lathem, T. L., Rodriguez-R, L. M., Barazesh, J. M., Anderson, B. E., Beyersdorf, A. J., ... Konstantinidis, K. T. (2013). Microbiome of the Upper Troposphere: Species Composition and Prevalence, Effects of Tropical Storms, and Atmospheric Implications. *Proceedings of the National Academy of Sciences*, *110*(7), 2575–2580.

East, R. (2013). Microbiome: Soil Science Comes to Life. *Nature*, *501*(7468), S18–S19.

Faust, K., & Raes, J. (2012). Microbial Interactions: From Networks to Models. *Nature Reviews Microbiology*, *10*(8), 538–550.

Fierer, N., & Jackson, R. B. (2006). The Diversity and Biogeography of Soil Bacterial Communities. *Proceedings of the National Academy of Sciences*, *103*(3), 626–631.

Gevers, M. (2005). Identification for Control: From the Early Achievements to the Revival of Experiment Design. *European Journal of Control*, *11*(4-5), 335–352.

Gilbert, J. A., Jansson, J. K., & Knight, R. (2014). The Earth Microbiome Project: Successes and Aspirations. *BMC Biology*, *12*, 1–4.

Karczewski, J., Poniedziałek, B., Adamski, Z., & Rzymski, P. (2014). The Effects of the Microbiota on the Host Immune System. *Autoimmunity*, *47*(8), 494–504.

Levine, S., Finn, C., Darrell, T., & Abbeel, P. (2016). End-to-end Training of Deep Visuomotor Policies. *The Journal of Machine Learning Research*, *17*(1), 1334–1373.

Lewis, F. L., Vrabie, D., & Vamvoudakis, K. G. (2012). Reinforcement Learning and Feedback Control: Using Natural Decision Methods to Design Optimal Adaptive Controllers. *IEEE Control Systems Magazine*, *32*(6), 76–105.

Liu, Y.-Y., Slotine, J.-J., & Barab'asi, A.-L. (2011). Controllability of Complex Networks. *Nature*, *473*(7346), 167–173.

Markovsky, I., & Rapisarda, P. (2008). Data-driven Simulation and Control. *International Journal of Control*, *81*(12), 1946–1959.

Pasqualetti, F., Zampieri, S., & Bullo, F. (2014). Controllability Metrics, Limitations and Algorithms for Complex Networks. *IEEE Transactions on Control of Network Systems*, *1*(1), 40–52.

Pepper, J. W., & Rosenfeld, S. (2012). The Emerging Medical Ecology of the Human Gut Microbiome. *Trends in Ecology & Evolution*, *27*(7), 381–384.

Proal, A. D., Lindseth, I. A., & Marshall, T. G. (2017). Microbe-microbe and Host-microbe Interactions Drive Microbiome Dysbiosis and Inflammatory Processes. *Discovery Medicine*, *23*(124), 51–60.

Qin, J., Li, R., Raes, J., Arumugam, M., Burgdorf, K. S., Manichanh, C., ... others (2010). A Human Gut Microbial Gene Catalogue Established by Metagenomic Sequencing. *Nature*, *464*(7285), 59–65.

Sender, R., Fuchs, S., & Milo, R. (2016). Revised Estimates for the Number of Human and Bacteria Cells in the Body. *PLoS Biology*, *14*(8), e1002533.

Shandilya, S. G., & Timme, M. (2011). Inferring Network Topology From Complex Dynamics. *New Journal of Physics*, *13*(1), 013004.

2 | Literature Review

2.1 Structured Overview

This chapter presents a comprehensive review of the existing literature spanning various domains pivotal to understanding and analyzing microbial communities. The literature review is systematically structured into four key areas, each addressing critical aspects of microbial network studies. The goal is to encapsulate the breadth of research conducted in this field, highlighting the evolution of methodologies and their implications.

1. **Microbial Communities Interactions and Network Model** — This section delves into the intricate nature of microbial communities. It focuses on elucidating the various types of interactions among microorganisms and how these relationships are intricately woven into network models. The discussion extends to the representation of these communities in network structures, providing insights into their ecological and biological dynamics.

2. **Microbial Network Reconstruction** — Here, we explore the diverse methodologies developed for reconstructing microbial networks. This includes an examination of both traditional and advanced computational approaches, assessing their efficacy, applications, and the constraints they present in different environmental and biological contexts.

3. **Causal Inference** — In a significant shift, the review narrows its focus to causal inference literature. It thoroughly examines ancestral graphical models and structural causal models, with particular emphasis on the role and implications of cycles within these frameworks. Additionally, this section delves into the nuances of do-calculus and interventional calculus.

4. **Network Control: Model-Based and Data-Driven Approaches** — The final segment reviews strategies for network control, contrasting model-based methods with data-driven approaches. Special emphasis is placed on rein-

forcement learning, discussing its theoretical underpinnings and practical applications in controlling and manipulating microbial networks.

Through this comprehensive literature review, we aim to provide a cohesive understanding of the current state of research in microbial network analysis, setting the stage for the subsequent chapters that build upon these foundational concepts.

2.2 Microbial Communities Interactions and Network Model

Microorganisms, encompassing a diverse range from bacteria and viruses to archaea and protists, do not exist solitarily but rather form intricate networks of ecological interactions. These interactions can be beneficial, detrimental, or neutral to the species involved. The nature of these interactions leads to various categorizations of relationship types. For example, in the realm of beneficial interactions, or mutualism, different bacterial species might synergize to degrade complex compounds like pollutants, enhancing their survival and functionality – a phenomenon observed in bioremediation processes (Adams, Fufeyin, Okoro, Ehinomen, et al., 2015). Another form of mutualism is observed in nitrogen-fixing bacteria in legume roots, where bacteria convert atmospheric nitrogen into forms usable by plants (MacLean, Finan, & Sadowsky, 2007).

On the other hand, in predator-prey dynamics, a classic loss-win interaction, certain fungi are known to predate on nematodes, thus controlling their population (Vidal-Diez de Ulzurrun & Hsueh, 2018). Bacteriophages preying on specific bacterial populations is another illustration of this interaction type. Conversely, amensalism is exemplified when fungal species release antibiotics, inhibiting nearby bacterial growth without any direct benefit to themselves (Adrio & Demain, 2003). Commensal relationships are exemplified in the human gut, where certain bacteria benefit from the environment without affecting the host significantly (Kelly, Conway, & Aminov, 2005). In contrast, competition, characterized as a loss-loss interaction, is prominently observed in various ecological settings. For instance, in the microbial

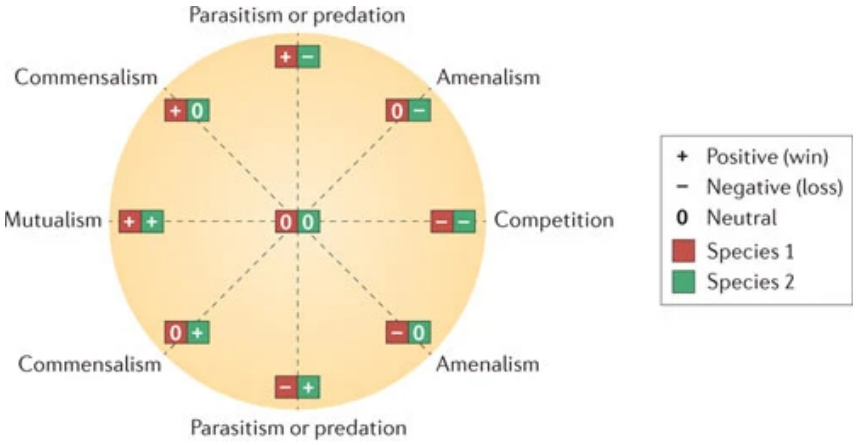


Figure 2.1: Summary of ecological interactions.

communities of soil, different fungal species often compete for essential resources like organic matter, leading to a dynamic where one species may dominate or suppress the other (Boddy & Hiscox, 2016). This competitive interaction aligns with the principle of competitive exclusion. A notable example in marine microbiology is the competition between different species of coral-associated bacteria, where limited resources on coral surfaces lead to the exclusion of less competitive species (Barott & Rohwer, 2012).

Interactions within ecological networks can yield various effects on the species involved, ranging from beneficial to detrimental or neutral. This spectrum of interactions leads to a diverse array of relationship types among species. An adapted model from Lidicker's original conceptualization (Lidicker Jr, 1979) uses a wheel display to categorize all possible pairwise interactions in a microbial community. An illustrative figure, as shown in (Faust & Raes, 2012), is depicted in Figure 2.1. In this model, each species in an interaction can experience one of three outcomes: positive (+), signifying a beneficial impact; negative (–), indicating a detrimental effect; or neutral (0), where the interaction has no significant effect on the species.

In exploring the intricate dynamics of bacterial communities, network theory emerges as a pivotal tool. In molecular biology, graph theory is employed to map the complex interactions within these communities (Koutrouli, Karatzas, Paez-Espino, & Pavlopoulos, 2020). Nodes in these graphs represent biological entities, while edges signify the nature of their interactions. These edges can be either directed, indicat-

ing asymmetrical relationships from one entity to another, or undirected, representing symmetrical links. Furthermore, edges can be weighted to reflect the intensity or strength of the connections. Essentially, these biological networks encapsulate a wide array of interactions among various elements involved in physiological processes, such as proteins, genes, or biomolecules.

Table 2.1 below illustrates various types of biological networks, their corresponding nodes, and the significance of their edges:

Table 2.1: Examples of biological networks

Biological Network Type	**Nodes Represent**	**Interpretation of Edges**
Gene Co-expression Networks	Gene sets	Levels of gene co-expression
Gene Regulatory Networks	Gene regulators and target genes	Patterns of gene regulation
Metabolic Pathway Networks	Metabolites	Interactions among microbes, environment, and hosts
Microbial Ecological Networks	Microbial taxa	Interactions in ecology and physiology
Protein Interaction Networks	Proteins	Similarities in protein structures

In the intricate framework of a microbial community, each constituent member is represented as a node within a network, and the interactions between these members are depicted as edges. These connections are discerned based on the abundance of different taxa, employing a variety of reverse engineering methods that draw upon principles from network theory (Butte & Kohane, 1999)(Herrero, D'ıaz-Uriarte, & Dopazo, 2003) (Sch"afer & Strimmer, 2005) (Basso et al., 2005). Notably, microbial networks may also incorporate nodes corresponding to key ecological or physiological factors, which exhibit notable correlations with the abundance levels of microorganisms. In the following section, we will delve into the current state of literature regarding the reconstruction of microbial networks. This part of the review will highlight seminal works and key developments in this field, providing a comprehensive overview of foundational research on the topic.

2.3 Microbial Network Reconstruction

2.3.1 Correlational Analysis Approach

In the domain of microbial network reconstruction, the elucidation of interactions among various taxa presents a complex analytical challenge. Correlation analysis, a widely employed method in this context, encounters intricacies due to the inherent compositional nature of microbial abundance data. In conventional approaches, direct application of correlation techniques to such datasets can inadvertently lead to spurious correlations or misrepresentative connections. To address this issue, it becomes crucial to implement appropriate data transformations on the count matrix prior to the application of correlation analysis. These preparatory steps are designed to align the correlation analysis more accurately with the true Euclidean relationships among the taxa. Methods falling under this category, known as correlational analysis approaches, generally culminate in the construction of an undirected network, representing the inferred inter-taxa relationships. This process highlights the importance of methodological precision in ensuring the reliability and accuracy of inferred microbial networks.

Recognizing the limitations of traditional correlation approaches, researchers have developed specialized methods tailored for compositional data, commonly found in microbiome studies. These advanced techniques aim to more accurately capture the dependencies among microbial species by accounting for the relative proportions inherent in the data. This advancement is crucial for constructing reliable microbiome correlation networks, which are pivotal in unraveling the complex interactions within microbial communities.

SparCC, short for Sparse Correlations for Compositional Data, stands as one of the pioneering methods addressing the compositional nature of microbial abundance data (Friedman & Alm, 2012). Developed by Friedman and Alm in 2012, SparCC is grounded in the concept of transforming compositional data through log-transformation to analyze absolute abundances. This method is particularly notable for its consideration of the compositionality inherent in microbial data, drawing

upon Aitchison's theory which utilizes the variance of log-ratios to quantify dependencies between compositional variables.

At its core, SparCC approximates the linear Pearson correlations between log-transformed components. This approximation is based on two critical assumptions: the presence of a large variety of operational taxonomic units (OTUs) and the sparsity of the actual correlation network. SparCC employs both a basic and an iterative procedure for inferring these correlations, with the iterative approach focusing on identifying pairs of components with high correlations at each step, maintaining network sparsity until a stable identification is achieved.

While SparCC offers a novel approach to mitigating compositional bias by focusing on log-transformed absolute abundances and excluding relative information, it is not without limitations. Firstly, since absolute abundances are unobservable, SparCC's effectiveness hinges on the quality of component count data. Additionally, the covariance matrix inferred by SparCC may not always be positive definite, and the correlation coefficients it yields can fall outside the standard range of [-1, 1]. Another concern is its computational efficiency, especially when dealing with datasets with high species diversity. It is worth noting, however, that Friedman and Alm found SparCC to be robust in simulations, even when the sparsity assumption was not entirely met. They also established a link between alpha diversity and compositional bias, suggesting that SparCC is best suited for datasets with lower diversity levels.

CCLasso, introduced by (Fang, Huang, Zhao, & Deng, 2015) in 2015, represents a significant advancement in analyzing compositional data, particularly in microbial network reconstruction. This method adeptly addresses the compositional nature of microbial abundance data by estimating correlations between log-transformed absolute abundances. It employs a weighted least-square regression with an L1 penalty on log-transformed relative abundance, ensuring that the estimated correlation matrix of latent variables is both positive definite and within the range of [-1, 1]. This approach contrasts with methods like SparCC, which do not inherently guarantee a positive semi-definite covariance matrix. CCLasso's alternating direction

algorithm, derived from the augmented Lagrangian method, effectively navigates the optimization challenges in this context. Despite its strengths, CCLasso shares some limitations with earlier methods, such as its reliance on linear relationships and the need for accurate component fraction estimation. Furthermore, its approach to handling excessive zeros in the data, through the use of pseudo-counts, and its performance in cases where taxa significantly outnumber samples, remain areas for further exploration and refinement.

REBACCA, introduced by (Ban, An, & Jiang, 2015) in 2015, presents an innovative approach for discerning significant co-occurrence patterns in microbial data. Distinct from SparCC, REBACCA constructs a linear system akin to log-ratio transformations and employs an L1-norm shrinkage method, focusing on sparse solutions under the assumption of sparsity. This methodology allows REBACCA to efficiently estimate pairwise relationships of log-transformed absolute abundances, thus circumventing compositional bias. The key differentiation from CCLasso lies in its parameter tuning approach, where REBACCA utilizes a stability resampling method. This method regulates the number of selected variables by assessing their selection probability, based on independent Lasso applications on two datasets derived from randomly splitting the original data. Despite its strengths, REBACCA faces challenges with scalability, particularly as the number of variables increases, leading to high computational requirements. Additionally, it lacks theoretical performance guarantees and is susceptible to inaccuracies in scenarios with small sample sizes or high noise levels.

2.3.2 Graphical Model Based Approach

In the realm of microbial network reconstruction, the graphical model-based approach represents a significant shift from simple correlation analysis to a more nuanced understanding of microbial interactions. This approach focuses on elucidating the complex web of relationships within microbial communities and try to disentangle the direct relationships between taxas.

At the heart of this methodology is the distinction between direct and indirect inter-

actions among taxa. Indirect interactions are those where the correlation between two taxa is influenced by external factors like environmental conditions or the presence of a third species. In contrast, direct interactions imply a more immediate and intrinsic connection between two taxa, devoid of such confounding elements. Given the more fundamental nature of these direct interactions, there has been a growing interest in methods that model the conditional dependencies among taxa, considering all other taxa in the network.

To achieve this, many researchers have turned to probabilistic graphical models. In these models, each node represents a taxon, and each edge signifies an undirected, conditional dependency between two taxa. The utility of these models is particularly pronounced when data follow a distribution within the exponential family, allowing for clear statistical interpretations of the relationships represented in the model. A prime example is the use of the multivariate Gaussian distribution, where the elements of the inverse covariance matrix (also known as the precision matrix) are key. Specifically, a zero element in this matrix indicates conditional independence between the corresponding variables, given all others.

However, the task of accurately estimating the precision matrix is not without its challenges, especially when dealing with real microbiome data. These data sets are often high-dimensional and sparse, leading to degenerate covariance matrices. Furthermore, the reliance on observed compositional data adds another layer of complexity to estimating the precision matrix in the context of absolute abundances. This approach, while complex, offers a more detailed and interconnected view of microbial networks, advancing our understanding beyond simple pairwise correlations.

SPIEC-EASI, developed by (Kurtz et al., 2015) in 2015, is a significant method in graphical model-based microbial network reconstruction, focusing on conditional independence (Kurtz et al., 2015). This method employs centered log-ratio (clr) transformation to approximate the covariance of log-ratio transformed absolute abundances. It is particularly effective when the number of taxa surpasses the number of subjects, adapting the Γ estimator as a proxy for the true covariance matrix,

Γ'.

The SPIEC-EASI method addresses the challenges of data sparsity and network identification by hypothesizing graph sparsity and employs a penalized maximum likelihood approach for estimating true abundances. The sparsity of the network's edges is regulated through the StARS selection algorithm.

While SPIEC-EASI is adept at integrating prior knowledge about data or network structure, it faces several limitations. It relies heavily on the condition number of the inverse covariance matrix and tends to overlook uncertainties in estimates. Additionally, its handling of excessive zeros in clr-transformed data is problematic, often not adhering to Gaussian or sub-Gaussian graphical models. The consistency of the estimated covariance matrix under SPIEC-EASI is also a concern, not always being guaranteed.

gCoda, developed by (Fang, Huang, Zhao, & Deng, 2017) in 2017, represents an innovative approach in microbial network reconstruction, primarily focusing on estimating sparse inverse covariance through penalized maximum likelihood. The method operates under two fundamental assumptions: the logarithm of absolute abundances adheres to a multivariate log-normal distribution, and the network exhibits sparse edge density. This framework allows gCoda to model the generation of relative abundances via a logistic normal distribution and subsequently estimate conditional dependence structures indicative of direct interactions within microbial communities. The method employs an L1 penalty to enforce sparsity in these estimates.

A notable feature of gCoda is its computational efficiency and stability, distinguishing it from other methods like SPIEC-EASI. It stands out as one of the first methodologies to derive the likelihood of compositional data for Gaussian graphical models. However, a critical limitation of gCoda lies in the nonconvex nature of its likelihood function, which raises challenges in optimization. Additionally, the consistency of its estimator is not always assured, representing a potential area for future refinement and improvement.

MPLasso, introduced (Lo & Marculescu, 2017) in 2017, offers a novel approach in microbial network reconstruction by integrating experimentally verified biological knowledge as a prior in the graphical Lasso framework. This method stands out for its use of automated text mining to extract microbial co-occurrence patterns and associations from a vast array of scientific literature. This prior information is crucial in narrowing down the search space, enhancing the reliability of interaction inference. In MPLasso, regularization is executed through an L1 penalty applied to the elements of the precision matrix, incorporating the co-occurrence matrix prior. Although this prior information doesn't dominate the results, it significantly constrains the search space.

A unique feature of MPLasso is its method for extracting prior knowledge. For each pair of taxa, it queries the PubMed database to gather the number of papers mentioning each taxon individually, both taxa, or neither. This data forms a contingency matrix, upon which Fisher's exact test with Bonferroni correction is applied to determine prior probabilities of associations. The Bayesian Information Criterion (BIC) in the Gaussian Graphical models context is then employed for optimal penalty parameter selection, striking a balance between likelihood maximization and edge number relative to sample size. While MPLasso's methodology is innovative, it is important to note that it has not been tested on dynamic models of microbial communities, indicating potential areas for future development and application.

MDiNE, introduced by (McGregor, Labbe, & Greenwood, 2020) in 2020, presents a Bayesian hierarchical model tailored for estimating co-occurrence networks within microbial communities, differentiated by a binary variable. This innovative approach models individual taxa counts through a logistic normal multinomial model, linking these counts to latent Gaussian random variables. The essence of MDiNE lies in inferring the co-occurrence network among taxa from a sparse precision matrix, which integrates all latent items. A distinctive feature of MDiNE is its capability to model and distinguish between networks for different groups, effectively mitigating the issue of reduced sample sizes in high taxa count scenarios.

The Hamiltonian Monte Carlo algorithm is employed in MDiNE for model estimation, facilitating interval estimations for both network parameters and the disparities between networks. However, MDiNE is not without limitations. It is prone to underestimating true associations in cases with small sample sizes and lacks a mechanism for zero-inflation, which can be problematic when dealing with data abundant in zeros. Moreover, the Markov Chain Monte Carlo (MCMC) sampling process in MDiNE can be considerably time-intensive and computationally burdensome, potentially impacting its applicability in larger datasets. Despite these constraints, MDiNE's unique approach to network estimation offers insightful perspectives into the variances among microbial communities categorized by binary variables.

Developed by (Baksi, Kuntal, & Mande, 2018), the TIME (Tool for Insights into Microbial Ecology) framework stands as a pivotal advancement in analyzing microbial ecology, particularly in understanding causal relationships within microbial communities. TIME leverages the Granger causality concept, a statistical tool initially conceived for economics and now extensively applied in neuroscience and microbiome studies. Granger causality operates on the premise that if a pair of time series, say X and Y, exhibit a scenario where X and Y's combined historical data better predict Y's current value than Y's history alone, then X is said to "Granger cause" Y. TIME incorporates this approach into a multifaceted toolkit that caters to various analytical needs in microbial ecology studies.

The framework offers three distinct methods for causal relationship analysis, enhancing its versatility. The first method relies on pairwise Granger causality inference for each taxa pair, contrasting two regression models. The second approach harnesses a Granger Lasso with regularization, adeptly managing high parameter volumes in multivariate settings. Lastly, the ensemble method of TIME merges insights from both preceding methods, pinpointing node pairs where causality is consistently indicated. TIME's ensemble approach thus provides a robust platform for identifying and visualizing causal relationships in complex microbial ecosystems, contributing significantly to the comprehension of dynamic interactions within these communities.

2.4 Causal Inference

The quest to understand causal relations and the principles that govern them has been a cornerstone of scientific endeavor since the advent of modern science in the seventeenth century. This section delves into the pivotal role of causal inference in scientific research. It concentrates on the evolution and utilization of directed graphical causal models (DGCMs) and diverse methodologies for deducing accurate causal representations from a variety of data sources. Our exploration includes the historical context in which experimentation and observational analysis were intertwined in discovering causal relationships. In contemporary science, particularly in molecular biology and systems biology, DGCMs have become prevalent in omics studies and neural connectivity research.

Traditionally, causal relations have been explored through interventions or randomized experiments. However, these methods often face practical limitations such as high costs and time constraints. As a result, there has been an increased focus on the analysis of observational data for causal discovery. This non-interventionist approach has advanced significantly in fields like genomics, ecology, epidemiology, and clinical medicine. We aim to explore both constraint-based and score-based methods for causal discovery, including algorithms such as PC and Fast Causal Inference (FCI). These methods are pivotal in deducing causal structures based on conditional independence relationships and are applicable across various data types and causal relations, provided that reliable conditional independence testing methods are available. Despite their broad utility, these methods typically result in equivalence classes rather than exhaustive causal information.

Moreover, this section will delve into score-based algorithms like the Greedy Equivalence Search (GES), which seek to identify causal structures by optimizing specific score functions. A particular emphasis will be placed on the causal inference literature that addresses the presence of cycles, a critical aspect relevant to our work in Chapter 4. Additionally, we will examine the integration of interventional experimental data in the field of causal inference, an area that closely aligns with the

methodologies and findings discussed in Chapter 5.

Directed Graphical Causal Models (DGCMs) serve as sophisticated tools in the domain of causal inference, offering a nuanced representation of causal relationships within a set of variables. These models encompass several key elements: (1) a collection of variables treated as stochastic entities, (2) directed links between variable pairs, suggesting causal associations under specific conditions, and (3) a comprehensive joint probability distribution encompassing all variables. These components facilitate the portrayal of variables as interconnected stochastic processes, which can be time-dependent and inclusive of unobserved variables. DGCMs accommodate various data types, including categorical, ordinal, and continuous, while addressing complexities like measurement errors and selection biases. Importantly, these models also incorporate unobserved individual variations, often labeled as 'noise' or 'disturbances.'

A prominent subclass of DGCMs, often represented as Structural Equation Models (SEMs) or Functional Causal Models (FCMs), posits that each variable's value is deterministically linked to its immediate causative factors in the model, alongside unmeasured disturbances. The functions connecting variables to their causative factors are versatile, with linear representations being the most prevalent. This class of models transcends traditional statistical models, including regression and factor models, and extends time series and latent class models, among others. Unlike systems governed by differential equations that necessitate initial or boundary conditions, DGCMs offer a more general framework for causal representation.

It is crucial to distinguish that not all directed graphical models inherently possess causal interpretations. However, the essence of a DGCM lies in its ability to imply causal associations. This is achieved through the model's second element, which asserts that an edge between two variables indicates causation under certain conditions of variable manipulation. Thus, if an edge exists from X_i to X_j in the directed graph, it implies that X_i is a direct cause of X_j. In essence, interventions on X_i would directly influence the distribution or value of X_j. This conceptualization aligns closely with Pearl's definition of a causal Bayesian network (Pearl

et al., 2000). The integration of directed graphs with joint probability distributions in DGCMs is not arbitrary but is governed by specific constraints. In acyclic directed graphs where there are no directed cycles, these constraints are manifested through the principle of *d-separation*, which must correspond to conditional independence within the probability distribution. The acyclicity is usually a convenient assumption in many of the literatures.

The concept of d-separation in directed acyclic graphs (DAGs) offers an insightful perspective on conditional independence, paralleling the more conceptually graspable but less pragmatically applicable Local Markov Condition. This condition posits that within a DAG, any given variable, X, maintains independence from its non-descendant variables, contingent on the variable's parent variables. These parent variables are defined as those directly influencing X through incoming edges. This notion echoes a well-established axiom in experimental research: stabilizing the influencers of a target variable X effectively mitigates the impact of indirect causative factors.

Graphs that share d-separation characteristics are often referred to as "Markov equivalent," indicating their implication of identical conditional independence relationships. A group of all DAGs sharing this Markov equivalence forms what is known as a Markov Equivalence Class (MEC). In linear systems, the principle of d-separation has been extended to encompass directed graphs that include cyclic paths, as expounded by (Spirtes, 2013). Furthermore, for systems represented by a graph that is a marginalized version of a larger system, a related concept known as m-separation emerges (Zhang, 2008). The Causal Markov Assumption applies when this Markov condition is presumed to hold true for both the causal graph and its corresponding population distribution, thus framing a foundational assumption in causal inference.

2.4.1 Constraint Based Approach

Constraint-based methods have emerged as a cornerstone in the domain of causal discovery, offering a nuanced approach to deciphering complex causal structures from observational data. These methods are predicated on the utilization of con-

ditional independence tests to uncover underlying causal relationships among variables within a given dataset.

The essence of constraint-based causal discovery lies in its ability to identify patterns of conditional independence, thereby inferring potential causal connections. This approach does not rely on pre-established models of the data-generating process, but instead, it discerns causal structures directly from the statistical properties inherent in the data.

The PC algorithm, named after its developers Peter Spirtes and Clark Glymour (Pearl et al., 2000), stands as a pivotal method in the landscape of constraint-based causal discovery. This algorithm serves as a fundamental tool for researchers aiming to elucidate causal structures from observational data, particularly in contexts where controlled experiments are impractical or impossible.

At its core, the PC algorithm is built upon the principles of conditional independence and the Causal Markov Condition. It assumes a framework wherein each variable within a dataset is conditionally independent of all other variables, given its direct causes. This principle is operationalized through an iterative process of edge deletion in a fully connected undirected graph, guided by conditional independence tests. The algorithm begins with a complete graph, where every pair of variables is connected, signifying potential causal links. It then systematically tests for conditional independence between variable pairs, given subsets of other variables. Edges between variables that are conditionally independent are removed, progressively simplifying the graph.

The PC algorithm can be divided into two distinct stages:

1. *Skeleton Identification*: In this initial phase, the algorithm identifies the 'skeleton' of the causal graph, which is the underlying undirected graph that captures the conditional independencies among the variables. The algorithm employs a series of conditional independence tests to iteratively remove edges from the initial fully connected graph, thereby refining the skeleton.

2. *Edge Orientation*: Once the skeleton is established, the algorithm proceeds

to the second stage, which involves orienting the edges to infer the directionality of the causal relationships. This process utilizes v-structures and other orientation rules to determine the direction of causal influence, thereby transforming the undirected skeleton into a Directed Acyclic Graph (DAG), to the extent possible.

The strength of the PC algorithm lies in its broad applicability across diverse types of data and its ability to operate without prior assumptions about the functional form of causal relationships. It is particularly effective in high-dimensional settings, where the number of variables is large relative to the sample size. However, the PC algorithm has its limitations. It assumes causal sufficiency, meaning it operates under the assumption that there are no unmeasured confounders. Latent confounders are unobserved variables that influence both the cause and effect in a study, potentially leading to spurious correlations and misleading conclusions about causal relationships. Additionally, while the algorithm is adept at revealing the structure of causal relationships, it does not inherently provide information on the strength or nature of these relationships. Furthermore, the accuracy of the PC algorithm is contingent on the reliability of the conditional independence tests employed, making it sensitive to the choice of statistical tests and the sample size of the data.

The Fast Causal Inference (FCI) algorithm, developed as an extension to the PC algorithm, represents a significant advancement in the field of causal discovery. It is particularly notable for its ability to handle the complexities introduced by latent variables and selection bias, which are common challenges in observational studies (Zhang, 2008).

The FCI algorithm was conceptualized to address a crucial limitation of the PC algorithm, its assumption of causal sufficiency, which presumes the absence of unmeasured confounders. In many real-world scenarios, this assumption is violated, leading to potential inaccuracies in causal inference. FCI extends the capabilities of causal discovery to scenarios where latent variables or hidden confounders are present, thus providing a more robust and realistic framework for uncovering causal structures from observational data.

FCI operates by initially following a procedure similar to the PC algorithm, constructing an undirected graph through conditional independence tests. However, it incorporates additional steps to detect and account for the presence of latent variables and selection bias. The algorithm uses a modified version of d-separation, called m-separation, to test for conditional independencies in the presence of latent variables. This allows FCI to infer a class of models known as Partial Ancestral Graphs (PAGs), which represent the Markov equivalence class of Directed Acyclic Graphs (DAGs) that could have generated the observed data.

While FCI offers significant advantages in handling complex data structures, it is not without challenges. The algorithm can be computationally intensive, especially for large datasets with many variables. Furthermore, the accuracy of FCI is contingent on the reliability of the conditional independence tests used, and the presence of a large number of latent variables can complicate the interpretation of the inferred PAGs.

In summary, the FCI algorithm marks a significant step forward in the field of causal discovery. It expands the boundaries of what can be inferred from observational data, accommodating the realities of unmeasured confounders and selection bias. Its contributions have been instrumental in enhancing the robustness and applicability of causal inference methods in various scientific and research contexts.

The development of the Fast Causal Inference (FCI) algorithm has spurred the creation of its variants, notably the Greedy Fast Causal Inference (GFCI) and Really Fast Causal Inference (RFCI) algorithms. These adaptations aim to address specific challenges and limitations of the original FCI, particularly concerning computational efficiency and scalability.

GFCI represents an innovative blend of the FCI algorithm with elements of the Greedy Equivalence Search (GES) method which is a famous score-based causal inference method and will be reviewed in the next section. This hybridization aims to enhance the efficiency and accuracy of causal discovery in large datasets.

The GFCI algorithm adopts a hybrid approach, effectively blending the conditional

independence tests of FCI with the score-based optimization from GES. This fusion results in an algorithm that leverages the strengths of both FCI and GES, providing a more nuanced and comprehensive approach to causal inference.

Efficiency and scalability are standout features of GFCI, especially beneficial when dealing with large datasets. The integration of GES's score-based search methodology enhances GFCI's computational efficiency, allowing it to process complex datasets more effectively.

GFCI also adeptly handles scenarios involving latent variables and selection bias, similar to its counterpart FCI. This capability is crucial for robust causal inferences, particularly in research environments where these complex factors play a significant role. GFCI's ability to address these challenges underscores its utility as a comprehensive tool for causal analysis.

RFCI, developed as a streamlined variant of FCI, focuses on significantly reducing computational complexity. This reduction is achieved while still maintaining the capability to discern causal structures, even in the context of latent variables and selection bias.

The primary advantage of RFCI is its simplified approach to conditional independence tests. By modifying some of the tests utilized in FCI, RFCI enables faster computational processes. This feature is particularly advantageous in extensive studies where efficient data processing is key.

Moreover, the design of RFCI makes it aptly suited for large-scale studies. The algorithm's simplifications cater to scenarios where computational resources are constrained, thereby facilitating effective causal analysis in large datasets.

However, there is a trade-off to be mindful of with RFCI. While it excels in speed and efficiency, it might compromise slightly on accuracy compared to FCI, particularly in datasets with complex causal structures. This factor necessitates a careful assessment of the benefits of increased speed against potential decreases in accuracy, especially in detailed causal analyses.

Both GFCI and RFCI are significant contributions to the field of causal discovery,

each with its unique strengths and suitable applications. GFCI is often preferred in scenarios where both accuracy and computational efficiency are paramount, especially in large but not overly complex datasets. RFCI, on the other hand, is advantageous in extremely large datasets where computational speed is a critical factor.

The efficacy of constraint-based methods is grounded in their theoretical underpinnings, with a robust framework that accommodates a wide range of data distributions and causal scenarios. However, these methods output Markov Equivalence Classes (MECs) rather than explicit causal diagrams, implying that they reveal sets of causal structures consistent with the observed conditional independencies, rather than pinpointing a single definitive causal model. This characteristic, while indicative of the methods' non-commital nature, also underscores the need for careful interpretation and potential further analysis to resolve ambiguities inherent in the MECs.

In summary, constraint-based methods of causal discovery represent a significant advancement in the field of causal inference, providing a powerful toolkit for scientists and researchers to unravel complex causal relationships from observational data, while navigating the challenges posed by latent variables and the inherent limitations of non-experimental data.

2.4.2 Score Based Approach

Score-based methods represent a pivotal category in the realm of causal inference, distinguished by their approach of using optimization techniques to identify the most likely causal structure from a given dataset. Unlike constraint-based methods, which primarily rely on conditional independence tests, score-based methods evaluate different causal models based on a specific scoring criterion, seeking the model that best fits the observed data.

The cornerstone of score-based methods lies in the scoring criteria used to evaluate the fit of a causal model. Commonly used criteria include the Bayesian Information Criterion (BIC), Akaike Information Criterion (AIC), and likelihood-based scores. Each criterion offers a different balance between model complexity and goodness

of fit, allowing for nuanced model selection.

Finding the optimal causal model requires navigating through the vast space of possible models. Score-based methods employ various search algorithms like hill-climbing, tabu search, or more sophisticated approaches like Greedy Equivalence Search (GES) to efficiently traverse this space.

GES (Chickering, 2002) is a prominent score-based algorithm that operates in two phases: the forward phase, which adds edges to increase the score, and the backward phase, which removes edges to improve the score. It directly searches over the space of Markov equivalence classes rather than individual graphs.

GES is particularly effective in scenarios where the number of variables is large, and the underlying causal structure is believed to be sparse. It has been successfully applied in various fields, including genomics and systems biology.

In the Bayesian paradigm, score-based methods evaluate causal models based on their posterior probabilities given the data. These methods incorporate prior knowledge about the likelihood of certain causal structures, providing a comprehensive framework for causal discovery.

Bayesian score-based methods are adept at handling uncertainty and integrating domain knowledge, making them powerful tools in complex causal discovery tasks.

While score-based methods have made significant strides in causal inference, they face challenges such as computational complexity, especially in high-dimensional data. Future research is geared towards developing more efficient algorithms, integrating machine learning techniques, and enhancing the robustness of these methods in the face of data limitations and complex causal mechanisms.

2.4.3 Cyclic Causal Inference

Cyclic causal inference is an advanced segment of causal analysis that challenges traditional assumptions of acyclicity in causal models. Recognizing that many real-world systems exhibit feedback loops and interdependencies, cyclic causal inference aims to elucidate the nature and dynamics of such complex causal structures.

In numerous domains, from biological systems to economic models, variables often influence each other in a circular manner, creating cycles. These cyclic interactions can lead to dynamic behaviors that are not captured by conventional acyclic models, making the study of cyclic causal structures essential for a more accurate understanding of these systems.

Cyclic causal inference presents both challenges and opportunities. While it complicates the causal discovery process, it also offers richer insights into the nature of causal relationships, especially in systems where feedback loops play a crucial role.

In the advanced landscape of causal inference, the adaptation of traditional causal models to include cyclic dependencies represents a pivotal evolution. Structural Causal Models (SCMs), foundational to this field, are undergoing transformative modifications to integrate these complex cyclic structures. This adaptation extends beyond mere inclusion of cycles, necessitating a reevaluation and reformulation of core concepts like d-separation, traditionally tailored for acyclic networks.

In the seminal work by (Bongers, Forr'e, Peters, & Mooij, 2021), a comprehensive framework for statistical causal modeling with cyclic SCMs is established. This groundbreaking research not only incorporates cycles and latent variables into SCMs but also embraces non-linear functional relationships between variables. This integration marks a significant departure from conventional causal models, bringing a nuanced understanding of causality in intricate systems characterized by feedback loops and hidden factors.

Recognizing the unique nature of cyclic structures, researchers have developed specialized algorithms and techniques to infer causal relationships in cyclic causal graphs. These methods navigate the complexities introduced by feedback loops and interdependencies.

A notable progression in causal modeling, particularly in handling cyclic dependencies, was introduced by in 2010 (Itani, Ohannessian, Sachs, Nolan, & Dahleh, 2010). This development addressed the limitations of traditional Directed Acyclic Graphs (DAGs) in scenarios where causal structures inherently contain cycles. The

method expanded the scope of Bayesian Network structure learning to include cyclic dependencies, leading to the creation of Generalized Bayesian Networks (GBNs) and Causal Generalized Bayesian Networks (CGBNs). These models are equipped to handle both cyclic dependencies and interventions, thus enhancing the modeling accuracy in complex systems. A unique aspect of this approach is the use of 'cycle breakers', which temporarily convert cyclic structures into acyclic forms, allowing for the application of conventional structure learning techniques. The practical application of this method in real-world scenarios has demonstrated its efficacy in unraveling complex causal relationships, establishing it as a significant breakthrough in the domain of causal inference and opening avenues for future research in various scientific fields.

The exploration of cyclic structures in causal inference was notably advanced by the seminal work of Eberhardt and Hyttinen in 2012 (Hyttinen, Eberhardt, & Hoyer, 2012). Their study delved into the intricate challenge of identifying and learning linear cyclic causal models, particularly in scenarios involving latent variables. A pivotal aspect of their research was the formulation and rigorous proof of necessary and sufficient conditions that determine the identifiability of linear cyclic causal models. Additionally, they introduced a complete search algorithm that further delineates the landscape of causal inference in complex systems.

In (Hyttinen, Hoyer, Eberhardt, & Jarvisalo, 2013), the introduction of a novel approach that utilizes d-separation constraints from overlapping observational or experimental datasets marks a significant stride forward. This method, distinct in its ability to handle both directed cycles (feedback loops) and latent variables, employs a logical representation of causal pathways. This representation allows for the incorporation of extensive background knowledge, while inference is conducted through a Boolean satisfiability (SAT) solver. Notably, this procedure is comprehensive, thoroughly utilizing available data to ascertain the presence or absence of specific edges in a causal model, resorting to 'unknown' only when the data is insufficient. This general approach encapsulates and extends beyond the capabilities of many existing constraint-based causal discovery algorithms, which are often limited by more

restrictive assumptions. Simulations conducted to assess this method have demonstrated its robustness and scalability, offering insights into how these assumptions influence the process of causal discovery in various contexts.

Recently, (Mooij & Claassen, 2020) has demonstrated that the Fast Causal Inference (FCI) algorithm, conventionally applied in acyclic scenarios, can accurately analyze observational data from systems involving feedback. This is underpinned by evidence that for data originating from a simple and σ faithful Structural Causal Model (SCM), FCI is both sound and complete. It effectively identifies the presence and absence of causal relations, direct causal connections, confounders, and specific (non-)cycles in the SCM's causal graph. This surprising versatility extends to other constraint-based causal discovery algorithms that integrate various forms of background knowledge, encompassing causally sufficient settings, like the PC algorithm, and Joint Causal Inference settings, exemplified by the FCI-JCI algorithm. These findings mark a significant advancement in the field of causal discovery, particularly in understanding and applying FCI and related algorithms in cyclic causal models.

The study of cyclic causal inference is still evolving, with ongoing research focused on refining methodologies, enhancing computational efficiency, and expanding its applicability. Future advancements are anticipated to unlock new possibilities in understanding and manipulating complex systems with cyclic causal dynamics.

2.4.4 Intervention Calculus

Causal inference extends beyond the mere identification of causal relationships to understanding the effects of potential interventions. This crucial aspect of causal analysis is encapsulated in the intervention calculus, a systematic method for predicting the outcomes of hypothetical or real interventions in a causal system. The intervention calculus provides a framework for not only determining how changes in one variable influence another but also quantifying the extent of these effects.

Central to this calculus is the concept of *do*-calculus, pioneered by (Pearl et al., 2000). This approach allows for the estimation of causal effects from observational

data under certain assumptions, enabling researchers to simulate interventions and predict their outcomes without the need for costly or impractical experiments. The intervention calculus is particularly valuable in fields where controlled experimentation is challenging, offering a robust tool for extracting actionable insights from complex causal systems.

In this section, we delve into the theoretical underpinnings of the intervention calculus and its applications in various domains. We explore how this methodology has revolutionized the understanding of causal mechanisms, providing a bridge between observational data and the realm of causal predictions. The discussion will cover key principles, methodologies, and breakthroughs in the field, shedding light on how intervention calculus has become an indispensable tool in causal inference research.

In the domain of acyclic causal models, (Eberhardt, Glymour, & Scheines, 2012) established critical boundaries for the minimal number of required experiments. Their study revealed that for identifying causal relationships among $N \geq 2$ variables, only $\log_2(N) + 1$ experiments are generally necessary.

(Hauser & B"uhlmann, 2014) have put forward two innovative strategies for the active learning of causal models using interventional data, catering to distinct learning goals. These methodologies encompass both batch and sequential intervention designs. The batch-intervention design clusters multiple interventions in a single experimental setup, whereas the sequential-intervention design employs a step-wise approach, focusing on progressively orienting undirected edges in the causal chain through successive experiments. This aligns with the active learning paradigm proposed by (He & Geng, 2008), where the goal is to systematically clarify causal relationships by orienting undirected edges via targeted interventions.

(Shanmugam, Kocaoglu, Dimakis, & Vishwanath, 2015) have expanded this field by establishing a lower bound on the number of experiments necessary for effective adaptive methods. This foundational work is predicated on the concept of separating systems and has significant implications for experimental design in causal inference. Despite their efficacy, these approaches predominantly operate under the

assumption of acyclicity, which, while prevalent in causal literature, can oversimplify complex systems with feedback loops or cyclic dependencies.

Finally, the COmbINE method, introduced in (Triantafillou & Tsamardinos, 2015), offers a solution to the challenges posed by statistical errors in multiple interventions over overlapping sets of variables. This method's ability to reconcile conflicts arising from such complexities marks a significant step forward in experimental design for causal inference.

2.5 Network Control: Model-Based and Data-Driven Approaches

The exploration of network control in complex systems, particularly within the domain of microbial communities, constitutes a pivotal area of contemporary research. This section delves into the diverse methodologies and paradigms underpinning network control, segregating them broadly into two categories: model-based and data-driven approaches.

Model-based approaches to network control rely on predefined models that encapsulate the underlying dynamics and interactions within a system. These models, often derived from theoretical frameworks or empirical observations, serve as the basis for designing control strategies. In microbial network control, these strategies are crucial for manipulating community composition or functionality, with potential applications ranging from disease treatment to environmental management.

In contrast, data-driven approaches eschew the reliance on preconceived models, instead leveraging advancements in computational techniques, particularly machine learning and artificial intelligence, to derive control strategies directly from data. This paradigm shift has been catalyzed by the burgeoning volume and complexity of data available in microbial studies, offering new vistas for control methodologies that are more adaptive and responsive to real-time changes within the system.

The intersection of model-based and data-driven approaches represents a fertile

ground for innovation in network control. As such, this section aims to provide a comprehensive overview of both approaches, highlighting their respective strengths, limitations, and potential for synergistic integration in the context of microbial community control.

The interrogation of networks to achieve controllability, which refers to the capability to steer a system from any initial state to any desired final state, is of paramount importance in understanding complex systems. The groundbreaking work by (Liu, Slotine, & Barab'asi, 2011) introduced a framework to determine the controllability of complex directed networks and identify driver nodes necessary for full network control. The driver nodes are pivotal for instigating state transitions within the network, governed by the system dynamics $\dot{x}(t) = Ax(t) + Bu(t)$, where $x(t)$ is the state vector, A is the interaction matrix, and B is the input matrix associated with the control signal $u(t)$. The criterion for controllability hinges on the rank condition of the controllability matrix $C = [B, AB, A^2B, \ldots, A^{N-1}B]$.

Within the realm of microbial communities, (Angulo, Moog, & Liu, 2019) put forth a strategy to ascertain the minimum set of driver species for structural accessibility by applying a maximum matching algorithm to the microbial interaction network. This innovation broadens the notion of structural controllability to encompass nonlinear dynamics, effectively merging the species interactions and controlled inputs within a comprehensive network graph.

Further advancing the control methodology, (Zañudo, Yang, & Albert, 2017) presented feedback vertex set control (FC), a technique that utilizes a subset of nodes intersecting every feedback loop, known as the feedback vertex set (FVS), to navigate the network's dynamics toward its inherent attractors. This method is delineated by the dynamics $\frac{d}{dt}x_i(t) = F_i(X_i, X_{I_i}, t)$, which underscore the potential of FC to attain any naturally occurring end state within the network.

In the study of complex network dynamics, a critical consideration is the control energy required to guide the system towards a desired state. (Lindmark & Altafini, 2018) address this concern, revealing that the proximity of a network's mode eigenvalues to the complex plane's imaginary axis significantly influences the neces-

sary control energy. Specifically, modes with eigenvalues nearer to the imaginary axis demand less energy for control. This inverse relationship is especially notable in networks with purely imaginary eigenvalues, like coupled harmonic oscillators, where the authors present constructive algorithms to identify driver nodes that ensure minimal control energy. Beyond these specific cases, they posit a broader heuristic for any directed network: control costs are reduced when efforts are concentrated on nodes exhibiting a high weighted outdegree to indegree ratio. This heuristic serves as a guiding principle for economical control across various networked systems, offering insights into energy-efficient control strategies.

Data-driven control strategies have gained momentum in the domain of dynamical systems, catalyzed by the burgeoning advancements in machine learning and artificial intelligence sectors. Traditionally, control systems have relied on an indirect approach wherein a system model is first inferred from data, followed by the computation of control parameters using the estimated model. This approach, however, can be hindered by the inaccuracies and extensive computational requirements of the identification algorithms. To circumvent these challenges, a spectrum of direct data-driven methods has been introduced, which forgo the preliminary identification step. This includes techniques such as (model-free) reinforcement learning, iterative learning control, and adaptive control.

In the domain of dynamic systems control, reinforcement learning (RL) presents a transformative approach, diverging from the classical model-based strategies. RL is predicated on the iterative process of policy optimization through environmental interaction, aiming to maximize the accumulation of rewards (Sutton & Barto, 2018). This method demonstrates its prowess in complex or poorly-defined environments where traditional control methods may flounder.

A salient feature of RL-based control is its inherent adaptability and proficiency in real-time learning and decision-making. It eschews the dependence on precise system models, instead dynamically refining its policy in response to environmental feedback, which culminates in enhanced performance over iterative cycles.

The integration of deep learning into RL, known as deep reinforcement learning,

has further expanded its capabilities, positioning it at the forefront of control system innovation (Li, 2017). Its successful deployment across various sectors, from autonomous vehicles to financial trading, underscores its potential as an efficacious control mechanism in unpredictable and intricate systems.

The burgeoning field of biomanufacturing has turned its gaze towards engineered synthetic microbial communities, whose enhanced productivity over monocultures stems from the partitioning of metabolic functions across different species. Despite these communities' benefits in efficiency and metabolic load reduction, their practical application has been impeded by the intricate task of community control. A cutting-edge solution hails from the realm of artificial intelligence, specifically, the use of reinforcement learning (RL) to regulate co-cultures within bioreactors (Treloar, Fedorec, Ingalls, & Barnes, 2020). Simulations have illuminated the efficacy of RL, where a suitably trained agent can proficiently maintain stable community populations. Remarkably, employing a model-free bang-bang control strategy via RL has outperformed traditional continuous proportional-integral controllers under conditions of infrequent sampling. Experiments deploying multiple bioreactors in unison have illustrated the swift policy optimization achievable by RL, often within a mere twenty-four hours. Further, RL has been successfully harnessed to directly amplify the productivity of co-culture bioprocesses, positing it as a formidable technique for microbial community management and control within bioreactors. The adaptability and effectiveness of deep reinforcement learning in these contexts underscore its potential as a transformative tool for the practical application of co-cultures in the field of biomanufacturing.

Deep reinforcement learning (DRL) emerges as a formidable approach in network science for pinpointing 'key players' within complex systems. These crucial nodes, upon activation or removal, drastically sway the network's overall functionality, with implications spanning from immunization strategies and epidemic mitigation to drug discovery and marketing campaigns. The challenge of these NP-hard problems is adeptly navigated by the DRL framework known as FINDER, which, trained on modestly-sized synthetic networks, generalizes to a broad array of practical con-

texts (Fan, Zeng, Sun, & Liu, 2020). Exhibiting a remarkable edge over traditional algorithms, FINDER not only excels in generating high-caliber solutions for complex network issues but also operates at a speed that outpaces existing techniques by several orders of magnitude for large networks. This novel methodology not only highlights the potential of deep learning to decipher complex networks' structural principles but also promises the strategic design of networks fortified against both targeted attacks and random failures.

2.6 Research Gaps

In the concluding section of this chapter, we identify the pivotal research gaps that are in the existing literature which have been targeted in throughout this doctoral study. A primary focus is laid on the comparative analysis of network reconstruction methodologies in the domain of microbial communities. We identify a critical gap: the challenges and ambiguities posed by cycles in microbial networks. Addressing this, we propose a novel approach where microbial interaction dynamics are encapsulated within the framework of causal inference. This perspective is novel, as it redefines microbial causality as the direct impact of one species on another within the network. Moreover, the steady-states data required in our study has less strict requirement than what is suggested in the current literature.

However, we acknowledge a limitation: the network inference derived solely from observational data results in a class of graphs, rather than a singular, definitive directed network. This arises from the fact that, although our inferred graph represents causal relationships, it is constructed from conditional independence tests. Consequently, a set of graphs forming a Markov Equivalence Class will exhibit identical observational data distributions. To refine the network structure, the acquisition of interventional data is imperative, transforming this challenge into an experimental design conundrum. Our work pioneers in providing a solution to identify a directed, potentially cyclic network within the class depicted by a Partial Ancestral Graph.

Another area of exploration is the field of microbial network control. We recognize

the untapped potential of reinforcement learning in enhancing control algorithm performance. Applying reinforcement learning to the specific context of microbial network control, we showcase an innovative foray into this realm. Our approach marks a step towards integrating advanced AI techniques in biological network management, opening new avenues for research and application.

References

Adams, G. O., Fufeyin, P. T., Okoro, S. E., Ehinomen, I., et al. (2015). Bioremediation, Biostimulation and Bioaugmention: a Review. *International Journal of Environmental Bioremediation & Biodegradation*, *3*(1), 28–39.

Adrio, J. L., & Demain, A. L. (2003). Fungal Biotechnology. *International Microbiology*, *6*, 191–199.

Angulo, M. T., Moog, C. H., & Liu, Y.-Y. (2019). A Theoretical Framework for Controlling Complex Microbial Communities. *Nature Communications*, *10*(1), 1045.

Baksi, K. D., Kuntal, B. K., & Mande, S. S. (2018). Time: A Web Application for Obtaining Insights Into Microbial Ecology Using Longitudinal Microbiome Data. *Frontiers in Microbiology*, *9*, 36.

Ban, Y., An, L., & Jiang, H. (2015). Investigating Microbial Co-occurrence Patterns Based on Metagenomic Compositional Data. *Bioinformatics*, *31*(20), 3322–3329.

Barott, K. L., & Rohwer, F. L. (2012). Unseen Players Shape Benthic Competition on Coral Reefs. *Trends in Microbiology*, *20*(12), 621–628.

Basso, K., Margolin, A. A., Stolovitzky, G., Klein, U., Dalla-Favera, R., & Califano, A. (2005). Reverse Engineering of Regulatory Networks in Human B Cells. *Nature Genetics*, *37*(4), 382–390.

Boddy, L., & Hiscox, J. (2016). Fungal Ecology: Principles and Mechanisms of Colonization and Competition by Saprotrophic Fungi. *Microbiology Spectrum*, *4*(6), 4–6.

Bongers, S., Forr'e, P., Peters, J., & Mooij, J. M. (2021). Foundations of Structural Causal Models With Cycles and Latent Variables. *The Annals of Statistics*, *49*(5), 2885–2915.

Butte, A. J., & Kohane, I. S. (1999). Mutual Information Relevance Networks: Functional Genomic Clustering Using Pairwise Entropy Measurements. In *Biocomputing 2000* (pp. 418–429). World Scientific.

Chickering, D. M. (2002). Optimal Structure Identification With Greedy Search.

Journal of Machine Learning Research, *3*(Nov), 507–554.

Eberhardt, F., Glymour, C., & Scheines, R. (2012). On the Number of Experiments Sufficient and in the Worst Case Necessary to Identify All Causal Relations Among N Variables. *ArXiv Preprint ArXiv:1207.1389*.

Fan, C., Zeng, L., Sun, Y., & Liu, Y.-Y. (2020). Finding Key Players in Complex Networks Through Deep Reinforcement Learning. *Nature Machine Intelligence*, *2*(6), 317–324.

Fang, H., Huang, C., Zhao, H., & Deng, M. (2015). CCLasso: Correlation Inference for Compositional Data Through Lasso. *Bioinformatics*, *31*(19), 3172–3180.

Fang, H., Huang, C., Zhao, H., & Deng, M. (2017). GCoda: Conditional Dependence Network Inference for Compositional Data. *Journal of Computational Biology*, *24*(7), 699–708.

Faust, K., & Raes, J. (2012). Microbial Interactions: From Networks to Models. *Nature Reviews Microbiology*, *10*(8), 538–550.

Friedman, J., & Alm, E. J. (2012). Inferring Correlation Networks From Genomic Survey Data. *PLoS Computational Biology*, *8*(9), e1002687.

Hauser, A., & B"uhlmann, P. (2014). Two Optimal Strategies for Active Learning of Causal Models From Interventional Data. *International Journal of Approximate Reasoning*, *55*(4), 926–939.

He, Y.-B., & Geng, Z. (2008). Active Learning of Causal Networks With Intervention Experiments and Optimal Designs. *Journal of Machine Learning Research*, *9*(Nov), 2523–2547.

Herrero, J., D'ıaz-Uriarte, R., & Dopazo, J. (2003). An Approach to Inferring Transcriptional Regulation Among Genes From Large-scale Expression Data. *Comparative and Functional Genomics*, *4*(1), 148–154.

Hyttinen, A., Eberhardt, F., & Hoyer, P. O. (2012). Learning Linear Cyclic Causal Models With Latent Variables. *The Journal of Machine Learning Research*, *13*(1), 3387–3439.

Hyttinen, A., Hoyer, P. O., Eberhardt, F., & Jarvisalo, M. (2013). Discovering Cyclic Causal Models With Latent Variables: A General SAT-based Procedure. *ArXiv Preprint ArXiv:1309.6836*.

Itani, S., Ohannessian, M., Sachs, K., Nolan, G. P., & Dahleh, M. A. (2010). Structure Learning in Causal Cyclic Networks. In *Causality: Objectives and Assessment* (pp. 165–176).

Kelly, D., Conway, S., & Aminov, R. (2005). Commensal Gut Bacteria: Mechanisms of Immune Modulation. *Trends in Immunology*, *26*(6), 326–333.

Koutrouli, M., Karatzas, E., Paez-Espino, D., & Pavlopoulos, G. A. (2020). A Guide to Conquer the Biological Network Era Using Graph Theory. *Frontiers in Bioengineering and Biotechnology*, *8*, 34.

Kurtz, Z. D., M"uller, C. L., Miraldi, E. R., Littman, D. R., Blaser, M. J., & Bonneau, R. A. (2015). Sparse and Compositionally Robust Inference of Microbial Ecological Networks. *PLoS Computational Biology*, *11*(5), e1004226.

Li, Y. (2017). Deep Reinforcement Learning: An Overview. *ArXiv Preprint ArXiv:1701.07274*.

Lidicker Jr, W. Z. (1979). A Clarification of Interactions in Ecological Systems. *Bioscience*, *29*(8), 475–477.

Lindmark, G., & Altafini, C. (2018). Minimum energy control for complex networks. *Scientific Reports*, *8*(1), 3188. doi: 10.1038/s41598-018-21398-7

Liu, Y.-Y., Slotine, J.-J., & Barab'asi, A.-L. (2011). Controllability of Complex Networks. *Nature*, *473*(7346), 167–173.

Lo, C., & Marculescu, R. (2017). MPLasso: Inferring Microbial Association Networks Using Prior Microbial Knowledge. *PLoS Computational Biology*, *13*(12), e1005915.

MacLean, A. M., Finan, T. M., & Sadowsky, M. J. (2007). Genomes of the Symbiotic Nitrogen-fixing Bacteria of Legumes. *Plant Physiology*, *144*(2), 615–622.

McGregor, K., Labbe, A., & Greenwood, C. M. (2020). MDiNE: a Model to Estimate Differential Co-occurrence Networks in Microbiome Studies. *Bioinformatics*, *36*(6), 1840–1847.

Mooij, J. M., & Claassen, T. (2020). Constraint-based Causal Discovery Using Partial Ancestral Graphs in the Presence of Cycles. In *Conference on Uncertainty in Artificial Intelligence* (pp. 1159–1168).

Pearl, J., et al. (2000). Models, Reasoning and Inference. *Cambridge, UK: CambridgeUniversityPress*, *19*(2), 3.

Sch"afer, J., & Strimmer, K. (2005). An Empirical Bayes Approach to Inferring Large-scale Gene Association Networks. *Bioinformatics*, *21*(6), 754–764.

Shanmugam, K., Kocaoglu, M., Dimakis, A. G., & Vishwanath, S. (2015). Learning Causal Graphs With Small Interventions. *Advances in Neural Information Processing Systems*, *28*.

Spirtes, P. L. (2013). Directed Cyclic Graphical Representations of Feedback Models. *ArXiv Preprint ArXiv:1302.4982*.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT press.

Treloar, N. J., Fedorec, A. J., Ingalls, B., & Barnes, C. P. (2020). Deep Reinforcement Learning for the Control of Microbial Co-cultures in Bioreactors. *PLoS Computational Biology*, *16*(4), e1007783.

Triantafillou, S., & Tsamardinos, I. (2015). Constraint-based Causal Discovery From Multiple Interventions Over Overlapping Variable Sets. *The Journal of Machine Learning Research*, *16*(1), 2147–2205.

Vidal-Diez de Ulzurrun, G., & Hsueh, Y.-P. (2018). Predator-prey Interactions of Nematode-trapping Fungi and Nematodes: Both Sides of the Coin. *Applied Microbiology and Biotechnology*, *102*, 3939–3949.

Zañudo, J. G. T., Yang, G., & Albert, R. (2017). Structure-based Control of Complex Networks With Nonlinear Dynamics. *Proceedings of the National Academy of Sciences*, *114*(28), 7234–7239.

Zhang, J. (2008). On the Completeness of Orientation Rules for Causal Discovery in the Presence of Latent Confounders and Selection Bias. *Artificial Intelligence*, *172*(16-17), 1873–1896.

3 | Comparative Analysis of Microbial Network Reconstruction Methods

Abstract

In the multifaceted realm of microbial ecosystems, unraveling the web of interactions within microbial communities is crucial for a deeper understanding of ecological systems' complex behaviors. This chapter offers a detailed review of the various methodologies currently in use for deducing microbiome interaction networks. These methodologies grapple with the unique challenges posed by microbial data, which are often compositional, high-dimensional, sparse, and heterogeneous. The surge in high-throughput sequencing technologies has catalyzed breakthroughs in microbial community studies, especially in investigating the gut microbiome and its significant impact on human health and disease. Nonetheless, the distinct properties of microbiome data pose substantial hurdles in reliably mapping ecological interaction networks.

In our analysis, we will categorize the methodologies into three primary approaches. Firstly, the correlation-based approaches focus on inferring undirected networks, capturing co-occurrence and co-exclusion patterns among microbial taxa. Secondly, the directed graphical model-based approaches aim to infer directed networks, revealing potential causal pathways and directional influences within the microbial communities. Lastly, we will delve into causal inference-related approaches, which seek to understand the underlying causal mechanisms driving microbial interactions. This comprehensive categorization aims to provide clarity in the landscape of microbial network inference methodologies and highlight the distinct features and applications of each approach.

3.1 Introduction

Microorganisms, the pioneering inhabitants of Earth, occupy diverse habitats ranging from extreme environments like hydrothermal vents (Van Dover, 2000) to the human gut, influencing health and disease (Qin et al., 2012; Turnbaugh et al., 2009). Their study has been revolutionized by 16S rRNA sequencing technology, which provides a window into the vast microbial diversity and their complex interactions (Woese & Fox, 1977; Caporaso et al., 2012).

Microbial interactions, varying from symbiotic to predatory, shape the dynamics of their communities and impact their environments. These interactions are modeled as networks, capturing the essence of microbial ecosystems (Gill et al., 2006; Cogen, Nizet, & Gallo, 2008; Fierer et al., 2012). Network models utilize weighted and directed edges to represent the intensity and directionality of interactions, with the potential to decipher complex microbial behaviors (Pace, 1997; Huttenhower et al., 2012).

Reconstructing microbial interaction networks from relative abundance data is a complex challenge. Traditional correlation-based methods are limited in differentiating direct from indirect interactions (Weiss et al., 2016). Directed graphical model network inference methods offer a more precise picture of microbial dynamics, while causal inference approaches seek to establish causality within these networks (Mainali, Bewick, Vecchio-Pagan, Karig, & Fagan, 2019).

This chapter navigates through the realm of microbial network reconstruction, evaluating methods from correlational to causal inference. It culminates in a discussion of the field's unresolved issues and future research trajectories, setting the stage for the next leap in understanding microbial community networks.

In conclusion, this chapter aims to delve deeper into the mathematical intricacies and distinctive characteristics of various methodologies for reconstructing microbial networks. While Chapter 2 laid the groundwork with an overview of the literature in microbial network reconstruction and causal inference, here we provide a more extensive survey and comparative analysis. We will dissect the mathematical

underpinnings of numerous methods, some of which were briefly introduced in Section 2.3 of the previous chapter. Additionally, we build upon the general discussion of causal inference presented in Section 2.4 to hone in on its specific interpretation and preliminary applications within microbial community networks. This chapter not only enriches the survey with a more detailed mathematical perspective but also spotlights the burgeoning role of causal inference in understanding the complex dynamics of microbial ecosystems.

3.2 Network Reconstruction in Microbial Ecosystems: Methodologies and Challenges

3.2.1 Varieties of Microbial Interaction Networks

In the realm of microbial community analysis, researchers encounter a variety of network structures, each offering unique insights into the complex interplays within these ecosystems. These network structures can be broadly categorized into three distinct types: undirected networks with edge weights, directed networks with edge weights, and variable-mediated networks. Here, we delve into the characteristics and implications of each type.

Undirected networks with edge weights are a fundamental form of representation in microbial community analysis. In these networks, microbes are denoted as nodes, and the interactions between them are represented as edges without a specific direction. The edge weights quantify the strength or magnitude of the interactions, providing a measure of how strongly two species are associated. This form of network is particularly useful for identifying potential synergistic or antagonistic relationships among species, but it does not provide insights into the directionality of these interactions. An illustration is represented in the first graph of Figure 3.1.

Directed networks with edge weights extend the undirected model by incorporating the directionality of interactions. Each edge in these networks not only carries a weight signifying the interaction's strength but also a direction, indicating the

source and target of the interaction. This directional aspect is crucial for understanding the causal or influential nature of relationships within the microbial community. For instance, if a directed edge points from species A to species B, it suggests that changes in A could have a direct impact on B. An illustration is represented in the second graph of Figure 3.1.

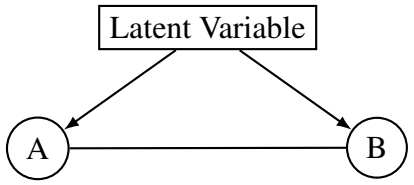
Variable-mediated networks introduce an additional layer of complexity, where interactions are influenced or mediated by external variables. These variables can be diverse, ranging from host-related factors in host-mediated networks to environmental conditions in environment-mediated networks. Such networks are instrumental in understanding how external factors shape the dynamics of microbial communities. They provide a more holistic view by linking the internal microbial interactions with external influences, offering insights into how changes in the environment or host can ripple through the microbial community. An illustration is represented in the graph of Figure 3.2.

Each of these network types plays a vital role in microbial ecology, offering distinct perspectives on the interactions and dynamics within microbial communities. Understanding the nuances of these networks is essential for a comprehensive analysis of microbial ecosystems.



Undirected Network Directed Network

Figure 3.1: Illustration of undirected and directed network



Variable-Mediated Network

Figure 3.2: Illustration of variable-mediated network such as metabolite-mediated, host-mediated, and environment-mediated interactions

3.2.2 Challenges

In the study of microbial ecosystems, the reconstruction of interaction networks within microbial communities is an endeavor of profound significance, one that offers insights into the intricate dynamics governing these complex biological systems. However, this task is replete with unique challenges and difficulties, largely stemming from the inherent properties of microbial data and the subtleties of microbial interactions. This section delves into the various obstacles and complexities that researchers encounter in this endeavor, highlighting the multifaceted nature of these challenges and their implications for the field.

The endeavor to reconstruct microbial community interaction networks is not a straightforward task. It involves deciphering a myriad of complex relationships among various microbial species, each influencing and being influenced by others in a dynamic environmental context. The nature of the data available, the methodologies employed, and the biological realities of these systems all contribute to the complexities faced in this domain.

In the realm of microbial community analysis, the sequencing of 16S rRNA gene data presents unique challenges due to its inherent compositional nature. This characteristic arises from both biological and technical aspects of the sequencing pro-

cess. One notable issue is the high sparsity of the sequencing count matrices; often, 70-95 percent of the values in these matrices are null. This sparsity is attributed to the limited sequencing depth achievable with current technologies and the nature of sample collection.

Another significant challenge is that the count data from sequencing do not directly reflect the absolute abundance of microbial taxa. Instead, these counts represent proportions of the total microbial community present in the sample. This is a critical distinction because an increase in the count of one taxon relative to the sequencing depth automatically results in a proportional decrease in the counts of other taxa. This phenomenon, known as compositional bias, can lead to misleading interpretations if not appropriately addressed.

Furthermore, the total number of counts obtained during a sequencing run often varies between samples, adding another layer of complexity to the analysis. The dependence of species abundances on a constant sum, such as the total sequencing depth, can inadvertently induce negative correlations when using relative abundances. This issue is particularly pronounced in spatial metagenomic survey data, where species abundances are often unevenly distributed. As a result, the negative bias, also referred to as the compositional effect, can lead to false predictions of interactions between species.

Addressing the compositional effect is vital for the accurate interpretation of microbial interactions inferred from sequencing data. Without proper correction for this effect, the use of correlation-based methods to infer interactions can be significantly hampered. Therefore, understanding and mitigating the limitations imposed by the compositional nature of sequencing data is a crucial challenge in the field of microbial community network reconstruction.

Another critical challenge in reconstructing microbial interaction networks is discerning true direct interactions from indirect dependencies. Indirect dependencies occur when the observed association between two microbial taxa is not due to a direct interaction between them, but rather a result of their interaction with one or more other taxa or environmental factors. This issue is a fundamental concern in

network analysis, as it can lead to the erroneous detection of false links, thereby distorting the true structure of the microbial community's interaction network.

Indirect dependencies can manifest in various forms. For example, two microbial species might appear to be correlated not because they interact directly with each other, but because they both respond similarly to an external environmental factor, such as pH or temperature. Similarly, a third microbial species acting as a common source or sink for resources might indirectly link two other species. These scenarios can give rise to spurious correlations that are mistaken for direct interactions, leading to a misrepresentation of the network structure.

Addressing this challenge requires analytical approaches that can distinguish between direct and indirect interactions. One such approach involves the use of conditional dependence methods, where the correlation between two species is evaluated while controlling for the effects of other species or environmental variables. This method can help in reducing the likelihood of detecting false positives due to indirect dependencies. Another strategy involves the application of network deconvolution techniques, which aim to infer direct interactions by removing indirect effects from the observed correlation matrix. These techniques are grounded in the assumption that the observed network is a convoluted mix of direct and indirect interactions, and that mathematical transformations can tease apart these components to reveal the underlying direct network.

Several computational tools and methodologies have been developed to tackle the challenge of indirect dependencies. For instance, methods such as SPIEC-EASI (SParse InversE Covariance Estimation for Ecological Association Inference) (Kurtz et al., 2015) utilize graphical models to estimate direct associations by inferring the inverse covariance matrix of microbial abundances. Other approaches, like CoNet (Faust & Raes, 2016) and MENA (Microbial Ecological Network Analysis) (Deng et al., 2012), use ensemble methods that combine multiple correlation and association measures to increase the robustness of inferred networks against false positives.

Despite the advancements in computational methods, the challenge of accurately identifying direct interactions in microbial networks remains an active area of re-

search. Recent studies emphasize the integration of multi-omics data and environmental variables to enhance the accuracy of network inference. For instance, integrating metagenomic with metatranscriptomic or metabolomic data can provide additional layers of information that help distinguish direct microbial interactions from indirect associations mediated by metabolic exchanges or shared responses to environmental cues.

The pursuit of understanding microbial community interactions confronts the dual challenge of high-dimensionality and sparsity in microbiome data, predominantly sourced from high-throughput sequencing technologies like 16S rRNA gene sequencing. These challenges manifest in datasets where the microbial taxa outnumber the samples, creating an analytical landscape fraught with complexity.

The high-dimensional nature of microbiome data inherently complicates analysis due to the 'curse of dimensionality' (Armstrong et al., 2022). This term encapsulates the phenomena arising in analyzing data with numerous dimensions, where the data required for reliable analysis grows exponentially with dimensionality. In such environments, conventional statistical methodologies falter, and models risk overfitting capturing noise rather than the underlying biological reality.

Complementing the high-dimensionality, the sparse nature of microbiome data adds another layer of complexity (Pan, 2021). Microbial communities often present a multitude of taxa that exist in low abundance or are undetectable in many samples. This sparsity can skew network inference, particularly impacting methods grounded in correlation or covariance.

3.3 Methods Overview

This section presents an extensive review of the current methodologies employed in the reconstruction of microbial community networks. These methods are systematically categorized based on the nature of the networks they reconstruct. We delve into the foundational principles underpinning each method and the specific challenges they are designed to address. Furthermore, this review provides a thorough

analysis of the advantages and limitations inherent to each approach, alongside discussions on the contexts in which they are optimally deployed.

Our exploration covers a spectrum of techniques, each distinct in its approach to network reconstruction. The primary focus is on elucidating how these methods conceptualize and tackle key issues in microbial community analysis. The discussion extends to assessing the strengths and weaknesses of each technique, offering a balanced perspective on their applicability and efficacy.

The categorization is centered around the type of network output generated by each method, ranging from undirected to directed networks, and including mixed network models that incorporate host or environmental factors. This classification facilitates a structured understanding of the varied approaches and their relevance to specific research objectives within the realm of microbial community studies.

In the subsequent sections, each method is examined in detail, shedding light on the core ideas driving these techniques, the types of data they utilize, and the solutions they offer to prevalent challenges in microbial network reconstruction. This analysis aims to provide readers with a comprehensive understanding of the current state of the art in microbial network modeling, highlighting the unique contributions and potential drawbacks of each approach.

3.3.1 Undirected Network Inference Approaches

In this section, we explore the realm of correlation-based methods employed for the reconstruction of undirected microbial networks. These methods are fundamental in elucidating the interplay among microbial taxa, particularly in deciphering the intricate associations within microbial communities. By leveraging statistical correlations, these techniques provide insights into the potential interactions between different microbial species.

Correlation-based methods for network reconstruction rely on the principle of identifying statistical associations between various microbial taxa. These methods are instrumental in inferring the presence and strength of interactions, albeit without

providing directional information. The primary focus is on understanding the pairwise relationships between species, which are represented as edges in an undirected graph.

Pearson Correlation: The Pearson correlation coefficient, often denoted as r, measures the linear relationship between two variables. In the context of microbial networks, it quantifies the linear association between the abundances of two microbial taxa. Mathematically, the Pearson correlation between two taxa X and Y is given by:

$$r_{XY} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}} \tag{3.1}$$

where n is the number of samples, X_i and Y_i are the abundances of taxa X and Y in the i-th sample, respectively, and $\bar{X}$ and $\bar{Y}$ are the mean abundances of X and Y.

Spearman Correlation: The Spearman correlation coefficient, denoted as ρ, is a non-parametric measure of rank correlation. It assesses how well the relationship between two variables can be described using a monotonic function. In microbial networks, it is particularly useful for capturing non-linear relationships between taxa abundances. The Spearman correlation between two taxa X and Y is calculated as the Pearson correlation between the rank values of X and Y:

$$\rho_{XY} = \frac{\sum_{i=1}^{n}(\text{rank}(X_i) - \bar{\text{rank}}(X))(\text{rank}(Y_i) - \bar{\text{rank}}(Y))}{\sqrt{\sum_{i=1}^{n}(\text{rank}(X_i) - \bar{\text{rank}}(X))^2}\sqrt{\sum_{i=1}^{n}(\text{rank}(Y_i) - \bar{\text{rank}}(Y))^2}} \tag{3.2}$$

where $\text{rank}(X_i)$ and $\text{rank}(Y_i)$ are the rank values of X_i and Y_i, and $\bar{\text{rank}}(X)$ and $\bar{\text{rank}}(Y)$ are the average ranks of X and Y.

Pearson and Spearman correlations are extensively used in microbial network reconstruction to identify potential interactions between taxa based on their abundance profiles across different samples. While Pearson is suitable for detecting linear relationships, Spearman can capture more complex, monotonic relationships. They form the basis of an undirected, weighted network where nodes represent taxa and edges represent significant correlations between them.

In consideration of the compositional attributes of microbial data, highlighted in the

previous section, conventional correlation metrics, despite their efficacy, are prone to producing spurious results. To mitigate this predicament, the *CCREPE* (Compositionality Corrected by REnormalization and PErmutation) package in R stands as an essential resource for compositional data analysis, particularly in microbial research (Schwager, Weingart, Bielski, & Huttenhower, 2014). This package skillfully tackles the challenge presented by the summation constraint in microbial abundance data, which can skew correlation interpretations. *CCREPE* effectively counters this by offering adjusted p-values, q-values, and Z-scores for pairwise correlations, thereby enabling more precise and reliable discernment of significant correlations in datasets where naive methods might otherwise be misled by compositional distortions.

CCREPE encompasses two main functions:

- `ccrepe`: Calculates similarity measures, p-values, and q-values using bootstrap and permutation matrices.
- `nc.score`: Extends the checkerboard score to ordinal data for species-level co-variation and co-exclusion analysis.

The algorithm, accessible at `http://huttenhower.sph.harvard.edu/ccrepe` (Schwager et al., 2014), represents a sophisticated approach to addressing the compositional effect in microbial data analysis. This algorithm discerns the significant divergence of a bootstrap-based distribution of correlation coefficients from a null distribution, which is ingeniously generated using uncorrelated species profiles normalized to mimic a compositional effect. The process involves repeated shuffling of a selected species' abundances, followed by normalization of the abundance matrix to relative abundances, and subsequent recalculation of correlations. Concurrently, the bootstrap distribution is formulated by sampling columns from the abundance matrix and recalculating the correlation coefficients. Notably, CCREPE is particularly conservative in cases of low signal-to-noise ratios, a feature that manifests as a wider, more centered bootstrap distribution, ensuring judicious correlation identification.

Continuing our exploration of network reconstruction tools, SparCC (Sparse Correlations for Compositional Data) (Friedman & Alm, 2012) emerges as a vital asset in the realm of microbiome research. SparCC has been explicitly designed to navigate the complexities of compositional data, which is a hallmark of microbiome studies. This tool stands out for its proficiency in discerning meaningful correlations amidst datasets where the emphasis lies on relative proportions rather than absolute numbers.

SparCC's methodology is predicated on the assumption of sparse interactions among microbial taxa, mirroring the intricate yet sparsely connected nature of microbial ecosystems. It employs an innovative iterative approach to approximate the covariance matrix, adeptly bypassing the usual compositional data challenge of spurious correlations. This iterative process enables SparCC to reliably pinpoint biologically significant correlations, efficiently mitigating compositional biases. A standout feature of SparCC is its adeptness in handling large and sparse datasets, a common characteristic in microbiome studies, thus making it an indispensable tool for researchers aiming to infer robust microbial networks.

As characterized in SparCC, the variance matrix based approaches are proposed firstly by (Aitchison, 1982), this literature also proposed the centered log-ratio (CLR) transformation as an alternative to tackle the problem of compositionality of data.

Centered log-ratio (CLR) transformation is a key technique in the analysis of compositional data, especially relevant in the context of microbial network studies. Compositional data, like microbial abundance data, is characterized by the constraint that the total sum of parts is constant. This can lead to spurious correlations if not correctly addressed. CLR transformation addresses this issue by converting compositional data into log-ratios, effectively normalizing the data and alleviating the compositional constraint.

Mathematically, the CLR transformation of a compositional vector $\mathbf{x} = (x_1, x_2, \ldots, x_D)$ is defined as:

$$\text{CLR}(\mathbf{x}) = \left(\log \frac{x_1}{g(\mathbf{x})}, \log \frac{x_2}{g(\mathbf{x})}, \ldots, \log \frac{x_D}{g(\mathbf{x})} \right) \tag{3.3}$$

where $g(\mathbf{x})$ is the geometric mean of the components of $\mathbf{x}$, given by:

$$g(\mathbf{x}) = \left(\prod_{i=1}^{D} x_i \right)^{\frac{1}{D}} \tag{3.4}$$

In this transformation, each component x_i of the vector $\mathbf{x}$ is divided by the geometric mean and log-transformed. This approach effectively normalizes the data, facilitating correlation and regression analyses in compositional contexts, such as microbiome studies. The CLR transformation ensures that the resulting data are free from the constraints of a constant sum, allowing a more accurate reflection of the underlying relationships among the components.

CCLasso is a novel approach to microbial network inference that leverages the centered log-ratio (CLR) transformation to effectively deal with the compositional nature of microbial abundance data (Fang, Huang, Zhao, & Deng, 2015). This method seeks to find a sparse inverse covariance matrix, which is crucial in revealing the underlying network structure among different microbial taxa. The sparsity is achieved using the LASSO (Least Absolute Shrinkage and Selection Operator) technique, a popular regularization method in statistical modeling and machine learning.

LASSO, introduced by (Tibshirani, 1996), is a regression analysis method that performs both variable selection and regularization to enhance the prediction accuracy and interpretability of the statistical model it produces. The LASSO approach modifies the traditional least squares estimation by imposing an L_1 penalty on the magnitude of the coefficients. Mathematically, the LASSO solution is defined as:

$$\hat{\beta}^{lasso} = \arg\min_{\beta} \left\{ \sum_{i=1}^{n} (y_i - \beta_0 - \sum_{j=1}^{p} x_{ij}\beta_j)^2 + \lambda \sum_{j=1}^{p} |\beta_j| \right\} \tag{3.5}$$

where y_i is the response variable, x_{ij} are the predictors, β_j are the coefficients to be estimated, β_0 is the intercept, n is the number of observations, p is the number of

predictors, and λ is a tuning parameter that controls the strength of the penalty. The L_1 penalty term, $\lambda \sum_{j=1}^{p} |\beta_j|$, encourages a sparse solution, effectively performing variable selection by driving some coefficients to be exactly zero.

In the context of CCLasso, the LASSO method is employed to identify a sparse structure in the inverse covariance matrix of CLR-transformed data, which represents the direct associations between microbial taxa.

Meta-Network (Yang, Yu, Cheng, & Ning, 2019) represents an advanced hybrid approach in microbial network analysis, blending the strengths of Pearson Correlation and graph-based methodologies. It is particularly designed to study indirect relationships and nonlinear associations within microbial communities.

At the core of Meta-Network lies the FS-Weight method, a graph-based technique that enhances the understanding of indirect relationships between microbial taxa. This aspect of the method helps to uncover more complex interactions that are not directly apparent, providing a deeper insight into the ecological dynamics of microbial communities. Additionally, Meta-Network employs the PCA-PMI (Part Mutual Information adjusted by Path Consistency Algorithm) method, tailored to capture nonlinear associations.

A table of summary is given in Table 3.1.

3.3.2 Directed Network Approaches

While correlation-based methods are instrumental in identifying potential associations within microbial networks, they inherently lack the precision to discern between direct and indirect interactions. This limitation often leads to the portrayal of an oversimplified network structure, potentially masking the true complexity of microbial interdependencies. To address this critical shortcoming, researchers have turned to methods that model conditional independence. These approaches, typically characterized by their more sophisticated computational frameworks, offer a nuanced perspective of microbial interactions.

Graphical models have emerged as a powerful tool to unravel the complexity of

these microbial networks. They offer a visual and mathematical framework to represent and analyze the dependencies and relationships between different microbial species.

Graphical models are a marriage of probability theory and graph theory. They provide a way to depict the structure of probabilistic models in a graphical form, offering clarity and insight into the data. In the realm of microbial network reconstruction, these models play a crucial role in elucidating the dynamics of microbial interactions. By representing each microbe as a node and the interactions as edges, graphical models transform complex biological data into a more understandable and analyzable format.

The primary advantage of graphical models lies in their ability to simplify and visualize complex relationships. They allow researchers to focus on the interactions that are most critical, filtering out the noise and reducing the dimensionality of the data. Furthermore, these models enable the incorporation of prior knowledge and biological constraints, enhancing the robustness of the inferences drawn.

Graphical models can be broadly classified into two categories: undirected and directed. Undirected models, such as Markov Random Fields, capture symmetrical relationships without implying causation. On the other hand, directed models where we discuss in this section, like Bayesian Networks, are effective in representing causal relationships, a crucial aspect when studying microbial communities.

Although conditional independence models provide a more detailed and accurate representation of microbial networks, they come with their own set of challenges. Notably, these methods demand significantly higher computational resources and longer processing times compared to their correlation-based counterparts. This trade-off underscores the need for a careful consideration of method selection, balancing the desire for detailed network insights with practical constraints on computational capacity and time.

This section highlights the transition from simpler correlation-based methods to more complex conditional dependence models in microbial network reconstruction.

The aim is to provide a nuanced understanding of the evolution of methodologies in this field, acknowledging both the advances they bring and the challenges they introduce.

SPIEC-EASI (SParse InversE Covariance Estimation for Ecological Association Inference) (Kurtz et al., 2015) represents a significant advancement in the field of microbial network reconstruction. This method is distinguished by its dual-strategy approach, designed to address the inherent compositional nature of microbial data. Central to its operation is the application of a Centered Log-Ratio (CLR) transformation, which effectively mitigates compositional biases, setting the stage for more accurate network inference.

The first strategy employed by SPIEC-EASI involves the generation of a graphical network via the estimation of a sparse inverse covariance matrix. This is achieved through sparse graphical model inference with the graphical LASSO (Glasso) algorithm. The graphical LASSO plays a pivotal role in identifying sparse structures within the microbial network, thus enabling the elucidation of direct microbial interactions from complex datasets.

The second strategy within SPIEC-EASI utilizes the Meinshausen-Buhlman method, a node-wise regression model. This approach offers an alternative perspective by iteratively constructing sparse linear models for each species, using the remaining species as covariates. This iterative process allows for a detailed exploration of potential ecological associations among microbial species.

A standout feature of SPIEC-EASI is its use of the Stability Approach to Regularization Selection. This approach determines the appropriate level of sparsity for the network, ensuring that the inferred network structure is neither overly simplistic nor unnecessarily complex. Furthermore, SPIEC-EASI differentiates itself from methods like CCLasso by directly employing the clr-transformed covariance matrix in place of a basis covariance matrix. This approximation becomes increasingly accurate with a growing number of species, offering a robust solution for large-scale microbial studies.

Following the footsteps of SPIEC-EASI, gCoda (Fang, Huang, Zhao, & Deng, 2017) emerges as a novel approach in the realm of microbial network reconstruction, specifically addressing the compositional and high-dimensional nature of microbiome data. This method introduces a refined perspective by modeling the underlying dynamics of microbiome data through a logistic normal distribution.

At the core of gCoda's methodology is the penalized maximum likelihood estimation, aimed at discerning the sparse structure of the inverse covariance matrix associated with latent normal variables. This aspect of gCoda is especially critical in managing the high dimensionality typical of microbiome data sets. To efficiently handle the computational demands of this task, gCoda employs an innovative Majorization-Minimization algorithm. This algorithm effectively navigates the optimization landscape, ensuring the reliable estimation of network structures. Through simulation studies, gCoda has demonstrated superior performance in edge recovery of inverse covariance matrices for compositional data, outclassing existing methodologies, including SPIEC-EASI, across various scenarios.

MInt (Biswas, Mcdonald, Lundberg, Dangl, & Jojic, 2016), standing for Microbial Interaction, represents a significant leap in the field of microbial network reconstruction, introducing a novel hierarchical model that adeptly addresses the intricacies of metagenomics sequencing data. The method adopts a Poisson-multivariate normal hierarchical framework specifically designed to discern direct interactions among microbial taxa.

The strength of MInt lies in its layered approach to data analysis. At the Poisson layer, MInt adeptly controls for confounding factors, a critical step in ensuring the accuracy of the inferred interactions. This layer effectively manages the count-based nature of the data, a common output format in standard metagenomics sequencing experiments. Progressing to the multivariate normal layer, MInt employs an $l1$ penalized precision matrix to capture the taxon-taxon interactions, focusing on the conditional correlation structure within the microbial community.

A unique aspect of MInt is its iterative conditional mode algorithm, which sequentially updates parameters based on the current estimates of others. This approach

ensures a structured optimization process, facilitating the accurate identification of direct microbial interactions. Moreover, MInt's Poisson-multivariate normal model structure allows for a flexible mean-variance framework, adeptly handling overdispersion in count data and enhancing the model's applicability in various scenarios.

However, MInt also demands careful consideration of potential confounding variables. Accurate measurement and inclusion of these variables are paramount to avoid misinterpreting indirect interactions as direct ones. While MInt does not explicitly address the compositional nature of microbiome data, its capability to incorporate confounders such as biological covariates and technical biases, including sequencing depth, positions it as a robust tool in the realm of microbial network inference.

HARMONIES (Jiang et al., 2020), standing for Hybrid Approach foR MicrobiOme Network Inferences via Exploiting Sparsity, introduces an innovative methodology for microbial network reconstruction, adeptly addressing the unique challenges of microbiome sequencing data. This approach distinct in its use of the zero-inflated negative binomial (ZINB) distribution coupled with a Dirichlet prior, specifically targeting the overdispersion and the prevalent zero counts characteristic of microbiome data. HARMONIES further refines its analysis through a graphical lasso approach to infer interactions, which has demonstrated superior performance compared to other methods like SPIEC-EASI and CCLasso, particularly in synthetic data scenarios enriched with zeros.

Unlike traditional normalization techniques that first normalize compositional data and then apply log-ratio transformation, HARMONIES adopts a model-based normalization strategy tailored for microbiome count data. By utilizing the ZINB model with a Dirichlet process prior (DPP), HARMONIES effectively normalizes data, setting the stage for accurate partial correlation estimations and direct association inferences among taxa. This approach provides a nuanced understanding of microbial interactions, showcasing the method's robustness in handling the distinctive challenges presented by microbiome sequencing data.

While HARMONIES excels in managing excessive zeros and over-dispersion, it faces limitations in scenarios of small sample sizes. This constraint can directly affect the reliability of the normalized matrix estimated from the ZINB model. Despite this, HARMONIES remains a significant tool in the field of microbial network inference, particularly for its model-based approach to normalization and sophisticated method for inferring direct associations.

A table of summary is given in Table 3.2.

3.3.3 Causality in Microbial Communities

Understanding causality in microbial communities is crucial for decoding the complex web of interactions that underpin these ecosystems. Causal inference, a field that is rapidly gaining prominence in the analysis of microbial networks, offers powerful insights into the mechanisms driving these interactions. By exploring the cause-and-effect relationships between microbial species, we can uncover the underlying dynamics that shape the structure and function of these communities. This understanding is not only fundamental for theoretical research but also has practical implications in fields ranging from environmental science to health and disease.

While there exists a notable overlap between the concepts of causality and directed network approaches in the study of microbial communities, it is essential to delineate these areas distinctly. Directed network approaches primarily focus on establishing the direction of influence between entities within a network, often through statistical correlations and dependencies. However, causality delves deeper, aiming to uncover the underlying mechanisms and forces driving these interactions. It transcends the mere identification of directionality in relationships, venturing into the realm of cause and effect. This distinction is crucial as it allows for a more nuanced understanding of microbial dynamics beyond mere directional associations. In this section, we place a specific emphasis on causal inference, exploring its unique challenges and methodologies. This specialized focus enables us to dissect the intricate causal webs within microbial ecosystems, providing insights that might remain obscured under traditional directed network analyses. By doing so, we aim to enrich

our understanding of microbial interactions at a fundamental level, paving the way for more accurate models and effective interventions in microbial community management and research.

The emergence of 'TIME' (Temporal Insights into Microbial Ecology) marks a significant advancement in the field of longitudinal microbiome studies (Baksi, Kuntal, & Mande, 2018). This web-based framework is specifically tailored for extracting meaningful insights from time-series microbiome data. 'TIME' is versatile, accepting a broad spectrum of popular data formats and accommodating multiple longitudinal time points with accompanying metadata. This tool extends beyond standard microbiome analytics by incorporating advanced time series analysis methods such as Dynamic Time Warping, Dickey-Fuller test, and notably, Granger causality, for generating interactive and insightful visualizations.

A cornerstone of 'TIME' is its implementation of Granger causality, a statistical concept initially introduced for economic data analysis (Seth, 2007), but now widely applied across various fields, including microbiome studies. Granger causality is grounded in the idea that if a time series X "Granger causes" another series Y, then the historical values of X (alongside Y) will provide a better prediction of Y's current value compared to using the history of Y alone.

Mathematically, Granger causality can be expressed through a set of linear regression models. For two time series X and Y, the hypothesis that X Granger causes Y is evaluated by comparing the following regressions:

$$Y_t = \alpha + \sum_{i=1}^{n} \beta_i Y_{t-i} + \epsilon_t \tag{3.6}$$

$$Y_t = \alpha' + \sum_{i=1}^{n} \beta_i' Y_{t-i} + \sum_{i=1}^{n} \gamma_i X_{t-i} + \epsilon_t' \tag{3.7}$$

Here, Y_t is the current value of Y, Y_{t-i} and X_{t-i} are past values of Y and X respectively, and ϵ_t and ϵ_t' are error terms. The first equation models Y_t using only its own past values, while the second incorporates past values of both X and Y. If the in-

clusion of X's past values provides a significantly better prediction of Y_t (typically assessed through statistical tests on the coefficients γ_i), then X is said to Granger cause Y.

In 'TIME', the Granger causality approach is extended to handle high-dimensional microbiome data. The tool offers options for pairwise Granger causality analysis, a multivariate Granger Lasso approach for dealing with numerous parameters, and an ensemble method that consolidates results from both pairwise and multivariate analyses. This multifaceted approach makes 'TIME' an invaluable asset for predicting causal relationships in complex microbial ecosystems, where controlled experimentation is not always feasible.

Moreover, (Mainali et al., 2019) also used Granger Causality to detect interaction network in the human microbiome. This work critically addresses the limitations of conventional correlation-based analyses in establishing true taxon interactions within the microbiome. Recognizing that correlation alone is insufficient to infer causation, the authors have innovatively applied conditional Granger causality to a comprehensive time series dataset of the human microbiome. This methodological shift allows for a more accurate disentanglement of the networks of correlation and causation among microbial genera.

The study reveals the nuanced nature of microbial interactions, highlighting a weak negative correlation between standard correlation metrics and biological causality. Notably, it uncovers that strong interspecific interactions in the microbiome are predominantly positive, suggesting mutualism or facilitation, whereas strong intraspecific interactions are largely negative, indicative of competition. These interactions also exhibit temporal specificity, with intraspecific interactions acting on shorter timescales. Furthermore, the research sheds light on the conservation of these interactions across different body sites, providing crucial insights into the underlying dynamics of the microbiome.

FlashWeave (Tackmann, Rodrigues, & von Mering, 2019) adeptly addresses the challenges of computational scalability and statistical resolution that have traditionally hampered ecological interaction prediction methods. FlashWeave incorporates

metadata integration, enhancing the interpretability of ecological patterns derived from vast metagenomic datasets.

FlashWeave employs a flexible Probabilistic Graphical Model framework that efficiently processes large-scale microbial abundance data, including datasets encompassing hundreds of thousands of samples. This method exhibits superior performance in runtime and accuracy compared to existing state-of-the-art approaches, as demonstrated in various benchmarking scenarios. A key innovation in FlashWeave is the utilization of a modified version of the semi-interleaved HITON-PC algorithm, which focuses on inferring the Markov blanket for each taxon. This feature ensures that a taxon is conditionally independent of all other taxa in the graph, given its Markov blanket. FlashWeave begins by identifying significantly associated taxa for each taxon, using either Pearson correlation or mutual information as a basis. The algorithm then retains only those taxa that exhibit conditional dependence on a given taxon when considering all combinations of other neighboring taxa. This heuristic-based approach, coupled with the optional inclusion of metadata, allows FlashWeave to disentangle direct microbial associations from confounding factors, particularly useful in cross-study analyses.

By leveraging FlashWeave, researchers have been able to analyze an extensive cross-study dataset comprising 69,818 human gut samples, resulting in one of the largest and most diverse networks of predicted direct gastrointestinal microbial interactions to date.

As we delve deeper into the causality-driven approaches in microbial network reconstruction, it is essential to explore the role of Bayesian networks in this domain. Bayesian networks provide a probabilistic perspective on causal relationships, making them a significant tool for understanding the intricate web of interactions within microbial communities.

Bayesian networks, rooted in probabilistic reasoning and Bayesian inference, offer a versatile framework for modeling the complex and often uncertain interactions in microbial ecosystems (Pearl, 2011). These networks allow for the representation of conditional dependencies and causal relationships between microbial taxa,

capturing the dynamics of microbial interactions in a probabilistic manner.

The application of Bayesian networks in causality analysis within microbial communities involves integrating prior biological knowledge with empirical data. This approach is particularly beneficial in microbial studies, where understanding the directionality and influence of interactions can lead to insights into ecological processes and disease mechanisms. Bayesian networks, through their inherent capacity to deal with uncertainty and incorporate prior information, provide a robust framework for inferring causal relationships in microbial networks. In the realm of microbial community analysis, the integration of Bayesian Networks with Co-occurrence Networks (CoNs) leads to the formulation of Signed Bayesian Networks (sBNs).

The study (Sazal, Mathee, Ruiz-Perez, Cickovski, & Narasimhan, 2020) presents an innovative methodology for constructing sBNs. The proposed approach hinges on the synergy between BNs and CoNs, providing insights into complex associations and colonization patterns within microbial communities. The key novelty lies in leveraging the directed edges of sBNs, in conjunction with negative correlations, to shed light on potential colonization orders.

The methodology employed in the study revolves around the analysis of abundance matrices without the requirement of temporal data. This aspect marks a significant departure from conventional approaches, which often rely heavily on temporal information. The findings of the study reveal a compelling association between the directed edges in sBNs and established colonization orders. This correlation underscores the efficacy of sBNs in unraveling the dynamics of microbial communities.

In (McGeachie et al., 2016), author showcases the application of Dynamic Bayesian Networks (DBNs) to analyze the microbial ecosystem of infants, offering new insights into the developmental patterns of gut microbiota. The study utilized sequencing data of the 16S rRNA gene from 58 preterm infants in a Neonatal Intensive Care Unit. Through the construction of DBNs from this dataset, the researchers successfully captured specific relationships and general trends in the gut microbial ecosystem. Notably, the DBN model highlighted patterns such as the increase of Clostridia, residual Bacilli, and the transition from Gammaproteobacteria

to Clostridia.

One of the striking aspects of this study was the predictive performance of the DBNs. The networks with fewer edges exhibited higher accuracy, particularly in predicting the composition of more challenging subjects. Additionally, the DBNs provided quantitative likelihood estimates for rare abrupt changes in microbiome composition.

This research (Sazal et al., 2020) employs advanced causal inference techniques to elucidate the causal interactions among various entities within a microbiome. The authors termed "disease network." , which integrates an outcome or "disease" variable, thereby enabling the identification of disease-relevant causal factors within the microbiome. The application of interventional techniques to this network facilitates the computation of causal effects of microbial taxa on specific outcomes or conditions. PC-stable algorithm was used to infer the skeleton based on the observational data samples.

3.4 Future Directions

As we have explored various methodologies for reconstructing microbial community networks, it is evident that significant challenges and opportunities for advancement remain. A prominent area for future development lies in addressing the implications of cycles within interaction networks, particularly in microbial communities. While Bayesian network approaches and other causal inference methods assume the absence of cycles, this assumption simplifies the complexity inherent in microbial interactions. Consequently, detecting causality in microbial communities, especially in the presence of cycles, emerges as a promising direction for future research.

The human microbiome forms a dense web of complex, context-dependent interactions. Current tools for inferring both undirected and directed interaction networks often generate inconsistent predictions, highlighting the need for more refined methods that are attuned to the unique statistical characteristics of microbiome data.

There is a demand for tools that are not only general but also accessible to researchers without extensive backgrounds in systems biology or statistics.

Different computational approaches, including species abundance patterns, metabolic reconstruction, and text mining, offer complementary insights into microbial interactions. However, a systematic benchmarking study and an integration scheme for these methods are still lacking. Addressing the dynamic nature of microbial interactions and integrating multiple data types and methods will likely provide a more accurate understanding of microbial dynamics.

Network analysis methods currently struggle with challenges such as compositionality bias, overdispersion, and a poor sample-to-feature ratio. The choice of analysis methods should be carefully considered based on the available computational resources and the specific biological questions being addressed. Furthermore, validation of these methods using universal benchmark datasets is crucial. While network analysis often yields plausible hypotheses, it falls short in inferring causality. Therefore, integrative methods utilizing diverse data types, including shotgun sequencing, metatranscriptomics, and metametabolomics, are needed to unravel the complexity of microbial interactions.

Efforts in computational method development must be matched with experimental studies of microbial communities, particularly in gnotobiology. This will enable the validation of computational findings and provide deeper insights into the complex interplay of microbial interactions.

Table 3.1: Summary of microbial network reconstruction methods of undirected network

Method	**Full Name**	**Implementation**	**Advantages**	**Drawbacks**
CCREPE	Compositionality Corrected by REnormalization and PErmutation	R via Bioconductor https://huttenhower.sph.harvard.edu/ccrepe	Handles compositional bias; Provides corrected p-values, q-values, Z-scores	Conservative in low signal-to-noise ratio scenarios
SparCC	Sparse Correlations for Compositional Data	Python https://bitbucket.org/yonatanf/sparcc	Effective for sparse data; Reduces false positives due to compositional effects	May not capture all true positives; Assumes linearity
CCLasso	Compositional Data Analysis with LASSO	R https://github.com/huayingfang/CCLasso	Identifies sparse basis covariance matrix; Suitable for high-dimensional data	Requires careful parameter tuning; Assumes sparsity
Meta-Network	N/A	R http://www.microbioinformatics.org/software/Meta-Network.html	Captures indirect and nonlinear correlations; Detects clusters and hubs	Complexity in implementation; May require extensive computational resources

Table 3.2: Summary of microbial network reconstruction methods of directed network

Method	**Full Name**	**Implementation**	**Advantages**	**Drawbacks**
SPIEC-EASI	Sparse InversE Covariance Estimation for Ecological Association Inference	R https://github.com/zdk123/SpiecEasi	Suitable for compositional data; Provides two network inference strategies	May struggle with extremely sparse data
gCoda	–	R https://github.com/huayingfang/gCoda	Effective for edge recovery; Addresses compositional nature of data	Limitations in high-dimensional data scenarios
HARMONIES	Hybrid Approach foR MicrobiOme Network Inferences via Exploiting Sparsity	R https://github.com/shuangj00/HARMONIES	Handles zeros and overdispersion; Accounts for sample heterogeneity	Affected by small sample sizes
MInt	–	R https://cran.r-project.org/src/contrib/Archive/MInt/	Controls confounders; Depicts taxon taxon interactions effectively	Potential confounders must be accurately measured

References

Aitchison, J. (1982). The Statistical Analysis of Compositional Data. *Journal of the Royal Statistical Society: Series B (Methodological)*, *44*(2), 139–160.

Armstrong, G., Rahman, G., Martino, C., McDonald, D., Gonzalez, A., Mishne, G., & Knight, R. (2022). Applications and Comparison of Dimensionality Reduction Methods for Microbiome Data. *Frontiers in Bioinformatics*, *2*, 821861.

Baksi, K. D., Kuntal, B. K., & Mande, S. S. (2018). âTimeâ: A Web Application for Obtaining Insights Into Microbial Ecology Using Longitudinal Microbiome Data. *Frontiers in Microbiology*, *9*, 36.

Biswas, S., Mcdonald, M., Lundberg, D. S., Dangl, J. L., & Jojic, V. (2016). Learning Microbial Interaction Networks From Metagenomic Count Data. *Journal of Computational Biology*, *23*(6), 526–535.

Caporaso, J. G., Lauber, C. L., Walters, W. A., Berg-Lyons, D., Huntley, J., Fierer, N., ... others (2012). Ultra-high-throughput Microbial Community Analysis on the Illumina HiSeq and MiSeq Platforms. *The ISME Journal*, *6*(8), 1621–1624.

Cogen, A., Nizet, V., & Gallo, R. (2008). Skin Microbiota: a Source of Disease or Defence? *British Journal of Dermatology*, *158*(3), 442–455.

Deng, Y., Jiang, Y.-H., Yang, Y., He, Z., Luo, F., & Zhou, J. (2012). Molecular Ecological Network Analyses. *BMC Bioinformatics*, *13*, 1–20.

Fang, H., Huang, C., Zhao, H., & Deng, M. (2015). CCLasso: Correlation Inference for Compositional Data Through Lasso. *Bioinformatics*, *31*(19), 3172–3180.

Fang, H., Huang, C., Zhao, H., & Deng, M. (2017). GCoda: Conditional Dependence Network Inference for Compositional Data. *Journal of Computational Biology*, *24*(7), 699–708.

Faust, K., & Raes, J. (2016). CoNet App: Inference of Biological Association Networks Using Cytoscape. *F1000Research*, *5*.

Fierer, N., Leff, J. W., Adams, B. J., Nielsen, U. N., Bates, S. T., Lauber, C. L., ... Caporaso, J. G. (2012). Cross-biome Metagenomic Analyses of Soil

Microbial Communities and Their Functional Attributes. *Proceedings of the National Academy of Sciences*, *109*(52), 21390–21395.

Friedman, J., & Alm, E. J. (2012). Inferring Correlation Networks From Genomic Survey Data. *PLoS Computational Biology*, *8*(9), e1002687.

Gill, S. R., Pop, M., DeBoy, R. T., Eckburg, P. B., Turnbaugh, P. J., Samuel, B. S., ... Nelson, K. E. (2006). Metagenomic Analysis of the Human Distal Gut Microbiome. *Science*, *312*(5778), 1355–1359.

Huttenhower, C., Gevers, D., Knight, R., Abubucker, S., Badger, J. H., Chinwalla, A. T., ... Consortium, T. H. M. P. (2012, Jun 01). Structure, Function and Diversity of the Healthy Human Microbiome. *Nature*, *486*(7402), 207-214. Retrieved from https://doi.org/10.1038/nature11234 doi: 10.1038/nature11234

Jiang, S., Xiao, G., Koh, A. Y., Chen, Y., Yao, B., Li, Q., & Zhan, X. (2020). HARMONIES: a Hybrid Approach for Microbiome Networks Inference Via Exploiting Sparsity. *Frontiers in Genetics*, *11*, 445.

Kurtz, Z. D., M"uller, C. L., Miraldi, E. R., Littman, D. R., Blaser, M. J., & Bonneau, R. A. (2015). Sparse and Compositionally Robust Inference of Microbial Ecological Networks. *PLoS Computational Biology*, *11*(5), e1004226.

Mainali, K., Bewick, S., Vecchio-Pagan, B., Karig, D., & Fagan, W. F. (2019). Detecting Interaction Networks in the Human Microbiome With Conditional Granger Causality. *PLoS Computational Biology*, *15*(5), e1007037.

McGeachie, M. J., Sordillo, J. E., Gibson, T., Weinstock, G. M., Liu, Y.-Y., Gold, D. R., ... Litonjua, A. (2016). Longitudinal Prediction of the Infant Gut Microbiome With Dynamic Bayesian Networks. *Scientific Reports*, *6*(1), 20359.

Pace, N. R. (1997). A Molecular View of Microbial Diversity and the Biosphere. *Science*, *276*(5313), 734–740.

Pan, A. Y. (2021). Statistical Analysis of Microbiome Data: the Challenge of Sparsity. *Current Opinion in Endocrine and Metabolic Research*, *19*, 35–40.

Pearl, J. (2011). Bayesian Networks.

Qin, J., Li, Y., Cai, Z., Li, S., Zhu, J., Zhang, F., ... others (2012). A Metagenome-wide Association Study of Gut Microbiota in Type 2 Diabetes. *Nature*,

490(7418), 55–60.

Sazal, M., Mathee, K., Ruiz-Perez, D., Cickovski, T., & Narasimhan, G. (2020). Inferring Directional Relationships in Microbial Communities Using Signed Bayesian Networks. *BMC Genomics*, *21*, 1–11.

Schwager, E., Weingart, G., Bielski, C., & Huttenhower, C. (2014). CCREPE: Compositionality Corrected by Permutation and Renormalization. *R/Bioconductor Https://doi. Org/10.18129 B*, *9*.

Seth, A. (2007). Granger Causality. *Scholarpedia*, *2*(7), 1667.

Tackmann, J., Rodrigues, J. F. M., & von Mering, C. (2019). Rapid Inference of Direct Interactions in Large-scale Ecological Networks From Heterogeneous Microbial Sequencing Data. *Cell Systems*, *9*(3), 286–296.

Tibshirani, R. (1996). Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *58*(1), 267–288.

Turnbaugh, P. J., Hamady, M., Yatsunenko, T., Cantarel, B. L., Duncan, A., Ley, R. E., ... others (2009). A Core Gut Microbiome in Obese and Lean Twins. *Nature*, *457*(7228), 480–484.

Van Dover, C. L. (2000). *The Ecology of Deep-sea Hydrothermal Vents*. Princeton University Press.

Weiss, S., Van Treuren, W., Lozupone, C., Faust, K., Friedman, J., Deng, Y., ... others (2016). Correlation Detection Strategies in Microbial Data Sets Vary Widely in Sensitivity and Precision. *The ISME Journal*, *10*(7), 1669–1681.

Woese, C. R., & Fox, G. E. (1977). Phylogenetic Structure of the Prokaryotic Domain: the Primary Kingdoms. *Proceedings of the National Academy of Sciences*, *74*(11), 5088–5090.

Yang, P., Yu, S., Cheng, L., & Ning, K. (2019). Meta-network: Optimized Species-species Network Analysis for Microbial Communities. *BMC Genomics*, *20*(2), 143–151.

4 | Microbial Causal Inference Framework for Community Network Reconstruction

Abstract

Inferring interactions between species in microbial communities is a critical step in understanding their underlying mechanisms and predicting their temporal behavior. Here, we introduce Microbial Causal Inference (MCI), a causal inference framework designed for inferring microbial species interactions. The framework is based on the theoretical foundation of structural causal models where equilibrium conditions are satisfied. Our method is extensively analyzed and validated for both acyclic and cyclic cases, including those with feedback loops. We employ both Fast Causal Inference (FCI) and Greedy Fast Causal Inference (GFCI) algorithms to infer the ancestral graph models and decipher the causal relationships they represent. Systematic validation of MCI is conducted using simulated data, with the framework accommodating data samples from communities possessing varying intrinsic parameters. By bridging the gap between network inference in microbial communities and the rapidly advancing field of causal inference, our work represents a significant step towards unraveling the complexity of microbial communities, such as the human gut microbiota. The novelty of this research lies in the integration of network reconstruction with causal inference, and the introduction of reinforcement learning for network control, providing new methodologies for accurately inferring and manipulating microbial interactions under complex conditions.

4.1 Introduction

Microbial communities are prevalent in the real world, spanning naturally occurring communities, such as in anaerobic digestion (Wade, Pattinson, Parker, & Dolfing, 2016), to synthetic microbial consortia (Song, Ding, Jia, Ma, & Yuan, 2014), which are considered as artificially designed microbial communities. In recent years, studies of microbial communities have become increasingly popular due to their vast

potential for applications in various fields, such as environmental enhancement, bio-engineering, and biomedicine, as discussed in (Pulz & Gross, 2004)(Zhou et al., 2014). Network models serve as instrumental tools for representing microbial species and their interactions in microbial communities, making them widely used in this area of study.

The primary motivation for employing network models in the context of microbial communities lies in their capacity to capture the inherent complexity and interconnectedness that characterizes microbial interactions (Sunagawa et al., 2015). Microorganisms engage in a wide spectrum of interactions, ranging from metabolic exchanges to competition, cooperation, and predation. These interactions, while diverse, profoundly shape the structure, function, and stability of microbial communities. By translating these intricate associations into a network framework, researchers gain the ability to visually and analytically dissect the intricate tapestry of relationships that govern the microbial world.

Microbial network reconstruction constitutes a valuable tool for comprehending the intricate interactions and dynamics within microbial communities. Various approaches have been developed in this domain, broadly correlation-based approaches and conditional dependence-based approaches. Many correlation-based methods, like variants of Pearson or Spearman correlation, are commonly used to estimate microbial interactions between pairs of taxa (Deng et al., 2012). A popular method employing this strategy is Sparse Correlations for Compositional Data (SparCC) (Friedman & Alm, 2012). SparCC employs iterative approximation and log-ratio transformed data to infer correlations between components. SparCC's advantage lies in its ability to mitigate spurious correlations, which are more likely with direct Pearson correlations. However, it comes with a trade-off of higher computational complexity. Correlation-based methods, as commonly utilized, encounter challenges in distinguishing between direct and indirect associations within microbial interactions. To address this limitation, a multitude of methods has been devised to model conditional dependence. A well-known method called SPIEC-EASI (SParse InversE Covariance Estimation for Ecological Association Inference) (Kurtz et al.,

2015), employed two distinct approaches subsequent to applying a CLR (centered log ratio transformation) to the data, addressing compositionality concerns. SPIEC-EASI adeptly determines the appropriate degree of network sparsity using the Stability Approach to Regularization Selection.

As in data persective, until now, network reconstruction has typically assumed the availability of temporal abundance data, which consists of time-series abundance data for each taxon in the microbial community. Many methods often assume a prior choice of population dynamic models to represent the inherent dynamics of the system. However, any choice may deviate from the true model in the real world, and it can be difficult to definitively justify such a choice (Jordan, 2003). Furthermore, while many methods assume the availability of sufficiently informative time-series data, these data are not always available in practice, as mentioned in (Cao, Gibson, Bashan, & Liu, 2017) (Angulo, Moreno, Lippner, Barab'asi, & Liu, 2017). In real-world scenarios, some microbial communities, such as the human gut microbiota, are stable and resilient to changes, meaning that collected temporal data may be heavily influenced by their steady-state behaviors (Consortium et al., 2012)(Lozupone, Stombaugh, Gordon, Jansson, & Knight, 2012). Acquiring sufficiently informative time-series data often requires assuming radical and frequent perturbations to the system, which can sometimes be impossible or even unethical, as in the case of human gut microbiota.

In this chapter, we aim to use directed graphical causal models (Pearl, 2009) to interpret microbial interactions in network models. A class of directed graphical causal models, commonly referred to as Structural Equation Models (SEMs), assumes that the value of each variable (node) is a deterministic function of its direct causes in the graph. More specifically, an edge exists between two variables if, when all other variables are fixed at certain values, the tail variable is exogenously varied. Symbolically, if $X_i \to X_j$ is present in the graph, then X_i is a direct cause of X_j; in other words, a change in X_i will directly affect the distribution of X_j. The microbes in microbial communities are represented by nodes, and typically their population dynamics can be described by a set of ordinary differential equations (ODEs) (Faust

& Raes, 2012), where x_i denotes the population of microbial species i:

$$\mathrm{d}x_i(t)/\mathrm{d}t = x_i(t) f_i(\boldsymbol{x}(t)), \quad i = 1, \ldots, N \tag{4.1}$$

A vast body of literature exists on learning causal graphical models from observational data, assuming that the model is represented by a directed acyclic graph (DAG). Every causal DAG implies a set of conditional independence relationships that can be inferred from the DAG using a concept known as d-separation (Spirtes, Glymour, & Scheines, 2000). Current methodologies encompass constraint-based approaches (Spirtes et al., 2000) (Pearl, 2009) and score-based approaches (Heckerman, Geiger, & Chickering, 1995) (Chickering, 2002). Among the earliest algorithms that consistently operate under i.i.d. sampling, assuming no latent confounders, is the PC algorithm (Spirtes et al., 2000). This algorithm offers a search framework that accommodates various statistical procedures for determining conditional independence. Since its introduction, numerous variations of the PC algorithm have been developed. A key generalization of this algorithm is the Fast Causal Inference (FCI) Algorithm (Zhang, 2008), which accommodates and uncovers unknown confounding variables. The algorithm's results have been demonstrated to be asymptotically accurate even when confounders are present. With the presence of latent and selection variables, maximal ancestral graphs(MAG) are introduced to encode causal relationships between the observed variables via edge marks and they encode conditional independence relationships among the observed variables via m-separation (Richardson & Spirtes, 2002), a generalisation of d-separation. Moreover, different MAGs which describe exactly the same conditional independence relationships form a Markov equivalence class which can be represented by a partial ancestral graph (PAG) (Ali, Richardson, & Spirtes, 2009). PAGs depict causal characteristics shared by every MAG within the Markov equivalence class, and therefore, by every DAG (potentially with latent and selection variables) that is compatible with the observable independence structure under the faithfulness assumption. Although most causal discovery literature typically assumes the system of interest is acyclic, however, in most real-life causal systems, feedback

loops are generally present, as they are necessary for stabilizing the system. Consequently, the causal directed graph associated with such systems exhibits cyclic patterns (P. Spirtes, 1995) (Hyttinen, Eberhardt, & Hoyer, 2012).

In this chapter, following the proof outlined in (Mooij, Janzing, & Sch"olkopf, 2013), we demonstrate that deterministic ordinary differential equations can be transformed within the context of structural causal models, even in cyclic cases. More specifically, we show that by using the generalized Lotka-Volterra model as an example of population dynamics, the mechanisms of microbial communities align seamlessly with structural causal models in the field of probability and statistics. Consequently, the interactions between species within the microbial communities of interest are equivalent to discovering the causal relationships between nodes. To address the network inference problem without being limited to acyclic cases, we adopt the σ-faithful Structural Causal Models introduced in (Mooij & Claassen, 2020). We expand our pipeline to use FCI and GFCI (Greedy Fast Causal Inference) (Ogarrio, Spirtes, & Ramsey, 2016), which is a combination of FCI and Greedy Equivalence Search(GES) (Chickering, 2003), as our main algorithms for deducing causal relationships among species. As such, the choice between these two algorithms is dictated by the specific requirements and constraints of each individual study.

Our approach first formalizes the mechanisms of microbial communities within the framework of structural causal models and then presents an acyclification of potential cyclic ground truth network models without altering the Markov equivalences of the specific directed graphical causal models. The latent projection is then processed in the cases where latent variables are present. We then apply the Fast Causal Inference algorithm and Greedy Fast Causal Inference to capture causal relations between species using steady-state data simulated with different parameters in the same system. Our method circumvents the problem of prior assumptions of population dynamics models, as it is purely data-driven and does not require any estimation of model parameters. Moreover, our approach avoids the issue of insufficient temporal data, as it only requires steady-state data. By connecting the classic network

reconstruction of microbial communities with the rapidly growing causal inference literature, our approach represents a key step towards improving microbial network reconstruction and a better understanding and explanation of complex microbial communities.

In this chapter, the following structure is followed to present the proposed framework of Microbial Causal Inference (MCI). We commence with a thorough exploration of the various methods and underlying principles involved in this process, encompassing preliminary concepts and the assumed dynamics underpinning these communities. A detailed discussion on the Structural Causal Model (SCM) lays the groundwork for understanding our approach. Subsequent sections introduce and elaborate on FCI and GRFCI, pivotal algorithms in our framework. We then discuss our methodology for generating the 'ground truth' in microbial network reconstruction, addressing partially observed systems and the challenge of latent confounding. The effectiveness of MCI is then examined through analysing simulation results in terms of accuracy, scalability, and the impact of sample sizes. A comparison with an existing method is also provided. Lastly, we discuss the implications, limitations, and potential future directions of our research.

4.2 Methods

4.2.1 Preliminaries

In this section, we introduce the terminology used throughout the chapter. We will define and describe both graphical concepts and their interpretations in the context of population dynamics of microbial communities. In the presented framework, nodes correspond to random variables, while links illustrate both conditional independence and parent-child connections.

Definitions Related to Graphical Models

Basically, a directed graph, denoted as $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$, consists of a vertex set $\mathcal{V} = \{X_1, \ldots, X_p\}$ and a directed edge set $\mathcal{E}$.

In the following, we will briefly explore different types of graphs used in our framework, such as directed mixed graphs, maximal ancestral graphs, and partial ancestral graphs. We will discuss their properties and relationships as they are understood in the causal discovery literature.

A *Directed Mixed Graph* (DMG) is a structure denoted as $\mathcal{G} = \langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$ composed of nodes $\mathcal{V}$ and two separate types of edges: *directed edges* $\mathcal{E} \subseteq \{(X_i, X_j) : X_i, X_j \in \mathcal{V}, X_i \neq X_j\}$, and *bidirected edges* $\mathcal{F} \subseteq \{(X_i, X_j) : X_i, X_j \in \mathcal{V}, X_i \neq X_j\}$. A directed edge $(X_i, X_j) \in \mathcal{E}$ is depicted as $X_i \to X_j$ or $X_j \leftarrow X_i$, with X_i being termed a *parent* of X_j. All parents of X_j in the graph $\mathcal{G}$ are denoted as $\mathrm{Pa}_{\mathcal{G}}(X_j) := \{X_i \in \mathcal{V} : X_i \to X_j \in \mathcal{E}\}$. Bidirected edges between X_i and X_j are denoted by $X_i \leftrightarrow X_j$. A sequence connecting two nodes $X_{i_0}, X_{i_n} \in \mathcal{V}$ can be denoted as a tuple $\langle X_{i_0}, e_1, X_{i_1}, e_2, X_{i_2}, \ldots, e_n, X_{i_n} \rangle$ and it is called a *walk* between them. This sequence alternates between nodes and edges in $\mathcal{G}$, where all $X_{i_0}, \ldots, X_{i_n} \in \mathcal{V}$ and all $e_1, \ldots, e_n \in \mathcal{E} \cup \mathcal{F}$. A *path* refers to a walk which contains each vertex at most once. A *directed path* refers to a path that every edge e_k on the path is a directed edge from node $X_{i_{k-1}}$ to node X_{i_k}, i.e. $X_{i_{k-1}} \to X_{i_k}$. A *directed loop* refers to a directed path starting from X_i and ending at X_j, and additionally includes an edge $X_j \to X_i \in \mathcal{E}$, forming a *directed cycle*. An *almost directed cycle* relaxes the additional edge from $X_j \to X_i \in \mathcal{E}$ to $X_j \leftrightarrow X_i \in \mathcal{F}$. The concept of ancestors can be extended by repeatedly tracing back to the parents of X_j in the form $\mathrm{An}_{\mathcal{G}}(X_j) := \{X_i \in \mathcal{V} : X_i = Xi_0 \to X_{i_1} \to \cdots \to X_{i_n} = X_j \text{ in } \mathcal{G}\}$. In a similar manner, we define the descendants of X_i as $\mathrm{De}_{\mathcal{G}}(X_i) := \{X_j \in \mathcal{V} : X_i = Xi_0 \to X_{i_1} \to \cdots \to X_{i_n} = X_j \text{ in } \mathcal{G}\}$. The set of nodes involved in directed cycles passing through $X_i \in \mathcal{V}$ comprise the *strongly connected component*, represented as $\mathrm{Sc}_{\mathcal{G}}(X_i) := \mathrm{An}_{\mathcal{G}}(X_i) \cap \mathrm{De}_{\mathcal{G}}(X_i)$. An *Acyclic Directed Mixed Graph*(ADMG) is a directed mixed graph with no directed cycles in it.

As defined in (Zhang, 2008), with the definition of colliders and inducing path, a directed *Maximal Ancestral Graph* is a directed mixed graph with all the following additional conditions satisfied: 1. There are no self-cycles and at most one

edge between two distinct nodes 2. There are no directed or almost directed cycles 3. Between two non-adjacent nodes, there is no inducing path. Two MAGs are *Markov equivalent*, as defined in Proposition 2 (Zhang, 2008) if they entail the same independence model with respect to *d-separation*, as defined in (Pearl, 2013). Partial Ancestral Graphs, introduced by (Zhang, 2008), offer a compact representation of a set of MAGs which are Markov equivalent, which is often useful for causal analysis. Given our assumption of no selection bias for the sake of simplicity, we will focus on Directed PAGs (DPAGs). These are PAGs which do not have undirected or circle-tail edges, essentially, edges shaped like $\circ\!\!-$, $-\!\!\circ$, and $-$. Here circle corresponds to a variant mark. Formally, a PAG is a graph with six possible edges: $\{\rightarrow, \leftarrow, \circ\!\!\rightarrow, \leftarrow\!\!\circ, \leftrightarrow, \circ\!\!-\!\!\circ\}$ such that it has the same adjacencies as any member MAG of Markov Equivalence Class (MEC) and every non-circle mark is an invariant mark in the MEC. The induced skeleton of a PAG is defined as an undirected graph which maintains the same nodes and includes an edge between any pair of nodes if these nodes are adjacent in the original DPAG.

4.2.2 Assumed Underlying Dynamics of Use

Following the methods of probabilistic graph models(Whittaker, 2009)(Jordan, 2003), each microbe is represented by a vertex, and an edge is present between two nodes if an interaction occurs between the two corresponding species. The adjacency matrix of the microbial network can capture the system's mechanisms. Furthermore, weights can be assigned to edges, with absolute values quantifying the connection strength; greater magnitudes indicate stronger interactions between species. A positive edge weight value implies a beneficial interaction, while a negative value indicates a detrimental interaction. In the context of probabilistic graphical models, the abundance values of species can be viewed as random variables. The population dynamics of microbial communities can be characterized using a series of ODEs:

$$\mathrm{d}x_i(t)/\mathrm{d}t = x_i(t) f_i(\boldsymbol{x}(t)), \quad i = 1, \ldots, N. \tag{4.2}$$

In this equation, $x_i(t)$ represents the absolute abundance of the i-th species at

a given time point, t. The vector $\boldsymbol{x}(t) = (x_1(t), \ldots, x_N(t))^{\mathrm{T}} \in \mathbb{R}^N$ is an N-dimensional vector, showing the abundances of all species at time t. $f_i(\boldsymbol{x}(t))$'s signifies undefined functions, with their functional forms dictating the inherent population dynamics of the microbial communities.

We examine population dynamics that take the following form:

$$\mathrm{d}\boldsymbol{x}(t)/\mathrm{d}t = \mathrm{diag}(\boldsymbol{x})\left[\boldsymbol{A}f_\theta(\boldsymbol{x}) + \boldsymbol{r}\right] \tag{4.3}$$

In this equation, $\boldsymbol{A} = (a_{ij}) \in \mathbb{R}^{N\times N}$ signifies the interaction matrix, encoding the interactions between different species within the microbial communities. A non-zero element, a_{ij}, indicates that species j has an impact on species i. The vector $\boldsymbol{r} = (r_1, \ldots, r_N)^{\mathrm{T}} \in \mathbb{R}^N$ represents the intrinsic growth rate vector.

The function $f_\theta : \mathbb{R}^N \to \mathbb{R}^N$ is defined as:

$$f_\theta\left(x_1, \cdots, x_N\right) = \left(\frac{x_1}{1+\theta_1 x_1}, \cdots, \frac{x_N}{1+\theta_N x_N}\right)^{\top} \tag{4.4}$$

The well-known Generalized Lotka-Volterra (GLV) model can be derived from Equation (4.3) by setting θ to be a zero vector. For $\theta \neq 0$, the function f_θ represents Holling Type II functional responses.

The objective of our study is to reconstruct the interaction network, the ground truth of which is represented by the interaction matrix. The data used for this reconstruction is derived from observations made within the microbial system.

4.2.3 The SCM

If we formulate a directed graph based on the interaction matrix, each node corresponds to a variable and a directed edge from x_i to x_j is present only if the rate of change of x_j with respect to time, $\mathrm{d}x_j(t)/\mathrm{d}t$, depends on x_i. We then label x_i as a parent of x_j and our dynamical system, expressed by the ODEs, can be described

as follows:

$$\mathrm{d}x_i(t)/\mathrm{d}t = x_i(t) f_i\left(\boldsymbol{x}_{\mathrm{Pa}(i)}\right), \quad i = 1, \ldots, N. \tag{4.5}$$

where $x_{\mathrm{Pa}(i)}$ represents the set of parent variables for x_i. At equilibrium, the left-hand side of equation (4.5) equals zero and we focus on non-trivial steady states, implying that none of the abundances are absent. As highlighted in (Mooij et al., 2013), these dynamics can generate equilibrium equations characterizing the equilibrium states. Consequently, the microbial ODEs can be converted into a system of Labeled Equilibrium Equations (LEEs) $\mathcal{E}$ for N variables, comprising N labeled equations of the following form:

$$\mathcal{E}_i: \quad 0 = f_i\left(\boldsymbol{x}_{\mathrm{Pa}_{\mathcal{E}}(i)}\right), \quad i = 1, \ldots, N. \tag{4.6}$$

Here, we follow the terminology of (Mooij et al., 2013) when considering the definition of Labeled Equilibrium Equations. A set of LEEs $\mathcal{E}$ for K variables $X_{i\{i\in\mathcal{I}\}}$ where $\mathcal{I}$ equals $1, \ldots, K$, is comprised of K labeled equations. Each of these equations follows the format:

$$\mathcal{E}_i: \quad 0 = g_i\left(\boldsymbol{X}_{\mathrm{Pa}_{\mathcal{E}}(i)}\right), \quad i \in \mathcal{I}, \tag{4.7}$$

In this formula, $\mathrm{Pa}_{\mathcal{E}}(i) \subseteq \mathcal{I}$ represents the set of parent variables' indices for the variable X_i. The function $g_i : \mathcal{R}_{\mathrm{Pa}_{\mathcal{E}}(i)} \to \mathcal{R}_i$ maps each set of parent variables to an outcome in the real number space. The configuration of a set of LEE $\mathcal{E}$ can be depicted as a directed graph $\mathcal{G}_{\mathcal{E}}$. This graph comprises individual nodes, each corresponding to a variable, and a directed edge that connects X_i to X_j (given that $j \neq i$). The presence of the directed edge is conditioned on whether $\mathcal{E}_i$ is dependent on X_j. A set of LEE $\mathcal{E}$ is deemed as *solvable* if a unique solution exists for the corresponding system of labeled equations, $\mathcal{E}_i$. A set of LEE $\mathcal{E}$ is considered *structurally solvable* if, for each $i \in \mathcal{I}$, the equation is solvable with respect to some set $I_i \subseteq \mathcal{I}$ that satisfies the condition $\mathrm{pa}_{\mathcal{E}}(i)\backslash\{i\} \subseteq I_i \subseteq \mathcal{I}\backslash\{i\}$.

Following the terminology of (Pearl, 2009), we present the definition of deterministic structural causal models. A Deterministic Structural Causal Model $\mathcal{M}$ with K

variables $\{X_i\}_{i\in\mathcal{I}}$ (where $\mathcal{I} := \{1, \ldots, K\}$) is composed of D structural equations, each taking the form:

$$X_i = h_i\left(\boldsymbol{X}_{\mathrm{Pa}_{\mathcal{M}}(i)}\right), \quad i \in \mathcal{I}, \tag{4.8}$$

In this representation, $\mathrm{Pa}_{\mathcal{M}}(i)$ signifies the set of indices for the parent variables of X_i, where this set is a subset of $\mathcal{I}\backslash\{i\}$. Each function h_i maps from $\mathcal{R}_{\mathrm{Pa}_{\mathcal{M}}(i)}$ to $\mathcal{R}_i$.

If the specified LEE (4.6) is structurally solvable, then we could construct an induced SCM $\mathcal{M}_{\mathcal{E}}$ with i'th structural equation defined as:

$$x_i = g_i\left(x_{\mathrm{pa}_{\mathcal{M}_{\mathcal{E}}}(i)}\right) \tag{4.9}$$

Note that this equation is equivalent to the labeled equation $\mathcal{E}_i$ since they have identical solution sets.

Applying the above generic theory to microbial communities, the dynamical differential equations can conceptually be converted into their corresponding SCMs. Subsequently, microbial interactions can be treated as causal relationships. We convert the dynamical differential equations characterizing microbial systems into their corresponding SCMs, and define microbial interactions in terms of causal relationships. Thus, deciphering the ecological network within microbial communities becomes a task equivalent to causal discovery within these SCMs. Note that the involvement of the dynamic differential equations in the above analysis is only for the purpose of establishing the theoretical framework to derive the intended method for microbial network inference; the application of the method does not require prior knowledge of specific differential equations for the microbial communities of interest.

While most of the literature on causal discovery presumes that the underlying causal system is acyclic, leveraging the benefits of ancestral relations, our case is different. Given that we are dealing with microbial communities where feedback loops are commonplace, it is likely that our underlying causal system could exhibit cyclic behavior.

Following the approach in (Mooij & Claassen, 2020), we represent the ground truth interaction network through directed mixed graphs, which can accommodate latent

variables. These graphs contain both directed edges $\rightarrow$ and bi-directed edges $\leftarrow\rightarrow$, and can be derived from directed graphs via a latent projection operation (Verma & Pearl, 2022).

When dealing with a cyclic causal system, we utilize an acyclification of the directed mixed graph to produce an ADMG without altering the σ-separation criterion as per (Mooij & Claassen, 2020). This criterion is essential since it determines Markov equivalence, given that our causal discovery methods are based on conditional independence.

To represent the corresponding MEC of the acquired ADMG, we use a PAG. It is worth noting that our approach relies solely on data, as we do not make any prior assumptions about the microbial communities. Furthermore, due to our reliance on conditional independence as the causal discovery criterion, we inevitably obtain a Markov equivalence class of a graph.

4.2.4 Fast Causal Inference (FCI) and Greedy Fast Causal Inference (RFCI)

First, when the assumption of causal sufficiency is satisfied, the system exhibits no latent variables, which are unmeasured common causes, and no selection variables. We work on the presumption that the causal relationships between variables can be depicted by a Directed Acyclic Graph (Spirtes et al., 2000)(Pearl, 2009). In this graphical representation, the nodes symbolize the random variables, while the edges indicate the direct causal influences. More specifically, X_i is a direct cause of X_j only if it is a parent of X_j and a cause of X_j only if it is an ancestor of X_j. However, in the context of microbial communities, the acyclicity of graphs could not always be satisfied with the presence of feedback loops.

Fast Causal Inference algorithm, which is a variant of PC algorithm (Spirtes et al., 2000), was introduced to learn a Markov equivalence class of Directed Acyclic Graphs that include latent and selection variables, solely from the conditional independence information among the observed variables. The soundness of FCI is

proven in (Spirtes et al., 2000) with the allowance of latent variables and selection variables and the completeness of FCI algorithm is proven by (Zhang, 2008). The output of FCI algorithm is a PAG which describes a Markov equivalence class of MAGs and the conditional independence relationships which can be read off from the MAG use a concept called m-separation, a generalisation of d-separation in DAGs. The ancestral graph models, which are introduced to solve the problem of latent and selection variables, could be utilized to represent causal relationships with possibly cyclic cases. FCI, which outputs PAGs, was surprisingly proven to be accurate when applied to observational data generated by such a system that involves feedback in (Mooij & Claassen, 2020). The d-separation in the context of DAGs and m-separation in the context of MAGs are relaxed to a more generalised setting, which is called σ-separation. With d-separation limited to systems with continuous variables having linear relationships, discrete variables, or those that fall under the acyclic case, The σ-separation Markov property appears to be applicable for a broad variety of cyclic structural causal models that possess non-linear functional relationships between non-discrete variables. This includes, for instance, structural causal models that correspond to equilibrium states of dynamic systems controlled by differential equations (Bongers & Mooij, 2018) (Bongers, Forr'e, Peters, & Mooij, 2021). With the assumed dynamics of use and the SCM under equilibrium introduced, FCI algorithm is the perfect fit for our context. We adopt FCI algorithm as one of the main algorithms to infer the network in our pipeline and introduce a high-level sketch of the algorithm below. The initial phase of the FCI algorithm parallels the PC algorithm. It begins with a fully connected undirected graph with all edges of the form $\circ\!\!-\!\!\circ$ and then estimates an initial skeleton. This process results in an order-independent initial structure. An edge exists between variables X_i and X_j if and only if these variables are conditionally dependent for all potential subsets S, which encompass the neighbours of both X_i and X_j. The algorithm iterates through all pairs (X_i, X_j) of adjacent nodes in the current graph. The edge between X_i and X_j is maintained only if the null hypothesis, which states that X_i and X_j are conditionally independent given the separation set, is rejected at a significance level $alpha$ for all subsets of it. We adopt the Gaussian independence test in our

framework. Following this, the algorithm redefines the $v - structures$ on the final skeleton using information from the separation set. Finally, to the greatest extent possible, the algorithm determines undetermined edge marks (represented as ○–○) by utilizing a subset of the 10 orientation rules proposed by (Zhang, 2008). Despite its label as "Fast Causal Inference" (FCI), the algorithm can be computationally demanding, especially when handling extensive graphs. Moreover, its performance may degrade with smaller sample sizes. To address these challenges, we incorporate the Greedy Fast Causal Inference (GFCI) methodology, especially when the underlying data distribution is compliant with the requirements of the GES component of GFCI. This hybrid technique tends to perform superiorly, particularly when dealing with a relatively smaller data sample set. However, if the data distribution deviates from the few that suit GES, FCI serves as a more comprehensive alternative. FCI relies solely on conditional independence tests for computing the graph's skeleton, thereby mitigating data requirements.
FCI algorithm is implemented using the R package "Pcalg" on https://cran.r-project.org/web/packages/pcalg/index.html. GFCI algorithm is implemented using R-package "RCausal" on https://github.com/bd2kccd/r-causal.

The pipeline of MCI is illustrated in Figure 4.1. More specifically,

(Inference Engine - Green):

1. For each network, equilibrium state data is generated by randomly selecting values for intrinsic growth rate and the theta vector.
2. The scale and number of data samples guide the choice of the causal inference algorithm. If the dataset is smaller, the GFCI algorithm is preferred over the FCI.
3. After inference, the resultant structure is an inferred Partial Ancestral Graph.
4. Causal information is then extracted from this graph, detailing direct causes and identifiable (non-)cycles.

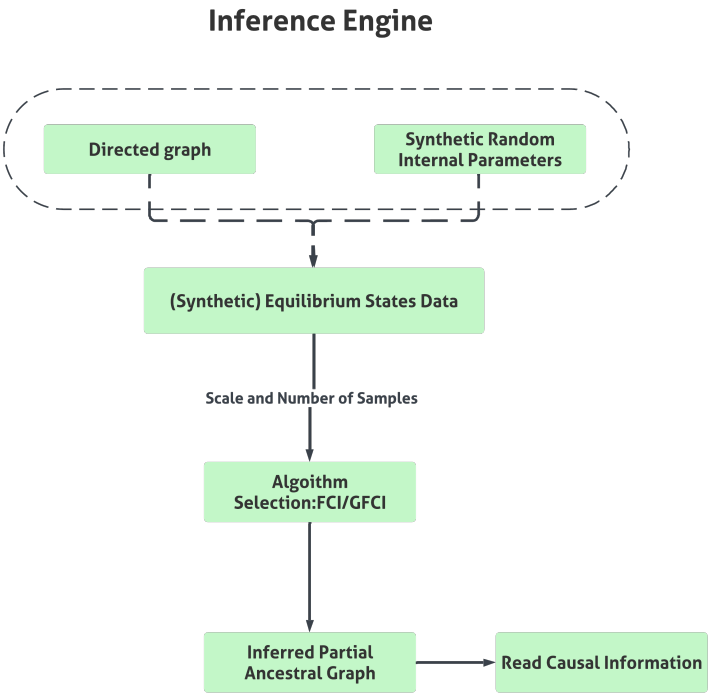


Figure 4.1: MCI pipeline

4.2.5 Generating the Ground Truth

To test the effectiveness of the proposed pipeline, its output needs to be compared with the ground truth which needs to be prepared for test cases.

In the acyclic setting, there can be multiple Directed Acyclic Graphs that encapsulate the same conditional independence relationships. Given the same conditional independence relationship $X_1 \perp X_2 \mid X_3$ for a simple three-variable graph X_1, X_2, X_3, there could be three distinct DAGs that encode the same information. Specifically, the DAGs $X_1 \rightarrow X_3 \rightarrow X_2, X_1 \leftarrow X_3 \leftarrow X_2$, and $X_1 \leftarrow X_3 \rightarrow X_2$ all represent this single conditional independence relationship and, consequently, they form a Markov Equivalence Class. In settings where cycles may occur and latent variables are allowed, we utilize directed mixed graphs and ancestral graph models, as introduced earlier. The underlying network of microbial community interactions, represented as a directed graph, needs to be converted into a form that represents a Markov Equivalence Class, which in this case, is a Partial Ancestral Graph.

The initial step involves applying the process of Acyclification to the ground truth directed graph. This procedure converts the potentially cyclic directed graph into an Acyclic Directed Mixed Graph. Drawing inspiration from the work of (Spirtes, 1994) and (Spirtes, 2013) that introduced the concept of "collapsed graph", we utilize a concept known as acyclification, specifically for a class of graphical causal models referred to as directed graphs with hyperedges (HEDGes), as discussed in (Mooij & Claassen, 2020). We employ this concept in a similar manner for directed mixed graphs.

Given a directed mixed graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$, where $\mathcal{V}$ denotes the vertex set, $\mathcal{E}$ denotes the set of directed edges and $\mathcal{F}$ denotes the set of bi-directed edges. An acyclification of $\mathcal{G}$ is an acyclic directed mixed graph $\mathcal{G}^{acy} = \langle \mathcal{V}^{acy}, \mathcal{E}^{acy}, \mathcal{F}^{acy} \rangle$ with

1. the two mixed graph share the same vertex set; i.e. $\mathcal{V} = \mathcal{V}^{acy}$
2. $i \to j \in \mathcal{E}^{acy}$ if and only if there is a vertex in the strongly connected component of j in graph $\mathcal{V}$ and there is a directed edge $i \to k \in \mathcal{E}$
3. for any vertex i in the strongly connected component of j in graph $\mathcal{V}$, connect $i \leftrightarrow j \in \mathcal{F}^{acy}$.

In simpler terms, rule 2 implies that we should maintain all edges leading into a node that is part of a strongly connected component, and these edges should be duplicated across all nodes within the same component. Rule 3 prescribes that all strongly connected components must be made fully interconnected using bi-directional edges specifically. It is important to note that there could be multiple ways to achieve acyclification as there are various ways to designate edges in order to completely connect strongly connected components, but we specifically choose to use all bi-directional edges in this case. The choice of acyclification does not impact the ground truth ancestral graph, since all acyclifications map back to the unique and consistent ancestral graph model. All these acyclifications maintain σ-separation equivalency to the original directed mixed graph, and their σ-independence models concur, as stated in (Mooij & Claassen, 2020).

An illustration of the acyclification process is presented in diagrams a and b of Figure 4.2, where x_1, x_2, and x_3 constitute a directed cycle in the initial directed mixed graph. From the standpoint of conditional independence, the edge leading from x_4 to x_1 is present in the ground truth, however, when we analyze from the data alone, it becomes impossible to pinpoint the exact node in the cycle that x_4 points to.

In this Figure, part c (highlighted in pink) is responsible for generating the ground truth based on the interaction network present in microbial communities. While part d (in green) is dedicated to inferring the network structure.

(Generating Ground Truth - Pink):

1. We start with a directed graphical model generator that creates directed graphs. These graphs can be either cyclic or acyclic.
2. If a graph is cyclic, an acyclification process is applied to transform it into an acyclic directed mixed graph. An illustrative example of this transformation process is shown.
3. In the presence of latent variables, a latent projection is performed on the ADMG, which results in a marginalized version of the graph.
4. The final step in this segment involves deriving the Markov Equivalence class from the ADMG to produce a Partial Ancestral Graph.

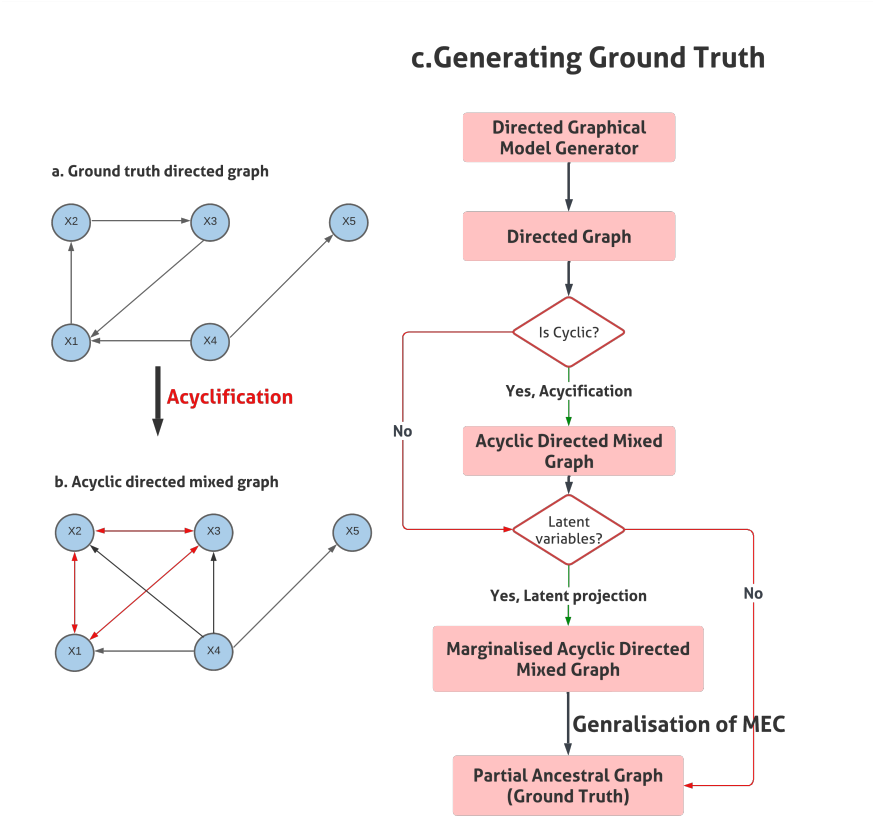


Figure 4.2: Example of acyclification and generating ground truth.

4.2.6 Partially Observed System and Latent Confounding

In real-world microbial communities, it is often the case that we are unable to collect abundance data for all species due to limitations in knowledge or technology. Furthermore, the species we're interested in may not interact directly with each other, but instead, through other unseen or unobserved species. Under such circumstances, it is critical to understand that the data collected might not completely reveal the underlying ecological network. Nevertheless, our method of transforming the network reconstruction problem into a causal graph inference problem can handle this issue. We can marginalize the conditional independence relations of unobserved variables by considering a causal model with latent variables based on the observed marginal distribution. The models we described earlier with hidden variables are appropriately represented by mixed graphs (Richardson, Robins, & Shpitser, 2017). While a mixed graph is normally obtained from a directed graph via latent projection, in our case, we must explicitly define the order of acyclification and latent projection. The process is described in the pipeline in Fig 4.2.

In addition, when an unseen species has simultaneous effects on two other species, this situation is referred to as latent confounding in the literature of causal inference. Latent confounding is a common and pervasive issue in the realm of causal inference. Essentially, a latent confounder is an unobserved factor that is a common cause of both the treatment and response variables.

When we reframe our network reconstruction problem as a task of inferring the corresponding causal graph, a direct impact from species X_i to species X_j would correspond to causal relationships where species X_i influences species X_j. In the results from our chosen causal inference algorithm, such as FCI, the output is presented as a partial ancestral graph. This graph includes bi-directed edges.

When we apply a causal interpretation to the graph, which is reasonable in the context of reconstructing microbial networks, a bi-directed edge between two nodes suggests the presence of at least one latent confounder for those two nodes.

A successful inference of a toy example of latent confounders is shown in Fig 4.3.

More specifically, in an ecosystem, the population of any species could potentially be influenced by the population levels of other species within the ecological network. We generate 200 data points simulating equilibrium states via the GLV mechanism with random internal parameters. The θ-vector and *r*-vector are drawn from a uniform distribution $U[-0.5, 0.5]$. In this network, the species X_5 affects both species X_2 and X_3 and acts as a latent variable. The ground truth directed network is shown in Fig 4.2 (a). We use the Pearson correlation to infer an indirect network of the remaining four species, setting a threshold of 0.5. The correlation plot and the inferred indirect network are displayed in Fig 4.2 (b). We observe that the edge between X_2 and X_3 is missing in this representation. Next, we input the data into our pipeline and select the FCI algorithm with a significance level of 0.05. The inferred PAG matches perfectly with the ground truth PAG, revealing a bi-directed edge between X_2 and X_3. This suggests the presence of a latent confounder.

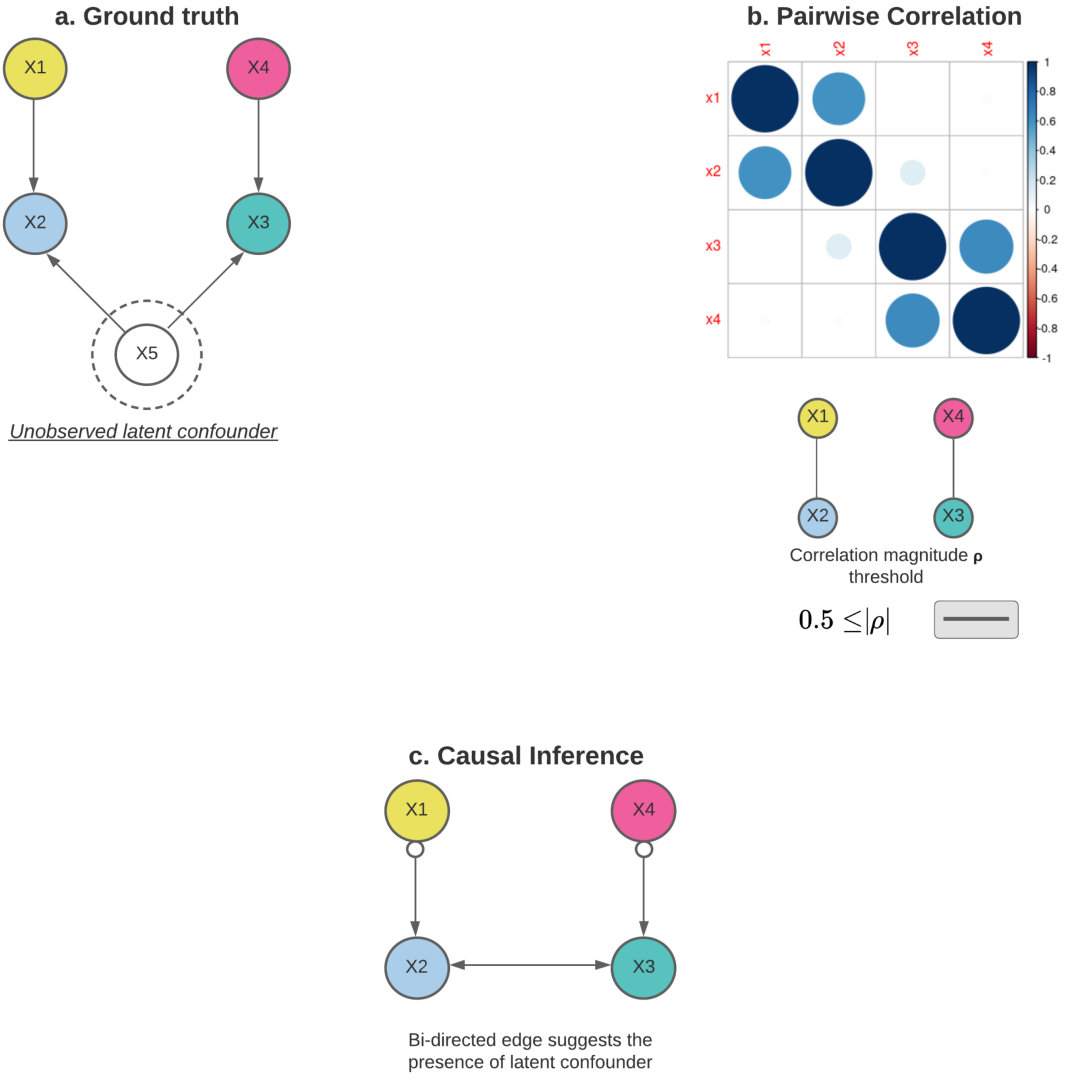


Figure 4.3: Confounder toy example: Microbial Causal Inference vs correlation analysis.

4.2.7 Inference of Direct Causes and Possible Cycles

In acyclic settings, the output PAG by FCI/GFCI represents a Markov equivalence class of MAGs which use the presence and type of edge markings to encode the causal relationships among the observed variables. A directed edge indicated by $X_i \to X_j$ in the graph signifies that X_i causes X_j, while also implying that X_2 is not a cause of X_1 in any possible ground truth Directed Acyclic Graphs. Partial Ancestral Graphs encapsulate the causal characteristics shared by each member of the Markov equivalence class, which includes all Maximal Ancestral Graphs. This is also extended to each Directed Acyclic Graph, potentially containing latent and selection variables, that aligns with the observed independence structure, assuming the condition of faithfulness is met (Colombo, Maathuis, Kalisch, Richardson, et al., 2012).

Complications arise when the underlying network is cyclic. We adopt the identification of direct non-causes and cycles from (Mooij & Claassen, 2020). In a PAG, a directed edge does not unequivocally represent a direct causal relationship. In the inferred PAG, if there's an edge represented as $X_i \leftarrow\!* \ X_j$ or no edge at all between X_i and X_j, it implies that in all possible underlying DMGs, there cannot be a directed edge from X_j to X_i. For discerning potential cycles, let's consider a directed mixed graph $\mathcal{G}$ and its corresponding complete partial ancestral graph ($\mathcal{P}$) output by the FCI algorithm. Let x and y be two distinct nodes in $\mathcal{P}$. If y is part of the strongly connected component of x, it implies that there's a circle edge between x and y ($x \circ\!\!-\!\!\circ y$). Moreover, for all nodes z we have the following property: $z \to x$ in $\mathcal{P}$ if and only if $z \to y$ in $\mathcal{P}$, and $z \leftrightarrow x$ in $\mathcal{P}$ if and only if $z \leftrightarrow y$ in $\mathcal{P}$, and $z \ *\!\!\to x$ in $\mathcal{P}$ if and only if $z \ *\!\!\to y$ in $\mathcal{P}$. This strict definition of strongly connected component of PAG is also necessary. These properties can help identify the potential presence of cyclic structures in the underlying system.

4.3 Results

4.3.1 Metrics

To verify the effectiveness of our approach, we systematically generate simulation data employing Holling Type II functional responses and the Generalised Lotka-Volterra model as foundational interaction mechanisms. As the output graph from our approach is classified as an ancestral graph, we juxtapose our results with the ancestral graph derived from the ground truth directed graph. In this context, we introduce four measurement criteria that we will use for the numerical validation of our method.

The initial metric is edges precision. It is characterized as the fraction of edges where both the presence and edgemarks of the edge (for each kind of edge $\circ\!\!\rightarrow, \circ\!\!-\!\!\circ, \rightarrow, \leftrightarrow$) have been correctly inferred, in relation to the total number of edges in the inferred graph. More precisely:

$$\text{Edges}_{\text{Precision}} = \frac{\text{Number of correctly inferred edges}}{\text{Total number of edges in inferred graph}} \tag{4.10}$$

The second metric we introduce is edges recall. This is calculated as the ratio of edges (for each kind of edge $\circ\!\!\rightarrow, \circ\!\!-\!\!\circ, \rightarrow, \leftrightarrow$) with correctly inferred edge presence and edgemarks to the total number of edges in the ground truth graph. The denominator is the only aspect that differentiates recall from precision. In more detail:

$$\text{Edges}_{\text{Recall}} = \frac{\text{Number of correctly inferred edges}}{\text{Total number of edges in the true graph}} \tag{4.11}$$

Due to the nature of the comparison between two ancestral graphs, the metrics mentioned above are rather stringent since all the edgemarks need to be consistent. Therefore, we consider another less stringent metric which gauges the accuracy of our inferred graph's adjacencies. The adjacencies/skeleton of a graph is an undirected graph derived from removing all edgemarks from the PAG. It is a concept widely utilized in the context of ancestral graphs; precisely we define both preci-

sion and recall as follows:

$$\text{Adjacencies}_{\text{Precision}} = \frac{\text{Number of correctly inferred edges in skeleton}}{\text{Total number of edges in the inferred graph skeleton}} \tag{4.12}$$

$$\text{Adjacencies}_{\text{Recall}} = \frac{\text{Number of correctly inferred edges in skeleton}}{\text{Total number of edges in the true graph skeleton}} \tag{4.13}$$

Simulating equilibrium data samples from ODE systems necessitates the assumption of system stability, a critical consideration for large, complex systems. We adhere to the stability criterion established by May, as detailed in his 1972 study "Will a Large Complex System be Stable?"(May, 1972). This criterion is taken into account when selecting a random community matrix for our simulations. All parameters employed in the subsequent simulation studies are carefully chosen to ensure compliance with this stability criterion. This diligent adherence to May's stability criterion bolsters the validity of our simulation and the reliability of the inferred results. Here's a more concise and structured version of the simulation pipeline:

We begin by generating a graph structure with the largest possible directed cycles denoted as 'k' and an edge is created between two nodes with probability c, which refers to *connectivity*, by visiting each node pair in order. Edge weights are then simulated uniformly within the range of $-\sigma$ to σ for those edges. For each data sample, we simulate a θ-vector uniformly within the range of 0 to θmax, where a θmax of zero indicates a GLV model, while nonzero values indicate Holling Type II functional responses. We simulate a random intrinsic growth vector uniformly within the range of 0 to 1. Each data sample is given a random initial state for the microbial community system, after which ODE computations are performed. We gather equilibrium data samples after a predetermined time span, which we have set to (0,100) for all subsequent simulation studies. The ground truth is generated using our transformations, as described in the "Generating the ground truth" section. Finally, the data samples are processed using two causal inference algorithms, GFCI and FCI, with a significance level of 0.01 selected for our ongoing study.

To facilitate replication and further exploration of our methodology, the codes utilized for this study have been made publicly available. The entire suite of scripts, written in the R programming language, can be accessed through the following GitHub repository: https://github.com/AFFFPupu/Microbial-Causal-Inference. This repository includes all necessary functions and workflows employed in our simulation studies and causal inference analyses, ensuring transparency and enabling other researchers to replicate, validate, and extend our work within their own investigations into microbial community dynamics.

4.3.2 Performance at Different Scales

To evaluate the scalability of our pipeline, we simulated datasets corresponding to networks of different scales, specifically networks with 5, 10, 20, and 50 nodes. For each network size, we generated 20 cases. We simulate the system dynamics using the GLV models introduced in Section 4.2.2, with Equations 4.3 and 4.4. In this model, the parameters θ (interaction saturation) and r (intrinsic growth rate) are drawn from a uniform random distribution, introducing variability in species interaction behaviors and growth rates. This approach reflects the diverse dynamics present in real ecosystems and allows for a comprehensive exploration of potential interaction and growth patterns. The performance and adaptability of our pipeline across these varying scales were then subjected to a comprehensive scalability analysis. The parameter setting for our scalability analysis is given in Table 4.1 and the evaluation is given in Figure 4.4. Our observations suggest that as the complexity and size of the system increase, the accuracy of recovery generally declines. Despite this, for adjacency metrics, the accuracy levels remain satisfactory, considering that the number of data samples does not scale with the number of nodes. When it comes to the more stringent edge metrics, GFCI tends to outperform FCI. Notably, we do not see a significant difference in performance between the GLV and Holling Type II models which is a bit surprising since the latter represents a more complex dynamics.

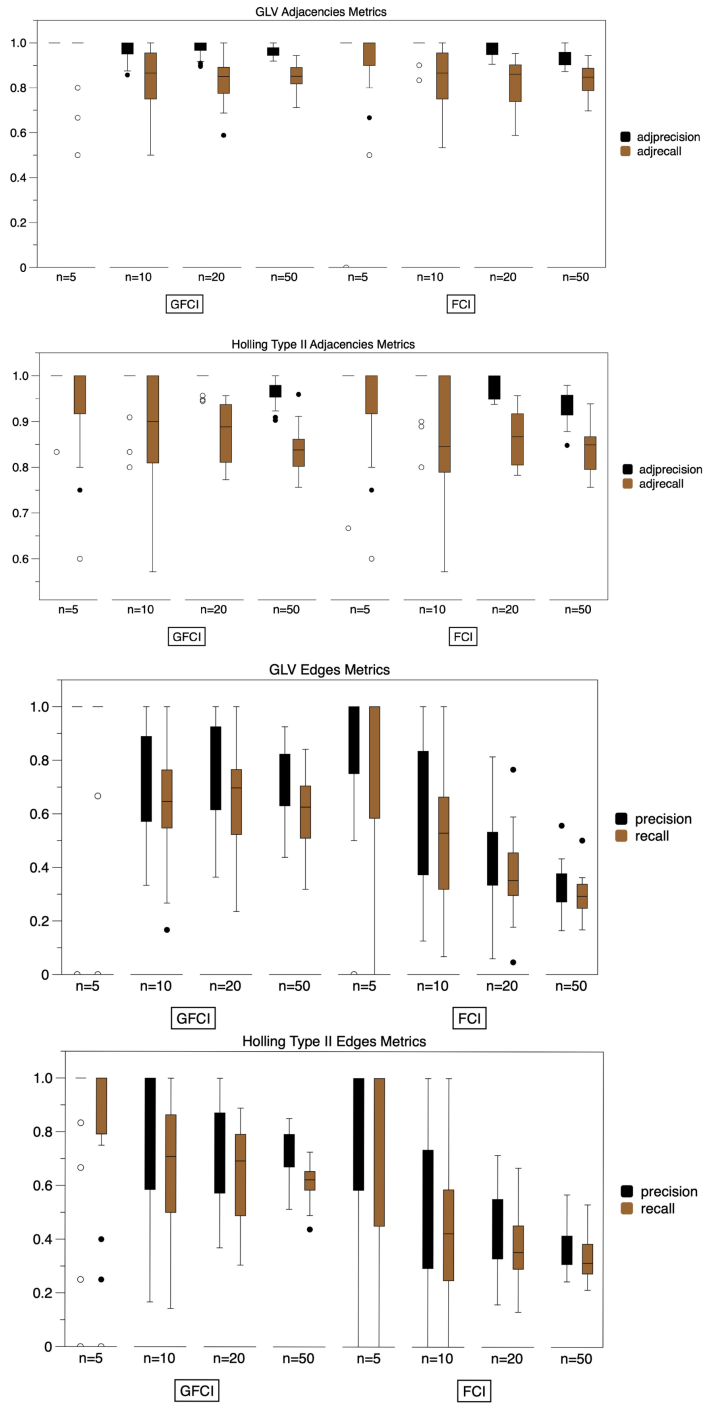


Figure 4.4: Scalability analysis of our pipeline MCI.

number of species	connectivity	θ_{max}	number of data points	σ
n=5	c=0.4	0 for GLV/ 0.2 for Holling Type II	1000	1.2
n=10	c=0.2	0 for GLV/ 0.2 for Holling Type II	1000	1.2
n=20	c=0.1	0 for GLV/ 0.2 for Holling Type II	1000	1.2
n=50	c=0.04	0 for GLV/ 0.2 for Holling Type II	1000	1.2

The connectivity of the system, defined as the ratio of the number of edges to the number of nodes, was deliberately set to adjust with the size of the system. This was done to maintain a consistent connectivity ratio irrespective of the number of species in the system.

Table 4.1: Parameters selection for scalability analysis.

4.3.3 Influence of Sample Sizes

In order to demonstrate the performance of our pipeline with varying amounts of available data samples, we conduct a comprehensive test using the same GLV models introduced in Section 4.2.2 with Equations 4.3 and 4.4. The parameter settings, given in Table 4.2, include θ_{max}, where θ is drawn from a uniform distribution as $\theta = \text{runif}(N) \times \theta_{\text{max}}$. This introduces variability in species interactions, reflecting the diverse behaviors observed in real ecosystems. The test results are shown in Figure 4.5. It is evident that the performance of both metrics improves significantly as the number of available data samples increases. Notably, GFCI consistently outperforms FCI, particularly when the number of data points is limited, as seen in the edge metrics. This observation aligns with the known characteristics of GFCI, given that it is a variant of FCI.

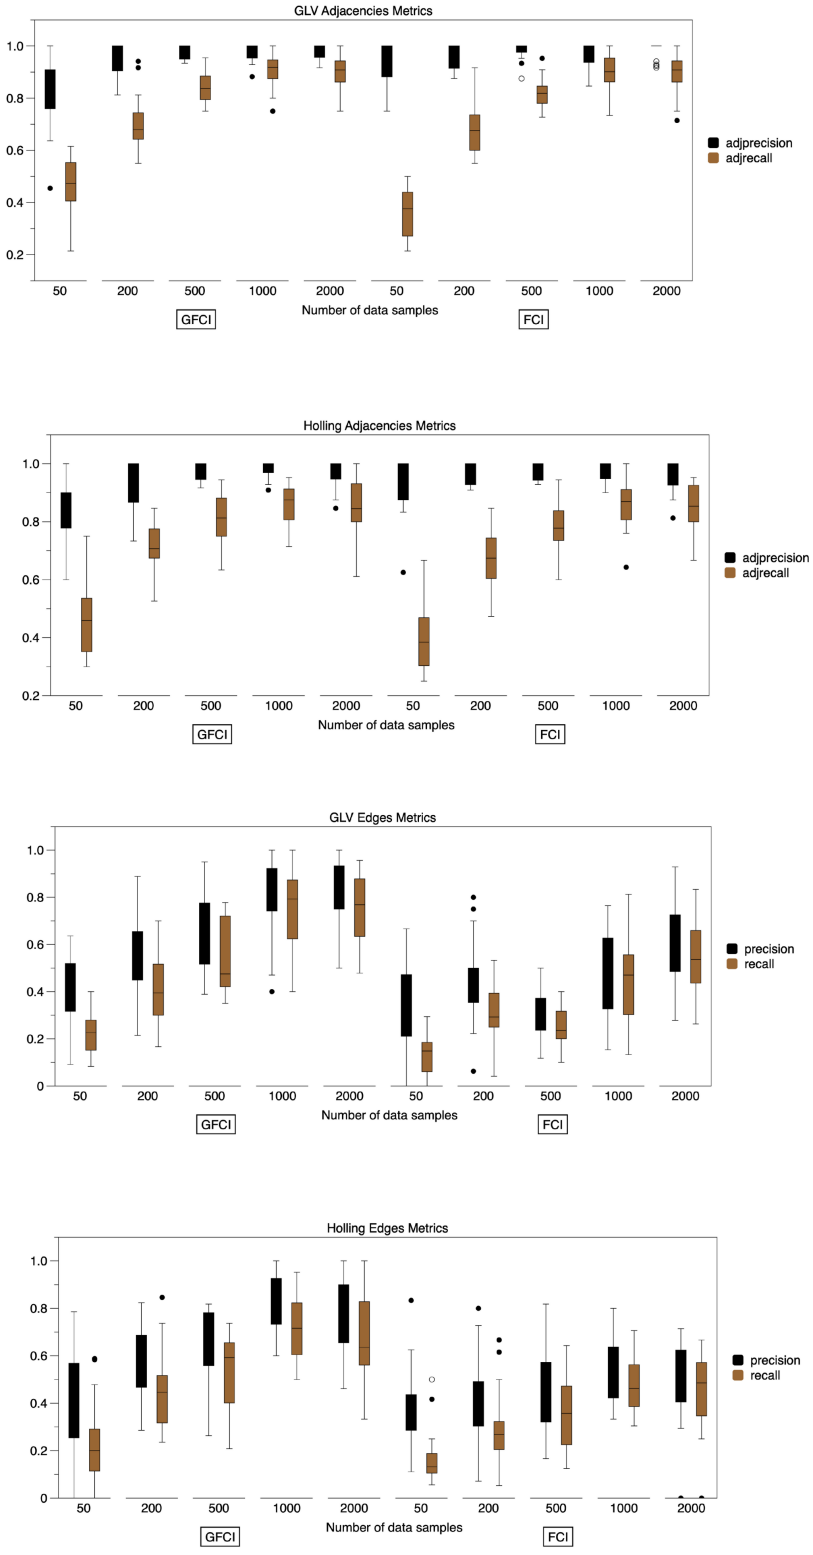


Figure 4.5: Sample size analysis of our pipeline MCI.

number of species	connectivity	θ_{max}	number of data points	σ
$n = 20$	$c = 0.1$	0 for GLV/ 0.2 for Holling Type II	50	1.2
$n = 20$	$c = 0.1$	0 for GLV/ 0.2 for Holling Type II	200	1.2
$n = 20$	$c = 0.1$	0 for GLV/ 0.2 for Holling Type II	500	1.2
$n = 20$	$c = 0.1$	0 for GLV/ 0.2 for Holling Type II	1000	1.2
$n = 20$	$c = 0.1$	0 for GLV/ 0.2 for Holling Type II	2000	1.2

All other parameters are held constant as we conduct our analysis on a reasonably sized system comprising of approximately 20 nodes and 20 edges.

Table 4.2: Parameters selection for sample size analysis.

4.3.4 Comparison of Methods

As proposed in (Xiao et al., 2017), a method was developed to leverage steady state data for mapping the ecological networks of microbial communities. Notably, this is the sole method that uses a similar type of data as our approach. We have carried out a systematic comparison between our method and theirs across different scales of species numbers in microbial communities. It is important to note that the numerical comparison was conducted using the same set of steady state data, generated using Holling Type II as the population dynamics model.

In their method, steady state data was collected from the same microbial system with fixed parameters. This type of data collection ensures that the difference between two corresponding sample pairs would generate a hyperplane orthogonal to it. The intersection of all these hyperplanes would then reveal the topology of the underlying network. However, when data comes from microbial communities with possibly distinct parameters, such as different intrinsic growth for each species, the difference between two steady state samples could result from the parameters' differences rather than the underlying interaction matrix. This assumption is not made in our method, and steady state data samples are only assumed to come from systems with the same interaction mechanism between microbial species. In the comparison, we will refer to their method as the Xiao's method.

The results, which we use the data from simulations suggested by Xiao's method, displayed in Figure 4.6, show that the performance of both methods declines as the number of species in the microbial communities increases. Here, the simulation of steady states data by manually setting the abundances of species to 0 in order. The brute force version (Xiao et al., 2017) of Xiao's method was used when the number of species was 5 or 10, and the heuristic version was employed when the species count was 15. Our method performs competitively in terms of precision, primarily because Xiao's method tends to produce more uncertain entries in the inferred interaction matrix. This uncertainty stems from the difficulty of finding the intersection among orthogonal hyperplanes.

However, it is important to clarify that while we've used comparable metrics to compare the two methods, our approach evaluates the match between two Partial Ancestral Graphs. These inherently represent some uncertainty due to the loss of information when transitioning from directed graphs to ancestral graphs. Conversely, Xiao's method's performance was evaluated between definite graphs, which could be considered more stringent to some extent.

Both our method and Xiao's method utilize data samples from systems at equilibrium. However, there are some key differences in the types of data each method can handle. Specifically, Xiao's method uses steady state data samples from a microbial community system with a single, fixed parameter: the intrinsic growth vector (the θ vector).

Our approach, on the other hand, is not restricted to data samples from systems with a fixed intrinsic growth vector. Our method can handle data from systems where the interaction matrix remains consistent, regardless of variations in internal parameters. This is significant as it allows our approach to deal with data from a wider range of realistic scenarios compared to Xiao's method.

To illustrate this, consider a system with the following dynamics:

$$\dot{x}_1 = x_1 (r_1 - x_1 + x_2) \tag{4.14}$$

$$\dot{x}_2 = x_2 (r_2 - x_2) \tag{4.15}$$

This system has three steady states: $(0, r_2)$, $(r_1, 0)$, and $(r_1 + r_2, r_2)$. Now, if we have a set of parameters where $r_1 = r_2 = 0.5$, we would have a steady state of $(0.5, 0)$. If we then change the parameters to $r_1 = r_2 = 0.25$, we end up with a steady state of $(0.25, 0.25)$. Although these steady states originate from the same interaction matrix, they have different internal parameters and would lead to incorrect conclusions if analyzed using Xiao's method.

Our method's ability to handle this type of variation in data makes it more versatile and applicable to a wider range of real-world situations.

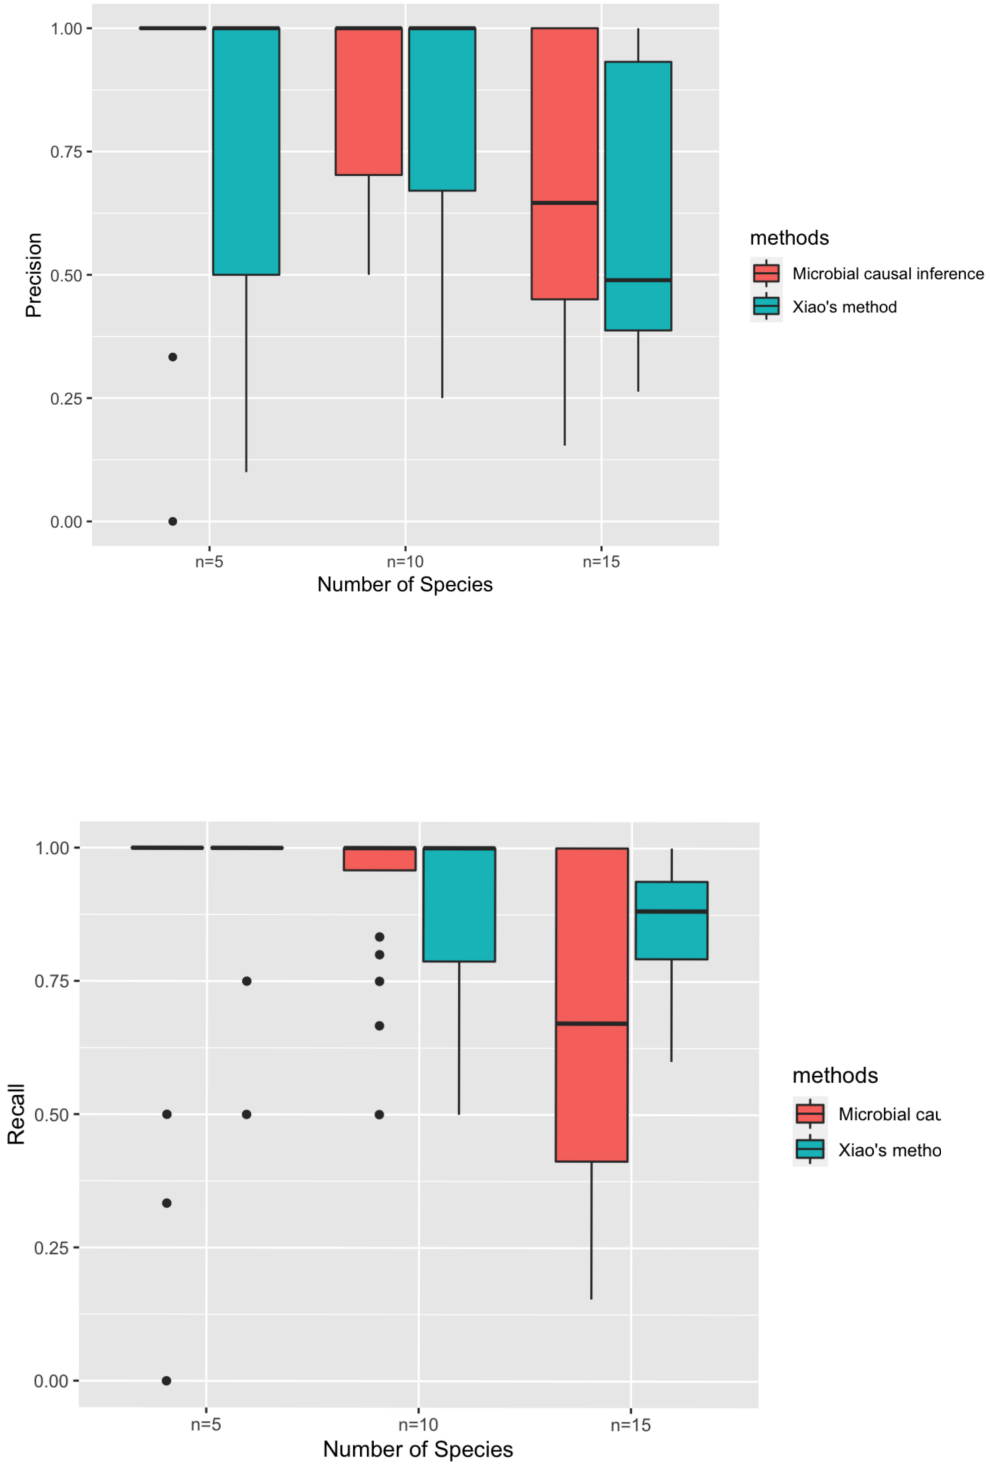


Figure 4.6: Comparison between methods: precision and recall for number of species 5,10,15

4.4 Conclusion

In this study, we presented a novel pipeline for reconstructing ecological networks within microbial communities, utilizing steady-state data. Our approach is uniquely capable of processing steady-state data from systems with varying parameters as long as the underlying ajacency matrix is qualitatively retained.

This pipeline has been crafted within the context of causal graphs, which suits well the interconnected nature of microbial communities under equilibrium conditions. We've successfully demonstrated the functionality of our pipeline using the population dynamics of the Generalized Lotka-Volterra and Holling Type II functional response models. Furthermore, our pipeline exhibits capabilities for detecting latent confounders and handling partially observed systems effectively.

A significant advantage of our pipeline is its compatibility with steady-state data derived from systems with different internal parameters. This is highly relevant in real-world scenarios where data could originate from systems with diverse internal parameters, such as varying health indices among patients. Being a purely statistical approach, our pipeline showcases the ability to scale, though it generally demands a larger sample size to reliably establish conditional dependencies and independencies.

Our pipeline outputs an ancestral graph, providing a more comprehensive perspective of the network than a simple directed network. The pipeline maintains impressive accuracy in recovering network skeletons, even in scenarios with a high number of nodes and dense networks. However, when considering edge marks, as indicated by the precision and recall metrics, there's potential for future improvements in delivering more accurate orientations, especially in the presence of directed cycles.

Our work brings together the fast-developing field of causal inference and microbial community analysis, opening up new possibilities for network reconstruction within microbial communities. The conventional focus of many causal inference pipelines on directed acyclic graphs has been a limitation when applied to scientific subjects like microbial communities, where feedback loops are common. By situating our

network reconstruction problem within the realm of causal graphs, we expand the applicability of such pipelines. As evidenced by our use of Greedy Fast Causal Inference (GFCI) and Fast Causal Inference (FCI), our work sets the stage for the deployment of more score-based pipelines for network reconstruction within microbial communities, heralding promising opportunities for further advancements in the field.

References

Ali, R. A., Richardson, T. S., & Spirtes, P. (2009). Markov Equivalence for Ancestral Graphs. *Ann. Statist.*.

Angulo, M. T., Moreno, J. A., Lippner, G., Barab'asi, A.-L., & Liu, Y.-Y. (2017). Fundamental Limitations of Network Reconstruction From Temporal Data. *Journal of the Royal Society Interface*, *14*(127), 20160966.

Bongers, S., Forr'e, P., Peters, J., & Mooij, J. M. (2021). Foundations of Structural Causal Models With Cycles and Latent Variables. *The Annals of Statistics*, *49*(5), 2885–2915.

Bongers, S., & Mooij, J. M. (2018). From Random Differential Equations to Structural Causal Models: The Stochastic Case. *ArXiv Preprint ArXiv:1803.08784*.

Cao, H.-T., Gibson, T. E., Bashan, A., & Liu, Y.-Y. (2017). Inferring Human Microbial Dynamics From Temporal Metagenomics Data: Pitfalls and Lessons. *BioEssays*, *39*(2), 1600188.

Chickering, D. M. (2002). Optimal Structure Identification With Greedy Search. *Journal of Machine Learning Research*, *3*(Nov), 507–554.

Chickering, D. M. (2003, mar). Optimal Structure Identification with Greedy Search. *J. Mach. Learn. Res.*, *3*(null), 507â554. Retrieved from https://doi.org/10.1162/153244303321897717 doi: 10.1162/153244303321897717

Colombo, D., Maathuis, M. H., Kalisch, M., Richardson, T. S., et al. (2012). Learning High-Dimensional Directed Acyclic Graphs With Latent and Selection Variables. *The Annals of Statistics*, *40*(1), 294–321. doi: 10.1214/11-AOS940

Consortium, H. M. P., et al. (2012). Structure, Function and Diversity of the Healthy Human Microbiome. *Nature*, *486*(7402), 207.

Deng, Y., Jiang, Y.-H., Yang, Y., He, Z., Luo, F., & Zhou, J. (2012). Molecular Ecological Network Analyses. *BMC Bioinformatics*, *13*, 1–20.

Faust, K., & Raes, J. (2012). Microbial Interactions: From Networks to Models.

Nature Reviews Microbiology, *10*(8), 538–550.

Friedman, J., & Alm, E. J. (2012). Inferring Correlation Networks From Genomic Survey Data.

Heckerman, D., Geiger, D., & Chickering, D. M. (1995). Learning Bayesian Networks: The Combination of Knowledge and Statistical Data. *Machine Learning*, *20*(3), 197–243.

Hyttinen, A., Eberhardt, F., & Hoyer, P. O. (2012). Learning Linear Cyclic Causal Models With Latent Variables. *Journal of Machine Learning Research*, *13*(Nov), 3387–3439.

Jordan, M. I. (2003). *An Introduction to Probabilistic Graphical Models.* preparation.

Kurtz, Z. D., M"uller, C. L., Miraldi, E. R., Littman, D. R., Blaser, M. J., & Bonneau, R. A. (2015). Sparse and Compositionally Robust Inference of Microbial Ecological Networks. *PLoS Computational Biology*, *11*(5), e1004226.

Lozupone, C. A., Stombaugh, J. I., Gordon, J. I., Jansson, J. K., & Knight, R. (2012). Diversity, Stability and Resilience of the Human Gut Microbiota. *Nature*, *489*(7415), 220–230.

May, R. M. (1972). Will a Large Complex System Be Stable? *Nature*, *238*(5364), 413–414.

Mooij, J. M., & Claassen, T. (2020). Constraint-based Causal Discovery Using Partial Ancestral Graphs in the Presence of Cycles. In *Conference on Uncertainty in Artificial Intelligence* (pp. 1159–1168).

Mooij, J. M., Janzing, D., & Sch"olkopf, B. (2013). From Ordinary Differential Equations to Structural Causal Models: the Deterministic Case. *ArXiv Preprint ArXiv:1304.7920*.

Ogarrio, J. M., Spirtes, P., & Ramsey, J. (2016). A Hybrid Causal Search Algorithm for Latent Variable Models. In *Conference on Probabilistic Graphical Models* (pp. 368–379).

Pearl, J. (2009). *Causality*. Cambridge University Press.

Pearl, J. (2013). *A Constraint Propagation Approach to Probabilistic Reasoning*.

Pulz, O., & Gross, W. (2004). Valuable Products From Biotechnology of Microal-

gae. *Applied Microbiology and Biotechnology*, *65*(6), 635–648.

Richardson, Robins, J. M., & Shpitser, I. (2017). Nested Markov Properties for Acyclic Directed Mixed Graphs. *ArXiv Preprint ArXiv:1701.06686*.

Richardson, & Spirtes. (2002). Ancestral Graph Markov Models. *Ann. Statist.*, *30*, 962–1030.

Song, H., Ding, M.-Z., Jia, X.-Q., Ma, Q., & Yuan, Y.-J. (2014). Synthetic Microbial Consortia: From Systematic Analysis to Construction and Applications. *Chemical Society Reviews*, *43*(20), 6954–6981.

Spirtes. (1994). Conditional Independence in Directed Cyclic Graphical Models for Feedback.

Spirtes. (2013). Directed Cyclic Graphical Representations of Feedback Models. *ArXiv Preprint ArXiv:1302.4982*.

Spirtes, Glymour, C., & Scheines, R. (2000). *Causation, Prediction, and Search* (2nd ed.). MIT press.

Spirtes, P. (1995). Directed Cyclic Graphical Representations of Feedback Models. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* (pp. 491–498).

Sunagawa, S., Coelho, L. P., Chaffron, S., Kultima, J. R., Labadie, K., Salazar, G., ... Velayoudon, D. (2015). Structure and function of the global ocean microbiome. *Science*, *348*. doi: 10.1126/science.1261359

Verma, T. S., & Pearl, J. (2022). Equivalence and Synthesis of Causal Models. In *Probabilistic and Causal Inference: The Works of Judea Pearl* (pp. 221–236).

Wade, M., Pattinson, R., Parker, N., & Dolfing, J. (2016). Emergent Behaviour in a Chlorophenol-mineralising Three-tiered Microbial food Web. *Journal of Theoretical Biology*, *389*, 171–186.

Whittaker, J. (2009). *Graphical Models in Applied Multivariate Statistics*. Wiley Publishing.

Xiao, Y., Angulo, M. T., Friedman, J., Waldor, M. K., Weiss, S. T., & Liu, Y.-Y. (2017). Mapping the Ecological Networks of Microbial Communities. *Nature Communications*, *8*(1), 1–12.

Zhang, J. (2008). On the Completeness of Orientation Rules for Causal Discovery

in the Presence of Latent Confounders and Selection Bias. *Artificial Intelligence*, *172*(16-17), 1873–1896.

Zhou, W., Chen, P., Min, M., Ma, X., Wang, J., Griffith, R., ... others (2014). Environment-enhancing Algal Biofuel Production Using Wastewaters. *Renewable and Sustainable Energy Reviews*, *36*, 256–269.

5 | Adaptive Learning and Experimental Design with Interventional Data

Abstract

In the exploration of dynamical systems such as microbial communities, the challenge lies in inferring the underlying causal structures from observational equilibrium data. This task is complicated by the presence of feedback loops, which obscure direct causal relationships. Partial Ancestral Graphs (PAGs) have emerged as a tool to address this issue, representing a specific Markov equivalence class of graphs capable of capturing causality within these complex systems. However, a significant gap remains in determining the complete causal structure, particularly in cyclic scenarios, as PAGs alone often leave edgemarks undetermined. Our work aims to bridge this gap by proposing a set of rules and algorithms tailored for designing effective experiments to gain interventional data. These experiments are crucial for transforming the PAG-inferred structure into a fully identified causal model. The novelty of our work lies in the development of these experimental design rules and algorithms that systematically convert PAGs into fully identified causal models, even in the presence of feedback loops. This adaptive learning approach not only sheds light on the intricate causal networks of microbial communities but also guides the experimental process to confirm or refine inferred causal relationships. Our methodology promises to enhance our understanding of microbial dynamics and offers a systematic path for future research in causal inference within complex biological systems.

5.1 Introduction

In numerous scientific disciplines, a fundamental objective is to elucidate causal relationships among variables within a given system (Pearl, 2009). These relationships are typically represented by a directed graph, where each directed edge symbolizes a direct causal link. For instance, a direct edge from variable X_1 to X_2

implies that X_1 is a direct cause of X_2 within the system. This concept is particularly relevant in the study of microbial communities, where the direct causation can be interpreted as the influence of one species on another.

In the preceding chapter, we tailored the concept of structural causal models to frame microbial systems, culminating in the development of Microbial Causal Inference. This framework is designed for inferring microbial interactions, accommodating potentially non-linear system dynamics and the presence of feedback loops. The output of our algorithm is a class of Markov equivalent graphs. Given that our approach is grounded in observational data alone, it inevitably faces limitations. Specifically, observational data alone cannot distinguish between multiple observationally equivalent structures, as highlighted in (Claassen & Mooij, 2023). To unravel the unique underlying directed graph that fully characterizes the system's interactions, the incorporation of interventional data is essential.

In this chapter, we tackle critical questions that arise from these constraints: (1) How can we accurately and fully recover the true structure of a system, especially when underlying feedback loops might introduce virtual edges, i.e. the edges created by the presence of cycles in the Markov Equivalence Class of observational data? (2) How do we determine the variant edgemarks within a Partial Ancestral Graph? To address these challenges, we will introduce an optimal adaptive learning algorithm, complemented by intervention experiments. This algorithm and its systematic validation with synthetic data represent a significant stride towards an improved approach to attain a more nuanced understanding of microbial interactions.

Interventional data are obtained by conducting experiments within a system, specifically targeting a subset of variables/nodes to observe the effects of such interventions. This involves comparing the distribution of relevant variables before and after the intervention. For directed acyclic graphs (DAGs), strategies for designing experiments to uniquely identify a causal graph from its Markov equivalence class have been thoroughly explored. In (Hauser & B"uhlmann, 2014), two optimal strategies for the active learning of causal models from interventional data are pro-

posed, tailored to different learning objectives. These include a batch-intervention design and a sequential-intervention design for an active learning approach, which aims to orient undirected edges in each chain component via separate intervention experiments (He & Geng, 2008). Furthermore, (Shanmugam, Kocaoglu, Dimakis, & Vishwanath, 2015) introduced a lower bound on the number of experiments required for adaptive methods, based on the concept of separating systems. However, these methods are primarily developed under the assumption of acyclicity, a common simplification in causal literature.

Studies have also considered experimental designs that account for potential cycles or feedback loops. (Shanmugam et al., 2015) recently proposed a unified experimental design approach applicable to both cyclic and acyclic causal models, extending the bounded-size experimental design from acyclic models. Hyttinen et al. developed a general Boolean satisfiability (SAT)-based solver for inferring cyclic causal models with latent variables (Hyttinen, Hoyer, Eberhardt, & Jarvisalo, 2013). For linear cyclic models with latent variables, a search algorithm has been proposed and validated using both simulated and real data (Hyttinen, Eberhardt, & Hoyer, 2012). To address conflicts arising from statistical errors in multiple interventions over overlapping variable sets, the COmbINE approach was proposed (Triantafillou & Tsamardinos, 2015).

In 2020, Mooij and Claassen demonstrated that the Fast Causal Inference (FCI) algorithm, originally designed for learning DAGs with latent variables, is also effective for graphs with cycles (Mooij & Claassen, 2020) and a joint causal inference framework has been proposed to include data from different context. (Meganck et al., 2006) presented a set of rules to guide experiments needed to transform a PAG into a completely oriented Semi-Markovian causal model (SMCM), although these rules do not consider cycles.

In this chapter, we address a notable research gap: refining the skeleton and further orienting the edgemarks in Partial Ancestral Graphs derived from potentially cyclic causal graphs. We introduce a novel two-stage algorithm to bridge this gap. Initially, our algorithm employs a minimal set covering technique to strategically

select intervention targets. This approach is designed to effectively disrupt cycles involving at least three nodes. Subsequently, we meticulously design experiments aimed at resolving the remaining undetermined edgemarks, with particular attention to two-node cycles. Our methodology rests on two fundamental assumptions: firstly, *causal sufficiency*, implying the absence of unmeasured common causes, or latent confounders, that might influence the measured variables; and secondly, the *faithfulness* condition, which we assume to be valid for both observational and post-interventional distributions. This assumption ensures that alterations in parameters representing quantitative relationships do not invalidate any observed independencies.

This chapter is organized as follows: We commence with introducing the concepts and theorems underpinning our methodology. The ensuing section provides necessary preliminaries, introducing notations and assumptions, and formally delineates our problem setting within the specified context. Subsequently, we present our algorithm in two distinct stages. In this segment, we delve into the intricacies of each stage, addressing the challenges we aim to surmount. Following this, we introduce an ideal version of our algorithm, wherein the accuracy of the Fast Causal Inference (FCI) and Greedy Fast Causal Inference (GFCI) algorithms is presumed. Our algorithm is then rigorously analyzed using synthetic data reflective of microbial communities, particularly focusing on scenarios with potential non-linear dynamics. The final part of this chapter encompasses a discussion where we contemplate the broader implications, inherent limitations, and prospective future directions of our research.

5.2 Problem Setting

5.2.1 Preliminaries in Graphs

In the preceding chapter, we introduced several fundamental preliminaries and notations pivotal to our study. In this chapter, we will revisit and emphasize those concepts that are particularly relevant to our ongoing discourse. Additionally, we

will present a series of essential definitions, theorems, and propositions that form the cornerstone of our work and are instrumental in advancing our analysis.

In graph theory, a directed graph (DG) is formally denoted as $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$. This structure comprises a vertex set $\mathcal{V} = \{X_1, \ldots, X_p\}$ and a directed edge set $\mathcal{E}$. A directed edge $(X_i, X_j) \in \mathcal{E}$ is depicted as $X_i \to X_j$ or $X_j \leftarrow X_i$, with X_i being designated as a *parent* of X_j. The collection of all parents of X_j in $\mathcal{G}$ is represented by $\mathrm{Pa}_{\mathcal{G}}(X_j) := \{X_i \in \mathcal{V} : X_i \to X_j \in \mathcal{E}\}$.

A *walk* between two nodes $X_{i_0}, X_{i_n} \in \mathcal{V}$ is defined as a sequence $\langle X_{i_0}, e_1, X_{i_1}, e_2, X_{i_2}, \ldots, e_n, X_{i_n} \rangle$, alternating between nodes and edges in $\mathcal{G}$. Each vertex in this walk is distinct, ensuring that the sequence contains each vertex at most once. A *directed path* in this context is a path where every edge e_k is a directed edge from $X_{i_{k-1}}$ to X_{i_k}, i.e., $X_{i_{k-1}} \to X_{i_k}$. A *directed loop* is a directed path starting from X_i and ending at X_j, inclusive of an additional edge $X_j \to X_i \in \mathcal{E}$, thus forming a *directed cycle*.

The concept of *ancestors* in a DG extends to tracing back the lineage of parents repeatedly, formalized as $\mathrm{An}_{\mathcal{G}}(X_j) := \{X_i \in \mathcal{V} : X_i = X_{i_0} \to X_{i_1} \to \cdots \to X_{i_n} = X_j \text{ in } \mathcal{G}\}$. Conversely, the *descendants* of X_i are defined as $\mathrm{De}_{\mathcal{G}}(X_i) := \{X_j \in \mathcal{V} : X_i = X_{i_0} \to X_{i_1} \to \cdots \to X_{i_n} = X_j \text{ in } \mathcal{G}\}$. Nodes that are part of directed cycles involving $X_i \in \mathcal{V}$ constitute the *strongly connected component*, denoted as $\mathrm{Sc}_{\mathcal{G}}(X_i) := \mathrm{An}_{\mathcal{G}}(X_i) \cap \mathrm{De}_{\mathcal{G}}(X_i)$. The *skeleton* of a DG is an undirected graph $\mathcal{G}' = \langle \mathcal{V}, \mathcal{E}' \rangle$, where an undirected edge $(X_1, X_2) \in \mathcal{E}'$ exists if either (X_1, X_2) or (X_2, X_1) is in $\mathcal{E}$. In an undirected graph, a *clique* refers to a subset of vertices such that every two distinct vertices within it are adjacent, forming an induced subgraph that is complete. A vertex X in the middle of a path within a causal graph is identified as a collider if the edges connected to X on the path direct towards it, marked by arrowheads pointing at X.

5.2.2 From ODEs to SCMs

In this chapter, the concepts of directed mixed graphs (DMG) and directed graphs will be frequently utilized. The primary distinction between these two lies in the

inclusion of bidirected edges within DMGs, representing the presence of latent variables. However, under the assumption of causal sufficiency in our work, where we consider microbial communities to be fully observable, the presence of latent variables is negated. Consequently, prior to the process of acyclification, which was introduced in Chapter 4, DMGs and DGs are essentially equivalent in our context.

The application of DMGs in our methodology is crucial, particularly as it aligns with the approach suggested in (Mooij & Claassen, 2020). In the context of a DMG, two nodes are defined as belonging to the same *strongly connected component* if and only if each node is an ancestor of the other. This mutual ancestral relationship establishes the criteria for nodes to be considered as part of the same strongly connected component within the DMG framework. This approach involves representing a (possibly cyclic) DMG by a directed partial ancestral graph. We aim to clarify this distinction to prevent any potential confusion or misuse of these terms.

The generative framework of our study is structured around the dynamic behavior of microbial communities, which can be effectively captured through a series of Ordinary Differential Equations (ODEs). In our investigation, we accommodate nonlinear population dynamics, a detailed exposition of which is provided in sections 4.2.2 and 4.2.3 of the preceding chapter. To model the causal architecture of these microbial systems, we employ Structural Causal Models (SCMs), a widely recognized approach for delineating causal processes within complex systems (see (Pearl et al., 2000)). Our methodology is particularly influenced by the advancements detailed in (Bongers, Forr'e, Peters, & Mooij, 2021), which offer robust solutions for addressing the intricate challenges associated with cyclic dependencies within these models. This approach enables a comprehensive analysis of the causal structures inherent in microbial community dynamics, ensuring a nuanced understanding of the underlying biological processes.

Definition 5.1. *A Structural Causal Model (SCM) is a tuple* $\mathcal{M}$*:*

$$\mathcal{M} := \langle \mathcal{U}, \mathcal{V}, \boldsymbol{\mathcal{X}}_{\mathcal{U}}, \boldsymbol{\mathcal{X}}_{\mathcal{V}}, \boldsymbol{F} \rangle \tag{5.1}$$

where:

- $\mathcal{U}$ *is a finite set indexing the endogenous variables,*
- $\mathcal{V}$ *is a finite set indexing the exogenous variables, disjoint from* $\mathcal{U}$*,*
- $\boldsymbol{\mathcal{X}}_{\mathcal{U}} = \prod_{u \in \mathcal{U}} \mathcal{X}_u$*, is the product of the domains of the* $\mathcal{U}$*, and each* $\mathcal{X}_u$ *is a standard measurable space,*
- $\boldsymbol{\mathcal{X}}_{\mathcal{V}} = \prod_{v \in \mathcal{V}} \mathcal{X}_v$*, is the product of the domains of the* $\mathcal{V}$*, and each* $\mathcal{X}_v$ *is also a standard measurable space,*
- $\boldsymbol{F} : \boldsymbol{\mathcal{X}}_{\mathcal{U}} \times \boldsymbol{\mathcal{X}}_{\mathcal{V}} \to \boldsymbol{\mathcal{X}}_{\mathcal{U}}$ *is a measurable causal function defining the relationships in the model, i.e. the causal mechanism*

Definition 5.2. *The* $\boldsymbol{F}$ *is specified by a set of equations called Structural Equations:*

$$x_i = f_i(\boldsymbol{x}, \boldsymbol{v}), \quad \boldsymbol{x} \in \boldsymbol{\mathcal{X}}_{\mathcal{U}}, \quad \boldsymbol{v} \in \boldsymbol{\mathcal{X}}_{\mathcal{V}} \tag{5.2}$$

where $i \in \mathcal{U}$*.*

A Structural Causal Model, denoted as $\mathcal{M}$, inherently gives rise to a directed mixed graph, symbolized as $\mathcal{G}(\mathcal{M})$. This graph is constituted by a vertex set $\mathcal{U}$, representing the endogenous variables, and an edge set $\mathcal{E}$. The edges in $\mathcal{E}$ include directed edges, indicated by $i \to j$, if and only if i and j are both elements of $\mathcal{U}$ and i acts as a parent to j. Additionally, bidirected edges, represented as $i \leftrightarrow j$, are included if and only if i and j are within $\mathcal{U}$, and there exists a common parent $v \in \mathcal{V}$ for both i and j. Under the assumption of causal sufficiency, where no latent confounders are present, this directed mixed graph simplifies to a directed graph.

In this chapter, we only consider a subclass of SCM which is simple SCM and a SCM $\mathcal{M}$ is called simple if it is uniquely solvable with respect to every subset $U \subset \mathcal{U}$.

5.2.3 Interventions

The intervention is both well-defined in our ODEs dynamics system and its corresponding generative structural causal models. We first define the perfect intervention which is introduced by (Pearl, 2009) of our SCMs here.

Definition 5.3. *The perfect intervention on an SCM is defined as follows:*

Let $\mathcal{M} = \langle \mathcal{U}, \mathcal{V}, \boldsymbol{\mathcal{X}}_{\mathcal{U}}, \boldsymbol{\mathcal{X}}_{\mathcal{V}}, \boldsymbol{F} \rangle$ *be an SCM,* $U \in \mathcal{U}$ *a subset and* $\xi_U \in \mathcal{X}_U$ *a fixed vector value. The perfect intervention is denoted as* $do(U, \xi_U)$ *and it is defined by a mapping from* $\mathcal{M}$ *to* $\mathcal{M}_{do(U,\xi_U)} := \mathcal{M} := \langle \mathcal{U}, \mathcal{V}, \boldsymbol{\mathcal{X}}_{\mathcal{U}}, \boldsymbol{\mathcal{X}}_{\mathcal{V}}, \boldsymbol{F'} \rangle$ *where* $\boldsymbol{F'}$ *is defined as*

$$f'_u(\boldsymbol{x}, \boldsymbol{v}) = \begin{cases} \xi_u, & u \in U, \\ f_u(\boldsymbol{x}, \boldsymbol{v}), & u \in \mathcal{U} \backslash U. \end{cases} \tag{5.3}$$

Executing a perfect intervention on a Structural Causal Model $\mathcal{M}$, resulting in $\mathcal{M}_{\text{do}(U,\xi_U)}$, has a corresponding effect on the induced directed mixed graph $\mathcal{G}(\mathcal{M})$, transforming it into $\mathcal{G}_{\text{do}(U,\xi_U)}(\mathcal{M})$. This relationship can be succinctly expressed as $(\mathcal{G} \circ \text{do}\,(U, \xi_U))\,(\mathcal{M}) = (\text{do}(U) \circ \mathcal{G})(\mathcal{M})$. In the graph operation, this intervention effectively removes all incoming edges to the vertices contained in the set U.

In scenarios where ODEs are analyzed under equilibrium conditions, they can be seamlessly transformed into SCMs. This transformation is achieved through a procedure known as 'labeled equilibrium equations', a concept extensively detailed in (Mooij, Janzing, & Sch"olkopf, 2013). For a comprehensive understanding of this process, readers are referred to section 4.2.3 of this thesis. The forthcoming discussion will delve into an example of microbial network(possibly cyclic) elucidating its transformation into a Structural Causal Model and exploring the dynamics of interventions within this framework.

Example 5.1. *Consider a microbial community with four species whose abun-*

dances are model by the following set of ODEs:

$$\begin{cases} \frac{dX_1}{dt} = f_1(X_1, X_2, X_4) \\ \frac{dX_2}{dt} = f_2(X_1, X_2) \\ \frac{dX_3}{dt} = f_3(X_2, X_3) \\ \frac{dX_4}{dt} = f_4(X_4) \end{cases} \tag{5.4}$$

The functions f_i represent the population dynamics within microbial communities. These dynamics can be modeled using various types of functional relationships, ranging from linear to non-linear forms. For instance, dynamics akin to the Lotka-Volterra model, Holling-Type II, or the DeAngelis-Beddington model are typical examples (Jost & Ellner, 2000).

When the system reaches equilibrium, the associated graphical representation can be observed in Fig. 5.1. A perfect intervention on the variable X_1 disrupts the cyclic interaction between X_1 and X_2, leading to an altered graph structure. This modified structure, which results from the intervention, is depicted in the second graph of Fig. 5.1. The system at equilibrium conditions would be naturally expressed in the form of SCM with causal mechanisms defined by f_i. The intervention on the ODE system also corresponds to the intervened SCM with the same induced graph.

In the context of ODEs describing the dynamics of the system, several approaches can implement a perfect intervention. A typical method involves incorporating terms like $\kappa\left(\xi_u - X_u\right)$, with $\kappa > 0$, into the differential equation for $\frac{dX_u(t)}{dt}$ for each $u \in U$:

$$\frac{dX_u(t)}{dt} = \begin{cases} f_u\left(\boldsymbol{X}_{\text{Pa}(u)}\right) + \kappa\left(\xi_u - X_u\right) & u \in U \\ f_u\left(\boldsymbol{X}_{\text{Pa}(u)}\right) & u \notin U \end{cases} \tag{5.5}$$

This modification effectively extends the system with components that monitor the values of $\{X_u\}_{u \in U}$, applying negative feedback when deviations from the target values $\{\xi_u\}_{u \in U}$ are detected. By considering the limit $\kappa \to \infty$, we approximate the ideal scenario where the intervention overrules the usual mechanisms determining

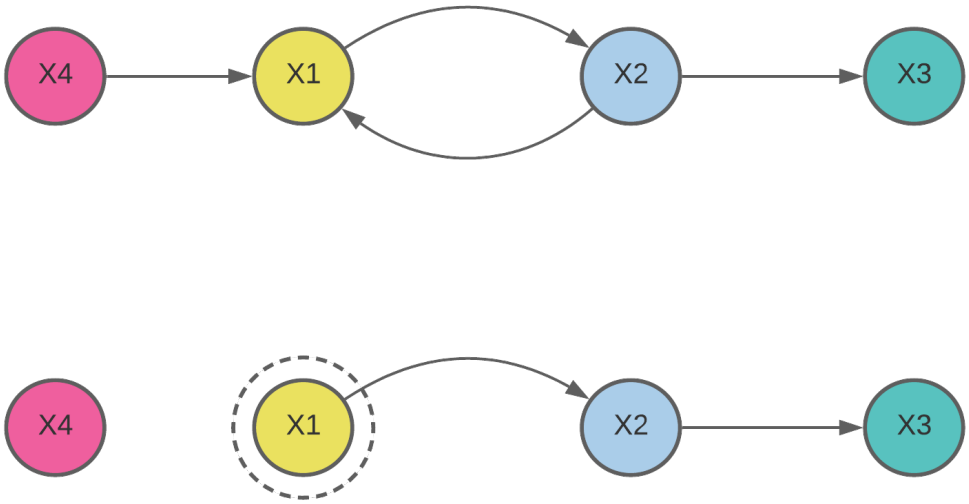


Figure 5.1: Graph of microbial community Example 5.1 and corresponding perfect intervention on X_1

$\boldsymbol{X}_U$'s value. Assuming boundedness of the functions $\{f_u\}_{u \in U}$, this limit leads to the intervened system:

$$\frac{dX_u(t)}{dt} = \begin{cases} 0 & u \in U \\ f_u\left(\boldsymbol{X}_{\mathrm{pa}_{\mathcal{D}}(u)}\right) & u \notin U \end{cases}, \quad X_u(0) = \begin{cases} \xi_u & u \in U \\ (\boldsymbol{X}_0)_u & u \notin U \end{cases} \tag{5.6}$$

5.2.4 Possible Cycles

In the preceding section, we established that the dynamics of a microbial population system, governed by a set of Ordinary Differential Equations, can be effectively modeled as a Structural Causal Model. Let's denote the system's underlying interaction graph as a directed mixed graph $\mathcal{G}$. Under the assumption of causal sufficiency, this graph can be reduced to a directed graph. We characterize the output of algorithms like Fast Causal Inference (FCI) (Zhang, 2008) and Greedy Fast Causal Inference (GFCI) (Ogarrio, Spirtes, & Ramsey, 2016) as $\mathcal{P}\left(\mathrm{IM}_\sigma(\mathcal{G})\right)$, which denotes a directed partial ancestral graph (DPAG). This DPAG represents a Markov Equivalence Class (MEC) of graphs. To precisely delineate the nature of this MEC,

it is necessary to first define the concepts of σ-separation within this framework.

Definition 5.4. *(Inducing path in DMG and σ-separation) Consider a directed mixed graph (DMG) $\mathcal{G} = \langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$. Define an inducing path (or walk) between two vertices $u, v \in \mathcal{V}$ as a path (or walk) within $\mathcal{G}$ connecting u and v, where every collider on this path (or walk) is a member of $\mathrm{An}_{\mathcal{G}}(\{u, v\})$. Additionally, every non-collider vertex on this path (or walk), except for u and v, should exclusively have outward directed edges to adjacent vertices on the path (or walk) that belong to the same strongly connected component in $\mathcal{G}$.*

When there is an inducing path in $\mathcal{G}$ between nodes u and v, we equivalently say that u and v cannot be σ-separated by any subset of nodes that does not contain either u or v.

The σ-separation is reduced to d-separation in the setting of directed acyclic graph or acyclic directed mixed graph.

Given a directed mixed graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E}, \mathcal{F} \rangle$, where $\mathcal{V}$ denotes the vertex set, $\mathcal{E}$ denotes the set of directed edges and $\mathcal{F}$ denotes the set of bi-directed edges. An acyclification of $\mathcal{G}$ is an acyclic directed mixed graph $\mathcal{G}^{acy} = \langle \mathcal{V}^{acy}, \mathcal{E}^{acy}, \mathcal{F}^{acy} \rangle$ with the following rules

1. the two mixed graph share the same vertex set; i.e. $\mathcal{V} = \mathcal{V}^{acy}$
2. $i \to j \in \mathcal{E}^{acy}$ if and only if there is a vertex in the strongly connected component of j in graph $\mathcal{V}$ and there is a directed edge $i \to k \in \mathcal{E}$
3. for any vertex i in the strongly connected component of j in graph $\mathcal{V}$, connect $i \leftrightarrow j \in \mathcal{F}^{acy}$.

Proposition 5.1. *(Mooij & Claassen, 2020) For any DMG $\mathcal{G}$ and the acyclification $\mathcal{G}^{acy}$ of $\mathcal{G}$ defined above, $\mathrm{IM}_{\sigma}(\mathcal{G}) = \mathrm{IM}_{\sigma}\left(\mathcal{G}^{acy}\right) = \mathrm{IM}_{d}\left(\mathcal{G}^{acy}\right)$.*

The acyclification process described, encapsulated in rules 2 and 3, possesses certain noteworthy properties. Specifically, rule 2 necessitates the preservation and duplication of all incoming edges to any node that forms part of a strongly connected component. Conversely, rule 3 mandates the complete interconnection of

all such components through the exclusive use of bi-directional edges. This process inevitably leads to the generation of additional, or spurious, edges. Consequently, it becomes evident that the skeleton of the original directed mixed graph $\mathcal{G}$, denoted as $Sk(\mathcal{G})$, is subsumed within the skeleton of its acyclified counterpart, $Sk(\mathcal{G}^{acy})$. Extending this logic further, it follows that $Sk(\mathcal{G})$ is also contained within $Sk(\mathcal{P}\left(\mathrm{IM}_{\sigma}(\mathcal{G})\right))$, which is equivalent to $Sk(\mathcal{P}\left(\mathrm{IM}_{\sigma}(\mathcal{G}^{acy})\right))$.

This property is pivotal as it implies that the skeleton of our primary directed graph is encompassed within its corresponding Directed Partial Ancestral Graph, which can be deduced from any Partial Ancestral Graph (PAG)-yielding inference algorithm such as Fast Causal Inference or Greedy Fast Causal Inference under ideal conditions. The significance of this result lies in establishing a foundational basis for using the inferred DPAG as a starting point for further refinement, including the determination of edges to be eliminated and the clarification of edge marks.

Proposition 5.2. *Identification of Possible Cycles from DPAG:*
Consider a Directed Mixed Graph denoted as $\mathcal{G}$ and its corresponding Directed Partial Ancestral Graph, represented by $\mathcal{P} = \mathcal{P}(\mathrm{IM}_{\sigma}(\mathcal{G}))$. It is proposed that any set of vertices forming a cycle of more than three nodes in $\mathcal{G}$ must result in a clique, with a size of at least three nodes, within the skeleton of the DPAG, $Sk(\mathcal{P}(\mathrm{IM}_{\sigma}(\mathcal{G})))$.

Sketch of proof:
Given a directed mixed graph $\mathcal{G}$, it is observed that no pair of nodes within any strongly connected component of $\mathcal{G}$ can be σ-separated. As a consequence, each strongly connected component in $\mathcal{G}$ transforms into a fully-connected component in its corresponding Directed Partial Ancestral Graph, denoted as $\mathcal{P}$. This phenomenon occurs due to the inherent properties of σ-separation in the context of acyclification.

For any two distinct nodes i and j within the same strongly connected component of $\mathcal{G}$, multiple acyclifications of $\mathcal{G}$ are conceivable. In one such acyclification (what we defined above is a possible one), there could be a directed edge from i to j, symbolized as $i \to j$, while in another, the direction could be reversed, represented as $i \leftarrow j$. Consequently, the edge between nodes i and j in the DPAG $\mathcal{P}$ must be

$i \circ\!\!-\!\!\circ j$.

5.2.5 Problem Description

Consider a microbial community whose population dynamics are governed by a set of Ordinary Differential Equations (ODEs), as detailed in previous sections. We observe an empirical distribution of endogenous variables, comprising equilibrium data from the microbial system. Let the underlying interaction network within this microbial community be represented by a graph $\mathcal{G}$, which, as established earlier, corresponds to the causal graph within the SCM framework.

Our primary objective, starting from a Partial Ancestral Graph $\mathcal{P}$ inferred using a PAG inference algorithm such as Fast Causal Inference is to formulate a set of experiments. These experiments aim to adaptively ascertain the actual directed graph, potentially cyclic, encapsulated within this PAG. This directed graph elucidates the causal mechanisms of the SCM instigated by the microbial interaction network, thereby uncovering the true underlying interaction network of the microbial communities.

The complexity of this task is heightened by the presence of cycles. In the forthcoming sections, we present a two-stage algorithm offering an experimental design in our context. This design aims to adaptively learn the genuine causal structure, employing interventional data to augment our understanding.

5.3 Experiment Design Strategy and Illustrative Examples

5.3.1 Breaking the Cycles

The initial phase in our experimental design involves identifying a set of interventions critical for disrupting potential cycles in the derived Directed Partial Ancestral Graph, denoted as $\mathcal{P}$. Proposition 5.2 elucidates that cycles in the underlying directed graph manifest distinct patterns within $\mathcal{P}$. In an ideal scenario, these patterns

are recognizable. However, solely from $\mathcal{P}$, it is not possible to definitively ascertain whether vertices u and v are constituents of a causal cycle. This ambiguity arises because nodes directing into a cycle in the underlying graph are perceived as part of the strongly connected component in the skeleton of $\mathcal{P}$.

Furthermore, the precise number of nodes in each causal cycle remains indeterminable from $\mathcal{P}$ alone. To tackle this, our strategy initially targets the disruption of cycles encompassing three or more nodes, postponing the resolution of two-node cycles to the subsequent stage of edge mark determination.

Given that we are examining cliques of size three or larger in $\mathcal{P}$'s skeleton, which is inherently an undirected graph, there may exist multiple overlapping cliques of. Our goal here is to identify a minimum set of nodes to break all cycles present in the graph. To optimize the intervention set and thereby reduce experimental costs, we employ the "Minimal Set Covering" technique, which is defined in the subsequent section. It is imperative to note that the selected nodes for intervention might not be part of the actual cycle. This stage of the algorithm proceeds until all cycles are effectively broken.

5.3.1.1 Cycle-Breaking in DPAG Using Greedy Algorithm for Minimal Set Cover

Before delving into the algorithm, it is essential to understand the fundamental concept of the Set Cover problem (Caprara, Toth, & Fischetti, 2000) which forms the basis of our approach.

Definition 5.5 (Set Cover Problem)**.** *Given a finite set $\mathcal{U}$ and a collection $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ of subsets of $\mathcal{U}$, where each subset $S_i \subseteq \mathcal{U}$, the Set Cover Problem is to find the smallest sub-collection $\mathcal{C} \subseteq \mathcal{S}$ such that the union of all sets in $\mathcal{C}$ equals $\mathcal{U}$. Formally, it is required that $\bigcup_{S_i \in \mathcal{C}} S_i = \mathcal{U}$, and the goal is to minimize the size of $\mathcal{C}$.*

The problem of identifying a minimal set of interventions for cycle-breaking in a Directed Partial Ancestral Graph can be viewed as a specialized case of the Set Cover Problem. In this adapted framework:

- The universe $\mathcal{U}$ is the collection of all nodes in the DPAG that are either part of cycles or connected to cycles.
- The sets $\mathcal{S}$ consist of cliques within the DPAG. Each clique represents a group of nodes interconnected in such a way that intervention on any node within the group has the potential to disrupt a cycle.
- The goal is to find a target set that minimally covers each clique in $\mathcal{S}$, ensuring that every clique contributes at least one node to the final intervention set. This aligns with the Set Cover Problem's objective of covering all elements in $\mathcal{U}$ with the fewest subsets.

Proposition 5.3 (Adapted Set Cover for DPAG Cycle-Breaking)**.** *For the DPAG cycle-breaking problem, the Set Cover Problem is defined as follows: given a DPAG with nodes $\mathcal{U}$ and a collection of cliques $\mathcal{S}$, find a minimal subset $\mathcal{C} \subseteq \mathcal{U}$ such that for each clique $S_i \in \mathcal{S}$, there is at least one element from S_i in $\mathcal{C}$. The objective is to minimize the size of $\mathcal{C}$, ensuring that it covers at least one node from each clique to break the cycles effectively.*

This specialized framing of the Set Cover Problem for cycle-breaking in DPAG highlights the unique aspects of our approach, where the focus is not just on covering all nodes but on strategically selecting nodes to disrupt cycles in the graph.

Given the special case of the Set Cover Problem tailored for cycle-breaking in DPAGs, we propose a greedy heuristic approach for a practical solution. This method iteratively constructs a minimal set of interventions, efficiently disrupting the cycles in the DPAG. The greedy heuristic strategy selects nodes in a manner that maximizes coverage of the cliques (sets) in each iteration, minimizing the overall number of interventions.

The algorithm operates on the principle of iteratively selecting nodes from the DPAG cliques, which cover the maximum number of uncovered nodes in each step. This process continues until all nodes in the DPAG that are part of or connected to cycles are covered. The key steps of the algorithm are:

- **Initialization:** Compile a list of all unique nodes across the cliques in the

DPAG and initialize an empty target set for interventions.

- **Iteration:** In each iteration, identify the clique with the maximum number of uncovered nodes and select one of its nodes for intervention.
- **Coverage Update:** After each selection, update the list of covered nodes.
- **Completion:** The process continues until all nodes connected to or part of cycles in the DPAG are covered.
- **Output:** The resulting target set forms the minimal set of interventions required for effective cycle-breaking in the DPAG.

However, the nodes selected in this way could not gurantee that they are from the cycle originally. Identifying nodes within a DPAG that contribute to cycles is challenging. The ambiguity in differentiating between nodes integral to cycles and those influencing them externally necessitates an iterative selection process including identifying cliques, selecting nodes for intervention using a greedy algorithm, gathering new data from these interventions, and learning a new DPAG structure. This process is repeated until the DPAG no longer exhibits significant cliques.

More specifically, the full overview of our algorithm is stated as:

Initialization:

- $\mathcal{P}$: The DPAG learned from observational data.
- $\mathcal{U}$: A set of unique nodes extracted from the cliques in $\mathcal{P}$.
- $\mathcal{C}$: The collection of cliques within $\mathcal{P}$.
- $\mathcal{T}$: An initially empty set, to hold the target nodes for intervention.
- $\mathcal{V}$: Another initially empty set, to track the covered nodes.

Greedy Set Covering:

- Iteratively identifies the clique with the maximum number of uncovered elements.
- Selects a node from this clique for $\mathcal{T}$.

- Updates $\mathcal{V}$ with covered nodes.

Data Collection and Learning:

- Gathers new interventional data based interventions on nodes in $\mathcal{T}$.
- Learns a new DPAG structure, $\mathcal{P}_{\text{new}}$, from this data.

Cycle Breaking and Repetition:

- Checks for remaining cliques of size ≥ 3 in $\mathcal{P}_{\text{new}}$.
- Repeats the process for new intervention nodes if cycles are not effectively broken.

Output:

- The final output is the set $\mathcal{T}$, representing the minimal set of intervention nodes that effectively breaks all significant cycles in the DPAG.

Algorithm Pseudocode

The following pseudocode outlines the greedy heuristic algorithm for cycle-breaking in DPAGs:

Algorithm 1 Greedy Heuristic for Minimal Set Covering in DPAGs

$\mathcal{P} \leftarrow$ DPAG learned from observational data

$\mathcal{U} \leftarrow$ unique nodes from DPAG cliques

$\mathcal{C} \leftarrow$ collection of cliques in $\mathcal{P}$

$\mathcal{T} \leftarrow$ empty set (target set for intervention)

$\mathcal{V} \leftarrow$ empty set (covered nodes)

while $|\mathcal{V}| < |\mathcal{U}|$ **do**

  $\mathcal{S} \leftarrow$ count uncovered elements in each clique in $\mathcal{C}$

  $i_{\text{max}} \leftarrow$ index of clique in $\mathcal{C}$ with max uncovered elements

  $e \leftarrow$ select a node from clique $\mathcal{C}[i_{\text{max}}]$

  $\mathcal{T} \leftarrow \mathcal{T} \cup \{e\}$

  $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{C}[i_{\text{max}}]$

end while

gather new interventional data based on intervention $\mathcal{T}$

learn new DPAG structure $\mathcal{P}_{\text{new}}$ from interventional data

if cliques in $\mathcal{P}_{\text{new}}$ have max size ≥ 3, **then**

  $\mathcal{U} \leftarrow \mathcal{U} \setminus \{$elements in $\mathcal{T}$ that failed to break cycles$\}$

  repeat above steps to find new $\mathcal{T}$

until all cycles are effectively broken in $\mathcal{P}_{\text{new}}$

return $\mathcal{T}$ ▷ Final set of intervention for cycle breaking

Example 5.2. *In this section, we consider a microbial system comprising fourteen nodes, featuring two distinct directed cycles. The first cycle encompasses nodes X_5, X_6, and X_7, highlighted in yellow, while the second cycle includes nodes X_8, X_9, X_{10}, and X_{11}, illustrated in pink. Additionally, nodes X_4 and X_3, marked in grey, are not part of any cycle, yet they play a significant role in the cliques due to their connections to the cycles. Here we simulate the edge strength with a uniform distribution, where interaction strengths are randomly drawn as either positive or negative values from 0.5 to 0.8. The population dynamics are simulated using the GLV model, introduced in Section 4.2.2 with Equations 4.3 and 4.4, where θ_{max} is set to 0 and 1000 data points are generated.*

The underlying directed graph of this microbial community, denoted as $\mathcal{G}$, is presented in part (a) of Figure 5.2. In part (b) of the figure, we display the skeleton of the DPAG, $\mathcal{P}$. The red dotted lines in part (b) represent the "spurious edges" within the DPAG, indicative of the Markov Equivalence Class, as elaborated in the preceding section detailing the problem setting.

In the context of our microbial system, the algorithm identifies several cliques within the DPAG, which are crucial for understanding the network's cycle structure. The identified cliques are as follows:

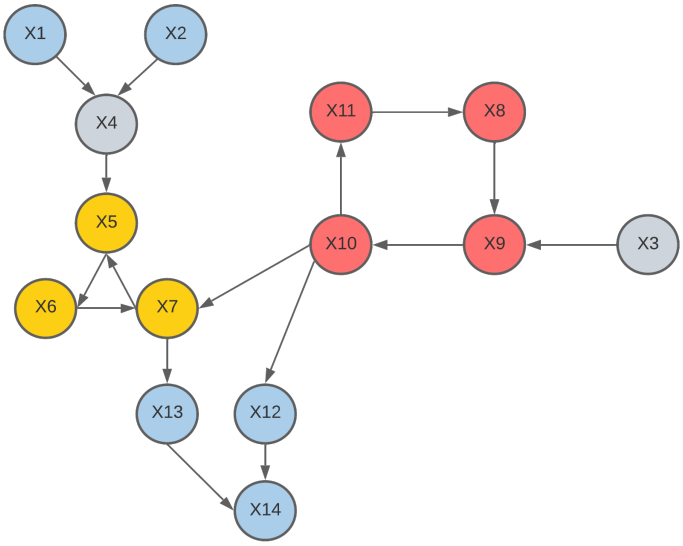
- The first clique includes nodes 6, 5, 7, and 4.
- The second clique comprises nodes 6, 5, 7, and 10.
- The third clique consists of nodes 10, 3, 11, 9, and 8.

Through the application of our minimal set covering algorithm, nodes 6 and 10 are selected as the targets for intervention. These nodes are pivotal in that they appear in multiple cliques, indicating their central role in the network's cycle structure. By intervening on nodes 6 and 10, we aim to effectively break the cycles and gain deeper insights into the causal dynamics of the microbial community.

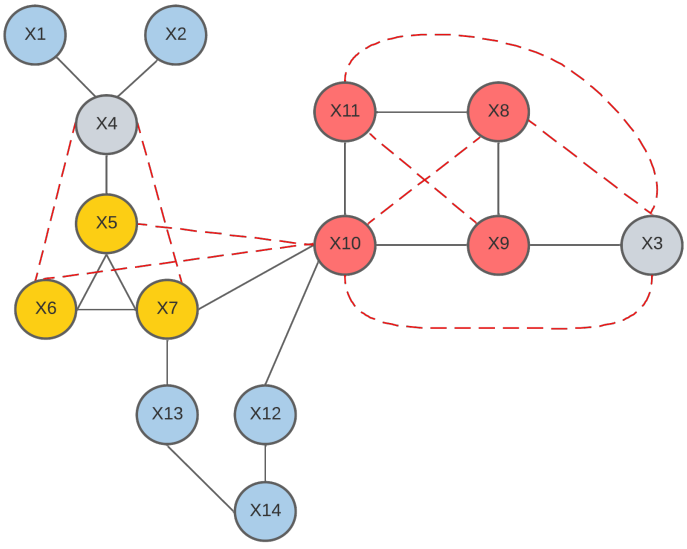
5.3.2 Determining Undetermined Edgemarks in DPAGs

In the subsequent phase of our algorithm, we concentrate on resolving the ambiguities present in DPAGs to effectively derive a Structural Causal Model that is representative of the corresponding Directed Acyclic Graph. The process involves disambiguating the edgemarks in DPAGs, particularly focusing on edges that possess at least one ambiguous endpoint, denoted by an 'o' mark. These edges in DPAGs are typically categorized into three types: directed edges ($\rightarrow$), partially directed edges ($\circ\!\!\rightarrow$), and non-directed edges ($\circ\!\!-\!\!\circ$), where the 'o' symbol represents an ambiguous mark that could either be a tail ('-') or a directed mark ('>').

Our objective in this stage is to systematically transform these ambiguous edgemarks into definitive directed edges, thereby transitioning from a DPAG to a more explicit representation in the form of an SCM. This transformation is crucial as it



a. Underlying Directed Graph



b. Skeleton of DPAG

Figure 5.2: Example in the case of cycle-breaking

allows us to construct a valid SCM that accurately reflects the causal structure of the underlying system.

In our framework, executing an intervention on a single variable X_i or a set of variables X is represented as experiment(X_i) or experiment(X), respectively. If these interventions are conditional upon another set of variables D, we notate them as experiment($X_i \mid D$) and experiment($X \mid D$).

In the context of causal analysis, experiments conducted on a variable X_i, akin to randomized controlled trials, serve to isolate X_i from the influence of other variables within the system. Such an experimental manipulation enforces a specific distribution on X_i, altering the joint distribution of all variables that are directly or indirectly influenced by X_i. Crucially, this experimental setup leaves the conditional distribution of other variables, given the values of X_i, unchanged. This principle allows us to discern the variables affected by X_i by observing the associations between X_i and other variables post-experimentation. Graphically, this equates to severing all incoming edges to X_i in the underlying Directed Acyclic Graph.

The experimental process keeps all parameters constant except for those related to the experimented variable, specifically $P\left(X_i \mid \text{Pa}\left(X_i\right)\right)$. We then observe the effects of this manipulation on other variables of interest to determine the post-interventional distribution.

To quantitatively assess the outcomes of the experiment, we perform a comparative analysis for each variable of interest X_j. This involves contrasting the original distribution P and the post-interventional distribution P_{Exp}, comparing $P\left(X_j\right)$ with $P_{\text{Exp}}\left(X_j\right) = P\left(X_j \mid \text{do}\left(X_i = x_i\right)\right)$.

In the scenario where an experiment is conducted on X_i and it impacts another variable X_j, we denote this relationship as $\text{Exp}\left(X_i\right) \rightsquigarrow X_j$. Conversely, if X_j remains unaffected by the experiment on X_i, we denote this as $\text{Exp}\left(X_i\right) \not\rightsquigarrow X_j$. This notation helps in elucidating the causal connections or the lack thereof between the experimented variable and other variables in the system.

To address the issue of ambiguous edges within the Partial Ancestral Graph, we

begin by resolving edges that have one end marked with a circle and the other with an arrowhead. Under the assumption of causal sufficiency, these edges can be resolved with greater certainty. Specifically, we can confidently orient such edges as directed edges, transforming the circle-marked end into an arrowtail. This process is grounded in our assumption that the system lacks latent confounders, allowing us to infer more direct causal relationships.

Next, we turn our attention to the edges marked with circles at both ends. These "circle-circle" edges represent more complex cases in our causal graph.

In the context of disambiguating "circle-circle" edges within a Partial Ancestral Graph, we employ a systematic approach to determine the directionalities of these edges. This approach involves conducting experiments on the nodes connected by such edges and analyzing the impact of these interventions. We detail this process as follows:

For an edge between two nodes X_i and X_j with both ends marked by circles, there is a potential for a directed cycle to exist between them. To resolve this ambiguity:

1. Experiment on X_i: We first perform an intervention on X_i (denoted as $\exp(X_i)$) and observe its effect on X_j. - If $\exp(X_i)$ has no effect on X_j (denoted as $\exp(X_i) \not\to X_j$), we can infer that there is no direct causal influence from X_i to X_j. In this scenario, the edge can be oriented from X_j to X_i. - Conversely, if $\exp(X_i)$ influences X_j (denoted as $\exp(X_i) \to X_j$), this suggests a causal relationship from X_i to X_j. However, to fully resolve the nature of the relationship, an additional step is required.

2. Further Experiment on X_j: Next, we conduct an intervention on X_j (denoted as $\exp(X_j)$) and observe its effect on X_i. - If $\exp(X_j)$ influences X_i (denoted as $\exp(X_j) \to X_i$), it indicates a directed cycle between X_i and X_j. - If $\exp(X_j)$ has no effect on X_i (denoted as $\exp(X_j) \not\to X_i$), this confirms a unidirectional causal relationship from X_i to X_j.

This methodical approach, involving a sequence of targeted experiments, allows for the accurate determination of causal relationships and the resolution of ambiguities

in the PAG structure. This process can be summarized in an algorithmic format to enhance clarity and facilitate implementation.

Algorithm 2 Resolving Circle-Circle Edges in DPAG

Input: DPAG with circle-circle edges, cycle-breaking intervention set
Output: Resolved edge directions in DPAG
for each edge $(X_i \circ\!\!-\!\!\circ X_j)$ in DPAG **do**
  if X_i or X_j is in the cycle-breaking intervention set **then**
    if X_i was intervened **then** orient edge as $X_i \to X_j$
    if X_j was intervened **then** orient edge as $X_j \to X_i$
  else
    Perform $\exp(X_i)$ and observe effect on X_j
    if $\exp(X_i) \not\to X_j$ **then**
      Orient edge as $X_j \to X_i$
    else if $\exp(X_i) \to X_j$ **then**
      Perform $\exp(X_j)$ and observe effect on X_i
      if $\exp(X_j) \to X_i$ **then**
        Identify directed cycle between X_i and X_j
      else if $\exp(X_j) \not\to X_i$ **then**
        Orient edge as $X_i \to X_j$
      end if
    end if
  end if
end for

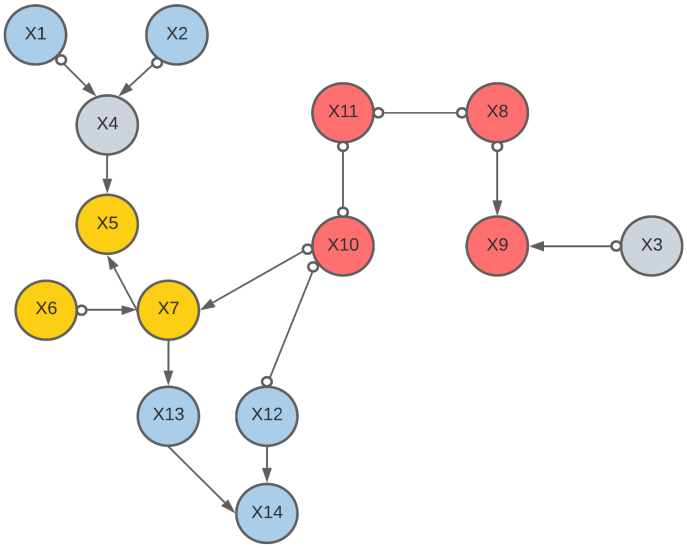
This algorithm provides a structured approach to systematically resolve ambiguous edges within a DPAG, ensuring a more accurate representation of the underlying causal relationships in the microbial community network.

Now we illustrate the approach by continuing Example 5.2. Following the cycle-breaking interventions on nodes X_6 and X_{10}, as discussed in Section 5.3.1, we observe notable modifications in the DPAG structure. This section examines the reorientation of partially directed edges within this DPAG. The resultant graph, depicted in Figure 5.3, serves as a basis for this analysis.

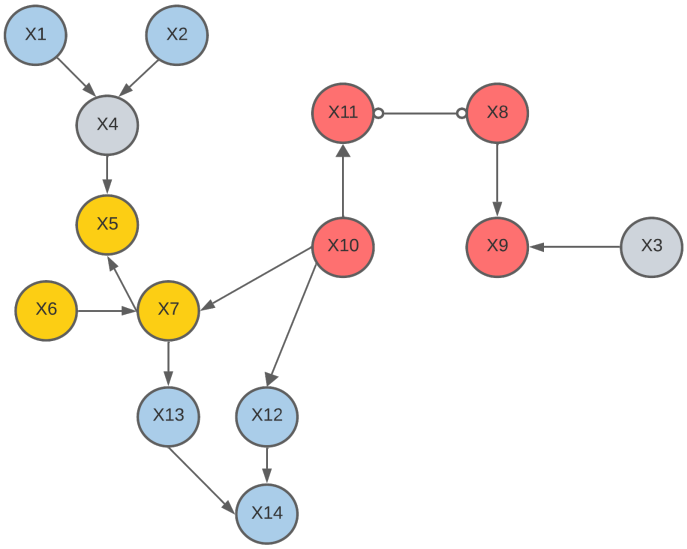
Figure 5.3 Analysis:

1. *Part a - Post-Cycle-Breaking DPAG:* Demonstrates the DPAG following the interventions on X_6 and X_{10}. It reveals the resulting changes in the graph structure, particularly the removal of specific edges and the redefinition of certain connections.

2. *Part b - Graph with Reoriented Edges:* Illustrates the DPAG after the orientation of partially directed "$\circ\!\!\rightarrow$" edges. This graph is derived by applying the algorithm for resolving circle-arrow edges.



a. DPAG after intervention on node X6 and X10



b. After orientation of partially directed edges

Figure 5.3: Example in the case of edgemarks determination

As we progress in refining the causal model of our microbial community, we then need to resolve ambiguities in "circle-circle" edges within the DPAG. A pertinent example of such an edge exists between nodes X_8 and X_{11}.

Experimental Approach: To elucidate the directionality of the edge between X_8 and X_{11}, we conducted a series of targeted experiments:

- **Experiment on X_{11}:** An intervention $\exp(X_{11})$ was performed to assess its influence on X_8. The observation from this experiment revealed that $\exp(X_{11})$ significantly affects X_8, suggesting a potential causal influence from X_{11} to X_8.
- **Experiment on X_8:** Subsequently, an intervention $\exp(X_8)$ was conducted to determine its impact on X_{11}. Interestingly, the results indicated that $\exp(X_8)$ had no discernible effect on X_{11}. This outcome implies the absence of a direct causal relationship from X_8 to X_{11}.

Edge Orientation Determination: Based on the experimental findings, we conclude the following:

- Since $\exp(X_{11})$ affects X_8 while $\exp(X_8)$ does not influence X_{11}, it is reasonable to infer a unidirectional causal relationship from X_{11} to X_8.
- Consequently, the previously ambiguous edge between X_8 and X_{11} in the DPAG is now oriented as $X_{11} \to X_8$.

Implications for the Causal Model: This resolution of the edge between X_8 and X_{11} significantly enhances the clarity and accuracy of the causal model of our microbial community. By systematically determining the directionality of ambiguous edges, we gain a more profound understanding of the intricate causal mechanisms governing microbial interactions.

This example illustrates the effectiveness of experimental interventions in disambiguating causal relationships within complex systems, thereby contributing to the robustness and reliability of the derived causal model.

5.3.3 Recovering the Cycles

In the concluding stage of our adaptive learning framework, we focus on recovering potential cycles that may have been disrupted during the initial cycle-breaking

interventions. This recovery process is essential to ensure the completeness and accuracy of our microbial community's interaction network. Our approach involves reintegrating edges that were previously eliminated to break cycles, thereby restoring the network's structure. We propose two methodologies for this recovery:

Our first method employs a heuristic approach to reevaluate the nodes within the cycle-breaking set. This process involves a detailed comparison between the neighboring nodes of a specific target node in both the observational DPAG and the current directed network post-interventions. The key steps are:

1. **Neighbor Comparison:** For each node in the cycle-breaking set, we compare its neighborhood in the original observational DPAG with its current neighborhood in the directed network post-intervention.
2. **Edge Assessment:** Based on this comparison, we identify potential edges that may have been erroneously omitted during the cycle-breaking process.
3. **Heuristic Edge Addition:** We then heuristically reintegrate these identified edges back into the target node, thereby attempting to restore the node's original neighborhood structure as observed in the DPAG.

This heuristic approach, as illustrated below, aims to reconstruct the network by carefully assessing and re-adding edges that were removed but are potentially significant for the accurate representation of microbial interactions.

The other methodology is more experimentally intensive but offers a direct and comprehensive approach to restoring the network. This method involves:

1. **Reapplying the Cycle-Breaking Algorithm:** We execute the cycle-breaking algorithm, as outlined in the initial stages, with a different selection of cycle-breaking sets. This variation aims to explore alternative configurations for cycle disruption.
2. **Gathering New Interventional Data:** Subsequent interventions based on the newly selected cycle-breaking sets generate fresh interventional data.
3. **Network Reconstruction:** We utilize this new set of interventional data to

reconstruct the network, paying particular attention to the edges connected to the nodes of interest.

4. **Edge Integration:** The newly inferred network structure, enriched with additional data, allows for the reintegration of previously omitted edges, thereby refining the overall network topology.

This experimental approach ensures a comprehensive reassessment of the network, capturing the intricate dynamics and dependencies within the microbial community. By employing both heuristic and experimental methodologies, we aim to recover a robust and accurate representation of the underlying microbial interaction network, ensuring that our model reflects the true complexity and interconnectedness of these biological systems.

In the specific example of our 14-node microbial system, as previously depicted, the reintegration of edges into nodes X_6 and X_{10} is approached through a heuristic approach. This section focuses on elucidating the appraoch and underlying rationale for this reintegration process.

Our heuristic approach to edge reintegration involves a careful examination of the network structure to align with the initial observational DPAG, $\mathcal{P}$. The primary steps in this process are:

1. **Identifying Potential Edges:** Given the nodes X_6 and X_{10}, we seek to determine the most plausible edges leading into these nodes that would be consistent with the original observational DPAG structure.

2. **Heuristic Selection of Edges:** For node X_6, the edge $X_5 \to X_6$ presents itself as a logical reintegration based on its presence in the observational DPAG. Similarly, for node X_{10}, the edge $X_9 \to X_{10}$ aligns with the observed DPAG structure.

3. **Network Consistency Check:** The reintegrated edges are then evaluated for consistency by comparing its ground truth DPAG image with the overall network structure, ensuring they do not introduce contradictions or implausible dependencies within the microbial community.

While our heuristic approach is guided by practical considerations and observable network characteristics, it is important to note the theoretical underpinnings of such methodologies. The search for edges within the directed graph (potentially cyclic) that is a part of the Markov Equivalence Class represented by our observational DPAG is fundamentally a process of network reconstruction.

Recent advancements in the field, such as those presented in (Claassen & Mooij, 2023), offer promising insights into establishing Markov equivalence in cyclic directed graphs. These theoretical developments could provide a more rigorous and generalizable framework for network reconstruction, including the identification and reintegration of edges in complex systems like microbial communities.

In conclusion, while our current methodologies predominantly rely on heuristic strategies, they demonstrate efficacy in the context of our specific microbial system. The reintegration of edges into nodes X_6 and X_{10} exemplifies this approach, highlighting our ability to recover critical network components. This process was implemented using the scripts available in our GitHub repository: `https://github.com/AFFFPupu/Microbial_PAG_with_Interventions`. Specifically, system dynamics simulations and graph interventions were handled by functions like `simulate_x_data_onematrix_julia()` and structural manipulations are available in the functions.R file. However, the field is continuously evolving, and the integration of more theoretical frameworks and methodologies holds the potential to further refine and enhance our understanding of microbial interaction networks. Such advancements will undoubtedly contribute to a deeper and more accurate representation of the complex dynamics governing these biological systems.

5.4 Conclusion and Future Directions

This chapter has outlined an attempt for adaptive learning and experimental design in the context of microbial community dynamics. We have introduced an algorithm that effectively integrates both observational and interventional data to infer the

true causal structure within microbial communities. Our approach addresses the inherent challenges posed by feedback loops and undetermined edgemarks in Partial Ancestral Graphs. By employing a minimal set covering technique and strategically designed experiments, we have made significant strides in accurately recovering the complete causal structure of these complex systems.

While our approach represents one step forward, it is not without limitations. The reliance on certain assumptions, such as causal sufficiency and faithfulness, may not always hold in real-world scenarios. Additionally, the complexity of microbial systems, coupled with the challenges of conducting controlled experiments, may limit the practical applicability of our algorithms in certain contexts. The computational demands of the algorithm, especially in large and complex networks, also pose a challenge.

Looking ahead, several avenues for future research emerge from our work. First, further refinement and optimization of the algorithm could enhance its applicability to even more complex microbial communities. Exploring ways to relax the assumptions of causal sufficiency and faithfulness could also broaden the scope of the algorithm applicability. Additionally, integrating our approach with advanced computational techniques such as machine learning could yield more efficient and scalable solutions. Finally, applying our methodology to real world microbial communities and comparing the results with empirical data will be crucial for validating and refining our approach.

References

Bongers, S., Forr'e, P., Peters, J., & Mooij, J. M. (2021). Foundations of Structural Causal Models With Cycles and Latent Variables. *The Annals of Statistics*, *49*(5), 2885–2915.

Caprara, A., Toth, P., & Fischetti, M. (2000). Algorithms for the Set Covering Problem. *Annals of Operations Research*, *98*(1-4), 353–371.

Claassen, T., & Mooij, J. M. (2023). Establishing Markov Equivalence in Cyclic Directed Graphs. In *Uncertainty in Artificial Intelligence* (pp. 433–442).

Hauser, A., & B"uhlmann, P. (2014). Two Optimal Strategies for Active Learning of Causal Models From Interventional Data. *International Journal of Approximate Reasoning*, *55*(4), 926–939.

He, Y.-B., & Geng, Z. (2008). Active Learning of Causal Networks With Intervention Experiments and Optimal Designs. *Journal of Machine Learning Research*, *9*(Nov), 2523–2547.

Hyttinen, A., Eberhardt, F., & Hoyer, P. O. (2012). Learning Linear Cyclic Causal Models With Latent Variables. *The Journal of Machine Learning Research*, *13*(1), 3387–3439.

Hyttinen, A., Hoyer, P. O., Eberhardt, F., & Jarvisalo, M. (2013). Discovering Cyclic Causal Models With Latent Variables: A General SAT-based Procedure. *ArXiv Preprint ArXiv:1309.6836*.

Jost, C., & Ellner, S. P. (2000). Testing for Predator Dependence in Predator-prey Dynamics: a Non-parametric Approach. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *267*(1453), 1611–1620.

Meganck, S., Maes, S., Leray, P., Manderick, B., Rouen, I., & St Etienne du Rouvray, F. (2006). Learning Semi-Markovian Causal Models Using Experiments. In *Probabilistic Graphical Models* (pp. 195–206).

Mooij, J. M., & Claassen, T. (2020). Constraint-based Causal Discovery Using Partial Ancestral Graphs in the Presence of Cycles. In *Conference on Uncertainty in Artificial Intelligence* (pp. 1159–1168).

Mooij, J. M., Janzing, D., & Sch"olkopf, B. (2013). From Ordinary Differen-

tial Equations to Structural Causal Models: the Deterministic Case. *ArXiv Preprint ArXiv:1304.7920*.

Ogarrio, J. M., Spirtes, P., & Ramsey, J. (2016). A Hybrid Causal Search Algorithm for Latent Variable Models. In *Conference on Probabilistic Graphical Models* (pp. 368–379).

Pearl, J. (2009). *Causality*. Cambridge University Press.

Pearl, J., et al. (2000). Models, Reasoning and Inference. *Cambridge, UK: CambridgeUniversityPress*, *19*(2).

Shanmugam, K., Kocaoglu, M., Dimakis, A. G., & Vishwanath, S. (2015). Learning Causal Graphs With Small Interventions. *Advances in Neural Information Processing Systems*, *28*.

Triantafillou, S., & Tsamardinos, I. (2015). Constraint-based Causal Discovery From Multiple Interventions Over Overlapping Variable Sets. *The Journal of Machine Learning Research*, *16*(1), 2147–2205.

Zhang, J. (2008). On the Completeness of Orientation Rules for Causal Discovery in the Presence of Latent Confounders and Selection Bias. *Artificial Intelligence*, *172*(16-17), 1873–1896.

6 | Reinforcement Learning in Network Control

Abstract

In the rapidly evolving landscape of network control, the integration of reinforcement learning (RL) techniques offers a transformative approach to optimizing complex systems. This chapter delves into the application of model-free discretized Q-learning, a subset of RL, in the realm of microbial network control. Microbial networks, characterized by their intricate interactions and dynamic nature, pose significant challenges for traditional control strategies. The main focus of this chapter is to explore how discretized Q-learning can effectively navigate the complexities of microbial networks to achieve optimal control.

The key contribution of this chapter is the development and implementation of a discretized Q-learning algorithm tailored specifically for microbial network control. This approach stands out due to its ability to operate without a predefined model of the network, thereby navigating uncertainties and variations inherent in microbial systems.

By bridging the gap between reinforcement learning theory and practical network control, this chapter presents a novel pathway for advancing microbial network management. It offers valuable insights for researchers and practitioners seeking to leverage the power of RL in complex biological systems, paving the way for more efficient, robust, and adaptive network control solutions.

6.1 Introduction

The pursuit of controlling complex directed networks, following their meticulous reconstruction, serves as the ultimate testament to our profound grasp of their intrinsic dynamics. In a landmark 2011 study (Liu, Slotine, & Barab'asi, 2011), researchers unveiled analytical tools designed to systematically assess the controlla-

bility of complex directed networks, delineating a set of critical driver nodes. These nodes, upon receiving specific perturbations, can shepherd the entire network to any desired state, hence defining system controllability. The study posited a structural resemblance in the controllability between nonlinear and linear dynamic systems, prompting an exploration of linear dynamics represented by:

$$\frac{dx(t)}{dt} = Ax(t) + Bu(t) \tag{6.1}$$

Here, $x(t) = (x_1(t), \ldots, x_N(t))^T$ denotes the state vector of a system with N nodes at time t, matrix A encapsulates the nodal interactions, and matrix B is an $N \times M$ input matrix, governing the nodes under external control through $u(t) = (u_1(t), \ldots, u_M(t))^T$. The system, governed by equation 6.1, achieves controllability if the controllability matrix $C = (B, AB, A^2B, \ldots, A^{N-1}B)$ exhibits full rank N.

Subsequently, the study introduced a 'maximum matching' method to ascertain the minimal number of driver nodes required for full network control. This method seeks the largest set of non-sharing links between start and end nodes. Intriguingly, in various networks from Erdos-Renyi models to the World Wide Web, driver nodes are often those with lower degrees, challenging conventional expectations.

Focusing on microbial communities, (Angulo, Moog, & Liu, 2019) have developed a theoretical framework for their control. The preliminary step is employing the maximum matching algorithm on the microbial network G, which calculates the minimal control inputs to ensure the structural accessibility of its complementary network, G^c. Structural accessibility, a nonlinear extension of structural controllability within network control, requires that each vertex in the complementary graph G^c is reachable via a path originating from a control input node, and all species nodes are enveloped by a disjoint union of cycles and paths. Subsequent steps involve calculating the control inputs to drive the system to a desired state, utilizing Model Predictive Control (MPC) (Holkar & Waghmare, 2010), with optimization problems often solved via algorithms like DIRECT (Jones, Perttunen, & Stuckman,

1993).

In microbial community management, two prevalent control paradigms exist: continuous control, employing prebiotics and bacteriostatic agents, and impulsive control, which integrates transplantation and bactericides at discrete interventions. Despite their operational differences, both paradigms can exert analogous effects on microbial communities, though impulsive control schemes are typically more straightforward to compute.

The theory of structural controllability has garnered extensive attention in the realm of network science, particularly due to its applicability across diverse domains, including biological and microbial ecosystems. In these contexts, the concept of full control is often nuanced, as the goal typically shifts towards guiding the network to naturally emergent end states, rather than any arbitrary configuration. Within the framework of biological systems, it is prudent to target these innate equilibria, reflective of the system's inherent stability and functionality.

In (Zañudo, Yang, & Albert, 2017), a nuanced control methodology, termed as Feedback Vertex Set Control (FVC), has been posited. This approach deviates from traditional strategies that rely on external control inputs. Instead, FVC posits that by strategically intervening at a particular subset of vertices-specifically, those that form a feedback vertex set (FVS)-one can effectively channel the network's trajectory towards its dynamic attractors. Herein, nodes are categorized into: (1) internal nodes, which are part of the network's interactive dynamics, and (2) source nodes, which operate autonomously of the internal node variables.

The dynamics governing these two node types are succinctly captured by the following equations:

$$\frac{d}{dt}x_i(t) = F_i(X_i, X_{I_i}, t) \tag{6.2}$$

$$\frac{d}{dt}S_j(t) = G_j(t) \tag{6.3}$$

Here, i ranges from 1 to $N - N_s$, denoting internal nodes, while j spans from $N - N_s + 1$ to N, indicating source nodes. The term N_s corresponds to the count

of source nodes within the network, and N signifies the total number of nodes. The set I_i comprises the predecessor nodes of node i, whose dynamics are characterized by $F_i(X_i, X_{I_i}, t)$, while the source node j's behavior is determined by $G_j(t)$.

The fundamental premise of FVC rests on the assertion that the state of a network can be directed towards any of its inherent attractors by manipulating the FVS, a subset of nodes intersecting all feedback loops in the network. Studies (Glass & Kauffman, 1973) (Thomas, 1978) illuminate that such an intervention strategy is not only feasible but also potent in steering the network's dynamics. In a robust validation of this theory, references (Fiedler, Mochizuki, Kurosawa, & Saito, 2013) (Mochizuki, Fiedler, Kurosawa, & Saito, 2013) provide a mathematical underpinning, asserting that in networks defined by nonlinear dynamics-as encapsulated by equation 6.2-the judicious control of the FVS ensures that the network will asymptotically converge to the desired naturally occurring state, irrespective of the explicit form of F_i.

The network control paradigms explored herein hinge on the availability and the fidelity of a mathematical model representing the underlying network. The quality of these models is paramount, as they serve as the foundation for any control strategy. However, the subsequent sections will delve into techniques that mitigate the reliance on such models or require minimal modeling, thereby expanding the toolbox available for network control in scenarios where high-quality models are scarce or altogether absent.

Machine learning, specifically Reinforcement Learning (RL) (Sutton & Barto, 2018), has emerged as a transformative approach in the domain of control systems. Reinforcement Learning, positioned at the intersection of machine learning and optimal control, involves an agent learning to make decisions by interacting with an environment to achieve a certain goal.

The crux of RL lies in the agent's ability to learn from experiences to maximize a cumulative reward. In the context of network control, RL can be employed to dynamically adapt control strategies based on the real-time feedback received from the network's state transitions. The agent, typically modeled as a controller, dis-

cerns through trial and error the optimal actions that drive the network towards the desired states, navigating the complexities of the network's dynamics without a comprehensive model.

Applications of RL in control have been profound and multifaceted (Bertsekas, 2019). For instance, in microbial network control, RL can facilitate the discovery of intervention strategies that are otherwise not apparent or intuitive. By iteratively interacting with the network and observing the outcomes of various control inputs, RL algorithms can infer patterns and develop strategies that optimize the microbial community's health and stability. This is particularly beneficial in scenarios where the network exhibits non-linear and unpredictable behavior, and traditional control methods may fall short.

Moreover, RL's model-free nature, where no explicit model of the environment is needed, is a significant advantage. This aligns seamlessly with the earlier discussion on control techniques that necessitate minimal reliance on detailed mathematical models. RL algorithms, through their iterative learning process, can implicitly capture the network dynamics and adjust the control policy accordingly, potentially leading to novel and effective control methodologies for complex networks.

In conclusion, the introduction has laid a foundational understanding of network control theories and methodologies, highlighting the promising avenue of Reinforcement Learning in complex system management. With its ability to adapt and learn from dynamic environments, RL stands as a pillar for the development of sophisticated control mechanisms in domains as intricate as microbial networks.

In the forthcoming sections, we will delve into the specifics of Reinforcement Learning, elucidating its principles, algorithms, and the nuances that make it particularly suited for microbial network control. We will dissect the RL framework, its learning process, and how it can be applied to achieve desired control objectives within these biological systems. Additionally, a working example will be presented to illustrate the practical application of Reinforcement Learning, providing insights into its operational effectiveness and the potential it holds for advancing the field of microbial network management.

Through this exploration, we aim to bridge the gap between theoretical models and real-world applications, showcasing how Reinforcement Learning can serve as a transformative tool in the quest to control and optimize microbial communities. The integration of these advanced control strategies promises to enhance our ability to manipulate and stabilize these complex networks, paving the way for innovative solutions to ecological and biomedical challenges.

6.2 Reinforcement Learning

6.2.1 General Philosophy

Reinforcement learning is a sub-field of machine learning which is much more focused on target-oriented from interaction (Sutton & Barto, 2018). It is a study of learning how to map situations to actions and the goal is to maximize a numerical reward signal. During the process of trial-and-error, the learner cumulates experience from not only immediate reward of its actions but all the subsequent rewards; delay rewards and trial-and-error are the two key features of reinforcement learning comparing to other learning processes. We define the key elements of reinforcement learning: states, actions, and rewards along with the interaction between the learning agent and its external environment in the framework of Markov decision process (MDP), where a Markov decision process is a discrete-time stochastic control process. At every time step, the agent receives a feedback, usually in the form of scalar reward, from the interacting environment and choose an action to respond to the signal. The environment corresponds to the agentâs action with a feedback given by a reward function. Usually, in the framework of MDP, the agent-environment interaction can be understood in the figure 6.1.

where $S_t \epsilon \mathcal{S}$ is environment state at time t, R_t is the reward received by the agent at time t, and $A_t \in \mathcal{A}$ is the action taken by the agent at time t after it receives reward R_t. Therefore, the Markov decision process and the agent along discrete time steps, $t = 0, 1, 2, 3, \ldots$, will give us a trajectory in the form of $S_0, A_0, R_1, S_1, A_1, R_2, \ldots$

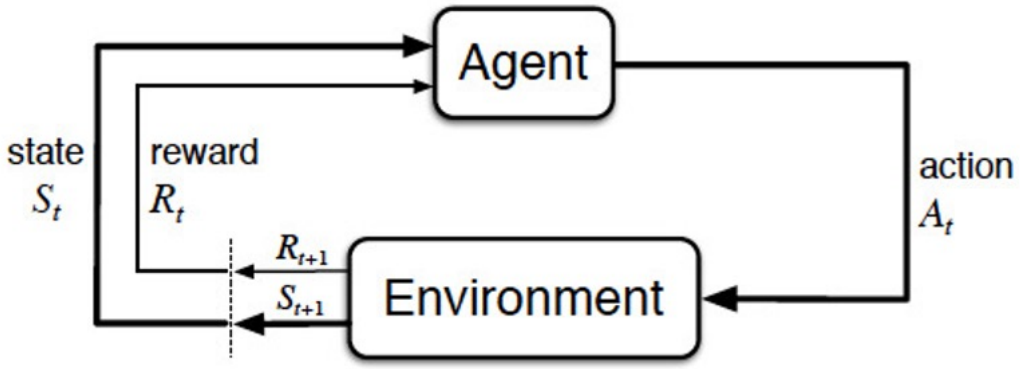


Figure 6.1: Agent-environment interaction

The dynamics of the MDP is given by the function below:

$$p(s', r \mid s, a) := \Pr\{S_t = s', R_t = r \mid S_{t-1} = s, A_{t-1} = a\} \tag{6.4}$$

for all $s, s' \in \mathcal{S}, a \in \mathcal{A}$, and $r \in \mathbb{R}$ and thus Markov property is satisfied. The flexibility of expressing the problem in the context of Markov decision process is guaranteed by the fact that both states and actions can be flexible in various forms. For example, we could have actions ranging from a next step in a GO game to turning left or right in a self-driving automobile. Similarly, the states can be the current situation in the GO game or the direct sensor readings of the self-driving automobile. As stated in (Sutton & Barto, 2018), actions can be any decision that matches the learning purpose, and states can be anything that can be observed which can also serve the learning purpose.

Reinforcement learning relates the learning purpose to scalar reward signal; in other words, the agents learns the decision by trying to maximize the total cumulative rewards it receives in all time steps. This is called the reward hypothesis of reinforcement learning. More specifically and formally, at time step t, the agent chooses the action from action space $\mathcal{A}$, to maximize the expected discounted return:

$$G_t := R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \cdots = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \tag{6.5}$$

where $0 \leq \gamma \leq 1$ is the discount factor. The discount factor gives the flexibility of how future rewards are seen at present. If $\gamma = 0$, the agent is extremely myopic and only concerns the present reward, disregarding what will happen in the future. If $\gamma = 1$, the agent is extremely farsighted since it takes the future rewards as important as the present reward.

We define a policy to be the guide that the agent follows to make decisions; more formally, a policy is defined as a mapping from states to probabilities of choosing different actions at that state, namely, $\pi(a \mid s)$. The state value function $v_\pi(s)$ and state-value function $q_\pi(s, a)$ under certain a policy π are defined as an estimate of (1) how good it is to be at state s and (2) how good it is to be at state s and take certain action a respectively. More specifically,

$$v_\pi(s) := \mathbb{E}_\pi \{G_t \mid S_t = s\} \quad \forall s \in \mathcal{S} \tag{6.6}$$

$$q_\pi(s, a) := \mathbb{E}_\pi \{G_t \mid S_t = s, A_t = a\} \quad \forall s \in \mathcal{S}, \forall a \in \mathcal{A} \tag{6.7}$$

By considering one step back in the Markov decision process, the Bellman Equation for state-value function is achieved, as in equation 6.8 below.

$$v_\pi(s) = \sum_a \pi(a \mid s) \sum_{s',r} p\left(s', r \mid s, a\right) \left[r + \gamma v_\pi\left(s'\right)\right] \quad \forall s \in \mathcal{S} \tag{6.8}$$

Our final goal is to learn how to make the best decision and thus is to find the best policy, namely optimal policy, π_*. Under the optimal policy, we have the associated optimal state-value function $v_*(s)$ and the optimal action-value function $q_*(s, a)$. The two optimal value functions can be solved by Bellman optimality equations as stated in equations 6.9 and 6.10.

$$v_*(s) = \max_a \sum_{s',r} p\left(s', r \mid s, a\right) \left[r + \gamma v_*\left(s'\right)\right] \quad \forall s \in \mathcal{S} \tag{6.9}$$

$$q_*(s,a) = \sum_{s',r} p\left(s', r \mid s, a\right) \left[r + \gamma \max_{a'} q_*\left(s', a'\right)\right] \tag{6.10}$$

The framework of Markov decision process naturally applies to control theory in engineering since the discrete-time optimal control problem can be easily formalized as MDP (Puterman, 2014). The policy we are trying to learn in reinforcement learning can be seen as a state feedback control law in control theory, and thus the optimal policy corresponds to the optimal control strategy.

6.2.2 Reinforcement Learning Algorithms

As seen in the Bellman optimality equations, if we have linear dynamics defining function p, then analytical solutions can be easily found by simply solving the equations. However, if the dynamics become non-linear, then it is impossible to solve the equations and achieve analytical solutions; instead, numerical algorithms should be used. One method called dynamic programming is basic but sometimes useful; it takes advantage of the properties of MDP that the Bellman equation gives recursive decomposition and the fact that and value function stores and reused solutions. Dynamic programming is extremely useful when optimal solution can be decomposed into subproblems and the subproblems are overlapping. Dynamic programming assumes full knowledge of the MDP. For control purpose, the input is a Markov decision process and the output will be an optimal value function and an optimal policy by policy iteration. Policy iteration is a method that turns the Bellman equation into an assignment process followed by policy improvement by acting greedily. The process can be clearly illustrated by the figure 6.2 below as in (Sutton & Barto, 2018).

Dynamic programming requires all knowledge about the model and dynamics; however, in most control problems, MDP model is unknown or too large to use. With the samples generated from experience, we can use model-free control to learn the optimal policy. The main reinforcement learning algorithms can be divided to two groups: on-policy learning and off-policy learning. Onpolicy learning is learning about policy π from experience sampled from π; off-policy learning is learning

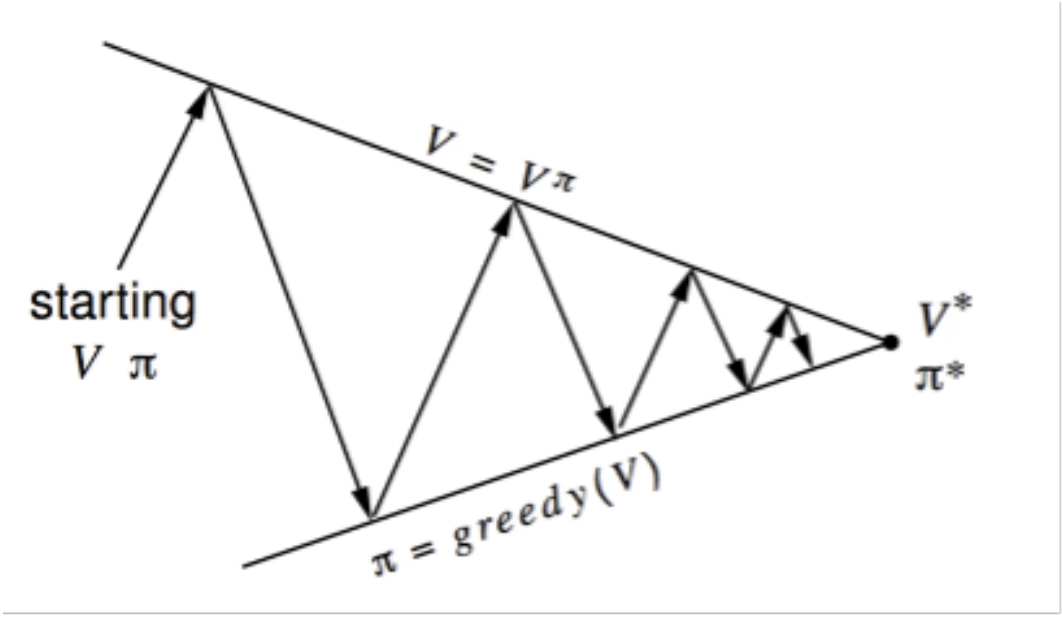


Figure 6.2: Policy Iteration

about policy π from experience sample from another policy μ. Monte-Carlo control is one on-policy learning method such that in every episode, we use Monte-Carlo policy evaluation followed by policy improvement using $\epsilon-$ greedy policy. Another class of algorithms is called temporal difference learning; it has advantages over Monte-Carlo learning since it will have lower variance and it is online. The algorithm below shows the pseudocode of the Sarsa algorithm for onpolicy control which is one of the most famous reinforcement learning algorithms.

Algorithm 3 Sarsa algorithm

Initialize $Q(s,a), \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$ arbitrarily, and $Q(\text{terminal-state}, \cdot) = 0$
repeat (for each episode)
    Initialize S
    Choose A from S using policy derived from Q (e.g., ε-greedy)
    repeat (for each step of episode)
        Take action A, observe R, S'
        Choose A' from S' using policy derived from Q (e.g., ε-greedy)
        $Q(S,A) \leftarrow Q(S,A) + \alpha\left[R + \gamma Q\left(S', A'\right) - Q(S,A)\right]$
        $S \leftarrow S'$
        $A \leftarrow A'$
    until S is terminal
until

The Q-learning algorithm can be used to learn the optimal policy no matter what

policy is actually applied to control the system; this off-policy learning algorithm is used extensively. The following algorithm shows the pseudocode of Q-learning.

Algorithm 4 Q-learning algorithm

Initialize $Q(s,a), \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$ arbitrarily, and $Q(\text{terminal-state}, \cdot) = 0$
repeat (for each episode)
    Initialize S
    repeat (for each step of episode)
        Choose A from S using policy derived from Q (e.g., ε-greedy)
        Take action A, observe R, S'
        $Q(S,A) \leftarrow Q(S,A) + \alpha \left[R + \gamma \max_a Q\left(S', a\right) - Q(S,A)\right]$
        $S \leftarrow S'$
    until S is terminal
until

In many temporal learning algorithms, two methods are often used to increase the convergence rate The first one is called eligibility trace (Singh & Sutton, 1996). Intuitively, the latest transition gives credit to the entire trajectory at time steps before; therefore, marking visited state-action pair along the trajectory would be helpful. Eligibility traces give more credit to more recent visited state-action pairs and the state-action pairs become less eligible as they move further into the past. The second method is called experience replay, which enables us to store all the transitions in memory and replay them when needed. Both methods guarantee the asymptotic convergence to the optimal policy in both Sarsa and Q-learning algorithms.

The methods we mentioned above require both states and actions to be discrete; however, in realworld control problems, for example impulsive network control strategy, states and actions are both continuous. In this case, we need to approximate the Q-function or policies. There are many function approximators available, such as linear combination of features and neural network. With the help of the gradient descent method, temporal difference learning algorithms with function approximators are called gradient TD and gradient Q-learning. Both methods follow true gradient c projected Bellman error. The fitted Q-iteration and Least-squares policy iteration are two popular offline approximate reinforcement learning algorithms which become model-free when samples ar obtained; still, knowledge of the system can be used to generate data samples.

In a way, reinforcement learning can be viewed as a process of information acquisition with the hope that a strategy can be developed effectively based on the information that is obtained about th underlying system. To some extent, the need for efficient information acquisition is also shared by other approaches, for example for building mathematical models to be used in a model-based control approach. In the following, we introduce the area of optimal experimental design, which is highly relevant to information acquisition as needed in various types of approaches for (microbial) network control.

6.3 Case Example: A Model-Free RL Approach to Network Control

In the following case example section, we illustrate a practical application of a model-free Reinforcement Learning (RL) approach to network control. This method is particularly pertinent for complex networks such as microbial communities, where the construction of an accurate model is fraught with challenges. By adopting a model-free RL strategy, we sidestep the need for precise system modeling, focusing instead on learning from the system's behavior to inform control decisions.

We revisit a test problem from the study (Angulo et al., 2019), involving a microbial community of three interacting species. The network is represented by three species x_1, x_2, and x_3, with respective colors blue, yellow, and green, and includes a time-dependent control input u influencing the third species (Figure 6.3).

In contrast to the original study, which utilized an ODE model (equations 6.11) to implement the control approach, our RL method does not assume prior knowledge of this model. Instead, the model is used to simulate the system's responses, acting as a virtual stand-in for the actual biological system during the RL agent's iterative learning process. This model-free approach opens new avenues for discovering effective control strategies in complex biological networks.

$$\begin{aligned}
\dot{x_1} &= 0.1 + x_1\left(1 - \frac{x_1}{5}\right)\left(\frac{x_1}{3} - 1\right) - \frac{0.1x_1x_3}{1+x_3} \\
\dot{x_2} &= 0.1 + x_2\left(1 - \frac{x_2}{4}\right)(x_2 - 1) - \frac{x_2x_3}{1+x_3} \\
\dot{x_3} &= x_3\left(1 - \frac{x_3}{2}\right)(x_3 - 1) + u
\end{aligned} \tag{6.11}$$

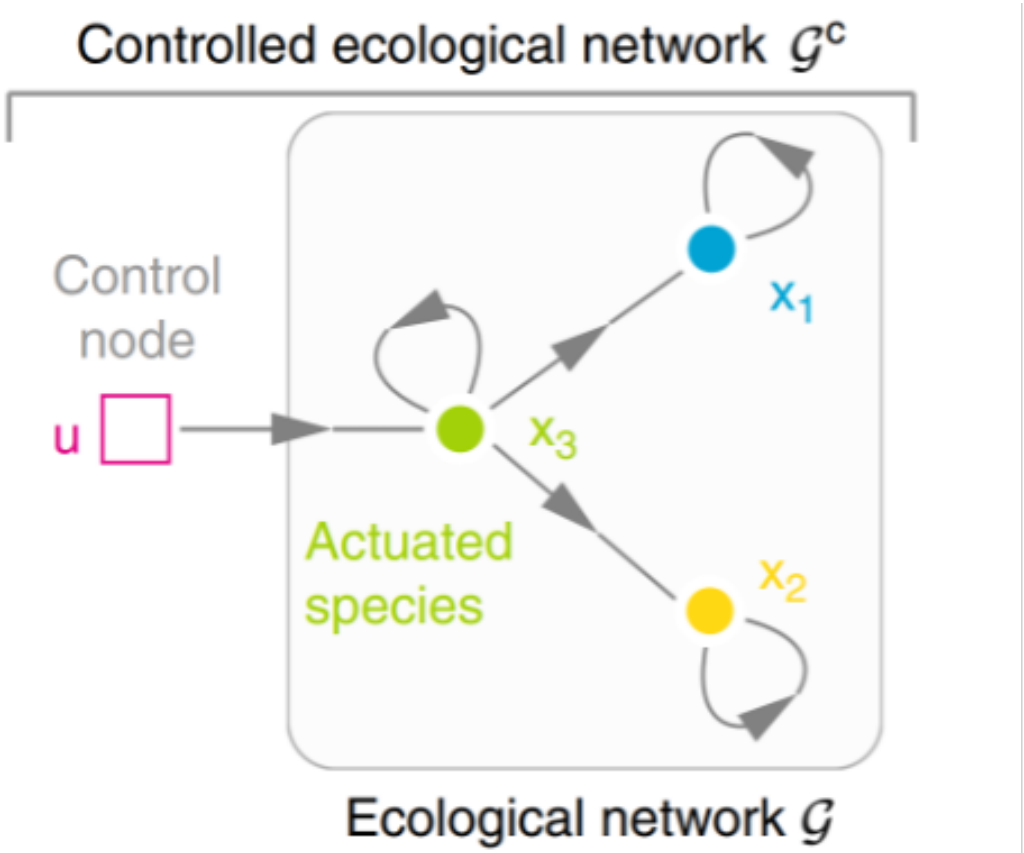


Figure 6.3: Controlled ecological network of test example

As showed by the model equations, one can apply continuous time-dependent signals to x_3 to affect the system. On the other hand, in this study we consider the impulsive control scheme, in which the control inputs are applied to the microbial community only at intervention time instants, say $\mathbb{T} = \{t_1, t_2, t_3, \ldots\}$, instantaneously changing the absolute abundance of controlled species. Therefore, the controlled species dynamics is almost the same as the continuous control scheme case except equation 6.11 is replaced by equation 6.12 below.

$$\dot{x_3} = x_3\left(1 - \frac{x_3}{2}\right)(x_3 - 1) \text{ and } x_3\left(t^+\right) = x_3(t) + u(t) \quad \text{if } t \in \mathbb{T} \tag{6.12}$$

As proved in (Angulo et al., 2019), by mathematical formalism, we can fully control this microbial community by applying control inputs only to the third species; in other words, the third species is a minimum set of driver species. We adopt the definition of control as in the definition in the paper of feedback vertex control (Zañudo et al., 2017). That is, since the test problem is in the framework of microbial networks, we tend to relax the definition of full control since it is reasonable to consider only the naturally occurring end states as desirable final states instead of any arbitrary state. In the framework of optimal control in reinforcement learning, we need to clearly define states, reward, actions and environment for our microbial network test example. Firstly, we define the states to be the abundance of all three species in the network and reward to be only positive when hitting the desired naturally occurring end state. Secondly, the environment is the microbial community modelled by the microbial network and actions to be impulsive control input which increases or deceases the abundance of controlled species, species 3, to a certain amount. The states take continuous values but we can discretize the state space by first restricting it to a finite interval. In this test example, we have three interacting species. The process of discretization will be decomposing a cube; the philosophy is shown in figure 6.4. Here we assume that, within every discretized interval or small cube, the agent should learn a consistent rule.

We then applied Q-learning to the discretized problem, which is schematically shown in figure 6.5.

In the initial test, the algorithm took a very long time for the agent to first reach the target interval, in which the target state lies in. The enormous time taken for the agent to hit the target interval was because the reward was too sparse comparing to the whole state space; the agent with completely no information at the first episode would keep wandering in the discretized space, especially when we discretize the interval into a great number of pieces. Since we are adopting the definition of control as in (Zañudo et al., 2017), we can make the reward much less sparse by

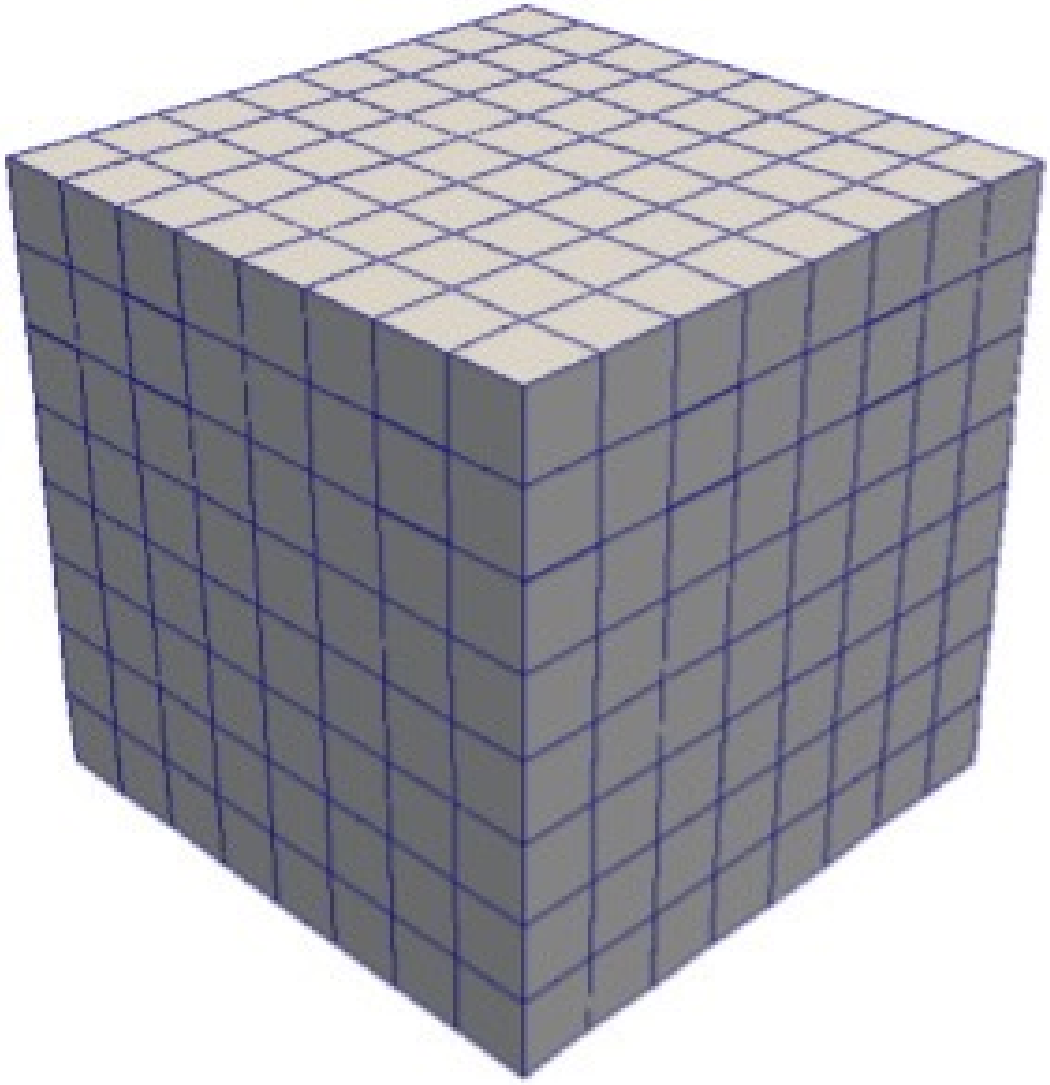


Figure 6.4: Discretization of state space: This figure demonstrates the discretization of a continuous 3D state space into a grid of smaller, discrete cubes. The original 3D cube is partitioned into multiple smaller cubes. Each small cube represents a distinct state that the reinforcement learning agent can occupy within the 3D space. This discretization reduces the complexity of the continuous state space, enabling the agent to efficiently learn and make decisions by navigating through a finite set of states.

putting a positive reward at each interval that will naturally evolve to the desired state without any further intervention and a much more satisfying result is obtained. With the adjustment of rewards, Q-learning converges rapidly in this toy example with multiple starting-desiring states tested.

First, with the state space discretized into 0.1*0.1*0.1 cubes, the agent could not achieve the reward, which is the desired end state, even in the first episode after

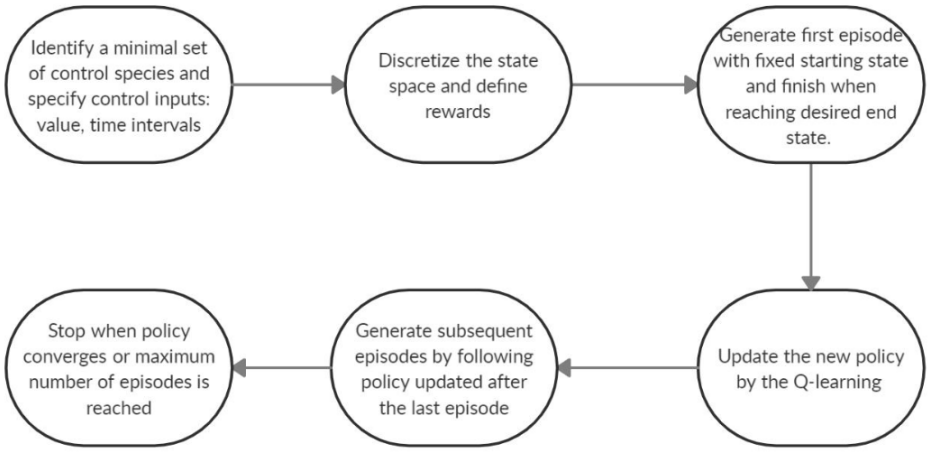


Figure 6.5: RL control scheme

wandering for more than a few hours. As a possible remedy, the state space was discretized into fewer intervals, but it was very hard to keep the assumption that the agent should learn consistent rules in every interval when the interval is not small enough. A few other algorithms were also tested (such as the Sarsa(Î») algorithm) and similar problems occurred. After some research, it was realized that the possible cause maybe the reward was too sparse compared to the size of state space. One of the possible solutions to the problem might be changing the initial policy in the first episode to a policy that is 'closer' to the optimal policy; however, in the context of Markov decision processes of network control, it is very hard to determine which policy is closer to the optimal policy since the differential equations of system dynamics provide no obvious insights.

Since we adopt the relaxed definition of microbial network control, which is driving any initial state to any one of the naturally occurring end states, the problem of sparse rewards could be solved by adding a process of examining the new states when the state was reached. We could place a reward on that state whenever it is reached and whenever it is identified that the state would naturally evolve to the desired end state without any further external manipulations to the system. In the toy example, we defined the actions to be instantaneously increasing or decreasing the abundance of species x_3 by 0.1 at each time interval of 5 standard units. After

restricting the domain of three species abundances to be (0,6) and discretizing the state space into 60^3 cubes, we applied the improved RL strategy for a problem with the starting state as (3.148,4.58,1) and the desired state as (5.137,4.033,0). The policy converged after a few episodes and the performance was reasonably satisfying. The experiments results are shown in figure 6.6.

```
0 [2.9, 4.8, 2.0]
0 [1.4, 4.8, 2.0]
0 [0.2, 4.8, 2.0]
0 [0.1, 4.8, 2.0]
0 [0.1, 4.8, 2.1]
0 [0.1, 4.8, 2.0]
0 [0.1, 4.8, 2.1]
0 [0.1, 4.8, 2.1]
1 [4.4, 4.1, 0.1]
1 [5.1, 4.1, 0.1]
1 [5.1, 4.1, 0.1]
2 [4.4, 4.1, 0.1]
2 [5.1, 4.1, 0.1]
2 [5.1, 4.1, 0.1]
3 [2.9, 4.8, 2.0]
3 [1.4, 4.8, 2.0]
3 [0.2, 4.8, 2.0]
3 [0.1, 4.8, 2.1]
3 [0.1, 4.8, 2.1]
4 [4.4, 4.1, 0.1]
4 [5.1, 4.1, 0.1]
4 [5.1, 4.1, 0.1]
5 [4.4, 4.1, 0.1]
5 [5.1, 4.1, 0.1]
5 [5.1, 4.1, 0.1]
6 [4.4, 4.1, 0.1]
6 [4.4, 4.1, 0.1]
7 [4.4, 4.1, 0.1]
7 [4.4, 4.1, 0.1]
8 [4.4, 4.1, 0.1]
8 [4.4, 4.1, 0.1]
9 [4.4, 4.1, 0.1]
9 [4.4, 4.1, 0.1]
10 [4.4, 4.1, 0.1]
10 [4.4, 4.1, 0.1]
11 [4.4, 4.1, 0.1]
11 [4.4, 4.1, 0.1]
12 [4.4, 4.1, 0.1]
12 [4.4, 4.1, 0.1]
13 [4.4, 4.1, 0.1]
13 [4.4, 4.1, 0.1]
```

Figure 6.6: Experiment results of test example. In each row, the first number denotes its episode number and $[x_1, x_2, x_3]$ denotes the states that agent has reached in that episode

It is clear that the policy converges in about 10 episodes and the corresponding policy is reducing the abundance of x_3 by 0.1 instantaneously as time 5 , which is very simple to our surprise. The effectiveness is confirmed by the simulated system dynamics, as shown in figure 6.7.

Compared to the control schemes calculated in (Angulo et al., 2019), this control

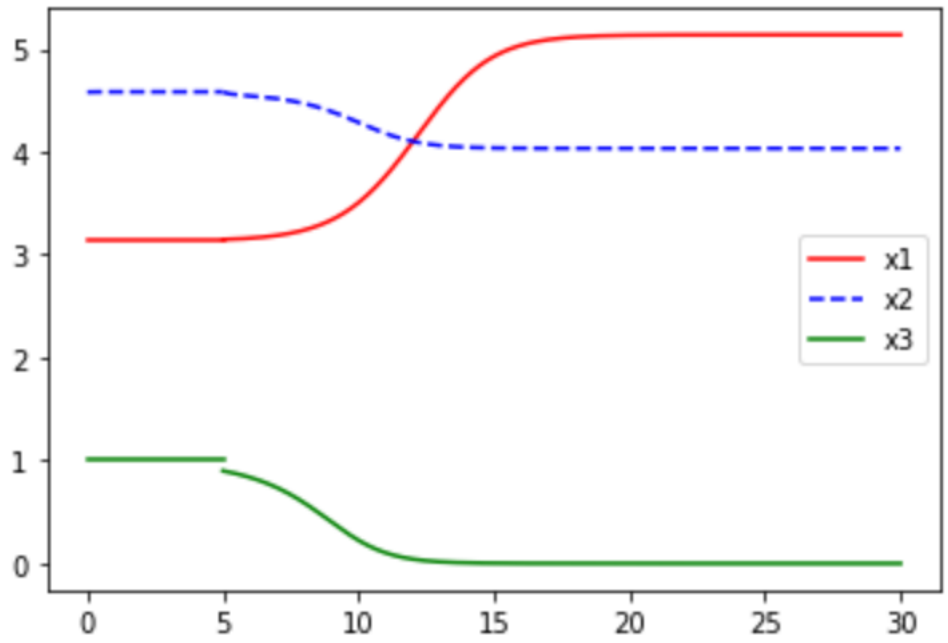


Figure 6.7: Simulated system dynamics to test the control policy learned from RL: the x-axis represents the time and the y-axis represents the absolute abundances of species

scheme is very simple and easy to use. It should be stressed that the mathematical model was used in the simulations to act as the actual (experimental) system for environment updates in the framework of reinforcement learning; this method does not depend on the knowledge of the model to compute the control schemes and is therefore model-free and has great potential when it comes to complex communities where the inferred network model is usually not accurate to some extent.

Moreover, more algorithms can be used for faster convergence in more complex network problems. Additionally, control actions are more flexible, and they can be part of the learning target when function approximation reinforcement learning is applied. For the full set of codes used in this research, please refer to my GitHub repository at:https://github.com/AFFFPupu/Microbial_RL_Control.git.

6.4 Conclusion

In summary, the exploration of reinforcement learning within the realm of network control has yielded promising outcomes, particularly when challenges such as

sparse rewards are effectively addressed. The key to success in these applications lies in the careful configuration of the reward system. The inherent flexibility of reinforcement learning allows for the design of versatile control schemes, adaptable to the continuous nature of state and action spaces often encountered in network control problems. This adaptability can be harnessed through function approximation methods or through strategic discretization of these spaces.

The work presented in this chapter demonstrates the feasibility of employing reinforcement learning for the control of microbial networks. It stands as a testament to the potential of RL in this field, while simultaneously acknowledging the necessity for further development. Future research should focus on refining RL methodologies, enhancing their effectiveness, and benchmarking their performance against alternative approaches. As we continue to advance our understanding and capabilities in this area, reinforcement learning promises to be a valuable asset in the pursuit of more sophisticated and robust network control strategies.

References

Angulo, M. T., Moog, C. H., & Liu, Y.-Y. (2019). A Theoretical Framework for Controlling Complex Microbial Communities. *Nature Communications*, *10*(1), 1045.

Bertsekas, D. (2019). *Reinforcement Learning and Optimal Control*. Athena Scientific.

Fiedler, B., Mochizuki, A., Kurosawa, G., & Saito, D. (2013). Dynamics and Control at Feedback Vertex Sets. I: Informative and Determining Nodes in Regulatory Networks. *Journal of Dynamics and Differential Equations*, *25*, 563–604.

Glass, L., & Kauffman, S. A. (1973). The Logical Analysis of Continuous, Nonlinear Biochemical Control Networks. *Journal of Theoretical Biology*, *39*(1), 103–129.

Holkar, K., & Waghmare, L. M. (2010). An Overview of Model Predictive Control. *International Journal of Control and Automation*, *3*(4), 47–63.

Jones, D. R., Perttunen, C. D., & Stuckman, B. E. (1993). Lipschitzian Optimization Without the Lipschitz Constant. *Journal of Optimization Theory and Applications*, *79*, 157–181.

Liu, Y.-Y., Slotine, J.-J., & Barab'asi, A.-L. (2011). Controllability of Complex Networks. *Nature*, *473*(7346), 167–173.

Mochizuki, A., Fiedler, B., Kurosawa, G., & Saito, D. (2013). Dynamics and Control at Feedback Vertex Sets. II: A Faithful Monitor to Determine the Diversity of Molecular Activities in Regulatory Networks. *Journal of Theoretical Biology*, *335*, 130–146.

Puterman, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.

Singh, S. P., & Sutton, R. S. (1996). Reinforcement Learning With Replacing Eligibility Traces. *Machine Learning*, *22*, 123–158.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT press.

Thomas, R. (1978). Logical Analysis of Systems Comprising Feedback Loops. *Journal of Theoretical Biology*, *73*(4), 631–656.

Zañudo, J. G. T., Yang, G., & Albert, R. (2017). Structure-based Control of Complex Networks With Nonlinear Dynamics. *Proceedings of the National Academy of Sciences*, *114*(28), 7234–7239.

7 | Conclusion and Future Work

7.1 Summary of Work and Key Contributions

This thesis has embarked on a comprehensive exploration of microbial communities, employing the lenses of network theory and control to unravel the complex interactions within these systems. The journey began with the aim to assess and refine methodologies for microbial network reconstruction, extend the frontiers of statistical causal inference within microbial dynamics, and innovate control strategies through reinforcement learning. The culmination of this endeavor has yielded several achievements, each aligning with the outlined objectives and collectively propelling the field of microbial network analysis forward.

Insights into Methodological Advancements A critical examination of existing microbial network reconstruction methods laid the foundation for this thesis. Through a comparative analysis, this work has illuminated the diverse methodological landscape, articulating the merits and demerits inherent in each approach. This scrutiny has not only served as a guide for researchers selecting appropriate methods for their studies but has also highlighted the urgent need for methods that adeptly handle feedback loops-a common feature in microbial networks yet often overlooked in traditional analyses.

Microbial Causal Inference (MCI) Framework At the core of this thesis is the Microbial Causal Inference (MCI) framework, a pioneering tool engineered to decipher the causal interplay among microbial species. The MCI framework stands as a testament to the synergy of theory and application, demonstrating its robustness across various scenarios, particularly those characterized by intricate feedback mechanisms. The development and validation of MCI epitomize the thesis's commitment to advancing our understanding of microbial ecosystems.

Experimental Design and Interventional Data The thesis has also made strides in experimental design strategies, advocating for the procurement of interventional data to refine network models. This forward-thinking approach underscores the im-

portance of empirical validation in strengthening theoretical models, ensuring that they are not only mathematically sound but also biologically relevant and applicable.

Reinforcement Learning for Network Control The thesis ventured into the domain of reinforcement learning, showcasing its efficacy in the strategic control of microbial networks. The innovative use of model-free discretized Q-learning has set a precedent for adaptive microbial network management, offering an avant-garde approach that promises transformative applications across biomedicine, agriculture, and environmental management.

7.2 Limitations and Future Work

This thesis has made several key contributions to the field of microbial network analysis; nonetheless, it is important to acknowledge certain limitations that suggest avenues for future research. Firstly, while the microbial network reconstruction methods presented herein are conceptually robust, they would benefit from more extensive computational validation, involving both simulated and real-world datasets. The diversity of underlying assumptions and intended applications for these methods renders a comprehensive computational comparison challenging, compounded by the current absence of a standard benchmark dataset. Secondly, the Microbial Causal Inference (MCI) framework proposed requires a substantial volume of data for effective causal inference, leading to considerable computational demands. Thirdly, the experimental designs aimed at elucidating the precise structure of directed networks within observational Markov Equivalence Classes may entail significant experimental resources, particularly when dealing with large-scale networks with numerous nodes and cycles. Lastly, in the realm of microbial network control, the adoption of continuous state control policies may offer a viable alternative to discretization, especially when managing large and complex networks. These considerations highlight the need for ongoing development and refinement of analytical tools in this domain.

The exploration of microbial communities through network theory and control has yielded promising results, yet it also unveils a frontier teeming with opportunities for further research. This thesis has laid the groundwork for a myriad of directions that future work may take to deepen and expand our understanding of microbial networks. The following are some of the key areas identified for prospective investigation:

Integration of Multi-omics Data Future studies could benefit from incorporating multi-omics data, such as metatranscriptomics, metaproteomics, and metabolomics, to provide a more comprehensive view of microbial communities. The integration of these data types will allow for a more nuanced understanding of microbial functions and their interactions, moving beyond mere compositional analysis.

Real-Time Network Control The application of reinforcement learning in microbial network control opens avenues for real-time control applications. Future work could focus on developing algorithms capable of adjusting control strategies dynamically, in response to changes within the microbial community, thereby enabling more responsive and adaptive network management.

Addressing the Complexity of Microbial Interactions Despite advancements, the complexity of microbial interactions remains a challenge. Future research should aim to develop novel computational tools that can handle the non-linearity, high-dimensionality, and sparsity inherent in microbial data. These tools would need to be scalable and robust, capable of deciphering complex ecological networks from large-scale datasets.

Experimental Validation of Theoretical Models There is a critical need for experimental validation of the theoretical models proposed by network inference methods. Future work should prioritize the design and implementation of experiments that can test the predictions made by these models, thereby closing the loop between theory and practice.

Accessibility and Usability of Computational Tools To maximize the impact of research in microbial network analysis, it is essential to develop user-friendly com-

putational tools that are accessible to researchers across various disciplines. Efforts should be made to create intuitive interfaces and comprehensive documentation to make advanced computational methods more accessible to non-specialists.

Benchmark Datasets for Method Validation The establishment of benchmark datasets, representative of various microbial ecosystems, is essential for method validation. Future work should focus on creating and curating these datasets, which would enable the comparison of different network inference methods and facilitate the standardization of methodological approaches.

These future directions not only promise to enhance the capabilities of microbial network analysis but also hold the potential to revolutionize our approach to understanding and manipulating microbial ecosystems for the benefit of human health, agriculture, and environmental sustainability.