

ABSTRACT

KAGY, BRYSON GRAHAM. Algebraic and Combinatorial Problems in Mathematical Phylogenetics. (Under the direction of Seth Sullivant).

Algebraic statistics uses tools from algebra, geometry, and combinatorics to study problems that are statistical in nature. Typically this takes the form of using the inherent combinatorial or algebrogeometric structure of a statistical model, such as a polytope or variety, to prove properties about the model. In this thesis, we will focus on problems in algebraic statistics originating from phylogenetics.

First we explore the combinatorial structure of phylogenetic split networks. Phylogenetic networks are generalizations of trees that allow for the modeling of non-tree like evolutionary processes. Split networks give a useful way to construct networks with intuitive distance structures induced from the associated split graph. We explore the polyhedral geometry of distance matrices built from circular split systems which have the added property of being equidistant. We give a characterization of the facet defining inequalities and the extreme rays of the cone of distances that arise from an equidistant network associated with any circular split network.

Next, we prove identifiability results for group-based phylogenetic mixture models. Phylogenetic mixture models are convex combinations of simpler phylogenetic models that can mimic the properties of more complex phylogenetic models. Conditions for the identifiability of phylogenetic mixture models under a general model of DNA evolution known as the General Markov Model (GMM) are given in [80]. We expand these conditions to make statements about identifiability of some specific group-based models, specifically the Jukes-Cantor (JC), Kimura 2 parameter model (K2P), Kimura 3 parameter model (K3P) and the Strand Symmetric model (SSM). These models are simpler than then GMM and so are favored by biologists for computational efficiency, but the increased symmetry makes the arguments more complicated.

Last, we explore the properties of a powerful new statistical test and apply it to phylogenetic models. We explore the problem of performing hypothesis testing for semialgebraic statistical models whose geometry is not conducive to the use of the likelihood ratio tests. In particular we focus on the application of the inference method from [89], (SDL test) to the network multi-species coalescent model and the Cavender-Farris-Neyman (CFN) 2-state model. We investigate how implementation choices such as constraint specification, kernel order, and decomposition into reducible components affect test performance. Our work offers practical insights for researchers applying the SDL test to semialgebraic models, particularly in biological settings where singularities are common.

© Copyright 2025 by Bryson Graham Kagy

All Rights Reserved

Algebraic and Combinatorial Problems in Mathematical Phylogenetics

by
Bryson Graham Kagy

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Mathematics

Raleigh, North Carolina
2025

APPROVED BY:

Nathan Reading

Laura Colmenarejo

Martin Helmer

Seth Sullivant
Chair of Advisory Committee

DEDICATION

To Anna.

BIOGRAPHY

Bryson Kagy grew up in Alpharetta, Georgia where he attended Milton High School. He received his undergraduate degrees in Mathematics and Physics from Georgia Institute of Technology in May 2019 and began his Ph.D. at North Carolina State University in August 2019. He will continue his mathematical career as a lecturer at Texas State University. In his spare time, Bryson enjoys board games, cooking, gardening, video games, Youtube videos, and spending time with his two cats Bean and Rosemary.

ACKNOWLEDGEMENTS

There are many people I would like to thank for their help along the way to my PhD. My biggest thanks goes to my advisor, Seth Sullivant. He has been a great mentor, helping me when I struggled, but also letting me find my own way through the jungle of mathematics. He has also given me great feedback on my work and has pushed me to become a better mathematician and writer. He has helped connect me with the wider algebraic statistics community, which has been very supportive and enriching to my career. I am very thankful for all of the support he has given me as an advisor and I'm very grateful that I chose him as an advisor.

I would also like to thank the algebraic statistics community for being so welcoming and for teaching me so much. I have had the great privilege of participating in two semester programs, one at IMSI and one at ICERM, that heavily involved people from the algebraic statistics community, and they have both been very enriching to my career. I would like to especially thank my collaborators David Barnhill, Joe Cummings, Maize Curiel, Bryan Currie, Danai Deligeorgaki, Marina Garrote-Lopez, Elizabeth Gross, Max Hill, Roser Homs Pons, Joseph Johnson, Aida Maraj, Pratik Misra, Udani Ranasinghe, John Rhodes, Seth Sullivant, Teresa Yu, and Joy Zhang. I have learned so much in our collaborations together and I hope to continue to learn more from them in future projects in my career. I would also like to thank my letter writers Elizabeth Gross, Stepan Paul, Nathan Reading, John Rhodes, Seth Sullivant, and my committee members Laura Colmenarejo, Martin Helmer, Nathan Reading, and Seth Sullivant.

I would also like to thank the many friends that I've made over my career in grad school, both at NC State and while participating in the semester programs. I cannot list everyone, but I would like everyone to know that I appreciate everything you've done to make my PhD enjoyable, from the board game nights to the long nights at the white board. I would like to in particular thank from NC State Will Anderson, Reeshad Arian, Ben Daniel, John Darges, Spencer Daughtery, and Joe Johnson and I would like to thank from the semester programs Joe Cummings, Maize Curiel, Danai Deligeorgaki, Aviva Englander, Max Hill, Ikenna Nometa, and Pardis Semnani.

I would also like to thank my long time friends from high school Allie Albright and Austyn Wohlers. Our calls over Discord and Zoom have helped keep me sane, especially during Covid.

I am very thankful for the support my family has given me. My parents have given me so much support and encouragement throughout the years and I really appreciate it. I'd also really like to thank my sister Meredith, and my Aunt Edie.

I would like to thank Bean and Rosemary for their boundless distractions.

Lastly, I would like to thank my partner Anna Johnson. Her love and support has made finishing this PhD so much easier. She has been an amazing friend and partner and I am eternally thankful for everything she does for me.

Part of this research was performed while the author was visiting the Institute for Mathematical and Statistical Innovation (IMSI), which is supported by the National Science Foundation (Grant No. DMS-1929348).

TABLE OF CONTENTS

List of Tables	vii
List of Figures	viii
Chapter 1 INTRODUCTION	1
1.1 Polyhedral Geometry	3
1.1.1 Polytopes and Polyhedra	3
1.1.2 The Minkowski-Weyl Theorem	6
1.1.3 Facets, vertices, extreme rays	6
1.2 Algebraic Geometry	8
1.2.1 Ideals and Varieties	8
1.2.2 Varieties from Maps	10
1.3 Statistics and statistical models	11
1.3.1 Statistical models	11
1.3.2 Identifiability	15
1.3.3 Hypothesis testing	18
1.4 Phylogenetic trees and networks	20
1.4.1 Trees and Networks	20
1.4.2 Distance methods for Trees and Networks	25
1.4.3 Phylogenetic models	30
1.5 Outline of Thesis	35
1.5.1 Equidistant Circular Split Networks	35
1.5.2 Identifiability of Group-Based Mixture Models	36
1.5.3 Semialgebraic Hypothesis Testing with Incomplete U-Statistics: Practical Issues	36
Chapter 2 Equidistant Circular Split Networks	37
2.1 Introduction	37
2.2 Split Networks	39
2.3 Distances from split networks	47
2.4 Facets of EDC_{KN_n}	53
2.5 Extreme Rays of EDC_{KN_n}	59
2.6 The Chan-Robbins-Yuen Polytope	73
Chapter 3 Identifiability of Large Phylogenetic Mixtures for Many Phylogenetic Model Structures	77
3.1 Introduction	77
3.2 Key Definitions and Statement of Main Results	79
3.3 Markov models on trees	82
3.4 The Discrete Fourier Transform	85
3.5 Tensors and Flattenings	88

3.6	A General Theorem on Identifiability of Phylogenetic Mixtures	91
3.7	Rank property, Extended Rank Property, and Standard Property for the Jukes-Cantor Model	98
3.8	The No Shuffling Property	106
Chapter 4 Methodological considerations for semialgebraic hypothesis testing with incomplete U-statistics 109		
4.1	Introduction	109
4.2	The SDL Test	112
4.2.1	Semialgebraic models and hypothesis testing	112
4.2.2	Overview of the SDL test	113
4.2.3	Kernel construction	116
4.2.4	SDL test parameters	117
4.3	Trinomial submodels	118
4.3.1	Basic examples	118
4.3.2	Semialgebraic descriptions of trinomial models	119
4.3.3	SDL rejection regions for trinomial submodels	120
4.3.4	The SDL test of trinomial submodels	121
4.4	Hypothesis Testing and Inference of Phylogenetic Trees	129
4.4.1	The CFN model and its semialgebraic descriptions	129
4.4.2	Data simulation	132
4.4.3	SDL test parameters and hypotheses	133
4.4.4	Hypothesis tests results	134
4.4.5	SDL-Based Phylogenetic Inference	137
4.5	Implementation Details and Computational Performance	140
4.6	Conclusion	141
References		144
APPENDICES		151
Appendix A	Coalescent models	152
A.1	The multispecies coalescent model	152
A.2	Deterministic tests	154
A.3	Additional Details on the CFN model	154
A.3.1	Generating sets for the CFN ideal	154
A.3.2	Additional results for Collection 1: Comparison of different constraint sets	160
A.3.3	Lack of long branch attraction bias	163
A.4	Technical Assumptions	165

LIST OF TABLES

Table A.1	Estimated tree topologies for the three methods SDL, MLE and SVD and the three topologies 12 34, 12 34 and 12 34 in the entire treespace of Fig. 4.9 and in the Felsenstein Zone. The <i>undecided</i> column reports the percentage of times that MLE fails to distinguish between topologies.	165
-----------	---	-----

LIST OF FIGURES

Figure 1.1	The two cones discussed in Example 1.1.6 and Example 1.1.9. Here $x = x_1$, $y = x_2$ and $z = x_3$. On the left is the cone C_1 which is defined by $x_1, x_2, x_3 \geq 0$ and $x_1 \geq x_2 \geq x_3$ and on the right is C_2 where $C_2 = \text{cone}((1, 1, 1), (1, 0, 1))$.	5
Figure 1.2	The graph described in Example 1.4.3 with the path (1, 5, 6, 7) highlighted in red.	21
Figure 1.3	The rooted binary tree described in Example 1.4.7 with root 1 and leaves 4, 5, 6, 7.	23
Figure 1.4	The X-tree with splits described in Example 1.4.10 with root 0 and leaves 1, 2, 3, 4, 5, 6.	24
Figure 1.5	The first step of creating an X-tree that realizes all of the splits in S as described in Example 1.4.10.	25
Figure 1.6	A phylogenetic X-tree with edge weights. The distance between leaves is calculated by summing over the weights on the path between leaves and it is described by δ as defined in Example 1.4.19.	27
Figure 1.7	An equidistant phylogenetic X-tree whose corresponding ultrametric tree metric is described in Example 1.4.25.	28
Figure 1.8	A split system with split weights as described in Example 1.4.28. Each of the different colors corresponds to a different split and the method to turn the split system into this graph is explained in Chapter 2.	30
Figure 1.9	A tree where the probability of observing states at the leaves is described in Example 1.4.35.	32
Figure 2.1	An example of a split graph whose corresponding functions are color-coded and described in Example 2.2.5. Here α is blue, β is orange, γ is red, and all other function values are black	41
Figure 2.2	The unique tree representing the split system $N = \{12 3456, 1265 34, 1234 56\} \cup \{i [6] \setminus i : i \in [6]\}$ in Example 2.2.9, which must exist by Theorem 2.2.8.	42
Figure 2.3	The tree representing the split system $N = \{12 3456, 1265 34, 1234 56\} \cup \{i [6] \setminus i : i \in [6]\}$ in Example 2.2.9. The process for constructing this tree is shown in Algorithm 2.2.11	43
Figure 2.4	A visualization of performing Algorithm 2.2.11 on the split system $N = \{12 3456, 16 2345, 123 456, 126 345\}$ from Example 2.2.12.	43
Figure 2.5	Applying Algorithm 2.2.11 to the split system $N = \{123 456, 126 345, 156 234\}$ from Example 2.2.13 by applying the splits in two different orders, resulting in different graphs	44
Figure 2.6	The Dual Polygon Representation of the Split System $N = \{01 2345, 12 0345, 0145 23, 0123 45\}$ from Example 2.2.16. Then Algorithm 2.2.11 is applied to create a split network representing N . Lastly, the split network is drawn rooted at 0.	46

Figure 2.7	The Hasse diagram for the poset that is formed from the split system $N = \{01 234, 12 034, 014 23, 04 123, 012 34\}$ from Example 2.2.18.	47
Figure 2.8	This weighted tree from Example 2.3.3	48
Figure 2.9	The circular split system from Example 2.3.9	50
Figure 2.10	The equidistant split network from Example 2.3.11	51
Figure 2.11	An example X diagram for a point δ described in Example 2.5.12.	66
Figure 2.12	The X -diagram on the left is not valid as it violates multiple implications as described in Example 2.5.16. The X -diagram on the right is valid and is obtained by starting with the one on the left and following all rules in Lemma 2.5.15.	68
Figure 2.13	An example of the method to find an extreme ray for an S such that $b_i r_\tau = 0$ for all $i \in S$ as described in Lemma 2.5.18.	70
Figure 3.1	A phylogenetic model with 3 leaves and two internal vertices a and b and a transition matrix	84
Figure 4.1	Parameter spaces (blue line segments) of four submodels of the trinomial model, with parameter space Δ^2 . The submodels capture the form of the quartet Concordance Factor if the species relationships have specific features, as described in the text.	118
Figure 4.2	SDL test behaviour for Model 1, with $m = 1, 5,$ and 15 (top row to bottom). The left column shows nominal vs. empirical sizes for the SDL and LR tests; the middle, histograms of p -value differences; and the right, SDL rejection regions.	121
Figure 4.3	Rejection regions for Model 2 under the SDL test using (L to R) a) the constraints $y - z \leq 0, z - y \leq 0,$ and $1/3 - x \leq 0$; b) replacing the last inequality by $2/3 - x - y \leq 0$; c) including $r = 10$ random convex combinations of the inequalities of (a) ; and d) including $r = 100$ random convex combinations.	123
Figure 4.4	Rejection regions for Model 3 under the SDL test using (L to R) $s = 1, 10,$ and 100 random permutations to partially symmetrize \check{h} . For all, $m = 15$	125
Figure 4.5	(L) Rejection region for Model 4 obtained from SDL test using semialgebraic description given above. (R) Rejection region for an Intersection-Union test using the SDL tests for the 3 irreducible components of Model 4 (each essentially Model 2).	126
Figure 4.6	Rejection regions for SDL tests of (L-R) (a) the Hardy-Weinberg 2-allele model defined by $y^2 - 4xz = 0,$ (b) a nodal cubic model defined by $(y - 1/3)^2 - 6(x - 2/5)^2(x - 1/9) = 0,$ (c) a cuspidal cubic model, defined by $(y - 1/3)^2 - (x - 1/3)^3 = 0.$	127

Figure 4.7	Rejection regions for SDL tests of the cuspidal cubic (L-R) with (a) constraints supplemented by $1/3 - x \leq 0$; (b) constraints supplemented by the inequality from (a) plus $r = 10$ random convex combinations of inequalities, and (c) constraints supplemented by 3 linear inequalities as described in the text and $r = 10$ random convex combinations of inequalities.	128
Figure 4.8	The 4-leaf binary tree topologies, with edge lengths t_i . The names $T_{x y zw}$ indicate the partition of leaves induced by the central edge.	130
Figure 4.9	Left: The tree $T_{12 34}$ with edge lengths $t_1 = t_3 = a$ and $t_2 = t_4 = t_5 = b$, in units of expected number of substitutions per site. Right: The tree space, with a, b varying from 0 to 1.2. In red, nine parameter pairs with $a, b \in \{0.05, 0.2, 0.8\}$. The dashed blue curve is the lower boundary of the Felsenstein zone, defined by $\theta(b)^2 - 2\theta(a) + \theta(a)^2 > 0$ for $\theta(t) = e^{-2t}$ [46].	133
Figure 4.10	Aggregated p -values for a test of the true null hypothesis $H_{12 34}$ (left) and a false null hypothesis $H_{13 24}$ (right) for datasets in Collection 1. Constraints sets CDM and PDM, and number of convex combinations $r = 0$ and 20 are varied.	135
Figure 4.11	p -values obtained from the SDL test on Collection 2 for different constraint sets: CDM (top 3 rows) and PDM (bottom 3 rows). The hypotheses tested are $H_{12 34}$ (left 3 columns) and $H_{13 23}$ (right 3 columns), with $r = 0$.	136
Figure 4.12	Histograms of p -values when testing $H_{12 34}$ (left two) and $H_{13 24}$ (middle two) showing the effect of including the internal edge constraint. The plot for $H_{14 23}$ is omitted because it is similar to that of $H_{13 24}$. Right: Histogram of $W^{(1)}, \dots, W^{(A)}$, approximating the test distribution \mathcal{T}_c for CDD (red) and CDD + inequality (blue) for aggregate data from 1000 trees with random parameters $a, b \in (0, 1.2]$	138
Figure 4.13	Performance of the SDL test for inferring the tree topology $T_{12 34}$ using different constraint sets and values of m . Left: CDM and PDM constraints with $m = 12$. Right: CDD constraints with $m = 12$, CDD with the inequality of Eq. (4.4.3) and $m = 12$, CDD with the inequality and $m = 30$. Rows vary the number of convex combinations, $r = 0$ and $r = 20$. Grey levels represent the frequency of correctly inferring the topology for edge length pairs (a, b) (black 100%, white 0%).	139
Figure 4.14	Performance of 3 methods of topological tree inference on data from Collection 1: (left) the SDL-based inference method using the CDD constraint set with the internal edge inequality, with $m = 30$ and $r = 20$, (middle) Maximum Likelihood [54], (right) the SVD method.	141
Figure A.1	Gene trees (in red) form within a species tree and network (black ‘tubes’)	153
Figure A.2	Rejection regions for Models 1, 2, 3, 4, and Hardy-Weinberg 2-alleles, using deterministic tests, as described in text, with sample size $n = 300$.	154

Figure A.3	Aggregated p -values for a test of the true null hypothesis $H_{12 34}$ from datasets in Collection 1. Columns correspond to choices of defining polynomials. Rows correspond to the value of r	161
Figure A.4	Aggregated p -values for a test of $H_{13 24}$ (a false null hypothesis) from datasets in Collection 1. Columns correspond to choices of defining polynomials. Rows correspond to the value of r . The test of $H_{14 23}$ produced similar results.	162
Figure A.5	Performance of the SDL test for inferring the tree topology $T_{12 34}$. Columns correspond to different CFN model constraints (CDD, CDM, CDR, PDM, PDR), and rows represent the number of convex combinations used, $r = 0$ and $r = 20$. Grey levels represent the frequency of correctly inferring the topology for edge length pairs (a, b) (black 100%, white 0%).	163
Figure A.6	Histogram of p -values for CDM (left) and PDM (right) for a tree in the Felsenstein zone ($a = 0.8$ and $b = 0.05$) with $n = 10000$ bp and $m = 12$. . .	164

CHAPTER

1

INTRODUCTION

Algebraic statistics is a relatively new area of study that takes tools from algebraic geometry and combinatorics to solve statistical problems. For instance, the study of an algebraic structure like a variety that arises from a statistical model can provide equations and inequalities that any observed probability distributions must follow, and then, be useful for statistical inference.

In this thesis we consider applications of algebraic statistics in the area of mathematical phylogenetics. Phylogenetics is the study of evolutionary relationships between species. One of the core problems in the field is: given information about living species, how can one determine what combinatorial structure, such as tree or graph, best represents the way those species evolve and relate to each other. As mathematicians, this problem is approached in generality: given any possible valid species information, what can be said about the possible evolutionary relationships? There are many approaches to this problem but three main ones will be of focus in this thesis: distance methods, Markov chain models, and the coalescent model.

In Chapter 2, we prove new results using distance methods on split networks. In distance methods, a notion of distance/ dissimilarity is given between the species by some biological means and the goal is to take those distances and determine the evolutionary structures that obey those distances. Polyhedral geometry can play a big part in this analysis, as one

can describe inequalities that the distance functions must obey in order to have a specific combinatorial structure, like a tree or network.

In Chapter 3, we analyze Markov mixture models and prove identifiability results about them. When using Markov chain models, we look at specific spot in the DNA sequence for a group of species. We can then assign a Markov transition matrix to every edge of a proposed combinatorial structure and use this Markov process to determine the probabilities of observing a set of DNA bases given that underlying combinatorial structure. There are many combinatorial properties of these collected matrices one can exploit for analysis and these calculated probabilities arise from a polynomial map, making varieties also useful for analysis.

In Chapter 4, we utilize a new hypothesis testing method on statistical models derived from the coalescent model. When using the coalescent model, we consider an underlying species tree or network and the observed gene structures of those species. Then we perform a recursive process to determine the probability of observing that specific gene structure from those species, given the underlying species tree or network. The calculated probabilities are polynomial equations and thus varieties are very useful in this kind of analysis.

In order to describe the problems considered in this thesis, we provide background material that will be used throughout the main body of the thesis. We start with a discussion of polyhedral geometry in Section 1.1.1 as well as a discussion of V and H representations in Section 1.1.2 and facets and faces in Section 1.1.3. There is a focus in these section on examples of polyhedral cones as all of this material will be used in Chapter 2 on a polyhedral cone. Next, we give background on algebraic geometry. In section 1.2.1, we go over ideals and varieties and in section 1.2.2, we describe how to define varieties from polynomial maps. The background in this section will be used in Chapter 3 and 4. Then we give background on statistics and statistical models in section 1.3.1 which will be useful for all of the proceeding chapters. We continue with two specific topics on statistical models, identifiability in section 1.3.2 and hypothesis testing in section 1.3.3 which will be used in Chapters 3 and 4 respectively. Lastly we give a description of mathematics used in phylogenetics, starting with definitions of phylogenetic trees and networks in section 1.4.1. We then give background on distance methods in Section 1.4.2 which will be used in Chapter 2 and background on Markov phylogenetic models in section 1.4.3 which will be used in Chapters 3 and 4. We end this chapter with a brief description of the research projects in the processing chapter in Section 1.5.

1.1 Polyhedral Geometry

1.1.1 Polytopes and Polyhedra

In this section we give background on polyhedral geometry. Polytopes and polyhedra arise naturally when working on discrete problems, as they generalize classic geometric structures like polygons and platonic solids. The study of polyhedral geometry is often motivated by optimization as the feasible region of a linear program is a polyhedron. Our motivation for studying polyhedra is to use them in a phylogenetics context to test if a set of species has a particular evolutionary structure, by seeing if they satisfy the defining inequalities for the polyhedron. In addition phylogenetic trees themselves can be naturally placed into what is called tree space, or BHV space which is a polyhedral cone complex [14]. In Chapter 2, we define a polyhedron by looking at the set of distance functions between leaves on phylogenetic networks. We first give a set of inequalities that hold on these distance functions and then we reduced to a facet description of the polyhedral cone and found an extreme ray description. We are also able to characterize the faces of this cone as determined by simpler phylogenetic networks. For more information on polytopes and polyhedra see [50, 96]. We start with definitions of convexity and cones, which will be necessary to define polyhedra.

Definition 1.1.1 A set $A \subseteq \mathbb{R}^d$ is convex if for all $x, y \in A$ and all $t \in [0, 1]$ we have $tx + (1-t)y \in A$. Given a set $S \subseteq \mathbb{R}^d$, the convex hull of S , $\text{conv}(S)$, is the smallest convex set containing S . We can represent all points in $\text{conv}(S)$ as a convex combination of the points in S , that is

$$\text{conv}(S) = \left\{ \sum_{i=1}^n \lambda_i x_i : n \in \mathbb{N}, x_i \in S, \lambda_i \geq 0, \sum_{i=1}^n \lambda_i = 1 \right\}.$$

Definition 1.1.2 A set $C \subseteq \mathbb{R}^d$ is a (convex) cone if for all $x_i \in C$, $\sum_{i=1}^n \alpha_i x_i \in C$ for all $\alpha_i \geq 0$. Given a set $C \subseteq \mathbb{R}^d$, the conical hull of C , $\text{cone}(C)$, is the smallest cone containing S . We can represent all points in $\text{cone}(C)$ as a conical combination of the points in S , that is

$$\text{cone}(C) = \left\{ \sum_{i=1}^n \alpha_i x_i : n \in \mathbb{N}, x_i \in C, \alpha_i \geq 0 \right\}.$$

Informally the convex hull of a set of points contains all of the points “in between” the original set of points. The conical hull of a set of points contains all of the points “in between” and all the positive multiples of the original set of points.

Definition 1.1.3 The Minkowski sum of two sets $P, Q \subseteq \mathbb{R}^d$ is

$$P + Q := \{x + y : x \in P, y \in Q\}.$$

A V -polyhedron is a set $P \subseteq \mathbb{R}^d$ that is the Minkowski sum of the convex hull of a finite set of points and the conical hull of a finite set of vectors. That is, a V -polyhedron P can be given in the form

$$P = \text{conv}(V) + \text{cone}(C)$$

for some finite set $V \in \mathbb{R}^d$ and finite set $C \in \mathbb{R}^d$. We call P a V -polyhedral cone if $P = \text{cone}(C)$.

Definition 1.1.4 An H -polyhedron is the intersection of finitely many closed half spaces. That is, an H -polyhedron is a set $P \subseteq \mathbb{R}^d$ of the form

$$P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq \mathbf{z}\}$$

for some $A \in \mathbb{R}^{m \times d}$, $\mathbf{z} \in \mathbb{R}^m$. Here $A\mathbf{x} \leq \mathbf{z}$ is shorthand for the system of inequalities $a_i \mathbf{x} \leq z_i$ where a_i is the i -th row of A and z_i the i -th component of \mathbf{z} . We call P a H -polyhedral cone if $\mathbf{z} = \mathbf{0}$.

The case where a polyhedron is bounded is well studied, so it gets a special name: polytope. Even when working with an unbounded polyhedron, studying related polytopes can often provide insight about the original polyhedron. For example, a polyhedral cone can be intersected with a simple hyperplane, such as the sum of the coordinates being 1, to create a polytope. Then the cone over this polytope is the original cone and understanding the structure of the polytope, tells one about the structure of the original cone. This technique is used in Chapter 2 to understand a polyhedral cone created by distance functions. Below we formally define both V - and H -polytopes.

Definition 1.1.5 An H -polytope is a bounded H -polyhedron. Similarly, a V -polytope is a V -polyhedron where $P = \text{conv}(V)$ for some finite set of vectors $V \in \mathbb{R}^d$.

Example 1.1.6 Let $C_1 \subseteq \mathbb{R}^3$ be a H -polyhedral cone defined by $x_1, x_2, x_3 \geq 0$ and $x_1 \geq x_2 \geq x_3$. Let $C_2 \subseteq \mathbb{R}^3$ be a V -polyhedral cone defined by $C_2 = \text{cone}((1, 1, 1), (1, 0, 1))$. The cones C_1 and C_2 are visualized in Figure 1.1.




Figure 1.1: The two cones discussed in Example 1.1.6 and Example 1.1.9. Here $x = x_1$, $y = x_2$ and $z = x_3$. On the left is the cone C_1 which is defined by $x_1, x_2, x_3 \geq 0$ and $x_1 \geq x_2 \geq x_3$ and on the right is C_2 where $C_2 = \text{cone}((1, 1, 1), (1, 0, 1))$.

Example 1.1.7 Let $A =$

$$\begin{bmatrix} 1 & 0 & 0 \\ -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & -1 \end{bmatrix} \text{ and } z = \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix}.$$

Let

$$P_1 = \{Ax \leq z\}.$$

Then P_1 is a H-polytope known as the unit cube. The matrix representation corresponds to the system of inequalities:

$$0 \leq x_1 \leq 1, \quad 0 \leq x_2 \leq 1, \quad 0 \leq x_3 \leq 1.$$

Let

$$P_2 = \text{conv} \left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\}.$$

Then, P_2 is a V-polytope and is, in fact, also the unit cube. These points are the 8 vertices of the cube. The fact that the cube can be both represented by a V-polytope and an H-polytope is an example of the main theorem from the next section.

1.1.2 The Minkowski-Weyl Theorem

A surprising fact is that V -polyhedra and H -polyhedra are actually the same. That is, anything that is a V -polyhedron is also an H -polyhedron and vice-versa. Henceforth we will often refer to V -polyhedra and H -polyhedra as polytopes or polyhedra when the situation does not require a specific representation.

Theorem 1.1.8 (Minkowski-Weyl) *A set $P \subseteq \mathbb{R}^d$ is a V -polytope if and only if it is an H -polytope. In addition, a set $P \subseteq \mathbb{R}^d$ is a V -polyhedron if and only if it is an H -polyhedron.*

For an in-depth proof, see Chapter 2 of [96]. For certain polyhedra it can be easier to write either the V -representation or the H -representation but finding the other representation is hard. However, since this correspondence always exists, converting to the other representation can lead to useful insights about the problem that defined the polyhedron. For instance, in Chapter 2, the polyhedral cone defined has a much easier to find H -representation but finding the V -representation took a significant amount of work. The extreme rays that were found give useful boundary cases for class of phylogenetic networks that define the polyhedral cone.

Example 1.1.9 *Consider C_1 and C_2 as defined in Example 1.1.6. The cone C_1 has the H -description $x_1 \geq x_2 \geq x_3 \geq 0$. These are not the only inequalities that hold on C_1 , for instance $x_1 + x_2 \geq x_3$ and $x_1 \geq 0$ also hold. Using the notation given in Definition 1.1.4, $C_1 = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{Ax} \leq \mathbf{z}\}$ where*

$$A = \begin{pmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & -1 \end{pmatrix} \text{ and } z = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

The cone C_1 has the V -description $C_1 = \text{cone}((1, 1, 1), (1, 1, 0), (1, 0, 0))$. Each of these rays can be seen as the equality of two of the inequalities from $x_1 \geq x_2 \geq x_3 \geq 0$, with $(1, 1, 1)$, for example, the intersection of $x_1 \geq x_2$ and $x_2 \geq x_3$.

The cone C_2 has the V -description $C_2 = \text{cone}((1, 1, 1), (1, 0, 1))$ and has the H -description $x_2 \geq 0, x_3 \geq x_2, x_3 \geq x_1, x_1 \geq x_3$. Note that $x_3 = x_1$, but this can be written as the two inequalities $x_3 \geq x_1, x_1 \geq x_3$ instead. Any H -description can be given as just a set of inequalities.

1.1.3 Facets, vertices, extreme rays

The following substructures of polyhedra are crucial to understanding their structure and proving results about them.

Definition 1.1.10 Let $P \subseteq \mathbb{R}^d$ be a polyhedron. An equality $\mathbf{c}\mathbf{x} \leq c_0$ is valid for P if it is satisfied for all points $\mathbf{x} \in P$. A face of P is a set F of the form:

$$F = P \cap \{\mathbf{x} \in \mathbb{R}^d : \mathbf{c}\mathbf{x} = c_0\}$$

for a valid inequality $\mathbf{c}\mathbf{x} \leq c_0$ on P . We call $\mathbf{c}\mathbf{x} \leq c_0$ the defining inequality for F .

Definition 1.1.11 The affine hull of S , $\text{aff}(S)$, is the smallest affine space containing S . Equivalently,

$$\text{aff}(S) = \left\{ \sum_{i=1}^n \lambda_i x_i : x_i \in S, \sum_{i=1}^n \lambda_i = 1 \right\}.$$

Definition 1.1.12 • The dimension of a face is the dimension of its affine hull.

- Faces of P of dimension 0 and $\dim(P) - 1$ are called vertices and facets, respectively.
- An extreme ray of a polyhedral cone C is a face of dimension 1.

The polyhedron P itself is considered a face and so is the empty set which are of dimension of $\dim(P)$ and -1 respectively. Since extreme rays are subsets of a polyhedral cone, which are unbounded, in addition to being 1-dimensional faces, they can be viewed as linearly independent vectors, which is more formally described in Proposition 1.1.14. In addition, there is a containment relation among the faces, with every face being contained in some facet, except P itself. In a polytope, every face contains some vertex, except the empty face. More generally, in a polyhedron every face contains a face of minimum degree, except the empty face.

Example 1.1.13 Consider C_1 as defined in Example 1.1.6. This cone has 9 faces. They are:

- The empty set
- The point $(0, 0, 0)$
- The rays over $(1, 1, 0)$, $(1, 1, 1)$, and $(1, 0, 0)$
- The facets defined by the 2-dimensional planes $x_3 = 0$, $x_2 = x_3$, and $x_1 = x_2$
- C_1 itself.

The faces of a polyhedron are core to its geometric structure. In fact, studying the highest and lowest dimensional faces gives a unique minimal way to generate the polyhedron, as can be seen in the following proposition.

Proposition 1.1.14 1. *Every polytope is the convex hull of its vertices.*

2. *Every polyhedral cone is the conical hull of its extreme rays.*

3. *Every polyhedron is the intersection of the halfspaces determined by the defining inequalities of its facets.*

4. *The extreme rays of a polyhedral cone, C , are the $v \in C$ such that there does not exist $u, w \in C$ linearly independent and $\lambda, \gamma > 0$ with $v = \lambda u + \gamma w$.*

The biggest consequence to Proposition 1.1.14 is that understanding the facets and vertices /extreme rays lets you characterize the polytope/polyhedron as a whole. Thus as will be seen in Chapter 2, special attention is given to understanding and proving what are the facet defining inequalities and vertices/extreme rays of polyhedra.

1.2 Algebraic Geometry

1.2.1 Ideals and Varieties

In this section we give background on algebraic geometry. Algebraic geometry takes tools from algebra to solve problems in geometry. The classic object of interest is the zero set of a collection of multivariate polynomial equations. This has a corresponding algebraic structure, the ideal of all polynomials that evaluate to zero on that set. See [74] for more details on algebraic geometry. Algebraic geometry provides tools to find polynomials that vanish on statistical models allowing the discovery of underlying geometric structure. In addition, algebraic geometry gives characterizations of generic points of statistical models. These tools have been used on phylogenetic models to help characterize when data arises from a particular model. The definitions and results discussed here will primarily be used in Chapter 3.

Definition 1.2.1 *Let \mathbb{K} be a field. Let $\mathbb{K}[\mathbf{p}] := \mathbb{K}[p_1, \dots, p_n]$ be the ring of polynomials with coefficients in \mathbb{K} and indeterminates p_1, \dots, p_n . The variety defined by $I \subseteq \mathbb{K}[\mathbf{p}]$ is:*

$$V(I) = \{a \in \mathbb{K}^n : f(a) = 0 \text{ for all } f \in I\}.$$

Example 1.2.2 Let $I = \langle p_1 \rangle \subseteq \mathbb{R}[p_1, p_2]$. Then $V(I)$ is the p_2 -axis of the Cartesian plane. Here we can also note that $V(I) = V(\langle p_1 \rangle) = V(\{p_1\})$, the set of points that vanish on p_1 will also vanish on the ideal generated by p_1 , since every polynomial in $\langle p_1 \rangle$ will have p_1 as a factor.

Just as every set of polynomials defines a variety, every set of points in \mathbb{K}^n defines an ideal.

Definition 1.2.3 Let $S \subseteq \mathbb{K}^n$. The vanishing ideal of S is:

$$I(S) = \{f \in \mathbb{K}[\mathbf{p}] : f(a) = 0 \text{ for all } a \in S\}.$$

Example 1.2.4 Let S be a unit circle in \mathbb{R}^2 . Then

$$I(S) = \langle p_1^2 + p_2^2 - 1 \rangle.$$

Hilbert's Nullstellensatz, Theorem 1.2.6, explains how varieties and vanishing ideals interact.

Definition 1.2.5 The radical of an ideal I , denoted \sqrt{I} , is

$$\sqrt{I} = \{f \in \mathbb{K}[\mathbf{p}] : f^k \in I \text{ for some } k \in \mathbb{N}\}.$$

An ideal is called radical if $\sqrt{I} = I$

Note that if $J = I(S)$ for some $S \subseteq \mathbb{K}^n$ then J is radical since if $f^k(a) = 0$ for some $a \in S$, $k \in \mathbb{N}$ then so does $f(a) = 0$.

Theorem 1.2.6 (Nullstellensatz) Let \mathbb{K} be an algebraically closed field. Then the vanishing ideal of the variety of an ideal is the radical of the ideal, i.e: $I(V(I)) = \sqrt{I}$.

Hilbert's Nullstellensatz also naturally motivates the following definition.

Definition 1.2.7 The set $V(I(S))$ is the Zariski closure of S , written \overline{S} . It is the smallest algebraic variety that contains S .

Example 1.2.8 Let S be the upper half of a unit circle with the point $(0, 1)$ removed. Since polynomials are continuous, any polynomial that vanishes on S will also vanish on $(0, 1)$. Thus the Zariski closure of S must contain the entire upper half of a unit circle. In addition, any point on the upper half of a unit circle must satisfy $p_2 = \sqrt{1 - p_1^2}$. This condition is not polynomial, but we can rewrite it as $p_2^2 = 1 - p_1^2$. Any polynomial that vanishes on the upper half of a unit circle must then have $p_2^2 + p_1^2 - 1$ as a factor and so the Zariski closure of S is the entire unit circle.

1.2.2 Varieties from Maps

Using the Zariski closure, one can always create a variety from a set in \mathbb{K}^n . The sets we use are not arbitrary sets but often come from a parametrization by a polynomial map. Algebraic statistical models are parametrized semi-algebraic sets with polynomial or rational functions. Thus the variety defining a statistical model can be obtained from its polynomial map directly which is done in Chapter 3.

Definition 1.2.9 *Let ϕ be a map:*

$$\begin{aligned}\phi : \mathbb{K}^n &\rightarrow \mathbb{K}^m \\ \theta &\mapsto (\phi_1(\theta), \dots, \phi_m(\theta))\end{aligned}$$

where the functions ϕ_i are all polynomials in the indeterminates $\theta = (\theta_1, \dots, \theta_n)$. The variety parameterized by ϕ is the Zariski closure of the image of ϕ and written as $V = \overline{\text{im}(\phi)}$. In addition, ϕ has a pullback map ϕ^* which is the \mathbb{K} -algebra homomorphism:

$$\begin{aligned}\phi^* : \mathbb{K}[\mathbf{p}] &\rightarrow \mathbb{K}[\mathbf{t}] \\ \mathbf{p}_i &\mapsto \phi_i(\mathbf{t}).\end{aligned}$$

The corresponding vanishing ideal of the variety parameterized by ϕ is $\ker(\phi^*)$.

Example 1.2.10 *Let $\phi : \mathbb{K}^6 \rightarrow \mathbb{K}^8$ be a map such that*

$$\phi(\alpha_1, \alpha_2, \beta_1, \beta_2, \beta_3, \beta_4) = (\alpha_1\beta_1, \alpha_1\beta_2, \alpha_1\beta_3, \alpha_1\beta_4, \alpha_2\beta_1, \alpha_2\beta_2, \alpha_2\beta_3, \alpha_2\beta_4).$$

The map ϕ has pullback $\phi^* : \mathbb{C}[p_{11}, p_{12}, p_{13}, p_{14}, p_{21}, p_{22}, p_{23}, p_{24}] \rightarrow \mathbb{C}[s_1, s_2, t_1, t_2, t_3, t_4]$ such that

$$\phi_{ij}^*(\mathbf{p}) = s_i t_j$$

where $i \in \{1, 2\}$, $j \in \{1, 2, 3, 4\}$. Calculating the vanishing ideal for the variety parametrized by ϕ gives $\ker(\phi^*) = \langle p_{1i}p_{2j} - p_{2i}p_{1j} : 1 \leq i < j \leq 4 \rangle$, the ideal generated by the 2×2 minors of

$$P = \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \end{bmatrix}.$$

The vanishing ideal for the variety parametrized by ϕ is the set of 2×2 minors of a matrix, and

a well-known equivalent condition to all 2×2 minors being equal to zero is that the matrix has rank ≤ 1 . In other words, a matrix has rank one or less if and only all of its 2×2 minors vanish.

We can see this equivalent condition through the map ϕ^* . A matrix $A \in \mathbb{K}^{m \times n}$ has rank ≤ 1 if and only if it can be written as the outer product of two vectors:

$$A = \mathbf{s} \cdot \mathbf{t}^T$$

for some vectors

$$\mathbf{s} = \begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_m \end{bmatrix}, \quad \mathbf{t} = \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_n \end{bmatrix}.$$

Thus the condition that $\phi_{ij}^*(\mathbf{p}) = s_i t_j$ transforms the matrix P into

$$P = \begin{bmatrix} s_1 t_1 & s_1 t_2 & s_1 t_3 & s_1 t_4 \\ s_2 t_1 & s_2 t_2 & s_2 t_3 & s_2 t_4 \end{bmatrix} = \mathbf{s} \cdot \mathbf{t}^T$$

where

$$\mathbf{s} = \begin{bmatrix} s_1 \\ s_2 \end{bmatrix}, \quad \mathbf{t} = [t_1 \quad t_2 \quad t_3 \quad t_4].$$

Thus, the definition of ϕ^* was equivalent to a 2 by 4 matrix having rank ≤ 1 . Since the resulting vanishing ideal was the set of 2×2 minors of a matrix, one could use this map to show that a matrix is rank one or less if and only all of its 2×2 minors vanish. This is a well-known fact but other maps can show much less obvious facts by analyzing the resulting vanishing ideal.

This vanishing ideal also vanishes exactly on the statistical the model of independence, the model shown in Example 1.3.12 in more detail.

1.3 Statistics and statistical models

1.3.1 Statistical models

In this section we give necessary background on statistics. The general goal in statistics is to take a set of observations and describe the underlying distribution that created those observations or draw inferences about the relations between the quantities studied. Typically observations are

thought of as coming from some underlying unknown statistical model that can be theoretically described based on the mathematical relations the original observations adhere to. Through this lens, the goal is to ascertain the underlying model from the observed data and its specific model parameters that best fit the data. There are numerous inference procedures in order to find these model parameters, and which one fits best varies based on the modeling context and the circumstances of the observed data. For example, maximum likelihood estimation finds the parameter values that maximize the likelihood of observed data. See [94] for more background on Statistics. Since we work in the realm of theoretical statistics, we consider the statistical models as abstract mathematical objects. In order to define some statistical models, we first define random variables, probability distributions, and the probability simplex as our statistical models will be based on random variables and are collections of probability distributions.

There are two main types of random variables, discrete and continuous. An example of a continuous variable is a Gaussian random variable. The random variables that will be considered in this thesis are discrete random variables.

Definition 1.3.1 *A discrete random variable is a function*

$$X : \Omega \rightarrow \mathbb{R}$$

defined on a sample space Ω , such that the set of possible values taken by X , denoted $\{x_1, x_2, x_3, \dots\}$, is countable. Each value x_i corresponds to an elementary event in Ω .

Throughout the rest of this thesis, we will only be considering finite state space, and so when referring to a discrete random variable, we will actually mean a random variable with a finite state space.

Definition 1.3.2 *A probability distribution of a discrete random variable X is a function*

$$p_X(x) = \mathbb{P}(X = x)$$

that assigns to each possible value x the probability that X takes that value. This function, called the probability mass function, satisfies:

- $p_X(x) \geq 0$ for all x ,
- $\sum_x p_X(x) = 1$, where the sum is over all values in the support of X .

We call the support of a random variable its state space, S . The notation $\mathbb{P}(X = i)$ for $i \in S$ means the probability that X equals state i . We now define the probability simplex, the space in which probability distributions reside.

Definition 1.3.3 *The $d - 1$ -dimensional simplex, $\Delta_{d-1} \subset \mathbb{R}^d$, is*

$$\Delta_{d-1} = \{(p_1, \dots, p_d) : p_1 + \dots + p_d = 1, p_i \geq 0\}.$$

A probability distribution with a state space S of size d is a point in Δ_{d-1} . Thus a discrete random variable is a map from a state space S of size d to Δ_{d-1} .

Example 1.3.4 *Consider a probability distribution for a fair 6-sided die. Then, the state space $S = \{1, 2, 3, 4, 5, 6\}$ is the 6 sides of the die and*

$$\mathbb{P}(X = 1) = \mathbb{P}(X = 2) = \mathbb{P}(X = 3) = \mathbb{P}(X = 4) = \mathbb{P}(X = 5) = \mathbb{P}(X = 6) = \frac{1}{6}.$$

Definition 1.3.5 *Let X, Y be discrete random variables with state spaces S, T , respectively. The notation $\mathbb{P}(X = i, Y = j)$ for $i \in S, j \in T$ signifies the probability that X is state i and Y is state j . We often abbreviate this as $\mathbb{P}(X = i, Y = j) = p_{ij}$. This is a joint probability distribution, a probability distribution given on multiple sets of state spaces. $\mathbb{P}(X = i, Y = j)$ is any probability distribution on the function $X \times Y \rightarrow \mathbb{R} \times \mathbb{R}$. The conditional probability of X being state i given that Y is state j is defined as:*

$$\mathbb{P}(X = i | Y = j) := \frac{\mathbb{P}(X = i, Y = j)}{\mathbb{P}(Y = j)}.$$

Definition 1.3.6 *We call two random variables X, Y independent if $\mathbb{P}(X = i | Y = j) = \mathbb{P}(X = i)$ for all $i \in S, j \in T$. Equivalently, $\mathbb{P}(X = i, Y = j) = \mathbb{P}(X = i)\mathbb{P}(Y = j)$ for all $i \in S, j \in T$. If X and Y are independent, we denote this $X \perp\!\!\!\perp Y$. We can think of independence of two random variables as those variables not affecting each other – knowing one variable's value does not affect our estimation of the value of the other.*

Before moving on to defining statistical models, we give some examples of random variables and probability distributions.

Example 1.3.7 *Let X_1, X_2 be discrete random variables with state space $S = [6] = \{1, 2, 3, 4, 5, 6\}$ as the sides of two different fair 6-sided dices. Then*

$$\mathbb{P}(X_1 = i_1) = \mathbb{P}(X_2 = i_2) = \frac{1}{6}$$

for all $i_1, i_2 \in S$. In addition since the two dice do not affect each other, their random variables are independent and thus:

$$\mathbb{P}(X_1 = i_1, X_2 = i_2) = \mathbb{P}(X_1 = i_1)\mathbb{P}(X_2 = i_2) = \frac{1}{36}.$$

Example 1.3.8 A Bernoulli random variables is one of the simplest random variables, taking a value of 1, a success, with probability θ and 0, a failure, with probability $1 - \theta$. That is:

$$\mathbb{P}_\theta(X = 1) = \theta, \mathbb{P}_\theta(X = 0) = 1 - \theta.$$

Example 1.3.9 Let X be a discrete random variable with $r+1$ states, $\{0, 1, \dots, r\}$ and let $\Theta = [0, 1]$. For all $\theta \in \Theta$ let

$$\mathbb{P}_\theta(X = i) = \binom{r}{i} \theta^i (1 - \theta)^{r-i}.$$

Then X is a binomial random variable with r samples and success probability parameter θ . This can be thought of as taking r samples of a Bernoulli variable with probability θ and looking at the probability that there are i successes.

Let us consider the specific case of a binomial random variable with $r = 3$. Then we have a Bernoulli variable with a success chance of θ and we are doing 3 trials. Thus we can have 0, 1, 2 or 3 successes. Then

$$\mathbb{P}_\theta(X = 0) = (1 - \theta)^3, \mathbb{P}_\theta(X = 1) = 3\theta(1 - \theta)^2, \mathbb{P}_\theta(X = 2) = 3\theta^2(1 - \theta), \mathbb{P}_\theta(X = 3) = \theta^3.$$

Definition 1.3.10 A statistical model, \mathcal{M} , is a collection of probability distributions. We call a model a parametric statistical model, \mathcal{M}_Θ if it is a mapping from a finite dimensional parameter space to a space of probability distributions. That is

$$P_\bullet : \Theta \subseteq \mathbb{R}^d \rightarrow \mathcal{M}_\Theta, \theta \mapsto p_\theta$$

where the parametric model is the image of this map, $\mathcal{M}_\Theta = \{p_\theta : \theta \in \Theta\}$.

Example 1.3.11 Let X be a binomial random variable with r samples and success probability parameter θ as described in Example 1.3.9. The binomial model, \mathcal{M}_Θ is all probability distributions arising from binomial random variables so

$$\mathcal{M}_\Theta = \left\{ \left((1 - \theta)^r, \binom{r}{1} \theta (1 - \theta)^{r-1}, \dots, \binom{r}{r-1} \theta^{r-1} (1 - \theta), \theta^r \right) \right\}$$

and $\mathcal{M}_\Theta \subseteq \Delta_r$.

Example 1.3.12 Let X_1, X_2 be discrete random variables with state spaces $[r_1]$ and $[r_2]$. The model of independence for X_1, X_2 , $\mathcal{M}_{X_1 \perp X_2}$ consists of all distributions $p \in \Delta_{r_1 r_2}$ such that

$$\mathbb{P}(X_1 = i_1, X_2 = i_2) = \mathbb{P}(X_1 = i_1)\mathbb{P}(X_2 = i_2)$$

for all $i_1 \in [r_1]$ and $i_2 \in [r_2]$. An example of two such random variables can be seen in Example 1.3.7. Here $\Theta = \Delta_{r_1-1} \times \Delta_{r_2-1}$ and we can parameterize this model using $\theta = (\alpha, \beta) \in \Theta$ in the following way:

$$\mathbb{P}_\theta(X_1 = i_1, X_2 = i_2) = \alpha_{i_1} \beta_{i_2}.$$

This is exactly the same parameterization from the map used in Example 1.2.10, with $r_1 = 2$ and $r_2 = 4$. Letting $\mathbb{P}_\theta(X_1 = i, X_2 = j) = p_{ij}$, we can see from the vanishing ideal that, for example, the following holds: $p_{ij}p_{kl} - p_{il}p_{kj} \in V(I)$ for $i < k \in [r_1]$ and $j < l \in [r_2]$ and so $p_{ij}p_{kl} - p_{il}p_{kj} = 0$. In other words using the fact that $\mathbb{P}(X_1 = i, X_2 = j) = \mathbb{P}(X_1 = i)\mathbb{P}(X_2 = j)$:

$$\begin{aligned} p_{ij}p_{kl} &= \mathbb{P}_\theta(X_1 = i, X_2 = j) \cdot \mathbb{P}_\theta(X_1 = k, X_2 = l) \\ &= \mathbb{P}_\theta(X_1 = i) \cdot \mathbb{P}_\theta(X_2 = j) \cdot \mathbb{P}_\theta(X_1 = k) \cdot \mathbb{P}_\theta(X_2 = l) \\ &= \mathbb{P}_\theta(X_1 = i) \cdot \mathbb{P}_\theta(X_2 = l) \cdot \mathbb{P}_\theta(X_1 = k) \cdot \mathbb{P}_\theta(X_2 = j) \\ &= \mathbb{P}_\theta(X_1 = i, X_2 = l) \cdot \mathbb{P}_\theta(X_1 = k, X_2 = j) \\ &= p_{il}p_{kj}. \end{aligned}$$

Thus $p_{ij}p_{kl} = p_{il}p_{kj}$.

1.3.2 Identifiability

A statistical model is identifiable if the parameters of the model can be recovered from observations. Identifiability is of interest theoretically and practically. Theoretically, a statistical model is structurally identifiable if its corresponding map from its parameter space to the probability simplex is one-to-one. Practical identifiability concerns whether or not one can recover specific information about parameters from observed noisy data. Structural identifiability is necessary for practical identifiability and is necessary for the consistency of statistical procedures. This thesis will only discuss structural identifiability and from now on, for brevity, will simply refer to it as identifiability. Our interest in identifiability comes from the identifiability of phylogenetic models, which are explained in more detail in Section 1.4. In mathematical phylogenetics,

theoretical identifiability means recovering the underlying evolutionary tree structure as well as the mutation rates from generic data about the living observable species. In Chapter 3 we prove the identifiability of types of phylogenetic mixture models, under certain conditions.

There are many kinds of identifiability, which we go over here now.

Definition 1.3.13 *Let $\phi : \Theta \rightarrow N$ be a rational map from the parameter space Θ to some set N , with statistical model $\mathcal{M} = \text{im } \phi$. Then the parameter θ is:*

- globally identifiable if ϕ is a one-to-one map;
- generically identifiable if $\phi^{-1}(\phi(\theta)) = \{\theta\}$ for almost all $\theta \in \Theta$.
- locally identifiable if $|\phi^{-1}(\phi(\theta))| < \infty$ for almost all $\theta \in \Theta$.

In some circumstances if the finite-to-one behavior of the locally identifiability is due to an intrinsic symmetry of the model, we may also call the model generically identifiable.

Example 1.3.14 *Let X_1, X_2 be discrete random variables with state spaces $[r_1]$ and $[r_2]$. Consider the model of independence for X_1, X_2 , $\mathcal{M}_{X_1 \perp X_2}$, defined in Example 1.3.12. This model is defined parametrically via the map*

$$\mathbb{P}_\theta(X_1 = i_1, X_2 = i_2) = \alpha_{i_1} \beta_{i_2}.$$

The model of independence is globally identifiable because we can recover α and β from the probability distribution $p \in \mathcal{M}_{X_1 \perp X_2}$. In fact letting $\mathbb{P}(X_1 = i_1, X_2 = i_2) = p_{i_1 i_2}$, we have that for $p \in \mathcal{M}_{X_1 \perp X_2}$:

$$\begin{aligned} \sum_j p_{ij} &= \sum_j \alpha_i \beta_j = \alpha_i \sum_j \beta_j = \alpha_i \\ \text{and} \quad \sum_i p_{ij} &= \sum_i \alpha_i \beta_j = \beta_j \sum_i \alpha_i = \beta_j \end{aligned}$$

since $\sum_j \beta_j = 1$ and $\sum_i \alpha_i = 1$. This model is also what is called rationally identifiable since all of the equations that recover the parameters are rational functions.

Example 1.3.15 *Suppose we have a model with two dice: A fair die where each face has probability $\frac{1}{6}$ and a biased die, which rolls a 6 with probability p , and each of $\{1, 2, 3, 4, 5\}$ with probability $\frac{1-p}{5}$. One of the two dice is rolled and $\pi \in [0, 1]$ is the probability of choosing the biased die. The die used for each roll is unobserved, but we observe the outcome of the roll. Let $X \in \{1, 2, 3, 4, 5, 6\}$ be the outcome of a single roll. Then:*

$$P(X = x) = \begin{cases} \pi \cdot \frac{1-p}{5} + (1-\pi) \cdot \frac{1}{6}, & \text{if } x \in \{1, 2, 3, 4, 5\} \\ \pi \cdot p + (1-\pi) \cdot \frac{1}{6}, & \text{if } x = 6. \end{cases}$$

If we observe the empirical probabilities:

$$f_6 = P(X = 6), \quad f_{-6} = P(X \in \{1, 2, 3, 4, 5\}) = 1 - f_6,$$

then from the model we know that:

$$\begin{aligned} f_6 &= \pi p + (1-\pi) \cdot \frac{1}{6} \\ f_{-6} &= \pi(1-p) + (1-\pi) \cdot \frac{5}{6}. \end{aligned}$$

Let:

$$(\pi_1, p_1) = \left(\frac{1}{3}, \frac{1}{2}\right), \quad (\pi_2, p_2) = \left(\frac{1}{6}, \frac{5}{6}\right).$$

Then $f_6 = \frac{5}{18}$ for both sets of parameters. Thus, since there are multiple sets of parameters that give the same probability distributions, the model is not identifiable. In fact, we can also solve for p in terms of π , to get:

$$p = \frac{f_6 - \frac{1}{6}}{\pi} + \frac{1}{6}$$

This defines a curve in parameter space (π, p) for any fixed $f_6 \in (\frac{1}{6}, 1)$, as long as $\pi \neq 0 \in (0, 1)$.

Hence, for almost all f_6 , there is a one dimensional subset of parameter pairs (π, p) yielding the same observable distribution, meaning that the model is not generically identifiable.

In general there can be both discrete structures and continuous model parameters that we are interested in understanding the identifiability of. By the parameter itself being identifiable, we mean that specific value of that parameter can be determined by the observable distribution. A continuous model parameter would be a parameter that can take real values, and a discrete model parameter would be some combinatorial object like a partition or graph structure. For example, in Chapter 3 we will be determining the identifiability of a tree structure (a discrete parameter) as well as continuous parameters from transition matrices on edges of the tree.

1.3.3 Hypothesis testing

In statistics, questions about collected data are often formed into statistical hypotheses. Typically this is phrased in the following way: given a set of independent, identically distributed samples, were these samples generated from a distribution in a particular model or not? In Chapter 4 we explore the intricacies of a new hypothesis testing method by applying it to a common phylogenetic model. We start by explaining a set of fundamental definitions about statistical hypotheses before then exploring with multiple examples.

Definition 1.3.16 • *The null hypothesis, H_0 , is the baseline assumption.*

- *The alternative Hypothesis (H_1 or H_a) is the claim being testing for.*
- *A function T is a test statistic if when evaluated on data that exactly fits the null hypothesis, T evaluates to 0.*
- *The critical value of a test statistic is the value of the test statistic that we are comparing samples against and it is used to determine whether to reject or not reject the null hypothesis.*
- *A p -value is the probability under the null hypothesis that the test statistic has a larger value than the given critical value.*

For example, given observed data x generated from an unknown distribution p , and model \mathcal{M} , a null hypothesis could be that $H_0 : p \in \mathcal{M}$. Then our alternative hypothesis would be that $H_1 : p \notin \mathcal{M}$. A small p -value gives evidence against the null hypothesis, while with a large p -value we cannot reject the null hypothesis. Typically we choose some value α that if the p -value calculated is less than α , we reject the null hypothesis. A traditional choice for α is 0.05. Also, typically we do not calculate a p -value exactly, but take random samples to approximate a p -value or use an asymptotic approximation. Below are two examples of p -value calculations using an asymptotic approximation, in Chapter 4 we will approximate p -values with random samples.

Example 1.3.17 *Suppose that a company claims the defect rate for their product is 3%. A sample of $n = 100$ items reveals $x = 6$ defects. This is a binomial model as described in Example 1.3.11. Here we have null and alternative hypothesizes as:*

$$H_0 : p = 0.03 \quad \text{vs.} \quad H_1 : p \neq 0.03.$$

We use the test statistic called a z -score:

$$z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}$$

Here \hat{p} is the observed failure rate for the binomial variable, p_0 is the hypothesized failure rate under the null hypothesis and n is the number of trials. In our example:

$$\hat{p} = \frac{6}{100} = 0.06, \quad p_0 = 0.03, \quad n = 100.$$

Also note that z evaluates to 0 if $\hat{p} = p_0$. Now evaluating our observation on the test statistic gives:

$$z = \frac{0.06 - 0.03}{\sqrt{\frac{0.03 \cdot 0.97}{100}}} \approx 1.76.$$

We can approximate the binomial distribution by a normal distribution under the Central Limit Theorem when the sample size n is large enough. Thus we can approximate the p -value as the following:

$$p\text{-value} = 2 * \mathbb{P}(Z > 1.76) \approx 2 * 0.0392 = 0.0784$$

This is the probability that a standard normal random variable exceeds 1.76, multiplied by 2 because we are performing a 2-sided test. If we choose $\alpha = 0.05$, we do not reject H_0 since $0.0784 > 0.05$. Thus there is insufficient evidence to conclude the defect rate is not 3%.

Example 1.3.18 Consider an experiment to determine if cat coat color and cat treat preferences are independent. That is, the goal is to determine whether the color of the cat's coat affects their preference for a specific type of treat. This is an example of the independence model shown in Example 1.3.14. The following table has our observations of 100 cats:

	Likes Treat	Dislikes Treat	Total
Black Coat	20	10	30
Gray Tabby Coat	30	40	70
Total	50	50	100

Here we have null and alternative hypotheses as:

H_0 : Cat coat color and Preferences are independent

H_1 : Cat coat color and Preferences are not independent.

Let O_{ij} be the observed counts for row i and column j of our table and E_{ij} the expected counts under the independence model. Then:

$$E_{11} = \frac{30 \cdot 50}{100} = 15, \quad E_{12} = 15, \quad E_{21} = 35, \quad E_{22} = 35.$$

Here we use a χ^2 -test statistic. Evaluating the test statistic on the data gives:

$$X^2 = \sum_{i,j} \frac{(O_{ij} - E_{ij})^2}{E_{ij}} = \frac{(20-15)^2}{15} + \frac{(10-15)^2}{15} + \frac{(30-35)^2}{35} + \frac{(40-35)^2}{35} \approx 4.76$$

For large enough n we can approximate our p value using a chi-squared distribution with 1 degree of freedom as described in Chapter 10 of [30]. Let $Z \sim \text{chisquare}(1)$:

$$p\text{-value} = \mathbb{P}(Z > 4.76) \approx 0.029.$$

If we choose $\alpha = 0.05$, $p\text{-value} < 0.05$, and so we can reject H_0 . Cat coat color and treat preference are not independent.

1.4 Phylogenetic trees and networks

1.4.1 Trees and Networks

In this section we give background on mathematical phylogenetics. This will be used through the thesis, in each of the main chapters. The goal motivating much of the work in mathematical phylogenetics is to infer an evolutionary history between species, from information from those species. Inferring an evolutionary history takes many forms from distance methods explored in Section 1.4.2 to using phylogenetic models explored in Section 1.4.3. We use trees and other graphs to represent the evolutionary process of a group of species, so we will start with a review of graph theory.

Definition 1.4.1 A graph G is a pair of sets $G = (V, E)$ where:

1. V is the set of vertices and
2. E is the set of edges where every edge $e \in E$ is an unordered set of two vertices, $e = \{v_1, v_2\}$ where $v_1, v_2 \in V$. Typically we do not allow $v_1 = v_2$ unless specified.

Two vertices $v_1, v_2 \in V$ are adjacent if $e = \{v_1, v_2\} \in E$.




Figure 1.2: The graph described in Example 1.4.3 with the path (1, 5, 6, 7) highlighted in red.

- Definition 1.4.2**
1. A walk on $G = (V, E)$ is an ordered list of n vertices $v_1, v_2 \dots v_{n-1}, v_n \in V$ such that for every $i \in \{1, \dots, n-1\}$, v_i is adjacent to v_{i+1} .
 2. A path on $G = (V, E)$ is a walk with n distinct vertices.
 3. A cycle is a walk where $v_1 = v_n$ and all other vertices are distinct and $n \geq 4$.
 4. A graph is connected if there exists a path between every pair of vertices in V .
 5. A graph is called acyclic if it has no cycles.
 6. A connected acyclic graph is called a tree.

Example 1.4.3 Let G be a graph, $G = (V, E)$ with vertex set $V = [10] = \{1, 2, \dots, 10\}$ and edge set:

$$E = \{(1, 2), (1, 3), (1, 4), (1, 5), (5, 6), (6, 7), (6, 9), (7, 8), (9, 10)\}.$$

The graph G is connected and contains no cycles and therefore is a tree. An example of a path in G is: $P = (1, 5, 6, 7, 8)$. The graph G is visualized in Figure 1.2 where each vertex is a circle with its numbered label inside, and there are lines between the vertices for the edges. The path P is highlighted in red.

In the context of phylogenetics, we use the vertices of graphs to represent species and have an edge between the species if there is a direct evolutionary relationship between those species, one evolved from the other. This naturally motivates defining a sense of direction on the tree, in order to show which species is the ancestor of which.

Definition 1.4.4 A directed graph G is a pair of sets $D = (V, E)$ where:

1. V is the set of vertices and

2. E is the set of edges where every edge $e \in E$ is an ordered set of two vertices, $e = (v_1, v_2)$ where $v_1, v_2 \in V$. Note that (v_1, v_2) is a distinct edge from (v_2, v_1) .

When using a directed graph for phylogenetics, the edge (v_1, v_2) indicates that v_1 is an ancestor of v_2 . There are many ways to “direct” an undirected graph, that is to define an order on each of its edges, turning a graph into a directed graph. One way common in phylogenetics, and many other contexts, is to specify a vertex as the “root”. Then, there is a natural way to direct the graph, where all edges point away from the root. Typically we will work with undirected graphs but when we want to utilize the directed nature of evolution, we specify a root, which defines our directing of the graph. A root is also relevant biologically as that vertex will be considered the evolutionary ancestor of all other species in the tree.

Definition 1.4.5 A rooted tree is a tree with a distinguished vertex called the root.

In a rooted phylogenetic tree, the vertices that have no edges originating from them are of particular interest as they represent species that are observable and alive today. A common way for new species to form is for one species to have a mutation and split into two different species. This naturally motivates a kind of “binary” rooted phylogenetic where if the root is placed at top, every vertex except the bottom ones will have exactly two vertices below it. We formally define both of these concepts below.

Definition 1.4.6 1. The degree of a vertex is the number of vertices adjacent to the vertex.

2. A leaf is a vertex of degree 1.

3. A binary rooted tree is a tree where every vertex is degree 3 except the root which is degree 2 and the leaves which are of degree 1.

Example 1.4.7 Let T be a directed graph $T = (V, E)$, where: $V = \{1, 2, 3, 4, 5, 6, 7\}$ and $E = \{(1, 2), (1, 3), (2, 4), (2, 5), (3, 6), (3, 7)\}$. T is visualized in Figure 1.3. We can see that T is a rooted binary tree with root 1 and leaves 4, 5, 6, 7.

Since only the leaves of a phylogenetic tree can be observed, the labels of those vertices have special importance. Thus we define a set of vertex labels that we will only assign to the leaves and the root. We do this by creating a set of leaf labels and then defining a map from these labels to the phylogenetic tree. Since these leaf labels are most important biologically, we will typically suppress the original labels and only show this new set of labels, as can be seen in Figure 1.4. This is formally defined below.




Figure 1.3: The rooted binary tree described in Example 1.4.7 with root 1 and leaves 4, 5, 6, 7.

Definition 1.4.8 Let X be a set of labels, $T = (V, E)$ a tree, and $\phi : X \rightarrow V$ an injective map where the image is exactly the set of leaves of T . Then we call (T, ϕ) a phylogenetic X -tree.

If one removes any edge from a phylogenetic X -tree, it separates the leaf labels into two disjoint sets. This naturally motivates the following definition.

Definition 1.4.9 Let $X = \{1, \dots, n\}$ be labels of an X -tree. A split $A|B$ is a partition of X into two nonempty sets. A set of splits is called a split system.

We consider the split obtained from an edge of an phylogenetic X -tree to be the two sets of labels from the two connected parts left behind by the edge removal. In Example 1.4.10 and all future examples, we will label the root vertex 0 and treat it the same as the other leaf labels, despite it not corresponding to a living species. The reason for this is that we want to be able to analyze the distance from the root to each of the leaves, which requires the root to have a label. This is discussed further in Section 1.4.2.

Example 1.4.10 Let T be a rooted tree as depicted in Figure 1.4 with the root labeled by 0. Then we have vertex labels $X = \{0, 1, 2, 3, 4, 5, 6\}$ with leaf labels [6] and root label 0. All other vertices are unlabeled. By removing each edge of this graph in turn, we can see that it has the following splits:

$$S = \{1|023456, 2|013456, 3|012456, 4|012356, 5|012346, 6|012345, \\ 12|03456, 34|01256, 1234|056, 01234|56\}.$$

The splits

$$1|023456, 2|013456, 3|012456, 4|012356, 5|012346, 6|012345$$




Figure 1.4: The X-tree with splits described in Example 1.4.10 with root 0 and leaves 1, 2, 3, 4, 5, 6.

are called trivial splits as they only separate one labeled vertex and are created by the edges above the leaves.

Since a split system can be obtained from an X-tree, a natural question is: given an arbitrarily split system how does one know if it originally came from an X-tree? The following definition fully characterizes when a split system is obtained from an X-tree.

Definition 1.4.11 A split system S is pairwise compatible if for every pair of splits $A_1|B_1, A_2|B_2 \in S$ at least one of the following is empty:

$$A_1 \cap A_2, A_1 \cap B_2, A_2 \cap B_1, B_1 \cap B_2.$$

Example 1.4.12 Consider the split system S described in Example 1.4.10. The splits $12|03456$, $34|01256$ are both in S . In addition we can see that these two splits are compatible with each other since:

$$A_1 \cap A_2 = \emptyset, A_1 \cap B_2 = \{12\}, A_2 \cap B_1 = \{34\}, B_1 \cap B_2 = \{056\}.$$

In fact, a set of splits is pairwise compatible if and only if it came from an phylogenetic X-tree. One can freely go between a pairwise compatible split system and a phylogenetic X-tree, as described in the following theorem:

Theorem 1.4.13 (Splits Equivalence Theorem [20]) Let N be a split system on X with trivial splits for every $i \in X$. Then there exists a phylogenetic X-tree T with N as the set of splits obtained from removing edges of T if and only if the splits in N are pairwise compatible. Furthermore, the tree T is uniquely determined.

Example 1.4.14 Consider the split system S described in Example 1.4.10. The X-tree that realizes this split system, shown in Figure 1.4 can be recreated by the following process. Start with a single




Figure 1.5: The first step of creating an X -tree that realizes all of the splits in S as described in Example 1.4.10.

vertex with all of the labels $0, 1, 2, 3, 4, 5, 6$. Then for each split add an edge and separate the two sets of leaf labels on either side of that edge. Adding an edge for every trivial split results in a star tree. An example of one of the non-trivial splits $12|03456$ being added is shown in Figure 1.5.

The focus in Chapter 2 will be on split systems that are not necessarily pairwise compatible, sometimes also referred to as *split networks*. There are many other conditions besides pairwise compatibility that one could consider for a split system. One characteristic of interest in this thesis is the circular condition because it has a nice description based on distances as will be discussed in Section 1.4.2.

Definition 1.4.15 A split system S is circular with respect to some cyclic ordering (x_1, \dots, x_n) of X if every split of S is of the form $x_{i+1}, x_{i+2}, \dots, x_{j-1}, x_j | x_{j+1}, x_{j+2}, \dots, x_{i-1}, x_i$, where $x_{n+1} = x_1$.

Example 1.4.16 The split system S described in Example 1.4.10 is circular with respect to the circular order $0, 1, 2, 3, 4, 5, 6$. An example of an split system that is not circular is $\{12|34, 13|24, 14|23\}$ as there is no circular order one can write all of those splits in.

1.4.2 Distance methods for Trees and Networks

Another common approach in mathematical phylogenetics is to start with a notion of distance between species and construct a phylogenetic tree or network that realizes that set of pairwise distances. This is a very natural approach from a biological prospective as there are numerous ways to generate a set of pairwise distances between a set of species and one would want to know what phylogenetic structure realizes those distances. Before continuing we introduce a new type of graph structure that will be crucial for distance methods.

Definition 1.4.17 An edge weighted graph is a pair (G, W) where $G = (V, E)$ is a graph and W is a set of edge weights where for each edge $\{x, y\} \in E$, there is a non-negative weight $w_{\{x,y\}} \geq 0 \in W$.

While it is possible to also consider negative edge weights, it is common to restrict to nonnegative weights while defining distances on graphs, as we have done in this thesis.

We now define two notions of distance commonly used in phylogenetics, dissimilarity maps and tree metrics.

Definition 1.4.18 *Let X be a set of labels. A function $\delta : X \times X \rightarrow \mathbb{R}_{\geq 0}$ is called a dissimilarity map if $\delta(x, y) = \delta(y, x) \geq 0$ for all $x, y \in X$. In addition, a dissimilarity map is a metric if it satisfies the triangle inequality on X , meaning for all $x, y, z \in X$,*

$$\delta(x, z) \leq \delta(x, y) + \delta(y, z).$$

A metric δ is a tree metric if there is some edge weighted tree (T, W) such that $\delta(x, y)$ is the sum of the edge weights on the path between x and y in T , for all $x, y \in X$.

Example 1.4.19 *Let δ be the following map where $\delta(i, j)$ is the (i, j) -th entry:*

$$\delta = \begin{bmatrix} 0 & 5 & 11 & 12 & 26 & 20 \\ & 0 & 12 & 13 & 27 & 21 \\ & & 0 & 3 & 27 & 21 \\ & & & 0 & 28 & 22 \\ & & & & 0 & 8 \\ & & & & & 0 \end{bmatrix}.$$

We will define this map so it is implicitly symmetric, and thus we do not write the lower triangular portion of the matrix. One can see that this map is a metric. For instance $\delta(3, 5) = 27 \leq 3 + 28 = \delta(3, 4) + \delta(4, 5)$. Now consider the phylogenetic X -tree from Example 1.4.10 but with edge weights, as shown in Example 1.6. We can see that δ is also a tree metric, since the distance between leave i and j is exactly $\delta(i, j)$. For example $\delta(1, 5) = 2 + 3 + 8 + 6 + 7 = 26$.

One can easily generate a tree metric by starting with an edge weighted tree, but a natural question is the opposite: when is a dissimilarity map also a tree metric? The following seminal theorem answers exactly that question.

Theorem 1.4.20 (Four-point condition, [21]) *A metric $\delta : X \times X \rightarrow \mathbb{R}_{\geq 0}$ is a tree metric if and only if for all $i, j, k, l \in X$*

$$\delta(i, j) + \delta(k, l) \leq \max(\delta(i, k) + \delta(j, l), \delta(i, l) + \delta(j, k))$$




Figure 1.6: A phylogenetic X -tree with edge weights. The distance between leaves is calculated by summing over the weights on the path between leaves and it is described by δ as defined in Example 1.4.19.

Example 1.4.21 Consider the tree metric defined in Example 1.4.19. We can verify that δ satisfies the four point condition. For instance, $\delta(1,2) + \delta(3,4) = 5 + 3 \leq \max(\delta(1,3) + \delta(2,4), \delta(1,4) + \delta(2,3)) = \max(11 + 13, 12 + 12) = 24$.

Theorem 1.4.20 gives a very nice if and only if statement characterizing when a dissimilarity map is a tree metric, allowing us to freely switch between a tree metric and a tree realizing that metric.

There are many conditions on a tree metric one might care about, for biological or mathematical reasons. Of great importance in the results in Chapter 2 are equidistant tree metrics which we formally define now.

Definition 1.4.22 Tree metrics are equidistant if the distance from the root to every leaf is equal. In other words: for all $i, j \in X$, $\delta(r, i) = \delta(r, j)$. In addition, we call a tree equidistant if it realizes an equidistant tree metric.

Biologically, equidistant tree metrics are of note because distance from the root can be thought of as a measure of total generations for a species to evolve. Thus equidistant tree metrics closely model a group of closely related species that have a similar generation size and started evolving from a common ancestor at the same time.

Given Theorem 1.4.20 and the definition of equidistant tree metrics, a natural question is the following: is there a similar characterization of exactly when a tree metric is a metric for an equidistant tree? The following definition and theorem give this desired characterization.

Definition 1.4.23 A metric $\delta : X \times X \rightarrow \mathbb{R}_{\geq 0}$ is an ultrametric if for all $x, y, z \in X$:

$$\delta(x, y) \leq \max(\delta(x, z), \delta(y, z)).$$




Figure 1.7: An equidistant phylogenetic X -tree whose corresponding ultrametric tree metric is described in Example 1.4.25.

Theorem 1.4.24 ([59]) *A metric is an ultrametric if and only if it can be represented by an equidistant tree.*

Example 1.4.25 *Consider the edge-weighted phylogenetic X -tree in Figure 1.7. This is the same tree as was discussed in Example 1.4.19 but with different weights so that the distance from the root to each leaf is 8. Let δ be the following map where $\delta(i, j)$ is the (i, j) -th entry:*

$$\delta = \begin{bmatrix} 0 & 6 & 10 & 10 & 16 & 16 \\ & 0 & 10 & 10 & 16 & 16 \\ & & 0 & 2 & 16 & 16 \\ & & & 0 & 16 & 16 \\ & & & & 0 & 4 \\ & & & & & 0 \end{bmatrix}.$$

We can see that δ is an ultrametric tree metric and the tree representing it is the tree in Figure 1.7. For instance, $\delta(1, 4) = 10 = 3 + 2 + 4 + 1 \leq \max(\delta(1, 6), \delta(4, 6)) = \max(16, 16) = 16$.

The main focus of Chapter 2 is split networks characterizing distance functions as was done for trees in Theorem 1.4.20 and equidistant trees in Theorem 1.4.24. We will first describe the state of what is known about distance functions for split systems. To start we need to define how distances between leaves on a split network are defined. Since a split system is just a set of splits that are potentially not compatible, it only makes sense to define “edge” weights on the splits themselves. There is in fact a way to turn these split systems into a non-tree graph where each of these splits corresponds to a set of edges, that we will describe in detail in Chapter 2.

Definition 1.4.26 *Let N be a split system, and $i, j \in [n]$. A split $A|B \in N$ separates i and j in N if $i \in A$ and $j \notin A$ or $i \in B$ and $j \notin B$.*

Definition 1.4.27 Let N be a split network on $[n]$. Let N have non-negative weights W such that for each split $A|B \in N$, there is $a_{A|B} \in W$ such that $a_{A|B} \in \mathbb{R}_{\geq 0}$ for $A|B \in N$. Then, define:

$$\mathbb{1}_{A|B}(i, j) = \begin{cases} 1, & \text{if } A|B \text{ separates } i \text{ and } j \\ 0, & \text{otherwise.} \end{cases}$$

Let $\delta : [n] \times [n] \rightarrow \mathbb{R}_{\geq 0}$ be a metric on N . We can calculate the distance between leaves $i, j \in [n]$ of N using the weights in W in the following way:

$$\delta(i, j) = \sum_{A|B \in N} a_{A|B} \mathbb{1}_{A|B}(i, j).$$

The distance between two leaves in a split network is the sum over all the weights of the splits that separate those two leaves. It is, equivalently, the sum over all of the edge weights on one of the shortest paths between the two leaves.

Example 1.4.28 Let N be the split system with root 0, leaves 1, 2, 3, 4, 5, and splits:

$$\{1|02345, 2|01345, 3|01245, 4|01235, 5|01234, \\ 12|0345, 01|2345, 23|0145, 45|0123\}.$$

Let the set of weights for the splits be:

$$a_{1|02345} = 3, a_{2|01345} = 2, a_{3|01245} = 3, a_{4|01235} = 1, a_{5|01234} = 2, \\ a_{12|0345} = 2, a_{01|2345} = 1, a_{23|0145} = 4, a_{45|0123} = 5.$$

The split system N with the above weights is depicted in Figure 1.8. We calculate distances between leaves, as described perviously. For example $\delta(1, 4) = 3 + 2 + 1 + 5 + 1 = 12$. Let δ be the following map where $\delta(i, j)$ is the (i, j) -th entry and distance between leave i and j :

$$\delta = \begin{bmatrix} 0 & 10 & 13 & 12 & 13 \\ & 0 & 7 & 14 & 15 \\ & & 0 & 13 & 14 \\ & & & 0 & 3 \\ & & & & 0 \end{bmatrix}.$$

Now, just as before, a natural question is: when is a metric also the metric for a split system?




Figure 1.8: A split system with split weights as described in Example 1.4.28. Each of the different colors corresponds to a different split and the method to turn the split system into this graph is explained in Chapter 2.

This is answered in the following theorem, for circular split systems.

Theorem 1.4.29 ([11]) *Let δ be a metric on X . Then δ is the distance function for a circular split system if and only if for every set of leaves $i < j < k < l \in X$ under some circular ordering π , both inequalities hold:*

$$\delta(i, j) + \delta(k, l) \leq \delta(i, k) + \delta(j, l) \quad \delta(i, l) + \delta(j, k) \leq \delta(i, k) + \delta(j, l).$$

Example 1.4.30 *Consider the metric defined on the split system in 1.4.28. Since N is a circular split system we can see it satisfies the conditions of Theorem 1.4.29. For example $\delta(1, 2) + \delta(3, 4) = 10 + 13 \leq \delta(1, 3) + \delta(2, 4) = 13 + 14$ and $\delta(1, 4) + \delta(2, 3) = 12 + 7 \leq \delta(1, 3) + \delta(2, 4) = 13 + 14$.*

The focus of Chapter 2 will be to characterize when a metric is also the distance function for equidistant circular split networks, as was done for equidistant trees in Theorem 1.4.24.

1.4.3 Phylogenetic models

Another approach to reconstructing phylogenetic trees is to use phylogenetic models based on Markov chains. Here we consider a tree structure with transition matrices at each edge signifying the change that a DNA nucleotide mutates into a different one. Then by looking at the observed probabilities at the leaves one can hopefully ascertain the mutation rate parameters that would produce that distribution. We start with a quick definition of Markov chains before defining the models that use them.

Definition 1.4.31 A Markov model is a set of κ states and probabilities of how those states change into each other, over time, where it is assumed that future states only depend on the present states.

There are many ways to represent a Markov chain but in this thesis we will focus on matrices describing the rates of change between the various states. We start with a definition of rate and transition matrices.

Definition 1.4.32 Let Q be a matrix in $\mathbb{R}^{\kappa \times \kappa}$ with $q_{ij} \geq 0$ for all $i \neq j$. Additionally, for each i we require:

$$\sum_{j=1}^{\kappa} q_{ij} = 0.$$

Then Q is called a rate matrix and the infinitesimal probability that a random variable jumps from state i to state j in a small amount of time dt is $q_{ij}dt$.

Definition 1.4.33 The probability that i ends up in state j after time unit t can be found with the Markov transition matrix $P(t)$ defined as:

$$P(t) = \exp(Qt) = I + Qt + \frac{Q^2 t^2}{2!} + \dots$$

We will be using Markov transition matrices to model evolution by having a set of DNA bases on the leaves as a given and then determining the probability of observing those states under different tree structures and edge transition matrix parameters. One can view this process as each of the vertices having an unknown DNA state and the transition matrices are the probability that the state mutates into a different base after a set amount of time. We define this formally below.

Definition 1.4.34 Let T be a n -leaf, rooted X -tree with root 0 . Let Int_T be the set of internal vertices of T . A κ -state phylogenetic Markov model on T is created by associating a κ -state random variable X_v to each vertex v of T . Let π the distribution of the random variable at the root.

We associate to each edge $e = (u, v)$ of T a $\kappa \times \kappa$ transition matrix M^e such that $M^e(i, j) = P(X_i = x_i | X_{pa(i)} = x_{pa(i)})$ where $pa(i)$ is the parent of i .

The probability of observing a configuration $(x_1, \dots, x_n) \in [\kappa]^n$ at the leaves is given by the formula:

$$P(X_1 = x_1, \dots, X_n = x_n) = \sum_{j \in [\kappa]^{\text{Int}(T)}} \pi_{j_0} \prod_{(u,v) \in E(T)} M^{(u,v)}(j_u, j_v).$$




Figure 1.9: A tree where the probability of observing states at the leaves is described in Example 1.4.35.

Here starting with the root at a state, we can break the cases of $P(X_1 = x_1)$ by going through each edge transition matrix to get to X_1 and multiplying every possible sets of transition matrix entries that end up with $X_1 = x_1$.

Example 1.4.35 Let $\kappa = 4$ and consider the tree shown in Figure 1.9 with phylogenetic Markov model associated to it. Thus there is a transition matrix of indeterminates associated to each edge and a 4-state random variable X_v to each vertex v . We can calculate the probability of observations states on the leaves 1, 2, 3 using:

$$P(X_1 = x_1, X_2 = x_2, X_3 = x_3) = \sum_{i=1}^4 \sum_{j=1}^4 \pi_i \cdot M^{(a \rightarrow 1)}(i, x_1) \cdot M^{(a \rightarrow b)}(i, j) \cdot M^{(b \rightarrow 2)}(j, x_2) \cdot M^{(b \rightarrow 3)}(j, x_3).$$

There are many different phylogenetic models defined in this same way, with a transition matrix associated to each edge. If there are no restrictions on the transition matrix associated to the edges as shown in the example above, this would define what is called the *General Markov Model*. Often simpler models that restrict the entries of the transition matrices are considered for their computational efficiency. We will define a few models here which will be the focus of Chapter 3. Below is the definition of all relevant models but note that these are traditionally defined by the restriction on the rate matrix instead of the transition matrix. In the context of mathematical phylogenetics, $\kappa = 4$ is a common choice, with the 4 states being the 4 DNA bases: A,G,C,T. There are other choices for k in phylogenetics such as $\kappa = 2$ for the Cavendar Farris-Neyman (CFN) model for 2-state sequence evolution where the DNA bases are just regarded as purines (A,G) and pyrimidines (C,T). The choice of $\kappa = 20$ is made when analyzing the 20 amino acids. For the rest of this section we will only consider $\kappa = 4$ and we will treat the DNA bases as simply the numbers 1, 2, 3, 4.

One of the simplest models for the mutation of DNA bases is the Jukes-Cantor model.

Definition 1.4.36 *The Jukes-Cantor model [60] has rate matrix:*

$$Q^{JC} = \begin{bmatrix} -3a & a & a & a \\ a & -3a & a & a \\ a & a & -3a & a \\ a & a & a & -3a \end{bmatrix}.$$

Let $\alpha = \frac{1}{4} - \frac{1}{4}e^{-4at}$ and $\beta = \frac{1}{4} + \frac{3}{4}e^{-4at}$. Then Jukes-Cantor model has transition matrix

$$M(t) = \exp(Q^{JC}(t)) = \begin{bmatrix} \beta & \alpha & \alpha & \alpha \\ \alpha & \beta & \alpha & \alpha \\ \alpha & \alpha & \beta & \alpha \\ \alpha & \alpha & \alpha & \beta \end{bmatrix}$$

This model gives one parameter defining the rate at which a DNA base mutates to anything but itself. Then the rate at which it mutates to itself is determined, since the entries of each row of the rate matrix sum to zero.

Definition 1.4.37 *The K2P and K3P [63] models have rate matrices:*

$$Q^{K2P} = \begin{bmatrix} -a-2b & b & a & b \\ b & -a-2b & b & a \\ a & b & -a-2b & b \\ b & a & b & -a-2b \end{bmatrix}$$

and

$$Q^{K3P} = \begin{bmatrix} -a-b-c & a & b & c \\ a & -a-b-c & c & b \\ b & c & -a-b-c & a \\ c & b & a & -a-b-c \end{bmatrix}$$

respectively.

The K2P can be thought of as having one mutation rate between purines (A,G) and another between pyrimidines (C,T). For the K3P model there is one variable for each mutation rate

between all 4 bases. Both of these more closely model real evolution, at the cost of more computational complexity. Let

$$\begin{aligned}\alpha &= \frac{1}{4} + \frac{1}{4}e^{-4at} - \frac{1}{2}e^{-2(a+b)t} \\ \beta &= \frac{1}{4} - \frac{1}{4}e^{-4bt} \\ \gamma &= \frac{1}{4} + \frac{1}{4}e^{-4bt} + \frac{1}{2}e^{-2(a+b)t}\end{aligned}$$

then the transition matrix for the K2P model is

$$M(t) = \exp(Q^{K2P}(t)) = \begin{bmatrix} \gamma & \beta & \alpha & \beta \\ \beta & \gamma & \beta & \alpha \\ \alpha & \beta & \gamma & \beta \\ \beta & \alpha & \beta & \gamma \end{bmatrix}.$$

Let

$$\begin{aligned}\alpha &= \frac{1}{4}(1 + e^{-2(a+b)t} + e^{-2(a+c)t} + e^{-2(b+c)t}) \\ \beta &= \frac{1}{4}(1 + e^{-2(a+b)t} + e^{-2(a+c)t} - e^{-2(b+c)t}) \\ \gamma &= \frac{1}{4}(1 - e^{-2(a+b)t} - e^{-2(a+c)t} + e^{-2(b+c)t}) \\ \delta &= \frac{1}{4}(1 + e^{-2(a+b)t} - e^{-2(a+c)t} - e^{-2(b+c)t}),\end{aligned}$$

then the transition matrix for the K3P model is

$$M(t) = \exp(Q^{K3P}(t)) = \begin{bmatrix} \alpha & \beta & \gamma & \delta \\ \beta & \alpha & \delta & \gamma \\ \gamma & \delta & \alpha & \beta \\ \delta & \gamma & \beta & \alpha \end{bmatrix}.$$

Definition 1.4.38 *The Strand symmetric model(SSM) [29] has transition matrix:*

$$\begin{bmatrix} a & b & c & d \\ e & f & g & h \\ h & g & f & e \\ d & c & b & a \end{bmatrix}.$$

Most of the aforementioned models, JC, K2P, and K3P are examples of what are called group-based models which we will define now.

Definition 1.4.39 *A phylogenetic model is group-based if for all transition matrices, M , associated to its edges there exists a function $f : G \rightarrow \mathbb{R}$ such that $M(g, h) = f(g - h)$ where G is an abelian group and $M(g, h)$ denotes the probability of going from state h given that the variable is in state g in the previous vertex.*

Group-based models are of interest for numerous reasons, because their structure naturally lends itself to algebraic tools. In addition, they allow for the use of the discrete Fourier transform, which greatly simplifies the models and makes analysis possible. This will be described in Chapter 3.

The SSM model is an *equivariant* model, meaning there is an underlying symmetry on on the transition matrix. The previous group-based models are also equivariant models.

1.5 Outline of Thesis

We now outline the rest of the chapters of the thesis.

1.5.1 Equidistant Circular Split Networks

In Chapter 2 we fully characterize when a metric is also an equidistant circular split network. The content of this chapter is joint work with Seth Sullivant and come from a paper that was published in the *SIAM Journal on Applied Algebra and Geometry* [61].

We start by providing background on circular split networks and their representations. Then we explain the known results on distance functions on circular split networks as well as define a complete circular split network which will be important for the later sections. Next we give a facet description of a cone of distance functions that make a complete circular network equidistant. Then we define a new combinatorial diagram in order to give an extreme ray description of this cone. All non-complete circular networks end up being a lower dimensional face of this cone. Lastly, we show a connection to a well studied polytope, the Chan-Robbins-Yuen polytope from geometric combinatorics.

1.5.2 Identifiability of Group-Based Mixture Models

In Chapter 3 we give identifiability conditions for group-based mixture models, specially the Jukes-Cantor, Kimura 2 parameter model, Kimura 3 parameter model and the Strand Symmetric model. The content of this chapter is also joint work with Seth Sullivant.

We start by giving definitions of the phylogenetic mixture models we are working with and type of identifiability. Then we give background on group-based models, the discrete Fourier transform, and tensors, all of which will be needed for the main theorem for identifiability. Next we define a series of properties that if a phylogenetic model possesses, we can prove its identifiability within certain bounds on its size. Lastly we prove that all of these models satisfy these conditions, utilizing the fact that some of these models are submodels of other models.

1.5.3 Semialgebraic Hypothesis Testing with Incomplete U-Statistics: Practical Issues

In Chapter 4 we explore the practical issues of a new hypothesis test by applying it to numerous phylogenetic models. The contents of this chapter are joint work with David Barnhill, Marina Garrote-Lopez, Elizabeth Gross, Max Hill, John Rhodes, and Joy Zhang and are in a preprint [13].

First we gave background on the incomplete U-statistics test that we investigate. We then give background on 4 phylogenetic models that we tested. Next in detail we explain the practical considerations one must think about when applying this test, including the size of the test parameters, the inequalities used to define your model, the symmetrization of the model, and breaking the model into irreducible components. Lastly we analyze this test on the Cavendar-Farris-Neyman (CFN) model for 2-state sequence evolution and detail the observation learned from that model.

CHAPTER

2

EQUIDISTANT CIRCULAR SPLIT NETWORKS

2.1 Introduction

Phylogenetics concerns uncovering evolutionary relationships between collections of species. Traditionally, these relationships are represented by trees. The combinatorics of rooted tree structures, and distances derived from trees are a staple of phylogenetic inference and at the heart of much of the mathematics of evolutionary biology. This is the perspective in classic books like [47, 84]. However, the presence of evolutionary processes that produce non-tree-like structures among species have been realized to play an important role in evolution.

Non-tree-like evolutionary processes include horizontal gene transfer, hybridization, and introgression. It is desirable to have phylogenetic structures that can encode these types of more complex, non-tree-like relationships. This has led to the creation and study of phylogenetic networks as a tool for phylogenetic inference, where the network structure can encode different types of non-tree-like relationships.

There are a few different approaches to making phylogenetic networks (see [86, Ch 10]).

Some of the choices of which network structure to use are based on which modeling paradigm is employed, others are based on which inference techniques are being used, and yet others are just based on whether the mathematics is interesting. See [64] for a discussion of the different classes of networks. In this chapter, we study the mathematics of distances based on phylogenetic networks. For this approach, one of the most natural network structures to study is based on split networks. This is because split networks are naturally tied to cut-semimetrics and the cut cone, so they naturally fit into the framework of those well-studied objects [39].

The most studied family of non-tree-like split networks is the family of circular split networks because they are the type of split network produced by the NeighborNet algorithm [19], a widely used algorithm, cited by 2318 papers on Google Scholar as of August 2024. The geometry of metrics associated with circular split systems is well-studied. The set of metrics compatible with a particular circular ordering are the Kalmanson metrics. Kalmanson metrics associated with the standard ordering $1, \dots, n$ are metrics δ on $[n]$ that satisfy the inequalities

$$\delta(i, j) + \delta(k, l) \leq \delta(i, k) + \delta(j, l) \quad \text{and} \quad \delta(i, l) + \delta(j, k) \leq \delta(i, k) + \delta(j, l)$$

whenever $i < j < k < l$. This condition is also famous in combinatorial optimization because the traveling salesman problem can be solved in polynomial time if distance constraints for the problem come from a Kalmanson metric [62]. There are a number of papers that explore the connection between Kalmanson metrics and phylogenetics [38, 69].

In this chapter, we explore a variation of metrics associated to circular split systems where we add the extra condition that the metric is equidistant. The equidistant condition means that the network has a special root vertex, and that each of the vertices $i \in [n]$ is at the same distance from the root. In the context of biology, an equidistant circular split network would represent a situation where there is a set of species that evolve at similar rates and are hybridizing with each other. The cone of equidistant circular split networks is the space of all possible distance functions that make an equidistant circular split network on those species. Thus if there is a set of closely related species that are suspected of hybridization, one could see, using the facet description, if the vector of pairwise distances are inside or near the cone of equidistant circular split networks. In addition, the study of the cone of equidistant circular split networks is of mathematical interest as a step towards creating a network version of UPGMA [75]. UPGMA is a popular algorithm for creating phylogenetic trees from pairwise distances which always produces equidistant trees. In the same way that NeighborNet is a generalization of the neighbor joining algorithm, understanding equidistant circular split networks will help in creating a similar algorithm generalization.

Our main results are a characterization of the inequalities and the extreme rays that define the cone of equidistant circular split networks. The resulting inequality system that arises is a kind of restriction of the Kalmanson conditions to take into account the equidistant condition shown in Theorem 2.4.1. On the other hand, while the general Kalmanson cone has only $\binom{n}{2}$ extreme rays, the general cone of equidistant circular metrics has $2^{n-1} - 1$ extreme rays shown, as in Theorem 2.5.22. In addition, we show that every face of the cone circular equidistant networks also corresponds to metrics for subnetworks, and characterize their inequalities and extreme rays. Finally, we show that the cone of equidistant circular split networks is closely related to the Chan-Robbins-Yuen polytope from geometric combinatorics.

2.2 Split Networks

We introduce the notion of split system and split networks. Split networks generalize phylogenetic trees by allowing for some limited cycles in the graph structure. This is inspired by hybridization and reticulation events in biology, which introduce cycles into phylogenetic trees. The material in this section is standard in the literature, and more background on split systems and split networks can be found in [86].

Definition 2.2.1 *Let X be a set of labels with $|X| = n$. A split $A|B$ is a partition of X into two nonempty sets. A split is a trivial split if one part of the partition has cardinality one. A set of splits is called a split system.*

Split graphs and split networks are visual tools used to represent a split system.

Definition 2.2.2 *Let $G = (V, E)$ be a connected bipartite finite graph, K a finite set of labels, and s a surjective map $s : E \rightarrow K$. The pairing (G, s) is a split graph if for all $u, v \in V$ and for each shortest path p between u and v , s maps the edges on p one-to-one to $S(u, v) \subseteq K$ with $S(u, v)$ the same for all such p .*

One feature of a split graph is that for every $k \in K$, removing all the edges with label k breaks the graph into two components, as seen in the following result.

Proposition 2.2.3 [57] *Let (G, s) be a split graph with $G = (V, E)$ and $s : E \rightarrow K$. For any $k \in K$, let $E_k = \{e \in E : s(e) = k\}$. Then the graph $(V, E \setminus E_k)$ has exactly two connected components for every $k \in K$.*

Now the notion of split graph can be combined with a split system to define the notion of a split graph representing a split system, which is called a split network.

Definition 2.2.4 *Let N be a split system on X . Let (G, s) be a split graph with $G = (V, E)$ and $s : E \rightarrow K$. Let $f : X \rightarrow V$ be a map such that for all $A|B \in N$ there exists a $k(A|B) \in K$ such that $f(A)$ and $f(B)$ are exactly in the two connected components of $(V, E - E_{k(A|B)})$. Furthermore, assume that each element in K corresponds to an element in N . Then (G, s, f) is a split network that represents N .*

Example 2.2.5 *Consider the bipartite graph G drawn in Figure 2.1. Let $s : E_G \rightarrow \{\alpha, \beta, \gamma, \delta, \epsilon, \zeta, \eta, \theta, \iota\}$ be such that*

$$\begin{aligned} s(\{c, d\}) &= s(\{f, g\}) = s(\{h, m\}) = s(\{d, i\}) = \alpha, \\ s(\{c, g\}) &= s(\{d, f\}) = s(\{k, m\}) = \beta, \\ s(\{g, m\}) &= s(\{f, h\}) = s(\{i, k\}) = \gamma, \\ s(\{a, c\}) &= \delta, s(\{b, d\}) = \epsilon, s(\{e, d\}) = \zeta, s(\{g, j\}) = \eta, s(\{k, l\}) = \theta, s(\{m, n\}) = \iota. \end{aligned}$$

Then the pairing (G, s) is a split graph. Additionally, consider the split system

$$N = \{1456|23, 1234|56, 1236|45\} \cup \{i|[6] \setminus i : i \in [6]\}.$$

Then if f is the map:

$$f(a) = 1, f(b) = 2, f(e) = 3, f(l) = 4, f(n) = 5, f(j) = 6,$$

(G, s, f) is a split network representing N .

A split system is represented by a split network if the split network has some edge or set of edges that realize every split in the split system. The following definition and theorem exactly characterize when a split system is represented by a tree.

Definition 2.2.6 *A split system N is pairwise compatible if for every pair of splits $A_1|B_1, A_2|B_2 \in N$ at least one of the following sets is empty:*

$$A_1 \cap A_2, A_1 \cap B_2, A_2 \cap B_1, B_1 \cap B_2.$$




Figure 2.1: An example of a split graph whose corresponding functions are color-coded and described in Example 2.2.5. Here α is blue, β is orange, γ is red, and all other function values are black

Definition 2.2.7 Let (G, s, f) be a split network on X . Let $\Sigma(G)$ be all of the splits of X that are induced by edge classes of G .

Theorem 2.2.8 (*Splits Equivalence Theorem*) Let N be a split system on X . Then there exists an X -tree T with $\Sigma(T) = N$ if and only if the splits in N are pairwise compatible. Furthermore, the tree T is uniquely determined.

For a detailed proof of Theorem 2.2.8 see [84]. In general split systems need not be pairwise compatible. Thus, general split systems extend trees to the case of non-compatible splits. In a general split network, each split is represented by a set of parallel edges that disconnect the graph according to that partition.

Example 2.2.9 Consider the split system N on 6 leaves

$$N = \{12|3456, 1265|34, 1234|56\} \cup \{i|[6] \setminus i : i \in [6]\}.$$

This set of splits is pairwise compatible. Thus by the Split Equivalence Theorem, there exists a unique tree with $\Sigma(T) = N$, shown in Figure 2.2,

A consequence of Theorem 2.2.8 is that a split system N will be pairwise compatible if and only if there is some tree T such that (T, s, f) represents N .

A generic split system will not necessarily have a split network that can be drawn as a planar graph. However, adding the following circular condition to the split system guarantees that the split network graph is planar. Aside from the planar nature of circular split systems, they also have the advantage of being easy to represent and have nice mathematical properties [11, 19, 38, 69].




Figure 2.2: The unique tree representing the split system $N = \{12|3456, 1265|34, 1234|56\} \cup \{i|[6] \setminus i : i \in [6]\}$ in Example 2.2.9, which must exist by Theorem 2.2.8.

Definition 2.2.10 A split system N with leaf labels X is circular with respect to some cyclic ordering (x_1, \dots, x_n) of X if every split of S is of the form

$$x_{i+1}, \dots, x_j | x_{j+1}, \dots, x_i$$

for some i and j , where the indices are considered cyclically modulo n (e.g. $x_{n+1} = x_1$).

The following algorithm constructs a split network that represents a circular split system N .

Algorithm 2.2.11 Circular Network Algorithm [42]

Let N be a split system with n leaves and the corresponding trivial split for each leaf, i.e. $\{i|[n] \setminus i : i \in [n]\} \in N$.

1. Construct a star graph with n leaves labeled $1, \dots, n$.
2. Let $i \dots j-1 | j \dots n \dots i-1 \in N$. Find a path p from i to $j-1$ that uses the least amount of edges. Let the size of the number of internal edges in p be g .
3. Let $p = i e_0 u_1 e_1 \dots u_g e_g (j-1)$ where e_k and u_k for $k \in [g]$ are internal edges and vertices along p respectively. Create a copy of $e_0 u_1 e_1 \dots u_g e_g$ which we will call $e'_0 u'_1 e'_1 \dots u'_g e'_g$.
4. Assume that the edges of u_k are $\{e_{k-1}, l_1, \dots, l_h, e_k, r_1, \dots, r_p\}$ where $l_1 \dots l_h$ are edges that move closer to the leaves $i, \dots, j-1$ and $r_1 \dots r_p$ are the edges that move closer to the leaves $j, \dots, n, \dots, i-1$. Let f_k be a new edge representing the split $i \dots j-1 | j \dots n \dots i-1$.




Figure 2.3: The tree representing the split system $N = \{12|3456, 1265|34, 1234|56\} \cup \{i|[6] \setminus i : i \in [6]\}$ in Example 2.2.9. The process for constructing this tree is shown in Algorithm 2.2.11




Figure 2.4: A visualization of performing Algorithm 2.2.11 on the split system $N = \{12|3456, 16|2345, 123|456, 126|345\}$ from Example 2.2.12.

Then change the edges of u_k to $\{e_{k-1}, l_1, \dots, l_h, e_k, f_k\}$ and the edges of u'_k to

$$\begin{cases} (f_k, e'_i, r_1, r_2, \dots, r_p) & \text{if } k = 1 \\ (e'_{k-1}, f_k, e'_k, r_1, r_2, \dots, r_p) & \text{if } 1 < k < g \\ (e'_{k-1}, f_k, r_1, r_2, \dots, r_p) & \text{if } k = g \end{cases} \quad (2.2.1)$$

5. Repeat Steps 2 through 4 for every split in N .

Constructing a split network from a split system can also be viewed as starting with a star graph and “pulling” the two sides of each additional split in a different direction, splitting in half any edges as necessary.

Example 2.2.12 Consider the split system N on 6 leaves with splits

$$\{12|3456, 16|2345, 123|456, 126|345\} \cup \{i|[6] \setminus i : i \in [6]\}.$$

These splits are not pairwise compatible. This can be seen with the splits $A_1|B_1 = 12|3456$ and




Figure 2.5: Applying Algorithm 2.2.11 to the split system $N = \{123|456, 126|345, 156|234\}$ from Example 2.2.13 by applying the splits in two different orders, resulting in different graphs

$A_2|B_2 = 16|2345$ and $A_1 \cap A_2 = 1$, $A_1 \cap B_2 = 2$, $B_1 \cap A_2 = 6$, $B_1 \cap B_2 = 34$, none of which are empty. See Figure 2.4 for a visualization of a split network representing N using Algorithm 2.2.11.

In addition, see Figure 2.3 for a visualization of a tree representing N in Example 2.2.9 using Algorithm 2.2.11.

For any particular drawing of a split network, the order in which the splits are drawn changes the resulting graph which can be seen in the following example.

Example 2.2.13 Consider the split system N on 6 leaves with splits

$$\{123|456, 126|345, 156|234\} \cup \{i|[6] \setminus i : i \in [6]\}.$$

If Algorithm 2.2.11 is performed by applying the nontrivial splits in the order $(123|456, 126|345, 156|234)$, it will result in a different graph than if they are applied in the order $(156|234, 123|456, 126|345)$. This is visualized in Figure 2.5.

Thus, unlike the tree case, there is a choice to be made for the order of the splits in a circular split network to apply Algorithm 2.2.11. Rather than fixing one particular ordering, the choice

made in this chapter was to view circular split networks as simply a set of splits, visualizing them using their dual polygon representation. This choice maintains an independence of the order of the circular split networks.

Definition 2.2.14 *Let N be a circular split system with leaves labeled $0, 1, \dots, n$. The dual polygon representation of N is constructed in the following way:*

Take an $n+1$ -gon and label the edges sequentially clockwise with $0, 1, \dots, n$. Label the vertices by the edge that is adjacent to it, clockwise. Let $i \dots j-1 | j \dots 0n \dots i-1$ where $i < j \in [n]$, be a non-trivial split in N . Then, $i \dots j-1 | j \dots 0n \dots i-1$ is represented by the diagonal of the $n+1$ -gon that connects the vertex i to the vertex j .

With the labeling above, the edges of the $n+1$ -gon are labeled by the numbers $0, 1, \dots, n$ in such a way that the diagonal corresponding to the nontrivial split $i \dots j-1 | j \dots 0n \dots i-1$ separates the edges into the two sets $\{i, \dots, j-1\}$ and $\{j \dots 0n \dots i-1\}$. The trivial splits of N correspond to the sides of the $n+1$ -gon.

We now introduce the notion of rooted graphs because the main focus in the rest of this chapter will be on rooted split systems, as that is key for the equidistant property.

Definition 2.2.15 *A graph G is rooted if one of its vertices has been specially designated as a root.*

In the context of biology, the root is the most recent common ancestor of all of the species in the network. In this chapter, for any rooted split system, the root will be labeled 0 and will be a leaf in the split network representing the rooted split system. The n other leaves will have the labels 1 through n . The root will be at the top of any picture, with its one leaf edge suppressed, and all of the rest of the leaves will be at the bottom.

Example 2.2.16 *Consider the split system N on 5 leaves and one root 0 with the following splits:*

$$N = \{01|2345, 12|0345, 0145|23, 0123|45\} \cup \{i|\{0, 1, 2, 3, 5\} \setminus i : i \in [5]\}.$$

The dual polygon representation for this split system is a hexagon with the edges labeled 0 through 5, the vertices labeled by the edge to their right, and the diagonals connecting the following vertices: $\{0-2, 1-3, 2-4, 4-0\}$. See Figure 2.6 for a visualization of this dual polygon representation. Additionally, in Figure 2.6 there is an application of Algorithm 2.2.11 to N and a rooting of the resulting split network.




Figure 2.6: The Dual Polygon Representation of the Split System $N = \{01|2345, 12|0345, 0145|23, 0123|45\}$ from Example 2.2.16. Then Algorithm 2.2.11 is applied to create a split network representing N . Lastly, the split network is drawn rooted at 0.

A dual polygon representation corresponds to a tree if and only if none of the diagonals intersect each other in their interiors. A representational split network can be constructed from the dual polygon representation by an application of the circular network algorithm to the underlying split system. If the dual polygon representation was rooted, to make the split network rooted, put the leaf labeled 0 at the top and draw the network descending from that vertex.

In addition to the root being the most recent common ancestor of every vertex in a rooted network, rooted split networks have a notion of a poset on the vertices. The join of two species is the most recent common ancestor of those species, making this poset of biological interest.

Definition 2.2.17 *Let G be the graph for a split network (G, s, f) that represents N , a rooted split system. The vertices of G form a poset. The root is the maximal element. Let P_x be the set of paths of minimal length to the root from the vertex x . A vertex y is above x if any of the paths in P_x pass through y . The join of two leaves, i and j will be referred to as the most recent common ancestor of i and j or $mcr a(i, j)$.*

Example 2.2.18 *Consider the split system N on 4 leaves and one root 0 with splits*

$$N = \{01|234, 12|034, 014|23, 04|123, 012|34\} \cup \{i|\{0, 1, 2, 3, 4\} \setminus i : i \in [4]\}.$$

The Hasse diagram for the poset that the split network representing N forms, can be seen in Figure 2.7.




Figure 2.7: The Hasse diagram for the poset that is formed from the split system $N = \{01|234, 12|034, 014|23, 04|123, 012|34\}$ from Example 2.2.18.

2.3 Distances from split networks

Now that we have defined the split systems that we will study in this chapter, we introduce distance matrices associated to the split systems, which generalize tree metrics. Our goal in this chapter will be to give polyhedral characterizations of the dissimilarity maps that can arise from rooted circular split systems. To start, we will define dissimilarity maps and show how they relate to tree metrics.

Definition 2.3.1 A function $\delta : X \times X \rightarrow \mathbb{R}_{\geq 0}$ is called a dissimilarity map if it is a real, symmetric, and nonnegative function. In addition, a dissimilarity map is a metric if it satisfies the triangle inequality on X , meaning for all $x, y, z \in X$,

$$\delta(x, z) \leq \delta(x, y) + \delta(y, z).$$

The following two theorems give necessary and sufficient conditions for when a dissimilarity map comes from a tree or a circular split system. For Theorem 2.3.2 see [84] for more details.

Theorem 2.3.2 (Four Point Condition) Let δ be a dissimilarity map on a finite set X . Then δ is the distance function for a tree if and only if it satisfies the four-point condition: for every four elements $i, j, k, l \in X$, two of the three terms

$$\delta(i, j) + \delta(k, l), \quad \delta(i, l) + \delta(j, k), \quad \delta(i, k) + \delta(j, l).$$

are equal and are greater than or equal to the third.




Figure 2.8: This weighted tree from Example 2.3.3

Example 2.3.3 Consider the dissimilarity map on the set $[5]$:

$$\delta = \begin{bmatrix} 0 & 5 & 15 & 12 & 17 \\ & 0 & 16 & 13 & 18 \\ & & 0 & 11 & 16 \\ & & & 0 & 7 \\ & & & & 0 \end{bmatrix}.$$

This is a distance function for a tree which can be seen by calculating $\delta(i, j) + \delta(k, l)$, $\delta(i, l) + \delta(j, k)$, $\delta(i, k) + \delta(j, l)$ for $i, j, k, l \in [5]$.

For example, for leaves 1, 2, 3, 4:

$$\delta(1, 2) + \delta(3, 4) = 16, \quad \delta(1, 3) + \delta(2, 4) = 28, \quad \delta(1, 4) + \delta(2, 3) = 28.$$

Since every set of 4 leaves satisfies the conditions of Theorem 2.3.2, δ is a distance function for a tree. A tree that represents the dissimilarity can be seen in Figure 2.8. See [86] for further details on this definition of distance.

This notion of distance on a tree can be generalized to split networks using the concept of a split separating leaves which is defined below:

Definition 2.3.4 Let N be a split system, and $i, j \in [n]$. A split $A|B \in N$ separates i and j in N if $i \in A$ and $j \notin A$ or $i \in B$ and $j \notin B$.

Definition 2.3.5 Let $A|B$ be a split and let the separation indicator function for $A|B$ be defined as:

$$\mathbb{1}_{A|B}(i, j) = \begin{cases} 1 & \text{if } A|B \text{ separates } i \text{ and } j \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\mathbb{1}_{A|B}$ is an example of a semimetric, since it is nonnegative, symmetric, and satisfies the triangle inequality. In the context of the theory of finite metric spaces, $\mathbb{1}_{A|B}$ is known as a cut semimetric [39].

Example 2.3.6 Consider the split $12|345$. This split separates 1 from 3, 1 from 4, 1 from 5, 2 from 3, 2 from 4, and 2 from 5. The function $\mathbb{1}_{A|B}$ satisfies $\mathbb{1}_{A|B}(1, 2) = \mathbb{1}_{A|B}(3, 4) = \mathbb{1}_{A|B}(3, 5) = \mathbb{1}_{A|B}(4, 5) = 0$, and $\mathbb{1}_{A|B}(1, 3) = \mathbb{1}_{A|B}(1, 4) = \mathbb{1}_{A|B}(1, 5) = \mathbb{1}_{A|B}(2, 3) = \mathbb{1}_{A|B}(2, 4) = \mathbb{1}_{A|B}(2, 5) = 1$.

Definition 2.3.7 Let N be a split system with n leaves. For each split $A|B \in N$, let $a_{A|B} \in \mathbb{R}_{\geq 0}$ be a weight. Let \mathbf{a} be the vector of weights for every split. The distance function between any two leaves $i, j \in [n]$ in the split system N with weights \mathbf{a} , denoted $\delta_{N, \mathbf{a}}$, is defined as follows:

$$\delta_{N, \mathbf{a}}(i, j) = \sum_{A|B \in N} a_{A|B} \mathbb{1}_{A|B}(i, j).$$

If it is clear in context which N and \mathbf{a} are being used we will write $\delta_{N, \mathbf{a}}(i, j) = \delta(i, j)$. If a split system is rooted, $\delta(0, i)$ will denote the distance from the root 0 to $i \in [n]$. Note that if N is a split network and the edges are labeled with the weight from the split they are realizing, then this distance between leaves is exactly the distance obtained from the sum of the edge weights along any shortest path between the leaves. Since sets of edges that all realize the same split have the same weight, the distance between any two leaves, $i, j \in [n]$, can be calculated in the dual polygon representation by drawing a line between the sides of the n -gon for those leaves and then taking the sum over weights of all diagonals that line intersects, as well as $a_{1 \dots (i-1)(i+1) \dots n | i}$ and $a_{1 \dots (j-1)(j+1) \dots n | j}$.

For circular split systems, there is a nice condition for when a metric represents the system, which can be seen in the following theorem:

Theorem 2.3.8 [36] Let δ be a metric on a finite set X . Then δ is the distance function for a circular split system N with weights \mathbf{a} if and only if it satisfies the Kalmanson condition with respect to some circular ordering π . That is, for every set of leaves $i < j < k < l$ under π , both inequalities hold:

$$\delta(i, j) + \delta(k, l) \leq \delta(i, k) + \delta(j, l) \quad \delta(i, l) + \delta(j, k) \leq \delta(i, k) + \delta(j, l).$$




Figure 2.10: The equidistant split network from Example 2.3.11

Definition 2.3.10 Let N be a rooted circular split system with n leaves, weights \mathbf{a} , and root labeled 0. The distance function $\delta_{N,\mathbf{a}}$ is equidistant if the distance from the root to every leaf is equal. That is, $\delta_{N,\mathbf{a}}$ is equidistant if for all $i, j \in [n]$, $\delta_{N,\mathbf{a}}(0, i) = \delta_{N,\mathbf{a}}(0, j)$.

The pair $(N, \delta_{N,\mathbf{a}})$ will be called an equidistant circular split system if $\delta_{N,\mathbf{a}}$ is equidistant. Note that a consequence of this condition is that the distance from every internal vertex to the leaves below it must also be equal.

Example 2.3.11 Consider the rooted split network from Example 2.2.16. If weights are applied to the splits of the split system as shown in Figure 2.10, this will be an equidistant circular split system where the distance from the root to each of the leaves is 8. The resulting dissimilarity map is

$$\delta = \begin{bmatrix} 0 & 12 & 16 & 16 & 16 \\ & 0 & 8 & 14 & 14 \\ & & 0 & 14 & 14 \\ & & & 0 & 4 \\ & & & & 0 \end{bmatrix}.$$

Now that equidistant split networks have been defined, we can consider the cone of all possible distance functions that make a particular split network equidistant. This cone, which is the main interest of this chapter, is the space of equidistant circular split networks and thus, membership in the cone indicates that a set of species has this particular hybridization structure.

Definition 2.3.12 Let N be a rooted circular split system. Then the Equidistant Cone of N is

$$EDC_N = \{\delta : \delta = \delta_{N,\mathbf{a}} \text{ for some } \mathbf{a} \text{ and } (N, \delta_{N,\mathbf{a}}) \text{ is equidistant}\} \subseteq \mathbb{R}^{\binom{n}{2}}. \quad (2.3.1)$$

The use of the word “cone” in the definition is made clear from the following proposition.

Proposition 2.3.13 *The space EDC_N is a polyhedral cone.*

Proof: Let N be a rooted circular split system on $[n]$ leaves and root 0. Consider the space

$$C_N = \{\delta : \delta = \delta_{N,\mathbf{a}} \text{ for some weight vector } \mathbf{a}\}.$$

Since

$$\delta_{N,\mathbf{a}}(i, j) = \sum_{A|B \in S} a_{A|B} \mathbb{1}_{A|B}(i, j),$$

each point in C_N is a positive combination of separation indicator functions, $\mathbb{1}_{A|B}$, and thus C_N is a polyhedral cone. Now EDC_N can be obtained by taking the cone C_N , intersecting it with the linear space where for all $i, j \in [n]$, $\delta_{N,\mathbf{a}}(0, i) = \delta_{N,\mathbf{a}}(0, j)$, and then projecting away the coordinates $\delta(0, i)$. Intersecting with a linear space and then projecting preserves being a polyhedral cone and so EDC_N is a polyhedral cone as well. \square

We wish to understand the structure of EDC_N for every circular split network, but first the following specific circular split system is introduced as every other circular split system is a subset of it. Thus studying the properties of this split system will provide insight into all other circular split systems as well.

Definition 2.3.14 *Let the complete circular split system, KN_n , on n leaves and root 0 be the rooted circular split system with every diagonal in the dual polygon representation and every trivial split. In other words it has splits:*

$$i \dots j | j+1 \dots n \ 0 \dots i-1$$

where $i < j \in [n]$ and

$$i | 01 \dots i-1 \ i+1 \dots n$$

where $i \in [n]$.

The complete circular split system has every possible circular split for a particular circular ordering. The complete circular split system KN_n has $\binom{n+1}{2} - 1$ splits, $\binom{n}{2} - 1$ of which come from diagonals of the $n+1$ -gon and n of which are trivial splits. Note we do not have a trivial split for the root 0 which is why KN_n has $\binom{n+1}{2} - 1$ splits instead of $\binom{n+1}{2}$.

2.4 Facets of EDC_{KN_n}

Our goal in this section and the next one is to give a complete polyhedral description of the cone EDC_{KN_n} . Our main result describes the facet defining inequalities of EDC_{KN_n} shown in Theorem 2.4.1. In Section 2.5, we determine the extreme rays of EDC_{KN_n} . All other circular split systems are contained in KN_n and thus EDC_{KN_n} can be used to characterize EDC_N for any circular rooted split system N .

Theorem 2.4.1 *Let EDC_{KN_n} be the equidistant cone of KN_n . Let $[i, j] = \{i, i + 1, \dots, j\}$. The facets of EDC_{KN_n} are*

$$\delta(1, i) \leq \delta(1, i + 1) \text{ where } i \in [2, n - 1], \quad (\text{left inequalities})$$

$$\delta(i, n) \leq \delta(i - 1, n) \text{ where } i \in [2, n - 1], \quad (\text{right inequalities})$$

$$\delta(i - 1, i + 1) \leq \delta(i - 1, i) + \delta(i, i + 1) \text{ where } i \in [2, n - 1], \quad (\text{triangle inequalities})$$

$$\delta(i, j) + \delta(i - 1, j + 1) \leq \delta(i, j + 1) + \delta(i - 1, j) \text{ where } i < j \in [2, n - 1]. \quad (\text{covering inequalities})$$

Note that triangle inequalities are the same as covering inequalities where $i = j$. The proof of this theorem will require the following lemmas and definitions.

Definition 2.4.2 *Let N be a split system and $I, J \subseteq [n]$. Let the set of separating splits for I and J , $S_{I|J}$, be the set of splits $A|B \in N$ for every $i \in I$ and $j \in J$, i is separated from j in $A|B$.*

This definition is important because

$$\delta(i, j) = \sum_{A|B \in S_{i|j}} a_{A|B}.$$

Example 2.4.3 *Consider the split system KN_6 . In this split system, $S_{2|0,5}$ consists of the following splits:*

$$\{03456|12, 013456|2, 0456|123, 0156|234, 056|1234, 01456|23\}$$

The following lemma provides equations to recover the weights \mathbf{a} from a distance vector δ in EDC_{KN_n} .

Lemma 2.4.4 Let KN_n be a complete circular network on n leaves and root 0. Let $\delta = \delta_{KN_n, a}$, then following equalities hold for $i, j \in [2, n-1]$:

$$\begin{aligned}\delta(1, i+1) - \delta(1, i) &= 2a_{12\dots i|i+1\dots n0}, \\ \delta(i-1, n) - \delta(i, n) &= 2a_{01\dots i-1|i+1\dots n}, \\ \delta(i-1, i) + \delta(i, i+1) - \delta(i-1, i+1) &= 2a_{1\dots(i-1)(i+1)\dots n0|i}, \\ \delta(i, j+1) + \delta(i-1, j) - \delta(i, j) - \delta(i-1, j+1) &= 2a_{1\dots i-1j+1\dots n0|i+1\dots j-1j}.\end{aligned}$$

Note that these equations give us formulas for every parameter, except the trivial splits for the leaves labeled 1 and n . These parameters are not necessary because of the equidistant condition, they can be written in terms of the other weights.

Proof: We go through the formulas one by one. First, for the difference $\delta(1, i+1) - \delta(1, i)$, since $\delta(1, i+1) = \sum_{A|B \in S_{1|i+1}} a_{A|B}$ and $\delta(1, i) = \sum_{A|B \in S_{1|i}} a_{A|B}$, the difference can be rewritten as:

$$\delta(1, i+1) - \delta(1, i) = \sum_{A|B \in S_{1|i+1,0}} a_{A|B} + a_{12\dots i|i+1\dots n0} - \sum_{C|D \in S_{1|i+1,0,1}} a_{C|D}.$$

Using the fact that the most recent common ancestor of i and $i+1$ must also be equidistant from i and $i+1$ the following equality must also be true:

$$\sum_{A|B \in S_{1|i+1,0}} a_{A|B} = \sum_{C|D \in S_{1|i+1,0}} a_{C|D} = a_{12\dots i|i+1\dots n0} + \sum_{C|D \in S_{1|i+1,0,1}} a_{C|D}.$$

Thus,

$$\delta(1, i+1) - \delta(1, i) = 2a_{12\dots i|i+1\dots n0} + \sum_{C|D \in S_{1|i+1,0,1}} a_{C|D} - \sum_{C|D \in S_{1|i+1,0,1}} a_{C|D} = 2a_{12\dots i|i+1\dots n0}.$$

Using a similar argument and symmetry, we can also derive:

$$\delta(i-1, n) - \delta(i, n) = 2a_{01\dots i-1|i+1\dots n}.$$

For the difference $\delta(i-1, i) + \delta(i, i+1) - \delta(i-1, i+1)$, since $\delta(i-1, i) = \sum_{A|B \in S_{i-1|i}} a_{A|B}$, $\delta(i, i+1) = \sum_{A|B \in S_{i|i+1}} a_{A|B}$ and $\delta(i-1, i+1) = \sum_{A|B \in S_{i-1|i+1}} a_{A|B}$, the difference can be rewritten as:

$$\sum_{A|B \in S_{i-1|i}} a_{A|B} + \sum_{C|D \in S_{i|i+1}} a_{C|D} - \sum_{E|F \in S_{i-1|i+1}} a_{E|F}.$$

Note both of $S_{i-1|i}$ and $S_{i|i+1}$ include the split $1 \dots (i-1)(i+1) \dots n0|i$, which $S_{i-1|i+1}$ does not. Thus,

$$\delta(i-1, i) + \delta(i, i+1) - \delta(i-1, i+1) = 2a_{1 \dots (i-1)(i+1) \dots n0|i}.$$

Lastly, for $\delta(i, j+1) + \delta(i-1, j) - \delta(i, j) - \delta(i-1, j+1)$ the terms $\delta(i, j) = \sum_{A|B \in S_{i|j}} a_{A|B}$ and $\delta(i-1, j+1) = \sum_{A|B \in S_{i-1|j+1}} a_{A|B}$. Note that both of $S_{i|j}$ and $S_{i-1|j+1}$ do not include the split $1 \dots i-1j+1 \dots n0|ii+1 \dots j-1j$. Similarly, the terms $\delta(i, j+1) = \sum_{A|B \in S_{i|j+1}} a_{A|B}$ and $\delta(i-1, j) = \sum_{A|B \in S_{i-1|j}} a_{A|B}$. Note that both of $S_{i|j+1}$ and $S_{i-1|j}$ do include the split $1 \dots i-1j+1 \dots n0|ii+1 \dots j-1j$. Thus the difference becomes:

$$\delta(i, j+1) + \delta(i-1, j) - \delta(i, j) - \delta(i-1, j+1) = 2a_{1 \dots i-1j+1 \dots n0|ii+1 \dots j-1j}.$$

□

The following related cone is parameterized by weights instead of pairwise distances. This cone is introduced because it will be easier to prove the last needed Lemma 2.4.6 in this other space.

Definition 2.4.5 Let EDC_N be the equidistant cone of N . For $\delta \in EDC_N$, let $\mathbf{a}_\delta \in \mathbb{R}^{\binom{n+1}{2}-1}$ be the vector of weights obtained from δ using the equations in Lemma 2.4.4. Then, $EDWC_N$, the weighted equidistant cone of N is :

$$\{\mathbf{a} : \mathbf{a} = \mathbf{a}_\delta \text{ for some } \delta \in EDC_N\} \subseteq \mathbb{R}^{\binom{n+1}{2}-1}.$$

Since the equations in Lemma 2.4.4 are a linear map of EDC_N , $EDWC_N$ is a polyhedral cone as well. The distance between leaves can be generalized to distance between arbitrary vertices in a split network using the previously stated formula

$$\delta(v, w) = \sum_{A|B \in S_{v|w}} a_{A|B}$$

where v and w are now arbitrary vertices in the split network and $S_{v|w}$ is the set of spits that are induced by the edge classes that any shortest path between u and v must have. This is the natural generalization of the previous definition of $S_{i|j}$, the set of splits that separates leaves i and j .

Lemma 2.4.6 Let $EDWC_{KN_n}$ be the weighted equidistant cone of KN_n . Let $\mathbf{a} \in EDWC_{KN_n}$. The number of non-redundant linear equalities between the weights that are the entries of the vector \mathbf{a} is $n-1$. That is, $EDWC_{KN_n}$ is contained in a linear space of codimension $n-1$.

Proof: First, we claim that every equality between the weights $a_{A|B} \in \mathbf{a}$ is generated by some internal vertex and its distance to the leaves below it. This is because $\mathbf{a} \in EDWC_{KN_n}$ and $EDWC_{KN_n}$ is an alternate parametrization of EDC_{KN_n} . Therefore, since EDC_{KN_n} is generated by the Kalmanson condition inequalities, shown in Theorem 2.3.8 and the equidistant condition (for all $i, j \in [n]$, $\delta_{N,\mathbf{a}}(0, i) = \delta_{N,\mathbf{a}}(0, j)$), $EDWC_{KN_n}$ is governed by these same inequalities and equalities. Translating these equations to be among the weights $a_{A|B} \in \mathbf{a}$, all of the equalities are generated by one of the equidistant conditions, since the Kalmanson condition itself has no equalities. Thus, since the equidistant condition can be instead thought of as a series of equalities generated by some internal vertex and its distance to the leaves below it, the claim is true.

Let $v_{i,j}$ be the most recent common ancestor of i and j with $i, j \in [n]$. The distance from $v_{i,j}$ to i is

$$\sum_{A|B \in S_{i|j,0}} a_{A|B}.$$

The distance from $v_{i,j}$ to j is

$$\sum_{C|D \in S_{j|i,0}} a_{C|D}.$$

Thus since $\delta(v_{i,j}, i) = \delta(v_{i,j}, j)$, we have:

$$\sum_{A|B \in S_{i|j,0}} a_{A|B} = \sum_{C|D \in S_{j|i,0}} a_{C|D}. \quad (2.4.2)$$

Consider the equality generated by each $v_{i,i+1}$ for all $i \in [n-1]$. They look like:

$$\sum_{A|B \in S_{i|i+1,0}} a_{A|B} = \sum_{C|D \in S_{i+1|i,0}} a_{C|D}.$$

There are $n-1$ of them and no combination of them makes any of the others redundant. Ordering these from vertices $v_{1,2}, v_{2,3}, \dots$, we see that adding a new equality to the list in order involves at least one new $a_{A|B}$ term that does not appear in prior equalities.

Thus it just remains to be shown that any equality between the weights $a_{A|B} \in \mathbf{a}$ can be written as a combination of those generated by $v_{i,i+1}$ for $i \in [n-1]$. Consider again the equality generated by $v_{i,j}$ which is:

$$\sum_{A|B \in S_{i|j,0}} a_{A|B} = \sum_{C|D \in S_{j|i,0}} a_{C|D}.$$

Then since,

$$S_{i|j,0} = \bigcup_{i'=i}^{i'=j-1} S_{i'|i'+1}$$

and

$$S_{j|i,0} = \bigcup_{i'=j}^{i'=i+1} S_{i'|i'-1},$$

the equality

$$\sum_{A|B \in S_{i|j,0}} a_{A|B} = \sum_{C|D \in S_{j|i,0}} a_{C|D},$$

is a combination of the equalities generated by $v_{i',i'+1}$ for $i' \in [i, j-1]$, namely:

$$\sum_{A|B \in S_{i|j,0}} a_{A|B} = \sum_{i'=i}^{j-1} \sum_{A|B \in S_{i'|i'+1}} a_{A|B} = \sum_{j'=i+1}^j \sum_{C|D \in S_{j'|j'-1}} a_{C|D} = \sum_{C|D \in S_{j|i,0}} a_{C|D}.$$

□

Proof: [Proof of Theorem 2.4.1] Let $c' \in \mathbb{R}^{\binom{n}{2}}$ and let

$$\boldsymbol{\delta} = \begin{bmatrix} \delta(1,2) \\ \delta(1,3) \\ \vdots \\ \delta(n-1,n) \end{bmatrix} \in EDC_{KN_n} \subseteq \mathbb{R}^{\binom{n}{2}} \quad (2.4.3)$$

be a vector of pairwise distances for a distance function δ . Let $c' \boldsymbol{\delta} \geq 0$ for all $\boldsymbol{\delta} \in EDC_{KN_n}$ be an arbitrary valid inequality on EDC_{KN_n} . Let M be a $\binom{n}{2}$ by $\binom{n+1}{2} - 1$ matrix whose rows are indexed by pairwise distances $\delta(i, j)$ with $i, j \in [n]$ and the columns are indexed by weights $a_{A|B}$ for a split $A|B \in KN_n$. Let the $(\delta(i, j), a_{A|B})$ th entry of M be $\mathbb{1}_{A|B}(i, j)$. The matrix M is a transformation from coordinates in pairwise distances $\delta(i, j)$ to coordinates in weights $a_{A|B}$ based on the parameterization of KN_n .

Let $\mathbf{a} \in EDWC_{KN_n}$. Let $c' M = c$, then $c \mathbf{a} \geq 0$ is a valid inequality for all $\mathbf{a} \in EDWC_{KN_n}$. $EDWC_{KN_n}$ can be defined by $\mathbf{A} \mathbf{a} = 0$ and $\mathbf{B} \mathbf{a} \geq 0$ where \mathbf{A} is a $n-1$ by $\binom{n+1}{2} - 1$ matrix whose rows are the coefficients of the non-redundant equidistant equalities of KN_n from Lemma 2.4.6 and $\mathbf{B} = I_{\binom{n+1}{2} - 1}$ which is from the non-negativity condition on the weights. Thus by the

Farkas Lemma, there exists a $\mu \in \mathbb{R}^{n-1}$ and a $\lambda \in \mathbb{R}_{\geq 0}^{\binom{n+1}{2}-1}$ such that

$$c = \mu A + \lambda B.$$

Consider $c - \mu A = \lambda B$. Since $B = I_{\binom{n+1}{2}-1}$, $\lambda B = \lambda \in \mathbb{R}_{\geq 0}^{\binom{n+1}{2}-1}$. Thus $(c - \mu A)\mathbf{a} = \lambda\mathbf{a}$ is a non-negative sum of the weights $a_{A|B}$. From Lemma 2.4.4 each of the facets in Theorem 2.4.1 can be reduced to be $2a_{A|B} \geq 0$ for some $A|B \in KN_n$. Thus $(c - \mu A)\mathbf{a}$ can be written as a non-negative sum of the facet inequalities.

Let N be a $\binom{n+1}{2}-1$ by $\binom{n}{2}$ matrix whose rows are indexed by weights $a_{A|B}$ for a split $A|B$ and whose columns are indexed by pairwise distances $\delta(i, j)$ with $i, j \in [n]$. Let the $(a_{A|B}, \delta(i, j))$ th entry of N be the coefficient of $\delta(i, j)$ in the facet inequality that can be reduced to $2a_{A|B} \geq 0$ by Lemma 2.4.4. Note that for $a_{2\dots n0|1}$ and $a_{12\dots n-10|n}$ there is no facet inequality but $\delta(1, 2) = 2a_{2\dots n0|1} \geq 0$ and $\delta(n-1, n) = 2a_{12\dots n-10|n} \geq 0$ can be used. This matrix is a transformation from coordinates in weights $a_{A|B}$ to coordinates in pairwise distances $\delta(i, j)$. Thus $cN = c'$. Now consider $(\mu AN)\boldsymbol{\delta}$ which reduces to

$$\mu AN\boldsymbol{\delta} = \mu A\mathbf{a} = 0.$$

Now consider again $(c - \mu A)\mathbf{a}$. It can be reduced to

$$(c - \mu A)\mathbf{a} = (c - \mu A)N\boldsymbol{\delta} = cN\boldsymbol{\delta} - \mu AN\boldsymbol{\delta} = cN\boldsymbol{\delta} = c'\boldsymbol{\delta}.$$

Thus since it was already shown that $(c - \mu A)\mathbf{a}$ can be written as a non-negative sum of the facet inequalities $c'\boldsymbol{\delta}$ can be as well. Since $c'\boldsymbol{\delta} \geq 0$ was an arbitrary inequality so the proof is done. \square

Example 2.4.7 Consider the cone EDC_{KN_n} . The facets of this cone are:

Left inequalities: $\delta(1, 2) \leq \delta(1, 3) \leq \delta(1, 4) \leq \delta(1, 5)$;

Right inequalities: $\delta(4, 5) \leq \delta(3, 5) \leq \delta(2, 5) \leq \delta(1, 5)$;

Triangle inequalities: $\delta(1, 3) \leq \delta(1, 2) + \delta(2, 3)$; $\delta(2, 4) \leq \delta(2, 3) + \delta(3, 4)$;

$\delta(3, 5) \leq \delta(3, 4) + \delta(4, 5)$; *Covering inequalities:* $\delta(2, 3) + \delta(1, 4) \leq \delta(1, 3) + \delta(2, 4)$;

$\delta(3, 4) + \delta(2, 5) \leq \delta(2, 4) + \delta(3, 5)$; $\delta(2, 4) + \delta(1, 5) \leq \delta(1, 4) + \delta(2, 5)$;

Now, using the characterization of the facets of EDC_{KN_n} , the equidistant cone for any

circular split system can be described. Every other split system, N , is a subset of KN_n and can be obtained by removing the splits not in N from KN_n .

Corollary 2.4.8 *Let N be a rooted circular split system on n leaves and root 0. Then EDC_N is a face of EDC_{KN_n} . Furthermore, every face of EDC_{KN_n} is of the form EDC_N for some subnetwork N .*

Proof: A way to construct N from KN_n is to start with KN_n and set all weights for splits not in N to 0. However, by Lemma 2.4.4, this means there is equality on all facets in KN_n that are not in N . Thus we can describe N by using all of the facet inequalities from KN_n but making the facets that can be rewritten using Lemma 2.4.4 as $2a_{A|B} \geq 0$ with $A|B \notin N$ into equalities. Then the cone EDC_N will be a face of EDC_{KN_n} as desired. \square

Example 2.4.9 *Consider the split system $N = \{04|123, 012|34, 014|23\} \cup \{i|[4] \setminus i : i \in [4]\}$. Then EDC_N is a face of EDC_{KN_4} and is generated by the following inequities and equalities:*

$$\text{Left inequalities: } \delta(1, 2) = \delta(1, 3) \leq \delta(1, 4);$$

$$\text{Right inequalities: } \delta(3, 4) \leq \delta(2, 4) = \delta(1, 4);$$

$$\text{Triangle inequalities: } \delta(1, 3) \leq \delta(1, 2) + \delta(2, 3); \quad \delta(2, 4) \leq \delta(2, 3) + \delta(3, 4);$$

$$\text{Covering inequalities: } \delta(2, 3) + \delta(1, 4) \leq \delta(1, 3) + \delta(2, 4);$$

Note that $\delta(1, 2) = \delta(1, 3)$ and $\delta(2, 4) = \delta(1, 4)$ since $034|12, 01|234 \notin N$, respectively.

2.5 Extreme Rays of EDC_{KN_n}

The goal of this section is to describe the extreme rays of EDC_{KN_n} . The extreme rays of EDC_{KN_n} can be described using a particular kind of set partition of $[n]$ which will be defined below. As an application, we also characterize which extreme rays lie on which facets of EDC_{KN_n} , which, given Corollary 2.4.8, will also characterize the facets and extreme rays for EDC_N for arbitrary networks.

Definition 2.5.1 *Let $\tau = t_1 | \dots | t_k$ be a set partition of $[n]$. We let r_τ be the vector whose coordi-*

nates, $r_\tau(i, j)$, are

$$r_\tau(i, j) = \begin{cases} 0 & \text{if there exists some } l \in [k] \text{ such that } i, j \in t_l \\ 1 & \text{otherwise.} \end{cases} \quad (2.5.1)$$

Definition 2.5.2 A fixed order set partition of $[n]$, $\tau = t_1 | \dots | t_k$, is a set partition such that for all $l \in [k]$, $t_l = [i, j] = \{i, i+1, \dots, j\}$ for some $i, j \in [n]$. Let

$$R_n = \{r_\tau \neq 0 : \tau \text{ is a fixed order set partition of } [n]\}.$$

Note that R_n has $2^{n-1} - 1$ elements since there are 2^{n-1} fixed order set partitions of n , and all but one of r_τ are nonzero.

Example 2.5.3 The fixed order set partitions of $[5]$ are the following:

$$\begin{aligned} &1|2345, 12|345, 123|45, 1234|5, 1|2|345, 1|23|45, 1|234|5, 12|3|45, 12|34|5, 123|4|5, \\ &1|2|3|45, 1|2|34|5, 1|23|4|5, 12|3|4|5, 1|2|3|4|5 \end{aligned}$$

We omit $\tau = 12345$, the partition with one block, since r_τ is just the zero vector. The set of rays, R_5 , corresponding to these set partitions are, respectively:

$$\begin{aligned} &\begin{bmatrix} 1 & 1 & 1 & 1 \\ & 0 & 0 & 0 \\ & & 0 & 0 \\ & & & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 0 & 0 \\ & & & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 & 1 \\ & 0 & 1 & 1 \\ & & 1 & 1 \\ & & & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & 1 \\ & 0 & 0 & 1 \\ & & 0 & 1 \\ & & & 1 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 0 & 0 \\ & & & 0 \end{bmatrix}, \\ &\begin{bmatrix} 1 & 1 & 1 & 1 \\ & 0 & 1 & 1 \\ & & 1 & 1 \\ & & & 0 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 & 1 \\ & 0 & 0 & 1 \\ & & 0 & 1 \\ & & & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 1 & 1 \\ & & & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 0 & 1 \\ & & & 1 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 & 1 \\ & 0 & 1 & 1 \\ & & 1 & 1 \\ & & & 1 \end{bmatrix}, \end{aligned}$$

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 1 & 1 \\ & & & 0 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 0 & 1 \\ & & & 1 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 & 1 \\ & 0 & 1 & 1 \\ & & 1 & 1 \\ & & & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 1 & 1 \\ & & & 1 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 & 1 \\ & 1 & 1 & 1 \\ & & 1 & 1 \\ & & & 1 \end{bmatrix}.$$

Lemma 2.5.4 *If τ is a fixed order set partition, then $r_\tau \in EDC_{KN_n}$.*

Proof: Let $\tau = t_1 | \dots | t_k$ be a fixed order set partition. We must show that r_τ satisfies all the inequalities from Theorem 2.4.1.

For the left inequalities, $\delta(1, i) \leq \delta(1, i+1)$ with $i \in [2, n-1]$ consider if there exists $l \in [k]$ such that $1, i \in t_l$. If this does exist then $\delta(1, i) = 0$ and so the facet will be $0 \leq \delta(1, i+1)$ which is always satisfied. If it does not exist then $i+1$ must also not be in the same block as 1, since τ is a fixed order set partition, and thus $\delta(1, i) = 1$ and $\delta(1, i+1) = 1$ so the facet will become $1 \leq 1$.

The right inequalities, $\delta(i, n) \leq \delta(i-1, n)$, follow by symmetry.

For the triangle inequalities, $\delta(i-1, i+1) \leq \delta(i-1, i) + \delta(i, i+1)$, consider whether there exists $l \in [k]$ such that $i-1, i+1 \in t_l$. If it exists then $\delta(i-1, i+1) = 0$. Additionally, it must be that, $i \in t_l$ as well so $\delta(i-1, i) = 0$ and $\delta(i, i+1) = 0$, and the facet becomes $0 \leq 0$. If it does not exist then $\delta(i-1, i+1) = 1$ and $\delta(i-1, i) + \delta(i, i+1)$ is 1 or 2, depending on if i is separated from both $i-1$ and $i+1$ or not. Thus the facet becomes $1 \leq 1$ or $1 \leq 2$.

For the covering inequalities $\delta(i, j) + \delta(i-1, j+1) \leq \delta(i, j+1) + \delta(i-1, j)$ with $i, j \in [2, n-1]$, because the r_τ only have 0 or 1 entries, each distance must be either 1 or 0. We can break up these distances into just the relevant entries of the vector r_τ . For this argument, it is useful to note that if some entries of r_τ equal 1, then certain others must equal 1 also. So for example, we can write $r_\tau(i-1, j+1) = \max(r_\tau(i-1, i), r_\tau(i, j), r_\tau(j, j+1))$. Using this observation, the inequality $\delta(i, j) + \delta(i-1, j+1) \leq \delta(i, j+1) + \delta(i-1, j)$, when applied to r_τ can be rewritten as

$$\begin{aligned} & r_\tau(i, j) + \max(r_\tau(i, j), r_\tau(i-1, i), r_\tau(j, j+1)) \\ & \leq \max(r_\tau(i, j), r_\tau(j, j+1)) + \max(r_\tau(i, j), r_\tau(i-1, i)). \end{aligned}$$

If $r_\tau(i, j) = 1$ then the inequality evaluates to $2 \leq 2$. If $r_\tau(i, j) = 0$ then the inequality evaluates to $\max(r_\tau(i-1, i), r_\tau(j, j+1)) \leq r_\tau(i-1, i) + r_\tau(j, j+1)$ which is always satisfied. This proves that the rays in $R_n \subseteq EDC_{KN_n}$. \square

Our next step will be to prove that $EDC_{KN_n} \subseteq \text{cone}(R_n)$. First, we will describe a general strategy to show that a given candidate set of extreme rays actually generates a cone.

Lemma 2.5.5 *Let $C = \{x : Ax \geq 0\}$ be a polyhedral cone where*

$$A = \begin{pmatrix} a_1 \\ \vdots \\ a_k \end{pmatrix} \quad (2.5.2)$$

and let $V = \{v_1, \dots, v_n\}$ be a set of vectors in \mathbb{R}^k . A subset $S \subseteq [k]$ is called *valid* for C if there exists some $y \in C$ such that $a_i y = 0$ for $i \in S$ and $a_j y > 0$ for all $j \notin S$.

If for each valid S there exists $v_l \in V$ such that $a_i v_l = 0$ for all $i \in S$, then $C \subseteq \text{cone}(V)$.

Proof: Suppose that we have an y in C , and let S be the associated set of indices. The proof is by induction on the size of $[k] \setminus S$. If $S = [k]$ there is nothing to show, since y must be the zero vector in that case. So suppose $S \neq [k]$ and let v_l be the element of V that is guaranteed to exist. Then we can compute $y' = y - \lambda v_l$ where $\lambda \geq 0$ is chosen as large as possible so that $y' \in C$. Since λ is as large as possible, y' must have an associated S' , which strictly contains S . By the inductive hypothesis, we can write $y' \in \text{cone}(V)$, so there are $\lambda_i \geq 0$ so that $y' = \sum_{i=1}^n \lambda_i v_i$. But then $y = \sum_{i=1}^n \lambda_i v_i + \lambda v_l$, so $y \in \text{cone}(V)$. \square

To use Lemma 2.5.5 on the cone EDC_{KN_n} , the next step will be to first characterize which S are valid for EDC_{KN_n} and then find an r_τ for each valid S . In order to do this, first we prove some useful non-facet inequalities on EDC_{KN_n} .

Lemma 2.5.6 *Let $\delta \in EDC_{KN_n}$ and let $\delta(i, j)$ be the (i, j) th coordinate of δ . Then for all $i < j < k \in [n]$, $\delta(i, j) \geq 0$, $\delta(i, j) \leq \delta(i, k)$, and $\delta(j, k) \leq \delta(i, k)$.*

Proof: The inequality $\delta(i, j) \geq 0$ is implied from δ being a dissimilarity map.

The inequality $\delta(1, j) \leq \delta(1, k)$ is implied by the facets $\delta(1, l) \leq \delta(1, l+1)$ for all $l \in [j, k-1]$. The inequality $\delta(j, n) \leq \delta(i, n)$ is implied by the facets $\delta(l, n) \leq \delta(l-1, n)$ for all $l \in [i+1, j]$.

The inequality $\delta(i, j) \leq \delta(i, k)$ is implied by the non-facet inequality $\delta(i-1, j) \leq \delta(i-1, k)$ and the covering facets:

$$\delta(i, l) + \delta(i-1, l+1) \leq \delta(i, l+1) + \delta(i-1, l)$$

for all $l \in [j, k-1]$. Incrementing l and canceling terms gives $\delta(i, j) \leq \delta(i, k)$ as desired. The non-facet inequality $\delta(i-1, j) \leq \delta(i-1, k)$ can be implied by facet inequalities using the same

argument as was just used for $\delta(i, j) \leq \delta(i, k)$. This will create a recursion of implications whose base case will be $\delta(1, j) \leq \delta(1, k)$ which was shown to be implied by facets above. Using a very similar argument the inequality $\delta(j, k) \leq \delta(i, k)$ is implied by the non-facet inequality $\delta(j, k+1) \leq \delta(i, k+1)$ and the facets:

$$\delta(l, k) + \delta(l-1, k+1) \leq \delta(l, k+1) + \delta(l-1, k)$$

for all $l \in [i+1, j]$. □

Now to characterize which S are valid for EDC_{KN_n} , we must see if equality on one facet or set of facets implies equality on any other facets. In order to clearly see which S are valid a new diagram will be introduced which will also provide additional insights into the structure of the space EDC_{KN_n} . Before defining the diagram itself we will need to define some components used in it.

Definition 2.5.7 *Let $\delta \in EDC_{KN_n}$ be represented as an $n \times n$ strictly upper triangular matrix. Let $\tilde{\delta}$ be an $n+1 \times n+1$ matrix obtained from δ where the bottom left n by n block is δ and first row and last column are all 1's. Because of the natural labeling of the rows and columns of δ , the rows of $\tilde{\delta}$ will be indexed from 0 to n and the columns of $\tilde{\delta}$ will be indexed from 1 to $n+1$.*

The following function is introduced because its possible values exactly correspond with valid S for EDC_{KN_n} .

Definition 2.5.8 *Let $k \in [0, n-2]$, $l \in [2, n]$ and $k < l$ such that if $k = 0$, $l \neq n$. The facet indicator function for δ , $f_{\delta}(k, l)$, is a Boolean function such that*

$$f_{\delta}(k, l) = \begin{cases} 1 & \text{if } \tilde{\delta}(k, l) + \tilde{\delta}(k+1, l+1) = \tilde{\delta}(k+1, l) + \tilde{\delta}(k, l+1) \\ 0 & \text{otherwise.} \end{cases}$$

Note that the facet indicator function is defined on each 2×2 sub-matrix of $\tilde{\delta}$ such that at least two of the entries are also in the upper triangular portion of δ . Additionally, notice that the facets of EDC_{KN_n} correspond directly with $f_{\delta}(k, l)$. The submatrix corresponding to $f_{\delta}(k, l)$ is the 2×2 matrix of the entries $\tilde{\delta}(k, l)$, $\tilde{\delta}(k+1, l)$, $\tilde{\delta}(k, l+1)$, $\tilde{\delta}(k+1, l+1)$. If two of the entries of the submatrix for $f_{\delta}(k, l)$ are 1's, the corresponding facet is a left or right inequality (depending on if the 1's are from the first row or last column respectively). If one of the entries of the submatrix is 0, the corresponding facet is a triangle inequality. If all four of the entries of the submatrix are also in δ then the corresponding facet is a covering inequality. Additionally,

if $f_{\tilde{\delta}}(k, l) = 1$ that corresponding facet will be satisfied with equality and if $f_{\tilde{\delta}}(k, l) = 0$ it will be a strict inequality. Thus, $f_{\tilde{\delta}}(k, l)$ gives a list of boolean outputs for which facets a particular point of EDC_{KN_n} lies on. Because of this property characterizing which f can arise as possible $f_{\tilde{\delta}}$ will characterize which S are valid for EDC_{KN_n} .

Definition 2.5.9 Let $k \in [0, n - 2]$, $l \in [2, n]$ and $k < l$ such that if $k = 0$, $l \neq n$. The facet corresponding to (k, l) is the facet $\tilde{\delta}(k, l) + \tilde{\delta}(k + 1, l + 1) \geq \tilde{\delta}(k + 1, l) + \tilde{\delta}(k, l + 1)$.

In order to help characterize for which f there exists a δ such that $f = f_{\delta}$, meaning that f is 1 exactly when a facet is satisfied by δ with equality, two more related Boolean functions for $\tilde{\delta}$ will be defined.

Definition 2.5.10 Let $i \in [0, n - 1]$, $j \in [n + 1]$. The vertical indicator function for $\tilde{\delta}$, $g_{\tilde{\delta}}$, is a Boolean function such that

$$g_{\tilde{\delta}}(i, j) = \begin{cases} 1 & \text{if } \tilde{\delta}(i, j) = \tilde{\delta}(i + 1, j) \\ 0 & \text{otherwise.} \end{cases}$$

Let $i' \in [0, n]$, $j' \in [n]$. The horizontal indicator function for $\tilde{\delta}$, $h_{\tilde{\delta}}$, is a Boolean function such that

$$h_{\tilde{\delta}}(i', j') = \begin{cases} 1 & \text{if } \tilde{\delta}(i', j') = \tilde{\delta}(i', j' + 1) \\ 0 & \text{otherwise.} \end{cases}$$

Note that the vertical and horizontal indicator functions are defined on each pair of vertically adjacent entries of $\tilde{\delta}$ and each pair of horizontally adjacent entries of $\tilde{\delta}$, respectively.

Definition 2.5.11 Let $\tilde{\delta} \in EDC_{KN_n}$. The grouping $(f_{\tilde{\delta}}, g_{\tilde{\delta}}, h_{\tilde{\delta}})$ is the X-diagram associated to $\tilde{\delta}$. More generally, a triple of boolean functions (f, g, h) where f has domain $\{(k, l) | k \in [0, n - 2], l \in [2, n] \text{ and } k < l\}$, g has domain $\{(i, j) | i \in [0, n - 1], j \in [n + 1]\}$, and h has domain $\{(i', j') | i' \in [0, n], j' \in [n]\}$ is called an X-diagram of size n . An X-diagram (f, g, h) is valid if there exists a $\tilde{\delta} \in EDC_{KN_n}$ such that $\tilde{\delta} \in EDC_{KN_n}$ with $(f, g, h) = (f_{\tilde{\delta}}, g_{\tilde{\delta}}, h_{\tilde{\delta}})$.

To draw an X-diagram (f, g, h) , take the following steps:

1. Draw a grid graph with vertices for each entry of $\tilde{\delta}$ (We draw the vertices as either boxes, or 0's and 1's if the entry of $\tilde{\delta}$ is always that number).
2. Remove the vertices of $\tilde{\delta}$ from the subdiagonal of δ and lower.

3. Take a subgraph of this grid graph, keeping edges for every two adjacent vertices only if the indicator function for corresponding entries of $\tilde{\delta}$ is 1. Here their indicator function is either horizontal or vertical indicator function depending on if they are horizontally or vertically adjacent.
4. Consider each 2×2 submatrix of $\tilde{\delta}$ where (k, l) in the domain of $f(k, l)$, i.e the entries $\tilde{\delta}(k, l), \tilde{\delta}(k+1, l), \tilde{\delta}(k, l+1)$, and $\tilde{\delta}(k+1, l+1)$. If $f(k, l) = 0$ place a dot in the center of the vertices of the grid graph corresponding to the entries of 2×2 submatrix. If $f(k, l) = 1$ place an X.

See Figure 2.11 for an example.

Example 2.5.12 Consider $\delta \in EDC_{KN_6}$, such that

$$\delta = \begin{bmatrix} 0 & 4 & 4 & 8 & 8 & 8 \\ & 0 & 1 & 6 & 6 & 7 \\ & & 0 & 3 & 5 & 6 \\ & & & 0 & 2 & 3 \\ & & & & 0 & 2 \\ & & & & & 0 \end{bmatrix}.$$

The X diagram for this δ can be seen in Figure 2.11. Additionally in this X-diagram we can see that

$$g_{\delta}(i, j) = \begin{cases} 1 & \text{if } i \in [0, 4] \text{ and } j = 7 \\ 0 & \text{otherwise;} \end{cases}$$

$$h_{\delta}(0, 2) = h_{\delta}(0, 3) = h_{\delta}(0, 4) = h_{\delta}(0, 5) = h_{\delta}(1, 2) = h_{\delta}(1, 4) = h_{\delta}(1, 5) = h_{\delta}(2, 4) = 1$$

and $h_{\delta}(i, j) = 0$ otherwise;

$$f_{\delta}(0, 2) = f_{\delta}(0, 4) = f_{\delta}(0, 5) = f_{\delta}(1, 4) = f_{\delta}(2, 5) = f_{\delta}(3, 5) = f_{\delta}(3, 4) = 1$$

and $f_{\delta}(i, j) = 0$ otherwise.

The following definition will be useful to characterize rays r_{τ} on an X-diagram and in the proof of Lemma 2.5.18.

Definition 2.5.13 An element of the form $\tilde{\delta}(i, i+1)$ with $1 \leq i \leq n$ is called a diagonal element.




Figure 2.11: An example X diagram for a point δ described in Example 2.5.12.

Notice that every $r_\tau \in R_n$ can be specified by which the diagonal elements are 0 and 1. Consider an $\tilde{\delta}$ as matrix for some X -diagram where only the diagonal elements of $\tilde{\delta}$ are specified. The other entries of the sub-matrix δ will be:

$$\delta(i, j) = \begin{cases} 0 & \text{if } \tilde{\delta}(k, k+1) = 0 \text{ for all } k \in [i, j] \\ 1 & \text{otherwise.} \end{cases}$$

This will give a δ such that $\delta = r_\tau$ for some $r_\tau \in R_n$ unless all of the diagonal entries are 0.

Example 2.5.14 Let $r_\tau = \delta \in EDC_{KN_6}$ with diagonal elements $\delta(1,2) = 0, \delta(2,3) = 1, \delta(3,4) = 0, \delta(4,5) = 0, \delta(5,6) = 0$. Then the other entries of δ are

$$\delta = \begin{bmatrix} 0 & 0 & 1 & 1 & 1 & 1 \\ & 0 & 1 & 1 & 1 & 1 \\ & & 0 & 0 & 0 & 0 \\ & & & 0 & 0 & 0 \\ & & & & 0 & 0 \\ & & & & & 0 \end{bmatrix}.$$

The extra information that comes from g and h in terms of the X -diagram will help to characterize the allowable f 's. Our next goal is to develop methods to detect which triples (f, g, h) are valid X -diagrams.

Lemma 2.5.15 Let (f, g, h) be a valid X -diagram.

1. If $f(k, l) = 1$ then $g(k, l) = g(k, l+1)$ and $h(k, l) = h(k+1, l)$.
2. If $g(k, l) = g(k, l+1) = 1$ then $f(k, l) = 1$. Similarly, if $h(k, l) = h(k+1, l) = 1$ then $f(k, l) = 1$.

3. Let $j > i$. Then if $g(i, j) = 1$, then $g(i, j + 1) = 1$. Similarly, if $h(i, j) = 1$, then $h(i - 1, j) = 1$.
4. If $g(i, j) = g(i, j + 1) = h(i, j) = 1$ then $h(i + 1, j) = 1$. Similarly, if $g(i, j + 1) = h(i, j) = h(i + 1, j) = 1$ then $g(i, j) = 1$.

Proof: For (1), if $f(k, l) = 1$ then,

$$\tilde{\delta}(k, l) + \tilde{\delta}(k + 1, l + 1) = \tilde{\delta}(k + 1, l) + \tilde{\delta}(k, l + 1).$$

If $g(k, l) = 1$ then $\tilde{\delta}(k, l) = \tilde{\delta}(k, l + 1)$ which implies that $\tilde{\delta}(k + 1, l + 1) = \tilde{\delta}(k + 1, l)$ which forces that $g(k, l + 1) = 1$. Similarly $g(k, l + 1) = 1$ forces that $g(k, l) = 1$. By the same logic $h(k, l) = 1$ forces $h(k + 1, l) = 1$ and vice versa. This proves (1).

For (2), if $g(k, l) = 1$ and $g(k, l + 1) = 1$, then $\tilde{\delta}(k, l) = \tilde{\delta}(k, l + 1)$ and $\tilde{\delta}(k + 1, l) = \tilde{\delta}(k + 1, l + 1)$. Thus $\tilde{\delta}(k, l) + \tilde{\delta}(k + 1, l + 1) = \tilde{\delta}(k + 1, l) + \tilde{\delta}(k, l + 1)$ so $f(k, l) = 1$. Similarly $h(k, l) = 1$ and $h(k + 1, l) = 1$ implies that $\tilde{\delta}(k, l) + \tilde{\delta}(k + 1, l + 1) = \tilde{\delta}(k + 1, l) + \tilde{\delta}(k, l + 1)$ so $f(k, l) = 1$, as desired.

For (3), Let $j > i$, and $g(i, j) = 1$. Thus $\tilde{\delta}(i, j) = \tilde{\delta}(i + 1, j)$. Consider the facet inequality

$$\tilde{\delta}(i, j) + \tilde{\delta}(i + 1, j + 1) \geq \tilde{\delta}(i + 1, j) + \tilde{\delta}(i, j + 1).$$

Since $\tilde{\delta}(i, j) = \tilde{\delta}(i + 1, j)$, this inequality reduces to

$$\tilde{\delta}(i + 1, j + 1) \geq \tilde{\delta}(i, j + 1).$$

By Lemma 2.5.6, $\tilde{\delta}(i + 1, j + 1) \leq \tilde{\delta}(i, j + 1)$, and thus $\tilde{\delta}(i + 1, j + 1) = \tilde{\delta}(i, j + 1)$, which implies that $g(i, j + 1) = 1$. A very similar argument will show that if $h(i, j) = 1$, then $h(i - 1, j) = 1$.

For (4), if $g(i, j) = g(i, j + 1) = h(i, j) = 1$ then $\tilde{\delta}(i, j) = \tilde{\delta}(i + 1, j)$, $\tilde{\delta}(i, j + 1) = \tilde{\delta}(i + 1, j + 1)$, and $\tilde{\delta}(i, j) = \tilde{\delta}(i, j + 1)$. Thus $\tilde{\delta}(i + 1, j) = \tilde{\delta}(i, j) = \tilde{\delta}(i, j + 1) = \tilde{\delta}(i + 1, j + 1)$ so $h(i + 1, j) = 1$. A similar argument will show that if $g(i, j + 1) = h(i, j) = h(i + 1, j) = 1$ then $g(i, j) = 1$. \square

Now, Lemma 2.5.15 can be used to classify which S are valid for EDC_{KN_n} .

Example 2.5.16 Consider the X -diagram on the left in Figure 2.12. Using the rules in Lemma 2.5.15, we can see that the underlying S for this X -diagram is not valid because $f(3, 6) = 1$ and $g(3, 7) = 1$ but $g(3, 6) = 0$ violating rule 1. In addition $g(1, 3) = 1$ and $g(1, 4) = 1$ but $f(1, 3) = 0$,




Figure 2.12: The X-diagram on the left is not valid as it violates multiple implications as described in Example 2.5.16. The X-diagram on the right is valid and is obtained by starting with the one on the left and following all rules in Lemma 2.5.15.

violating rule 2. Similarly, $h(3,3) = 1$ but $h(2,3) = 0$ violating rule 3. Lastly, $g(1,4) = h(1,4) = h(2,4) = 1$ but $g(1,5) = 0$ violating rule 4. The X-diagram on the right is the one obtained by following all of the rules in Lemma 2.5.15.

The following definition will be used to give the argument to find a r_τ for every valid S .

Definition 2.5.17 Let $\tilde{\delta}$ be the matrix for some $(\tilde{\delta}, f, g, h)$. The diagonal entries associated to (k, l) are the diagonal elements $\tilde{\delta}(k, k+1)$ and $\tilde{\delta}(l, l+1)$, if they exist. Note that all (k, l) have two diagonal elements associated to them, except for when $k=0$ or $l=n$, the ones corresponding to left and right facets, which only have one associated diagonal entry.

Lemma 2.5.18 Let $EDC_{KN_n} = \{x : Bx \geq 0\}$ with

$$B = \begin{bmatrix} b_1 \\ \vdots \\ b_k \end{bmatrix}. \quad (2.5.3)$$

For a given valid S for EDC_{KN_n} , there exists a $r_\tau \in R_n$ such that $b_i r_\tau = 0$ for all $i \in S$ and $b_i r_\tau > 0$ for all $i \notin S$.

Proof: First, notice that in a valid X-diagram, (f, g, h) , the only way that a $r_\tau \in R_n$ does not make $f(k, l) = 1$, is when the corresponding 2 by 2 submatrix has exactly the entries:

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

Thus if for a given valid S , if r_τ does not put $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ on any 2 by 2 sub matrix with upper left corner (k, l) and $f(k, l) = 1$ in the X -diagram corresponding to S , then $b_i r_\tau = 0$ for all $i \in S$. To show this is always possible, we will construct r_τ from the X -diagram corresponding to S by drawing a line from the left side of $\tilde{\delta}$ to the right side. Everything below the line will be 0 and everything on or above the line will be 1.

Starting at $(0, 0)$, increment the row index to its maximum before incrementing the column index. Let i_* be the index of the first column such that for some l_* , $f(l_*, i_*)$ has diagonal elements are that not 0. This must exist because if all diagonal elements are 0, the system S will just be equality everywhere, which can be handled by the ray $r_\tau = (1, \dots, 1)$. There are two cases, either $l_* = 0$ and $f(l_*, i_*)$ has one diagonal element, or $f(l_*, i_*)$ has two diagonal elements. Let (k', l') be the set of indices such that all $f(k', l')$ have $\tilde{\delta}(i_*, i_* + 1)$ as a diagonal element. Any (k', l') with any of its diagonal elements equal to 0, must have $f(k', l') = 1$ by Lemma 2.5.15 rules 2 and 3. However, if $f(k', l') = 1$ for all (k', l') in a row or column then $\tilde{\delta}(i_*, i_* + 1)$ would be 0 using the rule 1 from Lemma 2.5.15. Thus there must exist a (l_*, i_*) with $f(l_*, i_*) = 0$.

Draw the line starting from $\tilde{\delta}(l_*, 0)$ until reaching $\tilde{\delta}(l_*, i_* + 1)$ and then draw the line vertically until reaching $\tilde{\delta}(i_*, i_* + 1)$, which must be non-zero. Now consider the row of (i_*, l') who have $\tilde{\delta}(i_*, i_* + 1)$ as a diagonal element. One or both of two things are true:

1. there exists l_* with $f(i_*, l_*) = 0$ and both of its diagonal elements are not 0
2. $f(i_*, n) = 0$.

This is true because it cannot be that $f(i_*, l') = 1$ for all $l' \in [2, n]$ or by Lemma 2.5.15 rule 1, $\tilde{\delta}(i_*, i_* + 1)$ would 0, and the only possible (i_*, l') such that $f(i_*, l') \neq 1$ are those with non 0 diagonal blocks.

If the second case is true, draw the line horizontally to $\tilde{\delta}(i_*, n)$ and the line is finished. If only the first case is true then draw the line to $\tilde{\delta}(i_*, l_* + 1)$ and then draw it vertically down to the other non 0 diagonal block and continue the process. There are only finitely many entries so the line will eventually reach the right side. This line only puts $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ in its upper corners and the upper corners are only on (k, l) such that $f(k, l) = 0$ from the above argument. Thus $b_i r_\tau = 0$ for all $i \in S$ and $b_i r_\tau > 0$ for all $i \notin S$ as desired. \square

Example 2.5.19 Consider the valid X -diagram shown in Figure 2.13. The blue line is created by following the method described in Lemma 2.5.18. Thus the corresponding extreme ray for this S

P_{KN_n} as this will show that p_τ is a vertex of this polytope and thus r_τ will be an extreme ray in EDC_{KN_n} . Furthermore, we only need to check that the linear functional maximizes among all the points p_τ since $EDC_{KN_n} = \text{cone}(R_n)$, by previous results.

Consider the linear functional:

$$p_\tau^*(i, j) = \begin{cases} 1 & \text{if } p_\tau > 0 \\ -1 & \text{if } p_\tau = 0 \end{cases} \quad (2.5.5)$$

We will see this is maximized at p_τ . Note that evaluating p_τ^* on p_τ gives 1. But for any other $p_{\tau'}$ for $\tau' \neq \tau$, the highest it could sum to is 1 and it can not sum to 1 because that would mean $p_{\tau'}$ has exactly the same zero and non-zero entries as p_τ , contradicting that $\tau \neq \tau'$. \square

Theorem 2.5.22 *The extreme rays of EDC_{KN_n} are R_n .*

Proof: This follows from Lemmas 2.5.4, 2.5.20, and 2.5.21. \square

Using this same association we can describe which facets of EDC_{KN_n} each r_τ lies on.

Proposition 2.5.23 *Let KN_n be the complete rooted circular split network on n leaves. The extreme ray, r_τ , is contained in all facets of KN_n except the following:*

- *Left inequalities, $\delta(1, i) \leq \delta(1, i + 1)$ where i is separated from $i + 1$ and i is not separated from 1 in τ ,*
- *Right inequalities, $\delta(i, n) \leq \delta(i - 1, n)$ where i is separated from $i - 1$ and i is not separated from n in τ ,*
- *Triangle inequalities, $\delta(i - 1, i + 1) \leq \delta(i - 1, i) + \delta(i, i + 1)$ where i is separated from $i - 1$ and i is separated from $i + 1$ in τ ,*
- *Covering inequalities, $\delta(i, j) + \delta(i - 1, j + 1) \leq \delta(i, j + 1) + \delta(i - 1, j)$ where i is not separated from j but i is separated from $i - 1$ and j is separated from $j + 1$ and in τ .*

Proof: For left inequalities, $\delta(1, i) \leq \delta(1, i + 1)$, the only way there cannot be equality is if $\delta(1, i) = 0$ and $\delta(1, i + 1) = 1$. From the definition of r_τ , this happens precisely when i is separated from $i + 1$ and i is not separated from 1 in τ . The same argument applies by symmetry to the right inequalities.

For the triangle inequalities, $\delta(i - 1, i + 1) \leq \delta(i - 1, i) + \delta(i, i + 1)$, if $\delta(i - 1, i + 1) = 0$ that means that $i - 1$ and $i + 1$ are in the same block in τ . This implies that $\delta(i - 1, i) = 0$ and

$\delta(i, i+1) = 0$. Thus the only way to not have equality on this facet is if $\delta(i-1, i+1) = 1$ and $\delta(i-1, i) + \delta(i, i+1) = 2$. This happens precisely when i is separated from $i-1$ and i is separated from $i+1$ in τ .

For the covering inequalities, $\delta(i, j) + \delta(i-1, j+1) \leq \delta(i, j+1) + \delta(i-1, j)$, if $\delta(i, j) = 1$, that means that i and j are separated in τ . Then $\delta(i-1, j+1) = 1$ as well as $\delta(i, j+1) = 1$ and $\delta(i-1, j) = 1$. Thus the inequality becomes $2 \leq 2$. Thus in order for r_τ to not be on the covering inequality i and j must not be separated in τ . So assuming $\delta(i, j) = 0$, if $\delta(i-1, j+1) = 0$ as well, $i, j, i-1, j+1$ are all in the same block and so $\delta(i, j+1) = 0$ and $\delta(i-1, j) = 0$ as well, making the inequality $0 \leq 0$. Thus the only way to not get equality on this facet is if $\delta(i, j) = 0$, $\delta(i-1, j+1) = 1$, $\delta(i, j+1) = 1$, $\delta(i-1, j) = 1$, which is exactly when i is not separated from j but i is separated from $i-1$ and j is separated from $j+1$ and in τ . \square

Corollary 2.5.24 *Let N be a circular split network. The extreme rays of EDC_N are the subset of R_n that are contained in EDC_N .*

Proof: Since EDC_N is a face of EDC_{KN_n} by Corollary 2.4.8, its extreme rays must be a subset of R_n contained in the face. \square

Example 2.5.25 *Consider EDC_{KN_5} which has extreme rays associated to the following fixed order set partitions:*

1|2345, 12|345, 123|45, 1234|5, 1|2|345, 1|23|45, 1|234|5, 12|3|45, 12|34|5, 123|4|5,
1|2|3|45, 1|2|34|5, 1|23|4|5, 12|3|4|5, 1|2|3|4|5.

Using Proposition 2.5.23 we can characterize which facets of EDC_{KN_5} each of these rays lies on. For example:

1|2345 lies on all facets except $\delta(2, 5) \leq \delta(1, 5)$,
1|23|45 lies on all facets except $\delta(4, 5) \leq \delta(3, 5)$, $\delta(2, 3) + \delta(1, 4) \leq \delta(1, 3) + \delta(2, 4)$;
1|2|3|4|5 lies on all facets except $\delta(1, 3) \leq \delta(1, 2) + \delta(2, 3)$, $\delta(2, 4) \leq \delta(2, 3) + \delta(3, 4)$,
 $\delta(3, 5) \leq \delta(3, 4) + \delta(4, 5)$;

Example 2.5.26 *Consider the split system $N = \{01|2345, 12|0345, 0145|23, 0123|45\}$ together with the trivial splits, which are visualized in Figure 2.6. This split system is a subset of KN_5*

where the splits,

$$\{012|345, 05|1234, 045|123, 015|234, 0125|34\}$$

are not in N . So N can be obtained from KN_5 by setting the weights for those splits equal to 0. Specifically, this means that EDC_N will lie in the following equalities, corresponding to those splits

$$\begin{aligned}\delta(2, 5) - \delta(3, 5) &= 2a_{012|345} = 0 \\ \delta(1, 5) - \delta(1, 4) &= 2a_{05|1234} = 0; \delta(1, 4) - \delta(1, 3) = 2a_{045|123} = 0; \\ \delta(1, 4) + \delta(2, 5) - \delta(2, 4) - \delta(1, 5) &= 2a_{015|234} = 0; \\ \delta(2, 4) + \delta(3, 5) - \delta(3, 4) - \delta(2, 5) &= 2a_{0125|34} = 0;\end{aligned}$$

Using Corollary 2.5.24, EDC_N has the extreme rays associated to the following fixed order set partitions:

$$1|2345, 1|23|45, 12|3|45, 1|2|3|45, 1|23|4|5, 12|3|4|5, 1|2|3|4|5.$$

As mentioned in the introduction, EDC_N can be used to determine if a set of closely related species may have hybridization, using the facet description. While the extreme ray description of this cone is combinatorially nice, we did not see a clear biological interpretation of the extreme rays and this is a potential area of further study.

2.6 The Chan-Robbins-Yuen Polytope

The Chan-Robbins-Yuen Polytope (CRY_n) is a face of the Birkhoff polytope, and appears in other contexts as an example of a flow polytope [73]. It has generated interest in the combinatorics community because its normalized volume is a product of Catalan numbers, specifically:

$$\text{Vol}(CRY_n) = \prod_{i=1}^{n-2} \text{CAT}(i)$$

where $\text{CAT}(i) = \frac{1}{i+1} \binom{2i}{i}$. This polytope was first discussed in [32], and subsequently studied by many authors. In this section, we show a relation between the Chan-Robbins-Yuen polytope

and the cone EDC_{KN_n} .

Definition 2.6.1 *The Chan-Robbins-Yuen Polytope (CRY_n) is defined by the following set of equations and inequalities:*

$$CRY_n = \{x \in \mathbb{R}^{n \times n} : x_{ij} \geq 0 \text{ for all } i, j \in [n], \\ \sum_{i=1}^n x_{ij} = 1 \text{ for all } j \in [n], \sum_{j=1}^n x_{ij} = 1 \text{ for all } i \in [n], \text{ and} \\ x_{ij} = 0, \text{ if } i - j > 1\}.$$

Each element of CRY_n is an $n \times n$, doubly-stochastic matrix, which is zero below the first subdiagonal. The vertices of CRY_n are the permutation matrices that satisfy the condition that $x_{ij} = 0$, if $i - j > 1$. Note that there are exactly 2^{n-1} such permutation matrices in total.

To explain the relationship between the Chan-Robbins-Yuen polytope and equidistant network cone, we introduce a related polytope:

$$PEDC_n = EDC_{KN_n} \cap \{\delta \in \mathbb{R}^{n(n-1)/2} : \delta(1, n) \leq 1\}$$

Since all of the extreme rays of EDC_{KN_n} have a 1 in the $\delta(1, n)$ coordinate, this polytope has vertices consisting of all the vectors r_τ for all fixed ordered set partitions of $[n]$, including for the set partition with a single block $123 \cdots n$, which yields the origin. Note that $PEDC_n$ then also has 2^{n-1} vertices.

Theorem 2.6.2 *The polytopes $PEDC_n$ and CRY_n are affinely isomorphic. The isomorphism preserves the lattice spanned by the vertices in these polytopes, hence these polytopes have the same Ehrhart series and normalized volume.*

To prove Theorem 2.6.2 we will construct the affine isomorphism explicitly. Consider the map $\phi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n(n-1)/2}$, defined by

$$\phi_{k,l}(x) = 1 - \sum_{i=1}^k \sum_{j=l}^n x_{i,j}, \text{ for } 1 \leq k < l \leq n.$$

The inverse map $\psi : \mathbb{R}^{n(n-1)/2} \rightarrow \mathbb{R}^{n \times n}$ is provided by the following

$$\psi_{k,l}(\delta) = \begin{cases} 1 - \delta(1, n) & (k, l) = (1, n) \\ \delta(1, j+1) - \delta(1, j) & (k, l) = (1, j), 2 \leq j \leq n-1 \\ \delta(i, n) - \delta(i+1, n) & (k, l) = (i, n), 2 \leq i \leq n-1 \\ \delta(1, 2) & (k, l) = (1, 1) \\ \delta(n-1, n) & (k, l) = (n, n) \\ \delta(i-1, j) + \delta(i, j+1) - \delta(i, j) - \delta(i-1, j+1) & (k, l) = (i, j), 2 \leq i < j < n-1 \\ \delta(i-1, i) + \delta(i, i+1) - \delta(i-1, i+1) & (k, l) = (i, i), 2 \leq i \leq n-1 \\ 1 - \delta(i, i+1) & (k, l) = (i+1, i) \end{cases}$$

Proof: We want to show that the inequality system that defines the polytope $PEDC_n$, transforms into the inequality system for CRY_n when applying the map ϕ . That is, we consider the inequalities satisfied by $\delta \in PEDC_n$, and then make the substitution $\delta(k, l) = \phi_{k,l}(x)$, we should get the inequality system that describes CRY_n . We investigate each type of inequality that defines $PEDC_n$.

First, we consider the inequality $\delta(1, n) \leq 1$. Applying ϕ we get $\phi_{1,n}(x) = 1 - x_{1,n} \leq 1$, which is equivalent to $x_{1,n} \geq 0$.

Next we consider the inequalities $\delta(1, l) \leq \delta(1, l+1)$. Applying ϕ we get

$$1 - \sum_{i=l}^n x_{1,i} \leq 1 - \sum_{i=l+1}^n x_{1,i}$$

which is equivalent to $x_{1,l} \geq 0$. Similarly, the inequalities $\delta(k+1, n) \leq \delta(k, n)$ yields the inequality $x_{k,n} \geq 0$.

For the covering inequalities $\delta(i, j) + \delta(i-1, j+1) \leq \delta(i-1, j) + \delta(i, j+1)$, substitution and cancellation yields that this is equivalent to $x_{i,j} \geq 0$, using a simple inclusion and exclusion argument.

The most difficult to analyze are the triangle inequalities, $\delta(l-1, l+1) \leq \delta(l-1, l) + \delta(l, l+1)$. After making the substitution from ϕ , we get the inequality

$$\sum_{i=1}^l \sum_{j=l}^n x_{i,j} \leq 1 + x_{l,l}$$

We claim that this is equivalent to the inequality $x_{l,l} \geq 0$. To prove this, we use the double stochastic feature of the polytope CRY_n . In particular, we have the $\sum_{i=1}^{j+1} x_{i,j} = 1$ and $\sum_{j=i-1}^n x_{i,j} = 1$. In particular, the sum of the first l rows of x equals l , while the sum of the first $l-1$ columns is

$l - 1$. Since the first $l - 1$ columns are all zero below the subdiagonal, the difference between these two sums, on the one hand, is equal to 1, and on the other hand, is equal to $\sum_{i=1}^l \sum_{j=l}^n x_{i,j}$. Thus we get, $1 \leq 1 + x_{l,l}$, or $x_{l,l} \geq 0$.

At this point, we have shown that the inequalities that define $PEDC_n$ become the inequalities, $x_{i,j} \geq 0$ for all $1 \leq j < k \leq n$, and the inequalities $x_{i,i} \geq 0$ for $i = 2, \dots, n - 1$. This shows that $\phi(CRY_n) \subseteq PEDC_n$.

Next, we must verify that ψ is the inverse map of ϕ and that $\psi(PEDC_n) \subseteq CRY_n$. The fact that ψ is the inverse map can be checked directly by applying it coordinate by coordinate, using the different formulas for $\psi_{k,l}$. Most of this follows from the argument above. For instance, the proof that the covering inequality turns into the inequality $x_{i,j} \geq 0$, shows that $(\psi \circ \phi)_{i,j}(x) = x_{i,j}$ when $1 < i < j < n - 1$. A similar approach follows for the other coordinates. \square

It is worth noting how the maps ϕ and ψ transform vertices of the two polytopes. This is easiest to see on $\phi : CRY_n \rightarrow PEDC_n$. Each permutation matrix that is in CRY_n has the form of a block diagonal matrix, where each diagonal block is either a 1×1 block of a 1, or is a $k \times k$ block that has 1's on the subdiagonal and a 1 in the upper right corner. The map ϕ sends this permutation matrix to the 0/1 upper triangular array obtained by putting ones precisely in the part above and to the right of all the blocks of permutation matrices.

Here is an example of such a permutation matrix in CRY_9 and the vector r_τ that it maps to in $PEDC_9$ under ϕ :

$$\begin{pmatrix} 1 & & & & & & & & \\ & 0 & 0 & 1 & & & & & \\ & 1 & 0 & 0 & & & & & \\ & 0 & 1 & 0 & & & & & \\ & & & & 0 & 1 & & & \\ & & & & 1 & 0 & & & \\ & & & & & & 1 & & \\ & & & & & & & 0 & 1 \\ & & & & & & & & 1 & 0 \end{pmatrix} \mapsto \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ & & 0 & 1 & 1 & 1 & 1 & 1 \\ & & & 1 & 1 & 1 & 1 & 1 \\ & & & & 0 & 1 & 1 & 1 \\ & & & & & 1 & 1 & 1 \\ & & & & & & 1 & 1 \\ & & & & & & & 1 & 1 \\ & & & & & & & & & 0 \end{pmatrix}.$$

CHAPTER

3

IDENTIFIABILITY OF LARGE PHYLOGENETIC MIXTURES FOR MANY PHYLOGENETIC MODEL STRUCTURES

3.1 Introduction

Phylogenetic models represent the evolutionary relationship among a collection of taxa (short for taxonomic unit, which might be genes, species, or some other level of biological classification). Basic phylogenetic models are constructed by starting with a rooted tree where the taxa are vertices. Taxa closer to the root correspond to more ancestral units in the evolutionary history. Associated to each edge in the tree is a transition matrix. Each transition matrix encodes the substitution probabilities of the resulting characters along that edge. Specifying a substitution model amounts to placing restrictions on the transition matrices.

In the General Markov model there are no assumptions made on the entries of the transition matrices, whereas in the Jukes-Cantor model all off diagonal entries are equal so the transition matrices are highly constrained. Models like the Kimura 2-parameter, Kimura

3-parameter, and the strand symmetric model sit in between these two extremes.

When using any statistical model in practice, it is desirable if a distribution arising from the model must uniquely determine the parameters that produced it, i.e. the model should be *identifiable*. For basic phylogenetic models, these parameters are the tree parameter and the transition matrices. The tree parameter is a discrete parameter, in that there are only finitely many different trees on n leaves. On the other hand, the space of possible transition matrices forms a continuous parameter space. Questions of identifiability concern both the uniqueness of the discrete parameters and the continuous parameters as determined from probability distributions produced by the model. In this chapter, by identifiable we mean *generically identifiable*, that is, the parameters that produce a distribution are identifiable except possibly for a low dimensional subset of the parameter space. All the basic phylogenetic models are known to be generically identifiable by classic results in the literature [33].

This chapter concerns the identifiability of mixture models, which are more complex models which use the basic phylogenetic models as building blocks. In a mixture model, we do not assume that all characters evolve according to the same underlying phylogenetic tree with a fixed set of parameters. Rather there are (hidden) classes, and depending on which class a character belongs to, it evolves under the (basic) phylogenetic model according to that class. Classes might correspond to sites in DNA sequences that evolve at different rates (fast or slow). Across larger regions of DNA, classes could correspond to different genes which evolve according to different trees altogether. As mixture models have many more parameters than the basic phylogenetic models, questions of identifiability become increasingly more difficult to answer and their analysis requires more advanced tools.

Limited work has been done on the identifiability in group-based mixture models for specific numbers of classes and for completely general tree structures. For example, identifiability is known for 2 class mixtures for the Jukes-Cantor and Kimura 2-parameter models [4] for the Cavender-Farris-Neyman model and the Kimura 3-parameter model [55]. For 3 class mixtures identifiability is only known for the Jukes-Cantor model [70]. Work has been also been done on the identifiability of other types of mixture models such as Profile Mixture Models [95], which represent a useful submodel of general mixture models.

Probably the strongest result on identifiability of mixture models, the main theorem in [80] gives conditions under which mixture models under the General Markov model are identifiable. While this result does not allow for arbitrary mixtures with completely unrelated sets of trees, it does cover a large class of sets of trees, including the same tree mixtures which are the mostly commonly occurring. In most settings, these results allow for identifiability in a number of

mixture classes well above any that could be used in practice.

All the results mentioned so far consider generic identifiability, which, while a necessary consideration to use for mixture models, does not typically produce identifiability results for submodels. Hence, the main results of [80] only apply to the General Markov model, and not to any of the submodels: Jukes-Cantor (JC), Kimura 2 parameter model (K2P), Kimura 3 parameter model (K3P) and the Strand Symmetric model (SSM). These submodels are called *equivariant* because they have an underlying model symmetry to the assumptions they make on the transition matrices. The goal of this chapter is to broaden the tools developed from the main theorem in [80] to also deduce the identifiability of those important submodels.

The outline of this chapter is as follows. Before continuing onto the rest of the chapter, in the next section we give some important definitions necessary to state the main results of this chapter. Theorem 3.2.5 gives a general framework for proving identifiability results for phylogenetic mixture models and Corollary 3.2.6 establishes that the JC, K2P, K3P, and SSM all satisfy those conditions. In Section 3.3, background on phylogenetic trees and the Markov, group-based, and equivariant models is given. In Section 3.4, the discrete Fourier transform is introduced, which will be useful to determine the rank of certain matrices in our computations. In Section 3.5, Kruskal's theorem is introduced which gives bounds on the rank of tensors and will be crucial in the proof of the main theorem. The technical heart of the chapter is Sections 3.6, 3.7, and 3.8. Specifically, in Section 3.6, we prove our general identifiability result showing that if a Markov model μ satisfies four key properties, then the associated mixture models will be generically identifiable. Sections 3.7 and 3.8 then prove that the JC, K2P, K3P, and SSM models all satisfy these technical conditions to derive our main result.

3.2 Key Definitions and Statement of Main Results

The goal of this section is to give the necessary definitions so that we can state our main results. This section just covers definitions to help make sense of the results. Further definitions will appear in subsequent sections.

Let T be a single phylogenetic tree with n leaves. Let μ be a *Markov model*, (for example, JC, K2P, K3P, SSM, or GMM standing for Jukes-Cantor, Kimura 2-parameter, Kimura 3-parameter, Strand Symmetric model, or General Markov Model, respectively). Let S_T^μ be the parameter space of continuous parameters for the underlying model μ on tree T (which consists of transition matrices for each edge, and a root distribution). Let κ be the number of states of the random variables being considered ($\kappa = 4$ is most common in this chapter). Once we specify

these parts of a model, there exists a parametrization ψ_T^μ which gives the joint distribution of states on the leaves of T as a function of continuous parameters which specify the root distribution of the tree and all of the transition matrices on its edges. That is,

$$\psi_T^\mu : S_T^\mu \rightarrow \Delta^{\kappa^n - 1}$$

where $\Delta^{\kappa^n - 1} \subset [0, 1]^{\kappa^n - 1}$ is the probability simplex of non-negative real vectors summing to 1. Then the image of ψ_T^μ is the phylogenetic model M_T^μ .

Given a multiset of r phylogenetic trees T_1, \dots, T_r , a phylogenetic mixture model can be defined as the set of convex combinations of distributions from $M_{T_1}^\mu, \dots, M_{T_r}^\mu$.

Definition 3.2.1 *Let $\mathbf{T} = (T_1, \dots, T_r)$ be a r -tuple of trees each with n leaves. Let $S_{\mathbf{T}}^\mu = S_{T_1}^\mu \times \dots \times S_{T_r}^\mu \times \Delta^{r-1}$, each with underlying model μ . Then, the parametrization map of the mixture model is defined by*

$$\psi_{\mathbf{T}}^\mu : S_{\mathbf{T}}^\mu \rightarrow \Delta^{\kappa^n - 1}$$

where

$$\psi_{\mathbf{T}}^\mu(s_1, \dots, s_r, \pi) = \pi_1 \psi_{T_1}^\mu(s_1) + \dots + \pi_r \psi_{T_r}^\mu(s_r)$$

with $s_i \in S_{T_i}$, and $\pi \in \Delta^{r-1}$ is the vector of mixing parameters.

Since there is no order on the combination of trees in mixture model, identifiability on mixture model will have to be defined up to reordering of the trees.

Definition 3.2.2 *The tree parameters of an r -tree mixture model on $\mathbf{T} = (T_1, \dots, T_r)$ are generically identifiable if for generic choices of s_1, \dots, s_r and π with $s_i \in S_{T_i}^\mu$, and $\pi \in \Delta^{r-1}$,*

$$\psi_{\mathbf{T}}^\mu(s_1, \dots, s_r, \pi) = \psi_{\mathbf{T}}^\mu(s'_1, \dots, s'_r, \pi')$$

implies that there is a permutation $\sigma \in \mathfrak{S}_r$ (the Symmetric group on r letters) such that $\sigma \cdot \mathbf{T} = \mathbf{T}'$.

Definition 3.2.3 *The continuous/numerical parameters of an r -tree mixture model on $\mathbf{T} = (T_1, \dots, T_r)$ are generically identifiable if for generic choices of s_1, \dots, s_r and π with $s_i \in S_{T_i}^\mu$, and $\pi \in \Delta^{r-1}$,*

$$\psi_{\mathbf{T}}^\mu(s_1, \dots, s_r, \pi) = \psi_{\mathbf{T}}^\mu(s'_1, \dots, s'_r, \pi')$$

implies that there is a permutation on the r trees $\sigma \in S_r$ such that $\sigma \cdot \mathbf{T} = \mathbf{T}$, $s'_i = s_{\sigma(i)}$, and $\pi'_i = \pi_{\sigma(i)}$ for $i \in [r]$.

In the identifiability of the continuous parameters of a phylogenetic mixture model, we allow for label swapping, that permutes parameters between two classes of the mixture model that have the same underlying tree.

We must define the class of mixtures that we prove identifiability in. This amounts restricting to a subset of r -tuples of trees.

Definition 3.2.4 *Let $\mathcal{T}(r, n, k)$ be the set of r -tuples of n -leaf binary trees (T_1, \dots, T_r) such that there exists a tripartition of the leaves $A|B|C$ and for each $i \in [r]$ a vertex v_i in tree T_i such that the induced tripartition of the leaves in T_i induced by v_i is $A|B|C$. Furthermore, we assume that $\#A \geq \#B \geq \#C$ and $\#B \geq k$.*

We now state our main result, which depends on some technical definitions which will appear in later sections. However, this gives us the flavor of the results that can be achieved with these methods, and the fact that certain specific properties need to be proved for a model μ to deduce identifiability for a mixture model. These properties are *nontriviality, the rank property, the extended rank property, and the No Shuffling Property*.

Theorem 3.2.5 *Suppose that a phylogenetic model μ on $\kappa > 2$ states satisfies the following properties:*

1. μ is non-trivial,
2. μ has the rank property $RP(r, k)$
3. μ has the extended rank property $ERP(r, k)$
4. μ has the No Shuffling Property
5. $r \leq \kappa^{k-1}$, and
6. $n \geq 2k + 1$.

Then both tree parameters and the numerical parameters of $M_{\mathcal{T}}^{\mu}$ are generically identifiable in the class of trees $\mathcal{T}(r, n, k)$.

Corollary 3.2.6 *Suppose that μ is the Jukes-Cantor model, Kimura-2 or 3-parameter model, Strand symmetric model, or general Markov model. Suppose that $n \geq 2k + 1$ and $r \leq 2^k - k$. Then both tree parameters and the numerical parameters of $M_{\mathcal{T}}^{\mu}$ are generically identifiable in the class of trees $\mathcal{T}(r, n, k)$.*

Part of our goal in these proofs is to try to give a uniform argument to deduce identifiability for these models. A key component of the proof of Corollary 3.2.6 is to prove that μ has the rank property $RP(r, k)$ and extended rank property $ERP(r, k)$ for the specific values of r and k . To show this, we prove both properties for the Jukes-Cantor model for the particular values of r and k in the Theorem. Since the Jukes-Cantor transition matrices appear in all the other models, this implies the rank property for the larger models for those values of r and k . However, it is probably possible to derive stronger results for the other models by more carefully analyzing the rank property and extended rank property in other models. We discuss this in more detail in Section 3.7.

3.3 Markov models on trees

In this section, we review background on the combinatorial structure of trees, including splits and tripartitions associated to tree edges and vertices. We introduce the Markov models on a tree that are the main object of study in this chapter.

A tree $T = (V, E)$ is a graph that is connected and has no cycles. The degree of a vertex is the number of edges incident to the vertex. A vertex of degree 1 is called a leaf. Phylogenetic trees represent the evolution of taxa with the vertices each representing a different taxa. The closer a vertex is to the root vertex the older that taxa is evolutionarily. The leaves of the tree will represent taxa that are alive today. This gives the leaves special importance since the DNA of those species is observable. The following definitions of phylogenetic trees are standard (see [86]).

Definition 3.3.1 *Let T be a tree and X be a set of labels. Let ϕ be a map from X to the vertices of T . The pair (T, ϕ) is called a phylogenetic X -tree if the image of ϕ is exactly the set of leaves of T and ϕ is injective. A binary phylogenetic X -tree is a phylogenetic X -tree where all non-leaf vertices have degree 3. A rooted tree is a tree with a distinguished vertex called the root.*

Removing any edge of a phylogenetic tree separates the leaf labels into two parts, based on which connected component the leaf is in. This leads to following definition.

Definition 3.3.2 *A split $A|B$ of X is a partition of X into two nonempty disjoint sets. A split is considered valid for an X -tree T if it is obtained by removing an edge from T and taking A and B to be the two sets of leaf labels for the two component graphs.*

Similarly, a tripartition $A|B|C$ is obtained from a binary tree by removing an internal vertex and A and B and C being the three sets of leaf labels for the three component graphs.

Let T be a phylogenetic tree. We associate a random variable X_v to each vertex in T . Each of the random variables X_v has κ states for a fixed value κ . In general terms, the state space of X_v is $[\kappa] = \{1, 2, \dots, \kappa\}$. If $\kappa = 4$ this alphabet can be thought of as DNA bases $\{A, C, G, T\}$. There are other values for κ that are relevant biologically, such as $\kappa = 20$ for amino acids, $\kappa = 61$ for DNA codons, and $\kappa = 2$ for DNA where purines and pyrimidines are grouped together. This chapter will be primarily focusing on $\kappa = 4$.

Before defining Markov models on a tree, we will motivate their definition. Consider for each vertex a sequence using the characters of the alphabet. If $\kappa = 4$ that sequence on the vertex would be the DNA sequence for a gene and each vertex would represent a DNA sequence for a specific taxa. For simplicity, only point substitutions will be considered, i.e. mutations that change one DNA letter to another. Thus each sequence will be of the same length for all vertices in the graph, but possibly have different characters in each location. The Markov model on the tree is a model for how these sequences change over time. If we assume that each site in the sequence evolves independently under the same process, we can describe the model just by assuming we have sequences of length 1, or a single character. Thus, for each vertex we will associate a random variable with $\kappa = 4$ states where the internal non-leaf vertices will be hidden random variables (since they are unobserved).

To each edge in the graph $i \rightarrow j$ is associated a Markov transition matrix M^{ij} , which is a $\kappa \times \kappa$ matrix whose (x_i, x_j) entry is the conditional probability $P(X_j = x_j | X_i = x_i)$. That is:


$$M^{ij}(x_i, x_j) = P(X_j = x_j | X_i = x_i).$$

To be a Markov transition matrix M^{ij} , we require that all entries are nonnegative, and that the row sums are all equal to one.

Once a Markov transition matrix is specified for each edge, and a distribution of states $\pi \in \Delta_{r-1}$ is specified, we can write down the joint distribution of all states in the model. Let $p(x, y)$ denote the joint probability distribution of all random variables, where x corresponds to leaves and y corresponds to interior vertices. Let ρ denote the root vertex. Then

$$p(x, y) = \pi(y_\rho) \prod_{i \rightarrow j \in E(T)} M^{ij}(x_i, x_j).$$

Since we are typically only interested in the distribution of states at the leaves (because we do not have access to data from the taxa at interior nodes), then we have the probability



$$M^{b2} = \begin{pmatrix} 0.25 & 0.3 & 0.15 & 0.3 \\ 0.3 & 0.25 & 0.4 & 0.05 \\ 0.15 & 0.4 & 0.25 & 0.2 \\ 0.3 & 0.05 & 0.2 & 0.45 \end{pmatrix}$$

Figure 3.1: A phylogenetic model with 3 leaves and two internal vertices a and b and a transition matrix

distribution of interest is

$$p(x) = \sum_{y \in [\kappa]^{\text{Int}(T)}} p(x, y)$$

where $\text{Int}(T)$ denotes the interior vertices of T .

Example 3.3.3 Consider the rooted tree in Figure 3.1. Let the transition matrix written here be associated to the edge from $b \rightarrow 2$. Label the columns and rows of the matrix with A, C, G, T. Then the entries in this transition matrix represent probabilities of DNA substitution when going from the species represented in the vertex b to the one represented with the vertex 2. Explicitly the entry $M^{b2}(x_b, x_2)$ is the probability that the variable at X_b mutates from the state x_b to the state x_2 at vertex 2. So for example there is a 40% chance $X_b = C$ changes to $X_2 = G$.

Usually when defining a phylogenetic model, we put restrictions on the structure of transition matrices that are used in the model. We call such restriction on the transition matrices a *Markov model* μ . Examples of Markov models are the Jukes-Cantor model, Kimura 2 and 3 parameter models, the strand symmetric model, and the general Markov model. The resulting structures on the transition matrices are as follows. The matrix M^{K3P} denotes a generic matrix

for the Kimura 3-parameter model:

$$M^{K3P} = \begin{pmatrix} \alpha & \beta & \gamma & \delta \\ \beta & \alpha & \delta & \gamma \\ \gamma & \delta & \alpha & \beta \\ \delta & \gamma & \beta & \alpha \end{pmatrix}.$$

Note that this matrix only has 3 free parameters, because we require that $\alpha + \beta + \gamma + \delta = 1$. The Kimura 2-parameter model arises as a submodel obtained by requiring that $\beta = \delta$. And the Jukes-Cantor model is a further submodel where $\beta = \gamma = \delta$.

The Strand Symmetric Model is a larger model where the generic transition matrices have the following form:

$$M^{SSM} = \begin{pmatrix} \alpha & \beta & \gamma & \delta \\ \epsilon & \zeta & \eta & \theta \\ \theta & \eta & \zeta & \epsilon \\ \delta & \gamma & \beta & \alpha \end{pmatrix}.$$

This model has 6 free parameters associated to each transition matrix (again, because the row sums must equal 1). Note that this model contains the JC, K2P, and K3P models as submodels. Finally, the general Markov model makes no restrictions at all on the transition matrices, except that the matrices are transition matrices.

All of these models are examples of what are called *equivariant models*, a useful class of phylogenetic models with key symmetry properties that makes them easier to analyze [27]. A key family are the group-based models, of which the JC, K2P, and K3P are special cases.

Definition 3.3.4 *A phylogenetic model μ is group-based if for all transition matrices, M , associated to its edges there exists a function $f : G \rightarrow \mathbb{R}$ such that $M(g, h) = f(g - h)$ where G is an abelian group and $M(g, h)$ denotes the probability of going from state h given that the variable is in state g in the previous vertex.*

In the next section, we explain how the discrete Fourier transform can simplify the presentation of the group-based models.

3.4 The Discrete Fourier Transform

The discrete Fourier transform is a useful tool that, for certain phylogenetic models, can be applied to simplify the parametrization of the model, to get a parametrization in a new

coordinate system that involves polynomials with fewer terms. For the case of group-based models (including the Jukes-Cantor, Kimura 2 and 3 parameter models), it simplifies the phylogenetic model on a tree into a toric variety. For the strand symmetric model, the discrete Fourier transform does not produce a toric variety, but another variety that is considerably simpler than the standard parametrization.

Definition 3.4.1 *Let G be a finite abelian group and let $\hat{G} = \text{Hom}(G, \mathbb{C}^\times)$ be its dual group of one dimensional representations. Let f be a function $f : G \rightarrow \mathbb{C}$. The discrete Fourier transform of f is the function $\hat{f} : \hat{G} \rightarrow \mathbb{C}$ defined by*

$$\hat{f}(\chi) = \sum_{g \in G} f(g) \chi(g).$$

Definition 3.4.1 provides the Fourier transform for an arbitrary finite abelian group. We will see how this is applied and can simplify the parametrization for a group-based model and other equivariant models.

First we describe the transformation for the group-based models, and then we show how the parametrization looks of the group-based models in the Fourier coordinates. Let T is a rooted tree with m leaves, in a group based model, for each edge we have a transition matrix M^e that satisfies $M^e(g, h) = f^e(g - h)$ for some particular function $f^e : G \rightarrow \mathbb{R}$. Additionally there is a root distribution $\pi : G \rightarrow \mathbb{R}$. For the group-based models we consider here, we will assume that the root distribution π is always the uniform distribution.

Let $H(e)$ be the head vertex of edge e and $T(e)$ be its tail vertex. Let $\lambda(e)$ be the set of leaves below e and let Int be the set of interior vertices of T . Then the joint probability of observing (g_1, \dots, g_m) at the leaves is:

$$\begin{aligned} p(g_1, \dots, g_m) &= \sum_{g \in G^{Int}} \pi(g_r) \prod_{e \in T} M^e(g_{T(e)}, g_{H(e)}) \\ &= \sum_{g \in G^{Int}} \pi(g_r) \prod_{e \in T} f^e(g_{T(e)} - g_{H(e)}). \end{aligned}$$

See Chapter 15 of [91] for more details.

Now we compute the Fourier transform $p : G^m \rightarrow \mathbb{R}$ over the group G^m . Note that the dual group of a product group is the product of the dual groups. So, as a linear change of coordinates, this transformation is

$$\hat{p}(\chi_1, \dots, \chi_m) = \sum_{g_1, \dots, g_m \in G} \chi_1(g_1) \cdots \chi_m(g_m) p(g_1, \dots, g_m).$$

The Fourier transform of p can be described with the following product formula:

$$\hat{p}(\chi_1, \dots, \chi_m) = \hat{\pi} \left(\prod_{i=1}^m \chi_i \right) \prod_{e \in T} \hat{f}^e \left(\prod_{i \in \lambda(e)} \chi_i \right)$$

This was first discussed in [45, 53].

Example 3.4.2 Consider the rooted 3 leaf tree with underlying Jukes-Cantor model. Let $f^i(A) = b_i$ and $f^i(C) = f^i(G) = f^i(T) = a_i$. We make the identification that $A = (0, 0)$, $C = (0, 1)$, $G = (1, 0)$, and $T = (1, 1)$. The four characters of $\mathbb{Z}_2 \times \mathbb{Z}_2$ are the rows of the following character table:

	A	C	G	T
χ_A	1	1	1	1
χ_C	1	-1	1	-1
χ_G	1	1	-1	-1
χ_T	1	-1	-1	1

So the Fourier transform for a particular edge is

$$\hat{f}^i(\chi_A) = 3a_i + b_i \quad \text{and} \quad \hat{f}^i(\chi_C) = \hat{f}^i(\chi_G) = \hat{f}^i(\chi_T) = b_i - a_i.$$

Note that the Fourier transform of the uniform distribution has

$$\hat{\pi}(\chi_A) = 1, \quad \text{and} \quad \hat{\pi}^i(\chi_C) = \hat{\pi}^i(\chi_G) = \hat{\pi}^i(\chi_T) = 0$$

Let $x_i = 3a_i + b_i$ and $y_i = b_i - a_i$.

We use the shorthand $A = \chi_A, \dots$, and we write $q_{AAA} = \hat{p}(\chi_A, \chi_A, \chi_A), \dots$. Then we have

$$\begin{aligned} q_{AAA} &= x_1 x_2 x_3, & q_{ACC} &= q_{AGG} = q_{ATT} = x_1 y_2 y_3, \\ q_{CAC} &= q_{GAG} = q_{TAT} = y_1 x_2 y_3, & q_{CCA} &= q_{GGA} = q_{TTA} = y_1 y_2 x_3, \\ q_{CGT} &= q_{CTG} = q_{GCT} = q_{GTC} = q_{TCG} = q_{TGC} = y_1 y_2 y_3, \\ q_{\chi_1 \chi_2 \chi_3} &= 0 \quad \text{otherwise.} \end{aligned}$$

For a group-based model on the group $G = \mathbb{Z}_2 \times \mathbb{Z}_2$, we can express the parametrization in

the following form. We have Fourier parameters $a_g^{A|B}$ for each split $A|B$ in the tree, and each group element g . The parametrization in Fourier coordinates looks like

$$q_{g_1 \dots g_n} = \begin{cases} \prod_{A|B \in \Sigma(T)} a_{\sum_{a \in A} g_a}^{A|B} & \text{if } \sum_{i=1}^n g_i = (0, 0) \\ 0 & \text{otherwise.} \end{cases}$$

Then, depending on the specific model, there might be equalities that hold between some the Fourier parameters. In the Jukes-Cantor model, for each split $A|B \in \Sigma(T)$, we have that

$$a_{(0,1)}^{A|B} = a_{(1,0)}^{A|B} = a_{(1,1)}^{A|B}.$$

3.5 Tensors and Flattenings

In this section, we discuss how to think about the probability distribution associated to a phylogenetic mixture model as a multi-way tensor. The tensor perspective is useful, because we can also consider flattenings of our tensor to lower order tensors, and matrices, in order to use Kruskal's theorem and matrix ranks of flattenings to prove identifiability results. The approach us using Kruskal's theorem to prove identifiability results first appeared in [6].

For a tree with n leaves, and random variables with κ states, the joint probability distribution in our model is naturally considered as a $\kappa \times \dots \times \kappa$ n -tensor, inside of $\bigotimes_{i=1}^n \mathbb{R}^\kappa$. From this n -way tensor, we can also consider the tensor naturally as a lower order tensor, by grouping indices according to a set partition of the coordinates. Such a reorganization of the tensor is called a *flattening*.

Definition 3.5.1 *Let P be an order n tensor in $\bigotimes_{i=1}^n \mathbb{R}^{\kappa_i}$. Let $A_1 | \dots | A_k$ be a partition of $[n]$. The flattening of P according to $A_1 | \dots | A_k$, is an order- k tensor obtained by grouping indices according to the partition $A_1 | \dots | A_k$, and is denoted $\text{Flat}_{A_1 | \dots | A_k}(P)$. Specifically, $\text{Flat}_{A_1 | \dots | A_k}(P)$ is the image of P under the natural map*

$$\bigotimes_{i=1}^n \mathbb{R}^{\kappa_i} \rightarrow \bigotimes_{i=1}^k \left(\bigotimes_{a \in A_i} \mathbb{R}^{\kappa_a} \right).$$

In particular, $\text{Flat}_{A_1 | \dots | A_k}(P)$ is a $\prod_{a_1 \in A_1} \kappa_{a_1} \times \dots \times \prod_{a_k \in A_k} \kappa_{a_k}$ tensor.

Note that in the specific case when $\kappa_i = \kappa$ for all i , and if $k = 2$, then we have our partition $A|B$ and $\text{Flat}_{A|B}(P)$ is a matrix of size $\kappa^{\#A} \times \kappa^{\#B}$.

Example 3.5.2 Let T be a $3 \times 3 \times 3$ tensor with generic entries t_{ijk} , and consider the partition $A|B = 1|23$. Then

$$\text{Flat}_{1|23} = \begin{pmatrix} t_{111} & t_{112} & t_{113} & t_{121} & t_{122} & t_{123} & t_{131} & t_{132} & t_{133} \\ t_{211} & t_{212} & t_{213} & t_{221} & t_{222} & t_{223} & t_{231} & t_{232} & t_{233} \\ t_{311} & t_{312} & t_{313} & t_{321} & t_{322} & t_{323} & t_{331} & t_{332} & t_{333} \end{pmatrix}$$

One key advantage to considering probability distributions as tensors is that we can use tensor rank as a tool to answer questions of identifiability in phylogenetic models. To explain the key tool in this area, Kruskal's theorem, we need to review some concepts related to tensor rank.

For $j = 1, 2, 3$, let $\mathbf{m}^j = (m^j(1), \dots, m^j(\kappa_j)) \in \mathbb{R}^{\kappa_j}$. Let $\mathbf{m}^1 \otimes \mathbf{m}^2 \otimes \mathbf{m}^3$ denote the $\kappa_1 \times \kappa_2 \times \kappa_3$ tensor whose (u, v, w) entry is $m^1(u)m^2(v)m^3(w)$. The tensor $\mathbf{m}^1 \otimes \mathbf{m}^2 \otimes \mathbf{m}^3$ is called a *rank 1 tensor*. A tensor is said to have rank r if it can be written as the sum of r rank 1 tensors, and cannot be written as the sum of $r - 1$ rank 1 tensors. Kruskal's theorem concerns the uniqueness of the representation of a rank r tensor as the sum of r rank 1 tensors. To express this, we first introduce the triple product notation of matrices.

Definition 3.5.3 Let M_j be a $r \times \kappa_j$ matrix with i th row $\mathbf{m}_i^j = (m_i^j(1), \dots, m_i^j(\kappa_j))$. Then let $[M_1, M_2, M_3]$ be defined by

$$[M_1, M_2, M_3] = \sum_{i=1}^r \mathbf{m}_i^1 \otimes \mathbf{m}_i^2 \otimes \mathbf{m}_i^3.$$

The triple product $[M_1, M_2, M_3]$ is a $\kappa_1 \times \kappa_2 \times \kappa_3$ tensor and its (u, v, w) entry is

$$[M_1, M_2, M_3]_{(u,v,w)} = \sum_{i=1}^r m_i^1(u) \otimes m_i^2(v) \otimes m_i^3(w).$$

Note that there is a natural nonuniqueness that is always present. If Π is an $r \times r$ permutation matrix, and $r \times r$ diagonal matrices D_1, D_2, D_3 such that $D_1 D_2 D_3 = Id_r$, then

$$[M_1, M_2, M_3] = [\Pi D_1 M_1, \Pi D_2 M_2, \Pi D_3 M_3].$$

If

$$[M_1, M_2, M_3] = [N_1, N_2, N_3]$$

and there exist Π and D_1, D_2, D_3 such that

$$[N_1, N_2, N_3] = [\Pi D_1 M_1, \Pi D_2 M_2, \Pi D_3 M_3]$$

when we say that N_1, N_2, N_3 are related to M_1, M_2, M_3 by scaling and simultaneous permutation of the rows. If this is the only way we can have $[M_1, M_2, M_3] = [N_1, N_2, N_3]$ then we say that $[M_1, M_2, M_3]$ uniquely determines M_1, M_2, M_3 up to simultaneous permutation and scaling of the rows. Note that this means that the rank one tensors that appear in the decomposition are themselves unique.

The existence of a unique of the decomposition in the Kruskal theorem depends on the notion of the Kruskal rank of a matrix.

Definition 3.5.4 *Let M be a $r \times \kappa$ matrix. The Kruskal rank of a matrix M , written as $\text{rank}_K(M)$ is the largest k such that every subset of k rows of M is linearly independent.*

Note that $\text{rank}_K(M) \leq \text{rank}(M)$ but if M is of full row rank (that is, $\text{rank}(M) = r$, it will also be full Kruskal rank. Now we are ready to state Kruskal's Theorem.

Theorem 3.5.5 *(Kruskal's Theorem) [66] For $j = 1, 2, 3$, let M_j be an $r \times \kappa_j$ matrix and let $I_j = \text{rank}_K M_j$. If*

$$I_1 + I_2 + I_3 \geq 2r + 2$$

then $[M_1, M_2, M_3]$ uniquely determines M_1, M_2 , and M_3 up to simultaneous permutation and scaling of the rows.

The motivation for considering tensor rank comes from the fact that the probability distribution from a single tree can be represented as a triple product, when considered around a given vertex. Specifically, let T be a tree, let v be a vertex of degree 3 in T . We can assume that v is the root of the tree. The vertex v introduces a tripartition $A|B|C$ of the leaves. Let $\text{diag}(\pi)$ be a diagonal matrix with the root distribution. Let M_A be the $\kappa \times \kappa^{\#A}$ matrix which is the flattenings of the conditional distribution of X_A given X_v . Similarly, define M_B and \tilde{M}_C as conditional distributions of X_B given X_v and X_C given X_v , respectively. Then, if P is the joint distribution of states at the leaves of our tree T , we have

$$\text{Flat}_{A|B|C}(P) = [M_A, M_B, \text{diag}(\pi)\tilde{M}_C] = [M_A, M_B, M_C].$$

where we let $M_C = \text{diag}(\pi)\tilde{M}_C$. This representation was first developed in [4], and used to prove identifiability results in a number of hidden variable models including phylogenetic models.

Note that we could have absorbed the diagonal matrix $\text{diag}(\pi)$ into any of M_A, M_B or \tilde{M}_C . Since we are always working with generic distributions, this would not change the Kruskal rank or tensor rank. Hence we can consider M_A, M_B , and M_C as either conditional distributions given X_v , or just as joint distributions.

3.6 A General Theorem on Identifiability of Phylogenetic Mixtures

Our goal in this section is to prove our main general result about the identifiability of phylogenetic mixture models, Theorem 3.2.5. To do this, we must first define the rank property, extended rank property, the standard Markov model property, and the No Shuffling Property, and show how they are related to identifiability of mixture models.

Definition 3.6.1 *The Markov model μ is said to have the Rank property $RP(r, n)$ if the following holds: For any trees T_1, \dots, T_r on n leaves, and generic probability distributions $P_i \in M_{T_i}^\mu$, the matrix*

$$\left(\text{Flat}_{[n-1]n}(P_1) \quad \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \text{Flat}_{[n-1]n}(P_r) \right)$$

has rank $r\kappa$.

Definition 3.6.2 *The Markov model μ is said to have the extended rank property $ERP(r, n)$ if the following holds: For any tree T_1 on $n+1$ leaves that does not contain the split $[n-1]\{n, n+1\}$, and any trees T_2, \dots, T_r on n leaves, and generic probability distributions $P_i \in M_{T_i}^\mu$, the matrix*

$$\left(\text{Flat}_{[n-1]\{n, n+1\}}(P_1) \quad \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \text{Flat}_{[n-1]n}(P_r) \right)$$

has rank $(r-1)\kappa + \kappa^2$.

A useful fact about both the rank property and extended rank property is that they satisfy persistence properties in their values.

Proposition 3.6.3 *Let μ be a Markov model such that the rank property $RP(r, n)$ holds. Then the rank property $RP(s, m)$ holds for all $s \leq r$ and $m \geq n$. Similarly, if the extended rank property $ERP(r, n)$ holds, then the extended rank property $ERP(s, m)$ holds for all $s \leq r$ and $m \geq n$.*

Proof: Suppose that $RP(r, n)$ holds. Let

$$Q = \left(\text{Flat}_{[n-1]n}(P_1) \quad \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \text{Flat}_{[n-1]n}(P_r) \right).$$

Then the associated matrix Q' which is used to check $RP(s, n)$ for $s \leq r$ is obtained from Q by taking a submatrix of the columns of Q . Since Q had full column rank Q' must have full column rank as well.

Now suppose that $m \geq n$. Construct the matrix Q' to check $RP(r, m)$. Each of the distributions obtained to form Q is gotten by marginalizing the corresponding distribution from Q' . We can assume we are marginalizing some collection of random variables besides the last variable, m . On the level of matrices, Q is obtained from Q' by adding together rows of Q' . Since we assumed that Q has rank $r\kappa$ and Q is obtained as a linear transformation of Q' , and has exactly $r\kappa$ columns, it must also have had rank $r\kappa$, so μ satisfies $RP(r, m)$.

The same arguments work verbatim for the extended rank property. \square

One further property we need of our phylogenetic models is a certain notion of standardness of the set of transition matrices, as that will allow us to deduce that the matrices have a suitable Kruskal rank.

Definition 3.6.4 *A Markov model μ is standard if it satisfies the following conditions:*

- *the model is generically identifiable on trees*
- *the model contains the identity matrix, and*
- *the model contains two transition matrices M_1 and M_2 such that any pair of column vectors obtained from taking two columns of M_1 , or one column of M_1 and one column of M_2 are linearly independent.*

Proposition 3.6.5 *Let μ be a standard Markov model. Let T_1, \dots, T_r be trees on n leaves with $n \geq 2$. Let P_1, \dots, P_r be generic probability distributions $P_i \in M_{T_i}^\mu$. Then the matrix*

$$Q = \left(\text{Flat}_{[n-1]n}(P_1) \quad \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \text{Flat}_{[n-1]n}(P_r) \right)$$

has Kruskal rank ≥ 2 for any r .

Proof: By marginalizing, we can assume that all the trees are just a single edge with 2 leaves. To see that the matrices have Kruskal rank ≥ 2 , we just need to look at the case where $r = 1, 2$, and

taking any pair of columns from the resulting two transition matrices. However, the definition of standard specifically covers these cases, which shows that matrix Q has Kruskal rank at least 2. \square

Last we come to the No Shuffling Property, which says that parts of probability distributions on different trees cannot be combined to get a distribution on another tree.

Definition 3.6.6 *A Markov model μ is said to have the No Shuffling Property if the following holds: For any trees T_1, \dots, T_r with $n \geq 3$ leaves and generic probability distributions $P_i \in M_{T_i}^\mu$, let Q be the matrix*

$$Q = \left(\pi_1 \text{Flat}_{[n-1]n}(P_1) \quad \pi_2 \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \pi_r \text{Flat}_{[n-1]n}(P_r) \right)$$

where $\pi \in \Delta_r$ is generic. Form a new matrix Q' by taking κ not necessarily distinct columns of Q , and let P be the resulting tensor. If $P \in M_T^\mu$ for some T , then all columns of Q' must have come from the same P_i .

The importance of the No Shuffling Property is that it implies that if we permute the columns of Q , we can still tell which columns belong together (coming from the same distribution and same tree). This follows because it is not possible to get a distribution from a tree by taking a combination of some of the columns from different trees.

Our goal at this point is to show how a model that satisfies these properties can be made to follow the proof strategy from [80], closely following the outline of those ideas, but substituting in these new properties in place of directly using the General Markov model to prove the results.

Lemma 3.6.7 *Let μ be a Markov model that satisfies the extended rank property $ERP(r, k)$ and let $r \leq \kappa^{k-1}$. Let P be a probability that is a generic mixture of r generic distributions from trees T_1, \dots, T_r on n leaves. Suppose that $A|B$ is a split with $\#A > k$ and $\#B > k$.*

1. *If every tree T_1, \dots, T_r displays the split $A|B$ then $\text{rank}(\text{Flat}_{A|B}(P)) \leq \kappa r$.*
2. *If there is some tree T_i that does not display the split $A|B$ then $\text{rank}(\text{Flat}_{A|B}(P)) > \kappa r$.*

Proof: If $A|B$ is compatible with all trees in \mathbf{T} , then, by passing to binary resolutions of the T_i , let $A|B$ be associated to the edge $e_i = (a_i, b_i)$ in T_i . Then one sees that

$$\text{Flat}_{A|B}(P) = M_A^T Q M_B.$$

Here Q is the $\kappa r \times \kappa r$ block-diagonal matrix whose i th $\kappa \times \kappa$ block gives the joint probability distribution of states for the random variables at a_i and b_i , weighted by the component proportion π_i . The matrices M_A, M_B are stochastic, of sizes $\kappa r \times \kappa^{\#A}, \kappa r \times \kappa^{\#B}$, with entries in the i th block of κ rows giving probabilities of states of variables in A, B conditioned on states at a_i, b_i . Since Q is a $\kappa r \times \kappa r$ matrix it has rank at most κr , which implies that $\text{Flat}_{A|B}(P)$ has rank at most κr as well.

Suppose next that $A|B$ is not compatible with at least one of the trees in \mathbf{T} , say T_1 . To show that $\text{Flat}_{A|B}(P)$ generically has rank greater than κr , it is enough to give a single choice of parameters producing such a rank. Indeed, this follows from Proposition 3.2 in [80], applied to the model and the variety of matrices of rank at most κr .

For each T_i with $i > 1$ choose all Markov matrices for all internal edges of T_i to be the identity, I_κ . Since T_1 is not compatible with $A|B$, by Theorem 3.8.6 of [19], it has an edge $e = (c, d)$, with associated split $C|D$, such that all four sets $A \cap C, A \cap D, B \cap C, B \cap D$ are nonempty. For all internal edges of T_1 except e , choose Markov matrices to be I_κ as well. Since the effect of an identity matrix on an edge is the same as contracting that edge. For simplicity, we will refer to these contradicted trees by their original labeling and just work with these contracted edge trees. This reduces our model to the case where for $i > 1$, T_i is a star tree with central node a_i , and T_1 has the form of two star trees, on C and on D , that are joined at their central nodes by e .

Let $P = P_1 + P'$ where P_1 is the mixture component from T_1 , and P' the sum of the components on the star trees $T_2 = \dots = T_r$. Then one sees that

$$M_2 := \text{Flat}_{A|B}(P') = N_A^T R N_B,$$

with R an $\kappa(r-1) \times \kappa(r-1)$ diagonal matrix giving the distribution of states at a_i in components $2, \dots, r$ weighted by π_i and N_A, N_B are stochastic matrices of sizes $\kappa(r-1) \times \kappa^{\#A}, \kappa(r-1) \times \kappa^{\#B}$ with entries giving conditional probabilities of states of variables in A, B conditioned on states/components at the a_i . By choosing the positive root distributions at the nodes a_i , and positive π_i , R is ensured to have positive diagonal entries, and hence have full rank.

Consider P_1 where all matrices on pendant edges of T_1 are chosen to be I_κ . Also, both the root distribution at c and M_e are chosen to have all positive entries. Then let

$$M_1 := \text{Flat}_{A|B}(P_1) = N_{1,A}^T R_1 N_{1,B}.$$

where R_1 is a $\kappa^2 \times \kappa^2$ diagonal matrix with entries giving the joint distribution at c and d weighted by π_1 , and $N_{1,A}, N_{1,B}$ have all zero entries except for a single 1 in each row, and full

row rank. Thus M_1 has rank κ^2 . Moreover, it has at most one non-zero entry in each row and column, so both $\text{im}(M_1)$ and $\text{ker}(M_1)$ are coordinate subspaces.

Since $\text{Flat}_{A|B}(P) = M_1 + M_2$ and μ has the extended rank property $ERP(r, k)$, $\text{rank}(M_1 + M_2) > \kappa r$ and thus we are done. \square

Next goal is to show how to use the ranks of flattenings from Lemma 3.6.7 to identify the key features of a phylogenetic mixture model.

Lemma 3.6.8 *Suppose that $\mathbf{T} = (T_1, \dots, T_r) \in \mathcal{T}(r, n, k)$, and let $P \in M_{\mathbf{T}}^{\mu}$ be generic. Suppose that μ satisfies $RP(n, k)$ and $ERP(n, k)$. Then we can use ranks of flattenings to find a partition on $[n]$ into three sets A, B, C with $\#A \geq k$, $\#B \geq k$, such that $A|B \cup C$ and $B|A \cup C$ are valid splits for all the trees T_1, \dots, T_r .*

Proof: To find such an $A|B|C$, we simply test all the tripartitions that have $\#A \geq k$ and $\#B \geq k$. We know, by the definition of $\mathcal{T}(r, n, k)$, that there must exist a triple $A'|B'|C'$ where each of the splits $A'|B' \cup C'$, $B'|A' \cup C'$, and $C'|A' \cup B'$ holds for all the trees T_1, \dots, T_r . According to Lemma 3.6.7, a split $D|E$ has $\text{rankFlat}_{D|E}(P) \leq r\kappa$ if and only if the split $D|E$ appears in all the trees T_1, \dots, T_r , provided that $\kappa^{\#D} > r\kappa$ and $\kappa^{\#E} > r\kappa$. By our assumptions on k , this holds. So, the ranks of flattenings will find a tripartition $A|B|C$ of the desired type. \square

Note that just because we find $A|B \cup C$ and $B|A \cup C$ using the ranks of flattenings, it does not necessarily imply that $C|A \cup B$ is a valid split in all the trees. This is in spite of the fact that we know that there exists a triple $A'|B'|C'$ that is a common tripartition to all trees. Flattenings cannot necessarily find that tripartition alone. However, we can use Lemma 3.6.8 to prove our identifiability results anyways. This result follows the proof of Theorem 4.4 of [80].

Lemma 3.6.9 *Let μ be a standard Markov model that satisfies $RP(r, k)$ and the No Shuffling Property. Suppose that the trees $\mathbf{T} = (T_1, \dots, T_r)$ have a known common tripartition $A|B|C$ with $\#A \geq \#B \geq k$. Then both \mathbf{T} and the numerical parameters of the μ -mixture model on \mathbf{T} are generically identifiable.*

Proof: Since all the trees in \mathbf{T} share a common tripartition, we can write a distribution in the mixture model as a triple product

$$\text{Flat}_{A|B|C}(P) = [M_A, M_B, M_C]$$

where the matrices are M_A , M_B , and M_C are as described at the end of Section 3.5.

Since the model μ satisfies the rank property $RP(r, k)$, for generic choices of parameter the matrices M_A and M_B will each have rank $r\kappa$. The fact that μ is standard guarantees that M_C has Kruskal rank ≥ 2 .

The triple product representation expresses $\text{Flat}_{A|B|C}(P)$ as the sum of $r\kappa$ rank 1 tensors. But since

$$\text{rank}_K(M_A) + \text{rank}_K(M_B) + \text{rank}_K(M_C) \geq r\kappa + r\kappa + 2 = 2r\kappa + 2$$

we can apply Kruskal's theorem to see that the matrices M_A , M_B , and M_C can be recovered up to scaling and permuting the rows.

Each of the rows of the recovered matrices M_A , M_B , M_C will have entries from a scaled slice from a tree distribution on a subtree of one of the T_i (the subtree from the common vertex to the leaves A). We need to group these rows together by the mixture components they come from. However, since we have assumed that μ satisfies the No Shuffling Property, there is only one way to do this for M_A . Since the ordering of rows of M_A denotes the order of the rows of M_B and M_C , we can reassemble each scaled probability distribution $\pi_i P_i$ as the triple product $[M_{i,A}, M_{i,B}, M_{i,C}]$ for submatrices of M_A, M_B, M_C respectively.

From the scaled distribution $\pi_i P_i$ we recover the mixing weight via the sum

$$\pi_i = \sum_{(j_1, \dots, j_n) \in [k]^n} \pi_i P_i(j_1, \dots, j_n).$$

Then we can get the distribution P_i for the single tree T_i . We use that μ is a standard model, so the tree parameter T_i and numerical parameters are generically identifiable. \square

The next Lemma closely follows the proof of Theorem 4.6 in [80].

Lemma 3.6.10 *Let μ be a standard Markov model that satisfies $RP(r, k)$ and the No Shuffling Property. Suppose that the trees $\mathbf{T} = (T_1, \dots, T_r)$ have a known tripartition with $A|B|C$ with $\#A \geq \#B \geq k$ such that $A|B \cup C$ and $B|A \cup C$ are valid splits in all trees T_1, \dots, T_r . Then both \mathbf{T} and the numerical parameters of the μ -mixture model on \mathbf{T} are generically identifiable.*

Proof: Let P be a generic distribution of the μ mixture model. Fix some $c \in C$, let $D(c) = A \cup B \cup \{c\}$. Consider the marginalization of P to the set $D(c)$, and call this distribution P_c . This is the probability tensor for the induced r -tuple of trees

$$\mathbf{T}|_{D(c)} = (T_1|_{D(c)}, \dots, T_r|_{D(c)}).$$

Note that all the trees of $\mathbf{T}|_{D(c)}$ share the common tripartition $A|B|\{c\}$, which satisfies the conditions of Lemma 3.6.9, so that tree parameters and numerical parameters are generically identifiable on this subset of the leaves. We also have that

$$\text{Flat}_{A|B|\{c\}}(P_c) = [M_A, M_B, M_c].$$

Since μ satisfies the rank property $RP(r, k)$, we can assume that M_A has full rank. Thus, there is a right inverse matrix Q_A with the property that $M_A Q_A = I_{r\kappa}$, where $I_{r\kappa}$ is the $r\kappa$ identity matrix. Our goal is to use the matrix Q_A to finish the disentangling of distributions that go into P .

Now, since $A|B \cup C$ is a valid split in all the trees \mathbf{T} , we can write a factorization

$$\text{Flat}_{B \cup C|A}(P) = M_{B \cup C}^T \Pi \widetilde{M}_A$$

where \widetilde{M}_A and $M_{B \cup C}$ are stochastic matrices of probabilities of the states of the leaves in A and $B \cup C$ conditioned which tree T_i we are in and on the states at the root w_i in each tree where $T_i|_A$ attaches to the rest of T_i . The matrix Π is a diagonal matrix whose entries are the root distribution probabilities times the mixing weights. A key feature is that the order of all components can be taken so that all the parameters associated to a particular tree T_i can be assumed to be in the same block of rows.

Now we can write $M_A = R \widetilde{M}_A$ where R is a block diagonal matrix whose i th block gives the conditional probability of state changes from the root w_i to the adjacent vertex in T_i away from A . We can assume that this matrix is invertible since the blocks are just μ -transition matrices, and generically those are nonsingular since μ is standard.

Now we can compute

$$\text{Flat}_{B \cup C|A}(P) Q_A = M_{B \cup C}^T \Pi R^{-1} M_A Q_A = M_{B \cup C}^T \Pi R^{-1}.$$

This shows that taking the columns of $\text{Flat}_{B \cup C|A}(P) Q_A$ in blocks of κ we obtain entries associated to only one mixture component at a time. Multiplying a block of those columns by the associated corresponding rows of $M_A = R \widetilde{M}_A$ we obtain a single mixture component $\pi_i P_i$ from the single tree T_i , multiplied by the mixing weight π_i . We can identify π_i by summing all entries of this tensor, as in the proof of Lemma 3.6.9. Then the fact that the tree and numerical parameters are generically identified for the model μ for a single tree completes the proof. \square

Now we are in a position to combine all components to complete the proof of the main structural theory on identifiability of mixture models.

Proof: [Proof of Theorem 3.2.5] Let P be a generic distribution from a mixture model with some $\mathbf{T} = (T_1, \dots, T_r) \in \mathcal{T}(r, n, k)$. We need to show that from P alone, we can find \mathbf{T} and the numerical parameters of the model. According to Lemma 3.6.8, it is possible to use ranks of flattenings to find a tripartition $A|B|C$ of the leaf set such that $A|B \cup C$ and $B|A \cup C$ are valid splits in all the trees T_i , and both $\#A \geq k$ and $\#B \geq k$. Once those are identified, Lemma 3.6.10 shows that it is possible to identify the trees T_i , the mixing weights, and the numerical parameters for each tree. \square

3.7 Rank property, Extended Rank Property, and Standard Property for the Jukes-Cantor Model

In this section, we will prove that the Jukes-Cantor model satisfies the Rank Property, Extended Rank Property, and Standard condition with appropriate conditions on r and k . While this might seem narrow, a key observation is that if any of these properties are satisfied for a certain model μ , they are also satisfied for all models μ' that contain μ as submodels. Since the Jukes-Cantor model is contained in all the equivariant models as a submodel, this will prove that those three properties are also satisfied for those models.

First we will prove that the models under consideration in this chapter are all standard models, the most straightforward property to prove for a model.

Lemma 3.7.1 *The JC, K2P, K3P, SSM, and GMM models are all standard Markov models.*

Proof: All five models contain the identity matrix, and are known to be identifiable on trees (e.g. using the tensor rank arguments from [4]).

Recall that the Jukes-Cantor model consists of all transition matrices that have one value b for all off diagonal entries and a different value a for all diagonal entries. Clearly, the Jukes-Cantor model contains the identity matrix setting $b = 0$ and $a = 1$. For the second property related to Kruskal ranks we can take the identity matrix, together with any other matrix in the model that does not have rank 1 (e.g. take $b = \epsilon$ and $a = 1 - (\kappa - 1)\epsilon$). This pair of matrices will satisfy the condition on independence of column vectors. \square

A key fact that we will use (and that was also used in the proofs in [80]), is that matrices that are generalized Vandermonde matrices have full rank for generic choices of parameters.

Definition 3.7.2 *Let x^{u_1}, \dots, x^{u_r} be monomials in $\mathbb{C}[x_1, \dots, x_n]$. Let $v_1, \dots, v_s \in \mathbb{C}^n$ vectors. Then*

the matrix

$$V(u_1, \dots, u_r; v_1, \dots, v_s) = \begin{pmatrix} v_1^{u_1} & v_1^{u_2} & \cdots & v_1^{u_r} \\ \vdots & \vdots & \ddots & \vdots \\ v_s^{u_1} & v_s^{u_2} & \ddots & v_s^{u_r} \end{pmatrix}$$

is called the generalized Vandermonde matrix.

Proposition 3.7.3 *Let x^{u_1}, \dots, x^{u_r} be distinct monomials in $\mathbb{C}[x_1, \dots, x_n]$. Let $v_1, \dots, v_r \in \mathbb{C}^n$ generic vectors. Then the generalized Vandermonde matrix $V(u_1, \dots, u_r; v_1, \dots, v_r)$ has rank r .*

Proof: Since all vectors u_1, \dots, u_r are distinct, there is an assignment of variables $x_i = t^{d_i}$, so that the resulting powers of the variable $t^{u_1 d}, t^{u_2 d}, \dots$ are all distinct. Then the result follows from the fact that the standard Vandermonde matrix has full rank for generic values. \square

More generally, we have the following useful fact about the matroid that is determined by a generalized Vandermonde matrix.

Proposition 3.7.4 *Let x^{u_1}, \dots, x^{u_r} be not necessarily distinct monomials in $\mathbb{C}[x_1, \dots, x_n]$. Suppose that there are l distinct monomials among them. Let $v_1, \dots, v_s \in \mathbb{C}^n$ generic vectors. Then any nonzero vector in the row span of $V(u_1, \dots, u_r; v_1, \dots, v_s)$ has at least $l - s + 1$ nonzero entries.*

Proof: We can assume that there are no repeats in the list of monomials, since any repeats must necessarily yield repeated nonzero entries of vectors in the row span. Suppose that there is a nonzero vector x in the row span of $V = V(u_1, \dots, u_r; v_1, \dots, v_s)$ that had s or more nonzero entries. Let S be a set of exactly s of those entries. After permuting columns, we can assume that those are the first s entries of x . The fact that x is a nonzero vector of the row span means that there is an invertible $s \times s$ matrix M such that MV has x as the bottom row. Consider the square-submatrix $V' = V(u_1, \dots, u_s; v_1, \dots, v_s)$ obtained by taking the first s columns of V . By Proposition 3.7.3 V' is invertible, so MV' is also invertible. But MV' has a row of all zeroes (from taking the subvector of x). This shows that there must be at least $r - s + 1$ non-zero entries in x . \square

Lemma 3.7.5 *The Jukes-Cantor model satisfies the rank property $RP(r, k)$ for $r \leq 2^{k-1} - k + 1$ when $\kappa \geq 3$.*

Proof: Let r and k satisfy $r \leq 2^k - k + 1$. We must show that for any trees T_1, \dots, T_r on k leaves, and generic probability distributions $P_i \in M_{T_i}^\mu$, the matrix

$$\left(\text{Flat}_{[n-1]n}(P_1) \quad \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \text{Flat}_{[n-1]n}(P_r) \right)$$

has rank $r\kappa$. However, by the fact that rank of a matrix being $\leq \alpha$ is a closed condition, it suffices to show that there is a single choice of parameters that gives the rank $r\kappa$.

Suppose that T is a tree. Note that if we set the transition matrix on an edge to be the identity matrix (which is a transition matrix in the Jukes-Cantor model), that will give a probability distribution on a tree obtained from T by contracting the corresponding edge. Hence, if we set the transition matrices of all internal edges of the tree to be the identity matrix, this will give us distributions on the star tree. So distributions on the star tree appear as distributions in the model on any tree. Thus the result follows if we prove the Lemma when $T_1 = T_2 = \cdots = T_r$ are all the star trees.

To this end, let T_1, \dots, T_r be star trees with k leaves, and consider the matrix

$$M = \left(\text{Flat}_{[k-1]k}(P_1) \quad \text{Flat}_{[k-1]k}(P_2) \quad \cdots \quad \text{Flat}_{[k-1]k}(P_r) \right).$$

Our first step is to apply the Fourier transform to the parametrization. Let Q_i be the Fourier transformation of the probability distribution P_i . The Fourier transform is linear, and it transforms the matrix M into a new matrix M' which is

$$M' = \left(\text{Flat}_{[k-1]k}(Q_1) \quad \text{Flat}_{[k-1]k}(Q_2) \quad \cdots \quad \text{Flat}_{[k-1]k}(Q_r) \right).$$

The new matrix M' is obtained from M by row and column operations, so M and M' have the same rank.

Now we analyze the parameterization in the Fourier coordinates. For the Jukes-Cantor model on a star tree T_i we have that

$$q(g_1, \dots, g_k) = \begin{cases} \prod_{i=j}^k a_{g_j}^{(i,j)} & \sum_{j=1}^k g_j = 0 \\ 0 & \text{otherwise} \end{cases}$$

where we have a set of parameter $a_g^{(i,j)}$ for each tree T_i , each edge j , and each group element

$g \in G$. For the Jukes-Cantor model we have that

$$a_{(1,0)}^{(i,j)} = a_{(0,1)}^{(i,j)} = a_{(1,1)}^{(i,j)}$$

for all i and j . We can rearrange rows and columns of M' so that it has a block form, grouping all the columns by the value of g_k , and grouping the rows so that, for a fixed value of g_k we have all the (g_1, \dots, g_{k-1}) so $q(g_1, \dots, g_k) \neq 0$ are together.

After this rearrangement of rows and columns, M' will be a block diagonal matrix, with κ blocks, each block of size $\kappa^{k-2} \times r$. We need to show that each of these blocks has full rank, so we get that the total rank is $r \times \kappa$ as desired.

To show that that block matrices have the appropriate rank, we note that each such matrix is a generalized Vandermonde matrix. Indeed, each entry of the matrix is a monomial, and each column is an identical copy of the first column, but with new variables. Note, however, that the condition $a_{(1,0)}^{(i,j)} = a_{(0,1)}^{(i,j)} = a_{(1,1)}^{(i,j)}$ will yield repeated monomials in each column. So to complete the proof, we need to figure out how many distinct monomials there are, so we can apply Proposition 3.7.3.

Consider the map from

$$\phi : \mathbb{Z}_2 \times \mathbb{Z}_2 \rightarrow \{0, 1\}, \phi(g, h) = \begin{cases} 0 & \text{if } g = h = 0 \\ 1 & \text{otherwise.} \end{cases}$$

Then two monomials $q(g_1, \dots, g_k)$ and $q(h_1, \dots, h_k)$ are identical if and only if $\phi(g_i) = \phi(h_i)$ for all i . So we just need to count the number of equivalence classes for each fixed value of g_k .

This is straightforward to do: if $\phi(g_k) = 0$, then $(\phi(g_1), \dots, \phi(g_{k-1}))$ can be any string in $\{0, 1\}^{k-1}$ except the strings that have exactly one 1. There are $2^{k-1} - k + 1$ such strings. On the other hand, if $\phi(g_k) \neq 0$ then $(\phi(g_1), \dots, \phi(g_{k-1}))$ can be any string in $\{0, 1\}^{k-1}$ except the string with all zeroes. There are $2^{k-1} - 1$ such strings. We need to take the smaller of these two values to get a consistent rank across all the blocks. Hence this shows that the Jukes-Cantor model satisfies $RP(r, k)$ with $r \leq 2^{k-1} - k + 1$. \square

Corollary 3.7.6 *The Kimura 2-parameter model, the Kimura 3-parameter model, and the strand symmetric model all satisfy the rank property $RP(r, k)$ with $r \leq 2^{k-1} - k + 1$.*

Proof: We just need to show the existence of a single choice of parameters that give the desired rank condition. However, since the Jukes-Cantor model is a submodel of all of those

other models, the result of Lemma 3.7.5 gives the desired result. \square

Note that the bound $r \leq 2^{k-1} - k + 1$ is not best possible for those other models besides the Jukes-Cantor model. Each model would require a more careful analysis to improve the results. Following the proof, for the group-based models, it suffices to determine the number of distinct monomials of different types in the block structure of the matrix M' . Both Kimura 2-parameter and 3-parameter models will have significantly more distinct monomials than the Jukes-Cantor model, and so the rank property will hold for larger values of r .

Now we proceed to prove the Extended rank property for the Jukes-Cantor model. Again, that will also give a result for other models, though it is probably not the best possible for K2P, K3P, SSM.

Lemma 3.7.7 *The Jukes-Cantor models satisfies the extended rank property ERP(r, k) for $r \leq 2^{k-1} - k + 1$ when $\kappa \geq 3$.*

Proof: Let T_1 be a tree on $k + 1$ leaves that does not contain the split $[k - 1]||\{k, k + 1\}$, and any trees T_2, \dots, T_r on k leaves, and generic probability distributions $P_i \in M_{T_i}^\mu$. Consider the matrix the matrix

$$M = \left(\text{Flat}_{[k-1]||\{k, k+1\}}(P_1) \quad \text{Flat}_{[k-1]||k}(P_2) \quad \cdots \quad \text{Flat}_{[k-1]||k}(P_r) \right).$$

We must show that M has rank $(r - 1)\kappa + \kappa^2$, generically.

As in the proof of Lemma 3.7.5, it suffices to prove that there is a single choice of parameters that achieves the desired rank. Then we can set many parameters equal to identity matrices, and consider the resulting trees that arise by contracting those edges. To that end, we can assume that trees T_2, \dots, T_r are all k leaf star trees.

As in the proof of Lemma 3.7.5, we apply the Fourier transform to all probability distributions to get a matrix

$$M' = \left(\text{Flat}_{[k-1]||\{k, k+1\}}(Q_1) \quad \text{Flat}_{[k-1]||k}(Q_2) \quad \cdots \quad \text{Flat}_{[k-1]||k}(Q_r) \right).$$

Looking at the final blocks we have the matrix

$$\left(\text{Flat}_{[k-1]||k}(Q_2) \quad \cdots \quad \text{Flat}_{[k-1]||k}(Q_r) \right).$$

This is the same matrix we have seen in the proof of Lemma 3.7.5, with one fewer set of columns. So it has rank $(r - 1)\kappa$ since JC has the rank property $RP(r - 1, k)$ with these values of r and k . Furthermore, after reordering rows and columns, as in the proof of Lemma 3.7.5, it can be

broken into 4 blocks, each of which is a generalized Vandermonde matrix.

The 4 generalized Vandermonde matrices from the previous paragraph each have size $r - 1 \times 4^{k-2}$, and those generalized Vandermonde matrices each have full rank. We want to use Lemma 3.7.4 to complete the proof. In particular, we will show that there is a choice of parameters for the tree T_1 so that the resulting matrix $\text{Flat}_{[k-1]\{\{k, k+1\}\}}(Q_1)$ has the property that each of the four vectors that it contributes to the support of one of the four Vandermonde submatrices has at most 4 nonzero entries. This will insure that M' has the correct rank, by applying Lemma 3.7.4. With these observations in mind, we consider various restriction on the tree T_1 and matrices $\text{Flat}_{[k-1]\{\{k, k+1\}\}}(Q_1)$ that can be produced.

First of all, T_1 must have a split that is that is not compatible with $[k - 1]\{\{k, k + 1\}\}$. After relabeling the leaves, we can assume this split has the form

$$A_j|B_j = \{1, \dots, j - 1, k\}|\{j, \dots, k - 1, k + 1\}$$

for some j between 2 and $k - 1$. We can assume that T_1 is the tree with only this one split $A_j|B_j$ as an internal split, by setting all other internal edges to have an identity matrix as the transition matrix. We must consider a few scenarios based on the sizes of the sets A_j and B_j .

Case 1: Both $\#A_j$ and $\#B_j$ are even. In this case, we take all the transition matrices associated to pendant edges in the tree to be the identity matrix as well. Our assumption on T_1 yields a specific structure on the Fourier coordinates of the distributions in the model on T_1 . Considering a vector $(g_1, \dots, g_{k+1}) \in G^{k+1}$ we can break this into blocks $(\mathbf{h}_1, \mathbf{h}_2, g_k, g_{k+1})$, where

$$\mathbf{h}_1 = (g_1, \dots, g_{j-1}) \quad \mathbf{h}_2 = (g_j, \dots, g_{k-1}).$$

For a Fourier coordinate $q(\mathbf{h}_1, \mathbf{h}_2, g_k, g_{k+1})$ to be nonzero, we must have the following

1. $g_1 = g_2 = \dots = g_{j-1} = g_k$
2. $g_j = g_{j+1} = \dots = g_{k-1} = g_{k+1}$
3. $\sum_{j=1}^{k+1} g_j = (0, 0)$.

Conditions (1) and (2) are coming from the fact that the only nontrivial transition matrix for the model is the one corresponding to the split $A_j|B_j$. Condition (3) is the standard condition for Fourier coordinates in a group-based model. However, since both $\#A_j$ and $\#B_j$ are even, condition (3) is automatically satisfied by any vectors that satisfy conditions (1) and (2). Note that there are exactly 16 nonzero values of $q(\mathbf{h}_1, \mathbf{h}_2, g_k, g_{k+1})$, one for each of the possible pairs

$(g_k, g_{k+1}) \in G^2$. Since the columns of $\text{Flat}_{[k-1]\{k,k+1\}}(Q_1)$ are indexed by those pairs, we see that each column of $\text{Flat}_{[k-1]\{k,k+1\}}(Q_1)$ has exactly one nonzero entry, and they appear in different rows. Furthermore, the rearrangement of rows and columns so that

$$\left(\text{Flat}_{[k-1]k}(Q_2) \quad \cdots \quad \text{Flat}_{[k-1]k}(Q_r) \right)$$

is a block matrix has blocks indexed by the value of g_k . Hence, we have that each block of $\text{Flat}_{[k-1]\{k,k+1\}}(Q_1)$ contributes a four dimensional column space, all of which have at most 4 nonzero entries. This shows that M' will have rank $(r-1)4 + 4^2$, as desired.

Case 2: One of $\#A_j$ and $\#B_j$ is even, and one is odd. We can assume that $\#A_j$ is odd and $\#B_j$ is even. We set all the pendant edge parameters to the identity matrix except for the edge going to leaf $k+1$. As in Case 1, we consider a vector $(g_1, \dots, g_{k+1}) \in G^{k+1}$ we can break this into blocks $(\mathbf{h}_1, \mathbf{h}_2, g_k, g_{k+1})$, where

$$\mathbf{h}_1 = (g_1, \dots, g_{j-1}) \quad \mathbf{h}_2 = (g_j, \dots, g_{k-1}).$$

For a Fourier coordinate $q(\mathbf{h}_1, \mathbf{h}_2, g_k, g_{k+1})$ to be nonzero, we must have the following

1. $g_1 = g_2 = \cdots = g_{j-1} = g_k$
2. $g_j = g_{j+1} = \cdots = g_{k-1}$
3. $\sum_{j=1}^{k+1} g_j = (0, 0)$.

Note the change that we will not need g_{k+1} to be equal to the other values in Condition (2). Hence we have three groups of coordinates, each of which have an odd number of elements (the groups being $A_j, \{k+1\}$ and $B_j \setminus \{k+1\}$). Let

$$h_1 := g_1 = g_2 = \cdots = g_{j-1} = g_k \quad \text{and} \quad h_2 := g_j = g_{j+1} = \cdots = g_{k-1}$$

then we get a valid coordinate when

$$h_1 + h_2 + g_{k+1} = (0, 0).$$

This follows because each h_i is equal to the sum of the g_i 's in its group, because all are equal and the number of such elements is odd. We see that there are exactly 16 possibly solutions (choosing values for h_1 and h_2 arbitrarily forces a value for g_{k+1}). Furthermore, all possible pairs coordinates (g_k, g_{k+1}) are possible. Thus, as in Case 1, we see that each column of $\text{Flat}_{[k-1]\{k,k+1\}}(Q_1)$ has exactly one nonzero entry, and they appear in different rows. Furthermore, the rearrangement

of rows and columns so that

$$\left(\text{Flat}_{[k-1]||k}(Q_2) \cdots \text{Flat}_{[k-1]||k}(Q_r) \right)$$

is a block matrix has blocks indexed by the value of g_k . Hence, we have that each block of $\text{Flat}_{[k-1]||\{k,k+1\}}(Q_1)$ contributes a four dimensional column space, all of which have at most 4 nonzero entries. This shows that M' will have rank $(r-1)4 + 4^2$, as desired.

Case 3: Both $\#A_j$ and $\#B_j$ are odd. We set all the parameters corresponding to pendant edges to the identity matrix except for leaf 1, which we allow to be arbitrary. As in Case 1, we consider a vector $(g_1, \dots, g_{k+1}) \in G^{k+1}$, and we analyze which Fourier coordinates $q(g_1, \dots, g_{k+1})$ can be nonzero. With our assumption on the transition matrices, we must have

1. $g_2 = \dots = g_{j-1} = g_k$
2. $g_j = g_{j+1} = \dots = g_{k-1} = g_{k+1}$
3. $\sum_{j=1}^{k+1} g_j = (0, 0)$.

Note that g_1 does not appear in the first group. So we have three groups where we will have equal values

$$h_1 := g_2 = \dots = g_{j-1} = g_k \quad \text{and} \quad h_2 := g_j = g_{j+1} = \dots = g_{k-1} = g_{k+1}.$$

Since the first group $A_j \setminus \{1\}$ has an even number of elements, the sum of these elements will always be $(0, 0)$ regardless of what h_1 is. On the other hand, the second group has an odd number of elements, so h_2 equals the sum of all those elements. Then by condition (3), we are forced to have $g_1 = h_2$. Hence, there are 16 possible Fourier coordinates that have nonzero entries. As in Cases 1 and 2, they allow for all 16 different possibilities for the pairs (g_k, g_{k+1}) , and will hence give that M' has rank $(r-1)4 + 4^2$ as desired. \square

As in the case of the rank property, we also can see a similar result for the extended rank property, for the other equivariant models we have studied. We omit the proof which is the same as for the rank property. Again, these results are probably not best possible for those other models, and paying attention to the structure of the generalized Vandermonde matrices that arise in the other models can yield stronger results.

Corollary 3.7.8 *The Kimura 2-parameter model, the Kimura 3-parameter model, and the strand symmetric model all satisfy the extended rank property ERP(r, k) with $r \leq 2^{k-1} - k + 1$.*

3.8 The No Shuffling Property

In this section, we prove that the models JC, K2P, K3P, and SSM satisfy the No Shuffling Property. Like the results from the previous section on the rank property and extended rank property, we are able to just show the result for the JC model and immediately deduce it for all other models.

Recall the No Shuffling Property:

Definition 3.8.1 *A Markov model μ is said to have the No Shuffling Property if the following holds: For any trees T_1, \dots, T_r with $n \geq 3$ leaves and generic probability distributions $P_i \in M_{T_i}^\mu$, let Q be the matrix*

$$Q = \left(\pi_1 \text{Flat}_{[n-1]n}(P_1) \quad \pi_2 \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \pi_r \text{Flat}_{[n-1]n}(P_r) \right)$$

where $\pi \in \Delta_r$ be generic. Form a new matrix Q' by taking κ columns of Q , and let P be the resulting tensor. If $P \in M_T^\mu$ for some T , then all columns of Q' must have come from the same P_i .

Note the No Shuffling Property does not have any conditions on r . The results should hold for every r (and we can of course, stop at $r = \kappa$). Also, note that if we prove the No Shuffling Property for a Markov model μ just for trees with 3 leaves, this will prove that the property also holds for trees with n leaves with $n \geq 3$ as well. Otherwise, there is some n , and trees T_1, \dots, T_r and generic distributions $P_i \in M_{T_i}^\mu$, where we form the matrix

$$Q = \left(\pi_1 \text{Flat}_{[n-1]n}(P_1) \quad \pi_2 \text{Flat}_{[n-1]n}(P_2) \quad \cdots \quad \pi_r \text{Flat}_{[n-1]n}(P_r) \right)$$

where $\pi \in \Delta_r$ be generic. Then we can form a new matrix Q' by taking κ columns of Q , and let P be the resulting tensor. If $P \in M_T^\mu$ for some T , then we can marginalize to tree leaves and the same statement will be true for three leaves as well, contradicting that we had the No Shuffling Property for $n = 3$. Hence, for the rest of this section, we will restrict to three leaf trees.

To show the No Shuffling Property, we will use phylogenetic invariants that come from the general Markov model, specifically certain commutation invariants described in [8, 87].

Theorem 3.8.2 *Let P be a $\kappa \times \kappa \times \kappa$ tensor giving a distribution from the general Markov model with κ states on a 3 leaf tree with $\kappa > 2$. For $i = 1, \dots, \kappa$ let $P_{(i)}$ be the $\kappa \times \kappa$ matrix slice $P_{(i)} = (P(i, u, v))_{u,v}$. Then*

$$P_{(i)}(\text{adj}(P_{(j)}))P_{(k)} - P_{(k)}(\text{adj}(P_{(j)}))P_{(i)}$$

for any i, j, k .

Note that $\text{adj}(A)$ denotes the classical adjoint of A , which is $\det(A)A^{-1}$ for an invertible matrix, and is well-defined even for singular matrices (since the entries are all polynomials in the entries of A).

Lemma 3.8.3 *Let μ be any phylogenetic Markov model on $\kappa = 4$ states that contains the Jukes-Cantor model. Then μ satisfies the No Shuffling Property.*

Proof: We proved this result via a computation, and using the phylogenetic invariants from the general Markov model described in Theorem 3.8.2. Specifically, we generated three random tensors from the Jukes-Cantor model parameterization P^1, P^2, P^3 . Then from those $4 \times 4 \times 4$ tensors, we choose three matrix slices, and evaluate them in the polynomials from Theorem 3.8.2. We find that that this evaluates to zero only if the three matrix slices come from the same P^i .

The fact that the invariants from the General Markov Model evaluate to zero only when all three slices come from the same P^i , proves that all models that contain the Jukes-Cantor model will satisfy the No Shuffling Property. Indeed, if the No Shuffling Property is not satisfied for some model μ , then for all generic distributions from that model, there will be some choice of slices from different P^1, P^2, P^3 , that satisfy the invariants that come from that model. The fact that the invariants vanish for the generic distributions in the model, implies that they vanish for all distributions that come from the model. Since the General Markov Model contains every phylogenetic model, the invariants for the General Markov model are contained in the ideal of invariants for any other model. Since we have produced a distribution in the Jukes-Cantor model that does not satisfy the phylogenetic invariants for the General Markov model with any non-standard set of slices, this proves that no model that contains the Jukes-Cantor model will satisfy those invariants for General Markov model. Hence any model containing the Jukes-Cantor model satisfies the No Shuffling Property. \square

With the tools from Sections 3.7 and 3.8 in hand, we are ready to apply Theorem 3.2.5 to deduce Corollary 3.2.6.

Proof: [Proof of Corollary 3.2.6] According to Theorem 3.2.5, to deduce Corollary 3.2.6, we must show that all of the models JC, K2P, J3P, SSM, GMM:

1. are standard Markov models,
2. satisfy the rank property $RP(2^k - k, k)$,
3. satisfy the extended rank property $ERP(2^k - k, k)$, and

4. satisfy the No Shuffling property.

The fact that all the models are standard is the content of Lemma 3.7.1.

The JC model satisfies the rank property $RP(2^k - k, k)$ by Lemma 3.7.5. Any supermodel of the JC model will hence satisfy $RP(2^k - k, k)$ as well. The JC model satisfies the extended rank property $ERP(2^k - k, k)$ by Lemma 3.7.7. Any supermodel of the JC model will hence satisfy $ERP(2^k - k, k)$ as well. Finally, Lemma 3.8.3 shows that all the models that contain JC satisfy the No Shuffling Property. \square

CHAPTER

4

METHODOLOGICAL CONSIDERATIONS FOR SEMIALGEBRAIC HYPOTHESIS TESTING WITH INCOMPLETE U-STATISTICS

4.1 Introduction

Statistical models are typically described by a map from a parameter space to a set of distributions. Often the parameter space Θ can be identified with a full-dimensional subset of \mathbb{R}^d with submodels arising by restricting to a subset $\Theta_0 \subset \Theta$. In many instances Θ_0 is described by a set of polynomial equality and inequality constraints on \mathbb{R}^d , in which case we say the submodel is *semialgebraic*. (An *algebraic* model requires polynomial equality constraints only; the prefix *semi*- allows for inequalities.) Semialgebraic models are common in statistics, encompassing many log-linear models [40], latent class models [7, 51], discrete and Gaussian graphical models [68], as well as phylogenetic models [91]. The underlying algebraic structure of semi-algebraic

sets often yields valuable insights into model selection and inference [17, 83, 93, 22].

A semialgebraic set Θ_0 may be geometrically quite complicated. Singularities can occur where the dimension of Θ_0 collapses or it self-intersects. It may also have components of different dimensions, as well as boundaries. Such irregularities create difficulties for standard approaches to hypothesis testing. For instance, a likelihood ratio test using a χ^2 distribution is only justified through approximating the model by a tangent space. While some research has addressed such issues of model geometry [43, 76, 44], it is common for empirical studies to simply ignore the challenges irregularities pose due to the lack of available tools.

Recently, Sturma, Drton, and Leung [89], building on previous work [34, 35, 85], proposed a general hypothesis testing procedure based on randomized incomplete U-statistics [16, 18, 58] in order to overcome these problems. In addition to presenting the method and establishing its asymptotic behaviour, they provided a running example using the tetrad constraints of factor analysis and applied their method to a biological dataset, testing a semialgebraic Gaussian tree model.

In this chapter, we investigate the practical performance of the Sturma, Drton, and Leung (SDL) method through several other models, drawn from evolutionary biology. In particular we study how implementation choices such as constraint specification, kernel order, and decomposition into reducible components affect test performance. Our study offers practical insights for researchers applying the SDL method to semialgebraic models, particularly in biological settings where singularities are common.

Our first example models come from phylogenomics—the inference of species relationships from genomic-scale sequence data. These models are used to test whether biological species relationships are sufficiently described by an evolutionary tree or whether more complex depictions involving hybridization or gene flow are needed. These are semialgebraic submodels of the general trinomial model, allowing for 2-dimensional plotting of rejection regions, providing immediate visual insight into testing behaviour. Although more traditional deterministic tests have been developed for such models (see Appendix A.2), and we do not expect the SDL methodology to supplant them, comparison with those methods allows for better judgment of SDL performance.

We then consider the Cavender-Farris-Neyman (CFN) 2-state model of nucleotide substitution on a 4-taxon gene tree, a more complicated model in a higher dimensional space. After exploring the use of the SDL test for hypothesis testing when assuming a specific gene tree topology, we then adapt the test to present a novel inference procedure for topological gene trees. We emphasize that this procedure depends only on knowing semialgebraic de-

scriptions of the models for different trees without performing any likelihood computation or optimization.

These examples allow us to examine not only the general applicability of the SDL test to biologically meaningful models, but also the practical implications of certain parameter choices that must be made in order to implement the method. We explore the effects of user-specified options on statistical performance such as Type I and Type II errors. We also investigate the stochasticity of the test under different parameter regimes. Since SDL p -values have some randomness due to the test procedure, it is desirable to limit their variation when possible. While [89] suggests that the subsample size used in calculating the incomplete U-statistics should be small, moderately increasing it can greatly reduce variation while still controlling error.

Another user choice examined here is the specific constraints defining the semialgebraic parameter space Θ_0 , as these are not uniquely determined. We show that constraint choice can have a significant effect on the test's rejection region, and that using a redundant set of constraints is often desirable. We offer one approach which automatically produces a redundant set of constraints through convex combinations, making the test less dependent on the initial constraint choice. We also illustrate that redundant constraints not produced by our approach may be needed for better performance. A minimal set of constraints may lead to a highly conservative test, with performance improved by the introduction of valid but seemingly unrelated inequalities.

In addition, the intrinsic geometry of the model also plays a role in unexpected ways. If a model can be decomposed into irreducible components, doing so and using an intersection-union framework with the SDL test on each component can increase the test's statistical power, as one of our examples shows.

Finally, the SDL test procedure depends on a kernel function that must be symmetrized, although this can be computationally prohibitive. However, we found that a partial symmetrization, applying surprisingly few random permutations, is a highly effective substitute and can give good performance.

We emphasize that we ultimately obtained excellent performance of the SDL method for all models we considered. However, we believe that naive use for a specific model of interest, without exploration of the issues we found, is unlikely to achieve the best performance possible. While we give no new theoretical results in this chapter, we advance awareness of potential pitfalls thereby guiding users to better application of the methodology.

This chapter proceeds as follows. In Section 4.2, we introduce relevant background and

outline the methodology from [89]. In Section 4.3, we introduce four basic submodels of the trinomial model, with details of their biological motivation deferred to Appendix A. Section 4.3.4 is the main section of Chapter 4, presenting the issues and lessons learned through application of the hypothesis testing procedure to the four submodels. In Section 4.4, we apply the hypothesis test to the CFN model.

Our implementation of the test in R with the Rcpp package [79] is adapted from code used in the TestGGM package [88] shared by N. Sturma. Our code is freely available on GitHub at github.com/marinagarrote/Semialg-Hypothesis-Test-with-Incomplete-U-Stats, along with Supplementary Materials for this Chapter 4.

4.2 The SDL Test

We first outline the hypothesis testing methodology of [89] for semialgebraic models, henceforth referred to as the *SDL test*.

4.2.1 Semialgebraic models and hypothesis testing

A statistical model

$$\mathcal{M} := \{P_\theta : \theta \in \Theta\}$$

is *semialgebraic* if its parameter space Θ is a semialgebraic subset of \mathbb{R}^d , i.e., a finite union of sets, referred to as *basic* semialgebraic sets, defined by finitely many polynomial equalities and inequalities.

Semialgebraic statistical models arise frequently in applications. For example, the classical Hardy-Weinberg model for two alleles in equilibrium can be described by a single parameter $\theta \in (0, 1)$, with a parametrization map defined by

$$\phi(\theta) = (\theta^2, 2\theta(1-\theta), (1-\theta)^2),$$

possibly composed with a multinomial map for multiple samples. Alternatively, one may define the model by taking Θ to be the image of ϕ in the probability simplex Δ^2 . In this case, Θ is implicitly defined by the constraint $y^2 - 4xz = 0$, together with the linear constraints that define Δ^2 (namely, $x, y, z \geq 0$ and $x + y + z = 1$), and is thus semialgebraic.

To set notation in a hypothesis testing framework, we consider a model with parameter space $\Theta \subseteq \mathbb{R}^d$ (which need not be semialgebraic) and a semialgebraic submodel with parameter

space $\Theta_0 \subset \Theta$. Following [89], we assume throughout Θ_0 is a basic semialgebraic set. Noting that an equality is equivalent to two inequalities, we assume

$$\Theta_0 := \{\theta \in \mathbb{R}^d : f_i(\theta) \leq 0 \text{ for all } i = 1, \dots, p\}, \quad (4.2.1)$$

where the f_i are polynomials. Given data consisting of n independent and identically distributed (i.i.d.) samples and assuming that

$$X_1, \dots, X_n \sim P_\theta$$

for some $\theta \in \Theta$, we define the null and alternative hypotheses

$$H_0 : \theta \in \Theta_0 \quad \text{and} \quad H_1 : \theta \in \Theta \setminus \Theta_0, \quad (4.2.2)$$

respectively.

4.2.2 Overview of the SDL test

The SDL test uses randomized incomplete U-statistics and a Gaussian multiplier bootstrap approximation of the test distribution to perform hypothesis testing in the setting described in Section 4.2.1. We outline the main objects and steps of the method, focusing on computations. For full justification, see [89].

The kernel function

The incomplete U -statistic is defined using a kernel function to coarsely approximate $f(\theta)$. Let $f : \Theta \rightarrow \mathbb{R}^p$, $f(\theta) := (f_1(\theta), \dots, f_p(\theta))$, where the f_i are the constraint polynomials of Eq. (4.2.1). For some $m \geq 1$, let $h : \mathbb{R}^m \rightarrow \mathbb{R}^p$ be a *kernel function*, i.e., a measurable symmetric function satisfying $\mathbb{E}[h(X_1, \dots, X_m)] = f(\theta)$ for i.i.d. $X_i \sim P_\theta$. Section 4.2.3 gives details about the specific construction of such an h .

The quantity m —called the *order* of the kernel—is a user-specified choice of a subsample size. Given a random subsample X_{i_1}, \dots, X_{i_m} of the data, $h(X_{i_1}, \dots, X_{i_m})$ estimates $f(\theta)$, though perhaps poorly if m is small. The SDL method averages many such estimates to construct a better one: the randomized incomplete U-statistic.

The incomplete U-statistic and the SDL test statistic

Now that we have defined the kernel function, we can define the SDL test statistic. Let $I_{n,m}$ be the set of m -element subsets of $[n] = \{1, 2, \dots, n\}$, viewed as ordered m -tuples,

$$I_{n,m} := \{(i_1, \dots, i_m) \in \mathbb{Z}^m : 1 \leq i_1 < \dots < i_m \leq n\}.$$

Choose a *computational budget parameter* $N \leq \binom{n}{m}$. For each $\iota \in I_{n,m}$ let $Z_\iota \sim \text{Bernoulli}(N/\binom{n}{m})$, and define $\widehat{N} := \sum_{\iota \in I_{n,m}} Z_\iota$. The *randomized incomplete U-statistic* is

$$U'_{n,N} := \frac{1}{\widehat{N}} \sum_{\iota \in I_{n,m}} Z_\iota h(X_\iota), \quad (4.2.3)$$

where $X_\iota := (X_{i_1}, \dots, X_{i_m})$ if $\iota = (i_1, \dots, i_m)$.

The *SDL test statistic*, \mathcal{T} , is the maximum component of the studentization of $U'_{n,N}$:

$$\mathcal{T} := \max_{1 \leq j \leq p} \frac{\sqrt{n} U'_{n,N,j}}{\widehat{\sigma}_j}, \quad (4.2.4)$$

where $\widehat{\sigma}_j^2$ is a stochastic approximation of σ_j^2 , the variance of the j -th coordinate of $U'_{n,N}$ (see Section 4.2.2 for details on the computation of $\widehat{\sigma}_j^2$).

The critical threshold

A large value of \mathcal{T} is interpreted as evidence against H_0 . More precisely, \mathcal{T} is judged using an approximate distribution of a related statistic,

$$\mathcal{T}_c := \max_{1 \leq j \leq p} \frac{\sqrt{n} (U'_{n,N,j} - f_j(\theta))}{\widehat{\sigma}_j}. \quad (4.2.5)$$

Since $\mathbb{E}[U'_{n,N}] = f(\theta)$ for all $\theta \in \Theta_0$, \mathcal{T}_c differs from \mathcal{T} only in centring. Moreover, since the functions f_j are non-positive on the null model, $\mathcal{T} \leq \mathcal{T}_c$ whenever $\theta \in \Theta_0$. Thus, using the distribution of \mathcal{T}_c to assess \mathcal{T} would yield a conservative test. Although the exact distribution of \mathcal{T}_c is unknown, it can be approximated, as we describe next.

Let $U_{n,n_1}^\#$ be the Gaussian multiplier bootstrap of $\sqrt{n}(U'_{n,N} - f(\theta))$ presented in detail in the next section. The bootstrap statistic $U_{n,n_1}^\#$ has two independent sources of randomness: (1) the collection $\mathcal{D}_n = \{X_1, \dots, X_n\} \cup \{Z_\iota : \iota \in I_{n,m}\}$ and (2) a sample from $\binom{n}{m} + n_1$ independent

standard normal random variables

$$R = \{\xi'_l : l \in I_{n,m}\} \cup \{\xi_{i_1} : i_1 \in S_1\},$$

where S_1 is a pre-specified subset of $[n]$ and $n_1 = |S_1|$. Now let

$$W := \max_{1 \leq j \leq p} \frac{U_{n,n_1,j}^\#}{\hat{\sigma}_j}. \quad (4.2.6)$$

To estimate a p -value, we fix a large number A (chosen by the user), and then generate a sequence of random variables $W^{(1)}, \dots, W^{(A)}$ by evaluating W on each of A independent copies of R . The resulting p -value estimate is

$$\hat{p} := \frac{\#\{i \in [A] : W^{(i)} \geq \mathcal{T}\}}{A}.$$

The Gaussian bootstrap approximation.

The above procedure for estimating p -values is justified by [89, Corollary 2.10], which shows that, under technical assumptions, the conditional law of W given \mathcal{D}_n approximates \mathcal{T}_c for large n . As a consequence, the SDL test is asymptotically conservative [89, Corollary 3.1]. Nonetheless, it is important to understand how the approximation of \mathcal{T}_c depends on user-specified test parameters when n is bounded, as this can affect the p -value distribution and hence the statistical properties of the SDL test in practice.

The approximation proceeds in two steps: first the quantity $\sqrt{n}(U'_{n,N} - f(\theta))$ from Section 4.2.2 is approximated by a Gaussian random vector Y , and subsequently Y is approximated by a Gaussian bootstrap $U_{n,n_1}^\#$ defined in this section. By [89, Theorem 2.4], the expression $\sqrt{n}(U'_{n,N} - f(\theta))$ is well approximated asymptotically by the p -variate Gaussian

$$Y \sim \mathcal{N}_p(0, m^2\Gamma_g + \alpha_n\Gamma_h), \quad (4.2.7)$$

with

$$\alpha_n := \frac{n}{N}, \quad \Gamma_h := \text{Cov}[h(X_1, \dots, X_m)], \quad \text{and} \quad \Gamma_g := \text{Cov}[g(X_1)],$$

where $g := \mathbb{E}[h(x, X_2, \dots, X_m)]$ is the Hájek projection of h .

While the covariance matrix $m^2\Gamma_g + \alpha_n\Gamma_h$ of Y is typically unknown, since $Y = mY_g + \sqrt{\alpha_n}Y_h$ for independent $Y_g \sim \mathcal{N}_p(0, \Gamma_g)$ and $Y_h \sim \mathcal{N}_p(0, \Gamma_h)$, an approximation of Y can be obtained from approximating the distribution of these two normal random variables:

Y_h : To approximate Y_h , let $\{\xi'_\iota : \iota \in I_{n,m}\}$ be a collection of independent standard normal variables, and define the multiplier bootstrap

$$U_{n,h}^\# := \frac{1}{\sqrt{\widehat{N}}} \sum_{\iota \in I_{n,m}} \xi'_\iota \sqrt{Z_\iota} (h(X_\iota) - U'_{n,N}).$$

The distribution of $U_{n,h}^\#$ is used to approximate Y_h .

Y_g : Since g is not explicitly known, approximating Y_g is more complicated. Fix some $S_1 \subseteq [n]$ and let $n_1 = |S_1|$. For each $i_1 \in S_1$, partition $[n] \setminus \{i_1\}$ into $K := \lfloor \frac{n-1}{m-1} \rfloor$ disjoint subsets of size $m-1$: $S_{2,1}^{(i_1)}, S_{2,2}^{(i_1)}, \dots, S_{2,K}^{(i_1)}$. For each $i_1 \in S_1$, we estimate $g(X_{i_1})$ using the *divide-and-conquer estimator*

$$G_{i_1} := \frac{1}{K} \sum_{k=1}^K h(X_{i_1}, X_{S_{2,k}^{(i_1)}}).$$

With $\overline{G} := \frac{1}{n_1} \sum_{i_1 \in S_1} G_{i_1}$, define

$$U_{n_1,g}^\# := \frac{1}{\sqrt{n_1}} \sum_{i_1 \in S_1} \xi_{i_1} (G_{i_1} - \overline{G}), \quad (4.2.8)$$

where $\{\xi_{i_1} : i_1 \in S_1\}$ is a collection of n_1 independent standard normal variables. The distribution of $U_{n_1,g}^\#$ is used to approximate Y_g .

Finally, the combined Gaussian bootstrap used to approximate the distribution of Y is

$$U_{n,n_1}^\# := m U_{n_1,g}^\# + \sqrt{\alpha_n} U_{n,h}^\#.$$

Studentization.

For studentization of the statistics \mathcal{T} and W (Eqs. (4.2.4) and (4.2.6)) we estimate $\sigma_1^2, \dots, \sigma_p^2$. From the previous subsection, these can be obtained as $\widehat{\sigma}_j^2 := m^2 \widehat{\sigma}_{g,j}^2 + \alpha_n \widehat{\sigma}_{h,j}^2$, where

$$\widehat{\sigma}_{g,j}^2 := \frac{1}{n_1} \sum_{i_1 \in S_1} (G_{i_1,j} - \overline{G}_j)^2 \quad \text{and} \quad \widehat{\sigma}_{h,j}^2 := \frac{1}{\widehat{N}} \sum_{\iota \in I_{n,m}} Z_\iota (h_j(X_\iota) - U'_{n,N,j})^2.$$

4.2.3 Kernel construction

Now that we have laid out all the components of the SDL test, we discuss particulars about constructing a kernel function that satisfies the requirements of Section 4.2.2. For a semialgebraic

model, the following procedure for constructing a kernel h is suggested in [89, Section 4].

For each polynomial inequality $f_i(\theta) \leq 0$, $i \in \{1, \dots, p\}$ used in defining the model, write

$$f_i(\theta) = a_0 + \sum_{r=1}^s \sum_{\substack{j=(j_1, \dots, j_r) \\ j_i \in \{1, \dots, d\}}} a_j \theta_{j_1} \cdots \theta_{j_r}, \quad (4.2.9)$$

with $a_0, a_j \in \mathbb{R}$. Then the following steps construct a symmetric, unbiased estimator $h_i(X_1, \dots, X_m)$ of $f_i(\theta)$ from independent $X_i \sim P_\theta$, $\theta \in \Theta_0$:

1. For some $\eta \geq 1$, find functions $\hat{\theta}_1, \dots, \hat{\theta}_d : \mathbb{R}^\eta \rightarrow \mathbb{R}$ with $\mathbb{E}[\hat{\theta}_j(X_1, \dots, X_\eta)] = \theta_j$.
2. With $m = \eta \cdot \max_{1 \leq i \leq p} \{\deg(f_i)\}$, an unbiased estimator of $f_i(\theta)$ is $\check{h}_i(X_1, \dots, X_m)$, where

$$\check{h}_i(x_1, \dots, x_m) := a_0 + \sum_{r=1}^s \sum_{j \in J_r} a_j \prod_{z=1}^r \hat{\theta}_{j_z}(x_{(z-1)\eta+1}, x_{(z-1)\eta+2}, \dots, x_{z\eta}).$$

3. With S_m the symmetric group, the components of a symmetric kernel $h : \mathbb{R}^m \rightarrow \mathbb{R}^p$ are given by:

$$h_i(x_1, \dots, x_m) := \frac{1}{m!} \sum_{\pi \in S_m} \check{h}_i(x_{\pi(1)}, \dots, x_{\pi(m)}).$$

Note the symmetrization of step 3 is computationally expensive if $\deg(f_i)$ is large. In Section 4.3.4 we discuss this issue further.

4.2.4 SDL test parameters

Finally, we catalogue the different parameters that are needed for the SDL test, as these parameter choices will be explored in the context of our applications below. In addition to a semialgebraic description of a model, the SDL testing procedure requires four parameter values. They are listed here along with suggested values from [89].

- m : The order m of the kernel h is determined by the constraint degrees and the number of data points η used to estimate the θ_j . For the theoretical analysis of error bounds in [89, Theorem 2.4], it is assumed that $2 \leq m \leq \sqrt{n}$, while the bound itself depends quadratically on m . The authors suggest that m be small, as “larger m imply worse performance of the Gaussian approximation in terms of the required sample size” [89, Remark 2.6].

N : The computational budget parameter N specifies the average number of terms in the randomized incomplete U-statistic. The asymptotic error bounds of [89] require $N/|I_{n,m}| < 1/2$, but choosing $N = \mathcal{O}(n)$ is suggested as the error bounds vanish asymptotically under certain circumstances. Simulations in [89] suggest larger N provides more statistical power, but the authors warn too large an N may cause the test to perform poorly near model irregularities. Ultimately, they offer that $N = 2n$ was reasonable for their model simulations.

n_1 : The parameter n_1 specifies the number of terms used in the sum in Eq. (4.2.8) to estimate Y_g . In [89], a suggested value of $n_1 = n$, the maximum possible, is given so that bootstrap accuracy is maximized.

A : The final parameter, A , governs the number of samples W used in approximating their distribution via bootstrap, with a suggested value of $A = 1000$.

4.3 Trinomial submodels

Here we explore the behaviour of the SDL test on some simple null semialgebraic models that arise when considering the coalescent model in phylogenomics. Their small size, in terms of dimension, allows for rejection regions to be visualized and compared to those from other methodologies.

4.3.1 Basic examples




Figure 4.1: Parameter spaces (blue line segments) of four submodels of the trinomial model, with parameter space Δ^2 . The submodels capture the form of the quartet Concordance Factor if the species relationships have specific features, as described in the text.

Our first four example models are depicted inside the 2-simplex in Fig. 4.1. Each characterizes the frequencies of the three possible quartet gene tree topologies if four species are related by a tree or network with certain features. Appendix A.1 provides a more complete explanation, but knowledge of the application is not necessary for a reader primarily interested in the SDL test for other uses.

While each model is composed of line segments, they exhibit a variety of geometric features that may affect testing behaviours. Model 1 is regular. Model 2 has a boundary point in the interior of the simplex, causing a discontinuity in the asymptotic distribution of standard statistics. Model 3 has no boundary points in the simplex but exhibits a singularity (in the sense of algebraic geometry) at the centroid, where 3 lines cross. Again this causes a discontinuity in the asymptotic distribution, with slow convergence to it for parameters near the centroid. In Model 4 the centroid is both a singularity and a boundary of each of the component lines.

Because of their importance for testing whether biological data shows evidence for specific species relationships involving hybridization or other lateral gene flow, specialized test distributions for null hypotheses of Model 2 and 4 are derived in [76] and for Model 3 in [10]. Tests using these are implemented in the R package `MSCquartets` [81]. These improve on a naive use of a standard distribution such as a χ^2 that ignores the singularities and boundaries of the models. Model 1, of course, can be tested with a standard approach, as it lacks any irregularities. Thus for all these models we can compare SDL test behaviour to the behaviour of deterministic tests.

We also consider several other semialgebraic trinomial submodels that we do not depict here. These are the Hardy-Weinberg equilibrium model for 2 alleles (a regular model for which good deterministic test methods are established) and two artificial models chosen because of their specific algebraic nodal and cuspidal singularities. For these last two models we know of no other methods addressing their singularities, but they nonetheless illustrate important issues that may arise with general semialgebraic models.

Appendix A.2 presents rejection regions using current deterministic testing procedures for the null Models 1-4, as well as for the Hardy-Weinberg model, for a dataset of size 300.

4.3.2 Semialgebraic descriptions of trinomial models

Each of the models depicted in Fig. 4.1 is easily given a semialgebraic description. With the parameter space Θ for each of the models viewed as a subset of $\Delta^2 \subseteq \mathbb{R}^3$, we use coordinates (x, y, z) , with $x + y + z = 1$, $x, y, z \geq 0$, for simplex points.

Model 1 $y - z = 0$.

Model 2 $y - z = 0, \quad 1/3 - x \leq 0$.

Model 3 $(y - z)(x - y)(x - z) = 0$.

Model 4 $(x - y)(x - z)(y - z) = 0, \quad (x - z)^2(y - z)^2(1/3 - x) \leq 0,$
 $(x - y)^2(y - z)^2(1/3 - y) \leq 0, \quad (x - y)^2(x - z)^2(1/3 - z) \leq 0.$

Note that other semialgebraic descriptions of these models exist, and although these are ‘simple’ ones, we have no well-defined notion of a ‘simplest description’ in general. For instance, the linear inequality given above in the description of Model 2 could be replaced by others and the effect of changing this description is one issue with the SDL test that we investigate in Section 4.3.4.

4.3.3 SDL rejection regions for trinomial submodels

One way to understand a hypothesis test is through its rejection region at various test levels. For the models above, we considered all possible datasets (up to ordering) of size $n = 300$, that is all collections of 300 vectors each of which is a standard basis vector in \mathbb{R}^3 . The counts of the 3 basis vectors in such a dataset are then normalized (i.e., the mean of the vectors is computed) to give a point in the simplex. Applying the SDL test for a model to the dataset, this point can be coloured according to the dataset’s p -value, indicating rejection at various levels.

Note that rejection is based on the incomplete U-statistic of the data, which includes randomness, and the test distribution, which also includes randomness. Thus rejection region plots produced in this way may vary even though they are testing identical “data” and there is no well-defined “rejection region” in the simplex. Nonetheless, such plots, and the stochastic variation they show, give helpful insight into test behaviour.

In Section 4.3.4 we follow this procedure to colour the simplex for various models using nominal test levels of 0.10, 0.05, and 0.01 to delineate between purple, blue, green, and red colourings. Throughout, we use datasets of size $n = 300$. This size was chosen so that the rejection region plots were not overly pixelated, yet easily interpretable visually, since for very large n the size of the fail-to-reject region shrinks tightly around the model line segments.

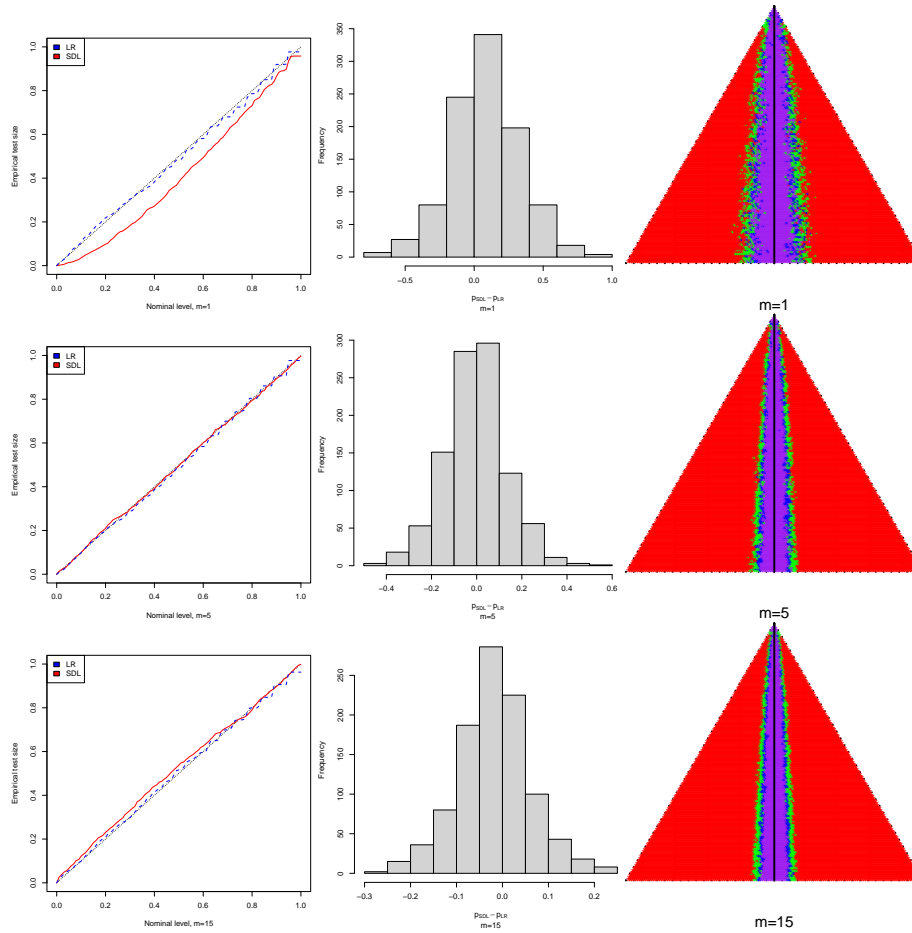


Figure 4.2: SDL test behaviour for Model 1, with $m = 1, 5$, and 15 (top row to bottom). The left column shows nominal vs. empirical sizes for the SDL and LR tests; the middle, histograms of p -value differences; and the right, SDL rejection regions.

4.3.4 The SDL test of trinomial submodels

For datasets of size $n = 300$, we fix parameters of the SDL test to $N = 1000$, $A = 1000$, and $n_1 = n = 300$ throughout, but vary m . Our values of A and n_1 follow suggestions of [89], since we only observed noticeable changes in performance with extreme variations from suggested values. Varying N or m has more impact. However, we found increasing N to 1000 reduced the randomness in our p -values and, with appropriately chosen m , still allowed us to ensure our tests were conservative. We therefore only vary m as [89] already illustrated the effects of varying N .

The order m of the kernel

As constructed in Section 4.2.3, the kernel function h depends on $m = \eta \cdot \max_i \{\deg(f_i)\}$ data points, with η the number per scalar parameter. While [89] suggests that m should be chosen to be small, we experimented with different choices of η and found that choosing a minimal value was generally *not* optimal as it could lead to both lower statistical power and increased stochasticity of the SDL p -values.

This conclusion is illustrated in Figure 4.2, which compares SDL p -values from Model 1 (which has only regular points) using $m = 1, 5$, and 15 (top to bottom), along with p -values from the standard Likelihood Ratio (LR) test. We estimated parameters with the mean of $\eta = m$ data values. The left column of Figure 4.2 compares the nominal level versus empirical test sizes of the SDL test (red) and LR test (blue) from 1000 simulated datasets of size 300, with model parameters $(1/3, 1/3, 1/3)$. The middle column histograms show the differences between the approximate p -values of the SDL test and the p -values computed with LR, for the same 1000 datasets. The right column depicts the SDL rejection region for all datasets of size $n = 300$.

Importantly, Figure 4.2 illustrates the danger of choosing m too large, since it impacts the conservativity of the SDL test. For $\eta = m = 1$, the test is highly conservative (top left), with SDL p -values tending to be larger than LR p -values (positive histogram mean). At $m = 5$, the test retained an acceptable size (middle left), and additional simulations with other parameters (not shown) indicate that $m = 5$ was a uniformly good choice. On the other hand, $m > 5$ resulted in invalid tests with an excess of small p -values. Figure 4.2 illustrates this for $m = 15$, with the leftmost plot exhibiting for most levels an excess in the test size, and the histogram a negative mean.

On the other hand, choosing $\eta = m$ very small (e.g., $m = 1$) is also suboptimal. For $m = 1$, the rejection region plot (top right) has a smaller rejection region than for the LR test (shown in Fig. A.2 of Appendix A.2), and its p -values exhibit substantial random variability. By contrast, increasing m had the benefit of increasing both the size of the rejection region and the precision of the SDL p -values (right column), with the latter observation also evident in the histograms, which concentrate with larger m . To quantify this, we also compute the variance of the SDL p -values from 100 test applications for each of 100 simulated datasets, and observed a decrease from 0.068 for $(m = 1)$ to 0.030 for $(m = 5)$.

We note that while choosing m minimally gives a conservative test here, in our examples below, and in [89], there are no theoretical assurances that this will be the case for all models. Regardless, varying m in the models we explored suggests a clear tradeoff between increasing m to reduce the stochasticity of p -values and type II errors, and keeping m small to reduce type

I errors. However, the value of m at which the test size exceeded the nominal level is dependent on the specific model, constraints used to describe it, and the model parameter $\theta \in \Theta_0$, and we were unable to develop any general rules to apply. Simulation at a number of model points seems to be the most informative approach.

In the following subsections, we use the largest m which simulations suggest gives a valid test size at a number of model points, including singularities and boundary points. For instance, we find that for Model 2 (discussed in the next subsection) $m = 5$ gave good performance for the boundary parameter point $(1/3, 1/3, 1/3)$, with empirical test size closely tracking the nominal level (plot not shown). However, for parameters $(2/3, 1/6, 1/6)$, this choice of m gives a conservative test for Model 2, and $m = 20$ gives a more powerful yet valid test at that point. We nonetheless consider $m = 5$ for Model 2 as the better choice overall.

Choice of model constraints

Semialgebraic models may have many different semialgebraic descriptions in which the polynomial equalities and inequalities differ. The choice of specific model constraints can impact the shape of the rejection region for the SDL test.




Figure 4.3: Rejection regions for Model 2 under the SDL test using (L to R) a) the constraints $y - z \leq 0$, $z - y \leq 0$, and $1/3 - x \leq 0$; b) replacing the last inequality by $2/3 - x - y \leq 0$; c) including $r = 10$ random convex combinations of the inequalities of (a) ; and d) including $r = 100$ random convex combinations.

For Model 2 with $m = 5$ data points in the kernel function, we illustrate this in Fig. 4.3. On the left we use the constraints given in the previous section. Note that the ‘flat bottom’ of the purple region reflects the horizontal boundary from the constraint $1/3 - x \leq 0$.

For the next plot in Fig. 4.3 the inequality $1/3 - x \leq 0$ is replaced by $2/3 - x - y \leq 0$, giving a different description of the same model. Again the shape of the rejection region reflects

the choice of the constraint. While both of these regions are valid in the sense of ensuring an acceptable rejection rate for data generated by the model, the fact that an arbitrary choice of constraints determines the shape of the rejection region is undesirable.

To be agnostic in terms of semi-algebraic description, it would be preferable to simultaneously use all possible constraints for the model. But by including only a small number of additional model constraints in a redundant model description, we found we could approximate that situation for Model 2.

In particular, after first converting the equality constraint to two inequalities, we created 10 and then 100 random convex combinations of the original three inequalities and included them in the SDL procedure. This gave the two rightmost plots in Fig. 4.3, with ‘rounded’ bottoms, approximately reflecting all the linear constraints that might be used to truncate the model line at the centroid of the simplex. Using more random combinations more consistently smooths the boundary, but at additional computational cost.

For this example, with a complete geometric view of the model in the ambient simplex, we could have chosen fewer specific combinations for the same effect. In more general settings, however, choosing randomly has the advantage of not requiring any detailed geometric understanding of the model.

However, it may be necessary to use many such combinations, especially when the model’s co-dimension is large. For a simple example, a model that is a half-line in a d -dimensional simplex is minimally described by $d - 1$ linear equalities and 1 linear inequality, or $2d - 1$ inequalities. Rejection region boundaries using such a set of constraints form a roughly polyhedral cylinder with opposite sides approximately parallel (due to the equality constraints), which is cut off by a hyperplane (from the inequality). If d is large, an adequate number of combinations to approximate a full set of inequality constraints might be quite large, but would give a rounder boundary.

For the work that follows, we introduce a new parameter, r , indicating the number of random convex combinations of the model’s specified inequality constraints to include as new constraints in the SDL testing procedure. By ‘random’ we mean that if the model is specified by D inequality constraints then the convex sum weight w for each new constraint is an independent random variable $w \sim \text{Dirichlet}(D; 1, 1, \dots, 1)$, meaning that w is drawn uniformly from Δ^{D-1} .

In Section 4.3.4 we consider a more complex situations in which supplying additional redundant constraints may be desirable.

Symmetrizing the kernel

As described in Section 4.2.3, we construct our kernel function h of m data points from the semi-algebraic model constraints by a process including symmetrization. Then the symmetrization occurs over the symmetric group \mathcal{S}_m .

For general semialgebraic models there is no upper bound on the degree of defining constraints, so even if η may be chosen to be small, $m = \eta \max_i \deg f_i$ may be large. Moreover, as was discussed in Section 4.3.4, performance of the method is sometimes improved by choosing η larger than its theoretical minimum. Thus m may be large in practice, and a full symmetrization may not be computationally feasible.

To investigate situations in which symmetrization of the kernel by summing over all data permutations is not feasible, we focus on Model 3 with $\eta = 5$ so $m = 15$. (Since this model is defined by a single equality constraint, convex combinations of the resulting inequalities would have no effect.) From the construction of the kernel we already have symmetry within the 5-element blocks of data points which are averaged to estimate each parameter. Thus full symmetrization would only require

$$\frac{15!}{(5!)^3} \approx 7.5 \times 10^5$$

permutations, though this is already computationally excessive. We therefore explore summing only over a relatively small number, s , of permutations, chosen uniformly at random. We sample these permutations anew each time the kernel must be evaluated, both for computing the test statistic and for estimating the distribution by which it is judged.




Figure 4.4: Rejection regions for Model 3 under the SDL test using (L to R) $s = 1, 10,$ and 100 random permutations to partially symmetrize \check{h} . For all, $m = 15$.

In Fig. 4.4 (left) we see that even a single ($s = 1$) random permutation produces an appropriately symmetric rejection region, though that region is quite small. With even $s = 10$ permutations used (middle) the rejection region grows considerably. This trend continues through $s = 100$ permutations, although the gain between these last two is not large.

While our explorations indicate that this random partial symmetrization scheme can be effective, theory justifying its use is currently lacking. The incomplete U-statistics already incorporate two sources of randomness — the data and the subsampling/bootstrapping of the test procedure — and random partial symmetrization brings in a third which is not considered in [89]. Moreover, our simulations are all low-dimensional and we did not explore thoroughly how increasing dimension may affect the number of random permutations needed. While in Section 4.4 we explore one higher dimensional case, extension of the underlying theory of the SDL test is needed.

Irreducible components and an intersection-union test

Some natural semialgebraic statistical models are formed as the union of several components, such as the intersecting line segments that comprise Models 3 and 4. More specifically, in algebro-geometric terms, a model may be Zariski-dense in a variety with several irreducible components. Although for these examples the irreducible components are simply lines, more generally irreducible components may be of higher degree but will have degree at most that of the full model. Computational algebra software can be used to calculate equality constraints of the components.

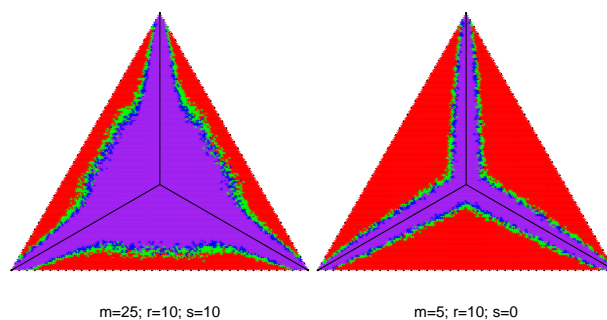


Figure 4.5: (L) Rejection region for Model 4 obtained from SDL test using semialgebraic description given above. (R) Rejection region for an Intersection-Union test using the SDL tests for the 3 irreducible components of Model 4 (each essentially Model 2).

In addition to performing the SDL test directly for Model 4 using the constraints given above, we performed an intersection-union test by applying the SDL test to each irreducible component, rejecting the full null hypothesis if we reject it for each of the component null hypotheses. Thus we take the maximum of the p -values from the irreducible component tests as an overall p -value.

Fig. 4.5 shows comparison plots for Model 4, using the standard SDL test and the intersection-union variant. In both we used $\eta = 5$ data points to estimate individual model parameters, giving $m = 25$ and 5, respectively, due to the different degrees of the constraints. Note the intersection-union test led to both a larger rejection region and less randomness in its boundary. Indeed, the direct SDL test for Model 4 remained conservative for all values of m we tried (up to 45) and in particular the null hypothesis was never rejected in a very large central region of the simplex. In addition to having much greater power, the intersection-union test was faster to compute, and showed less random behaviour.

Model 3 can similarly be decomposed, with an SDL intersection-union test showing better performance than was obtained in Section 4.3.4. We suspect that similar gains can be achieved for other reducible models

Higher degree irreducible models

As seen for Models 3 and 4, the degree of the model's constraints seems to affect the power of the test, particularly around singularities, but somewhat for points far from these. If the model can be decomposed into irreducible components of lower degree, an intersection-union approach may ameliorate the behaviour. To investigate the effect of degree further, we considered several irreducible models of degree 2 and 3.




Figure 4.6: Rejection regions for SDL tests of (L-R) (a) the Hardy-Weinberg 2-allele model defined by $y^2 - 4xz = 0$, (b) a nodal cubic model defined by $(y - 1/3)^2 - 6(x - 2/5)^2(x - 1/9) = 0$, (c) a cuspidal cubic model, defined by $(y - 1/3)^2 - (x - 1/3)^3 = 0$.

The Hardy-Weinberg 2-allele model, whose SDL rejection region is shown in Fig. 4.6(a), is a quadratic model with no irregularities. The rejection region for $n = 300$ is close to that for the standard chi-squared test of the model (Fig. A.2 of Appendix A.2) with the added stochastic variation inherent in uses of SDL. Note the low value of $m = 4$ here; higher values produced excesses in small p -values

Fig. 4.6(b) shows results for a nodal cubic model (chosen for its degree and geometry rather than any application) with a single crossing singularity. The higher degree seems to result in both less power than seen in previous models, and more stochastic variation at the boundary of the rejection region, at least for the same choices of test parameters used for previous models.

In Fig. 4.6(c) the SDL test is applied to a cuspidal cubic model. Note the large region (extending downward and right from the cusp) on which the test fails to reject the model. In that region the equality constraint is nearly met, with the polynomial taking on small values, resulting in an inability of the SDL approach to reject the model. This is an important feature to note, since it shows that a minimal set of model constraints may fail to adequately distinguish between points on the model and some off the model for an SDL test.




Figure 4.7: Rejection regions for SDL tests of the cuspidal cubic (L-R) with (a) constraints supplemented by $1/3 - x \leq 0$; (b) constraints supplemented by the inequality from (a) plus $r = 10$ random convex combinations of inequalities, and (c) constraints supplemented by 3 linear inequalities as described in the text and $r = 10$ random convex combinations of inequalities.

In Fig. 4.7(a) we see that adding a single linear inequality which is satisfied on the model expands the rejection region, and increases the test's power. This further reinforces the point of Section 4.3.4 that ideally one would use all semialgebraic constraints satisfied on the model. However, the linear constraint introduced here is not suggested by the model's defining equation, and it is unclear how one might determine a good finite set of supplementary constraints in an automated way. Through human agency, doing so would be facilitated by a thorough

understanding of the model geometry, but particularly in high dimensional settings that may be difficult to obtain.

Fig. 4.7(b), which uses the same inequality constraint as in (a), illustrates an instance of the random convex combination approach of Section 4.3.4 failing to have much impact. For Fig. 4.7(c) we included two additional linear inequalities, with bounding lines stretching from the cusp to the points at which the model intersects the simplex boundary. These improve performance, though note the slight bulge in the non-rejection region to the right of the cusp. Adding additional non-linear constraints, with appropriate concavity, can remove this bulge, though such an approach is *ad hoc*.

The conservative nature of the SDL test near model singularities may be partially explained by the vanishing of the gradients of the equality constraints at such points. This implies the constraints will be nearly satisfied at nearby points off the model, and (if there are only equality constraints) the incomplete U-statistics may be close to 0 as well. Notice this is quite different from the behaviour at non-singular boundary points of a model as in Section 4.3.4.

Finally, note in Fig. 4.7(a) the reduced stochasticity of the rejection region boundary for the linear constraint vs. the cubic. This suggests that using low degree constraints (when possible) is preferable.

4.4 Hypothesis Testing and Inference of Phylogenetic Trees

We next explore the performance of the SDL method for testing and inference of phylogenetic tree topologies through *phylogenetic invariants*. Introduced in [31, 67], phylogenetic invariants are polynomials vanishing on pattern distributions in genetic alignments. They have been widely studied and used to establish parameter identifiability for various models [e.g., 1, 90, 2, 9, 41, 24], and underlie several inference methods [23, 37, 48, 5, 27]. (See [91] for a general introduction.) Viewing invariants as equality constraints on the data distribution, the SDL method offers a new statistical approach for their use.

4.4.1 The CFN model and its semialgebraic descriptions

We focus on the Cavender-Farris-Neyman (CFN) model for 2-state sequence evolution on 4-leaf binary trees [84, Chapter 8]), a higher-dimensional model than those considered in previous sections. The two states 0, 1 usually represent purines (A,G) and pyrimidines (C,T) in DNA sequences.




Figure 4.8: The 4-leaf binary tree topologies, with edge lengths t_i . The names $T_{x|yzw}$ indicate the partition of leaves induced by the central edge.

Let T be one of the leaf-labeled trees of Fig. 4.8. Arbitrarily introducing a tree root representing the common ancestor of 1, 2, 3, 4, the CFN base substitution process on an edge of length t is given by a two-state, continuous-time, time-reversible Markov chain, with equal state transition rates and expected number of transitions t , proceeding from the parent to child node. Time reversibility ensures this model is independent of root location.

The CFN model on T is the marginal distribution of states on leaves, as internal tree states are hidden, represented by the $2 \times 2 \times 2 \times 2$ tensor

$$p = (p_{ijkl})_{i,j,k,l \in \{0,1\}}, \quad p_{ijkl} = \mathbb{P}[X_1 = i, X_2 = j, X_3 = k, X_4 = l],$$

where X_i is the state at leaf i . This models a single site in a sequence alignment arising from tree T , with all aligned sites viewed as i.i.d. samples. Fixing the topology of T , but varying edge lengths gives a parametrized family of models on T . Reparametrizing with $\theta_i := e^{-2t_i}$ gives this family a polynomial parameterization:

$$\phi_T : (0, 1]^5 \rightarrow \Delta^{15} \subset \mathbb{R}^{16}.$$

By the term *CFN model* on a topological tree T , we mean the parametrized family of statistical models given by the image \mathcal{M} of this map. As the polynomial image of a semialgebraic set, \mathcal{M} is a semialgebraic subset of Δ^{15} . The polynomials vanishing on this set, and thus all polynomial equality constraints for the model, form an ideal I_T , which can be computed using Gröbner basis techniques with computational algebra software such as Macaulay2 [49].

The set of points on which the polynomials in I_T vanish form an algebraic variety $V_T \supset \mathcal{M}$. Both V_T and \mathcal{M} are of dimension 5, matching the number of numerical edge length parameters on T . I_T is finitely generated, and any choice of generators gives sufficient equality constraints to define V_T .

For $T = T_{12|34}$, one set of defining equations for V_T is the 2 quadratic constraints:

$$f_1 = \det \begin{pmatrix} q_{0000} & q_{0011} \\ q_{1100} & q_{1111} \end{pmatrix} = 0, \quad f_2 = \det \begin{pmatrix} q_{0101} & q_{1001} \\ q_{0110} & q_{1010} \end{pmatrix} = 0, \quad (4.4.1)$$

where $q_{0000} := p_{0000} + p_{0001} + p_{0010} + p_{0011} + p_{0100} + p_{0101} + p_{0110} + p_{0111}$,
 $q_{1111} := p_{0000} - p_{0001} - p_{0010} + p_{0011} - p_{0100} + p_{0101} + p_{0110} - p_{0111}$,
 $q_{0011} := p_{0000} - p_{0001} - p_{0010} + p_{0011} + p_{0100} - p_{0101} - p_{0110} + p_{0111}$,
 $q_{1100} := p_{0000} + p_{0001} + p_{0010} + p_{0011} - p_{0100} - p_{0101} - p_{0110} - p_{0111}$,
 $q_{0101} := p_{0000} + p_{0001} - p_{0010} - p_{0011} + p_{0100} + p_{0101} - p_{0110} - p_{0111}$,
 $q_{0110} := p_{0000} - p_{0001} + p_{0010} - p_{0011} - p_{0100} + p_{0101} - p_{0110} + p_{0111}$,
 $q_{0110} := p_{0000} + p_{0001} - p_{0010} - p_{0011} - p_{0100} - p_{0101} + p_{0110} + p_{0111}$, and
 $q_{1001} := p_{0000} - p_{0001} + p_{0010} - p_{0011} + p_{0100} - p_{0101} + p_{0110} - p_{0111}$,

along with the 9 linear equations:

$$\begin{aligned} (p_{0000} + p_{0001} + p_{0010} + \dots + p_{1111}) - 1 &= 0, \\ p_{0000} - p_{1111} &= 0, \quad p_{0001} - p_{1110} = 0, \quad p_{0010} - p_{1101} = 0, \quad p_{0011} - p_{1100} = 0, \\ p_{0100} - p_{1011} &= 0, \quad p_{0101} - p_{1010} = 0, \quad p_{0110} - p_{1001} = 0, \quad p_{0111} - p_{1000} = 0. \end{aligned} \quad (4.4.2)$$

The linear polynomials are *model invariants*, since they are zero for any of the 3 topological trees, and the quadratics are *topology invariants*, as they are not zero for some tree [28].

Computation shows (see Appendix A) that V_T 's singularities are

$$(V_T)_{\text{sing}} = \{ \phi_T(\theta_1, \dots, \theta_5) : \theta_1, \dots, \theta_5 \in [0, 1] \text{ and } \theta_1 = \theta_2 = 0 \text{ or } \theta_3 = \theta_4 = 0 \text{ or } \theta_5 = 0 \}.$$

Since $\theta_i = 0$ corresponds to $t_i = \infty$, which produce sequence data that is uncorrelated at the ends of an edge, such singularities are unlikely to be relevant to empirical analyses.

For the stochastic model $\theta_i \in (0, 1]$, one finds $\mathcal{M} \subsetneq V_T \cap \Delta^{15}$, but imposing additional polynomial inequalities restricts from V_T to \mathcal{M} [71, 65]. In particular, the quadratic inequality

$$q_{0101}q_{1010} + q_{1001}q_{0110} - 2(q_{0011}q_{1100}) \leq 0 \quad (4.4.3)$$

expresses $t_5 \geq 0$, with similar inequalities for the pendant edges. We consider only the inequality in Eq. (4.4.3), as it is the only one that changes for different tree topologies.

While Eqs. (4.4.1) and (4.4.2) give one set of equality constraints for V_T , others are equally natural. We say that a topology invariant $F \in I_T$ is *partially distinguishing* if there exists a tree $T' \neq T$ on the same taxa such that $F \in I_{T'}$, as well. If F is not partially distinguishing, we say

that it is *completely distinguishing*. We consider the following five specific choices of quadratic topology invariants that, together with the linear invariants, generate I_T . Explicit formulas are given in Appendix A.3.1.

- (CDD) *Completely Distinguishing Determinantal*: These are derived from the determinantal polynomials in Eq. (4.4.1) together with Eq. (4.4.2) (see Appendix A.3.1 for the explicit construction).
- (PDR) *Partially Distinguishing Rank*: These constraints are indirectly obtained from 3×3 minors of a certain flattening of the tensor p described in [9].
- (PDM) *Partially Distinguishing Minimal*: This is a minimal basis obtained by applying the mingens function of Macaulay2 to the kernel of ϕ_T .
- (CDR) *Completely Distinguishing Rank*: These two polynomials are the sum and difference of the polynomials of PDR).
- (CDM) *Completely Distinguishing Minimal*: These two polynomials are the sum and difference of the invariants of PDM.

4.4.2 Data simulation

To evaluate the SDL test on the CFN model, we focused on datasets from the trees studied in [56], shown in Fig. 4.9 (left), where tree $T_{12|34}$ has edge lengths $t_1 = t_3 = a$ and $t_2 = t_4 = t_5 = b$, for varying $a, b > 0$.

A dataset consists of n independent samples drawn from the multinomial distribution with parameter

$$\bar{p} = (\bar{p}_{xxx}, \bar{p}_{xxy}, \bar{p}_{xyx}, \bar{p}_{xyy}, \bar{p}_{yxx}, \bar{p}_{yxy}, \bar{p}_{yyx}, \bar{p}_{yyy}) \in \Delta^7, \quad (4.4.4)$$

where x, y represent distinct states in $\{0, 1\}$, and the coordinates of \bar{p} are $\bar{p}_{xxx} = p_{0000} + p_{1111}$, $\bar{p}_{xxy} = p_{0001} + p_{1110}$, and so forth, where $p = \phi_T(\theta_1(a), \theta_2(b), \theta_3(a), \theta_4(b), \theta_5(b))$ and $\theta_i(t) = e^{-2t}$. We thus assume *a priori* that the linear constraints of Eq. (4.4.2) hold, allowing us to reduce the length of the data vector of length 16 to 8, and subsequently ignore those equalities.

We consider two collections of datasets:

1. *Collection 1*. We generated 30 datasets of size $n = 1000$ site samples for each pair of parameters (a, b) with a and b ranging from 0 to 1.2 in increments of 0.05.




Figure 4.9: Left: The tree $T_{12|34}$ with edge lengths $t_1 = t_3 = a$ and $t_2 = t_4 = t_5 = b$, in units of expected number of substitutions per site. Right: The tree space, with a, b varying from 0 to 1.2. In red, nine parameter pairs with $a, b \in \{0.05, 0.2, 0.8\}$. The dashed blue curve is the lower boundary of the Felsenstein zone, defined by $\theta(b)^2 - 2\theta(a) + \theta(a)^2 > 0$ for $\theta(t) = e^{-2t}$ [46].

2. *Collection 2.* We selected nine parameter pairs to be analysed in greater detail, with $a, b \in \{0.05, 0.2, 0.8\}$. We generated 1000 datasets for each choice of parameters, with each dataset consisting of $n = 1000$ site samples.

Collection 1 samples from throughout the tree space of Fig. 4.9 (right). The upper left region is the “Felsenstein zone,” leading to datasets susceptible to long branch attraction, which makes accurate tree inference by standard methods difficult [46, 52]. The nine parameter choices underlying Collection 2 are indicated in red dots in the figure.

4.4.3 SDL test parameters and hypotheses

To apply the SDL test we must choose its test parameters, m, n_1, N, A as well as a partial symmetrization level s . For our data sets of size $n = 1000$, preliminary investigations led us to use

$$m = 12, \quad N = 1000, \quad n_1 = 80, \quad A = 5000, \quad s = 100.$$

Large values of n_1, N, s lead to substantial computation, but the values above gave a good balance between performance and runtime. For example, no major impact on the test results was observed compared with $n_1 = 500$ and $N = 5000$.

We consider each of the five different sets of quadratic equality constraints presented in Section 4.4.1. We also increased the number of polynomial constraints by adding $r = 20$ random convex combinations of the original ones.

We denote by $H_{12|34}$ the hypothesis that the true tree topology is $T_{12|34}$, and similarly $H_{13|24}$ and $H_{14|23}$. Constraints for tests of $H_{13|24}$ and $H_{14|23}$ can be found by permuting taxon labels from those for $H_{12|34}$, and are given in Ap. Since our simulated data is always sampled from a $T_{12|34}$ tree, in our experiments $H_{12|34}$ is always the true hypothesis and the other two are false.

4.4.4 Hypothesis tests results

We compute p -values from simulated data to test several different null hypotheses.

Collection 1

As an initial exploration of the behaviour of the SDL test, we examined the distribution of all p -values from Collection 1 for each of the three hypotheses $H_{12|34}$, $H_{13|24}$, and $H_{14|23}$. Aggregating p -values across a wide range of parameter values (a , b) in a single histogram gives insight into the overall behaviour of the test.

Since varying the model constraints can affect test behaviour (Section 4.3.4), we created histograms for the five sets of quadratic equality constraints in Section 4.4.1. No other constraints, including Eq. (4.4.3), were used. For each constraint set, we also added $r = 20$ random convex combinations of the resulting inequalities.

Fig. 4.10 presents aggregated p -values for each of 4 conditions (CDM and PDM, $r = 0$ and 20), for the true null hypothesis $H_{12|34}$ (left) and a false null hypothesis $H_{13|24}$ (right). (See Appendix A.3.2 for all five constraint sets.)

For $r = 0$ and the true $H_{12|34}$, the PDM set shows anti-conservative behaviour, with an excess of small p -values. The CDM constraints, on the other hand, shows conservative behaviour, with an excess of large p -values. For the false $H_{13|24}$, the CDM constraints gave a greater concentration of p -values near zero compared to the PDM constraints, suggesting greater power.

Increasing r did not substantially change the behaviour of the test with the CDM set. However, for the PDM set, increasing r had two important and beneficial effects: first, it decreased the number of small p -values when testing $H_{12|34}$, and second, it increased the number of small p -values when testing $H_{13|23}$. This suggests for PDM, the addition of convex combination constraints simultaneously made the test more conservative as well as increased

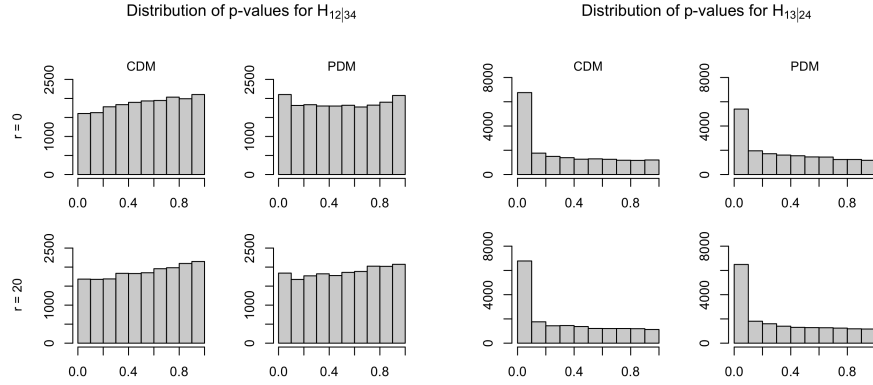


Figure 4.10: Aggregated p -values for a test of the true null hypothesis $H_{12|34}$ (left) and a false null hypothesis $H_{13|24}$ (right) for datasets in Collection 1. Constraints sets CDM and PDM, and number of convex combinations $r = 0$ and 20 are varied.

its statistical power. However, the effect of adding convex combinations constraints highly depends on the choice of starting constraints, as we discuss in Appendix A.3.2.

Although these effects of increasing r might appear relatively small, they are based on aggregated p -values from a large parameter regime, and it is possible specific regions of the parameter space might exhibit more substantial effects. In Appendix A.3.3 we show this is the case, by analysing a particular choice of parameters (a, b) within the Felsenstein zone (a region of particular interest for the phylogenetics community).

Collection 2

We next examine the performance of the SDL test more closely, for the 9 particular edge parameters shown in Fig. 4.9. Fig. 4.11 shows histograms of 1000 p -values, with the test differing only in use of the topology constraints CDM and PDM; in both cases the internal branch inequality Eq. (4.4.3) is not used.

Despite this seemingly small difference, the SDL test with the CDM polynomials tends to be both more conservative and more powerful than when compared to the PDM polynomials. Fig. 4.11 illustrates that when testing the true hypothesis $H_{12|34}$ PDM is more likely have p -values close to zero for 8 out of the 9 choices of model parameters. On the other hand, when testing the false $H_{13|24}$, both CDM and PDM constraints produce small p -values for $a, b \in \{0.05, 0.2\}$. However, for 4 of the remaining 5 choices for (a, b) , the test utilizing the CDM constraints gave small p -values for the incorrect null hypothesis substantially more often than the test utilizing the PDM constraints. Results for $H_{14|23}$ (not shown) were similar to those for $H_{13|24}$.

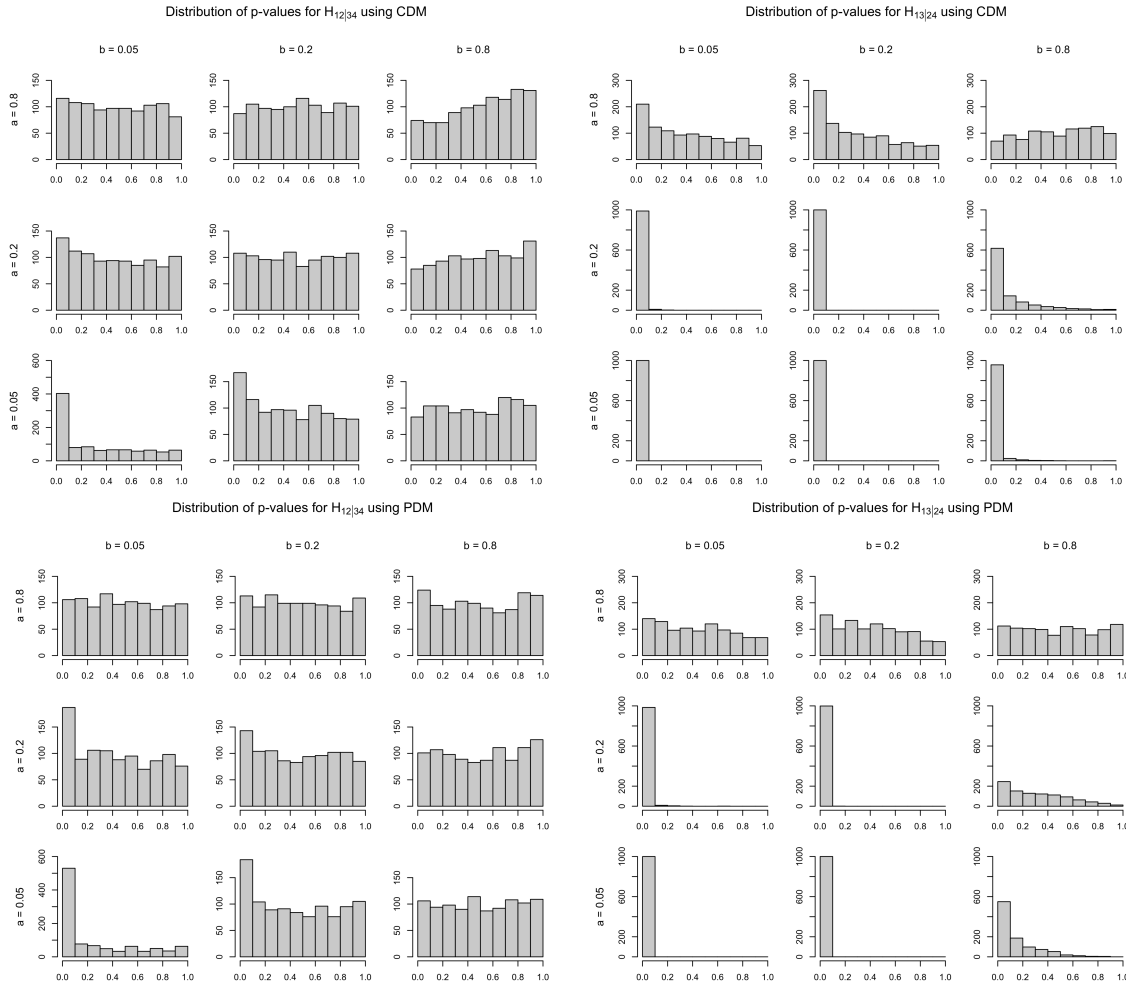


Figure 4.11: p -values obtained from the SDL test on Collection 2 for different constraint sets: CDM (top 3 rows) and PDM (bottom 3 rows). The hypotheses tested are $H_{12|34}$ (left 3 columns) and $H_{13|23}$ (right 3 columns), with $r = 0$.

The SDL test performed quite poorly when testing the correct model hypothesis $H_{12|34}$ for trees with short edge parameters. For example, when $(a, b) = (0.05, 0.05)$ the test produced far too many small p -values, regardless of whether the polynomials were the CDM or PDM sets, though worse for the second.

Effect of internal edge constraint

We also investigated the effect on the SDL test of augmenting the CDD set with the inequality of Eq. (4.4.3), expressing that the tree's internal edge length is non-negative. Fig. 4.12 compares the

distribution of p -values using the CDD generating set with the internal edge inequality verses without it, amalgamating all test results for Collection 1 on a true $H_{12|34}$ and false $H_{13|24}$ null hypothesis. Tests of the false $H_{14|23}$ were similar, and are omitted.

Based on the aggregated p -values, including the internal edge inequality appears to make the test more conservative, with no appreciable change in power. These results were essentially unchanged for $r = 10$ and 20 . An analogous analysis (not shown) considered test results for the datasets of Collection 2, not amalgamating over different parameters. We observed a similar behaviour as in Fig. 4.12: Testing $H_{12|34}$ gave an increase in the number of large p -values and a decrease in the number of very small p -values. In fact, for 8 of the 9 parameters, there was a reduction in the number of p -values less than 0.1 , with the exceptional case, $(a, b) = (0.2, 0.2)$, showing no difference.

This effect of adding a constraint may seem counter-intuitive. By further restricting the model, one might think the test would be more inclined to reject a true hypothesis $H_{12|34}$. Indeed, the test statistic \mathcal{T} is defined in Eq. (4.2.4) as a maximum over all constraints, so an additional constraint can only lead to larger \mathcal{T} values. However, the critical threshold \mathcal{T}_c , as well as the quantities $W^{(1)}, \dots, W^{(A)}$ used to approximate it (see Eqs. (4.2.5) and (4.2.6)) also correspondingly increase. For our simulations, we did not observe a significant increase in the value of \mathcal{T} when the new inequality was included, but we did observe a shift in the distribution of W to larger values across many parameter choices. This is clearly shown in Fig. 4.12 (right) comparing the amalgamated distribution of W with and without the internal branch inequality for aggregate data from 1000 trees drawn randomly from the treespace shown in Fig. 4.9. Similar tests with data drawn from fixed trees support this conclusion.

For the false $H_{13|24}$, the aggregate histogram plots in Fig. 4.12 (middle two plots) shows no effect from including the internal branch inequality. However, in testing $H_{13|24}$ and $H_{14|23}$ on Collection 2 (not shown), we observed an effect dependent on the region of the parameter space. When $a, b \leq .2$, the inclusion of the internal edge inequality had no appreciable effect on the observed distribution of p -values, which were overwhelmingly concentrated near zero regardless. However, for $(a, b) \in \{(.8, .05), (.8, .2), (.2, .8)\}$, including the internal branch inequality increased the number of small p -values. However $(a, b) = (.05, .8)$ with $H_{13|24}$ was exceptional, showing almost no difference.

4.4.5 SDL-Based Phylogenetic Inference

We next investigate the potential of the SDL test as an inference method for tree topology from sequence data. Standard statistical approaches for this depend on repeated calculation of a




Figure 4.12: Histograms of p -values when testing $H_{12|34}$ (left two) and $H_{13|24}$ (middle two) showing the effect of including the internal edge constraint. The plot for $H_{14|23}$ is omitted because it is similar to that of $H_{13|24}$. Right: Histogram of $W^{(1)}, \dots, W^{(A)}$, approximating the test distribution \mathcal{T}_C for CDD (red) and CDD + inequality (blue) for aggregate data from 1000 trees with random parameters $a, b \in (0, 1.2]$.

likelihood function depending not only on the tree topology but also its edge lengths, with a search or MCMC exploration over all these parameters. The SDL methodology suggests a procedure to avoid consideration of the edge length “nuisance parameters” and likelihood calculations by first calculating SDL p -values for each possible tree topology and then selecting the tree with the highest p -value. We implemented this procedure for the CFN model with 4 taxa.

Performance for differing constraint sets

Fig. 4.13 shows results of this method applied to p -values from Collection 1, following a standard graphical depiction introduced in [56]. The columns of plots correspond to different choices of constraint sets, and the rows to $r = 0$ and 20. Within each plot, each pixel corresponds to a pair (a, b) of edge length parameters, as in Fig. 4.9. Grey levels indicate the frequency of inferring the true tree topology (black=100%, white=0%). The red curves demarcate a region of good performance where correct inference occurs with frequency at least 90%. This region forms a right-skewed hump along the horizontal axis, similar to those produced by other well-performing methods [56, 48, 27]. In this region, the SDL method performs well in part due to the p -values for $H_{13|24}$ and $H_{14|23}$ being highly concentrated near zero (see Figs. 4.10 and 4.11).

The left two columns of Fig. 4.13 compare the use of the constraint sets CDM and PDM ($m = 12$), as in Fig. 4.10. For $r = 0$, CDM outperformed PDM both in terms of raw success percentage and overall shape and size of the dark region. However, this advantage was diminished with improved performance of PDM when $r = 20$ convex combination constraints were included. A similar pattern was observed for CDR and PDR as we show in Appendix A.3.2.

This observation that for $r = 0$, the use of CDM gives better performance for model selection

than PDM is consistent with our conclusions from Section 4.4.4 on hypothesis testing. However, when we increase the number r of convex combinations this performance gap almost entirely disappears, suggesting that the use of convex combinations may be a powerful general-purpose tool to improve performance of the SDL test, especially when the geometry of a model is not fully understood.

We next investigated whether the performance of the SDL-based inference method improved with the inclusion of the internal edge inequality, Eq. (4.4.3). Fig. 4.13 (right) presents results for the CDD constraints, showing the more complete semi-algebraic model description expands the region of good performance. This reinforces previous observations about the importance of using the full semi-algebraic description for phylogenetic model selection [26, 25]. We also found that increasing m from 12 to 30 resulted in a larger region of good performance. Despite theoretical reasons to prefer smaller values of m , for model selection choosing $m = 30$ resulted in a better performance, even though the p -value distributions showed little difference.

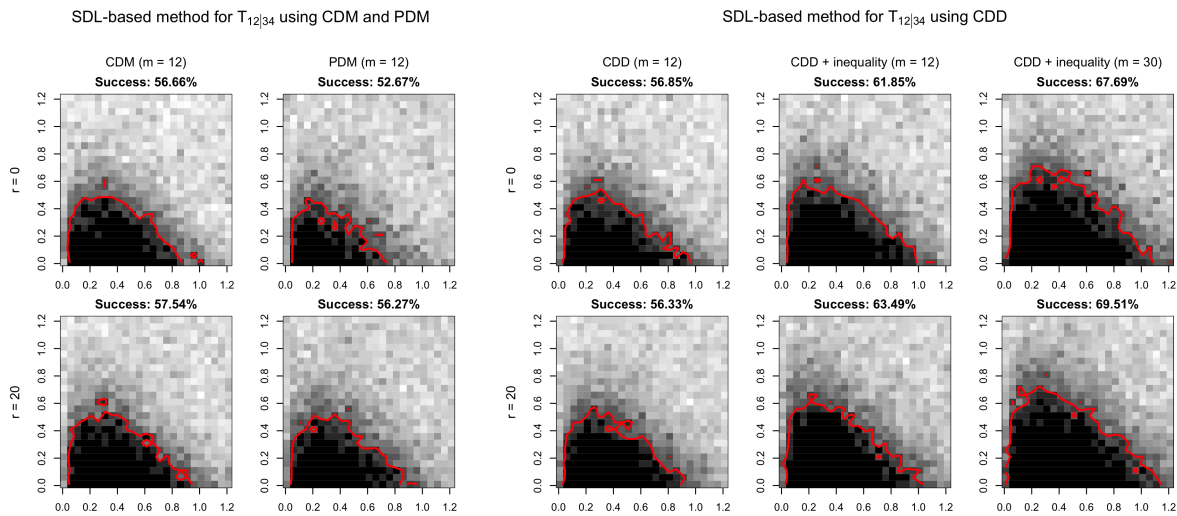


Figure 4.13: Performance of the SDL test for inferring the tree topology $T_{12|34}$ using different constraint sets and values of m . Left: CDM and PDM constraints with $m = 12$. Right: CDD constraints with $m = 12$, CDD with the inequality of Eq. (4.4.3) and $m = 12$, CDD with the inequality and $m = 30$. Rows vary the number of convex combinations, $r = 0$ and $r = 20$. Grey levels represent the frequency of correctly inferring the topology for edge length pairs (a, b) (black 100%, white 0%).

Comparison with other inference methods

We compared the performance of SDL-based inference to that of two other phylogenetic reconstruction methods, Maximum Likelihood and the Singular Value Decomposition (SVD) method. The SVD approach is also motivated by polynomial model constraints, as it relies on the fact that a certain matrix flattening of the probability tensor p , determined by the tree topology, must have rank 2. Although based on essentially the same constraints as PDR, it uses the Singular Value Decomposition of an estimate of p to measure its closeness (in Frobenius norm) to one of rank 2, choosing the tree topology minimizing this. SVD-based inference has been exploited for empirical inference several in phylogenetic settings [5, 37, 48].

Fig. 4.14 shows the performance of these three methods on identical simulated data. For the SDL approach we use the CDD constraint set together with the internal edge inequality, $m = 30$ and $r = 20$. For the gold standard maximum likelihood estimation (MLE), calculations used the Julia package `FourLeafMLE.jl` [54]. An important conclusion of Fig. 4.14 is that with well-chosen user-specified parameters, the SDL method can achieve overall performance approaching Maximum Likelihood, and better than the SVD approach most often used in algebraic approaches to inference.

Of special note is the performance of the SDL test for tree parameters in the Felsenstein zone (see Fig. 4.9) in which correct inference is difficult for all methods. The SDL test achieved a success rate of 60.2%, compared to 71.65% for MLE and 37.13% for SVD. Thus while performance declined in this region, for SDL the decline was considerably less than for SVD. We also observed that the SDL test substantially reduced (especially compared to SVD) the bias toward a specific false hypothesis (i.e., long branch attraction) in the Felsenstein zone, as is common for other methods. For more details see Appendix A.3.3.

However the SDL approach is by far the more computationally intensive than the other two methods. The computational time producing this figure for the SDL-based approach was 12.57 hours (using an R and C++ implementation) versus 53.2 minutes for MLE (in Julia) and 11.25 seconds for SVD (in R) (see Section 4.5 for more details).

4.5 Implementation Details and Computational Performance

The code used in our simulations is primarily written in R (version 4.2.2), with performance-critical parts implemented in C++ and integrated using the Rcpp package (version 1.0.12). The code, which builds on the original implementation from [89], is available at:

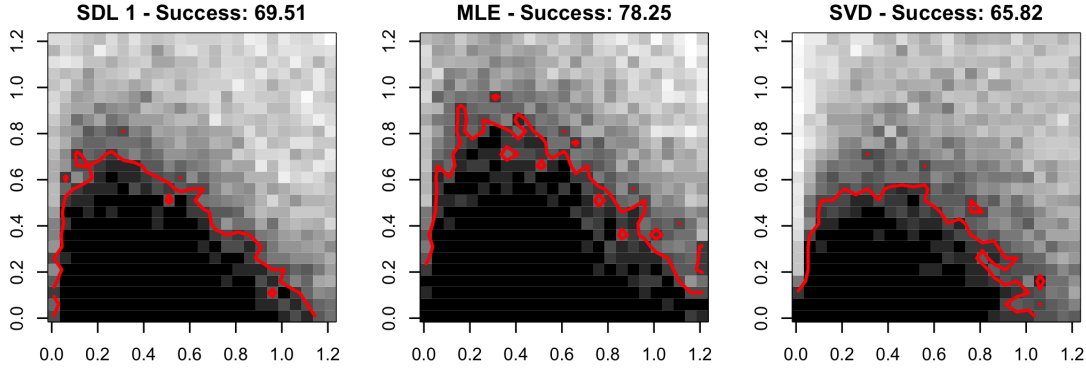


Figure 4.14: Performance of 3 methods of topological tree inference on data from Collection 1: (left) the SDL-based inference method using the CDD constraint set with the internal edge inequality, with $m = 30$ and $r = 20$, (middle) Maximum Likelihood [54], (right) the SVD method.

github.com/marinagarrote/Semialg-Hypothesis-Test-with-Incomplete-U-Stats.

All computations were performed with an Intel(R) Core(TM) i5-10400 CPU @ 2.90GHz Processor equipped with 64 GB RAM, running Debian 12.5.

Average computation times for the trinomial models presented in Section 4.3 are as follows. For Model 1 computing a single p -value took an average of 0.21 seconds when $m = 1$, 0.09 seconds when $m = 5$, and 0.06 seconds when $m = 15$. In the case of Model 2, the average time to compute a single p -value was 0.09 seconds for $r = 0$, 0.12 seconds for $r = 10$, and 0.3 seconds for $r = 100$. Model 3 p -values required an average of 0.11 seconds for $s = 1$, 0.36 seconds for $s = 10$, and 2.88 seconds for $s = 100$.

To efficiently run simulations for the CFN model in Section 4.4, we used the `parallel` package in R (version 4.2.2) and 6 cores. For a fixed choice of parameters (a, b) as defined in Section 4.4.2, the average runtime for a single p -value was approximately 0.78 seconds when $r = 0$ and 0.97 seconds when $r = 20$. The chosen constraint set of polynomials had negligible effect on these runtimes.

Finally, the MLE computations presented in Section 4.4.5 were carried out using Julia (version 1.10.3).

4.6 Conclusion

The SDL method offers a general-purpose framework for hypothesis testing for models defined by polynomial constraints. It is a strong and much needed technique, especially in settings

where traditional frameworks are not available, such as when models have singularities or boundaries. Indeed, as illustrated by the trinomial submodels in Section 4.3, the method's performance can closely match that of traditional deterministic tests, such as the likelihood ratio or χ^2 , where they are justified, but is more widely applicable. By focusing on two well-studied types of algebraic models used in phylogenetic inference, our investigation confirms that the method performs well across different settings. While no alternative method matches its generality, our results emphasize that thoughtful implementation choices, particularly around the key elements of constraint specification, kernel construction, and symmetrization, are necessary to enhance test performance.

In the case of the multispecies coalescent trinomial submodels, the SDL method not only recovers rejection regions that closely match those of conventional tests when available, but also remains valid at boundary and singular points, such as line crossings and nodes. However, our simulations show that near singularities, the complement of the rejection region widens, making rejection more difficult. This indicates that the behaviour of the test is influenced not just by the zero set of the defining polynomials, but also by the size of constraint polynomial values near that set. As we see in those models, adding redundant constraints, especially near singularities and boundaries, can increase the power of the test, but how to choose these in a general manner requires further investigation.

For the CFN model, we illustrated how the SDL method can be used both for hypothesis testing and selection among non-nested models. This is especially useful in situations such as phylogenetic tree inference where the three possible four-leaf topologies give rise to intersecting semi-algebraic sets. In addition, the CFN model highlighted how the choice of generating polynomials for the defining ideal plays a key role. Generators that are completely distinguishing for the tree topology of interest lead to better-calibrated tests than partially distinguishing ones. Furthermore, this case study illustrated how the addition of convex combinations or extra constraints can have mixed effects, sometimes improving and sometimes degrading performance. Both of these issues raises the question of whether it is possible to develop a principled method for constraint choice.

Both types of models that we explore in this chapter are relatively small, in terms of ambient dimension, in terms of the number of constraints, and in terms of the constraint degrees. The method presents computational challenges for moving to higher dimensional, and higher degree, settings. In particular, full symmetrization is infeasible for large degree constraints, which requires larger m , but our results indicate that partial symmetrization using a modest number of permutations performs well in practice. This raises an important theoretical question: How

many permutations are sufficient to approximate the fully symmetrized kernel, and how does this number scale with dimension and degree?

While our case studies were chosen from evolutionary biology, they highlight that the SDL method fills a critical methodological gap in statistics for any semi-algebraic model. However, its performance is intimately tied to both algebraic and geometric aspects of the model. Future work under the lens of algebraic geometry would be helpful to develop a more complete theoretical understanding of how types of singularities and constraint choices influence the behaviour of the method, especially in higher dimensional settings in which visualisation is difficult. Such developments would further enhance the utility of the SDL method for both hypothesis testing and model selection in phylogenetics and other fields.

REFERENCES

- [1] Elizabeth Allman and John Rhodes. Quartets and parameter recovery for the general markov model of sequence mutation. *AMRX Applied Mathematics Research eXpress*, 2004, 01 2004.
- [2] Elizabeth Allman and John Rhodes. The identifiability of tree topology for phylogenetic models, including covarion and mixture models. *Journal of Computational Biology*, 13:1101–13, 07 2006.
- [3] Elizabeth S. Allman, Hector Baños, Jonathan D. Mitchell, and John A. Rhodes. The tree of blobs of a species network: identifiability under the coalescent. *Journal of Mathematical Biology*, 86(1):10, 2023.
- [4] Elizabeth S. Allman, James H. Degnan, and John A. Rhodes. Identifying the rooted species tree from the distribution of unrooted gene trees under the coalescent. *Journal of Mathematical Biology*, 62(6):833–862, 2011.
- [5] Elizabeth S. Allman, Laura S. Kubatko, and John A. Rhodes. Split scores: a tool to quantify phylogenetic signal in genome-scale data. *Systematic Biology*, 66(4):620–636, 01 2017.
- [6] Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *Ann. Statist.*, 37(6A):3099–3132, 2009.
- [7] Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 37(6A):3099–3132, 2009.
- [8] Elizabeth S Allman and John A Rhodes. Phylogenetic invariants for the general markov model of sequence mutation. *Math. Biosci.*, 186(2):113–144, December 2003.
- [9] Elizabeth S. Allman and John A. Rhodes. Phylogenetic ideals and varieties for the general markov model. *Advances in Applied Mathematics*, 40:127–148, 2 2008.
- [10] E.S. Allman, H. Baños, J.D. Mitchell, and J.A. Rhodes. Tinnik: Inference of the tree of blobs of a species network under the coalescent. *Algorithms in Molecular Biology*, 19(1):??, 2024.
- [11] Hans-Jürgen Bandelt and Andreas W.M. Dress. Split decomposition: A new and useful approach to phylogenetic analysis of distance data. *Molecular Phylogenetics and Evolution*, 1(3):242–252, 1992.
- [12] Hector Baños. Identifying species network features from gene tree quartets under the coalescent model. *Bulletin of Mathematical Biology*, 81:494–534, 2019.

- [13] David Barnhill, Marina Garrote-López, Elizabeth Gross, Max Hill, Bryson Kagy, John A. Rhodes, and Joy Z. Zhang. Methodological considerations for semialgebraic hypothesis testing with incomplete u-statistics, 2025.
- [14] Louis J. Billera, Susan P. Holmes, and Karen Vogtmann. Geometry of the space of phylogenetic trees. *Adv. Appl. Math.*, 27(4):733–767, November 2001.
- [15] Marianne B. Bjorner, Erin K. Molloy, Colin N. Dewey, and Claudia Solis-Lemus. Detectability of varied hybridization scenarios using genome-scale hybrid detection methods. *Bulletin of the Society of Systematic Biologists*, 3(1), 10 2024.
- [16] Gunnar Blom. Some properties of incomplete u-statistics. *Biometrika*, pages 573–580, 1976.
- [17] Tobias Boege, Kaie Kubjas, Pratik Misra, and Liam Solus. Colored gaussian dag models. *arXiv preprint arXiv:2404.04024*, 2024.
- [18] BM Brown and DG Kildea. Reduced u-statistics and the hodes-lehmann estimator. *The Annals of Statistics*, pages 828–835, 1978.
- [19] David Bryant and Vincent Moulton. Neighbor-Net: An Agglomerative Method for the Construction of Phylogenetic Networks. *Molecular Biology and Evolution*, 21(2):255–265, 02 2004.
- [20] Peter Buneman. The recovery of trees from measures of dissimilarity. In *Mathematics the the Archeological and Historical Sciences*, pages 387–395, United Kingdom, 1971. Edinburgh University Press.
- [21] Peter Buneman. A note on the metric properties of trees. *Journal of Combinatorial Theory, Series B*, 17(1):48–50, 1974.
- [22] Ruichu Cai, Zhiyi Huang, Wei Chen, Zhifeng Hao, and Kun Zhang. Causal discovery with latent confounders based on higher-order cumulants. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org, 2023.
- [23] M Casanellas and J Fernández-Sánchez. Performance of a new invariants method on homogeneous and nonhomogeneous quartet trees. *Molecular Biology and Evolution*, 24(1):288–293, 10 2006.
- [24] Marta Casanellas and Jesús Fernández-Sánchez. Geometry of the kimura 3-parameter model. *Advances in Applied Mathematics*, 41(3):265–292, 2008.
- [25] Marta Casanellas, Jesus Fernandez-Sanchez, and Marina Garrote-Lopez. Saq: Semi-algebraic quartet reconstruction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 18:2855–2861, 11 2021.

- [26] Marta Casanellas, Jesús Fernández-Sánchez, and Marina Garrote-López. Distance to the stochastic part of phylogenetic varieties. *Journal of Symbolic Computation*, 104:653–682, 5 2021.
- [27] Marta Casanellas, Jesús Fernández-Sánchez, Marina Garrote-López, and Marc Sabaté-Vidales. Designing weights for quartet-based methods when data are heterogeneous across lineages. *Bulletin of Mathematical Biology*, 85:68, 7 2023.
- [28] Marta Casanellas and Mike Steel. Phylogenetic mixtures and linear invariants for equal input models. *Journal of Mathematical Biology*, 74:1107–1138, 4 2017.
- [29] Marta Casanellas and Seth Sullivant. *The Strand Symmetric Model*, page 305–321. Cambridge University Press, 2005.
- [30] George Casella and Roger L Berger. *Statistical inference*, volume 2. Duxbury Pacific Grove, CA, 2002.
- [31] James A. Cavender and Joseph Felsenstein. Invariants of phylogenies in a simple case with discrete states. *Journal of Classification*, 4(1):57–71, 1987.
- [32] Clara S. Chan, David P. Robbins, and David S. Yuen. On the volume of a certain polytope. *Experiment. Math.*, 9(1):91–99, 2000.
- [33] Joseph T. Chang. Full reconstruction of markov models on evolutionary trees: Identifiability and consistency. *Mathematical Biosciences*, 137(1):51–73, 1996.
- [34] Xiaohui Chen. Gaussian and bootstrap approximations for high-dimensional u-statistics and their applications. *The Annals of Statistics*, 46(2), 2018.
- [35] Xiaohui Chen and Kengo Kato. Randomized incomplete U -statistics in high dimensions. *The Annals of Statistics*, 47(6):3127 – 3156, 2019.
- [36] Victor Chepoi and Bernard Fichet. A note on circular decomposable metrics. *Geom. Dedicata*, 69(3):237–240, 1998.
- [37] Julia Chifman and Laura Kubatko. Quartet inference from snp data under the coalescent model. *Bioinformatics*, 30:3317–3324, 12 2014.
- [38] Satyan L. Devadoss and Samantha Petti. A space of phylogenetic networks. *SIAM J. Appl. Algebra Geom.*, 1(1):683–705, 2017.
- [39] Michel Marie Deza and Monique Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer, Heidelberg, 2010. First softcover printing of the 1997 original [MR1460488].

- [40] Adrian Dobra, Stephen E. Fienberg, Alessandro Rinaldo, Aleksandra Slavkovic, and Yi Zhou. *Algebraic Statistics and Contingency Table Problems: Log-Linear Models, Likelihood Estimation, and Disclosure Limitation*, pages 63–88. Springer New York, New York, NY, 2009.
- [41] Jan Draisma and Jochen Kuttler. On the ideals of equivariant tree models. *Mathematische Annalen*, 344(3):619–644, December 2008.
- [42] Andreas Dress and Daniel Huson. Constructing splits graphs. *IEEE Trans. Comput. Biol. Bioinformatics*, 1:1109–1115, 01 2004.
- [43] Mathias Drton. Likelihood ratio tests and singularities. *Ann. Statist.*, 37(1):979–1012, 2009.
- [44] Robin J Evans. Model selection and local geometry. *The Annals of Statistics*, 48(6):3513–3544, 2020.
- [45] Steven N. Evans and T. P. Speed. Invariants of some probability models used in phylogenetic inference. *The Annals of Statistics*, 21(1):355–377, 1993.
- [46] Joseph Felsenstein. Cases in which parsimony or compatibility methods will be positively misleading. *Systematic Zoology*, 27:401, 12 1978.
- [47] Joseph Felsenstein. *Inferring Phylogenies*. Sinauer, 2003.
- [48] Jesús Fernández-Sánchez and Marta Casanellas. Invariant versus classical quartet inference when evolution is heterogeneous across sites and lineages. *Systematic Biology*, 65:280–291, 3 2016.
- [49] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at <http://www2.macaulay2.com>.
- [50] Branko Grünbaum, Volker Kaibel, Victor Klee, and Günter M. Ziegler. *Convex polytopes*. Springer, New York, 2003.
- [51] Yuqi Gu and Gongjun Xu. Partial identifiability of restricted latent class models. *The Annals of Statistics*, 48(4):pp. 2082–2107, 2020.
- [52] Michael D. Hendy and David Penny. A framework for the quantitative study of evolutionary trees. *Systematic Zoology*, 38:297, 12 1989.
- [53] Michael D. Hendy and David Penny. Complete families of linear invariants for some stochastic models of sequence evolution, with and without molecular clock assumption. *J. Comput. Biol.*, 3(1):19–31, 1996.
- [54] Max Hill and Jose Israel Rodriguez. A maximum likelihood estimator for quartets under the cavender-farris-neyman model. *ACM Communications in Computer Algebra*, 58:35–38, 6 2024.

- [55] Benjamin Hollering and Seth Sullivant. Identifiability in phylogenetics using algebraic matroids. *J. Symbolic Comput.*, 104:142–158, 2021.
- [56] John P. Huelsenbeck. Performance of phylogenetic methods in simulation. *Systematic Biology*, 44(1):17–48, 03 1995.
- [57] Daniel H. Huson, Regula Rupp, and Celine Scornavacca. *Phylogenetic Networks: Concepts, Algorithms and Applications*. Cambridge University Press, 2010.
- [58] Svante Janson. The asymptotic distributions of incomplete u-statistics. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 66(4):495–505, 1984.
- [59] C.J. Jardine, N. Jardine, and R. Sibson. The structure and construction of taxonomic hierarchies. *Mathematical Biosciences*, 1(2):173–179, 1967.
- [60] Cantor C. Jukes T.H. Evolution of protein molecules. *Mammalian Protein Metabolism*, pages 21–32, 1969.
- [61] Bryson Kagy and Seth Sullivant. Equidistant circular split networks. *To appear in SIAM Journal on Applied Algebra and Geometry*, 2024. arXiv:2402.11032.
- [62] Kenneth Kalmanson. Edgeconvex circuits and the traveling salesman problem. *Canadian J. Math.*, 27(5):1000–1010, 1975.
- [63] Motoo Kimura. A simple method for estimating evolutionary rates of base substitutions through comparative studies of nucleotide sequences. *Journal of Molecular Evolution*, 16:111–120, 1980.
- [64] Sungsik Kong, Joan Pons, Laura Kubatko, and Kristina Wicke. Classes of explicit phylogenetic networks and their biological and mathematical significance. *Journal of Mathematical Biology*, 84, 05 2022.
- [65] Dimitra Kosta and Kaie Kubjas. Maximum likelihood estimation of symmetric group-based models via numerical algebraic geometry. *Bulletin of Mathematical Biology*, 81(2):337–360, October 2018.
- [66] Joseph Kruskal. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra Appl.*, 18(2):95–138, 1977.
- [67] James A Lake. A rate-independent technique for analysis of nucleic acid sequences: evolutionary parsimony. *Molecular biology and evolution*, 4(2):167–191, 1987.
- [68] Steffen L. Lauritzen. *Graphical Model*. Oxford University Press, 1996.
- [69] Dan Levy and Lior Pachter. The neighbor-net algorithm. *Adv. in Appl. Math.*, 47(2):240–258, 2011.

- [70] Colby Long and Seth Sullivant. Identifiability of 3-class Jukes-Cantor mixtures. *Adv. in Appl. Math.*, 64:89–110, 2015.
- [71] Frederick A. Matsen. Fourier transform inequalities for phylogenetic trees. *IEEE/ACM transactions on computational biology and bioinformatics*, 6(1):89–95, 2008.
- [72] C. Meng and L.S. Kubatko. Detecting hybrid speciation in the presence of incomplete lineage sorting using gene tree incongruence: a model. *Theoretical Population Biology*, 75(1):35–45, 2009.
- [73] Karola Mészáros and Alejandro H. Morales. Flow polytopes of signed graphs and the Kostant partition function. *Int. Math. Res. Not. IMRN*, 2015(3):830–871, 2013.
- [74] Mateusz Michałek and Bernd Sturmfels. *Invitation to nonlinear algebra*, volume 211 of *Graduate studies in mathematics*. American Mathematical Society, Providence, R.I., 2021.
- [75] Charles D. Michener and Robert R. Sokal. A quantitative approach to a problem in classification. *Evolution*, 11(2):130–162, 1957.
- [76] Jonathan D. Mitchell, Elizabeth S. Allman, and John A. Rhodes. Hypothesis testing near singularities and boundaries. *Electronic Journal of Statistics*, 13(1):2150, 2019.
- [77] Pekka Pamilo and Masatoshi Nei. Relationships between gene trees and species trees. *Mol. Biol. Evol.*, 5(5):568–583, 1988.
- [78] Sarah L. Parks and Nick Goldman. Maximum likelihood inference of small trees in the presence of long branches. *Systematic Biology*, 63(5):798–811, 07 2014.
- [79] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2023.
- [80] John Rhodes and Seth Sullivant. Identifiability of large phylogenetic mixture models. *Bull. Math. Biol.*, 74(1):212–231, 2012.
- [81] John A. Rhodes, Hector Baños, Jonathan D. Mitchell, and Elizabeth S Allman. MSCquartets 1.0: quartet methods for species trees and networks under the multispecies coalescent model in R. *Bioinformatics*, 37(12):1766–1768, 10 2020.
- [82] Joseph P. Rusinko and Brian Hipp. Invariant based quartet puzzling. *Algorithms for Molecular Biology*, 7:1–9, 2012.
- [83] Daniela Schkoda, Elina Robeva, and Mathias Drton. Causal discovery of linear non-Gaussian causal models with unobserved confounding. *arXiv:2408.04907*, 2024.
- [84] Charles Semple and Mike Steel. *Phylogenetics*. Oxford lecture series in mathematics and its applications. Oxford University Press, 2003.

- [85] Yanglei Song, Xiaohui Chen, and Kengo Kato. Approximating high-dimensional infinite-order U -statistics: Statistical and computational guarantees. *Electronic Journal of Statistics*, 13(2):4794 – 4848, 2019.
- [86] Mike Steel. *Phylogeny—discrete and random processes in evolution*, volume 89 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2016.
- [87] V. Strassen. Rank and optimal computation of generic tensors. *Linear Algebra Appl.*, 52/53:645–685, 1983.
- [88] Nils Sturma. *TestGGM: Testing Gaussian Graphical Models*, 2021. R package version 1.0.
- [89] Nils Sturma, Mathias Drton, and Dennis Leung. Testing many constraints in possibly irregular models using incomplete U -statistics. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, page qkae022, 03 2024.
- [90] Bernd Sturmfels and Seth Sullivant. Toric ideals of phylogenetic invariants. *Journal of Computational Biology*, 12:457–481, 5 2005.
- [91] Seth Sullivant. *Algebraic statistics*, volume 194 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2018.
- [92] Edward Susko and Andrew J Roger. Long branch attraction biases in phylogenetics. *Systematic Biology*, 70(4):838–843, 02 2021.
- [93] Y. Samuel Wang and Mathias Drton. High-dimensional causal discovery under Non-Gaussianity. *Biometrika*, 107(1):41–59, 2019.
- [94] L. Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer Texts in Statistics. Springer New York, 2010.
- [95] Samaneh Yourdkhani, Elizabeth Allman, and John Rhodes. Parameter identifiability for a profile mixture model of protein evolution. *J. Comput. Biol.*, 28(6):570–586, 2021.
- [96] Günter M. Ziegler. *Lectures on polytopes*. Springer-Verlag, New York, 1995.

APPENDIX

APPENDIX

A

COALESCENT MODELS

A.1 The multispecies coalescent model

The (network) multispecies coalescent (MSC) [77, 72] models the formation of gene trees within species trees or networks, for example as in Fig. A.1. A gene tree describes the history of a single genetic locus drawn from individuals in several extant species, as lineages trace back through individuals in the ancestral species populations, coalescing at common ancestors. While constrained by the species relationships, a gene tree may differ from them significantly, due to multiple gene lineages remaining distinct in an ancestral population until coalescence between less closely related species becomes possible. This effect, called *incomplete lineage sorting*, is most pronounced when edges in the species tree or network are short (in number of generations) or population sizes are large (since bottlenecks promote coalescence).

Considering only trees or networks relating four species, a *quartet Concordance Factor (CF)* for a fixed network is the vector of probabilities of the 3 possible unrooted topological gene trees shown in Fig. 4.8 that may arise under the coalescent model. To be precise, we fix the order

$$CF = (p_{12|34}, p_{13|24}, p_{14|23}),$$




Figure A.1: Gene trees (in red) form within a species tree and network (black 'tubes')

for some fixed designation of species 1, 2, 3, 4.

Under the MSC model, the form of CFs arising from metric species networks with certain topological structures has been studied in several papers, leading to the four submodels of Δ^2 depicted in Fig. 4.1. Model 1 is all CFs that may arise from a species network with a cut edge separating species 1, 2 from 3, 4 [4]. Model 2 is all CFs that arise from a species tree with the same species separation [4, 76]. Model 3 is all CFs from a network with a cut edge separating the species into some pair of sets of two [3], and Model 4 all CFs from a tree with such a cut edge [76]. Models 3 and 4 are obtained from Models 1 and 2 by considering the union of models obtained by permuting CF entries. It is also known that all points in Δ^2 arise as CFs of some networks [12], so rejecting these models in a hypothesis test is a natural way to find evidence for gene flow or hybridization [15].

From genomic sequences, one may infer many gene trees and from them estimate frequencies of the three possible quartet gene tree topologies. A hypothesis test with one of the above null models can then, give insight into an unknown network structure. For instance, rejection of Model 3 suggests that the data did not arise on a tree, so hybridization or introgression occurred among the species. Specialized test distributions for null hypotheses of Model 2 and 4 are derived in [76] and for Model 3 in [10] that improve upon a naive use of a standard distribution that ignores the singularities and boundaries of the models. (Model 1 can be tested with a standard distribution, as it lacks any irregularities.) However, these models are all semialgebraic, and the SDL approach offers an alternative testing framework without the need for such detailed work for each model.

A.2 Deterministic tests

For comparison to the rejection region plots produced by the SDL tests in Section 4.3.4 we show those for deterministic tests for models 1-4 and Hardy-Weinberg with sample size $n = 300$. For Model 1 this is a standard Likelihood Ratio test; for Models 2, 3, and 4 we use the tests implemented in MSCquartets [81] as “T1”, “cut”, and “T3”. These last all use non-standard test distributions for the Likelihood ratio statistic, to deal with the boundaries and singularities of these models. For the Hardy-Weinberg 2-allele model we use a standard chi-squared test.




Figure A.2: Rejection regions for Models 1, 2, 3, 4, and Hardy-Weinberg 2-alleles, using deterministic tests, as described in text, with sample size $n = 300$.

A.3 Additional Details on the CFN model

A.3.1 Generating sets for the CFN ideal

This section details the derivation of generating sets for the ideal associated with the 4-taxon Cavender-Farris-Neyman (CFN) model presented in Sections 4.4.1 and 4.4.2, and provides explicit formulas for them.

Although initially presented in $2^4 = 16$ dimensions, using pairwise equalities of certain pattern probabilities, the model can also be presented in 8 dimensions. Specifically, let

- $I_T \subset \mathbb{C}[p_{ijkl}]$ be the full phylogenetic ideal in the ring of 16 site pattern probabilities p_{ijkl} ,

$i, j, k, l \in \{0, 1\}$.

- $\bar{I}_T \subset \mathbb{C}[\bar{p}_{xyzt}]$, the ideal in the ring of 8 symmetrized pattern probabilities \bar{p}_{xyzt} (e.g., $\bar{p}_{xxxx} = p_{0000} + p_{1111}$).

Our data consists of n independent multinomial samples, with parameter

$$\bar{p} = (\bar{p}_{xxxx}, \bar{p}_{xxxy}, \bar{p}_{xxyx}, \bar{p}_{xxyy}, \bar{p}_{xyxx}, \bar{p}_{xyxy}, \bar{p}_{xyyx}, \bar{p}_{xyyy}) \in \Delta^7, \quad (\text{A.3.1})$$

and we seek generators of \bar{I}_T in the \bar{p} coordinates.

With $\mathbb{C}[p] := \mathbb{C}[p_{ijkl}]$, and $\mathbb{C}[\bar{p}] := \mathbb{C}[\bar{p}_{xyzt}]$, the relationship between I_T and \bar{I}_T is given by the ring homomorphism $\psi : \mathbb{C}[\bar{p}] \rightarrow \mathbb{C}[p]$ that substitutes each \bar{p}_{xyzt} with its definition as a sum of two p_{ijkl} :

$$\psi(\bar{p}_{xxxx}) = p_{0000} + p_{1111}, \quad \psi(\bar{p}_{xxxy}) = p_{0001} + p_{1110}, \quad \dots \quad \psi(\bar{p}_{xyyy}) = p_{0111} + p_{1000}.$$

Let $L_{\text{symm}} \subset \mathbb{C}[p]$ be the ideal generated by the 8 linear symmetry relations,

$$p_{0000} - p_{1111} = 0, \quad p_{0001} - p_{1110} = 0, \quad \dots, \quad p_{0111} - p_{1000} = 0.$$

Then

$$I_T = \psi(\bar{I}_T) + L_{\text{symm}}. \quad (\text{A.3.2})$$

We will show that \bar{I}_T is generated by the linear polynomial, $\ell := (\bar{p}_{xxxx} + \bar{p}_{xxxy} + \bar{p}_{xxyx} + \bar{p}_{xxyy} + \bar{p}_{xyxx} + \bar{p}_{xyxy} + \bar{p}_{xyyx} + \bar{p}_{xyyy}) - 1$, along with a set of quadratic polynomials. It then follows from Equation A.3.2 that I_T is generated by the symmetry ideal L_{symm} , the linear polynomial $\psi(\ell) = (\sum_{ijkl} p_{ijkl}) - 1$, and the ψ -images of the aforementioned quadratic polynomials (which remain quadratic in the p_{ijkl} coordinates). The sets CDD, CDM, CDR, PDM and PDR consist of variations of these quadratics in the \bar{p}_{ijkl} coordinates.

For the tree $T = T_{12|34}$, we calculate generators for \bar{I}_T using Macaulay2 (version 1.21). Below, we code the parametrization of the \bar{p}_{xyzt} in terms of transformed edge lengths $\theta_i = e^{-2t_i}$, with p_{xxxx} , p_{xxxy} , etc., corresponding to the coordinates \bar{p}_{xxxx} , \bar{p}_{xxxy} , etc..

```

i1 : R = QQ[θ1, θ2, θ3, θ4, θ5]
i2 : Sp = QQ[pxxxx, pxxxy, pxxyx, pxxyy, pxyxx, pxyxy, pxyyx, pxyyy]

i3 : β = θ -> (1-θ)/2
i4 : α = θ -> (1+θ)/2

i5 : Pxxxx = α(θ1)*α(θ2)*α(θ3)*α(θ4)*α(θ5) + α(θ1)*α(θ2)*β(θ3)*β(θ4)*β(θ5) +
          α(θ3)*α(θ4)*β(θ1)*β(θ2)*β(θ5) + α(θ5)*β(θ1)*β(θ2)*β(θ3)*β(θ4)
          -- p0000 + p1111
i6 : Pxxxy = α(θ1)*α(θ2)*α(θ3)*α(θ5)*β(θ4) + α(θ1)*α(θ2)*α(θ4)*β(θ3)*β(θ5) +
          α(θ3)*β(θ1)*β(θ2)*β(θ4)*β(θ5) + α(θ4)*α(θ5)*β(θ1)*β(θ2)*β(θ3)
          -- p0001 + p1110
i7 : Pxyyx = α(θ1)*α(θ2)*α(θ3)*β(θ4)*β(θ5) + α(θ1)*α(θ2)*α(θ4)*α(θ5)*β(θ3) +
          α(θ3)*α(θ5)*β(θ1)*β(θ2)*β(θ4) + α(θ4)*β(θ1)*β(θ2)*β(θ3)*β(θ5)
          -- p0010 + p1101
i8 : Pxyyy = α(θ1)*α(θ2)*α(θ3)*α(θ4)*β(θ5) + α(θ1)*α(θ2)*α(θ5)*β(θ3)*β(θ4) +
          α(θ3)*α(θ4)*α(θ5)*β(θ1)*β(θ2) + β(θ1)*β(θ2)*β(θ3)*β(θ4)*β(θ5)
          -- p0011 + p1100
i9 : Pxyxx = α(θ1)*α(θ3)*α(θ4)*α(θ5)*β(θ2) + α(θ1)*β(θ2)*β(θ3)*β(θ4)*β(θ5) +
          α(θ2)*α(θ3)*α(θ4)*β(θ1)*β(θ5) + α(θ2)*α(θ5)*β(θ1)*β(θ3)*β(θ4)
          -- p0100 + p1011
i10 : Pxyxy = α(θ1)*α(θ3)*α(θ5)*β(θ2)*β(θ4) + α(θ1)*α(θ4)*β(θ2)*β(θ3)*β(θ5) +
          α(θ2)*α(θ3)*β(θ1)*β(θ4)*β(θ5) + α(θ2)*α(θ4)*α(θ5)*β(θ1)*β(θ3)
          -- p0101 + p1010
i11 : Pxyyx = α(θ1)*α(θ3)*β(θ2)*β(θ4)*β(θ5) + α(θ1)*α(θ4)*α(θ5)*β(θ2)*β(θ3) +
          α(θ2)*α(θ3)*α(θ5)*β(θ1)*β(θ4) + α(θ2)*α(θ4)*β(θ1)*β(θ3)*β(θ5)
          -- p0110 + p1001
i12 : Pxyyy = α(θ1)*α(θ3)*α(θ4)*β(θ2)*β(θ5) + α(θ1)*α(θ5)*β(θ2)*β(θ3)*β(θ4) +
          α(θ2)*α(θ3)*α(θ4)*α(θ5)*β(θ1) + α(θ2)*β(θ1)*β(θ3)*β(θ4)*β(θ5)
          -- p0111 + p1000

i13 : P = {Pxxxx, Pxxxy, Pxyyx, Pxyyy, Pxyxx, Pxyxy, Pxyyx, Pxyyy};
i14 : f = map(R, Sp, P);

```

Minimal generating sets To compute the *Partially Distinguishing Minimal* (PDM) generating set for the ideal \bar{I}_T in the \bar{p}_{xyzt} coordinates, we compute a minimal generating set for the kernel of the homomorphism f .

```

i15 : I = kernel f;
i16 : M = mingens I;
i17 : netList flatten entries M
+-----+
o17 = |pxxxx + pxxxxy + pxxxyx + pxxxyy + pxyxx + pxyxy + pxyyx + pxyyy - 1 |
+-----+
|pxxxy*pxyxx - pxxxy*pxyxy + pxxxy*pxyyx + pxxxy*pxyyx + pxxxy*pxyyx + |
|pxyxx*pxyyx + pxyxy*pxyyx + pxyyx*pxyyx + pxxxy*pxyyy + pxyyx*pxyyy - pxyyx |
+-----+
|pxxxy*pxyxx + pxxxy*pxyxy + pxxxy*pxyxy + pxxxy*pxyxy + pxyxx*pxyxy + |
|pxyxy*pxyxy - pxxxy*pxyyx + pxyxy*pxyyx + pxxxy*pxyyy + pxyxy*pxyyy - pxyxy |
+-----+

```

The kernel computation yields three generators for \bar{I}_T : one linear, $\ell = \sum \bar{p}_{xyzt} - 1$, and two quadratic. The PDM set consists of the quadratics:

$$\begin{aligned} \bar{h}_1 &= \bar{p}_{xxyx}\bar{p}_{xyxx} - \bar{p}_{xxyy}\bar{p}_{xyxy} + \bar{p}_{xxyx}\bar{p}_{xyyx} + \bar{p}_{xxyx}\bar{p}_{xyyx} + \bar{p}_{xxyy}\bar{p}_{xyyx} + \bar{p}_{xyxx}\bar{p}_{xyyx} + \\ &\quad \bar{p}_{xyxy}\bar{p}_{xyyx} + \bar{p}_{xyyx}\bar{p}_{xyyx} + \bar{p}_{xxyx}\bar{p}_{xyyx} + \bar{p}_{xyyx}\bar{p}_{xyyx} - \bar{p}_{xyyx}, \text{ and} \\ \bar{h}_2 &= \bar{p}_{xxyx}\bar{p}_{xyxx} + \bar{p}_{xxyx}\bar{p}_{xyxy} + \bar{p}_{xxyx}\bar{p}_{xyxy} + \bar{p}_{xxyy}\bar{p}_{xyxy} + \bar{p}_{xyxx}\bar{p}_{xyxy} + \bar{p}_{xyxy}\bar{p}_{xyxy} - \\ &\quad \bar{p}_{xxyy}\bar{p}_{xyyx} + \bar{p}_{xyxy}\bar{p}_{xyyx} + \bar{p}_{xxyx}\bar{p}_{xyyx} + \bar{p}_{xyxy}\bar{p}_{xyyx} - \bar{p}_{xyxy}. \end{aligned}$$

The *Completely Distinguishing Minimal* (CDM) generating set is formed by the linear combinations $\bar{h}_1 + \bar{h}_2$ and $\bar{h}_1 - \bar{h}_2$.

Completely Distinguishing Determinantal generating set For group-based models such as the CFN, applying a linear change of coordinates (a Fourier or Hadamard transformation [52, 90]) is often advantageous. The new coordinates q_{xyzt} simplify the parametrization and the description of \bar{I}_T . For the CFN model on the tree $T = T_{12|34}$, this change of coordinates is as follows:

```

i18 : Sq = QQ[qxxxx, qxxyy, qxyxy, qxyyx, qyxxxy, qyxyx, qyyxx, qyyyy];

i19 : Qxxxx = Pxxxx + Pxxxxy + Pxxxyx + Pxxxyy + Pxyxxx + Pxyxy + Pxyyx + Pxyyy
o19 = 1
i20 : Qxxyy = Pxxxx - Pxxxxy - Pxxxyx + Pxxxyy + Pxyxxx - Pxyxy - Pxyyx + Pxyyy
o20 =  $\theta^3 \theta^4$ 
i21 : Qxyxy = Pxxxx - Pxxxxy + Pxxxyx - Pxxxyy - Pxyxxx + Pxyxy - Pxyyx + Pxyyy
o21 =  $\theta^2 \theta^4 \theta^5$ 
i22 : Qxyyx = Pxxxx + Pxxxxy - Pxxxyx - Pxxxyy - Pxyxxx - Pxyxy + Pxyyx + Pxyyy
o22 =  $\theta^2 \theta^3 \theta^5$ 
i23 : Qyxxxy = Pxxxx - Pxxxxy + Pxxxyx - Pxxxyy + Pxyxxx - Pxyxy + Pxyyx - Pxyyy
o23 =  $\theta^1 \theta^4 \theta^5$ 
i24 : Qyxyx = Pxxxx + Pxxxxy - Pxxxyx - Pxxxyy + Pxyxxx + Pxyxy - Pxyyx - Pxyyy
o24 =  $\theta^1 \theta^3 \theta^5$ 
i25 : Qyyxx = Pxxxx + Pxxxxy + Pxxxyx + Pxxxyy - Pxyxxx - Pxyxy - Pxyyx - Pxyyy
o25 =  $\theta^1 \theta^2$ 
i26 : Qyyyy = Pxxxx - Pxxxxy - Pxxxyx + Pxxxyy - Pxyxxx + Pxyxy + Pxyyx - Pxyyy
o26 =  $\theta^1 \theta^2 \theta^3 \theta^4$ 

```

The generating set for the ideal \bar{I}_T in the q_{xyzt} coordinates is found by computing the kernel of g .

```

i27 : Q = {Qxxxx, Qxxyy, Qxyxy, Qxyyx, Qyxxxy, Qyxyx, Qyyxx, Qyyyy};
i28 : g = map(R, Sq, Q);
i29 : netList entries gens kernel g
+-----+-----+-----+-----+
o29 = |qxxxx - 1|qxxyy*qyyxx - qyyyy|qxyyx*qyxxxy - qxyxy*qyxyx|
+-----+-----+-----+-----+

```

Transforming back to the \bar{p}_{xyzt} probability coordinates, the linear polynomial $q_{xxxx} - 1$ becomes $\ell = \sum \bar{p}_{xyzt} - 1$, and the two quadratics yield the *Completely Distinguishing Determinantal* (CDD) set.

```

i30 : qxxxx = 1; -- pxxxx + pxxxxy + pxxxyx + pxxxyy + pxxyxx + pxxyxy + pxxyyx + pxxyyy
i31 : qxxyy = pxxxx - pxxxxy - pxxxyx + pxxxyy + pxxyxx - pxxyxy - pxxyyx + pxxyyy;
i32 : qxyxy = pxxxx - pxxxxy + pxxxyx - pxxxyy - pxxyxx + pxxyxy - pxxyyx + pxxyyy;
i33 : qxyyx = pxxxx + pxxxxy - pxxxyx - pxxxyy - pxxyxx - pxxyxy + pxxyyx + pxxyyy;
i34 : qyxxy = pxxxx - pxxxxy + pxxxyx - pxxxyy + pxxyxx - pxxyxy + pxxyyx - pxxyyy;
i35 : qyxxy = pxxxx + pxxxxy - pxxxyx - pxxxyy + pxxyxx + pxxyxy - pxxyyx - pxxyyy;
i36 : qyyxx = pxxxx + pxxxxy + pxxxyx + pxxxyy - pxxyxx - pxxyxy - pxxyyx - pxxyyy;
i37 : qyyyy = pxxxx - pxxxxy - pxxxyx + pxxxyy - pxxyxx + pxxyxy + pxxyyx - pxxyyy;

i38 : M1 = matrix{{qxxxx, qxxyy},
                  {qyyxx, qyyyy}}

i39 : M2 = matrix{{qxyxy, qyxxy},
                  {qxyyx, qyxxy}}

i40 : F1 = det(M1)
o40 = - pxxxx^2 + pxxxxy^2 + 2pxxxxy*pxxxyx + pxxxyx^2 - 2pxxxx*pxxxyy - pxxxyy^2 -
      2pxxxxy*pxxyxx - 2pxxxxy*pxxyxx + pxxyxx^2 + 2pxxxx*pxxyxy + 2pxxxxy*pxxyxy -
      pxxyxy^2 + 2pxxxx*pxxyyx + 2pxxxxy*pxxyyx - 2pxxyxy*pxxyyx - pxxyyx^2 -
      2pxxxxy*pxxyyy - 2pxxxxy*pxxyyy + 2pxxyxx*pxxyyy + pxxyyy^2 + pxxxx - pxxxxy -
      pxxxyx + pxxxyy - pxxyxx + pxxyxy + pxxyyx - pxxyyy

i41 : F2 = det(M2)
o41 = - 4(pxxxxy*pxxyxx - pxxxyx*pxxyxx - pxxxx*pxxyxy + pxxxyy*pxxyxy +
      pxxxx*pxxyyx - pxxxyy*pxxyyx - pxxxxy*pxxyyy + pxxxyx*pxxyyy)

```

Note that for the F_i defined in lines i38 and i39 of the code above $\langle \psi(F_i) \rangle + L_{\text{symm}} = \langle f_i \rangle + L_{\text{symm}}$, where f_1, f_2 are the polynomials of Eq. (4.4.1). In other words, up to the symmetries in L_{symm} and a constant factor, $\psi(F_i)$ is the same as f_i , $i = 1, 2$.

Rank generating sets The probabilities p_{ijkl} for the tree T can be arranged into a 4×4 matrix according to the partition 12|34 of its leaves, where rows are indexed by the states of leaves 1, 2 and columns by the states of 3, 4:

$$\text{Flat}_{12|34}(p) = \begin{pmatrix} p_{0000} & p_{0001} & p_{0010} & p_{0011} \\ p_{0100} & p_{0101} & p_{0110} & p_{0111} \\ p_{1000} & p_{1001} & p_{1010} & p_{1011} \\ p_{1100} & p_{1101} & p_{1110} & p_{1111} \end{pmatrix} = \frac{1}{2} \begin{pmatrix} \bar{p}_{xxxx} & \bar{p}_{xxxxy} & \bar{p}_{xxyyx} & \bar{p}_{xxyyy} \\ \bar{p}_{xyxx} & \bar{p}_{xyxy} & \bar{p}_{xyyx} & \bar{p}_{xyyy} \\ \bar{p}_{xyyy} & \bar{p}_{xyyx} & \bar{p}_{xyxy} & \bar{p}_{xyxx} \\ \bar{p}_{xxyy} & \bar{p}_{xxyx} & \bar{p}_{xxxxy} & \bar{p}_{xxxx} \end{pmatrix}.$$

The matrix $\text{Flat}_{12|34}(p)$ has rank at most 2, reflecting a conditional independence statement holding for leaves separated by the central edge of T [9]. Therefore, its 3×3 minors are polynomials in the CFN ideal I_T . We use this to construct the *Partially Distinguishing Rank* (PDR) set, working from the matrix $2 \cdot \text{Flat}_{12|34}(p)$ expressed in the \bar{p} variables.

```
i44 : Flat1234 = matrix{{pxxxx, pxxxy, pxxyx, pxxyy},
                        {pxyxx, pxyxy, pxyyx, pxyyy},
                        {pxyyy, pxyyx, pxyxy, pxyxx},
                        {pxxyy, pxxyx, pxxxy, pxxxx}};

i45 : I = minors(3, Flat1234);
i46 : netList primaryDecomposition I
+-----+
o46 = |ideal (pxxyx*pxyxx - pxxyy*pxyxy - pxxxx*pxyyx + pxxxy*pxyyy, |
|      (pxxyy*pxyxx - pxxxx*pxyxy - pxxyy*pxyyx + pxxyx*pxyyy) |
+-----+
|ideal (pxyxy - pxyyx, pxyxx - pxyyy, pxxxy - pxxyx, pxxxx - pxxyy) |
+-----+
|ideal (pxyxy + pxyyx, pxyxx + pxyyy, pxxxy + pxxyx, pxxxx + pxxyy) |
+-----+
```

The ideal generated by all 3×3 minors of $2 \cdot \text{Flat}_{12|34}(p)$ (Flat1234 in the code) is not prime. The first component in the primary decomposition o46 corresponds to the CFN model. The quadratic polynomials from this component form the PDR set:

$$\begin{aligned}\bar{g}_1 &= \bar{p}_{xxyx}\bar{p}_{xyxx} - \bar{p}_{xxyy}\bar{p}_{xyxy} - \bar{p}_{xxxx}\bar{p}_{xyyx} + \bar{p}_{xxxy}\bar{p}_{xyyy}, \\ \bar{g}_2 &= \bar{p}_{xxxy}\bar{p}_{xyxx} - \bar{p}_{xxxx}\bar{p}_{xyxy} - \bar{p}_{xxyy}\bar{p}_{xyyx} + \bar{p}_{xxyx}\bar{p}_{xyyy}.\end{aligned}$$

The *Completely Distinguishing Rank* (CDR) set consists of the polynomials $\bar{g}_1 + \bar{g}_2$ and $\bar{g}_1 - \bar{g}_2$.

A.3.2 Additional results for Collection 1: Comparison of different constraint sets

We provide additional results on the performance of the SDL test on data from Collection 1, supplementing Sections 4.4.4 and 4.4.5 of the main text.

Aggregated p -value histograms

We analyse the performance of the five different choices of model constraints introduced in Section 4.4.1 by aggregating p -values across Collection 1. Fig. A.3 and Fig. A.4 are analogous to the left and right parts of Fig. 4.10 in the main text, but also include the CDD, CDR and PDR constraints. These figures further support that the test behaviour is affected by the choice of model description.




Figure A.3: Aggregated p -values for a test of the true null hypothesis $H_{12|34}$ from datasets in Collection 1. Columns correspond to choices of defining polynomials. Rows correspond to the value of r .

In the $r = 0$ case, Fig. A.3 shows that the partially distinguishing sets PDR and PDM do not produce conservative tests due to an excess of small p -values when testing $H_{12|34}$. On the other hand, for both CDD and CDR, the p -value distribution appears to be close to uniform, and CDM gives an especially conservative test, with an excess of large p -values (as previously seen in Section 4.4.4). Overall, completely distinguishing polynomials seem to produce a conservative test when $r = 0$. In Fig. A.4 we observe that for $r = 0$, the completely distinguishing constraints had slightly greater concentration of p -values near zero compared to the partially distinguishing constraints, similarly to what we observed in Section 4.4.4.

The effects of varying r in both figures are relatively minor, and whether the effect was beneficial or not depended on whether the initial choice of constraints was partially distinguishing or completely distinguishing. When only partially distinguishing constraints were used, adding convex combinations improved performance by increasing the number of small p -values when testing the wrong model parameter (see Fig. A.4). The beneficial effect observed in Fig. 4.10

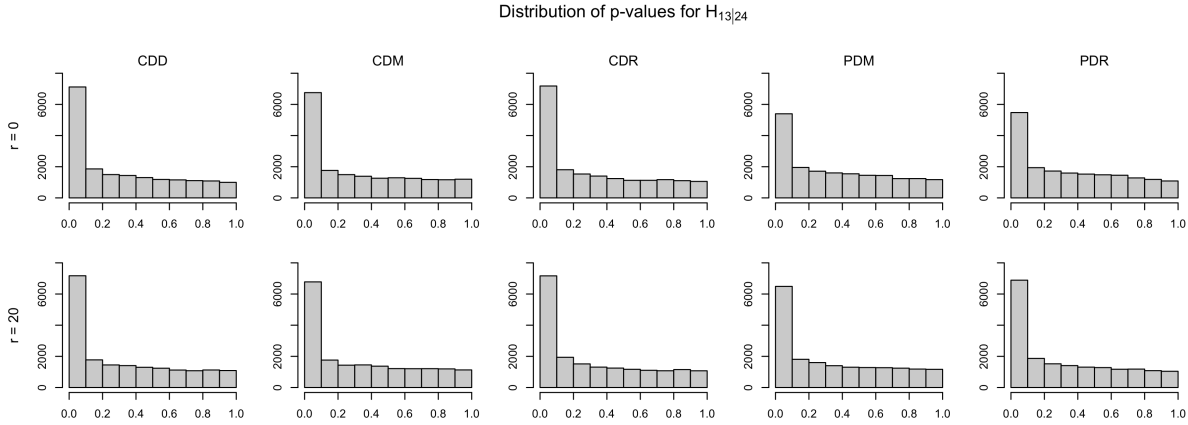


Figure A.4: Aggregated p -values for a test of $H_{13|24}$ (a false null hypothesis) from datasets in Collection 1. Columns correspond to choices of defining polynomials. Rows correspond to the value of r . The test of $H_{14|23}$ produced similar results.

that increasing r made the test more conservative for PDM was not similarly observed for PDR. On the other hand, when completely distinguishing constraints were used, adding convex combinations constraints risks negatively affecting the quality of the p -values. Evidence for this can be seen in Fig. A.3, which shows that for CDR, the test appears to be conservative when $r = 0$, but not when $r = 20$ due to an increased proportion of small p -values.

SDL-based Phylogenetic Inference

We analyse the performance of the SDL phylogenetic inference method for all five constraint sets in Fig. A.5, which is analogous to the left part of Fig. 4.13, but includes the additional sets CDD, CDR, and PDR.

The conclusions from this figure are comparable to those of Section 4.4.5. First, in the case $r = 0$, the use of completely distinguishing constraint sets yields better performance than partially distinguishing sets (viz., columns 1,2,3, which have larger dark region and higher success percentages than columns 4,5). The completely distinguishing sets CDD, CDM, and CDR all performed similarly: the differences in p -value distributions among them observed in Fig. A.3 appeared to have no bearing on their performance for inference in this setting.

The second important conclusion from Fig. A.5 is that the performance of the partially distinguishing generating sets PDR and PDM increased when r was increased from 0 to 20. Indeed, as a result of this improvement, all five sets performed comparably in the $r = 20$ case. This improvement in performance for PDM and PDR is consistent with our observations in

Fig. A.4, that — at least for partially distinguishing constraints — increasing r appeared to increase the power of the test.

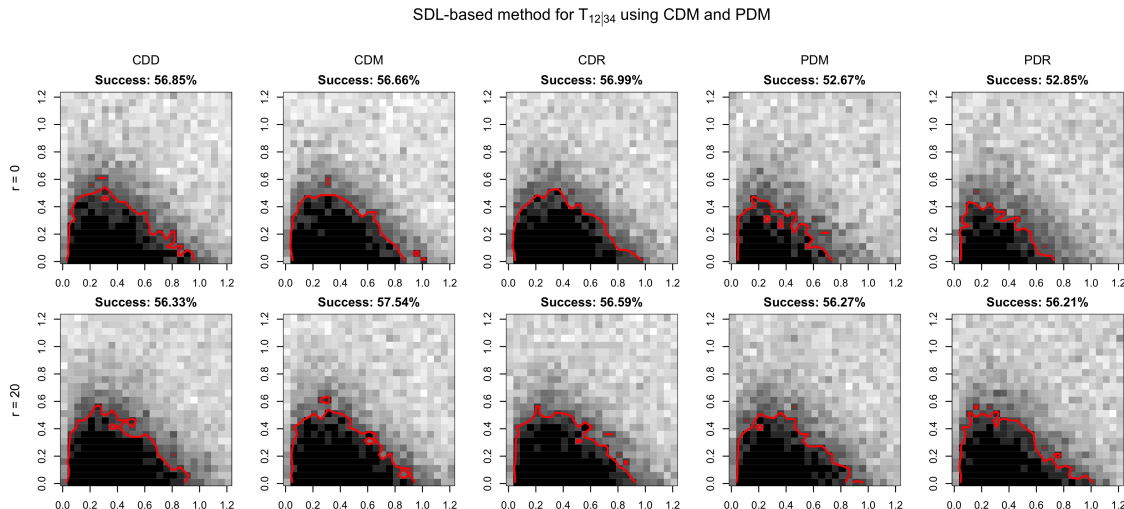


Figure A.5: Performance of the SDL test for inferring the tree topology $T_{12|34}$. Columns correspond to different CFN model constraints (CDD, CDM, CDR, PDM, PDR), and rows represent the number of convex combinations used, $r = 0$ and $r = 20$. Grey levels represent the frequency of correctly inferring the topology for edge length pairs (a, b) (black 100%, white 0%).

Note that variations in performance using different algebraic constraint sets for inference were previously observed [82], with symmetrizing ideal generators improving model selection.

A.3.3 Lack of long branch attraction bias

In this section we analyse the SDL test's behaviour for trees in the Felsenstein zone (see Fig. 4.9), showing it differs from that of common methods used for phylogenetic inference. In particular, maximum parsimony [46] exhibits a *long-branch attraction* bias in this region, in which the false topology $T_{13|24}$, pairing the two taxa on long pendent edges, is most frequently inferred. Similar bias is observed for maximum likelihood [92, 78] and previous algebraic methods [48].

In Fig. A.6, we present p -values obtained from the SDL test using data generated from one tree with Felsenstein zone parameters $a = 0.8$, $b = 0.05$, with $n = 10,000$. We compared the SDL test using two different sets of constraints: CDM (left plots) and PDM (right plots); in both cases the internal edge inequality of Eq. (4.4.3) was also used. We tested the three null hypotheses $H_{12|34}$, $H_{13|23}$, and $H_{14|23}$ (plot columns) for $r = 0$ and 20 (plot rows).

The choice of the CDM versus PDM constraints produces a marked discrepancy in test behaviour, especially for $r = 0$. The first row of Fig. A.6 ($r = 0$) shows that the SDL test is much more likely to reject $H_{13|24}$ than $H_{14|23}$ for small test levels when using the CDM constraints; on the other hand, the two false hypotheses are rejected at roughly equal frequency with the PDM constraints. Both of these behaviours are in contrast with classical phylogenetic inference methods, which would tend to strongly support $H_{13|24}$ over $H_{14|23}$. Constraints CDM and PDM produce almost-uniform distributions of p -values when testing $H_{12|34}$.

The second row of Fig. A.6 shows that the addition of $r = 20$ convex combinations for both the CDM and PDM constraints reduced the asymmetry between test results of $H_{13|24}$ and $H_{14|23}$, and gave a more powerful test. Moreover, the test remained conservative for all values of r . However, a slight bias for $H_{13|24}$ appears, but only for the PDM constraints.

In contrast to Fig. 4.10, which showed increasing r had little effect on the aggregated p -value distribution over a larger set of parameters, Fig. A.6 indicates that for certain parameter values, incorporating convex combinations can have a major effect — in particular, by increasing the power of the SDL test.

The general lack of bias toward $H_{13|24}$, together with the overall conservativeness of the test, indicates that the SDL test can perform quite well in the Felsenstein zone. Furthermore, the differing p -value distributions between CDM and PDM underscore how the choice of constraint sets can significantly impact SDL test performance.

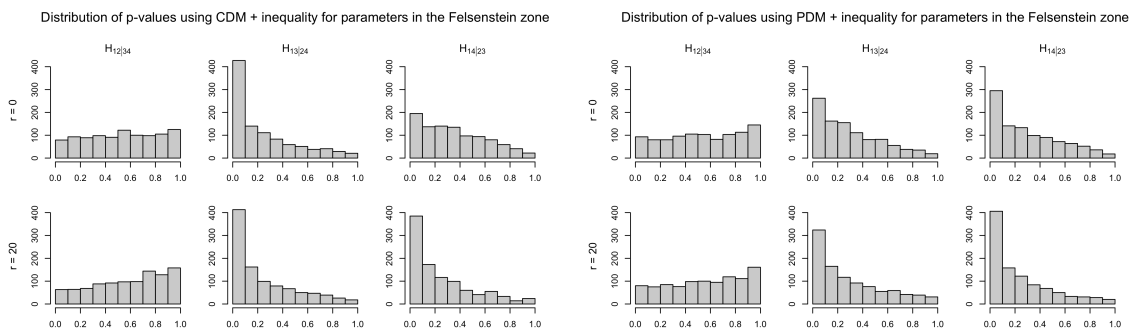


Figure A.6: Histogram of p -values for CDM (left) and PDM (right) for a tree in the Felsenstein zone ($a = 0.8$ and $b = 0.05$) with $n = 10000$ bp and $m = 12$.

The reduced long branch attraction bias for SDL is not unique to the parameters used for Fig. A.6, but persists across the Felsenstein zone. In Table A.1, we present the percentage of times that each of the three possible quartet topologies is inferred by the SDL, MLE and SVD

methods from data in Collection 1, both across the whole treespace shown in Fig. 4.9 and only across the Felsenstein zone.

These results show that for all three inference methods the topology $T_{13|24}$ was inferred more frequently than $T_{14|23}$ across the full parameter space, but especially in the Felsenstein zone. However, the SDL-based method showing the least susceptibility to this preference. In the Felsenstein zone, there is an extreme bias for the SVD method, with $T_{13|24}$ inferred 46.07% of the time, even more frequently than the 37.13% for the true $T_{12|34}$. For MLE, the effect was less pronounced, although $T_{13|24}$ was inferred noticeably more often than $T_{14|23}$ (15.74% vs. 10.35%). For the SDL-based method (using the CDD constraints with $r = 20$), the imbalance was proportionally smallest among the three methods (21.78% vs. 18.02%).

	Treespace				Felsenstein Zone			
	12 34	13 24	14 23	<i>undecided</i>	12 34	13 24	14 23	<i>undecided</i>
SDL	69.51%	15.8%	14.69%	-	60.2%	21.78%	18.02%	-
MLE	78.25%	8.86%	6.95%	5.94%	71.88%	15.74%	10.36%	2.01%
SVD	65.82%	19.64%	14.54%	-	37.13%	46.07%	16.8%	-

Table A.1: Estimated tree topologies for the three methods SDL, MLE and SVD and the three topologies 12|34, 12|34 and 12|34 in the entire treespace of Fig. 4.9 and in the Felsenstein Zone. The *undecided* column reports the percentage of times that MLE fails to distinguish between topologies.

A.4 Technical Assumptions

In order for the SDL test to be asymptotically valid for a particular hypothesis testing problem, there are a number of technical assumptions which need to be satisfied. In this section, we state the six conditions assumed in [89], and verify that they hold for the models considered in Chapter 4. Despite their technical nature, these conditions are all straightforward to verify for the models we consider.

To state the conditions, let $X_1, \dots, X_m \sim P_\theta$ be iid random variables, and let $\mu = (\mu_1, \dots, \mu_p)^\top := \mathbb{E}[h(X_1, \dots, X_m)]$. In addition, define $\sigma_{h,j}^2 := \mathbb{E}[(h_j(X_1, \dots, X_m) - \mu_j)^2]$ and $\sigma_{g,j}^2 := \mathbb{E}[(g_j(X_1, \dots, X_m) - \mu_j)^2]$. For any positive β , define the function $\psi_\beta(x) = \exp(x^\beta) - 1$, and for any random variable Y define $\|Y\|_{\psi_\beta} := \inf\{t > 0 : \mathbb{E}[\psi_\beta(|Y|/t)] \leq 1\}$.

The theoretical results in [89] assume that there exists a constant $\beta \in (0, 1]$ and a sequence $D_1, D_2, \dots \geq 1$ such that:

$$(C1) \quad \mathbb{E}[|h_j(X_1, \dots, X_m) - \mu_j|] \leq \sigma_{h,j}^2 D_n^l \text{ for all } j = 1, \dots, p \text{ and } l = 1, 2.$$

$$(C2) \quad \|h_j(X_1, \dots, X_m) - \mu_j\|_{\psi_\beta} \leq D_n \text{ for all } j = 1, \dots, p.$$

$$(C3) \quad \text{There exists } \underline{\sigma}_h^2 > 0 \text{ such that } \underline{\sigma}_h^2 \leq \min_{1 \leq j \leq p} \sigma_{h,j}^2.$$

$$(C4) \quad \text{There exists } \underline{\sigma}_{g^{(1)}}^2 > 0 \text{ such that } \underline{\sigma}_{g^{(1)}}^2 \leq \min_{1 \leq j \leq p_1} \sigma_{g,j}^2 \text{ for some positive integer } p_1 < p.$$

$$(C5) \quad \text{There exists } k \text{ such that } \|g_j(X_1) - \mu_j\|_{\psi_\beta} \leq n^{-k} D_n \text{ for all } j = p_1 + 1, \dots, p.$$

$$(C6) \quad \mathbb{E}[|g_j(X_1) - \mu_j|^{2+l}] \leq \sigma_{g,j}^2 D_n^l \text{ for all } j = 1, \dots, p \text{ and } l = 1, 2.$$

In the above conditions, it is furthermore assumed that $2 \leq m \leq \sqrt{n}$, $n \geq 4$, $p \geq 3$. (Note that in Model 1 of Section 4.3, we have only $p = 2$, but the assumption that $p \geq 3$ is not strictly necessary; for more detailed discussion of these assumptions, see [89, Section 2.1]).

Next, we check that conditions (C1)-(C6) hold:

- First observe that condition (C3) is satisfied whenever $h_j(X_1, \dots, X_m)$ is not almost surely constant, which is straightforward to check for all the examples considered in Chapter 4, since in all our examples X_1 takes the form of a multinomial random variable with a single trial.
- Second, for the examples considered in Chapter 4, the state space \mathcal{S} of X_1 is always a finite set, and hence $h_j(X_1, \dots, X_m)$ is almost surely bounded. Together with (C3), this implies that we can choose finite D_n satisfying

$$D_n \geq \left(\max_{\substack{1 \leq j \leq p \\ x_1, \dots, x_m \in \mathcal{S}}} \frac{|h_j(x_1, \dots, x_m) - \mu_j|}{\sigma_{h,j}^2} \right) \vee 1$$

for all $n \geq 1$. Moreover, for this choice of D_n , condition (C1) holds

- Next we show that by possibly making each D_n larger, it is possible to find D_n large enough that (C6) is also satisfied. On the one hand, if $\sigma_{g,j}^2 = 0$ then the inequality in (C6) holds trivially with both sides equal to zero. On the other hand, for j with $\sigma_{g,j}^2 > 0$, the inequality

in (C6) is satisfied if

$$D_n \geq \max_{l \in \{1,2\}} \left(\max_{j: \sigma_{g,j}^2 > 0} \max_{x_1 \in \mathcal{S}} \frac{|g_j(x_1) - \mu_j|^{2+l}}{\sigma_{g,j}^2} \right)^{\frac{1}{l}},$$

and without loss of generality we can assume this inequality holds since right-hand side is finite (due to the maximums being taken over finite sets).

- Furthermore, we will show that the terms of the sequence D_1, D_2, \dots can also be chosen large enough to satisfy (C2). To see this, write $Y = h_j(X_1, \dots, X_m)$ and observe that since Y has finite state space, there exists finite C_j such that $|Y| \leq C_j$ almost surely. It then follows by definition of $\|\cdot\|_{\psi_\beta}$ that $\|Y\|_{\psi_\beta} \leq C_j / \sqrt[\beta]{\log(2)}$. Again, without loss of generality, D_n may be chosen so that $D_n \geq C_j / \sqrt[\beta]{\log(2)}$ which is sufficient to imply (C2).
- Finally, it remains to consider conditions (C4) and (C5), which together are referred to as the *mixed degeneracy* conditions in [89]. In fact, there is nothing to show: for all the examples considered in Chapter 4, we have $p = O(1)$ as $n \rightarrow \infty$, and as a consequence of this, conditions (C4) and (C5) hold trivially, as discussed in [89, Section 2.1].