ABSTRACT

BERNSTEIN, DANIEL IRVING. Matroids in Algebraic Statistics. (Under the direction of Seth Sullivant.)

Algebraic statistics is a relatively new field of research, broadly concerned with connections between algebraic geometry and statistics. This dissertation addresses problems in three subfields of this emerging area: low-rank matrix completion, phylogenetics, and discrete hierarchical models. The unifying theme among the problems addressed is that the behavior we aim to understand is governed by a matroid.

In a low-rank matrix completion problem, one observes a subset of entries of a matrix and wishes to reconstruct the missing entries such that the rank of the completed matrix is minimized, or equal to some fixed number. In Chapter 2, we lay some theoretical groundwork for using algebraic geometry to solve these types of problems. In particular, we characterize the algebraic matroid underlying the variety of $m \times n$ matrices of rank at most two. This characterization is a consequence of a characterization we provide of the algebraic matroid underlying the variety of $n \times n$ skew-symmetric matrices of rank at most two. To obtain this skew-symmetric characterization, we use tropical geometry to reduce the problem to a question about tree metrics which we then solve.

A fundamental problem of phylogenetics is to infer the evolutionary history among a set of species. In the distance-based approach to this problem, the data consists of some measure of distance between each pair of species and the outputted evolutionary history may be the tree metric or ultrametric nearest to the dataset according to some norm. Due to the tropical structure of the sets of tree metrics and ultrametics, the l^∞-norm is a natural choice. However, there can be multiple tree metrics and ultrametrics l^∞-nearest to a given dataset. We study this phenomenon in Chapter 3. Non-uniqueness of l^∞-nearest tree metrics and ultrametrics is partially due to the fact that the point in a linear subspace $L \subseteq \mathbb{R}^n$ which is l^∞-nearest to a given $x \in \mathbb{R}^n$ is not always unique. Hence we demonstrate how the oriented matroid underlying a linear subspace $L \subseteq \mathbb{R}^n$ can be used to compute the dimension of the subset of L consisting of points l^∞-nearest to a given $x \in \mathbb{R}^n$. A consequence is that the point in $L \subseteq \mathbb{R}^n$ which is l^∞-nearest to a given $x \in \mathbb{R}^n$ is unique for all such x if and only if the matroid underlying L is uniform.

The discrete hierarchical models form a class of log-linear models that are indexed by the pairs $(\mathcal{C}, \mathbf{d})$ consisting of a simplicial complex $\mathcal{C}$ with vertex weights $\mathbf{d}$. In Chapter 4,

we classify the $(\mathcal{C}, \mathbf{d})$ whose corresponding hierarchical model is unimodular. We begin by classifying the unimodular simplicial complexes, which are the simplicial complexes $\mathcal{C}$ such that the hierarchical model corresponding to $(\mathcal{C}, (2, \ldots, 2))$ is unimodular.

Our classification of the unimodular simplicial complexes is given constructively, and in terms of forbidden minors. We identify a handful of operations that transform a unimodular simplicial complex into a larger unimodular simplicial complex. We then identify three families of simplicial complexes from which all unimodular simplicial complexes can be constructed via these operations. A minor of a simplicial complex $\mathcal{C}$ is a simplicial complex $\mathcal{D}$, obtainable from $\mathcal{C}$ via a sequence of vertex links and vertex deletions. We show that any minor of a unimodular simplicial complex is unimodular, and identify which minors are forbidden for unimodular simplicial complexes. We extend our classification of the unimodular $\mathcal{C}$ to classify the pairs $(\mathcal{C}, \mathbf{d})$ giving rise to unimodular hierarchical models. We use our constructive characterization of the unimodular hierarchical models to describe the Graver basis of any unimodular hierarchical model.

Matroids in Algebraic Statistics

by
Daniel Irving Bernstein

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

2018

APPROVED BY:

_____   _____
          Dávid Papp   Agnes Szanto

_____   _____
        Cynthia Vinzant   Seth Sullivant
                                              Chair of Advisory Committee

DEDICATION

To my parents

BIOGRAPHY

Daniel Irving Bernstein graduated from Davidson College in Spring 2013 and began his Ph.D. at North Carolina State University the following fall. He will continue his mathematical career as an NSF Mathematical Sciences Postdoctoral Research Fellow at the Massachusetts Institute of Technology, and a Semester Postdoctoral Fellow at the Institute for Computational and Experimental Research in Mathematics at Brown University. When he is not doing mathematics, Daniel enjoys playing music, reading, powerlifting, and running.

ACKNOWLEDGEMENTS

I have so many people to thank for helping me reach this point. First and foremost is my advisor, Seth Sullivant. His high expectations, detailed feedback, and generosity with his time, attention, and resources have driven me to produce far more results and of a much higher quality than the me of five years ago thought I was capable of.

The applied algebraic geometry community is extremely welcoming to its new members. I was fortunate enough to enjoy this generosity of spirit early on, not just from Seth, but from the postdocs and older graduate students in Seth's group. I would especially like to thank Ruth Davidson, Elizabeth Gross, Colby Long, and Nikki Meshkat for all the conversations we had about the unique stresses associated with beginning a career in academia. Other mathematicians whose mentoring has been particularly helpful are Louis Theran, who introduced me to low-rank matrix completion, and Cynthia Vinzant, who provided me with valuable feedback on my teaching. I am also greatly indebted to my letter-writers - Greg Blekherman, Jesus De Loera, Seth Sullivant, Louis Theran, Cynthia Vinzant, and Josephine Yu - and my committee, Dávid Papp, Seth Sullivant, Agnes Szanto, and Cynthia Vinzant.

My coauthors have played no small role in my mathematical development. They have introduced me to new problems, taught me clever mathematical tricks, and provided me with the inspiration necessary to finish our projects. I would especially like to thank Greg Blekherman, Colby Long, Christopher O'Neill, Rainer Sinn, Katherine St. John, and Seth Sullivant.

Without the mentoring I received from the mathematics department at Davidson College, it is possible that I would not have decided to go to graduate school. I owe a particular debt of gratitude to Timothy Chartier, Richard Neidinger, and Carl Yerger, who each mentored me on an undergraduate research project.

Of the friends I made while in graduate school, I owe a particularly big thanks to my former officemate and lifting buddy Alan Liddell and his wife Jency. They helped ease the growing pains of early adulthood I experienced after first moving to Raleigh, a city where I knew nobody.

I am grateful to my family for their constant support. My parents provided me with every possible opportunity to succeed without ever putting pressure on me. My father, whose journey towards becoming a physician included an unexpected eight-year detour

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

Introduction

Algebraic statistics is the research area concerned with identifying and exploiting connections between algebraic geometry and statistics. This is a relatively new research area, with its first paper [24] being published in 1998. Since its inception, algebraic statistics has lead not only to new algorithms and research directions in statistics, but in algebraic geometry and and related areas of mathematics as well. For a current sampling of the algebraic statistics landscape, see the textbook-in-press [60].

This thesis solves problems coming from low-rank matrix completion, distance-based phylogenetic reconstruction, and discrete log-linear models. While these applications may seem disjointed, the problems solved therein are unified by the fact that some important underlying structure is a matroid.

Low-rank matrix completion

Chapter 2 concerns the low-rank matrix completion problem. The content of this chapter was published in *Linear Algebra and its Applications* [6].

The classic low-rank matrix completion problem begins with a matrix where only a subset of the entries are known and asks for the values that the missing entries should take if the rank is to be minimized. Perhaps the most famous instance of this is the "Netflix problem" which is a special instance of *collaborative filtering* (see e.g. [16]). Variations exist where one wishes to complete a partial matrix to a given fixed rank r. For example, by completing a certain partial matrix to rank two or three, one can solve instances of the *global positioning problem* [54, 14, 55] and the *structure from motion problem* [61, 19].

If M is an $n \times n$ matrix of rank r satisfying a genericity assumption called "incoher-

ence," then with high probability, one can use semidefinite programming to recover M from just $\Theta(n^{1.25} r \log n)$ entries chosen uniformly at random [17, 18]. These assumptions are valid in many applications but another approach is needed for when they are not valid. This motivated Király, Theran, and Tomioka to develop a new approach using methods of algebraic geometry and matroid theory in 2015 [41], building on a connection with rigidity theory noted by Singer and Cucuringu in 2010 [55]. Improving their approach requires solutions to a slew of interesting mathematical problems, the most basic of which is to characterize the algebraic matroids underlying certain determinantal varieties. Chapter 2 gives such a characterization for the case of non-symmetric and skew-symmetric matrices of rank two (Theorems 2.3.2 and 2.3.4). This characterization is obtained by using tropical geometry to translate our question about characterizing this algebraic matroid into a question about phylogenetic trees.

Phylogenetics and linear spaces

Chapter 3 concerns distance-based phylogenetic reconstruction in the l^∞-norm and a related problem about linear spaces. The content of this chapter was joint work with Colby Long and it was published in *SIAM Journal on Discrete Mathematics* [8].

Distance-based methods for phylogenetic reconstruction aim to infer the evolutionary relationships among a set of species from the set of all observed "distances" between each pair. Certain cases admit a geometric interpretation wherein one views their dataset of observed distances as a point in some high-dimensional space and wishes to find the "closest" point that lies within a certain polyhedral complex. Of course, one can choose any metric in which to find this "closest" point. Connections between phylogenetics and tropical geometry [4, 5, 56] suggest that one should investigate the l^∞-metric in this context. However, a peculiarity that arises here is that the closest point within this polyhedral complex often fails to be unique. Chapter 3 investigates this failure of uniqueness with the aim of quantifying what is possible.

To begin this investigation, we consider a simpler mathematical question. Namely, given a linear subspace $L \subseteq \mathbb{R}^n$ and a point $x \in \mathbb{R}^n$, what is the dimension of the subset of L consisting of points l^∞-nearest to x? We then use the (oriented) matroid underlying L to give a polyhedral decomposition of $\mathbb{R}^n$ such that any two points in the same cell have the same dimension of their set of l^∞-nearest neighbors in L (Theorem 3.1.7). A consequence of this decomposition is that every $x \in \mathbb{R}^n$ has a unique l^∞-closest point in

L if and only if the matroid underlying L is uniform (Theorem 3.1.9).

Unimodular hierarchical models

Chapter 4 concerns the classification of the unimodular hierarchical models. Part of the content of this chapter was joint work with Seth Sullivant, published in *Journal of Combinatorial Theory, Series B* [12] and the rest was joint work with Chris O'Neill, published in *Journal of Algebraic Statistics* [10].

The *hierarchical models* form a family of discrete log-linear models that are useful for categorical data analysis. In particular, they include the family of discrete graphical models [63]. Such models are naturally indexed by pairs $(\mathcal{C}, \mathbf{d})$ where $\mathcal{C}$ is a simplicial complex whose ground set is in bijection with the random variables in the model, and $\mathbf{d}$ is a vector giving the number of states of each random variable. One can ask: what useful properties of a hierarchical model can be inferred from the combinatorics of the pair $(\mathcal{C}, \mathbf{d})$?

Chapter 4 investigates one particular property that such a model may satisfy: unimodularity. The main result here is a complete classification of *all* unimodular hierarchical models in terms of the combinatorics of the pairs $(\mathcal{C}, \mathbf{d})$ (Theorem 4.9.1). This classification enables us to give a combinatorial description of the Graver basis of any unimodular hierarchical model (Remark 4.10.4). For unimodular hierarchical models, a description of the Graver basis is equivalent to a description of the oriented matroid underlying a particular matrix associated to such a model.

Outline

The remaining sections in this chapter provide the elementary background on a number of concepts that play a prominent role in this thesis. Readers familiar with any of the concepts therein should feel comfortable skipping the corresponding section.

Section 1.1 reviews some basic concepts from algebraic geometry. Its content is required for all the later introductory sections aside from Section 1.4. As the title indicates, matroids, which are introduced in Section 1.2, play an important role in this thesis. However, the introductory material we provide is only really needed for Chapters 2 and 3. Oriented matroids play a major role in Chapter 3 and they are introduced in Section 1.3. They are also used in Chapter 4, but only at the end. Tree metrics are introduced in

Section 1.4 and used in Chapters 2 and 3. Tropical geometry is introduced in Section 1.5 and used in Chapter 2. Toric ideals are introduced in Section 1.6 and used in Chapter 4.

1.1   Ideals and varieties

This section gives the required background on algebraic geometry. We expect that the reader has seen most, if not all, of the concepts described here so we proceed rather quickly. All rings in this thesis are assumed to be commutative and unitial. Let $\mathbb{K}$ be a field and let $\mathbb{K}[x_1, \ldots, x_n]$ denote the polynomial ring over $\mathbb{K}$ in n indeterminates. We will use the shorthand $\mathbf{x}^u := x_1^{u_1} \cdots x_n^{u_n}$ to denote monomials in $\mathbb{K}[x_1, \ldots, x_n]$ and $\mathbb{K}[\mathbf{x}] := \mathbb{K}[x_1, \ldots, x_n]$ for the ring itself. Each polynomial $f \in \mathbb{K}[\mathbf{x}]$ defines a function $f : \mathbb{K}^n \to \mathbb{K}$, sending $a := (a_1, \ldots, a_n) \in \mathbb{K}^n$ to $f(a) := f(a_1, \ldots, a_n)$.

Definition 1.1.1. Given a polynomial $f \in \mathbb{K}[x_1, \ldots, x_n]$, the *hypersurface defined by f*, denoted $V(f)$, is defined to be the subset of $\mathbb{K}^n$ consisting of points that evaluate to zero when plugged in to f. That is,

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

An *affine variety* is a (possibly empty) intersection of sets of the form $V(f)$.

We will often drop the qualifier "affine" as all varieties considered will be assumed so unless otherwise stated. We use the shorthand $V(f_1, \ldots, f_r)$ to denote $V(f_1) \cap \cdots \cap V(f_r)$. Moreover, in cases of a small number of variables, we may substitute different letters for $x_1, \ldots, x_n$. For example for $n = 3$, we may write $\mathbb{K}[x, y, z]$ instead of $\mathbb{K}[x_1, x_2, x_3]$.

Definition 1.1.2. An *ideal* in a ring R is a subset $I \subseteq R$ such that

1. if $f, g \in I$ than $f + g \in I$, and
2. if $f \in I$ and $h \in R$, then $hf \in I$.

Given a subset F of a ring R, we denote by $\langle F \rangle$ the ideal generated by F. That is,

$$\langle F \rangle := \{g_1 f_1 + \cdots + g_r f_r : f_1, \ldots, f_r \in F, g_1, \ldots, g_r \in R\}.$$

When $F = \{f_1, \ldots, f_r\}$ is a finite set, we may write $\langle f_1, \ldots, f_r \rangle := \langle F \rangle$. For a fixed ideal $I \subseteq \mathbb{K}[\mathbf{x}]$, we refer to any subset $F \subseteq \mathbb{K}[\mathbf{x}]$ such that $\langle F \rangle = I$ as a *generating set for I*.

4

The following theorem says that every ideal $I \subseteq \mathbb{K}[\mathbf{x}]$ has a finite generating set. A proof can be found in e.g. [32].

Theorem 1.1.3 (Hilbert Basis Theorem)**.** *Every ideal $I \subseteq \mathbb{K}[x_1, \ldots, x_n]$ can be expressed $I = \langle f_1, \ldots, f_r \rangle$ for some $f_1, \ldots, f_r \in I$.*

Given a subset $S \subseteq \mathbb{K}^n$, the *vanishing ideal of S*, denoted $I(S)$, is the set of polynomials vanishing on S. That is,

$$I(S) := \{f \in \mathbb{K}[x_1, \ldots, x_n] : f(a) = 0 \text{ for all } a \in S\}.$$

It is not difficult to see that $I(S)$ is always an ideal. An ideal I in a ring R is said to be *prime* if whenever $fg \in I$, then either $f \in I$ or $g \in I$.

Example 1.1.4. The ideal $\langle x^2 + 1 \rangle$ is prime as an ideal in $\mathbb{R}[x]$. To see this, note that $x^2 + 1$ is irreducible as a polynomial in $\mathbb{R}[x]$ and so if $fg \in \langle x^2 + 1 \rangle$, then either f or g must be a multiple of $x^2 + 1$. In other words, either $f \in \langle x^2 + 1 \rangle$ or $g \in \langle x^2 + 1 \rangle$. However, $\langle x^2 + 1 \rangle$ is not prime as an ideal in $\mathbb{C}[x]$ since $x^2 + 1 = (x + i)(x - i) \in \langle x^2 + 1 \rangle$ but neither $x + i$ nor $x - i$ is a member of $\langle x^2 + 1 \rangle$.

A variety $V \subseteq \mathbb{K}^n$ is said to be *irreducible* if V cannot be expressed as a nontrivial union of subvarieties. That is, whenever $V = V_1 \cup V_2$, either $V_1 = V$ or $V_2 = V$. If V is reducible (i.e. not irreducible) then we can express $V = V_1 \cup \cdots \cup V_k$ where where each V_i is irreducible. Up to reindexing, this representation is unique and the irreducible varieties $V_1, \ldots, V_k$ are called the *irreducible components of V*.

It is not hard to see that V is irreducible if and only if $I(V)$ is prime. Moreover, if $V = V_1 \cup V_2$, then $I(V) = I(V_1) \cap I(V_2)$ and so if the irreducible components of V are $V_1, \ldots, V_k$, then $I(V) = \bigcap_k V_k$.

Example 1.1.5. We consider the geometric side of Example 1.1.4. Note that $V(x^2 + 1) \subseteq \mathbb{C}^1$ is a reducible variety since $V(x^2 + 1) = \{-i, i\} = V(x+i) \cup V(x-i)$. Moreover, $\{-i\}$ and $\{i\}$ are the irreducible components of $V(x^2 + 1)$. Also note that $\langle x^2 + 1 \rangle = \langle x+i \rangle \cap \langle x-i \rangle$ and that $I(\{-i\}) = \langle x + i \rangle$ and $I(\{i\}) = \langle x - i \rangle$, both of which are prime. Over the reals however, $V(x^2 + 1) \subseteq \mathbb{R}^1$ is irreducible since $V(x^2 + 1) = \emptyset$ which clearly cannot be expressed as a nontrivial union of subvarieties.

Many sets whose geometry we wish to study are not varieties. However, one can still gain insight into such a set by considering its *Zariski closure*, the smallest variety that contains it. The precise definition is below.

Definition 1.1.6. Let $S \subseteq \mathbb{K}^n$ be a set. Then the *Zariski closure of S* is the set $V(I(S))$. Given a variety $V \subseteq \mathbb{C}^n$ and a subset $S \subseteq V$, if $V(I(S)) = V$, then we say that S is *Zariski dense in V*.

1.1.1 Tangent space and dimension

In this subsection, any unspecified field $\mathbb{K}$ will be assumed to be either $\mathbb{R}$ or $\mathbb{C}$. Given a point $a \in \mathbb{K}^n$ and $f \in \mathbb{K}[x_1, \ldots, x_n]$ we define $D_a(f) \in \mathbb{K}[x_1, \ldots, x_n]$ to be the linear form given as

$$D_a(f)(x_1, \ldots, x_n) := \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}(a) x_i.$$

The affine function $L_a(x_1, \ldots, x_n) := f(a) + D_a(f)(x_1 - a_1, \ldots, x_n - a_n)$ is often known as the *first Taylor approximation to f at a* or the *linearization of f at a*.

Definition 1.1.7. Given an irreducible variety $V \subseteq \mathbb{K}^n$ and a point $a \in V$, the *tangent space to V at a* is the variety $T_a V$ defined as follows

$$T_a V := V(\{D_a(f) : f \in I(V)\}).$$

Recall that $V = I(f_1, \ldots, f_r)$ for some $f_1, \ldots, f_r \in \mathbb{K}[x_1, \ldots, x_n]$ and so each $f \in I(V)$ can be expressed as $f = \sum_i g_i f_i$. It follows from this, linearity of the partial derivative operator, the product rule, and $f_i(a) = 0$ for $a \in V$ that

$$T_a V = \bigcap_{i=1}^{r} V(D_a(f_i)).$$

Since each $V(D_a(f_i))$ is a hyperplane, $T_a V$ is a vector subspace of $\mathbb{K}^n$.

Definition 1.1.8. Let $\mathbb{K}$ be either $\mathbb{C}$ or $\mathbb{R}$ and let $V \subseteq \mathbb{K}^n$ be an irreducible variety. Then the *dimension of V*, denoted $\dim(V)$, to be the smallest dimension of a tangent space of V as a vector subspace of $\mathbb{K}^n$. That is,

$$\dim(V) := \min_{a \in V} \dim(T_a V)$$

where $\dim(T_a V)$ is the dimension of $T_a V$ as a vector subspace of $\mathbb{K}^n$. When V is reducible, we define the dimension of V to be the maximum among the dimensions of the irreducible components of V. That is, if $V = V_1 \cup \cdots \cup V_k$ with each V_i irreducible, then

$$\dim V := \max_i \dim V_i.$$

The dimension of a variety is more commonly defined as the Krull dimension of its ring of regular functions. This turns out to be equivalent to Definition 1.1.8 - see [28, Chapters 10 and 16].

Example 1.1.9. We now discuss the intuition behind this definition of the dimension of an irreducible variety using $V := V(x^3 - y^2) \subseteq \mathbb{C}^2$ as an example (note that it is irreducible). Let $a_t := (t^2, t^3)$ be a point on V and denote $f := x^3 - y^2$. Then $D_{a_t}(f)$ is the linear form on $\mathbb{C}^2$ corresponding to the following row vector

$$\begin{pmatrix} 3t^2 & 2t^3 \end{pmatrix}.$$

Therefore $T_{a_t} V$ is defined by the vanishing of this linear form. When $t \neq 0$, the vanishing of $D_{a_t}(V)$ defines a 1-dimensional subspace of $\mathbb{C}^2$ and when $t = 0$, the vanishing defines the 2-dimensional subspace. Since $1 < 2$, V has dimension one. This of course matches our intuition about what the dimension of a curve, such as V, should be.

In Example 1.1.13, we will see why in general it is appropriate to define the dimension of an irreducible variety to be the minimum dimension of any tangent space.

1.1.2 Genericity

In this section, we make repeated use of the following basic fact from linear algebra.

Proposition 1.1.10. *Given a matrix $A \in \mathbb{K}^{m \times n}$, the rank of A is r if and only if all $(r+1) \times (r+1)$ subdeterminants of A vanish and some $r \times r$ subdeterminant is nonzero.*

Definition 1.1.11. Given an irreducible variety $V \subseteq \mathbb{K}^n$, a property P is said to hold *for a generic point in V* if P is true for all points in V that lie in the complement of a particular subvariety.

Example 1.1.12. Let $V \subseteq \mathbb{K}^{2 \times 2}$ be the variety defined by the vanishing of the determinant. That is,

$$V := \left\{ \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{pmatrix} : x_{11}x_{22} - x_{12}x_{21} = 0 \right\} \subseteq \mathbb{K}^{2 \times 2}.$$

Given a matrix $A \in V$, the rank of A is either 0 or 1. The rank is 0 if and only if A lies in the subvariety U defined by the four additional equations $x_{11} = x_{12} = x_{21} = x_{22} = 0$. Hence any matrix in $V \setminus U$ has rank 1 and so we say that a generic point in V is a matrix of rank 1.

More generally, let $V \subseteq \mathbb{K}^{m \times n}$ be the variety defined by the vanishing of all $(r+1) \times (r+1)$ subdeterminants of an $m \times n$ matrix of variables. Then the elements of V are the $m \times n$ matrices of rank r or less. A generic point of V is a matrix of rank r since in order for a point $A \in V$ to have rank $r-1$ or less, A must lie in the subvariety of V defined by the vanishing of all $r \times r$ subdeterminants of an $m \times n$ matrix of variables.

Example 1.1.13. When invoking genericity, one is not usually explicit about the particular subvariety they are excluding. For example, consider the following statement:

$$\text{For generic } a \in V, \ \dim T_a V = \dim V.$$

Typically when such a statement is made, the intended audience can easily see that it is true within the complement of some variety. We walk through the above statement to illustrate this. Our definition of the dimension of a variety requires that for any $a \in V$, $\dim T_a V \geq \dim V$. Moreover, we established that if $V = V(f_1, \ldots, f_r)$ then $T_a V$ can be computed by intersecting the hypersurfaces defined by the vanishing of the linear forms $D_a(f_i)$. In other words, if $D_a(f_1, \ldots, f_r)$ denotes the $r \times n$ matrix whose ij entry is given by the coefficient of x_j in $D_a(f_i)$, then $T_a V$ is simply the kernel of $D_a(f_1, \ldots, f_n)$. Denoting $d := \dim V$, we can see that the maximum possible rank of rank $D_a(f_1, \ldots, f_r)$ is $n - d$ and that this upper bound is achieved for some $a \in V$. Moreover, for any $a \in V$ satisfying $\dim T_a V > \dim V$, we must have rank $D_a(f_1, \ldots, f_r) < n - d$. In other words, $\dim T_a V > \dim V$ if and only if a lies in the subvariety of V defined by the vanishing of the $(n-d) \times (n-d)$ subdeterminants of $D_x(f_1, \ldots, f_r)$.

1.2 Matroids

Every chapter in this thesis involves some sort of system that depends on several variables that are intertwined in some way. We will often seek a concise way to describe which subsets of variables constrain each other and which subsets are free. The mathematical object most well-suited for our needs here is the *matroid*.

Before defining matroids, we give an example demonstrating the type of structure a matroid captures. Consider the matrices A and B shown below, each with columns labeled a, b, c, d

$$
A := \begin{pmatrix} a & b & c & d \\ 1 & 2 & 1 & -1 \\ 0 & 0 & -1 & 1 \end{pmatrix} \qquad B := \begin{pmatrix} a & b & c & d \\ 2 & 4 & 2 & 4 \\ -1 & -2 & 3 & 6 \end{pmatrix}.
$$

Not only are these matrices different, their rows span different linear subspaces of $\mathbb{R}^4$. However, they have the same combinatorial structure in the following sense: no columns are zero and the pairs of column labels corresponding to spanning sets of $\mathbb{R}^2$ are ac, ad, bc, and bd. Equivalently, in the subspace of $\mathbb{R}^4$ spanned by the rows of each, the pairs of coordinates that satisfy a nontrivial linear constraint are those corresponding to the label pairs ab and cd. This underlying combinatorial structure that these two matrices share is what is known as a matroid.

Definition 1.2.1. Let E be a finite set and let $\mathcal{I}$ be a set of subsets of E. The pair $\mathcal{M} := (E, \mathcal{I})$ is called a *matroid* if the following three conditions are satisfied:

1. $\emptyset \in \mathcal{I}$
2. if $J \subset I$ and $I \in \mathcal{I}$, then $J \in \mathcal{I}$
3. if $I, J \in \mathcal{I}$ with $|I| = |J| + 1$, then there exists some $e \in I \setminus J$ such that $J \cup \{e\} \in \mathcal{I}$.

In this case E is called the *ground set of* $\mathcal{M}$ and the elements of $\mathcal{I}$ are called the *independent sets of* $\mathcal{M}$. Perhaps the simplest matroids are *uniform matroids*, described below.

Example 1.2.2. Let $U_{d,n} := (E, I)$ where E is some finite set and I is the set of all subsets of E of size d or less. Then $U_{d,n}$ is a matroid. Matroids of the form $U_{d,n}$ are called *uniform*.

Every matrix, graph, and irreducible variety has an underlying matroid. This is something we now explore in detail.

Proposition 1.2.3. *Let $A \in \mathbb{K}^{m \times n}$ be a matrix with entries in a field $\mathbb{K}$. Let $E = \{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$ be the set of columns of A and let $\mathcal{I}$ be the subsets of E that are linearly independent. Then $(E, \mathcal{I})$ is a matroid.*

Proof. It is clear that $(E, \mathcal{I})$ satisfies the first two requirements to be a matroid so we now show that it satisfies the third. Let $I, J \subseteq E$ be linearly independent sets with $|I| = |J| + 1$. Let L, M be the linear subspaces of $\mathbb{K}^m$ with bases given by I and J. For a given $e \in I \setminus J$, the set $J \cup \{e\}$ is dependent if and only if e lies in M. But if every $e \in I$ lies in M, then $L \subseteq M$ which cannot happen because $\dim(L) = \dim(M) + 1$. $\qquad\square$

Definition 1.2.4. For a given matrix A, the matroid that arises as in Proposition 1.2.3 is called the *(linear) matroid underlying A*, denoted $\mathcal{M}_A$. A matroid $\mathcal{M}$ satisfying $\mathcal{M} = \mathcal{M}_A$ for some $\mathbb{K}$-matrix A is called $\mathbb{K}$-*representable*, or *representable over $\mathbb{K}$*.

There do exist matroids that are not representable over any field $\mathbb{K}$ (see e.g. [49, Proposition 2.2.26]) but we will not be concerned with such matroids in this thesis.

Example 1.2.5. Every uniform matroid $U_{d,n}$ is $\mathbb{Q}$-representable. In particular, if $A \in \mathbb{Q}^{d \times n}$ is a $d \times n$ matrix such that no $d \times d$ sub-determinant vanishes, then $\mathcal{M}_A = U_{d,n}$. However, a given $U_{d,n}$ may not be $\mathbb{K}$-representable for certain finite fields $\mathbb{K}$. For example, $U_{2,4}$ is not $\mathbb{F}_2$-representable. If it were, then suppose $A \in \mathbb{F}_2^{d \times 4}$ had $\mathcal{M}_A = U_{2,4}$. Then we could assume $d = 2$ since the rank of A must be 2. But there are exactly four vectors in $\mathbb{F}_2^2$, so either a column of A gets repeated, or $(0,0)$ is a column. Either way, the underlying matroid will not be uniform.

Proposition 1.2.6. *Let $G = (V, E)$ be a graph with edge set E. Let $\mathcal{I}$ be the subsets of E that do not have any cycle. Then $(E, \mathcal{I})$ is a matroid.*

Proof. Fix a total ordering $\prec$ on the vertex set V. Let $A_G^{\prec}$ be the matrix whose rows are indexed by V, whose columns are indexed by E, and whose (v, e) entry is 0 if e is a loop, 1 if v is the $\prec$-minimal vertex of e, -1 is the $\prec$-maximal vertex of v, and 0 otherwise. We now show that $\mathcal{I}$ is the set of independent sets in $\mathcal{M}_{A_G^{\prec}}$.

Given $S \subseteq E$, let A_S denote the submatrix of $A_G^{\prec}$ consisting of the columns indexed by S. Now let $S \in \mathcal{I}$; i.e. let $S \subseteq E$ be a subset of edges with no cycle. There is a row

of A_S that has exactly one non-zero entry corresponding to a vertex s of degree one in the graph (V, S). This implies that any vector x giving linear dependence $A_S x = 0$ must satisfy $x_s = 0$. Therefore the columns of A_S are linearly dependent if and only if the columns of $A_{S \setminus \{s\}}$ are too. However, the columns of $A_{S \setminus \{s\}}$ are linearly independent by induction.

It now remains to show that if $S \subseteq E$ has a cycle, then A_S has linearly dependent columns. So assume S contains a cycle $C = v_1, v_2, \ldots, v_k, v_1$. Define $x \in \mathbb{K}^S$ so that $x_e = 0$ if e does not appear in the cycle, $x_e = 1$ if $e = \{v_i, v_j\}$ with $i < j$ and $v_i \prec v_j$ and $x_e = -1$ if $e = \{v_i, v_j\}$ with $i < j$ but $v_i \succ v_j$. Then $A_S x = 0$. $\qquad\square$

Definition 1.2.7. For a given graph G, the matroid that arises as in Proposition 1.2.6 is called the *(polygon) matroid underlying G*, denoted $\mathcal{M}_G$. A matroid $\mathcal{M}$ satisfying $\mathcal{M} = \mathcal{M}_G$ for some graph G is called *graphic*.

It is clear from the proof of Proposition 1.2.6 that each graphic matroid is representable over any field $\mathbb{K}$. However, there do exist matroids that are representable, but not graphic. For example, $U_{2,4}$ is not graphic and this follows from the fact that it is not $\mathbb{F}_2$-representable.

Example 1.2.8. Let G be the graph shown below and let $\prec$ order the vertices according to their integer labels. then the independent sets of $\mathcal{M}_G$ are all proper subsets of $\{b, c, d\}$. One can check that these are also the independent sets of $\mathcal{M}_{A_G^{\prec}}$.
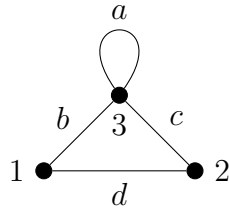
$$A_G^{\prec} = \begin{array}{c} 1 \\ 2 \\ 3 \end{array}\begin{pmatrix} a & b & c & d \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & -1 \\ 0 & -1 & -1 & 0 \end{pmatrix}.$$

The last class of matroids we will encounter come from irreducible varieties. In the following proposition we use the notation $\mathbb{K}^E$ to denote the $\mathbb{K}$-vector space whose coordinates are in some natural bijection with E, which will be a finite set.

Proposition 1.2.9. *Let $\mathbb{K}$ be $\mathbb{R}$ or $\mathbb{C}$ and let E be a finite set. Let $V \subseteq \mathbb{K}^E$ be an irreducible variety and let $\mathcal{I}$ be the subsets $S \subseteq E$ such that the projection of V onto $\mathbb{K}^S$ is full-dimensional (i.e. $|S|$-dimensional). Then $(E, \mathcal{I})$ is a matroid.*

Proof. For a given $x \in V$, let $T_x V$ denote the tangent space of V at x. For generic $x \in V$, the projection of V onto $\mathbb{K}^S$ is full-dimensional if and only if the projection of $T_x V$ onto $\mathbb{K}^S$ is full-dimensional. Therefore it suffices to prove the proposition in the case that V is a linear space. Now let V be the row span of some matrix $A \in \mathbb{K}^{n \times |E|}$ whose columns are indexed by E. Then the projection of V onto $\mathbb{K}^S$ is full-dimensional if and only if the columns of A corresponding to the elements of S are linearly independent. It then follows from Proposition 1.2.3 that $(E, \mathcal{I})$ is a matroid. $\qquad\square$

Definition 1.2.10. For a given irreducible variety V in $\mathbb{R}^n$ or $\mathbb{C}^n$, the matroid that arises as in Proposition 1.2.9 is called the *(algebraic) matroid underlying V*, denoted $\mathcal{M}_V$.

In the proof of Proposition 1.2.9, we saw that the algebraic matroid underlying a real or complex irreducible variety is the same as the algebraic matroid underlying a generic tangent space. Moreover, the algebraic matroid underlying a linear space is simply the column matroid of a certain matrix. Hence $\mathcal{M} = \mathcal{M}_V$ for some irreducible real or complex variety V if and only if $\mathcal{M}$ is $\mathbb{R}$- or $\mathbb{C}$-representable, respectively.

Example 1.2.11. Let $E = \{a, b, c, d\}$. If V is the variety in the space with coordinates labeled by E defined by the two polynomials below, then $\mathcal{M}_V$ is the same as the matroid $\mathcal{M}_G$ from Example 1.2.8

$$a - 1 = 0 \qquad b^2 + 3c^3 - d + 1 = 0.$$

The algebraic matroid underlying a variety is usually formulated in terms of algebraic independence in the corresponding function field (see e.g. Section 6.7 in [49]). For varieties defined over $\mathbb{R}$ or $\mathbb{C}$, either definition leads to the same class of matroids, which as we already remarked, are simply the $\mathbb{R}$- or $\mathbb{C}$-representable matroids. However, in order to define the algebraic matroid underlying a variety over a finite field, one must use this definition in terms of algebraic independence in the function field. For finite fields, this actually leads to a larger class of matroids, but we chose to present the algebraic matroid in terms of projections because it is more geometrically intuitive, and we will only be considering $\mathbb{R}$ and $\mathbb{C}$ as base fields anyway.

In many applications of matroid theory, the independent sets of a given matroid may not be as interesting as the (minimal) dependent sets. For this reason, the following definition is made.

Definition 1.2.12. Let $\mathcal{M} = (E, \mathcal{I})$ be a matroid. A *circuit of* $\mathcal{M}$ is a subset $C \subseteq E$ such that C is not independent in $\mathcal{M}$, but every proper subset is.

Example 1.2.13. Let $G = (V, E)$ be a graph. The circuits in the matroid $\mathcal{M}_G$ underlying G are precisely the subsets of E that support a cycle.

By definition, the independent sets of a matroid determine the circuits. The converse is also true. Namely, given the ground set E of a matroid $\mathcal{M}$ and its set $\mathcal{C}$ of circuits, the independent sets are simply the subsets of E that do not contain a circuit. In this way, the circuits and the independent sets of a matroid carry the same structure. In fact, due to the following proposition, one can even axiomatize matroid theory in terms of circuits instead of independent sets.

Proposition 1.2.14. *Let E be a finite set and let $\mathcal{C} \subseteq 2^E$ be a set of subsets of E. Then $\mathcal{C}$ is the set of circuits of a matroid on ground set E if and only if*

1. *$\emptyset \notin \mathcal{C}$*
2. *if $C_1 \subseteq C_2$ then $C_1 = C_2$*
3. *if $C_1 \neq C_2$, then for any $e \in C_1 \cap C_2$, there exists $C_3 \in \mathcal{C}$ such that $C_3 \subseteq (C_1 \cup C_2) \setminus \{e\}$.*

For a proof of Proposition 1.2.14, see [49, Chapter 1.1]. We end this section by noting several other structures one often associates with a matroid. A *basis of* $\mathcal{M}$ is an independent set of maximum cardinality. The *rank function of* $\mathcal{M}$ is the function $\rho : 2^E \to \mathbb{N}$ that sends a given $S \subseteq E$ to the cardinality of the largest independent set contained in S. For a given $S \subseteq E$, we often refer to $\rho(S)$ as *the rank of S*. The *rank of* $\mathcal{M}$ is simply $\rho(E)$, the rank of the ground set. A *spanning set of* $\mathcal{M}$ is a superset of a basis.

Example 1.2.15. Let $\mathcal{M}$ be the matroid shown in Examples 1.2.8 and 1.2.11. The bases of $\mathcal{M}$ are $\{b, c\}$, $\{b, d\}$ and $\{c, d\}$, and the circuits are $\{a\}$ and $\{b, c, d\}$. The sets of rank zero are $\emptyset$ and $\{a\}$. The sets of rank one are $\{b\}, \{c\}, \{d\}, \{a, b\}, \{a, c\}, \{a, d\}$. The other nine subsets of $\{a, b, c, d\}$ have rank 2 and are spanning sets.

Just as in the case of circuits, one can infer the independent sets of a matroid $\mathcal{M}$ given only the bases, the rank function, or the spanning sets of $\mathcal{M}$. Moreover, one can axiomatize matroid theory in terms of bases, rank functions, or spanning sets. The same is true for several other structures associated with a matroid. This phenomenon is often known as *matroid cryptomorphism.*

1.3 Oriented matroids

Given a matroid coming from a matrix with entries in an ordered field (usually $\mathbb{R}$) or a graph with directed edges, there is some extra structure that one can add on to the underlying matroid. A matroid equipped with this extra structure is called an *oriented matroid*. We begin with some examples. Consider the two following real matrices

$$\begin{array}{ccc} a & b & c \end{array}$$
$$\begin{pmatrix} 1 & 1 & -1 \\ 1 & -1 & 0 \end{pmatrix} \quad \text{and} \quad \begin{array}{ccc} a & b & c \end{array} \begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & 0 \end{pmatrix}.$$

Both have the same underlying matroid; in particular, $\{a, b, c\}$ is the unique circuit of both. Any linear dependence corresponding to this circuit in the first matrix is in the subspace spanned by $(1, 1, 2)$, while such a dependence for the second is in that of $(1, 1, -2)$. In particular, any linear dependence of the first matrix has sign pattern $(+, +, +)$ or $(-, -, -)$ whereas any linear dependence of the second matrix has sign pattern $(+, +, -)$ or $(-, -, +)$. The oriented matroids associated to each of these matrices is what captures this difference.

One can also define an oriented matroid corresponding to a directed graph which captures information that the underlying matroid misses. For example, consider the following two directed graphs.
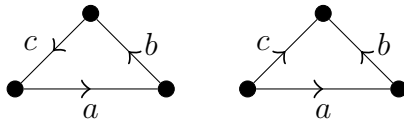
As in the previous example with matrices, both have the same underlying matroid with unique circuit $\{a, b, c\}$. If we traverse this circuit counterclockwise in the first graph, we will be going in accordance with the orientation at all three edges. This can be represented with the sign vector $(+, +, +)$, with the first $+$ indicating that we traversed a according to its orientation, the second indicating that for b and the third for c. If we traverse it clockwise, then we will be going against the orientation at all three edges which similarly will be represented $(-, -, -)$. Traversing the circuit in the second graph counterclockwise gives us the sign vector $(+, +, -)$ and clockwise gives us $(-, -, +)$. The oriented matroids associated to each of these directed graphs captures this information about orientation that the matroid by itself does not.

We now introduce the formal definitions needed to define an oriented matroid. Let E be a finite set. A *sign vector* is an element of the set $\{0, +, -\}^E$; that is, an assignment of either $0, +$ or $-$ to each element of E. Given $u \in \mathbb{R}^E$, $\text{sign}(u)$ denotes the sign vector obtained by replacing each entry of u with its sign. We will often assume an ordering on the elements of E, in which case can write a sign vector as $(s_1, \ldots, s_{|E|})$. Given a sign vector $\sigma = (\sigma_1, \ldots, \sigma_{|E|})$, we define $-\sigma := (-\sigma_1, \ldots, -\sigma_{|E|})$ according to the rules $-- = +$, $-+ = -$, and $-0 = 0$. The *composition* of two sign vectors σ and τ, denoted $\sigma \circ t$ is defined by

$$(\sigma \circ \tau)_e = \begin{cases} \sigma_e & \text{if } \sigma_e \neq 0 \\ \tau_e & \text{otherwise} \end{cases}.$$

Note that for $u, v \in \mathbb{R}^n$, $\text{sign}(u + \varepsilon v) = \text{sign}(u) \circ \text{sign}(v)$ when ε is sufficiently small. The *separation set* for sign vectors σ, τ is defined by

$$S(\sigma, \tau) := \{e \in E : \sigma_e = -\tau_e \neq 0\}.$$

For $f \in S(\sigma, \tau)$, we say that ρ *eliminates* f *between* σ *and* τ if

$$\rho_f = 0 \quad \text{and} \quad \rho_e = (\sigma \circ \tau)_e \quad \text{for all} \quad e \notin S(\sigma, \tau).$$

Note that for $u, v \in \mathbb{R}^E$ satisfying $u_e = 1$ and $v_e = -1$, $\text{sign}(u + v)$ eliminates i across $\text{sign}(u)$ and $\text{sign}(v)$.

Definition 1.3.1. Let $\mathcal{V}$ be a set of sign vectors for E. Then the pair $\mathcal{O} := (E, \mathcal{V})$ is said to be an *oriented matroid* if the following four conditions are satisfied:

1. $(0, \ldots, 0) \in \mathcal{V}$
2. if $\sigma \in \mathcal{V}$ then $-\sigma \in \mathcal{V}$
3. if $\sigma, \tau \in \mathcal{V}$ then $\sigma \circ \tau \in \mathcal{V}$
4. if $\sigma, \tau \in \mathcal{V}$ with $f \in S(\sigma, \tau)$, there exists $\rho \in \mathcal{V}$ that eliminates f between σ and τ.

In this case, the set $\mathcal{V}$ of sign vectors is called the *vectors* of $\mathcal{O}$ and E is called the *ground set*. These four conditions are often called the *(co)vector axioms of an oriented matroid*. We will often abuse notation and terminology, identifying the set of vectors of an oriented matroid with the oriented matroid itself. For example, we will write things like $\sigma \in \mathcal{O}$ to mean that σ is a vector of $\mathcal{O}$. When $\mathcal{O}$ and $\mathcal{P}$ are oriented matroids on

the same ground set, we may similarly write $\mathcal{O} \subseteq \mathcal{P}$ to mean that every vector of the oriented matroid $\mathcal{O}$ is also a vector of the oriented matroid $\mathcal{P}$.

We now detail how one can obtain an oriented matroid from any linear subspace of $\mathbb{R}^n$. For any real number $r \in \mathbb{R}$, $\text{sign}(r) \in \{+, -, 0\}$ is the *sign of r*. For a linear functional $c \in (\mathbb{R}^m)^*$, $\text{sign}(c) \in \{+, -, 0\}^m$ is defined by $\text{sign}(c)_i = \text{sign}(c_i)$. Given a sign vector $\sigma \in \{+, -, 0\}^m$, we define $|\sigma| := |\text{supp}(\sigma)|$.

Proposition 1.3.2. *Let $L \subseteq \mathbb{R}^m$ be a linear subspace. Let $\mathcal{V}$ be the set of sign vectors $\sigma \in \{+, -, 0\}^m$ such that $\sigma = \text{sign}(c)$ for some linear functional $c \in (\mathbb{R}^m)^*$ vanishing on L. Then $([m], \mathcal{V})$ is an oriented matroid.*

Proof. It is clear that the first two vector axioms are satisfied. Let c^σ and c^τ be linear functionals vanishing on L such that $\text{sign}(c^\sigma) = \sigma$ and $\text{sign}(c^\tau) = \tau$. For small enough $\varepsilon > 0$, $\text{sign}(c^\sigma + \varepsilon c^\tau) = \sigma \circ \tau$ thus showing that the third vector axiom is satisfied. Choose $e \in S(\sigma, \tau)$. By scaling appropriately, we can choose $c^\sigma_e = -c^\tau_e$. Then $\text{sign}(c^\sigma + c^\tau)$ eliminates e between σ and τ thus showing that the fourth vector axiom is satisfied. $\square$

For a linear space $L \subseteq \mathbb{R}^m$ we refer to the oriented matroid of Proposition 1.3.2 as the *oriented matroid associated to L*, denoted $\mathcal{O}_L$. Let $A \in \mathbb{R}^{m \times n}$ be a matrix. Then we may refer to $\mathcal{O}_{\text{row } A}$, the oriented matroid underlying the row-span of A, as the *oriented matroid underlying the columns of A*. This is reasonable because it consists of all sign patterns of dependencies among the columns of A. When $\mathcal{O} = \mathcal{O}_L$ for some L, we simplify notation and write $\mathcal{M}_L$ instead of $\mathcal{M}_{\mathcal{O}_L}$. Note that this is consistent with our notation for the algebraic matroid underlying an irreducible variety (linear spaces are irreducible varieties).

Example 1.3.3. Let $A = \begin{pmatrix} 1 & 1 & 0 \end{pmatrix}$ and let $L = \text{row } A \subseteq \mathbb{R}^3$. Denote the coordinates of the underlying space $\mathbb{R}^3$ as x, y and z. Linear functionals that vanish on L include $x - y$, $3z$, and $2x - 2y - z$. It is not hard to see that the vectors in $\mathcal{O}$, the oriented matroid underlying L (equivalently, the oriented matroid underlying the columns of A), are the following

$$
\begin{array}{ccc}
(+, -, 0) & (+, -, +) & (+, -, -) \\
(-, +, 0) & (-, +, +) & (-, +, -) \\
(0, 0, 0) & (0, 0, +) & (0, 0, -).
\end{array}
$$

Let $\prec^*$ be the partial order on $\{+,-,0\}$ given by $0 \prec^* +$ and $0 \prec^* -$ with $+$ and $-$ incomparable. Then $\prec$ is the partial order on $\{+,-,0\}^m$ that is the Cartesian product of $\prec^*$ m times. The *(signed) circuits* of an oriented matroid $\mathcal{O}$ are the nonzero vectors of $\mathcal{O}$ that are minimal with respect to $\prec$. A consequence of the following proposition is that in order to completely specify an oriented matroid, it suffices to specify its set of circuits.

Proposition 1.3.4. *Let $\mathcal{O}$ be an oriented matroid. Given $\sigma \in \mathcal{O}$, there exist signed circuits $\tau^1, \ldots, \tau^k$ such that $\sigma = \tau^1 \circ \cdots \circ \tau^k$.*

Proof. Let $\{\tau^1, \ldots, \tau^k\}$ be the set of all signed circuits satisfying $\tau \prec \sigma$. For each $i = 1, \ldots, k$ and each $e \in E$, either $\tau_e^i = \sigma_e$ or $\tau_e^i = 0$. Therefore $\tau^1 \circ \cdots \circ \tau^k \prec \sigma$. It now suffices to show that for each $e \in \mathrm{supp}(\sigma)$, there exists some circuit $\tau \prec \sigma$ satisfying $\tau_e = \sigma_e$. Let $\tau \in \mathcal{O}$ be the support-minimal vector such that $\tau \prec \sigma$ and $\tau_e = \sigma_e$. If τ is not a circuit, then there exists some circuit $\rho \prec \tau$ with $\rho_e = 0$. Then let η be a vector that eliminates some $f \in S(\tau, -\rho)$ between τ and $-\rho$. Then $\rho \circ \eta$ contradicts the minimality of τ; note that $(\rho \circ \eta)_e = \eta_e \neq 0$, $\rho \circ \eta \prec \tau$ and $\rho \circ \eta \neq \tau$ since $(\rho \circ \eta)_f = 0 \neq \tau_f$. $\qquad \square$

Example 1.3.5. The oriented matroid $\mathcal{O}_L$ from Example 1.3.3 has four circuits which are $(0, 0, +), (+, -, 0)$, and their negatives. Note that we can represent any non-zero non-circuit as a composition of two circuits. For example, $(+, -, +) = (+, -, 0) \circ (0, 0, +)$.

The *support* of a sign vector σ, denoted $\mathrm{supp}(\sigma)$, is defined to be the set $\{e \in E : \sigma_e \neq 0\}$. Lemma 1.3.6 below will be used to prove Proposition 1.3.7, which justifies the name "oriented matroid."

Lemma 1.3.6. *Given a sign vector $\tau \in \{+,-,0\}^E$, if there exists some $\sigma \in \mathcal{O}$ with $\sigma \neq \pm\tau$ and $\mathrm{supp}(\sigma) \subseteq \mathrm{supp}(\tau)$, then τ is not a signed circuit.*

Proof. If $\tau \notin \mathcal{O}$, then τ is trivially not a signed circuit, so assume $\tau \in \mathcal{O}$. We proceed by induction on $\min|S(\pm\sigma, \tau)|$, which we assume without loss of generality is $|S(\sigma, \tau)|$. For the base case, note that since $\mathrm{supp}(\sigma) \subseteq \mathrm{supp}(\tau)$, $S(\sigma, \tau) = \emptyset$ implies that $\sigma \prec \tau$, thus showing that τ is not a circuit. So assume there exists $e \in S(\sigma, \tau)$ and let $\rho \in \mathcal{O}$ eliminate e between σ and τ. Then ρ satisfies $\mathrm{supp}(\rho) \subseteq \mathrm{supp}(\tau)$ and $S(\rho, \tau) \subsetneq S(\sigma, \tau)$. So by induction on $\min|S(\pm\sigma, \tau)|$, τ is not a circuit. $\qquad \square$

Proposition 1.3.7. *If $\mathcal{O} = (E, \mathcal{V})$ is an oriented matroid with signed circuits $\mathcal{C}$, then $\{\mathrm{supp}(\sigma) : \sigma \in \mathcal{C}\}$ is the set of circuits of a matroid on ground set E.*

Proof. We proceed by showing that $\{\operatorname{supp}(\sigma) : \sigma \in \mathcal{C}\}$ satisfies the conditions given in Proposition 1.2.14. That $\{\operatorname{supp}(\sigma) : \sigma \in \mathcal{C}\}$ satisfies the first condition follows from the fact that circuits are defined to be nonzero. That they satisfy the second condition follows from Lemma 1.3.6.

We now show the third condition. Let $C_1 = \operatorname{supp}(\sigma)$ and $C_2 = \operatorname{supp}(t)$ and assume $C_1 \neq C_2$ with $e \in C_1 \cap C_2$. Since $\operatorname{supp}(\sigma) = \operatorname{supp}(-\sigma)$, we may assume that $\sigma_e = -\tau_e$. Then the fourth axiom of oriented matroids implies that there exists some $\rho \in \mathcal{V}$ that eliminates between e between σ and τ. Let $\eta \preceq \rho$ be a signed circuit. Then $\operatorname{supp}(\eta) \subseteq (C_1 \cup C_2) \setminus \{e\}$ and so we can take $C_3 = \operatorname{supp}(\eta)$. $\qquad\square$

We call the matroid of Proposition 1.3.7 the *matroid underlying* $\mathcal{O}$. The *rank* of any sign vector σ with respect to $\mathcal{O}$, is the rank of $\operatorname{supp}(\sigma)$ in $\mathcal{M}_\mathcal{O}$. Note that this is well defined for any sign vector, not just those in $\mathcal{O}$.

Example 1.3.8. We return once again to the oriented matroid from Example 1.3.3. The matroid underlying $\mathcal{O}_L$ has three independent sets, which are $\emptyset$, $\{1\}$ and $\{2\}$. The ranks of the three sign vectors in the bottom row of the table in Example 1.3.3 are all zero, and the ranks of the other six are all one.

We now describe how one can obtain an oriented matroid from a directed graph. Let $G = (V, A)$ be a directed graph with vertex set V and arc set A. Let C be a cycle in the undirected graph underlying G. We will now show how C gives rise to two sign vectors σ^C and τ^C. For each arc $a \in A$ whose corresponding undirected edge does not lie in C, set $\sigma_a = \tau_a = 0$. There are exactly two cyclic orientations of C which we denote C^+ and C^-. For each arc $a \in A$ whose corresponding edge e lies in C, set $\sigma_a = -\tau_a = +$ if the orientation C^+ imposes on e agrees with a, and set $\sigma_a = -\tau_a = -$ otherwise. Define $\mathcal{C}_G := \bigcup_C \{\sigma^C, \tau^C\}$.

Proposition 1.3.9. *For a directed graph $G = (V, A)$, the set $\mathcal{C}_G$ is the set of circuits of an oriented matroid $\mathcal{O}_G$.*

Proof. Let M_G be the matrix whose rows are indexed by V, whose columns are indexed by A, and whose entry corresponding to a pair (v, a) is 1 if a emanates from v, -1 if a points towards v, and 0 otherwise. Then note that the set of signed circuits of the oriented matroid underlying the row-span of M_G is $\mathcal{C}_G$. Therefore $\mathcal{C}_G$ is the set of circuits of an oriented matroid by Proposition 1.3.2. $\qquad\square$

For a directed graph G, we refer to the oriented matroid $\mathcal{O}_G$ from Proposition 1.3.9 as the *oriented matroid underlying* G.

Example 1.3.10. Consider the directed graph G shown below
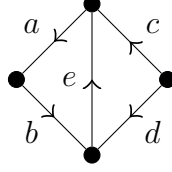
Ordering the arcs as (a, b, c, d, e), one can check that the circuits of $\mathcal{O}_G$ are the following sign vectors and their negatives:

$$(+, +, 0, 0, +) \qquad (+, +, +, -, 0) \qquad (0, 0, +, -, -).$$

1.3.1 Oriented matroid duality

Definition 1.3.11. Let E be a finite set. Given two sign vectors $\sigma, \tau \in \{+, -, 0\}^E$, we say that σ, τ are *orthogonal* and write $\sigma \cdot \tau = 0$ if one of the following holds

1. for each $e \in E$, $\sigma_e = 0$ or $\tau_e = 0$, or
2. there exist $e, f \in E$ such that $\sigma_e = \tau_e \neq 0$ and $\sigma_f = -\tau_f \neq 0$.

Given a set of sign vectors $S \subseteq \{+, -, 0\}^E$, we define the *orthogonal complement of S* as

$$S^{\perp} := \{\tau \in \{+, -, 0\}^E : \sigma \cdot \tau = 0 \text{ for all } \sigma \in S\}.$$

Example 1.3.12. Consider the following two examples

$$(+, 0, 0, -) \cdot (+, 0, -, +) = 0 \qquad \text{but} \qquad (+, 0, 0, +) \cdot (+, 0, -, +) \neq 0.$$

For the first case, note that if $x = (1, 0, 0, -1)$ and $y = (1, 0, -1, 1)$, then $x \cdot y = 0$, $\text{sign}(x) = (+, 0, 0, -)$, and $\text{sign}(y) = (+, 0, -, +)$. However, for the second case, note that if x, y satisfy $\text{sign}(x) = (+, 0, 0, +)$ and $\text{sign}(y) = (+, 0, -, +)$, then $x \cdot y > 0$. More generally, one can see that two sign vectors $\sigma, \tau \in \{+, -, 0\}^E$ satisfy $\sigma \cdot \tau = 0$ if and only if there exist $x, y \in \mathbb{R}^E$ satisfying $x \cdot y = 0$, $\text{sign}(x) = \sigma$, and $\text{sign}(y) = \tau$.

Orthogonality of sign vectors gives us a notion of duality for oriented matroids.

Definition 1.3.13. Let $\mathcal{O} = (E, \mathcal{V})$ be an oriented matroid. Then the *dual oriented matroid of* $\mathcal{O}$ is

$$\mathcal{O}^* := (E, \mathcal{V}^\perp).$$

Example 1.3.14. Consider the oriented matroid $\mathcal{O}$ given in Example 1.3.3. Then the dual oriented matroid $\mathcal{O}^*$ contains exactly two nonzero vectors, both of which are signed circuits. They are $(+, +, 0)$ and $(-, -, 0)$.

For a proof of the following proposition, see e.g. [15].

Proposition 1.3.15. *Given an oriented matroid* $\mathcal{O} = (E, \mathcal{V})$, *the dual oriented matroid* $\mathcal{O}^*$ *is an oriented matroid. Moreover,* $\mathcal{O}^{**} = \mathcal{O}$.

When an oriented matroid $\mathcal{O}$ underlies a linear subspace $L \subseteq \mathbb{R}^n$ or a directed graph G, then the dual oriented matroids $\mathcal{O}_L^*$ and $\mathcal{O}_G^*$ can be constructed explicitly. We begin with linear spaces.

Definition 1.3.16. Let $A \in \mathbb{R}^{r \times n}$ be a matrix of rank r. A *Gale dual* of A is a matrix $B \in \mathbb{R}^{(n-r) \times n}$ of rank $n - r$ such that $AB^T = 0$.

Note that if B is a Gale dual of A, then A is a Gale dual of B. We now show that if A and B are Gale duals, then $\mathcal{O}_A^* = \mathcal{O}_B$.

Proposition 1.3.17. *Let* $A \in \mathbb{R}^{r \times n}$ *be a matrix of rank* r *and let* B *be a Gale dual of* A. *Then* $\mathcal{O}_A^* = \mathcal{O}_B$.

Proof. Let $\sigma \in \mathcal{O}_B$ and let $x \in \mathbb{R}^n$ such that $Bx = 0$ and $\mathrm{sign}(x) = \sigma$. Now let $y \in \mathbb{R}^n$ such that $Ay = 0$. Since the columns of B^T are a basis of $\ker A$, $y = B^T z$ for some $z \in \mathbb{R}^{n-r}$. Then $x^T y = x^T B^T z = 0$ which implies $\mathrm{sign}(x) \cdot \sigma = 0$. This implies that $\sigma \in \mathcal{O}_A^*$ and therefore $\mathcal{O}_B \subseteq \mathcal{O}_A^*$.

By making an analogous argument with the roles of A and B reversed we see that $\mathcal{O}_A \subseteq \mathcal{O}_B^*$. Note that for any two collections of sign vectors $\mathcal{S}, \mathcal{T}$, if $\mathcal{S} \subseteq \mathcal{T}$, then $\mathcal{T}^\perp \subseteq \mathcal{S}^\perp$. Therefore we have $\mathcal{O}_A^* \subseteq \mathcal{O}_B^{**} = \mathcal{O}_B$. $\qquad\square$

Example 1.3.18. Let A be the matrix from Example 1.3.3. We display A alongside a Gale dual B below

$$A = \begin{pmatrix} 1 & 1 & 0 \end{pmatrix} \qquad B = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

We saw in Example 1.3.14 that the dual oriented matroid $\mathcal{O}_A^*$ has exactly two nonzero vectors which are $(+, +, 0)$ and $(-, -, 0)$. This is exactly the oriented matroid underlying B, whose kernel is spanned by A^T.

We now discuss the combinatorial interpretation of the circuits in oriented matroids that are dual to those underlying directed graphs. When dealing with oriented matroids underlying graphs, it tends to be more natural to work with circuits than with vectors so we begin with the following proposition.

Proposition 1.3.19. *Let $\mathcal{O} = (E, \mathcal{V})$ be an oriented matroid and let $\mathcal{C} \subseteq \mathcal{V}$ denote its set of circuits. Then $\mathcal{V}^\perp = \mathcal{C}^\perp$.*

Proof. It is clear that $\mathcal{V}^\perp \subseteq \mathcal{C}^\perp$ so we prove the other direction. Let $\sigma \in \mathcal{C}^\perp$ and let $\tau \in \mathcal{V}$. Assume $\tau \notin \mathcal{C}$ so that $\tau = \tau^1 \circ \cdots \circ \tau^k$. As $\tau_1 \in \mathcal{C}$, $\sigma \cdot \tau^1 = 0$. There are two cases to consider. The first case is that $\sigma_e = 0$ whenever $\tau_e^1 \neq 0$. From this and the inductive hypothesis, it follows that $(\tau^2 \circ \cdots \circ \tau^k) \cdot \sigma = 0$, which in turn implies $\tau \cdot \sigma = 0$. The second case is that there exist $e, f \in E$ such that $\sigma_e = \tau_e^1 \neq 0$ and $\sigma_f = -\tau_f^1 \neq 0$. This immediately implies $\tau \cdot \sigma = 0$. $\square$

Definition 1.3.20. Let $G = (V, E)$ be a graph on vertex set V and edge set E. A *cut of G* is a subset $C \subseteq E$ of edges such that the subgraph $(V, E \setminus C)$ has strictly more connected components than G. A *bond of G* is an inclusion-minimal cut.

Note that any bond C of an undirected graph $G = (V, E)$ will separate exactly one connected component of G into two connected components in $(V, E \setminus C)$. We denote the subgraphs induced on these connected components by $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$. So each edge $e \in C$ is incident to exactly one vertex in each V_i. When G is the undirected graph underlying some directed (V, A), we create a sign vector $\sigma \in \{+, -, 0\}^A$ such that $\sigma_e = 0$ for all $e \notin C$, $\sigma_e = +$ if $e \in C$ and e is oriented from V_1 to V_2, and $\sigma_e = -$ if $e \in C$ and e is oriented from V_2 to V_1. Note that changing the roles of V_1 and V_2 negates the corresponding sign vector.

Definition 1.3.21. For a directed graph $G = (V, A)$, any sign vector in $\{+, -, 0\}^A$ obtained as described above is called a *signed bond of G*.

Proposition 1.3.22. *Let $G = (V, A)$ be a directed graph. The set of signed bonds of G is the set of circuits in $\mathcal{O}_G^*$.*

Proof. Let σ be a signed bond of G. Since signed bonds are support-minimal by definition, it only remains to show that $\sigma \in \mathcal{O}_G^*$. By Proposition 1.3.19, it suffices to show that $\sigma \cdot \tau = 0$ for every signed circuit τ of G. So let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be the new connected components of $(V, A \setminus \mathrm{supp}(\sigma))$. Assume σ was specified so that $\sigma_e = +$ whenever e points from V_1 into V_2. Let $\tau \in \mathcal{O}_G$. Then if $\sigma_e = \tau_e = +$ (without loss of generality), this means that as we traverse the cycle $\mathrm{supp}(\tau)$ according to τ, we are moving from V_1 into V_2 at the edge e, which means we must move back from V_2 into V_1 at some other edge f. If f points from V_1 to V_2 we would have $\tau_f = -$ and $\sigma_f = +$ and if f points from V_2 into V_1, then we would have $\tau_f = +$ and $\sigma_f = -$. Either way, $\sigma \cdot \tau = 0$.

Now we must show that the circuits of $\mathcal{O}_G^*$ are signed bonds. So let $\sigma \in \mathcal{O}_G^*$. First we claim that $\mathrm{supp}(\sigma)$ is a cut in the undirected graph underlying G. To see this, note that for any cycle C of this graph, $|\mathrm{supp}(\sigma) \cap C| \neq 1$. Now if $e \in \mathrm{supp}(\sigma)$ has endpoints v, w, then every path between v and w in the undirected graph underlying G must have an edge in $\mathrm{supp}(\sigma)$. Otherwise, if $v, x_1, x_2, \ldots, x_k, w$ is a path from v to w with $x_i \notin \mathrm{supp}(\sigma)$, then the cycle obtained by adding the edge e back to this path is a cycle that intersects $\mathrm{supp}(\sigma)$ in exactly one element, namely e.

Now since the support of each $\sigma \in \mathcal{O}_G^*$ is a cut, $\mathrm{supp}(\sigma)$ must contain a bond. We have already seen that each bond of G is the support of some vector (moreover a circuit) in $\mathcal{O}_G^*$, so if we assume that σ is a circuit, then Lemma 1.3.6 implies that σ is a signed bond of G. $\qquad\square$

Example 1.3.23. Let G be the directed graph from Example 1.3.10. Again, ordering the arcs as (a, b, c, d, e), we see that the signed bonds of G are the vectors shown below, along with their negatives. Note that these are the circuits of the dual oriented matroid $\mathcal{O}_G^*$.

$$(+, -, 0, 0, 0) \qquad (+, 0, -, 0, -) \qquad (0, 0, +, +, 0)$$
$$(0, +, 0, +, -) \qquad (0, +, -, -, 0) \qquad (+, 0, 0, +, -)$$

1.4 Tree metrics

Tree metrics are a fundamental object in distance-based phylogenetics. They also play an important role in tropical geometry. The first half of this section collects the basic facts we will need about tree metrics, and the second half describes the polyhedral geometry

of the set of all tree metrics.

1.4.1 Basic definitions

Let X be a finite set. A *dissimilarity map* on X is a function $d : X \times X \to \mathbb{R}$ such that $d(x, y) = d(y, x)$ and $d(x, x) = 0$. We will usually associate each dissimilarity map d on X with the point $d \in \mathbb{R}^{\binom{[n]}{2}}$ given by $d_{xy} = d(x, y)$.

Definition 1.4.1. A *tree metric on X* is a dissimilarity map $d \in \mathbb{R}^{\binom{X}{2}}$ satisfying the following *four point condition* for all distinct $x, y, z, w \in X$

$$d_{xy} + d_{zw} \leq \max\{d_{xz} + d_{yw}, d_{xw} + d_{yz}\}.$$

This is equivalent to the condition that the maximum of $d_{xy} + d_{zw}, d_{xz} + d_{yw}$, and $d_{xw} + d_{yz}$ is attained twice.

Other authors often define tree metrics to be the dissimilarity map satisfying the four point condition, not just for sets of four *distinct* points, but additionally in cases where there is overlap among them. The reason for this is that the four point condition becomes the triangle inequality when only three of the four points are distinct, and it ensures non-negativity of each d_{xy} when only two are. Therefore, this more restrictive definition ensures that tree metrics are indeed metrics. In spite of this, we take this less restrictive definition because it makes the set of all tree metrics into the tropicalization of a particular variety, a fact we later exploit. Proposition 1.4.3 justifies the terminology "tree metric."

Definition 1.4.2. A *tree* is a connected graph $T = (V, E)$ with no cycles. A *leaf* of a tree is a vertex of degree one and an *internal edge* is an edge that is not incident to any leaf. Given a finite set X, an *X-tree* is a tree with leaf set X that has no vertices of degree two. A tree is *binary* if all internal vertices have degree three.

Proposition 1.4.3 ([53, Chapter 7]). *Let $T = (V, E)$ be an X-tree and let $\omega : E \to \mathbb{R}$ be a weighting of the edges of T such that $\omega(e) \geq 0$ for all internal edges of T. Let $d_T^\omega : X \times X \to \mathbb{R}$ be the dissimilarity map defined by*

$$d_T^\omega(x, y) = \sum_{e \in P} \omega(e)$$
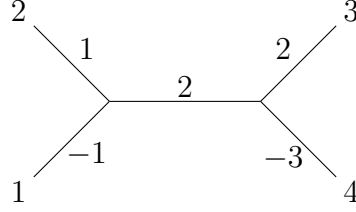
Figure 1.1: An edge-weighted tree with positive internal edge weights whose leaves are labeled by $\{1, 2, 3, 4\}$ showing that $d = (d_{12}, d_{13}, d_{14}, d_{23}, d_{24}, d_{34}) = (0, 3, -2, 5, 0, -1)$ is a tree metric.

where the sum is taken over the edges in the unique path from x to y in T. Then d is a tree metric. Moreover, for every tree metric d, there is a unique X-tree $T = (V, E)$ and edge weighting ω satisfying $\omega_e \neq 0$ for all $e \in E$ such that $d = d_T^\omega$.

When one takes the alternate definition of tree metric that requires the four point condition for all $x, y, z, w \in X$, and not just those that are distinct, then Proposition 1.4.3 is still true, but *all* edge weights are required to be nonnegative, not just the internal ones. The uniqueness of the tree specified in the second part of Proposition 1.4.3 gives us a combinatorial invariant of a tree metric.

Definition 1.4.4. Let d be a tree metric on X. The unique tree T from Proposition 1.4.3 is called the *(tree) topology of d*.

Example 1.4.5. Let $X = [4]$ and let d be the dissimilarity map given as

$$(d_{12}, d_{13}, d_{14}, d_{23}, d_{24}, d_{34}) = (0, 3, -2, 5, 0, 1).$$

Note that $d_{12} + d_{34} = 1$, and $d_{13} + d_{24} = d_{14} + d_{23} = 3$ and so the maximum of the three sums is attained twice. Hence d satisfies the four point condition and is therefore a tree metric. Moreover, if T is the tree displayed in Figure 1.1 with the indicated edge weighting ω, then $d = d_T^\omega$.

In Chapter 3, we will be concerned with a particular subset of tree metrics called *ultrametrics* or *rooted tree metrics*. We defer discussion of ultrametrics until Section 3.2.2.

1.4.2 Polyhedral geometry of the space of tree metrics

A *polyhedral cone* is the intersection of finitely many linear half-spaces. That is, an intersection of finitely many sets of the form $\{x \in \mathbb{R}^n : ax \geq 0\}$. Note that the intersection of any two polyhedral cones is also a polyhedral cone. Given a polyhedral cone $C \subseteq \mathbb{R}^n$, a *face* is a subset $F \subseteq C$ of the form

$$F = \{x \in C : cx = 0\}$$

for some $c \in (\mathbb{R}^n)^*$ such that $cx \geq 0$ for all $x \in C$. Note that every face of a polyhedral cone is itself a polyhedral cone. Setting $c = 0$, we see that the entire cone C is a face of itself. A *proper face* of C is any face other than C itself An inclusion-maximal proper face is called a *facet*.

Definition 1.4.6. A *polyhedral fan* is a set $\mathcal{C}$ of polyhedral cones in $\mathbb{R}^n$ satisfying the following two conditions

1. if $C \in \mathcal{C}$ and F is a face of C, then $F \in \mathcal{C}$
2. if $C, D \in \mathcal{C}$, then $C \cap D$ is a face of both C and D.

The elements of $\mathcal{C}$ are called the *faces of* $\mathcal{C}$ and the inclusion maximal faces are called *facets*.

Example 1.4.7. For $i = 1, 2, 3, 4$, let Q_i denote the i^{th} quadrant of $\mathbb{R}^2$ and let A_i denote the rotation of the ray $\{(x, 0) : x \geq 0\}$ by $90(i - 1)$ degrees counter-clockwise about the origin. Define

$$\mathcal{C} := \{Q_1, Q_2, Q_3, Q_4, A_1, A_2, A_3, A_4, \{0, 0\}\}$$

and

$$\mathcal{D} := \{Q_1, Q_2, Q_3 \cup Q_4, A_1, A_2, A_3, A_1 \cup A_3, \{0, 0\}\}.$$

Both $\mathcal{C}$ and $\mathcal{D}$ satisfy the first condition required of a polyhedral fan. Figure 1.2 displays the maximal elements of both $\mathcal{C}$ and $\mathcal{D}$. However, only $\mathcal{C}$ satisfies the second - note that $(Q_3 \cup Q_4) \cap Q_1 = A_1$ which is a face of Q_1 but not of $(Q_3 \cap Q_4)$.

The *support* of a polyhedral fan $\mathcal{C}$ is $S := \bigcup_{C \in \mathcal{C}} C$. In this situation, we say that S *supports* $\mathcal{C}$. The polyhedral fan structure supported on a given $S \subseteq \mathbb{R}^n$ is in general not unique. For example, the polyhedral fan $\mathcal{C}$ from Example 1.4.7 is supported on $\mathbb{R}^2$, but
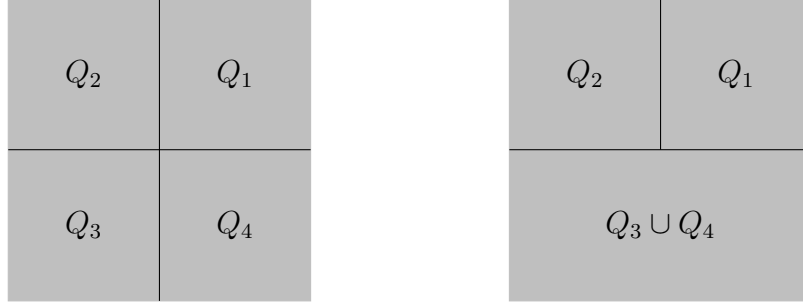
Figure 1.2: The figure on the left shows $\mathcal{C}$, which is a polyhedral complex. The figure on the right shows $\mathcal{D}$, which fails to be a polyhedral complex because Q_1 and $Q_3 \cup Q_4$ intersect along a non-face of $Q_3 \cup Q_4$.

so is the polyhedral fan $\{\mathbb{R}^2\}$. The set of all tree metrics on $[n]$ naturally supports a polyhedral fan structure which we now describe.

Definition 1.4.8. Let $\mathcal{T}_n$ denote the polyhedral fan structure supported on the set of all tree metrics on $[n]$ that has a face C_T for every $[n]$-tree T, consisting of all tree metrics with topology T' such that T' can be obtained from T via a (possibly empty) sequence of internal edge contractions.

The facets of $\mathcal{T}_n$ are the faces corresponding to the binary $[n]$-trees. We will often abuse notation and identify the polyhedral fan $\mathcal{T}_n$ with its support.

Example 1.4.9. We consider the polyhedral structure of $\mathcal{T}_4$ in detail. A tree metric on leaf set $\{1, 2, 3, 4\}$ can have one of four topologies, displayed in Figure 1.3. The binary tree topologies will be denoted T_1, T_2, and T_3 and the nonbinary tree topology will be denoted S. Thus $\mathcal{T}_4$ has three facets, each consisting of all the tree metrics with topology S or T_i for some $i = 1, 2, 3$. The intersection of any two facets is the set of all tree metrics with topology S.

1.5 Tropical geometry

Tropical geometry gives us a way to transform an irreducible variety $V \subseteq \mathbb{C}^n$ into a polyhedral complex $\mathrm{trop}(V) \subseteq \mathbb{R}^n$, called the *tropicalization of V*, that encodes many essential properties of V. One such property is the algebraic matroid.
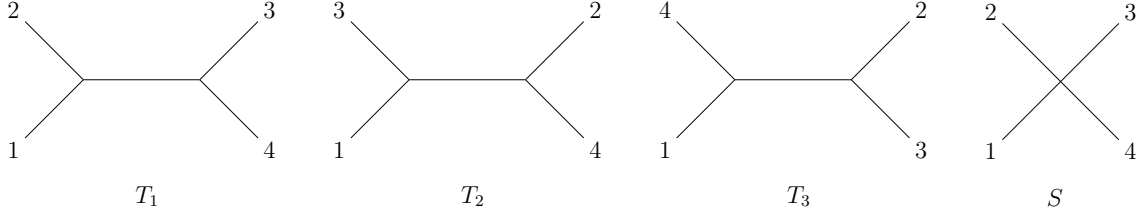
Figure 1.3: Possible tree topologies for a tree metric on four leaves.

We consider two equivalent ways one can define the tropicalization of a variety $V \subseteq \mathbb{C}$. For one of these ways, the first step is to consider V over a certain field extension of $\mathbb{C}$, called the *field of Puiseux series*. This field is naturally endowed with a *valuation*, which is a particular kind of function into $\mathbb{Q}$. By applying this function coordinate-wise to every point in V, one obtains a rational subset of $\mathbb{R}^n$, the Euclidean closure of which is $\text{trop}(V)$. This view of $\text{trop}(V)$ is well-suited for proving certain theorems, and in particular, that $\text{trop}(V)$ caries the algebraic matroid structure of V. One can equivalently define $\text{trop}(V)$ to be the set of points $\omega \in \mathbb{R}^n$ such that the initial ideal $in_\omega(I(V))$ contains no monomials. This definition is well-suited for computing tropical varieties.

1.5.1   Puiseux series

Definition 1.5.1. A *Puiseux series* is a formal power series with rational exponents with coefficients in $\mathbb{C}$ such that the set of all exponents of the formal variable is bounded below, and expressible over a common denominator. That is, a Puiseux series is a formal power series of the form $\sum_{\alpha \in \mathcal{J}_{d,k}} c_\alpha t^\alpha$ where $d \in \mathbb{N}$ and $k \in \mathbb{Z}$ and $\mathcal{J}_{d,k} = \{n/d : n \in \mathbb{Z}, n \geq k\}$. Addition and multiplication of formal power series induces the structure of an algebraically closed field on the set of Puiseux series [45, Theorem 2.1.5]. We denote this field by $\mathbb{C}\{\{t\}\}$. The *valuation map* is the function val $: \mathbb{C}\{\{t\}\} \to \mathbb{Q}$ that sends a Puiseux series to its minimum exponent.

Example 1.5.2. Consider the following three formal power series

$$f := \sum_{n=-8}^{\infty} n t^{n/3} \qquad g := \sum_{n=1}^{\infty} t^{-n} \qquad h := \sum_{n=1}^{\infty} t^{1/n}.$$

Among them, only f is a Puiseux series. Since the smallest exponent of f is $-8/3$, we

27

have that $\mathrm{val}(f) = -8/3$. The issue with g is that its set of exponents is not bounded below. The issue with h is that its exponents cannot all be expressed over a common denominator.

Definition 1.5.3. Let $V \subseteq \mathbb{C}^n$ be a variety and let $(\mathbb{C}\{\{t\}\} \otimes V) \subseteq (\mathbb{C}\{\{t\}\})^n$ denote the variety defined by the ideal $I(V)\mathbb{C}\{\{t\}\}[x_1, \ldots, x_n]$. Let $\mathrm{val} : (\mathbb{C}\{\{t\}\})^n \to \mathbb{R}^n$ denote the map sending each $(a_1, \ldots, a_n) \in (\mathbb{C}\{\{t\}\})^n$ to $(\mathrm{val}(a_1), \ldots, \mathrm{val}(a_n))$. Then the *tropicalization of V* is
$$\mathrm{trop}(V) := \overline{(-1) \cdot \mathrm{val}(\mathbb{C}\{\{t\}\} \otimes V)},$$
the Euclidean closure of the image of $\mathbb{C}\{\{t\}\} \otimes V$ under the negation of the coordinate-wise valuation map val. Sets of the form $\mathrm{trop}(V)$ are called *tropical varieties*.

One often defines the tropicalization of a variety in almost the same way, but without taking negatives at the end (see e.g. the introductory text [45]). Both conventions result in the same theory, modulo some trivial sign and order differences. Our motivation for choosing the convention that takes negatives is that it makes the connection with phylogenetics cleaner.

Example 1.5.4. Let $f(x, y) = x^2 + xy + y - 1$. Note that the variety $\mathbb{C}\{\{t\}\} \otimes V(f(x, y))$ contains the points $(t^{-1}, -t^{-1} + 1)$ and $(-1, t^2)$. Applying the coordinate-wise valuation map val to each point gives us $(-1, -1)$ and $(0, 2)$. Therefore, their negations $(1, 1)$ and $(0, -2)$ lie in the tropical variety $\mathrm{trop}(V(f(x, y))$. In the next section, we will develop tools that will enable us to describe the entire set $\mathrm{trop}(V(f(x, y))$.

1.5.2 Initial ideals

Definition 1.5.5. Let $f = \sum_{\alpha \in \mathcal{J}} c_\alpha \mathbf{x}^\alpha$ be a polynomial in $\mathbb{K}[x_1, \ldots, x_n]$ and let $\omega \in \mathbb{R}^n$. Define
$$\mathcal{J}_\omega := \{\alpha \in \mathcal{J} : \omega \cdot \alpha \text{ is maximized over } \mathcal{J}\}.$$
The *initial form of f with respect to ω* is the polynomial
$$in_\omega f := \sum_{\alpha \in \mathcal{J}_\omega} c_\alpha \mathbf{x}^\alpha.$$

Given $\omega \in \mathbb{R}^n$ and an ideal $I \subseteq \mathbb{K}[\mathbf{x}]$, the *initial ideal of I with respect to ω is*

$$in_\omega I := \langle in_\omega f : f \in I \rangle.$$

If $I = \langle f_1, \ldots, f_r \rangle$, it is clear that $\langle in_\omega f_1, \ldots, in_\omega f_r \rangle \subseteq in_\omega I$. However, the converse is false in general as the following example shows.

Example 1.5.6. Define $f_1(x_1, x_2, x_3) := x_1 - x_2$ and $f_2(x_1, x_2, x_3) := x_1 - x_3$ and let I denote the ideal they generate in $\mathbb{K}[x_1, x_2, x_3]$. Let $\omega = (3, 2, 1)$. Then $in_\omega f_1 = in_\omega f_2 = x_1$ and so $\langle in_\omega f_1, in_\omega f_2 \rangle = \langle x_1 \rangle$. However, this cannot be the whole initial ideal because $x_2 \in in_\omega I$ as $in_\omega(x_2 - x_3) = x_2$ and $x_2 - x_3 \in I$.

Definition 1.5.7. Let $I \subseteq \mathbb{K}[x_1, \ldots, x_n]$ be an ideal and let $\omega \in \mathbb{R}^n$. A set of polynomials $f_1, \ldots, f_r \in I$ is called a *Gröbner basis for ω* if

$$in_\omega I = \langle in_\omega f_1, \ldots, in_\omega f_r \rangle.$$

One might be concerned about the fact that Gröbner bases are not required by definition to generate the ideal they correspond to. However, this property is an easy consequence of the definition. Given an arbitrary generating set of an ideal, one can compute a Gröbner basis with respect to any $\omega \in \mathbb{R}^n$ using *Buchberger's algorithm*. For more details, see e.g. [32, Chapter 2]. Initial ideals and Gröbner bases are particularly important for tropical geometry due to the following theorem which gives an alternate characterization of the tropicalization of a variety.

Theorem 1.5.8 ([45], Theorem 3.2.3). *Let $V \subseteq \mathbb{C}^n$ be a variety. Then the tropicalization of V is the set of all $\omega \in \mathbb{R}^n$ such that the corresponding initial ideal of $in_\omega I(V)$ contains no monomials. That is,*

$$\mathrm{trop}(V) = \{\omega \in \mathbb{R}^n : in_\omega I(V) \text{ contains no monomials}\}.$$

Example 1.5.9. Consider the polynomial $f(x, y) = x^2 + xy + y - 1$ from Example 1.5.4. Given $\omega \in \mathbb{R}^2$, the initial ideal $in_\omega \langle f(x, y) \rangle$ contains no monomials if and only if $in_\omega f(x, y)$ is itself not a monomial. By way of case analysis over all pairs of monomials of $f(x, y)$, one can see that $in_\omega f(x, y)$ is not a monomial when $\omega_1 = 0$, or $\omega_1 \geq 0$ and $\omega_2 = \omega_1$, or $\omega_1 \leq 0$. Theorem 1.5.8 then implies that $\mathrm{trop}(V(f(x, y))$ is the union of the ω_2-axis, the ray $\omega_1 = \omega_2 \geq 0$, and the ray $\omega_1 \leq 0$.

29

1.5.3 Algebraic matroids and tropical geometry

As we noted earlier, the tropicalization of an irreducible variety V carries with it the structure of the algebraic matroid underlying V. This section makes that statement precise. All unspecified fields $\mathbb{K}$ in this subsection are assumed to be $\mathbb{R}$ or $\mathbb{C}$.

Definition 1.5.10. Let $V \subseteq \mathbb{K}^E$ be a set. The *independence complex of V*, denoted $M(V)$, is the set of subsets $S \subseteq E$ such that the coordinate projection of S onto $\mathbb{K}^S$ is full-dimensional in $\mathbb{K}^S$.

When $V \subseteq \mathbb{K}^E$ is an irreducible variety, $M(V)$ is the set of independent sets of $\mathcal{M}_V$, the algebraic matroid underlying V. When V is an arbitrary set, $M(V)$ need not be the independent sets in any matroid.

Proposition 1.5.11 ([64], Lemma 2)**.** *Let $V \subseteq \mathbb{C}^E$ be an irreducible variety. Then $S \subseteq E$ is independent in the algebraic matroid $\mathcal{M}_V$ underlying V if and only if $S \in M(\mathrm{trop}(V))$.*

The utility of Proposition 1.5.11 is that computing the dimension of a projection of a tropical variety may be easier than computing the dimension of the projection of the variety itself. This is because a tropical variety is a highly structured polyhedral complex. We will wait until Chapter 2 to make this statement precise. For now, we give an example.

Example 1.5.12. Let $E = \binom{[4]}{2}$ be the set consisting of all pairs of integers $1, 2, 3, 4$. Let $V \subseteq \mathbb{C}^E$ be the variety defined by the following polynomial

$$p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23}.$$

It is not hard to see directly that the algebraic matroid underlying V is simply the uniform matroid $U_{5,6}$, but we prove this using Proposition 1.5.11.

Since $I(V)$ is generated by a single polynomial, an initial ideal $in_\omega I(V)$ contains a monomial if and only if $in_\omega(p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23})$ is a monomial. Moreover, such an initial form is a monomial if and only if ω satisfies the four point condition. Therefore, the tropicalization $\mathrm{trop}(V)$ is $\mathcal{T}_4$, the set of all tree metrics on leaf set $\{1, 2, 3, 4\}$. There are three distinct binary tree topologies on four leaves and each corresponds to a maximal cone in $\mathcal{T}_4$. The linear hull of each such maximal cone is the linear space spanned by the rows of one of the following three matrices

$$
A = \begin{array}{c} \begin{array}{cccccc} 12 & 13 & 14 & 23 & 24 & 34 \end{array} \\ \left(\begin{array}{cccccc} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{array} \right) \end{array}
\qquad
B = \begin{array}{c} \begin{array}{cccccc} 12 & 13 & 14 & 23 & 24 & 34 \end{array} \\ \left(\begin{array}{cccccc} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{array} \right) \end{array}
$$

$$
C = \begin{array}{c} \begin{array}{cccccc} 12 & 13 & 14 & 23 & 24 & 34 \end{array} \\ \left(\begin{array}{cccccc} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \end{array} \right) \end{array}.
$$

Hence by Proposition 1.5.11, a subset $S \subseteq E$ is independent in the algebraic matroid $\mathcal{M}_V$ if and only if the projection of the row space of one of these matrices onto S is full-dimensional. For a particular choice of the matrix M, this is equivalent to the condition that the column submatrix M_S have rank $|S|$. Since the ranks of A, B, C are all five, it follows that E is not itself independent in $\mathcal{M}_V$. Now let $S \subseteq E$ such that $|S| = 5$. Then two of A_S, B_S, C_S have rank five: if $12 \notin S$ or $34 \notin S$, then B_S and C_S do if $13 \notin S$ or $24 \notin S$, then A_S and C_S do, and if $14 \notin S$ or $23 \notin S$, then A_S and B_S do. Hence S is independent in $\mathcal{M}_V$.

1.5.4 Tree metrics and tropical geometry

The Grassmannian $\mathrm{Gr}(2, n)$, defined below, is of fundamental importance in many different areas of mathematics since its points naturally parameterize the set of all 2-dimensional linear subspaces of $\mathbb{C}^n$. Motivated by applications to low-rank matrix completion, we will characterize the algebraic matroid underlying this variety in Chapter 2. Theorem 1.5.14 and Proposition 1.5.11 turn out to be key there.

Definition 1.5.13. For $1 \le i < j < k < l \le n$, define $f_{ijkl} \in \mathbb{C}[p_{ab} : 1 \le a < b \le n]$ to be

$$f_{ijkl} := p_{ij}p_{kl} - p_{ik}p_{jl} + p_{il}p_{jk}.$$

Then $\mathrm{Gr}(2, n)$, the *Grassmannian of planes in affine n-space*, is the variety defined by the vanishing of all polynomials of the form f_{ijkl}.

Theorem 1.5.14 ([56])**.** *The tropicalization of the Grassmannian of planes in affine n-space is the set of tree metrics on leaf set $[n]$. That is, $\mathrm{trop}(\mathrm{Gr}(2, n)) = \mathcal{T}_n$.*

We will not prove Theorem 1.5.14 here, but we can easily show $\mathrm{trop}(\mathrm{Gr}(2, n)) \subseteq \mathcal{T}_n$. So let $\omega \in \mathrm{trop}(\mathrm{Gr}(2, n))$. Then the condition that $in_\omega f_{ijkl}$ is not a monomial is equivalent to the condition that the maximum of $\omega_{ij} + \omega_{kl}$, $\omega_{ik} + \omega_{jl}$, and $\omega_{il} + \omega_{jl}$ is attained at least twice. Since this condition must be satisfied for all quadruples $ijkl$, ω satisfies the four point condition.

1.6 Toric ideals

Toric varieties, the varieties defined by toric ideals, have been extensively studied by algebraic geometers as their combinatorial nature makes them useful objects for testing conjectures and building intuition [22]. As we will see in Section 1.6, many models that are used in discrete multivariate analysis can be defined in terms of a toric ideal. Because of this, toric ideals offer a natural bridge between algebraic geometry and statistics. We now continue with some preliminaries.

Let L be a subgroup of $\mathbb{Z}^n$ - that is, L is a *lattice*. Let $\mathbb{K}$ be a field. The *lattice ideal of L* is the ideal I_L in the polynomial ring $\mathbb{K}[x_1, \ldots, x_n]$ defined as

$$I_L := \langle \mathbf{x}^u - \mathbf{x}^v : u - v \in L \rangle.$$

Given an integer matrix $A \in \mathbb{Z}^{m \times n}$, we let $\ker_\mathbb{Z} A$ denote the kernel of A as a map from $\mathbb{Z}^n$ to $\mathbb{Z}^m$.

Proposition 1.6.1. *The lattice ideal I_L is prime if and only if $L = \ker_\mathbb{Z} A$ for some integer matrix $A \in \mathbb{Z}^{m \times n}$.*

Proof. Theorem 7.4 in [46] implies that I_L is prime if and only if $\mathbb{Z}^n / L$ is free abelian. If $L = \ker_\mathbb{Z} A$ for some $A \in \mathbb{Z}^m \times n$ then $\mathbb{Z}^n / L$ is isomorphic to the lattice generated by the

columns of A, which is free abelian. Conversely, if $\mathbb{Z}^n/L$ is free abelian, then we have an isomorphism $\mathbb{Z}^n/L \to \mathbb{Z}^m$ for some m. This induces a linear map $\mathbb{Z}^n \to \mathbb{Z}^m$ with kernel L. Then any matrix $A \in \mathbb{Z}^{m \times n}$ representing this linear map in the standard bases of $\mathbb{Z}^m$ and $\mathbb{Z}^n$ satisfies the property that $\ker_{\mathbb{Z}} A = L$. $\qquad\square$

Prime lattice ideals are called *toric ideals*. Each integer matrix $A \in \mathbb{Z}^{m \times n}$ gives rise to the toric ideal $I_{\ker_{\mathbb{Z}} A}$ which we denote by I_A. Thus if two matrices A and B have the same integer kernel then they define the same toric ideal. The variety defined by the vanishing of I_A is called the *toric variety associated to A* which we denote by $V_A \subseteq \mathbb{C}^n$. As V_A is irreducible, it has an associated algebraic matroid which turns out to be the linear matroid associated to the columns of A (Proposition 1.6.3).

Example 1.6.2. Let $L \subseteq \mathbb{R}^4$ be the one dimensional lattice consisting of all integer combinations of $u := (1, -1, -1, 1)$. We can construct a matrix A satisfying $L = \ker_{\mathbb{Z}} A$ as follows

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

Proposition 1.6.1 implies that $I_L = I_A$ is prime. It is the principal binomial ideal below

$$I_A = \langle x_1 x_4 - x_2 x_3 \rangle.$$

Proposition 1.6.3. *The algebraic matroid underlying a toric variety V_A is the linear matroid underlying the columns of A.*

Proof. This follows from the following two facts. First, the dimension of the toric variety V_A is equal to the rank of A [57, Lemma 4.2]. Second, the coordinate projection of a toric variety onto coordinates corresponding to a subset of columns of A is equal to the toric variety defined by the corresponding column submatrix. $\qquad\square$

1.6.1 Toric ideals in algebraic statistics

We now explain the importance of toric ideals for algebraic statistics. Let X be a discrete random variable with states $\{1, \ldots, n\}$ and let p_i denote the probability that $X = i$. To set things up geometrically, we view $p = (p_1, \ldots, p_n)$, the probability distribution on X,

as a point in $\mathbb{R}^n$. Since p is a probability distribution, we know

$$p \in \Delta^{n-1} := \{(x_1, \ldots, x_n) : x_i \geq 0 \text{ and } \sum_{i=1}^{n} x_i = 1\} \subset \mathbb{R}^n.$$

The set Δ^{n-1} is called the *probability simplex*. Any model (i.e. set of assumptions) we impose on X restricts p to lie in a subset of Δ^{n-1}. For many widely used statistical models, such a subset is simply the intersection of Δ^{n-1} with a toric variety.

Definition 1.6.4. Let $A \in \mathbb{Z}^{m \times n}$ be an integer matrix. The associated *log-linear model* $\mathcal{M}_A$ is the intersection of the probability simplex with the toric variety associated to A. In symbols,

$$\mathcal{M}_A := V_A \cap \Delta^{n-1}.$$

Understanding the structure of the toric ideal I_A corresponding to a log-linear model $\mathcal{M}_A$ offers many practical benefits. Perhaps the most salient example of this is detailed in [27, Section 1.1] where it is shown how one can use a set of binomial generators of I_A to test how well the model $\mathcal{M}_A$ fits a given dataset. We now give a simple example of a log-linear model.

Example 1.6.5. Let R be the random variable representing whether or not it rains tomorrow in Raleigh and let L be the random variable representing whether or not I pack a lunch. We will denote the two states of R by r and d for "rain" and "dry," and the two states of L by l and n for "lunch" and "no lunch." We will let p_r, p_d, p_l and p_n denote the probabilities of rain, no rain, lunch, and no lunch respectively.

From R and L we can make a new random variable X, called the *joint random variable*, which has four states, each corresponding to a possible outcome of the pair (R, L). We will denote these four states by rl, rn, dl, and dn where ij is the event that $R = i$ and $L = j$. So for example, the state rl is the event that it rains and I do pack a lunch. We will denote the corresponding probabilities as $p_{rl}, p_{rn}, p_{dl}, p_{dn}$.

One might assume that whether or not it rains has no effect on whether or not I pack a lunch, and vice versa. This corresponds to the random variables R and L being independent which means that the joint probabilities satisfy the polynomial

$$p_{rl}p_{dn} - p_{rn}p_{dl} = 0.$$

34

This is equivalent to the condition that the point $(p_{rl}, p_{rn}, p_{dl}, p_{dn})$ lies in the variety defined by the vanishing of the ideal below, which is the toric ideal from Example 1.6.2

$$\langle x_1 x_4 - x_2 x_3 \rangle.$$

The model of independence, described in Example 1.6.5, is a *hierarchical model*. Hierarchical models form a large subclass of log-linear models and they are able to capture many assumptions one might wish to make about categorical data. In particular, this class subsumes the class of *discrete graphical models* which are used in many applications [63]. The hierarchical models are naturally indexed by pairs $(\mathcal{C}, \mathbf{d})$ where $\mathcal{C}$ is a simplicial complex on $\{1, 2, \ldots, n\}$ and $\mathbf{d} = (d_1, \ldots, d_n)$ is an integer vector satisfying $d_i \geq 2$ for each i. One of the main results of Chapter 4 is a classification of the pairs $(\mathcal{C}, \mathbf{d})$ whose corresponding hierarchical model has a *unimodular* toric ideal.

1.6.2 Unimodularity

In this section we discuss unimodularity, which is a property that a toric ideal may satisfy. One of the main results of Chapter 4 is a classification of all hierarchical models whose corresponding toric ideals are unimodular. We defer the discussion of unimodularity's importance for hierarchical models to Chapter 4.

Let $u \in \mathbb{R}^n$ be a vector. Then we can write $u = u^+ - u^-$ where u^+ and u^- have nonnegative entries and disjoint support. Define a partial order $\preceq$ on $\mathbb{R}^n$ such that $u \preceq v$ if $u^+ \leq v^+$ and $u^- \leq v^-$. We write $u \prec v$ if $u \preceq v$ but $u \neq v$. Let $L \subseteq \mathbb{Z}^n$ be a lattice.

Definition 1.6.6. Let L be a lattice. The *Graver basis of L*, denoted $\mathrm{Gr}(L)$, is the set of all $u \in L$ that are minimal with respect to $\prec$.

A sum of vectors $v + w$ is said to be a *conformal sum* if both $v \prec v + w$ and $v \prec v + w$. Thus one can equivalently define the Graver basis of a lattice to be the set of $u \in L$ such that u cannot be expressed as a nontrivial conformal sum of $v, w \in L$.

Definition 1.6.7. Let $u \in L$. If u has relatively prime entries and minimal support among elements of L, then we say that u is a *circuit of L*. The set of all circuits of L is denoted $\mathcal{C}(L)$.

Note that if $L = \ker_{\mathbb{Z}} A$ for some $A \in \mathbb{Z}^{m \times n}$, then $u \in L$ is a circuit of L if and only if $\mathrm{sign}(u)$ is a signed circuit of the oriented matroid $\mathcal{O}_A$. It is easy to see that the set of

circuits is a subset of the Graver basis. However, the converse is not true in general. When $L = \ker_{\mathbb{Z}} A$ for a matrix A, we may write A in the place of $\ker_{\mathbb{Z}} A$. So for example, we may write $\mathrm{Gr}(A)$ instead of $\mathrm{Gr}(\ker_{\mathbb{Z}} A)$. We will identify binomials of the form $\mathbf{x}^{u^+} - \mathbf{x}^{u^-}$ with the corresponding vector $u = u^+ - u^-$ and say that $\mathbf{x}^{u^+} - \mathbf{x}^{u^-}$ is a circuit or lies in the Graver basis if u does.

Definition 1.6.8. A toric ideal I_A and its corresponding toric variety V_A are said to be *unimodular* if the Graver basis of A contains only $\{0, 1, -1\}$-vectors.

Example 1.6.9. Consider the 2×4 matrices A and B shown below

$$A := \begin{pmatrix} 1 & -1 & 0 & 1 \\ 0 & 0 & 1 & -1 \end{pmatrix} \qquad B := \begin{pmatrix} 1 & -1 & 0 & 1 \\ 1 & 0 & 1 & -1 \end{pmatrix}.$$

We claim that $\mathrm{Gr}(A)$ consists of the six vectors shown below. Since they are all $\{0, 1, -1\}$-vectors, it will follow that the toric ideal I_A and the toric variety V_A are unimodular.

$$\begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} -1 \\ -1 \\ 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 1 \\ 0 \\ -1 \\ -1 \end{pmatrix}, \quad \begin{pmatrix} -1 \\ 0 \\ 1 \\ 1 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ -1 \\ -1 \\ -1 \end{pmatrix}.$$

It is clear that these six vectors are contained in the Graver basis of A as one can check by hand that they are $\prec$-minimal elements of $\ker_{\mathbb{Z}} A$. To see that they constitute all of $\mathrm{Gr}(A)$, let $y = (a, b, c, d)^T \in \mathrm{Gr}(A)$. Then we must have $c = d$ and $a - b + c = 0$. If $a > 0$, then $b - c = a > 0$. In this case, y must be one of the two $\prec$-minimal vectors satisfying these two conditions, which are $(1, 1, 0, 0)$ and $(1, 0, -1, -1)$. Similarly, if $a < 0$ then y must be $(-1, -1, 0, 0)$ or $(-1, 0, 1, 1)$. If $a = 0$, then $b = c = d$ so y must be one of the two $\prec$-minimal vectors satisfying these conditions, which are $(0, 1, 1, 1)$ and its negative.

The matrix B is not unimodular. To see this it is sufficient to note that the vector $(1, 2, 0, 1)^T$ is in $\mathrm{Gr}(B)$.

Our general method for proving that a given matrix A is unimodular will be to compute its entire Graver basis using the software package 4ti2 [2] and directly verifying that every element is a $\{0, 1, -1\}$-vector. To prove that a given matrix is not unimodular we will use 4ti2 to identify a single Graver basis element that has an entry with absolute

value at least two. The following proposition gives us alternative ways to characterize unimodularity.

Proposition 1.6.10. *Let $L \subseteq \mathbb{Z}^n$ be a lattice. If $L = \ker_\mathbb{Z} A$ for some matrix $A \in \mathbb{Z}^{m \times n}$ then the following are equivalent.*

(1) *The toric ideal $I_A = I_L$ is unimodular*

(2) *Every circuit of L is a $\{0, 1, -1\}$ vector*

(3) *For any matrix A' obtained from $r := \operatorname{rank} A$ linearly independent rows of A, there exists some λ such that each $r \times r$ minor of A' is 0 or $\pm\lambda$*

(4) *For any b in the affine semigroup generated by the columns of A, the polyhedron $P_{A,b} = \{x \in \mathbb{R}^s : Ax = b, x \geq 0\}$ has all integral vertices.*

Moreover, every element of $\operatorname{Gr}(L)$ is a circuit; that is, $\mathcal{C}(L) = \operatorname{Gr}(L)$.

Proof. See [51], Theorem 19.2 for the equivalence of (3) and (4). We can see that (2) is equivalent to (3) because the circuits can be computed via a determinantal formula using Cramer's rule [57, p. 35]. Given the equivalence of (2) and (3), (2) and (1) are equivalent by [57], Propositions 4.11 and 8.11. The final remark follows by Proposition 8.11 in [57]. $\qquad\square$

Proposition 1.6.10(4) implies that when A is unimodular, integer programming problems over polyhedra $P_{A,b}$ can be solved via linear relaxation. In light of this, as a property of the matrix A, unimodularity can be viewed as a kernel-invariant generalization of total unimodularity (which requires that every minor is either 0 or ± 1). The final statement of Proposition 1.6.10 and item (2) tell us that classifying the Graver basis of a unimodular toric ideal I_A is equivalent to characterizing the oriented matroid underlying the columns of A.

Example 1.6.11. Let A and B be the matrices as defined in Example 1.6.9. Not surprisingly, each of the equivalent conditions given in Proposition 1.6.10 are satisfied by A but not by B. For example, note that each of the six 2×2 submatrices of A has determinant $0, 1$, or -1 whereas the submatrix formed by the first and last column of B has determinant -2.

Proposition 1.6.12. *Let $A \in \mathbb{Z}^{m \times n}$ be unimodular. Then the following are also unimodular matrices.*

(1) A matrix obtained by reordering the columns of A

(2) A matrix obtained by multiplying a column of A by -1

(3) Any matrix that has the same rowspace as A

(4) Any matrix formed by a subset of columns of A.

Proof. Let ϕ_i denote the map that sends A to the matrix specified in item (i) above. Each ϕ_i induces a map $\phi_i^* : \ker_{\mathbb{Z}} A \to \ker_{\mathbb{Z}} \phi(A)$. In each case, we will see that ϕ_i^*, and therefore ϕ_i, preserves unimodularity. In particular, ϕ_1^* permutes coordinates, ϕ_2^* multiplies a single coordinate by -1, ϕ_3^* is the identity, and ϕ_4^* is a coordinate projection. It is clear that the first three maps preserve unimodularity. To see that coordinate projections also preserve unimodularity, let L be a unimodular lattice and let πL be a coordinate projection. If $\pi(u) \in \pi L$ lies in the Graver basis of πL but has a coordinate i such that $|\pi(u)_i| \geq 2$, then u cannot lie in the Graver basis of L and so we can express $u = a + b$ as a conformal sum with $a, b \in L$. But then $\pi u = \pi a + \pi b$, and $\pi a, \pi b \in \pi L$ and the sum is conformal. So πu is not in the Graver basis of L. $\qquad\square$

Chapter 2

Low-rank matrix completion

The material in this chapter was published in the journal *Linear Algebra and its Applications* [6].

Given a matrix where only some of the entries are known, the low-rank matrix completion problem is to determine the missing entries under the assumption that the matrix has some low rank r. One can also assume additional structure such as (skew) symmetry or positive definiteness of the matrix. Practical applications of the low-rank matrix completion problem abound. A well-known example is the so-called "Netflix Problem" of predicting an individual's movie preferences from ratings given by several other users. A brief survey of other applications appears in [16].

Singer and Cucuringu [55] show how ideas from rigidity theory can be applied to this problem. Jackson, Jordán, and Tanigawa [35, 36] further develop these ideas. Király, Theran, and Tomioka [39] incorporate ideas from algebraic geometry into this rigidity-theoretic framework and Király, Theran and Rosen [40] further develop these ideas. We add tools from tropical geometry to this picture.

Let V be a determinantal variety over some algebraically closed field $\mathbb{K}$. The results in this chapter concern the cases where $V = \mathcal{S}_r^n(\mathbb{K})$, the collection of $n \times n$ skew-symmetric $\mathbb{K}$-matrices of rank at most r, or $V = \mathcal{M}_r^{m \times n}(\mathbb{K})$, the collection of $m \times n$ $\mathbb{K}$-matrices of rank at most r. A *masking operator* corresponding to some $S \subseteq \binom{[n]}{2}$ in the skew symmetric case, or $S \subseteq [m] \times [n]$ in the rectangular case, is a map $\Omega_S : V \to \mathbb{K}^S$ that projects a matrix M onto the entries specified by S. In the case of skew-symmetric $n \times n$ matrices, we view S as the edge set of a graph on vertex set $[n]$, which we denote $G(S)$. In the case of rectangular matrices, we view S as the edge set of a bipartite graph on

partite sets of size m and n which we also denote $G(S)$. Context will make the proper interpretation of $G(S)$ clear.

A low-rank matrix completion problem can now be phrased as: given $\Omega_S(M)$ can we recover M if we know $M \in V$? For generic M the answer to this question only depends on the observed entries S and not the particular values observed. Namely, given $\Omega(M)$ for generic $M \in V$, M may be recovered up to finitely many choices if and only if S is a spanning set of the algebraic matroid associated to V. Hence it is useful to find combinatorial descriptions of the algebraic matroids associated to various determinantal varieties. We obtain such combinatorial descriptions for the cases where $V = \mathcal{S}_2^n(\mathbb{C})$ and $V = \mathcal{M}_2^{m \times n}(\mathbb{C})$. The most natural way to phrase our characterization is in terms of the independent sets of V. Note that a subset S of entries is independent in the algebraic matroid underlying V if and only if $\Omega_S(V)$ is Zariski dense in $\mathbb{C}^S$. The salient feature for independent sets is that $\mathbb{C} \setminus \Omega_S(V)$ has Lebesgue measure zero. We also note that our result here answers a question of Kalai, Nevo, and Novik in [38] to find a combinatorial classification of what they call "minimally $(2,2)$-rigid graphs" (posed in the paragraph after the proof of their Example 5.5). Using our language, these are the maximal independent sets in the algebraic matroid underlying $\mathcal{M}_2^{m \times n}(\mathbb{C})$.

We will state the main result of this chapter, but first we must make a definition. An alternating closed trail in a directed graph is a walk $v_0, v_1, \ldots, v_k$ such that each edge appears at most once, $v_k = v_0$, and adjacent edges $v_{i-1}v_i$ and $v_i v_{i+1}$ have opposite orientations (indices taken modulo $k+1$). The main result of this chapter is the following.

Theorems 2.3.2 and 2.3.4. *Let $V = \mathcal{S}_2^n(\mathbb{C})$ be the variety of skew-symmetric $n \times n$ matrices of rank at most 2, or $V = \mathcal{M}_2^{m \times n}(\mathbb{C})$ be the variety of rectangular $m \times n$ matrices of rank at most 2. A subset of observed entries $S \subseteq \binom{[n]}{2}$ (skew symmetric case) or $S \subseteq [m] \times [n]$ (rectangular case) is independent in the algebraic matroid underlying V if an only if there exists some acyclic orientation of $G(S)$ that has no alternating closed trail.*

Using techniques of [39], one can see that deciding whether a given $S \subseteq \binom{[n]}{2}$ or $S \subseteq [m] \times [n]$ is independent in the algebraic matroid underlying $\mathcal{S}_2^n(\mathbb{C})$ or $\mathcal{M}_2^{m \times n}(\mathbb{C})$ is in the complexity class RP. Hence both decision problems are also in NP. We provide an explicit combinatorial certificate of this fact. Existence of a polynomial time algorithm for solving either decision problem remains open.

The key to our approach is to use tropical geometry to allow us to reduce to the

following easier question: Which entries of a dissimilarity map may be arbitrarily specified such that the resulting partial dissimilarity map may be completed to a tree metric? Our result here is as follows.

Theorem 2.2.5. *Let $S \subseteq \binom{[n]}{2}$. Any partial dissimilarity map whose known distances are given by S can be completed to a tree metric, regardless of what those specified values are, if and only if there exists some acyclic orientation of $G(S)$ that has no alternating closed trail.*

As in the case of partial matrices, we give an explicit combinatorial certificate showing that the corresponding decision problem is in NP but we do not know whether a polynomial time algorithm exists.

The problem of deciding whether a *particular* partial dissimilarity map is completable to a tree metric was shown to be NP-complete in [29]. Note that this is distinct from our decision problem here, because we are not setting values of the observed entries. Special cases that allow a polynomial time algorithm were investigated in [31, 30]. Questions about whether a partial dissimilarity map can be completed to a tree metric with a particular topology have been addressed in [25, 26].

The outline of this chapter is as follows. Section 2.1 lays out some general theory for using tropical geometry to characterize algebraic matroids. Section 2.2 contains our results relating to completion of tree metrics. Section 2.3 shows how our results on partial matrices are easily obtained from our results on trees.

2.1 Completion and tropical varieties

We begin with the necessary preliminaries from tropical geometry. For a more leisurely treatment of this material, see Section 1.5. The most important parts of this section are Lemmas 2.1.4 and 2.1.5 which enable us to use the polyhedral structure of a tropical variety to gain insight into the corresponding algebraic matroid.

Denote by $\mathbb{C}\{\{t\}\}$ the field of complex formal Puiseux series. That is, $\mathbb{C}\{\{t\}\}$ is the set of all formal sums $\sum_{\alpha \in J} c_\alpha t^\alpha$ where $J \subset \mathbb{Q}$ such that J has a smallest element and the elements of J can be expressed over a common denominator. The valuation map $\mathrm{val}(\cdot) : \mathbb{C}\{\{t\}\} \to \mathbb{Q}$ sends $\sum_{\alpha \in J} c_\alpha t^\alpha$ to $\min\{\alpha \in J : c_\alpha \neq 0\}$. For any affine variety V

over $\mathbb{C}\{\{t\}\}$, the corresponding *tropical variety* is

$$\mathrm{trop}(V) = \overline{\{(-\,\mathrm{val}(x_1), \ldots, -\,\mathrm{val}(x_n)) : (x_1, \ldots, x_n) \in V\}} \subseteq \mathbb{R}^n$$

where the overline indicates closure in the Euclidean topology on $\mathbb{R}^n$. We sometimes refer to $\mathrm{trop}(V)$ as *the tropicalization of V*.

We can also tropicalize varieties over $\mathbb{C}$ by lifting to $\mathbb{C}\{\{t\}\}$ and tropicalizing there. More specifically, let $V \subseteq \mathbb{C}^n$ be an affine variety over $\mathbb{C}$ with ideal $I \subseteq \mathbb{C}[x_1, \ldots, x_n]$. By lifting this ideal into $\mathbb{C}\{\{t\}\}[x_1, \ldots, x_n]$ we obtain a variety $V' \subseteq (\mathbb{C}\{\{t\}\})^n$. The tropical variety $\mathrm{trop}(V)$ corresponding to V is simply $\mathrm{trop}(V')$.

As we discussed in Section 1.5.3, tropicalizing an irreducible complex variety preserves the underlying algebraic matroid structure. Before restating this precisely in Lemma 2.1.1, we review some definitions from the introduction.

Let $\mathbb{K}$ be a field and let $V \subset \mathbb{K}^n$. We let $M(V)$ denote the *independence complex* of S which we define to be the collection of subsets S of $\{1, \ldots, n\}$ such that the projection of V onto the coordinates indicated by S is full-dimensional in $\mathbb{K}^S$. When V is an irreducible variety, then $M(V)$ is called the *algebraic matroid underlying V*.

Lemma 2.1.1 ([64], Lemma 2). *Let V be an irreducible affine variety over either $\mathbb{C}$ or $\mathbb{C}\{\{t\}\}$. Then the independence complex of V and $\mathrm{trop}(V)$ are the same.*

Many matrix completion problems ask for a combinatorial description of the algebraic matroid associated to a particular irreducible affine variety. Lemma 2.1.1 says that we can tackle such problems by looking at the corresponding tropical variety instead. The advantage in this is that tropical varieties have a useful polyhedral structure which we now describe.

Definition 2.1.2. Let Σ be a rational fan in $\mathbb{R}^n$ of pure dimension d. We say that Σ is *balanced* if we can associate a positive integer $m(\sigma)$, called the *multiplicity*, to each top-dimensional cone σ in a way such that for each cone $\tau \in \Sigma$ of dimension $d-1$,

$$\sum_{\sigma \supsetneq \tau} m(\sigma) v_\sigma \in \mathrm{span}(\tau)$$

where v_σ is the first lattice point on the ray $\sigma/\mathrm{span}(\tau)$. We say that Σ is *connected through codimension one* if for any d-dimensional cones $\sigma, \rho \in \Sigma$, there exists a sequence of d-dimensional cones $\sigma = \sigma_0, \sigma_1, \ldots, \sigma_k = \rho \in \Sigma$ such that $\sigma_i \cap \sigma_{i+1}$ has dimension $d-1$.

The well-known structure theorem for tropical varieties applied to the special case where the defining equations have constant coefficients gives us the following.

Theorem 2.1.3 ([45], Theorem 3.3.5). *Let V be an irreducible d-dimensional affine variety over $\mathbb{C}$. Then $\mathrm{trop}(V)$ is the support of a balanced fan of pure dimension d that is connected through codimension one.*

The following two lemmas give us ways that one can use the polyhedral structure of $\mathrm{trop}(V)$ to gain insight into the structure of $M(V)$. When we refer to a cone in $\mathrm{trop}(V)$, we mean in a polyhedral subdivision of $\mathrm{trop}(V)$ that is balanced and connected through codimension one.

Lemma 2.1.4. *Let $V \subseteq \mathbb{C}^n$ be an irreducible affine variety of dimension d. Then $S \subseteq [n]$ is independent in $M(V)$ if and only if S is independent in $M(\mathrm{span}(\sigma))$ for some d-dimensional cone σ in $\mathrm{trop}(V)$.*

Proof. Let S be independent in $M(V)$. By Lemma 2.1.1, the projection of $\mathrm{trop}(V)$ onto $\mathbb{R}^S$ is full dimensional. In particular, the projection of some maximal dimensional cone $\sigma \in \mathrm{trop}(V)$ onto $\mathbb{R}^S$ is full dimensional and therefore S is independent in the matroid $M(\mathrm{span}(\sigma))$.

Now let S be independent in $M(\mathrm{span}(\sigma))$ for some maximal cone $\sigma \in \mathrm{trop}(V)$. Then the projection of $\mathrm{span}(\sigma)$ onto $\mathbb{R}^S$ is full dimensional. Therefore the same holds for σ and therefore $\mathrm{trop}(V)$. So S is in the independence complex of $\mathrm{trop}(V)$ and therefore independent in $M(V)$ by Lemma 2.1.1. $\square$

Lemma 2.1.5. *Let $V \subseteq \mathbb{C}^n$ be an irreducible d-dimensional affine variety. Let τ be a $d-1$ dimensional cone of $\mathrm{trop}(V)$ and let $\sigma_1, \ldots, \sigma_k$ be the d-dimensional cones in $\mathrm{trop}(V)$ containing τ. If $B \subseteq [n]$ is a basis of $M(\mathrm{span}(\sigma_1))$ then B is also a basis of $M(\mathrm{span}(\sigma_i))$ for some $i \neq 1$.*

Proof. Let $L = \mathrm{span}(\tau)$ and $L_i = \mathrm{span}(\sigma_i)$. By Theorem 2.1.3 there exist $v_i \in \sigma_i \setminus L$ such that $\sum_{i=1}^{k} v_i \in L$. Let B be a basis of L_1. Then the projection of L_1 onto $\mathbb{R}^B$ is all of $\mathbb{R}^B$ and the projection of L onto $\mathbb{R}^B$ is some hyperplane $\{x \in \mathbb{R}^B : c^T x = 0\}$. By padding with extra zeros, we can extend c to an element of $\mathbb{R}^n$. Then $L = L_1 \cap \{x \in \mathbb{R}^n : c^T x = 0\}$ and $c^T v_i \neq 0$. Since $\sum_{i=1}^{k} v_i \in L$, we must have $\sum_{i=1}^{k} c^T v_i = 0$. Therefore there must exist some $i \neq 1$ such that $c^T v_i \neq 0$. Since $L_i = L + \mathrm{span}(v_i)$, the projection of L_i onto $\mathbb{R}^B$ is all of $\mathbb{R}^B$. So B is a basis of L_i. $\square$

We end this section by noting a nice feature about projections of tropical varieties which was noted by Yu in [64].

Proposition 2.1.6. *Let $\mathbb{K}$ be either $\mathbb{C}$ or $\mathbb{C}\{\{t\}\}$ and let $V \subseteq \mathbb{K}^n$ be an irreducible affine variety. If S is independent in $M(V)$ then the projection of $\mathrm{trop}(V)$ into $\mathbb{R}^S$ is all of $\mathbb{R}^S$.*

2.2 Tree metrics and tree matroids

In this section we determine which entries of a dissimilarity map can be arbitrarily specified while still allowing completion to a tree metric. We begin by reviewing the necessary preliminaries on tree metrics. Proposition 2.2.1 describes the tropical and polyhedral structure on the set of tree metrics. Lemma 2.2.2 reduces our question about completion to an arbitrary tree metric to the question about completion to a tree metric whose topology is a caterpillar tree. We answer this simpler question in Proposition 2.2.3 and give the section's main result in Theorem 2.2.5.

We now quickly review the background about tree metrics that was given in Section 1.4. Let X be a set. A *dissimilarity map on X* is a function $\delta : X \times X \to \mathbb{R}$ such that $\delta(x, x) = 0$ and $\delta(x, y) = \delta(y, x)$ for all $x, y \in X$. If $X = \{x_1, \ldots, x_n\}$ and $[n] = \{1, \ldots, n\}$ then there is a natural bijection between dissimilarity maps on X and points in $\mathbb{R}^{\binom{[n]}{2}}$. Namely, denoting the entry of a point $d \in \mathbb{R}^{\binom{[n]}{2}}$ corresponding to $\{i, j\}$ by d_{ij} or d_{ji}, we associate a dissimilarity map δ on X with the point $d \in \mathbb{R}^{\binom{[n]}{2}}$ such that $d_{ij} = \delta(x_i, x_j)$. Hence we will often speak of dissimilarity maps as if they were merely points in $\mathbb{R}^{\binom{[n]}{2}}$.

A tree with no vertices of degree 2 with leaf set X is called an *X-tree*. Let T be an X-tree and let w be an edge-weighting on T such that $w(e) > 0$ for all internal edges e of T. Some authors require that $w(e) > 0$ for *any* edge e of T, but we do not. The triple (T, X, w) gives rise to a dissimilarity map d where d_{ij} is the sum of the edge weights given by w along the unique path from x_i to x_j in T. Any dissimilarity map d that arises from a triple in this way is called a *tree metric*. For example, if $X = \{1, 2, 3, 4\}$, then the dissimilarity map $d = (d_{12}, d_{13}, d_{14}, d_{23}, d_{24}, d_{34}) = (0, 3, -2, 5, 0, -1)$ is a tree metric because it can be displayed on a tree as in Figure 1.1. Given a tree metric d on X, there is a unique X-tree T that realizes d as such [53, Theorem 7.1.8]. This tree T is called the *topology* of d. A tree T is *binary* if all internal vertices have degree three.

Authors who require that all edge weights be positive in the definition of a tree metric have a similar four point condition. The only difference is that i, j, k, l need not

44
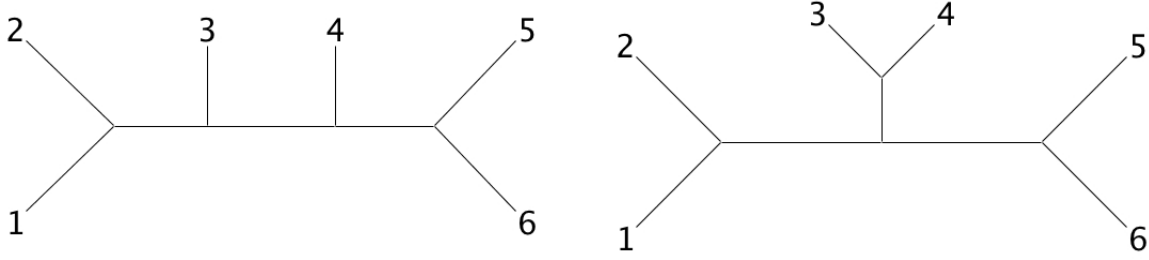
Figure 2.1: The tree on the left has cherries 12 and 56 hence it is a caterpillar. The tree on the right has cherries 12, 34 and 56 and is therefore *not* a caterpillar.

be distinct. We denote the set of all tree metrics on a set of size n by $\mathcal{T}_n$.

Let T be a tree with leaf labels $[n]$. A *cherry* on a binary $[n]$-tree T is a pair of leaves ij such that i and j are adjacent to the same vertex of degree 3. A caterpillar tree on $n \geq 4$ vertices is a binary tree with exactly two cherries. See Figure 2.1 for an illustration.

The following proposition summarizes many known facts about the polyhedral and tropical structure of $\mathcal{T}_n$ (references are given in the proof). In particular, it tells us how to realize $\mathcal{T}_n$ as a tropical variety along with the polyhedral subdivision guaranteed to exist by Theorem 2.1.3. We note that it would not be true had we taken the more restrictive definition of tree metric that required all (as opposed to just internal) edge weights to be positive. Recall that the Grassmannian $\mathrm{Gr}(k,n) \subset \mathbb{K}^{\binom{[n]}{k}}$ is the irreducible affine variety parameterized by the $k \times k$ minors of a $k \times n$ matrix over $\mathbb{K}$.

Proposition 2.2.1. *The space of phylogenetic trees $\mathcal{T}_n$ is $\mathrm{trop}(\mathrm{Gr}(2,n))$. We can give $\mathcal{T}_n$ a balanced fan structure connected through codimension one as follows. Each open cone is the collection of tree metrics whose topology is T for a particular tree T. Such a cone is maximal if and only if T is binary. The cone has codimension one if and only if T can be obtained from a binary tree by contracting exactly one edge. Given a cone τ of codimension one corresponding to tree T, the cones containing τ correspond to the three binary trees that can be contracted to T.*

Proof. Let L_n be the set of tree metrics whose topology is a star tree - that is, a tree with no internal edges. Since an edge-weighting of a leaf-labeled tree gives rise to a tree metric if and only if all *internal* edge weights are nonnegative, L_n is the lineality space of $\mathcal{T}_n$. Theorem 3.4 in [56] implies that $\mathcal{T}_n/L_n = \mathrm{trop}(\mathrm{Gr}(2,n))/L_n$. But as L_n is also the lineality space of $\mathrm{Gr}(2,n)$ (c.f. remarks following Corollary 3.1 in [56]), this implies that

45

$$
\begin{array}{c c c c c c c}
 & 12 & 13 & 14 & 23 & 24 & 34 \\
a & 1 & 1 & 1 & 0 & 0 & 0 \\
b & 1 & 0 & 0 & 1 & 1 & 0 \\
c & 0 & 1 & 1 & 1 & 1 & 0 \\
d & 0 & 1 & 0 & 1 & 0 & 1 \\
e & 0 & 0 & 1 & 0 & 1 & 1
\end{array}
$$
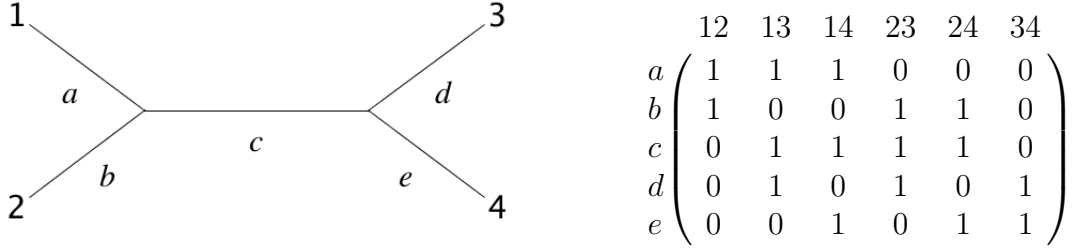
Figure 2.2: Let T be the tree on the left with leaves $\{1,2,3,4\}$ and edges labels $\{a, b, c, d, e\}$. The matrix on the right is A_T. Its columns are the $\lambda_T(ij)$'s.

$\mathcal{T}_n = \mathrm{trop}(\mathrm{Gr}(2, n))$.

The polyhedral structure is given in [13]. Connectedness through codimension one follows from the fact that any two binary trees on the same leaf set can be reached from one another via a finite sequence of nearest-neighbor-interchanges (see [62, Theorem 2]). We can see that this polyhedral fan is balanced by assigning multiplicity 1 to each maximal cone (see Remark 4.3.11 and Theorem 3.4.14 in [45]). □

Let T be a tree without vertices of degree two whose leaves are labeled by $[n]$. Denote its edge set by E and let $\mathbb{R}^E$ denote the vector space of edge-weightings of T. For each pair $ij \in \binom{[n]}{2}$, define $\lambda_T(ij) \in \mathbb{R}^E$ by $\lambda_T(ij)_e = 1$ whenever e is on the unique path from the leaf labeled i to the leaf labeled j and 0 otherwise. We denote the linear matroid underlying these $\lambda_T(ij)$ by $M(T)$ and call it the *matroid associated to T*. We will say that a set $S \subseteq \binom{[n]}{2}$ is independent in $M(T)$ to mean that $\{\lambda_T(ij)\}_{ij \in S}$ is independent in $M(T)$. Let A_T be the matrix whose columns are the $\lambda_T(ij)$s. An example is given in Figure 2.2. Note that $\mathrm{row}(A_T)$ is the linear span of the cone in $\mathcal{T}_n$ containing all tree metrics with topology T. Moreover, $M(T)$ is the matroid of this linear space. These matroids were introduced and studied in [25].

Lemma 2.2.2. *A set $B \subseteq \binom{[n]}{2}$ is a basis of $M(\mathcal{T}_n)$ if and only if B is a basis in the matroid associated to some caterpillar tree.*

Proof. Lemma 2.1.4 implies that S is independent in $M(\mathcal{T}_n)$ if and only if S is independent in $M(T)$ for some binary tree T. We proceed by showing that if B is a basis of $M(T)$ for a non-caterpillar binary tree T, then B is also the basis of $M(T')$ for some binary tree T' with one fewer cherry. It will follow by induction that B is a basis of $M(C)$ for some caterpillar C.

So let T be a binary tree with three distinct cherries ij, $i'j'$, $i''j''$ with corresponding degree-3 vertices p, p', p''. There is a single vertex q in the common intersection of the three paths p to p', p' to p'', and p to p''. If we delete q from T we get three trees U, V, W which contain the cherries ij, $i'j'$, and $i''j''$ respectively.

We now induct on the number of leaves in U. The base case is where U consists of the two leaves i and j. We can visualize this as in Figure 2.3a letting U' be i and U'' be j. Lemma 2.1.5 and Proposition 2.2.1 imply that B must also be a basis of one of the trees T_1 or T_2, depicted in Figures 2.3b and 2.3c respectively. Note that T_1 and T_2 each have fewer cherries than T.

Now assume U has more than two leaves. Let s be the first node on the path from q to p. Deleting s from U splits it into two subtrees which we denote U' and U''. This is depicted in Figure 2.3a. Assume ij belongs to the subtree U'. Using Lemma 2.1.5 and Proposition 2.2.1 as before, we see that B must also be a basis of one of the trees T_3 or T_4, depicted in Figures 2.3b and 2.3c respectively.
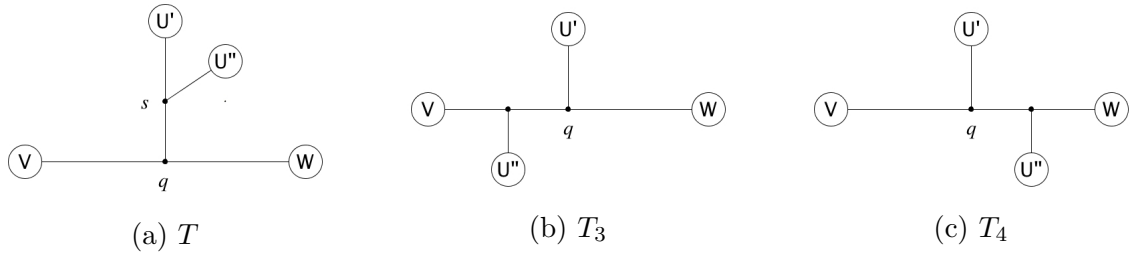
Figure 2.3: Breaking T into subtrees.

We can now repeat the process with U' taking the place of U. In the case that B is a basis of T_3 we expand V to include U'', and in the case of T_4 we instead expand W to contain U''. The result has strictly fewer leaves in U and so we induct. $\qquad\square$

Lemma 2.2.2 reduces the problem of describing the independent sets of $M(\mathcal{T}_n)$ to the simpler problem of describing the independent sets of $M(T)$ where T is a caterpillar tree. Luckily, such independent sets have a simple combinatorial description. Given a subset $S \subseteq \binom{[n]}{2}$, we let $G(S)$ denote the graph with vertex set $[n]$ and edge set S. For each n, we let $\mathrm{Cat}(n)$ denote the caterpillar with cherries $1, 2$ and $n-1, n$ such that the other leaf labels increase from 3 to $n-2$ along the path from the $1, 2$ cherry to the $n-1, n$

cherry. This is depicted in Figure 2.4. Recall that a *closed trail* in a graph is a sequence of vertices $v_0, \ldots, v_k$ such that each (v_i, v_{i+1}) is an edge (indices taken mod $k+1$), no edge is repeated, and $v_0 = v_k$.
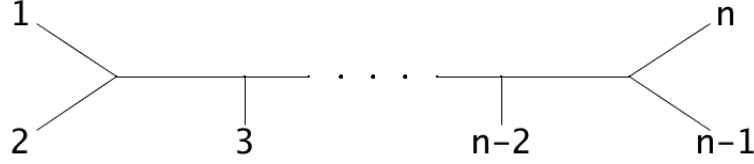
Figure 2.4: The caterpillar $\mathrm{Cat}(n)$.

Proposition 2.2.3. *Let $S \subseteq \binom{[n]}{2}$. Then S is independent in $M(\mathrm{Cat}(n))$ if and only if $G(S)$ contains no closed trail with alternating vertices. That is, S contains no closed trail of the form $v_0, v_1, \ldots, v_{2k-1}$ where $v_{2i-1}, v_{2i+1} < v_{2i}$ for each $i = 0, \ldots, k-1$ where indices are taken mod $2k$.*

Proof. Let $A_{\mathrm{Cat}(n)}$ be the matrix whose columns are given by the $\lambda_{\mathrm{Cat}(n)}(ij)$s. We claim that if $x \in \mathbb{R}^{\binom{[n]}{2}}$ then $x \in \mathrm{row}(A_{\mathrm{Cat}(n)})$ if and only if $x_{ik} + x_{jl} - x_{il} - x_{jk} = 0$ for all $i < j < k < l$. To see this, note that $\mathrm{row}(A_{\mathrm{Cat}(n)})$ is equal to the linear hull of the collection of tree metrics with topology $\mathrm{Cat}(n)$. If x is a tree metric with topology $\mathrm{Cat}(n)$, then each quartet $i < j < k < l$ must satisfy $x_{ij} + x_{kl} < x_{ik} + x_{jl} = x_{il} + x_{jk}$. Since the topology of a tree metric is determined by its quartets [53] the relations $x_{ik} + x_{jl} - x_{il} - x_{jk} = 0$ define the linear hull of the set of tree metrics with topology $\mathrm{Cat}(n)$.

Now let B be the vertex-edge incidence matrix of a complete bipartite graph on n vertices. That is, the columns of B are indexed by the set $[n] \times [n]$, the rows are indexed by two disjoint copies of $[n]$, and the column corresponding to (i, j) has 1s at the entries corresponding to i in the first copy of $[n]$ and j in the second copy, and 0s elsewhere. Then $x \in \mathrm{row}(B)$ if and only if $x_{(i,k)} + x_{(j,l)} - x_{(i,l)} - x_{(j,k)} = 0$ for all i, j, k, l.

Notice that any functional vanishing on $\mathrm{row}(A_{\mathrm{Cat}(n)})$ can be associated to a functional vanishing on $\mathrm{row}(B)$ in a coordinate-wise way. In particular, for $i < j$ we associate each coordinate ij in $\mathrm{row}(A_{\mathrm{Cat}(n)})$ with the coordinate (i, j) in $\mathrm{row}(B)$. This means that $M(\mathrm{Cat}(n))$, which is the column matroid of $A_{\mathrm{Cat}(n)}$, is the restriction of the column matroid of B to ground set $\{x_{(i,j)}\}$ where $i < j$.

The column matroid of B is the polygon matroid on the complete bipartite graph on two disjoint copies of $[n]$. To realize $M(\mathrm{Cat}(n))$ as a submatroid, we restrict to the collection of edges (i, j) from $\{1, \ldots, n-1\}$ to $\{2, \ldots, n\}$ such that $i < j$. Closed trails supported on such edges in this bipartite graph correspond exactly to alternating closed trails in the complete graph on vertex set $[n]$. $\square$

Remark 2.2.4. We now have a purely combinatorial proof that the problem of deciding whether a given $S \subseteq \binom{[n]}{2}$ is independent in $M(\mathcal{T}_n)$ is in the complexity class NP. Namely, let H_n be the bipartite graph on partite sets $A = \{1, \ldots, n-1\}$ and $B = \{2, \ldots, n\}$ where (i, j) is an edge from A to B if and only if $i < j$. By Proposition 2.2.3, $S \subseteq \binom{[n]}{2}$ is independent in $M(\mathrm{Cat}_n)$ if and only if the corresponding edges in H_n form no cycles. Therefore, by Lemma 2.2.2, a polynomial-time verifiable certificate that a given $S \subseteq \binom{[n]}{2}$ is independent in $M(\mathcal{T}_n)$ is a permutation σ of $[n]$ such that the edges $\{(\sigma(i), \sigma(j)), (i, j) \in S\}$ of H_n induce no cycles.

Recall that an acyclic orientation of a graph is an assignment of directions to each edge in a way such that produces no directed cycles. An alternating closed trail in a directed graph is a closed trail $v_0, \ldots, v_k$ such that each pair of adjacent edges $v_{i-1}v_i$ and $v_i v_{i+1}$ have opposite orientations. We remind the reader of Proposition 2.1.6 which implies that if $S \subseteq \binom{[n]}{2}$ is independent in $M(\mathcal{T}_n)$, then one may arbitrarily specify the distances among pairs in S and still be able to extend the result to a tree metric. We are now ready to prove Theorem 2.2.5 which we restate below using the language of this section.

Theorem 2.2.5. *Let $S \subseteq \binom{[n]}{2}$. Then S is independent in $M(\mathcal{T}_n)$ if and only if some acyclic orientation of $G(S)$ has no alternating closed trails.*

Proof. Assume that some acyclic orientation of $G(S)$ has no alternating closed trails. Fix such an acyclic orientation. Choose a permutation $\sigma : [n] \to [n]$ such that edge ij is oriented from i to j if and only if $\sigma(i) < \sigma(j)$. Then by Proposition 2.2.3, applying σ to the vertices of $G(S)$ gives an independent set in $M(\mathrm{Cat}(n))$. Since independence in $M(\mathcal{T}_n)$ is invariant under permutation of the leaves, this implies that S is independent in $M(\mathcal{T}_n)$.

Now assume that every acyclic orientation of $G(S)$ has an alternating closed trail. Then there does not exist any permutation of $[n]$ that produces an independent set of

$M(\mathrm{Cat}(n))$ when applied to S. Therefore S is not independent in $M(T)$ for any caterpillar T. Lemma 2.2.2 then implies that S is not independent in $M(\mathcal{T}_n)$. □

2.3 Rank two matrices

We show how Theorem 2.2.5 immediately characterizes the algebraic matroid underlying the set of $n \times n$ skew-symmetric matrices of rank at most 2. We then use this to characterize the algebraic matroid underlying the set of general $m \times n$ matrices of rank at most 2.

Denote by $\mathcal{S}_r^n(\mathbb{K})$ the collection of $n \times n$ skew-symmetric $\mathbb{K}$-matrices of rank at most r and denote by $\mathcal{M}_r^{m \times n}(\mathbb{K})$ the collection of $m \times n$ $\mathbb{K}$-matrices of rank at most r. Both are irreducible algebraic varieties. We will only be concerned with these sets when $\mathbb{K}$ is $\mathbb{C}$ or $\mathbb{C}\{\{t\}\}$. Lemma 2.1.1 implies that limited to these choices, the algebraic matroid does not depend on the ground field and so we will suppress $\mathbb{K}$ from notation.

Proposition 2.3.1. *The tropicalization of the set of $n \times n$ skew symmetric $\mathbb{C}\{\{t\}\}$-matrices of rank at most 2 is the set of all tree metrics on n leaves. That is,* $\mathrm{trop}(\mathcal{S}_2^n(\mathbb{C})) = \mathcal{T}_n$.

Proof. It is shown in [56] that $\mathrm{trop}(\mathrm{Gr}(2, n)) = \mathcal{T}_n$. Therefore the proposition follows from the claim that the projection of $\mathcal{S}_2^n(\mathbb{C})$ onto the upper-triangular coordinates is $\mathrm{Gr}(2, n)$. This claim seems to be a well-known fact, but it is difficult to find a precise reference so we give a proof here.

Let $\{x_{ij} : 1 \leq i < j \leq n\}$ be indeterminates and let M be the $n \times n$ skew-symmetric matrix whose ij entry is x_{ij} whenever $i < j$. Let $I_{2,n} \subseteq \mathbb{C}\{\{t\}\}[x_{ij} : 1 \leq i < j \leq n]$ be the radical of the ideal generated by the 3×3 minors of M. The variety of $I_{2,n}$ is the projection onto the upper-triangular coordinates of the set of all $n \times n$ skew-symmetric matrices of rank at most 2. Moreover, $I_{2,n}$ contains all principal 4×4 minors of M, and these polynomials are the squares of the canonical generating set of $\mathrm{Gr}(2, n)$. Therefore $\mathrm{Gr}(2, n)$ contains the variety of $I_{2,n}$. Since $\mathrm{Gr}(2, n)$ is irreducible, as it is parameterized by the 2×2 determinants of a generic $2 \times n$ matrix, it can be seen to be equal to the variety of $I_{2,n}$ by showing that the variety of $n \times n$ skew-symmetric matrices of rank at most 2 has dimension at least that of $\mathrm{Gr}(2, n)$. Since the dimension of $\mathcal{T}_n$ is $2n - 3$ (see e.g. [53]), Theorem 2.1.3 implies that $\mathrm{Gr}(2, n)$ has dimension $2n - 3$ as well. We can see

that the set of skew-symmetric matrices of rank at most 2 has at least this dimension because in constructing such a matrix we may arbitrarily specify all but the first entry of the first row and all but the first and second entries of the second row. $\qquad\square$

Theorem 2.3.2. *Let $S \subseteq \binom{[n]}{2}$. Then S is independent in $M(\mathcal{S}_2^n)$ if and only if some acyclic orientation of $G(S)$ has no alternating closed trails.*

Proof. This follows from Lemma 2.1.1 and Proposition 2.3.1 and Theorem 2.2.5. $\qquad\square$

Remark 2.3.3. The same combinatorial certificate as in Remark 2.2.4 can be verified in polynomial time to show that that a given $S \subseteq \binom{[n]}{2}$ is independent in $M(\mathcal{S}_2^n)$.

The dimension of $\mathcal{S}_2^n$ is $2n - 3$. Therefore the rank of the matroid $M(\mathcal{S}_2^n)$ is also $2n - 3$. Also note that $M(\mathcal{S}_2^k)$ is a restriction of $M(\mathcal{S}_2^n)$ whenever $k < n$. It follows that if S is a basis of $M(\mathcal{S}_2^n)$, then $G(S)$ has $2n - 3$ edges and each subgraph of size k has at most $2k - 3$ edges. It is a famous theorem of Laman [42] that graphs satisfying these constraints are exactly the minimal graphs which are generically infinitesimally rigid in the plane. These graphs are often called "Laman graphs." One might wonder whether all Laman graphs are bases of $M(\mathcal{S}_2^n)$ but this is not the case. A counterexample is $K_{3,3}$, the complete bipartite graph on two partite sets of size three (see Figure 2.5). This is a Laman graph but it is dependent in $M(\mathcal{S}_2^n(\mathbb{C}))$ because the coordinates specified by $K_{3,3}$ are the entries in a 3×3 submatrix. Such a submatrix must be singular in a matrix of rank 2 and so those entries satisfy a polynomial relation. Alternatively, one could appeal to Theorem 2.3.2 and check that every acyclic orientation of $K_{3,3}$ induces an alternating cycle.
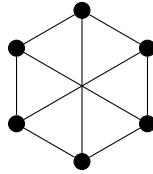
Figure 2.5: $K_{3,3}$, a Laman graph that is not a basis of $M(\mathcal{S}_2^n)$.

As we will see in the proof of the following theorem, $\mathcal{M}_2^{m \times n}$ can be realized as a coordinate projection of $\mathcal{S}_2^k$ for any $k \geq m + n$. This implies that the matroid $M(\mathcal{M}_2^{m \times n})$

can be realized as the restriction of such a $M(\mathcal{S}_2^k)$. We can use this fact to give a char-acterization of the algebraic matroid of $M(\mathcal{M}_2^{m \times n})$, whose bases are also the minimally $(2,2)$-rigid graphs in [38]. The ground set of $M(\mathcal{M}_2^{m \times n})$ can be associated with the edges in the complete bipartite graph with partite vertex sets $[m]$ and $[n]$, denoted $K_{m,n}$.

Theorem 2.3.4. *Let S be a collection of edges of $K_{m,n}$. Then S is independent in $M(\mathcal{M}_2^{m \times n})$ if and only if some acyclic orientation of $G(S)$ has no alternating closed trails.*

Proof. Define $k := m + n$. We claim that $\mathcal{M}_2^{m \times n}$ is a coordinate projection of $\mathcal{S}_2^k$. To see this, first let $A \in \mathcal{M}_2^{m \times n}$. Then choose $u_1, u_2 \in \mathbb{C}^{m \times 1}$ and $v_1, v_2 \in \mathbb{C}^{1 \times n}$ such that $A = u_1 v_1 - u_2 v_2$. Then define the following $k \times k$ matrix

$$B := \begin{pmatrix} u_1 \\ v_2^T \end{pmatrix} \begin{pmatrix} u_2^T & v_1 \end{pmatrix} - \begin{pmatrix} u_2 \\ v_1^T \end{pmatrix} \begin{pmatrix} u_1^T & v_2 \end{pmatrix}.$$

Note that B is skew-skew symmetric of rank at most 2, and its upper-right block is equal to A. Thus the claim is proven and therefore $M(\mathcal{M}_2^{m \times n})$ is a restriction of $M(\mathcal{S}_2^k)$ The result then follows from Theorem 2.3.2. $\qquad \square$

Remark 2.3.5. The content of Remarks 2.2.4 and 2.3.3 can be adapted to give a combi-natorial polynomial-time verifiable certificate that a given $S \subseteq [m] \times [n]$ is independent in $M(\mathcal{M}_2^{m \times n})$. Namely, assuming $m > n$, we can translate S into a subset T of $\binom{[m+n]}{2}$ by mapping each $(i, j) \in S$ to the pair $\{i + m, j\}$. Then the construction given in Remark 2.2.4 can be applied to T.

Chapter 3

Phylogenetics and linear spaces

The material in this chapter comes from joint work with Colby Long. It was published in *SIAM Journal of Discrete Mathematics* [8].

One approach to phylogenetic reconstruction is to use *distance-based* methods. Given a distance matrix consisting of the pairwise distances between n species, a distance-based method returns a tree metric or equidistant tree metric (ultrametric) that best fits the data. Typically, the distance matrix is constructed from biological data. It has been shown that both the set of equidistant tree metrics and the set of tree metrics have close connections to tropical geometry [4, 5, 56]. Because addition in the tropical semiring is defined as taking the maximum of two elements, the l^∞-metric offers an appealing choice as a measure of best fit for phylogenetic reconstruction.

Computing a closest tree metric to a given distance matrix using the l^∞-metric is NP-hard [3]. However, there exists a polynomial-time algorithm for computing an l^∞-closest *equidistant* tree metric [21]. Although the algorithm gives us a way to compute a closest equidistant tree metric to an arbitrary point in $\mathbb{R}^{\binom{n}{2}}$ quickly, the set of closest equidistant tree metrics is not in general a singleton. Indeed, it may be of high dimension or contain points corresponding to trees with entirely different topologies. Thus, for phylogenetic reconstruction, there may be several different trees that explain the data equally well from the perspective of the l^∞-metric. Recent work has studied the properties of equidistant tree space with the l^∞-metric [4, 43, 44] but to our knowledge the dimensions and topologies of the sets of l^∞-closest equidistant tree metrics have not been examined. Similarly, one might ask all of the same questions for tree metrics. Thus, we are motivated by the following problem.

Problem 3.0.1. Given a dissimilarity map $x \in \mathbb{R}^{\binom{n}{2}}$, describe the set of (equidistant) tree metrics that are closest to x in the l^∞-metric.

For both equidistant tree metrics and tree metrics, we obtain results concerning the dimensions of these sets as well as the tree topologies involved. Since the set of tree metrics and the set of equidistant tree metrics on n species are both polyhedral complexes, we begin by addressing the following problem as a stepping stone. The results obtained may be of independent interest to those studying combinatorics or optimization.

Problem 3.0.2. Given a point $x \in \mathbb{R}^m$ and a linear space $L \subseteq \mathbb{R}^m$, describe the subset of L consisting of points that are closest to x in the l^∞-metric.

Just as with tree metrics, the l^∞-closest point in a linear space is not unique in general. We give a polyhedral decomposition of $\mathbb{R}^m$ based on the dimension of the set of points in L that are l^∞-closest to x. One particularly nice implication of this decomposition is Theorem 3.1.9 which says that the l^∞-closest point in a linear space $L \subseteq \mathbb{R}^n$ to a given $x \in \mathbb{R}^n$ is unique for all such x if and only if the matroid underlying L is uniform.

The set of (equidistant) tree metrics on a fixed set of species is a polyhedral fan. Each open cone in this fan is the set of (equidistant) tree metrics corresponding to a particular tree topology. For many dissimilarity maps, optimizing to the set of (equidistant) tree metrics will be equivalent to optimizing to the linear hull of one such maximal cone. The equations defining the linear hulls of these cones are highly structured and the corresponding matroids are not uniform. Therefore, Theorem 3.1.9 implies the existence of dissimilarity maps with a positive-dimensional set of l^∞-closest (equidistant) tree metrics. For example, we show that there is a full-dimensional set of dissimilarity maps in $\mathbb{R}^{\binom{n}{2}}$ for which the set of l^∞-closest equidistant tree metrics has dimension $n-2$. Our construction shows that we can often obtain many l^∞-closest equidistant tree metrics to a dissimilarity map by adjusting branch lengths in an equidistant tree representing one such metric. We will also see that there are dissimilarity maps for which the set of l^∞-closest equidistant tree metrics contains equidistant tree metrics representing different tree topologies. In the case of 4-leaf trees, we provide a decomposition of $\mathbb{R}^{\binom{4}{2}}$ according to the topologies represented.

We begin in Section 3.1 with our results on l^∞-optimization to linear spaces. In particular, we give a natural way to assign a combinatorial type to each $x \in \mathbb{R}^m$ with respect to some linear subspace $L \subseteq \mathbb{R}^m$. We show that this combinatorial type gives a

polyhedral decomposition of $\mathbb{R}^m$ based on the dimension of the set of l^∞-closest points in L, from which Theorem 3.1.9 follows. Section 3.2 applies the results and ideas for linear spaces to phylogenetics. We investigate questions that would be of practical interest for phylogenetic reconstruction such as the dimension and corresponding tree topologies in the set of closest ultrametrics. We conclude by exploring the l^∞-metric as a distance-based method for reconstructing tree metrics.

3.1 l^∞-optimization to Linear Spaces

Given a linear space $L \subseteq \mathbb{R}^m$, we demonstrate a way to associate a sign vector in $\{+,-,0\}^m$ to each $x \in \mathbb{R}^m$. The associated sign vectors are then precisely the elements of the oriented matroid associated to L. For each $x \in \mathbb{R}^m$, this vector will encode information about the dimension of the set of l^∞-closest points to x in L.

We now quickly recall from Section 1.3 the necessary background on oriented matroids associated to linear subspaces of $\mathbb{R}^n$. For any real number $r \in \mathbb{R}$, $\operatorname{sign}(r) \in \{+,-,0\}$ is the *sign of r*. For a linear functional $c \in (\mathbb{R}^m)^*$, $\operatorname{sign}(c) \in \{+,-,0\}^m$ is defined by $\operatorname{sign}(c)_i = \operatorname{sign}(c_i)$. Given a sign vector $\sigma \in \{+,-,0\}^m$, we define $|\sigma| := \#\{i : \sigma_i \neq 0\}$. For a linear space $L \subseteq \mathbb{R}^m$ the *oriented matroid associated to L*, denoted $\mathcal{O}_L$, is the set of all sign vectors s in $\{+,-,0\}^m$ such that $s = \operatorname{sign}(c)$ for some linear functional $c \in (\mathbb{R}^m)^*$ that vanishes on L. The elements of an oriented matroid $\mathcal{O}$ are the *signed vectors* of $\mathcal{O}$. Let $\prec^*$ be the partial order on $\{+,-,0\}$ given by $0 \prec^* +$ and $0 \prec^* -$ with $+$ and $-$ incomparable. Then $\prec$ is the partial order on $\{+,-,0\}^m$ that is the cartesian product of $\prec^*$ m times. The *signed circuits* of an oriented matroid $\mathcal{O}$ are the signed vectors of $\mathcal{O}$ that are minimal with respect to $\prec$.

An oriented matroid can also be derived from a *zonotope*, the image of a cube under an affine map. Let $C_\delta(x) \subseteq \mathbb{R}^m$ denote the cube of side length 2δ centered at x. That is,

$$C_\delta(x) = \{y \in \mathbb{R}^m : |y_i - x_i| \leq \delta, i = 1, \ldots, m\}.$$

To each face F of $C_\delta(x)$, associate a sign vector $\text{sign}(F) \in \{+, -, 0\}^m$ as follows

$$\text{sign}(F)_i = \begin{cases} + & \text{if } y_i = x_i + \delta \quad \text{for all } y \in F \\ - & \text{if } y_i = x_i - \delta \quad \text{for all } y \in F \\ 0 & \text{otherwise.} \end{cases}$$

Figure 3.1 gives an illustration of the sign vectors associated to a square.
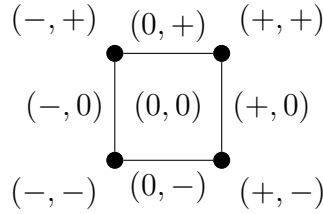
Figure 3.1: Sign vectors corresponding to faces of a square.

Let $V \in \mathbb{R}^{(m-d) \times m}$ be a matrix of full row rank and let $\pi : \mathbb{R}^m \to \mathbb{R}^{m-d}$ be the affine map given by $x \mapsto Vx$. For fixed $x \in \mathbb{R}^m$ and $\delta > 0$, $\pi(C_\delta(x)) \subset \mathbb{R}^{m-d}$ is a polytope called the *zonotope associated to* V. The inverse image of each face of the zonotope is a face of $C_\delta(x)$. Thus, for each face G of $\pi(C_\delta(x))$, we define $\text{sign}(G) := \text{sign}(\pi^{-1}(G))$. Proposition 3.1.1 below tells us that the collection of all such sign vectors is an oriented matroid which only depends on the matrix V and so we denote it $\mathcal{O}_V$. Recall from Section 1.3 that the oriented matroid $\mathcal{O}_L$ underlying a linear subspace $L \subseteq \mathbb{R}^m$ has as its set of vectors the set of all sign vectors $\text{sign}(c)$ for functionals $c \in (\mathbb{R}^m)^*$ that vanish on L.

Proposition 3.1.1 ([65],Corollary 7.17)**.** *Let $V \in \mathbb{R}^{(m-d) \times m}$ be a matrix of full row rank. Then $\mathcal{O}_V = \mathcal{O}_{\ker V}$.*

Given a linear space $L \subseteq \mathbb{R}^m$, we demonstrate a way to associate a sign vector in $\{+, -, 0\}^m$ to each $x \in \mathbb{R}^m$. As we let x vary, the associated sign vectors that arise are then precisely the vectors of the oriented matroid associated to L. For each $x \in \mathbb{R}^m$, this vector will encode information about the dimension of the set of l^∞-closest points to x in L.

In the rest of this section, we will use the language of matroids to state our main

56

results for linear spaces. Before we begin, we establish some notation.

Definition 3.1.2. Let $S \subseteq \mathbb{R}^m$ be an arbitrary set and let $x, z \in \mathbb{R}^m$. We denote the l^∞-distance from x to z by $d(x, z)$, the l^∞-distance from x to S by $d(x, S)$ and the set of all points in S closest to x by $C(x, S)$. That is

$$d(x, z) := \sup_i |x_i - z_i| \qquad d(x, S) := \inf_{y \in S} d(x, y) \qquad C(x, S) := \{y \in S : d(x, y) = d(x, S)\}.$$

Note that $C(x, S) = C_{d(x,S)}(x) \cap S$. Furthermore, when S is a linear space, there exists a unique minimal face F of $C_{d(x,S)}(x)$ that contains $C(x, S)$. We use the sign vector $\text{sign}(F)$ to give each $x \in \mathbb{R}^m$ a combinatorial type as in the following definition.

Definition 3.1.3. Let L be a linear space and F the minimal face of $C_{d(x,L)}(x)$ containing $C(x, L)$. The *type of x with respect to L* is $\text{type}_L(x) := \text{sign}(F)$.

Example 3.1.4. Consider linear spaces $L_1 = \{(t, t) \in \mathbb{R}^2 : t \in \mathbb{R}\}$ and $L_2 = \{(t, 0) \in \mathbb{R}^2 : t \in \mathbb{R}\}$ and let $x = (-3, -1)$ and $y = (5, 3)$. Then $\text{type}_{L_1}(x) = (+, -)$, $\text{type}_{L_1}(y) = (-, +)$, $\text{type}_{L_2}(x) = (0, +)$, and $\text{type}_{L_2}(y) = (0, -)$. See Figure 3.2 for an illustration.

Figure 3.2: Types of x and y with respect to L_1 and L_2.

We will see that the sign vectors that can arise as the type of a point with respect to L are precisely the elements of the oriented matroid associated to L. To aid in the proof we introduce the following convention for generating a vector with a given sign signature.

Definition 3.1.5. For $\sigma \in \{+, -, 0\}^m$, $u(\sigma)$ is the vector in $\mathbb{R}^m$ with

$$u(\sigma)_i := \begin{cases} 1 & \text{if} \quad \sigma_i = + \\ -1 & \text{if} \quad \sigma_i = - \\ 0 & \text{if} \quad \sigma_i = 0 \end{cases}.$$

Lemma 3.1.6. *Let $L \subseteq \mathbb{R}^m$ be a linear space. Then the sign vectors that can arise as the type of a point with respect to L are precisely the elements of the oriented matroid associated to L. That is,*

$$\mathcal{O}_L = \{\text{type}_L(x) : x \in \mathbb{R}^m\}.$$

Proof. First, we will show that $\mathcal{O}_L \subseteq \{\text{type}_L(x) : x \in \mathbb{R}^m\}$. Let $\sigma \in \mathcal{O}_L$, we will show that the type of $-u(\sigma)$ with respect to L is equal to σ. Since $\sigma \in \mathcal{O}_L$, by the definition of $\mathcal{O}_L$, there must exist a linear functional $c \in (\mathbb{R}^m)^*$ that vanishes on L with $\text{sign}(c) = \sigma$. Now we claim that $d(-u(\sigma), L) = 1$. Since $\mathbf{0} \in L$ and $d(-u(\sigma), \mathbf{0}) = 1$, $d(-u(\sigma), L) \leq 1$. If $x \in \mathbb{R}^m$ such that $d(-u(\sigma), x) < 1$, then for each index i with $\sigma_i \neq 0$, $\text{sign}(x_i) = -\sigma_i$. Therefore, $cx < 0$ which implies $x \notin L$. Thus, $d(x, L) = 1$.

Next, we claim that $\text{type}_L(-u(\sigma)) = \sigma$. Observe that

$$F = \{y \in C_1(-u(\sigma)) : y_i = 0 \text{ whenever } \sigma_i \neq 0\}$$

is a face of $C_1(-u(\sigma))$ and that $\text{sign}(F) = \sigma$. Therefore, it will suffice to show that F is the minimal face of $C_1(-u(\sigma))$ containing $C(-u(\sigma), L)$. So let $x \in C(-u(\sigma), L)$. We have already shown that this implies that $x \in C_1(-u(\sigma))$. Moreover, for all i, either $x_i = 0$ or $\text{sign}(x_i) = \text{sign}(-u(\sigma)_i) = -\sigma_i$. It must be the case then that if $\sigma_i \neq 0$ then $x_i = 0$. Otherwise, $cx < 0$, which is impossible, since c vanishes on L. Therefore, $x \in F$, and so F contains $C(-u(\sigma), L)$. Finally, all that remains to show is that F is the *minimal* face of $C_1(-u(\sigma))$ containing $C(-u(\sigma), L)$. If not, then there must exist j with $\sigma_j = 0$ such that $C(-u(\sigma), L)$ is contained in a facet of $C_1(-u(\sigma))$ of the form $\{y \in C_1(-u(\sigma)) : y_j = 1\}$ or $\{y \in C_1(-u(\sigma)) : y_j = -1\}$. But this is impossible, since we have already shown that $\mathbf{0} \in C(-u(\sigma), L)$. Hence, $\text{type}_L(-u(\sigma)) = \sigma$.

We now show $\{\text{type}_L(x) : x \in \mathbb{R}^m\} \subseteq \mathcal{O}_L$. Let $x \in \mathbb{R}^m$. We will show that $\text{type}_L(x) \in \mathcal{O}_L$. Assume L has dimension d and let $V \in \mathbb{R}^{(m-d) \times m}$ be a matrix whose rows form a basis for $L^\perp$. Let $\pi : \mathbb{R}^m \to \mathbb{R}^{m-d}$ be the map $x \mapsto Vx$. Let F be the minimal face of $C_{d(x,L)}(x)$ that contains $C(x, L)$ so that $\text{type}_L(x) = \text{sign}(F)$. The image of $C_{d(x,L)}(x)$ under π is the

zonotope associated to V in $\mathbb{R}^{m-d}$. Our goal will be to show that F is the inverse image of one of the faces of this zonotope. By the definition of the oriented matroid associated to a zonotope, this implies that $\operatorname{sign}(F)$ is an element of $\mathcal{O}_V$. Proposition 3.1.1 shows that the oriented matroids $\mathcal{O}_V$ and $\mathcal{O}_{\ker V}$ are equal. Since V is specifically constructed so that $\ker V = L$, this will also imply that $\operatorname{type}_L(x) = \operatorname{sign}(F) \in \mathcal{O}_L$.

By the hyperplane separation theorem, there exists a hyperplane separating L and the interior of $C_{d(x,L)}(x)$. Observe that any such hyperplane must contain L and intersect F in its interior. Therefore, we may choose $c \in (\mathbb{R}^m)^*$ in the row-span of V such that the hyperplane $\mathcal{H}_c = \{y \in \mathbb{R}^m : cy = 0\}$ is a face-defining hyperplane for F and write $c = bV$ for some $b \in (\mathbb{R}^{m-d})^*$.

Since $\mathcal{H}_c$ is a face defining hyperplane of $C_{d(x,L)}(x)$, $\mathcal{H}_b = \{z \in \mathbb{R}^{m-d} : bz = 0\}$ must be a face-defining hyperplane of $\pi(C_{d(x,L)}(x))$. Clearly, $\pi(F)$ is contained in the face of $\pi(C_{d(x,L)}(x))$ defined by $\mathcal{H}_b$. In fact, we have equality. If $z \in \pi(C_{d(x,L)}(x))$ then $z = Vy$ for some $y \in C_{d(x,L)}(x)$. So if $bz = 0$, then $bz = (bV)y = cy = 0$. This implies that $y \in F$ and so $z \in \pi(F)$. Similarly, if $\pi(y) \in \pi(F)$ then $0 = b(\pi(y)) = b(Vy) = (bV)(y) = cy$, which implies $y \in F$. Thus, we have just shown that $\pi(F)$ is a face of the zonotope and that $\pi^{-1}(\pi(F)) = F$. Therefore, $\pi(F)$ inherits its sign from F, and so $\operatorname{sign}(F) \in \mathcal{O}_V$. As noted, this implies that $\operatorname{type}_L(x) = \operatorname{sign}(F) \in \mathcal{O}_L$. □

As we show in the following theorem, the dimension of $C(x, L)$ depends entirely on $\operatorname{type}_L(x)$. Recall from Section 1.3 that for any $\sigma \in \mathcal{O}_L$, the *rank of σ in $\mathcal{O}_L$*, denoted $\operatorname{rank}(\sigma)$, is the rank of the support of σ in the matroid underlying $\mathcal{O}_L$. So $\operatorname{rank}(\sigma)$ is the smallest number k such that there exists indices $i_1, \ldots, i_k$ with $\sigma_{i_j} \neq 0$ such that for all $y \in L$, if $y_{i_1} = \cdots = y_{i_k} = 0$, then $y_j = 0$ for $\sigma_j \neq 0$.

Theorem 3.1.7. *Let $L \subset \mathbb{R}^m$ be a linear space of dimension d and let $\sigma \in \mathcal{O}_L$ be a sign vector in the oriented matroid associated to L. If $x \in \mathbb{R}^m$ has $\operatorname{type}_L(x) = \sigma$, then the collection of l^∞-closest points to x in L has dimension $d - \operatorname{rank}(\sigma)$.*

Proof. Let $L(\sigma)$ denote the linear space obtained by intersecting L and the $|\sigma|$ hyperplanes $\{x \in \mathbb{R}^m : x_i = 0\}$ for $\sigma_i \neq 0$. We claim that if $x \in \mathbb{R}^m$ with $\operatorname{type}_L(x) = \sigma$, then $\dim C(x, L) = \dim L(\sigma)$.

Suppose $\operatorname{type}_L(x) = \sigma$, and let F be the minimal face of $C_{d(x,L)}(x)$ containing $C(x, L)$. Let y be a point in $C(x, L)$ that is also in the interior of F. Then given any point $z \in L(\sigma)$, it is possible to choose ε so that $y + \varepsilon z \in F$ and hence in $C(x, L)$. Therefore,

$\dim C(x, L) \geq \dim L(\sigma)$. Moreover, any two points in $C(x, L)$ are contained in $F \cap L$ and so differ only by an element of $L(\sigma)$. Therefore, $\dim C(x, L) \leq \dim L(\sigma)$.

We now show that $\dim(L(\sigma)) = d - \operatorname{rank}(\sigma)$. Let $k := \operatorname{rank}(\sigma)$ and let $i_1, \ldots, i_k$ be indices such that for all $y \in L$, $y_{i_1} = \cdots = y_{i_k} = 0$ implies $y_j = 0$ when $\sigma_j \neq 0$. So $L(\sigma)$ can be expressed as the intersection of L with the hyperplanes $\{x \in \mathbb{R}^m : x_{i_j} = 0\}$, and by minimality of rank, this is not true of any subset of these hyperplanes. So $\dim(L(\sigma)) = d - \operatorname{rank}(\sigma)$. $\qquad \square$

Example 3.1.8. Let $L := \{(t, t, 0) \in \mathbb{R}^3 : t \in \mathbb{R}\}$. Consider the points $x = (0, 0, -1)$ and $y = (6, 4, 0)$. Then $\operatorname{type}_L(x) = (0, 0, +)$ and $\operatorname{type}_L(y) = (-, +, 0)$. Since $\operatorname{rank}(0, 0, -) = 0$ and $d = 1$, Theorem 3.1.7 tells us that $\dim(C(x, L)) = 1$. Since $\operatorname{rank}(+, -, 0) = 1$, Theorem 3.1.7 tells us that $\dim(C(y, L)) = 0$. Figure 3.3 shows x and y each surrounded by a cube of side length 2 (colored red and light blue, respectively). The intersections with L are $C(x, L)$ and $C(y, L)$.
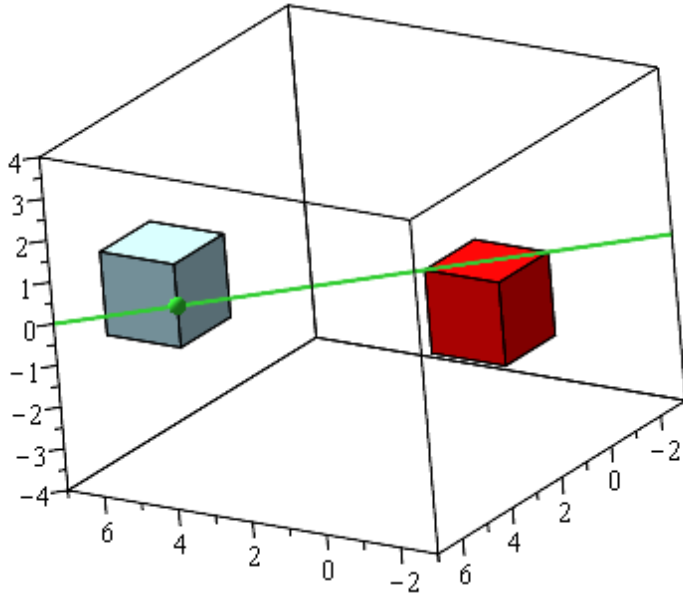
Figure 3.3: L and cubes around $(0, 0, -1)$ and $(6, 4, 0)$.

We can use the structure of the matroid $\mathcal{M}_L$ to glean information about possible values of $\dim(C(x, L))$. Let $U_{d,m}$ denote the uniform matroid of rank d on ground set

$\{1, \ldots, m\}$; that is, the circuits of $U_{d,m}$ are all $d + 1$-element subsets of $\{1, \ldots, m\}$.

Theorem 3.1.9. *Let $L \subseteq \mathbb{R}^m$ be a linear space. Then the l^∞-closest point to x in L is unique for all $x \in \mathbb{R}^m$ if and only if the matroid underlying L is uniform.*

Proof. Let d be the dimension of L. If $\mathcal{M}_L$ is *not* uniform, then $\mathcal{O}_L$ has a circuit σ with $|\sigma| \leq d$, so $\operatorname{rank}(\sigma) \leq d - 1$. Then Lemma 3.1.6 and Theorem 3.1.7 imply that there exists a point $x \in \mathbb{R}^m$ such that $\dim C(x, L) = d - \operatorname{rank}(\sigma) \geq 1$.

If $\mathcal{M}_L = U_{d,m}$ then $\operatorname{rank}(\sigma) = d$ for all $\sigma \in \mathcal{O}_L$. Theorem 3.1.7 implies that $\dim(C(x, L)) = 0$ for all $x \in \mathbb{R}^m$. $\qquad\square$

Lemma 3.1.6 enables us to give a partition of $\mathbb{R}^m$ by type with respect to L.

Proposition 3.1.10. *Let $L \subseteq \mathbb{R}^m$ be a linear space and let $\sigma \in \mathcal{O}_L$ be a sign vector in the oriented matroid associated to L. The set of all points in $\mathbb{R}^m$ with type σ with respect to L is the Minkowski sum of L and the interior of the conical hull of $\{-u(\tau) : \sigma \preceq \tau\}$. That is,*

$$\{x \in \mathbb{R}^m : \operatorname{type}_L(x) = \sigma\} = L + \operatorname{int}(\operatorname{cone}(\{-u(\tau) : \sigma \preceq \tau\})).$$

Proof. Let $\sigma \in \mathcal{O}_L$. Define $\mathcal{V}_\sigma := \{-u(\tau) : \sigma \preceq \tau\}$. First, we will show that everything in $L + \operatorname{int}(\operatorname{cone}(\mathcal{V}_\sigma))$ has type σ. Since adding an element of L to a point does not change its type with respect to L, it will suffice to show everything in $\operatorname{int}(\operatorname{cone}(\mathcal{V}_\sigma))$ has type σ.

Let $x \in \operatorname{int}(\operatorname{cone}(\mathcal{V}_\sigma))$. Then there exists $\alpha > 0$ such that if $\sigma_i = +$ or $\sigma_i = -$, then $|x_i| = \alpha$ and $|x_i| < \alpha$ otherwise. By Lemma 3.1.6, there exists $c \in (\mathbb{R}^m)^*$ such that $\operatorname{sign}(c) = \sigma$ and $cy = 0$ for all $y \in L$. Let $\mathcal{H}_c := \{y \in \mathbb{R}^m : cy = 0\}$ be the hyperplane defined by c. It is clear that $d(x, \mathcal{H}_c) = \alpha$, and any $y \in C(x, \mathcal{H}_c)$ must have $y_i = 0$ if $|x_i| = \alpha$. Since $L \subseteq \mathcal{H}_c$, the same is true for each $y \in C(x, L)$. Therefore, if $\sigma_i = +$ or $\sigma_i = -$, then $\sigma_i = \operatorname{type}_L(x)_i$.

Sine $d(x, \mathcal{H}_c) = \alpha$ and $L \subseteq \mathcal{H}_c$, $d(x, L) \geq \alpha$. Since $d(x, \mathbf{0}) = \alpha$ and $\mathbf{0} \in L$, this implies $d(x, L) = \alpha$ and $\mathbf{0} \in C(x, L)$. If $\sigma_i = 0$ then $|x_i - 0| = |x_i| < \alpha$, which implies that if $\sigma_i = 0$, $\operatorname{type}_L(x)_i = 0$. Thus, $\operatorname{type}_L(x) = \sigma$.

To see that everything of type σ is contained in $L + \operatorname{int}(\operatorname{cone}(\mathcal{V}_\sigma))$, let x be such that $\operatorname{type}_L(x) = \sigma$. By definition of type, this means that σ is the unique sign vector such that if F is the face of the unit cube $C_1(\mathbf{0})$ with type σ, then there exists some $y \in \operatorname{int}(F)$ such that $x + \lambda y \in L$ for some $\lambda > 0$. So there exists some $l \in L$ such that $x = l + \lambda(-y)$ thus showing that $x \in L + \operatorname{int}(\operatorname{cone}(\mathcal{V}_\sigma))$. $\qquad\square$

Modulo the lineality space L, the closures of the cones in Proposition 3.1.10 form the face fan of the zonotope obtained by projecting the cube $C_{d(x,L)}(x)$ onto $L^\perp$.

Corollary 3.1.11. *Let $x \in \mathbb{R}^m$ and let V be a matrix whose rows span $L^\perp$. Then $\mathrm{type}_L(x)$ is equal to the sign of the unique face F of $Z(V)$ such that the conic hull of the interior of F contains Vx.*

The signs of the facets of $Z(V)$ correspond to circuits of $\mathcal{O}_L$ [65, Corollary 7.17]. This implies that the full dimensional cones of the partition correspond to circuits. Hence $\mathrm{type}_L(x)$ is generically a circuit of $\mathcal{O}_L$.

3.2 Applications to Phylogenetics

In this section, we will consider how the results from Section 3.1 can be applied to phylogenetic reconstruction using the l^∞-metric. We will address Problem 3.0.1, concerning the structure of the set of l^∞-closest points to the set of equidistant tree metrics. In particular, we show that there can be many (equidistant) tree metrics that are equally close to a given dissimilarity map, and they can represent many different tree topologies. We also decompose the space of dissimilarity maps on 3 elements and on 4 elements according to the tree topologies represented in the set of l^∞-closest equidistant tree metrics. Finally, we investigate optimizing to the set of tree metrics and show how many of the results for equidistant tree metrics carry over.

3.2.1 Rooted trees and ultrametrics

We begin by giving the necessary definitions related to ultrametrics. As we will see, there is a relationship between ultrametrics and rooted leaf-labeled trees that is similar to the relationship between tree metrics and non-rooted leaf-labeled trees, discussed in Sections 1.4 and 2.2. Let $RP(n)$ be the set of all n-leaf rooted trees with leaves labeled by $[n] := \{1, \ldots, n\}$. Following the convention of [52, Section 2.2], we call the elements of $RP(n)$ *rooted phylogenetic $[n]$-trees*. We will also consider the set of *rooted binary phylogenetic $[n]$-trees* which we will denote $RB(n)$. A *polytomy* of a non-binary tree is a vertex with degree greater than three - that is, a witness to the property of being non-binary. To represent the topology of $\mathcal{T} \in RP(n)$ we use the notation $(S_1(S_2))$ to indicate that the leaves labeled by the set S_1 and S_2 are on opposite sides of the root in $\mathcal{T}$. We

apply this notation recursively to give the topology of the rooted subtree in $\mathcal{T}$ induced by the labels in S_1 and S_2. Thus, for example, we can express the topology of the the rooted tree in Figure 3.4 by $(D(C(AB)))$.

Let $\mathcal{T} \in RP(n)$ and assign a positive weighting to the edges of $\mathcal{T}$. This naturally induces a metric δ on the leaves of $\mathcal{T}$ where $\delta(i, j)$ is the sum of the edge weights on the unique path between i and j. If we further assume that the distance from each leaf vertex to the root is the same then δ is an *ultrametric*.

Definition 3.2.1. [52, Definition 7.2.1] A dissimilarity map $\delta : X \times X \to \mathbb{R}$ is called an *ultrametric on X* if for every three distinct elements $i, j, k \in X$,

$$\delta(i, j) \leq \max\{\delta(i, k), \delta(j, k)\}.$$

An equidistant edge weighting of a rooted tree is a weighting of the edges where the distance from each leaf to the root is the same and where the weight of every internal edge is positive. Note that this allows the possibility of nonpositive weights on leaf edges. Given any ultrametric u on $[n]$, there exists a unique $\mathcal{T} \in RP(n)$ and an equidistant weighting w such that the ultrametric induced by $(\mathcal{T} : w)$ is equal to u [52, Theorem 7.2.8]. We call $(\mathcal{T} : w)$ an *equidistant representation of u* and say that $\mathcal{T}(u) := \mathcal{T}$ is the *topology of u*.

We can also convert an equidistant edge weighting of a tree into a *vertex weighting* of that same tree [52, Theorem 7.2.8]. Given any internal vertex v in an equidistant representation of an ultrametric u, $u(i, j)$ is constant over all pairs of leaves i, j having v as their most recent common ancestor. We obtain a vertex weighting from an edge weighting by labeling each internal vertex by this constant value. For what follows, we will represent dissimilarity maps on n elements as points in $\mathbb{R}^{\binom{n}{2}}$ by letting $\delta_{ij} = \delta(i, j)$ and use $U_n \subseteq \mathbb{R}^{\binom{n}{2}}$ to denote the set of all ultrametrics on n elements. Many of our examples will involve dissimilarity maps on 4 elements, in which case we let $(x_{AB}, x_{AC}, x_{AD}, x_{BC}, x_{BD}, x_{CD})$ be the coordinates of an arbitrary point in $\mathbb{R}^{\binom{4}{2}}$. We will also use the notation e_{ij} to denote the standard basis vector with $x_{ij} = 1$ and all other entries equal to zero.

Example 3.2.2. Consider the ultrametric $u = (5, 7, 9, 7, 9, 9) \in \mathbb{R}^6$. Figure 3.4 shows two equivalent ways of representing u: with a vertex weighting on the left and an equidistant

edge weighting on the right.
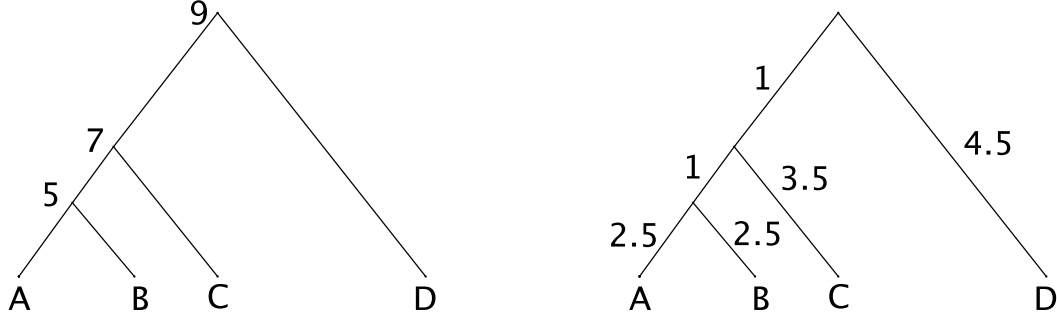
Figure 3.4: Two different representations of $u = (5, 7, 9, 7, 9, 9)$.

3.2.2 l^∞-optimization to the set of Ultrametrics

Given a dissimilarity map δ, we let δ_U be the unique coordinate-wise maximum ultra-metric which is coordinate-wise less than δ. This is called the *subdominant ultrametric of δ*. For a proof of the existence and uniqueness of the subdominant ultrametric, and a polynomial-time algorithm for computing it, see [52, Chapter 7].

Our interest in the subdominant ultrametric is that it gives us a way to determine an l^∞-closest ultrametric to a dissimilarity map δ [21]. We first compute the subdominant ultrametric δ_U and then use that

$$d(\delta, U_n) = \frac{1}{2} d(\delta, \delta_U).$$

We define δ_c, the *canonical closest ultrametric to δ*, by $\delta_c(i, j) = \delta_U(i, j) + \frac{1}{2} d(\delta, \delta_U)$.

As noted in the introduction, the set of l^∞-closest ultrametrics is in general not a single point. Moreover, in many instances, the set of l^∞-closest ultrametrics to δ, $C(\delta, U_n)$, will contain ultrametrics representing different topologies. Thus, there may be several different trees that explain the data equally well from the perspective of the l^∞-metric.

Example 3.2.3. Figure 3.5 depicts three ultrametrics in the set of closest ultrametrics to $\delta = (2, 4, 6, 8, 10, 12)$. The subdominant ultrametric is $\delta_U = (2, 4, 6, 4, 6, 6)$, $d(\delta, U_4) = 3$, and the canonical closest ultrametric (pictured far left) is $\delta_c = (5, 7, 9, 7, 9, 9)$.
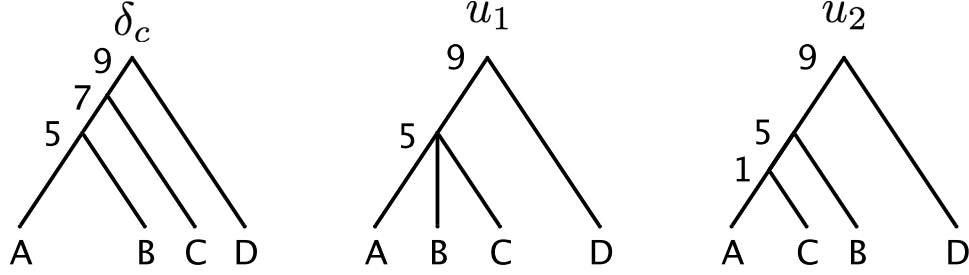
Figure 3.5: Three ultrametrics in $C(\delta, U_4)$ for $\delta = (2, 4, 6, 8, 10, 12)$.

One can easily verify that the canonical closest ultrametric inherits dominance from the subdominant ultrametric. That is, for any $\delta \in \mathbb{R}^{\binom{n}{2}}$, every ultrametric in $C(\delta, U_n)$ is coordinate-wise less than δ_c. Thus, we can construct closest ultrametrics by "sliding down" vertices of δ_c so long as the l^∞-distance between δ and the new ultrametric does not exceed $d(\delta, \delta_c)$. In Figure 3.5, we obtain u_1 from δ_c by sliding the middle internal vertex until it reaches the lowest one. We obtain u_2 by continuing to slide this vertex until we can do so no more. Observe that in this case, the root vertex must remain fixed.

Example 3.2.3 shows that it is possible for the set of l^∞-closest ultrametrics to a point to contain different topologies. Below, we consider what sets of topologies are represented in $C(\delta, U_n)$ for an arbitrary point $\delta \in \mathbb{R}^{\binom{n}{2}}$. The idea behind most of these proofs is to find a linear space that contains the ultrametrics for many different tree topologies and apply the constructions for linear spaces developed in Section 3.1.

Definition 3.2.4. Let $\delta \in \mathbb{R}^{\binom{n}{2}}$ and $u \in U_n$. Define

$$Top(\delta) := \{\mathcal{T}(u) : \ u \in C(\delta, U_n)\}.$$

In [23], the authors study the geometry of the set of dissimilarity maps around a polytomy with respect to the Euclidean norm. They showed that locally this space could be partitioned according to the closest tree topology. Proposition 3.2.5 shows the contrast between that situation and using the l^∞-metric.

Proposition 3.2.5. *Let $\mathcal{T}$ be a rooted phylogenetic $[n]$-tree with a polytomy. Assume that $\mathcal{T}$ is not the star tree. Then there exists $\delta \in \mathbb{R}^{\binom{n}{2}}$ such that $Top(\delta)$ contains $\mathcal{T}$ and all of its resolutions.*

65

Proof. Let u be an ultrametric with $\mathcal{T}(u) = \mathcal{T}$. Since $\mathcal{T}$ is not the star tree, there exist three leaves $\{i, j, k\}$ such that $u_{ij} < u_{ik} = u_{jk}$. Define $\delta := u + \varepsilon(e_{ik} - e_{jk})$ for some $0 < \varepsilon < u_{ik} - u_{ij}$. Then u is in $C(\delta, U_n)$ and so are all possible resolutions of the polytomy. $\qquad\square$

Example 3.2.6. Let $(x_{AB}, x_{AC}, x_{AD}, x_{BC}, x_{BD}, x_{CD})$ be the coordinates of a point in $\mathbb{R}^{\binom{4}{2}}$ and consider $u = (5, 5, 10, 5, 10, 10) \in \mathbb{R}^{\binom{4}{2}}$. The topology of the ultrametric u is the rooted tree $(D(ABC))$ with an unresolved tritomy.

Note that $u_{BC} < u_{BD} = u_{CD}$. Choose $\varepsilon = 1$ and let

$$\delta = u + \varepsilon(e_{CD} - e_{BD}) = (5, 5, 10, 5, 9, 11).$$

The canonical closest ultrametric $\delta_c = (6, 6, 10, 6, 10, 10)$ also has an unresolved tritomy and $C(\delta, U_4)$ contains ultrametrics corresponding to each different resolution. For example, $(4, 6, 10, 6, 10, 10), (6, 4, 10, 6, 10, 10)$, and $(6, 6, 10, 4, 10, 10)$ are all elements of $C(\delta, U_4)$.

We obtain the following corollary by choosing a tree with a single resolved triple in the proof of Proposition 3.2.5.

Corollary 3.2.7. *There exist points in $\mathbb{R}^{\binom{n}{2}}$ for which $Top(\delta) \cap RB(n)$ contains $(2n - 3)!!/3$ different tree topologies.*

We will also see from our decomposition of $\mathbb{R}^{\binom{4}{2}}$ that there are actually 6-dimensional polyhedral cones in which every point in the interior has five l^∞-closest binary tree topologies.

Even when all l^∞-closest ultrametrics to some given $\delta \in \mathbb{R}^{\binom{n}{2}}$ have the same topology, the dimension of the set of l^∞-closest ultrametrics can be high. The affine hull of each maximal cone of U_n is a linear space defined by relations of the form $x_{ik} - x_{jk} = 0$ where $(k(ij))$ is a triple compatible with the corresponding tree. As before, we can find points where optimizing to U_n is equivalent to optimizing to such a linear space and so our results from Section 3.1 can be applied.

Proposition 3.2.8. *Let $\mathcal{T} \in RB(n)$. There exists $\delta \in \mathbb{R}^{\binom{n}{2}}$ such that $\dim(C(\delta, U_n)) = n - 2$ and every ultrametric in $C(\delta, U_n)$ has topology $\mathcal{T}$.*

Proof. Let u be an ultrametric and $(k(ij))$ a triple compatible with $\mathcal{T}(u)$. For $\varepsilon > 0$, let $\delta = u + \varepsilon(e_{ik} - e_{jk})$. If ε is sufficiently small, $C(\delta, U_n) = C(\delta, L)$ where L is the affine hull

of the maximal cone of U_n containing u. The type of x relative to L is the sign vector σ where $\sigma_{ij} = +$, $\sigma_{ik} = -$, and all other entries are zero. The rank of σ in $\mathcal{O}_L$ is one, and thus by Theorem 3.1.7, $\dim(C(\delta, U_n)) = (n-1) - 1 = n - 2$. $\square$

Example 3.2.9. Let $(x_{AB}, x_{AC}, x_{AD}, x_{BC}, x_{BD}, x_{CD})$ be the coordinates of a point in $\mathbb{R}^{\binom{4}{2}}$. Choose $u = (5, 7, 9, 7, 9, 9)$, the ultrametric corresponding to the tree at far left in Figure 3.4. We will perturb u to construct a dissimilarity map δ where the set of l^∞-closest points to δ has dimension two. The triple $(C(AB))$ is compatible with $\mathcal{T}(u)$ and so we let

$$\delta = (5, 7, 9, 7, 9, 9) + (e_{AC} - e_{BC}) = (5, 8, 9, 6, 9, 9).$$

The subdominant ultrametric $\delta_U = (5, 6, 9, 6, 9, 9)$ and the canonical ultrametric $\delta_c = (6, 7, 10, 7, 10, 10)$. We have two degrees of freedom that come from adjusting the values of $\{(\delta_c)_{AD}, (\delta_c)_{BD}, (\delta_c)_{CD}\}$ (sliding down the root) or $\{(\delta_c)_{AB}\}$ (sliding down the most recent common ancestor of A and B).

3.2.3 The Decomposition for 3-Leaf and 4-Leaf Trees

The following definition makes formal the idea of partitioning the points in $\mathbb{R}^{\binom{n}{2}}$ according to their sets of l^∞-closest trees.

Definition 3.2.10. Let $\{\mathcal{T}_1, \ldots, \mathcal{T}_k\} \subseteq RP(n)$. The *district* of $\{\mathcal{T}_1, \ldots, \mathcal{T}_k\}$ is the set

$$D(\{\mathcal{T}_1, \ldots, \mathcal{T}_k\}) := \{\delta \in \mathbb{R}^{\binom{n}{2}} : Top(\delta) = \{\mathcal{T}_1, \ldots, \mathcal{T}_k\}\}.$$

We can represent a dissimilarity map on three elements as a point $(x_{12}, x_{13}, x_{23}) \in \mathbb{R}^3$. There are three maximal cones of U_3 corresponding to the three elements of $RB(3)$. Modulo the common lineality space of each of these cones, $\text{span}\{(1, 1, 1)\}$, we can fix the first coordinate at zero and represent the space of dissimilarity maps on three elements in the plane.

Figure 3.6 depicts a polyhedral subdivision of $\mathbb{R}^3$ according to districts. There are seven cones in this subdivision. The labels $(1(23)), (3(12))$, and $(2(13))$ label the image of the set of ultrametrics for each topology. These labels also label the areas between the dashed lines which are the 2-dimensional images of the three 3-dimensional districts $D\{(1(23))\}$, $D\{(2(13))\}$, and $D\{(3(12))\}$. The dotted lines themselves are the images of the 2-dimensional cones whose interiors form the districts $D\{(123), (1(23)), (2(13))\}$,
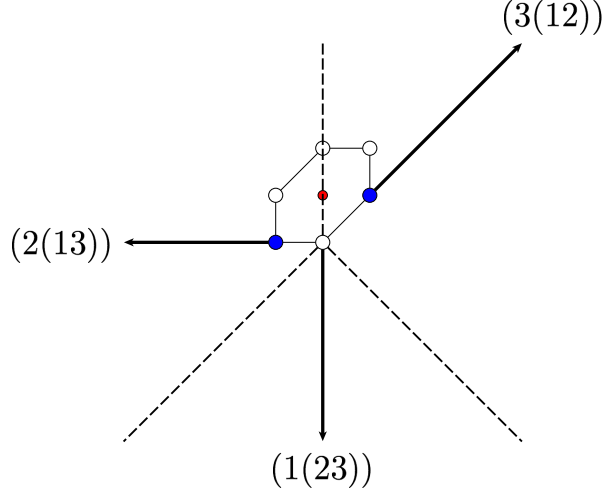
Figure 3.6: A 2-dimensional representation of the polyhedral subdivision of $\mathbb{R}^3$ according to district.

$D\{(123),(1(23)),(3(12))\}$, and $D\{(123),(2(13)),(3(12))\}$. The origin represents the image of $\text{span}\{(1,1,1)\}$ which is the district of the 3-leaf claw tree, $D\{(123)\}$.

Example 3.2.11. The image of the dissimilarity map $\delta = (1,1,3)$ after modding out by U_3's lineality space is pictured in Figure 3.6. Note that $d(\delta, U_3) = 1$. The hexagon surrounding it is the zonotope that is the image of the cube $C_1(\delta)$. The filled vertices of the zonotope correspond to the fully resolved l^∞-closest ultrametrics $(2,1,2)$ and $(1,2,2)$. The origin corresponds to the l^∞-closest ultrametric $(2,2,2)$ and $\delta \in D\{(123),(2(13)),(3(12))\}$.

The decomposition for 4-leaf trees is much more complicated. The supplemental materials, located at [9] contain a Maple [1] file for computing a polyhedral subdivision of $\mathbb{R}^{\binom{4}{2}}$ into a fan consisting of 723 maximal polyhedral cones labeling 37 different districts. Each maximal cone is labeled by a set $\{\mathcal{T}_1, \ldots, \mathcal{T}_k\} \subseteq RP(n)$, meaning that each point in the interior of the cone is in $D(\{\mathcal{T}_1, \ldots, \mathcal{T}_k\})$.

The computations rely heavily on the functionality of the package `PolyhedralSets` (available in Maple2015 and later versions). The fan is computed by first considering each of the fifteen different trees in $RB(4)$ individually. For each $\mathcal{T} \in RB(4)$, we construct a fan with support $\mathbb{R}^{\binom{4}{2}}$, where all of the points in the interior of each maximal cone in the fan satisfy either $\mathcal{T} \in Top(\delta)$ or $\mathcal{T} \notin Top(\delta)$. The resulting polyhedral subdivision is the common refinement of these fifteen fans. Note that there are far more than 37 districts

since our construction only labels the 6-dimensional districts.

Based on the 3-leaf case, one might hope that districts are easily described or possess some nice properties. For example, the 3-leaf districts are all convex and tropically convex. However, the 4-leaf case shows that many of these properties do not hold in general. For the rest of this section, we let $(x_{12}, x_{13}, x_{14}, x_{23}, x_{24}, x_{34})$ denote an arbitrary point in $\mathbb{R}^{\binom{4}{2}}$. We include some results for those familiar with tropical geometry and its connections to phylogenetics without the requisite background that would take us too far afield.

Proposition 3.2.12. *Districts are not necessarily convex nor tropically convex.*

Proof. We offer the following counterexample in $\mathbb{R}^{\binom{4}{2}}$. Let $\delta^1 = (10, 20, 21, 23, 25, 27)$ and $\delta^2 = (10, 23, 21, 20, 25, 27)$. Not only is $Top(\delta^1) = Top(\delta^2) = \{(4(3(12)))\}$, but in fact $\delta^1_U = \delta^2_U = (10, 20, 21, 20, 21, 21)$. The point $\delta^3 = \frac{1}{2}\delta^1 + \frac{1}{2}\delta^2$ lies on the line between these two points but $Top(\delta^3) = \{(3(4(12)))\}$.

Similarly, using the operations of the max-plus algebra, the point $\delta^4 = (0 \odot \delta^1) \oplus (-\frac{3}{2} \odot \delta^2)$ lies on the tropical line between these two points but $Top(\delta^4) = \{(3(4(12)))\}$. $\square$

Let δ and δ' be dissimilarity maps in $\mathbb{R}^{\binom{n}{2}}$. From the algorithm for computing the subdominant ultrametric, it is clear that δ and δ' will have the same subdominant ultrametric topology if they have the same relative ordering of coordinates - that is, $\delta_{ij} \leq \delta_{kl}$ if and only if $\delta'_{ij} \leq \delta'_{kl}$ and $\delta_{ij} < \delta_{kl}$ if and only if $\delta'_{ij} < \delta'_{kl}$. For 3-leaf trees, relative ordering also completely determines district. The example below demonstrates that for trees with more than three leaves this is not the case.

Example 3.2.13. Let $\delta^1 = (4, 8, 12, 9, 21, 22)$ and $\delta^2 = (4, 8, 12, 9, 13, 14)$. Both dissimilarity maps satisfy $\delta^i_{12} < \delta^i_{13} < \delta^i_{23} < \delta^i_{14} < \delta^i_{24} < \delta^i_{34}$. However, $Top(\delta^1) = \{(4(3(12)))\}$ and $Top(\delta^2) = \{(4(3(12))), (4(2(13))), (4(1(23)))\}$.

It does not appear possible to simplify the given subdivision of $\mathbb{R}^{\binom{4}{2}}$ much further by combining cones. Consider for example the forty maximal cones that constitute the district $D(\{(4(3(12)))\})$. Any five element subset of these cones contains a pair whose convex hull has full dimensional intersection with the interior of a maximal cone from another district. Therefore, by combining these cones the best we could hope for is to represent this district as the union of ten maximal convex cones. While a few can be patched together the final description does not appear any simpler.

This polyhedral subdivision was constructed by examining each possible 4-leaf subdominant ultrametric topology and writing out inequalities to determine when we could

obtain a new topology. It is certainly possible, though likely much more difficult, to do the same thing for trees with any fixed number of leaves. It is unclear how to generalize our approach to an arbitrary number of leaves and so the following problem remains open.

Problem 3.2.14. Give a polyhedral decomposition of $\mathbb{R}^{\binom{n}{2}}$ according to districts.

3.2.4 Tree Metrics

We end with a note about l^∞-optimization to the set of tree metrics. A tree metric δ on $[n]$ is a metric induced by a positive edge weighting of an n-leaf tree (no longer rooted nor equidistant). Note that this is a slight modification from the definition of a tree metric given in Sections 1.4 and 2.2 in that here we require all edges weights to be positive (not just internal). The pair $(\mathcal{T} : w)$ that realizes this metric is called a *tree metric representation of δ*. A metric δ is a tree metric if and only if it satisfies the four-point condition [52, Theorem 7.2.6] .

Definition 3.2.15. [52, Definition 7.2.1] A dissimilarity map $\delta : X \times X \to \mathbb{R}$ satisfies the four-point condition if for every four (not necessarily distinct) elements $w, x, y, z \in X$,

$$\delta(w, x) + \delta(y, z) \leq \max\{\delta(w, y) + \delta(x, z), \delta(w, z) + \delta(x, y)\}.$$

We use the notation $\mathcal{T}_n \subset \mathbb{R}^{\binom{n}{2}}$ to denote the set of all tree metrics on $[n]$. If we insist that the points in Definition 3.2.15 are distinct, then the set of metrics satisfying the *distinct* 4-point condition is the tropical Grassmannian [56]. Thus, the problem of finding the closest tree metric is closely related to the problem of l^∞-optimization to this tropical variety.

Although there is no subdominant tree metric, we can still compute the l^∞-distance from an arbitrary point to the set of tree metrics. The set of binary phylogenetic trees with label set $[n]$ is $B(n)$. For each $\mathcal{T} \in B(n)$, the distance to the set of tree metrics with topology $\mathcal{T}$ can be found by solving a linear program. Taking the minimum of these $(2n - 5)!!$ individually computed distances gives us the distance to the set of tree metrics. Proposition 3.2 in [7] show that $C(\delta, U_n)$ is always a tropical polytope and thus connected. This is a nice property from the perspective of phylogenetic reconstruction, as it means we have a set of closest ultrametrics any of which can be obtained from another

by shrinking and growing branch lengths without ever leaving the set $C(\delta, U_n)$. The same does not hold for tree metrics.

Proposition 3.2.16. *There exists $\delta \in \mathbb{R}^{\binom{6}{2}}$ such that $C(\delta, \mathcal{T}_6)$ and $C(\delta, \mathcal{G}_{2,6})$ are not connected.*

Proof. Let
$$\delta = (35, 22, 32, 49, 42, 26, 34, 23, 32, 39, 41, 34, 46, 49, 32)$$
be the metric in $\mathbb{R}^{\binom{6}{2}}$ with coordinates $(\delta_{12}, \delta_{13}, \delta_{14}, \dots, \delta_{45}, \delta_{46}, \delta_{56})$. Then $d(\delta, \mathcal{T}_6) = d(\delta, \mathcal{G}_{2,6}) = 5$. The set $C(\delta, \mathcal{T}_6)$ is the union of two disjoint polyhedra. One is four-dimensional and corresponds to the 6-leaf tree with nontrivial splits $13|2456, 134|256$ and $25|1346$ and the other is six-dimensional and corresponds to the 6-leaf tree with nontrivial splits $14|2356, 134|256$ and $56|1234$. In this instance, $C(\delta, \mathcal{T}_6) = C(\delta, \mathcal{G}_{2,6})$. $\square$

Unfortunately, many of the less than desirable properties exhibited in the ultrametric case hold for tree metrics. Simple modifications to the constructions for ultrametics give analogous results for tree metrics and unrooted trees to the results in Propositions 3.2.5 and 3.2.8, and Corollary 3.2.7. We conclude with one such example about the possible dimension of the set of l^∞-closest tree metrics to a point.

Proposition 3.2.17. *Let $\mathcal{T} \in B(n)$. There exists $\delta \in \mathbb{R}^{\binom{n}{2}}$ such that $\dim(C(\delta, \mathcal{T}_n)) = 2n - 6$ and every tree metric in $C(\delta, \mathcal{T}_n)$ has $\mathcal{T}$ as a tree metric representation.*

Proof. Let z be a tree metric with tree metric representation $\mathcal{T}$. Let L be the affine hull of the cone of tree metrics corresponding to $\mathcal{T}$. This is the linear space of dimension $2n - 3$ defined by all of the equalities of the form $x_{ik} + x_{jl} - x_{il} - x_{kj} = 0$ where $ij|kl$ is an induced quartet of $\mathcal{T}$.

Choose i, j, k and l so that $ij|kl$ is a quartet of $\mathcal{T}$. For $\varepsilon > 0$, let $\delta = z + \varepsilon(e_{ik} + e_{jl} - e_{il} - e_{kj})$. Since $\mathcal{T}$ is binary, if ε is sufficiently small, $C(\delta, \mathcal{T}_n) = C(\delta, L)$. The type of z relative to L is the signed vector σ where $\sigma_{il} = +, \sigma_{kj} = +, \sigma_{ik} = -, \sigma_{jl} = -$, and all other entries are zero. The rank of σ in $\mathcal{O}_L$ is three, and thus by Theorem 3.1.7, $\dim(C(\delta, U_n)) = (2n - 3) - 3 = 2n - 6$. $\square$

71

Chapter 4

Unimodular Hierarchical Models

Part of the material in this chapter comes from joint work with Seth Sullivant that was published in *Journal of Combinatorial Theory, Series B* [12]. The remaining material comes from joint work with Chris O'Neill that was published in *Journal of Algebraic Statistics* [10].

This chapter studies a particular subclass of the discrete log-linear models called *hierarchical models*. Such models are naturally indexed by pairs $(\mathcal{C}, \mathbf{d})$ where $\mathcal{C}$ is a simplicial complex on ground set $\{1, \ldots, n\}$ and $\mathbf{d} = (d_1, \ldots, d_n)$ is an integer vector satisfying $d_i \geq 2$ for each i. The main results of this chapter are Theorem 4.9.1 and Remark 4.10.4. Theorem 4.9.1 gives a complete classification of the pairs $(\mathcal{C}, \mathbf{d})$ such that the corresponding hierarchical model has a unimodular toric ideal. Remark 4.10.4 gives a combinatorial classification of all the elements in the Graver basis of any unimodular toric ideal coming from a hierarchical model.

4.1 Preliminaries

Definition 4.1.1. A *(abstract) simplicial complex* $\mathcal{C}$ is a pair $(V, \mathcal{F})$ where V is a finite set and $\mathcal{F}$ is a set of subsets of V such that if $G \subset F$ and $F \in \mathcal{F}$, then $G \in \mathcal{F}$.

Given a simplicial complex $(V, \mathcal{F})$, the set V is called the *ground set* of $\mathcal{C}$ and each $F \in \mathcal{F}$ is called a *face* of $\mathcal{C}$. Inclusion-wise maximal faces of $\mathcal{C}$ are called *facets*. We will use the notation ground($\mathcal{C}$) and facet($\mathcal{C}$) to denote the ground set and facets of a simplicial complex $\mathcal{C}$, respectively.

Let $X_1, \ldots, X_n$ be discrete random variables such that X_i has d_i states. Their joint distribution is a $d_1 \times \cdots \times d_n$ tensor P. Let $\mathcal{C}$ be a simplicial complex on ground set $[n]$. Each $F \in \mathcal{C}$ corresponds to the tensor giving the marginal distribution on $\{X_i : i \in F\}$. Denote $\mathbf{d} := (d_1, \ldots, d_n)$ and let $\mathcal{L}_{\mathcal{C}, \mathbf{d}} \subset \mathbb{R}^{d_1} \otimes \cdots \otimes \mathbb{R}^{d_n}$ be the lattice consisting of $d_1 \times \cdots \times d_n$-way integer tables whose marginals indexed by $\mathcal{C}$ are all zero. We will soon see that the lattice $\mathcal{L}_{\mathcal{C}, \mathbf{d}}$ can be realized as the kernel of a matrix $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ and so $I_{\mathcal{L}_{\mathcal{C}, \mathbf{d}}} = I_{\mathcal{A}_{\mathcal{C}, \mathbf{d}}}$ is toric (see Section 1.6 for the necessary background on toric ideals). The *hierarchical model* corresponding to $(\mathcal{C}, \mathbf{d})$ is the set of joint probability tensors P such that $\log(P) \in V_{\mathcal{A}_{\mathcal{C}, \mathbf{d}}}$ Such models are used in categorical data analysis [63, Sections 1.2 and 7.2]. Each pair $(\mathcal{C}, \mathbf{d})$ consisting of a simplicial complex and an integer weighting of its ground set will be called an *HM pair*. Definition 4.1.2 gives a construction of a matrix $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ such that $\ker_{\mathbb{Z}} \mathcal{A}_{\mathcal{C}, \mathbf{d}} = \mathcal{L}_{\mathcal{C}, \mathbf{d}}$.

Definition 4.1.2. Fix an HM pair $(\mathcal{C}, \mathbf{d})$. Define $\mathbf{d}_F := [d_{i_1}] \times \cdots \times [d_{i_k}]$ for each facet $F = \{i_1, \ldots, i_k\}$ of $\mathcal{C}$. We write $\mathbb{R}^{\mathbf{d}_F}$ for the vector space with coordinates indexed by $\mathbf{j} \in \mathbf{d}_F$ (whose coordinates are in turn indexed by the vertices of F). For $\mathbf{i} \in [d_1] \times \cdots \times [d_n]$, define $a^{\mathbf{i}} \in \bigoplus_{F \in \mathrm{facet}(C)} \mathbb{R}^{\mathbf{d}_F}$ such that

$$
a_{F, \mathbf{j}}^{\mathbf{i}} = \begin{cases} 1 & \text{whenever } \mathbf{i}_k = \mathbf{j}_k \text{ for each } k \in F \\ 0 & \text{otherwise} \end{cases}
$$

and let $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ denote the matrix with columns $a^{\mathbf{i}}$ as $\mathbf{i}$ ranges over $[d_1] \times \cdots \times [d_n]$.

This matrix $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ can be seen as the representation in standard bases of the linear map that takes a $d_1 \times \cdots \times d_n$-way table to its marginals indexed by $\mathcal{C}$. It then follows easily from the definition of $\mathcal{L}_{\mathcal{C}, \mathbf{d}}$ that $\ker_{\mathbb{Z}} \mathcal{A}_{\mathcal{C}, \mathbf{d}} = \mathcal{L}_{\mathcal{C}, \mathbf{d}}$. We give an example of this construction.

Example 4.1.3. Consider the simplicial complex $\mathcal{C}$ with ground set $[3]$ and facets $\{1, 2\}$ and $\{2, 3\}$. The matrix $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ for $\mathbf{d} = (3, 2, 2)$ as in Definition 4.1.2 is displayed in Figure 4.1 with row and column labels.

We remind the reader of some terms and notation introduced in 1.6.2. Given a vector $u \in \mathbb{R}^n$, we can write $u = u^+ - u^-$ where u^+ and u^- have disjoint support and nonnegative entries. Define a partial order $\prec$ on $\mathbb{R}^n$ such that $u \prec v$ if $u^+ \leq u^-$ and $v^+ \leq v^-$.

		1 1 1	1 1 2	1 2 1	1 2 2	2 1 1	2 1 2	2 2 1	2 2 2	3 1 1	3 1 2	3 2 1	3 2 2
{1,2}	(1,1)	1	1	0	0	0	0	0	0	0	0	0	0
	(1,2)	0	0	1	1	0	0	0	0	0	0	0	0
	(2,1)	0	0	0	0	1	1	0	0	0	0	0	0
	(2,2)	0	0	0	0	0	0	1	1	0	0	0	0
	(3,1)	0	0	0	0	0	0	0	0	1	1	0	0
	(3,2)	0	0	0	0	0	0	0	0	0	0	1	1
{2,3}	(1,1)	1	0	0	0	1	0	0	0	1	0	0	0
	(1,2)	0	1	0	0	0	1	0	0	0	1	0	0
	(2,1)	0	0	1	0	0	0	1	0	0	0	1	0
	(2,2)	0	0	0	1	0	0	0	1	0	0	0	1

Figure 4.1: $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ for the HM pair $(\mathcal{C}, \mathbf{d})$ in Example 4.1.3.

Definition 4.1.4. The *Graver basis* of an integer matrix $A \in \mathbb{Z}^{m \times n}$, denoted $\mathrm{Gr}(A)$, is the set of $\prec$-minimal elements of $\ker_{\mathbb{Z}} A$.

Definition 4.1.5. An integer matrix $A \in \mathbb{Z}^{m \times n}$ and its corresponding toric ideal I_A and toric variety V_A are said to be *unimodular* if any/all of the following four equivalent conditions (c.f. Proposition 1.6.10) hold:

(1) every element of the Graver basis of A is a $\{0, 1, -1\}$ vector

(2) every circuit of A is a $\{0, 1, -1\}$ vector

(3) For any matrix A' obtained from $r := \mathrm{rank}\, A$ linearly independent rows of A, there exists some λ such that each $r \times r$ minor of A' is 0 or $\pm\lambda$

(4) For any b in the affine semigroup generated by the columns of A, the polyhedron $P_{A,b} = \{x \in \mathbb{R}^s : Ax = b, x \geq 0\}$ has all integral vertices.

4.1.1 Problem statements, motivation, and outline

We say that an *HM pair* $(\mathcal{C}, \mathbf{d})$ *is unimodular* to mean that the toric ideal $I_{\mathcal{A}_{\mathcal{C},\mathbf{d}}}$ is unimodular (c.f. Definition 1.6.8). This chapter solves the following problem.

Problem 4.1.6. Classify the unimodular HM pairs $(\mathcal{C}, \mathbf{d})$.

From the perspective of hypothesis testing, unimodularity is perhaps the most desirable property a hierarchical model could satisfy. In particular, the integer programming problems that arise in sequential importance sampling are easy to solve for unimodular models [20]. If a hierarchical model is unimodular, then it also satisfies a property called *normality* which was identified by Rauh and Sullivant in [50] as helpful when one uses toric fiber products to construct Markov bases of hierarchical models. Moreover, classifying the unimodular hierarchical models is a good first step towards the problem of classifying the normal hierarchical models since these families coincide for HM pairs $(\mathcal{C}, (2, \ldots, 2))$ when $\mathcal{C}$ has a facet containing all but one vertex. See [11] for more about the normality question. This chapter also solves the following problem.

Problem 4.1.7. Combinatorially describe the Graver bases of all unimodular hierarchical models.

The motivation for solving Problem 4.1.7 is that the algorithm described in [24] can be used to perform hypothesis tests on a hierarchical model, but only if a Markov basis of the corresponding toric ideal is known. Since Graver bases are Markov bases [57], the solution to Problem 4.1.7 has algorithmic consequences. Moreover, efficient generation of random elements of such a Markov basis is also necessary, and having a combinatorial description of the Markov basis can be leveraged for this purpose. Knowing the Graver basis of unimodular hierarchical models can also be used to generate Markov bases of more general hierarchical models via the toric fiber product construction [59].

We solve Problem 4.1.6 in two major steps. The first step is to classify the simplicial complexes $\mathcal{C}$ such that $(\mathcal{C}, (2, \ldots, 2))$ is unimodular. This is done in Sections 4.2 through 4.5. The remaining sections solve the general problem. Our solution to Problem 4.1.7 is almost an immediate consequence of our solution to Problem 4.1.6.

4.2 Constructions of Unimodular Complexes

Given a simplicial complex $\mathcal{C}$, we say that $\mathcal{C}$ *is unimodular* if the HM pair $(\mathcal{C}, (2, \ldots, 2))$ is unimodular. In this section we describe some operations on simplicial complexes that preserve unimodularity. In particular we show that unimodularity is preserved when passing to induced subcomplexes, to the Alexander dual of a simplicial complex, and to the link of a face of a simplicial complex. We can also build new unimodular complexes

from old ones by adding cone vertices, ghost vertices, or taking a Lawrence lifting. We also construct the basic examples of unimodular complexes. These tools together go in to our constructive description of unimodular complexes.

Proposition 4.2.1. *All induced sub-complexes of a unimodular complex are unimodular.*

Proof. Let $\mathcal{C}'$ be the induced subcomplex of $\mathcal{C}$ obtained by restricting to some vertex set $F \subseteq [n]$. Let B be the matrix obtained by taking the columns of $\mathcal{A}_{\mathcal{C}}$ corresponding to the elements $i = (i_1, \ldots, i_n) \in \prod_{j \in [n]} [d_j]$ such that $i_j = 1$ for all $j \in [n] \setminus F$. Let B' be the matrix obtained from B by removing rows of all zeros. Then $B' = \mathcal{A}_{\mathcal{C}'}$. So $\mathcal{A}_{\mathcal{C}'}$ and B have the same Graver basis, which is a subset of the Graver basis of $\mathcal{A}_{\mathcal{C}}$. So unimodularity of $\mathcal{A}_{\mathcal{C}}$ implies unimodularity of $\mathcal{A}_{\mathcal{C}'}$. $\square$

Next we will show that unimodularity is preserved under the Alexander duality operation.

Definition 4.2.2. Let $\mathcal{C}$ be a simplicial complex on $[n]$. The *Alexander dual* of $\mathcal{C}$, denoted $\mathcal{C}^*$ is defined as:
$$\mathcal{C}^* = \{S \subseteq [n] : [n] \setminus S \notin \mathcal{C}\}.$$

Note that if $\mathcal{C}$ has d faces, then $\mathcal{C}^*$ has $2^n - d$ faces and that the facets of $\mathcal{C}^*$ are the complements of the *minimal* non-faces of $\mathcal{C}$.

Since we are now restricting attention to the binary case, where $d_i = 2$ for each $i \in [n]$, entries of vectors in $\mathbb{R}^{\mathbf{d}}$ are indexed by elements of $\{1, 2\}^n$ and entries of vectors in $\bigoplus_{f \in \text{facet}(\mathcal{C})} \mathbb{R}^{\mathbf{d}_F}$ are indexed by pairs $(S, \mathbf{j})$ where $S \in \text{facet}(\mathcal{C})$ and $\mathbf{j} \in \{1, 2\}^S$. However, the statement and proof of Proposition 4.2.6 become cleaner if we replace $\{1, 2\}^n$ and $\{1, 2\}^S$ by $\{0, 1\}^n$ and $\{0, 1\}^S$, respectively. So we adapt this convention for the remainder of this section. We now introduce some notation that we use in Propositions 4.2.4 and 4.2.6.

Definition 4.2.3. Let $S \subseteq [n]$ and $\mathbf{j} \in \{0, 1\}^S$ and $\mathbf{i} \in \{0, 1\}^{[n] \setminus S}$. Then $e_{\mathbf{i}, \mathbf{j}}$ denotes the element $v \in \mathbb{R}^{\{0, 1\}^n}$ such that $v_{\mathbf{k}} = 1$ if $k_t = j_t$ when $t \in S$ and $k_t = i_t$ when $t \in [n] \setminus S$, and $v_{\mathbf{k}} = 0$ otherwise.

We are now ready to give an explicit description of $\mathcal{A}_{\mathcal{C}^*}$.

Proposition 4.2.4. *Let $\mathcal{C}$ be a simplicial complex on $[n]$. Let M be a matrix with the following set of columns:*

$$\left\{ \sum_{\mathbf{j} \in \{0,1\}^S} e_{\mathbf{i},\mathbf{j}} : S \text{ is a minimal non-face of } \mathcal{C}, \mathbf{i} \in \{0,1\}^{[n] \backslash S} \right\}.$$

Then $\mathcal{A}_{\mathcal{C}^*} = M^T$.

Proof. We index the columns of M by pairs $(S, \mathbf{i})$ where S is a minimal non-face of $\mathcal{C}$, and $\mathbf{i} \in \{0,1\}^{[n] \backslash S}$. The rows of M are indexed by elements of $\{0,1\}^n$. By construction, $M(\mathbf{k}, (S, \mathbf{i})) = 1$ if and only if $\mathbf{k}|_{[n] \backslash S} = \mathbf{i}$. Similarly, when F is a facet of $\mathcal{C}^*$, $\mathcal{A}_{\mathcal{C}^*}((F, \mathbf{i}), \mathbf{k}) = 1$ if and only $\mathbf{k}|_F = \mathbf{i}$. The result follows since S is a minimal non-face of $\mathcal{C}$ if and only if $[n] \backslash S$ is a facet of $\mathcal{C}^*$. $\qquad\square$

In order to relate the unimodularity of $\mathcal{A}_{\mathcal{C}}$ and $\mathcal{A}_{\mathcal{C}^*}$, we need two propositions. The first is a standard result from the theory of matroid duality, though we could not find a precise reference.

Proposition 4.2.5. *Let $A \in \mathbb{K}^{r \times n}$ have rank r. Let $B \in \mathbb{K}^{(n-r) \times n}$ have rank $n - r$ such that $AB^T = 0$. Then there exists a non-zero scalar $\lambda \in \mathbb{K}^*$ such that for all $S \subseteq [n]$ with $\#S = r$,*

$$\det(A_S) = \pm \lambda \det(B_{[n] \backslash S}).$$

Proof. Note that $AB^T = 0$ if and only if $CAB^T D^T = 0$ for all $C \in GL_r(\mathbb{K})$ and $D \in GL_{n-r}(K)$. Moreover, the $r \times r$ minors of A and CA are constant scalar multiples of each other as are the $(n-r) \times (n-r)$ minors of B and DB. Therefore, after multiplying by an element of $GL_r(\mathbb{K})$ and permuting columns, we can assume A has the form $A = (I_r \; M)$ where I_r is an $r \times r$ identity matrix. Similarly, we can also suppose B has the form $B = (-M^T \; I_{n-r})$. Then any $\det(A_S)$ is a minor of M with row indices $[r] \backslash S$ and column indices $S \backslash [r]$. The determinant $\det(B_{[n] \backslash S})$ is the same minor of M, up to a sign. $\qquad\square$

Proposition 4.2.6. *Let $\mathcal{C}$ be a simplicial complex on $[n]$. Then the following set of vectors spans $\ker(\mathcal{A}_{\mathcal{C}})$:*

$$K = \left\{ \sum_{\mathbf{j} \in \{0,1\}^S} (-1)^{\|\mathbf{j}\|_1} e_{\mathbf{i},\mathbf{j}} : S \text{ is a minimal non-face of } \mathcal{C}, \mathbf{i} \in \{0,1\}^{[n] \backslash S} \right\}.$$

Furthermore, if we view the entries of K as the columns of a matrix M', then we can multiply some set of rows and columns of M' by -1 and get M, as defined in Proposition 4.2.4.

Proof. First, note that for any $x \in K$, $\mathcal{A}_{\mathcal{C}} x = 0$, so $K \subset \ker(\mathcal{A}_{\mathcal{C}})$. If we view the entries of K as the columns of a matrix M', we are done if we show that $\text{rank}(M') = \dim(\ker(\mathcal{A}_{\mathcal{C}}))$. We proceed by proving the second statement of the lemma. The first statement will follow since Proposition 4.2.4 implies $\text{rank}(M) = \text{rank}(\mathcal{A}_{\mathcal{C}^*})$ and we know $\text{rank}(\mathcal{A}_{\mathcal{C}^*}) = 2^n - \#\mathcal{C} = \dim(\ker(\mathcal{A}_{\mathcal{C}}))$ ([33, Thm 2.6]).

As before, we index the columns of M' by pairs $(S, \mathbf{i})$ where S is a minimal non-face of $\mathcal{C}$, and we index the rows of M' by the binary n-tuples. Let $M'_{(S,\mathbf{i})}$ be a column of M'. For any $\mathbf{k} \in \{0,1\}^n$, the entry $M'_{(S,\mathbf{i}),\mathbf{k}}$ is nonzero if and only if $\mathbf{k}|_{[n]\backslash S} = \mathbf{i}$ - i.e. for each column, the $\mathbf{k}$ for each nonzero entry has fixed $\mathbf{k}|_{[n]\backslash S} = \mathbf{i}$. Now in this case, recall that $M'_{(S,\mathbf{i}),\mathbf{k}} = +1$ if $\|\mathbf{j}\|_1$ is even, and $M'_{(S,\mathbf{i}),\mathbf{k}} = -1$ if $\|\mathbf{j}\|_1$ is odd. Therefore, when $\|\mathbf{i}\|_1$ is even, $M'_{(S,\mathbf{i}),\mathbf{k}} = +1$ if $\|\mathbf{k}\|_1$ is even and $M'_{(S,\mathbf{i}),\mathbf{k}} = -1$ when $\|\mathbf{k}\|_1$ is odd. And when $\|\mathbf{i}\|_1$ is odd, $M'_{(S,\mathbf{i}),\mathbf{k}} = +1$ when $\|\mathbf{k}\|_1$ is odd and $M'_{(S,\mathbf{i}),\mathbf{k}} = -1$ when $\|\mathbf{k}\|_1$ is even. So for a fixed column $M'_{(S,\mathbf{i})}$, the sign of a nonzero entry $M'_{(S,\mathbf{i}),\mathbf{k}}$ depends only on the parity of $\mathbf{k}$. So if we multiply all the rows $M'_{,\mathbf{k}}$ such that $\|\mathbf{k}\|_1$ is odd by -1, the entries in each column will all have the same sign. Then we can multiply all the negative columns by -1, to arrive at the matrix from Proposition 4.2.4. $\square$

Now we are ready to show that Alexander duality preserves unimodularity.

Proposition 4.2.7. *Let $\mathcal{C}$ be a simplicial complex on ground set $[n]$. Then $\mathcal{C}$ is unimodular if and only if $\mathcal{C}^*$ is unimodular.*

Proof. Consider the set K from Proposition 4.2.6 as a matrix of column vectors. Since K spans the kernel of $\mathcal{A}_{\mathcal{C}}$, Proposition 4.2.5 implies that K^T is unimodular if and only if $\mathcal{A}_{\mathcal{C}}$ is unimodular. Then we can multiply the appropriate rows and columns of K by -1 to get M as in Proposition 4.2.4. Since these operations preserve the absolute values of full rank determinants, K^T is unimodular if and only if M^T is. Proposition 4.2.4 says that $\mathcal{A}_{\mathcal{C}^*} = M^T$. $\square$

Taking Alexander duals and induced complexes gives rise to another unimodularity preserving operation which we now define.

Definition 4.2.8. Let $S \in \mathcal{C}$ be a face of $\mathcal{C}$. Then the link of S in $\mathcal{C}$ is the new simplicial complex

$$\text{link}_S(\mathcal{C}) = \{F \setminus S : F \in \mathcal{C} \text{ and } S \subseteq F\}.$$

When $S = \{v\}$, we simply write $\text{link}_v(\mathcal{C}) := \text{link}_{\{v\}}(\mathcal{C})$.

Note that we can obtain $\text{link}_S(\mathcal{C})$ by repeatedly taking links with respect to vertices. That is if S is a face of $\mathcal{C}$ and $\#S \geq 2$ and $v \in S$, then

$$\text{link}_S(\mathcal{C}) = \text{link}_v(\text{link}_{S \setminus \{v\}}(\mathcal{C})).$$

Proposition 4.2.9. *If $\mathcal{C}$ is a simplicial complex and S is a face of $\mathcal{C}$, then $\text{link}_S(\mathcal{C}) = (\mathcal{C}^* \setminus S)^*$.*

Proof. By definition, we have:

$$(\mathcal{C}^* \setminus S)^* = \{R \subseteq [n] \setminus S : ([n] \setminus S) \setminus R \notin \mathcal{C}^* \setminus S\}.$$

Then we have the following chain of equivalences on some $R \subseteq [n] \setminus S$:

$$([n] \setminus S) \setminus R \notin \mathcal{C}^* \setminus S$$
$$\iff ([n] \setminus S) \setminus R \notin \mathcal{C}^*$$
$$\iff [n] \setminus (([n] \setminus S) \setminus R) \in \mathcal{C}$$

and since $[n] \setminus (([n] \setminus S) \setminus R) = R \cup S$, we have $R \cup S \in \mathcal{C}$. $\square$

Corollary 4.2.10. *Let $\mathcal{C}$ be a unimodular simplicial complex on ground set $[n]$. Then for any face S of $\mathcal{C}$, $\text{link}_S(\mathcal{C})$ is unimodular.*

Proof. Proposition 4.2.9 implies that $\text{link}_S(\mathcal{C})$ can be obtained via Alexander duality and passing to an induced subcomplex. Unimodularity then follows from Propositions 4.2.1 and 4.2.7. $\square$

Now we turn to operations for taking a complex that is unimodular and constructing larger unimodular complexes. If $\mathcal{C}$ has a vertex v that lies in each facet of $\mathcal{C}$, we say that v is a *cone vertex* of $\mathcal{C}$. The following proposition tells us that unimodularity is invariant under adding or removing a cone vertex.

Proposition 4.2.11. *Let $\mathcal{C}$ be a simplicial complex on $[n]$. We define $\mathcal{C}'$ on $[n+1]$ to be the simplicial complex with the following facets:*

$$\{F \cup \{n+1\} : F \text{ is a facet of } \mathcal{C}\}.$$

Then $\mathcal{A}_\mathcal{C}$ is unimodular if and only if $\mathcal{A}_{\mathcal{C}'}$ is.

Proof. We will index the columns of $\mathcal{A}_{\mathcal{C}'}$ by the binary $n+1$ tuples such that those with the $n+1$ coordinate equal to 1 come before those with $n+1$ coordinate equal to 2. Then we get the following block form:

$$\mathcal{A}_{\mathcal{C}'} = \left(\begin{array}{c|c} \mathcal{A}_\mathcal{C} & \mathbf{0} \\ \hline \mathbf{0} & \mathcal{A}_\mathcal{C} \end{array} \right).$$

The Graver basis of $\mathcal{A}_{\mathcal{C}'}$ is therefore $\{(u,0), (0,u) : u \in Gr_A\}$. Hence the Graver basis of $\mathcal{A}_{\mathcal{C}'}$ consists of $0, \pm 1$ elements if and only if the Graver basis of $\mathcal{A}_\mathcal{C}$ consists of $0, \pm 1$ elements. $\qquad\square$

By induction, adding or removing multiple cone vertices from a simplicial complex does not affect unimodularity. We introduce the following notation to denote this.

Definition 4.2.12. Let $\mathcal{C}$ be a simplicial complex on vertex set $[n]$. Then we define $\mathrm{cone}^p(\mathcal{C})$ to be the simplicial complex on $[n] \cup \{u_1, \ldots, u_p\}$ with the following facets

$$\{F \cup \{u_1, \ldots, u_p\} : F \text{ is a facet of } \mathcal{C}\}.$$

When $p = 0$, we define $\mathrm{cone}^p(\mathcal{C}) = \mathcal{C}$.

Definition 4.2.13. For any matrix $A \in \mathbb{R}^{s \times t}$, we define the *Lawrence lifting* of A to be the matrix

$$\Lambda(A) \quad = \quad \begin{pmatrix} A & \mathbf{0} \\ \mathbf{0} & A \\ \mathbf{1} & \mathbf{1} \end{pmatrix} \quad \in \quad \mathbb{R}^{(2s+t) \times 2t}$$

where $\mathbf{0}$ denotes the $s \times t$ matrix of all zeroes and $\mathbf{1}$ denotes a $t \times t$ identity matrix.

By Theorem 7.1 in [57], $\Lambda(A)$ is unimodular if and only if A is unimodular. This gives rise to another unimodularity-preserving operation on simplicial complexes.

Definition 4.2.14. Let $\mathcal{C}$ be a simplicial complex on $[n]$. We define the *Lawrence lifting* of $\mathcal{C}$ to be the simplicial complex $\Lambda(\mathcal{C})$ on $[n+1]$ that has the following set of facets:

$$\{[n]\} \cup \{F \cup \{n+1\} : F \text{ is a facet of } \mathcal{C}\}.$$

In this case, we refer to the facet $[n]$ as a *big facet*.

If a complex $\mathcal{C}$ on $[n+1]$ has big facet $[n]$, then $\mathcal{C} = \Lambda(\text{link}_{n+1}(\mathcal{C}))$. In particular, $\mathcal{C} = \Lambda(\mathcal{D})$ for some complex $\mathcal{D}$ if and only if $\mathcal{C}$ has a big facet.

Proposition 4.2.15. *The simplicial complex $\mathcal{C}$ is unimodular if and only if $\Lambda(\mathcal{C})$ is unimodular.*

Proof. Note that by construction $\Lambda(\mathcal{A}_\mathcal{C}) = \mathcal{A}_{\Lambda(\mathcal{C})}$. Hence Theorem 7.1 in [57] implies the proposition. $\qquad\square$

If a simplicial complex $\mathcal{C}$ can be expressed as $\Lambda(\mathcal{C}')$, then we say that $\mathcal{C}$ is of Lawrence type. A simplicial complex on n vertices is of Lawrence type if and only if it has a facet containing $n-1$ vertices. We will refer to any facet of $\mathcal{C}$ that has $n-1$ elements as a *big facet*.

We now give our final unimodularity preserving operation.

Definition 4.2.16. Let $\mathcal{C}$ be a simplicial complex on ground set $[n]$. Let $G^p(\mathcal{C})$ denote the same simplicial complex but on ground set $[n+p]$. Note that the vertices $n+1, \ldots, n+p$ are not contained in any face of $G^p(\mathcal{C})$. In this case we say that $n+1, \ldots, n+p$ are *ghost vertices*. When $p = 1$ we drop the superscript; i.e. we just write $G(\mathcal{C})$.

Proposition 4.2.17. *A simplicial complex $\mathcal{C}$ is unimodular if and only if $G^p(\mathcal{C})$ is unimodular.*

Proof. This is true because

$$\mathcal{A}_{G^p(\mathcal{C})} = \begin{pmatrix} \mathcal{A}_\mathcal{C} & \cdots & \mathcal{A}_\mathcal{C} \end{pmatrix}. \qquad\square$$

Note that adding a ghost vertex to a complex is Alexander dual to taking the Lawrence lifting of a simplicial complex, that is $G(\mathcal{C})^* = \Lambda(\mathcal{C}^*)$.

We now state a useful fact about the interaction of these operations.

Proposition 4.2.18. *The operation* $\text{cone}^p(\cdot)$ *commutes with the operations of taking links, duals, induced sub-complexes and adding ghost vertices.*

For any vertex v of $\mathcal{C}$, we let $\mathcal{C} \setminus v$ denote the induced subcomplex on $\mathcal{C} \setminus \{v\}$. Then for vertices $v \neq u$ of $\mathcal{C}$, the operations $\cdot \setminus v$ and $\text{link}_u(\cdot)$ commute. So if $\mathcal{D}$ can be obtained from $\mathcal{C}$ by applying a series of deletions, and taking links, then we can write

$$\mathcal{D} = \text{link}_R(\mathcal{C} \setminus S)$$

where $F \in \mathcal{C}$ and S, R are a subset of vertices of $\mathcal{C}$ such that $S \cap R = \emptyset$. This gives rise to the following definition:

Definition 4.2.19. Let $\mathcal{C}$ be a simplicial complex. Then $\mathcal{D}$ is a *minor* of $\mathcal{C}$ if $\mathcal{D} = \text{link}_R(\mathcal{C} \setminus S)$ where S and R are subsets of vertices of $\mathcal{C}$ such that $S \cap R = \emptyset$, and R is a face of $\mathcal{C}$.

We warn the reader that although it is natural to view simplicial complexes as a generalization of graphs, our definition of simplicial complex minor does *not* generalize the usual notion of graph minor. Alternatively, a simplicial complex can be seen as a generalization of a matroid, and our definition of minor is a generalization of the usual notion of a matroid minor [37]. This definition of simplicial complex minor is useful for our purposes because of Proposition 4.2.20.

Proposition 4.2.20. *If $\mathcal{C}$ is a unimodular simplicial complex, then every minor of $\mathcal{C}$ is unimodular.*

Proof. This follows immediately from Propositions 4.2.1 and 4.2.10. □

The fundamental example of a unimodular simplicial complex is the disjoint union of two simplices. Here a simplex Δ_n is the simplicial complex on an $n + 1$ element set with a single facet consisting of all the elements. We denote the irrelevant complex $\{\emptyset\}$ by Δ_{-1} and the void complex $\{\}$ by Δ_{-2}.

Proposition 4.2.21. *Let $\mathcal{C} = \Delta_m \sqcup \Delta_n$ be the disjoint union of two simplices, for $m, n \geq 0$. Then $\mathcal{C}$ is unimodular.*

Proof. The matrix $\mathcal{A}_{\mathcal{C}}$ in this case is the vertex-edge incidence matrix of a complete bipartite graph with 2^{m+1} and 2^{n+1} vertices in the two parts of the partition. Such

vertex-edge incidence matrices are examples of network matrices and are hence totally unimodular [51, Ch. 19]. □

A second family of fundamental examples comes from taking the duals of the disjoint union of two simplices. We use $D_{m,n}$ to denote the dual of a disjoint union of an m-simplex and an n-simplex, i.e.

$$D_{m,n} := (\Delta_m \sqcup \Delta_n)^*.$$

We close this section by giving a workable description of $D_{m,n}$. We can divide the vertices of $D_{m,n}$ into disjoint sets M, N such that M contains the vertices of the Δ_m in $D_{m,n}^*$ and N contains the vertices of the Δ_n in $D_{m,n}^*$. Then, the facets of $D_{m,n}$ are precisely the subsets of $M \sqcup N$ that leave out exactly one element of M and one element of N. Notice that the complexes induced on M, N are $\partial \Delta_m$ and $\partial \Delta_n$ respectively. Also, note that for any $v \in M$, $\text{link}_v(D_{m,n}) = D_{m-1,n}$ and for any $v \in N$, $\text{link}_v(D_{m,n}) = D_{m,n-1}$.

4.3 β-avoiding Simplicial Complexes

Part of our main result is a forbidden minor classification of unimodular simplicial complexes. In this section, we identify these forbidden minors and prove various properties about the complexes that avoid them.

Proposition 4.3.1. *The following simplicial complexes are minimal nonunimodular simplicial complexes:*

1. *P_4, the path on 4 vertices*
2. *O_6, the boundary of the octahedron or its dual O_6^**
3. *J_1, the complex on $\{1, 2, 3, 4, 5\}$ with facets $12, 15, 234, 345$ or its dual J_1^*, with facets $134, 235, 245$.*
4. *J_2, complex on $\{1, 2, 3, 4, 5\}$ with facets $12, 235, 34, 145$*
5. *For $n \geq 1$, $\partial \Delta_n \sqcup \{v\}$, the disjoint union of the boundary of an n-simplex and a single vertex.*

Note that P_4, J_2 and $\partial \Delta_n \sqcup \{v\}$ are isomorphic to their own duals.

Proof. We can check that the complexes $P_4, O_6, O_6^*, J_1, J_1^*$, and J_2 are not unimodular by using the software 4ti2 [2] to compute the Graver basis and looking for entries that are not
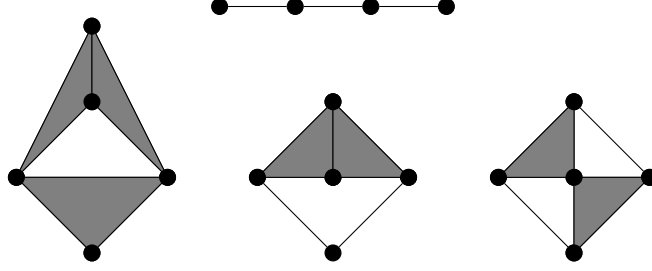
Figure 4.2: The complexes P_4 (top), J_1 (center), J_1^* (left), and J_2 (right).

$0, \pm 1$. In the case of O_6 and O_6*, these are too large to compute the entire Graver basis. However, selecting sufficiently large random subsets of the columns produced Graver basis elements of the desired form. A Macaulay2 script for doing these computations can be found on my website [9]. For the infinite family $\partial \Delta_n \sqcup \{v\}$ where $n \geq 1$, examples of a non-squarefree Graver basis element appear in [58]. These results show that these examples are not unimodular. To see that they are minimal, note that every subcomplex obtained by deleting a single vertex or taking the link at a vertex produced a unimodular complex in all cases. □

Definition 4.3.2. A simplicial complex $\mathcal{C}$ is β-*avoiding* if it does not contain any of the complexes from Proposition 4.3.1 as minors.

Since none of the complexes from Proposition 4.3.1 is unimodular, Proposition 4.2.20 implies that if $\mathcal{C}$ is unimodular, then $\mathcal{C}$ is β-avoiding. The converse of this is our forbidden minor classification of unimodular complexes which we prove in Theorem 4.5.3.

We now give some technical results about β-avoiding complexes. Before beginning, we remind the reader that $\partial \Delta_1 \sqcup \{v\}$ is an independent set on 3 vertices, and so no β-avoiding complex can have an independent set of size 3.

Proposition 4.3.3. *Let $\mathcal{C}$ be a β-avoiding simplicial complex. Then $\mathcal{C}^*$ is also β-avoiding.*

Proof. The list of prohibited minors of β-avoiding complexes is closed under taking duals. The proposition follows when we note that if $\mathcal{D}$ is a minor of $\mathcal{C}$, then $\mathcal{D}^*$ is a minor of $\mathcal{C}^*$. This is true because if $\mathcal{D}^* = \text{link}_S(\mathcal{C}^* \setminus R)$, then two applications of Proposition 4.2.9 give

$$\mathcal{D}^* = ((\mathcal{C}^* \setminus R)^* \setminus S)^*$$
$$= (\text{link}_R(\mathcal{C}) \setminus S)^*$$

84

and therefore $\mathcal{D} = \text{link}_R(\mathcal{C}) \setminus S$. So $\mathcal{D}$ is a minor of $\mathcal{C}$. $\qquad\square$

Proposition 4.3.4. *Let $\mathcal{C}$ be a β-avoiding simplicial complex that has C_4 induced. Then the complex induced on the non-ghost vertices of $\mathcal{C}$ is $\text{cone}^p(C_4)$ for some p.*

Proof. Assume $\mathcal{C}$ has no ghost vertices. Let $u_1, \ldots, u_4$ be the vertices that induce C_4 in $\mathcal{C}$. If these are the only vertices, we are done, so let v be another vertex in $\mathcal{C}$. Let $\mathcal{D}$ denote the complex induced on $v, u_1, \ldots, u_4$. Any minor of $\mathcal{D}$ is a minor of $\mathcal{C}$, so $\mathcal{D}$ must also be β-avoiding. Vertex v cannot be disconnected from $u_1, \ldots, u_4$ since otherwise $\mathcal{D}$ has an independent set of size 3. Furthermore, v must connect to $u_1, \ldots, u_4$ with a triangle. Otherwise if v connected to u_i with an edge, then u_i has edge degree 3 and so $\text{link}_v(\mathcal{D})$ is an independent set on 3 vertices. There are 4 possible triangles that v can join with $u_1, \ldots, u_4$ in $\mathcal{D}$ so that $u_1, \ldots, u_4$ induce C_4. If v is only in one such triangle, wlog $\{v, u_1, u_2\}$, then $\{v, u_1, u_3, u_4\}$ is an induced P_4. The two non-isomorphic complexes on 5 vertices that have two triangles and an induced C_4 are J_1 and J_2. If $\mathcal{D}$ has 3 of the possible triangles, then $\text{link}_v(\mathcal{D})$ is P_4. So $\mathcal{D}$ must have all four - i.e. it must be $\text{cone}^1(C_4)$.

Any vertices $v, v' \notin \{u_1, u_2, u_3, u_4\}$ must connect to each other. Otherwise, the induced complex on $v, v', u_1, \ldots, u_4$ is O_6 which is not unimodular. Now we denote the vertices of $\mathcal{C}$ that are not any of the u_is as $v_1, \ldots, v_k$. Since v_k is a cone point over the C_4 on $u_1, \ldots, u_4$ and since v_k connects to all $v_j, j < k$, $\text{link}_{v_k}(\mathcal{C})$ is a β-avoiding complex on $v_1, \ldots, v_{k-1}, u_1, \ldots, u_4$ that has C_4 induced on $u_1, \ldots, u_4$. So, by induction on the number of vertices, $\text{link}_{v_k}(\mathcal{C})$ is an iterated cone over C_4, $\text{cone}^{k-1}(C_4)$. But also the induced complex on $v_1, \ldots, v_{k-1}, u_1, \ldots, u_4$ is $\text{cone}^{k-1}(C_4)$ by induction. This implies that $\mathcal{C} = \text{cone}^k(C_4)$. $\qquad\square$

The following Proposition generalizes Proposition 4.3.4. Its proof is an induction argument that uses Proposition 4.3.4 as a base case.

Proposition 4.3.5. *Let $\mathcal{C}$ be a β-avoiding complex that has $D_{m,n}$ induced for some $m, n \geq 1$. Then the complex induced on the non-ghost vertices of $\mathcal{C}$ is $\text{cone}^p(D_{m,n})$ for some p.*

Proof. Assume $\mathcal{C}$ has no ghost vertices. Just as in the remarks at the end of section 4.2, we divide the vertices of $D_{m,n}$ into disjoint sets M, N such that M contains the vertices of the Δ_m in $D_{m,n}^*$ and N contains the vertices of the Δ_n in $D_{m,n}^*$. Now, we proceed by induction on m and n. Note that $D_{1,1} = C_4$, so the base case is handled by Proposition 4.3.4. Assume $m \geq 2$ without loss of generality.

Let $v \in \mathcal{C} \setminus D_{m,n}$. The vertex v must connect to some $u \in M$ to avoid inducing $\partial\Delta_m \sqcup \{v\}$. Then, $\text{link}_u(\mathcal{C})$ contains v and $D_{m-1,n}$, and so by induction, the complex induced on v and the vertices of $D_{m-1,n}$ in $\text{link}_u(\mathcal{C})$ is a cone over $D_{m-1,n}$ with v as a cone vertex. This means that v is in every facet that contains u. Since $m \geq 2$, u is connected to every other vertex in $D_{m,n}$ and thus v is attached to every vertex in $D_{m,n}$. Now we apply the same argument to each of the vertices in the set M to see that v is in every facet that contains any vertex in set M. Every facet of $D_{m,n}$ contains some element of M. So this implies that the induced complex on v and $D_{m,n}$ must be the cone over $D_{m,n}$.

Now assume v, v' are both vertices in $\mathcal{C} \setminus D_{m,n}$. If they were not connected by an edge, then the induced complex on v, v' and $D_{m,n}$ contains a minor which is isomorphic to O_6. This means that v, v' are connected by an edge. Taking the link $\text{link}_v(\mathcal{C})$ produces a smaller complex with an induced $D_{m,n}$ hence it must be a cone by induction on the number of vertices not in the $D_{m,n}$. This reduces us the case where $\text{link}_v(\mathcal{C})$ is $\text{cone}^{p-1}(D_{m,n})$ and $\mathcal{C} \setminus v$ is $\text{cone}^{p-1}(D_{m,n})$ which implies that $\mathcal{C}$ is $\text{cone}^p(D_{m,n})$. $\qquad\square$

Proposition 4.3.6. *Let $\mathcal{C}$ be a β-avoiding simplicial complex on vertices $u_1, \ldots, u_{k+1}, v$ such that the complex induced on $\{u_1, \ldots, u_{k+1}\}$ is $\partial\Delta_k$. Then v must be in a k-simplex with some subset of the u_i's.*

Proof. We proceed by induction on k. For the base case $k = 1$, note that $\partial\Delta_1$ is two isolated points. In this case, v must connect to u_1 or u_2 as a 1-simplex (an edge) to avoid inducing $\partial\Delta_1 \sqcup \{v\}$.

Now assume $k > 1$. The vertex v must attach to some u_i to avoid inducing $\partial\Delta_k \sqcup \{v\}$. Then, $\text{link}(u_i)$ has v and $\partial\Delta_{k-1}$, so by induction, v must form a $k-1$ simplex with some collection $u_1, \ldots, \hat{u}_i, \ldots, \hat{u}_j, \ldots, u_{k+1}$. So in $\mathcal{C}$, v must be in a k simplex with $u_1, \ldots, \hat{u}_j, \ldots, u_{k+1}$. $\qquad\square$

Proposition 4.3.7. *Let $\mathcal{C}$ be a β-avoiding simplicial complex on $m+n+2$ vertices that does not have a facet of dimension $m+n$. Assume $D_{m,n} \subseteq \mathcal{C}$. Then $\mathcal{C} = D_{m,n}$.*

Proof. Note that $\mathcal{C}^* \subseteq \Delta_m \sqcup \Delta_n$ because $D_{m,n} \subseteq \mathcal{C}$. Since $\mathcal{C}$ does not have a facet of dimension $m+n$, $\mathcal{C}^*$ has no ghost vertices. Since $\mathcal{C}^*$ has no ghost vertices, if $\mathcal{C}^* \subsetneq \Delta_m \sqcup \Delta_n$, then $\mathcal{C}^*$ would have an induced $\{v\} \sqcup \partial\Delta_k$ for some $k \geq 1$. In this case $\mathcal{C}^*$ would not be β-avoiding and so by Proposition 4.3.3, neither would $\mathcal{C}$. $\qquad\square$

4.4 The 1-Skeleton of a β-avoiding Complex

In this section we prove Lemma 4.4.3 which gives a complete characterization of the 1-skeleton of a β-avoiding simplicial complex. This is a crucial technical lemma in the proof of Theorem 4.5.3. We start with a technical proposition about graphs.

Proposition 4.4.1. *Let H be a connected graph that avoids K_3 and P_4 as induced subgraphs. Then H is a complete bipartite graph.*

Proof. Let $u \in V(H)$, let $N(u)$ denote the neighbors of u, and let $M(u)$ denote the non-neighbors of u (this set includes u). The bipartition of the vertices of H will be $M(u)$ and $N(u)$. Let $v \in M(u) \backslash \{u\}$. Since H is connected, there exists a path $u = u_1, u_2, \ldots, u_k = v$. Assume k is minimal. Since $v \neq u$, $k > 1$. We cannot have $k = 2$ since u, v are non-neighbors. We cannot have $k = 4$, since in order to avoid an induced P_4, we would need an edge (u_i, u_{i+2}) inducing a K_3, or an edge (u, v) contradicting that u, v are non-neighbors. If $k \geq 5$, there must exist an edge (u_1, u_4) to avoid an induced P_4, contradicting minimality of k. So we have $k = 3$. So for any $v \in M(u) \setminus \{u\}$, there exists a path u, a, v.

Now we show that H is bipartite with bipartition $M(u)$ and $N(u)$. It is clear that $N(u)$ is an independent set of vertices, for if $v, w \in N(u)$ had an edge between them, there would be a K_3 on u, v, w. Now we show that $M(u)$ is an independent set of vertices. Assume $w, v \in M(u)$. If either w, v is u, there is no edge between them by definition of $M(u)$, so assume $w, v \neq u$. Then by the above, we have paths u, a, v and u, b, w. So an edge (v, w) would induce the 5-cycle u, a, v, w, b, u. This 5-cycle must have a chord to avoid inducing P_4, but any chord in a 5-cycle induces a K_3.

Now we show that H is complete bipartite. Let $x \in M(u)$ and $y \in N(u)$. Since H is connected, there is a path $x = u_1, ..., u_k = y$. We may without loss of generality assume $k \leq 3$ since otherwise we could shorten the path using the edge (u_1, u_4) required to avoid a P_4. The sets $M(u)$ and $N(u)$ are disjoint, so $k \neq 1$. If $k = 3$, then there is an induced P_4 x, u_2, y, u. So $k = 2$, and so (x, y) is an edge. $\square$

For a graph G we define the complement graph G^c on the same set of vertices such that (u, v) is an edge of G^c if and only if (u, v) is not an edge in G. Now we can use Proposition 4.4.1 to give a strong restriction on the structure of G^c whenever G is the 1-skeleton of a β-avoiding complex.

Proposition 4.4.2. *If G is the 1-skeleton of a β-avoiding simplicial complex $\mathcal{C}$, then each connected component of G^c is complete bipartite.*

Proof. Note that $\{v\} \sqcup \partial \Delta_1$ consists of three disconnected vertices, its complement graph is K_3. The path P_4 is its own complement. Since $\mathcal{C}$ is β-avoiding, G avoids P_4 and an independent set of size three as induced subgraphs. So each connected component of G^c avoids P_4 and K_3 as induced subgraphs. Proposition 4.4.1 therefore implies that each connected component of G^c is complete bipartite. $\qquad \square$

Now we are ready to characterize the 1-skeleton of a β-avoiding simplicial complex.

Lemma 4.4.3. *Let $\mathcal{C}$ be a β-avoiding simplicial complex and let G denote its 1-skeleton. Then G is one of the following*

1. *K_N*
2. *Two complete graphs glued along a (possibly empty) common clique*
3. *The iterated cone over a 4-cycle.*

Proof. By Proposition 4.4.2, we know that each connected component of G^c is complete bipartite. Let H denote the induced subgraph of G^c that removes all isolated vertices. If H is empty, then G^c is an independent set of vertices and therefore $G = K_N$. So assume H is nonempty. We claim that if H is neither $K_{m,n}$ nor $K_2 \sqcup K_2$ for $m, n \geq 1$, then H induces either $K_2 \sqcup K_2 \sqcup K_2$ or $P_3 \sqcup K_2$. To prove the claim, first assume that H avoids $K_2 \sqcup K_2 \sqcup K_2$. Since H is not $K_{m,n}$ and has no isolated vertices, this implies that H has exactly two components. Since H is not $K_2 \sqcup K_2$, some connected component of H has at least three vertices. So we have a P_3 induced by this component, and K_2 induced by the other component and so the claim is proven.

The complement graphs of $K_2 \sqcup K_2 \sqcup K_2$ or $P_3 \sqcup K_2$ are shown in Figure 4.3. Both graphs have C_4 induced, but neither is the 1-skeleton for an iterated cone over C_4 since neither graph has a suspension vertex. Proposition 4.3.4 therefore implies that neither is the 1-skeleton of a β-avoiding simplicial complex and so G may not induce either. So G^c may not induce $K_2 \sqcup K_2 \sqcup K_2$ nor $P_3 \sqcup K_2$ and therefore neither may H. The claim then implies that H must be either $K_{m,n}$ or $K_2 \sqcup K_2$. Assume G^c has p isolated vertices. If $H = K_{m,n}$ then G is a K_{m+p} and a K_{n+p} glued along a common K_p. If $H = K_2 \sqcup K_2$ then G is an iterated cone over a 4-cycle. $\qquad \square$

As we see in the following proposition, the 1-skeleton of $\mathcal{C}$ completely determines $\mathcal{C}$ when it is obtained by gluing two complete graphs along an empty clique.
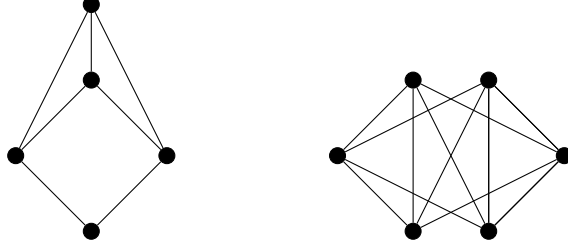
Figure 4.3: The complement graphs of $P_3 \sqcup K_2$ and $K_2 \sqcup K_2 \sqcup K_2$.

Proposition 4.4.4. *Let $\mathcal{C}$ be a β-avoiding simplicial complex such that the 1-skeleton of $\mathcal{C}$ is the disjoint union of cliques K_m and K_n. Then $\mathcal{C} = \Delta_m \sqcup \Delta_n$.*

Proof. Let $v_1, \ldots, v_k$ be vertices of the K_m. If $\{v_1, \ldots, v_k\}$ is a minimal non-face of $\mathcal{C}$, then if we let u be a vertex in the K_n, then the complex indued on $v_1, \ldots, v_k, u$ is $\partial\Delta_k \sqcup \{u\}$. $\quad\square$

4.5 The Main Theorem

The goal of this section is to give a proof of Theorem 4.5.3 which gives a complete characterization of the unimodular binary hierarchical models.

We begin this section by defining a *nuclear* complex. A nuclear complex is a complex that can be obtained from a disjoint union of two simplices by adding cone vertices, adding ghost vertices, taking Lawrence liftings and taking Alexander duals. Since $\Delta_m \sqcup \Delta_n$ is unimodular and these operations all preserve unimodularity, nuclear complexes are unimodular. Part of our main result is the converse - unimodular complexes are nuclear.

Definition 4.5.1. A simplicial complex $\mathcal{C}$ is *nuclear* if one of the following is true

(1) $\mathcal{C} = \Lambda(\mathcal{D})$ where $\mathcal{D}$ is nuclear
(2) $\mathcal{C} = G(\mathcal{D})$ where $\mathcal{D}$ is nuclear
(3) $\mathcal{C} = \mathrm{cone}^p(\Delta_m \sqcup \Delta_n)$ for $p, m, n \geq 0$
(4) $\mathcal{C} = \mathrm{cone}^p(D_{m,n})$ for $p \geq 0$ and $m, n \geq 1$
(5) $\mathcal{C} = \Delta_k$ for $k \geq -2$.

Every nuclear complex $\mathcal{C}$ can be constructed by applying the operations $\mathrm{cone}^p(\cdot)$, $G(\cdot)$, and $\Lambda(\cdot)$ to a complex $\mathcal{D}$ where $\mathcal{D}$ is of the form $\Delta_m \sqcup \Delta_n$, $D_{m,n}$, or Δ_k. We refer to $\mathcal{D}$ as the *nucleus* of $\mathcal{C}$.

Note that $D_{m,0}$ has a ghost vertex for all m. This is why we have $m, n \geq 1$ in ((4)). We note that the collection of nuclear complexes is closed under Alexander duality.

Proposition 4.5.2. *If $\mathcal{C}$ is nuclear then so is $\mathcal{C}^*$.*

Proof. Possibilities (1) and (2) are dual to each other, as are (3) and (4). On k vertices, Δ_k and Δ_{-2} are dual to each other. □

We now state our main result.

Theorem 4.5.3. *Let $\mathcal{C}$ be a simplicial complex. Then the following are equivalent*

(1) $\mathcal{C}$ is unimodular
(2) $\mathcal{C}$ is β-avoiding
(3) $\mathcal{C}$ is nuclear.

We defer the proof of Theorem 4.5.3 until the end of the section, but we give a roadmap here. It is immediate from results in previous sections that nuclear complexes are unimodular and that unimodular complexes are β-avoiding.

We show that any β-avoiding complex $\mathcal{C}$ is nuclear by choosing a particular vertex v of $\mathcal{C}$ and using induction on the number of vertices to conclude that $\text{link}_v(\mathcal{C})$ is nuclear. From here, we have five cases to consider - one for each of the ways that $\text{link}_v(\mathcal{C})$ can be nuclear. In four of these cases, it is relatively easy to show that $\mathcal{C}$ is nuclear. The difficulty lies in the case where $\text{link}_v(\mathcal{C}) = D_{m,n}$, which is handled in Lemma 4.5.8. The proof of Lemma 4.5.8 is split into two main parts and the second part is further split into six cases. Proposition 4.5.4 helps with the first part. Propositions 4.5.5, 4.5.6, and 4.5.7 handle the hard cases in the second part.

Proposition 4.5.4. *Let $\mathcal{C}$ be a β-avoiding simplicial complex. Assume that all proper minors of $\mathcal{C}$ are nuclear. Assume $\mathcal{C}$ has a vertex v such that $\mathcal{C} \setminus v = \text{cone}^p(\Delta_m \sqcup \Delta_n)$ with $m, n \geq 1$ and $p \geq 1$. Assume $\mathcal{C}$ is C_4-free. Then $\mathcal{C} = \text{cone}^{p+1}(\Delta_m \sqcup \Delta_n)$, $\mathcal{C} = \text{cone}^p(\Delta_{m+1} \sqcup \Delta_n)$, or $\mathcal{C} = \text{cone}^p(\Delta_m \sqcup \Delta_{n+1})$.*

Proof. We induct on p. For the base case, take $p = 1$. If we let u denote the cone vertex in $\mathcal{C} \setminus v$ then v must connect to one of the simplices in the $\Delta_m \sqcup \Delta_n \subset \mathcal{C} \setminus u$. Otherwise we have an independent set with three vertices induced in $\mathcal{C} \setminus u$ and this complex is not nuclear. If v connects to only one such simplex, then v must also connect to u to avoid inducing P_4 which is not nuclear. If v connects to both such simplices, and v does not

connect to u, then $\mathcal{C}$ has C_4 induced which contradicts our hypothesis. So assume v connects to u and to at least one simplex. Since u is a cone vertex in $\mathcal{C} \setminus v$, $\text{link}_u(\mathcal{C})$ has no ghost vertices and by hypothesis, it is nuclear. Furthermore, $\text{link}_u(\mathcal{C} \setminus v) = \Delta_m \sqcup \Delta_n$ and so $\text{link}_u(\mathcal{C})$ is either $\Delta_{m+1} \sqcup \Delta_n$ or $\text{cone}^1(\Delta_m \sqcup \Delta_n)$ (note that these are the only nuclear complexes that become disconnected upon removing a vertex). In the first case, $\mathcal{C} = \text{cone}^1(\Delta_{m+1} \sqcup \Delta_n)$ and in the second case $\mathcal{C} = \text{cone}^2(\Delta_m \sqcup \Delta_n)$.

Now assume $p \geq 2$. Let $u_1, \ldots, u_p$ denote the cone vertices of $\mathcal{C} \setminus v$. Let M, N denote the vertex sets of the Δ_m, Δ_n respectively. Then v must connect to at least one of M and N to avoid inducing an independent set on three vertices. Furthermore, v must connect to each u_i - if v only connects to one of M or N, then this is required to avoid P_4, and if v connects to both then we need this to ensure that $\mathcal{C}$ is C_4-free. Furthermore, the set $\{u_1, \ldots, u_p, v\}$ is a facet of $\mathcal{C}$. This is clear when $p = 1$ since otherwise v doesn't connect to u_1. Induction on p gives that $\{u_1, \ldots, u_{p-1}, v\}$ is a facet of $\text{link}_{u_p}(\mathcal{C})$ and therefore $\{u_1, \ldots, u_p, v\}$ is a facet of $\mathcal{C}$.

Now, we can see that $\text{link}_{u_p}(\mathcal{C})$ has no ghost vertices and is C_4-free. Furthermore, $\text{link}_{u_p}(\mathcal{C}) \setminus v$ and $\mathcal{C} \setminus \{u_p, v\}$ are both equal to $\text{cone}^{p-1}(\Delta_m \sqcup \Delta_n)$. So by induction on p, each of $\text{link}_{u_p}(\mathcal{C})$ and $(\mathcal{C} \setminus u_p)$ can either be $\text{cone}^p(\Delta_m \sqcup \Delta_n)$ or $\text{cone}^{p-1}(\Delta_{m+1} \sqcup \Delta_n)$. If $\text{link}_{u_p}(\mathcal{C}) = \mathcal{C} \setminus u_p$, then u_p is a cone vertex in $\mathcal{C}$. So if both equal $\text{cone}^p(\Delta_m \sqcup \Delta_n)$ then $\mathcal{C} = \text{cone}^{p+1}(\Delta_m \sqcup \Delta_n)$. If both equal $\text{cone}^{p-1}(\Delta_{m+1} \sqcup \Delta_n)$ then $\mathcal{C} = \text{cone}^p(\Delta_{m+1} \sqcup \Delta_n)$.

Since $\text{link}_{u_p}(\mathcal{C}) \subseteq \mathcal{C} \setminus u_p$, it is impossible to have $\text{link}_{u_p}(\mathcal{C}) = \text{cone}^p(\Delta_m \sqcup \Delta_n)$ and $\mathcal{C} \setminus u_p = \text{cone}^{p-1}(\Delta_{m+1} \sqcup \Delta_n)$. It is also impossible to have $\mathcal{C} \setminus u_p = \text{cone}^p(\Delta_m \sqcup \Delta_n)$ and $\text{link}_{u_p}(\mathcal{C}) = \text{cone}^{p-1}(\Delta_{m+1} \sqcup \Delta_n)$. For if this were the case, then v would be a cone vertex in $\mathcal{C} \setminus u_p$ and it would be part of the Δ_{m+1} in $\text{link}_{u_p}(\mathcal{C})$. Since $n \geq 1$, $\#N \geq 2$, so given $b_1, b_2 \in N$ and $a \in M$ the complex induced on a, b_1, b_2, u_p, v has facets $\{v, a, u_p\}, \{v, b_1, b_2\}, \{u_p, b_1, b_2\}$. This is J_1^*, and so $\mathcal{C}$ is not β-avoiding. $\qquad\square$

Proposition 4.5.5. *Let $\mathcal{C}$ be the simplicial complex on q vertices with a vertex v such that $\text{link}_v(\mathcal{C}) = D_{m,n}$ with $m, n \geq 1$ and $(\mathcal{C} \setminus v)^* = G^{m+n-m'-n'}(\Delta_{m'} \sqcup \Delta_n)$ with $-1 \leq m' < m$ and $0 \leq n' \leq n$. Unless $m' = -1$ and $n' = n$, $\mathcal{C}$ has a proper minor that is not nuclear.*

Proof. Let M, N denote the sets of vertices such that in $(\text{link}_v(\mathcal{C}))^*$, the complex induced on M is Δ_m and the complex induced on N is Δ_n. Then some vertex $x \in N$ is a non-ghost in $(\mathcal{C} \setminus v)^*$. We show that $\text{link}_x(\mathcal{C})$ is not nuclear.

First we describe the facets of $\text{link}_x(\mathcal{C})$. We claim that the facets of $\text{link}_x(\mathcal{C})$ are precisely the sets containing $N - 1 - 2$ vertices, where the omitted vertex pair is one of

the following

1. some $a \in M$ and some $b \in N \setminus \{x\}$ or

2. v and some vertex g that is a ghost in $(\mathcal{C} \setminus v)^*$.

We now prove the claim. The facets of $\mathrm{link}_x(\mathcal{C})$ that contain v are the sets whose complements are pairs of the first type since $\mathrm{link}_v(\mathcal{C}) = D_{m,n}$. The facets of $\mathrm{link}_x(\mathcal{C})$ that do not contain v are the sets whose complements are minimal non-faces of $(\mathcal{C} \setminus v)^*$ that do not become faces of $\mathrm{link}_x(\mathcal{C})$ upon adding v. These minimal non-faces of $(\mathcal{C} \setminus v)^*$ are the ghost vertices. Therefore the facets of $\mathrm{link}_x(\mathcal{C})$ that do not contain v are the sets of the second type.

We now show that $\mathrm{link}_x(\mathcal{C})$ is not nuclear. The complex $\mathrm{link}_x(\mathcal{C})$ cannot be $\Lambda(\mathcal{D})$ for any complex $\mathcal{D}$ since no facet lacks fewer than two vertices. Nor can $\mathrm{link}_x(\mathcal{C})$ have any ghost vertices - $m \geq 1$ implies $\#M \geq 2$ and so every vertex in M is in a facet of the first type, every vertex of N is in a facet of the second type, and v is in every facet of the first type. It is clear that no vertex is in every facet, so $\mathrm{link}_x(\mathcal{C})$ cannot be $\mathrm{cone}^p(D_{k,l})$ nor $\mathrm{cone}^p(\Delta_k \sqcup \Delta_l)$ for any $p \geq 1$. It is clear that $\mathrm{link}_x(\mathcal{C})$ has more than 2 facets, so $\mathrm{link}_x(\mathcal{C})$ cannot be $\Delta_k \sqcup \Delta_l$, nor Δ_k.

The only case left to check is $\mathrm{link}_x(\mathcal{C}) = D_{k,l}$. If this is the case then we can partition the vertices of $\mathrm{link}_x(\mathcal{C})$ into disjoint sets C_1, C_2 such that the facets of $\mathrm{link}_x(\mathcal{C})$ are precisely the subsets of $C_1 \cup C_2$ that omit exactly one from each C_1 and C_2.

Assume $n' = n$ and $m' \geq 0$. Let $g \in M$ be a ghost vertex and without loss of generality assume $g \in C_1$. Facets of the first type imply $M \subseteq C_1$. Facets of the second type imply $v \in C_2$. Since $m' \geq 0$, there exists some $a \in M$ that is not a ghost in $(\mathcal{C} \setminus v)^*$. Snd so $(C_1 \cup C_2) \setminus \{a, v\}$ is a facet of $\mathrm{link}_x(\mathcal{C})$. But this does not fit our earlier description of the facets of $\mathrm{link}_x(\mathcal{C})$.

Now assume $-1 \leq m' < m$ and $0 \leq n' < n$. Then there exist $g' \in M$ and $g'' \in N$ that are ghosts in $(\mathcal{C} \setminus v)^*$. But this contradicts our earlier description of the facets of $\mathrm{link}_x(\mathcal{C})$ - facets of the first kind imply that g', g'' are in different C_is and facets of the second kind imply that they are in the same C_i. $\qquad \square$

Proposition 4.5.6. *Let $\mathcal{C}$ be a simplicial complex on q vertices. Let v be a vertex in $\mathcal{C}$ such that $\mathrm{link}_v(\mathcal{C}) = D_{m,n}$ with $m, n \geq 1$. Assume $(\mathcal{C} \setminus v)^* = G^{m+1}(\Delta_n)$. Then $\mathcal{C} = D_{m,n+1}$.*

Proof. Let M, N denote the sets of vertices such that in $(\mathrm{link}_v(\mathcal{C}))^*$, the complex induced on M is Δ_m and the complex induced on N is Δ_n. Then all the vertices of M and none

of the vertices of N are ghosts in $(\mathcal{C} \setminus v)^*$. We claim that each facet of $\mathcal{C}$ contains $N - 2$ vertices, and the omitted pair is one of the following

1. some $a \in M$ and some $b \in N$
2. v and some $a \in M$.

To see this, note that if a facet contains v, then it must be of the first form since $\mathrm{link}_v(\mathcal{C}) = D_{m,n}$. If a facet does not contain v, then its complement must be v, along with a minimal non-face of $(\mathcal{C} \setminus v)^*$. These minimal non faces are precisely the vertices in M. So $\mathcal{C} = D_{m,n+1}$ where M and $N \cup \{v\}$ are the sets of vertices for the Δ_m and Δ_{n+1} in $\mathcal{C}^*$ respectively. $\qquad\square$

Proposition 4.5.7. *Let $\mathcal{C}$ be a simplicial complex on q vertices. Let v be a vertex in $\mathcal{C}$ such that $\mathrm{link}_v(\mathcal{C}) = D_{m,n}$ with $m, n \geq 1$ and has no ghost vertices. Let M, N denote the sets of vertices such that in $(\mathrm{link}_v(\mathcal{C}))^*$, the complex induced on M is Δ_m and the complex induced on N is Δ_n. Assume $(\mathcal{C} \setminus v)^* = G^{m+1}(\mathcal{D})$ where $\mathcal{D}$ is a nuclear complex on vertex set N with more than one facet. Then $\mathcal{C}$ has a proper minor that is not nuclear.*

Proof. Let $x \in M$. We show that $\mathrm{link}_x(\mathcal{C})$ is not nuclear. We start by showing it is not Lawrence type. Let F be a facet of $\mathrm{link}_x(\mathcal{C})$ that does not contain v. Then $F \cup \{x\}$ is a facet of $\mathcal{C} \setminus v$ so it must also lack some $a \in M \setminus \{x\}$. If F is a facet of $\mathrm{link}_x(\mathcal{C})$ that contains v, then $F \cup \{x\} \setminus \{v\}$ is a facet of $\mathrm{link}_v(\mathcal{C})$ and therefore lacks some $a \in M$ and $b \in N$. So $\mathrm{link}_x(\mathcal{C})$ is not $\Lambda(\mathcal{F})$ for any complex $\mathcal{F}$.

We can see that $\mathrm{link}_x(\mathcal{C})$ has more than two facets - since $m, n \geq 1$, there are at least two that include v and since $(C \setminus v)^* = G^{m+1}(\mathcal{D})$, there is at least one that does not. $\mathrm{link}_x(\mathcal{C})$ cannot be $\mathrm{cone}^p(\Delta_k \sqcup \Delta_l)$ for any $p \geq 0$.

All of the vertices of M are ghost vertices in $(\mathcal{C} \setminus v)^*$, and $(\mathcal{C} \setminus v)^*$ has a minimal non-face $S \subseteq N$ with at least two vertices. This implies that $\mathcal{C}$ contains a facet F such that $M \subseteq F$, there exist $b_1, b_2 \in N$ such that $b_1, b_2 \notin F$ and $v \notin F$ (since $\mathrm{link}_v(\mathcal{C}) = D_{m,n}$). Facets in $\mathrm{link}_x(\mathcal{C})$ that contain v have $N - 1 - 2$ vertices - they lack some $a \in M \setminus \{x\}$ and some $b \in N$. Furthermore, $F \setminus \{x\}$ is a facet of $\mathrm{link}_x(\mathcal{C})$ containing at most $N - 1 - 3$ vertices. This means that $\mathrm{link}_x(\mathcal{C})$ is not pure, and therefore not $\mathrm{cone}^p(D_{k,l})$ for any $p \geq 0$. So $\mathrm{link}_x(\mathcal{C})$ is a proper minor of $\mathcal{C}$ that is not nuclear. $\qquad\square$

Lemma 4.5.8. *Let $\mathcal{C}$ be a β-avoiding simplicial complex on q vertices. Assume all proper minors of $\mathcal{C}$ are nuclear. If there exists a vertex v of $\mathcal{C}$ such that $\mathrm{link}_v(\mathcal{C}) = \mathrm{cone}^p(D_{m,n})$ for $p \geq 0$ and $m, n \geq 1$, then $\mathcal{C}$ is nuclear.*

Proof. We split this into two cases. For the first case, assume $p \geq 1$. Here, Propositions 4.2.9 and 4.2.18 give $\mathcal{C}^* \setminus v = \text{cone}^p(\Delta_m \sqcup \Delta_n)$. By Proposition 4.3.3, $\mathcal{C}^*$ is β-avoiding so in $\mathcal{C}^*$, v must connect to one of the simplices in the $\Delta_m \sqcup \Delta_n \subset \mathcal{C}^*$ to avoid inducing an independent set on three vertices. Now let u be a cone vertex. If v connects to only one such simplex, then v must also connect to u to avoid inducing P_4. If v connects to both such simplices, and v does not connect to u, then $\mathcal{C}^*$ has C_4 induced. In this case Proposition 4.3.4 implies $\mathcal{C}^* = \text{cone}^k(C_4)$ and so $\mathcal{C} = \text{cone}^k(\Delta_1 \sqcup \Delta_1)$. So we can assume v connects to at least one simplex, and to all cone vertices in $\mathcal{C}^* \setminus v$. From this we can see that $\mathcal{C}^*$ has no induced C_4. Since every proper minor of $\mathcal{C}^*$ is nuclear, Proposition 4.5.4 implies $\mathcal{C}^* = \text{cone}^{p+1}(\Delta_m \sqcup \Delta_n)$ or $\mathcal{C}^* = \text{cone}^p(\Delta_{m+1} \sqcup \Delta_n)$. Proposition 4.5.2 then implies that $\mathcal{C}$ is nuclear.

For the second case, assume $p = 0$. So $\text{link}_v(\mathcal{C}) = D_{m,n}$ with $m, n \geq 1$. Then $D_{m,n} \subseteq \mathcal{C} \setminus v$ and so $(\mathcal{C} \setminus v)^* \subseteq \Delta_m \sqcup \Delta_n$. Let M, N denote the sets of vertices of $\mathcal{C}$ such that in $(\text{link}_v(\mathcal{C}))^*$, the vertices of the Δ_m are M, and the vertices of the Δ_n are N. The non-ghost vertices of $(\mathcal{C} \setminus v)^*$ must form a nuclear complex, and the only disconnected nuclear complexes are of the form $\Delta_s \sqcup \Delta_t$. From this it follows that $(\mathcal{C} \setminus v)^*$ must be one of the following forms (without loss of generality)

1. $\Delta_m \sqcup \Delta_n$
2. $G^{m+n-m'-n'}(\Delta_{m'} \sqcup \Delta_{n'})$ with $-1 \leq m' < m$ and $0 \leq n' \leq n$, but not both $m' = -1$ and $n' = n$
3. $G^{m+1}(\Delta_n)$
4. $G^{m+1}(\mathcal{D})$ where $\mathcal{D}$ is a nuclear complex on N with more than one facet
5. $G^{m+n+2}(\{\emptyset\})$
6. $G^{m+n+2}(\{\})$.

We handle each possibility separately.

Case 1. Assume $(\mathcal{C} \setminus v)^* = \Delta_m \sqcup \Delta_n$. Then $(\mathcal{C} \setminus v)^*$ has no ghost vertices, and so $\mathcal{C} \setminus v$ is not of Lawrence type. So, by Lemma 4.3.7, $\mathcal{C} \setminus v = D_{m,n}$ and so $\mathcal{C} = \text{cone}^1(D_{m,n})$ with v as the cone point.

Case 2. Assume $(\mathcal{C} \setminus v)^* = G^{m+n-m'-n'}(\Delta_{m'} \sqcup \Delta_{n'})$ with $-1 \leq m' < m$ and $0 \leq n' \leq n$ but not both $m' = -1$ and $n' = n$. By Proposition 4.5.5, $\mathcal{C}$ has a proper minor M that is not nuclear. By induction, M is not β-avoiding and so neither is $\mathcal{C}$.

Case 3. Assume $(\mathcal{C} \setminus v)^* = G^{m+1}(\Delta_n)$. Then Proposition 4.5.6 implies that $\mathcal{C} = D_{m,n+1}$.

Case 4. Assume $(\mathcal{C} \setminus v)^* = G^{m+1}(\mathcal{D})$ where $\mathcal{D}$ is a nuclear complex on N with more than one facet. By Proposition 4.5.7, $\mathcal{C}$ has a proper minor M that is not nuclear. By induction, M is not β-avoiding and so neither is $\mathcal{C}$.

Case 5. Assume $(\mathcal{C} \setminus v)^* = G^{m+n+2}(\{\emptyset\})$. Then $\mathcal{C} \setminus v = \partial \Delta_{m+n+1}$. But $\mathrm{link}_v \mathcal{C} = D_{m,n}$ and so v is not in any $m + n - 1$ simplex with the vertices of M and N. By Proposition 4.3.6, $\mathcal{C}$ is not β-avoiding.

Case 6. Assume $(\mathcal{C} \setminus v)^* = G^{m+n+2}(\{\})$. Then $\mathcal{C} \setminus v = \Delta_{m+n+1}$ and so $\mathcal{C} = \Lambda(D_{m,n})$ with v as the added vertex. $\qquad\qquad\square$

Now we have all the necessary tools to prove Theorem 4.5.3.

Proof of Theorem 4.5.3. By Proposition 4.2.20, (1) implies (2). Each nuclear complex can be obtained by applying the unimodularity-preserving operations from Section 4.2 to $\Delta_n \sqcup \Delta_m$, which is unimodular by Proposition 4.2.21. Therefore (3) implies (1). We now show (2) implies (3).

Let $\mathcal{C}$ be a β-avoiding complex on q vertices. We show that $\mathcal{C}$ is nuclear by induction on q. For the base case, note that all simplicial complexes on 2 or fewer vertices are both β-avoiding and nuclear. We may assume that $\mathcal{C}$ has no ghost vertices since adding and removing ghost vertices does not affect the properties of being nuclear nor β-avoiding. In light of Propositions 4.3.4 and 4.4.4 and Lemma 4.4.3, we only need to consider the cases where the 1-skeleton of $\mathcal{C}$ is K_q, or the union of two complete graphs, glued along a common *nonempty* clique. Therefore, we can choose a vertex v of $\mathcal{C}$ that is in a facet with every other vertex. Since $\mathcal{C}$ is β-avoiding, so is $\mathrm{link}_v(\mathcal{C})$. By induction, $\mathrm{link}_v(\mathcal{C})$ is also nuclear. There are five main cases.

Case 1. This is the case where $\mathrm{link}_v(\mathcal{C}) = \Lambda(\mathcal{D})$ for a nuclear $\mathcal{D}$. Let F denote the big facet of $\mathcal{C}$ and let u be the vertex of $\mathcal{C}$ that is not in F. Then $\mathcal{C} = \Lambda(\mathcal{C} \setminus u)$, and $\mathcal{C} \setminus u$ is nuclear by induction.

Case 2. This is the case where $\mathrm{link}_v(\mathcal{C})$ has ghost vertices. Since v connects to all other vertices, this is impossible.

Case 3. This is the case where $\mathrm{link}_v(\mathcal{C}) = \mathrm{cone}^p(\Delta_m \sqcup \Delta_n)$ with $p, m, n \geq 0$. We split this into the sub cases, the first where $m, n \geq 1$ and the second where $m = 0$.

Case 3.1 This is the sub case where $m, n \geq 1$. Proposition 4.2.9 gives $\mathcal{C}^* \setminus v = \mathrm{cone}^p(D_{m,n})$ and so $\mathcal{C}^*$ has $D_{m,n}$ induced. Proposition 4.3.5 then implies that $\mathcal{C}^*$ is either $\mathrm{cone}^{p+1}(D_{m,n})$, or $\mathrm{cone}^p(D_{m,n})$ with v as a ghost vertex. In the first case, $\mathcal{C} = \mathrm{cone}^{p+1}(\Delta_m \sqcup \Delta_n)$. In the second case, $\mathcal{C} = \Lambda(\mathrm{cone}^p(\Delta_m \sqcup \Delta_n))$.

Case 3.2 This is the sub case where $m = 0$. Proposition 4.2.9 gives $\mathcal{C}^* \setminus v = \text{cone}^p(D_{0,n})$ which has a ghost vertex u. Then $\mathcal{C}^* \setminus u$ is nuclear by induction and since u is still a ghost vertex in $\mathcal{C}^*$, this implies $\mathcal{C}^*$ is nuclear. Proposition 4.5.2 implies that $\mathcal{C}$ is nuclear.

Case 4. This is the case where $\text{link}_v(\mathcal{C}) = \text{cone}^p(D_{m,n})$ with $p \geq 0$ and $m, n \geq 1$. By induction, all proper minors of $\mathcal{C}$ are nuclear. Then Lemma 4.5.8 implies that $\mathcal{C}$ is nuclear.

Case 5. This is the case where $\text{link}_v(\mathcal{C}) = \Delta_k$ with $k \geq -2$. Since $\text{link}_v(\mathcal{C})$ has no ghost vertices $k = q - 2$. So $\mathcal{C} = \Delta_{q-1}$. $\qquad \square$

4.6 Operations on non-binary HM pairs

Proposition 4.6.1. *Assume* $(\mathcal{C}, \mathbf{d})$ *is unimodular. Then for all* $\mathbf{d}'$ *such that* $\mathbf{d}' \leq \mathbf{d}$ *componentwise,* $(\mathcal{C}, \mathbf{d}')$ *is unimodular.*

Proof. Note that $\mathcal{A}_{\mathcal{C}, \mathbf{d}'}$ can be realized as a subset of columns of $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$. $\qquad \square$

The special case where $\mathbf{d}' = \mathbf{2}$ gives us the following useful corollary.

Corollary 4.6.2. *If* $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ *is unimodular, then* $\mathcal{C}$ *is nuclear.*

So in order to classify all unimodular hierarchical models we only need to consider nuclear complexes. Therefore we can approach the general classification problem by looking at each nuclear $\mathcal{C}$ and identifying the minimal values of $\mathbf{d}$ that give rise to non-unimodular $(\mathcal{C}, \mathbf{d})$. Before proceeding with this, we show that the class of unimodular $(\mathcal{C}, \mathbf{d})$ is still closed under taking minors of $\mathcal{C}$. However Corollary 4.2.7 fails in the non-binary case; we cannot freely take the Alexander dual of $\mathcal{C}$. Because of this, our proof of Corollary 4.2.10 is not valid in the non-binary case. We will give an alternate proof that the class of unimodular $(\mathcal{C}, \mathbf{d})$ is closed under taking links in $\mathcal{C}$. We start with a useful proposition.

Proposition 4.6.3. *Let* $A \in \mathbb{R}^{r \times n}$ *be a unimodular matrix with columns* $\{\mathbf{a}_i\}_{i=1}^n$. *Assume* $\mathbf{a}_n$ *is nonzero. Let* A' *be the matrix that results when we project* A *onto the hyperplane orthogonal to* $\mathbf{a}_n$. *Then* A' *is unimodular.*

Proof. We may assume A has rank r since otherwise we can delete unnecessary rows. Let $B \in \mathbb{R}^{n \times (n-r)}$ have rank $n - r$ such that $AB^T = 0$. Denote the columns of B as $\{\mathbf{b}_i\}_{i=1}^n$.

By Proposition 4.2.5, B is also unimodular. Denote the columns of A' as $\{\mathbf{a}'_i\}_{i=1}^{n-1}$. Then

$$\mathbf{a}'_i = \mathbf{a}_i - \frac{\langle \mathbf{a}_i, \mathbf{a}_n \rangle}{\|\mathbf{a}_n\|^2} \mathbf{a}_n.$$

Let B' be the matrix with columns $\{b_i\}_{i=1}^{n-1}$. Then B' has rank $n-r$ or $n-r-1$. The second case implies that $\mathbf{b}_n$ is a coloop in the matroid underlying B. Since the matroids underlying A and B are duals, this would imply that $\mathbf{a}_n = 0$. But then $A' = A$ which we assumed to be unimodular. So we can assume that $\text{rank}(B') = \text{rank}(B) = n-r$ and that $\mathbf{a}_n \neq 0$. Therefore $\text{rank}(A') = r-1$ and so the dimension of its kernel is $n-r$, which is the rank of B'. We claim that $A'(B')^T = 0$. From this it follows by Proposition 4.2.5 that A' is unimodular.

Now we prove the claim. Let $\mathbf{b}_i$ be a column of B' (so $1 \leq i \leq n-1$). Letting b_{ji} denote the jth entry of $\mathbf{b}_i$, we have

$$A'\mathbf{b}_i^T = \sum_{j=1}^{n-1} b_{ji} \left(\mathbf{a_j} - \frac{\langle \mathbf{a}_i, \mathbf{a}_n \rangle}{\|\mathbf{a}_n\|^2} \mathbf{a}_n \right).$$

Note that $b_{jn}\left(\mathbf{a}_n - \frac{\langle \mathbf{a}_n, \mathbf{a}_n \rangle}{\|\mathbf{a}_n\|^2} \mathbf{a}_n \right) = 0$, so we can add the nth term to the above sum. This enables us to break it up as follows

$$\sum_{j=1}^{n} b_{ji}\mathbf{a}_j - \frac{1}{\|\mathbf{a}_n\|^2} \sum_{j=1}^{n} b_{ji}\langle \mathbf{a}_j, \mathbf{a}_n \rangle \mathbf{a}_n.$$

The first term is $A\mathbf{b}_i^T = 0$. The second term is $\left(\frac{1}{\|\mathbf{a}_n\|^2} \mathbf{a}_n^T A \mathbf{b}_i^T \right) \cdot \mathbf{a}_n = 0$. So the claim is proven. $\qquad \square$

Corollary 4.6.4. *Assume $(\mathcal{C}, \mathbf{d})$ is unimodular. Let v be a vertex of $\mathcal{C}$ and let $\mathbf{d}'$ denote the vector obtained by deleting the entry for v from $\mathbf{d}$. Then $(\text{link}_v \mathcal{C}, \mathbf{d}')$ is unimodular.*

Proof. Let $A = \mathcal{A}_{\text{link}_v \mathcal{C}, \mathbf{d}'}$ and let $B = \mathcal{A}_{\mathcal{C} \setminus v, \mathbf{d}'}$. By the remark following Lemma 2.2 in

[34] and the lemma itself, we can write

$$
\mathcal{A}_{\mathcal{C},\mathbf{d}} = \begin{pmatrix}
A & \mathbf{0} & \dots & \mathbf{0} \\
\mathbf{0} & A & \dots & \mathbf{0} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0} & \mathbf{0} & \dots & A \\
B & B & \dots & B
\end{pmatrix}.
$$

Assume $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{l \times n}$ and so $\mathcal{A}_{\mathcal{C},\mathbf{d}} \in \mathbb{R}^{dm+l,dn}$. By Proposition 1.6.12((3)), we may remove rows from $\mathcal{A}$ to make $A, B, \mathcal{A}$ all have full row rank. So assume that they do. Let $\mathcal{A}'$ denote the matrix that results when we project $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ onto the subspace orthogonal to the last $(d-1)n$ columns. Then if $1 \le i \le n$, the ith column of $\mathcal{A}'$ can be expressed as the ith column of $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ minus a linear combination of the last $(d-1)n$ columns of $\mathcal{A}_{\mathcal{C},\mathbf{d}}$, all of which are 0 in the top m rows. So this means that the top m rows of $\mathcal{A}'$ are

$$
\begin{pmatrix} A & \mathbf{0} & \dots & \mathbf{0} \end{pmatrix}.
$$

Furthermore, since A and $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ both have full row rank, the final $(d-1)n$ columns of $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ have rank $(d-1)m+l$. Therefore $\mathcal{A}'$ has rank m. Since A also has rank m, we may delete the bottom $(d-1)m+l$ rows of $\mathcal{A}'$ without affecting the rowspace and therefore unimodularity. So the matrix $\begin{pmatrix} A & \mathbf{0} & \dots & \mathbf{0} \end{pmatrix}$, and therefore A, is unimodular. □

Our proof of Proposition 4.2.1 is still valid in the non-binary case. Therefore, we have the following.

Proposition 4.6.5. *Assume $\mathcal{C}'$ is a minor of $\mathcal{C}$. If $(\mathcal{C}, \mathbf{d})$ is unimodular, then so is $(\mathcal{C}', \mathbf{d}')$ where $\mathbf{d}'$ is the restriction of $\mathbf{d}$ to the vertices that are in $\mathcal{C}'$.*

An HM pair $(\mathcal{D}, \mathbf{d}')$ is a *minor* of $(\mathcal{C}, \mathbf{d})$ if $\mathcal{D}$ can be obtained from $\mathcal{C}$ via a (possibly empty) sequence of vertex deletions and vertex links and $d'_v \le d_v$ for every vertex of $\mathcal{D}$. The pair $(\mathcal{C}, \mathbf{d})$ is said to be *unimodular* if the matrix $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ is unimodular, and *binary* if $\mathbf{d} = \mathbf{2}$. In view of Proposition 4.6.6 below, we say that an HM pair is *minimally nonunimodular* if it is not unimodular but every minor is.

Proposition 4.6.6. *Minors of unimodular HM pairs are unimodular.*

Proof. This follows from Propositions 4.6.1 and 4.6.5. □

Proposition 4.6.6 tells us how we can create a "smaller" unimodular HM pair from a larger one. We now discuss how to go the other way. Given a simplicial complex $\mathcal{C}$, we say $v \in \operatorname{ground}(\mathcal{C})$ is

- a *cone vertex* if it appears in every facet of $\mathcal{C}$,
- a *ghost vertex* if it does not appear in any facet of $\mathcal{C}$, and
- and a *Lawrence vertex* if its complement in the ground set of $\mathcal{C}$ is a facet.

We denote by $\operatorname{cone}(\mathcal{C})$ and $G\mathcal{C}$ the complex obtained by adding a cone vertex and ghost vertex to $\mathcal{C}$, respectively, and we denote by $\Lambda\mathcal{C}$ the complex obtained by adding a Lawrence vertex v to $\mathcal{C}$ such that $\operatorname{link}_v(\mathcal{C}) = \mathcal{C}$. Iterated application of each aforementioned operation will be denoted by superscript. For example, $G^5\mathcal{C}$ denotes the complex obtained by adding five ghost vertices to $\mathcal{C}$. The proofs of Propositions 4.2.11, 4.2.17 generalize to the non-binary setting and so we have the following.

Proposition 4.6.7. *If the pair $(\mathcal{C}, \mathbf{d})$ gives rise to unimodular $\mathcal{A}_{\mathcal{C},\mathbf{d}}$, so do $(\operatorname{cone}^p(\mathcal{C}), \mathbf{d}')$ and $(G(\mathcal{C}), \mathbf{d}'')$ where $\mathbf{d}' = \begin{pmatrix} \mathbf{d} & c_1 & \cdots & c_p \end{pmatrix}$ and $\mathbf{d}'' = \begin{pmatrix} \mathbf{d} & c \end{pmatrix}$ for any $c, c_1, \ldots, c_p \geq 2$.*

Proof. In the case of a cone, the matrix $\mathcal{A}_{\operatorname{cone}^p(\mathcal{C}),\mathbf{d}'}$ is a block diagonal matrix:

$$
\mathcal{A}_{\operatorname{cone}^p(\mathcal{C}),\mathbf{d}'} = \begin{pmatrix} \mathcal{A}_{\mathcal{C},\mathbf{d}} & 0 & \cdots & 0 \\ 0 & \mathcal{A}_{\mathcal{C},\mathbf{d}} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathcal{A}_{\mathcal{C},\mathbf{d}} \end{pmatrix},
$$

with $c_1 c_2 \cdots c_p$ blocks down the diagonal. In the case of adding a ghost vertex, the matrix $\mathcal{A}_{G(\mathcal{C}),\mathbf{d}'}$ has repeated columns:

$$
\mathcal{A}_{G(\mathcal{C}),\mathbf{d}''} = \begin{pmatrix} \mathcal{A}_{\mathcal{C},\mathbf{d}} & \mathcal{A}_{\mathcal{C},\mathbf{d}} & \cdots & \mathcal{A}_{\mathcal{C},\mathbf{d}} \end{pmatrix}
$$

with c blocks. $\qquad\square$

We can also extend Proposition 4.2.21.

Proposition 4.6.8. *Let $\mathcal{C} = \Delta_m \sqcup \Delta_n$ and $\mathbf{d} \in \mathbb{Z}_{\geq 2}^{m+n+2}$. Let M, N denote the vertex sets of $\mathcal{C}$ in the Δ_m, Δ_n respectively. Let $\mathcal{D}$ be the complex with facets $\{1, 2\}$ and let $\mathbf{e} = (e_1, e_2)$ where*

$$
e_1 = \prod_{v \in M} d_v \qquad e_2 = \prod_{v \in N} d_v.
$$

Then $\mathcal{A}_{\mathcal{C},\mathbf{d}} = \mathcal{A}_{\mathcal{D},\mathbf{e}}$. *This matrix is unimodular.*

Proof. We can see that $\mathcal{A}_{\mathcal{C},\mathbf{d}} = \mathcal{A}_{\mathcal{D},\mathbf{e}}$ by inspection. The matrix $\mathcal{A}_{\mathcal{D},(e_1,e_2)}$ is the vertex edge incidence matrix of a complete bipartite graph with e_1 and e_2 vertices in each set of the partition. Such vertex-edge incidence matrices are examples of network matrices and are hence totally unimodular [51, Ch. 19]. □

We also note that Proposition 4.2.15 holds in a slightly more general setting.

Proposition 4.6.9. *Let* $(\mathcal{C}, \mathbf{d})$ *be such that* $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ *is unimodular. Then* $\mathcal{A}_{\Lambda(\mathcal{C}),\mathbf{d}'}$ *is also unimodular if* $\mathbf{d}' = \begin{pmatrix} \mathbf{d} & 2 \end{pmatrix}$.

Proof. The proof here is similiar to the proof for Proposition 4.2.15. Note that $\Lambda(\mathcal{A}_{\mathcal{C},\mathbf{d}}) = \mathcal{A}_{\Lambda(\mathcal{C}),\mathbf{d}'}$. Hence Theorem 7.1 in [57] implies the proposition. □

4.7   Minimally non-unimodular HM pairs

We say that an HM pair $(\mathcal{C}, \mathbf{d})$ is *minimally nonunimodular* if $(\mathcal{C}, \mathbf{d})$ is not unimodular, but every proper minor is. This section proves that a particular list of HM pairs are minimally nonunimodular. We will later see that this list is exhaustive. Let us begin with a useful observation.

Remark 4.7.1. *Different HM pairs may yield the same matrix. Let* $\mathcal{C}$ *be a complex on ground set* V. *Assume* $\mathcal{C}$ *has a face* E *such that any for any facet* F, $E \cap F \neq \emptyset$ *implies* $E \subseteq F$. *Let* $\mathcal{C}'$ *denote the complex on ground set* $V' := (V \cup \{v_0\}) \setminus E$ *with facets*

$$\text{facet}(\mathcal{C}') = \{F \in \text{facet}(\mathcal{C}) : F \cap E = \emptyset\} \cup \{F \cup \{v_0\} : E \subset F\}.$$

Let $\mathbf{d} \in \mathbb{Z}_{\geq 2}^V$ *and define* $\mathbf{d}' \in \mathbb{Z}_{\geq 2}^{V'}$ *such that* $\mathbf{d}'_v = \mathbf{d}_v$ *for all* $v \in V \cap V'$ *and* $\mathbf{d}'_{v_0} = \prod_{v \in F} \mathbf{d}_v$. *Then* $\ker_{\mathbb{Z}} \mathcal{A}_{\mathcal{C},\mathbf{d}} = \ker_{\mathbb{Z}} \mathcal{A}_{\mathcal{C}',\mathbf{d}'}$.

We are now ready to list the minimally nonunimodular HM pairs.

Proposition 4.7.2. *The following HM pairs are minimally nonunimodular:*

(1) *any of the complexes listed in Proposition 4.3.1 with* $\mathbf{d} = \mathbf{2}$,
(2) $\Lambda(\Delta_0 \sqcup \Delta_0)$ *with facets* $\{12, 23, 13\}$ *and* $\mathbf{d} = (3, 3, 3)$,
(3) $\Lambda(\Delta_1 \sqcup \Delta_1)$ *with facets* $\{125, 345, 1234\}$ *and* $\mathbf{d} = (2, 2, 2, 2, 3)$,

(4) $\Lambda(\Delta_1 \sqcup \Delta_0)$ *with facets* $\{124, 34, 123\}$ *and* $\mathbf{d} = (2, 2, 3, 3)$,

(5) $D_{1,1}$ *with facets* $\{12, 23, 34, 14\}$ *and* $\mathbf{d} = (2, 2, 2, 3)$,

(6) $\Lambda G(\Delta_0 \sqcup \Delta_0)$ *with facets* $\{12, 13, 234\}$ *and* $\mathbf{d} = (4, 2, 2, 2)$,

(7) $\Lambda D_{1,1}$ *with facets* $\{1234, 125, 235, 345, 145\}$ *and* $\mathbf{d} = (2, 2, 2, 2, 3)$, *and*

(8) $\Lambda \Lambda G(\Delta_0 \sqcup \Delta_0)$ *with facets* $\{1234, 1235, 145, 245\}$ *and* $\mathbf{d} = (2, 2, 2, 3, 3)$.

Proof. All minors of the given HM pairs are easily seen to be unimodular by building them up from $(\Delta_m \sqcup \Delta_n, (2, \ldots, 2))$ via the operations that are known to preserve unimodularity. Hence it suffices to show that each HM pair is not unimodular.

Proposition 4.3.1 verifies nonunimodularity of the complexes implied by (1). Macaulay2 scripts using 4ti2 [2] to verify nonunimodularity of (5),(7), and (8) can be found on my website [9]. In light of Remark 4.7.1, we see that the matrix for (3) has the same integer kernel as the matrix for the HM pair $(\{12, 23, 13\}, (4, 4, 3))$ and the matrix for (4) has the same integer kernel as the matrix for the HM pair $(\{12, 23, 13), (4, 3, 3)\}$. Both have (2) as a minor, the minimal nonunimodularity of which is given in [48, Table 1]. Remark 4.7.1 can also be applied to see that the matrix for (6) has the same integer kernel as the matrix for $(J_1^*, (2, \ldots, 2))$ where J_1^* is a complex shown to be minimally nonunimodular in Proposition 4.3.1. $\qquad\square$

Before we can prove our complete classification, we will need the unimodularity preserving operation given in the next section. However, in the special case that $\mathcal{C}$ is nuclear with nucleus $D_{m,n}$, we are ready to classify the unimodular HM pairs $(\mathcal{C}, \mathbf{d})$.

Proposition 4.7.3. *Let $\mathcal{C}$ be a nuclear complex with nucleus $D_{m,n}$ with $m, n \geq 1$. Then $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ is unimodular if and only if*

1. *$d_v = 2$ for all v in the original $D_{m,n}$ and*
2. *$d_v = 2$ for all v corresponding to a Lawrence lifting.*

Proof. The fact that all such $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ are unimodular follows from applying Propositions 4.6.7 and 4.6.9. We now show that these are the only unimodular $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ with this type of complex $\mathcal{C}$.

Note that since taking cone vertices commutes with adding ghost vertices and Lawrence liftings and this always preserves unimodularity, we can assume there are no cone vertices. Now if $\mathcal{C}$ is a nuclear complex obtained by successively adding ghost vertices and Lawrence vertices to $D_{m,n}$, then the induced subcomplex on just the $D_{m,n}$ and Lawrence

vertices is an iterated Lawrence lifting of $D_{m,n}$. If any of the vertices of v that are Lawrence vertices have $d_v > 2$, then by taking suitable links, one obtains (3) in Proposition 4.7.2. On the other hand, if $d_v > 2$ for some v in the original $D_{m,n}$, then by taking suitable links, one obtains (5) in Proposition 4.7.2. $\qquad \square$

4.8 A new unimodularity-preserving operation

In this section we describe a new unimodularity-preserving operation that is key in proving our classification of the unimodular HM pairs. We also describe how this operation acts on the Graver basis in the case that is relevant to us. The results in this section are about general integer matrices and do not depend on any specific properties of hierarchical models.

Given a matrix $A \in \mathbb{Z}^{d \times n}$, denote by $G_q A$ and $\Lambda_p A$ the matrices

$$G_q A = \begin{pmatrix} A & \dots & A \end{pmatrix} \qquad \text{and} \qquad \Lambda_p A = \begin{pmatrix} A & 0 & \dots & 0 & 0 \\ 0 & A & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & A & 0 \\ I & I & \dots & I & I \end{pmatrix}.$$

Note that $G_q \mathcal{A}_{\mathcal{C},\mathbf{d}} = \mathcal{A}_{G\mathcal{C},(\mathbf{d} \ q)}$ and $\ker_{\mathbb{Z}}(\Lambda_p \mathcal{A}_{\mathcal{C},\mathbf{d}}) = \ker_{\mathbb{Z}}(\mathcal{A}_{\Lambda\mathcal{C},(\mathbf{d} \ p)})$. The operations Λ_2 and G_q for $q \geq 1$ are unimodularity-preserving, in the sense that applying them to a unimodular matrix produces a unimodular matrix (for Λ_2, this follows from [57, Theorem 7.1]). In this section, we add a new unimodularity preserving operation to the list: inserting a ghost vertex operation immediately before a Lawrence lift (Proposition 4.8.1). This operation provides the last crucial step in generalizing Theorem 4.5.3 to arbitrary unimodular HM pairs.

Proposition 4.8.1. *Let $A \in \mathbb{Z}^{d \times n}$ be a matrix and let $p \geq 2$ be a fixed integer. Then $\Lambda_p A$ is unimodular if and only if $\Lambda_p G_q A$ is unimodular for all integers $q \geq 2$.*

Proof. Since unimodularity of A is a property of $\ker_{\mathbb{Z}} A$, we may assume A has full row-rank d. It follows that $\Lambda_p G_q A$ and $\Lambda_p A$ do as well. So, it suffices to show that any maximal square submatrix of $\Lambda_p G_q A$ has determinant ± 1 or 0 whenever any maximal square submatrix of $\Lambda_p A$ does. We proceed by showing that the possible values of a

determinant of such a sub-matrix is independent of q. To this effect, we claim that the absolute value of any such determinant is equal to the absolute value of a determinant of the form

$$
\begin{array}{c}
\\
\\
\\
\\
\{1\} \\
\vdots \\
R \\
S \\
T \\
\vdots \\
[p-1]
\end{array}
\left(
\begin{array}{c|cccc|cccc|c|cccc}
\cdots & A_1^R & 0 & \cdots & 0 & A_1^S & 0 & \cdots & 0 & \cdots & B_1 & 0 & \cdots & 0 \\
\cdots & 0 & A_2^R & \cdots & 0 & 0 & A_2^S & \cdots & 0 & \cdots & 0 & B_2 & \cdots & 0 \\
\cdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \cdots & \vdots & \vdots & \ddots & \vdots \\
\cdots & 0 & 0 & \cdots & A_{p-1}^R & 0 & 0 & \cdots & A_{p-1}^S & \cdots & 0 & 0 & \cdots & B_{p-1} \\
\hline
\cdots & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\
\cdots & 0 & 0 & \cdots & 0 & \vdots & \vdots & \ddots & \vdots & \cdots & 0 & 0 & \cdots & 0 \\
\cdots & I_1^R & I_2^R & \cdots & I_{p-1}^R & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\
\cdots & 0 & 0 & \cdots & 0 & I_1^S & I_2^S & \cdots & I_{p-1}^S & \cdots & 0 & 0 & \cdots & 0 \\
\cdots & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\
\cdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \cdots & \vdots & \vdots & \cdots & \vdots \\
\cdots & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0
\end{array}
\right)
$$

where in the block corresponding to each $S \subseteq [p-1]$, if $i \notin S$ then $A_i^S = I_i^S$ is the unique 0×0 matrix, and if $i \in S$, then A_i^S is some constant matrix A^S and I_i^S is the identity matrix with the same number of columns as A^S. To see this, note that a square column submatrix of $\Lambda_p G_q A$ will have a (possibly empty) submatrix of an identity matrix for its final columns. We can find the determinant by applying Laplace expansion about these columns. Each column of the resulting matrix is obtained by padding a column of A either with all zeros, or all zeros and a single 1. We move all columns of the former description to the rightmost side of the matrix and rearrange them to match the right hand part of the block structure shown above. For the rest of the block structure, note that these remaining columns are naturally partitioned into $p-1$ blocks of the form $\begin{pmatrix} 0 & A_i^T & 0 & I \end{pmatrix}^T$ where $i = 1, \ldots p-1$ and A_i is a column submatrix of A. Each column a_j of A appears in some subset of these blocks which we can index by a subset of $[p-1]$. We can then organize the columns of this submatrix according to this subset indexing and then organize the bottom block of rows to get the desired block structure.

Now if S is a singleton, then every row in the block of rows labeled S has exactly one 1 and all the remaining entries 0. Therefore, we can remove all those blocks using Laplace expansion along these rows. Using the identity matrices in the bottom half of our matrix,

we can apply row operations to turn each row-block of the form $\begin{pmatrix} 0 & \cdots & 0 & A^S \end{pmatrix}$ into $\begin{pmatrix} -A^S & \cdots & -A^S & 0 \end{pmatrix}$. This leaves several columns that have a single 1 and all other entries 0. Applying Laplace expansion along these columns leaves the matrix

$$\begin{pmatrix} \cdots & A^R & \cdots & 0 & A^S & \cdots & 0 & \cdots & B_1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \cdots & 0 & \cdots & A^R & 0 & \cdots & A^S & \cdots & 0 & \cdots & B_{p-2} & 0 \\ \cdots & -A^R & \cdots & -A^R & -A^S & \cdots & -A^S & \cdots & 0 & \cdots & 0 & B_{p-1} \end{pmatrix}.$$

Note that each A^S is either the 0×0 matrix, or a column submatrix of A (possibly with repeated columns). So the set of possible nonzero determinants is independent of q. $\quad\square$

We conclude this section by demonstrating how to recover the Graver basis of $\Lambda_3 G_q A$ from the Graver basis of $\Lambda_3 A$, for use in Remark 4.10.4.

In what follows, we write each $v \in \ker \Lambda_3 G_q A$ in the form $v = (a, b, c)$ with $a, b, c \in \ker G_q A$. Additionally, we let $v_i = (a_i, b_i, c_i) \in \mathbb{Z}^{3q}$ denote the restriction of v to coordinates corresponding to the i-th column of A (in particular, these columns are not sequential above).

Proposition 4.8.2. *Let $A \in \mathbb{Z}^{d \times n}$ be a matrix and fix $q \geq 1$. Assume $\Lambda_3 G_q A$ is unimodular. Then the Graver basis of $\Lambda_3 G_q A$ consists of vectors $v = (a, b, c)$ obtained in the following ways (up to scaling by -1 and permutation of a, b and c):*

(a) *for some $i \leq n$ and $j, k \leq q$, $a_{i,j} = b_{i,k} = 1$ and $a_{i,k} = b_{i,j} = -1$ are the only nonzero entries in v;*

(b) *for every $i \leq n$, each vector a_i, b_i, and c_i has at most 1 nonzero entry, and writing a_i', b_i' and c_i' for the sum of the entries of a_i, b_i, and c_i respectively, the vector $(a_1', \ldots, a_n', b_1', \ldots, b_n', c_1', \ldots, c_n')$ lies in the Graver basis of $\Lambda_3 A$; or*

(c) *for some $i \leq n$ and $j, k \leq q$, $a_{i,j} = 1$, $a_{i,k} = -1$, $b_{i,j} = -1$, and $c_{i,k} = 1$, and the vector $v' = (a', b', c')$ with all coordinates the same as in v, aside from*

$$a_{i,j}' = a_{i,k}' = 0, \quad c_{i,j}' = 1, \quad and \quad c_{i,k}' = 0$$

also lies in the Graver basis of $\Lambda_3 G_q A$.

We clarify the statement of Proposition 4.8.2 with an example before giving the proof.

Example 4.8.3. Consider the matrix

$$A = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix},$$

whose Graver basis consists of $v = \mathbf{e}_1 - \mathbf{e}_2 + \mathbf{e}_3$ and its negative. One can compute the Graver basis of $\Lambda_3 A$ using the software 4ti2 [2] to check unimodularity. It follows from Proposition 4.8.1 that $\Lambda_3 G_q A$ is unimodular for all q. Up to reordering, every vector in the Graver basis of $\Lambda_3 A$ has the form $(v, -v, 0)$. Using Proposition 4.8.2, we can obtain Graver basis vectors for $\Lambda_3 G_3 A$ in the following ways.

- Vectors of the form $(a, -a, 0)$ for some a in the Graver basis of $G_3 A$, yielding type (i) vectors such as $((\mathbf{e}_1, -\mathbf{e}_1, 0), (-\mathbf{e}_1, \mathbf{e}_1, 0), (0, 0, 0))$, and type (ii) vectors such as

$$((\mathbf{e}_1 + \mathbf{e}_3, -\mathbf{e}_2, 0), (-\mathbf{e}_1 - \mathbf{e}_3, \mathbf{e}_2, 0), (0, 0, 0))$$

  obtained by "spreading" the vector v across $G_3 A$. Writing the latter vector in the form (a, b, c) and using the notation introduced above Proposition 4.8.2, we have $a_1 = (1, 0, 0)$, $b_1 = (-1, 0, 0)$ and $c_1 = (0, 0, 0)$, whose entries correspond to the columns of $\Lambda_3 G_3 A$ containing the first column of A.
- Vectors obtained from another Graver basis vector (a, b, c) by (up to reordering of a, b and c) moving a single nonzero entry of a to an unoccupied entry of a corresponding to the same column of A, and then adding a 1 and a -1 to c appropriately. This yields Graver basis vectors such as

$$((\mathbf{e}_1, -\mathbf{e}_1, 0), (-\mathbf{e}_1, 0, \mathbf{e}_1), (0, \mathbf{e}_1, -\mathbf{e}_1)) \quad \text{and}$$

$$((\mathbf{e}_1 + \mathbf{e}_3, -\mathbf{e}_2, 0), (-\mathbf{e}_3, \mathbf{e}_2, -\mathbf{e}_1), (-\mathbf{e}_1, 0, \mathbf{e}_1))$$

  obtained from the each of the vectors above. Notice that vectors obtained in this way are always of type (iii), and that this process applied can be applied to each type (i) and type (ii) vector at most once for each column of A.

Proof of Proposition 4.8.2. Suppose v lies in the Graver basis of $\Lambda_3 G_q A$. Since $\Lambda_3 G_q A$ is unimodular, the first qn rows of $\Lambda_3 G_q A$ ensure that $\{a_{i,j}, b_{i,j}, c_{i,j}\}$ is $\{0, 0, 0\}$ or $\{1, 0, -1\}$ for all $i \leq n$ and $j \leq q$. We first prove that a_i, b_i, and c_i each have no repeated nonzero entries for $i \leq n$. To this end, suppose after appropriately scaling v and relabeling a, b and

c that $a_{i,j} = a_{i,k} = 1$, $b_{i,j} = -1$, and $c_{i,j} = 0$ for some $j, k \leq q$. Then $\{b_{i,k}, c_{i,k}\} = \{0, -1\}$, and in either case, the vector $v' = (a', b', c')$ obtained from v by setting

$$a'_{i,j} = a_{i,j} + a_{i,k} = 2, \quad b'_{i,j} = b_{i,j} + b_{i,k}, \quad c'_{i,j} = c_{i,j} + c_{i,k}, \quad \text{and} \quad a'_{i,k} = b'_{i,k} = c'_{i,k} = 0,$$

is a conformal sum of nonzero vectors in $\ker_{\mathbb{Z}} A$ if and only if v is, which contradicts the unimodularity of $\Lambda_3 G_q A$.

Next, if for all i, the vectors a_i, b_i, and c_i each have at most one nonzero entry, then we are in case (ii) above. Otherwise, after appropriate scaling of v and relabeling of a, b and c, we have $a_{i,j} = 1$, $a_{i,k} = -1$, $b_{i,j} = -1$, and $c_{i,j} = 0$ for some $i \leq n$ and $j, k \leq q$. If $b_{i,k} = 1$, then these 4 nonzero entries form a primitive vector in the kernel of $\Lambda_3 G_q A$, so v has no other nonzero entries and we are in case (i) above. In all remaining cases, $b_{i,k} = 0$ and $c_{i,k} = 1$, meaning we are in case (iii) above. Note that the vector v' constructed from v still yields 0 in the j-th and k-th rows of the identity blocks portion of $\Lambda_3 G_q A$, as well as each row of $\Lambda_3 G_q A$ consisting of copies of A. As such, v' still lies in the kernel of $\Lambda_3 G_q A$. This completes the proof. $\qquad\square$

4.9 The classification

In this section, we present the complete classification of the unimodular HM pairs. Just as in the characterization of unimodular simplicial complexes in Theorem 4.5.3, this classification comes in two forms: a recipe for constructing any unimodular HM pair, and a list of forbidden minors.

Theorem 4.9.1. *The following are equivalent:*

(a) $(\mathcal{C}, \mathbf{d})$ is unimodular;

(b) $(\mathcal{C}, \mathbf{d})$ contains no minor isomorphic to any HM pair listed in Proposition 4.7.2; and

(c) $\mathcal{C}$ is nuclear. If $\mathcal{C}$ has nucleus $D_{m,n}$, then $\mathbf{d}_v = 2$ for each Lawrence vertex v and each vertex v from the nucleus $D_{m,n}$. Otherwise, we can choose the vertices of $\mathcal{C}$ to make nucleus $\Delta_m \sqcup \Delta_n$ so that either

    (1) $\mathbf{d}_v = 2$ for each Lawrence vertex v, or

    (2) $\min\{m, n\} = 0$, $\mathbf{d}_v = 2$ for the unique vertex v of Δ_0 and $\mathbf{d}_v \leq 3$ for each Lawrence vertex v with equality attained at most once.

Proof. The implication (a) $\implies$ (b) follows from Propositions 4.7.2 and 4.6.6. Now, if $(\mathcal{C}, \mathbf{d})$ satisfies (b), then via the minors in (1), Theorem 4.5.3 implies $\mathcal{C}$ is nuclear. So (c) follows once we show that the remaining minors in Proposition 4.7.2 ensure $\mathbf{d}$ satisfies the necessary requirements.

If $\mathcal{C}$ has nucleus $D_{m,n}$, then minors (5) and (7) ensure that (c) is satisfied, so assume $\mathcal{C}$ has nucleus $\Delta_m \sqcup \Delta_n$ and that $\mathcal{C}$ has a Lawrence vertex v such that $\mathbf{d}_v \geq 3$. If $\mathbf{d}_v \geq 4$, then minor (6) ensures that in the iterative construction of $\mathcal{C}$ as a nuclear complex, v must have been added before any ghost vertices. Therefore $\mathcal{C}$ is built up from $\Lambda^k(\Delta_m \sqcup \Delta_m)$. Permuting the order in which we add a Lawrence vertex does not change the resulting vertex-labeled complex so we may assume $k = 1$. Minor (3) ensures we can assume without loss of generality that $n = 0$. So at this point, we know $(\mathcal{C}, \mathbf{d})$ has been built from an HM pair $(\mathcal{C}', \mathbf{d}')$ where $\mathcal{C}' = \Lambda(\Delta_m \sqcup \Delta_0)$. Minor (4) ensures that $\mathbf{d}'_u = 2$ for the vertex u from Δ_0. However, we can realize this same complex where u is the Lawrence vertex and v is the unique vertex in Δ_0, so we could have chosen a different set of vertices to play the role of the nucleus that would not require v to be added as a Lawrence vertex.

To complete the proof of (b) $\implies$ (c), it remains to consider the case where $\mathbf{d}_v \in \{2,3\}$ for each Lawrence vertex v. Using a similar argument as before, we can see that if all the Lawrence vertices with $\mathbf{d}_v = 3$ are added before any ghost vertices in the construction of $\mathcal{C}$ as a nuclear complex, then by choosing different vertices to play the role of the nucleus we could have $\mathbf{d}_v = 2$ for all Lawrence vertices v. So assume all non-binary Lawrence vertices are added after a ghost vertex. Minor (8) ensures that at most one Lawrence vertex v can have $\mathbf{d}_v = 3$, and minors (2), (3), and (4) imply the remaining conditions.

It remains to show (c) $\implies$ (a) so assume $(\mathcal{C}, \mathbf{d})$ satisfies (c). If $\mathcal{C}$ has nucleus $D_{m,n}$, then Proposition 4.7.3 implies $(\mathcal{C}, \mathbf{d})$ is unimodular if and only if it satisfies (c). Since the operation cone commutes with Λ and G and unimodularity of $(\mathcal{C}, \mathbf{d})$ is independent of $\mathbf{d}_v$ for all cone vertices v, we may assume that $\mathcal{C}$ can be obtained without using the cone operation. We may also assume that G is never applied twice in a row, as doing so yields the same matrix as simply adding a single ghost vertex with a larger vertex label.

If, on the other hand, $\mathcal{C}$ has nucleus Δ_m, then since $\Lambda G \Delta_m = \text{cone}^{m+1}(\Delta_0 \sqcup \Delta_0)$, we may restrict attention to the final case, namely where $\mathcal{C}$ has nucleus $\Delta_m \sqcup \Delta_n$. In particular,

$$\mathcal{C} = \Lambda^{k_0} G \Lambda^{k_1} G \Lambda^{k_2} G \cdots G \Lambda^{k_l}(\Delta_m \sqcup \Delta_n).$$

If $(\mathcal{C}, \mathbf{d})$ satisfies the first case of (c) then unimodularity of $(\mathcal{C}, \mathbf{d})$ follows from uni-

modularity of $(\Delta_m \sqcup \Delta_n, \mathbf{d}')$ and the fact that adding Lawrence vertices with label 2 and ghost vertices of any vertex label preserves unimodularity. So it only remains to show unimodularity of pairs $(\mathcal{C}, \mathbf{d})$ that satisfy the second case of (c).

Let $(\mathcal{C}', \mathbf{d}')$ be a hierarchical model. We claim that if $(\Lambda\mathcal{C}', (\mathbf{d}', 3))$ is unimodular, then $(\Lambda G \mathcal{C}', (\mathbf{d}', q, 3))$ and $(\Lambda\Lambda\mathcal{C}', (\mathbf{d}', 2, 3))$ are unimodular as well. Indeed, unimodularity of $(\Lambda G \mathcal{C}', (\mathbf{d}', q, 3))$ follows from Proposition 4.8.1 and $\ker_{\mathbb{Z}}(\mathcal{A}_{\Lambda G \mathcal{C}', (\mathbf{d}', q, 3)}) = \ker_{\mathbb{Z}}(\Lambda_3 G_q \mathcal{A}_{\mathcal{C}', \mathbf{d}'})$ and $\ker_{\mathbb{Z}}(\mathcal{A}_{\Lambda\mathcal{C}', (\mathbf{d}', 3)}) = \ker_{\mathbb{Z}}(\Lambda_3 \mathcal{A}_{\mathcal{C}', \mathbf{d}'})$. Additionally, $(\Lambda\Lambda\mathcal{C}', (\mathbf{d}', 2, 3))$ has the same defining matrix as $(\Lambda\Lambda\mathcal{C}', (\mathbf{d}', 3, 2))$, which is unimodular if and only if $(\Lambda\mathcal{C}', (\mathbf{d}', 3))$ is. At this point, unimodularity of $(\mathcal{C}, \mathbf{d})$ follows by induction if we show unimodularity of the HM pair $(\Lambda(\Delta_m \sqcup \Delta_0), \mathbf{e})$ where $\mathbf{e}_v = 3$ for the Lawrence vertex and $\mathbf{e}_v = 2$ for the vertex of Δ_0. Letting p be the product of the vertex labels in Δ_m, this HM pair has the same matrix as the HM pair $(\{12, 13, 23\}, (3, 2, p))$ which was shown to be unimodular in [47]. $\qquad\square$

As a corollary of Theorem 4.9.1, we obtain a classification of the unimodular *discrete undirected graphical models*, that is, hierarchical models whose simplicial complex is the clique complex of a graph. For a given graph G, we let $\mathcal{C}(G)$ denote the clique complex of G. A *suspension vertex* of G is a vertex that shares an edge with every other vertex. Let $S^k G$ denote the graph obtained by adding k suspension vertices to G.

Corollary 4.9.2. *The matrix $\mathcal{A}_{\mathcal{C}(G), \mathbf{d}}$ is unimodular if and only if one of the following holds:*

(a) *G is a complete graph;*
(b) *$G = S^k C_4$, where C_4 is the four-cycle and $\mathbf{d}_v = 2$ for each vertex from C_4; or*
(c) *G is obtained by gluing two complete graphs along a (possibly empty) common clique.*

Proof. The constraints on G follow immediately from [12, Lemma 5.3]. When $G = S^k C_4$, Theorem 4.9.1 implies the constraints on $\mathbf{d}$. $\qquad\square$

4.10 The Graver basis of a unimodular hierarchical model

To conclude this section, we present a combinatorial characterization of the Graver basis of any unimodular hierarchical model's defining matrix. Following the constructive

characterization of unimodular hierarchical models in Theorem 4.9.1, our characterization comes in two steps: (i) a description of the Graver basis of each nucleus, and (ii) a description of how to obtain the Graver basis of a matrix produced by one of the unimodularity preserving operations allowed by Theorem 4.9.1(c) given the Graver basis of its input matrix.

Proposition 4.10.1 characterizes the Graver basis of each possible nucleus. We have already observed that for any unimodular matrix, the Graver basis consists of vectors with entries in $\{0, 1, -1\}$ and coincides with the set of circuits (Proposition 1.6.10). As such, the characterization presented in Proposition 4.10.1 simply describes the signed circuits of the oriented matroid underlying each given hierarchical model. We remind the reader how to construct oriented matroids from signed circuits and cuts of a directed graph; for a more thorough introduction to oriented matroids, see [15].

A *signed circuit* of a directed graph G is a bipartition of the set of edges in a simple cycle $v_1, \ldots, v_n, v_1$ of the undirected graph underlying G according to whether or not the edge $v_i v_{i+1}$ agrees with G's orientation. The edges that agree are called *positive* while those that disagree are called *negative*. A *signed bond* of G is a bipartition of the set of edges in a bond (i.e. minimal cut) of G that splits G into connected components A and B according to whether or not the edge e is directed from A to B. The edges pointing from A to B are called *positive* while those pointing from B to A are called *negative*. The set of signed circuits and the set of signed bonds of G each are the signed circuits of an oriented matroid. These two oriented matroids are dual to each other.

Let $K_{2^{m+1}, 2^{n+1}}$ denote the complete bipartite graph on partite sets $2^{[m+1]}$ and $2^{[n+1]}$. Label the vertices in each partite set by the binary $(m+1)$- and $(n+1)$-tuples in $\{1, 2\}^{m+1}$ and $\{1, 2\}^{n+1}$, respectively. Each edge is naturally labeled with the $(m + n + 2)$-tuple obtained by concatenating the labels of its vertices. Let $G_1^{m,n}$ be the directed graph with underlying undirected graph $K_{2^{m+1}, 2^{n+1}}$ where all edges are directed from the partite set $2^{[m+1]}$ to the partite set $2^{[n+1]}$. Let $G_2^{m,n}$ be the directed graph obtained from $G_1^{m,n}$ by reversing the orientation of the edges whose $(m+n+2)$-tuple has an odd number of twos.

Proposition 4.10.1. *The set of signed circuits of the oriented matroid underlying the columns of $\mathcal{A}_{\Delta_m}$ is empty. The signed circuits of the oriented matroid underlying the columns of $\mathcal{A}_{\Delta_m \sqcup \Delta_n, \mathbf{2}}$ are the signed circuits of $G_1^{m,n}$. The signed circuits of the oriented matroid underlying the columns of $\mathcal{A}_{D_{m,n}, \mathbf{2}}$ are the signed bonds of $G_2^{m,n}$.*

Proof. For the first statement, note that the matrix $\mathcal{A}_{\Delta_m}$ is square with determinant 1 and

therefore has trivial kernel. For the second statement, recall that row operations do not affect the underling oriented matroid. Then note that after multiplying the appropriate rows of $\mathcal{A}_{\Delta_m \sqcup \Delta_n, \mathbf{2}}$ by -1, we obtain the vertex-arc incidence matrix of $G_1^{m,n}$.

By Proposition 4.2.6 and its proof, $\mathcal{A}_{D_{m,n}, \mathbf{2}}$ can be obtained from a particular Gale dual B of $\mathcal{A}_{\Delta_m \sqcup \Delta_n, \mathbf{2}}$ by negating the columns of B corresponding to the binary $(m + n + 2)$-tuples with an odd number of twos, then negating all negative rows of the resulting matrix. The oriented matroid underlying the columns of B is dual to the oriented matroid underlying the columns of $\mathcal{A}_{\Delta_m \sqcup \Delta_n, \mathbf{2}}$. Therefore, the signed circuits of the oriented matroid underlying the columns of B are the signed bonds of $G_1^{m,n}$. On the oriented matroid level, the process of turning B into $\mathcal{A}_{D_{m,n}, \mathbf{2}}$ by negating the appropriate rows and columns has the effect of reversing the orientation of the edges of $G_1^{m,n}$ corresponding to binary $(m + n + 2)$-tuples with an odd number of twos. This gives us $G_2^{m,n}$. □

Example 4.10.2. We illustrate Proposition 4.10.1 in the case $m = 1$ and $n = 0$. We can draw the relevant simplicial complexes as follows. Note that $D_{1,0}$ has a ghost vertex which we indicate pictorially with an open circle.

Figure 4.4 depicts $\mathcal{A}_{\Delta_1 \sqcup \Delta_0, \mathbf{2}}$ and $\mathcal{A}_{D_{1,0}, \mathbf{2}}$ alongside the directed graphs $G_1^{1,0}$ and $G_2^{1,0}$.

The edges corresponding to the binary 3-tuples $221, 222, 212, 211$ form a cycle in $G_1^{1,0}$ with 221 and 212 having positive orientation, and 222 and 211 both having negative orientation. Therefore the vector $e_{221} + e_{212} - e_{222} - e_{211}$ is in the Graver basis of $\mathcal{A}_{\Delta_1 \sqcup \Delta_0, \mathbf{2}}$. The edges corresponding to the binary 3-tuples $221, 211, 121, 111$ form a minimal cut in $G_2^{1,0}$, partitioning the vertices into $\{1\}$ and its complement. Calling these sets A and B respectively, note that the edges corresponding to 221 and 111 point from A into B whereas the edges corresponding to 211 and 121 point from B into A. Therefore, the vector $e_{221} + e_{111} - e_{211} - e_{121}$ lies in the Graver basis of $\mathcal{A}_{D_{1,0}, \mathbf{2}}$.

We now state the obvious extension of Proposition 4.10.1 to the case where $\mathcal{C} = \Delta_m \sqcup \Delta_n$ but $\mathbf{d} = \mathbf{2}$ need not hold.

$$\mathcal{A}_{\Delta_1 \sqcup \Delta_0, \mathbf{2}} = \begin{array}{cc}
 & \begin{matrix} 1 & 1 & 1 & 1 & 2 & 2 & 2 & 2 \\ 1 & 1 & 2 & 2 & 1 & 1 & 2 & 2 \\ 1 & 2 & 1 & 2 & 1 & 2 & 1 & 2 \end{matrix} \\
\begin{matrix} \{1,2\} & (1,1) \\ & (1,2) \\ & (2,1) \\ & (2,2) \\ \hline \{3\} & (1) \\ & (2) \end{matrix} & \left(\begin{matrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{matrix}\right)
\end{array}$$

$$\mathcal{A}_{D_{1,0}, \mathbf{2}} = \begin{array}{cc}
 & \begin{matrix} 1 & 1 & 1 & 1 & 2 & 2 & 2 & 2 \\ 1 & 1 & 2 & 2 & 1 & 1 & 2 & 2 \\ 1 & 2 & 1 & 2 & 1 & 2 & 1 & 2 \end{matrix} \\
\begin{matrix} \{1\} & (1) \\ & (2) \\ \hline \{2\} & (1) \\ & (2) \end{matrix} & \left(\begin{matrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \end{matrix}\right)
\end{array}$$

$G_1^{1,0}$ graph with vertices labeled 22, 21, 1, 2, 12, 11.

$G_2^{1,0}$ graph with vertices labeled 22, 21, 1, 2, 12, 11.

Figure 4.4: The matrices $\mathcal{A}_{\Delta_1 \sqcup \Delta_0, \mathbf{2}}$ and $\mathcal{A}_{D_{1,0}, \mathbf{2}}$, along with their respective graphs $G_1^{1,0}$ and $G_2^{1,0}$, from Example 4.10.2.

Proposition 4.10.3. *Let $\mathcal{C} = \Delta_m \sqcup \Delta_n$ with vertex set $\{1, \ldots, m + n + 2\}$ and facets $\{1, \ldots, m+1\}$ and $\{m+2, \ldots, m+n+2\}$. Fix $\mathbf{d} \in \mathbb{Z}^{m+n+2}$ and let $G_{\mathbf{d}}^{m,n}$ be the complete bipartite graph with partite vertex sets $[d_1] \times \cdots \times [d_{m+1}]$ and $[d_{m+2}] \times \cdots \times [d_{m+n+2}]$, with edges oriented so that they all point towards the same partite set. Label each edge by the element of $[d_1] \times [d_{m+n+2}]$ obtained by concatenating its vertices. Then the signed circuits of the oriented matroid underlying the columns of $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ are the signed circuits of $G_{\mathbf{d}}^{m,n}$.*

Having now characterized the Graver basis of each nucleus, it remains to describe how each of the operations $\mathrm{cone}(-)$, G, Λ_2 and Λ_3 affect the Graver basis of a unimodular matrix $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$. Adding ghost vertices changes $\mathcal{A}_{\mathcal{C}, \mathbf{d}}$ by $\mathcal{A}_{G\mathcal{C}, (\mathbf{d}\ q)} = G_q(\mathcal{A}_{\mathcal{C}, \mathbf{d}})$. Therefore, every element in the Graver basis of $\mathcal{A}_{G\mathcal{C}, (\mathbf{d}\ q)}$ is either of the form

$$(0 \quad \cdots \quad 0 \quad e_i \quad 0 \quad \cdots \quad 0 \quad -e_i \quad 0 \quad \cdots \quad 0)$$

111

where e_i is the i^{th} standard basis vector, or $\begin{pmatrix} u_1 & u_2 & \ldots & u_q \end{pmatrix}$ where $u_1 + u_2 + \cdots + u_q$ is a conformal sum which lies in the Graver basis of $\mathcal{A}_{\mathcal{C},\mathbf{d}}$. It is easy to see that the kernel of $\mathcal{A}_{\text{cone}(\mathcal{C}),(\mathbf{d}\ \ q)}$ is the same as the kernel of the following matrix

$$\begin{pmatrix} \mathcal{A}_{\mathcal{C},\mathbf{d}} & 0 & \ldots & 0 \\ 0 & \mathcal{A}_{\mathcal{C},\mathbf{d}} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \mathcal{A}_{\mathcal{C},\mathbf{d}} \end{pmatrix}.$$

The Graver basis of $\mathcal{A}_{\text{cone}(\mathcal{C}),(\mathbf{d},q)}$ consists of elements of the form $\begin{pmatrix} 0 & \cdots & 0 & u & 0 & \cdots & 0 \end{pmatrix}$ where u is in the Graver basis of $\mathcal{A}_{\mathcal{C},\mathbf{d}}$. The kernel of $\mathcal{A}_{\Lambda(\mathcal{C}),(\mathbf{d}\ \ p)}$ is the same as the kernel of $\Lambda_p \mathcal{A}_{\mathcal{C},\mathbf{d}}$. From this it easily follows that the Graver basis of $\mathcal{A}_{\Lambda(\mathcal{C}),(\mathbf{d}\ \ 2)}$ consists of elements of the form $\begin{pmatrix} u & -u \end{pmatrix}$ where u is in the graver basis of $\mathcal{A}_{\mathcal{C},\mathbf{d}}$.

This leaves the operation Λ_3. Theorem 4.9.1 tells us that we only need to consider Λ_3 when applied to $\mathcal{A}_{\mathcal{C},\mathbf{d}}$ where $\mathcal{C}$ is nuclear with nucleus $\Delta_m \sqcup \Delta_0$ and $\mathbf{d}_v = 2$ for the vertex of the Δ_0 and every Lawrence vertex. The operation $\text{cone}(-)$ commutes with each other operation, and Λ_2 commutes with Λ_3, so it suffices to describe the Graver basis of complexes of the form

$$\mathcal{C} = \Lambda_3 G \Lambda^{k_1} G_{q_1} \Lambda_2^{k_2} G_{q_2} \cdots \Lambda_2^{k_l} G_{q_l} \Lambda_2^{k_{l+1}} (\Delta_m \sqcup \Delta_0).$$

Using Proposition 4.8.2, a Graver basis for the matrix corresponding to $\mathcal{C}$ can be obtained inductively, starting with the Graver basis for the matrix corresponding to the labeled complex $\Lambda_3(\Delta_m \sqcup \Delta_0)$, whose matrix is the same as for the HM pair $(\{12, 13, 23\}, (3, 2, p))$.

Remark 4.10.4. The Graver basis of any unimodular HM pair can be obtained by way of Propositions 4.10.1 and 4.10.3 and the discussion thereafter.

REFERENCES

[1] Maple 2015. Maplesoft, a division of Waterloo Maple Inc. Waterloo, Ontario.

[2] 4ti2 team. 4ti2—a software package for algebraic, geometric and combinatorial problems on linear spaces. Available at www.4ti2.de.

[3] Richa Agarwala, Vineet Bafna, Martin Farach, Mike Paterson, and Mikkel Thorup. On the approximability of numerical taxonomy (fitting distances by tree metrics). *SIAM Journal on Computing*, 28(3):1073–1085, 1998.

[4] Federico Ardila. Subdominant matroid ultrametrics. *Annals of Combinatorics*, 8:379–389, 2004.

[5] Federico Ardila and Caroline J. Klivans. The Bergman complex of a matroid and phylogenetic trees. *Journal of Combinatorial Theory, Series B*, 96(1):38 – 49, 2006.

[6] Daniel Irving Bernstein. Completion of tree metrics and rank 2 matrices. *Linear Algebra and its Applications*, 533:1–13, 2017.

[7] Daniel Irving Bernstein. L-infinity optimization to bergman fans of matroids with an application to phylogenetics. *arXiv preprint arXiv:1702.05141*, 2017.

[8] Daniel Irving Bernstein and Colby Long. L-infinity optimization to linear spaces and phylogenetic trees. *SIAM Journal on Discrete Mathematics*, 31(2):875–889, 2017.

[9] Daniel Irving Bernstein, Colby Long, Christopher O'Neill, and Seth Sullivant. Supplementary materials.
www4.ncsu.edu/~dibernst/thesis/ThesisSupplement.html, 2017.

[10] Daniel Irving Bernstein and Christopher ONeill. Unimodular hierarchical models and their Graver bases. *Journal of Algebraic Statistics*, 8(2), 2017.

[11] Daniel Irving Bernstein and Seth Sullivant. Normal binary hierarchical models. *Experimental Mathematics*, 26(2):153–164, 2017.

[12] Daniel Irving Bernstein and Seth Sullivant. Unimodular binary hierarchical models. *Journal of Combinatorial Theory, Series B*, 123:97–125, 2017.

[13] Louis J Billera, Susan P Holmes, and Karen Vogtmann. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics*, 27(4):733–767, 2001.

[14] Pratik Biswas, Tzu-Chen Lian, Ta-Chung Wang, and Yinyu Ye. Semidefinite programming based algorithms for sensor network localization. *ACM Transactions on Sensor Networks (TOSN)*, 2(2):188–220, 2006.

[15] Anders Björner, Michel Las Vergnas, Neil White, Bernd Sturmfels, and Gunter M Ziegler. Oriented matroids. 1999.

[16] Emmanuel J Candes and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

[17] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.

[18] Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.

[19] Pei Chen and David Suter. Recovering the missing components in a large noisy low-rank matrix: Application to SFM. *IEEE transactions on pattern analysis and machine intelligence*, 26(8):1051–1063, 2004.

[20] Yuguo Chen, Ian H. Dinwoodie, and Seth Sullivant. Sequential importance sampling for multiway tables. *Ann. Statist.*, 34(1):523–545, 2006.

[21] V. Chepoi and B. Fichet. l_∞-approximation via subdominants. *Journal of Mathematical Psychology*, 44:600–616, 2000.

[22] David A Cox, John B Little, and Henry K Schenck. *Toric varieties.* American Mathematical Soc., 2011.

[23] Ruth Davidson and Seth Sullivant. Distance-based phylogenetic algorithms around a tritomy. *IEEE/ACM Transactions in Computational Biology and Bioinformatics*, 11:325–335, 2014.

[24] Persi Diaconis and Bernd Sturmfels. Algebraic algorithms for sampling from conditional distributions. *Ann. Statist.*, 26(1):363–397, 1998.

[25] Andreas Dress, Katharina Huber, and Mike Steel. A matroid associated with a phylogenetic tree. *Discrete Mathematics and Theoretical Computer Science*, 16(2):41–55, 2014.

[26] Andreas WM Dress, Katharina T Huber, and Mike Steel. Lassoing a phylogenetic tree I: basic properties, shellings, and covers. *Journal of mathematical biology*, 65(1):77–105, 2012.

[27] Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on algebraic statistics*, volume 39. Springer Science & Business Media, 2008.

[28] David Eisenbud. *Commutative Algebra: with a view toward algebraic geometry*, volume 150. Springer Science & Business Media, 2013.

[29] Martin Farach, Sampath Kannan, and Tandy Warnow. A robust model for finding optimal evolutionary trees. *Algorithmica*, 13(1-2):155–179, 1995.

[30] Alain Guénoche and Bruno Leclerc. The triangles method to build x-trees from incomplete distance matrices. *RAIRO-Operations Research*, 35(2):283–300, 2001.

[31] Alain Guénoche, Bruno Leclerc, and Vladimir Makarenkov. On the extension of a partial metric to a tree metric. *Discrete mathematics*, 276(1):229–248, 2004.

[32] Brendan Hassett. *Introduction to algebraic geometry*. Cambridge University Press, 2007.

[33] Serkan Hoşten and Seth Sullivant. Gröbner bases and polyhedral geometry of reducible and cyclic models. *J. Combin. Theory Ser. A*, 100(2):277–301, 2002.

[34] Serkan Hoşten and Seth Sullivant. A finiteness theorem for Markov bases of hierarchical models. *J. Combin. Theory Ser. A*, 114(2):311–321, 2007.

[35] Bill Jackson, Tibor Jordán, and Shin-ichi Tanigawa. Combinatorial conditions for the unique completability of low-rank matrices. *SIAM Journal on Discrete Mathematics*, 28(4):1797–1819, 2014.

[36] Bill Jackson, Tibor Jordán, and Shin-ichi Tanigawa. Unique low rank completability of partially filled matrices. *Journal of Combinatorial Theory, Series B*, 121:432–462, 2016.

[37] Tomáš Kaiser. Minors of simplicial complexes. *Discrete Applied Mathematics*, 157(12):2597 – 2602, 2009. Second Workshop on Graph Classes, Optimization, and Width Parameters.

[38] Gil Kalai, Eran Nevo, and Isabella Novik. Bipartite rigidity. *Transactions of the American Mathematical Society*, 368(8):5515–5545, 2016.

[39] Franz Király, Louis Theran, and Ryota Tomioka. The algebraic combinatorial approach for low-rank matrix completion. *Journal of Machine Learning Research*, 16:1391–1436, 2015.

[40] Franz J Király, Zvi Rosen, and Louis Theran. Algebraic matroids with graph symmetry. *arXiv preprint arXiv:1312.3777*, 2013.

[41] Franz J Király, Louis Theran, and Ryota Tomioka. The algebraic combinatorial approach for low-rank matrix completion. *Journal of Machine Learning Research*, 16:1391–1436, 2015.

[42] Gerard Laman. On graphs and rigidity of plane skeletal structures. *Journal of Engineering mathematics*, 4(4):331–340, 1970.

[43] Bo Lin, Bernd Sturmfels, Xiaoxian Tang, and Ruriko Yoshida. Convexity in tree spaces. *https://arxiv.org/abs/1510.08797*, 2016.

[44] Bo Lin and Ruriko Yoshida. Tropical Fermat-Weber points. *https://arxiv.org/abs/1604.04674*, 2016.

[45] Diane Maclagan and Bernd Sturmfels. *Introduction to tropical geometry*, volume 161. American Mathematical Soc., 2015.

[46] Ezra Miller and Bernd Sturmfels. *Combinatorial commutative algebra*, volume 227. Springer Science & Business Media, 2004.

[47] Hidefumi Ohsugi and Takayuki Hibi. Toric ideals arising from contingency tables.

Commutative Algebra and Combinatorics, Ramanujan Mathematical Society Lecture Note Series, (4):91–115, 2007.

[48] Hidefumi Ohsugi and Takayuki Hibi. Non-very ample configurations arising from contingency tables. *Annals of the Institute of Statistical Mathematics*, 62(4):639–644, 2010.

[49] James Oxley. *Matroid Theory*. Oxford University Press, second edition, 2011.

[50] Johannes Rauh and Seth Sullivant. Lifting markov bases and higher codimension toric fiber products. `arXiv:1404.6392`, 2014.

[51] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons, Ltd., Chichester, 1986. A Wiley-Interscience Publication.

[52] Charles Semple and Mike Steel. *Phylogenetics*. Oxford University Press, Oxford, 2003.

[53] Charles Semple and Mike A Steel. *Phylogenetics*, volume 24. Oxford University Press on Demand, 2003.

[54] Amit Singer. A remark on global positioning from local distances. *Proceedings of the National Academy of Sciences*, 105(28):9507–9511, 2008.

[55] Amit Singer and Mihai Cucuringu. Uniqueness of low-rank matrix completion by rigidity theory. *SIAM Journal on Matrix Analysis and Applications*, 31(4):1621–1641, 2010.

[56] David Speyer and Bernd Sturmfels. The tropical Grassmannian. *Advances in Geometry*, 4(3):389–411, 2004.

[57] Bernd Sturmfels. *Gröbner Bases and Convex Polytopes*, volume 8 of *University Lecture Series*. American Mathematical Society, Providence, RI, 1996.

[58] Seth Sullivant. Small contingency tables with large gaps. *SIAM Journal of Discrete Mathematics*, 18(4):787–793, 2005.

[59] Seth Sullivant. Toric fiber products. *Journal of Algebra*, 316(2):560–577, 2007.

[60] Seth Sullivant. Algebraic statistics (in press). `http://www4.ncsu.edu/~smsulli2/Pubs/asbook.html`, 2018.

[61] Carlo Tomasi and Takeo Kanade. Shape and motion from image streams under orthography: a factorization method. *International Journal of Computer Vision*, 9(2):137–154, 1992.

[62] Michael S Waterman and Temple F Smith. On the similarity of dendrograms. *Journal of Theoretical Biology*, 73(4):789–800, 1978.

[63] Joe Whittaker. *Graphical models in applied multivariate statistics*. Wiley Publishing, 2009.

[64] Josephine Yu. Algebraic matroids and set-theoretic realizability of tropical varieties. *arXiv preprint arXiv:1506.01427*, 2015.

[65] Günter Ziegler. *Lectures on Polytopes*. Springer, 2000.