

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

From Voronoi Cells to Algebraic Statistics

Permalink

<https://escholarship.org/uc/item/0qv7p6mk>

Author

Alexandr, Yulia

Publication Date

2023

Peer reviewed|Thesis/dissertation

From Voronoi Cells to Algebraic Statistics

by

Yulia Alexandr

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Bernd Sturmfels, Co-chair

Professor Serkan Hoşten, Co-chair

Professor Nikhil Srivastava

Professor Vadim Gorin

Fall 2023

From Voronoi Cells to Algebraic Statistics

Copyright 2023

by

Yulia Alexandr

Abstract

From Voronoi Cells to Algebraic Statistics

by

Yulia Alexandr

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Bernd Sturmfels, Co-chair

Professor Serkan Hoşten, Co-chair

Algebraic statistics is a relatively young field, which explores how algebra and statistics interact, thus fostering a meaningful dialogue between theory and applications. Many statistical models are naturally made up of distributions whose coordinates satisfy polynomial equations. Viewing these models as algebraic varieties, we may use tools from algebraic geometry to gain additional insight into their properties. Alternatively, we may revisit familiar algebraic and geometric objects in the setting of probability distributions.

In the first part of this thesis, I develop the theory of *logarithmic Voronoi cells*. These are convex sets used to divide experimental data based on which point in the model each sample most likely came from. For finite models, linear models, toric models, and models of maximum likelihood degree one, I prove that these sets are polytopes and characterize them combinatorially. I then use their structure to maximize information divergence to linear and toric models. For the latter family, I present a new algorithm for computing maximizers using vertices of logarithmic Voronoi polytopes. For non-polytopal logarithmic Voronoi cells, I develop a method to compute them via the framework of numerical algebraic geometry.

The second and third parts of this dissertation focus on conditional independence. In the second part, I study *context-specific independence* and introduce the family of decomposable CSmodels. I prove that these models mirror many of the algebraic and combinatorial properties that characterize decomposable graphical models, and hence they are good candidates for decomposable models in the context-specific setting. I give the strongest possible algebraic characterization of decomposable CSmodels by describing their prime ideals. In the third part, I focus on *nonparametric algebraic statistics*. I study dimensions, defining polynomials, and degrees of the moment varieties of conditionally independent mixture distributions on \mathbb{R}^n . The last chapter features both symbolic and numerical computational methods.

To my family

Contents

Contents	ii
List of Figures	iv
1 Introduction	1
1.1 Discrete models	1
1.2 Continuous models	20
1.3 Statement of contributions	30
I Logarithmic Voronoi cells	32
2 Discrete setting	33
2.1 When are logarithmic Voronoi cells polytopes?	33
2.2 The chaotic universe model	40
2.3 Numerical algebraic geometry	48
3 Linear models	52
3.1 Combinatorial types	53
3.2 On the boundary	57
3.3 Information divergence	60
3.4 Partial linear models	63
4 Toric models	67
4.1 Critical points of information divergence	69
4.2 The chamber complex and the algorithm	72
4.3 Reducible models	82
4.4 Models of ML degree one	90
5 Gaussian setting	101
5.1 Linear concentration models	102
5.2 Directed graphical models	107
5.3 Covariance models	109

II	Context-specific models	117
6	Decomposable context-specific models	118
6.1	Introduction to CStrees	120
6.2	Decomposable CSmodels in three variables	127
6.3	Combinatorial properties of balanced CStrees	130
6.4	Algebraic characterization	136
III	Nonparametric algebraic statistics	142
7	Moment varieties for mixtures of products	143
7.1	Familiar varieties	145
7.2	Finiteness	147
7.3	Toric combinatorics	149
7.4	Secant varieties	152
	Bibliography	160

List of Figures

1.1	Probability simplices we can visualize.	2
1.2	From left to right: the Hardy-Weinberg curve, the twisted cubic curve, and the independence model of two binary random variables.	3
1.3	Logarithmic Voronoi cells of the twisted cubic.	13
1.4	Nodal cubic in \mathbb{R}^2 (left) and the bivariate correlation model in PD_3 (right).	22
1.5	The unrestricted correlation model for $m = 3$	24
1.6	Logarithmic Voronoi cells (in pink) of the model in Example 1.2.8 plotted on the (x_1, x_2) -plane and (y_1, y_2) -plane, respectively.	28
2.1	Logarithmic Voronoi cell (green) inside its log-normal polytope (pink) for a given point (yellow) in the model from Example 2.3.3.	34
2.2	One-dimensional log-normal polytopes at various points	36
2.3	Nonlinear boundary arising from two disjoint linear models	39
2.4	The fibres and image of the moment map for the Segre of Example 2.1.5	40
2.5	Logarithmic Voronoi cells (rhombic dodecahedra) of interior points for $n = 4$, $d = 9$ (on the left) and $d = 10$ (on the right).	41
3.1	Partition of the tetrahedron Δ_3 into triangles (left) and quadrilaterals (right).	52
3.2	Sampled points on the linear model corresponding to $B = [1 \ -5 \ 2 \ 2]$ and triangular logarithmic Voronoi cells.	56
3.3	Linear model given by $B = [-2, -1, 1, 2]^T$	63
3.4	Polytopes Q_p and \overline{Q}_p for a point on a facet of a 2-dimensional model in Δ_3	65
3.5	Logarithmic Voronoi cells at sampled points on the model defined as the convex hull of three points.	66
4.1	The chamber complex of a pentagon.	74
4.2	Chamber complex of the 2×3 independence model (left and middle) and the middle chamber (right).	82
4.3	Twisted Veronese models $\mathcal{V}_{2,3}$ and $\mathcal{V}_{3,2}$, respectively.	92
4.4	Chamber complex of the box model $\mathcal{B}_{3,3,2}^{(1)}$	96
4.5	The chamber complex and two logarithmic Voronoi polytopes that yield maximum divergence.	98

5.1	Pillow-shaped logarithmic Voronoi spectrahedra	104
5.2	The logarithmic Voronoi cell at Σ' of $1 \rightarrow 2 \rightarrow 4 \leftarrow 3$	108
5.3	The logarithmic Voronoi cell at $\Sigma_{1/2}$ for the bivariate correlation model.	114
6.1	A CStree for $p = 3$ and its minimal context DAGs.	122
6.2	All CStrees with $p = 3$, which do not represent a DAG.	128
6.3	A balanced CStree with a non-perfect minimal context.	130
6.4	The context subtree of the tree in Figure 6.5 for the context $X_3 = 0$, and its minimal context DAGs.	132
6.5	Balanced CStree on five binary random variables whose empty context DAG is not perfect.	137

Acknowledgments

I am immensely grateful to my advisors, Bernd Sturmfels and Serkan Hoşten, whose expertise and guidance have been instrumental in shaping my academic journey, and whose belief in my abilities has driven me forward. Bernd, thank you for your boundless encouragement, inspirational energy, and hands-on mentorship. I feel very lucky to be a part of your wonderful academic family. Serkan, thank you for your infinite patience, contagious enthusiasm, and invaluable advice. Working with you has been a source of motivation and wisdom.

Thank you to my collaborators: Eliana Duarte, Alexander Heaton, Serkan Hoşten, Joe Kileel, Bernd Sturmfels, Sascha Timme, and Julian Vill. I had a lot of fun doing math together, and I learned so much from each of you. Thank you, Vadim Gorin and Nikhil Srivastava, for taking the time to be on my dissertation committee. Thank you, Vera Serganova, for the interesting classes you taught during my time at Berkeley. Thank you to Vicky, Jon, Clay, Christian, and Saskia for all your help with the tricky administrative tasks. Thank you to all the algebraic statistics folks, whose contributions have motivated and influenced my own work. I am proud to be a member of such a warm and welcoming mathematical community.

Thank you to my dear friends, who made my PhD journey so memorable. Yelena, thank you for always being there for me, for the adventures we shared all over the world, for the many vegan meals and true-crime episodes we enjoyed, and for the countless hours we spent talking, crying, and laughing together. You are the best friend, (academic) sister, and roommate I could ever hope for. Thank you, Adam, for always understanding me, for your infectious sense of humor, and for the many healing, deep, and silly conversations we have had over Boichik bagels. Thank you, AJ, for your ability to turn the most routine errands into exciting adventures, for our regular peaceful walks and tofu soup dinners. Thank you, Roshni, for being the best climbing partner, and for our much-needed late-night milkshake escapades followed by bad TV. Thank you, Yana, for making our Oakland apartment feel like home, for our open-minded conversations and the variety of teas we enjoyed during our three years of living together. Thank you, Maksym, for being a fun office mate and supportive friend. Thank you, Misha, for sharing many joyful moments with me in Middletown, San-Francisco, Oakland, Hawai'i, and Dublin. Thank you, Cameron, for helping me believe anything was possible. Thank you, Rikhav, Alisar, Ruwan, Alex, Pi, and Krishnan, for being fantastic friends, and for the memories we created together. Thank you to my academic siblings: Chiara, Claudia, and Yassine for helping me navigate this unique experience.

Thank you, Eric, for always supporting my dreams, your unshakeable belief in me during times of doubt, and for the many adventures we have shared and those on the horizon. Thank you, Andrew-the-cat, for your emotional support and abuse. Finally, an immense thank you to my parents and my family for your continuous encouragement and the shared amusement at the idea of me becoming a doctor.

Chapter 1

Introduction

In this chapter, we introduce and motivate the main themes in this dissertation. We describe what a statistical model is to algebraic statisticians and what questions they are interested in studying. In particular, we define the maximum likelihood estimation problem, which inspires the study of logarithmic Voronoi cells (Part I). We also describe the concept of conditional independence, which plays a major role in the study of discrete context-specific models (Part II) and of the moment varieties in the continuous nonparametric setting (Part III). Section 1.3 summarizes all contributions in this dissertation and acknowledges joint work.

1.1 Discrete models

In this section, we focus on discrete statistical models. We define a probability simplex and an algebraic statistical model, and give motivating examples. We then describe the maximum likelihood estimation problem for such models and define logarithmic Voronoi cells. Finally, we focus on conditional independence and the corresponding algebraic equations that arise in the discrete setting.

Variety of distributions

Discrete statistical models arise when we work with random variables that have finite state spaces. Such models live inside the *probability simplex*, defined as

$$\Delta_{n-1} = \{(p_1, p_2, \dots, p_n) \in \mathbb{R}_{\geq 0}^n : p_1 + p_2 + \dots + p_n = 1\}.$$

Here, n is the total number of possible outcomes, while p_i is the probability of observing state i for all $i \in [n]$. The probability simplex Δ_{n-1} is a polytope of dimension $n - 1$ in \mathbb{R}^n . We may visualize Δ_{n-1} for $n = 1, 2, 3, 4$ inside the hyperplane $p_1 + p_2 + \dots + p_n = 1$, as in Figure 1.1. In this thesis, we will often work with the open probability simplex Δ_{n-1}° , where we assume that $p_i > 0$ for all $i \in [n]$.




Figure 1.1: Probability simplices we can visualize.

A *statistical model* is defined to be any subset of the probability simplex. Most of the models we will consider in this thesis will be naturally made up of distributions whose coordinates satisfy polynomial equations. Hence, we will be working with *algebraic statistical models* [114], given as the intersection of Δ_{n-1} with an algebraic variety or as the image of a rational map φ . In the latter case, we say that a model admits a parametric description and we may write $\mathcal{M} = \varphi(\Theta) \subseteq \Delta_{n-1}$ for some parameter space Θ . Important examples of parametric models include linear and toric models, discussed in detail in later chapters. The *Zariski closure* of a model $\mathcal{M} \in \Delta_n \subseteq \mathbb{C}^n$ is the smallest complex variety containing \mathcal{M} . The *dimension* of an algebraic statistical model \mathcal{M} is the dimension of the intersection of its Zariski closure with the hyperplane defining the simplex $p_1 + p_2 + \cdots + p_n = 1$. The process of recovering the ideal of the Zariski closure of the image of φ from the parametric description is called *implicitization* [87, Chapter 4]. This problem can be computationally challenging for large models, as the standard approach relies on Gröbner basis computations [87, Chapter 1]. We demonstrate the extent of this difficulty for mixture models in Chapter 7. For some classes of models, the defining equations can be obtained by observing special combinatorial properties of their parametrization. In Chapter 6, we characterize the ideals of decomposable context-specific models by utilizing the structure of their graphical representation.

Remark 1.1.1. Note that even though we defined a discrete statistical model as a subset of the probability simplex, we will often identify our models with the underlying complex varieties. In other words, we will often work with the Zariski closure of a statistical model in the complex space. This will be particularly useful in Definition 1.1.8 and when we treat exponential families in Chapter 4.

A natural example of a parametric statistical model arises when we consider an experiment given by a coin flip.

Example 1.1.2 (Coin flips). Suppose we have a biased coin that lands tails up with probability t . When we flip it two times and record the number of heads, there are three possible

outcomes: 0 head, 1 heads, and 2 heads. Our model \mathcal{M} is a curve in Δ_2 , parametrized by

$$\varphi : (0, 1) \rightarrow \Delta_2 : t \mapsto p(t) = (t^2, 2t(1-t), (1-t)^2) \in \mathcal{M}.$$

Statisticians know it as the *Hardy-Weinberg curve*. To obtain the implicit equations, we need to compute the *elimination ideal* in *Macaulay2* [60] by running the commands below.

```
R=QQ[t,s,p1,p2,p3];
I=ideal(p1-t^2, p2-2*t*s, p3-s^2);
J=eliminate(I,{t,s})
```

This returns the principal ideal generated by the quadric $p_2^2 - p_1p_3$. It defines a hypersurface in \mathbb{R}^3 , which intersected with the simplex gives us the model \mathcal{M} . If we were to flip the same coin three times instead of two, we would obtain the model known as the *twisted cubic curve* inside Δ_3 , which we will revisit several times in this dissertation. Both of the models are plotted in Figure 1.2 (left and middle, respectively).




Figure 1.2: From left to right: the Hardy-Weinberg curve, the twisted cubic curve, and the independence model of two binary random variables.

Another example we consider is the independence model of two random variables, where the statistical independence relation translates into the algebraic rank condition on the matrix of joint probabilities.

Example 1.1.3 (Independence model). Let X and Y be two random variables with the state spaces $[r_1]$ and $[r_2]$, respectively. The probability that X takes some value $i \in [r_1]$, while Y takes some value $j \in [r_2]$ at the same time is the *joint probability*

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

The collection of all joint probabilities for some general random variables X and Y can be represented as the set of all $r_1 \times r_2$ real matrices (p_{ij}) whose entries are nonnegative and sum to one. Equivalently, these can be identified with the probability simplex $\Delta_{r_1 r_2 - 1}$. Note that from the joint probabilities, we can always recover the probability of an individual random variable taking a certain value. The probabilities

$$p_{i+} := \mathbb{P}(X = i) = \sum_{j=1}^{r_2} p_{ij} \quad \text{and} \quad p_{+j} := \mathbb{P}(Y = j) = \sum_{i=1}^{r_1} p_{ij}$$

are known as the *marginal probabilities*. Two variables X and Y are said to be *independent*, denoted by $X \perp\!\!\!\perp Y$, if the joint probabilities factor as products of marginal probabilities, i.e. we have $p_{ij} = p_{i+}p_{+j}$ for every $(i, j) \in [r_1] \times [r_2]$. An *independence model* $\mathcal{M}_{X \perp\!\!\!\perp Y} \subseteq \Delta_{r_1 r_2 - 1}$ is the set of all probability distributions that admit such a factorization. That is, $\mathcal{M}_{X \perp\!\!\!\perp Y}$ is parametrized as

$$\varphi : \Delta_{r_1 - 1} \times \Delta_{r_2 - 1} \rightarrow \Delta_{r_1 r_2 - 1} : ((s_1, \dots, s_{r_1}), (t_1, \dots, t_{r_2})) \mapsto (p_{ij} = s_i t_j : (i, j) \in [r_1] \times [r_2]).$$

Proposition 1.1.4. [47, Proposition 1.1.2.] The random variables X and Y are independent if and only if the matrix of their joint probabilities $p = (p_{ij})$ has rank 1.

Proof. Suppose X and Y are independent, so every entry of p can be re-written as $p_{ij} = p_{i+}p_{+j}$. This means that $p = vw^T$ is the outer product of two vectors $v = (p_{i+}) \in \mathbb{R}^{r_1}$ and $w = (p_{+j}) \in \mathbb{R}^{r_2}$. Therefore, it has rank 1. Conversely, if p has rank 1 and has nonnegative entries, we may write $p = vw^T$ for some $v \in \mathbb{R}_{\geq 0}^{r_1}$ and $w \in \mathbb{R}_{\geq 0}^{r_2}$. Then $p_{ij} = v_i w_j$, so we have $p_{i+} = v_i w_+$ and $p_{+j} = v_+ w_j$ where v_+ and w_+ are sums of the entries in v and w , respectively. Since (p_{ij}) is a probability distribution, all the entries in p sum to one. It follows that:

$$p_{ij} = v_i w_j = v_i(1)w_j = v_i(p_{++})w_j = v_i(v_+ w_+)w_j = (v_i w_+)(v_+ w_j) = p_{i+}p_{+j},$$

as desired. □

Hence, the implicit description of the model $\mathcal{M}_{X \perp\!\!\!\perp Y}$ is given by the ideal whose generators are the 2-minors of the matrix p . In Chapter 6, we revisit similar equations for *conditionally* independent random variables in the context-specific case. When both X and Y are binary random variables, we have $X \perp\!\!\!\perp Y$ whenever the determinant $p_{11}p_{22} - p_{12}p_{21}$ vanishes. Hence, $\mathcal{M}_{X \perp\!\!\!\perp Y}$ is a surface in the tetrahedron Δ_3 , illustrated in Figure 1.2 on the right.

For models of three or more random variables, the independence relation similarly translates to the algebraic condition that the *tensor* of joint probabilities has rank 1. We return to the model of three independent random variables in Example 1.1.7.

Mixture models. Another important example of statistical models we will encounter in this thesis are mixture models. Such models are a subclass of hidden variable models, which

arise in the presence of data that we would like to measure but are not able to do so directly. These hidden variables, however, may have an impact on the quantities we do observe and how these quantities interact.

Formally, let $\mathcal{M} \subset \Delta_{n-1}$ be a discrete statistical model. Let X be a random variable modeled by \mathcal{M} and let H be a hidden random variable with the state space $[k]$. Suppose that $\mathbb{P}(H = i) = \pi_i$ and $\mathbb{P}(X = j | H = i) = p_j^i$ for some $\pi = (\pi_1, \dots, \pi_k) \in \Delta_{k-1}$ and some $\mathbf{p}^i = (p_1^i, \dots, p_n^i) \in \mathcal{M}$. Since H is a hidden variable we cannot measure, we are interested in the distribution of X only. Note that

$$\mathbb{P}(X = j) = \sum_{i=1}^k \mathbb{P}(X = j \text{ and } H = i) = \sum_{i=1}^k \mathbb{P}(X = j | H = i) \mathbb{P}(H = i) = \sum_{i=1}^k \pi_i p_j^i,$$

so X has a distribution given by the convex combination $\pi_1 \mathbf{p}^1 + \dots + \pi_k \mathbf{p}^k$ of distributions in \mathcal{M} . This motivates the definition of the k th mixture model of $\mathcal{M} \subset \Delta_{n-1}$ as

$$\text{Mixt}^k(\mathcal{M}) = \{\pi_1 \mathbf{p}^1 + \dots + \pi_k \mathbf{p}^k : \pi \in \Delta_{k-1}, \mathbf{p}^1, \dots, \mathbf{p}^k \in \mathcal{M}\}.$$

Algebraically, such models are related to secant varieties, upon taking Zariski closure. A k th secant variety [114, Chapter 14] of some variety V , denoted by $\sigma_k(V)$ is the Zariski closure of all points that lie on affine linear planes that are spanned by k points in V . Namely, it is

$$\sigma_k(V) := \overline{\{\alpha_1 v^1 + \dots + \alpha_k v^k : \sum \alpha_i = 1 \text{ and } v^i \in V \text{ for all } i\}}.$$

One may also generalize the definition of a mixture model above by relaxing the requirement that all the distributions $\mathbf{p}^1, \dots, \mathbf{p}^k$ belong to the same underlying model \mathcal{M} . Such mixture models are algebraically related to more general join varieties.

Bob's coin flips we will see in Example 2.3.3 give rise to a mixture model, since the outcome of the first flip is not recorded, yet may impact the overall result of the experiment. Mixtures of products of continuous random variables are studied extensively in Chapter 7. The next example illustrates the non-trivial geometric aspects that arise in the study of mixture models, even when the underlying model is the independence model. A matrix is said to be *nonnegative* if all of its entries are nonnegative. The *nonnegative rank* of a nonnegative matrix M is the minimal number r of nonnegative rank one matrices that sum to M .

Example 1.1.5 (Mixtures of independence models). Consider the third mixture of the independence model of two variables $\text{Mixt}^3(\mathcal{M}_{X \perp Y})$ where X and Y take values in $[r_1]$ and $[r_2]$, respectively. In Example 1.1.3, we saw that the ideal of the independence model $\mathcal{M}_{X \perp Y}$ consists of probability matrices $p \in \Delta_{r_1 r_2 - 1}$ that have rank 1. The set of all convex combinations of three rank one $r_1 \times r_2$ probability matrices is precisely the set of $r_1 \times r_2$ probability matrices of nonnegative rank at most 3. This set is our model $\text{Mixt}^3(\mathcal{M}_{X \perp Y})$

by definition. The third secant variety of the Zariski closure of our model, denoted by $\sigma_3(\overline{\mathcal{M}_{X \perp Y}})$, consists of all $r_1 \times r_2$ matrices whose rank is at most 3. However, the matrix

$$\frac{1}{8} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 \end{pmatrix}$$

has rank 3 and nonnegative rank 4, see [117]. Hence $\text{Mixt}^3(\mathcal{M}_{X \perp Y}) \neq \sigma_3(\overline{\mathcal{M}_{X \perp Y}}) \cap \Delta_{r_1 r_2 - 1}$. But since the rank of a matrix is always bounded above by its nonnegative rank, we have the “ \subseteq ” containment. In particular, this means that the mixture model $\text{Mixt}^3(\mathcal{M}_{X \perp Y})$ has a nontrivial boundary [114, Example 14.1.6].

Maximum likelihood estimation

Example 1.1.6 (Finite model). You and your friends are at an algebraic statistics conference in Hawai’i. Close to the venue, there are five stands that sell ice cream cones of $n = 3$ flavors: chocolate, strawberry, and vanilla. Each of the stands has some distribution of flavors $(p_1, p_2, p_3) \in \Delta_2$. Hence, our statistical model consists of five such points in the simplex Δ_2 , perhaps

$$\mathcal{M} = \{(0.5, 0.25, 0.25), (0.7, 0.1, 0.2), (0.15, 0.35, 0.6), (0.3, 0.4, 0.3), (0.1, 0.1, 0.9)\}.$$

A friend approaches you with $N = 10$ cones that she claims to have picked randomly at one of the stands. You both look at the cones and observe the following data:



Counting the number of cones of each flavor, one may choose to represent the data above mathematically as the vector $u = (2, 5, 3) \in \mathbb{N}^3$. As statisticians, one question we may ask is which one of the five stands did the ten cones most likely come from? To answer this, we would need to maximize the likelihood function

$$L_u(p) = \frac{10!}{2! 5! 3!} p_1^2 p_2^5 p_3^3$$

over the five distributions in our model. Comparing the values of $L_u(p)$ at the five points in \mathcal{M} , we find that our 10 cones most likely came from the third stand, with the distribution of flavors $(0.15, 0.35, 0.6)$.

Moving beyond the triangle of ice cream flavors, let us fix a discrete statistical model $\mathcal{M} \in \Delta_{n-1}$. Suppose that we observe some data vector $u \in \mathbb{N}^n$ of sample size $N = u_1 + \cdots + u_n$ and we wish to find out which point q best explains this particular data under the assumed

statistical model \mathcal{M} . Note first that, dividing by the total number of observations N , we may assume that u is a distribution in Δ_{n-1} . Then, our task is to maximize the *likelihood function* $L_u(p) = \prod_{i=1}^n p_i^{u_i}$ over all points $p \in \mathcal{M}$. Note that we may ignore the multinomial coefficient we saw in the previous example, since it does not impact which distribution maximizes the function. For the same reason, it is convenient to take the logarithm of the objective function. This leads us to introduce the *log-likelihood function* given by u as

$$\ell_u(p) = u_1 \log p_1 + u_2 \log p_2 + \cdots + u_n \log p_n.$$

The *maximum likelihood estimation* problem is the following optimization problem:

$$\text{Maximize } \ell_u(p) \text{ subject to } p \in \mathcal{M}.$$

Any global maximizer of $\ell_u(p)$ over \mathcal{M} is called a *maximum likelihood estimate (MLE)* of u .

Maximum likelihood estimation is a powerful statistical method of estimating parameters, given some observed data. It is widely used for statistical inference in data analysis.

Likelihood equations. For parametric statistical models, finding MLE can be done by solving a simultaneous system of algebraic equations, known as likelihood equations. Let \mathcal{M} be a d -dimensional statistical model with a parametrization

$$\varphi : \Theta \rightarrow \Delta_{n-1} : \theta = (\theta_1, \theta_2, \dots, \theta_d) \mapsto (g_1(\theta), g_2(\theta), \dots, g_n(\theta)), \quad (1.1)$$

where each g_i is a rational function in θ . We may express the likelihood function as a function of θ , namely $\ell_u(\theta) = \sum_{i=1}^n u_i \log g_i(\theta)$. In order to find all critical points of the function, we set all partial derivatives to zero. This way, we obtain a system of d equations of the form

$$\sum_{i=1}^n \frac{u_i}{g_i(\theta)} \cdot \frac{\partial g_i(\theta)}{\partial \theta_j} = 0 \quad \text{for } j = 1, \dots, d. \quad (1.2)$$

called the *likelihood equations* of \mathcal{M} .

Example 1.1.7. Consider the independence model $\mathcal{M}_{X \perp\!\!\!\perp Y \perp\!\!\!\perp Z}$ of three binary random variables X, Y , and Z . This is a parametric model given by the map

$$\varphi : \Delta_1 \times \Delta_1 \times \Delta_1 \rightarrow \Delta_7 : (\alpha, \beta, \gamma) \mapsto \left(g_{ijk}(\alpha, \beta, \gamma) = \alpha_i \beta_j \gamma_k : i, j, k = 1, 2 \right).$$

Suppose we collected data, which we summarized in a tensor of counts $u \in \mathbb{N}^{2 \times 2 \times 2}$. The log-likelihood function in this case becomes

$$\begin{aligned} \ell_u(\alpha, \beta, \gamma) &= \sum_{i,j,k=1,2} u_{ijk} \log(\alpha_i \beta_j \gamma_k) \\ &= \sum_{i,j,k=1,2} u_{ijk} \log \alpha_i + \sum_{i,j,k=1,2} u_{ijk} \log \beta_j + \sum_{i,j,k=1,2} u_{ijk} \log \gamma_k \\ &= \sum_{i=1,2} u_{i++} \log \alpha_i + \sum_{j=1,2} u_{+j+} \log \beta_j + \sum_{k=1,2} u_{++k} \log \gamma_k. \end{aligned}$$

where $u_{i++} = \sum_{j,k=1,2} u_{ijk}$ is the marginal probability with respect to X , and the other marginal probabilities are defined similarly. Since $\alpha_2 = 1 - \alpha_1$, $\beta_2 = 1 - \beta_1$, and $\gamma_2 = 1 - \gamma_1$, the likelihood equations are

$$\frac{u_{1++}}{\alpha_1} - \frac{u_{2++}}{1 - \alpha_1} = 0, \quad \frac{u_{+1+}}{\beta_1} - \frac{u_{+2+}}{1 - \beta_1} = 0, \quad \frac{u_{++1}}{\gamma_1} - \frac{u_{++2}}{1 - \gamma_1} = 0.$$

Clearing the denominators, we find the following unique solution,

$$\hat{\alpha}_1 = \frac{u_{1++}}{u_{+++}}, \quad \hat{\beta}_1 = \frac{u_{+1+}}{u_{+++}}, \quad \hat{\gamma}_1 = \frac{u_{++1}}{u_{+++}}.$$

Plugging these parameters in φ , we obtain the MLE of u .

Note that in the above example, we were able to write the MLE of a general data u as a rational function in the coordinates of u . This is not a coincidence, and can be explained by the fact that independence models have maximum likelihood degree one. The formal definition is as follows.

Definition 1.1.8. The *maximum likelihood degree (ML degree)* of an algebraic statistical model \mathcal{M} is the number of complex critical points of ℓ_u on the Zariski closure of \mathcal{M} for generic data $u \in \Delta_{n-1}$.

For a parametric statistical model, the ML degree is the number of complex solutions to the likelihood equations for generic data. It was proven in [48, 66] that a statistical model has ML degree one if and only if it has a rational maximum likelihood estimator. Generalizing Example 1.1.7, we may conclude that independence models have ML degree one. In fact, independence models are a subfamily of decomposable hierarchical log-linear models that are known to have ML degree one [114, Chapter 2]. We study these models (as well as more general reducible models) in Section 4.3.

Example 1.1.9. Consider the curve in Δ_3 parametrized as

$$\varphi : t \mapsto s(1, t, t^2, t^3),$$

where $s = \frac{1}{1+t+t^2+t^3}$. This curve is also sometimes referred to as the twisted cubic curve. Note that the coefficients in this parametrization are different from *our* twisted cubic in Example 1.1.2, which affects the ML degree. Indeed, given $u \in \Delta_3$, the log-likelihood function becomes:

$$\begin{aligned} \ell_u(t) &= u_1 \log(s) + u_2 \log(st) + u_3 \log(st^2) + u_4 \log(st^3) \\ &= (u_1 + u_2 + u_3 + u_4) \log s + (u_2 + 2u_3 + 3u_4) \log t. \end{aligned}$$

Hence, the likelihood equation is

$$\frac{\partial \ell_u}{\partial t} = \frac{u_2 + 2u_3 + 3u_4}{t} - \frac{1 + 2t + 3t^2}{1 + t + t^2 + t^3} = 0.$$

Clearing denominators, we obtain the polynomial

$$3(1 - u_4)t^3 + 2(1 - u_3)t^2 + (1 - u_2)t - (u_2 + 2u_3 + 3u_4) = 0$$

of degree 3 in t . Hence the ML degree of this model is 3.

The ML degree measures the algebraic complexity of maximum likelihood estimation for a given model. For some real-world data, it means choosing the best estimators from potentially many critical points. For statistical models with complicated geometry, it can be a challenging task. In particular, while *clearing denominators* was easy in Examples 1.1.7 and 1.1.9 due to the rational functions g_i being special, in general it can be very complicated. If the polynomials g_i in the parametrization (1.1) are generic, then after clearing denominators in the likelihood equations (1.2), we get the system

$$\sum_{i=1}^n u_i g_1 \cdots g_{i-1} g_{i+1} \cdots g_n \frac{\partial g_i}{\partial \theta_j} = 0 \quad \text{for } j = 1, \dots, d.$$

If θ is a parameter such that $g_i(\theta) = g_k(\theta)$ for some $k \in [n]$, then θ is a solution to the above system, but not a solution to the system of original likelihood equations. Such extraneous solutions arise, for example, in random censoring; see [47, Example 2.1.3].

Implicit models. Whenever we only have the implicit equations of the model, the likelihood equations are no longer defined. However, we may still perform maximum likelihood estimation by constructing the likelihood ideal, which we discuss next [26, Chapter 11], [114, Chapter 7]. For this, we will be working with a model defined by a projective variety $V \subseteq \mathbb{P}^{n-1}$. Suppose that the defining ideal $I_V = \langle f_1, \dots, f_k \rangle$ of our model \mathcal{M} is generated by k homogeneous polynomials and also consider the inhomogeneous polynomial $f_0 = p_1 + \cdots + p_n - 1$ defining the simplex. Let $c = \text{codim}(V)$ as a projective variety in \mathbb{P}^{n-1} .

Maximum likelihood estimation is an optimization problem, so we may use the method of Lagrange multipliers to solve it. Let $\mathcal{J} = (\partial f_i / \partial p_j)$ denote the $(k+1) \times n$ Jacobian matrix. Fixing data $u \in \Delta_{n-1}$, construct the *augmented Jacobian* from \mathcal{J} by prepending the row $\nabla \ell_u = (u_1/p_1, \dots, u_n/p_n)$. Next, clear denominators in this new matrix, which amounts to multiplying the i th column by p_i , and denote the resulting matrix by $\mathcal{A}\mathcal{J}$. If the gradient of the objective function were to lie in the normal space of the model, the matrix $\mathcal{A}\mathcal{J}$ would have rank at most $c+1$. This leads us to define the ideal

$$J_V = \langle (c+2) \times (c+2) \text{ minors of } \mathcal{A}\mathcal{J} \rangle + I_V + \langle f_0 \rangle.$$

However, we must remove the extraneous solutions coming from the singular locus of V and from the points in the hyperplane arrangement $\mathcal{H} = \{p \in \mathbb{P}^{n-1} : p_1 \cdots p_n(p_1 + \cdots + p_n) = 0\}$, where the log-likelihood function is undefined. Let

$$Q = \langle (c+1) \times (c+1) \text{ minors of } \mathcal{J} \rangle \cdot p_1 \cdots p_n(p_1 + \cdots + p_n).$$

The *likelihood ideal* L is obtained by saturating J_V by Q , i.e.

$$L = J_V : Q^\infty.$$

The degree of the likelihood ideal for a generic u is the ML degree of the model. The solutions of L are the complex critical points of the maximum likelihood estimation.

Example 1.1.10. Let's pretend that the parametrization of the curve in Example 1.1.9 is not known, but instead we know the defining equations of this model, which are

$$I_V = \langle p_3^2 - p_2p_4, p_2p_3 - p_1p_4, p_2^2 - p_1p_3 \rangle.$$

Upon clearing the denominators, the augmented Jacobian takes the form

$$\mathcal{AJ} = \begin{pmatrix} u_1 & u_2 & u_3 & u_4 \\ p_1 & p_2 & p_3 & p_4 \\ 0 & -p_2p_4 & 2p_3^2 & -p_2p_4 \\ -p_1p_4 & p_2p_3 & p_2p_3 & -p_1p_4 \\ -p_1p_3 & 2p_2^2 & -p_1p_3 & 0 \end{pmatrix}.$$

Let $u = (13/57, 4/11, 3/13, 1447/8151)$. We may compute the likelihood ideal in Macaulay2 with respect to u using the following code.

```
needsPackage "EigenSolver";
R=QQ[p1,p2,p3,p4];
u1=13/57; u2=4/11; u3=3/13; u4=1-u1-u2-u3;
I=ideal(p3^2-p2*p4,p2*p3-p1*p4,p2^2-p1*p3);
M=matrix {{u1, u2, u3, u4}, {p1, p2, p3, p4}, {0, -p2*p4, 2*p3^2, -p2*p4},
          {-p1*p4, p2*p3, p2*p3, -p1*p4}, {-p1*p3, 2*p2^2, -p1*p3, 0}};
N=submatrix(M, {1,2,3,4}, {0,1,2,3});
Q=minors(3,N)*(p1*p2*p3*p4)*(p1+p2+p3+p4);
J=minors(4,M)+I+ideal(p1+p2+p3+p4-1);
L=saturate(J,Q);
```

```
degree L
zeroDimSolve L
```

We find that L is a zero-dimensional ideal of degree 3 (the ML degree of the model). It has only one solution in the simplex $(.131028, .190386, .276633, .401953)$, which is the MLE of u . The other two solutions are complex.

The fact that the example above had a unique critical point on the model is not a surprise. This is a property of toric models in general. For such models, it is much more efficient to use *Birch's theorem* to find the maximum likelihood estimate, which takes advantage of the convexity of the log-likelihood function. We state and use this important theorem in Chapters 2 and 4, where toric models are discussed in detail.

Kullback-Leibler divergence. So far, we have been maximizing the log-likelihood function to determine a point on the model that best explains our data. However, we could be minimizing the Kullback-Leibler divergence instead. For any two distributions u and w in Δ_{n-1} , the *Kullback-Leibler (KL) divergence*, also known as *information divergence*, from u to w is defined as $D(u||w) = \sum_{i=1}^n u_i \log(u_i/w_i)$. We use the convention that $0 \log 0 = 0 \log(0/0) = 0$ and $D(p || q) = +\infty$ if $\text{supp}(p) \not\subseteq \text{supp}(q)$. This function is non-negative and is zero if and only if $u = w$ [26, Chapter 11]. The maximum likelihood estimation may be equivalently phrased in terms of information divergence, as opposed to the log-likelihood function.

Proposition 1.1.11. The maximum likelihood estimation with respect to fixed data u over a model \mathcal{M} is equivalent to

$$\text{Minimize } D(u||p) \text{ subject to } p \in \mathcal{M}.$$

Proof. Re-write the KL divergence as

$$D(u||p) = \sum_{i=1}^n u_i \log u_i - \sum_{i=1}^n u_i \log p_i = \sum_{i=1}^n u_i \log u_i - \ell_u(p).$$

For fixed data u , the term $\sum_{i=1}^n u_i \log u_i$ is constant, hence minimizing KL divergence is equivalent to minimizing $-\ell_u(p)$, which is equivalent to maximizing $\ell_u(p)$. \square

We call the minimizer of the KL divergence, $D_{\mathcal{M}}(u) := \min_{p \in \mathcal{M}} D(u || p)$, the *information divergence* (or just divergence) from u to \mathcal{M} . In Section 3.3 and Chapter 4, we study $D(\mathcal{M}) = \max_{p \in \Delta_{n-1}} D_{\mathcal{M}}(p)$ and the points which achieve $D(\mathcal{M})$ when \mathcal{M} is a linear or a discrete exponential (toric) model. In other words, we consider the problem of *maximizing* KL divergence from a point u to its MLE over a linear or a toric model \mathcal{M} . This problem is of interest to information geometers and finds applications in bio-neural networks.

In what follows, it is also useful to define the relation known as the *likelihood correspondence* [67, Definition 1.5]. Namely, for $\mathcal{M} \subset \Delta_{n-1}$, we define the relation $\Phi \subset \Delta_{n-1} \times \mathcal{M}$ by

$$(u, p) \in \Phi \iff p \in \text{argmax}_{q \in \mathcal{M}} \{\ell_u(q) : q \in \mathcal{M}\}.$$

If $(u, p) \in \Phi$, then we write $\Phi(u) = p$. We write $\Delta_{n-1}^{\mathcal{M}}$ for the set of $u \in \Delta_{n-1}$ such that $\Phi(u)$ exists. Describing the set $\Delta_{n-1}^{\mathcal{M}}$ and how it extends to the boundary of Δ_{n-1} is an active area of research, especially with respect to zeros in the data [52, 61]. MLE existence is also connected to polystable and stable orbits in invariant theory [11]. For the important family of log-linear (toric) models, studied in Chapters 2 and 4 of the current thesis, [51] shows that positive data $u \in \Delta_{n-1}$ guarantees existence, and in general the MLE exists exactly when the observed margins belong to the relative interior of a certain polytope. See also [114, Theorem 8.2.1].

Finally, we note that for models with more complicated geometry, $\Phi(u)$ cannot always be computed by finding critical points of ℓ_u restricted to manifold points of \mathcal{M} . Chapter 2 of this dissertation provides the first step of computing logarithmic Voronoi cells for models where critical points of ℓ_u succeed in finding the MLE, as well as some interesting finite models. More complicated examples include models of nonnegative rank r matrices [77].

What is a logarithmic Voronoi cell?

For any subset $X \subset \mathbb{R}^n$, the *Voronoi cell* of a point $p \in X$ consists of all points of \mathbb{R}^n which are closer to p than to any other point of X in the Euclidean metric. Voronoi cells find use in many fields of science where analyzing spatially distributed data is of interest. In the first part of this thesis, we discuss the analogous *logarithmic Voronoi cells* which find applications in statistics. Informally, logarithmic Voronoi cells are used to divide experimental data based on which point in the statistical model each sample most likely came from. Information geometry [17] considers maximum likelihood estimation in the context of the Kullback-Leibler divergence of probability distributions, sending data to the nearest point with respect to a Riemannian metric on Δ_{n-1} . Algebraic statistics [47] considers the case where \mathcal{M} can be described as either the image or kernel of algebraic maps. Recent work in metric algebraic geometry [26, 39, 45, 62] concerns the properties of real algebraic varieties that depend on a distance metric. Logarithmic Voronoi cells are natural objects of interest in all three subjects.

Let us informally showcase the mathematical objects defined later in this section for the twisted cubic curve, first introduced in Section 1.1.

Example 1.1.12. Consider flipping a biased coin three times. There are four possible outcomes, 3 heads (hhh), 2 heads (hht,hth,thh), 1 head (htt,tht,tth), and 0 heads (ttt). This is the *twisted cubic* we encountered in Example 1.1.2, which has the parametrization

$$t \mapsto p(t) = (t^3, 3t^2(1-t), 3t(1-t)^2, (1-t)^3) \in \mathcal{M}.$$

For this model's many lives, see [79]. We compute and plot logarithmic Voronoi cells $\log \text{Vor}_{\mathcal{M}}(p(t))$ for sampled parameter values

$$t \in \left\{ \frac{1}{25}, \frac{2}{25}, \dots, \frac{24}{25} \right\}$$

which live inside the simplex $\Delta_3 \subset \mathbb{R}^4$, and whose orthogonal projections into 3-space are shown in Figure 1.3. In this case, the logarithmic Voronoi cells are polytopes, and we get both triangles and quadrilaterals, depending on the point $p(t) \in \mathcal{M}$. The fact that these polytopes are equal to the logarithmic Voronoi cells will follow from Theorem 2.1.10 in Section 2.1. The $\Delta/\square/\Delta$ combinatorial type change is explained by the chamber complex of the model, discussed in detail in Section 4.2. Chapters 2 and 4 will describe how to compute these polytopes for any parameter value t .




Figure 1.3: Logarithmic Voronoi cells of the twisted cubic.

In general, let $\mathcal{M} \subset \Delta_{n-1}$ be a discrete statistical model, as in Section 1.1 and let $p \in \mathcal{M}$ be a distribution in the model. The *logarithmic Voronoi cell* at p is

$$\log \text{Vor}_{\mathcal{M}}(p) := \{u \in \Delta_{n-1} : \Phi(u) = p\}.$$

Whenever $p \in \mathcal{M} \subset \mathbb{R}^n$ admits a tangent space at the point p , we denote by $N_p\mathcal{M}$ its orthogonal complement with respect to the Euclidean inner product on \mathbb{R}^n . Similarly, we are also interested in the *log-normal space* at the point $p \in \mathcal{M}$, defined by

$$\log N_p\mathcal{M} := \{u \in \mathbb{R}^n : \nabla \ell_u(p) \in N_p\mathcal{M}\}.$$

Here, $\nabla \ell_u(p)$ is the vector whose entries are given by the partial derivatives of ℓ_u with respect to each of the variables p_1, \dots, p_n . We are interested in the log-normal space since, in many cases, it will contain the logarithmic Voronoi cell. In Section 2.1 we will see several different situations where the logarithmic Voronoi cell is equal to the intersection of the log-normal space with the probability simplex.

Lemma 1.1.13. The log-normal space $\log N_p\mathcal{M}$ is a linear subspace of \mathbb{R}^n .

Proof. The normal space $N_p\mathcal{M}$ is a linear subspace. Arrange a basis of the normal space as the rows of a matrix. Adjoin another row with entries u_i/p_i , the partial derivatives of $\ell_u(p)$

with respect to each p_i . The maximal minors of the resulting matrix are linear equations in the variables u_i and therefore cut out a linear space of such $u \in \mathbb{R}^n$. This space is the log-normal space at p . \square

By Lemma 1.1.13, the intersection of the log-normal space at a point $p \in \mathcal{M}$ with the probability simplex Δ_{n-1} is a polytope $\log \text{Poly}(p)$, which we call its *log-normal polytope*. In Section 2.1, we will investigate when a logarithmic Voronoi cell equals its log-normal polytope. The former is always contained in the latter.

Conditional independence

In this section, we introduce conditional independence for discrete random variables and describe the algebraic conditions that arise in this case. We define conditional independence ideals and briefly discuss their algebraic properties. We then focus on graphical models by introducing Markov properties for directed acyclic graphs and explaining how they relate to distributions in the corresponding model. We also explain how directed graphical models are parametrized via the recursive factorization property and what is known about their ideals, thus motivating the technical algebraic discussion in Section 6.4.

The independence assumption is often too strong in practice. Sometimes random variables are not entirely independent, like in Example 1.1.3, but are unrelated given the knowledge of the values of some other random variables.

Example 1.1.14. One amusing classical example from [89] inquires whether watching soccer causes hair loss. Consider the random variable X_1 that measures the length of a person's hair. We assume X_1 is discrete and has four states: bald, short, medium, and long. Also consider the ternary random variable X_2 , which measures how often a person watches soccer. The three states of X_2 are: never, sometimes, and often. While hair length and soccer watching habits seem unrelated, they are not independent! If we were to sample 300 people, then record their hair length and how often they watch soccer, we would find that people with shorter hair tend to watch soccer more often. Indeed, there is a third random variable G involved, namely gender. If we were to restrict our survey to only individuals of a particular gender, we would see that it is reasonable to assume the independence of X_1 and X_2 . We say that X_1 and X_2 are conditionally independent given G , denoted by $X_1 \perp\!\!\!\perp X_2 \mid G$. If the random variable G is not observed (hidden), we obtain a mixture model from Section 1.1.

In general, we start with m discrete random variables X_1, \dots, X_m . We assume for all $i \in [m]$, the random variable X_i is d_i -ary, taking values in the set $[d_i]$. We then consider the random vector $X = (X_1, \dots, X_m)$ with the state space $\mathcal{R} = \prod_{i=1}^m [d_i]$. For any subset $A \subset [m]$, we may define $X_A = (X_i : i \in A)$ and $\mathcal{R}_A := \prod_{a \in A} [d_a]$. As before, for any $i \in \mathcal{R}$, we let $p_i = \mathbb{P}(X = i)$. Moreover, similarly to Example 1.1.3, we let $p_{i_A+} = \mathbb{P}(X_A = i_A)$. For any pairwise disjoint subsets $A, B, C \subset [m]$, we will use the convention that the string of indices

$i \in \mathcal{R}$ is grouped into the string $(i_A, i_B, i_C, j_{[m] \setminus A \cup B \cup C})$. We say that the random vector X_A is *conditionally independent* of X_B given X_C whenever $p_{i_A i_B i_C} = p_{i_A + i_C + j_{[m] \setminus A \cup B \cup C}}$. We denote it by $X_A \perp\!\!\!\perp X_B \mid X_C$, and often abbreviate it as $A \perp\!\!\!\perp B \mid C$. Such relations between sets of random variables are called *conditional independence (CI) statements*.

The following four important axioms follow from the definition of conditional independence [114, Proposition 4.1.4]. They state that if $A, B, C, D \subset [m]$ are pairwise disjoint then:

1. $X_A \perp\!\!\!\perp X_B \mid X_C \implies X_B \perp\!\!\!\perp X_A \mid X_C$ (symmetry);
2. $X_A \perp\!\!\!\perp X_{B \cup D} \mid X_C \implies X_B \perp\!\!\!\perp X_A \mid X_C$ (decomposition);
3. $X_A \perp\!\!\!\perp X_{B \cup D} \mid X_C \implies X_A \perp\!\!\!\perp X_B \mid X_{C \cup D}$ (weak union);
4. $X_A \perp\!\!\!\perp X_B \mid X_{C \cup D}$ and $X_A \perp\!\!\!\perp X_D \mid X_C \implies X_A \perp\!\!\!\perp X_{B \cup D} \mid X_C$ (contraction).

For positive distributions only, an additional intersection axiom holds [114, Proposition 4.1.5]:

$$X_A \perp\!\!\!\perp X_B \mid X_{C \cup D} \text{ and } X_A \perp\!\!\!\perp X_C \mid X_{B \cup D} \implies X_A \perp\!\!\!\perp X_{B \cup C} \mid X_D.$$

In Chapter 6, we will study even more refined independence relations, which depend on the specific values of the random variables in the conditioning set. For those *context-specific* independence relations, two more axioms will be important: absorption and specialization [50].

Conditional independence relations, just like independence relations in Proposition 1.1.4, can be described by algebraic equations.

Proposition 1.1.15. [114, Proposition 4.1.6] Let $A, B, C \subset [m]$ be pairwise disjoint. Then $X_A \perp\!\!\!\perp X_B \mid X_C$ if and only if

$$p_{i_A i_B i_C} \cdot p_{j_A j_B i_C} - p_{i_A j_B i_C} \cdot p_{j_A i_B i_C} = 0 \tag{1.3}$$

for all $i_A, j_A \in \mathcal{R}_A$, $i_B, j_B \in \mathcal{R}_B$ and $i_C \in \mathcal{R}_C$.

Proof. Note that marginalization with respect to $A \cup B \cup C$ is present both in the definition of conditional independence and in the binomials (1.3) so we may assume $A \cup B \cup C = [m]$. Moreover, since the value i_C of X_C is fixed, we may also assume that C is empty. Aggregating the states indexed by A and also the states indexed by B , we may further work with the independence statement $X_1 \perp\!\!\!\perp X_2$. By Proposition 1.1.4, this statement holds if and only if the matrix of joint probabilities of X_1 and X_2 has rank one. After accounting for marginalization, conditioning, and state aggregation, the 2×2 minors of this matrix are precisely the equations in (1.3), and the claim follows. \square

The *conditional independence ideal* $I_{A \perp\!\!\!\perp B | C}$ is the ideal generated by all quadrics in (1.3). This is the ideal of all 2×2 minors of $|\mathcal{R}_C|$ matrices of size $|\mathcal{R}_A| \times |\mathcal{R}_B|$. This ideal is always prime. A *conditional independence model* is a family of distributions that satisfy a collection of CI statements $\mathcal{C} = \{A_1 \perp\!\!\!\perp B_1 | C_1, A_2 \perp\!\!\!\perp B_2 | C_2, \dots\}$, where A_k, B_k, C_k are pairwise disjoint for each k . The conditional independence ideal of the collection \mathcal{C} is $I_{\mathcal{C}} = I_{A_1 \perp\!\!\!\perp B_1 | C_1} + I_{A_2 \perp\!\!\!\perp B_2 | C_2} + \dots$. These ideals $I_{\mathcal{C}}$ can be used to study implications between CI statements in \mathcal{C} , via primary decomposition [47, Chapter 3].

Example 1.1.16. Consider the model given by two CI relations between three random variables

$$\mathcal{C} = \{X_1 \perp\!\!\!\perp X_3, X_1 \perp\!\!\!\perp X_3 | X_2\}.$$

These are the same CI statements that will define the Gaussian model in Example 1.2.8. For now, however, we will assume that all three variables are binary taking values in $\{1, 2\}$. Using Proposition 1.1.15, we get

$$\begin{aligned} I_{\mathcal{C}} &= I_{X_1 \perp\!\!\!\perp X_3} + I_{X_1 \perp\!\!\!\perp X_3 | X_2} \\ &= \langle (p_{111} + p_{121})(p_{212} + p_{222}) - (p_{112} + p_{122})(p_{211} + p_{221}), \\ &\quad p_{111}p_{212} - p_{112}p_{211}, p_{121}p_{222} - p_{122}p_{221} \rangle. \end{aligned}$$

Computing the primary decomposition of $I_{\mathcal{C}}$, we find that it is given as an intersection of two prime ideals. Upon further inspection, we recognize the two prime components as conditional independence ideals $I_{1 \perp\!\!\!\perp \{2,3\}}$ and $I_{\{1,2\} \perp\!\!\!\perp 3}$, i.e.

$$I_{\mathcal{C}} = I_{1 \perp\!\!\!\perp \{2,3\}} \cap I_{\{1,2\} \perp\!\!\!\perp 3}.$$

This, in turn, implies $\mathcal{V}(I_{\mathcal{C}}) = \mathcal{V}(I_{1 \perp\!\!\!\perp \{2,3\}}) \cup \mathcal{V}(I_{\{1,2\} \perp\!\!\!\perp 3})$ on the level of varieties. Therefore, we conclude that $X_1 \perp\!\!\!\perp X_3$ and $X_1 \perp\!\!\!\perp X_3 | X_2$ together imply that either $X_1 \perp\!\!\!\perp X_{\{2,3\}}$ or $X_{\{1,2\}} \perp\!\!\!\perp X_3$.

In general, however, the primary decomposition of $I_{\mathcal{C}}$ is not guaranteed to contain components that are conditional independence ideals. Sometimes, these components will define varieties on the boundary of the simplex. Moreover, the components in the decomposition need not be prime ideals, only primary. This often makes it challenging to interpret and extract information about CI statement implications. See [47, 114] for more examples.

Graphical models. Chapter 6 of this dissertation will be concerned with a special class of conditional independence models, called graphical models. These models are defined by conditional independence constraints according to the non-adjacencies of a graph G . These constraints are known as Markov properties of G . One reason for studying graphical representations is that the associated statistical models often have a natural parametrization, induced by the combinatorial structure of the underlying graph. Moreover, since such models are parametrized, the corresponding varieties are irreducible and their ideals are prime.

Understanding how these prime ideals relate to the Markov properties of the graph helps us characterize discrete graphical models both combinatorially and algebraically.

While undirected graphical models are an important subclass of hierarchical log-linear models, studied in Section 4.3, we will not explicitly focus on their ideals in this dissertation. However, the combinatorics and algebra of graphical models corresponding to directed acyclic graphs (DAGs) plays a big role in the analysis of decomposable context-specific models in Section 6.1. Next, we introduce these models and state their main properties.

Let $G = (V, E)$ be a directed graph with no directed cycles. Such a graph is called a *directed acyclic graph*, abbreviated by DAG. Consider a random vector $X = (X_v)_{v \in V}$ that is indexed by the nodes of the graph. For any vertex $v \in V$, let $\text{pa}(v)$ denote the set of *parents* of v , i.e. all the vertices $w \in V$ such that $(w, v) \in E$. Let $\text{an}(v)$ denote the *ancestors* of v , i.e. all $w \in V$ such that there is a directed path from w to v in E . Similarly, let $\text{de}(v)$ denote the *descendants* of v , i.e. all $w \in V$ such that there is a directed path from v to w in E . Finally, the *non-descendants* of v are $\text{nd}(v) = V \setminus (\{v\} \cup \text{de}(v))$.

The *local Markov property* of G is the collection of CI statements

$$\text{local}(G) = \{X_v \perp\!\!\!\perp X_{\text{nd}(v)} \mid X_{\text{pa}(v)} : v \in V\}.$$

These statements reflect the expected independence structure if the edges of the DAG represent parent-child or cause-effect relationships. However, they may imply other CI statements. Hence, to study more general implications of these local constraints, we must introduce the global Markov property, defined via the technical notion of d -separation in a DAG.

Let $\pi = (v_0, \dots, v_k)$ be an undirected path of edges in G . For some $i \in [k - 1]$, we say v_i is a *collider* if the edges incident to it on the path π are of the form $v_{i-1} \rightarrow v_i \leftarrow v_{i+1}$. The induced directed subpath is called a v -*structure*. We say that two vertices v and w of a DAG G are d -*connected* given $C \subset V \setminus \{v, w\}$ if there is an undirected path π from v to w such that the following two conditions are satisfied:

1. all colliders of π are in $\text{an}(u)$ for some $u \in C$;
2. no non-collider on π is in C .

For pairwise disjoint $A, B, C \subset V$ with A and B nonempty, we say that C d -*separates* A and B if no pair of vertices $a \in A$ and $b \in B$ are d -connected given C . The *global Markov property* of the DAG G is the collection

$$\text{global}(G) = \{X_A \perp\!\!\!\perp X_B \mid X_C : C \text{ } d\text{-separates } A \text{ and } B\}.$$

For example, in the DAG $X_1 \rightarrow X_2 \rightarrow X_3$, we see that X_2 d -separates X_1 and X_3 . On the other hand, in the v -structure $X_1 \rightarrow X_2 \leftarrow X_3$, the empty set separates X_1 and X_3 .

The global Markov property of G often includes more statements than the local one. In general, we have the relationship $I_{\text{local}(G)} \subseteq I_{\text{global}(G)}$ for ideals and the reverse containment for varieties. However, on the level of probability distributions, the two Markov properties are equivalent.

We say that a joint distribution obeys a Markov property if it exhibits the conditional independence statements in the list of the associated CI constraints.

Theorem 1.1.17. [114, Theorem 13.1.11] A joint distribution of the random vector X obeys the directed local Markov property for the directed acyclic graph G if and only if it obeys the directed global Markov property for G .

Therefore, one could define a statistical model associated to a DAG G to be the set of all probability distributions that obey $\text{local}(G)$ or $\text{global}(G)$. However, we may parametrize the model $\mathcal{M}(G)$ as follows. Let $f \in \Delta_{|\mathcal{R}|-1}$ be a distribution, and let $f(i_A | i_B) = \mathbb{P}(X_A = i_A | X_B = i_B)$ denote conditional probabilities. We say that f is *Markov to G* if it satisfies the following recursive factorization property for all $i \in \mathcal{R}$:

$$p_i = f(i) = \prod_{v \in V} f(i_v | i_{\text{pa}(v)}). \tag{1.4}$$

The *DAG model* $\mathcal{M}(G)$ is the set of all the distributions in $\Delta_{|\mathcal{R}|-1}^\circ$ that factorize as the product of conditionals in (1.4). Models associated to DAGs are also called *Bayesian networks*. Such parametrized models are given by irreducible varieties and prime ideals. Finding generators of the defining prime ideal P_G is the strongest possible characterization of a model. The ideal P_G can also be described as the kernel of a map φ that sends the coordinates p_i in the distribution f to their parametrization in (1.4). We discuss this map in detail for stage tree models in Chapter 6.

Example 1.1.18. Consider the simple DAG $G = X_1 \rightarrow X_2$ where both X_1 and X_2 are binary random variables. The factorization property states that the distributions $p = (p_{11}, p_{12}, p_{21}, p_{22})$ in this model satisfies $p_{i_1 i_2} = f(i_1 i_2) = f(i_1) f(i_2 | i_1)$. Letting $s_i = \mathbb{P}(X_1 = i)$ and $t_{jk} = \mathbb{P}(X_2 = k | X_1 = j)$, we see that $\mathcal{M}(G)$ is given as the image of the following rational map

$$\Delta_1 \times \Delta_1 \times \Delta_1 \rightarrow \Delta_3 : ((s_1, s_2), (t_{11}, t_{12}), (t_{21}, t_{22})) \mapsto (s_1 t_{11}, s_1 t_{12}, s_2 t_{21}, s_2 t_{22}).$$


It is easily checked that the map is surjective, so $\mathcal{M}(G) = \Delta_3$. On the other hand, if we take the graph $H = X_1 \quad X_2$ with no edges, then the parametrization (1.4) shows that

$$p_{i_1 i_2} = f(i_1 i_2) = f(i_1) f(i_2) = p_{i_1 + i_2}.$$

Hence, $\mathcal{M}(H)$ is the familiar independence model from Example 1.1.3.

In this section we introduced three ideals: $I_{\text{local}(G)}$, $I_{\text{global}(G)}$, and P_G , and it is natural to inquire how they are related. In general, the prime ideal P_G need not be equal to $I_{\text{local}(G)}$ or $I_{\text{global}(G)}$. However, it always appears as a minimal prime in the primary decompositions of *both* $I_{\text{local}(G)}$ and $I_{\text{global}(G)}$ [53, Theorem 8]. Section 6.4 dives deeper into the relationship between these ideals for decomposable context-specific models. It also highlights the importance of the ideal corresponding to all saturated conditional independence statements.

Example 1.1.19. [47, Example 3.3.11] Consider the DAG



where all four variables are binary. The local Markov property of G consists of two CI statements

$$\text{local}(G) = \{X_1 \perp\!\!\!\perp X_4 \mid X_{\{2,3\}}, X_2 \perp\!\!\!\perp X_3 \mid X_1\}.$$

In this case, the ideal $I_{\text{local}(G)}$ is itself prime of dimension 9 and the variety $\mathcal{V}(I_{\text{local}(G)})$ is irreducible with no components on the boundary of the simplex. In this case, both ideals are equal:

$$\begin{aligned} I_{\text{local}(G)} = I_{\text{global}(G)} = \langle & (p_{1111} + p_{1112})(p_{1221} + p_{1222}) - (p_{1121} + p_{1122})(p_{1211} + p_{1212}), \\ & (p_{2111} + p_{2112})(p_{2221} + p_{2222}) - (p_{2121} + p_{2122})(p_{2211} + p_{2212}), \\ & p_{1111}p_{2112} - p_{1112}p_{2111}, p_{1121}p_{2122} - p_{1122}p_{2121}, \\ & p_{1211}p_{2212} - p_{1212}p_{2211}, p_{1221}p_{2222} - p_{1222}p_{2221} \rangle. \end{aligned}$$

The recursive parametrization property induces a map ψ that sends each coordinate p_{ijkl} to its parametrization. For example, $\psi(p_{1212}) = a(1 - b_1)c_1(1 - d_{21})$ where $a = \mathbb{P}(X_1 = 1)$, $b_1 = \mathbb{P}(X_2 = 1 \mid X_1 = 1)$, $c_1 = \mathbb{P}(X_3 = 1 \mid X_1 = 1)$, and $d_{21} = \mathbb{P}(X_3 = 1 \mid X_2 = 2, X_3 = 1)$, and so on. Computing the kernel of ψ in $\mathbb{R}[p_{ijkl} : i, j, k, \ell \in \{1, 2\}]$, we find that

$$P_G = \ker \psi = I_{\text{local}(G)} + \left\langle \sum_{i,j,k,\ell} p_{ijkl} - 1 \right\rangle$$

describes all algebraic relations among coordinates of the probability distributions in $\mathcal{M}(G)$.

1.2 Continuous models

This section follows a similar structure as Section 1.1, except we focus on continuous distributions. We define Gaussian models, both generally and inside the positive definite cone. We then describe the maximum likelihood estimation problem for such models and define logarithmic Voronoi cells. Finally, we focus on conditional independence and the corresponding algebraic equations that arise in the continuous setting.

Whenever a random variable follows a distribution that has a density function, it is called *continuous*. In Part III of this thesis, we will focus on general continuous random variable on \mathbb{R} , with no assumptions on the distribution. In this sense, our setup will be nonparametric, and we will give all relevant definitions and motivations in Chapter 7. However, multivariate normal (Gaussian) distributions are more well-studied in algebraic statistics. Hence, this section will focus on the central concepts we will need to study logarithmic Voronoi cells for Gaussian models in Chapter 5. Our exposition in this chapter follows [5, 47, 114].

Gaussian models

The motivation for studying Gaussian models comes from several fields of science. For example, Gaussian mixture models find use in phylogenetics, while Gaussian graphical models are used in computational biology and economics, where leveraging their geometric structure aids parameter estimation. Furthermore, Gaussian distributions exhibit maximum entropy among all real-valued distributions with a specified mean and covariance. Consequently, the Gaussianity assumption places the fewest structural constraints beyond those given by the first and second moments [119].

Let $X = (X_1, \dots, X_m)$ be an m -dimensional Gaussian random vector, which has the density function

$$p_{\mu, \Sigma}(x) = \frac{1}{(2\pi)^{m/2}(\det \Sigma)^{1/2}} \exp \left\{ -\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu) \right\}, \quad x \in \mathbb{R}^m$$

with respect to two parameters: the mean vector $\mu \in \mathbb{R}^m$ and the covariance matrix $\Sigma \in \text{PD}_m$. Here, PD_m is the cone of real symmetric positive definite $m \times m$ matrices. Here, the mean vector μ determines the center of the distribution, while the covariance matrix Σ determines the distribution's spread. Such X is said to be distributed according to the *Gaussian distribution*, denoted by $\mathcal{N}(\mu, \Sigma)$. When $m = 1$, we recover the familiar normal random variable, geometrically represented by the famous bell curve.

For $\Theta \subseteq \mathbb{R}^m \times \text{PD}_m$, the statistical model

$$\mathcal{P}_\Theta = \{\mathcal{N}(\mu, \Sigma) : \theta = (\mu, \Sigma) \in \Theta\}$$

is called a *Gaussian model*. Since the parameter space Θ completely determines the model, we will use Θ and \mathcal{P}_Θ interchangeably. The special model that fills the whole ambient space

is called the *saturated Gaussian model* $\Theta = \mathbb{R}^m \times \text{PD}_m$. General Gaussian models can have a complicated geometry, and the maximum likelihood estimation for them (defined in the next section) can be quite challenging. Moreover, note that every Gaussian model $\Theta \subseteq \mathbb{R}^m \times \text{PD}_m$ has two components: the Euclidean component contained in \mathbb{R}^m and the component contained in the positive definite cone PD_m . Specializing to each of these components, we will instead consider simpler Gaussian models of the forms $\Theta = \Theta_1 \times \{I_m\}$ and $\Theta = \mathbb{R}^m \times \Theta_2$, where I_m denotes the $m \times m$ identity matrix. Chapter 5 is primarily devoted to the Gaussian models of the second kind.

Example 1.2.1 (Nodal cubic). Consider the Gaussian model $\Theta = \Theta_1 \times \{I_2\}$, where $\Theta_1 = \{(t^2 - 1, t(t^2 - 1)) \in \mathbb{R}^2 : t \in \mathbb{R}\}$. Then Θ is the nodal cubic curve in \mathbb{R}^2 , plotted on the left in Figure 1.4. It has a singularity at $t = \pm 1$. We will see in the next section that maximum likelihood estimation on such models is equivalent to solving a least-squares problem.

Example 1.2.2 (Bivariate correlation). Now consider a Gaussian model $\Theta = \mathbb{R}^m \times \Theta_2$ where

$$\Theta_2 = \left\{ \Sigma_x := \begin{pmatrix} 1 & x \\ x & 1 \end{pmatrix} : x \in (-1, 1) \right\} \subseteq \text{PD}_2.$$

This model is the *bivariate correlation model*. It is a one-dimensional model inside the three-dimensional positive definite cone

$$\text{PD}_2 = \left\{ \begin{pmatrix} y & x \\ x & z \end{pmatrix} \in \mathbb{R}^{2 \times 2} : y > 0 \text{ and } yz - x^2 > 0 \right\}.$$

It is plotted on the right of Figure 1.4, using Mathematica [68]. Maximum likelihood estimation for Θ was studied in [12]. Despite the simplicity of the above description, it is a nontrivial task. Section 5.3 of this thesis is devoted to completely characterizing logarithmic Voronoi cells of the bivariate correlation model.

Maximum likelihood estimation

Let $\Theta \subseteq \mathbb{R}^m \times \text{PD}_m$ be a Gaussian model. Suppose we sampled some data and recorded it in n vectors $X^{(1)}, \dots, X^{(n)} \in \mathbb{R}^m$. We then define the sample mean and the sample covariance matrix as

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X^{(i)} \quad \text{and} \quad S = \frac{1}{n} \sum_{i=1}^n (X^{(i)} - \bar{X})(X^{(i)} - \bar{X})^T,$$

respectively. We will identify each sample $X^{(1)}, \dots, X^{(n)}$ with the tuple (\bar{X}, S) and consider any two samples whose sample mean and sample covariance are equal to be the same. Throughout this section, we fix a positive integer n , which denotes the sample size. Given our n sampled data vectors, the *log-likelihood function*, up to an additive constant, is defined as

$$\ell_n(\mu, \Sigma) = -\frac{n}{2} \log \det \Sigma - \frac{n}{2} \text{tr} (S \Sigma^{-1}) - \frac{n}{2} (\bar{X} - \mu)^T \Sigma^{-1} (\bar{X} - \mu). \quad (1.5)$$




Figure 1.4: Nodal cubic in \mathbb{R}^2 (left) and the bivariate correlation model in PD_3 (right).

For a fixed model Θ , the sample mean \bar{X} , and the sample covariance S , the *maximum likelihood estimation* is defined similarly to the discrete case:

$$\text{Maximize } \ell_n(\mu, \Sigma) \text{ subject to } (\mu, \Sigma) \in \Theta.$$

Any global maximizer $\hat{\theta} = (\hat{\mu}, \hat{\Sigma})$ of $\ell_n(\mu, \Sigma)$ over Θ is a *maximum likelihood estimate (MLE)* of (\bar{X}, S) . When $\Theta = \mathbb{R}^m \times \text{PD}_m$ is a saturated model, every sample (\bar{X}, S) lies on the model and is its own MLE. For the models of the form $\Theta = \Theta_1 \times \{I_m\} \subseteq \mathbb{R}^m \times \text{PD}_m$, the maximum likelihood estimation turns into a least-squares problem.

Proposition 1.2.3. [47, Proposition 2.1.10] Let $\Theta = \Theta_1 \times \{I_m\}$ be a Gaussian model. Then the maximum likelihood estimate of the mean parameter $\bar{X} \in \mathbb{R}^m$ is the point $\hat{\mu} \in \Theta_1$ that is closest to a sample \bar{X} in the Euclidean metric.

Proof. When $\Sigma = I_m$, the log-likelihood function becomes

$$\ell_n(\mu) = -\frac{n}{2} \text{tr}(S) - \frac{n}{2} (\bar{X} - \mu)^T (\bar{X} - \mu) = -\frac{n}{2} \text{tr}(S) - \|\bar{X} - \mu\|^2$$

where $\|\cdot\|$ denotes the L^2 -norm. Hence, maximizing $\ell_n(\mu)$ over Θ is equivalent to minimizing the Euclidean distance from \bar{X} to Θ_1 in \mathbb{R}^m . \square

Example 1.2.4 (Nodal cubic continued). Consider the nodal cubic from Example 1.2.1. Let $\bar{X} = (x_1, x_2) \in \mathbb{R}^2$ be a general point. To compute the maximum likelihood estimate of \bar{X} , we maximize

$$\ell(t) = \|(x_1, x_2) - ((t^2 - 1), t(t^2 - 1))\|^2 = (x_2 - t^3 + t)^2 + (x_1 - t^2 + 1)^2.$$

over the model. Differentiating with respect to t and setting to 0, we get the irreducible equation

$$3t^5 - 2t^3 - 3t^2x_2 - 2tx_1 - t + x_2 = 0.$$

The solutions to this equations are the critical points of the maximum likelihood estimation problem. Hence the *maximum likelihood degree* of this model is 5, and it is equal to the *Euclidean distance degree* [45]. Evaluating $\ell(t)$ at all real solutions for a given \bar{X} , we find a parameter \hat{t} giving the maximum value. Then the point $((\hat{t}^2 - 1), \hat{t}(\hat{t}^2 - 1)) \in \Theta_1$ is the MLE of \bar{X} . Note that since the singularities of the curve at $t = \pm 1$ are not solutions to the critical equation above for general data, the presence of singularities did not affect our ML degree computations. However, the ML degree may drop for models with singularities of a different nature. One such example is the cuspidal cubic; see [114, Example 7.1.7].

In practice, when we work with Gaussian models, we often set the mean to be either zero or the sample mean. Hence, most of the Gaussian models in this thesis will have parameter spaces of the form $\Theta = \mathbb{R}^m \times \Theta_2$, where $\Theta_2 \subseteq \text{PD}_m$. Hence, we may identify a Gaussian model Θ with a subset of the positive definite cone PD_m . Moreover, we will only work with *algebraic* Gaussian models, i.e. those models that can be written as $\Theta = \mathcal{V} \cap \text{PD}_m$ for some variety $\mathcal{V} \subseteq \mathbb{C}^{\binom{m+1}{2}}$ in the entries of $\Sigma = (\sigma_{ij}) \in \text{PD}_m$.

Given a sample covariance matrix S and setting $\mu = \hat{X}$, the log-likelihood function in (1.5) simplifies to

$$\ell_n(\Sigma, S) = -\frac{n}{2} \log \det \Sigma - \frac{n}{2} \text{tr}(S\Sigma^{-1}). \quad (1.6)$$


The *maximum likelihood degree* of a Gaussian model $\Theta \subseteq \text{PD}_m$ is the number of nonsingular complex critical points of $\ell_n(\Sigma, S)$ for generic S on the Zariski closure of Θ in the space of complex symmetric $m \times m$ matrices.

Example 1.2.5 (The samosa). Consider the Gaussian model

$$\Theta = \left\{ \Sigma \in \text{PD}_3 : \Sigma = \begin{pmatrix} 1 & x & z \\ x & 1 & y \\ z & y & 1 \end{pmatrix} \right\},$$

known as the *unrestricted correlation model* with $m = 3$, further discussed in Example 5.3.4. It is often referred to as the *elliptope* or *samosa*, due to its geometric resemblance to the Indian pastry; see Figure 1.5. Note that this Gaussian model is parametric, since every matrix $\Sigma \in \Theta$ is defined by three parameters: x, y , and z . Given a generic sample covariance matrix S , we may find its MLE Σ similarly to the discrete case, by constructing the likelihood equations. With respect to the first parameter x , we get the first likelihood equation:

$$\frac{\partial}{\partial x} (-\log \det \Sigma - \text{tr}(S\Sigma^{-1})) = -\frac{\frac{\partial}{\partial x} \det \Sigma}{\det \Sigma} - \frac{\partial}{\partial x} \text{tr}(S\Sigma^{-1}) = 0.$$

Figure 1.5: The unrestricted correlation model for $m = 3$.

Clearing the denominators, the equation becomes $\frac{\partial}{\partial x} \det \Sigma + \det \Sigma \frac{\partial}{\partial x} \text{tr}(S\Sigma^{-1}) = 0$. However, we are not done clearing denominators. The inverse matrix Σ^{-1} contains rational entries. Hence we re-write $\Sigma^{-1} = \frac{1}{\det \Sigma} M$ where

$$M = \begin{pmatrix} -y^2 + 1 & yz - x & xy - z \\ yz - x & -z^2 + 1 & xz - y \\ xy - z & xz - y & -x^2 + 1 \end{pmatrix}$$

is the adjugate matrix. Then the likelihood equation becomes

$$\frac{\partial}{\partial x} \det \Sigma + \det \Sigma \frac{\partial}{\partial x} \text{tr} \left(\frac{SM}{\det \Sigma} \right) = \frac{\partial}{\partial x} \det \Sigma + \frac{\partial}{\partial x} \text{tr}(SM) - \frac{\text{tr}(SM)}{\det \Sigma} \frac{\partial}{\partial x} \det \Sigma = 0.$$

Clearing the denominators again, we get the critical equation

$$\frac{\partial}{\partial x} \det \Sigma (\det \Sigma - \text{tr}(SM)) + \det \Sigma \frac{\partial}{\partial x} \text{tr}(SM) = 0. \quad (1.7)$$

Repeating the above steps for y and z , we get two more critical equations of the form (1.7), which we want to solve. Thus, we form the ideal I generated by these three equations. However, since we multiplied by $\det(\Sigma)$ while clearing denominators, we must saturate I by $\langle \det(\Sigma) \rangle$ to remove extraneous solutions. The resulting ideal is the likelihood ideal with respect to the covariance matrix S . We compute it using the following Macaulay2 code.

```
loadPackage "EigenSolver";
R = QQ[x,y,z];
Sigma = matrix{{1,x,z}, {x,1,y}, {z,y,1}};
detSigma = det(Sigma);
A = random(ZZ^3, ZZ^3);
S = A*transpose(A); --generates a random positive definite matrix
```

```

M = matrix({{1-y^2, -(x-y*z), x*y-z},
           {-(x-y*z), 1-z^2, -(y-x*z)},
           {x*y-z, -(y-x*z), 1-x^2}});
trSM = trace(S*M);
I = ideal((detSigma-trSM)*diff(x,detSigma)+detSigma*diff(x,trSM),
          (detSigma-trSM)*diff(y,detSigma)+detSigma*diff(y,trSM),
          (detSigma-trSM)*diff(z,detSigma)+detSigma*diff(z,trSM));
L = saturate(I,detSigma);
degree L
zeroDimSolve L

```

For a random covariance matrix S , the likelihood ideal L is zero-dimensional of degree 15. Therefore, the ML degree of the ellipsope is 15. One of the positive definite solutions of L that maximizes $\ell_n(S, \Sigma)$ will yield the MLE of S .

A general ML degree formula for unrestricted correlation models is not known. Furthermore, ML degrees are not known for the unrestricted correlation models when $m \geq 7$. The ML degrees for $m = 4, 5, 6$ are 109, 1077, and 13695, respectively; these were computed in [12].

Logarithmic Voronoi cells

We now introduce logarithmic Voronoi cells for Gaussian models. We also prove several basic facts about the logarithmic Voronoi cells for the saturated model and Gaussian models of the form $\Theta = \Theta_1 \times \{I_m\}$. For the models of the type $\Theta = \mathbb{R}^m \times \Theta_2$, we define the log-normal spectrahedron and motivate the in-depth study of their logarithmic Voronoi cells in Chapter 5 with the analysis of a conditional independence model.

For a point $\theta = (\mu, \Sigma)$ on a Gaussian model $\Theta \subseteq \mathbb{R}^m \times \text{PD}_m$, we define its *logarithmic Voronoi cell* $\log \text{Vor}_\Theta(\mu, \Sigma)$ to be the set of all $X^{(1)}, \dots, X^{(n)} \in \mathbb{R}^m$ with sample mean \bar{X} and sample covariance S such that the log-likelihood function ℓ_n with respect to this sample is maximized at θ . In this thesis, we will study logarithmic Voronoi cells only at nonsingular points of Gaussian models.

Proposition 1.2.6. [5] Let $\Theta = \mathbb{R}^m \times \text{PD}_m$ be the saturated Gaussian model. For any point in this model, its logarithmic Voronoi cell is the point itself.

Proof. For any given sample (\bar{X}, S) , its maximum likelihood estimate $(\hat{\mu}, \hat{\Sigma})$ is the point (\bar{X}, S) itself [47, Section 2.1]. Therefore, for any given point $(\mu, \Sigma) \in \Theta$, its logarithmic Voronoi cell is $\log \text{Vor}_\Theta(\mu, \Sigma) = \{(\mu, \Sigma)\}$, as desired. \square

Recall that for any $U \subseteq \mathbb{R}^m$ and $p \in U$, the *Euclidean Voronoi cell* at p is the set of all points in \mathbb{R}^m that are closer to p than any other point in U with respect to the Euclidean

metric. Euclidean Voronoi cells of varieties were studied in [32] and are a topic in *metric algebraic geometry* [26, 39, 45, 121]. In general, logarithmic Voronoi cells are not equal to Euclidean Voronoi cells. However, it turns out they are the same for the next model.

Proposition 1.2.7. [5] Consider the Gaussian model with parameter space $\Theta = \Theta_1 \times \{\text{Id}_m\}$ for some $\Theta_1 \subseteq \mathbb{R}^m$. For any point in this model, its logarithmic Voronoi cell is equal to its Euclidean Voronoi cell.

Proof. We may identify the parameter space Θ with the subset of real vectors $\Theta_1 \subseteq \mathbb{R}^m$. For any sample $X^{(1)}, \dots, X^{(n)}$ with sample mean \bar{X} , the maximum likelihood estimate is the point in the model $\hat{\mu} \in \Theta_1$ that is closest to \bar{X} in the Euclidean metric, by Proposition 1.2.3. So, for any point $\mu \in \Theta_1$ in the model, the logarithmic Voronoi cell at μ is the set of all sample means $\bar{X} \in \mathbb{R}^m$ that are closer to μ than any other point in \mathbb{R}^m . This is precisely the Euclidean Voronoi cell at μ . \square

Besides the above relatively simple cases, logarithmic Voronoi cells of Gaussian models are fairly complex convex sets. In the rest of the thesis we will consider Gaussian models given by parameter spaces of the form $\Theta = \mathbb{R}^m \times \Theta_2$ where $\Theta_2 \subseteq \text{PD}_m$, which can be identified with subsets of PD_m . Hence, for any point $\Sigma \in \Theta$, its logarithmic Voronoi cell $\log \text{Vor}_\Theta(\Sigma)$ is the set of all matrices $S \in \text{PD}_m$ such that Σ is a maximizer of $\ell_n(\Sigma, S)$, viewed as a function of Σ .

Before exploring logarithmic Voronoi cells for such models further, we demonstrate the main themes in the following example.

Example 1.2.8 (Conditional independence model). Consider the model Θ that is given as the intersection of the algebraic variety

$$\{\Sigma = (\sigma_{ij}) : \sigma_{13} = 0, \sigma_{12}\sigma_{23} - \sigma_{22}\sigma_{13} = 0\} = \{\Sigma = (\sigma_{ij}) : \sigma_{13} = 0, \sigma_{12}\sigma_{23} = 0\}$$

with the cone PD_3 of positive definite symmetric 3×3 matrices. This is the conditional independence model given by $X_1 \perp\!\!\!\perp X_3$ and $X_1 \perp\!\!\!\perp X_3 \mid X_2$, and it is the union of two linear planes of dimension four. We may write

$$\Theta = \left\{ \begin{pmatrix} t_1 & 0 & 0 \\ 0 & t_2 & t_3 \\ 0 & t_3 & t_4 \end{pmatrix} \succ 0 : t_i \in \mathbb{R} \right\} \cup \left\{ \begin{pmatrix} s_1 & s_2 & 0 \\ s_2 & s_3 & 0 \\ 0 & 0 & s_4 \end{pmatrix} \succ 0 : s_i \in \mathbb{R} \right\}.$$

Let Θ_1 and Θ_2 denote the two components above, respectively. Given a matrix $\Sigma \in \Theta$, the set of sample covariance matrices $S \in \text{PD}_3$ that have Σ as their maximum likelihood estimate form the *logarithmic Voronoi cell* at Σ . The set of all matrices $S \in \text{PD}_3$ that have Σ as a critical point while optimizing the log-likelihood function with respect to S over Θ is the *log-normal spectrahedron* at Σ . The log-normal spectrahedron at a general matrix

$\Sigma \in \Theta_1 \setminus \Theta_2$ is two-dimensional, parametrized as

$$\left\{ \begin{pmatrix} t_1 & x_1 & x_2 \\ x_1 & t_2 & t_3 \\ x_2 & t_3 & t_4 \end{pmatrix} \succ 0 : x_1, x_2 \in \mathbb{R} \right\}.$$

This spectrahedron is a semi-algebraic set, defined by the two inequalities

$$-x_1^2 + t_1 t_2 > 0 \text{ and } -t_2 x_2^2 + 2t_3 x_1 x_2 - t_4 x_1^2 - t_1 t_3^2 + t_1 t_2 t_4 > 0.$$

Since Σ is assumed to be positive definite, for any choice of t_i , the log-normal spectrahedron at Σ is an ellipse. By symmetry the same is true of any $\Sigma \in \Theta_2 \setminus \Theta_1$. For a point $\Sigma = \text{diag}(\sigma_1, \sigma_2, \sigma_3) \in \Theta_1 \cap \Theta_2$, the log-normal spectrahedron is three-dimensional, given as

$$\left\{ (x, y, z) \in \mathbb{R}^3 : \begin{pmatrix} \sigma_1 & x & y \\ x & \sigma_2 & 0 \\ y & 0 & \sigma_3 \end{pmatrix} \succ 0 \text{ and } \begin{pmatrix} \sigma_1 & 0 & y \\ 0 & \sigma_2 & z \\ y & z & \sigma_3 \end{pmatrix} \succ 0 \right\}.$$

The maximum likelihood degree of Θ is two with one critical point in each linear component. Namely, for a general matrix $S = (s_{ij}) \in \text{PD}_3$, the two critical points on the model are $\Sigma_1 \in \Theta_1$, given by $t_1 = s_{11}, t_2 = s_{22}, t_3 = s_{23}, t_4 = s_{33}$, and $\Sigma_2 \in \Theta_2$, given by $s_1 = s_{11}, s_2 = s_{12}, s_3 = s_{22}, s_4 = s_{33}$. Now consider a general matrix $\Sigma \in \Theta_1 \setminus \Theta_2$. The logarithmic Voronoi cell at Σ is a subset of its log-normal ellipse, and it can be written as

$$\left\{ S = \begin{pmatrix} t_1 & x_1 & x_2 \\ x_1 & t_2 & t_3 \\ x_2 & t_3 & t_4 \end{pmatrix} \succ 0 : \ell_n(\Sigma, S) \geq \ell_n(\Sigma', S) \right\} \quad (1.8)$$

where ℓ_n is the log-likelihood function and $\Sigma' = \begin{pmatrix} t_1 & x_1 & 0 \\ x_1 & t_2 & 0 \\ 0 & 0 & t_4 \end{pmatrix}$. Writing out the inequality in (1.8), we find that it is equivalent to

$$-t_3 \sqrt{t_1/t_4} \leq x_1 \leq t_3 \sqrt{t_1/t_4}. \quad (1.9)$$

Thus, the logarithmic Voronoi cell at $\Sigma \in \Theta_1 \setminus \Theta_2$ is the log-normal ellipse at Σ intersected with the strip defined by (1.9). In particular, it is a semi-algebraic set. We plot the logarithmic Voronoi cell for $t_1 = 1, t_2 = 2, t_3 = 1, t_4 = 3$ in Figure 1.6 (on the left). Similarly, one checks that the logarithmic Voronoi cell at $\Sigma \in \Theta_2 \setminus \Theta_1$ is the semi-algebraic set

$$\left\{ S = \begin{pmatrix} s_1 & s_2 & y_1 \\ s_2 & s_3 & y_2 \\ y_1 & y_2 & s_4 \end{pmatrix} \succ 0 : -s_2 \sqrt{s_4/s_1} < y_2 < s_2 \sqrt{s_4/s_1} \right\}.$$

We plot the logarithmic Voronoi cell for $s_1 = 2, s_2 = 1, s_3 = 3, s_4 = 4$ in Figure 1.6 (on the right). Thus, the logarithmic Voronoi cell at a general point of Θ is not equal to its log-normal ellipse.




Figure 1.6: Logarithmic Voronoi cells (in pink) of the model in Example 1.2.8 plotted on the (x_1, x_2) -plane and (y_1, y_2) -plane, respectively.

Despite strict containment of the logarithmic Voronoi cells in log-normal spectrahedra, the former is still a semi-algebraic set. This phenomenon is surprising, since the inequality in (1.8) that defines the logarithmic Voronoi cell together with the positive definiteness condition involves the log-likelihood function, which is not a polynomial function. Also note that at the singular points $\Sigma \in \Theta_1 \cap \Theta_2$ the logarithmic Voronoi cells equal the log-normal spectrahedra which are three-dimensional.

We can now generalize the concepts introduced in the previous example. For a point $\Sigma \in \Theta$, we define the *log-normal matrix space* at Σ , denoted by $\mathcal{N}_\Sigma\Theta$, to be the set of all symmetric $m \times m$ matrices S such that Σ appears as a critical point when optimizing $\ell_n(\Sigma, S)$. This is the set of all points such that the gradient $\nabla \ell_n(\Sigma, S)$ with respect to Σ lies in the normal space of the model Θ at Σ . This condition is linear in S , so the log-normal matrix space is an affine linear space. Intersecting it with PD_m , we obtain a spectrahedron $\mathcal{K}_\Theta(\Sigma) = \text{PD}_m \cap \mathcal{N}_\Sigma\Theta$, which we call the *log-normal spectrahedron* at Σ . We immediately obtain the following.

Proposition 1.2.9. Each logarithmic Voronoi cell $\log \text{Vor}_\Theta(\Sigma)$ is contained in the log-normal spectrahedron $\mathcal{K}_\Theta(\Sigma)$. In particular,

$$\log \text{Vor}_\Theta \Sigma = \{S \in \mathcal{K}_\Theta \Sigma : \ell_n(\Sigma, S) \geq \ell_n(\Sigma', S) \text{ for all critical points } \Sigma'\}.$$

The reverse containment does not hold in general, as we have seen in Example 1.2.8. This is typical, and we will see more instances of this phenomenon. Most of Chapter 5 is devoted to comparing and contrasting log-normal spectrahedra and logarithmic Voronoi cells for different Gaussian models.

We end this section with the table, which juxtaposes the analogous concepts defined so far in the discrete vs. Gaussian case.

	discrete	Gaussian
ambient space	probability simplex Δ_{n-1}	positive definite cone PD_m
model	$\mathcal{M} \subseteq \Delta_{n-1}$	$\Theta \subseteq \text{PD}_m$
log-likelihood function	$\sum_{i=1}^n u_i \log p_i$	$-\log \det \Sigma - \text{tr}(S\Sigma^{-1})$
log-normal linear space	space $\log N_p \mathcal{M}$	matrix space $\mathcal{N}_\Sigma \Theta$
log-normal convex set	polytope $\log \text{Poly}(p)$	spectrahedron $\mathcal{K}_\Theta(\Sigma)$

Conditional independence

In this section, we introduce conditional independence for Gaussian random variables and describe the arising algebraic conditions on the entries of covariance matrices.

Conditional independence for Gaussian models is defined similarly to the discrete case. We again start with a multivariate normal random vector $X = (X_1, X_2, \dots, X_m) \sim \mathcal{N}(\mu, \Sigma)$ with the Gaussian density function f . As before, for any subset $A \subset [m]$, we may define the subvector $X_A = (X_a : a \in A)$ and the *marginal density function* $f_A(x_A)$ with respect to A by integrating $f(x)$ over the coordinate subspace defined by $[m] \setminus A$. The conditional density $f_{A|B}(x_A|x_B)$ is defined as

$$f_{A|B}(x_A|x_B) = \frac{f_{A \cup B}(x_{A \cup B})}{f_B(x_B)}.$$

We say that X_A is *conditionally independent* of X_B given X_C for pairwise disjoint $A, B, C \subset [m]$ whenever

$$f_{A \cup B|C}(x_{A \cup B}|x_C) = f_{A|C}(x_A|x_C) f_{B|C}(x_B|x_C)$$

for all x_A, x_B , and x_C such that $f_C(x_C) > 0$. The following proposition translates the notion of conditional independence into the language of algebra.

Proposition 1.2.10. [47, Proposition 3.1.13] The conditional independence statement $X_A \perp\!\!\!\perp X_B | X_C$ holds for a multivariate normal random vector $X \sim \mathcal{N}(\mu, \Sigma)$ if and only if the submatrix $\Sigma_{A \cup C, B \cup C}$, induced by the rows indexed by $A \cup C$ and columns indexed by $B \cup C$, has rank $|C|$.

Example 1.2.11. In Example 1.2.8, the two CI statements $X_1 \perp\!\!\!\perp X_3$ and $X_1 \perp\!\!\!\perp X_3 | X_2$ translate into the condition that two minors of the covariance matrix vanish:

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{12} & \sigma_{22} & \sigma_{23} \\ \sigma_{13} & \sigma_{23} & \sigma_{33} \end{pmatrix}.$$

These are precisely the defining equations of the model: $\sigma_{13} = 0$ and $\sigma_{12}\sigma_{23} - \sigma_{12}\sigma_{13} = 0$.

1.3 Statement of contributions

In *Chapter 2: Discrete setting*, we study logarithmic Voronoi cells for discrete statistical models. In particular, we investigate sufficient criteria for models to have polytopes for logarithmic Voronoi cells. We prove that finite models, linear models, toric models, and models of ML degree one always satisfy this property. We completely characterize logarithmic Voronoi polytopes for the finite model of all distributions of a fixed sample size combinatorially. For the models whose logarithmic Voronoi cells are more complicated convex sets, we develop a method to reliably compute them via the framework of numerical algebraic geometry. This chapter is based on joint work with Alexander Heaton, published in *Algebraic Statistics* [3]. The code for this chapter is available at Mathrepo¹, a repository of the Max Planck Institute for Mathematics in the Sciences in Leipzig, dedicated to mathematical research data [34].

In *Chapter 3: Linear models*, we completely describe the combinatorics of logarithmic Voronoi polytopes for linear models, both at the interior points and on the boundary of the simplex. We then use logarithmic Voronoi cells to study the problem of maximizing information divergence from linear models. Finally, we investigate the combinatorics of logarithmic Voronoi polytopes for partial linear models, where the points on the boundary of the model are especially of interest. This chapter is largely based on a solo paper, to appear in *Algebraic Statistics* [1].

In *Chapter 4: Toric models*, we study the question of maximizing divergence from toric models from a new perspective using logarithmic Voronoi polytopes. We present an algorithm that combines the combinatorics of the chamber complex with numerical algebraic geometry. We pay special attention to reducible models and models of maximum likelihood degree one. This chapter, as well as Section 3.3 in Chapter 3, are based on joint work with Serkan Hoşten, submitted for publication and available as a preprint [6]. The code for this chapter is available on Github².

In *Chapter 5: Gaussian setting*, we extend the theory of logarithmic Voronoi cells from Chapter 2 to Gaussian models. We show that logarithmic Voronoi cells are equal to log-normal spectrahedra for models of ML degree one and linear concentration models. We also study covariance models, for which logarithmic Voronoi cells are, in general, strictly contained in log-normal spectrahedra. We give an explicit semi-algebraic description of logarithmic Voronoi cells for the bivariate correlation model. Finally, we state a conjecture that logarithmic Voronoi cells for unrestricted correlation models, such as the ellipse in Example 1.2.5, are not semi-algebraic. This chapter is based on joint work with Serkan Hoşten, published in *Journal of Symbolic Computation* [5]. The code is available on Github³.

¹<https://mathrepo.mis.mpg.de/logarithmicVoronoi>

²<https://github.com/yuliaalexandr/maximizing-divergence>

³<https://github.com/yuliaalexandr/gaussian-log-voronoi>

In *Chapter 6: Decomposable context-specific models*, we study how algebra can be used to model causality and capture finer forms of independence, such as context-specific independence. We describe the construction of *decomposable context-specific models* from the subclass of staged tree models known as CStree models. We study algebraic and combinatorial properties and characterizations of such models, including context-specific independence relations, graphical representations, and their prime ideals. This chapter is based on joint work with Eliana Duarte and Julian Vill, submitted for publication and available as a preprint [2].

In *Chapter 7: Moment varieties for mixtures of products*, we study moment varieties of conditionally independent mixture distributions on \mathbb{R}^n and their images under certain coordinate projections. These are the secant varieties of toric varieties that express independence in terms of univariate moments. We prove several results about dimensions, degrees, and defining polynomials of these varieties. We also investigate finiteness properties of their defining ideals. This work includes several computational results, featuring both symbolic and numerical methods. This chapter is based on joint work with Joe Kileel and Bernd Sturmfels, published in the proceedings of the *ACM International Symposium on Symbolic and Algebraic Computation* [7]. The code is available on Mathrepo⁴.

⁴<https://mathrepo.mis.mpg.de/MomentVarietiesForMixturesOfProducts>

Part I

Logarithmic Voronoi cells

Chapter 2

Discrete setting

This chapter takes the first step of characterizing and computing logarithmic Voronoi cells, defined in Section 1.1 in the discrete setting. First, we show that logarithmic Voronoi cells are always convex sets. However, for certain families of models, such as finite, linear, toric, and models of ML degree one, they are polytopes. We describe logarithmic Voronoi polytopes for the finite model consisting of all empirical distributions of a fixed sample size. These polytopes are dual to the logarithmic root polytopes of Lie type A, and we characterize their faces. Finally, we present a computation of non-polytopal logarithmic Voronoi cells of a mixture model using numerical algebraic geometry. This chapter is based on [3].

2.1 When are logarithmic Voronoi cells polytopes?

In this section we prove several basic results about logarithmic Voronoi cells. The main focus is to explore criteria which ensure that the logarithmic Voronoi cells are polytopes, rather than more general convex sets, as in Proposition 2.1.3 below.

Proposition 2.1.1. Let \mathcal{M} be any finite statistical model. Then the logarithmic Voronoi cells $\log \text{Vor}_{\mathcal{M}}(p)$ are polytopes for each $p \in \mathcal{M}$.

Proof. Fix $p \in \mathcal{M}$. The set of all points $u \in \Delta_{n-1}$ such that $\ell_u(p) \geq \ell_u(q)$ for all $q \in \mathcal{M}$ is the logarithmic Voronoi cell at p . For any $q \in \mathcal{M}$, $\ell_u(p) \geq \ell_u(q)$ becomes the condition that

$$\sum_{i=1}^n u_i \log \left(\frac{p_i}{q_i} \right) \geq 0.$$

But this is linear in u and so defines a closed halfspace. Since there are finitely many points in \mathcal{M} , we see that the logarithmic Voronoi cell is an intersection of finitely many closed halfspaces (including those defining Δ_{n-1}). Therefore it is a polytope. \square

For infinite models, the logarithmic Voronoi cells are, in general, not polytopes. However, if the model is smooth at p , the logarithmic Voronoi cell will be contained in the log-normal polytope. Figure 2.1 shows a logarithmic Voronoi cell for $p \in \mathcal{M} \subset \Delta_5 \subset \mathbb{R}^6$ which is not a polytope, but is contained in a polytope. In this case, $\log \text{Poly}(p) = \log N_p \mathcal{M} \cap \Delta_5$ is a hexagon. Since the log-normal space is 2-dimensional, by choosing an orthonormal basis agreeing with this subspace we can visualize the logarithmic Voronoi cell, despite it living in \mathbb{R}^6 . We discuss this example in detail in Section 2.3. For more on finite models, see Section 2.2.




Figure 2.1: Logarithmic Voronoi cell (green) inside its log-normal polytope (pink) for a given point (yellow) in the model from Example 2.3.3.

Lemma 2.1.2. Let $\Phi(u) = p$ for some $p \in \mathcal{M} \subset \Delta_{n-1}$ such that $U \cap \mathcal{M}$ is a manifold for some p -neighborhood U in \mathbb{R}^n . Then u lies in the logarithmic normal space $\log N_p \mathcal{M}$ and

$$\log \text{Vor}_{\mathcal{M}}(p) \subset \log \text{Poly}_{\mathcal{M}}(p).$$

Proof. Note that $\ell_u(x) = \sum u_i \log(x_i)$ is a smooth function on any neighborhood of $p \in \mathcal{M}$ contained in Δ_{n-1} . Consider the gradient $\nabla \ell_u(p)$. Note that $\mathbb{R}^n = T_p \mathcal{M} \oplus N_p \mathcal{M}$ and if $\nabla \ell_u(p)$ had any nonzero tangential component then there would exist some $q \in \mathcal{M}$ such that $\ell_u(q) > \ell_u(p)$, contradicting the fact that $\Phi(u) = p$. \square

Proposition 2.1.3. Logarithmic Voronoi cells are convex sets.

Proof. As in the proof of Proposition 2.1.1, the logarithmic Voronoi cell at p is defined by the inequalities $\sum_{i \in [n]} u_i \log(p_i/q_i) \geq 0$ for every $q \in \mathcal{M}$, each linear in u . Hence, the logarithmic Voronoi cell at p is an intersection of (possibly infinitely many) closed half-spaces, and the result follows. \square

Recall from Section 1.1 that for an algebraic statistical model \mathcal{M} , the *ML degree* is the number of complex critical points of ℓ_u on the Zariski closure of \mathcal{M} for generic data $u \in \Delta_{n-1}$ [114, p. 140]. The following theorem concerns algebraic models with ML degree 1. These were characterized in [66] and studied further in [48]. They include, for example, Bayesian networks, decomposable graphical models studied in Chapter 4, and CStree models studied in Chapter 6.

Recall that $\Delta_{n-1}^{\mathcal{M}}$ denotes the set of $u \in \Delta_{n-1}$ such that $\Phi(u)$ exists.

Theorem 2.1.4. Let \mathcal{M} be any algebraic model with ML degree 1 which is smooth on Δ_{n-1} . Then the logarithmic Voronoi cell at every $p \in \mathcal{M}$ equals its log-normal polytope on $\Delta_{n-1}^{\mathcal{M}}$.

Proof. We will show that $\log \text{Vor}_{\mathcal{M}}(p) = \log N_p \mathcal{M} \cap \Delta_{n-1}^{\mathcal{M}}$. Let $u \in \Delta_{n-1}$ be an element of $\log \text{Vor}_{\mathcal{M}}(p)$. Then $\Phi(u) = p$ and since \mathcal{M} is smooth, $u \in \log N_p \mathcal{M} \cap \Delta_{n-1}^{\mathcal{M}}$ by Lemma 2.1.2. For the reverse direction, let $u \in \log N_p \mathcal{M} \cap \Delta_{n-1}^{\mathcal{M}}$. Recall that $\Phi(u)$ is the argmax of $\ell_u(q)$ over all points $q \in \mathcal{M}$. Since $\Phi(u)$ exists and \mathcal{M} is smooth, this argmax must be among the critical points of ℓ_u restricted to \mathcal{M} , which include p . But since the ML degree is 1, there is only one complex critical point, and hence $\Phi(u) = p$. Therefore u is in the logarithmic Voronoi cell at p , and the result follows. \square

Example 2.1.5. Consider $\mathcal{M} = V(f)$ for $f : \mathbb{C}^4 \rightarrow \mathbb{C}^2$ given by the polynomial system

$$f(x) = \begin{bmatrix} x_1 x_4 - x_2 x_3 \\ x_1 + x_2 + x_3 + x_4 - 1 \end{bmatrix} : \mathbb{C}^4 \rightarrow \mathbb{C}^2$$

A parametrization of this model is given by

$$(t_1, t_2) \mapsto (t_1 t_2, t_1(1 - t_2), (1 - t_1)t_2, (1 - t_1)(1 - t_2)).$$

This is the *independence model* from Example 1.1.3 on two binary random variables, and also the Segre embedding of $\mathbb{P}^1 \times \mathbb{P}^1$. The points of this 2-dimensional model live in the 3-dimensional hyperplane $\sum x_i = 1$ inside \mathbb{R}^4 , so we can choose a basis agreeing with this hyperplane to plot them.

For each $x \in \mathcal{M}$, we construct an $(m + 1) \times n$ matrix $A(x)$ by augmenting the row $\nabla \ell_u$ to the Jacobian matrix df :

$$A(x) = \begin{bmatrix} x_4 & -x_3 & -x_2 & x_1 \\ 1 & 1 & 1 & 1 \\ u_1/x_1 & u_2/x_2 & u_3/x_3 & u_4/x_4 \end{bmatrix}.$$

Since our model has codimension two, the 3-minors of $A(x)$ give linear equations describing the log-normal space.

$$\begin{aligned} u_2 - u_3 - \frac{u_1 x_2}{x_1} + \frac{u_1 x_3}{x_1} + \frac{u_2 x_4}{x_2} - \frac{u_3 x_4}{x_3} &= 0 \\ u_1 - u_4 - \frac{u_2 x_1}{x_2} + \frac{u_1 x_3}{x_1} - \frac{u_4 x_3}{x_4} + \frac{u_2 x_4}{x_2} &= 0 \\ u_1 - u_4 + \frac{u_1 x_2}{x_1} - \frac{u_3 x_1}{x_3} - \frac{u_4 x_2}{x_4} + \frac{u_3 x_4}{x_3} &= 0 \\ u_2 - u_3 + \frac{u_2 x_1}{x_2} - \frac{u_3 x_1}{x_3} - \frac{u_4 x_2}{x_4} + \frac{u_4 x_3}{x_4} &= 0. \end{aligned}$$

Restricting this space to its intersection with the simplex $u_1 + u_2 + u_3 + u_4 - 1 = 0$ to compute the log-normal polytope, we find that the polytopes are line segments. We plot them for various points on the model in Figure 2.2. Since \mathcal{M} has ML degree 1, Theorem 2.1.4 tells us that log-Voronoi cells equal log-normal polytopes, so they are also line segments.




Figure 2.2: One-dimensional log-normal polytopes at various points

Example 2.1.6 below shows that the ML degree 1 condition in Theorem 2.1.4 is sufficient, but not necessary, for the equality of logarithmic Voronoi cells with the interior of their log-normal polytopes. Consider the independence model of two identically distributed binary random variables. The natural parametrization in a statistical context leads to the Hardy-Weinberg curve defined by $x_2^2 - 4x_1x_3$, which has ML degree 1 [67]. A similar-looking model, which has been called the cousin of the Hardy-Weinberg curve [63], is defined by the polynomial $f = x_2^2 - x_1x_3$. It turns out that the ML degree of this model is 2 [63, p. 394]. It was demonstrated in [63] that ML degree is extremely sensitive to scaling of the coordinates, so the difference between the ML degrees of the Hardy-Weinberg curve and its cousin is not surprising. The effect of scaling on the ML degree of toric varieties has been studied in [8].

Example 2.1.6. The algebraic model defined by the polynomial $f = x_2^2 - x_1x_3$ has ML degree 2, yet the logarithmic Voronoi cells are equal to their log-normal polytopes.

Although this follows from the later Theorem 2.1.10, we will first prove it more explicitly. Calculate the Jacobian matrix of Lemma 1.1.13 by taking the gradients of $f = x_2^2 - x_1x_3$ and

$g = x_1 + x_2 + x_3 - 1$, augmenting this matrix with an additional row of the u_i/x_i . Consider the equation of the plane given by the determinant of this matrix. Note that \mathcal{M} is a curve in Δ_2 , so the log-normal space at each point is defined by the vanishing of the determinant at that point. This plane has normal vector given by

$$\begin{pmatrix} 2x_1x_2^2 - x_1^2x_3 + x_1x_2x_3 + x_1x_3^2 \\ -2x_1x_2^2 - 2x_1x_2x_3 - 2x_2^2x_3 \\ x_1^2x_3 + x_1x_2x_3 + 2x_2^2x_3 - x_1x_3^2 \end{pmatrix}$$

where (x_1, x_2, x_3) is any point in the common zero set of f and g . Consider the cross-product of this vector with the all ones vector, which will give us the direction vector of the log-normal polytope at (x_1, x_2, x_3) . Computing and simplifying each coordinate in the quotient ring

$$\mathbb{Q}[x_1, x_2, x_3] / \langle x_1 + x_2 + x_3 - 1, -x_2^2 + x_1x_3 \rangle,$$

we find that this cross product is given by

$$\begin{pmatrix} -(x_2 + x_3 - 1)x_3 \\ 2(x_2 + x_3 - 1)x_3 \\ -(x_2 + x_3 - 1)x_3 \end{pmatrix} = x_1x_3 \begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}.$$

This means that regardless of the point on the curve, the log-normal polytopes will be line segments whose direction vector is $(-1, 2, -1)$. We claim that for any distinct $p, q \in \mathcal{M}$ the corresponding line segments are disjoint. Consider the tangent space at some point x in the intersection of Δ_2 and the common zero set of f and g . Applying Gaussian elimination to the 2×3 Jacobian matrix, it can be shown that if $2x_2 + x_3 \neq 0$ then all tangent vectors are multiples of

$$\left(\frac{x_3 - x_1}{2x_2 + x_3} - 1, \frac{x_1 - x_3}{2x_2 + x_3}, 1 \right), \tag{2.1}$$

while if $2x_2 + x_3 = 0$ then all tangent vectors are multiples of $(-1, 1, 0)$. In neither case is it possible that a tangent vector is parallel to $(1, -2, 1)$. For $(-1, 1, 0)$ this is obvious, but for (2.1), a contradiction can be derived by showing that if the vector is parallel to $(1, -2, 1)$ the first and the last coordinates in (2.1) are equal, forcing $x_1 + 4x_2 + x_3 = 0$. But on Δ_2 all coordinates are positive. Thus no line parallel to $(1, -2, 1)$ meets the model in two distinct points. We conclude the log-normal polytopes are disjoint, and the result follows from Lemma 2.1.7 below.

Lemma 2.1.7. Let \mathcal{M} be any smooth model in Δ_{n-1} . If all log-normal polytopes for each point $p \in \mathcal{M}$ are disjoint, then the logarithmic Voronoi cells equal log-normal polytopes on $\Delta_{n-1}^{\mathcal{M}}$.

Proof. We will show that $\log \text{Vor}_{\mathcal{M}}(p) = \log N_p \mathcal{M} \cap \Delta_{n-1}^{\mathcal{M}}$. The \subset direction follows from Lemma 2.1.2. For the reverse direction, let $u \in \log N_p \mathcal{M} \cap \Delta_{n-1}^{\mathcal{M}}$. Recall that $\Phi(u)$ is the

argmax of $\ell_u(q)$ over all points $q \in \mathcal{M}$. Since $\Phi(u)$ exists and \mathcal{M} is smooth, this argmax must be among the critical points of ℓ_u restricted to \mathcal{M} , which include p . If $\Phi(u)$ were not equal to p then u would be in the intersection of Δ_{n-1} with the log-normal space to the point $\Phi(u) \in \mathcal{M}$. But the log-normal polytopes were assumed to be disjoint by the hypothesis. Therefore $\Phi(u) = p$, which means that $u \in \log \text{Vor}_{\mathcal{M}}(p)$, and the result follows. \square

Let $f_1(\theta), \dots, f_r(\theta)$ be nonzero linear polynomials in θ such that $\sum_{i=1}^r f_i(\theta) = 1$. Let Θ be the set such that $f_i(\theta) > 0$ for all $\theta \in \Theta$ and suppose that $\dim \Theta = d$. The model $\mathcal{M} = f(\Theta) \subseteq \Delta_{r-1}$ is called a *discrete linear model* [114, p.152]. Linear models appear in [94, Section 1.2]. An example is DiaNA's model in Example 1.1 of [94].

Theorem 2.1.8. Let \mathcal{M} be a linear model. Then the logarithmic Voronoi cells are equal to their log-normal polytopes.

Proof. We will show that $\log \text{Vor}_{\mathcal{M}}(p) = \log N_p \mathcal{M} \cap \Delta_{n-1}$. The \subset direction follows from Lemma 2.1.2 since an affine linear subspace intersected with Δ_{n-1} is smooth. For the reverse direction, let $u \in \log N_p \mathcal{M} \cap \Delta_{n-1}$. We must show $\Phi(u) = p$. Since ℓ_u is strictly concave on Δ_{n-1} , it is strictly concave when restricted to any convex subset, such as the affine-linear subspace \mathcal{M} . Therefore there is only one critical point. Since \mathcal{M} is smooth, u must be in the log-normal space of $\Phi(u)$, and so $\Phi(u)$ must be p . \square

Example 2.1.9. From the above results, one might hope that finite unions of linear models would admit logarithmic Voronoi cells which are polytopes. However, this is not the case. The log-normal spaces from two disjoint linear models can meet in such a way that the boundary created on a logarithmic Voronoi cell is nonlinear. For an explicit example with two linear models, $\mathcal{M} := \mathcal{M}_1 \cup \mathcal{M}_2 \subset \Delta_3 \subset \mathbb{R}^4$, see Figure 2.3. Here

$$\begin{aligned} \mathcal{M}_1 &= \{(1-s)p_1 + sp_2 : s \in [0, 1] \subset \mathbb{R}\}, \\ \mathcal{M}_2 &= \{(1-t)q_1 + tq_2 : t \in [0, 1] \subset \mathbb{R}\}, \\ p_1 &= (1/5, 1/5, 3/5, 0), \\ p_2 &= (1/7, 3/7, 0, 3/7), \\ q_1 &= (1/13, 9/13, 3/13, 0), \\ q_2 &= (4/13, 4/13, 0, 5/13). \end{aligned}$$

We sampled 3000 points from the log-normal polytope at a given point

$$p = (19/105, 29/105, 2/5, 1/7) \in \mathcal{M}_1$$

and colored them blue or red depending on if their MLE was p or if their MLE was located on \mathcal{M}_2 . Therefore, the blue convex set is the logarithmic Voronoi cell at $p \in \mathcal{M}$.¹

¹Computations for this example can be found at <https://mathrepo.mis.mpg.de/logarithmicVoronoi>




Figure 2.3: Nonlinear boundary arising from two disjoint linear models

Next we consider log-linear, or toric, models. These include many important families of statistical models, such as undirected graphical models [55], independence models [114], and others. For an $m \times n$ integer matrix A with $\mathbf{1} \in \text{rowspan}(A)$, the corresponding *log-linear model* \mathcal{M}_A is defined to be the set of all points $p \in \Delta_{n-1}$ such that $\log(p) \in \text{rowspan}(A)$ [114, p. 122].

Theorem 2.1.10. Let $A \in \mathbb{Z}^{m \times n}$ be an integer matrix such that $\mathbf{1} \in \text{rowspan } A$. Let \mathcal{M} be the associated log-linear (toric) model. Then for any point $p \in \mathcal{M}$, the log-Voronoi cell at p is equal to the log-normal polytope at p .

Proof. We will show that $\log \text{Vor}_{\mathcal{M}}(p) = \log N_p \mathcal{M} \cap \Delta_{n-1}$. The forward direction follows from Lemma 2.1.2, since these models are smooth off the coordinate hyperplanes (see [114, p.150] and [8]). For the reverse direction, let $u \in \log N_p \mathcal{M}$. Although the log-likelihood function can have many complex critical points, it is strictly concave on log-linear models \mathcal{M} for positive u , in particular for $u \in \Delta_{n-1}$. This means that there is exactly one critical point in the positive orthant, and it is the unique solution $p \in \mathcal{M}$ to the linear system $Ap = Au$. [47, Prop. 2.1.5]. This is known as *Birch's Theorem*. It follows that $\Phi(u) = p$, as desired. \square

As a corollary, the polytopes shown in Figure 1.3 and Figure 2.2 are logarithmic Voronoi cells. Following [87], define the map sending a point in projective space to a convex combination of the columns a_i of A , so that the image is a polytope, namely

$$\begin{aligned} \phi_A : \mathbb{P}_{\mathbb{C}}^{n-1} &\rightarrow \mathbb{R}^m \\ z &\mapsto \frac{1}{\sum_{i=1}^n |z_i|} \sum_{i=1}^n |z_i| a_i. \end{aligned}$$

This restricts to what [87, p.120] calls the *algebraic moment map* $\phi_A|_{\mathcal{M}_A} = \mu_A : \mathcal{M}_A \rightarrow \mathbb{R}^m$, where \mathcal{M}_A is the projective toric variety associated to A . The maximum likelihood estimator, then, is the map $\mu_A^{-1} \circ \phi_A$ restricted to Δ_{n-1} , identified as a subset of $\mathbb{P}_{\mathbb{C}}^{n-1}$ by extending scalars and using the quotient map defining projective space. The fact [87, Corollary 8.24] that there is a unique preimage, allowing the definition of μ_A^{-1} , played a crucial role in Theorem 2.1.10. Thus we have the following

Corollary 2.1.11. For toric models, logarithmic Voronoi cells are the preimages $\phi_A^{-1}(\mu_A(p))$ intersected with Δ_{n-1} . Thus, $\phi_A|_{\Delta_{n-1}}$ is a map whose image is a polytope and whose fibres are also polytopes.

For the Segre embedding of Example 2.1.5, the image is a square and the fibres are line segments, depicted in Figure 2.4, which adjoins our Figure 2.2 with [87, Figure 2, p.121]. For more on the algebraic moment map, see [107].




Figure 2.4: The fibres and image of the moment map for the Segre of Example 2.1.5

2.2 The chaotic universe model

Consider running experiments with sample size d and choosing the model defined by

$$\mathcal{M} := \frac{\mathbb{Z}^n \cap d \cdot \Delta_{n-1}}{d}.$$

Philosophically, \mathcal{M} is the *chaotic universe model*. Adopting this model is to abandon the idea that experiments tell us about some simpler underlying truth, since the experimental data will always lie exactly on the model. In this section we investigate the Euclidean and logarithmic Voronoi cells for $p \in \mathcal{M}_{n,d}$. Our motivation to study this model is historical, since Georgy Voronoi was interested in lattices and the partitions of Euclidean space they induce from the closest-point map. These became *Voronoi cells*. It led us to study the logarithmic Voronoi cells coming from maximum likelihood estimation for a lattice intersected with the

probability simplex. In doing so, we found interesting connections with root polytopes of type A and were able to generalize [31, Theorem 1] to our context, finding a complete combinatorial description of the face structure of the logarithmic Voronoi cells for $\mathcal{M}_{n,d}$. We give more historical context in the end of the section.

For convenience we work with the scaled set $d \cdot \Delta_{n-1}$ since all polytopes considered will be combinatorially equivalent to those we could define in Δ_{n-1} . Then we define $\mathcal{M}_{n,d}$ as the $N := \binom{n+d-1}{d}$ nonnegative integer vectors summing to d . Thus $(p_1, p_2, \dots, p_n) = p \in \mathcal{M}_{n,d}$ has all coordinates $p_i \in \mathbb{N}$. These vectors can be used to create a projective toric variety, the d th Veronese embedding of \mathbb{P}^{n-1} into \mathbb{P}^{N-1} [87, Chapter 8], but instead we treat them as the model itself. By Proposition 2.1.1, the logarithmic Voronoi cells for $p \in \mathcal{M}_{n,d}$ are polytopes. For any $p \in \mathcal{M}_{n,d}$ such that all coordinates $p_i > 1$, we will provide a full characterization of the faces of the corresponding logarithmic root polytopes in Theorem 2.2.2. Theorem 2.2.4 shows that these logarithmic root polytopes are dual to the logarithmic Voronoi cells. These are the main results of the section. Again using orthogonal projection from \mathbb{R}^4 , Figure 2.5 shows all the logarithmic Voronoi cells for interior points of $\mathcal{M}_{4,9}$ and $\mathcal{M}_{4,10}$.




Figure 2.5: Logarithmic Voronoi cells (rhombic dodecahedra) of interior points for $n = 4$, $d = 9$ (on the left) and $d = 10$ (on the right).

The Euclidean Voronoi cells for $p \in \mathcal{M}_{n,d}$ are the duals of *root polytopes* of type A_{n-1} , i.e. the facets are defined by inequalities whose normal vectors are $\{e_i - e_j : i \neq j\}$. Root polytope often refers to the convex hull of the origin and the positive roots $\{e_i - e_j : i < j\}$. These were studied in [56] in terms of their relationship to certain hypergeometric functions. However, we define root polytopes to be the convex hull of *all* roots, as studied in [31]. We also note that these polytopes are *Young orbit polytopes* for the partition $(n - 1, 1)$ and find

application in combinatorial optimization [93].

Denote the $(n - 1)$ -dimensional root polytope by $P_n \subset \mathbb{R}^n$, so that the Euclidean Voronoi cells of $p \in \mathcal{M}_{n,d}$ are the dual P_n^* . The volume of P_n is equal to $\frac{n}{(n-1)!}C_{n-1}$, where C_{n-1} is a Catalan number. Every nontrivial face of P_n is a Cartesian product of two simplices, and corresponds to a pair of nonempty, disjoint subsets $I, J \subset [n]$. Every m -dimensional face of P_n is the convex hull of the vectors $\{e_i - e_j : i \in I, j \in J\}$ with $|I| + |J| = m + 2$, so there is a bijection between nontrivial faces and the set of ordered partitions of subsets of $[n]$ with two blocks [31, Theorem 1]. This result is related to the face description of Π_{n-1} , the permutahedron, since P_n is a *generalized permutahedron* and can be obtained by collapsing certain faces of Π_{n-1} .

In the logarithmic setting, analogous polytopes $\log P_n(p)$ exist, playing the same role as the root polytopes in the Euclidean case. However, their details are more complicated. The correct modifications motivate the following definition.

Definition 2.2.1. The *logarithmic root polytope* for $p \in \mathcal{M}_{n,d}$ is defined as the convex hull of the $2\binom{n}{2}$ vertices v_{ij} for $i \neq j \in [n]$ given by the formulas

$$v_{ij} := \frac{1}{b_j p_j - a_i p_i} \left[a_i e_i - b_j e_j - \frac{(a_i - b_j)}{n} \mathbf{1} \right]$$

where

$$a_i := \log\left(\frac{p_i+1}{p_i}\right) \quad b_j := \log\left(\frac{p_j}{p_j-1}\right)$$

and where $\mathbf{1} := \sum_{k \in [n]} e_k$. Note that $a_i, b_j > 0$ are always positive real numbers and all vectors v_{ij} are orthogonal to $\mathbf{1}$. We denote the polytope by $\log P_n(p)$.

The statement and proof of the following Theorem 2.2.2 was inspired by and closely follows [31, Theorem 1]. However, significant details needed to be modified. For example, the linear functional

$$g = (1, 0, -1, 1, -1, 0, -1)$$

is replaced by

$$\begin{bmatrix} -a_1 a_4 b_3 b_5 p_1 - a_1 a_4 b_3 b_7 p_1 - a_1 a_4 b_5 b_7 p_1 - a_1 b_3 b_5 b_7 p_1 + a_4 b_3 b_5 b_7 p_3 + a_4 b_3 b_5 b_7 p_4 + a_4 b_3 b_5 b_7 p_5 + a_4 b_3 b_5 b_7 p_7 \\ 0 \\ a_1 a_4 b_5 b_7 p_1 - a_1 a_4 b_3 b_5 p_3 - a_1 a_4 b_3 b_7 p_3 - a_1 b_3 b_5 b_7 p_3 - a_4 b_3 b_5 b_7 p_3 + a_1 a_4 b_5 b_7 p_4 + a_1 a_4 b_5 b_7 p_5 + a_1 a_4 b_5 b_7 p_7 \\ a_1 b_3 b_5 b_7 p_1 + a_1 b_3 b_5 b_7 p_3 - a_1 a_4 b_3 b_5 p_4 - a_1 a_4 b_3 b_7 p_4 - a_1 a_4 b_5 b_7 p_4 - a_4 b_3 b_5 b_7 p_4 + a_1 b_3 b_5 b_7 p_5 + a_1 b_3 b_5 b_7 p_7 \\ a_1 a_4 b_3 b_7 p_1 + a_1 a_4 b_3 b_7 p_3 + a_1 a_4 b_3 b_7 p_4 - a_1 a_4 b_3 b_5 p_5 - a_1 a_4 b_5 b_7 p_5 - a_1 b_3 b_5 b_7 p_5 - a_4 b_3 b_5 b_7 p_5 + a_1 a_4 b_3 b_7 p_7 \\ 0 \\ a_1 a_4 b_3 b_5 p_1 + a_1 a_4 b_3 b_5 p_3 + a_1 a_4 b_3 b_5 p_4 + a_1 a_4 b_3 b_5 p_5 - a_1 a_4 b_3 b_7 p_7 - a_1 a_4 b_5 b_7 p_7 - a_1 b_3 b_5 b_7 p_7 - a_4 b_3 b_5 b_7 p_7 \end{bmatrix}.$$

This linear functional plays the same role for the logarithmic root polytope at $(p_1, \dots, p_7) \in \mathcal{M}_{7,d}$ as g plays for the usual root polytope in the proof of [31, Theorem 1].

Theorem 2.2.2. For $m \in \{0, 1, \dots, n-2\}$, every m -dimensional face of the logarithmic root polytope for $p \in \mathcal{M}_{n,d}$ is given by the convex hull of the vertices v_{ij} for $i \in I, j \in J$, where I, J are disjoint nonempty subsets of $[n]$ such that $|I| + |J| = m + 2$. Thus there is a bijection between nontrivial faces and the set of ordered partitions of subsets of $[n]$ with two blocks, where the dimension of the face corresponding to (I, J) is $|I| + |J| - 2$.

Proof. Each face of a polytope can be described as the subset of the polytope maximizing a linear functional. Recall that we have fixed some $p \in \mathcal{M}_{n,d}$ with all $p_k > 1$ and that

$$a_i := \log\left(\frac{p_i+1}{p_i}\right) \quad \text{and} \quad b_j := \log\left(\frac{p_j}{p_j-1}\right).$$

In our formula (2.2) we use a shorthand for writing square-free monomials in the a_1, a_2, \dots, a_n and the b_1, b_2, \dots, b_n . For example if $I = \{1, 2, 4\}$ then $a^I = a_1 a_2 a_4$, while if $J = \{3, 5\}$ then $b^J = b_3 b_5$. For a pair of disjoint nonempty subsets I, J of $[n]$ we define the linear functional $g_{IJ} = (g_1, g_2, \dots, g_n) \in (\mathbb{R}^n)^*$ by the formulas

$$\begin{aligned} \text{If } \ell \in I, \quad g_\ell &= \sum_{i \in I \setminus \ell} a^{I \setminus \{\ell, i\}} b^J (a_i p_i - a_\ell p_\ell) + \sum_{j \in J} a^{I \setminus \ell} b^{J \setminus j} (b_j p_j - a_\ell p_\ell) \\ \text{If } \ell \in J, \quad g_\ell &= \sum_{i \in I} a^{I \setminus i} b^{J \setminus \ell} (a_i p_i - b_\ell p_\ell) + \sum_{j \in J \setminus \ell} a^I b^{J \setminus \{\ell, j\}} (b_j p_j - b_\ell p_\ell) \\ \text{Else,} \quad g_\ell &= 0. \end{aligned} \tag{2.2}$$

Then the convex hull of the vectors $\{v_{ij} : i \in I, j \in J\}$ is the face maximizing g_{IJ} . To see this, first note that $g_{IJ} \cdot \mathbf{1} = 0$. Because of this fact we can ignore the component of v_{ij} in the $\mathbf{1}$ direction. Recall that

$$v_{ij} := \frac{1}{b_j p_j - a_i p_i} \left[a_i e_i - b_j e_j - \frac{(a_i - b_j)}{n} \mathbf{1} \right],$$

so that to evaluate g_{IJ} on v_{ij} it is enough to evaluate on

$$\frac{1}{b_j p_j - a_i p_i} [a_i e_i - b_j e_j].$$

Recalling that the a_i and b_j are always positive and that the $p_k > 1$, it can be seen that g_{IJ} takes equal value on every vertex v_{rs} for $r \in I, s \in J$, and strictly less on every other vertex. We omit the details of the admittedly lengthy calculation, but note that the common maximum value attained on all vertices v_{rs} for $r \in I, s \in J$, is equal to

$$\sum_{i \in I} a^{I \setminus i} b^J + \sum_{j \in J} a^I b^{J \setminus j}.$$

Conversely, given an arbitrary linear functional $f = (f_1, f_2, \dots, f_n)$ determining a nontrivial face F , collect the indices where its components are nonnegative in a set I and the indices where its components are negative in a set J . Then (I, J) is a partition of $[n]$ and we refer

to the same formulas (2.2) as above in order to define the sets (I', J') as follows. If $I \neq \emptyset$ and $J \neq \emptyset$ then let

$$\begin{aligned} I' &:= \{i : f_i/g_i = \max(f_\ell/g_\ell : \ell \in I)\} \\ J' &:= \{j : f_j/g_j = \max(f_\ell/g_\ell : \ell \in J)\}. \end{aligned}$$

If $I = \emptyset$ then let

$$\begin{aligned} I' &:= \{i : f_i/g_i = \min(f_\ell/g_\ell : \ell \in J)\} \\ J' &:= \{j : f_j/g_j = \max(f_\ell/g_\ell : \ell \in J)\}, \end{aligned}$$

while if $J = \emptyset$ then let

$$\begin{aligned} I' &:= \{i : f_i/g_i = \max(f_\ell/g_\ell : \ell \in I)\} \\ J' &:= \{j : f_j/g_j = \min(f_\ell/g_\ell : \ell \in I)\}. \end{aligned}$$

Note that the face F is the convex hull of the vectors $\{v_{ij} : i \in I', j \in J'\}$ and hence (I', J') are determined independently of the choice of linear functional which maximizes the given face.

Now we show that the dimension of the face corresponding to disjoint nonempty sets I, J of $[n]$ is $|I| + |J| - 2$. Let $I = \{i_1, \dots, i_{|I|}\}$ and $J = \{j_1, \dots, j_{|J|}\}$. Then

$$X = \{v_{i_1, j_\ell} : \ell = 1, \dots, |J|\} \cup \{v_{i_\ell, j_1} : \ell = 2, \dots, |I|\}$$

is a maximal linearly independent subset of $|I| + |J| - 1$ of the vectors $v_{ij}, i \in I, j \in J$. In addition, for any $i \in I, j \in J$ either $v_{ij} \in X$ or we can write it as an affine combination (coefficients sum to 1) of vectors in X , namely

$$v_{i,j} = \left(\frac{b_{j_1} p_{j_1} - a_i p_i}{b_j p_j - a_i p_i} \right) v_{i, j_1} - \left(\frac{b_{j_1} p_{j_1} - a_{i_1} p_{i_1}}{b_j p_j - a_i p_i} \right) v_{i_1, j_1} + \left(\frac{b_j p_j - a_{i_1} p_{i_1}}{b_j p_j - a_i p_i} \right) v_{i_1, j}.$$

Hence, X is an affine basis of the face corresponding to I, J , whose dimension is $|X| - 1$, which is $|I| + |J| - 2$ as desired. This completes the proof. \square

Example 2.2.3. Let $n = 6$, $I = \{1, 4\}$, $J = \{2, 3, 5\}$ and $p = (2, 15, 3, 5, 9, 6)$. We implemented the formulas (2.2) in floating point arithmetic (due to the logarithms) and obtain (shown to only three digits)

$$g_{IJ} = (0.00415, -0.00200, -0.00398, 0.00474, -0.00291, -0.000).$$

We can evaluate this linear functional on the vertices v_{ij} for $i \neq j$ where $i, j \in [6]$ and obtain the following values, which attain their maximum on $v_{12}, v_{13}, v_{15}, v_{42}, v_{43}, v_{45}$, as expected.

0.008135843945	$v(1, 2) = (1.56, -0.559, -0.251, -0.251, -0.251, -0.251)$
0.008135843948	$v(1, 3) = (1.00, 0.000, -1.00, 0.000, 0.000, 0.000)$
0.002052114856	$v(1, 4) = (1.23, -0.0997, -0.0997, -0.832, -0.0997, -0.0997)$
0.008135843948	$v(1, 5) = (1.43, -0.192, -0.192, -0.192, -0.665, -0.192)$
0.005950315119	$v(1, 6) = (1.30, -0.132, -0.132, -0.132, -0.132, -0.776)$
-0.007192386292	$v(2, 1) = (-1.41, 0.405, 0.250, 0.250, 0.250, 0.250)$
0.005982647332	$v(2, 3) = (0.229, 0.488, -1.40, 0.229, 0.229, 0.229)$
-0.008044880930	$v(2, 4) = (0.179, 0.615, 0.179, -1.33, 0.179, 0.179)$
0.002322216671	$v(2, 5) = (0.0963, 0.796, 0.0963, 0.0963, -1.18, 0.0963)$
-0.001027169161	$v(2, 6) = (0.156, 0.669, 0.156, 0.156, 0.156, -1.29)$
-0.007691322875	$v(3, 1) = (-1.20, 0.129, 0.679, 0.129, 0.129, 0.129)$
-0.005863205380	$v(3, 2) = (-0.213, -0.615, 1.47, -0.213, -0.213, -0.213)$
-0.008723741580	$v(3, 4) = (-0.0426, -0.0426, 1.10, -0.926, -0.0426, -0.0426)$
-0.004075725208	$v(3, 5) = (-0.144, -0.144, 1.32, -0.144, -0.742, -0.144)$
-0.004962538041	$v(3, 6) = (-0.0762, -0.0762, 1.17, -0.0762, -0.0762, -0.867)$
-0.004242519680	$v(4, 1) = (-1.28, 0.179, 0.179, 0.563, 0.179, 0.179)$
0.008135843941	$v(4, 2) = (-0.153, -0.713, -0.153, 1.33, -0.153, -0.153)$
0.008135843947	$v(4, 3) = (0.122, 0.122, -1.21, 0.720, 0.122, 0.122)$
0.008135843947	$v(4, 5) = (-0.0723, -0.0723, -0.0723, 1.15, -0.865, -0.0723)$
0.004743470845	$v(4, 6) = (0.000, 0.000, 0.000, 1.00, 0.000, -1.00)$
-0.007271750954	$v(5, 1) = (-1.36, 0.224, 0.224, 0.224, 0.464, 0.224)$
-0.001944541355	$v(5, 2) = (-0.0700, -0.866, -0.0700, -0.0700, 1.15, -0.0700)$
0.004878535171	$v(5, 3) = (0.186, 0.186, -1.32, 0.186, 0.579, 0.186)$
-0.008151512920	$v(5, 4) = (0.117, 0.117, 0.117, -1.21, 0.745, 0.117)$
-0.002105123850	$v(5, 6) = (0.0880, 0.0880, 0.0880, 0.0880, 0.811, -1.16)$
-0.006239195419	$v(6, 1) = (-1.31, 0.195, 0.195, 0.195, 0.195, 0.528)$
0.001256424608	$v(6, 2) = (-0.129, -0.756, -0.129, -0.129, -0.129, 1.27)$
0.005540018448	$v(6, 3) = (0.144, 0.144, -1.25, 0.144, 0.144, 0.672)$
-0.005547164875	$v(6, 4) = (0.0602, 0.0602, 0.0602, -1.11, 0.0602, 0.867)$
0.002536892813	$v(6, 5) = (-0.0448, -0.0448, -0.0448, -0.0448, -0.915, 1.09)$

Theorem 2.2.4. The logarithmic Voronoi cells for $p \in \mathcal{M}_{n,d}$ with all $p_i > 1$ are the dual polytopes $(\log P_n(p))^*$ of the logarithmic root polytopes $\log P_n(p)$.

Proof. Given a point $p \in \mathcal{M}_{n,d}$, the logarithmic Voronoi cell can be defined as the intersection of $d \cdot \Delta_{n-1}$ with all the halfspaces $H_q(u) \geq 0$ for all points $q \in \mathcal{M}_{n,d}$ with $q \neq p$, where

$$H_q(u) := \sum_{i \in [n]} u_i \log \left(\frac{p_i}{q_i} \right).$$

We say that this system of inequalities is *sufficient* to define the logarithmic Voronoi cell. However, not all of these inequalities are necessary. Lemma 2.2.5 shows that a certain set of $2 \binom{n}{2}$ inequalities is sufficient for all $n \in \mathbb{Z}_{\geq 2}$. These are the inequalities $H_q(u) \geq 0$ for $q = p + e_i - e_j$ for $i \neq j$. We avoid logarithms of zero since $p_k > 1$ and we are away from the simplex boundary. In other words, we get one inequality from every point q reachable from p by moving along a root of type A_{n-1} .

These $H_q(u) \geq 0$ inequalities are linear, with constant term zero. However, projecting the normal vectors of these hyperplanes along the all ones vector $\mathbf{1}$ and viewing p as the origin

of a new coordinate system, we obtain inequalities with nonzero constant terms. These inequalities describe the same logarithmic Voronoi polytope on the hyperplane $\sum_k u_k = d$. Dividing each inequality by the constant terms we obtain a system of inequalities which is of the form $Au \leq \mathbf{1}$, following the notation of [126], where the rows of A are exactly the vectors v_{ij} . By [126, Theorem 2.11], the dual polytope is given by the convex hull of these v_{ij} . \square

Lemma 2.2.5. Let $p \in \mathcal{M}_{n,d}$ with every entry $p_i > 1$. A sufficient system of inequalities defining the logarithmic Voronoi cell is given by the $2\binom{n}{2}$ halfspaces $u \in \mathbb{R}^n$ such that $H_\delta(u) \geq 0$ for $\delta \in R := \{e_i - e_j : i \neq j, i, j \in [n]\}$ and the affine plane $\sum u_i = d$, where

$$H_\delta(u) := \sum_{i \in [n]} u_i \log \left(\frac{p_i}{p_i + \delta_i} \right).$$

Proof. We prove that the $2\binom{n}{2}$ inequalities $H_\delta(u) \geq 0$ for $\delta \in R$ are sufficient. Fix $p \in \mathcal{M}$ with all $p_i > 1$. Let $u \in \mathbb{R}^n$ such that $H_\delta(u) \geq 0$ for all $\delta \in R$. Fix some $q = p + \delta + \delta'$ where $\delta, \delta' \in R$, and assume that $\delta + \delta' \notin R$. We wish to show $H_q(u) = \sum_i u_i \log \frac{p_i}{q_i} \geq 0$. Consider several cases. First, if $\delta = \delta' = e_j - e_k$, it suffices to show that

$$u_j \log \frac{p_j}{p_j + 2} + u_k \log \frac{p_k}{p_k - 2} \geq 0.$$

We claim that

$$u_j \log \frac{p_j}{p_j + 2} + u_k \log \frac{p_k}{p_k - 2} \geq 2u_j \log \frac{p_j}{p_j + 1} + 2u_k \log \frac{p_k}{p_k - 1}, \quad (2.3)$$

which would be sufficient, since the right-hand side of the above equation is ≥ 0 by assumption. We show that

$$u_j \log \frac{p_j}{p_j + 2} \geq 2u_j \log \frac{p_j}{p_j + 1} \quad \text{and} \quad u_k \log \frac{p_k}{p_k - 2} \geq 2u_k \log \frac{p_k}{p_k - 1}. \quad (2.4)$$

Observe:

$$\begin{aligned} u_j \log \frac{p_j}{p_j + 2} \geq 2u_j \log \frac{p_j}{p_j + 1} &\iff p_j^2 + 2p_j + 1 \geq p_j^2 + 2p_j, \\ u_k \log \frac{p_k}{p_k - 2} \geq 2u_k \log \frac{p_k}{p_k - 1} &\iff p_k^2 - 2p_k + 1 \geq p_k^2 - 2p_k. \end{aligned}$$

Thus (2.4) holds, and we conclude that (2.3) is true in this case, as desired. If $\delta \neq \delta'$, but they share both indices, then $p = q$, and we're done. If they do not share any indices, then we have that $H_q(u) = H_\delta(u) + H_{\delta'}(u) \geq 0$ by assumption. Suppose $\delta \neq \delta'$, and δ and δ' share one index, j . If $\delta = e_i - e_j$ and $\delta' = e_j - e_k$ for $i \neq j \neq k$, then $\delta + \delta' = e_i - e_k$, a

contradiction to the assumption $\delta + \delta' \notin R$. Similarly when $\delta = e_j - e_i$ and $\delta' = e_k - e_j$. Suppose then that $\delta = e_i - e_j$ and $\delta' = e_i - e_k$. We wish to show that

$$u_i \log \frac{p_i}{p_i + 2} + u_j \log \frac{p_j}{p_j - 1} + u_k \log \frac{p_k}{p_k - 1} \geq 0.$$

Note then that

$$u_i \log \frac{p_i}{p_i + 2} \geq 2u_i \log \frac{p_i}{p_i + 1} \iff p_i^2 + 2p_i + 1 \geq p_i^2 + 2p_i,$$

and the last inequality always holds for positive p_i , so the lemma is true for this case. The case when $\delta = e_j - e_i$ and $\delta' = e_k - e_i$ is proved similarly. Since $H_q(u) \geq 0$ in all of the cases we considered, and the cases are exhaustive, we conclude that the lemma holds. \square

A family of polytopes. Using the face characterization of Theorem 2.2.2, we may compute the f -vectors of the logarithmic Voronoi cells for any n . We can also numerically calculate the f -vector for the logarithmic Voronoi cell at any specific point $p \in \mathcal{M}_{n,d}$ with $p_i > 1, \forall i$ by explicitly constructing the polytope using inequalities. Of course, both of these calculations match. Below we list the f -vectors (which we computed in both ways) for small values $n \in \{2, 3, 4, 5, 6, 7\}$ to give the reader a sense for their behavior. The logarithmic Voronoi cells for every $\mathcal{M}_{n,d}$ are combinatorially isomorphic to the dual of the corresponding root polytope, exactly as in the Euclidean case.

$$\begin{aligned} n = 2 & (1, 2, 1) \\ n = 3 & (1, 6, 6, 1) \\ n = 4 & (1, 14, 24, 12, 1) \\ n = 5 & (1, 30, 70, 60, 20, 1) \\ n = 6 & (1, 62, 180, 210, 120, 30, 1) \\ n = 7 & (1, 126, 434, 630, 490, 210, 42, 1) \end{aligned}$$

Therefore we have a family of Euclidean Voronoi polytopes that tile \mathbb{R}^{n-1} and a family of logarithmic Voronoi polytopes that tile the open simplex Δ_{n-1} . This family begins

$$\begin{array}{ccccccc} n - 1 = 1 & n - 1 = 2 & n - 1 = 3 & \dots & & & \\ \text{line segment} & \text{hexagon} & \text{rhombic dodecahedron} & \dots & & & \end{array}$$

Root polytopes of type A have connections to tropical geometry. The rhombic dodecahedron is a polytrope which has been called the 3-pyrope because of the mineral $\text{Mg}_3\text{Al}_2(\text{SiO}_4)_3$ whose pure crystal can take the same shape. For more on root polytopes, tropical geometry, and polytropes, see [72].

Historical comments. Georgy Voronoi devoted many years of his life to studying properties of 3-dimensional parallelohedra, convex polyhedra that tessellate 3-dimensional Euclidean space. His paper on the subject called *Recherches sur les paralléloèdres primitifs*

[120] was a result of his twelve-year work. In a cover letter to the manuscript, he wrote: “I noticed already long ago that the task of dividing the n -dimensional analytical space into convex congruent polyhedra is closely related to the arithmetic theory of positive quadratic forms” [116]. Indeed, Voronoi was interested in studying cells of lattices in \mathbb{Z}^n with the aim of applying them to the theory of quadratic forms. This motivated us to study a lattice intersected with the probability simplex, the topic of our current section. Today, Voronoi decomposition finds applications to the analysis of spatially distributed data in many fields of science, including mathematics, physics, biology, archaeology, and even cinematography. In [124], the author uses Voronoi cells to optimize search paths in an attempt to improve the final 6-minute scene of Andrei Tarkovsky’s *Offret (the Sacrifice)*. Voronoi diagrams are so versatile they even found their way into baking: Ukrainian pastry chef Dinara Kasko uses Voronoi diagrams to 3D-print silicone molds which she then uses to make cakes [74].

2.3 Numerical algebraic geometry

In case the logarithmic Voronoi cell is not a polytope, techniques from numerical nonlinear algebra can still be used to compute it effectively. In this section we demonstrate these methods. In particular, we explain how to set up the randomization that must be used in case the algebraic variety is defined by more polynomials than its codimension. The relevant equation (2.5) is explained below, and then used in Theorem 2.3.1. Numerical algebraic geometry [20, 106] can be used to efficiently find all isolated solutions of a square system of polynomial equations (square means equal number of equations and variables). The system of equations used in Theorem 2.3.1 formulates our problem specifically to take advantage of these tools.

Let f be the $1 \times m$ row vector whose entries are the polynomials f_1, \dots, f_m in the variables x_1, \dots, x_n . We assume that the first polynomial defines the simplex, i.e. $f_1 = \sum_{i=1}^n x_i - 1$. Let the algebraic set defined by f_1, \dots, f_m have codimension c . Let df denote the $m \times n$ Jacobian matrix whose rows are the gradients of f_1, \dots, f_m . Let A be a $c \times m$ matrix whose entries are chosen randomly from independent normal distributions. Let B be a similarly chosen random $(m - c) \times (n + c)$ matrix. Let $[\lambda - 1]$ be the row vector of length $c + 1$ whose first c entries are variables $\lambda_1, \dots, \lambda_c$ and whose last entry is -1 and let I_{n+c} be the identity matrix of size $n + c$. We are interested in the following vector equation whose components give $n + c$ polynomial equations in $n + c$ unknowns:

$$\underbrace{\begin{bmatrix} [\lambda - 1] & \begin{bmatrix} A \cdot df \\ \nabla \ell_u \end{bmatrix} & f \end{bmatrix}}_{1 \times (n+m)} \begin{bmatrix} I_{n+c} \\ B \end{bmatrix} = \underbrace{[0 \ \dots \ 0]}_{1 \times (n+c)}. \quad (2.5)$$

Theorem 2.3.1. Let \mathcal{M} be the intersection of Δ_{n-1} and an irreducible algebraic model given by the polynomial map $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Let $u \in \Delta_{n-1}$ be fixed and generic. With

probability 1, all points $p \in \mathcal{M}$ such that $u \in \log N_p \mathcal{M}$ are among the finitely many isolated solutions to the square system of equations given in (2.5).

Proof. As a consequence of [67, Theorem 1.6], if $u \in \Delta_{n-1}$ is generic, then with probability 1 there will be finitely many critical points of ℓ_u restricted to \mathcal{M} . If the algebraic set defined by f has codimension c then the dimension of the rowspace of df will be equal to c and the rows will span $N_x \mathcal{M}$ for any generic $x \in \mathcal{M}$ [104, p.93]. With probability 1, multiplying by the random matrix A will result in a $c \times n$ matrix of full row rank, whose rows also span $N_x \mathcal{M}$. Appending the row $\nabla \ell_u$ and multiplying the resulting matrix by the row vector $[\lambda - 1]$ produces n polynomials which evaluate to zero whenever $\nabla \ell_u$ is in the normal space $N_x \mathcal{M}$. Appending the polynomials f_1, \dots, f_m gives a $1 \times (n + m)$ row vector of polynomials evaluating to zero whenever $x \in \mathcal{M}$ and $\nabla \ell_u$ lies in the normal space $N_x \mathcal{M}$. However, this system of equations is overdetermined. Applying Bertini's theorem [20, Theorem 9.3] or [106, Theorem A.8.7] we can take random linear combinations of these polynomials using I_{n+c} and B , and with probability 1, the isolated solutions of the resulting square system of polynomials will contain all isolated solutions of the original system of equations. The result follows. \square

Remark 2.3.2. If we are interested in computing the logarithmic Voronoi cell at a specific point $p \in \mathcal{M}$, then we can generate a generic point $u_0 \in \log N_p \mathcal{M}$ by taking a random linear combination of the gradients of f_1, \dots, f_m . Using this point u_0 we can formulate our system of equations (2.5), one of whose solutions we already know, namely p . Using monodromy, we can quickly find many other solutions p' by perturbing our parametrized system of equations through a loop in parameter space. For more details, see [4]. This is especially useful in the case where the ML degree is known a priori, since we can stop our monodromy search after finding ML degree many solutions. This process yields an optimal start system for homotopy continuation, allowing us to almost immediately compute solutions for other data points since we need only follow the ML degree-many solution paths via homotopy continuation.

In the next example, we utilize the formulation in Theorem 2.3.1 to numerically compute a logarithmic Voronoi cell in a larger example of statistical interest. Namely, we consider a mixture of two binomial distributions, also known as a secant variety.

Example 2.3.3. Bob has three biased coins, one in each pocket, and one in his hand. He flips the coin in his hand, and depending on the outcome, chooses either the coin in his left or right pocket, which he then flips 5 times, recording the total number of heads in the last 5 flips. To estimate the biases of Bob's coins, Alice treats this situation as a 3-dimensional statistical model $\mathcal{M} \subset \Delta_5 \subset \mathbb{R}^6$. Using implicitization [87, Section 4.2], Alice derives the

following algebraic equations describing \mathcal{M} :

$$f(x) = \begin{bmatrix} 20x_1x_3x_5 - 10x_1x_4^2 - 8x_2^2x_5 + 4x_2x_3x_4 - x_3^3 \\ 100x_1x_3x_6 - 20x_1x_4x_5 - 40x_2^2x_6 + 4x_2x_3x_5 + 2x_2x_4^2 - x_3^2x_4 \\ 100x_1x_4x_6 - 40x_1x_5^2 - 20x_2x_3x_6 + 4x_2x_4x_5 + 2x_3^2x_5 - x_3x_4^2 \\ 20x_2x_4x_6 - 8x_2x_5^2 - 10x_3^2x_6 + 4x_3x_4x_5 - x_4^3 \\ x_1 + x_2 + x_3 + x_4 + x_5 + x_6 - 1 \end{bmatrix}.$$

For a concrete example, consider the point which arises by setting the biases of the coins to $b_1 = \frac{7}{11}, b_2 = \frac{3}{5}, b_3 = \frac{3}{7}$. Explicitly this point $p \in \mathcal{M}$ is

$$p = \left(\frac{518}{9375}, \frac{124}{625}, \frac{192}{625}, \frac{168}{625}, \frac{86}{625}, \frac{307}{9375} \right).$$

The log-normal space $\log N_p \mathcal{M}$ is 3-dimensional, becoming a 2-dimensional polytope when intersected with $\Delta_5 \subset \mathbb{R}^6$. This intersection is the log-normal polytope, in this case, a hexagon. In fact, this hexagon is the (2-dimensional) convex hull of the following six vertices:

$$\begin{aligned} & \left(0, \frac{651}{1625}, 0, \frac{30569}{58500}, \frac{43}{2250}, \frac{3377}{58500} \right) \\ & \left(0, \frac{124}{375}, \frac{88}{375}, \frac{77}{375}, \frac{86}{375}, 0 \right) \\ & \left(\frac{8288}{76875}, 0, \frac{3176}{5125}, 0, \frac{1376}{5125}, \frac{307}{76875} \right) \\ & \left(\frac{259}{1875}, 0, \frac{52}{125}, \frac{91}{250}, 0, \frac{307}{3750} \right) \\ & \left(\frac{518}{76875}, \frac{1984}{5125}, 0, \frac{2779}{5125}, 0, \frac{4912}{76875} \right) \\ & \left(\frac{2849}{29250}, \frac{31}{1125}, \frac{8734}{14625}, 0, \frac{903}{3250}, 0 \right). \end{aligned}$$

By choosing an orthonormal basis agreeing with $\log N_p \mathcal{M}$ we can plot this hexagon, though it lives in \mathbb{R}^6 . Figure 2.1 shows the log-normal polytope and our numerical approximation of the logarithmic Voronoi cell (which is not a polytope) surrounding the point p . By rejection sampling, we computed 60000 points $u_1, u_2, \dots, u_{60000} \in \log N_p \mathcal{M} \cap \Delta_{n-1}$ in the log-normal polytope. By a result in [63], we know that the ML degree of this model is 39. Using the formulation presented in Theorem 2.3.1, we successfully computed all 39 complex critical points for each $\ell_{u_i}, i \in \{1, 2, \dots, 60000\}$ restricted to \mathcal{M} . We easily find each $\Phi(u_i)$ by comparing the 39 values, choosing the maximum. If $p = \Phi(u_i)$ then $u_i \in \log \text{Vor}_{\mathcal{M}}(p)$ and we color that point green in Figure 2.1, while if $p \neq \Phi(u_i)$ we color the point pink. The repeated computations of each set of 39 critical points were accomplished using the software `HomotopyContinuation.jl` [27], which can efficiently compute the isolated solutions to systems of polynomial equations using homotopy continuation [20, 106]. A full description of the `Julia` code needed to compute this example can be found online at [4].

Conclusion. In this chapter, we proved that for finite models, linear models, toric models, and models of ML degree one, logarithmic Voronoi cells are equal to log-normal polytopes. We also characterized the combinatorial types of these polytopes for the finite chaotic universe model. For logarithmic Voronoi cells that are not polytopes, we described an approach to reliably compute them numerically via homotopy continuation and monodromy. The next two chapters dive deeper into the combinatorics of logarithmic Voronoi polytopes of linear and toric models, respectively. In particular, we will study how the combinatorial type of these polytopes changes as we vary points on the model, and how this helps us maximize information divergence, introduced in Section 1.1.

Chapter 3

Linear models

In this chapter we study partitions of the probability simplex Δ_{n-1} into combinatorially equivalent polytopes. These partitions are induced by linear statistical models. In Figure 3.1, such models are given by the dotted line segments. Specifically, we focus on the combinatorics of logarithmic Voronoi cells of linear models, which are polytopes by Theorem 2.1.8. First, we describe the vertices of these polytopes in terms of certain co-circuits. We then show that logarithmic Voronoi polytopes at all points on a linear model have the same combinatorial type. We also study these polytopes at the points on the boundary of a linear model. In particular, we give a condition for them to have the same combinatorial type as those at the interior points. Finally, we focus on logarithmic Voronoi cells of partial linear models, where the points on the boundary of the model are especially of interest. This chapter is based on [1] and [6].




Figure 3.1: Partition of the tetrahedron Δ_3 into triangles (left) and quadrilaterals (right).

Let us now recall the formal definition of a linear model from Section 2.1.

Definition 3.0.1. Let f_1, \dots, f_n be linear polynomials in $\theta = (\theta_1, \dots, \theta_d)$ with $\sum_{i=1}^n f_i(\theta) = 1$. Let $\Theta \subseteq \mathbb{R}^d$ be the d -dimensional parameter space where $f_i(\theta) > 0$ for all $\theta \in \Theta$ and $i = 1, \dots, n$. The image of the map

$$f : \Theta \rightarrow \Delta_{n-1} : \theta \mapsto (f_1(\theta), \dots, f_n(\theta))$$

is called a *discrete linear model* [94, Section 1.2],[114, Section 7.2].

Hence, a linear model is a polytope, obtained by intersecting an affine linear space with a probability simplex. Assume this affine linear space contains a point all of whose coordinates are positive. The dimension of a linear model is the dimension of the corresponding linear space. Given a point $u \in \Delta_{n-1}$, the log-likelihood function ℓ_u is strictly concave on the simplex and hence on any convex subset of the simplex, such as a linear model. Hence, the MLE will always exist and be unique for every $u \in \Delta_{n-1}^\circ$ and their logarithmic Voronoi polytopes partition the probability simplex.

Fix a linear model \mathcal{M} as in the definition. Recall from Chapter 2 that for linear models, logarithmic Voronoi cells are polytopes. Hence, we will refer to their logarithmic Voronoi cells as *logarithmic Voronoi polytopes*. For any point $p \in \mathcal{M}$, we denote the logarithmic Voronoi polytope at p by $\log \text{Vor}_{\mathcal{M}}(p)$. We call a point $p = (p_1, \dots, p_n) \in \mathcal{M}$ *interior* if $p_i > 0$ for all $i \in [n]$. Our hypotheses in Definition 3.0.1 imply that any linear model can be written as

$$\mathcal{M} = \{c - Bx : x \in \Theta\}$$

where B is a $n \times d$ matrix, each of whose columns sums to 0, and $c \in \mathbb{R}^n$ is a vector, whose coordinates sum to 1.

3.1 Combinatorial types

In this section, we give an explicit combinatorial description of logarithmic Voronoi polytopes at interior points on the linear model \mathcal{M} . We find that these polytopes have the same combinatorial type. The next proposition gives a formula for the vertices of those polytopes. Slightly abusing the notation in [126, Chapter 6], we define a *cocircuit* of the matrix B to be any vector $v \in \mathbb{R}^n$ of minimal support such that $vB = 0$.

Proposition 3.1.1. For any interior point $p \in \mathcal{M}$, the vertices of $\log \text{Vor}_{\mathcal{M}}(p)$ are of the form $v \cdot \text{diag}(p)$ where v is any positive cocircuit of B such that $\sum_{i=1}^n v_i p_i = 1$.

Proof. The log-likelihood function of a point $u \in \Delta_{n-1}$ is

$$\ell_u(x) = \sum_{i=1}^n u_i \log(c_i - b_i \cdot x)$$

where b_1, \dots, b_n are the rows of B . The likelihood equations [47, Chapter 2] have the form

$$\frac{u_1}{p_1} \cdot b_{1j} + \dots + \frac{u_n}{p_n} \cdot b_{nj} = 0 \quad \text{for each } j \in [d].$$

Since the log-likelihood function ℓ_u is strictly concave, it has a unique critical point on the model. Thus, the points in the logarithmic Voronoi polytope at p are the distributions u on which likelihood equations vanish. Equivalently, $\log \text{Vor}_{\mathcal{M}}(p)$ is the set of all distributions that satisfy the linear equations $(u_1/p_1, u_2/p_2, \dots, u_n/p_n) \cdot B = 0$. Hence, we may write:

$$\begin{aligned} \log \text{Vor}_{\mathcal{M}}(p) &= \{u \in \Delta_{n-1} : u \cdot \text{diag}(p)^{-1} B = 0\} \\ &= \left\{ r \cdot \text{diag}(p) \in \mathbb{R}^n : rB = 0, r \geq 0, \sum_{i=1}^n r_i p_i = 1 \right\}. \end{aligned}$$

Now consider an $n \times (d+1)$ matrix M obtained from B by adjoining $(p_1, \dots, p_n)^T$ as the first column. Then the logarithmic Voronoi polytope at p can be identified with the feasible region of a linear program, namely $r^T M = (1, 0, \dots, 0), r \geq 0$. From the simplex method [23, Chapter 3], we know that the vertices of such polytope are the basic feasible solutions, i.e. minimal support vectors in the region. Those basic solutions are precisely the positive cocircuits v of B for which $\sum_{i=1}^n v_i p_i = 1$. Since p is interior, $v \cdot \text{diag}(p)$ has the same support as v . Thus the vertices of the logarithmic Voronoi polytope at p are precisely the points $v \cdot \text{diag}(p)$ where v is a positive cocircuit of B for which $\sum_{i=1}^n v_i p_i = 1$. \square

Now we shall describe logarithmic Voronoi cells at interior points combinatorially. We use the formalism of Gale diagrams, as described in [126, Chapter 6]. For two polytopes P_1, P_2 , we will write $P_1 \sim P_2$ to mean that P_1 and P_2 are combinatorially equivalent. We will denote the polar dual of a polytope P by P^Δ .

Given our linear model $\mathcal{M} = \{c - Bx : x \in \Theta\}$, note that the configuration b_1, \dots, b_n of row vectors of B is totally cyclic, i.e. $\sum_{i=1}^n b_i = 0$, since each column of B sums to 0. Hence, B is a Gale transform of some affine configuration of n vectors $\{v_1, \dots, v_n\}$ in \mathbb{R}^{n-d-1} [126, Section 6.4]. Since a Gale transform uniquely determines the configuration up to an affine transformation, we may assume that $0 \in \text{conv}\{v_1, \dots, v_n\}$. Note that this configuration is not necessarily in convex position; however, its dual is a polytope. This polytope will have the same combinatorial type as the logarithmic Voronoi cells at interior points of \mathcal{M} , as shown in the next theorem.

Theorem 3.1.2. For any interior point p of the linear model \mathcal{M} , the logarithmic Voronoi polytope at p is combinatorially equivalent to the dual of the polytope obtained by taking the convex hull of a vector configuration with Gale diagram B .

Proof. As discussed above, let $\{v_1, \dots, v_n\}$ be a vector configuration whose Gale diagram is B . Let $P = \text{conv}\{v_1, \dots, v_n\}$ and assume $0 \in P$. We wish to show that $\log \text{Vor}_{\mathcal{M}}(p) \sim P^\Delta$.

Define

$$Q := \left\{ r \in \mathbb{R}^n : rB = 0, r \geq 0, \sum_{i=1}^n r_i = 1 \right\}.$$

Then $\log \text{Vor}_{\mathcal{M}}(p) \sim Q$, since multiplication by $\text{diag}(p)$ is an affine transformation and does not change the combinatorial type of $\log \text{Vor}_{\mathcal{M}}(p)$. It then suffices to show $Q \sim P^\Delta$. Let V be the matrix whose column vectors are v_1, \dots, v_n , and let

$$A = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ v_1 & v_2 & \cdots & v_n \end{bmatrix}.$$

We may assume that the rows of A are linearly independent. Since the Gale diagram of B^T is A^T , we know that $\ker B^T = \text{im } A^T$.

We will show $Q \sim P^\Delta$. Observe

$$\begin{aligned} Q = \left\{ r \in \mathbb{R}^n : rB = 0, r \geq 0, \sum_{i=1}^n r_i = 1 \right\} &= \left\{ xA \in \mathbb{R}^n : xA \geq 0, \sum_{i=1}^n (xA)_i = 1 \right\} \\ &\sim \left\{ x \in \mathbb{R}^{n-d} : xA \geq 0, \sum_{i=1}^n (xA)_i = 1 \right\}. \end{aligned}$$

The last equivalence follows from the fact that the rows of A^T are linearly independent. Therefore, the cone over Q is $C(Q) = \{x \in \mathbb{R}^{n-d} : xA \geq 0\}$.

For the dual of the polytope P , we may write:

$$\begin{aligned} P^\Delta &= \{z \in \mathbb{R}^{n-d-1} : V^T z \leq 1\} = \{z \in \mathbb{R}^{n-d-1} : (1, -z)A \geq 0\} \\ &\sim \{x \in \mathbb{R}^{n-d} : x_1 = 1, xA \geq 0\}. \end{aligned}$$

Hence the cone over P^Δ is also $\{x \in \mathbb{R}^{n-d} : xA \geq 0\}$. Note that this is a pointed cone at the origin, and both polytopes Q and P^Δ are obtained by intersecting this cone with a hyperplane that doesn't contain the origin. Hence, all the extreme rays are intersected by both hyperplanes. It follows that Q and P^Δ have the same combinatorial type by [126, Proposition 2.4]. Therefore, $\log \text{Vor}_{\mathcal{M}}(p)$ is indeed combinatorially equivalent to P^Δ . \square

The next theorem immediately follows from Theorem 3.1.2.

Theorem 3.1.3. The Logarithmic Voronoi polytopes at all interior points in a linear model have the same combinatorial type.

Example 3.1.4. The points $\{v_1, \dots, v_n\}$ in the previous statement need not be in convex position, but the dual of their configuration is. For example, consider a 1-dimensional linear

model inside the 3-simplex, given by $c = (1/4, 1/4, 1/4, 1/4)^T$ and $B = [1 \ -5 \ 2 \ 2]^T$. The parameter space is the interval $\Theta = [-1/20, 1/8] \subseteq \mathbb{R}$ and the model is parametrized

$$\Theta \rightarrow \Delta_3, \quad x \mapsto (-x + 1/4, 5x + 1/4, -2x + 1/4, -2x + 1/4).$$

The 4×1 matrix B is a Gale transform of the non-convex configuration $\{(-1, -1), (1, 1), (3, 0), (0, 3)\}$. Its convex hull is the triangle $\text{conv}\{(-1, -1), (3, 0), (0, 3)\}$, a self-dual polytope. The logarithmic Voronoi polytope at any interior point $p = c - Bx$ is also a triangle, with the vertices

$$\frac{1}{7}(0, 40x + 2, 0, -40x + 5), \frac{1}{7}(0, 40x + 2, -40x + 5, 0), \frac{1}{6}(-20x + 5, 20x + 1, 0, 0).$$

for the corresponding parameter $x \in (-1/20, 1/8)$. This is demonstrated in Figure 3.2.




Figure 3.2: Sampled points on the linear model corresponding to $B = [1 \ -5 \ 2 \ 2]$ and triangular logarithmic Voronoi cells.

If we take B to be the matrix $[1 \ 5 \ -3 \ -3]^T$, which is a Gale diagram of a convex 4-gon, the logarithmic Voronoi polytopes at the interior points on this model would be quadrilaterals in Δ_3 . So, (a dual of) any 2-dimensional convex polytope on 4 vertices is a logarithmic Voronoi polytope for some 1-dimensional model in Δ_3 . In fact, this holds in general.

Proposition 3.1.5. Every $(n - d - 1)$ -dimensional polytope with at most n facets appears as a logarithmic Voronoi polytope of a d -dimensional linear model inside Δ_{n-1} .

Proof. Let P be a polytope of dimension $n - d - 1$ with at most n facets. By Prop. 6.3 in [95], any polytope in \mathbb{R}^{n-d-1} with n facets is combinatorially equivalent to an intersection of Δ_{n-1} with an affine space of co-dimension d . The same is true for a $(n - d - 1)$ -dimensional

polytope with less than n facets, as we could first intersect Δ_{n-1} with an affine hyperplane to obtain Δ_{n-2} , and apply induction. That is, we may write our polytope P as

$$P = \left\{ x \in \mathbb{R}^n : Mx = b, \sum_{i=1}^n x_i = 1, x \geq 0 \right\}$$

where $M = (m_{ij}) \in \mathbb{R}^{d \times n}$ and $b \in \mathbb{R}^d$. Changing coordinates on \mathbb{R}^n , such that $x_n = 1$, we may re-write our polytope as $P = \{x \in \mathbb{R}^n : Nx = 0, x_n = 1, x \geq 0\}$. Here, N is the $d \times n$ matrix obtained from M by subtracting $(m_{in}, \dots, m_{in}, b_i)$ from the i th row. The cone of P is then $C(P) = \{x \in \mathbb{R}^n : Nx = 0, x \geq 0\}$. Let $x' = (x'_1, \dots, x'_n) \in C(P)$. Scaling the i th column of N by x'_i , we get a new matrix N' . Then the cone $C'(P) = \{x \in \mathbb{R}^n : N'x = 0, x \geq 0\}$ contains the all-ones vector. This guarantees that each row of N' sums to 0, and letting $B := (N')^T$, we see that the cone $C'(P)$ is equal to the cone of the logarithmic Voronoi polytope at an interior point of a model associated to B . As B is an $n \times d$ matrix, this model is d -dimensional in Δ_{n-1} . Thus, P is combinatorially equivalent to a logarithmic Voronoi polytope, as desired. \square

3.2 On the boundary

In this section we study logarithmic Voronoi polytopes at the points of a linear model that lie on the boundary of the simplex, where the log-likelihood function is undefined. The next example demonstrates that the combinatorial type of logarithmic Voronoi polytopes at the points on the boundary of Δ_{n-1} will depend on the positioning of the linear model inside the simplex. Namely, if the intersection of the affine linear space defining the model with Δ_{n-1} is not general, logarithmic Voronoi polytopes at the boundary points will degenerate.

The definition of the log-likelihood function can be extended to the boundary of the simplex by considering each boundary component of the model as a linear model inside a smaller simplex. Namely, let \mathcal{M} be a d -dimensional linear model inside Δ_{n-1} , and let f be a face of \mathcal{M} that lies on the boundary of Δ_{n-1} . Then the relative interior of f lies in the interior of some Δ_{k-1} , which is on the boundary of Δ_{n-1} . We may then treat f as its own linear model inside Δ_{k-1} , and the log-likelihood function is defined for all interior points of f .

Example 3.2.1. Consider a polytope, combinatorially isomorphic to the 3-dimensional cube. According to Proposition 3.1.5, this polytope appears as a logarithmic Voronoi cell at an interior point on some 2-dimensional linear model in Δ_5 . One such model \mathcal{M} is given by

$$B = \begin{bmatrix} -10 & -2 & -4 & 6 & 6 & 4 \\ 3 & 2 & 1 & -1 & -2 & -3 \end{bmatrix}^T \text{ and } c = (1/12, 1/3, 1/6, 1/12, 1/6, 1/6).$$

It is a triangle with the vertices $(1/168, -1/21)$, $(1/24, 1/6)$, and $(-1/24, -1/9)$. The logarithmic Voronoi polytopes at the interior points of \mathcal{M} are combinatorially equivalent to the 3-dimensional cube. The vertices are

$$\begin{pmatrix} 0, & 0, & 16x - 4y + 2/3, & -8x + 4/3y + 1/9, & -8x + 8/3y + 2/9, & 0 &) \\
 (0, & 0, & 84/5x - 21/5y + 7/10, & -72/5x + 12/5y + 1/5, & 0, & -12/5x + 9/5y + 1/10 &) \\
 (0, & 4/3x - 4/3y + 2/9, & 32/3x - 8/3y + 4/9, & 0, & -12x + 4y + 1/3, & 0 &) \\
 (0, & 4x - 4y + 2/3, & 2x - 1/2y + 1/12, & 0, & 0, & -6x + 9/2y + 1/4 &) \\
 (40/17x - 12/17y + 1/51, & 72/17x - 72/17y + 12/17, & 0, & 0, & 0, & -112/17x + 84/17y + 14/51 &) \\
 (30x - 9y + 1/4, & 0, & 0, & -6x + y + 1/12, & -24x + 8y + 2/3, & 0 &) \\
 (240/11x - 72/11y + 2/11, & 12/11x - 12/11y + 2/11, & 0, & 0, & -252/11x + 84/11y + 7/11, & 0 &) \\
 (35x - 21/2y + 7/24, & 0, & 0, & -27x + 9/2y + 3/8, & 0, & -8x + 6y + 1/3 &)
 \end{pmatrix}$$

for the parameters (x, y) . Given a point in \mathcal{M} on the boundary of Δ_5 , parametrized by (\hat{x}, \hat{y}) , the vertices of its logarithmic Voronoi polytope are obtained by plugging (\hat{x}, \hat{y}) into the equations above. One checks that at all the points on the boundary of \mathcal{M} , the logarithmic Voronoi polytopes are also combinatorially equivalent to the 3-dimensional cube.

On the other hand, consider the model given by

$$B = \begin{bmatrix} -10 & -3 & -20 & 6 & 6 & 21 \\ 1 & 3 & 2 & -1 & -3 & -2 \end{bmatrix}^T \text{ and } c = (1/6, 1/12, 1/3, 1/6, 1/12, 1/6).$$

It is a quadrilateral in Δ_5 with the vertices parametrized by $(1/153, -1/68)$, $(2/171, 3/76)$, $(-5/324, 1/81)$, and $(-7/288, -11/144)$. The logarithmic Voronoi polytopes at the interior points are also combinatorially equivalent to the 3-dimensional cube. However, at the vertex parametrized by $(-5/324, 1/81)$, the logarithmic Voronoi polytope is no longer a cube: it degenerates to a 2-dimensional quadrilateral. This is explained by the fact that the vertex lies on a 2-dimensional face of the simplex (as opposed to a 3-dimensional face).

In general, whenever each vertex of a d -dimensional linear model lies on a $(n - d - 1)$ -dimensional face of Δ_{n-1} , the combinatorial type of the logarithmic Voronoi cell at a boundary point is the same as at the interior points. Before proving this result, we first fix some notation.

Notation: Let $\mathcal{M} = \{c - Bx : x \in \Theta\}$ and let z be a cocircuit of B with support S such that $\sum_{i=1}^n z_i f_i(x) = 1$, where $f : \Theta \rightarrow \mathbb{R}^n$ is a parametrization of \mathcal{M} . Let $V_z(x) : \Theta \rightarrow \mathbb{R}^n$ be the vertex of the logarithmic Voronoi polytope determined by z , as a function of $x \in \Theta$. That is, $V_z(x) = (z_1(c_1 - \langle b_1, x \rangle), \dots, z_n(c_n - \langle b_n, x \rangle)) \in \Delta_{n-1}$. If $w = f(\hat{x}) \in \mathcal{M}$ is a point on the boundary of the simplex, then the vertices of the logarithmic Voronoi polytope at w are given as limits of the vertices $V_z(y^{(i)})$ where $\{f(y^{(i)})\}$ is a sequence of interior points converging to w . Let M be the $n \times (d+1)$ matrix obtained by concatenating c_i to the i th row of B , for all $i \in [n]$. If U and V are two sets of the same cardinality in $[n]$ and $[d+1]$, respectively, we denote by $M_{U,V}$ the submatrix of M , whose rows are indexed by U and whose columns are indexed by V . We define $B_{U,V}$ similarly. Assume, without loss of generality, that the last k columns of $M_{S,[d+1]}$ are linearly independent. We have the following technical lemma.

Lemma 3.2.2. Let $v = f(\hat{x})$ be a vertex of \mathcal{M} with support I and let z be a cocircuit of B . The i th coordinate of $V_z(\hat{x})$ is zero if and only if $\det M_{([n] \setminus I) \cup \{i\}, [d+1]} = 0$.

Proof. Since \mathcal{M} is d -dimensional, each vertex of the model is determined by the vanishing of precisely d coordinates, i.e. $c_i = \langle b_i, \hat{x} \rangle \forall i \in [n] \setminus I$ and $c_i > \langle b_i, \hat{x} \rangle \forall i \in I$. Without loss of generality, assume $I = \{d+1, \dots, n\}$, so v is determined by the vanishing of the first d coordinates. Then $\hat{x} = (\hat{x}_1, \dots, \hat{x}_d)$ is a solution to the linear system $B_{[d],[d]}x = (c_1, \dots, c_d)$. We may assume $\det B_{[d],[d]} \neq 0$. By Cramer's rule, we may then write $\hat{x}_i = (-1)^{d-i} \frac{\det M_{[d],[d+1] \setminus \{i\}}}{\det B_{[d],[d]}}$ for all $i \in [d]$. Let S denote the support of the cocircuit z and suppose $|S| = k$. Let z' be the projection of z onto its support. Since z is a cocircuit of B such that $\sum_{i=1}^n z_i(c_i - \langle b_i, x \rangle) = 1$, it satisfies the equation $c_1 z_1 + \dots + c_n z_n = 1$. Thus, z' is a solution to the system $y M_{S,[d+1]} = (0, \dots, 0, 1) \in \mathbb{R}^d$. If $d+1 \geq k$, then $d+1-k$ equations in this system must be redundant. From our assumption, the first $d+1-k$ equations are redundant, so removing them, we get a $k \times k$ linear system with a unique solution. Using Cramer's rule again, we find that for any $i \in S$, $z_i = (-1)^{k+i'} \frac{\det B_{S \setminus \{i\}, [d] \setminus [d+1-k]}}{\det M_{S,[d+1] \setminus [d+1-k]}}$, where i' is the index of z_i in z' . If $i \notin S$, the i th coordinate of $V_z(\hat{x})$ is 0. If $i \in S$, we have the i th coordinate of $V_z(\hat{x})$ is given by

$$\frac{\det B_{S \setminus \{i\}, [d] \setminus [d+1-k]}}{\det M_{S,[d+1] \setminus [d+1-k]} B_{[d],[d]}} \left[c_i \det B_{[d],[d]} - ((-1)^{d-1} b_{i1} \det M_{[d],[d+1] \setminus \{1\}} + \dots + b_{id} \det M_{[d],[d+1] \setminus \{d\}}) \right].$$

Note the expression in square brackets is $(-1)^{k+i'} \det M_{[d] \cup \{i\}, [d+1]}$, so the i th coordinate of $V_z(\hat{x})$ vanishes if and only if $\det M_{[d] \cup \{i\}, [d+1]} = 0$. The case $d+1 < k$ is not possible, as it would imply the existence of a cocircuit whose support is strictly contained in S . \square

Theorem 3.2.3. Let \mathcal{M} be a d -dimensional linear model obtained by intersecting the affine linear space L with Δ_{n-1} . Let $w \in \mathcal{M}$ be a point on the boundary of the simplex. If L intersects Δ_{n-1} transversally, then the logarithmic Voronoi polytope at w has the same combinatorial type as those at the interior points of \mathcal{M} .

Proof. It suffices to show that the combinatorial type of the logarithmic Voronoi polytopes at the vertices of the model is the same as at the interior points. Let $v = f(\hat{x})$ be a vertex of \mathcal{M} and without loss of generality assume that it has support $\{d+1, \dots, n\}$. By Lemma 3.2.2, if $i \in S \cap \{d+1, \dots, n\}$, the logarithmic Voronoi vertex $V_z(\hat{x})$ degenerates to the vertex with 0 in the i th coordinate if and only if $\det M_{[d] \cup \{i\}, [d+1]} = 0$. This condition translates to v lying on a face of Δ_{n-1} of dimension less than $n-d-1$, namely the one spanning the affine space $\{x \in \mathbb{R}^n : x_j = 0 \text{ for all } j \in [d] \cup \{i\}\}$. This means that the affine space L does not intersect Δ_{n-1} transversally, a contradiction. Thus, the logarithmic Voronoi polytope at any vertex of \mathcal{M} has the same combinatorial type as at the interior points. \square

The next example gives a concrete formula for the vertices of logarithmic Voronoi polytopes when the linear model is one-dimensional. The compact description follows from the fact that cocircuits are easy to compute in this case. A one-dimensional model will intersect the simplex transversally if and only if the $1 \times n$ matrix B has no repeated entries.

Example 3.2.4 ($d = 1$). Let $\mathcal{M} = \{c - Bx : x \in \Theta\}$ be a 1-dimensional linear model inside the simplex Δ_{n-1} . Let $B = [b_1, \dots, b_m, b_{m+1}, \dots, b_n]^T$, and without loss of generality assume

$b_i > 0$ for $i = 1, \dots, m$ and $b_i < 0$ for $i = m + 1, \dots, n$. Then $\Theta \subseteq \mathbb{R}$ is a closed interval $[x_\ell, x_r]$, where $x_\ell = c_\ell/b_\ell$ for some $\ell > m$ and $x_r = c_r/b_r$ for some $r \leq m$. Rotating the simplex, if necessary, we may ensure that $r = 1$. Note that any positive cocircuit z of B has support $\{i, j\}$ of size two, where $b_i > 0$ and $b_j < 0$. So, we find the logarithmic Voronoi polytope at x_r is the polytope at the boundary of Δ_{n-1} with the vertices

$$\{e_j : b_j < 0\} \cup \left\{ \frac{(c_i - b_i(c_1/b_1))b_j}{b_j c_i - b_i c_j} e_i - \frac{(c_j - b_j(c_1/b_1))b_i}{b_j c_i - b_i c_j} e_j : \begin{matrix} i \neq 1, \\ b_i > 0, \\ b_j < 0 \end{matrix} \right\}.$$

The logarithmic Voronoi polytope at x_ℓ is described similarly. Figure 3.1 plots logarithmic Voronoi polytopes at sampled points on 1-dimensional linear models in general position given by $c = (1/4, 1/4, 1/4, 1/4)$, $B = [1, -5, 3, 1]^T$ and $B = [-2, -1, 1, 2]^T$, respectively.

Example 3.2.5 (Moduli spaces). The moduli space $\mathcal{M}_{g,m}$ is the space of genus g curves with m marked points. The moduli space $\mathcal{M}_{0,m}$ is the space of m marked points in \mathbb{P}^1 and can be viewed as a linear statistical model of dimension $m - 3$ inside the simplex Δ_{n-1} , where $n = m(m - 3)/2$. The connection between particle physics and algebraic statistics via moduli spaces has been studied in [110]. The model $\mathcal{M}_{0,6}$ is a 3-dimensional linear model (a tetrahedron) inside the 8-dimensional simplex. It is parametrized by

$$(x, y, z) \mapsto \left(\frac{5x}{9}, \frac{y}{3}, \frac{z}{9}, \frac{y-x}{9}, \frac{z-x}{9}, \frac{y-z}{9}, \frac{1-x}{3}, \frac{1-y}{3}, \frac{1-z}{3} \right).$$

Logarithmic Voronoi polytopes at the interior points on this model are 5-dimensional with the f -vectors $(7, 19, 26, 19, 7)$.

The affine space defining this model does not intersect the simplex transversally; furthermore, none of the four vertices lie on the interior of a 5-dimensional face of Δ_8 . Two of the vertices lie on 4-dimensional faces of Δ_8 and the other two vertices lie on 2-dimensional faces of Δ_8 . The logarithmic Voronoi polytopes at these vertices degenerate into 4-dimensional and 2-dimensional polytopes, respectively. These polytopes are the entire faces of Δ_8 that contain the corresponding vertices in their relative interior.

3.3 Information divergence

In this section, we study the problem of maximizing information divergence, defined in Section 1.1, from a linear model \mathcal{M} . This problem of studying the maximum and the maximizers of the divergence function was first posed by Ay [16], but only for exponential families (toric models). This section investigates the same problem for linear models. Our motivation comes from potential applications in bio-neural networks, where we are interested in maximizing the mutual information between the input and output of each layer.

Let \mathcal{M} be a d -dimensional linear model in Δ_{n-1} given by an $n \times d$ matrix B and a vector $c \in \Delta_{n-1}$ as before. That is, \mathcal{M} is the image of the linear map

$$f : \Theta \rightarrow \Delta_{n-1} : (x_1, \dots, x_d) \mapsto (c_1 - \langle b_1, x \rangle, \dots, c_n - \langle b_n, x \rangle).$$

We want to find

$$D(\mathcal{M}) = \max_{p \in \Delta_{n-1}} D_{\mathcal{M}}(p)$$

and its maximizers, where $D_{\mathcal{M}}(p) = \min_{q \in \mathcal{M}} D(p||q)$ and $D(p||q) = \sum_{i=1}^n p_i \log(p_i/q_i)$.

Proposition 3.3.1. Let $\mathcal{M} \subset \Delta_{n-1}$ be a linear model and let $q \in \mathcal{M}$. Then the maximum of $D_{\mathcal{M}}(u)$ restricted to the logarithmic Voronoi polytope $\log \text{Vor}_{\mathcal{M}}(q)$ is achieved at a vertex of this polytope. The maximizers are a subset of the vertices in $\log \text{Vor}_{\mathcal{M}}(q)$.

Proof. Note that $D(u||q)$ is strictly convex in u over Δ_{n-1} ; see for instance [100, Proposition 2.14 (iii)]. The result follows since $\log \text{Vor}_{\mathcal{M}}(q)$ is a convex polytope. \square

We wish to find $D(\mathcal{M})$ and the points $p \in \Delta_{n-1}$ at which the information divergence $D_{\mathcal{M}}(p)$ from the linear model is maximized. By Proposition 3.3.1,

$$D(\mathcal{M}) = \max_{q \in \mathcal{M}} \max_{p \in \log \text{Vor}_{\mathcal{M}}(q)} D(p||q).$$

Hence, the maximum is achieved at some of the vertices of the logarithmic Voronoi cell $\log \text{Vor}_{\mathcal{M}}(q)$ at q . The vertices of $\log \text{Vor}_{\mathcal{M}}(q)$ at $q = f(x)$ are given by the co-circuits of B and can be expressed as functions in q (or the parameters x). Here, by a co-circuit of B we mean a nonzero $z \in \mathbb{R}^n$ of minimal support so that $z^T B = 0$. Each co-circuit z of B such that $\langle z, q \rangle = \sum_{i=1}^n z_i q_i = 1$ defines a vertex $V_z(q) = (z_1 q_1, \dots, z_n q_n)$ of $\log \text{Vor}_{\mathcal{M}}(q)$. Note that the choice of the co-circuit representative does not depend on the point q , i.e. we may always choose the representative z such that $\langle z, q \rangle = 1$ for all $q \in \mathcal{M}$ simultaneously. Indeed, let y be some co-circuit of B . We wish to find $k \in \mathbb{R}$ such that $z = ky$ has the property $\langle z, q \rangle = 1$ for all $q \in \mathcal{M}$. Since $q = c - Bx$ for some $x \in \Theta$, we have that

$$1 = \langle z, q \rangle = k \langle y, c - Bx \rangle = k \langle y, c \rangle.$$

Hence, $z = ky$ where $k = 1/\langle y, c \rangle$ is the desired co-circuit representative. For every such co-circuit we wish to maximize the information divergence over all $q \in \mathcal{M}$. We then compare the maximum divergences over all such co-circuits to find the global maximum.

Lemma 3.3.2. Let \mathcal{M} be a linear model defined by the matrix B and the vector c . For a fixed co-circuit z of B , the information divergence $D(V_z(q)||q)$ is linear in $q \in \mathcal{M}$.

Proof. Note that for a fixed co-circuit z of B , we get:

$$D(V_z(q)||q) = \sum_{i=1}^n (z_i q_i) \log(z_i q_i / q_i) = \sum_{i=1}^n (z_i \log(z_i)) q_i.$$

Therefore, $z_i \log(z_i)$ is a constant for each i , and so the divergence function is indeed linear in the coordinates of q . \square

Hence, for each co-circuit z , we are maximizing a linear function over the polytope \mathcal{M} . We summarize this in the following result.

Theorem 3.3.3. The maximum divergence of a linear model \mathcal{M} is always achieved at a vertex of the logarithmic Voronoi polytope $\log \text{Vor}_{\mathcal{M}}(q)$ where q itself is a vertex of \mathcal{M} .

Remark 3.3.4. A particular kind of discrete exponential family that is also a linear model is a partition model. The information divergence from partition models have been studied in [85]. A result similar to Theorem 3.3.3 is Proposition 2 in this reference.

Theorem 3.3.3 can be used to obtain compact formulas for maximum divergence for special families of linear models, such as the one below.

Corollary 3.3.5. Let \mathcal{M} be a one-dimensional linear model in Δ_3 given by the matrix $B = [-a, -b, b, a]^T$, $a, b > 0$ and $c = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Then $D(\mathcal{M}) = \log\left(\frac{4 \max\{a, b\}}{a+b}\right)$, maximized at two vertices of Δ_3 .

Proof. Without loss of generality, assume that $a > b$. Then the model is parametrized as $f : x \mapsto (ax + 1/4, bx + 1/4, -bx + 1/4, -ax + 1/4)$. The two vertices of the model are $v_1 = f(-\frac{1}{4a})$ and $v_2 = f(\frac{1}{4a})$. Each logarithmic Voronoi polytope is a quadrangle, so the matrix B has four co-circuits which parameterize the four vertices of this polytope at a general point $q = f(x)$:

$$\begin{aligned} V_1(x) &= (0, 2bx + 1/2, -2bx + 1/2, 0) \\ V_2(x) &= (0, (4abx + a)/(a + b), 0, (b - 4abx)/(a + b)) \\ V_3(x) &= (2ax + 1/2, 0, 0, -2ax + 1/2) \\ V_4(x) &= ((4abx + b)/(a + b), 0, (a - 4abx)/(a + b), 0). \end{aligned}$$

Note that $D(V_1(x)||f(x)) = D(V_3(x)||f(x)) = \log(2)$ for all $x \in [-\frac{1}{4a}, \frac{1}{4a}]$. On the other hand,

$$\begin{aligned} D\left(V_2\left(-\frac{1}{4a}\right) \| v_1\right) &= \frac{(a-b) \log\left(\frac{4a}{a+b}\right) + 2b \log\left(\frac{4b}{a+b}\right)}{a+b} < \log\left(\frac{4a}{a+b}\right) = D\left(V_2\left(\frac{1}{4a}\right) \| v_2\right) \\ D\left(V_4\left(-\frac{1}{4a}\right) \| v_1\right) &= \log\left(\frac{4a}{a+b}\right) > \frac{(a-b) \log\left(\frac{4a}{a+b}\right) + 2b \log\left(\frac{4b}{a+b}\right)}{a+b} = D\left(V_4\left(\frac{1}{4a}\right) \| v_2\right) \end{aligned}$$

so the maximum divergence $\log\left(\frac{4a}{a+b}\right)$ is achieved at the two vertices $V_2\left(\frac{1}{4a}\right) = (0, 1, 0, 0)$ and $V_4\left(-\frac{1}{4a}\right) = (0, 0, 1, 0)$. The proof for $b > a$ is identical. \square

Example 3.3.6. Consider the 1-dimensional linear model \mathcal{M} inside Δ_3 given by $B = [-2, -1, 1, 2]^T$ and $c = (1/4, 1/4, 1/4, 1/4)$. It is a line segment in Δ_3 with the vertices $v_1 = f(-1/8) = (0, 1/8, 3/8, 1/2)$ and $v_2 = f(1/8) = (1/2, 3/8, 1/8, 0)$. The global maximum divergence $\log(8/3)$ is achieved at $V_4(-1/8) = (0, 0, 1, 0)$ and $V_2(1/8) = (0, 1, 0, 0)$. This is illustrated in Figure 3.3.




Figure 3.3: Linear model given by $B = [-2, -1, 1, 2]^T$.

3.4 Partial linear models

A *partial* linear model of dimension d is a statistical model given by a d -dimensional polytope inside the probability simplex Δ_{n-1} , such that not all facets of the polytope lie on the boundary of the simplex.

Let \mathcal{M} be a partial linear model of dimension d inside Δ_{n-1} . The intersection of the affine span of the polytope \mathcal{M} with the simplex Δ_{n-1} is a d -dimensional linear model \mathcal{M}' . We say \mathcal{M}' *extends* \mathcal{M} . As before, $\mathcal{M}' = \{c - Bx : x \in \Theta'\}$ for some appropriate c, B , and parameter space $\Theta' \subseteq \mathbb{R}^d$. Since \mathcal{M}' extends \mathcal{M} , it follows that we may also write

$$\mathcal{M} = \{c - Bx : x \in \Theta\}$$

for some $\Theta \subseteq \Theta'$. Note that both Θ and Θ' are polytopes.

Theorem 3.4.1. Let \mathcal{M} be a partial linear model of dimension d with extension \mathcal{M}' . If p is a point in the relative interior of \mathcal{M} , then $\log \text{Vor}_{\mathcal{M}}(p) = \log \text{Vor}_{\mathcal{M}'}(p)$.

Proof. We show these sets are contained in each other. First, let $u \in \log \text{Vor}_{\mathcal{M}'}(p)$. Then $\ell_u(x)$ is maximized at p in \mathcal{M}' . Since $\mathcal{M} \subseteq \mathcal{M}'$, and $p \in \mathcal{M}$ as well, it follows that $\ell_u(x)$ will also be maximized at p in \mathcal{M} . Thus, $u \in \log \text{Vor}_{\mathcal{M}}(p)$.

Now, let $u \in \log \text{Vor}_{\mathcal{M}}(p)$. If $u \notin \log \text{Vor}_{\mathcal{M}'}(p)$, then over \mathcal{M}' , $\ell_u(x)$ is maximized at some other point $q \in \mathcal{M}' \setminus \mathcal{M}$. Then the line segment $[p, q]$ must intersect the boundary of the model \mathcal{M} . Note that any point on $[p, q]$ can be written as $a_x = (1 - x)p + xq$ for some $x \in [0, 1]$. Recall that the log-likelihood function ℓ_u is strictly concave on the simplex and hence on any convex subset of the simplex, such as our model \mathcal{M}' . So, for any $x \in (0, 1)$, we have

$$\ell_u(a_x) = \ell_u((1 - x)p + xq) > (1 - x)\ell_u(p) + x\ell_u(q) > \ell_u(p),$$

where the last inequality follows from the assumptions $\ell_u(q) > \ell_u(p)$ and $x > 0$. But since p is in the relative interior of the polytope \mathcal{M} , this implies that there is another interior point r on the line segment $[p, q]$ such that $\ell_u(r) > \ell_u(p)$. This is a contradiction to u 's inclusion in $\log \text{Vor}_{\mathcal{M}}(p)$. Therefore, $u \in \log \text{Vor}_{\mathcal{M}'}(p)$, as desired. \square

The theorem above tells us that the logarithmic Voronoi polytopes at points in the interior of the polytope \mathcal{M} are the same as those in the full linear extension \mathcal{M}' . The points $u \in \Delta_{n-1}$ that are not in $\log \text{Vor}_{\mathcal{M}}(p)$ for any p in the interior of \mathcal{M} will be mapped to the points on the boundary of \mathcal{M} via the MLE map. Note that for each point q on the boundary of \mathcal{M} , we still have $\log \text{Vor}_{\mathcal{M}'}(q) \subseteq \log \text{Vor}_{\mathcal{M}}(q)$. However, in general, this containment will be strict.

Given a facet F of \mathcal{M} , let p be a point in the relative interior of F (i.e. p does not lie on any lower-dimensional face). Treating F as its own partial linear model with extension F' inside Δ_{n-1} , we know that $\log \text{Vor}_F(p) = \log \text{Vor}_{F'}(p)$ is an $(n-d)$ -dimensional polytope. Moreover, it is clear that $\log \text{Vor}_{\mathcal{M}'}(p) \subseteq \log \text{Vor}_F(p)$. Observe that $\log \text{Vor}_{\mathcal{M}'}(p)$ has dimension $n-d-1$ and the boundary of this polytope is included in the boundary of $\log \text{Vor}_F(p)$, since these logarithmic Voronoi polytopes are the intersections of affine linear spaces with the simplex. Hence, $\log \text{Vor}_{\mathcal{M}'}(p)$ divides the polytope $\log \text{Vor}_F(p)$ into two $(n-d)$ -dimensional polytopes. Since p is on the boundary of the polytope \mathcal{M} , one of those polytopes will intersect the relative interior of \mathcal{M} .

Notation: Denote the two polytopes defined above by Q_p and \overline{Q}_p . Assume \overline{Q}_p is the polytope that intersects the relative interior of \mathcal{M} .

Theorem 3.4.2. Let p be a point in the relative interior of some facet F of \mathcal{M} . Let Q_p be as above. Then $Q_p = \log \text{Vor}_{\mathcal{M}}(p)$.

Proof. Let $u \in \log \text{Vor}_{\mathcal{M}}(p)$. Since $F \subseteq \mathcal{M}$ and p is in the interior of F , we have

$$u \in \log \text{Vor}_{\mathcal{M}}(p) \subseteq \log \text{Vor}_{F'}(p) = \log \text{Vor}_F(p) = Q_p \cup \overline{Q}_p.$$

If $u \notin Q_p$, then $u \in \overline{Q}_p \setminus Q_p$. In particular, $u \notin \log \text{Vor}_{\mathcal{M}'}(p)$. But since $\log \text{Vor}_{\mathcal{M}'}(p) \subseteq \log \text{Vor}_{\mathcal{M}}(p)$ and $\log \text{Vor}_{\mathcal{M}}(p)$ is convex, we also have $\text{conv}\{u, \log \text{Vor}_{\mathcal{M}'}(p)\} \subseteq \log \text{Vor}_{\mathcal{M}}(p)$. But then by construction of \overline{Q}_p , the convex hull above will contain an interior point in \mathcal{M} . This is a contradiction, as logarithmic Voronoi polytopes at distinct points on the model cannot intersect; thus $u \in Q_p$.




Figure 3.4: Polytopes Q_p and \bar{Q}_p for a point on a facet of a 2-dimensional model in Δ_3 .

For the other direction, note that the polytope Q_p has two types of points: the points in the polytope $\log \text{Vor}_{\mathcal{M}'}(p)$ and the points not in $\log \text{Vor}_{\mathcal{M}'}(p)$. Since $\log \text{Vor}_{\mathcal{M}'}(p) \subseteq \log \text{Vor}_{\mathcal{M}}(p)$, it suffices to show that $w \in \log \text{Vor}_{\mathcal{M}}(p)$ for each point $w \in Q_p$ of the second type. We show that $w \in \log \text{Vor}_{\mathcal{M}}(p)$. Note that we may assume that w is in the interior of Δ_{n-1} , since taking the closure would preserve the containment. Note that $\ell_w(x)$ is a strictly concave function on the simplex, so its super-level sets

$$C_\alpha = \{x \in \Delta_{n-1} : \ell_w(x) \geq \alpha\}$$

are convex $(n - 1)$ -dimensional sets. Since the maximum of $\ell_w(x)$ on F' is achieved at p , we know that it is given by $\ell_w(p) = \max\{\alpha : C_\alpha \cap F' \neq \emptyset\}$. Note that F' divides the linear extension \mathcal{M}' into two polytopes, S_1 and S_2 , where S_1 is the polytope containing the model \mathcal{M} . If $w \notin \log \text{Vor}_{\mathcal{M}}(p)$, then $\Phi_{\mathcal{M}}(w) = q \neq p$, where $q \notin F$. So, $q \in S_1$ lies on some other facet of \mathcal{M} . Moreover, $m \notin F'$, since p is the maximizer over F' and $\ell_w(m) > \ell_w(p)$.

Case 1: Suppose $m \in S_1$. Note that $R = \bigcup_{r \in F'} \log \text{Vor}_{\mathcal{M}'}(r)$ is an $(n - 2)$ -dimensional hypersurface inside Δ_{n-1} , obtained by intersecting a ruled hypersurface in \mathbb{R}^n with the simplex. Thus, R subdivides the simplex into two full-dimensional parts. By construction, w and m are on different sides of R . Since logarithmic Voronoi cells are convex sets, the line $[w, m] \subseteq \log \text{Vor}_{\mathcal{M}'}(m)$, and this line intersects R . This is a contradiction, since logarithmic Voronoi cells at two distinct points on the same model cannot intersect.

Case 2: If $m \in S_2$, then since $\ell_w(m) > \ell_w(p)$ and $\ell_w(q) > \ell_w(p)$, there exists some α such that $C_\alpha \subsetneq C_{\ell_w(p)}$ and such that C_α contains q and m , but does not contain p . Since super-level sets are convex, the line segment $[q, m]$ between q and m is contained in C_α . But since $q \in S_1$ and $m \in S_2$, the line $[q, m]$ intersects F' in some point $s \neq p$. But then $\ell_w(s) > \ell_w(p)$, a contradiction.

We conclude that $w \in \log \text{Vor}_{\mathcal{M}}(p)$. Since logarithmic Voronoi cells are closed sets, the closure of all such points w is also contained in $\log \text{Vor}_{\mathcal{M}}(p)$, and the conclusion follows. \square

Now suppose F is a face of \mathcal{M} of dimension $d - k$ for some $k \geq 2$. Then F is the intersection of at least k faces of dimension $d - k + 1$. Denote those faces by $\{G_1, \dots, G_m\}$, where $m \geq k$. For each $i \in [m]$, $\log \text{Vor}_{G'_i}(p)$ subdivides $\log \text{Vor}_{F'}(p)$ into two polytopes. Exactly one of these polytopes will intersect the face G_i at an interior point; call such polytope \overline{Q}_i . Call the other polytope Q_i . We present the following conjecture.

Conjecture 3.4.3. Let p be a point in the relative interior of the face F of \mathcal{M} . Then $\bigcap_{i \in [m]} Q_i = \log \text{Vor}_{\mathcal{M}}(p)$. In particular, if \mathcal{M} is in general position, $\dim \log \text{Vor}_{\mathcal{M}}(p) = (n - 1) - \dim F$.

Example 3.4.4. Let $d = 2$, $n = 4$, and consider the model \mathcal{M} defined as the convex hull of $(\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{2}{5})$, $(\frac{1}{5}, \frac{1}{5}, \frac{2}{5}, \frac{1}{5})$, and $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Below we plot the logarithmic Voronoi cells at interior points, edges, and vertices consecutively.




Figure 3.5: Logarithmic Voronoi cells at sampled points on the model defined as the convex hull of three points.

Conclusion. In this chapter, we proved that logarithmic Voronoi cells at the interior point of a linear model have the same combinatorial type. We also showed that the logarithmic Voronoi polytopes at the boundary points have the same combinatorial type as those at the interior points, as long as the linear model intersects the simplex transversally. Finally, we generalized this analysis to partial linear models, where the points on the boundary of the model are especially of interest. The next chapter studies how the combinatorial type of logarithmic Voronoi polytopes changes over a toric model.

Chapter 4

Toric models

In this chapter, we study the problem of maximizing information divergence to toric models. Recall that this means finding

$$D(\mathcal{M}) = \max_{p \in \Delta_{n-1}} D_{\mathcal{M}}(p)$$

and its maximizers, where \mathcal{M} is a toric model, $D_{\mathcal{M}}(p) = \min_{q \in \mathcal{M}} D(p||q)$, and $D(p||q) = \sum_{i=1}^n p_i \log(p_i/q_i)$. Recall that in Section 3.3, we paid close attention to the vertices of logarithmic Voronoi polytopes of linear models, and that having polytopes of just one combinatorial type was crucial in our analysis. This phenomenon will carry over to the toric case, except we will need to account for the fact that the logarithmic Voronoi polytopes have more than one (but finitely many) combinatorial types. In Section 4.1, we will review results that will be useful in locating vertices of logarithmic Voronoi polytopes of the same combinatorial type that potentially maximize the information divergence. Then in Section 4.2 we will see how to parameterize the different combinatorial types and how this helps develop an algorithm to compute $D(\mathcal{M})$. This chapter is based on [6].

Related prior work

The problem of determining $D(\mathcal{E})$ and studying the maximizers of the divergence function from an exponential family $\mathcal{E} \subset \Delta_{n-1}$ was first posed by Ay [16] who computed the gradient of $D_{\mathcal{E}}(p)$. The exponential family \mathcal{M} of probability distributions of independent random variables X_i , $i = 1, \dots, m$ with state spaces $[d_i] := \{1, \dots, d_i\}$ is known as an independence model. In this case, $D_{\mathcal{M}}(p)$ is the multi-information, and $D(\mathcal{M}) \leq \sum_{i=1}^{m-1} \log(d_i)$ where $2 \leq d_1 \leq d_2 \leq \dots \leq d_m$ [18]. In the same work, the structure of the global maximizers of the multi-information when the above bound is achieved was also determined. Subsequently, Matúš has computed the optimality conditions for $D_{\mathcal{E}}(p)$ for any exponential model $\mathcal{E} \subset \Delta_{n-1}$ [84]. We will use these conditions heavily. Rauh's dissertation [100] as well as his work in [101] gave algorithms to compute $D(\mathcal{M})$ for a discrete exponential family \mathcal{M} . These algorithms have two components: a combinatorial step followed by an algebraic step, both

of which can be challenging. Nevertheless, they were capable of computing the maximum multi-information to an independence model with $d_1 = 2$ and $d_2 = d_3 = 3$, the smallest case where the aforementioned bound is not attained. This chapter will provide another algorithm in the same spirit with combinatorial and algebraic steps. Finally, the literature contains results on the maximum divergence from certain hierarchical models [83], partition models [102], naive Bayes models and restricted Boltzmann machines [91].

Preliminaries

Let \mathfrak{X} be a finite set of cardinality n and let A be a $d \times n$ matrix with entries in \mathbb{R} . With respect to the reference measure $\omega(x)$, $x \in \mathfrak{X}$, the exponential family $\mathcal{E} = \mathcal{E}_{\omega, A}$ consists of the positive probability distributions in Δ_{n-1} of the form

$$P_{\theta}(x) = \frac{\omega(x)}{Z_{\theta}} \exp \left(\sum_{i=1}^d \theta_i A_{i,x} \right),$$

where A_i is the i th row of A and Z_{θ} is the normalizing constant. Here $\theta_i \in \mathbb{R}$ and $\bar{\mathcal{E}}$, the Euclidean closure of \mathcal{E} in Δ_{n-1} , will be referred to as the extended exponential family. Usually we will identify \mathfrak{X} with $[n]$ and write p_i and ω_i instead of $P(x)$ and $\omega(x)$, respectively.

In this chapter, we consider *discrete* exponential families because of the bridge to toric geometry and algebraic statistics [47, 114]. This means that A is a matrix with *integer* entries. Since without loss of generality we can assume that the row span of A

$$A = \begin{pmatrix} a_1 & a_2 & \cdots & a_n \end{pmatrix}$$

contains $(1, 1, \dots, 1)$, we will take the columns $a_j \in \mathbb{N}^d$, $j = 1, \dots, n$ and fix the first row of A to be the row of all ones. The toric variety $X_{\omega, A}$ is the Zariski closure in \mathbb{C}^n of the image of the algebraic torus $(\mathbb{C}^*)^d$ under the monomial map $\Psi : (\mathbb{C}^*)^d \rightarrow \mathbb{C}^n$ given by

$$z = (z_1, \dots, z_d) \mapsto (\omega_1 z^{a_1}, \omega_2 z^{a_2}, \dots, \omega_n z^{a_n}).$$

Because of the assumption on the first row of A , we can also view $X_{\omega, A}$ as a toric variety in the projective space \mathbb{P}^{n-1} . The following theorem connects exponential families and toric varieties.

Theorem 4.0.1. [55, Theorem 3.2] The extended exponential family $\bar{\mathcal{E}}_{\omega, A}$ equals $X_{\omega, A} \cap \Delta_{n-1}$.

Therefore, we will refer to discrete exponential families as toric models. We will denote them by $\mathcal{M}_{\omega, A}$ or just \mathcal{M}_A .

Given a toric model \mathcal{M}_A and a fixed $u \in \Delta_{n-1}$, we know from Proposition 1.1.11 that the minimum $D_{\mathcal{M}_A}(u)$ is attained at a unique point $q \in \mathcal{M}_A$, which the maximum likelihood

estimate of u . Birch's Theorem (see [47, Proposition 2.1.5], [78, Theorem 4.8], [94, Theorem 1.10]) states that the maximum likelihood estimate of u is equal to the unique point in the intersection

$$\mathcal{M}_A \cap \{p \in \Delta_{n-1} : Au = Ap\}.$$

The second term in this intersection is the polytope $Q_u := \{p \in \Delta_{n-1} : Au = Ap\}$. If $q \in \mathcal{M}_A$ is the MLE of u , by Birch's Theorem $Q_q = Q_u$. The polytope Q_q is precisely the logarithmic Voronoi polytope at q , which we previously denoted by $\log \text{Vor}_{\mathcal{M}}(q)$. Since we will use logarithmic Voronoi polytopes often in this chapter, we will use the less cumbersome notation Q_q moving forward. The analogue of Proposition 3.3.1 holds in the toric case.

Proposition 4.0.2. Let $\mathcal{M} \subset \Delta_{n-1}$ be a toric model and let $q \in \mathcal{M}$. Then the maximum of $D_{\mathcal{M}}(u)$ restricted to the logarithmic Voronoi polytope Q_q is achieved at a vertex of Q_q . The maximizers are a subset of the vertices in Q_q .

Proof. It is again true that $D(u \parallel q)$ is strictly convex in u over Δ_{n-1} . The result follows since Q_q is a convex polytope. \square

Corollary 4.0.3. [16, Proposition 3.2] Let $\mathcal{M}_A \subseteq \Delta_{n-1}$ be a toric model where $A \in \mathbb{N}^{d \times n}$ and $\text{rank}(A) = d$. If p is a maximizer of the information divergence then $|\text{supp}(p)| \leq d = \dim(\mathcal{M}_A) + 1$.

Proof. If $q \in \mathcal{M}_A$ is the MLE of p , then p is a vertex of $Q_q = \{u \in \Delta_{n-1} : Au = Aq\}$. Any vertex of Q_q is a basic feasible solution to the system $Au = Aq$. In other words, it is of the form $p = (p_B, p_N)$ where $p_N = 0$ and $Bp_B = Aq$ with B a $d \times d$ invertible submatrix of A . This shows $|\text{supp}(p)| \leq d$. \square

4.1 Critical points of information divergence

For a face F of a given polytope Q , we define the support of F as the union of the supports of the vertices on F and denote it by $\text{supp}(F)$. This section starts with a definition that will pave the path for characterizing the critical points of the function $D_{\mathcal{M}}(\cdot)$.

Definition 4.1.1. Let Q_q be a logarithmic Voronoi polytope at a point on a toric model $\mathcal{M}_A \subset \Delta_{n-1}$. A vertex v of Q_q is *complementary* if there exists a face F of Q_q such that $\text{supp}(F) = [n] \setminus \text{supp}(v)$. We call F the complementary face of v .

Definition 4.1.2. Let $\mathcal{M}_A \subset \Delta_{n-1}$ be a toric model and let p be a point in Δ_{n-1} whose MLE is q with $\text{supp}(q) = [n]$. We say that p is a projection point if

$$p_i = \begin{cases} \frac{q_i}{\sum_{j \in \text{supp}(p)} q_j} & \text{if } i \in \text{supp}(p) \\ 0 & \text{otherwise.} \end{cases}$$

Remark 4.1.3. We can relax the condition for the full support of the MLE in the above definition. In this case, we need to consider MLEs that are in the extended exponential family, namely, those that are on \mathcal{M}_A and on a proper face Γ of Δ_{n-1} . However, these can be separately treated by focusing on the toric model $\mathcal{M}_{A_\Gamma} \subset \Gamma$ where A_Γ consists of the columns a_i of A with $i \in \text{supp}(\Gamma)$.

Theorem 4.1.4. If p is a local maximizer of $D_{\mathcal{M}_A}$ then p is a projection point. Moreover, every such projection point is a complementary vertex of Q_q where q is the MLE of p . A complementary vertex v of Q_q with the complementary face F is a projection point if and only if the line passing through v and q intersects the relative interior of F .

Proof. The first statement is proved in [84, Theorem 5.1]. Since p is a local maximizer it needs to be a vertex of Q_q . The point \tilde{p} defined by

$$\tilde{p}_i = \begin{cases} \frac{q_i}{\sum_{j \notin \text{supp}(p)} q_j} & \text{if } i \notin \text{supp}(p) \\ 0 & \text{otherwise} \end{cases}$$

is obtained by $\tilde{p} = p + \frac{1}{\sum_{j \notin \text{supp}(p)} q_j} [q - p]$ where $[q - p]$ is a vector parallel to the line through p and q . The support of \tilde{p} is precisely $[n] \setminus \text{supp}(p)$ and therefore it is contained in the interior of a face F with identical support. Hence p is a complementary vertex and the last statement follows. \square

Example 4.1.5. The binomial model of size 3 is precisely the twisted cubic curve from Example 1.1.12. It is the set of probability distributions on $\mathfrak{X} = \{0, 1, 2, 3\}$ parametrized as

$$q_j = \binom{3}{j} \theta^j (1 - \theta)^{3-j}, \quad j = 0, 1, 2, 3.$$

This is a one-dimensional toric model that describes the experiment of flipping a coin with the bias θ three times. The matrix A can be taken to be

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 \end{pmatrix}.$$

For $u = (u_0, u_1, u_2, u_3) \in \Delta_3$ the MLE is given by

$$\begin{aligned} q_0 &= \frac{1}{27} (3u_0 + 2u_1 + u_2)^3, \\ q_1 &= \frac{1}{9} (u_1 + 2u_2 + 3u_3) (3u_0 + 2u_1 + u_2)^2, \\ q_2 &= \frac{1}{9} (u_1 + 2u_2 + 3u_3)^2 (3u_0 + 2u_1 + u_2), \\ q_3 &= \frac{1}{27} (u_1 + 2u_2 + 3u_3)^3. \end{aligned}$$

The logarithmic Voronoi polytopes are of the form $Q_b = \{u \in \Delta_3 : u_1 + 2u_2 + 3u_3 = b\}$ where $0 < b < 3$. For $0 < b < 1$ and $2 < b < 3$ these polytopes are triangles. The first kind has vertices with supports $\{0, 1\}$, $\{0, 2\}$, and $\{0, 3\}$. The vertices of the second kind have supports $\{0, 3\}$, $\{1, 3\}$, and $\{2, 3\}$. None of these triangles have a complementary vertex. When $b = 1$ and $b = 2$, Q_b is still a triangle: the supports of the vertices of Q_1 are $\{1\}$, $\{0, 2\}$, and $\{0, 3\}$. Those of Q_2 are $\{0, 3\}$, $\{1, 3\}$, and $\{2\}$. In Q_1 , the vertex $(0, 1, 0, 0)$ is a projection point with divergence $\log \frac{9}{4}$. In Q_2 , the vertex $(0, 0, 1, 0)$ is a projection point with the same divergence. The logarithmic Voronoi polytopes for $1 < b < 2$ are quadrangles with vertex supports $\{0, 2\}$, $\{0, 3\}$, $\{1, 2\}$, and $\{1, 3\}$. Therefore each vertex is a complementary vertex where the corresponding complementary face F is a vertex itself. Among all these, we find projection points only when $b = 3/2$. The vertices of $Q_{\frac{3}{2}}$ are $(\frac{1}{4}, 0, \frac{3}{4}, 0)$, $(\frac{1}{2}, 0, 0, \frac{1}{2})$, $(0, \frac{1}{2}, \frac{1}{2}, 0)$, and $(0, \frac{3}{4}, 0, \frac{1}{4})$. All are projection points with the MLE $q = (\frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8})$ which is the intersection of the diagonals of the quadrangle. The divergences from each vertex to this binomial model are $\log(2)$, $2\log(2)$, $2\log(2) - \log(3)$, and $\log(2)$, respectively. Therefore $(\frac{1}{2}, 0, 0, \frac{1}{2})$ is the unique global maximizer attaining $D(\mathcal{M}) = 2\log(2)$.

Corollary 4.1.6. [101, Section VI] Let \mathcal{M}_A be a codimension one toric model in Δ_{n-1} , i.e., let $\text{rank}(A) = d = n - 1$. Then there are exactly two projection points and at most two global maximizers of $D_{\mathcal{M}_A}$.

Proof. The toric variety X_A is defined by a single equation which we can assume is of the form $x_1^{u_1} x_2^{u_2} \cdots x_r^{u_r} - x_{r+1}^{u_{r+1}} \cdots x_n^{u_n}$ where $\sum_{i=1}^r u_i = \sum_{j=r+1}^n u_j$. The one-dimensional $\ker(A)$ is spanned by $(u_1, \dots, u_r, -u_{r+1}, \dots, -u_n)$, and all logarithmic Voronoi polytopes are one-dimensional whose affine span is parallel to $\ker(A)$. Since each such polytope has exactly two vertices, the line through these vertices always intersects \mathcal{M}_A . Hence, for these vertices to be projection points, we only need to make sure that they have complementary support. This can only happen if the vertices are $p = (p_1, \dots, p_r, 0, \dots, 0)$ and $\tilde{p} = (0, \dots, 0, \tilde{p}_{r+1}, \dots, \tilde{p}_n)$ where $p_i = \frac{u_i}{\sum_{i=1}^r u_i}$ for $i = 1, \dots, r$ and $\tilde{p}_j = \frac{u_j}{\sum_{j=r+1}^n u_j}$ for $j = r+1, \dots, n$. Both points are projection points and either one or both of them are global maximizers of $D_{\mathcal{M}_A}$. \square

Example 4.1.7. Let X and Y be two independent binary random variables. The set of joint probability distributions $q_{ij} = \mathbb{P}(X = i, Y = j)$ with $i, j \in \{0, 1\}^2$ is parametrized by $q_{ij} = a_i b_j$. This toric model $\mathcal{M}_A \subset \Delta_3$ has codimension one and can be given by the matrix

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

The kernel of A is generated by $(1, -1, -1, 1)$, and the only two projection points are $(\frac{1}{2}, 0, 0, \frac{1}{2})$ and $(0, \frac{1}{2}, \frac{1}{2}, 0)$ with the MLE $q = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Since the information divergence from both projection points is $\log(2)$ they are both global maximizers.

We finish this section with a result that will be useful later.

Theorem 4.1.8. [90, Lemma 3.2] Let A_1, A_2, \dots, A_k be $d_i \times n_i$ matrices, $i = 1, \dots, k$ with nonnegative entries and with the corresponding all ones vector as their first row. Let

$$A = \begin{pmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_k \end{pmatrix}.$$

Then $D(\mathcal{M}_A) = \max\{D(\mathcal{M}_{A_1}), \dots, D(\mathcal{M}_{A_k})\}$.

Proof. Let $n = \sum_{i=1}^k n_i$ and $d = \sum_{i=1}^k d_i$. The toric variety X_A as an affine variety is $X_{A_1} \times \cdots \times X_{A_k}$ and the defining toric ideal is $I_A = I_{A_1} + \cdots + I_{A_k}$. Without loss of generality we can assume that $D(\mathcal{M}_{A_1})$ attains the maximum among the maximum information divergences for $\mathcal{M}_{A_1}, \dots, \mathcal{M}_{A_k}$. Let $p^{(1)} \in \Delta_{n_1-1}$ be a global maximizer with the associated MLE $q^{(1)}$. Setting $p = (p^{(1)}, 0, \dots, 0)$ and $q = (q^{(1)}, 0, \dots, 0)$, we get $p \in \Delta_{n-1}$ and $q \in \mathcal{M}_A = \Delta_{n-1} \cap X_A$. Since $D_{\mathcal{M}_A}(p||q) = D_{\mathcal{M}_{A_1}}(p^{(1)}||q^{(1)})$ we conclude that $D_{\mathcal{M}_A} \geq D_{\mathcal{M}_{A_1}}$. Conversely, let $p = (p^{(1)}, \dots, p^{(k)})$ be a global maximizer of $D_{\mathcal{M}_A}$ with the MLE $q = (q^{(1)}, \dots, q^{(k)})$. Set $p_+^{(i)} = \sum_{j=1}^{n_i} p_j^{(i)}$ and $q_+^{(i)} = \sum_{j=1}^{n_i} q_j^{(i)}$. Note that $A_i p^{(i)} = A_i q^{(i)}$, so $p_+^{(i)} = q_+^{(i)}$, and $q^{(i)} \in X_{A_i}$. Moreover $\sum_{i=1}^k p_+^{(i)} = 1$. Now let $\tilde{p}^{(i)} = \frac{1}{p_+^{(i)}} p^{(i)}$ and $\tilde{q}^{(i)} = \frac{1}{q_+^{(i)}} q^{(i)}$. We see that $\tilde{p}^{(i)} \in \Delta_{n_i-1}$, and $\tilde{q}^{(i)} \in \mathcal{M}_{A_i}$ is the MLE of $\tilde{p}^{(i)}$. Since $D(\tilde{p}^{(i)}||\tilde{q}^{(i)}) = \frac{1}{p_+^{(i)}} D(p^{(i)}||q^{(i)})$ we conclude that

$$D(\mathcal{M}_A) = D(p||q) = \sum_{i=1}^k D(p^{(i)}||q^{(i)}) = \sum_{i=1}^k p_+^{(i)} D(\tilde{p}^{(i)}||\tilde{q}^{(i)}) \leq \max\{D(\mathcal{M}_{A_1}), \dots, D(\mathcal{M}_{A_k})\},$$

as desired. □

4.2 The chamber complex and the algorithm

We devote this section to describing an algorithm to compute $D(\mathcal{M}_A)$ and the corresponding global maximizers for a toric model \mathcal{M}_A . We first introduce the chamber complex of A : a polytopal complex \mathcal{C}_A that is supported on the convex hull of (the columns of) A . This combinatorial object parametrizes all logarithmic Voronoi polytopes for the model \mathcal{M}_A . In particular, the finitely many faces (chambers) of \mathcal{C}_A correspond to all possible combinatorial types of these logarithmic Voronoi polytopes. It appears that, in order to locate all global maximizers of $D_{\mathcal{M}_A}$, one needs to examine the vertices of all logarithmic Voronoi polytopes. With the help of \mathcal{C}_A we will reduce this task to examining vertices of each combinatorial type where we essentially do an algebraic computation for each chamber in \mathcal{C}_A . For any omitted details in the definition and computation of \mathcal{C}_A as well as its properties we refer to [38, Chapter 5].

Recall that A is a $d \times n$ matrix with nonnegative integer entries and $\text{rank}(A) = d$. We also assume that the first row of A is the vector of all ones. This means that the convex hull of the columns of A , $\text{conv}(A)$, is a polytope of dimension $d - 1$ whose set of vertices is a subset of the columns of A . For a nonempty $\sigma \subset [n]$ we let $A_\sigma = \{a_i : i \in \sigma\}$. When $|\sigma| = d$ and A_σ is invertible, $\text{conv}(A_\sigma)$ is a $(d - 1)$ -dimensional simplex. We will also use σ to denote $\text{conv}(A_\sigma)$. By Carathéodory's theorem [126, Proposition 1.15], $\text{conv}(A)$ is the union of all such simplices.

Definition 4.2.1. For $b \in \text{conv}(A)$ let $C_b := \bigcap_{\sigma \ni b} \sigma$. The chamber complex of A is

$$\mathcal{C}_A := \{C_b : b \in \text{conv}(A)\}.$$

We note that \mathcal{C}_A is a polytopal complex supported on $\text{conv}(A)$, and each C_b is a face of \mathcal{C}_A . Each such face of \mathcal{C}_A is called a chamber. For every $b \in \text{conv}(A)$ the set $Q_b = \{p \in \Delta_{n-1} : Ap = b\}$ is a logarithmic Voronoi polytope. The polytope Q_b has the maximum dimension $n - d - 1$ if and only if b is in the relative interior of $\text{conv}(A)$.

Theorem 4.2.2. Let C be a chamber of the chamber complex \mathcal{C}_A . Then for each b in the relative interior of C , the vertices of Q_b are in bijection with $\sigma \subset [n]$ such that C is contained in the relative interior of $\text{conv}(A_\sigma)$ where the columns of A_σ are linearly independent. The support of the vertex corresponding to such σ is precisely σ . More generally, each face F of Q_b is of the form $F = Q_b \cap \bigcap_{i \notin \text{supp}(F)} \{p_i = 0\}$. As b varies in the relative interior of C , the support of each face of Q_b as well as the combinatorial type of Q_b does not change.

Proof. The polytope Q_b is a polyhedron in standard form. Hence, $v \in \Delta_{n-1}$ is a vertex of Q_b if and only if $Av = b$ where there exists $\sigma \subset [n]$ such that the columns of A_σ are linearly independent, and $i \notin \sigma$ implies $v_i = 0$; see [23, Theorem 2.4]. This is equivalent to $C \subset \text{conv}(A_\sigma)$. The extra condition that C is contained in the relative interior of $\text{conv}(A_\sigma)$ is equivalent to $\text{supp}(v) = \sigma$. More generally, each face F of Q_b is defined by some subset of coordinate hyperplanes $p_i = 0$. Since $\text{supp}(F)$ is the union of the supports of all the vertices on F we conclude that $F = Q_b \cap \bigcap_{i \notin \text{supp}(F)} \{p_i = 0\}$. By the first part of this theorem, as b varies in the relative interior of C , the support of each vertex does not change, and hence the support of each face does not change. Since each face is determined by the set of vertices contained in that face this implies that the face lattice of Q_b is constant, i.e. every Q_b has the same combinatorial type. \square

Example 4.2.3. Let

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 & 2 \\ 1 & 0 & 0 & 1 & 2 \end{pmatrix},$$

where we denote the columns of A by a, b, c, d , and e . Here $\text{conv}(A)$ is a pentagon which, together with its chamber complex \mathcal{C}_A , can be seen in Figure 4.1. This chamber complex consists of 10 vertices, 20 edges, and 11 two-dimensional chambers. For some chambers C we have depicted the logarithmic Voronoi polytopes Q_b where b is in the relative interior of C . For instance, the horizontal (red) edge of the pentagonal chamber supports logarithmic Voronoi polytopes that are quadrangles. The supports of their vertices are $\{a, d\}, \{a, c, e\}, \{b, c, e\}$, and $\{b, d, e\}$ because C is contained in the relative interiors of $\text{conv}(A_{a,d}), \text{conv}(A_{a,c,e}), \text{conv}(A_{b,c,e})$, and $\text{conv}(A_{b,d,e})$.




Figure 4.1: The chamber complex of a pentagon.

Remark 4.2.4. Although each chamber of \mathcal{C}_A gives rise to logarithmic Voronoi polytopes of \mathcal{M}_A that have the same combinatorial type, different chambers might yield identical combinatorial types. For instance, in Example 4.2.3 we see that there are multiple chambers

that support logarithmic Voronoi polytopes that are triangles or quadrangles. In fact, \mathcal{C}_A parametrizes these polytopes according to a finer invariant, namely, the normal fan of each polytope. We will not directly need this finer differentiation, though we will use the fact from Theorem 4.2.2 that the supports of the faces of Q_b given by b in a fixed chamber are constant.

Remark 4.2.5. In the algorithm we present, first we have to compute the chamber complex \mathcal{C}_A . Using Definition 4.2.1 for this computation is highly inefficient. Here is an outline for a more efficient way. First, one computes a Gale transform B of A where B is a $(n-d) \times n$ matrix whose rows form a basis for the kernel of A . Then the secondary fan Π_B of B is computed. This is a complete fan in \mathbb{R}^n in which each cone consists of weight vectors that induce the same regular subdivision of the vector configuration given by the n columns of B . The cones of the secondary fan Π_B are in bijection with the chambers of \mathcal{C}_A . More precisely, if u_1, \dots, u_k are the generators of a cone in Π_B the corresponding chamber in \mathcal{C}_A is the convex hull of Au_1, \dots, Au_k . The details can be found in [38, Section 5.4]; in particular, for the claimed bijection see Theorem 5.4.5 in the same reference. We used *Gfan* [70] to compute Π_B which can also be accessed via *Macaulay 2* [60].

Example 4.2.6. As the matrix A gets larger, all of these computations become challenging. To give an idea, we consider the toric model \mathcal{M}_A that is the independence model of a binary and two ternary random variables. It is a 5-dimensional model in Δ_{17} . The f -vector of the 5-dimensional polytope $\text{conv}(A)$ is $(18, 45, 48, 27, 8)$, i.e., this polytope has 18 vertices, 45 edges, etc. The chamber complex \mathcal{C}_A that was computed via the methods outlined in Remark 4.2.5 has the f -vector

$$(3503407, 33084756, 105341820, 151227738, 100828884, 25361616).$$

The computation took about two days on a standard laptop, and it could only be done after taking into account the symmetries of $\text{conv}(A)$. We note that, luckily, this computation needs to be done only once, and once \mathcal{C}_A is computed, its chambers have to be processed as we will explain in our algorithm. This processing can be shortened by considering the symmetries of the chamber complex (if there are any) as well as by using a few simple observations on the structure of the supports of the vertices of the logarithmic Voronoi polytopes. We will outline these ideas below.

According to Theorem 4.1.4, given a logarithmic Voronoi polytope Q_b where $b \in \text{conv}(A)$, we need to identify complementary vertices of Q_b and decide whether any of these vertices are projection points. These, in turn, are potential local and global maximizers of $D_{\mathcal{M}_A}$. The following proposition gives a way to decide whether a complementary vertex is a projection point.

Proposition 4.2.7. Let v be a complementary vertex of the logarithmic Voronoi polytope Q_b of a toric model \mathcal{M}_A with the complementary face F . Let $\mathcal{L}_{v,F}$ be the collection of the

lines passing through v and each point on F . Then v is a projection point if and only if $\mathcal{L}_{v,F}$ intersects \mathcal{M}_A .

Proof. By Birch's theorem, Q_b intersects \mathcal{M}_A in a single point, namely, the MLE q of any point p in Q_b . The vertex v is a projection point if and only if one of the lines in $\mathcal{L}_{v,F}$ passes through q . This happens if and only if $\mathcal{L}_{v,F}$ intersects \mathcal{M}_A in the only possible point q . \square

In light of Proposition 4.2.7, to check whether a complementary vertex v of Q_b is a projection point reduces to an algebraic computation. Let F be the complementary face of dimension k . Then $\overline{\mathcal{L}_{v,F}}$, the Zariski closure of $\mathcal{L}_{v,F}$, is an affine subspace of dimension $k + 1$ whose defining equations can easily be computed. For instance, if v_1, \dots, v_{k+1} are vertices of F that are affinely independent, then $\overline{\mathcal{L}_{v,F}}$ is the image of the map

$$(s, t_1, \dots, t_{k+1}) \mapsto sv + (1 - s)(t_1v_1 + \dots + t_{k+1}v_{k+1})$$

where $t_1 + \dots + t_{k+1} = 1$. To intersect $\overline{\mathcal{L}_{v,F}}$ with \mathcal{M}_A we use the equations of $\overline{\mathcal{L}_{v,F}}$ and the binomial equations defining the toric variety X_A . Since $\overline{\mathcal{L}_{v,F}}$ is contained in the affine span of Q_b , and since the latter affine subspace intersects X_A in finitely many complex points (see Definition 4.4.1), $\overline{\mathcal{L}_{v,F}}$ intersects X_A also in finitely many points. They can be computed using a numerical algebraic geometry software such as *Bertini* [19] or *HomotopyContinuation.jl* [27]. Finally, one checks whether this finite set contains a point with positive coordinates.

Example 4.2.8. We use Example 4.2.3. The point $b = (1, 7/4, 1)$ is the midpoint of the horizontal (red) edge of the pentagonal chamber. The vertex $v = (5/12, 0, 0, 7/12, 0)$ of Q_b is complementary to another vertex $v_1 = (0, 1/4, 1/4, 0, 1/2)$. The toric variety X_A is defined by the equations

$$p_2^2 p_4^2 - p_3^3 p_5 = p_1 p_3^3 - p_2^3 p_4 = p_1 p_4 - p_2 p_5 = 0.$$

The affine subspace spanned by v and v_1 is just a line defined by

$$12p_4 + 14p_5 - 7 = 2p_3 - p_5 = 2p_2 - p_5 = 12p_1 + 10p_5 - 5 = 0.$$

The intersection of X_A with $\overline{\mathcal{L}_{v,\{v_1\}}}$ is empty. Hence, we conclude that v is not a projection point.

The above discussion describes a way of checking whether a complementary vertex of a *fixed* logarithmic Voronoi polytope Q_b is a projection point. Next, we describe how to accomplish the same task for a complementary vertex of Q_b as b varies in the interior of a fixed chamber C in the chamber complex \mathcal{C}_A . By Theorem 4.2.2 each such Q_b has the same combinatorial type and the support of any face of Q_b stays constant. Now let $(v(b), F(b))$ be a pair of a complementary vertex and its corresponding complementary face in Q_b where b is in the relative interior of a chamber C . Let w_1, \dots, w_m be the vertices of C . Then $b = \sum_{i=1}^m r_i w_i$ where $\sum_{i=1}^m r_i = 1$ and $r_i \geq 0$ for all $i = 1, \dots, m$. This means that the coordinates of $v(b)$

and those of the vertices $v_1(b), \dots, v_z(b)$ of $F(b)$ are linear functions of r_1, \dots, r_m . Next, we parametrize a general point $w(b)$ on $F(b)$ via $w(b) = \sum_{i=1}^z t_i v_i(b)$ where $\sum_{i=1}^z t_i = 1$ and $t_i \geq 0$ for all $i = 1, \dots, z$. Finally, the line segment between $v(b)$ and $w(b)$ is parametrized by $sv(b) + (1-s)w(b)$ where $0 \leq s \leq 1$. The last expression gives points in Δ_{n-1} where each coordinate is a polynomial in the parameters $r_1, \dots, r_m, t_1, \dots, t_z$, and s , and it defines the map

$$\Psi_{v,F} : \Delta_{m-1} \times \Delta_{z-1} \times \Delta_1 \longrightarrow \Delta_{n-1}.$$

Proposition 4.2.7 implies that $v(b)$ is a projection point for some $b \in C$ if and only if the image of $\Psi_{v,F}$ intersects \mathcal{M}_A . Again, this boils down to an algebraic computation. We substitute the coordinates of $sv(b) + (1-s)w(b)$ into the equations defining X_A , check whether this system of equations has solutions in \mathbb{C}^{m+z+1} , and if there are any, compute $\text{im}\Psi_{v,F} \cap \mathcal{M}_A$ by imposing the positivity constraints on the solution set. The resulting semi-algebraic set is then the feasible region over which $D_{\mathcal{M}_A}$ can be maximized to identify local maximizers with support equal to the support of $v(b)$. Finally, we locate the global maximizer(s) among these local maximizers contributed by each chamber C of the chamber complex \mathcal{C}_A that supports projection points. We summarize this in a high-level algorithm.

Algorithm:

Input: $A \in \mathbb{N}^{d \times n}$ that defines a toric model $\mathcal{M}_A \subset \Delta_{n-1}$ of dimension $d - 1$.

Output: All maximizers of $D_{\mathcal{M}_A}$.

1. Compute the equations of the toric variety X_A .
2. Compute the chamber complex \mathcal{C}_A .
3. For each chamber C in \mathcal{C}_A do:
 - a) for any fixed \hat{b} in the relative interior of C compute the face lattice of $Q_{\hat{b}}$ and identify complementary vertex/face pairs (v, F) ;
 - b) for each (v, F) do:
 - i. compute the parametrization $\Psi_{v,F}$ and substitute it into the equations of X_A ;
 - ii. if the resulting algebraic set in \mathbb{C}^{m+z+1} is nonempty then
 - * compute the semi-algebraic set $\text{im}\Psi_{v,F} \cap \mathcal{M}_A$ by imposing positivity constraints on the parameters in $\Psi_{v,F}$;
 - * find the maximizers $D_{C,v,F}$ of $D_{\mathcal{M}_A}$ over $\text{im}\Psi_{v,F} \cap \mathcal{M}_A$.
4. Identify global maximizers of $D_{\mathcal{M}_A}$ by comparing all $D_{C,v,F}$.

Example 4.2.9. We illustrate this algorithm using the toric model of Example 4.2.3. The equations of X_A are the three polynomials computed in Example 4.2.8. The chamber complex \mathcal{C}_A is the polytopal complex in Figure 4.1. The chambers which support complementary vertices are the (relative interior of) the boundary edges of the pentagonal chamber. Step 3

is executed only for these chambers. For instance, the horizontal edge is the convex hull of its vertices $(3/2, 1)$ and $(2, 1)$, and the unique complementary vertex/face pair (v, F) is given by vertex v with support $\{a, d\}$ and the vertex F with support $\{b, c, e\}$. We note that for such pair of complementary *vertices* $(v, \{w\})$ we do not need to consider the pair $(w, \{v\})$ in the next computation. The parametrization $\Psi_{v,F}$ is given by

$$(r_1, s) \mapsto \left(s\left(\frac{1}{6}r_1 + \frac{1}{3}\right), (1-s)\frac{r_1}{2}, (1-s)\frac{1-r_1}{2}, s\left(-\frac{1}{6}r_1 + \frac{2}{3}\right), \frac{1-s}{2} \right),$$

where we are parametrizing b on this edge by $r_1(3/2, 1) + (1-r_1)(2, 1)$. Substituting $\Psi_{v,F}$ into the equations of X_A results in

$$\begin{aligned} s^2r_1^2 + 7s^2r_1 - 8s^2 - 18sr_1 + 9r_1 &= 0 \\ 197s^4r_1 - 194s^4 - 1401s^3r_1 - 3sr_1^3 + 1014s^3 + 4398s^2r_1 + 246sr_1^2 - 2094s^2 - 5837sr_1 - 81r_1^2 + 2s + 2349r_1 &= 0 \\ 885s^4 - 31312s^3r_1 - 294sr_1^3 + 32392s^3 + 179435s^2r_1 + 17016sr_1^2 - 117350s^2 - 295438sr_1 - 6165r_1^2 + 2560s + 129141r_1 - 591 &= 0. \end{aligned}$$

This is a zero-dimensional system that has 11 solutions which we have computed using *Bertini*. Four of these are complex and seven are real. There is a unique real solution where $0 < r_1, s < 1$, namely

$$r_1 = 0.4702953126494577 \text{ and } s = 0.4106301713351522.$$

The corresponding KL-divergence at the vertex v is 0.890062259952966. At the vertex w , the divergence is 0.528701425022976. For each of the remaining four edges of this pentagonal chamber we also get a pair of projection vertices with corresponding KL-divergences equal to

$$\begin{aligned} &0.729916767214609 \text{ and } 0.657681783609608 \\ &0.736523721240758 \text{ and } 0.651574202843057 \\ &0.927851227501820 \text{ and } 0.503192212618303 \\ &0.856820834934792 \text{ and } 0.552532602066626. \end{aligned}$$

The global maximizer is the vertex

$$v = (0, 0.6722451790633609, 0, 0, 0.3277548209366391)$$

corresponding to the divergence value 0.927851227501820. It is a vertex of the logarithmic Voronoi polytope Q_b where $b = (1.3277548209366392, 0.6555096418732783)$ lies on the edge of the pentagonal chamber contained in the line segment between $(1, 0)$ and $(2, 2)$.

The basic algorithm above can be improved on many fronts. We will now present some ideas for such improvements.

For Step 1, one could replace the equations of X_A , which could be challenging to compute for large models, with $n - d$ equations corresponding to a basis of $\ker_{\mathbb{Z}}(A)$. Let B be an

$(n - d) \times n$ matrix whose rows b_i , $i = 1, \dots, n - d$ form such a basis. The lattice basis ideal

$$I_B = \left\langle \prod_{j=1}^n p_j^{b_{ij}^+} - \prod_{j=1}^n p_j^{b_{ij}^-}, \quad i = 1, \dots, n - d \right\rangle$$

where $b_i = b_i^+ - b_i^-$ with $b_i^+, b_i^- \geq 0$ and $\text{supp}(b_i^+) \cap \text{supp}(b_i^-) = \emptyset$ defines a variety Y_B containing X_A . In fact, Y_B is the union of X_A together with varieties contained in various coordinate subspaces defined by setting a subset of coordinates equal to zero (see [109, Section 8.3] and [64]). This means that $\mathcal{M}_A^{>0} = X_A \cap \Delta_{n-1}^\circ$ is equal to $Y_B \cap \Delta_{n-1}^\circ$. This is what is ultimately needed in Step 3.b.ii.

For Step 2, Example 4.2.6 illustrated that computing \mathcal{C}_A might be out of reach due to the combinatorial explosion in the number of chambers. However, one does not need to compute \mathcal{C}_A all at once. It can be computed one chamber at a time. This is how a software like *Gfan* [70] internally computes \mathcal{C}_A based on reverse search enumeration [15]. In this case, Step 3 can be executed as chambers get computed.

In Step 3, not all chambers need to be considered. For instance, any chamber that is contained in the boundary of $\text{conv}(A)$ can be skipped: if b is in such a chamber, the logarithmic Voronoi polytope Q_b is contained in the boundary of Δ_{n-1} . Such Q_b does not contribute global maximizers of $D_{\mathcal{M}_A}$. There are also ways to eliminate chambers since they cannot contain complementary vertices. We present a few ways this can be done.

Proposition 4.2.10. Suppose $\text{conv}(A)$ is a simplicial polytope where each column of A is a vertex. Let C be a chamber that intersects the boundary as well as the interior of $\text{conv}(A)$. Then for any b that is in the relative interior of C , the logarithmic Voronoi polytope Q_b does not contain complementary vertices.

Proof. The intersection of C with the boundary of $\text{conv}(A)$ is a simplex spanned by a subset of columns of A , say A_{i_1}, \dots, A_{i_k} . Then the support of every vertex of Q_b contains $\{i_1, \dots, i_k\}$. This disallows the existence of complementary vertices. \square

Note, for instance, in our running Example 4.2.3, it is enough to consider the pentagonal chamber and its faces by the above proposition. In fact, the interior of this chamber does not have to be considered either for the following reason.

Proposition 4.2.11. Let C be a chamber of dimension k where $k + 1 > n/2$. Then for any b that is in the relative interior of C , the logarithmic Voronoi polytope Q_b does not contain complementary vertices.

Proof. By Theorem 4.2.2, each of the vertices of Q_b has support of size at least $k + 1$. If a vertex v of Q_b is complementary there must exist a vertex w such that $\text{supp}(v) \cap \text{supp}(w) = \emptyset$. Such two vertices can only exist when $2(k + 1) \leq n$. \square

Proposition 4.2.12. If (v, F) is a pair of a complementary vertex and its complementary face F where both v and F are contained in the same facet F' of Q_b , then v cannot be a projection point.

Proof. The line segments from v to the points in F are entirely contained in F' which is in the boundary of Δ_{n-1} . Then v cannot be a projection point since no such line segment can intersect the toric model \mathcal{M}_A in the interior of Δ_{n-1} . \square

Again, we note that, Proposition 4.2.12 rules out the zero-dimensional chambers that are the vertices of the pentagonal chamber in Example 4.2.3 since they give rise to complementary pairs (v, F) lying in the same facet of their logarithmic Voronoi polytope.

Proposition 4.2.13. Let C be a chamber in the chamber complex \mathcal{C}_A . If no two vertices of the logarithmic Voronoi polytope Q_b corresponding to points b in the relative interior of C have disjoint supports, then the same is true for any chamber C' containing C .

Proof. The supports of vertices of $Q_{b'}$ where b' is in the relative interior of C' are in bijection with σ' where columns of $A_{\sigma'}$ are affinely independent and the relative interior of $\text{conv}(A_{\sigma'})$ contains the relative interior of C' . Since C is a face of C' , for any σ such that the columns of A_{σ} are affinely independent and the relative interior of $\text{conv}(A_{\sigma})$ contains the relative interior of C , there is (possibly multiple) $\sigma' \supset \sigma$ as above. Hence if no two vertices of Q_b have disjoint supports, the same is true for $Q_{b'}$. \square

Corollary 4.2.14. Let $A \in \mathbb{N}^{3 \times 5}$ such that $\text{conv}(A)$ is a planar pentagon. If a logarithmic Voronoi polytope Q_b contains a projection point then b is in the interior of an edge of the pentagonal chamber. Moreover, each such edge contributes either finitely many projection points or for every b on the edge, Q_b has a projection point.

Proof. Propositions 4.2.10, 4.2.11, and 4.2.12 imply the first statement. Any logarithmic Voronoi polytope Q_b where b is on an edge of the pentagonal chamber has a pair of complementary vertices (v, w) . The Zariski closure of the image of $\Psi_{v,w} : \Delta_1 \times \Delta_1 \rightarrow \Delta_4$ in $\mathbb{P}_{\mathbb{C}}^4$ is a two-dimensional irreducible surface. Since X_A is also two-dimensional and irreducible, and it is never equal to the former Zariski closure, their intersection has either finitely many points (this is the generic case) or it is an algebraic curve. This means that $\text{im}\Psi_{v,w} \cap \mathcal{M}_A$ has either finitely many points or contains the positive real part of an algebraic curve. In the second case, the projection of the preimage of this positive real part under $\Psi_{v,w}$ to the first Δ_1 in the domain of $\Psi_{v,w}$ must be all of Δ_1 . Hence, for every b on this edge, Q_b has a projection point. \square

Our final remark about the algorithm concerns the step where $D_{\mathcal{M}_A}$ needs to be maximized over the semi-algebraic set $\text{im}\Psi_{v,F} \cap \mathcal{M}_A$. Of course, this is a challenging step. However, generically one expects this set to be finite. In that case, numerical algebraic geometry tools

perform well to compute each point in this finite intersection. Another relatively easier case is when the maximum likelihood degree of X_A is one; see Definition 4.4.1. There are two advantages in this case. First, the intersection of $\text{im}\Psi_{v,F}$ with X_A is guaranteed to be in Δ_{n-1} since the affine span of each logarithmic Voronoi polytope Q_b intersects X_A in exactly one point, namely the unique maximum likelihood estimator $q(b)$ in Q_b . Second, $q(b)$ is a rational function of $v(b)$ – an equivalent condition for an algebraic statistical model to have maximum likelihood degree equal to one. In other words, both $v(b)$ and $q(b)$ are rational functions of the parameters $(r_1, \dots, r_m) \in \Delta_{m-1}$. In turn, $D_{\mathcal{M}_A}$ restricted to the potential projection points $v(b)$ is a greatly simplified function of the same parameters. Now, one needs to optimize $D_{\mathcal{M}_A}(r_1, \dots, r_m)$ over Δ_{m-1} .

Example 4.2.15 (Independence model 2×3). Consider the independence model of two random variables, binary X and ternary Y . Similar to Example 4.1.7, this is a 3-dimensional toric model inside Δ_5 given by the matrix

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

The polytope $\text{conv}(A)$ is a triangular prism with six vertices that is highly symmetrical due to the action of the group $S_2 \times S_3$ on the states of X and Y . This, in turn, induces a partition of the elements in the chamber complex \mathcal{C}_A into symmetry classes. This way, 18 full-dimensional chambers are split into 5 classes, 44 ridges are split into 7 classes, 36 edges are split into 6 classes, and 11 vertices are split into 3 classes. Figure 4.2 demonstrates this division of full-dimensional chambers: any chambers that share a color are in the same symmetry class. The red middle chamber is a bipyramid with a triangular base and is the only one in its class.

To run the algorithm, note that 9 out of the 11 vertices are on the boundary of $\text{conv}(A)$, and hence do not contribute any projection vertices by Proposition 4.2.12. Call the two interior vertices b_1 and b_2 . They are both in the same symmetry class and lie on the middle red full-dimensional chamber. The logarithmic Voronoi polytope at a point corresponding to b_1 is a triangle with no complementary vertices, and the same is true of b_2 by symmetry. Hence, no vertices of the chamber complex will contribute any projection points. Next, out of 36 edges 21 are on the boundary. Moreover, 6 of the remaining edges contain the vertex b_1 and by symmetry another 6 contain the vertex b_2 , so we do not need to check these edges by Proposition 4.2.13. This leaves us with three edges e_1 , e_2 , and e_3 on the base of the red bipyramid. We will treat them in the next paragraph. Out of 44 ridges, 14 are on the boundary, 12 contain vertex b_1 , and another 12 contain vertex b_2 . The remaining 6 are in the same symmetry class. Logarithmic Voronoi polytopes corresponding to these are quadrilaterals with supports like $\{1234, 1345, 1246, 156\}$ that contain no complementary




Figure 4.2: Chamber complex of the 2×3 independence model (left and middle) and the middle chamber (right).

vertices. Hence, none of the ridges will contribute projection points. Finally, none of the three-dimensional chambers will contribute any projection points by Proposition 4.2.11.

Hence, we only need to run step 3 of our algorithm on the edges e_1 , e_2 , and e_3 . By symmetry, it suffices to run it on e_1 only. A point b on this edge can be parametrized as $r_1(1/2, 1/2, 0) + (1 - r_1)(1/2, 0, 1/2)$. The only vertex-face pair we need to consider is the pair of complementary vertices $(v, \{w\})$, where $v = 1/2(1, 0, 0, 0, r_1, 1 - r_1)$ and $w = 1/2(0, r_1, 1 - r_1, 1, 0, 0)$. The parametrization $\Psi_{v, \{w\}}$ of the line between them gives rise to the single equation $(s - 1)^2 - s^2 = 0$. Therefore $s = 1/2$, while r_1 is a free variable between 0 and 1. Upon substituting $s = 1/2$ into $D(v, \text{in}(\Psi_{v, \{w\}}))$, we get the constant value $\log 2$. Therefore, the divergence at *every* point b of the edge e_1 is $\log 2$, attained at the two vertices of the logarithmic Voronoi polytope v and w . By symmetry, the same is true of e_2 and e_3 . We conclude that the maximum divergence from this model is $\log 2$ and there are infinitely many maximizers which we completely characterized above. These maximizers were also studied and visualized in [18].

4.3 Reducible models

This section is devoted to logarithmic Voronoi polytopes of toric models that are known as reducible hierarchical log-linear models [47, 65, 78]. Besides giving a structural result about these polytopes, we also prove results relating the maximum information divergence to such models with those that are obtained by certain marginalizations. For similar work see [83].

A *simplicial complex* is a set $\Gamma \subseteq 2^{[m]}$ such that if $F \in \Gamma$ and $S \subseteq F$, then $S \in \Gamma$. The elements of Γ are called *faces*. We refer to inclusion-maximal faces of Γ as *facets*. It is sufficient to list the facets to describe a simplicial complex. For example, $\Gamma = [12][13][23]$ will denote the simplicial complex $\Gamma = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}\}$.

Let X_1, \dots, X_m be discrete random variables. For each $i \in [m]$, assume that X_i has the state space $[d_i]$ for some $d_i \in \mathbb{N}$. Let $\mathcal{R} = \prod_{i=1}^m [d_i]$ be the state space of the random vector (X_1, \dots, X_m) . For each $i = (i_1, \dots, i_m) \in \mathcal{R}$ and $F = \{f_1, f_2, \dots\} \subseteq [m]$, we will denote $i_F = (i_{f_1}, i_{f_2}, \dots)$. Moreover, each such subset $F \subseteq [m]$ gives rise to the random vector $X_F = (X_f)_{f \in F}$ with the state space $\mathcal{R}_F = \prod_{f \in F} [d_f]$.

Definition 4.3.1. Let $\Gamma \subseteq 2^{[m]}$ be a simplicial complex and let $d_1, \dots, d_m \in \mathbb{N}$. For each facet $F \in \Gamma$, introduce $|\mathcal{R}_F|$ parameters $\theta_{i_F}^{(F)}$, one for each $i_F \in \mathcal{R}_F$. The *hierarchical log-linear model* associated with Γ and $\mathbf{d} = (d_1, \dots, d_m)$ is defined to be

$$\mathcal{M}_{\Gamma, \mathbf{d}} = \left\{ p \in \Delta_{|\mathcal{R}|-1} : p_i = \frac{1}{Z(\theta)} \prod_{F \in \text{facets}(\Gamma)} \theta_{i_F}^{(F)} \text{ for all } i \in \mathcal{R} \right\}$$

where $Z(\theta)$ is the normalizing constant defined as

$$Z(\theta) := \sum_{i \in \mathcal{R}} \prod_{F \in \text{facets}(\Gamma)} \theta_{i_F}^{(F)}.$$

If $u \in \mathbb{N}^{|\mathcal{R}|}$ is a $d_1 \times \dots \times d_m$ contingency table containing data for the random vector (X_1, \dots, X_m) and $F = \{f_1, f_2, \dots\} \subseteq [m]$, let u_F denote the $d_{f_1} \times d_{f_2} \times \dots$ table with $(u_F)_{i_F} = \sum_{j \in \mathcal{R}_{[m] \setminus F}} u_{i_F, j}$. Such table u_F is called the *F-marginal* of u . For simplicity, we will denote the simplex in which $\mathcal{M}_{\Gamma, \mathbf{d}}$ lives by $\Delta_{\Gamma, \mathbf{d}}$.

Proposition 4.3.2. [47, Prop. 1.2.9] Hierarchical log-linear models are toric models. For any simplicial complex $\Gamma \subseteq 2^{[m]}$ and positive integers $\mathbf{d} = (d_1, \dots, d_m)$, the model $\mathcal{M}_{\Gamma, \mathbf{d}}$ is realized by the 0/1 matrix $A_{\Gamma, \mathbf{d}}$ representing the marginalization map

$$\varphi(u) = (u_{F_1}, u_{F_2}, \dots)$$

where F_1, F_2, \dots are the facets of Γ . In other words, $\mathcal{M}_{\Gamma, \mathbf{d}} = \mathcal{M}_{A_{\Gamma, \mathbf{d}}}$.

Here we wish to point out that for any point $q \in \Delta_{\Gamma, \mathbf{d}}$ (in particular, for $q \in \mathcal{M}_{\Gamma, \mathbf{d}}$) the logarithmic Voronoi polytope Q_q^Γ consists of all $p \in \Delta_{\Gamma, \mathbf{d}}$ such that $\varphi(p) = \varphi(q)$.

Definition 4.3.3. A simplicial complex Γ on $[m]$ is called *reducible* with decomposition (Γ_1, S, Γ_2) if there exist sub-complexes Γ_1, Γ_2 of Γ and a subset $S \subseteq [m]$ such that $\Gamma = \Gamma_1 \cup \Gamma_2$ and $\Gamma_1 \cap \Gamma_2 = 2^S$. We say Γ is *decomposable* if it is reducible and each of the Γ_1, Γ_2 is either decomposable or a simplex. A hierarchical log-linear model associated to a reducible (decomposable) simplicial complex is called *reducible (decomposable)*.

Decomposition theory of logarithmic Voronoi polytopes

Let Γ be a reducible simplicial complex on $[m]$ with decomposition (Γ_1, S, Γ_2) and $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$. Suppose Γ_1 has the vertex set $\alpha = \{\alpha_1, \dots, \alpha_k\}$ and Γ_2 has the vertex set $\beta = \{\beta_1, \dots, \beta_s\}$. Then $S = \{\alpha_1, \dots, \alpha_k\} \cap \{\beta_1, \dots, \beta_s\}$. We also let $\mathbf{d}_\alpha = (d_{\alpha_1}, \dots, d_{\alpha_k})$, with analogous definitions for \mathbf{d}_β and \mathbf{d}_S . Let p be a point in $\Delta_{\Gamma, \mathbf{d}}$ and consider the maps

$$\pi_1 : \Delta_{\Gamma, \mathbf{d}} \rightarrow \Delta_{\Gamma_1, \mathbf{d}_\alpha} \quad p \mapsto p_1 = p_{\{\alpha_1, \dots, \alpha_k\}}.$$

$$\pi_2 : \Delta_{\Gamma, \mathbf{d}} \rightarrow \Delta_{\Gamma_2, \mathbf{d}_\beta} \quad p \mapsto p_2 = p_{\{\beta_1, \dots, \beta_s\}}.$$

More precisely,

$$(\pi_1(p))_{i_\alpha} = \sum_{j \in \mathcal{R}: j_\alpha = i_\alpha} p_j \quad \text{and} \quad (\pi_2(p))_{i_\beta} = \sum_{j \in \mathcal{R}: j_\beta = i_\beta} p_j.$$

Lemma 4.3.4. Let Γ be a reducible simplicial complex on $[m]$ with decomposition (Γ_1, S, Γ_2) and $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$. Let $q \in \mathcal{M}_{\Gamma, \mathbf{d}}$ so that $q_1 = \pi_1(q)$ and $q_2 = \pi_2(q)$. Furthermore, consider the maps

$$\pi'_1 : \Delta_{\Gamma_1, \mathbf{d}_\alpha} \rightarrow \Delta_{2^S, \mathbf{d}_S} \quad p \mapsto p_S$$

$$\pi'_2 : \Delta_{\Gamma_2, \mathbf{d}_\beta} \rightarrow \Delta_{2^S, \mathbf{d}_S} \quad p \mapsto p_S$$

defined by

$$(\pi'_1(p))_{i_S} = \sum_{j \in \mathcal{R}_\alpha: j_S = i_S} p_j \quad \text{and} \quad (\pi'_2(p))_{i_S} = \sum_{j \in \mathcal{R}_\beta: j_S = i_S} p_j.$$

Then $q_1 \in \mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}$ and $q_2 \in \mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}$, and the following diagram commutes:

$$\begin{array}{ccc}
 & q \in \mathcal{M}_{\Gamma, \mathbf{d}} & \\
 \swarrow \pi_1 & & \searrow \pi_2 \\
 q_1 \in \mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha} & & q_2 \in \mathcal{M}_{\Gamma_2, \mathbf{d}_\beta} \\
 \searrow \pi'_1 & & \swarrow \pi'_2 \\
 & q_3 \in \mathcal{M}_{2^S, \mathbf{d}_S} &
 \end{array}$$

Proof. By the definitions of the maps, $\pi'_1 \circ \pi_1 = \pi'_2 \circ \pi_2$. Also, since $\mathcal{M}_{2^S, \mathbf{d}_S} = \Delta_{2^S, \mathbf{d}_S}$ it is clear that $q_3 \in \mathcal{M}_{2^S, \mathbf{d}_S}$. We just need to show $q_1 \in \mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}$ and $q_2 \in \mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}$. We prove the first claim since the second one requires the same argument. Let $t \in \mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}$ be the MLE of q_1 and $r \in \mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}$ be the MLE of q_2 . We will show that $q_1 = t$. Note that $q \in \mathcal{M}_{\Gamma, \mathbf{d}}$, so it is its own MLE in the model. Since t is in the same logarithmic Voronoi polytope as

q_1 and r is in the same logarithmic Voronoi polytope as q_2 , we see that $\pi'_1(t) = t_S = q_3$ and $\pi'_2(r) = r_S = q_3$. Then by [78, Prop 4.1.4]

$$q_{i_1, \dots, i_m} = \frac{(t_{i_\alpha}) \cdot (r_{i_\beta})}{(r_S)_{i_S}},$$

where $i_\alpha = (i_{\alpha_1}, \dots, i_{\alpha_k})$ and $i_\beta = (i_{\beta_1}, \dots, i_{\beta_s})$. Then observe that for any i_α , we get

$$(q_1)_{i_\alpha} = \sum_{j \in \mathcal{R}: j_\alpha = i_\alpha} \frac{(t_{j_\alpha}) \cdot (r_{j_\beta})}{(r_S)_{j_S}} = \frac{t_{i_\alpha}}{(r_S)_{i_S}} \sum_{j_\beta: j_S = i_S} r_{j_\beta} = \frac{t_{i_\alpha} \cdot (r_S)_{i_S}}{(r_S)_{i_S}} = t_{i_\alpha}.$$

Since i_α was arbitrary, we get that $q_1 = t$. \square

Now we are ready to prove the main result of this section. From the discussion so far we see that π_1 and π_2 restrict to logarithmic Voronoi polytopes, i.e., $\pi_1 : Q_p^\Gamma \rightarrow Q_{p_1}^{\Gamma_1}$ and $\pi_2 : Q_p^\Gamma \rightarrow Q_{p_2}^{\Gamma_2}$ where $p_1 = \pi_1(p)$ and $p_2 = \pi_2(p)$ for $p \in \Delta_{\Gamma, \mathbf{d}}$. In fact, we can take $q = p \in \mathcal{M}_{\Gamma, \mathbf{d}}$ so that q_1 and q_2 are in $\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}$ and $\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}$, respectively, by the above lemma. The next theorem reconstructs Q_p^Γ from the logarithmic Voronoi polytopes $Q_{p_1}^{\Gamma_1}$ and $Q_{p_2}^{\Gamma_2}$.

Theorem 4.3.5. Let Γ be a reducible simplicial complex on $[m]$ with a decomposition (Γ_1, S, Γ_2) and $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$. Let $\psi : \Delta_{\Gamma, \mathbf{d}} \rightarrow \Delta_{\Gamma_1, \mathbf{d}_\alpha} \times \Delta_{\Gamma_2, \mathbf{d}_\beta}$ be the map $\psi(u) = (\pi_1(u), \pi_2(u))$. Then for any $q \in \mathcal{M}_{\Gamma, \mathbf{d}}$, we have

$$Q_q^\Gamma = \left[\left\{ u \in \Delta_{\Gamma, \mathbf{d}} : u_{i_1, \dots, i_m} = \frac{v_{i_\alpha} \cdot w_{i_\beta}}{(q_S)_{i_S}} \text{ for } v \in Q_{q_1}^{\Gamma_1} \text{ and } w \in Q_{q_2}^{\Gamma_2} \right\} + \ker(\psi) \right] \cap \Delta_{\Gamma, \mathbf{d}}$$

Proof. We proceed by double containment. To show that the right-hand side is contained in Q_q^Γ , let $u = u^{(1)} + u^{(2)} \in \Delta_{\Gamma, \mathbf{d}}$ where $u_{i_1, \dots, i_m}^{(1)} = \frac{v_{i_\alpha} \cdot w_{i_\beta}}{(q_S)_{i_S}}$ for $v \in Q_{q_1}^{\Gamma_1}$, $w \in Q_{q_2}^{\Gamma_2}$ and $u^{(2)} \in \ker(\psi)$. Let $F = \{f_1, \dots, f_k\}$ be any facet of Γ . Then F is either in Γ_1 or in Γ_2 . Without loss of generality, assume F is in Γ_1 . Then for any $i_F = (i_{f_1}, \dots, i_{f_k})$, we have

$$\begin{aligned} ((u^{(1)})_F)_{i_F} &= \sum_{\{j_\alpha: j_F = i_F\}} \left(\sum_{\{j \in \mathcal{R}: j_\alpha = i_\alpha, j_F = i_F\}} \frac{v_{j_\alpha} \cdot w_{j_\beta}}{(q_S)_{j_S}} \right) = \sum_{\{j_\alpha: j_F = i_F\}} \frac{v_{j_\alpha} \cdot (w_S)_{j_S}}{(q_S)_{j_S}} \\ &= \sum_{\{j_\alpha: j_F = i_F\}} v_{j_\alpha} = (v_F)_{i_F} = (q_F)_{i_F}. \end{aligned}$$

Hence $u^{(1)} \in Q_q^\Gamma$. But since $u^{(2)} \in \ker(\psi)$, it has a zero F -marginal for every facet of Γ . Thus, $u = u^{(1)} + u^{(2)} \in Q_q^\Gamma$, as desired.

To show the reverse containment, let $u \in Q_q^\Gamma$ and let $v \in \Delta_{\Gamma, \mathbf{d}}$ be the point defined by

$$v_{i_1, \dots, i_m} = \frac{(u_\alpha)_{i_\alpha} \cdot (u_\beta)_{i_\beta}}{(qs)_{i_S}}$$

for all (i_1, \dots, i_m) . We write $u = v + (u - v)$. Since $u_\alpha = \pi_1(u) \in Q_{q_1}^{\Gamma_1}$ and $u_\beta = \pi_2(u) \in Q_{q_2}^{\Gamma_2}$, it suffices to show that $u - v \in \ker \psi$. That is, we must show that $[(u - v)_\alpha]_{i_\alpha} = 0$ and $[(u - v)_\beta]_{i_\beta} = 0$. For any i_α , $[(u - v)_\alpha]_{i_\alpha}$ is equal to

$$(u_\alpha)_{i_\alpha} - \sum_{\{j \in \mathcal{R}: j_\alpha = i_\alpha\}} \frac{(u_\alpha)_{j_\alpha} \cdot (u_\beta)_{j_\beta}}{(qs)_{j_S}} = (u_\alpha)_{i_\alpha} - \frac{(u_\alpha)_{i_\alpha}}{(qs)_{i_S}} \sum_{\{j_\beta: j_S = i_S\}} (u_\beta)_{j_\beta} = (u_\alpha)_{i_\alpha} - \frac{(u_\alpha)_{i_\alpha} \cdot (u_S)_{i_S}}{(qs)_{i_S}} = 0.$$

Similarly, one shows that $(u - v)_{i_\beta} = 0$ as well. Thus, $u - v \in \ker \psi$, and this concludes the proof. \square

In this theorem, the first summand in the Minkowski sum that appears in the decomposition of Q_q^Γ is an interesting object. It is nonlinear and captures a portion of Q_q^Γ .

Definition 4.3.6. Let Γ be a reducible simplicial complex on $[m]$ with decomposition (Γ_1, S, Γ_2) and $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$. Let $p \in \Delta_{\Gamma, \mathbf{d}}$ and $p_i = \pi_i(p)$ for $i = 1, 2$. Then the product of $Q_{p_1}^{\Gamma_1}$ and $Q_{p_2}^{\Gamma_2}$ is defined as

$$Q_{p_1}^{\Gamma_1} \otimes_p Q_{p_2}^{\Gamma_2} = \left\{ u \in \Delta_{\Gamma, \mathbf{d}} : u_{i_1, \dots, i_m} = \frac{v_{i_\alpha} \cdot w_{i_\beta}}{(ps)_{i_S}} \text{ for } v \in Q_{p_1}^{\Gamma_1} \text{ and } w \in Q_{p_2}^{\Gamma_2} \right\}.$$

Remark 4.3.7. If $p' \in Q_p^\Gamma$ we get the equality of the logarithmic Voronoi polytopes $Q_{p'}^\Gamma = Q_p^\Gamma$. Moreover, since $p'_i = \pi_i(p') \in Q_{p'_i}^{\Gamma_i}$ for $i = 1, 2$, we see that $Q_{p'_i}^{\Gamma_i} = Q_{p_i}^{\Gamma_i}$. Therefore, $Q_{p_1}^{\Gamma_1} \otimes_p Q_{p_2}^{\Gamma_2} = Q_{p'_1}^{\Gamma_1} \otimes_{p'} Q_{p'_2}^{\Gamma_2}$. In other words, the product depends only on the logarithmic Voronoi polytope Q_p^Γ and not on the individual points in the polytope.

Example 4.3.8. Consider the complex $\Gamma = [12][13][23][24][34]$ for $m = 4$. Suppose both X_1, X_2, X_3 , and X_4 are binary random variables, i.e., $d_i = 2$ for all $i = 1, \dots, 4$. Let $\Gamma_1 = [12][13][23]$ and $\Gamma_2 = [23][24][34]$, so $S = \{2, 3\}$. The logarithmic Voronoi polytopes $Q_{p_1}^{\Gamma_1}$ and $Q_{p_2}^{\Gamma_2}$ have dimension one whereas the dimension of Q_p^Γ is six. This is consistent with Theorem 4.3.5 since $Q_{p_1}^{\Gamma_1} \otimes_p Q_{p_2}^{\Gamma_2}$ is a two-dimensional surface in Δ_{15} and $\ker(\psi)$ has dimension four. More explicitly, if $v = (v_{ijk})$ and $w = (w_{jkl})$ are points in $Q_{p_1}^{\Gamma_1}$ and $Q_{p_2}^{\Gamma_2}$, respectively, where in particular $v_{+jk} = v_{1jk} + v_{2jk}$ and $w_{jk+} = w_{jk1} + w_{jk2}$ are equal to each other for all $j, k = 1, 2$, then $Q_{p_1}^{\Gamma_1} \otimes_p Q_{p_2}^{\Gamma_2}$ consists of points $u = (u_{ijkl})$ where

$$u_{ijkl} = \frac{v_{ijk} \cdot w_{jkl}}{v_{+jk}}.$$

Comparing divergences

Since a reducible model $\mathcal{M}_{\Gamma, \mathbf{d}}$ associated to a simplicial complex Γ on $[m]$ with decomposition (Γ_1, S, Γ_2) has the two associated models $\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}$ and $\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}$ it is natural to ask how the divergences from these three models are related. Before we present our contributions we wish to cite two results of Matúš that are relevant.

Proposition 4.3.9. [83, Lemma 3] For any $p \in \Delta_{\Gamma, \mathbf{d}}$ and a reducible model $\mathcal{M}_{\Gamma, \mathbf{d}}$,

$$D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(p) = D_{\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}}(\pi_1(p)) + D_{\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}}(\pi_2(p)) + H(\pi_1(p)) + H(\pi_2(p)) - H(p) - H((\pi'_1 \circ \pi_1)(p))$$

where $H(\cdot)$ is the entropy.

Proposition 4.3.10. [83, Corollary 3] For a hierarchical log-linear model $\mathcal{M}_{\Gamma, \mathbf{d}}$ we have

$$D(\mathcal{M}_{\Gamma, \mathbf{d}}) \leq \min_{F \text{ facet of } \Gamma} \left\{ \sum_{i \notin F} \log d_i \right\}.$$

With regards to Proposition 4.3.9 we point out that the four entropy terms together give a nonnegative quantity because of the strong subadditivity property of entropy. Therefore, for a reducible model we get the inequality $D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(p) \geq D_{\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}}(\pi_1(p)) + D_{\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}}(\pi_2(p))$. We state and prove a similar inequality in Corollary 4.3.13. In the case when the point p lives in the product portion of its logarithmic Voronoi polytope (as in Theorem 4.3.5), we recover the equality below.

Proposition 4.3.11. Let Γ be a reducible simplicial complex on $[m]$ with decomposition (Γ_1, S, Γ_2) and $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$. Let $p \in \Delta_{\Gamma, \mathbf{d}}$ and $p_i = \pi_i(p)$ for $i = 1, 2$. If $u = v \otimes_p w$ where $v \in Q_{p_1}^{\Gamma_1}$ and $w \in Q_{p_2}^{\Gamma_2}$, then $D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(u) = D_{\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}}(v) + D_{\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}}(w)$.

Proof. Let $t \in \mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}$ and $r \in \mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}$ be the respective maximum likelihood estimators of v and w . Similarly, let $q \in \mathcal{M}_{\Gamma, \mathbf{d}}$ be the maximum likelihood estimator of u . By [78, Prop 4.1.4], $q = t \otimes_p r$ and

$$D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(u) = \sum_{i \in \mathcal{R}} u_i \log(u_i/q_i) = \sum_{i \in \mathcal{R}} u_i \log(v_{i_\alpha}/t_{i_\alpha}) + \sum_{i \in \mathcal{R}} u_i \log(w_{i_\beta}/r_{i_\beta}).$$

Since $(u_\alpha)_{i_\alpha} = \sum_{j \in \mathcal{R}: j_\alpha = i_\alpha} u_j = v_{i_\alpha}$ and $(u_\beta)_{i_\beta} = \sum_{j \in \mathcal{R}: j_\beta = i_\beta} u_j = w_{i_\beta}$ (see the proof of Lemma 4.3.4) we conclude that

$$D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(u) = \sum_{i_\alpha} v_{i_\alpha} \log(v_{i_\alpha}/t_{i_\alpha}) + \sum_{i_\beta} w_{i_\beta} \log(w_{i_\beta}/r_{i_\beta}) = D_{\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}}(v) + D_{\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}}(w).$$

□

Corollary 4.3.12. Let Γ be a reducible simplicial complex on $[m]$ with decomposition (Γ_1, S, Γ_2) and $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$. Let $p \in \Delta_{\Gamma, \mathbf{d}}$ and $p_i = \pi_i(p)$ for $i = 1, 2$. Suppose $v \in Q_{p_1}^{\Gamma_1}$ maximizes the divergence to $\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}$ over all points in $Q_{p_1}^{\Gamma_1}$. Similarly, suppose $w \in Q_{p_2}^{\Gamma_2}$ and $p' \in Q_p^\Gamma$ be such maximizers. Then $D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(p') \geq D_{\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}}(v) + D_{\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta}}(w)$.

Proof. Use Proposition 4.3.11 with $D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(p') \geq D_{\mathcal{M}_{\Gamma, \mathbf{d}}}(u)$ where $u = v \otimes_p w$. □

The corollary has the following implication for the maximum divergence to a reducible model.

Corollary 4.3.13. Let Γ be a reducible simplicial complex on $[m]$ with decomposition (Γ_1, S, Γ_2) and $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$. Let $p \in \Delta_{\Gamma, \mathbf{d}}$ be a point that attains the maximum divergence $D(\mathcal{M}_{\Gamma, \mathbf{d}})$ and $p_i = \pi_i(p)$ for $i = 1, 2$. If $Q_{p_1}^{\Gamma_1}$ and $Q_{p_2}^{\Gamma_2}$ contain points which attain the maximum divergence $D(\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha})$ and $D(\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta})$, respectively, then $D(\mathcal{M}_{\Gamma, \mathbf{d}}) \geq D(\mathcal{M}_{\Gamma_1, \mathbf{d}_\alpha}) + D(\mathcal{M}_{\Gamma_2, \mathbf{d}_\beta})$.

Independence and related models

We have already encountered an independence model in Example 4.2.15. More generally, for discrete random variables X_1, \dots, X_m with respective state spaces $[d_i]$, the independence model is the hierarchical log-linear model on $[m]$ associated to the simplicial complex Γ consisting of just the m vertices $\Gamma = [1][2] \cdots [m]$. This is a reducible (in fact decomposable) model. Proposition 4.3.10 immediately implies the following (see also [18]).

Corollary 4.3.14. Let \mathcal{M} be the independence model of m discrete random variables X_1, \dots, X_m with state spaces $[d_i]$, respectively, where $d_1 \leq d_2 \leq \dots \leq d_m$. Then

$$D(\mathcal{M}) \leq \log d_1 + \dots + \log d_{m-1}.$$

Those independence models which achieve the upper bound in this result have been characterized [18, Theorem 4.4]. For instance, when $m = 2$ as well as in the case of $d_1 = \dots = d_m$, the upper bound is achieved. Since we will use it later we record a precise result regarding the latter case. For this, let S_{d+1} denote the group of permutations on $\{0, 1, \dots, d\}$ and let δ denote the Dirac delta (indicator) function. Namely, δ_v returns a vector in \mathbb{R}^n that has 1 in the coordinate indexed by v and 0 everywhere else.

Theorem 4.3.15. [18] Let \mathcal{M} be the independence model of m $(d+1)$ -ary random variables. Then the maximum divergence from \mathcal{M} is $D(\mathcal{M}) = (m-1) \log(d+1)$. This maximum value is achieved at vertices of the logarithmic Voronoi polytope at the unique point $q = \left(\frac{1}{(d+1)^m}, \dots, \frac{1}{(d+1)^m} \right)$. Each maximizer has the form $\frac{1}{d+1} \sum_{j=0}^d \delta_{j, \sigma_2(j), \dots, \sigma_m(j)}$ where $\sigma_i \in S_{d+1}$ for all $i = 2, \dots, m$.

Now we wish to illustrate the utility of Corollary 4.3.13 for independence models in two examples.

Example 4.3.16. The independence model \mathcal{M}_{222} of three binary variables is a 3-dimensional model in Δ_7 . We denote the coordinates of the points in Δ_7 by p_{ijk} where $i, j, k = 1, 2$. By using Corollary 4.3.13 we can show that $D(\mathcal{M}_{222}) = 2 \log 2$. Note that this model is reducible with $\Gamma = (\Gamma_1, S, \Gamma_2)$ where $\Gamma_1 = [1][2]$, $\Gamma_2 = [2][3]$, and $S = \{2\}$. The models \mathcal{M}_{Γ_1} and \mathcal{M}_{Γ_2} are themselves independence models of two binary variables. In Example 4.1.7 we saw that $D(\mathcal{M}_{\Gamma_1}) = D(\mathcal{M}_{\Gamma_2}) = \log 2$ where there are exactly two maximizers

$$v = (v_{11}, v_{12}, v_{21}, v_{22}) = \left(\frac{1}{2}, 0, 0, \frac{1}{2}\right) \text{ and } w = (w_{11}, w_{12}, w_{21}, w_{22}) = \left(0, \frac{1}{2}, \frac{1}{2}, 0\right).$$

We will view v and w as elements of the logarithmic Voronoi polytopes $Q_v^{\Gamma_1}$ and $Q_w^{\Gamma_2}$, respectively. These two polytopes are *compatible* in the sense that $\pi'_1(v) = (v_{+1}, v_{+2}) = \left(\frac{1}{2}, \frac{1}{2}\right)$ is equal to $\pi'_2(w) = (w_{1+}, w_{2+}) = \left(\frac{1}{2}, \frac{1}{2}\right)$. In other words, there exists a logarithmic Voronoi polytope Q_p^Γ such that $Q_v^{\Gamma_1} \otimes Q_w^{\Gamma_2} \subset Q_p^\Gamma$. Here $p = v \otimes w$ where $p_{ijk} = \frac{v_{ij}w_{jk}}{v_{+j}}$, and we see that $p_{112} = p_{221} = \frac{1}{2}$. Since $D_{\mathcal{M}_{222}}(p) = D(\mathcal{M}_{\Gamma_1}) + D(\mathcal{M}_{\Gamma_2}) = \log 2 + \log 2$ we conclude that $D(\mathcal{M}_{222}) = 2 \log 2$.

Example 4.3.17. Now we consider the independence model \mathcal{M}_{233} . Corollary 4.3.14 states that $D(\mathcal{M}_{233}) \leq \log 2 + \log 3$, but this bound cannot be attained by [18, Theorem 4.4]. We wish to provide a rationale based on Corollary 4.3.13. The model \mathcal{M}_{Γ_1} is the independence model of a binary and a ternary random variables, and the model \mathcal{M}_{Γ_2} is the independence model of two ternary random variables. By Example 4.2.15, there are six types of divergence maximizers for \mathcal{M}_{Γ_1} . If we denote the points in Δ_5 in which \mathcal{M}_{Γ_1} is contained by $v = (v_{11}, v_{12}, v_{13}; v_{21}, v_{22}, v_{23})$ these maximizers are

$$\begin{aligned} \left(\frac{1}{2}, 0, 0; 0, \frac{r}{2}, \frac{1-r}{2}\right) & \quad \left(0, \frac{r}{2}, \frac{1-r}{2}; \frac{1}{2}, 0, 0\right) \\ \left(0, \frac{1}{2}, 0; \frac{r}{2}, 0, \frac{1-r}{2}\right) & \quad \left(\frac{r}{2}, 0, \frac{1-r}{2}; 0, \frac{1}{2}, 0\right) \\ \left(0, 0, \frac{1}{2}; \frac{r}{2}, \frac{1-r}{2}, 0\right) & \quad \left(\frac{r}{2}, \frac{1-r}{2}, 0; 0, 0, \frac{1}{2}\right) \end{aligned}$$

where $0 < r < 1$. According to Theorem 4.3.15, there are six divergence maximizers of $\mathcal{M}_{\Gamma_2} \subset \Delta_8$. If we denote the points in Δ_8 by $w = (w_{jk} : j, k = 1, 2, 3)$ these maximizers are w^σ for each $\sigma \in S_3$ given by

$$\begin{aligned} w_{11}^{id} = w_{22}^{id} = w_{33}^{id} &= \frac{1}{3} & w_{12}^{(12)} = w_{21}^{(12)} = w_{33}^{(12)} &= \frac{1}{3} \\ w_{13}^{(13)} = w_{22}^{(13)} = w_{31}^{(13)} &= \frac{1}{3} & w_{11}^{(23)} = w_{23}^{(23)} = w_{32}^{(23)} &= \frac{1}{3} \\ w_{12}^{(123)} = w_{23}^{(123)} = w_{31}^{(123)} &= \frac{1}{3} & w_{13}^{(132)} = w_{21}^{(132)} = w_{32}^{(132)} &= \frac{1}{3} \end{aligned}$$

Now we see that $\pi'_1(v) = (v_{+1}, v_{+2}, v_{+3})$ and $\pi'_2(w^\sigma) = (w_{1+}^\sigma, w_{2+}^\sigma, w_{3+}^\sigma) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ are not equal to each other for any choice of the maximizer v of \mathcal{M}_{Γ_1} and w^σ of \mathcal{M}_{Γ_2} . This means

that $Q_v^{\Gamma_1}$ and $Q_w^{\Gamma_2}$ are not compatible. In other words, it is impossible to apply Corollary 4.3.13. Indeed, the bound cannot be attained as it was explicitly shown in [101, Example 20]. The maximum divergence is equal to $\log(3 + 2\sqrt{2}) < \log 6 = \log 2 + \log 3$. Up to symmetry there is a unique global maximizer given by

$$p_{111} = \sqrt{2} - 1, \quad p_{222} = p_{233} = 1 - \frac{\sqrt{2}}{2}.$$

We believe that for reducible models induced by $\Gamma = (\Gamma_1, S, \Gamma_2)$ finding *compatible* logarithmic Voronoi polytopes $Q_v^{\Gamma_1}$ and $Q_w^{\Gamma_2}$ that will give meaningful bounds for $D(\mathcal{M}_\Gamma)$ is worth exploring.

We close our discussion of reducible hierarchical log-linear models with a result involving conditional independence. For this we consider three random variables X_1, X_2 , and X_3 with state spaces $[d_1], [d_2]$, and $[d_3]$, respectively. We let $\mathbf{d} = (d_1, d_2, d_3)$. The simplicial complex we will use is $\Gamma = [12][23]$ with the decomposition $([12], \{2\}, [23])$. The toric model $\mathcal{M}_{\Gamma, \mathbf{d}}$ consists of the joint probability distributions where $X_1 \perp\!\!\!\perp X_3 \mid X_2$.

Proposition 4.3.18. Let Γ be the reducible simplicial complex on $[3]$ with decomposition $([12], \{2\}, [23])$ and $\mathbf{d} = (d_1, d_2, d_3) \in \mathbb{N}^3$. Then $D(\mathcal{M}_{\Gamma, \mathbf{d}}) = \min(\log d_1, \log d_3)$.

Proof. Proposition 4.3.10 implies that $D(\mathcal{M}_{\Gamma, \mathbf{d}}) \leq \min(\log d_1, \log d_3)$. The 0/1 matrix $A_{\Gamma, \mathbf{d}}$ defining the model can be organized as follows. Recall that this model is defined by the parametrization $p_{ijk} = a_{ij}b_{jk}$ with $i \in [d_1], j \in [d_2]$, and $k \in [d_3]$. We order the indices (i, j, k) lexicographically as follows: $(i, j, k) < (i', j', k')$ if $j < j'$, or if $j = j'$ and $i < i'$, or if $j = j'$ and $i = i'$ and $k < k'$. We sort the columns of $A_{\Gamma, \mathbf{d}}$ with respect to this ordering. We will also sort the rows of the matrix into d_2 blocks where in block j we list first the rows corresponding to the parameters a_{ij} with $i = 1, \dots, d_1$ and then the parameters b_{jk} with $k = 1, \dots, d_3$. Then

$$A_{\Gamma, \mathbf{d}} = \begin{pmatrix} A_{d_1, d_3} & 0 & \cdots & 0 \\ 0 & A_{d_1, d_3} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_{d_1, d_3} \end{pmatrix}$$

where A_{d_1, d_3} is the matrix defining an independence model of two random variables with state spaces $[d_1]$ and $[d_3]$. Since the maximum divergence from such a model is $\min(\log d_1, \log d_3)$, and each block gives the same maximum divergence, Theorem 4.1.8 implies the result. \square

4.4 Models of ML degree one

I start by introducing the following version of the definition of ML degree, specifically for toric models.

Definition 4.4.1. The maximum likelihood degree (*ML degree*) of a toric model $\mathcal{M}_A \subset \Delta_{n-1}$ is the number of points in \mathbb{C}^n that are in the intersection of the toric variety X_A and the affine span of the logarithmic Voronoi polytope Q_b , where $b = Au$ for a generic $u \in \Delta_{n-1}$.

The definition above is equivalent to the standard Definition 1.1.8. Moreover, a model has ML degree one if and only if the maximum likelihood estimate of any data $u \in \Delta_{n-1}$ can be expressed as a rational function of the coordinates of u [67]. We study the maximum divergence to such models in this section. Two dimensional models (toric surfaces) of ML degree one were classified in [37] along with some families of three dimensional models. We treat these families and the generalizations of some of them.

Multinomial distributions

We start by considering m independent identically distributed $(d + 1)$ -ary random variables X_1, \dots, X_m with state spaces $\{0, \dots, d\}$. Let s_j be the probability of state j , and let $p_{i_0 \dots i_d}$ denote the probability of observing exactly i_j occurrences of state j for each $j \in \{0, 1, \dots, d\}$. Thus, $p_{i_0 \dots i_d} = \binom{m}{i_0, \dots, i_d} s_0^{i_0} \dots s_d^{i_d}$ and we have the d -dimensional toric model parametrized as

$$\varphi : \Delta_d \rightarrow \Delta_{n-1} : (s_0, \dots, s_d) \mapsto (p_{i_0 \dots i_d} : \sum i_j = m)$$

where $n = \binom{m+d}{d}$. We refer to the the Zariski closure of the image of φ as the *twisted Veronese model* and denote it by $\mathcal{V}_{d,m}$. The columns of the matrix A corresponding to the parametrization are the nonnegative integer solutions to $i_0 + \dots + i_d = m$. Note that A has the constant vector $n\mathbf{1}$ in its rowspan. Moreover, it has the same rowspan as the matrix A' whose columns are of the form $(1, v)$ where $v \in \mathbb{R}^d$ is a nonnegative integer solution to the inequality $i_1 + \dots + i_d \leq m$. Thus, geometrically, the model is given by all lattice points in the convex hull of $\{0, me_1, \dots, me_d\} \subseteq \mathbb{R}^d$, a d -dimensional simplex dilated by a factor of m . Hence $\mathcal{V}_{d,m}$ is isomorphic to the Veronese variety, except the weights are modified so that $\mathcal{V}_{d,m}$ has ML degree one. That is, if we were to change all multinomial coefficients in the definition of $p_{i_0 \dots i_d}$ to 1, we would recover the usual Veronese variety.

Since the ML degree of $\mathcal{V}_{d,m}$ is one, the MLE can be expressed as a rational function of the data. Fix $u \in \Delta_{n-1}$ and suppose u is in the logarithmic Voronoi polytope Q_p at some unknown $p = (p_1, \dots, p_n) \in \mathcal{V}_{d,m}$. Let $Ap = b$ where $b = (b_0, \dots, b_d)$ is the point corresponding to p in $\text{conv}(A)$. Then each coordinate of the MLE of u can be expressed as

$$q_{i_0, \dots, i_d} = \binom{m}{i_0, \dots, i_d} \left(\frac{b_0}{m}\right)^{i_0} \dots \left(\frac{b_d}{m}\right)^{i_d}. \tag{4.1}$$

Example 4.4.2 ($d = 2, m = 3$). Consider the twisted Veronese variety $\mathcal{V}_{2,3}$. As discussed above, the defining matrices A and A' can be written as either

$$A = \begin{pmatrix} 3 & 2 & 2 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 2 & 1 & 0 & 3 & 2 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 2 & 0 & 1 & 2 & 3 \end{pmatrix} \text{ or } A' = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 2 & 1 & 0 & 3 & 2 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 2 & 0 & 1 & 2 & 3 \end{pmatrix}.$$

In our computations, we will usually use the matrix A . The polytope associated to $\mathcal{V}_{2,3}$ is plotted in Figure 4.3 on the left.




Figure 4.3: Twisted Veronese models $\mathcal{V}_{2,3}$ and $\mathcal{V}_{3,2}$, respectively.

The maximum divergence from $\mathcal{V}_{2,3}$ is $2 \log 3$, achieved at the unique point $v \in \Delta_9$, uniformly supported on $3e_1, 3e_2$, and $3e_3$, i.e. $v_{300} = v_{030} = v_{003} = 1/3$ and all other coordinates of v are 0. Note that v is a vertex of the logarithmic Voronoi polytope Q_b corresponding to the centroid b of $\text{conv}(A)$, i.e. $Aq = b = (1, 1, 1)$. The point q can be computed using (4.1), so $q = (1/3)^3(1, 3, 3, 3, 6, 3, 1, 3, 3, 1)$. The maximum divergence is

$$\begin{aligned} D(v||q) &= v_{300} \log(v_{300}/q_{300}) + v_{030} \log(v_{030}/q_{030}) + v_{003} \log(v_{003}/q_{003}) \\ &= 1/3 \log(3^3/3) + 1/3 \log(3^3/3) + 1/3 \log(3^3/3) = 2 \log 3. \end{aligned}$$

Example 4.4.3 ($d = 3, m = 2$). For the twisted Veronese model $\mathcal{V}_{3,2}$, the maximum divergence is $\log 4$. It is achieved at 10 different vertices of the logarithmic Voronoi polytope at the point $q = 1/16(1, 2, 1, 2, 2, 1, 2, 2, 2, 1)$. Note that $Aq = b = (1/2, 1/2, 1/2, 1/2)$, which is again the centroid of $\text{conv}(A)$. One of such vertices is $v = (1/4, 0, 1/4, 0, 0, 0, 0, 0, 1/2, 0)$, so the divergence is $D(v||q) = 1/4 \log 4 + 1/4 \log 4 + 1/2 \log 4 = \log 4$. The polytope $\text{conv}(A)$ for this model is shown in Figure 4.3 on the right. Each of the 10 maximizers arises from one of the 10 permutations in S_4 of order at most two. This will follow from the proof of Theorem 4.4.4 in the next section.

The formula for the maximum divergence from a general model $\mathcal{V}_{d,m}$ as well as the full description of maximizers were given in [73]. We summarize these results in the theorem below. A more detailed discussion about the maximizers will be presented in the next section.

Theorem 4.4.4. [73, Theorem 1.1] The maximum divergence to $\mathcal{V}_{d,m}$ equals $(m - 1) \log(d + 1)$. It is achieved at some vertices of the unique logarithmic Voronoi polytope Q_b where $b = (\frac{m}{d+1}, \dots, \frac{m}{d+1})$. There is a unique vertex of this polytope maximizing divergence if and only if $m > 2$.

Box model

In this section we will consider a generalization of the twisted Veronese model. Suppose we have $k > 1$ groups of random variables, with a_i independent identically distributed $(d + 1)$ -ary random variables with state space $\{0, \dots, d\}$ in the i th group for $i \in [k]$. Let $s_{i\ell}$ be the probability of state ℓ in the group i and let $p_{(j_{10}\dots j_{1d}), \dots, (j_{k0}\dots j_{kd})}$ denote the probability of observing exactly $j_{i\ell}$ occurrences of state ℓ in the group i . Hence, $p_{(j_{10}\dots j_{1d}), \dots, (j_{k0}\dots j_{kd})} = \prod_{i=1}^k \binom{a_i}{j_{i0}\dots j_{id}} s_{i0}^{j_{i0}} \dots s_{id}^{j_{id}}$ and we have a kd -dimensional toric model parametrized as

$$\Delta_d \times \dots \times \Delta_d \rightarrow \Delta_{n-1} : ((s_{10}, \dots, s_{1d}), \dots, (s_{k0}, \dots, s_{kd})) \mapsto (p_{(j_{10}\dots j_{1d}), \dots, (j_{k0}\dots j_{kd})} : 0 \leq j_{i\ell} \leq a_i),$$

where $n = \prod_{i=1}^k \binom{a_i+d}{d}$. The columns of the corresponding matrix A are naturally identified with the nonnegative integer solutions to the linear system $j_{i0} + \dots + j_{id} = a_i$ for $i \in [k]$. We will refer to this model as the *box model* motivated by the shape of $\text{conv}(A)$ when $d = 1$. We denote these models by $\mathcal{B}_{a_1, \dots, a_k}^{(d)}$. The special case when $d = 1, k = 2$ was studied in [37]. The box model also has ML degree one, and hence the MLE can be written as a rational function of data. For $u \in \Delta_{n-1}$ such that $Au = b = ((b_{10}, \dots, b_{1a_1}), \dots, (b_{k0}, \dots, b_{ka_k}))$, the MLE is

$$q_{(j_{10}\dots j_{1d}), \dots, (j_{k0}\dots j_{kd})} = \prod_{i=1}^k \binom{a_i}{j_{i0}, \dots, j_{id}} \left(\frac{b_{i0}}{a_i}\right)^{j_{i0}} \dots \left(\frac{b_{id}}{a_i}\right)^{j_{id}}.$$

Theorem 4.4.5. The maximum divergence to the model $\mathcal{B}_{a_1, \dots, a_k}^{(d)}$ equals $(a_1 + \dots + a_k - 1) \log(d+1)$. It is achieved at $[(d+1)!]^{k-1}$ vertices of the unique logarithmic Voronoi polytope Q_b such that $b = ((\frac{a_1}{d+1}, \dots, \frac{a_1}{d+1}), \dots, (\frac{a_k}{d+1}, \dots, \frac{a_k}{d+1}))$.

Our proof of Theorem 4.4.5 relies heavily on the methods used to prove Theorem 4.4.4 in [73]. Before we present both proofs, we outline the general theory below.

Let \mathcal{F} be a model inside the simplex Δ_{N-1} . Let S_N be the symmetric group of all permutations on $[N]$. This group acts on Δ_{N-1} by permuting the coordinates of the points in the simplex. Let G be a subgroup of S_N .

Definition 4.4.6. The model \mathcal{F} is said to be *G-symmetrical* if for all $\sigma \in G$ and all $p \in \mathcal{F}$, we have $\sigma p \in \mathcal{F}$. A point $p \in \Delta_{N-1}$ is said to be *G-exchangeable* if for all $\sigma \in G$, we have $\sigma p = p$.

Let \mathcal{F} be a *G-symmetrical* model and let \mathcal{F}/G denote the set of all orbits of \mathcal{F} under the action of G . Let

$$\gamma_G : \mathcal{F} \rightarrow \mathcal{F}/G : p \mapsto \{\sigma p : \sigma \in G\}$$

be the map that sends an element in \mathcal{F} to its orbit. Denote by \mathcal{E} the closure of all *G-exchangeable* distributions in Δ_{N-1} and let $\mathcal{M} = \mathcal{F} \cap \mathcal{E}$ denote the induced model of all exchangeable distributions in \mathcal{F} . The following theorem holds in general.

Theorem 4.4.7. [73, Corollary 2.6] Let G be a subgroup of S_N and let $\mathcal{F} \subseteq \Delta_{N-1}$ be a G -symmetrical family of distributions. If there exists a maximizer of $D_{\mathcal{F}}$ that is exchangeable, then $D(\gamma_G(\mathcal{M})) = D(\mathcal{F})$ and all maximizers of $D_{\gamma_G(\mathcal{M})}$ are of the form $\gamma_G(v)$ where v is an exchangeable maximizer of $D_{\mathcal{F}}$.

Proof of Theorem 4.4.4 [73]. Let \mathcal{F} be the independence model of m $(d+1)$ -ary random variables induced by $\Delta_d \times \cdots \times \Delta_d$. Then $\mathcal{F} \subseteq \Delta_{N-1}$ where $N = (d+1)^m$ and it is given by the following parametrization

$$\varphi : ((x_{10}, \dots, x_{1d}), \dots, (x_{m0}, \dots, x_{md})) \mapsto (p_{j_1, \dots, j_m} = x_{1j_1} x_{2j_2} \cdots x_{mj_m} : j_1, \dots, j_m \in \{0, \dots, d\}). \quad (4.2)$$

Let G be the subgroup of S_N given by

$$G = \{\sigma_\rho \in S_N : \rho \in S_m \text{ and } \sigma_\rho((x_{10}, \dots, x_{1d}), \dots, (x_{m0}, \dots, x_{md})) = ((x_{\rho(1)0}, \dots, x_{\rho(1)d}), \dots, (x_{\rho(m)0}, \dots, x_{\rho(m)d}))\}.$$

Note that $G \cong S_m$ and it acts on each coordinate of $p \in \mathcal{F}$ as $\sigma p_{i_1 \dots i_m} = p_{i_{\sigma(1)} \dots i_{\sigma(m)}}$. Under this action, note that \mathcal{F} is G -symmetrical and that the set of all G -exchangeable distributions in \mathcal{F} is $\mathcal{M} = \{\varphi(x, x, \dots, x) : x \in \Delta_d\}$. Then the twisted Veronese model $\mathcal{V}_{d,m}$ can be identified with the set of all orbits of \mathcal{F} coming from exchangeable distributions, i.e. $\mathcal{V}_{d,m} = \gamma_G(\mathcal{M})$.

Since \mathcal{F} is an independence model, we know that $D(\mathcal{F}) = (m-1) \log(d+1)$ and all of its maximizers are given in Theorem 4.3.15. Denote each such maximizer by $v_{\sigma_2, \dots, \sigma_m} := \frac{1}{d+1} \sum_{j=0}^d \delta_{j, \sigma_2(j), \dots, \sigma_m(j)}$. Note that $v = v_{\text{id}, \dots, \text{id}}$ is the distribution in \mathcal{F} such that $v_{jj \dots j} = \frac{1}{d+1}$ for each $j \in \{0, \dots, d\}$ and 0 otherwise. By Theorem 4.4.7, it then follows that $D(\mathcal{V}_{d,m}) = D(\gamma_G(\mathcal{M})) = D(\mathcal{F}) = (m-1) \log(d+1)$, as desired. Moreover, $w = \gamma_G(v)$ is a maximizer of $D_{\mathcal{V}_{d,m}}$. Explicitly, it is given as $w_x = \frac{1}{d+1}$ if $x = m e_j$ for $j \in \{0, \dots, d\}$ and 0 otherwise.

If $m > 2$, we claim that w is the unique maximizer. Indeed, let $v = v_{\sigma_2, \dots, \sigma_m}$ be another exchangeable maximizer of \mathcal{F} . Without loss of generality, assume $\sigma_2 \neq \text{id}$, so there is some j such that $\sigma_2(j) \neq j$. Since $(j, \sigma_2(j), j_3, \dots, j_m) \in \text{supp}(v)$ for some $j_3, \dots, j_m \in \{0, \dots, d\}$, it has to be the case that $(\sigma_2(j), j, j_3, \dots, j_m) \in \text{supp}(v)$, since v is exchangeable and has to have the same value in both coordinates. But since $m > 2$, it then follows that $j_k = \sigma_k(j) = \sigma_k(\sigma_2(j))$, so σ_k is not injective for every $k \geq 3$, a contradiction.

If $m = 2$, then let $\sigma = \sigma_2$ and note that if $\sigma^2 \neq \text{id}$, then $v_{jk} \in \text{supp}(v)$, but $v_{jk} = 0$ for some $j, k \in \{0, \dots, d\}$, which would contradict exchangeability of v . Hence, every maximizer of $\mathcal{V}_{d,2}$ is of the form $w = \frac{1}{d+1} \sum_{j=0}^d \delta_{e_j} + \delta_{e_{\sigma(j)}}$ for some $\sigma \in S_{d+1}$ of order at most two. Every nonzero coordinate of w is thus either $\frac{1}{d+1}$ or $\frac{2}{d+1}$. The number of maximizers in this case is the number of permutations in S_{d+1} of order at most two. Note that for each of the maximizers, we have $Av = \left(\frac{m}{d+1}, \dots, \frac{m}{d+1}\right)$ and hence they all lie in the same logarithmic Voronoi polytope at $q \in \mathcal{V}_{d,m}$, corresponding to the centroid of $\text{conv}(A)$. \square

Proof of Theorem 4.4.5. Let \mathcal{F} be the independence model of $a_1 + \dots + a_k$ $(d+1)$ -ary random variables divided into k groups, induced by

$$\underbrace{(\Delta_d \times \dots \times \Delta_d)}_{a_1} \times \dots \times \underbrace{(\Delta_d \times \dots \times \Delta_d)}_{a_k}.$$

Then $\mathcal{F} \subseteq \Delta_{N-1}$ where $N = (d+1)^{a_1 + \dots + a_k}$ and has the parametrization φ like (4.2), except each probability p_\bullet factors as a product of $a_1 + \dots + a_k$ parameters. Let G be a subgroup of S_N defined as

$$G = \{ \sigma_{\rho_1, \dots, \rho_k} \in S_N : \rho_i \in S_{a_i} \text{ and } \sigma_{\rho_1, \dots, \rho_k}((y_1^{(1)}, \dots, y_{a_1}^{(1)}), \dots, (y_1^{(k)}, \dots, y_{a_k}^{(k)})) = ((y_{\rho_1(1)}^{(1)}, \dots, y_{\rho_1(a_1)}^{(1)}), \dots, (y_{\rho_k(1)}^{(k)}, \dots, y_{\rho_k(a_k)}^{(k)})) \},$$

where $y_j^{(i)} = (x_{j_1}^{(i)}, \dots, x_{j_d}^{(i)}) \in \Delta_d$. Note that $G \cong S_{a_1} \times \dots \times S_{a_k}$.

Under this action, \mathcal{F} is G -symmetrical and the set of all G -exchangeable distributions in \mathcal{F} is $\mathcal{M} = \{ \varphi((x^{(1)}, \dots, x^{(1)}), \dots, (x^{(k)}, \dots, x^{(k)})) : x^{(i)} \in \Delta_d \text{ for all } i \in [k] \}$. The box model $\mathcal{B}_{a_1, \dots, a_k}^{(d)}$ is then identified with the set of all orbits of \mathcal{F} coming from exchangeable distributions, i.e. $\mathcal{B}_{a_1, \dots, a_k}^{(d)} = \gamma_G(\mathcal{M})$.

Since \mathcal{F} is again an independence model, we know that $D(\mathcal{F}) = (a_1 + \dots + a_k - 1) \log(d+1)$ from Theorem 4.3.15. Denote each maximizer by

$$v_{\sigma_2^{(1)}, \dots, \sigma_{a_1}^{(1)}, \dots, \sigma_1^{(k)}, \dots, \sigma_{a_k}^{(k)}} := \frac{1}{d+1} \sum_{j=0}^d \delta_{j, \sigma_2^{(1)}(j), \dots, \sigma_{a_1}^{(1)}(j), \dots, \sigma_1^{(k)}(j), \dots, \sigma_{a_k}^{(k)}(j)}.$$

First let $\sigma_2^{(1)} = \dots = \sigma_{a_1}^{(1)} = \pi_1 = \text{id}$ and $\sigma_1^{(i)} = \dots = \sigma_{a_i}^{(i)} = \pi_i$ for some $\pi_i \in S_{d+1}$ for all $i > 1$. Then v is a G -exchangeable maximizer of \mathcal{F} , and is explicitly given as $v_{(j \dots j), (\pi_2(j) \dots \pi_2(j)), \dots, (\pi_k(j) \dots \pi_k(j))} = \frac{1}{d+1}$ for any choice of $j \in \{0, \dots, d\}$ and 0 otherwise. The image of this maximizer under γ is then $w = \frac{1}{d+1} \sum_{j=0}^d \delta_{\times_{i=1}^k a_i e_{\pi_i(j)}}$, where \times denotes the Cartesian product of vectors in \mathbb{R}^{d+1} . By

Theorem 4.4.7, w is a maximizer of $\mathcal{B}_{a_1, \dots, a_k}^{(d)}$. There are exactly $[(d+1)]^{k-1}$ such maximizers: one for every choice of $(\pi_2, \dots, \pi_k) \in S_{d+1} \times \dots \times S_{d+1}$. Note also that for each such maximizer w , we have $Aw = ((\frac{a_1}{d+1}, \dots, \frac{a_1}{d+1}), \dots, (\frac{a_k}{d+1}, \dots, \frac{a_k}{d+1}))$, so all of them are the vertices of the logarithmic Voronoi polytope at the point $q \in \mathcal{B}_{a_1, \dots, a_k}^{(d)}$ corresponding to the centroid of $\text{conv}(A)$.

We claim that there are no other maximizers of $\mathcal{B}_{a_1, \dots, a_k}^{(d)}$. Indeed, if $a_i > 2$ for all $i \in [k]$, then there are no other G -exchangeable maximizers of \mathcal{F} by the proof of Theorem 4.4.4. Indeed, if $a_1 = a_2 = 1$ and $k = 2$, then all maximizers are of the form discussed in the previous paragraph. If $a_i = 2$ for some $i \in [k]$, without loss of generality assume that $a_1 = 2$ and that v is a G -exchangeable maximizer of \mathcal{F} with $\pi(j) = \sigma_2^{(1)}(j) \neq j$ for some j . If v is G -exchangeable, it has to be the case that there are some values $j_3, \dots, j_{a_1 + \dots + a_k}$ such that both $(j, \pi(j), j_3, \dots, j_{a_1 + \dots + a_k})$ and $(\pi(j), j, j_3, \dots, j_{a_1 + \dots + a_k})$ are in $\text{supp}(v)$. But then $j_3 = \sigma_1^{(2)}(j) = \sigma_1^{(2)}(\pi(j))$, a contradiction to the injectivity of π . Hence, there are no other maximizers. \square

When $d = 1$, the maximizers of the box model $\mathcal{B}_{a_1, \dots, a_k}^{(1)}$ have the following nice geometric interpretation.

Corollary 4.4.8. The maximum divergence from the box model $\mathcal{B}_{a_1, \dots, a_k}^{(1)}$ equals $(a_1 + \dots + a_k - 1) \log 2$. The maximum divergence is achieved at 2^{k-1} vertices of the unique logarithmic Voronoi polytope. These vertices correspond to the main diagonals of $\text{conv}(A)$.

Example 4.4.9 ($d = 1, k = 3$). Consider the box model $\mathcal{B}_{3,3,2}^{(1)}$. It is a 3-dimensional model inside Δ_{47} . The columns of the corresponding matrix A can be identified with the lattice points $\{(i, j, k) \in \mathbb{Z}^3 : 0 \leq i, j \leq 3, 0 \leq k \leq 2\}$. The maximum divergence of this model is $7 \log 2$ and it is achieved at four vertices of the logarithmic Voronoi polytope Q_b corresponding to the point $b = (3/2, 3/2, 1)$ in $\text{conv}(A)$. This is illustrated in Figure 4.4.




Figure 4.4: Chamber complex of the box model $\mathcal{B}_{3,3,2}^{(1)}$.

The four vertices giving $D(\mathcal{B}_{3,3,2}^{(1)})$ are supported on the main diagonals of $\text{conv}(A)$. Explicitly, the four maximizers are

$$w_1 = \frac{1}{2} (\delta_{(0,0,0)} + \delta_{(3,3,2)}), \quad w_2 = \frac{1}{2} (\delta_{(0,0,2)} + \delta_{(3,3,0)}),$$

$$w_3 = \frac{1}{2} (\delta_{(0,3,0)} + \delta_{(3,0,2)}), \quad w_4 = \frac{1}{2} (\delta_{(3,0,0)} + \delta_{(0,3,2)}).$$

Trapezoid model

One interesting extension of the box model $\mathcal{B}_{a_1, a_2}^{(1)}$ is the trapezoid model, which we discuss in this section. It also has ML degree one [37]. Fix some positive integers a, b, d . Suppose we have two coins, with the probabilities of flipping heads being s and t , respectively. First, we flip the second coin b times, and record the number of heads $j \in \{0, \dots, b\}$. Then we flip the

first coin a times and record the number of heads, and then flip the first coin again $d(b-j)$ times and record the number of heads. The probability $p_{r,j}$ of getting exactly r heads from the first coin and exactly j heads from the second coin is then

$$p_{r,j} = c_{r,j} s^r (1-s)^{a+d(b-j)-r} t^j (1-t)^{b-j} \quad (4.3)$$

where

$$c_{r,j} = \binom{b}{j} \sum_{\substack{0 \leq i \leq a \\ 0 \leq k \leq d(b-j) \\ i+k=r}} \binom{a}{i} \binom{d(b-j)}{k}.$$

Geometrically, this model is given by all lattice points inside the trapezoid with the vertices $\{(0,0), (0,b), (a,b), (a+db,0)\}$ with the weight of each point (r,j) given by $c_{r,j}$. Hence we will call this model the *trapezoid model* and denote it by $\mathcal{T}_{a,b,d}$. This is a 2-dimensional model inside Δ_{n-1} , parametrized by $(s,t) \mapsto (p_{r,j} : 0 \leq j \leq b, 0 \leq r \leq a+d(b-j))$ where $n = \sum_{j=0}^b \sum_{r=0}^{a+d(b-j)} r$.

The MLE for any data point $u \in \Delta_n$ is a rational function of u . If $Au = (1, b_1, b_2)$, the MLE of u is the point $q \in \mathcal{T}_{a,b,d}$ such that $Aq = (1, b_1, b_2)$. This point is given as

$$q_{r,j} = c_{r,j} \left(\frac{b_1}{a+d(b-b_2)} \right)^r \left(1 - \frac{b_1}{a+d(b-b_2)} \right)^{a+d(b-j)-r} \left(\frac{b_2}{b} \right)^j \left(1 - \frac{b_2}{b} \right)^{b-j}.$$

Example 4.4.10 ($a = b = d = 1$). Consider the simplest nontrivial trapezoid model with $a = b = d = 1$. It is a 2-dimensional toric model in Δ_4 where

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 2 \\ 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$

and the chamber complex is shown in the middle of Figure 4.5. Note that two-dimensional chambers will not contribute any projection points by Proposition 4.2.11. There are only three interior vertices: $(1/2, 1/2)$, $(2/3, 2/3)$ and $(1, 1/2)$. All of them have triangles for logarithmic Voronoi polytopes, with supports $\{14, 25, 234\}$, $\{23, 14, 125\}$, and $\{34, 25, 145\}$, respectively. Therefore, all potential projection points will come from the edges. Running our algorithm on the ten interior edges, we find that there are exactly four projection vertices, corresponding to two different points on the model $\mathcal{T}_{1,1,1}$. The first point maps to $b_1 = (2\sqrt{5}/5, -\sqrt{5}/5 + 1) \in \text{conv}(A)$, and the two projection vertices are

$$\left(\sqrt{5}/5, 0, 0, -\sqrt{5}/5 + 1, 0 \right) \text{ and } \left(0, -\sqrt{5}/5 + 1, 3\sqrt{5}/5 - 1, 0, -2\sqrt{5}/5 + 1 \right).$$

The latter vertex yields the maximum divergence $\log 2 + \log \left(\frac{1}{3-\sqrt{5}} \right)$. Similarly, the second point on the model maps to $b_2 = (-\sqrt{5}/5 + 1, -\sqrt{5}/5 + 1)$, and the two projection vertices are

$$\left(0, \sqrt{5}/5 + 1, 0, 0, \sqrt{5}/5 \right) \text{ and } \left(-2\sqrt{5}/5 + 1, 0, 3\sqrt{5}/5 - 1, -\sqrt{5}/5 + 1, 0 \right).$$

The former vertex yields the same maximum divergence $\log 2 + \log\left(\frac{1}{3-\sqrt{5}}\right)$.

The logarithmic Voronoi polytopes corresponding to both b_1 and b_2 are quadrilaterals. After projecting onto two-dimensional planes, we plot both polytopes in Figure 4.5.




Figure 4.5: The chamber complex and two logarithmic Voronoi polytopes that yield maximum divergence.

Theorem 4.4.11. The divergence to the trapezoid model $\mathcal{T}_{a,b,d}$ is bounded above by $(a + bd + b) \log 2$.

Proof. Fix a model $\mathcal{T}_{a,b,d}$ in Δ_{n-1} , parametrized by s and t . Let $u = (u_{r,j}) \in \Delta_{n-1}$ be a general data vector. Then the log-likelihood function is

$$\begin{aligned} \ell_u(p) = \sum u_{r,j} \log(p_{r,j}) &= \sum u_{r,j} \log c_{r,j} + \sum u_{r,j} [r \log s + (a + d(b-j) - r) \log(1-s)] \\ &\quad + \sum u_{r,j} [j \log(t) + (b-j) \log(1-t)]. \end{aligned}$$

Taking the partial derivatives and solving for the parameters, we get that $\hat{s} = \frac{\sum u_{r,j} r}{\sum u_{r,j} (a + d(b-j))}$ and $\hat{t} = \frac{\sum u_{r,j} j}{b}$. The MLE q is obtained by plugging in these parameters into (4.3). Hence, the divergence function from the general point u to the model $\mathcal{T}_{a,b,d}$ is

$$\begin{aligned} D(u || \mathcal{T}_{a,b,d}) &= \sum u_{r,j} \log(u_{r,j}/q_{r,j}) = \underbrace{-H(u) - \sum u_{r,j} \log(c_{r,j})}_{\leq 0} \\ &\quad - \sum u_{r,j} r \log(\hat{s}) - \sum u_{r,j} (a + (b-j) - r) \log(1 - \hat{s}) \\ &\quad - \sum u_{r,j} j \log(\hat{t}) - \sum u_{r,j} (b-j) \log(1 - \hat{t}), \end{aligned}$$

where $H(u) = -\sum_{i=1}^n u_i \log(u_i)$ is the entropy. Let $h(p)$ denote the entropy of a binary random variable with the probability of success p , i.e. $h(p) = -p \log(p) - (1-p) \log(1-p)$.

Note that $h(p)$ always attains its maximum value at $p = 1/2$. Therefore, we have








$$\begin{aligned} D(u||\mathcal{T}_{a,b,d}) &\leq \left(\sum u_{r,j}(a + d(b - j)) \right) h(\hat{s}) + bh(\hat{t}) \\ &\leq \left(\sum u_{r,j}(a + d(b - j) + b) \right) h(1/2) \\ &= (a + db + b) \log 2 - d \sum u_{r,j}j \\ &\leq (a + db + b) \log 2, \end{aligned}$$

as desired. □

Note that the polytope of the trapezoid model $\mathcal{T}_{a,b,d}$ sits in-between the polytopes of the two box models $\mathcal{B}_{a,b}^{(1)}$ and $\mathcal{B}_{a+db,b}^{(1)}$. However, the weights assigned to the lattice points are different. This presents the question of whether $D(\mathcal{T}_{a,b,d})$ is bounded by $D(\mathcal{B}_{a,b}^{(1)})$ and $D(\mathcal{B}_{a+db,b}^{(1)})$ below and above, respectively. Note that this is indeed the case for Example 4.4.10. We present the following conjecture.

Conjecture 4.4.12. The divergence from the trapezoid model $\mathcal{T}_{a,b,d}$ is at least $(a+b-1) \log 2$ and at most $(a + bd + b - 1) \log 2$. This upper bound is sharp if and only if $d = 0$.

In [37], the authors present several families of 3-dimensional models that have ML degree one. We compute the maximum divergence for the simplest nontrivial examples in those families in the table below.

polytope	$\text{conv}(A)$	$D(\mathcal{M})$	notes
	$\text{conv}\{(0,0,0), (0,0,1), (0,1,0), (0,1,1), (1,1,0), (2,0,0)\}$	$2 \log 2$	conjectured
	$\text{conv}\{(0,0,0), (0,0,1), (0,1,0), (1,0,1), (1,1,0), (2,0,0)\}$	$\log 2 + \log\left(\frac{1}{3-\sqrt{5}}\right)$	
	$\text{conv}\{(0,0,0), (0,0,1), (0,1,0), (1,1,0), (2,0,0)\}$	$\log 2 + \log\left(\frac{1}{3-\sqrt{5}}\right)$	boundary
	$\text{conv}\{(0,0,0), (0,0,1), (0,1,0), (0,1,1), (2,0,0), (2,1,0)\}$	$\log 3$	conjectured
	$\text{conv}\{(0,0,0), (0,0,1), (0,1,0), (1,0,0), (1,1,0)\}$	$\log 2$	boundary
	$\text{conv}\{(0,0,0), (0,0,1), (0,1,0), (1,0,0), (0,0,1), (1,0,1), (1,1,0), (1,1,1)\}$	$2 \log 2$	box model
	$\text{conv}\{(0,0,0), (0,0,1), (0,1,0), (1,0,0), (0,0,1), (1,0,1)\}$	$\log 2$	2×3 independence

Interestingly, the second example in the table has infinitely many maximizers. For the third and the fifth models, all the maximizers of information divergence lie on the boundary of the simplex. For the conjectured examples, we were able to compute most of the ideals in step 3 of the algorithm, but not all. Some of the higher-dimensional ideals that arise in those cases are very complicated and we were not able to solve them using numerical tools.

Conclusion. In this chapter, we revisited the problem of maximizing information divergence to toric models from the new perspective using logarithmic Voronoi polytopes. We presented an algorithm for locating maximizers that combines the combinatorics of the chamber complex with numerical algebraic geometry. We also paid special attention to reducible models and models of maximum likelihood degree one. In particular, we provided a way to reconstruct logarithmic Voronoi polytopes of a reducible model \mathcal{M} from the logarithmic Voronoi polytopes of the models induced by the the reduction of \mathcal{M} . We then explained how to use this decomposition to obtain and bound information divergence to reducible models. We also established the maximum divergence and characterized the set of maximizers for the box model. Finally, we established an upper bound for divergence to the trapezoid model.

Chapter 5

Gaussian setting

In this chapter, we study logarithmic Voronoi cells for different families of Gaussian models. We start by proving that logarithmic Voronoi cells are always convex sets, like in the discrete case. Then we prove several basic results about when logarithmic Voronoi cells coincide with the log-normal spectrahedra, in which they are contained. We then investigate the geometry and combinatorics of logarithmic Voronoi cells for linear concentration models in Section 5.1, directed graphical models in Section 5.2, and covariance models in Section 5.3. This chapter is based on [5].

Proposition 5.0.1. For a Gaussian model $\Theta \subseteq \text{PD}_m$ and $\Sigma \in \Theta$, the logarithmic Voronoi cell $\log \text{Vor}_\Theta \Sigma$ is a convex set.

Proof. The logarithmic Voronoi cell at Σ is

$$\log \text{Vor}_\Theta(\Sigma) = \{S \in \text{PD}_m : \ell_n(\Sigma, S) \geq \ell_n(\Sigma', S) \text{ for all } \Sigma' \in \Theta\}.$$

Since $\ell_n(\Sigma, S)$ is linear in S , each inequality $\ell_n(\Sigma, S) \geq \ell_n(\Sigma', S)$ defines a closed halfspace. Therefore the logarithmic Voronoi cell at Σ is the intersection of these halfspaces for each $\Sigma' \in \Theta$ and the convex cone PD_m . \square

In Proposition 1.2.9, we saw that log-Voronoi cells are contained in log-normal spectrahedra. More precisely, we may write

$$\log \text{Vor}_\Theta \Sigma = \{S \in \mathcal{K}_\Theta \Sigma : \ell_n(\Sigma, S) \geq \ell_n(\Sigma', S) \text{ for all critical points } \Sigma'\}. \quad (5.1)$$

Example 1.2.8 illustrated that the reverse containment doesn't hold in general. However, the two convex sets are equal if the log-likelihood function has a unique optimum on the model Θ , and more strongly, if the maximum likelihood degree of the Gaussian model is one. We get the following two corollaries.

Corollary 5.0.2. If $\ell_n(\Sigma, S)$ has a unique maximum Σ over the Gaussian model $\Theta \subseteq \text{PD}_m$, then $\log \text{Vor}_\Theta(\Sigma) = \mathcal{K}_\Theta(\Sigma)$.

Proof. Since Σ is the unique maximum the inequalities in (5.1) are superfluous. \square

Corollary 5.0.3. If the ML degree of a Gaussian model $\Theta \subseteq \text{PD}_m$ is one then $\log \text{Vor}_\Theta(\Sigma) = \mathcal{K}_\Theta(\Sigma)$ for every $\Sigma \in \Theta$.

Proof. Since Σ is the unique critical point, the result follows from Corollary 5.0.2. \square

5.1 Linear concentration models

In a multivariate Gaussian distribution, the inverse of the covariance matrix $K = \Sigma^{-1}$ is known as the *concentration matrix*. Linear concentration models [14] are given by concentration matrices which form a linear subspace. Let \mathcal{L} be a d -dimensional linear subspace of $m \times m$ real symmetric matrices. Then a *linear concentration model* is given by

$$\Theta = \{\Sigma \in \text{PD}_m : K = \Sigma^{-1} \in \mathcal{L}\}.$$

The log-likelihood function equals

$$-\frac{n}{2} \log \det K - \frac{n}{2} \text{tr}(SK),$$

and it is a strictly concave function on $\mathcal{L} \cap \text{PD}_m$. If K_1, \dots, K_d are a basis of \mathcal{L} and S is a sample covariance matrix, the maximizer $\hat{\Sigma} = \hat{K}^{-1}$ of the log-likelihood function is the unique solution to

$$\text{tr}(\hat{\Sigma}K_j) = \text{tr}(SK_j), \quad j = 1, \dots, d. \quad (5.2)$$

This follows from writing $K = \sum_{i=1}^d \lambda_i K_i$ and taking partial derivatives of the log-likelihood function with respect to λ_j , $j = 1, \dots, d$; see [112]. Therefore, we immediately get the following.

Proposition 5.1.1. Let Θ be a linear concentration model given by $\mathcal{L} = \text{span}\{K_1, \dots, K_d\}$, and let $\Sigma \in \Theta$. Then

$$\log \text{Vor}_\Theta(\Sigma) = \mathcal{K}_\Theta(\Sigma) = \{S \in \text{PD}_m : \text{tr}(SK_j) = \text{tr}(\Sigma K_j), \quad j = 1, \dots, d\}.$$

Proof. The equality of the logarithmic Voronoi cell and the log-normal spectrahedron follows from Corollary 5.0.2. The linear description of these convex sets follows from (5.2). \square

Corollary 5.1.2. Let Θ be a one-dimensional linear concentration model spanned by $K \in \text{PD}_m$. For $\lambda > 0$ and $\Sigma = \frac{1}{\lambda}K^{-1}$, the logarithmic Voronoi cell at Σ is $\log \text{Vor}_\Theta(\Sigma) = \{S \in \text{PD}_m : \text{tr}(SK) = \frac{m}{\lambda}\}$. Therefore, it is the intersection of a translate of \mathcal{L}^\perp with PD_m where $\mathcal{L} = \text{span}(K)$.

Proof. Since $\text{tr}(\Sigma K) = \frac{m}{\lambda}$, the result follows Proposition 5.1.1. \square

Corollary 5.1.3. When $m = 2$, the logarithmic Voronoi cells of one-dimensional concentration models are convex regions defined by ellipses.

Proof. Let $K = \begin{pmatrix} a & b \\ b & c \end{pmatrix} \succ 0$, $\lambda > 0$, and $\Sigma = \frac{1}{\lambda}K^{-1}$. With $S = \begin{pmatrix} s_{11} & s_{12} \\ s_{12} & s_{22} \end{pmatrix}$, from (5.2) we get $as_{11} + 2bs_{12} + cs_{22} = \frac{2}{\lambda}$. Then $\log \text{Vor}_{\Theta}(\Sigma)$ is defined by the inequalities

$$\frac{1}{a} \left(\frac{2}{\lambda} - 2bs_{12} - cs_{22} \right) s_{22} - s_{12}^2 \geq 0, \quad \frac{1}{a} \left(\frac{2}{\lambda} - 2bs_{12} - cs_{22} \right) \geq 0, \quad s_{22} \geq 0.$$

The quadric defines an ellipse since its Δ -invariant and δ -invariant [57, Section 5.2] are $\frac{\lambda^2}{a^2} \neq 0$ and $\frac{ac-b^2}{a^2} > 0$, respectively. Finally, the nonnegativity of the quadric implies the other inequalities for positive definite S . \square

We would like to point out that, despite the concavity of $\ell_n(\Sigma, S)$ on a linear concentration model Θ , the maximum likelihood degree of Θ is much bigger than one. This was first studied in [112] which included conjectures on the ML degree of such models. Most of these conjectures were settled in [82] and [86]. See also [10] and [71] for related work.

Undirected graphical models

Let $G = (V, E)$ be a simple undirected graph with $|V(G)| = m$. A *concentration model* of G is

$$\Theta(G) = \{ \Sigma \in \text{PD}_m : (\Sigma)_{ij}^{-1} = 0 \text{ if } ij \notin E(G) \text{ and } i \neq j \}.$$

Concentration models of undirected graphs are examples of linear concentration models. Thus, their logarithmic Voronoi cells are equal to the log-normal spectrahedra. Following Proposition 5.1.1, we can describe logarithmic Voronoi cells explicitly as

$$\log \text{Vor}_{\Theta(G)}(\Sigma) = \{ S \in \text{PD}_m : \Sigma_{ij} = S_{ij} \text{ for all } ij \in E(G) \text{ and } i = j \}.$$

Example 5.1.4. Consider the graphical model associated to the undirected path $1-2-3-4$ on four vertices. This model is

$$\Theta(G) = \{ \Sigma \in \text{PD}_4 : (\Sigma^{-1})_{13} = (\Sigma^{-1})_{14} = (\Sigma^{-1})_{24} = 0 \}.$$

Let $\Sigma' = \begin{pmatrix} 6 & 1 & 1/7 & 1/28 \\ 1 & 7 & 1 & 1/4 \\ 1/7 & 1 & 8 & 2 \\ 1/28 & 1/4 & 2 & 9 \end{pmatrix}$. Then the logarithmic Voronoi cell at Σ' is

$$\log \text{Vor}_{\Theta(G)}(\Sigma') = \left\{ (x, y, z) : M_{x,y,z} = \begin{pmatrix} 6 & 1 & x & y \\ 1 & 7 & 1 & z \\ x & 1 & 8 & 2 \\ y & z & 2 & 9 \end{pmatrix} \succ 0 \right\}.$$

We plot the algebraic boundary of this spectrahedron on the left of Figure 5.1. It is defined by the quartic $\det(M_{x,y,z})$. The right figure is the spectrahedron itself where “ears” are removed by the quadric that is the third principal minor of $M_{x,y,z}$.




Figure 5.1: Pillow-shaped logarithmic Voronoi spectrahedra

The boundary of $\log \text{Vor}_{\Theta(G)}(\Sigma')$ consists of matrices of rank at most three. This spectrahedron has four singular points that have rank two. Computations reveal that for *any* matrix Σ on the model $\Theta(G)$, the general combinatorial type of the logarithmic Voronoi spectrahedron at Σ will be the same as that of Σ' ; informally, $\log \text{Vor}_{\Theta(G)}(\Sigma)$ is pillow-shaped for all $\Sigma \in \Theta(G)$. An interesting problem would be to study the logarithmic Voronoi cells of concentration models and how their combinatorial type changes as points on the model vary. In the discrete setting, it was done for linear models [1]. In the Gaussian setting, combinatorial types of spectrahedra can be described using *patches*; see [33, 98].

Decomposition of logarithmic Voronoi cells

In the theory and practice of graphical models, *reducible* and *decomposable* models play a significant role [78, 114]. They provide a recursive structure that can be exploited, for instance, in maximum likelihood estimation. In this subsection, we develop a decomposition theory of the logarithmic Voronoi cells for such models.

Let $G = (V, E)$ be an undirected graph with the vertex set labeled by $[m]$. A clique of G is a subset $C \subseteq [m]$ such that $ij \in E(G)$ for every $i, j \in C$. We say that a clique in G is *maximal* if the subgraph it induces does not embed into a larger clique of G . Let $\mathcal{C}(G)$ denote the

set of all cliques of G . Note that $\mathcal{C}(G)$ is a simplicial complex on $[m]$, whose facets are the maximal cliques of G .

A simplicial complex $\Gamma \subseteq [m]$ is called *reducible* with decomposition (Γ_1, T, Γ_2) if there exist sub-complexes Γ_1, Γ_2 of Γ and a subset $T \subseteq [m]$ such that $\Gamma = \Gamma_1 \cup \Gamma_2$ and $\Gamma_1 \cap \Gamma_2 = 2^T$. Moreover, we assume that $\Gamma_i \neq 2^T$ for $i = 1, 2$. We say Γ is *decomposable* if it is reducible and each of the Γ_1, Γ_2 is either decomposable or a simplex. A graphical model associated to an undirected graph G is reducible (resp. decomposable) if its complex of cliques $\mathcal{C}(G)$ is reducible (resp. decomposable).

Given a graph G on m vertices with the complex of cliques $\mathcal{C}(G)$, let $\Theta(G)$ denote the associated graphical model. Suppose $\Theta(G)$ is reducible with a decomposition (Γ_1, T, Γ_2) of $\mathcal{C}(G)$. Note that the simplicial complex Γ_i is the clique complex of a subgraph $G_i \subset G$ for $i = 1, 2$, and the intersection of G_1 and G_2 is the complete graph on the vertex set T . We will denote the vertex set of G_1 by U and the vertex set of G_2 by W . Associated to these subgraphs we have graphical models $\Theta(G_1)$ and $\Theta(G_2)$. When A is an $m \times m$ matrix whose rows and columns are indexed by $[m]$, we let A_{IJ} denote the submatrix of A , whose rows are indexed by $I \subseteq [m]$ and whose columns are indexed by $J \subseteq [m]$. For any $|I| \times |J|$ matrix $B = (b_{ij})_{i \in I, j \in J}$, we define $[B]^{[m]}$ to be the matrix obtained from B by filling in zero entries to obtain a $m \times m$ matrix, i.e.

$$([B]^{[m]})_{ij} := \begin{cases} b_{ij} & \text{if } i \in I, j \in J \\ 0 & \text{otherwise} \end{cases}.$$

The maximum likelihood estimate of $S \in \text{PD}_m$ in $\Theta(G)$ can be computed as follows.

Proposition 5.1.5. [78, Proposition 5.6] Let $\Theta(G)$ be a reducible graphical model on the undirected graph G with a decomposition (Γ_1, T, Γ_2) . Let $S \in \text{PD}_m$, and let $\hat{\Sigma}_{UU}$ be the MLE of S_{UU} in $\Theta(G_1)$ and let $\hat{\Sigma}_{WW}$ be the MLE of S_{WW} in $\Theta(G_2)$. Let $\hat{\Sigma}_{TT} := S_{TT}$. The maximum likelihood estimate $\hat{\Sigma}$ of S in $\Theta(G)$ is given by

$$\hat{\Sigma}^{-1} = [(\hat{\Sigma}_{UU})^{-1}]^{[m]} + [(\hat{\Sigma}_{WW})^{-1}]^{[m]} - [(\hat{\Sigma}_{TT})^{-1}]^{[m]}.$$

For a graph G and a matrix $\Sigma \in \Theta(G)$, we denote the logarithmic Voronoi cell at Σ by $\log \text{Vor}_G(\Sigma)$. For reducible graphical models, we have the following decomposition theorem.

Theorem 5.1.6. Let $\Theta(G)$ be a reducible graphical model on the undirected graph G with a decomposition (Γ_1, T, Γ_2) . For any matrix $\Sigma \in \Theta(G)$, the logarithmic Voronoi cell $\log \text{Vor}_G(\Sigma)$ equals

$$\left(\left\{ ([S_1^{-1}]^{[m]} + [S_2^{-1}]^{[m]} - [\Sigma_{TT}^{-1}]^{[m]})^{-1} : S_1 \in \log \text{Vor}_{G_1}(\Sigma_{UU}) \text{ and } S_2 \in \log \text{Vor}_{G_2}(\Sigma_{WW}) \right\} + \ker(\psi) \right) \cap \text{PD}_m,$$

where $\psi : \text{Sym}(\mathbb{R}^m) \rightarrow \text{Sym}(\mathbb{R}^U) \times \text{Sym}(\mathbb{R}^W)$ is the map

$$\psi : M \mapsto (M_{UU}, M_{WW}).$$

Proof. First, observe that the projections Σ_{UU} and Σ_{WW} are in $\Theta(G_1)$ and $\Theta(G_2)$, respectively. This follows from the Schur complement formula for matrix inverses. Let $A := U \setminus T$ and $B := W \setminus T$. First consider the matrix S given by

$$S^{-1} = [S_1^{-1}]^{[m]} + [S_2^{-1}]^{[m]} - [\Sigma_{TT}^{-1}]^{[m]}$$

where $S_1 \in \log \text{Vor}_{G_1}(\Sigma_{UU})$ and $S_2 \in \log \text{Vor}_{G_2}(\Sigma_{WW})$. We will show that $S \in \log \text{Vor}_G(\Sigma)$. Recall that the logarithmic Voronoi cell at Σ is the set

$$\log \text{Vor}_G(\Sigma) = \{S \in \text{PD}_m : \Sigma_{ij} = S_{ij} \text{ for all } ij \in E(G) \text{ and } i = j\}.$$

Hence, it suffices to show that $S_{CC} = \Sigma_{CC}$ for every clique C of G . Note first that we may write S^{-1} in the block form as follows:

$$S^{-1} = \begin{bmatrix} (S_1^{-1})_{AA} & (S_1^{-1})_{AT} & 0 \\ (S_1^{-1})_{TA} & (S_1^{-1})_{TT} + (S_2^{-1})_{TT} - (\Sigma_{TT}^{-1}) & (S_2^{-1})_{TB} \\ 0 & (S_2^{-1})_{BT} & (S_2^{-1})_{BB} \end{bmatrix}.$$

Using Schur complements one checks that $S_{UU} = ((S^{-1})^{-1})_{UU} = (S_1^{-1})^{-1} = S_1$ and $S_{WW} = S_2$. Now, let C be a clique in G , so either $C \subseteq \Gamma_1$ or $C \subseteq \Gamma_2$. Without loss of generality, assume $C \subseteq \Gamma_1$. Then

$$S_{CC} = (S_{UU})_{CC} = (S_1)_{CC} = (\Sigma_{UU})_{CC} = \Sigma_{CC},$$

and we conclude that $S \in \log \text{Vor}_G(\Sigma)$.

Now let $M = (m_{ij}) \in \ker(\psi)$, i.e., $m_{ij} = 0$ for all $ij \in E(G)$ or $i = j$. In particular, $M_{CC} = 0$ for every clique C of G . Thus, if $S + M$ is positive definite, we have $S + M \in \log \text{Vor}_G(\Sigma)$, as desired.

For the other direction, let $S \in \log \text{Vor}_G(\Sigma)$. Define $S_1 := S_{UU}$ and $S_2 := S_{WW}$. Note that for any clique $C \subseteq \Gamma_1$, we have $(S_1)_{CC} = (S_{UU})_{CC} = S_{CC} = \Sigma_{CC} = (\Sigma_{UU})_{CC}$, so $S_1 \in \log \text{Vor}_{G_1}(\Sigma_{UU})$. Similarly, $S_2 \in \log \text{Vor}_{G_2}(\Sigma_{WW})$. Let $L = [S_1^{-1}]^{[m]} + [S_2^{-1}]^{[m]} - [\Sigma_{TT}^{-1}]^{[m]}$, and let $M := S - (L^{-1})$. Note that $S = L^{-1} + (S - L^{-1}) = L^{-1} + M$, so it suffices to show that $M \in \ker(\psi)$. We observe that $(L^{-1})_{UU} = ((S_1^{-1})^{-1})_{UU} = S_1 = S_{UU}$, so $M_{UU} = S_{UU} - (L^{-1})_{UU} = 0$. Similarly, we find that $M_{WW} = 0$. Hence, indeed $M \in \ker(\psi)$, and this concludes the proof. \square

Remark 5.1.7. The analogous decomposition theorem holds for discrete hierarchical models associated to a reducible simplicial complex, as discussed in Chapter 4. In both cases, the decomposition of logarithmic Voronoi cells is interesting: $\log \text{Vor}_G(\Sigma)$ as well as $\log \text{Vor}_{G_1}(\Sigma_{UU})$ and $\log \text{Vor}_{G_2}(\Sigma_{WW})$ are spectrahedra, but the first term in the Minkowski sum in Theorem 5.1.6 is a nonlinear object. How the geometry and combinatorics of the spectrahedra $\log \text{Vor}_{G_1}(\Sigma_{UU})$ and $\log \text{Vor}_{G_2}(\Sigma_{WW})$ affect that at $\log \text{Vor}_G(\Sigma)$ via this decomposition is worthwhile to study in a future project.

5.2 Directed graphical models

In this section we turn to Gaussian models defined by *directed acyclic graphs* (DAGs). A DAG G consists of a vertex set V of cardinality m and a set E of directed edges (i, j) without a directed cycle. We will assume that $(i, j) \in E$ implies $i < j$. Such a topological ordering of the vertices can always be achieved. For each vertex $j \in G$ there is a normal random variable X_j such that $X_j = \sum_{k \in \text{pa}(j)} \lambda_{jk} X_k + \varepsilon_j$. Here $\text{pa}(j)$ denotes the set of parents of the vertex j . The coefficients λ_{jk} are real parameters, known as *regression coefficients*. The term ε_j is a random variable that has a univariate normal distribution. This model can be summarized by the identity

$$X = \Lambda^T X + \varepsilon$$

where $\Lambda = (\lambda_{kj})$ is an upper triangular matrix with $\lambda_{kk} = 0$ for $k = 1, \dots, m$. The joint random variable $X = (X_1, \dots, X_m)^T$ has a Gaussian distribution with covariance matrix Σ . We also denote the diagonal covariance matrix of ε by Ω . With this $\Sigma = (I - \Lambda)^{-T} \Omega (I - \Lambda)^{-1}$, and the maximum likelihood estimation aims to estimate the $|E|$ and m parameters in Λ and Ω , respectively. The concentration matrix $K = \Sigma^{-1}$ is equal to $(I - \Lambda) \Omega^{-1} (I - \Lambda)^T$. The maximum likelihood estimate can be found by solving a sequence of independent least squares problems for each vertex in the graph coming from the gradient of the log-likelihood function: Given n independent observations of the random variable X , we collect them in a $n \times m$ matrix Y . Then the log-likelihood function is

$$\ell_n(\Sigma, Y^T Y) = \log \det K - \text{tr}(Y^T Y K) = - \sum_{k=1}^m \log \Omega_{kk} - \sum_{k=1}^m \frac{1}{\Omega_{kk}} [(Y(I - \Lambda))^T (Y(I - \Lambda))]_{kk}.$$

One observes the solutions to $\nabla \ell_n = 0$ are obtained by first minimizing $\|Y_k - \sum_{j=1}^{k-1} \lambda_{jk} Y_j\|^2$ for $k = 1, \dots, m$ independently. These are least squares problems with unique solutions. This development leads to the following.

Theorem 5.2.1. [78, p. 154], [122] The maximum likelihood degree of a Gaussian model on a directed acyclic graph is one. Therefore, the logarithmic Voronoi cell at Σ on such a model is equal to its log-normal spectrahedron.

For algebraic computations a convenient parametrization for Gaussian models on DAGs based on the *trek rule* exists [113]. In this parametrization, for each directed edge $(i, j) \in E$ there is λ_{ij} and for each vertex $i \in [m]$ there is a_i . For each pair of vertices $i, j \in [m]$, we let $T(i, j)$ be the set of paths from i to j which do not contain colliders where a collider is a pair of edges (s, t) and (u, t) with the same head. Such a path without colliders is called a trek. Every trek P from i to j is a sequence of edges from i up to $\text{top}(P)$, the “top” vertex on the path, and then a sequence of edges down to j . With this the parametrization of the entries of the covariance matrix Σ reads as follows:

$$\sigma_{ij} = \sum_{P \in T(i, j)} a_{\text{top}(P)} \prod_{(k, l) \in P} \lambda_{kl}.$$

We note that $\sigma_{ii} = a_i$, and if $T(i, j) = \emptyset$ then $\sigma_{ij} = 0$.

Example 5.2.2. Consider the DAG $1 \rightarrow 2 \rightarrow 4 \leftarrow 3$. The associated graphical model Θ is seven-dimensional. We may express $\Sigma = (\sigma_{ij}) \in \Theta$ parametrically as

$$\begin{aligned} \sigma_{ii} &= a_i \text{ for } i = 1, 2, 3, 4, \\ \sigma_{12} &= a_1 \lambda_{12}, \sigma_{13} = 0, \sigma_{14} = a_1 \lambda_{12} \lambda_{24}, \sigma_{23} = 0, \sigma_{24} = a_2 \lambda_{24}, \sigma_{34} = a_3 \lambda_{34}. \end{aligned}$$

The logarithmic Voronoi cell and hence the log-normal spectrahedron of a general $\Sigma \in \Theta$ is three-dimensional, given as

$$\left\{ \begin{pmatrix} a_1 & a_1 \lambda_{12} & x & y \\ a_1 \lambda_{12} & a_2 & z & a_2 \lambda_{24} + \lambda_{34} z \\ x & z & a_3 & a_3 \lambda_{34} + \lambda_{24} z \\ y & a_2 \lambda_{24} + \lambda_{34} z & a_3 \lambda_{34} + \lambda_{24} z & 2 \lambda_{24} \lambda_{34} z + a_4 \end{pmatrix} \succ 0 : x, y, z \in \mathbb{R} \right\}.$$

For the matrix Σ' given by the parameters

$$a_1 = 1, a_2 = 2, a_3 = 3, a_4 = 4, \lambda_{12} = 1/2, \lambda_{24} = 1, \lambda_{34} = 1/2,$$

the spectrahedron $\log \text{Vor}_\Theta(\Sigma')$ is the intersection of a quadric, defining a cylinder, and a quartic, defining a surface with five components. The intersection is the middle component of the quartic surface. We plot the quadric surface (left), the quartic surface (middle) and their intersection (right) in Figure 5.2.

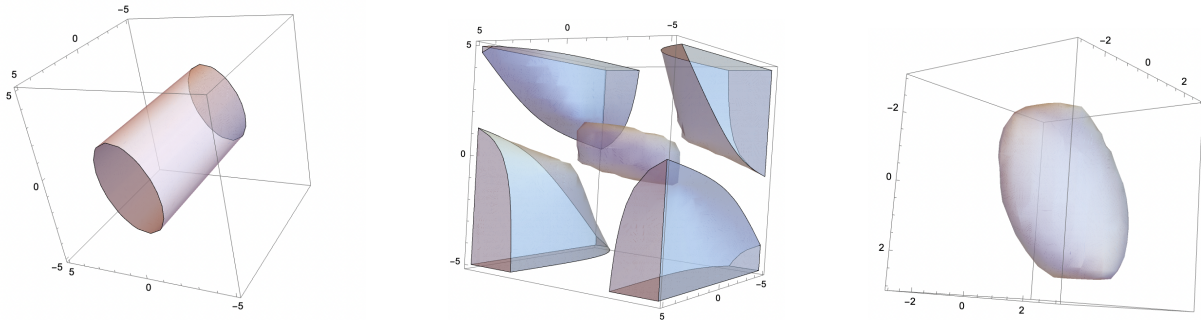


Figure 5.2: The logarithmic Voronoi cell at Σ' of $1 \rightarrow 2 \rightarrow 4 \leftarrow 3$.

We close this section with a simple decomposition result for logarithmic Voronoi cells when the underlying graph is the disjoint union of two graphs.

Proposition 5.2.3. Let G be a DAG with vertex set $[m]$ such that G is a disjoint union of two graphs G_1 and G_2 with vertex sets U and $W = [m] \setminus U$, respectively. Then

$$\Theta(G) = \left\{ \Sigma = \begin{pmatrix} \Sigma_1 & 0_{UW} \\ 0_{WU} & \Sigma_2 \end{pmatrix} : \Sigma_1 \in \Theta(G_1), \Sigma_2 \in \Theta(G_2) \right\},$$

and for $\Sigma \in \Theta(G)$

$$\log \text{Vor}_{\Theta(G)}(\Sigma) = \left\{ \begin{pmatrix} S_1 & S_{UW} \\ S_{WU} & S_2 \end{pmatrix} \succ 0 : S_1 \in \log \text{Vor}_{\Theta(G_1)}(\Sigma_1), S_2 \in \log \text{Vor}_{\Theta(G_2)}(\Sigma_2) \right\}.$$

Proof. The first statement is a direct consequence of Proposition 3.6 in [113]. The second statement follows from the observation that $\ell_n(\Sigma, S) = \ell_n(\Sigma_{UU}, S_{UU}) + \ell_n(\Sigma_{WW}, S_{WW})$. \square

5.3 Covariance models

Let $A \in \text{PD}_m$ and let \mathcal{L} be a linear subspace of $\text{Sym}(\mathbb{R}^m)$. Then $A + \mathcal{L}$ is an affine subspace of $\text{Sym}(\mathbb{R}^m)$. Models defined by $\Theta = (A + \mathcal{L}) \cap \text{PD}_m$ are called *covariance models*. For such models, a necessary condition for $\Sigma \in \text{PD}_m$ to be a local maximum of the log-likelihood function is $\Sigma - A \in \mathcal{L}$ and $K - KSK \in \mathcal{L}^\perp$ where $K = \Sigma^{-1}$; see [12] and [111]. From this, one can describe the log-normal spectrahedron at Σ in the model Θ explicitly.

Proposition 5.3.1. The log-normal spectrahedron $\mathcal{K}_\Sigma(\Theta)$ at $\Sigma \in \Theta$ on a covariance model $\Theta = (A + \mathcal{L}) \cap \text{PD}_m$ is equal to $\mathcal{N}_\Sigma \Theta \cap \text{PD}_m$ where

$$\mathcal{N}_\Sigma \Theta = \{S \in \text{Sym}(\mathbb{R}^m) : K - KSK \in \mathcal{L}^\perp\}.$$

The log-likelihood function $\ell_n(\Sigma, S)$ is generally not concave on a covariance model, and the maximum likelihood degree of such models can be arbitrarily high [111]. Therefore, in general, the logarithmic Voronoi cells are strictly contained in log-normal spectrahedra. On the other hand, we would like to point out the following interesting result.

Proposition 5.3.2. [127, Proposition 3.1] Let $\Theta \subseteq \text{PD}_m$ be a Gaussian covariance model and let $S \in \text{PD}_m$. The log-likelihood function $\ell_n(\Sigma, S)$ is strictly concave on the convex set $\Delta_{2S} = \{\Sigma \in \text{PD}_m : 0 \prec \Sigma \prec 2S\}$, and hence it is strictly concave on $\Delta_{2S} \cap \Theta$.

This proposition immediately implies the following.

Corollary 5.3.3. Let $\Theta \subseteq \text{PD}_m$ be a Gaussian covariance model and let $\Sigma \in \Theta$. Then we have the following containments:

$$\{S \in \mathcal{K}_\Theta(\Sigma) : 0 \prec \Sigma \prec 2S\} \subseteq \log \text{Vor}_\Theta(\Sigma) \subseteq \mathcal{K}_\Theta(\Sigma). \quad (5.3)$$

In general, both containments may be strict, as demonstrated in the next example.

Example 5.3.4. Consider the covariance model given by

$$\Theta = \left\{ \Sigma \in \text{PD}_3 : \Sigma = \begin{pmatrix} 1 & x & z \\ x & 1 & y \\ z & y & 1 \end{pmatrix} \right\}.$$

This is an unrestricted correlation model. Its ML degree is 15. We can represent each matrix $\Sigma \in \Theta$ by the triple $(\Sigma_{12}, \Sigma_{23}, \Sigma_{13})$. To see that the two containments in (5.3) are strict, first consider the matrix $\Sigma' = (1/2, 1/3, 1/4) \in \Theta$, and let

$$S = \begin{pmatrix} 1211/4560 & -217/3420 & 1/30 \\ -217/3420 & 827/2565 & 1/9 \\ 1/30 & 1/9 & 1 \end{pmatrix} \in \mathcal{K}_\Theta(\Sigma').$$

The log-likelihood function $\ell_n(\Sigma, S)$ has 15 critical points, three of which are real. The real points are given numerically by

$$\{(1/2, 1/3, 1/4), (-0.73841, 0.213623, -0.0580265), (0.182141, 0.316592, 0.190067)\}.$$

The values of the log-likelihood function are, respectively

$$\begin{aligned} & -1.53844955693696, \\ & -1.24750351572487, \\ & -1.55375020617405. \end{aligned}$$

We see that the maximum is achieved at the second point, meaning $S \notin \log \text{Vor}_\Theta(\Sigma')$. This shows that the second containment in (5.3) is strict. To see that the first containment is strict, let

$$S = \begin{pmatrix} 813/304 & 103/76 & 1/2 \\ 103/76 & 85/57 & 1/3 \\ 1/2 & 1/3 & 1/3 \end{pmatrix} \in \mathcal{K}_\Theta(\Sigma').$$

The matrix $2S - \Sigma'$ is not positive definite. However, $\ell_n(\Sigma, S)$ has only one real critical point, namely Σ' . Thus $S \in \log \text{Vor}_\Theta(\Sigma')$, which shows that the first containment is also strict.

Bivariate correlation model

A *bivariate correlation model* is an affine covariance model given parametrically as

$$\Theta = \left\{ \Sigma_x := \begin{pmatrix} 1 & x \\ x & 1 \end{pmatrix} : x \in (-1, 1) \right\}.$$

Maximum likelihood estimation of this model has been studied extensively in [12]. In this section we give an explicit description of its logarithmic Voronoi cells and show that they are semialgebraic sets. This is extremely surprising. As the development below will demonstrate, the potential constraints which define the boundary of logarithmic Voronoi cells of these one-dimensional models are very complicated. In particular, they are not algebraic. Nevertheless, one recovers a semi-algebraic description.

Given a sample covariance matrix S , the derivative of the log-likelihood function $\ell_n(\Sigma, S)$ with respect to x is $\frac{2}{(1-x^2)^2} \cdot f(x)$ where

$$f(x) = x(x^2 - 1) - S_{12}(1 + x^2) + x(S_{11} + S_{22}).$$

This polynomial has at least one real root in the interval $(-1, 1)$, which corresponds to a positive definite covariance matrix in the model. This tells us that the MLE always exists, and hence the logarithmic Voronoi cells fill the cone PD_2 . Letting $a = (S_{11} + S_{22})/2$ and $b = S_{12}$, the polynomial f can be re-written as $f(x) = x^3 - bx^2 - x(1 - 2a) - b$. This polynomial has either one or three real roots in the interval $(-1, 1)$. In the first case, there is a unique positive definite matrix that appears as a critical point when optimizing $\ell_n(\Sigma, S)$. In the second case, there are three possible positive definite critical points. As shown in [12], the latter happens if and only if $\Delta_f(b, a) > 0$ and $a < 1/2$, where

$$\Delta_f(b, a) = -4[b^4 - (a^4 + 8a - 11)b^2 + (2a - 1)^3]$$

is the discriminant of f .

Fix $c \in (-1, 1)$. We wish to compute the logarithmic Voronoi cell at Σ_c . Note that for a sample covariance matrix S to have Σ_c as a critical point, c must be a root of $f(x)$. Substituting c for x in f , we get an equation $f(c) = 0$ in a and b . From this equation, we may express a in terms of b and c :

$$a = \frac{bc^2 - c^3 + b + c}{2c}. \quad (5.4)$$

Only $S \in \text{PD}_2$ that satisfy this equation will have Σ_c as a critical point when maximizing $\ell_n(\Sigma, S)$. If for such S we have $\Delta_f(b, a) \leq 0$ or $a \geq 1/2$, then S has Σ_c as the MLE and thus $S \in \log \text{Vor}_\Theta(\Sigma_c)$. If $\Delta_f(b, a) > 0$ and $a < 1/2$, then we must compare the value that ℓ_n takes on Σ_c to the values that it takes on the two matrices Σ_1, Σ_2 corresponding to the other two real roots of $f(x)$, for the fixed a and b . Given the relationship between a and b as in (5.4) we find all three roots of $f(x)$ in terms of b and c . They are

$$c_1 = \frac{bc - c^2 - \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc}}{2c}$$

$$c_2 = \frac{bc - c^2 + \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc}}{2c}$$

and, of course, c itself.

Let

$$S_{b,k} = \begin{pmatrix} k & b \\ b & 2a - k \end{pmatrix} \succ 0, \quad 0 < k < 2a$$

denote a general matrix in PD_2 that has Σ_c as a critical point when computing the MLE. In particular, the relation (5.4) is satisfied. This set of matrices forms the log-normal spectrahedron $\mathcal{K}_\Theta(\Sigma_c)$ of Σ_c .

Theorem 5.3.5. Let Θ be the bivariate correlation model and let $\Sigma_c \in \Theta$. If $c > 0$, then $\log \text{Vor}_\Theta(\Sigma_c) = \{S_{b,k} \in \mathcal{K}_\Theta(\Sigma_c) : b \geq 0\}$. If $c < 0$, then $\log \text{Vor}_\Theta(\Sigma_c) = \{S_{b,k} \in \mathcal{K}_\Theta(\Sigma_c) : b \leq 0\}$. If $c = 0$, then $\log \text{Vor}_\Theta(\Sigma_c) = \{\text{diag}(k, 2a - k) : a \geq 1/2, 0 < k < 2a\}$. In particular, logarithmic Voronoi cells of Θ are semi-algebraic sets.

Proof. First, suppose that $c > 0$. Since we only consider the positive definite matrices $S_{b,k}$, we are working in the cone $a > |b|$. This gives us the restriction $b > c(c - 1)/(c + 1)$. Note that

$$\Delta_f(b, c) = \frac{(b^2c - 2bc^2 - 4b + c^3)(bc^2 - 2c^3 - b)^2}{c^3}.$$

Thus, $\Delta_f \leq 0$ if and only if

$$\frac{c^2 - 2\sqrt{c^2 + 1} + 2}{c} \leq b \leq \frac{c^2 + 2\sqrt{c^2 + 1} + 2}{c}. \quad (5.5)$$

Moreover, since $a = \frac{bc^2 - c^3 + b + c}{2c}$, we also have $a \geq 1/2$ if and only if $b \geq \frac{c^3}{c^2 + 1}$. Since for $c > 0$, we always have

$$\frac{c^3}{c^2 + 1} \leq \frac{c^2 + 2\sqrt{c^2 + 1} + 2}{c},$$

it follows that

$$\left\{ S_{b,k} \in \mathcal{K}_\Theta(\Sigma_c) : b \geq \frac{c^2 - 2\sqrt{c^2 + 1} + 2}{c} \right\} \subseteq \log \text{Vor}_\Theta(\Sigma_c).$$

Now, suppose $b \leq \frac{c^2 - 2\sqrt{c^2 + 1} + 2}{c}$. Such sample covariance matrices $S_{b,k}$ will have three positive definite roots when optimizing $\ell_n(\Sigma, S)$, namely Σ_{c_1} , Σ_{c_2} , and Σ_c . In order for such matrix $S_{b,k}$ to be in $\log \text{Vor}_\Theta(\Sigma_c)$, it has to be the case that

$$\ell_n(\Sigma_c, S) \geq \ell_n(\Sigma_{c_i}, S) \text{ for } i = 1, 2.$$

A computation (in SAGE [103]) shows that both inequalities above are inequalities in b only. The only constraints on k are given by the positive definiteness of S . The values of the

log-likelihood function are

$$\begin{aligned} \ell_n(\Sigma_{c_1}, S_{b,k}) &= -\frac{1}{D} \left[4bc^2 + (2bc^2 - c^3 - (b^2 - 2)c + 2b) \log \left(\frac{2bc^2 - c^3 - (b^2 - 2)c + \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc(b-c) + 2b}}{2c} \right) \right. \\ &\quad \left. + \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc} \left((b-c) \log \left(\frac{2bc^2 - c^3 - (b^2 - 2)c + \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc(b-c) + 2b}}{2c} \right) + 2b \right) + 2b - 2c^3 - 2(b^2 - 1)c \right] \\ \ell_n(\Sigma_{c_2}, S_{b,k}) &= -\frac{1}{D} \left[4bc^2 + (2bc^2 - c^3 - (b^2 - 2)c + 2b) \log \left(\frac{2bc^2 - c^3 - (b^2 - 2)c - \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc(b-c) + 2b}}{2c} \right) \right. \\ &\quad \left. - \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc} \left((b-c) \log \left(\frac{2bc^2 - c^3 - (b^2 - 2)c - \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc(b-c) + 2b}}{2c} \right) + 2b \right) + 2b - 2c^3 - 2(b^2 - 1)c \right] \\ \ell_n(\Sigma_c, S_{b,k}) &= -\frac{c \log(-c^2 + 1) + b + c}{c}, \end{aligned}$$

where $D = 2bc^2 - c^3 - (b^2 - 2)c + \sqrt{b^2c^2 - 2bc^3 + c^4 - 4bc(b-c) + 2b}$. Note that for fixed c , the last function is linear in b , with the negative slope $-1/c$. At $b = 0$, we always have $\ell_n(\Sigma_{c_1}, S_{0,k}) = \ell_n(\Sigma_c, S_{0,k}) > \ell_n(\Sigma_{c_2}, S_{0,k})$, so $S_{0,k} \in \log \text{Vor}_\Theta(\Sigma_c)$ for $0 < k < 2a$. Since logarithmic Voronoi cells are convex sets by Proposition 5.0.1, this means that the containment $\{S_{b,k} \in \mathcal{K}_\Theta(\Sigma_c) : b \geq 0\} \subseteq \log \text{Vor}_\Theta(\Sigma_c)$ holds. For the other containment, let $g(b) = \ell_n(\Sigma_{c_1}, S_{b,k}) - \ell_n(\Sigma_c, S_{b,k})$ and consider its Taylor expansion $g(b) = g(0) + g'(0)b + \dots$. Note that $g(0) = 0$ and $g'(0) < 0$ for all $0 < c < 1$. Thus, for all $b^* < 0$ with $|b^*|$ sufficiently small, the term $g'(0)b^*$ is positive and dominating in the expansion. This means that for such $b^* < 0$, we have $\ell_n(\Sigma_{c_1}, S_{b^*,k}) > \ell_n(\Sigma_c, S_{b^*,k})$, and so $S_{b^*,k} \notin \log \text{Vor}_\Theta(\Sigma_c)$. Thus, $\{S_{b,k} \in \mathcal{K}_\Theta(\Sigma_c) : b \geq 0\} = \log \text{Vor}_\Theta(\Sigma_c)$ by convexity of logarithmic Voronoi cells. The proof for $c < 0$ is similar. For $c = 0$, we have that $b = 0$ and the log-normal spectrahedron is given by $\{S_{a,k} := \text{diag}(k, 2a - k) : 0 \leq k \leq 2a\}$. The two other critical points, besides 0, are given by $c_1 = \sqrt{1 - 2a}$ and $c_2 = -\sqrt{1 - 2a}$. These are real if $a \leq 1/2$. In this case, the values of the log-likelihood function are as follows:

$$\ell_n(\Sigma_{c_1}, S_{a,k}) = \ell_n(\Sigma_{c_2}, S_{a,k}) = -\log(2a) - 1, \quad \ell_n(\Sigma_c, S_{a,k}) = -2a.$$

Note that $\ell_n(\Sigma_{c_1}, S_{a,k})$ is a monotone decreasing strictly convex function and $\ell_n(\Sigma_c, S_{a,k})$ is a linear function with slope -2 , tangent to $\ell_n(\Sigma_{c_1}, S_{a,k})$ at $a = 1/2$. Thus, the only time Σ_0 is the MLE in this regime is when $a = 1/2$. If $a > 1/2$, $c = 0$ is the only real critical point and gives the MLE. We conclude that $\log \text{Vor}_\Theta(\Sigma_0) = \{\text{diag}(k, 2a - k) : a \geq 1/2, 0 < k < 2a\}$. \square

Remark 5.3.6. Note that since the bivariate correlation model is a compact set inside PD_2 , its log-normal spectrahedra for general matrices as well as logarithmic Voronoi cells are unbounded. In general, for $0 < c < 1$, the part of the log-normal spectrahedron at Σ_c that is not in the logarithmic Voronoi cell at Σ_c is small. This is due to the fact that $c(c-1)/(c+1)$ is a negative number with small magnitude. As $c \rightarrow 1$, the logarithmic Voronoi cell converges to the log-normal spectrahedron. In Figure 5.3, we plot the logarithmic Voronoi cell at $c = 1/2$ as the intersection of the pink log-normal spectrahedron and the blue half-space $b \geq 0$. Similar statement is true of $-1 < c < 0$.




Figure 5.3: The logarithmic Voronoi cell at $\Sigma_{1/2}$ for the bivariate correlation model.

Equicorrelation models

An *equicorrelation model* is given by the parameter space

$$\Theta_m = \{\Sigma_x \in \text{Sym}(\mathbb{R}^m) : \Sigma_{ii} = 1, \Sigma_{ij} = x \text{ for } i \neq j, i, j \in [m], x \in \mathbb{R}\} \cap \text{PD}_m.$$

Note that this model is an instance of the affine covariance model, with $A = I_m$ and $\mathcal{L} = \text{span}_{\mathbb{R}}\{\mathbf{1}\mathbf{1}^T - I_m\}$, where $\mathbf{1}$ denotes the all-ones vector in \mathbb{R}^m . Note also that Θ_2 is precisely the bivariate correlation model. For a symmetric matrix $\Sigma_x = (1 - x)I_m + x\mathbf{1}\mathbf{1}^T$ to be positive definite, $-\frac{1}{m-1} < x < 1$ must hold.

Given $c \in \mathbb{R}$ such that $-\frac{1}{m-1} < c < 1$, we wish to describe the logarithmic Voronoi cell at $\Sigma_c \in \Theta_m$. Let $S \in \text{PD}_m$ be a sample covariance matrix. Following [12], define the *symmetrized sample covariance matrix* to be the matrix

$$\bar{S} = \frac{1}{m!} \sum_{P \in S_m} P S P^T$$

where S_m denotes the group of all $m \times m$ permutation matrices. Let \mathcal{N} denote the space of all symmetrized sample covariance matrices. Note that for $i, j \in [m]$, we have $\bar{S}_{ii} = a$ and $\bar{S}_{ij} = b$ whenever $i \neq j$. From Lemma 5.2 in [12], we have $\langle S, \Sigma_c^{-1} \rangle = \langle \bar{S}, \Sigma_c^{-1} \rangle$, so optimizing $\ell_n(\Sigma, S)$ is equivalent to optimizing $\ell_n(\Sigma, \bar{S})$. Hence, we may fully recover the logarithmic Voronoi cells at Σ_c with the matrices \bar{S} for which $\ell_n(\Sigma, \bar{S})$ is maximized at Σ_c .

From Theorem 5.4 in [12], we know that the ML degree of the equicorrelation model is 3 and the critical points for a general \bar{S} with $\bar{S}_{ii} = a$ and $\bar{S}_{ij} = b$ for $i \neq j$ are given by the points Σ_r where r is a root of the cubic

$$f_m(x) = (m - 1)x^3 + ((m - 2)(a - 1) - (m - 1)b)x^2 + (2a - 1)x - b.$$

Since we are interested in the matrices \bar{S} that have Σ_c as a critical point, c must be a root of $f_m(x)$. Then, the equation $f_m(c) = 0$ becomes an equation expressing the relationship between a and b , namely

$$a = -\frac{(b+2)c^2 - c^3 - ((b+1)c^2 - c^3)m - b - c}{c^2m - 2c^2 + 2c}.$$

All positive definite matrices \bar{S} satisfying the above relationship are the set $\mathcal{K}_{\Theta_m}(\Sigma_c) \cap \mathcal{N}$. The matrices in this set depend only on the parameter b , so we will denote them by \bar{S}_b . Not all such matrices may be in the logarithmic Voronoi cell at Σ_c . The matrices \bar{S}_b for which the discriminant $\Delta_{f,m}(b, a)$ is negative will be in the logarithmic Voronoi cell, since for such points, f_m has only one real root. When $\Delta_{f,m}(b, a) \geq 0$, the situation is more complicated. The good news is that most such matrices \bar{S}_b satisfying $f_m(c) = 0$ will still have only one positive definite critical point in the model, namely Σ_c . However, some matrices may have two additional critical points. In such cases, we have to evaluate $\ell_n(\Sigma, \bar{S}_b)$ on the other two roots of $f_m(x)$, denoted by c_1 and c_2 . If $\ell_n(\Sigma_c, \bar{S}_b)$ is the largest, then \bar{S}_b would be in the logarithmic Voronoi cell at Σ_c . Precisely, we have that

$$\log \text{Vor}_{\Theta_m}(\Sigma_c) \cap \mathcal{N} = \{\bar{S}_b \in \mathcal{K}_{\Theta_m}(\Sigma_c) \cap \mathcal{N} : \ell_n(\Sigma_c, \bar{S}_b) \geq \ell_n(\Sigma_{c_i}, \bar{S}_b), i = 1, 2\}.$$

For fixed c , the two inequalities defining the above set are inequalities in one variable b . Thus, the set $\log \text{Vor}_{\Theta_m}(\Sigma_c) \cap \mathcal{N}$ is one-dimensional. We have the following theorem.

Theorem 5.3.7. Let $\Sigma_c \in \Theta_m$. The logarithmic Voronoi cell at Σ_c is given as

$$\log \text{Vor}_{\Theta_m}(\Sigma_c) = \{S \in \text{PD}_m : \psi(S) = \bar{S}, \bar{S} \in \mathcal{N} \cap \log \text{Vor}_{\Theta_m}(\Sigma_c)\},$$

where $\psi : \text{PD}_m \rightarrow \mathcal{N} : S \mapsto \bar{S}$.

Proof. This follows from the equality $\langle S, \Sigma_c^{-1} \rangle = \langle \bar{S}, \Sigma_c^{-1} \rangle$. □

Note that the pre-image of any symmetrized covariance matrix \bar{S} under ψ has dimension $\binom{m+1}{2} - 2$, and so the logarithmic Voronoi cell at any generic $\Sigma_c \in \Theta_m$ has dimension $\binom{m+1}{2} - 2 + 1 = \binom{m+1}{2} - 1$, i.e. co-dimension 1, as expected.

We also observe that when m increases, the number of matrices \bar{S}_b that have two other positive-definite critical points besides Σ_c decreases. Moreover, in statistical practice, such matrices \bar{S}_b are rare, even for small sample sizes [12]. This means that for practical purposes, we may say that the logarithmic Voronoi cell at $\Sigma_c \in \Theta_m$ is *approximately* its log-normal spectrahedron.

Transcendentality of logarithmic Voronoi cells

In this chapter we have introduced logarithmic Voronoi cells for Gaussian models. In the case of models that are also well-known in algebraic statistics we have proved that the logarithmic

Voronoi cells are spectrahedra with explicit descriptions. These include linear concentration models such as Gaussian models on undirected graphs as well as Gaussian models on DAGs. The spectrahedra we have identified deserve further study.

The case of bivariate correlation models is quite interesting since they provide the first small instance where logarithmic Voronoi cells need not be semi-algebraic. However, we showed that even in this case we get semialgebraicity even though the logarithmic Voronoi cells are not equal to the log-normal spectrahedra. The bivariate correlation models fit into a larger class of models known as unrestricted correlation models. Such a model is given by the parameter space

$$\Theta = \{\Sigma \in \text{Sym}(\mathbb{R}^m) : \Sigma_{ii} = 1, i \in [m]\} \cap \text{PD}_m.$$

The ML degree of these models for $m \leq 6$ was computed in [12]. The case $m = 2$ is the bivariate correlation model whose ML degree is 3.

When $m = 3$, the model is a compact spectrahedron known as the *elliptope* in convex algebraic geometry literature. We have encountered this model with ML degree 15 in Example 5.3.4. The logarithmic Voronoi cells of the elliptope are unbounded 3-dimensional convex sets. We found it quite challenging to give a good description for them besides the one coming from its definition. We venture the following conjecture.

Conjecture 5.3.8. The logarithmic Voronoi cells for general points on the elliptope are not semi-algebraic; in other words, their boundary is defined by transcendental functions.

Conclusion. In this chapter, we showed that for Gaussian models of ML degree one and linear covariance models, logarithmic Voronoi cells and log-normal spectrahedra coincide. In particular, they are equal for both directed and undirected graphical models. We introduced a decomposition theory of logarithmic Voronoi cells for the latter family. We also gave an explicit description of logarithmic Voronoi cells for the bivariate correlation model and showed that they are semi-algebraic sets. Finally, we stated an important conjecture that logarithmic Voronoi cells for unrestricted correlation models are not semi-algebraic.

Part II

Context-specific models

Chapter 6

Decomposable context-specific models

In this chapter, we are interested in how algebra can be used to model causality and capture finer forms of independence; specifically, context-specific independence. In Section 6.1, we present the necessary background on CStree models and their connection to discrete DAG models discussed in Chapter 1. Decomposable CSmodels are defined in 6.2, where we illustrate the nature of CStrees and decomposable CSmodels by presenting a classification of all CStree models in three random variables. A highlight from this section is Theorem 6.2.1, which states that if the number of random variables is three then a CStree model is decomposable if and only if all of its minimal context DAGs are perfect. This is no longer true for four variables, see Example 6.2.4. In Section 6.3, we establish several combinatorial properties for balanced CStrees. Finally, Section 6.4 contains the proofs of the main algebraic results. This chapter is based on [2].

From Chapter 1, we know that a discrete graphical model $\mathcal{M}(G)$ is a set of joint probability distributions for a vector of discrete random variables (X_1, \dots, X_p) that satisfy conditional independence (CI) relations according to the non-adjacencies of a graph $G = ([p], E)$ where $[p] := \{1, \dots, p\}$. The type of graph used to encode CI relations is typically a directed acyclic graph (DAG) or an undirected graph (UG), although other types of graphs are possible [78].

Graphical models are widely used in several fields of science, such as artificial intelligence, biology, and epidemiology [76, 81, 96]. However, in some applications it is useful to consider models that encode a finer form of independence. Context-specific independence (CSI) is a generalization of conditional independence where the conditional independence between the random variables only holds for particular outcomes of the variables in the conditioning set. The classical graphical models based on DAGs or UGs are no longer able to capture these more refined relations. Several extended graphical representations of CSI models have been proposed in the literature, [25, 30, 97, 99, 105]. Apart from its usage to encode model assumptions more accurately, context-specific independence is also important in the study of structural causal models because the presence of more refined independence can improve

the identifiability of causal links [118].

We have seen graphical models associated to directed acyclic graphs in Section 1.1. A graphical model $\mathcal{M}(G)$ associated to an undirected graph G is called *decomposable* if G is chordal. Decomposable graphical models play a prominent role among graphical models because they exhibit optimal properties for probabilistic inference. There are several characterizations of decomposable models in terms of their algebraic, combinatorial, and geometric properties [49, 54, 55, 78]. This chapter generalizes this class of models to the context-specific setting via *decomposable context-specific models* (see Section 6.2) and shows that they mirror many of the properties that characterize decomposable graphical models. For brevity, we will refer to these models as decomposable CSmodels. Previous work on context-specific versions of decomposable models defines them by introducing labels to the edges of a decomposable undirected graph and taking care of preserving the properties such as perfect elimination ordering and clique factorization [35, 69]. Our approach in this chapter will be different in that we will define decomposable CSmodels to be the set of log-linear models contained inside the more general context-specific models known as CStrees. Thus each decomposable CSmodel $\mathcal{M}(\mathcal{T})$ is represented by a CStree \mathcal{T} .

Similar to discrete DAG models in Chapter 1, there are two ways to define a decomposable CSmodel $\mathcal{M}(\mathcal{T})$. The first approach uses a recursive factorization property according to \mathcal{T} , while the second one uses the local CSI relations implied by \mathcal{T} . From an algebro-geometric point of view, the recursive factorization property is a polynomial parametrization of $\mathcal{M}(\mathcal{T})$, while the polynomials associated to the local CSI statements in \mathcal{T} define the model $\mathcal{M}(\mathcal{T})$ implicitly. An important open problem that arises in the study of context-specific models is characterizing the set of all CSI statements implied by the local CSI statements defining the model [25, 36]. This problem is especially amenable to algebraic techniques because any CSI relation that holds in the model can be represented by a collection of polynomials. Stated in algebraic terms, this problem is equivalent to that of finding a prime polynomial ideal that defines $\mathcal{M}(\mathcal{T})$ implicitly [53]. Moreover, for log-linear models, the generators of the ideal that defines $\mathcal{M}(\mathcal{T})$ form a Markov basis [40]. Our first main theorem is an algebraic characterization of the distributions that belong to a decomposable CSmodel. This theorem is similar to the Hammersley-Clifford theorem for undirected graphical models and its generalization [55, Theorem 4.2]. Note that we do not assume positivity of the distribution.

Theorem 6.0.1 (Context-specific Hammersley-Clifford). A distribution f factorizes according to a decomposable CSmodel $\mathcal{M}(\mathcal{T})$ if and only if the polynomials associated to saturated CSI statements in \mathcal{T} vanish at f . Moreover, the polynomials associated to the saturated CSI statements of a decomposable CSmodel form its Markov basis.

Every decomposable CSmodel $\mathcal{M}(\mathcal{T})$ can also be represented by a collection of *minimal context DAGs*. We also prove that the saturated CSI statements in Theorem 6.0.1 are

obtained as the union of saturated d -separation statements that hold in each of the minimal context DAGs that represent the model $\mathcal{M}(\mathcal{T})$ (Corollary 6.4.5).

This work also informs us about the set of CSI statements that hold for certain context-specific models known as LDAGs [97]. Briefly, an LDAG is a context-specific model represented by a DAG with edge labels, where these labels encode the extra CSI relations that hold for the model. Every CStree is an LDAG and every LDAG is a staged tree [50]. Whenever the LDAG is represented by a balanced CStree, Theorem 6.0.1 gives a complete characterization of the CSI statements that hold for the LDAG. In general, however, describing all CSI statements that hold for LDAGs is coNP-hard [36].

6.1 Introduction to CStrees

In this section, we introduce notation used in this chapter and generalize conditional independence statements to the context-specific case. We then introduce staged trees, CStrees, and explain how they may be represented by DAGs. Finally, we define balanced CStrees and discuss their importance.

Notation

Similar to Chapter 1, we consider a vector of discrete random variables $X_{[p]} = (X_1, \dots, X_p)$ where $[d_i]$ is the state space of X_i and $\mathcal{R} = \prod_{i \in [p]} [d_i]$ is the state space of $X_{[p]}$. To be consistent with the literature, we will denote elements in \mathcal{R} by sequences $\bar{x} = (x_1, \dots, x_p) = x_1 \cdots x_p$ where $x_k \in [d_k]$ for every $k \in [p]$. Note that these sequences are the same sequences as in Section 1.1, except we denoted their coordinates by i_j as opposed to x_j . For $A \subset [p]$, X_A is a vector of discrete random variables with indices in A . A probability distribution f for $X_{[p]}$ is a tuple $(f(\bar{x}) : \bar{x} \in \mathcal{R})$ where $f(\bar{x}) > 0$ and $\sum_{\bar{x} \in \mathcal{R}} f(\bar{x}) = 1$, $f(\bar{x})$ is the probability of the outcome $\bar{x} \in \mathcal{R}$. All distributions for $X_{[p]}$ live inside $\Delta_{|\mathcal{R}|-1}^\circ$.

To define a subvariety of the probability simplex we use the polynomial ring $\mathbb{R}[D] := \mathbb{R}[p_{\bar{x}} : \bar{x} \in \mathcal{R}]$. For any subset $H \subset \mathbb{R}[D]$ we denote the algebraic variety $\mathcal{V}(H) = \{x \in \mathcal{C}^{|\mathcal{R}|} : g(x) = 0 \text{ for all } g \in H\}$, so the intersection $\mathcal{V}(H) \cap \Delta_{|\mathcal{R}|-1}^\circ$ is a statistical model. In our situation, statistical models can also be defined as closed images of rational maps intersected with the probability simplex. A main question in Algebraic Statistics is to find the implicit equations that define the parametrized model. In this statistical setting, the defining algebraic equations translate into restrictions on the distributions which encode model assumptions.

Example 6.1.1. Consider the case $p = 3$ and X_1, X_2, X_3 are binary random variables. The graph $G = ([3], 1 \rightarrow 2 \rightarrow 3)$ imposes conditional independence relations among the X_i , $i \in [3]$. The state space of (X_1, X_2, X_3) is $\mathcal{R} = \{0, 1\}^3$ because all random variables are binary, hence $|\mathcal{R}| = 8$. The polynomial ring is $\mathbb{R}[D] = \mathbb{R}[p_{000}, p_{001}, p_{010}, p_{011}, p_{100}, p_{101}, p_{110}, p_{111}]$.

From Chapter 1, we know that G tells us that X_1 is independent of X_3 , given X_2 , as there is no edge from 1 to 3. This CI relation translates into two equations, one for each outcome of the conditioning variable X_2 :

$$p_{100}p_{001} - p_{101}p_{000}, \quad p_{110}p_{011} - p_{111}p_{010}.$$

These two equations, together with the hyperplane $\sum_{\bar{x} \in \mathcal{R}} p_{\bar{x}} = 1$, define a variety in the affine 8-dimensional space. Taking into account positivity conditions yields the model inside Δ_7° . By ignoring the sum-to-one hyperplane we immediately see that this model defines a toric variety in \mathbb{P}^7 as it is cut out by a prime binomial ideal.

Context-specific conditional independence statements

Let A, B, C, S be disjoint subsets of $[p]$ and $\bar{x}_C \in \mathcal{R}_C$, a distribution $f \in \Delta_{|\mathcal{R}|-1}^\circ$ satisfies the *context-specific conditional independence statement* (CSI statement) $X_A \perp\!\!\!\perp X_B | S, X_C = \bar{x}_C$ if for all outcomes $(\bar{x}_A, \bar{x}_B, \bar{x}_S) \in \mathcal{R}_A \times \mathcal{R}_B \times \mathcal{R}_S$

$$f(\bar{x}_A | \bar{x}_B, \bar{x}_S, \bar{x}_C) = f(\bar{x}_A | \bar{x}_S, \bar{x}_C).$$

When C is empty we recover the notion of conditional independence $X_A \perp\!\!\!\perp X_B | X_S$. To each CSI statement $X_A \perp\!\!\!\perp X_B | X_S, X_C = \bar{x}_C$ we associate the collection of polynomials

$$p_{\bar{x}_A \bar{x}_B \bar{x}_S \bar{x}_C} + p_{\bar{y}_A \bar{y}_B \bar{x}_S \bar{x}_C} - p_{\bar{x}_A \bar{y}_B \bar{x}_S \bar{x}_C} + p_{\bar{y}_A \bar{x}_B \bar{x}_S \bar{x}_C} \quad (6.1)$$

for every $\bar{x}_A, \bar{y}_A \in \mathcal{R}_A$, $\bar{x}_B, \bar{y}_B \in \mathcal{R}_B$ and $\bar{x}_S \in \mathcal{R}_S$ where

$$p_{\bar{x}_A \bar{x}_B \bar{x}_S \bar{x}_C} = \sum_{\bar{z} \in \mathcal{R}_{[p] \setminus (A \cup B \cup C)}} p_{\bar{x}_A \bar{x}_B \bar{x}_S \bar{x}_C \bar{z}}.$$

Note that these polynomials are the 2-minors of the matrix $(p_{\bar{x}_A \bar{x}_B \bar{x}_S \bar{x}_C})_{\bar{x}_A \in \mathcal{R}_A, \bar{x}_B \in \mathcal{R}_B}$ for all outcomes $\bar{x}_S \in \mathcal{R}_S$. We define the ideal $I_{X_A \perp\!\!\!\perp X_B | X_S, X_C = \bar{x}_C}$ in $\mathbb{R}[D]$ to be the ideal generated by all the polynomials in (6.1).

Given a collection \mathcal{C} of CSI statements, we define the CSI ideal generated by the polynomials associated to all CSI statements in \mathcal{C} , i.e.

$$I_{\mathcal{C}} = \sum_{X_A \perp\!\!\!\perp X_B | X_S, X_C = \bar{x}_C \in \mathcal{C}} I_{X_A \perp\!\!\!\perp X_B | X_S, X_C = \bar{x}_C}.$$

Graphical models are a widely used class of CI models. Recall that the DAG model $\mathcal{M}(G)$ is the set of all the distributions in $\Delta_{|\mathcal{R}|-1}^\circ$ that satisfy the recursive factorization property according to G . However, DAG models can also be defined implicitly by using the polynomials associated to the CI statements obtained via the different Markov properties. Any CI




Figure 6.1: A CStree for $p = 3$ and its minimal context DAGs.

relation that holds for the model $\mathcal{M}(G)$ is obtained from G via d -separation statements [78] and the set of all d -separation statements, denoted by $\text{global}(G)$, defines the global Markov property on G . The corresponding ideal $I_{\text{global}(G)}$ is not prime in general, as we discussed in the introduction. It is however prime and generated by binomials if the graph is perfect (see Definition 6.1.7).

While DAG models are well-suited to encode CI statements, they cannot encode CSI statements. There are several models one could use instead to encode CSI statements. In this chapter we focus on staged tree models, first introduced in [105], which we define next.

Staged Trees and CStrees

To have control of the CSI statements that hold in the model, it is convenient to represent the outcome space of $X_{[p]}$ as an event tree. Let $\pi_1 \cdots \pi_p$ be an ordering of $[p]$ and let $\mathcal{T} = (V, E)$ be a rooted tree with $V := \{\text{root}\} \cup \bigcup_{j \in [p]} \mathcal{R}_{\{\pi_1 \cdots \pi_j\}}$ and set of edges

$$E := \{\text{root} \rightarrow x_{\pi_1} : x_{\pi_1} \in [d_{\pi_1}]\} \cup \{x_{\pi_1} \cdots x_{\pi_{k-1}} \rightarrow x_{\pi_1} \cdots x_{\pi_k} : x_{\pi_1} \cdots x_{\pi_{k-1}} \in \mathcal{R}_{\pi_1, \dots, \pi_{k-1}}, x_k \in [d_k], k \in [p]\}.$$

The level of a node $v \in V$ is the number of edges in the unique path in \mathcal{T} from the root to v . The k th level of \mathcal{T} is the set of all nodes in \mathcal{T} at level k and is denoted by L_k . For \mathcal{T} as defined, we see that L_k is in bijection with the space of outcomes of the random vector $X_{\{\pi_1, \dots, \pi_k\}}$. Hence we associate the variable X_{π_k} with the level L_k and denote this association by $(L_1, \dots, L_p) \sim (X_{\pi_1}, \dots, X_{\pi_p})$. For any tree \mathcal{T} we write $V_{\mathcal{T}}$ and $E_{\mathcal{T}}$ for its sets of vertices and edges, respectively. We write $E(v)$ to denote the set of all outgoing edges from v . Without loss of generality, throughout this chapter we will assume $\pi_i = i$ for all $i \in [p]$. In particular, this implies $(L_1, \dots, L_p) \sim (X_1, \dots, X_p)$. The tree in Figure 6.1 represents an event tree for a vector (X_1, X_2, X_3) of binary random variables. To illustrate part of the notation, in this tree we have $L_2 = \{00, 01, 10, 11\}$ and $E(0) := \{0 \rightarrow 00, 0 \rightarrow 01\}$.

Let $\mathcal{T} = (V, E)$ be a rooted tree with levels $(L_1, \dots, L_p) \sim (X_1, \dots, X_p)$, \mathcal{L} a finite set of labels and $\theta : E \rightarrow \mathcal{L}$ a labeling of the edges. The pair (\mathcal{T}, θ) is a *staged tree* if

- (1) $|\theta(E(v))| = |E(v)|$ for all $v \in V$, and
- (2) for any pair $v, w \in V$, $\theta(E(v))$ and $\theta(E(w))$ are either equal or disjoint.

Two vertices v, w in (\mathcal{T}, θ) are in the same stage if and only if $\theta(E(v)) = \theta(E(w))$. In this case we write $v \sim w$. The equivalence relation \sim on the set V induces a partition of V called the *staging* of \mathcal{T} . We refer to each set in this partition as a *stage*. When depicting staged trees, such as in Figure 6.1, we use colors in the vertices to indicate that two vertices are in the same stage, except for white vertices which always represent singleton stages. Intuitively, the vertices in a stage S in level L_{k-1} represent outcomes $x_1 \cdots x_{k-1}$ for which the conditional distributions $f(X_k | x_1 \cdots x_{k-1}, x_1 \cdots x_{k-1} \in S)$ are all equal.

Definition 6.1.2. The staged tree (\mathcal{T}, θ) is a *CStree* if

- (1) $\theta(E(v)) \neq \theta(E(w))$ if v, w are in different levels,
- (2) $\theta(x_1 \cdots x_{k-1} \rightarrow x_1 \cdots x_{k-1} x_k) = \theta(y_1 \cdots y_{k-1} \rightarrow y_1 \cdots y_{k-1} x_k)$ whenever the nodes $x_1 \cdots x_{k-1}$ and $y_1 \cdots y_{k-1}$ are in the same stage.
- (3) For every stage $S \subset L_{k-1}, k \in [p]$, there exists $C \subset [k-1]$ and $\bar{x}_C \in \mathcal{R}_C$ such that

$$S = \bigcup_{\bar{y} \in \mathcal{R}_{[k-1] \setminus C}} \{\bar{x}_C \bar{y}\}.$$

Condition (1) ensures that two nodes in different levels do not share the same conditional distribution. Condition (2) forces that edges with the same label must point to the same outcome of X_k . The CStree model $\mathcal{M}(\mathcal{T})$ has the space of parameters

$$\Theta_{\mathcal{T}} = \left\{ x \in \mathcal{R}^{|\mathcal{L}|} : \forall e \in E, x_{\theta(e)} \in (0, 1) \text{ and } \forall v \in V, \sum_{e \in E(v)} x_{\theta(e)} = 1 \right\}$$

and is defined to be the image of the map

$$\Psi_{\mathcal{T}}: \Theta_{\mathcal{T}} \rightarrow \Delta_{|\mathcal{R}|-1}^{\circ}, \quad x \mapsto \left(\prod_{e \in E(\text{root} \rightarrow \bar{x})} x_{\theta(e)} \right)_{\bar{x} \in \mathcal{R}}$$

where $E(\text{root} \rightarrow \bar{x})$ denotes the set of all edges on the path from the root to \bar{x} . We say a distribution $f \in \Delta_{|\mathcal{R}|-1}^{\circ}$ factors according to \mathcal{T} if $f \in \text{im}(\Psi_{\mathcal{T}})$.

To describe a CStree model $\mathcal{M}(\mathcal{T})$ as an algebraic variety intersected with the open probability simplex we consider the ring homomorphism

$$\psi_{\mathcal{T}}: \mathbb{R}[D] \rightarrow \mathbb{R}[\Theta_{\mathcal{T}}], \quad p_{\bar{x}} \mapsto \prod_{e \in E(\text{root} \rightarrow \bar{x})} \theta(e)$$

where $\mathbb{R}[\Theta_{\mathcal{T}}] := \mathbb{R}[\theta(e) : e \in E] / \langle \theta - 1 \rangle$ and $\langle \theta - 1 \rangle := \langle \sum_{e \in E(v)} \theta(e) - 1 : v \in V \rangle$ is the ideal representing the sum-to-one conditions on the parameter space. The ring map $\psi_{\mathcal{T}}$ is the algebraic counterpart of the parametrization $\Psi_{\mathcal{T}}$ of the model $\mathcal{M}(\mathcal{T})$. Using $\psi_{\mathcal{T}}$, we can write the CStree model as

$$\mathcal{M}(\mathcal{T}) = \mathcal{V}(\ker(\psi_{\mathcal{T}})) \cap \Delta_{|\mathcal{R}|-1}^{\circ}.$$

Remark 6.1.3. The DAG model $\mathcal{M}(G)$ is the set of all the distributions in $\Delta_{|\mathcal{R}|-1}^{\circ}$ that satisfy the recursive factorization property according to G . By [50, Section 2.1] any DAG model $\mathcal{M}(G)$ is a CStree model. The level $k-1$ of the CStree \mathcal{T}_G representing G has one stage $S_{\bar{x}_{\text{pa}(k)}} = \{\bar{x}_{\text{pa}(k)} \bar{y} : \bar{y} \in \mathcal{R}_{[k-1] \setminus \text{pa}(k)}\}$ for each outcome $\bar{x}_{\text{pa}(k)} \in \mathcal{R}_{\text{pa}(k)}$.

For a DAG G , the map $\psi_{\mathcal{T}_G}$ is the well-known recursive factorization according to G . It is important to know that the ideal $I_{\text{global}(G)}$ is not always equal to the prime ideal $\ker(\psi_{\mathcal{T}_G})$. [53] contains several examples where equality holds as well as numerous counterexamples. The strongest possible algebraic characterization of a model is to find the generators for $\ker(\psi_{\mathcal{T}})$. For most graphical models, discrete and Gaussian, it is an open question to find generators of $\ker(\psi_{\mathcal{T}})$. A recent overview of the state of the art is presented in [81, Chapter 3]. For discrete decomposable DAG models such characterization can be found in [55, Theorem 4.4]. Our Theorem 6.4.3 characterizes $\ker(\psi_{\mathcal{T}})$ in terms of CSI statements for all balanced CStree models as defined in 6.1.8.

The next lemma describes the type of CSI statements encoded by a CStree; they are a consequence of condition (3) in Definition 6.1.2.

Lemma 6.1.4. [50, Lemma 3.1] Let \mathcal{T} be a CStree with levels $(L_1, \dots, L_p) \sim (X_1, \dots, X_p)$. Then for any $f \in \mathcal{M}(\mathcal{T})$ and stage $S \subset L_{k-1}$, we have that f entails the CSI statement $X_k \perp\!\!\!\perp X_{[k-1] \setminus C} | X_C = \bar{x}_C$ where C is the set of all indices ℓ such that all elements in S have the same outcome for X_{ℓ} . Hence, $\bar{x}_C = \bar{y}_C$ for any $\bar{y} \in S$.

For any stage S the context $X_C = \bar{x}_C$ in Lemma 6.1.4 is called the *stage-defining context* of the stage S . Given a stage defining context $X_C = \bar{x}_C$ for a stage S in level L_{k-1} we recover the stage from the statement $X_k \perp\!\!\!\perp X_{[k-1] \setminus C} | X_C = \bar{x}_C$ as $S = \bigcup_{\bar{y} \in \mathcal{R}_{[k-1] \setminus C}} \{\bar{x}_C \bar{y}\}$.

CStrees as collections of context DAGs

Let $\mathcal{J}(\mathcal{T})$ be the set of all CSI statements implied by applying the CSI axioms [50, Section 3.2] to the statements in Lemma 6.1.4. By the absorption axiom, there exists a collection $\mathcal{C}_{\mathcal{T}} := \{X_C = \bar{x}_C\}$ of contexts such that for any $X_A \perp\!\!\!\perp X_B | X_S, X_C = \bar{x}_C \in \mathcal{J}(\mathcal{T})$ with $X_C = \bar{x}_C \in \mathcal{C}_{\mathcal{T}}$, there is no subset $T \subseteq C$ for which $X_A \perp\!\!\!\perp X_B | X_{S \cup T}, X_{C \setminus T} = \bar{x}_{C \setminus T} \in \mathcal{J}(\mathcal{T})$. We call such $X_C = \bar{x}_C \in \mathcal{C}_{\mathcal{T}}$ a *minimal context* of \mathcal{T} . Note that

$$\mathcal{J}(\mathcal{T}) = \bigcup_{X_C = \bar{x}_C \in \mathcal{C}_{\mathcal{T}}} \mathcal{J}(X_C = \bar{x}_C)$$

where $\mathcal{J}(X_C = \bar{x}_C)$ is the set of all CI statements of the form $X_A \perp\!\!\!\perp X_B | X_S$ that hold in the context $X_C = \bar{x}_C$. To each such $X_C = \bar{x}_C \in \mathcal{C}_{\mathcal{T}}$, we can associate a *minimal context DAG* with set of nodes $[p] \setminus C$, denoted by $G_{X_C = \bar{x}_C}$, via the I-MAP of $\mathcal{J}(X_C = \bar{x}_C)$ [50, Section 3.2]. An example of a CStree and its corresponding collection of context DAGs is depicted in Figure 6.1.

Each context DAG $G_{X_C = \bar{x}_C}$ is in particular a DAG, thus by Remark 6.1.3, it has a staged tree representation which we denote by $\mathcal{T}_{G_{X_C = \bar{x}_C}}$. To relate $\mathcal{T}_{G_{X_C = \bar{x}_C}}$ to the original tree \mathcal{T} , we define a *context subtree* $\mathcal{T}_{X_C = \bar{x}_C}$ for each context $X_C = \bar{x}_C$. Let $x_1 \cdots x_k \in \mathcal{T}$ and denote by $\mathcal{T}_{x_1 \cdots x_k}$ the directed subtree of \mathcal{T} with root node $x_1 \cdots x_k$. For $C \subset [p]$ and $\bar{x}_C \in \mathcal{R}_C$ we construct the context subtree $\mathcal{T}_{X_C = \bar{x}_C} = (V_{X_C = \bar{x}_C}, E_{X_C = \bar{x}_C})$ by deleting all subtrees $\mathcal{T}_{x_1 \cdots x_k}$ and all edges $x_1 \cdots x_{k-1} \rightarrow x_1 \cdots x_k$ with $x_k \neq \bar{x}_{C \cap k}$, and then contracting the edges $x_1 \cdots x_{k-1} \rightarrow x_1 \cdots x_{k-1}(\bar{x}_C)_k$ for all $x_1 \cdots x_{k-1} \in \mathcal{R}_{[k-1]}$, for all $k \in C$. The single node resulting from this contraction is labeled $x_1 \cdots x_{k-1}(\bar{x}_C)_k$ and it is in the same stage as $x_1 \cdots x_{k-1}\bar{x}_{C \cap \{k\}}$ in \mathcal{T} . All of the other nodes in $\mathcal{T}_{X_C = \bar{x}_C}$ inherit their staging from \mathcal{T} . Note that the context subtree $\mathcal{T}_{X_C = \bar{x}_C}$ is itself a CStree and $\mathcal{T}_{x_1 \cdots x_k}$ is the context subtree $\mathcal{T}_{X_{[k]} = x_1 \cdots x_k}$.

Moreover, let $X_C = \bar{x}_C$ be a minimal context of \mathcal{T} . Then by the construction of $\mathcal{T}_{X_C = \bar{x}_C}$ and $G_{X_C = \bar{x}_C}$, we have

$$\mathcal{J}(\mathcal{T}_{X_C = \bar{x}_C}) = \{X_A \perp\!\!\!\perp X_B | X_S, X_D = \bar{x}_D : X_A \perp\!\!\!\perp X_B | X_S, X_D = \bar{x}_D, X_C = \bar{x}_C \in \mathcal{J}(\mathcal{T})\}$$

with $A, B, S, D \subset [p] \setminus C$, and the CI statements implied by $G_{X_C = \bar{x}_C}$ are

$$\mathcal{J}(X_C = \bar{x}_C) = \{X_A \perp\!\!\!\perp X_B | X_S : X_A \perp\!\!\!\perp X_B | X_S, X_C = \bar{x}_C \in \mathcal{J}(\mathcal{T})\}.$$

This shows that the CI statements implied by the DAG $G_{X_C = \bar{x}_C}$ are exactly the CI statements implied by the CStree $\mathcal{T}_{X_C = \bar{x}_C}$. In general, the CStrees $\mathcal{T}_{G_{X_C = \bar{x}_C}}$ and $\mathcal{T}_{X_C = \bar{x}_C}$ are different, since the CStree $\mathcal{T}_{X_C = \bar{x}_C}$ may imply more CSI statements (see Example 6.3.1). If $\emptyset \in \mathcal{C}_{\mathcal{T}}$ then G_{\emptyset} is a DAG that captures the CI relations implied by \mathcal{T} . When $\emptyset \notin \mathcal{C}_{\mathcal{T}}$, then \mathcal{T} entails no CI relations, in this case we associate to \mathcal{T} the complete DAG on $[p]$ nodes whose directed arrows are in agreement with the causal ordering of \mathcal{T} , we also denote this DAG by G_{\emptyset} .

Example 6.1.5. Consider the binary CStree \mathcal{T} on three random variables given in Figure 6.1. Since the two nodes in level 1 are in the same stage (represented by the same colors), this CStree implies the CI statement $X_1 \perp\!\!\!\perp X_2$. As the nodes 00 and 10 are in the same stage but 01 and 11 are not, we get the CSI statement $X_3 \perp\!\!\!\perp X_1 | X_2 = 0$. Therefore, $\mathcal{J}(\mathcal{T}) = \{X_1 \perp\!\!\!\perp X_2, X_1 \perp\!\!\!\perp X_3 | X_2 = 0\}$. In particular, we see that the minimal contexts are \emptyset and $X_2 = 0$.

Example 6.1.6. For the sake of intuition we present an example of a collection of context DAGs that do *not* define a CStree. Consider the two DAGs $G_{\emptyset} = ([3], \{2 \rightarrow 3\})$ and

$G_{X_1=0} = (\{2, 3\}, \emptyset)$ and assume all random variables are binary. These two DAGs imply the CI relation $X_1 \perp\!\!\!\perp X_{2,3}$ and the CSI relation $X_2 \perp\!\!\!\perp X_3 | X_1 = 0$.

Let \mathcal{T} be a staged tree representing the outcomes of the vector $X_{[3]}$. We will see that the staging of \mathcal{T} cannot be a CStree. The vertices 10 and 00 are in the same stage because in the empty context DAG, G_\emptyset , $f(X_3 | X_{1,2} = 10) = f(X_3 | X_{1,2} = 00)$. Moreover, since $X_2 \perp\!\!\!\perp X_3 | X_1 = 0$, we also have $f(X_3 | X_{1,2} = 00) = f(X_3 | X_{1,2} = 01)$, thus 00 and 01 are also in the same stage. This implies 10, 00, 01 are all in the same stage, which by definition of CStree implies that so are all vertices 00, 01, 10, 11. Thus the CI statement $X_3 \perp\!\!\!\perp X_{1,2}$ holds in the CStree. However, this statement is not implied by the two DAGs.

Balanced CStrees

Decomposable graphical models are a set of graphical models for which the undirected and directed Markov properties coincide. These are characterized in many different ways, combinatorially as chordal UGs or as perfect DAGs, and geometrically as those DAG models that are discrete exponential families [54], also known as toric varieties in the Algebraic Statistics literature [114]. The article [50] suggests the family of balanced staged tree models as a suitable generalization of decomposable DAG models because these models are discrete exponential families. Furthermore, a DAG is perfect if and only if its CStree representation is balanced. Thus we identify the class of balanced CStrees as a good candidate for decomposable models in the context-specific setting. Our main goal is to explore to which extent the properties of decomposable DAG models carry over to the context-specific case.

Definition 6.1.7. A DAG $G = ([p], E)$ is *perfect* if the skeleton of the induced subgraph on the vertices $\text{pa}(k)$ is a complete graph for all $k \in [p]$.

There are several equivalent ways to define a perfect DAG. Another way to characterize a perfect DAG G is to require that its skeleton is chordal and there is no triple u, v, w of vertices such that $u \rightarrow w, v \rightarrow w$ are edges in G but u and v are not adjacent. One can also characterize a perfect DAG via its moral graph, see [78].

Let G be a DAG and let \mathcal{T}_G be the staged tree representation of G . A characterization of perfect DAGs is also available via balanced CStrees.

Definition 6.1.8. Let \mathcal{T} be a CStree. For any vertex $v = x_1 \cdots x_{k-1} \in \mathcal{T}$ we define the polynomial

$$t(v) := \sum_{\bar{z} \in \mathcal{R}_{[p] \setminus [k-1]}} \left(\prod_{e \in E(v \rightarrow v\bar{z})} \theta(e) \right) \in \mathbb{R}[\theta(e) : e \in E].$$

A pair of vertices $v = x_1 \cdots x_{k-1}$ and $w = y_1 \cdots y_{k-1}$ in the same stage is *balanced* if for all $s, r \in [d_k]$, we have the equality

$$t(vs)t(wr) = t(vr)t(ws)$$

in the polynomial ring $\mathbb{R}[\theta(e) : e \in E]$. The tree \mathcal{T} is *balanced* if every pair of vertices in the same stage is balanced.

Remark 6.1.9. The polynomial $t(\text{root})$ in the previous definition is called the *interpolating polynomial* of \mathcal{T} . Such polynomial is useful to study equivalence classes of staged tree models [59] and enumerating the trees in the equivalence class [58] of any given staged tree.

Theorem 6.1.10. [49, Theorem 3.1] The DAG G is perfect if and only if \mathcal{T}_G is balanced if and only if $\mathcal{M}(G)$ is decomposable.

One could hope that the direct generalization of Theorem 6.1.10 is true for balanced CStrees. Namely that a CStree \mathcal{T} is balanced if and only all of its minimal context DAGs are perfect. This equivalence only holds for three random variables ($p = 3$). For $p = 4$, Example 6.2.4 provides a counterexample. In general, only one implication holds, namely, the CStree model $\mathcal{M}(\mathcal{T})$ is balanced whenever all minimal contexts are perfect (Theorem 6.3.5).

6.2 Decomposable CSm models in three variables

Our subject of study from this point forward are balanced CStree models, the combinatorics of their context-specific DAG representations and the properties of their defining equations. The results in Section 6.4 show that the properties of these models closely mirror those of decomposable DAG models. Therefore we will refer to balanced CStree models as *decomposable context-specific models (decomposable CSm models)*. Consistent with our previous notation, we will denote such models by $\mathcal{M}(\mathcal{T})$, where \mathcal{T} denotes the associated balanced CStree.

The goal of this section is to provide a complete description of CStree models in three random variables and to prove the generalization of Theorem 6.1.10 for $p = 3$. That is, we prove the following result.

Theorem 6.2.1. A CStree \mathcal{T} with $p = 3$ is balanced, i.e. $\mathcal{M}(\mathcal{T})$ is a decomposable CSm model, if and only if all minimal context DAGs of \mathcal{T} are perfect.

Before proving the above theorem, we classify all possible CStrees on three random variables along with their minimal contexts, taking advantage of the small value of p .

Example 6.2.2. We provide a list of all CStrees which are not staged tree representations of a DAG with causal ordering 123 (see Figure 6.2). In this case there are four families of CStrees.




Figure 6.2: All CStrees with $p = 3$, which do not represent a DAG.

Consider the two pairs of DAGs $\{([3], \{1 \rightarrow 2, 1 \rightarrow 3, 2 \rightarrow 3\}), ([3], \{1 \rightarrow 3, 2 \rightarrow 3\})\}$ and $\{(\{1, 3\}, \emptyset), (\{2, 3\}, \emptyset)\}$. Each of the four families is defined as follows: Choose two graphs G_1, G_2 , one from each pair. Let $i \in \{1, 2\}$ such that i does not appear in the vertex set of G_2 and let $I \subsetneq [d_i]$. Now, consider the CStree defined by taking G_1 as its empty context DAG and G_2 as the minimal context DAG for the contexts $X_i = j$ for every $j \in I$. Depending on the choice of G_1, G_2 we get exactly the four families in Figure 6.2. The first family for example corresponds to choosing G_1 to be the complete graph, $G_2 = (\{1, 3\}, \emptyset)$ and some $I \subsetneq [d_2]$.

Note that any such choice does define a CStree and all contexts will be minimal contexts (except the empty context if the DAG is chosen to be the complete graph). If one would take $I = [d_i]$ this would no longer be true and the CStree would in fact be the staged tree representation of a DAG. The staged trees on the left of Figure 6.2 are examples in which all random variables are binary and such that the minimal context which is not the empty context is either $X_1 = x_1^1$ or $X_2 = x_2^1$.

In the first two families the empty context is not a minimal context as there are no CI relations that hold. We still draw the complete graph in this example for consistency.

Moreover, we can check that the first two families are balanced CStrees whereas the latter two are not. In the first two cases we also see that all minimal context DAGs are perfect which again is not the case for the latter two (cf. Theorems 6.3.5, 6.2.1).

We will not give a proof here that this is in fact a complete list of CStrees that are not DAGs for $p = 3$. However, as we are especially interested in balanced CStrees, we show in the next theorem that the first two families are in fact the only CStrees on $p = 3$ vertices that are balanced and are not staged tree representations of DAGs.

Lemma 6.2.3. Let \mathcal{T} be a balanced CStree with $p = 3$. If $X_1 \perp\!\!\!\perp X_2$, then \mathcal{T} represents a DAG.

Proof. Since $X_1 \perp\!\!\!\perp X_2$, all vertices in the first level of \mathcal{T} are in the same stage and $\emptyset \in \mathcal{C}_{\mathcal{T}}$. We claim that there are no other minimal contexts, besides the empty one. Since \mathcal{T} is balanced by assumption, for any two vertices v and u in level 1 and $v_i, v_j \in \text{ch}_{\mathcal{T}}(v), u_i, u_j \in \text{ch}_{\mathcal{T}}(u)$ with $\theta(u \rightarrow u_\ell) = \theta(v \rightarrow v_\ell)$ for $\ell \in [d_2]$, we have $t(v_i)t(u_j) = t(v_j)t(u_i)$. Since $p = 3$, we have one of the following cases:

- 1) $t(v_i) = t(v_j) \implies v_i \sim v_j$ and $u_i \sim u_j$
- 2) $t(v_i) \neq t(v_j) \implies u_i \sim v_i$ and $u_j \sim v_j$,

where \sim denotes the equivalence relation of being in the same stage. Since \mathcal{T} is a CStree, the first case implies that all children of any vertex v in level 1 are in the same stage. But then $X_2 \perp\!\!\!\perp X_3 | X_1$, so $X_1 = \ell$ is not a minimal context for any $\ell \in [d_1]$. In the second case, we get that for any two vertices v and u in level 1, we have $v' \sim w'$ for some $v' \in \text{ch}_{\mathcal{T}}(v)$ and $w' \in \text{ch}_{\mathcal{T}}(w)$. But then $X_1 \perp\!\!\!\perp X_3 | X_2$, so again $X_2 = \ell \notin \mathcal{C}_{\mathcal{T}}$ for any $\ell \in [d_2]$. We conclude $\mathcal{C}_{\mathcal{T}} = \{\emptyset\}$, so indeed \mathcal{T} represents a DAG. \square

We are now ready to prove Theorem 6.2.1.

Proof of Theorem 6.2.1. We show in Theorem 6.3.5 that (for any p) if all minimal context DAGs of \mathcal{T} are perfect, then \mathcal{T} is balanced. Hence, it suffices to show the other implication. Let \mathcal{T} be a balanced CStree with $p = 3$. Assume there exists a minimal context DAG G that is not perfect. This has to be the empty context DAG as other minimal context DAGs can only have two vertices. Hence, the empty minimal context DAG is $1 \rightarrow 3 \leftarrow 2$ which implies $X_1 \perp\!\!\!\perp X_2$. Using Lemma 6.2.3 it follows that $\mathcal{T} = \mathcal{T}_G$ is the staged tree representation of G . However, such a CStree is unbalanced by Theorem 6.1.10. \square

We have just observed that every balanced CStree has only perfect minimal contexts DAGs when $p = 3$. This is no longer true for $p \geq 4$, as illustrated by the next example.

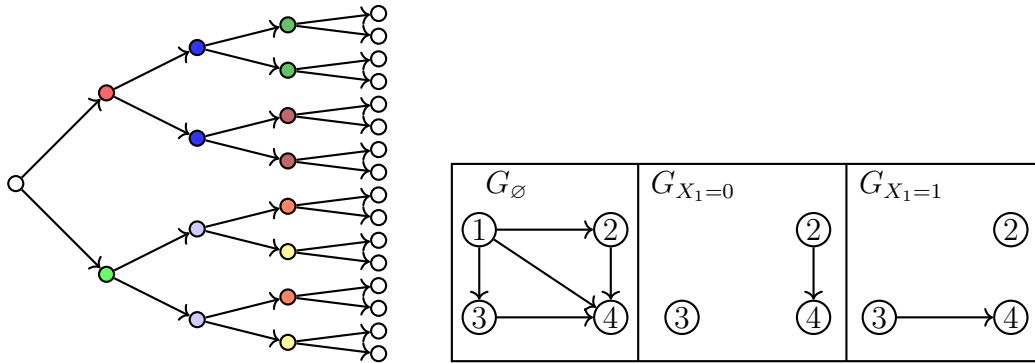


Figure 6.3: A balanced CStree with a non-perfect minimal context.

Example 6.2.4. We consider the binary CStree in Figure 6.3 on $p = 4$ binary random variables which is equivalently given by its three minimal context DAGs. This CStree is balanced as one can check using Definition 6.1.8, but the empty minimal context DAG G_\emptyset is not perfect as the parents of 4 do not form a complete graph. Therefore, a straightforward generalization of Theorem 6.2.1 is not true.

This example can also be generalized to get more counterexamples for any $p \geq 4$ and with an arbitrarily large number of minimal context. We may note that the statement $X_2 \perp\!\!\!\perp X_3 | X_1$ (which prevents the parents of 4 from forming a complete graph) is implied by the other two minimal contexts using absorption. We reveal why this happens in Section 6.4.

Remark 6.2.5. In the case $p = 4$ the CStree in Figure 6.3 is essentially the only binary balanced CStree with a non-perfect minimal context (up to swapping the outcomes of X_1 in the minimal contexts). If we do not restrict to binary CStrees there exists a family of such CStrees with a non-perfect context DAG, it can be constructed similarly to Example 6.2.2.

Remark 6.2.6. Another characterizing property of decomposable graphical is in terms of its maximum likelihood estimator (MLE). Decomposable graphical models are the only class of undirected graphical models whose MLE has closed form [55, Theorem 4.4]. Decomposable CSmodels also have closed form for their MLE, which follows from the fact that they are a subclass of staged tree models [48].

6.3 Combinatorial properties of balanced CStrees

This section studies context subtrees of CStrees to understand which properties of CStrees are preserved when restricting to specific contexts. It turns out that any context subtree of a balanced CStree is itself balanced (Theorem 6.3.4) which can be seen as a generalization of the fact that removing a vertex from a perfect DAG results in a perfect DAG. Moreover, we

saw in Example 6.2.4 that a CStree can be balanced without its minimal context DAGs being perfect. The reverse implication does hold, i.e. if all minimal context DAGs are perfect, then the CStree is balanced (Theorem 6.3.5). The proof is mostly combinatorial in nature and does not make use of algebraic methods other than the definition of balancedness. This section also establishes Proposition 6.3.6 which will be used in the main proof of the last section. It gives an interpretation of the staging of a CStree in terms of CSI statements, as well as combinatorial conditions on minimal context DAGs for stagings to exist.

Context subtrees

While Section 6.1 provides the formal definition of a context subtree, an illustrative example is given below. For any context $X_C = \bar{x}_C$, the subtree $\mathcal{T}_{X_C = \bar{x}_C}$ is a CStree, the DAG $G_{X_C = \bar{x}_C, \emptyset}$ denotes the empty context DAG of $\mathcal{T}_{X_C = \bar{x}_C}$.

Example 6.3.1. We consider the CStree \mathcal{T} in Figure 6.5 and construct the context subtree $\mathcal{T}_{X_3=0}$ given in Figure 6.4. We remove all subtrees with root x_1x_21 and $x_1, x_2 \in \{0, 1\}$ and contract the edges $x_1x_2 \rightarrow x_1x_20$. The stage of the node resulting from this contraction is the stage of the node x_1x_21 . The stages of level 2 do not exist anymore and they do not have any meaning in the construction of the context subtree. We could now construct the minimal context DAGs from this CStree. However, we will instead do this from the minimal context DAGs of the full tree. We check if any minimal context is invalid in the case $X_3 = 0$, i.e. is only valid for $X_3 = 1$, and discard this DAG. This however is not the case here. Now we remove the node 3 from any minimal context DAG, resulting in the collection of DAGs in Figure 6.4. In this case all non-empty contexts are in fact minimal contexts of the context subtree $\mathcal{T}_{X_3=0}$, however this is not true in general.

Moreover, we see that this context subtree $\mathcal{T}_{X_3=0}$ is different from the tree $\mathcal{T}_{G_{X_3=0, \emptyset}}$ (the staged tree representation of the empty context DAG) of which every stage is a singleton.

Lemma 6.3.2. Suppose \mathcal{T} is a CStree with levels $(L_1, \dots, L_p) \sim (X_1, \dots, X_p)$ and let $X_C = \bar{x}_C$ be a context.

- (1) Every stage in $\mathcal{T}_{G_\emptyset}$ is a subset of a stage in \mathcal{T} .
- (2) Suppose $C \subset [p]$ is a context with maximum index k and let $v = x_1 \dots x_q \in V_{\mathcal{T}}$ with $k \leq q$. Then for any $\bar{x}_C \in \mathcal{R}_C$, $(\bar{x}_C)_i = x_i$ ($i \in C$), $\mathcal{T}_v = (\mathcal{T}_{X_C = \bar{x}_C})_v$. Since $t_{\mathcal{T}}(v)$ only depends on the subtree \mathcal{T}_v we see $t_{\mathcal{T}}(v) = t_{\mathcal{T}_{X_C = \bar{x}_C}}(v) \in \mathbb{R}[\Theta_{\mathcal{T}_v}]$.

Proof. (1) Let S be a stage in $\mathcal{T}_{G_\emptyset}$ immediately preceding the level of the variable X_k , $k \in [p]$. Since $\mathcal{T}_{G_\emptyset}$ represents a DAG, the stage defining context $X_A = \bar{x}_A$ of S satisfies $A = \text{pa}_{G_\emptyset}(k)$




Figure 6.4: The context subtree of the tree in Figure 6.5 for the context $X_3 = 0$, and its minimal context DAGs.

and $\bar{x}_A \in \mathcal{R}_A$. Thus, as a subset of vertices in \mathcal{T} ,

$$S = \bigcup_{\bar{y} \in \mathcal{R}_{[k-1] \setminus A}} \{\bar{x}_A \bar{y}\}$$

for some $\bar{x}_A \in \mathcal{R}_A$. By the ordered Markov property in G_\emptyset , G_\emptyset encodes the CI relation $X_k \perp\!\!\!\perp X_{[k-1] \setminus A} | X_A$. This CI statement in G_\emptyset corresponds to the CI statement $X_k \perp\!\!\!\perp X_{[k-1] \setminus A} | X_A$ in \mathcal{T} . Thus, by [50, Theorem 3.3] $X_k \perp\!\!\!\perp X_{[k-1] \setminus A} | X_A$ holds in \mathcal{T} . By specialization to $X_A = \bar{x}_A$, the statement $X_k \perp\!\!\!\perp X_{[k-1] \setminus A} | X_A = \bar{x}_A$ holds in \mathcal{T} . The fact that this latter statement holds in \mathcal{T} , implies that the nodes in S must be a subset of a stage in \mathcal{T} .

(2) The vertices of the two trees are the same. A stage in the tree \mathcal{T}_v is defined by a statement $X_j \perp\!\!\!\perp X_{[j-1] \setminus ([q] \cup D)} | X_D = \bar{x}_D$ for some $j > q$ and $D \subset [j-1] \setminus [q]$. A stage in the tree $(\mathcal{T}_{X_C = \bar{x}_C})_v$ is defined by exactly the same kind of statement since $C \subset [k] \subset [q]$. \square

Lemma 6.3.2 (1) says that every CStree \mathcal{T} is a coarsening of the CStree $\mathcal{T}_{G_\emptyset}$, as every stage of \mathcal{T} is the union of possibly several stages in $\mathcal{T}_{G_\emptyset}$. This coarsening is a result of other minimal context DAGs entailing more CSI statements. Hence, if $\mathcal{T} = \mathcal{T}_{G_\emptyset}$ all CSI statements implied by \mathcal{T} are specializations of CI statements also implied by \mathcal{T} .

We recall a useful lemma to prove balancedness.

Lemma 6.3.3 ([49, Lemma 3.2]). Let $G = ([p], E)$ be a DAG and assume $\pi = 12 \cdots p$ is a linear extension of G . Then \mathcal{T}_G is balanced if and only if for every pair of vertices in the same stage with $v = x_1 \cdots x_i, w = x'_1 \cdots x'_i \in \mathcal{R}_{\{i\}}$, there exists a bijection

$$\begin{aligned} \Phi : \mathcal{R}_{[p] \setminus [i+1]} \times \mathcal{R}_{[p] \setminus [i+1]} &\rightarrow \mathcal{R}_{[p] \setminus [i+1]} \times \mathcal{R}_{[p] \setminus [i+1]} \\ (y_{i+2} \cdots y_p, y'_{i+2} \cdots y'_p) &\mapsto (z_{i+2} \cdots z_p, z'_{i+2} \cdots z'_p) \end{aligned}$$

such that for all $k \geq i + 2$ and all $s \neq r \in [d_{i+1}]$

$$\begin{aligned} & f(y_k \mid (x_1 \cdots x_i, s, y_{i+2} \cdots y_p)_{\text{pa}(k)}) f(y'_k \mid (x'_1 \cdots x'_i, r, y'_{i+2} \cdots y'_p)_{\text{pa}(k)}) \\ &= f(z_k \mid (x'_1 \cdots x'_i, s, z_{i+2} \cdots z_p)_{\text{pa}(k)}) f(z'_k \mid (x_1 \cdots x_i, r, z'_{i+2} \cdots z'_p)_{\text{pa}(k)}). \end{aligned}$$

Theorem 6.3.4. If a CStree \mathcal{T} is balanced, then so is $\overline{\mathcal{T}} := \mathcal{T}_{X_C = \overline{x}_C}$ for every context $X_C = \overline{x}_C$.

Proof. Let $k \in [p] \setminus C$ and suppose $v = x_1 \cdots x_{k-1}$ and $w = y_1 \cdots y_{k-1}$ are two vertices in the same stage in $\overline{\mathcal{T}}$ with $x_i = y_i$ for $i \in C \cap [k-1]$. Note that v, w are also in the same stage in \mathcal{T} since $k \notin C$.

Then their children in \mathcal{T} and $\overline{\mathcal{T}}$ are

$$\begin{aligned} \text{ch}(v) &= \{x_1 \cdots x_{k-1}s : s \in [d_k]\}, \\ \text{ch}(w) &= \{y_1 \cdots y_{k-1}s : s \in [d_k]\}. \end{aligned}$$

Let $v_1, v_2 \in \text{ch}(v)$ and $w_1, w_2 \in \text{ch}(w)$ be such that $\theta(v \rightarrow v_i) = \theta(w \rightarrow w_i)$, ($i = 1, 2$). Since CStrees are compatibly labeled, then

$$\begin{aligned} v_1 &= x_1 \cdots x_{k-1}s, & v_2 &= x_1 \cdots x_{k-1}r, \\ w_1 &= y_1 \cdots y_{k-1}s, & w_2 &= y_1 \cdots y_{k-1}r, \end{aligned}$$

for some $s, r \in [d_k]$. We want to show $t_{\overline{\mathcal{T}}}(v_1)t_{\overline{\mathcal{T}}}(w_2) = t_{\overline{\mathcal{T}}}(w_1)t_{\overline{\mathcal{T}}}(v_2)$. Choose a monomial on the left-hand-side. This corresponds to a product of edge labels of two paths λ'_1 and λ'_2 in $\overline{\mathcal{T}}$, λ'_1 is a path from v_1 to a leaf and λ'_2 is a path from w_2 to a leaf. Each leaf in $\overline{\mathcal{T}}$ is also a leaf in \mathcal{T} (Section 6.1). In \mathcal{T} there exists a directed path λ_1 from v_1 to the leaf using all edges in λ'_1 and a directed path λ_2 from w_2 to the other leaf using λ'_2 .

Since v, w are in the same stage in $\overline{\mathcal{T}}$, they are also in the same stage in \mathcal{T} . The balanced condition in \mathcal{T} implies

$$t_{\mathcal{T}}(x_1 \cdots x_{k-1}s)t_{\mathcal{T}}(y_1 \cdots y_{k-1}r) = t_{\mathcal{T}}(x_1 \cdots x_{k-1}r)t_{\mathcal{T}}(y_1 \cdots y_{k-1}s). \quad (6.2)$$

Choose the product of monomials on the left hand side of this equation that is the product of the edge labels in the concatenation of paths $\lambda_1\lambda_2$ and denote it by $\theta(\lambda_1)\theta(\lambda_2)$. Since \mathcal{T} is balanced, it follows from the bijection in Lemma 6.3.3 that there exists a product $\theta(\lambda_3)\theta(\lambda_4)$ corresponding to paths λ_3, λ_4 in \mathcal{T} from v_2 to a leaf and w_1 to a leaf on the right-hand side of (6.2) such that

$$\theta(\lambda_1)\theta(\lambda_2) = \theta(\lambda_3)\theta(\lambda_4). \quad (6.3)$$

We claim that the paths λ_3, λ_4 are paths in $\overline{\mathcal{T}}$, i.e. the nodes in the paths λ_3, λ_4 contract to nodes in $\overline{\mathcal{T}}$. Let $j \in [p] \setminus [k]$ and denote by $e_{i,j}$ the edge of the path i , ($i = 1, 2, 3, 4$) from level j to level $j + 1$. The fact that \mathcal{T} is stratified and (6.3) holds, implies

$$\theta(e_{1,j})\theta(e_{2,j}) = \theta(e_{3,j})\theta(e_{4,j}) \text{ for all } j \in [p] \setminus [k]. \quad (6.4)$$

If $j+1 \in C$ then the edges $e_{1,j}, e_{2,j}$ point to the same outcome $(\bar{x}_C)_{j+1}$. From (6.4) and since \mathcal{T} is compatibly labeled, different outcomes can never have equal edge labels, thus $e_{3,j}, e_{4,j}$ must also point to the outcome $(\bar{x}_C)_{j+1}$. This shows λ_3, λ_4 are paths in $\bar{\mathcal{T}}$.

Finally, if $j+1 \notin C$ then (6.4) implies that the product of the edge labels of the restrictions λ'_3, λ'_4 of λ_3, λ_4 to paths in $\bar{\mathcal{T}}$ is equal to the product of the edge labels of λ'_1, λ'_2 . This establishes a bijection between terms on the right-hand side and the left-hand side of $t_{\bar{\mathcal{T}}}(v_1)t_{\bar{\mathcal{T}}}(w_2) = t_{\bar{\mathcal{T}}}(w_1)t_{\bar{\mathcal{T}}}(v_2)$, which means $\bar{\mathcal{T}}$ is balanced. \square

Decomposable DAG models and decomposable CSmodels

We start by proving the one-sided generalization of Theorem 6.2.1 and Theorem 6.1.10 to CStrees.

Theorem 6.3.5. Let \mathcal{T} be a CStree with only perfect minimal contexts. Then \mathcal{T} is balanced.

Proof. Let $v = x_1 \dots x_{k-1}, w = y_1 \dots y_{k-1} \in V_{\mathcal{T}}$ be two vertices in the same stage S in \mathcal{T} . Then S has a stage defining context C that entails the CSI relation

$$X_k \perp\!\!\!\perp X_{[k-1] \setminus C} | X_C = \bar{x}_C$$

for some $\bar{x}_C \in \mathcal{R}_C$. By definition $C \subset [k-1]$, and from [50, Lemma 3.2] there exists a minimal context $C' \subset C$ such that

$$X_k \perp\!\!\!\perp X_{[k-1] \setminus C} | X_{C \setminus C'}, X_{C'} = \bar{x}_{C'}$$

holds with $\bar{x}_{C'} = (\bar{x}_C)_{C'}$. Every node in S contains the context $\bar{x}_{C'}$, hence every node in S appears in $\mathcal{T}_{X_{C'} = \bar{x}_{C'}}$, and by construction S is a stage in $\mathcal{T}_{X_{C'} = \bar{x}_{C'}}$. We claim that v and w are also in the same stage in $\mathcal{T}_{G_{X_{C'} = \bar{x}_{C'}}}$:

By [49, Proposition 2.2], this holds if and only if $(v)_{\text{pa}_{G_{X_{C'} = \bar{x}_{C'}}}(k)} = (w)_{\text{pa}_{G_{X_{C'} = \bar{x}_{C'}}}(k)}$. That is, the entries of v and w agree for the indices in $\text{pa}_{G_{X_{C'} = \bar{x}_{C'}}}(k)$. Let $i \in \text{pa}_{G_{X_{C'} = \bar{x}_{C'}}}(k)$ then $X_k \not\perp\!\!\!\perp X_i | X_{C'} = \bar{x}_{C'}$. Therefore $i \notin [k-1] \setminus C$, i.e. $i \in C \setminus C'$ because we are in the context $X_{C'} = \bar{x}_{C'}$. Since C is the stage defining context of S , we have $x_i = y_i$.

Since $G_{X_{C'} = \bar{x}_{C'}}$ is perfect by assumption, the nodes v, w are balanced in the CStree $\bar{\mathcal{T}} := \mathcal{T}_{G_{X_{C'} = \bar{x}_{C'}}}$ by [49, Theorem 3.1]. This means that for any $v_1, v_2 \in \text{ch}_{\bar{\mathcal{T}}}(v)$ and $w_1, w_2 \in \text{ch}_{\bar{\mathcal{T}}}(w)$ with $\theta_{\bar{\mathcal{T}}}(v \rightarrow v_i) = \theta_{\bar{\mathcal{T}}}(w \rightarrow w_i)$, ($i = 1, 2$) the following equation holds

$$t_{\bar{\mathcal{T}}}(v_1)t_{\bar{\mathcal{T}}}(w_2) = t_{\bar{\mathcal{T}}}(v_2)t_{\bar{\mathcal{T}}}(w_1)$$

in $\mathbb{R}[\Theta_{\bar{\mathcal{T}}}]$. Since there is a surjective ring homomorphism $\mathbb{R}[\Theta_{\bar{\mathcal{T}}}] \rightarrow \mathbb{R}[\Theta_{\mathcal{T}_{X_{C'} = \bar{x}_{C'}}}]$ the same equation

$$t_{\mathcal{T}_{X_{C'} = \bar{x}_{C'}}}(v_1)t_{\mathcal{T}_{X_{C'} = \bar{x}_{C'}}}(w_2) = t_{\mathcal{T}_{X_{C'} = \bar{x}_{C'}}}(v_2)t_{\mathcal{T}_{X_{C'} = \bar{x}_{C'}}}(w_1)$$

holds in $\mathbb{R}[\Theta_{\mathcal{T}_{X_{C'}=\bar{x}_{C'}}}]$. By Lemma 6.3.2 (ii) we have $t_{\mathcal{T}}(v) = t_{\mathcal{T}_{X_{C'}=\bar{x}_{C'}}}(v)$ and hence the equality

$$t_{\mathcal{T}}(v_1)t_{\mathcal{T}}(w_2) = t_{\mathcal{T}}(v_2)t_{\mathcal{T}}(w_1)$$

holds in $\mathbb{R}[\Theta_{\mathcal{T}}]$, i.e. the nodes v, w are balanced. \square

Proposition 6.3.6. Let \mathcal{T} be a CStree. Let $A, B, C \subset [p]$ pairwise disjoint with $A \cup B \cup C = [k-1]$. Let $\bar{x}_A \in \mathcal{R}_A, \bar{x}_B \in \mathcal{R}_B, \bar{x}_C \in \mathcal{R}_C$, then the following rule holds for the CSI statements in \mathcal{T} :

$$X_k \perp\!\!\!\perp X_A | X_{B \cup C} = \bar{x}_B \bar{x}_C \text{ and } X_k \perp\!\!\!\perp X_B | X_{A \cup C} = \bar{x}_A \bar{x}_C \Rightarrow X_k \perp\!\!\!\perp X_{A \cup B} | X_C = \bar{x}_C.$$

Proof. In level $k-1$ we have the two stages

$$S_1 = \bigcup_{\bar{y}_A \in \mathcal{R}_A} \{\bar{y}_A \bar{x}_B \bar{x}_C\}, \quad S_2 = \bigcup_{\bar{y}_B \in \mathcal{R}_B} \{\bar{x}_A \bar{y}_B \bar{x}_C\}.$$

However, these are both contained in a single stage: Both contain the element $\bar{x}_A \bar{x}_B \bar{x}_C$. But different stages cannot intersect, hence the two are contained in a single stage S .

Let $\bar{y}_A \neq \bar{x}_A$ and $\bar{y}_B \neq \bar{x}_B$. The elements $\bar{x}_A \bar{y}_B \bar{x}_C$ and $\bar{y}_A \bar{x}_B \bar{x}_C$ are contained in S and therefore $\bar{z}_A \bar{z}_B \bar{x}_C \in S$ for every $\bar{z}_A \in \mathcal{R}_A, \bar{z}_B \in \mathcal{R}_B$. But this means $X_k \perp\!\!\!\perp X_{A \cup B} | X_C = \bar{x}_C$. \square

In terms of context DAGs the last lemma says the following: If in a context DAG $G_{X_C=\bar{x}_C, \emptyset}$ there is an edge $i \rightarrow j$, i.e. $X_i \not\perp\!\!\!\perp X_j | X_C = \bar{x}_C$, but $X_i \perp\!\!\!\perp X_j | X_C = \bar{x}_C, X_{C'} = \bar{x}_{C'}$, then for every $v \in C'$ there is an edge $v \rightarrow j$. The lemma can also be understood as a stronger but slightly different version in CStreets of the intersection axiom

$$X_A \perp\!\!\!\perp X_B | X_{S \cup D}, X_C = \bar{x}_C, X_A \perp\!\!\!\perp X_D | X_{S \cup B}, X_C = \bar{x}_C \Rightarrow X_A \perp\!\!\!\perp X_{B \cup D} | X_S, X_C = \bar{x}_C$$

as it only requires the first two CSI statements to each hold in one context $X_D = \bar{x}_D$ and $X_B = \bar{x}_B$.

Example 6.3.7. Consider a CStree \mathcal{T} with empty minimal context DAG $G_{\emptyset} = ([4], \{1 \rightarrow 2, 2 \rightarrow 3, 2 \rightarrow 4\})$. Lemma 6.3.6 implies that this CStree is in fact the staged tree representation of the DAG G_{\emptyset} . Indeed, there is no vertex with at least two incoming edges which implies that any CSI statement in \mathcal{T} is already a specialization of a CI statement.

Using these observations, one can see that the CStreets given in Example 6.2.2 are in fact all CStreets on $p=3$ variables. To generalize the other implication of Theorem 6.1.10 we use an algebraic approach presented in the next section.

6.4 Algebraic characterization

The core of this section is Theorem 6.4.3 as it provides a complete characterization of the CSI statements that hold in a decomposable CSmodel. It states that for a balanced CStree \mathcal{T} , the polynomials associated to saturated CSI statements are a generating set of the prime ideal $\ker(\psi_{\mathcal{T}})$ that defines $\mathcal{M}(\mathcal{T})$ implicitly. This is precisely the case for perfect DAG models, see [55, Theorem 4.4], which once again highlights the important role that decomposable CSmodels play in generalizing the algebraic properties of single DAGs to collections of DAGs in the context-specific setting. The proof of this result uses the algebraic notion of the toric fiber product, first introduced in [115].

For any collection \mathcal{C} of CSI statements in a CStree \mathcal{T} , recall that the ideal $I_{\mathcal{C}}$ is generated by the polynomials associated to all CSI statements in \mathcal{C} , as defined in Section 6.1.

Setup

Let \mathcal{T} be a balanced CStree, $\overline{\mathcal{T}}$ the subtree of \mathcal{T} up to level $p - 1$ and S_1, \dots, S_r the stages in \mathcal{T} in level $p - 1$. Let $\mathcal{T}_p = \bigcup_{i \in [r]} \mathcal{B}_i$, where each \mathcal{B}_i is a one-level tree together with its edge labels as in [13, Section 3]. Consider the rings

$$\begin{aligned} \mathbb{R}[\overline{\mathcal{T}}] &:= \mathbb{R}[p_{\overline{x}}^i : i \in [r], \overline{x} \in S_i], \\ \mathbb{R}[\mathcal{T}_p] &:= \mathbb{R}[p_k^i : i \in [r], k \in [d_p]], \\ \mathbb{R}[\mathcal{T}] &:= \mathbb{R}[p_{\overline{x}k}^i : i \in [r], \overline{x} \in S_i, k \in [d_p]] \end{aligned}$$

with multigrading $\deg(p_{\overline{x}}^i) = \deg(p_k^i) = \deg(p_{\overline{x}k}^i), i \in [r], \overline{x} \in S_i, k \in [d_p]$ where $\mathcal{A} = \{e_1, \dots, e_r\}$ and e_i is the i -th standard unit vector in \mathbb{Z}^r . Note that the rings $\mathbb{R}[\mathcal{T}]$ and $\mathbb{R}[D]$ are the same, except the former is multigraded. Consider the ring homomorphism

$$\mathbb{R}[\mathcal{T}] \rightarrow \mathbb{R}[\overline{\mathcal{T}}] \otimes \mathbb{R}[\mathcal{T}_p], \quad p_{\overline{x}k}^i \mapsto p_{\overline{x}}^i \otimes p_k^i \quad (i \in [r], \overline{x} \in S_i, k \in [d_p]).$$

Following [115], we call the kernel of this map Quad. It is given by

$$\text{Quad} = \langle p_{\overline{x}k_1}^i p_{\overline{y}k_2}^i - p_{\overline{x}k_2}^i p_{\overline{y}k_1}^i : k_1 \neq k_2 \in [d_p], \overline{x}, \overline{y} \in S_i, i \in [r] \rangle.$$

Note that the generators of Quad are the 2×2 minors of the matrices $(p_{\overline{x}k}^i)_{\overline{x} \in S_i, k \in [d_p]}$ for all $i \in [r]$. Now, consider the ring homomorphism

$$\mathbb{R}[\mathcal{T}] \rightarrow \mathbb{R}[\overline{\mathcal{T}}] / \ker(\psi_{\overline{\mathcal{T}}}) \otimes \mathbb{R}[\mathcal{T}_p], \quad p_{\overline{x}k}^i \mapsto p_{\overline{x}}^i \otimes p_k^i \quad (i \in [r], \overline{x} \in S_i, k \in [d_p]).$$

The kernel of this map is the *toric fiber product* of $\ker(\psi_{\overline{\mathcal{T}}})$ and the zero ideal $\langle 0 \rangle \subseteq \mathbb{R}[\mathcal{T}_p]$, and is denoted by $\ker(\psi_{\overline{\mathcal{T}}}) \times_{\mathcal{A}} \langle 0 \rangle$. By [13, Proposition 3.5], this toric fiber product is equal to $\ker(\psi_{\mathcal{T}})$ when \mathcal{T} is balanced. The generators of $\ker(\psi_{\mathcal{T}})$ are obtained from two sets, namely $\ker(\psi_{\mathcal{T}}) = \langle \text{Quad}, \text{Lift}(F) \rangle$, where F is a set of generators of $\ker(\psi_{\overline{\mathcal{T}}})$ and

$$\text{Lift}(F) := \{p_{\overline{x}_1 k_1}^i p_{\overline{y}_1 k_2}^j - p_{\overline{x}_2 k_1}^i p_{\overline{y}_2 k_2}^j : \overline{x}_1, \overline{x}_2 \in S_i, \overline{y}_1, \overline{y}_2 \in S_j, k_1, k_2 \in [d_p], p_{\overline{x}_1}^i p_{\overline{y}_1}^j - p_{\overline{x}_2}^i p_{\overline{y}_2}^j \in F\}.$$

Note that the construction above relies heavily on the fact that \mathcal{T} is balanced since this is the only case in which $\ker(\psi_{\mathcal{T}})$ and $\ker(\psi_{\overline{\mathcal{T}}})$ are toric and \mathcal{A} -homogeneous.




Figure 6.5: Balanced CStree on five binary random variables whose empty context DAG is not perfect.

Main results

In what follows saturated CSI statements will be the main actors. Let \mathcal{T} be a CStree with p levels and let $\mathcal{M}(\mathcal{T})$ be the associated model. A *saturated CSI statement* is a CSI statement of the form $X_A \perp\!\!\!\perp X_B | X_S, X_C = \bar{x}_C$, where $A \cup B \cup C \cup S = [p]$. If \mathcal{C} is a collection of saturated CSI statements then the ideal $I_{\mathcal{C}}$ is generated by binomials. Any ideal that is generated by binomials and in addition is prime is a *toric ideal*.

Definition 6.4.1. Let \mathcal{C} be any collection of CSI statements of random variables X_1, \dots, X_p . We define $\text{Sat}(\mathcal{C})$ to be the set of all saturated CSI statements in \mathcal{C} . For a CStree \mathcal{T} let $\text{Sat}(\mathcal{T}) := \text{Sat}(\mathcal{J}(\mathcal{T}))$ where $\mathcal{J}(\mathcal{T})$ denotes the set of all CSI statements that hold in \mathcal{T} . For a DAG G we define $\text{Sat}(G) := \text{Sat}(\mathcal{T}_G)$. Since $\mathcal{J}(\mathcal{T}_G) = \text{global}(G)$, we also get $\text{Sat}(G) = \text{Sat}(\text{global}(G))$.

The proofs of the results in this section rely heavily on the toric fiber product construction. We motivate these results with the following concrete example.

Example 6.4.2 (Decomposable CSmodel with a non-perfect empty context). Consider the decomposable CSmodel given by the CStree in Figure 6.5 for $p = 5$. It has four minimal contexts, namely,

$$\mathcal{C}_{\mathcal{T}} = \{\emptyset, X_1 = 1, X_1X_2 = 01, X_2 = 0\}.$$

Only the non-empty minimal contexts are perfect, yet the tree is balanced. The CSI statements corresponding to the four minimal contexts are, respectively,

$$X_3 \perp\!\!\!\perp X_4 | X_1X_2, X_4 \perp\!\!\!\perp X_5 | X_2X_3, X_1 = 1, X_4 \perp\!\!\!\perp X_5 | X_3, X_1X_2 = 01, X_4 \perp\!\!\!\perp X_5 | X_1X_3, X_2 = 0.$$

The last three statements, corresponding to the three perfect minimal contexts, are saturated. Contracting each of these statements with the statement $X_3 \perp\!\!\!\perp X_4 | X_1 X_2$ (corresponding to the empty context), we get the following three saturated statements

$$X_3 X_5 \perp\!\!\!\perp X_4 | X_2, X_1 = 1, \quad X_3 X_5 \perp\!\!\!\perp X_4 | X_1 X_2 = 01, \quad X_3 X_5 \perp\!\!\!\perp X_4 | X_1, X_2 = 0.$$

These three saturated statements give rise to 24 polynomials, 8 of which coincide with stage-defining statements for level 5. These 8 polynomials, one of which is

$$p_{00000}^1 p_{00011}^1 - p_{00001}^1 p_{00010}^1,$$

are precisely the polynomials in Quad. The remaining 16 polynomials, one of which is

$$p_{00000}^1 p_{00110}^2 - p_{00010}^1 p_{00100}^2,$$

are the polynomials in $\text{Lift}(F)$, defined above. Hence, the 24 polynomials associated to the saturated statements are precisely the generators of $\ker(\psi_{\mathcal{T}})$.

Turns out, the phenomenon in the example above can be generalized to all decomposable CSmodels. The next theorem is the technical foundation of this chapter. It demonstrates the important role that saturated CSI statements play in the algebra of decomposable CSmodels and it also contains most of the technical work in its proof.

Theorem 6.4.3. If $\mathcal{M}(\mathcal{T})$ is a decomposable CSmodel, then $\ker(\psi_{\mathcal{T}})$ is generated by the quadratic binomials associated to all saturated CSI statements in $\mathcal{J}(\mathcal{T})$, i.e.

$$\ker(\psi_{\mathcal{T}}) = I_{\text{Sat}(\mathcal{T})}.$$

Proof. The containment $I_{\text{Sat}(\mathcal{T})} \subset \ker(\psi_{\mathcal{T}})$ holds because all polynomials associated to statements in $\mathcal{J}(\mathcal{T})$ belong to $\ker(\psi_{\mathcal{T}})$. In particular, all binomials coming from saturated CSI statements are in $\ker(\psi_{\mathcal{T}})$.

For the other containment we proceed by induction on the number of random variables in \mathcal{T} . The statement is trivially true for $p = 1, 2$. Suppose that \mathcal{T} has p levels, and any balanced CStree with less than p levels satisfies the statement. Let $\overline{\mathcal{T}}$ be the subtree of \mathcal{T} up to level $p - 1$. Then $\overline{\mathcal{T}}$ is balanced, thus by induction hypothesis $\ker(\psi_{\overline{\mathcal{T}}})$ is generated by a set F of binomials associated to saturated CSI statements in the variables $X_{[p-1]}$. Moreover,

$$\ker(\psi_{\mathcal{T}}) = \ker(\psi_{\overline{\mathcal{T}}}) \times_{\mathcal{A}} \langle 0 \rangle = \langle \text{Quad}, \text{Lift}(F) \rangle.$$

Hence, it suffices to prove that Quad and $\text{Lift}(F)$ are polynomials associated to saturated CSI statements in the variables $X_{[p]}$. Let S_1, \dots, S_r be the stages of level $p - 1$ in \mathcal{T} . For all $i \in [r]$, let $X_{C_i} = \overline{x}_{C_i}$ be the stage defining context of the stage S_i . Recall that

$$\text{Quad} = \langle p_{\overline{x}_{k_1} k_2}^i p_{\overline{y}_{k_2} k_1}^i - p_{\overline{x}_{k_2} k_1}^i p_{\overline{y}_{k_1} k_2}^i : k_1 \neq k_2 \in [d_p], \overline{x}, \overline{y} \in S_i, i \in [r] \rangle.$$

which, by (6.1), is precisely the set of binomials associated to the saturated CSI statements $X_p \perp\!\!\!\perp X_{[p-1]\setminus C_i} | X_{C_i} = \bar{x}_{C_i}$ for all $i \in [r]$.

Since \mathcal{T} is balanced, F is a set of \mathcal{A} -homogeneous binomials. Let $g \in F$, then it is associated to a CSI statement $X_A \perp\!\!\!\perp X_B | X_D, X_C = \bar{x}_C$ with $A \cup B \cup C \cup D = [p-1]$ and $\bar{x}_C \in \mathcal{R}_C$ as in (6.1).

Choose $\bar{y}_A, \bar{y}'_A \in \mathcal{R}_A$ and $\bar{y}_B, \bar{y}'_B \in \mathcal{R}_B$ such that for all $i \in A$, $(\bar{y}_A)_i \neq (\bar{y}'_A)_i$, and for all $i \in B$, $(\bar{y}_B)_i \neq (\bar{y}'_B)_i$. Consider the polynomial

$$h = p_{\bar{y}_A \bar{y}_B \bar{x}_C \bar{x}_D}^k p_{\bar{y}'_A \bar{y}'_B \bar{x}_C \bar{x}_D}^\ell - p_{\bar{y}'_A \bar{y}_B \bar{x}_C \bar{x}_D}^m p_{\bar{y}_A \bar{y}'_B \bar{x}_C \bar{x}_D}^n$$

associated to the same CSI statement as g and its lift

$$h_{z_1, z_2} = p_{\bar{y}_A \bar{y}_B \bar{x}_C \bar{x}_D z_1}^k p_{\bar{y}'_A \bar{y}'_B \bar{x}_C \bar{x}_D z_2}^\ell - p_{\bar{y}'_A \bar{y}_B \bar{x}_C \bar{x}_D z_1}^m p_{\bar{y}_A \bar{y}'_B \bar{x}_C \bar{x}_D z_2}^n, \quad z_1, z_2 \in [d_p].$$

Since h is \mathcal{A} -homogeneous, either $(k, \ell) = (m, n)$ or $(k, \ell) = (n, m)$. Assume it is the former. By the assignment of the grading, it follows that $(\bar{y}_A \bar{y}_B \bar{x}_C \bar{x}_D z_1)_{C_k} = (\bar{y}'_A \bar{y}'_B \bar{x}_C \bar{x}_D z_1)_{C_k}$ because they are in the same stage, therefore $C_k \cap A = \emptyset$. The CSI statement associated to the stage S_k is $X_p \perp\!\!\!\perp X_{[p-1]\setminus C_k} | X_{C_k} = \bar{x}_{C_k}$, this entails $X_p \perp\!\!\!\perp X_A | X_B = \bar{y}_B, X_{C \cup D} = \bar{x}_C \bar{x}_D$ because $C_k \cap A = \emptyset$.

For every $\bar{z}_B, \bar{z}'_B \in \mathcal{R}_B$ there exist $\alpha, \beta \in [r]$ such that the grading is either

$$p_{\bar{y}_A \bar{z}_B \bar{x}_C \bar{x}_D z_1}^\alpha p_{\bar{y}'_A \bar{z}'_B \bar{x}_C \bar{x}_D z_2}^\beta - p_{\bar{y}'_A \bar{z}_B \bar{x}_C \bar{x}_D z_1}^\alpha p_{\bar{y}_A \bar{z}'_B \bar{x}_C \bar{x}_D z_2}^\beta$$

or

$$p_{\bar{y}_A \bar{z}_B \bar{x}_C \bar{x}_D z_1}^\alpha p_{\bar{y}'_A \bar{z}'_B \bar{x}_C \bar{x}_D z_2}^\beta - p_{\bar{y}'_A \bar{z}_B \bar{x}_C \bar{x}_D z_1}^\beta p_{\bar{y}_A \bar{z}'_B \bar{x}_C \bar{x}_D z_2}^\alpha.$$

Case 1: For every \bar{z}_B and \bar{z}'_B entry-wise different, we have the first grading. Then by the same argument as for \bar{y}_B, \bar{y}'_B we get the saturated CSI statement $X_p \perp\!\!\!\perp X_A | X_B = \bar{z}_B, X_{C \cup D} = \bar{x}_C \bar{x}_D$ for all \bar{z}_B . Hence, by absorption we get

$$X_p \perp\!\!\!\perp X_A | X_B, X_{C \cup D} = \bar{x}_C \bar{x}_D.$$

Contracting this statement with $X_A \perp\!\!\!\perp X_B | X_D, X_C = \bar{x}_C$ we get the saturated CSI statement

$$X_A \perp\!\!\!\perp X_{B \cup \{p\}} | X_{C \cup D} = \bar{x}_C \bar{x}_D.$$

This statement entails all binomials in $\text{Lift}(g)$, equivalently $\text{Lift}(g) \subset I_{X_A \perp\!\!\!\perp X_{B \cup \{p\}} | X_{C \cup D} = \bar{x}_C \bar{x}_D}$.

Case 2: There exists a pair \bar{z}_B, \bar{z}'_B , entry-wise different, such that the binomial has the second grading. Using the same argument as above with B instead of A , this implies the statement $X_p \perp\!\!\!\perp X_B | X_A = \bar{y}_A, X_{C \cup D} = \bar{x}_C \bar{x}_D$. Combining this statement with $X_p \perp\!\!\!\perp X_A | X_B = \bar{y}_B, X_{C \cup D} = \bar{x}_C \bar{x}_D$ and Proposition 6.3.6 we get

$$X_p \perp\!\!\!\perp X_{A \cup B} | X_{C \cup D} = \bar{x}_C \bar{x}_D.$$

By the weak union axiom, we get $X_p \perp\!\!\!\perp X_A | X_B, X_{C \cup D} = \bar{x}_C \bar{x}_D$. As in Case 1, we obtain the CSI statement

$$X_A \perp\!\!\!\perp X_{B \cup \{p\}} | X_{C \cup D} = \bar{x}_C \bar{x}_D.$$

and the conclusion $\text{Lift}(g) \subset I_{X_A \perp\!\!\!\perp X_{B \cup \{p\}} | X_{C \cup D} = \bar{x}_C \bar{x}_D}$ follows. The proof for the second choice of grading $(k, \ell) = (n, m)$ of h is analogous, by swapping the roles of A and B in the above arguments. \square

For the rest of this section, we can relax the assumption of working with minimal contexts. Let \mathcal{T} be a CStree and \mathcal{C} be any collection of contexts with associated DAGs $G_{X_C = \bar{x}_C}$, $X_C = \bar{x}_C \in \mathcal{C}$, such that $\mathcal{J}(\mathcal{T}) = \cup_{X_C = \bar{x}_C \in \mathcal{C}} \text{global}(G_{X_C = \bar{x}_C})$. That is, assume that the CSI statements that hold in \mathcal{C} are the same CSI statements that hold in \mathcal{T} . The collection of minimal contexts is one such choice for \mathcal{C} .

Corollary 6.4.4. Let $\mathcal{M}(\mathcal{T})$ be a decomposable CSmodel. The ideal $\ker(\psi_{\mathcal{T}})$ is generated by the binomials associated to all saturated CSI statements that hold in the context DAGs $G_{X_C = \bar{x}_C}$, $X_C = \bar{x}_C \in \mathcal{C}$, i.e.

$$\ker(\psi_{\mathcal{T}}) = \sum_{X_C = \bar{x}_C \in \mathcal{C}} I_{\text{Sat}(G_{X_C = \bar{x}_C})}.$$

Proof. This follows from the fact that $\mathcal{J}(\mathcal{T}) = \cup_{X_C = \bar{x}_C \in \mathcal{C}} \mathcal{J}(X_C = \bar{x}_C)$ and Theorem 6.4.3. \square

Corollary 6.4.5. Let $\mathcal{M}(\mathcal{T})$ be a decomposable CSmodel. Then

$$\ker(\psi_{\mathcal{T}}) = \sum_{X_C = \bar{x}_C \in \mathcal{C}} I_{\text{global}(G_{X_C = \bar{x}_C})}.$$

Proof. We show the following chain of inclusions

$$\ker(\psi_{\mathcal{T}}) = \sum_{X_C = \bar{x}_C \in \mathcal{C}} I_{\text{Sat}(G_{X_C = \bar{x}_C})} \subseteq \sum_{X_C = \bar{x}_C \in \mathcal{C}} I_{\text{global}(G_{X_C = \bar{x}_C})} \subseteq \ker(\psi_{\mathcal{T}}),$$

which implies the theorem. The equality follows from Corollary 6.4.4 and the middle inclusion follows from the containment $\text{Sat}(G_{X_C = \bar{x}_C}) \subseteq \text{global}(G_{X_C = \bar{x}_C})$ for all $X_C = \bar{x}_C \in \mathcal{C}$. For the last inclusion, let $J := \sum_{X_C = \bar{x}_C \in \mathcal{C}} I_{\text{global}(G_{X_C = \bar{x}_C})}$. From [50, Theorem 3.3], we know the equality

$$\mathcal{V}(J) \cap \Delta_{|\mathcal{R}|-1}^{\circ} = \mathcal{V}(\ker(\psi_{\mathcal{T}})) \cap \Delta_{|\mathcal{R}|-1}^{\circ} = \mathcal{M}(\mathcal{T}).$$

Since $\ker \psi_{\mathcal{T}}$ is a prime ideal, this implies that

$$\begin{aligned} J &\subseteq \mathcal{I}(\mathcal{V}(J) \cap \Delta_{|\mathcal{R}|-1}^{\circ}) \cap \mathbb{R}[D] \\ &= \mathcal{I}(\mathcal{V}(\ker(\psi_{\mathcal{T}})) \cap \Delta_{|\mathcal{R}|-1}^{\circ}) \cap \mathbb{R}[D] \\ &= \mathcal{I}(\mathcal{V}(\ker \psi_{\mathcal{T}})) \cap \mathbb{R}[D] = \ker \psi_{\mathcal{T}} \end{aligned}$$

where $\mathcal{I}(V)$ denotes the set of polynomials in $\mathbb{C}[D]$ that vanish on a set $V \subseteq \mathcal{C}^{|\mathcal{R}|}$. \square

Conclusion. In this chapter, we introduced a family of discrete decomposable context-specific models, which we constructed from the subclass of staged tree models known as CStree models. We gave an algebraic and combinatorial characterization of all context specific independence relations that hold in a decomposable context-specific model, which yields a Markov basis. More generally, we established that several algebraic, combinatorial, and geometric properties of decomposable context-specific models generalize those of decomposable graphical models to the context-specific setting.

Part III

Nonparametric algebraic statistics

Chapter 7

Moment varieties for mixtures of products

In this chapter, we introduce moment varieties of conditionally independent mixture distributions on \mathbb{R}^n . There will be no assumptions on the constituent distributions in the product, so the setup in this chapter is *nonparametric*. First, we introduce these moment varieties and discuss how some familiar varieties from statistics arise as examples. We then present several results revolving around the dimensions, degrees, defining polynomials, and finiteness properties of our general moment varieties. We focus on both toric varieties and their secants. This chapter is based on [7].

Consider n independent random variables X_1, X_2, \dots, X_n on the line \mathbb{R} . We make no assumptions about the X_k other than that their moments $\mu_{ki} = \mathbb{E}(X_k^i)$ exist. Then, by [24, Theorem 30.1], the random variable X_k is uniquely characterized by its sequence of moments $\mu_{k1}, \mu_{k2}, \mu_{k3}, \dots$. These moments satisfy the Hamburger moment condition, which states that

$$\begin{bmatrix} \mu_{k0} & \mu_{k1} & \mu_{k2} & \cdots \\ \mu_{k1} & \mu_{k2} & \mu_{k3} & \cdots \\ \mu_{k2} & \mu_{k3} & \mu_{k4} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} \text{ is positive semi-definite for all } k. \quad (7.1)$$

For us, the μ_{ki} are unknowns. The only equations we require are $\mu_{k0} = 1$ for $k = 1, 2, \dots, n$.

We write $m_{i_1 i_2 \dots i_n}$ for the moments of the random vector $X = (X_1, X_2, \dots, X_n)$. The moments are the expected values of the monomials $X_1^{i_1} X_2^{i_2} \cdots X_n^{i_n}$. By independence, we have

$$m_{i_1 i_2 \dots i_n} = \mathbb{E}(X_1^{i_1} X_2^{i_2} \cdots X_n^{i_n}) = \mathbb{E}(X_1^{i_1}) \mathbb{E}(X_2^{i_2}) \cdots \mathbb{E}(X_n^{i_n}) = \mu_{1i_1} \mu_{2i_2} \cdots \mu_{ni_n}. \quad (7.2)$$

We examine this squarefree monomial parametrization for the moments with $i_1 + i_2 + \cdots + i_n = d$. Its image is a toric variety $\mathcal{M}_{n,d}$ in the projective space $\mathbb{P}^{\binom{n+d-1}{d}-1}$ of symmetric tensors.

Example 7.0.1 ($n = d = 3$). The moment variety $\mathcal{M}_{3,3}$ is defined by the monomial parametrization $m_{i_1 i_2 i_3} = \mu_{1i_1} \mu_{2i_2} \mu_{3i_3}$ for $i_1 + i_2 + i_3 = 3$. We find that $\mathcal{M}_{3,3}$ is a cubic hypersurface in the space \mathbb{P}^9 of symmetric $3 \times 3 \times 3$ tensors. It is defined by $m_{012} m_{120} m_{201} = m_{021} m_{210} m_{102}$.

This chapter studies mixtures of r independent distributions. The associated moment variety $\sigma_r(\mathcal{M}_{n,d})$ is the r th secant variety of the toric variety $\mathcal{M}_{n,d}$. It is parametrized by

$$m_{i_1 i_2 \dots i_n} = \sum_{j=1}^r \mu_{1i_1}^{(j)} \mu_{2i_2}^{(j)} \cdots \mu_{ni_n}^{(j)} \quad \text{where } i_1, i_2, \dots, i_n \geq 0 \text{ and } i_1 + i_2 + \cdots + i_n = d. \quad (7.3)$$

These are the moment varieties in this chapter. Mixture weights can be omitted in (7.3) since we work in projective geometry.

We study these and their images under certain coordinate projections $\mathbb{P}^{\binom{n+d-1}{d}-1} \dashrightarrow \mathbb{P}^{|N_\lambda|-1}$. Here λ is any partition of d , and N_λ is the set of moments $m_{i_1 i_2 \dots i_n}$ where $\{i_1, i_2, \dots, i_n\} \setminus \{0\}$ equals λ as a multiset. The images in $\mathbb{P}^{|N_\lambda|-1}$ of the restricted parametrizations (7.2) and (7.3) are denoted by $\mathcal{M}_{n,\lambda}$ and $\sigma_r(\mathcal{M}_{n,\lambda})$. The restricted varieties make sense for statistics because they refer to subclasses of moments that are natural when inferring parameters. We note that $\mathcal{M}_{n,\lambda}$ is also a toric variety and $\sigma_r(\mathcal{M}_{n,\lambda})$ is its r th secant variety.

Example 7.0.2 ($n = 5, d = 3$). There are three partitions $\lambda = (111)$, $\lambda = (21)$ and $\lambda = (3)$. The number of moments $m_{i_1 i_2 i_3 i_4 i_5}$ equals $\binom{5+3-1}{3} = 35$. This is the sum of $N_{(111)} = 10$, $N_{(21)} = 20$ and $N_{(3)} = 5$. The following three toric varieties have dimensions 4, 8, 4 respectively:

$$\begin{aligned} \mathcal{M}_{5,(111)} \subset \mathbb{P}^9 & : & m_{11100} = \mu_{11} \mu_{21} \mu_{31}, m_{11010} = \mu_{11} \mu_{21} \mu_{41}, \dots, m_{00111} = \mu_{31} \mu_{41} \mu_{51}. \\ \mathcal{M}_{5,(21)} \subset \mathbb{P}^{19} & : & m_{21000} = \mu_{12} \mu_{21}, m_{12000} = \mu_{11} \mu_{22}, m_{20100} = \mu_{12} \mu_{31}, \dots, m_{00012} = \mu_{41} \mu_{52}. \\ \mathcal{M}_{5,(3)} = \mathbb{P}^4 & : & m_{30000} = \mu_{13}, m_{03000} = \mu_{23}, \dots, m_{00003} = \mu_{53}. \end{aligned}$$

Combining these parametrizations yields the 14-dimensional moment variety $\mathcal{M}_{5,3} \subset \mathbb{P}^{34}$. We will discuss the ideals of these toric varieties and their secant varieties later on.

This chapter is a sequel to the work of Zhang and Kileel in [125]. That article takes an applied data science perspective and it offers numerical algorithms for learning the parameters $\mu_{ki}^{(j)}$ from empirical moments $m_{i_1 i_2 \dots i_n}$. The primary focus in [125] lies on numerical tensor methods for this recovery task. A key ingredient for their approach is identifiability, which means that the dimension of the moment variety $\sigma_r(\mathcal{M}_{n,\bullet})$ matches the number of free parameters.

The work in this chapter lies at the interface of computer algebra and nonparametric statistics. Our set-up is nonparametric in the sense that no model assumptions are made on the

constituent random variables on \mathbb{R} . Conditional independence arises by passing via (7.3) to multivariate distributions on \mathbb{R}^n . This imposes semialgebraic constraints on the moments $m_{i_1 i_2 \dots i_n}$. We disregard the inequalities in (7.1) and focus on polynomial equations. This leads us to projective varieties, as is customary in algebraic statistics [114]. Their defining equations provide test statistics for mixtures of products [46, Section 3].

7.1 Familiar varieties

This section demonstrates the wide range of interesting models that are featured in this chapter. The scope includes themes from the early days of algebraic statistics: factor analysis [46] and permutation data [40, Section 6]. For the special partition $\lambda = (1, 1, \dots, 1)$ one obtains the toric ideals associated with hypersimplices.

The study of highly structured projective varieties is a main theme in algebraic statistics. This includes varieties of discrete probability distributions as well as moment varieties of continuous distributions; see e.g. [9, 75]. Note that Veronese varieties fall into both categories. In this section we match some of our moment varieties $\sigma_r(\mathcal{M}_{n,\bullet})$ with the existing literature.

We begin with $d = 2$, so each (i_1, i_2, \dots, i_n) has at most two non-zero entries. We change notation so that the second moments are the entries of the $n \times n$ covariance matrix (m_{ij}) .

Example 7.1.1 ($n = 5, d = 2$). The parametrization of the toric variety $\mathcal{M}_{5,2}$ given in (7.2) is written in matrix form as

$$\begin{bmatrix} m_{11} & m_{12} & m_{13} & m_{14} & m_{15} \\ m_{12} & m_{22} & m_{23} & m_{24} & m_{25} \\ m_{13} & m_{23} & m_{33} & m_{34} & m_{35} \\ m_{14} & m_{24} & m_{34} & m_{44} & m_{45} \\ m_{15} & m_{25} & m_{35} & m_{45} & m_{55} \end{bmatrix} = \begin{bmatrix} \mu_{12} & \mu_{11}\mu_{21} & \mu_{11}\mu_{31} & \mu_{11}\mu_{41} & \mu_{11}\mu_{51} \\ \mu_{11}\mu_{21} & \mu_{22} & \mu_{21}\mu_{31} & \mu_{21}\mu_{41} & \mu_{21}\mu_{51} \\ \mu_{11}\mu_{31} & \mu_{21}\mu_{31} & \mu_{32} & \mu_{31}\mu_{41} & \mu_{31}\mu_{51} \\ \mu_{11}\mu_{41} & \mu_{21}\mu_{41} & \mu_{31}\mu_{41} & \mu_{42} & \mu_{41}\mu_{51} \\ \mu_{11}\mu_{51} & \mu_{21}\mu_{51} & \mu_{31}\mu_{51} & \mu_{41}\mu_{51} & \mu_{52} \end{bmatrix}.$$

This shows that $\mathcal{M}_{5,2}$ is the join in \mathbb{P}^{14} of the projective space $\mathcal{M}_{5,(2)} = \mathbb{P}^4$ with coordinates $m_{ii} = \mu_{i2}$ and the 4-dimensional toric variety $\mathcal{M}_{5,(11)} \subset \mathbb{P}^9$ given by the off-diagonal entries. The toric fourfold $\mathcal{M}_{5,(11)}$ has degree 11. This is visualized in [108, Figure 9-2]. Its ideal is generated by ten binomial quadrics, like $m_{12}m_{34} - m_{13}m_{24}$, as shown in [108, Equation (9.2)].

Passing from $r = 1$ to $r = 2$, we note that the secant variety has the join decomposition

$$\sigma_2(\mathcal{M}_{5,2}) = \sigma_2(\mathbb{P}^4 \star \mathcal{M}_{5,(11)}) = \mathbb{P}^4 \star \sigma_2(\mathcal{M}_{5,(11)}) \subset \mathbb{P}^4 \star \mathbb{P}^9 = \mathbb{P}^{14}.$$

The star denotes the join of projective varieties, which arises from the Minkowski sum of the corresponding affine cones.

The prime ideal of the model $\sigma_2(\mathcal{M}_{5,2})$ is found by eliminating the diagonal entries m_{ii} from the ideal of 3×3 -minors of the symmetric 5×5 matrix (m_{ij}) . The elimination ideal is principal, and its generator is the polynomial

$$\begin{aligned} & m_{12}m_{13}m_{24}m_{35}m_{45} - m_{12}m_{13}m_{25}m_{34}m_{45} - m_{12}m_{14}m_{23}m_{35}m_{45} + m_{12}m_{14}m_{25}m_{34}m_{35} \\ & + m_{12}m_{15}m_{23}m_{34}m_{45} - m_{12}m_{15}m_{24}m_{34}m_{35} + m_{13}m_{14}m_{23}m_{25}m_{45} - m_{13}m_{14}m_{24}m_{25}m_{35} \\ & - m_{13}m_{15}m_{23}m_{24}m_{45} + m_{13}m_{15}m_{24}m_{25}m_{34} + m_{14}m_{15}m_{23}m_{24}m_{35} - m_{14}m_{15}m_{23}m_{25}m_{34}. \end{aligned} \quad (7.4)$$

This quintic is known as the *pentad*, and it plays an important role in *factor analysis* [46]. In conclusion, the 8-dimensional moment variety $\sigma_2(\mathcal{M}_{5,(11)})$ is already familiar to statisticians.

We now generalize to the toric variety defined by $\lambda = (1^d) = (1, 1, \dots, 1)$ for any $n > d$.

Remark 7.1.2. The moment variety $\mathcal{M}_{n,(1^d)}$ is the toric variety associated with the hypersimplex

$$\Delta(n, d) = \text{conv}\{e_{l_1} + e_{l_2} + \dots + e_{l_d} : 1 \leq l_1 < l_2 < \dots < l_d \leq n\}.$$

The variety $\mathcal{M}_{n,(1^d)}$ has dimension $n - 1$, it lives in $\mathbb{P}^{\binom{n}{d}-1}$ and its degree is the Eulerian number $A(n, d)$. This is the number of permutations of $[n] = \{1, 2, \dots, n\}$ which have exactly d descents. Indeed, the degree of any projective toric variety equals the normalized volume of the associated polytope [108, Theorem 4.16], and the formula $\text{Vol}(\Delta(n, d)) = A(n, d)$ is well-known in algebraic combinatorics. In [80, Theorem 2.2] it is attributed to Laplace.

In the case of the second hypersimplex, one can compute the prime ideal of $\sigma_r(\mathcal{M}_{n,(11)})$ by eliminating the diagonal entries m_{ii} from the ideal of $(r+1) \times (r+1)$ minors of the covariance matrix (m_{ij}) . This elimination problem is tough. For some instances see [46, Table 1].

Example 7.1.3. The moment variety $\sigma_5(\mathcal{M}_{9,(11)})$ is a hypersurface of degree 54 in \mathbb{P}^{35} . The representation of its equation by means of resultants is explained in [46, Example 24].

We now turn to a scenario that played a pivotal role in launching algebraic statistics in the 1990s, namely the spectral analysis of permutation data, as described in [40, Section 6]. This is based on the toric ideal associated with the Birkhoff polytope, whose vertices are the $n!$ permutation matrices of size $n \times n$. For an algebraic discussion see [108, Section 14.B].

Proposition 7.1.4. The moment variety $\mathcal{M}_{n,\lambda}$ for the partition $\lambda = (n - 1, n - 2, \dots, 2, 1)$ is the toric variety of the Birkhoff polytope, which lives in $\mathbb{P}^{n!-1}$ and has dimension $(n - 1)^2$.

Proof. The moment coordinates $m_{i_1 i_2 \dots i_n}$ for $\mathcal{M}_{n,\lambda}$ are indexed by the $n!$ permutations of $\{0, 1, 2, \dots, n - 1\}$. The monomials on the right hand side of (7.2) have degree $n - 1$ in $n(n - 1)$ distinct parameters μ_{ki} . The exponent vectors of these monomials can be identified with the permutation matrices from which the first row has been removed. This removal is an affine isomorphism, so it preserves the Birkhoff polytope, which has dimension $(n - 1)^2$. \square

Example 7.1.5 ($n = 4$). The toric variety $\mathcal{M}_{4,(321)} \subset \mathbb{P}^{23}$ has dimension 9 and degree 352. Its ideal is minimally generated by 18 quadrics like $m_{0123}m_{1032} - m_{0132}m_{1023}$ and 160 cubics like $m_{0123}m_{1203}m_{2013} - m_{0213}m_{2103}m_{1023}$. The latter are induced from the $n = 3$ case in Example 7.0.1. See [108, Example 14.7] for a Gröbner basis and [40, Section 6.1] for a statistical perspective.

Example 7.1.6 ($n = 5$). This appears in the last paragraph of [40, Section 6.1]. The toric variety $\mathcal{M}_{5,(4321)} \subset \mathbb{P}^{119}$ has dimension 16 and degree 4718075. Its ideal is minimally generated by 1050 quadrics and 28840 cubics.

Yamaguchi, Ogawa and Takemura [123] showed that the toric ideal for the variety in Proposition 7.1.4 is always generated in degree two and three. Theorem 7.3.6 generalizes this result.

The degrees reported in Examples 7.1.5 and 7.1.6 are the volume of the Birkhoff polytope. This volume is known up to $n = 10$ [21, 22].

7.2 Finiteness

This section shows that our moment varieties exhibit finiteness up to symmetry, in the sense of Draisma and collaborators [28, 42, 43, 44]. Namely, if r, d, λ are fixed and n is unbounded then finitely many S_n -orbits of polynomials suffice to cut out the varieties $\sigma_r(\mathcal{M}_{n,\bullet})$. Therefore, computer algebra can be useful for high-dimensional data analysis in the setting of [125].

We consider the projective variety $\sigma_r(\mathcal{M}_{n,d}) \subset \mathbb{P}^{\binom{n+d-1}{d}-1}$ defined by (7.3). This parametrization can be understood as follows without any reference to probability or statistics. Namely, $m_{i_1 i_2 \dots i_n}$ is the coefficient of the monomial $x_1^{i_1} x_2^{i_2} \cdots x_n^{i_n}$ in the expansion of the polynomial

$$F(x_1, x_2, \dots, x_n) = \sum_{j=1}^r f_1^{(j)}(x_1) f_2^{(j)}(x_2) \cdots f_n^{(j)}(x_n), \quad (7.5)$$

where $f_k^{(j)}(x) = 1 + \mu_{k1}^{(j)}x + \mu_{k2}^{(j)}x^2 + \cdots + \mu_{kd}^{(j)}x^d$ are rn unknown univariate polynomials. For any partition λ of d , let N_λ be the subset of coefficients $m_{i_1 i_2 \dots i_n}$ where $\{i_1, i_2, \dots, i_n\} \setminus \{0\}$ equals λ as a multiset. The variety $\sigma_r(\mathcal{M}_{n,\lambda})$ is the closure of the image of $\sigma_r(\mathcal{M}_{n,d})$ under the map $\mathbb{P}^{\binom{n+d-1}{d}-1} \dashrightarrow \mathbb{P}^{|N_\lambda|-1}$. The polynomial in (7.5) is the truncated moment generating function. Taking $r = 1$ we obtain the toric varieties $\mathcal{M}_{n,\bullet}$.

The equations for these varieties satisfy *finiteness up to symmetry* when r, d, λ are fixed and n grows. Here symmetry refers to the action of the symmetric group S_n on our varieties, their parametrizations (7.3), and their prime ideals. These ideals satisfy natural inclusions

$$I(\sigma_r(\mathcal{M}_{n,\bullet})) \subset I(\sigma_r(\mathcal{M}_{n+1,\bullet})), \quad \text{where } \bullet \in \{d, \lambda\},$$

by appending a zero to the indices of every coordinate. In symbols, $m_{i_1 i_2 \dots i_n} \mapsto m_{i_1 i_2 \dots i_n 0}$. If we iterate these inclusions and let the big symmetric group act, then we obtain inclusions

$$\langle S_n I(\sigma_r(\mathcal{M}_{n_0, \bullet})) \rangle \subseteq I(\sigma_r(\mathcal{M}_{n, \bullet})) \quad \text{for } n > n_0. \quad (7.6)$$

Ideal-theoretic finiteness means that there exists n_0 such that equality holds for all $n > n_0$. The weaker notion of *set-theoretic finiteness* means that equality holds in (7.6) after the left ideal is enlarged to its radical. The smallest possible n_0 , if it exists, is a function of r, d, λ .

Example 7.2.1 (Pentads and 3×3 minors). If $r = 2$ and $\lambda = (11)$ then ideal-theoretic finiteness holds with $n_0 = 6$. This was proved by Brouwer and Draisma in [28, Theorem 1.7] in response to [46, Conjecture 26]. The prime ideal of $\sigma_2(\mathcal{M}_{n, (11)})$ is generated by the $\binom{n}{5}$ pentads and the $5 \binom{n}{6}$ off-diagonal 3×3 minors of a symmetric $n \times n$ matrix. See [28, Section 3] for an equivariant Gröbner basis. If $r = 1$ then ideal-theoretic finiteness holds with $n_0 = 4$, by the Gröbner basis in [108, Theorem 9.1] for the toric ideal of the second hypersimplex.

In recent years, there has been considerable progress on commutative algebra in infinite polynomial rings with an action of the infinite symmetric group, or of rings over the category FI of finite sets with injections. The following result reflects the state of the art on that topic.

Theorem 7.2.2. Given any partition $\lambda \vdash d$ and integer $r \geq 1$, set-theoretic finiteness holds for the varieties $\sigma_r(\mathcal{M}_{n, d})$ and $\sigma_r(\mathcal{M}_{n, \lambda})$. Ideal-theoretic finiteness holds in the toric case $r = 1$.

Proof. The statement about toric varieties ($r = 1$) follows from [44, Theorem 1.1]. For $r \geq 2$, we apply the main theorem in [43]. First we consider the varieties $\sigma_r(\mathcal{M}_{n, d})$. Use the map (7.3) between two polynomial rings in countably many variables. This is a morphism of FI-algebras as in [43, Section 1.1]. In the formulation of [43, Corollary 1.1.2], the parametrization takes the $rd \times \mathbb{N}$ matrix whose entries are $\mu_{ki}^{(j)}$ to the $\mathbb{N} \times \dots \times \mathbb{N}$ (d times) tensor whose entries are $m_{i_1 \dots i_n}$ viewed as degree- d moments in n dimensions. The closure of the image is topologically $\text{Sym}(\mathbb{N})$ -Noetherian, which yields set-theoretic finiteness for $\sigma_r(\mathcal{M}_{n, d})$. The case of $\sigma_r(\mathcal{M}_{n, \lambda})$ is similar. \square

Remark 7.2.3. It is conjectured in [43, Conjecture 1.1.3] that the main result in [43] holds ideal-theoretically. This would imply ideal-theoretic finiteness for $\sigma_r(\mathcal{M}_{n, d})$ and $\sigma_r(\mathcal{M}_{n, \lambda})$.

Whenever ideal-theoretic finiteness holds, one can try to use equivariant Gröbner bases [28] for computing the desired finite generating set. An implementation for the toric case is described in [44], but we found this to be quite slow. The case $\lambda = (11)$ is covered by Example 7.2.1.

Example 7.2.4 (Cycles in bipartite graphs). If $r = 1$ and $\lambda = (21)$ then ideal-theoretic finiteness holds with $n_0 = 4$. Namely, the toric ideal of $\mathcal{M}_{n, (21)}$ is generated by $6 \binom{n}{4}$ quadrics

and $\binom{n}{3}$ cubics. This follows from [92, Lemma 1.1]. Indeed, the above binomials correspond to the chordless cycles in the bipartite graph that is obtained from $K_{n,n}$ by removing the n edges $(1, 1), (2, 2), \dots, (n, n)$. For $n > 4$, every such chordless cycle is supported on a bipartite subgraph of the same kind with $n_0 = 4$.

Example 7.2.5 (Hypersimplex). As seen in Section 7.1, when $r = 1$ and $\lambda = (1^d)$, the moment variety $\mathcal{M}_{n,\lambda}$ is the variety associated to the hypersimplex $\Delta(n, d)$. Its ideal is generated by quadrics [108, Section 14A]. The indices occurring in each quadratic binomial are 1 in at most $2d$ of the n coordinates. Therefore, ideal-theoretic finiteness holds with $n_0 = 2d$.

We close with a corollary that generalizes the previous two examples. Its proof rests on a forward reference to the next section, where we derive various results for our toric ideals.

Corollary 7.2.6. Fix a partition λ with e nonzero parts, fix $r = 1$, and suppose that n increases. The toric varieties $\mathcal{M}_{n,\lambda}$ satisfy ideal-theoretic finiteness for some $n_0 \leq 3e$ where e is the length of λ .

Proof. Theorem 7.3.6 says that the ideal of $\mathcal{M}_{n,\lambda}$ is generated by binomials of degree at most 3. Each of the two monomials in such a binomial is a product of two or three variables $m_{i_1 i_2 \dots i_n}$. The two monomials have the same A -degree, where A is the matrix representing (7.2). This implies that the slots $\ell \in \{1, 2, \dots, n\}$ where a nonzero index i_ℓ occurs are the same in both monomials. The total number of such slots is at most $3e$. This yields the bound $n_0 \leq 3e$. \square

Remark 7.2.7. Every partition $\lambda \vdash d$ satisfies $e \leq d$, and equality holds only for $\lambda = (1^d)$, as in Example 7.2.5. For $\lambda = (21)$ in Example 7.2.4, we have $e = 2$, and this yields $n_0 = 2e = 4$.

7.3 Toric combinatorics

This section is a detailed study of the toric varieties $\mathcal{M}_{n,\bullet}$. In particular, we study their dimensions, polytopes, and toric ideals. The ideal for $\mathcal{M}_{n,\lambda}$ is generated by quadrics and cubics, but the ideal for $\mathcal{M}_{n,d}$ is more complicated.

With each such toric variety we associate a 0-1 matrix A as in [108] whose columns correspond to the monomials in (7.2). The rank of A is one more than the dimension of the projective toric variety. We first show that $\mathcal{M}_{n,d}$ has the expected dimension, namely the number of parameters minus one.

Theorem 7.3.1. The dimension of the moment variety $\mathcal{M}_{n,d}$ is

$$\min \left\{ nd - 1, \binom{n+d-1}{d} - 1 \right\}.$$

Proof. First assume $n > d$. We will show that the A -matrix associated to the moment variety has rank nd by displaying a nonzero $nd \times nd$ minor. Consider the d special partitions

$$(1, 1, \dots, 1), (2, 1, \dots, 1), \dots, (d-2, 1, 1), (d-1, 1), (d). \quad (7.7)$$

Each of these partitions induces (by permutation) at least n columns in the A -matrix. For each $(k, 1, \dots, 1) \vdash d$, pick n of these columns such that k appears in each of the n spots. The principal submatrix of A induced by all these columns is an $nd \times nd$ matrix of the form

$$B = \left[\begin{array}{c|c|c|c|c|c} M & * & * & \dots & * & 0 \\ \hline 0 & I_n & 0 & \dots & 0 & 0 \\ \hline 0 & 0 & I_n & \dots & 0 & 0 \\ \hline 0 & 0 & 0 & \ddots & 0 & 0 \\ \hline 0 & 0 & 0 & \dots & I_n & 0 \\ \hline 0 & 0 & 0 & \dots & 0 & I_n \end{array} \right],$$

where the d row blocks are labeled $(\mu_{k1} : k \in [n]), \dots, (\mu_{kd} : k \in [n])$ and the d column blocks are (7.7). The matrix M gives a column basis for the A -matrix of the hypersimplex variety $\mathcal{M}_{n,(1,1,\dots,1)}$, so it is invertible. We conclude $\det B = \det M \neq 0$, and so $\text{rank}(A) = nd$.

Now suppose $n \leq d$. Index the columns of the A -matrix by permutations of (i_1, \dots, i_n) with $i_1 + \dots + i_n = d$ ordered reverse-lexicographically. Index the rows by $\mu_{11}, \mu_{12}, \dots, \mu_{nd}$. The principal submatrix on the first $2d + 1$ rows and columns is invertible, so the first $2d + 1$ columns of A are linearly independent. From the remaining columns, we pick $d(n - 2) - 1$ of them such that for every $j = 2d + 1, \dots, nd$ exactly one has 1 in the j th coordinate. In this way we obtain nd linearly independent columns of A . Therefore, A has full rank. \square

Given a partition $\lambda \vdash d$ padded by zeroes to have length n , we define a partition ν , called the *reduction* of λ . Let $k_0 \geq \dots \geq k_s$ be the multiplicities of the distinct parts in λ . Then

$$\nu = \left(\underbrace{s, \dots, s}_{k_s}, \underbrace{s-1, \dots, s-1}_{k_{s-1}}, \dots, \underbrace{1, \dots, 1}_{k_1}, \underbrace{0, \dots, 0}_{k_0} \right). \quad (7.8)$$

We write s for the largest part of ν , so $s + 1$ is the number of distinct parts of λ . For example, the partitions $(8, 5, 5, 4)$ and $(7, 7, 3, 0)$ have the same reduction $\nu = (2, 1, 0, 0)$, with $s = 2$.

Lemma 7.3.2. If ν is the reduction of λ then $|N_\lambda| = |N_\nu|$ and $\mathcal{M}_{n,\lambda} = \mathcal{M}_{n,\nu}$ in $\mathbb{P}^{|N_\nu|-1}$.

Proof. Let the μ_{k_0} be unknowns in the monomial parametrization (7.2). The image of this altered map also equals $\mathcal{M}_{n,\lambda}$. The toric variety $\mathcal{M}_{n,\nu}$ has the same parametrization, after changing the index $i \in \lambda$ in each parameter μ_{ki} to the corresponding entry in ν . \square

Example 7.3.3 (Hypersimplex). If $s = 1$ and $\lambda = (1^d)$ with $n/2 < d < n$ then $\nu = (1^{n-d})$ in Lemma 7.3.2, and we recover the identification of the hypersimplices $\Delta(n, d)$ and $\Delta(n, n-d)$.

Theorem 7.3.4. The moment variety $\mathcal{M}_{n,\lambda} = \mathcal{M}_{n,\nu}$ has dimension $(n - 1)s$, for ν in (7.8).

Proof. We must show that the A -matrix of $\mathcal{M}_{n,\nu}$ has rank $(n - 1)s + 1$. We proceed by induction on s , the base case $s = 1$ being the hypersimplex. We partition the rows of A into s blocks $(\mu_{k_1} : k \in [n]), \dots, (\mu_{k_s} : k \in [n])$. The rows of A in the i th block sum to the constant vector (k_i, k_i, \dots, k_i) . Hence, the rank of A is bounded above by $(n - 1)s + 1$. We will show that this is also a lower bound by displaying an invertible submatrix of this size.

First, assume that $k_s > 1$. Consider the columns of the A -matrix indexed by $m_{i_1 \dots i_n}$ such that $i_1 = s$. By induction on n , these columns induce a submatrix of rank $(n - 2)s + 1$. Hence, we may pick $ns - 2s + 1$ linearly independent columns from this set. Next, for each $j = 0, \dots, s - 1$, pick a column indexed by some $m_{i_1 \dots i_n}$ with $i_1 = j$. By construction, adding these s columns does not introduce dependence relations. We have constructed a set of $ns - 2s + 1 + s = (n - 1)s + 1$ linearly independent columns of A , so A has the desired rank.

Now consider the case when $k_s = 1$. Again, consider the columns of the A -matrix indexed by $m_{i_1 \dots i_n}$ such that $i_1 = s$. By induction, but now also on s , these columns induce a submatrix of rank $(n - 2)(s - 1) + 1$. Next, add $n - 1$ columns that are indexed by $m_{i_1 \dots i_n}$ where $i_1 = s - 1$ and such that for each $j = 2, \dots, n$ there is an index with $i_j = s$. Finally, add $s - 1$ columns indexed by $m_{i_1 \dots i_n}$ such that for each $j = s - 2, \dots, 0$ there is an index with $i_1 = j$ and $m_{i_2} \neq s, s - 1$. This way we obtain $(n - 2)(s - 1) + 1 + (n - 1) + (s - 1) = (n - 1)s + 1$ columns, which are linearly independent by construction. Therefore, A has the desired rank. \square

The toric variety $\mathcal{M}_{n,d}$ is an aggregate of the $\mathcal{M}_{n,\lambda}$ for $\lambda \vdash d$, but there is no easy transition. For instance, ideal generators for $\mathcal{M}_{n,d}$ do not restrict to ideal generators for $\mathcal{M}_{n,\lambda}$.

Example 7.3.5 ($n = d = 4$). The partitions $\lambda = (4), (31), (22), (211), (1111)$ have the reductions $\nu = (1), (21), (11), (21), ()$ with $s = 1, 2, 2, 2, 0$. Two nontrivial varieties $\mathcal{M}_{4,\nu}$ are given by the off-diagonal entries of 4×4 -matrices. The variety $\mathcal{M}_{4,4}$ has dimension 15 and degree 1072 in \mathbb{P}^{34} , and its ideal is generated by 52 quadrics and 28 cubics. The subset which involves the twelve unknowns $m_{2110}, \dots, m_{0112}$ does not suffice to cut out $\mathcal{M}_{4,(211)}$ in \mathbb{P}^{11} . The ideal of $\mathcal{M}_{4,(211)}$ is generated by 6 quadrics and 4 cubics, namely the cycles in Example 7.2.4.

The toric ideals for individual partitions are very nice. Our next result builds upon [123].

Theorem 7.3.6. For any partition λ , the ideal of $\mathcal{M}_{n,\lambda}$ is generated by quadrics and cubics.

Proof. For a partition $\lambda = (\lambda_1, \dots, \lambda_e)$, set $R = \mathbb{R}[m_{i_1 i_2 \dots i_n} : \{i_1, \dots, i_n\} \text{ is } \lambda]$. Let $I \subset R$ be the toric ideal defining $\mathcal{M}_{n,\lambda}$. Note that $\eta = (e, e - 1, \dots, 1)$ is a partition of the same length but possibly with a different sum. Let $J \subset S$ be the toric ideal defining $\mathcal{M}_{n,\eta}$ where $S = \mathbb{R}[m_{j_1 j_2 \dots j_n} : \{j_1, \dots, j_n\} \text{ is } \eta]$. Define a surjective ring homomorphism $\varphi : S \rightarrow R$ by mapping $m_{j_1 j_2 \dots j_n}$ to $m_{i_1 i_2 \dots i_n}$ where $\{i_1, \dots, i_n\}$ is obtained from $\{j_1, \dots, j_n\}$ by replacing e

with λ_1 , replacing $e - 1$ with λ_2 , etc. We claim that $\varphi(J) = I$. This implies the theorem because J is generated by quadrics and cubics [123, Theorem 2.1] and φ preserves \mathbb{N} -degree.

Since I and J are toric, we may verify $\varphi(J) = I$ on binomials. To show $\varphi(J) \subseteq I$, fix a binomial in J , say of degree δ , written as $\mathbf{b} = \prod_{\alpha=1}^{\delta} m_{j_{\alpha 1} \cdots j_{\alpha n}} - \prod_{\alpha=1}^{\delta} m_{j'_{\alpha 1} \cdots j'_{\alpha n}}$. Encode this by two $\delta \times n$ matrices $B = (j_{\alpha\beta})$ and $C = (j'_{\alpha\beta})$. Membership in J means that substituting the parametrization (7.2) into \mathbf{b} gives the result 0, and this is equivalent to the multiset of entries in corresponding columns of B and C being equal. This property is preserved after replacing e by λ_1 , replacing $e - 1$ by λ_2 etc. throughout B and C . So $\varphi(\mathbf{b}) \in I$ as desired.

To prove $I \subseteq \varphi(J)$, let $\mathbf{c} = \prod_{\alpha=1}^{\delta} m_{i_{\alpha 1} \cdots i_{\alpha n}} - \prod_{\alpha=1}^{\delta} m_{i'_{\alpha 1} \cdots i'_{\alpha n}}$ be a binomial in I encoded by matrices $D = (i_{\alpha\beta})$ and $E = (i'_{\alpha\beta})$. We will construct a binomial $\mathbf{d} \in J$ such that $\varphi(\mathbf{d}) = \mathbf{c}$. In terms of matrices D and E , in each of their rows we must choose one element that equals λ_1 and replace it by e , then choose another element that equals λ_2 and replace it by $e - 1$, and so forth until the set of nonzero elements in each row has been replaced by $[e]$, in such a way so that the multiset of entries in corresponding columns of the transformed matrices D and E are equal. To achieve this it suffices to consider distinct values in λ one at a time.

Without loss of generality, assume $\lambda = (1^e)$. Now D and E have e ones and $n - e$ zeros per each row. To choose the elements to replace by e , we consider a bipartite multigraph between the rows of D and the rows of E , where an edge is drawn between a row in D and a row in E for every column in which there is a 1 in both rows. A perfect matching would give a valid choice of elements to replace by e . Such a matching exists by Hall's Marriage Theorem. Indeed, for any subset W of rows in D their neighborhood must contain at least $|W|$ rows in E . Otherwise, there exists a column in D with more ones than the corresponding column in E , since each row contains the same number of ones. But this contradicts $\mathbf{c} \in I$. Similarly, we carry out the subsequent replacements. Thus a suitable binomial \mathbf{d} exists. It follows $I \subseteq \varphi(J)$. Combining with the preceding paragraph, we conclude $\varphi(J) = I$. \square

By contrast, the ideals for $\mathcal{M}_{n,d}$ appear to be more complicated. We conjecture that there does not exist a uniform degree bound for their generators that is independent of n, d .

Example 7.3.7 ($n=3, d=7$). The toric variety $\mathcal{M}_{3,7}$ has dimension 20 and degree 14922 in \mathbb{P}^{35} . Its ideal is minimally generated by 46 cubics, 168 quartics, 135 quintics and 18 sextics.

7.4 Secant varieties

In this section we inquire about the identifiability of the secant varieties $\sigma_r(\mathcal{M}_{n,\bullet})$ for $r \geq 2$ and present what we know about their dimensions. The main results are Theorems 7.4.3 and 7.4.8. These rest on integer programming and tropical geometry. Moreover, the parametrization (7.3) represents a challenging implicitization problem. We also report on some compu-

tational results, featuring both symbolic and numerical methods. For most examples in this chapter, symbolic computations were used.

Dimension

Theorems 7.3.1 and 7.3.4 gave the dimensions of our moment varieties for $r = 1$. We next focus on $r \geq 2$, where $\sigma_r(\mathcal{M}_{n,d})$ and $\sigma_r(\mathcal{M}_{n,\lambda})$ are no longer toric. We begin with an example.

Example 7.4.1 ($n = 5, d = 3$). The toric variety $\mathcal{M}_{5,3}$ and its secant variety $\sigma_2(\mathcal{M}_{5,3})$ live in the projective space \mathbb{P}^{34} of symmetric $5 \times 5 \times 5$ tensors; see Example 7.0.2. By Theorem 7.3.1, we have $\dim(\mathcal{M}_{5,3}) = 14$. The expected dimension of the secant variety $\sigma_2(\mathcal{M}_{5,3})$ would be $2 \cdot 14 + 1 = 29$. However, we must subtract $5 = 4 + 1$ because $\mathcal{M}_{5,3}$ is a cone with apex $\sigma_2(\mathcal{M}_{5,3}) = \mathbb{P}^4$. Therefore, $\sigma_2(\mathcal{M}_{5,3})$ has dimension 24. The prime ideal of $\sigma_2(\mathcal{M}_{5,3})$ will be presented in Proposition 7.4.11.

Our first result explains the drop in dimension seen in the example above.

Proposition 7.4.2. The dimension of the moment variety satisfies the upper bound

$$\dim(\sigma_r(\mathcal{M}_{n,d})) \leq \min\{rnd - rn + n - 1, \binom{n+d-1}{d} - 1\}. \quad (7.9)$$

Proof. The given toric variety is a cone over the projective space $\mathcal{M}_{n,(d)} = \mathbb{P}^{n-1}$. In symbols, $\mathcal{M}_{n,d} = \mathbb{P}^{n-1} \star \widetilde{\mathcal{M}}_{n,d}$, where $\widetilde{\mathcal{M}}_{n,d}$ is the toric variety given by all $\binom{n+d-1}{d} - n$ moments that involve more than one coordinate. By counting parameters, we find $\dim(\widetilde{\mathcal{M}}_{n,d}) \leq n(d-1) - 1$. We obtain the secant variety of the big toric variety as the join of the apex with the reduced toric variety: $\sigma_r(\mathcal{M}_{n,d}) = \mathbb{P}^{n-1} \star \sigma_r(\widetilde{\mathcal{M}}_{n,d})$. The dimension of the right-hand side is bounded above by $n + r(\dim(\widetilde{\mathcal{M}}_{n,d})) + r - 1 \leq n + r \cdot (n(d-1) - 1) + r - 1$. This yields (7.9). \square

We found the inequality (7.9) to be strict when $r \geq n$. The following sharper bound holds. (To see it is sharper, consider $S = [d]$ and $S = \{d\}$ in (7.10).)

Theorem 7.4.3. The dimension of the secant variety $\sigma_r(\mathcal{M}_{n,d})$ is bounded above by the optimal value of the following integer linear programming problem:

$$\begin{aligned} \text{maximize } c_1 + c_2 + \cdots + c_d - 1 & \quad \text{subject to } 0 \leq c_i \leq nr \text{ for } i \in [d] \\ & \quad \text{and } \sum_{i \in S} c_i \leq \sum_{\lambda \cap S \neq \emptyset} |N_\lambda| \text{ for } S \subseteq [d]. \end{aligned} \quad (7.10)$$

The last sum ranges over partitions $\lambda \vdash d$ of length $\leq n$ having nonempty intersection with S .

Proof. The secant variety $\sigma_r(\mathcal{M}_{n,d})$ is parametrized by the polynomial map (7.3). Therefore its dimension is one less than the maximal rank assumed by the differential of (7.3). This Jacobian matrix has size $\binom{n+d-1}{d} \times nrd$, where the rows are labeled by $m_{i_1 i_2 \dots i_n}$ such that $i_1, \dots, i_n \geq 0$ and $i_1 + \dots + i_n = d$, and the columns are labeled by $\mu_{1i}^{(j)}, \dots, \mu_{ni}^{(j)}$ for $i = 1, \dots, d$

and $j = 1, \dots, r$. We view this as a block matrix, where the rows are grouped according to the partition λ given by (i_1, \dots, i_n) and the columns are grouped according to the degree i . Notice that the matrix is sparse, in that a block labeled by (λ, i) is nonzero only if $i \in \lambda$.

Let \mathcal{C} be a set of linearly independent columns in the Jacobian matrix, with c_i columns labeled by i . The integers c_i satisfy $0 \leq c_i \leq nr$ for $i = 1, \dots, d$. Let $S \subseteq [d]$ and \mathcal{C}' the subset of columns in \mathcal{C} that are labeled by elements of S . Since \mathcal{C}' is linearly independent, the number of rows which are nonzero in \mathcal{C} exceeds $|\mathcal{C}'|$. By the aforementioned sparsity,

$$\sum_{i \in S} c_i \leq \sum_{\substack{\lambda \vdash d, \lambda \cap S \neq \emptyset, \\ \text{length of } \lambda \text{ is at most } n}} |N_\lambda|. \quad (7.11)$$

We conclude that $|\mathcal{C}| - 1$ is bounded above by the maximum value in (7.10), as desired. \square

Solving an integer linear program is expensive in general. However, the integer linear program in (7.10) has a special structure which allows for a greedy solution that is optimal.

Theorem 7.4.4. We construct a feasible solution for (7.10) greedily, starting with $\mathbf{c}^{(0)} = 0$ in \mathbb{Z}^d . For $t = 1, \dots, r$, choose $\mathbf{c}^{(t)} \in \mathbb{Z}^d$ such that $c_i^{(t-1)} \leq c_i^{(t)} \leq c_i^{(t-1)} + n$ for all $i \in [d]$, and if $c_i^{(t)} < c_i^{(t-1)} + n$ then there exists $S \subseteq [d]$ containing i such that $\sum_{j \in S} c_j^{(t)} = \sum_{\lambda \cap S \neq \emptyset} |N_\lambda|$. Then $\mathbf{c}^{(r)} \in \mathbb{Z}^d$ is optimal for the integer linear program (7.10).

Proof. We claim that $\mathbf{c}^{(r)}$ is optimal for the linear program (7.10), with integrality constraints dropped. The dual linear program has variables y_S for $\emptyset \neq S \subseteq [d]$ and z_i for $i \in [d]$. This dual linear program equals:

$$\begin{aligned} & \text{minimize } nr(z_1 + \dots + z_d) + \sum_{S \subseteq [d]} \left(\sum_{\lambda \cap S \neq \emptyset} |N_\lambda| \right) y_S - 1 \\ & \text{subject to } \mathbf{y} \in \mathbb{R}_{\geq 0}^{2^d - 1}, \mathbf{z} \in \mathbb{R}_{\geq 0}^d \text{ and } z_i + \sum_{S \ni i} y_S \geq 1 \text{ for } i \in [d]. \end{aligned}$$

It suffices to find a dual feasible point at which the dual objective equals the primal objective evaluated at $\mathbf{c}^{(r)}$. We call a set $S \subseteq [d]$ *saturated* if equality holds in (7.11) for $\mathbf{c}^{(r)}$. We define

$$y_S^{(r)} = \begin{cases} 1 & \text{if } S \subseteq [d] \text{ is saturated and maximal such set} \\ 0 & \text{otherwise;} \end{cases}$$

$$z_i^{(r)} = \begin{cases} 1 & \text{if } \nexists S \subseteq [d] \text{ s.t. } i \in S \text{ and } S \text{ is saturated} \\ 0 & \text{otherwise.} \end{cases}$$

The vector $(\mathbf{y}^{(r)}, \mathbf{z}^{(r)})$ is dual feasible. Further, we claim that there is a unique maximal saturated subset of $[d]$, possibly empty. Suppose that S and T are saturated. Then

$\sum_{i \in T \setminus S} c_i^{(r)} + \sum_{i \in T \cap S} c_i^{(r)} = \sum_{\lambda \cap T \neq \emptyset} |N_\lambda|$ and $\sum_{i \in T \cap S} c_i^{(r)} \leq \sum_{\lambda \cap (T \cap S) \neq \emptyset} |N_\lambda|$ since T is saturated and $\mathbf{c}^{(r)}$ is primal feasible. Subtracting these, we find

$$\sum_{i \in T \setminus S} c_i^{(r)} \geq \sum_{\lambda \cap T \neq \emptyset} |N_\lambda| - \sum_{\lambda \cap (S \cap T) \neq \emptyset} |N_\lambda| = \sum_{\lambda \cap T \neq \emptyset, \lambda \cap (S \cap T) = \emptyset} |N_\lambda|.$$

Adding $\sum_{i \in S} c_i^{(r)} = \sum_{\lambda \cap S \neq \emptyset} |N_\lambda|$ implies $\sum_{i \in S \cup T} c_i^{(r)} \geq \sum_{\lambda \cap (S \cup T) \neq \emptyset} |N_\lambda|$. Hence $S \cup T$ is saturated by primal feasibility. Thus there is a unique maximal saturated subset of $[d]$. It follows that the dual objective evaluated at $(\mathbf{y}^{(r)}, \mathbf{z}^{(r)})$ equals the primal objective evaluated at $\mathbf{c}^{(r)}$. This completes the proof. \square

We conjecture that the integer linear program (7.10) computes the correct dimension:

Conjecture 7.4.5. If $d \geq 3$ then the bound for $\dim(\sigma_r(\mathcal{M}_{n,d}))$ in Theorem 7.4.3 is tight.

Informally, the conjecture says that the secant variety has the maximal dimension possible given the sparsity pattern of its parametrization (7.3). This has been verified in many cases.

Example 7.4.6. Let $n = 4, d = 12$. The inclusion $\sigma_r(\mathcal{M}_{4,12}) \subset \mathbb{P}^{454}$ is strict for $r \leq 11$. The dimensions are 47, 91, 135, 175, 215, 255, 291, 327, 363, 399, 431. This was found correctly by Theorem 7.4.3. Compare this to the sequence 47, 91, 135, 179, 223, 267, 311, 355, 399, 443, 454, which is the upper bound $\min\{44r + 3, 454\}$ for $\dim(\sigma_r(\mathcal{M}_{4,12}))$ given in Proposition 7.4.2.

The question of finding the dimension is equally intriguing if we replace the parameter d by one specific partition $\lambda \vdash d$. Of particular interest is the partition $\lambda = (1, 1, 1, \dots, 1) = (1^d)$. This toric variety has dimension $n - 1$, and hence we have the trivial upper bound

$$\dim(\sigma_r(\mathcal{M}_{n,(1^d)})) \leq r(n - 1) + r - 1 = nr - 1. \quad (7.12)$$

Based on extensive computations, we conjecture that equality holds outside the matrix case:

Conjecture 7.4.7. Secant varieties of hypersimplices, other than the second hypersimplex, have the expected dimension. In symbols, if $3 \leq d \leq n - 3$ then $\dim(\sigma_r(\mathcal{M}_{n,(1^d)})) = nr - 1$.

Theorem 5.1 in [125] implies $\sigma_r(\mathcal{M}_{n,(1^d)})$ is strongly identifiable for $r \lesssim n^{\lfloor (d-1)/2 \rfloor}$. In particular, the secant variety has the expected dimension. Our next result is that the secant variety also has the expected dimension if $r \lesssim n^{d-2}$. The proof relies on tropical geometry [41].

Theorem 7.4.8. The secant variety of the hypersimplex has the expected dimension if

$$\left(1 + d(n - d) + \binom{d}{2} \binom{n - d}{2}\right) (r - 1) < \binom{n}{d}. \quad (7.13)$$

Proof. Assume (7.13) holds. Let $S = \Delta(n, d) \cap \mathbb{Z}^n$. From [41, Lemma 3.8] and [41, Corollary 3.2], it suffices to show that there exist points $v_1, \dots, v_r \in \mathbb{R}^n$ such that each of the Voronoi cells

$$\text{Vor}_i(v) := \{\alpha \in S : \|\alpha - v_i\|_2 < \|\alpha - v_j\|_2 \text{ for all } i \neq j\}$$

spans an affine space of dimension $n - 1$ inside \mathbb{R}^n . To argue this, choose $v_1 \in S$ arbitrarily and set $N_1 = \{\alpha \in S : \|\alpha - v_1\|_2 \leq 2\}$. Next choose $v_2 \in S \setminus N_1$ arbitrarily and set $N_2 = \{\alpha \in S : \|\alpha - v_2\|_2 \leq 2\}$. Next choose $v_3 \in S \setminus (N_1 \cup N_2)$ arbitrarily and define N_3 . We continue until $S \setminus (N_1 \cup N_2 \cup \dots) = \emptyset$. Note that the parenthesized sum on the left-hand side of (7.13) equals the size of each set N_i , while the right-hand side gives the size of S . Thus, (7.13) guarantees that we choose at least r points in S . Furthermore, these points differ pairwise in at least 3 coordinates by construction. So, the i th Voronoi cell contains all elements of S that differ from v_i in at most one coordinate. That is, it contains v_i and all vertices in the hypersimplex adjacent to v_i . Hence $\text{Vor}_i(v)$ has the same affine span as $\Delta(n, d)$. \square

Implicitization

We verified the dimensions in Section 7.4 with numerical methods for fairly large instances, by computing the rank of the Jacobian matrix of the parametrization (7.2). For this we employed *Maple*, *Julia*, and the numerical *Macaulay2* package in [29]. We found it much more difficult to solve the implicitization problem, that is, to compute the defining polynomials of our moment varieties. The pentad (7.4) suggests that such polynomials can be quite interesting. This section offers more examples of equations, along with the degrees for our varieties.

Remark 7.4.9. It is preferable to work with birational parametrizations when numerically computing the degree of a variety [88]. However the map (7.2) is d -to-1: if ω is a primitive d th root of unity then we can replace μ_{ki} by $\mu_{ki}\omega^i$ without changing $m_{i_1 i_2 \dots i_n}$. This implies that the map (7.3) has fibers of size at least $r!d^r$. We set $\mu_{2,1} = 1$ and let $\mu_{1,0}$ be an unknown to make (7.2) into a birational parametrization of $\mathcal{M}_{n,d}$. Likewise, we turn (7.3) into a parametrization of $\sigma_r(\mathcal{M}_{n,d})$ that is expected to be $r!$ -to-1, by setting $\mu_{2,1}^{(j)}$ to 1 and using unknowns for $\mu_{1,0}^{(j)}$.

Let us now present a case study for implicitization, focused on the hypersimplex $\Delta(6, 3)$.

Example 7.4.10 ($n = 6, d = 3$). The 5-dimensional toric variety $\mathcal{M}_{6,(111)}$ lives in \mathbb{P}^{19} , and it has degree $A(6, 3) = 66$ by Remark 7.1.2. Its toric ideal is minimally generated by 69 binomial quadrics. These quadrics are the 2×2 minors that are visible (i.e. do not involve

any stars) in the following masked Hankel matrix:

$$\begin{bmatrix} \star & \star & \star & \star & \star & m_{123} & m_{124} & m_{125} & m_{126} & m_{134} & m_{135} & m_{136} & m_{145} & m_{146} & m_{156} \\ \star & m_{123} & m_{124} & m_{125} & m_{126} & \star & \star & \star & \star & m_{234} & m_{235} & m_{236} & m_{245} & m_{246} & m_{256} \\ m_{123} & \star & m_{134} & m_{135} & m_{136} & \star & m_{234} & m_{235} & m_{236} & \star & \star & \star & m_{345} & m_{346} & m_{356} \\ m_{124} & m_{134} & \star & m_{145} & m_{146} & m_{234} & \star & m_{245} & m_{246} & \star & m_{345} & m_{346} & \star & \star & m_{456} \\ m_{125} & m_{135} & m_{145} & \star & m_{156} & m_{235} & m_{245} & \star & m_{256} & m_{345} & \star & m_{356} & \star & m_{456} & \star \\ m_{126} & m_{136} & m_{146} & m_{156} & \star & m_{236} & m_{246} & m_{256} & \star & m_{346} & m_{356} & \star & m_{456} & \star & \star \end{bmatrix}$$

The rows are labeled by $i \in \{1, 2, \dots, 6\}$, and the columns by pairs $\{j, k\}$ of such indices. The entry is m_{ijk} if these are disjoint, and it is \star otherwise. Here we use $m_{123} = m_{111000}$, $m_{124} = m_{110100}, \dots, m_{456} = m_{000111}$. Our toric ideal is generated by all 2×2 minors without \star .

The 6×15 matrix has twenty 3×3 minors without \star , and these vanish on $\sigma_2(\mathcal{M}_{6,(111)})$. In addition to these cubics, the ideal contains 12 pentads (7.4), one for each facet $\Delta(5, 2)$ of the hypersimplex $\Delta(6, 3)$. Our ideal for $r = 2$ is generated by these 20 cubics and 12 quintics. Numerical degree computations using Remark 7.4.9 with *HomotopyContinuation.jl* [27] reveal

$$\deg(\sigma_2(\mathcal{M}_{6,(111)})) = 465 \text{ and } \deg(\sigma_3(\mathcal{M}_{6,(111)})) = 80. \tag{7.14}$$

Symbolic computations for $r = 3$ are challenging. Our secant variety has codimension 2 in \mathbb{P}^{19} . There are no quadrics or cubics vanishing on $\sigma_3(\mathcal{M}_{6,(111)})$, but there is a unique quartic:

$$\begin{aligned} & m_{123}m_{145}m_{246}m_{356} - m_{123}m_{145}m_{256}m_{346} - m_{123}m_{146}m_{245}m_{356} \\ & + m_{123}m_{146}m_{256}m_{345} + m_{123}m_{156}m_{245}m_{346} - m_{123}m_{156}m_{246}m_{345} \\ & - m_{124}m_{135}m_{236}m_{456} + m_{124}m_{135}m_{256}m_{346} + m_{124}m_{136}m_{235}m_{456} \\ & - m_{124}m_{136}m_{256}m_{345} - m_{124}m_{156}m_{235}m_{346} + m_{124}m_{156}m_{236}m_{345} \\ & + m_{125}m_{134}m_{236}m_{456} - m_{125}m_{134}m_{246}m_{356} - m_{125}m_{136}m_{234}m_{456} \\ & + m_{125}m_{136}m_{246}m_{345} + m_{125}m_{146}m_{234}m_{356} - m_{125}m_{146}m_{236}m_{345} \\ & - m_{126}m_{134}m_{235}m_{456} + m_{126}m_{134}m_{245}m_{356} + m_{126}m_{135}m_{234}m_{456} \\ & - m_{126}m_{135}m_{245}m_{346} - m_{126}m_{145}m_{234}m_{356} + m_{126}m_{145}m_{235}m_{346} \\ & + m_{134}m_{156}m_{235}m_{246} - m_{134}m_{156}m_{236}m_{245} - m_{135}m_{146}m_{234}m_{256} \\ & + m_{135}m_{146}m_{236}m_{245} + m_{136}m_{145}m_{234}m_{256} - m_{136}m_{145}m_{235}m_{246}. \end{aligned} \tag{7.15}$$

Note the beautiful combinatorics in this polynomial: the role of the 5-cycle for the pentad is now played by the *quadrilateral set*, i.e. the six intersection points of four lines in the plane.

We conclude this article with the smallest non-trivial secant varieties. Here “non-trivial” means $r \geq 2$, the variety does not fill its ambient projective space, and the ambient dimension is as small as possible. The next two results feature all cases where $\binom{n+d-1}{d} \leq 50$. The list consists of $(r, n, d) = (2, 5, 3)$ from Example 7.4.1 and $(r, n, d) = (2, 4, 4)$ from Example 7.3.5. We state these as propositions because they represent case studies that are of independent interest for experimental mathematics, especially in the ubiquitous setting of tensor decompositions.

Proposition 7.4.11. The secant variety $\sigma_2(\mathcal{M}_{5,3})$ has dimension 24 and degree 3225 in \mathbb{P}^{34} . Its prime ideal is generated by 313 polynomials, namely 10 cubics, 283 quintics, 10 sextics and 10 septics. These ideal generators are obtained by elimination from the ideal of 3×3 minors of the 5×15 matrix

$$\begin{bmatrix} a_{23} & a_{24} & a_{25} & a_{34} & a_{35} & a_{45} & \star & \star & \star & \star & \star & b_{21} & b_{31} & b_{41} & b_{51} \\ a_{13} & a_{14} & a_{15} & \star & \star & \star & a_{34} & a_{35} & a_{45} & \star & b_{12} & \star & b_{32} & b_{42} & b_{52} \\ a_{12} & \star & \star & a_{14} & a_{15} & \star & a_{24} & a_{25} & \star & a_{45} & b_{13} & b_{23} & \star & b_{43} & b_{53} \\ \star & a_{12} & \star & a_{13} & \star & a_{15} & a_{23} & \star & a_{25} & a_{35} & b_{14} & b_{24} & b_{34} & \star & b_{54} \\ \star & \star & a_{12} & \star & a_{13} & a_{14} & \star & a_{23} & a_{24} & a_{34} & b_{15} & b_{25} & b_{35} & b_{45} & \star \end{bmatrix}. \quad (7.16)$$

Proposition 7.4.11 is important in that it displays a general technique of obtaining equations for varieties of low rank structured symmetric tensors from masked Hankel matrices.

Notation and Proof. The visible entries in the masked matrix (7.16) are 30 of the 35 moments $m_{i_1 i_2 i_3 i_4 i_5}$. The 10 moments for $\lambda = (111)$ are denoted $a_{12} = m_{00111}, a_{13} = m_{01011}, \dots, a_{35} = m_{111010}, a_{45} = m_{111100}$, and the 20 moments for $\lambda = (12)$ are denoted $b_{12} = m_{21000}, b_{13} = m_{20100}, \dots, b_{21} = m_{12000}, \dots, b_{53} = m_{00102}, b_{54} = m_{00012}$. The 25 stars are distinct new unknowns, and these are being eliminated. The matrix contains ten 3×3 -submatrices with no stars. Their determinants are the ten cubics mentioned in Proposition 7.4.11.

The ideal of $\sigma_2(\mathcal{M}_{5,3})$ is homogeneous in the bigrading given by a and b . Among the generators, we find 1, 55, 110, 90, 27 quintics of bidegrees $(5, 0), (3, 2), (2, 3), (1, 4), (0, 5)$. The quintic of bidegree $(5, 0)$ is the pentad of the symmetric 5×5 -matrix (a_{ij}) . One of the quintics of bidegree $(2, 3)$ is

$$\begin{aligned} & a_{13}a_{45}b_{25}b_{41}b_{53} - a_{13}a_{45}b_{25}b_{43}b_{51} - a_{14}a_{34}b_{21}b_{43}b_{54} \\ & + a_{14}a_{34}b_{23}b_{41}b_{54} - a_{14}a_{35}b_{23}b_{45}b_{51} + a_{14}a_{35}b_{25}b_{43}b_{51} \\ & + a_{14}a_{45}b_{24}b_{45}b_{51} - a_{14}a_{45}b_{25}b_{41}b_{54} + a_{15}a_{34}b_{21}b_{45}b_{53} \\ & - a_{15}a_{34}b_{25}b_{41}b_{53} - a_{34}a_{45}b_{24}b_{45}b_{53} + a_{34}a_{45}b_{25}b_{43}b_{54}. \end{aligned}$$

The ten sextics have bidegrees $(4, 2)$ and $(0, 6)$, five each. All ten septics have bidegree $(3, 4)$. The $27 + 5$ generators of bidegrees $(0, 5)$ and $(0, 6)$ generate the prime ideal of $\sigma_2(\mathcal{M}_{5,(21)})$. They arise from the 5×5 matrix (b_{ij}) by eliminating the diagonal. The degree 3225 was first found numerically, and later confirmed symbolically by *Macaulay2*. \square

Our final result concerns tensors of format $4 \times 4 \times 4 \times 4$.

Proposition 7.4.12. The secant variety $\sigma_2(\mathcal{M}_{4,4})$ has dimension 27 and degree 8650 in \mathbb{P}^{34} . Its prime ideal has only three minimal generators in degrees at most six. These are the cubics

$$\det \begin{bmatrix} m_{2200} & m_{2110} & m_{2020} \\ m_{1201} & m_{1111} & m_{1021} \\ m_{0202} & m_{0112} & m_{0022} \end{bmatrix}, \det \begin{bmatrix} m_{2200} & m_{2101} & m_{2002} \\ m_{1210} & m_{1111} & m_{1012} \\ m_{0220} & m_{0121} & m_{0022} \end{bmatrix},$$

$$\det \begin{bmatrix} m_{2020} & m_{2011} & m_{2002} \\ m_{1120} & m_{1111} & m_{1102} \\ m_{0220} & m_{0211} & m_{0202} \end{bmatrix}. \quad (7.17)$$

Proposition 7.4.12 is proved by direct computation. The degree 8650 was found with *HomotopyContinuation.jl* using the method in [88]. The absence of minimal generators in degrees 4, 5, 6 was verified by solving the linear equations for each A -degree in that range. Each solution was found to be in the ideal (7.17). At present we know of no ideal generators for $\sigma_2(\mathcal{M}_{4,4})$ that involve the moments $m_{3100}, m_{1300}, \dots, m_{0013}$. What is the smallest degree in which we can find such generators?

Conclusion. In this chapter, we defined the moment varieties of conditionally independent mixture distributions on \mathbb{R}^n . We focused on computing their dimensions, defining polynomials, and degrees. Our future research in this area would extend the results to block-wise independence structures, systematically generate equations by resultants, and numerically compute high degrees.

Bibliography

- [1] Yulia Alexandr. “Logarithmic Voronoi polytopes for discrete linear models”. In: *Algebraic Statistics* (2023). To appear.
- [2] Yulia Alexandr, Eliana Duarte, and Julian Vill. “Decomposable context-specific models”. 2023. arXiv: 2210.11521.
- [3] Yulia Alexandr and Alexander Heaton. “Logarithmic Voronoi cells”. In: *Algebraic Statistics* 12.1 (2021), pp. 75–95.
- [4] Yulia Alexandr, Alexander Heaton, and Sascha Timme. “Computing a logarithmic Voronoi cell”. www.JuliaHomotopyContinuation.org/examples/logarithmic-voronoi/.
- [5] Yulia Alexandr and Serkan Hoşten. “Logarithmic Voronoi cells for Gaussian models”. In: *Journal of Symbolic Computation* 122 (2024).
- [6] Yulia Alexandr and Serkan Hoşten. “Maximum information divergence from linear and toric models”. 2023. arXiv: 2308.15598.
- [7] Yulia Alexandr, Joe Kileel, and Bernd Sturmfels. “Moment Varieties for Mixtures of Products”. In: *Proceedings of the 2023 International Symposium on Symbolic and Algebraic Computation*. Association for Computing Machinery, 2023, pp. 53–60.
- [8] Carlos Améndola, Nathan Bliss, Isaac Burke, Courtney R. Gibbons, Martin Helmer, Serkan Hoşten, Evan D. Nash, Jose Israel Rodriguez, and Daniel Smolkin. “The maximum likelihood degree of toric varieties”. In: *J. Symbolic Comput.* 92 (2019), pp. 222–242.
- [9] Carlos Améndola, Jean-Charles Faugère, and Bernd Sturmfels. “Moment varieties of Gaussian mixtures”. In: *Journal of Algebraic Statistics* 7 (2016), pp. 14–28.
- [10] Carlos Améndola, Lukas Gustafsson, Kathlén Kohn, Orlando Marigliano, and Anna Seigal. “The maximum likelihood degree of linear spaces of symmetric matrices”. In: *Matematiche (Catania)* 76.2 (2021), pp. 535–557.
- [11] Carlos Améndola, Kathlén Kohn, Philipp Reichenbach, and Anna Seigal. “Invariant Theory and Scaling Algorithms for Maximum Likelihood Estimation”. In: *SIAM Journal on Applied Algebra and Geometry* 5.2 (2021), pp. 304–337.

- [12] Carlos Améndola and Piotr Zwiernik. “Likelihood geometry of correlation models”. In: *Matematiche (Catania)* 76.2 (2021), pp. 559–583.
- [13] Lamprini Ananiadi and Eliana Duarte. “Gröbner bases for staged trees”. In: *Algebraic Statistics* 12.1 (2021), pp. 1–20.
- [14] Theodore W. Anderson. “Estimation of covariance matrices which are linear combinations or whose inverses are linear combinations of given matrices”. In: *Essays in Probability and Statistics*. Univ. North Carolina Press, Chapel Hill, N.C., 1970, pp. 1–24.
- [15] David Avis and Komei Fukuda. “Reverse search for enumeration”. In: *Discrete Appl. Math.* 65.1-3 (1996). First International Colloquium on Graphs and Optimization (GOI), 1992 (Grimentz), pp. 21–46.
- [16] Nihat Ay. “An information-geometric approach to a theory of pragmatic structuring”. In: *Ann. Probab.* 30.1 (2002), pp. 416–436.
- [17] Nihat Ay, Jürgen Jost, Hông Vân Lê, and Lorenz Schwachhöfer. *Information Geometry*. Springer Verlag, New York, 2017.
- [18] Nihat Ay and Andreas Knauf. “Maximizing multi-information”. In: *Kybernetika* 42.5 (2006), pp. 517–538.
- [19] Daniel J. Bates, Jonathan D. Hauenstein, Andrew J. Sommese, and Charles W. Wampler. “Bertini: Software for Numerical Algebraic Geometry”. Available at bertini.nd.edu with permanent doi: dx.doi.org/10.7274/R0H41PB5.
- [20] Daniel J. Bates, Andrew J. Sommese, Jonathan D. Hauenstein, and Charles W. Wampler. *Numerically Solving Polynomial Systems with Bertini*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2013.
- [21] Matthias Beck and Dennis Pixton. “The Ehrhart Polynomial of the Birkhoff Polytope”. In: *Discrete & Computational Geometry* 30 (2003), pp. 623–637.
- [22] Matthias Beck and Dennis Pixton. “The volume of the 10th Birkhoff polytope”. arXiv: [math/0305332](https://arxiv.org/abs/math/0305332).
- [23] Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to linear optimization*. Athena Scientific, 1997.
- [24] Patrick Billingsley. *Probability and measure*. New York: Wiley-Interscience, 1995.
- [25] Craig Boutilier, Nir Friedman, Moises Goldszmidt, and Daphne Koller. “Context-specific independence in Bayesian networks”. In: *Uncertainty in artificial intelligence (Portland, OR, 1996)*. Morgan Kaufmann, San Francisco, CA, 1996, pp. 115–123.
- [26] Paul Breiding, Kathlén Kohn, and Bernd Sturmfels. *Metric Algebraic Geometry*. Oberwolfach Seminars, Birkhäuser, Basel, 2024.

- [27] Paul Breiding and Sascha Timme. “HomotopyContinuation.jl: A Package for Homotopy Continuation in Julia”. In: *Mathematical Software – ICMS 2018*. Ed. by James H. Davenport, Manuel Kauers, George Labahn, and Josef Urban. Springer International Publishing, 2018, pp. 458–465.
- [28] Andries E. Brouwer and Jan Draisma. “Equivariant Gröbner bases and the Gaussian two-factor model”. In: *Mathematics of Computation* 80 (2011), pp. 1123–1133.
- [29] Justin Chen and Joe Kileel. “Numerical implicitization”. In: *Journal of Software for Algebra and Geometry* 9 (2019), pp. 55–63.
- [30] David M. Chickering, David Heckerman, and Christopher Meek. “A Bayesian Approach to Learning Bayesian Networks with Local Structure”. In: *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*. Providence, Rhode Island: Morgan Kaufmann Publishers Inc., 1997, pp. 80–89.
- [31] Soojin Cho. “Polytopes of roots of type A_n ”. In: *Bull. Austral. Math. Soc.* 59.3 (1999), pp. 391–402.
- [32] Diego Cifuentes, Kristian Ranestad, Bernd Sturmfels, and Madeleine Weinstein. “Voronoi cells of varieties”. In: *Journal of Symbolic Computation* 109 (2022), pp. 351–366. ISSN: 0747-7171.
- [33] Daniel Ciripoi, Nidhi Kaihnsa, Andreas Löhne, and Bernd Sturmfels. “Computing convex hulls of trajectories”. In: *Rev. Un. Mat. Argentina* 60.2 (2019), pp. 637–662.
- [34] Christiane Görden Claudia Fevola. “The Mathematical Research-Data Repository MathRepo”. In: *Der Computeralgebra-Rundbrief* (2022).
- [35] Jukka Corander. “Labelled Graphical Models”. In: *Scandinavian Journal of Statistics* 30.3 (2003), pp. 493–508.
- [36] Jukka Corander, Antti Hyttinen, Juha Kontinen, Johan Pensar, and Jouko Väänänen. “A Logical Approach to Context-Specific Independence”. In: *Logic, Language, Information, and Computation*. Springer Berlin Heidelberg, 2016, pp. 165–182.
- [37] Isobel Davies, Eliana Duarte, Irem Portakal, and Miruna-Ştefana Sorea. “Families of Polytopes with Rational Linear Precision in Higher Dimensions”. In: *Foundations of Computational Mathematics* (2022).
- [38] Jesús A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations*. Vol. 25. Algorithms and Computation in Mathematics. Structures for algorithms and applications. Springer-Verlag, Berlin, 2010, pp. xiv+535.
- [39] Sandra Di Rocco, David Eklund, and Madeleine Weinstein. “The Bottleneck Degree of Algebraic Varieties”. In: *SIAM J. Appl. Algebra Geom.* 4.1 (2020), pp. 227–253.
- [40] Persi Diaconis and Bernd Sturmfels. “Algebraic algorithms for sampling from conditional distributions”. In: *Annals of Statistics* 26 (1998), pp. 363–397.
- [41] Jan Draisma. “A tropical approach to secant dimensions”. In: *Journal of Pure and Applied Algebra* 212 (2008), pp. 349–363.

- [42] Jan Draisma. “Finiteness for the k -factor model and chirality varieties”. In: *Advances in Mathematics* 223 (2010), pp. 243–256.
- [43] Jan Draisma, Rob H. Eggermont, Azhar Farooq, and Leandro Meier. “Image closure of symmetric wide-matrix varieties”. arXiv: 2212.12458.
- [44] Jan Draisma, Rob H. Eggermont, Robert Krone, and Anton Leykin. “Noetherianity for infinite-dimensional toric varieties”. In: *Algebra Number Theory* 9 (2015), pp. 1857–1880.
- [45] Jan Draisma, Emil Horobeţ, Giorgio Ottaviani, Bernd Sturmfels, and Rekha R. Thomas. “The Euclidean distance degree of an algebraic variety”. In: *Found. Comput. Math.* 16.1 (2016), pp. 99–149.
- [46] Mathias Drton, Bernd Sturmfels, and Seth Sullivan. “Algebraic factor analysis: Tetrads, pentads and beyond”. In: *Probability Theory and Related Fields* 138 (2007), pp. 463–493.
- [47] Mathias Drton, Bernd Sturmfels, and Seth Sullivan. *Lectures on algebraic statistics*. Vol. 39. Oberwolfach Seminars. Birkhäuser Verlag, Basel, 2009, pp. viii+171.
- [48] Eliana Duarte, Orlando Marigliano, and Bernd Sturmfels. “Discrete statistical models with rational maximum likelihood estimator”. In: *Bernoulli* 27.1 (2021), pp. 135–154.
- [49] Eliana Duarte and Liam Solus. “A new characterization of discrete decomposable models”. In: *Proceedings of the American Mathematical Society* 151 (2023), pp. 1325–1338.
- [50] Eliana Duarte and Liam Solus. “Representation of context-specific causal models with observational and interventional data”. In: *arXiv:2101.09271* (2022).
- [51] Nicholas Eriksson, Stephen E. Fienberg, Alessandro Rinaldo, and Seth Sullivan. “Polyhedral conditions for the nonexistence of the MLE for hierarchical log-linear models”. In: *J. Symbolic Comput.* 41.2 (2006), pp. 222–233.
- [52] Stephen E. Fienberg. “Quasi-independence and maximum likelihood estimation in incomplete contingency tables”. In: *J. Amer. Statist. Assoc.* 65 (1970), pp. 1610–1616.
- [53] Luis David Garcia, Michael Stillman, and Bernd Sturmfels. “Algebraic geometry of Bayesian networks”. In: *Journal of Symbolic Computation* 39.3 (2005), pp. 331–355.
- [54] Dan Geiger, David Heckerman, Henry King, and Christopher Meek. “Stratified exponential families: Graphical models and model selection”. In: *Ann. Statist.* 29.2 (2001), pp. 505–529.
- [55] Dan Geiger, Christopher Meek, and Bernd Sturmfels. “On the toric algebra of graphical models”. In: *Ann. Statist.* 34.3 (2006), pp. 1463–1492.
- [56] Israel M. Gelfand, Mark I. Graev, and Alexander Postnikov. “Combinatorics of hypergeometric functions associated with positive roots”. In: *The Arnold-Gelfand mathematical seminars*. Birkhäuser Boston, Boston, MA, 1997, pp. 205–221.

- [57] Christopher G. Gibson. *Elementary Geometry of Algebraic Curves: An Undergraduate Introduction*. Cambridge University Press, 1998.
- [58] Christiane Görgen, Anna Bigatti, Eva Riccomagno, and Jim Q. Smith. “Discovery of statistical equivalence classes using computer algebra”. In: *Internat. J. Approx. Reason.* 95 (2018), pp. 167–184.
- [59] Christiane Görgen and Jim Q. Smith. “Equivalence classes of staged trees”. In: *Bernoulli* 24.4A (2018), pp. 2676–2692.
- [60] Daniel R. Grayson and Michael E. Stillman. “Macaulay2, a software system for research in algebraic geometry”. Available at <http://www.math.uiuc.edu/Macaulay2/>.
- [61] Elizabeth Gross and Jose Israel Rodriguez. “Maximum likelihood geometry in the presence of data zeros”. In: *ISSAC 2014—Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*. ACM, New York, 2014, pp. 232–239.
- [62] Emil Horobeţ and Madeleine Weinstein. “Offset hypersurfaces and persistent homology of algebraic varieties”. In: *Comput. Aided Geom. Design* 74 (2019), pp. 101767, 14.
- [63] Serkan Hoşten, Amit Khetan, and Bernd Sturmfels. “Solving the likelihood equations”. In: *Found. Comput. Math.* 5.4 (2005), pp. 389–407.
- [64] Serkan Hoşten and Jay Shapiro. “Primary decomposition of lattice basis ideals”. In: *J. Symbolic Comput.* 29.4-5 (2000). Symbolic computation in algebra, analysis, and geometry (Berkeley, CA, 1998), pp. 625–639.
- [65] Serkan Hoşten and Seth Sullivant. “Gröbner bases and polyhedral geometry of reducible and cyclic models”. In: *J. Combin. Theory Ser. A* 100.2 (2002), pp. 277–301.
- [66] June Huh. “Varieties with maximum likelihood degree one”. In: *J. Algebr. Stat.* 5.1 (2014), pp. 1–17.
- [67] June Huh and Bernd Sturmfels. “Likelihood geometry”. In: *Combinatorial algebraic geometry*. Vol. 2108. Lecture Notes in Math. Springer, Cham, 2014, pp. 63–117.
- [68] Wolfram Research Inc. “Mathematica, Version 13.3”. Champaign, IL, 2023. URL: <https://www.wolfram.com/mathematica>.
- [69] Tomi Janhunen, Martin Gebser, Jussi Rintanen, Henrik Nyman, Johan Pensar, and Jukka Corander. “Learning discrete decomposable graphical models via constraint optimization”. In: *Statistics and Computing* 27.1 (2017), pp. 115–130.
- [70] Anders N. Jensen. “Gfan, a software system for Gröbner fans and tropical varieties”. Available at <http://home.imf.au.dk/jensen/software/gfan/gfan.html>.
- [71] Yuhan Jiang, Kathlén Kohn, and Rosa Winter. “Linear spaces of symmetric matrices with non-maximal maximum likelihood degree”. In: *Matematiche (Catania)* 76.2 (2021), pp. 461–481.

- [72] Michael Joswig and Katja Kulas. “Tropical and ordinary convexity combined”. In: *Adv. Geom.* 10.2 (2010), pp. 333–352.
- [73] Josef Juriček. “Maximization of the information divergence from multinomial distributions”. In: *Acta Universitatis Carolinae. Mathematica et Physica* 52.1 (2011), pp. 27–35.
- [74] Dinara Kasko. *Pastry Art*. Accessed: 2020-06-09. URL: <http://www.dinarakasko.com/>.
- [75] Kathlén Kohn, Boris Shapiro, and Bernd Sturmfels. “Moment varieties of measures on polytopes”. In: *Annali della Scuola Normale Superiore di Pisa* 21 (2020), pp. 739–770.
- [76] Daphne Koller and Nir Friedman. *Probabilistic graphical models*. Adaptive Computation and Machine Learning. Principles and techniques. MIT Press, Cambridge, MA, 2009, pp. xxxvi+1231.
- [77] Kaie Kubjas, Elina Robeva, and Bernd Sturmfels. “Fixed points EM algorithm and nonnegative rank boundaries”. In: *Ann. Statist.* 43.1 (2015), pp. 422–461.
- [78] Steffen L. Lauritzen. *Graphical models*. Vol. 17. Oxford Statistical Science Series. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1996.
- [79] John B. Little. “The many lives of the twisted cubic”. In: *Amer. Math. Monthly* 126.7 (2019), pp. 579–592.
- [80] Ricky Ini Liu. “Laurent polynomials, Eulerian numbers, and Bernstein’s theorem”. In: *Journal of Combinatorial Theory, Series A* 124 (2014), pp. 244–250.
- [81] Marloes Maathuis, Mathias Drton, Steffen Lauritzen, and Martin Wainwright. *Handbook of Graphical Models*. First. CRC Press, Inc., 2018.
- [82] Laurent Manivel, Mateusz Michałek, Leonid Monin, Tim Seynnaeve, and Martin Vodička. “Complete quadrics: Schubert calculus for Gaussian models and semidefinite programming”. In: *J. Eur. Math. Soc.* (2022).
- [83] František Matúš. “Divergence from factorizable distributions and matroid representations by partitions”. In: *IEEE Trans. Inform. Theory* 55.12 (2009), pp. 5375–5381.
- [84] František Matúš. “Optimality conditions for maximizers of the information divergence from an exponential family”. In: *Kybernetika (Prague)* 43.5 (2007), pp. 731–746.
- [85] František Matúš and Nihat Ay. “On maximization of the information divergence from an exponential family”. In: *Proceedings of 6th workshop on uncertainty processing: Hejnice, September 24-27, 2003*. [Praha]: Oeconomica, 2003, pp. 199–204.
- [86] Mateusz Michałek, Leonid Monin, and Jarosław A. Wiśniewski. “Maximum likelihood degree, complete quadrics, and \mathbb{C}^* -action”. In: *SIAM J. Appl. Algebra Geom.* 5.1 (2021), pp. 60–85.

- [87] Mateusz Michałek and Bernd Sturmfels. *Invitation to Nonlinear Algebra*. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2021.
- [88] Laura Brustenga i Moncusí, Sascha Timme, and Madeleine Weinstein. “96120: The degree of the linear orbit of a cubic surface”. In: *Matematiche (Catania)* 75 (2020), pp. 425–437.
- [89] David Mond, Smith Jim, and van Straten Duco. “Stochastic factorizations, sandwiched simplices and the topology of the space of explanations”. In: *Proc. R. Soc. Lond. A* 459 (2003), pp. 2821–2845.
- [90] Guido Montúfar, Johannes Rauh, and Nihat Ay. “Expressive Power and Approximation Errors of Restricted Boltzmann Machines”. In: *Advances in Neural Information Processing Systems*. Ed. by J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger. Vol. 24. Curran Associates, Inc., 2011.
- [91] Guido Montúfar, Johannes Rauh, and Nihat Ay. “Maximal Information Divergence from Statistical Models defined by Neural Networks”. In: *Geometric Science of Information, Lecture Notes in Computer Science*. Vol. 8085. Springer, 2013, pp. 759–766.
- [92] Hidefumi Ohsugi and Takayuki Hibi. “Toric ideals generated by quadratic binomials”. In: *Journal of Algebra* 218 (1999), pp. 509–527.
- [93] Shmuel Onn. “Geometry, complexity, and combinatorics of permutation polytopes”. In: *J. Combin. Theory Ser. A* 64.1 (1993), pp. 31–49.
- [94] Lior Pachter and Bernd Sturmfels, eds. *Algebraic statistics for computational biology*. Cambridge University Press, New York, 2005, pp. xii+420.
- [95] Athanase Papadopoulos and Marc Troyanov. *Handbook of Hilbert Geometry*. IRMA Lectures in Mathematics and Theoretical Physics. European Mathematical Society, 2014.
- [96] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.
- [97] Johan Pensar, Henrik Nyman, Timo Koski, and Jukka Corander. “Labeled directed acyclic graphs: a generalization of context-specific independence in directed graphical models”. In: *Data Mining and Knowledge Discovery* 29.2 (2015), pp. 503–533.
- [98] Daniel Plaumann, Rainer Sinn, and Jannik Wesner. “Families of faces and the normal cycle of a convex semi-algebraic set”. In: *Beiträge zur Algebra und Geometrie / Contributions to Algebra and Geometry* 64 (2023), pp. 851–875.
- [99] David Poole and Nevin L. Zhang. “Exploiting contextual independence in probabilistic inference”. In: *J. Artificial Intelligence Res.* 18 (2003), pp. 263–313.
- [100] Johannes Rauh. “Finding the Maximizers of the Information Divergence from an Exponential Family”. PhD thesis. Universität Leipzig, 2011.

- [101] Johannes Rauh. “Finding the maximizers of the information divergence from an exponential family”. In: *IEEE Trans. Inform. Theory* 57.6 (2011), pp. 3236–3247.
- [102] Johannes Rauh. “Optimally approximating exponential families”. In: *Kybernetika* 49.2 (2013), pp. 199–215.
- [103] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 9.4)*. <https://www.sagemath.org>. 2021.
- [104] Igor R. Shafarevich. *Basic algebraic geometry. 1*. Russian. Varieties in projective space. Springer, Heidelberg, 2013, pp. xviii+310.
- [105] Jim Q. Smith and Paul E. Anderson. “Conditional independence and chain event graphs”. In: *Artificial Intelligence* 172.1 (2008), pp. 42–68.
- [106] Andrew J. Sommese and Charles W. Wampler II. *The numerical solution of systems of polynomials*. Arising in engineering and science. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2005, pp. xxii+401.
- [107] Frank Sottile. “Toric ideals, real toric varieties, and the moment map”. In: *Contemp. Math.* 334 (2003), pp. 225–240.
- [108] Bernd Sturmfels. “Gröbner bases and convex polytopes”. In: *American Mathematical Society, University Lectures Series, No 8* (1996).
- [109] Bernd Sturmfels. *Solving systems of polynomial equations*. Vol. 97. CBMS Regional Conference Series in Mathematics. Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2002.
- [110] Bernd Sturmfels and Simon Telen. “Likelihood equations and scattering amplitudes”. In: *Algebraic Statistics* 12.2 (2021), pp. 167–186.
- [111] Bernd Sturmfels, Sascha Timme, and Piotr Zwiernik. “Estimating linear covariance models with numerical nonlinear algebra”. In: *Algebr. Stat.* 11.1 (2020), pp. 31–52.
- [112] Bernd Sturmfels and Caroline Uhler. “Multivariate Gaussian, semidefinite matrix completion, and convex algebraic geometry”. In: *Ann. Inst. Statist. Math.* 62.4 (2010), pp. 603–638.
- [113] Seth Sullivant. “Algebraic Geometry of Gaussian Bayesian Networks”. In: *Adv. Appl. Math.* 40.4 (2008), pp. 482–513.
- [114] Seth Sullivant. *Algebraic statistics*. Vol. 194. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2018, pp. xiii+490.
- [115] Seth Sullivant. “Toric fiber products”. In: *J. Algebra* 316.2 (2007), pp. 560–577.
- [116] Halyna Syta and Rien van de Weygaert. “Life and Times of Georgy Voronoi”. 2009. arXiv: 0912.3269.
- [117] L. B. Thomas. “Rank Factorization of Nonnegative Matrices (A. Berman)”. In: *SIAM Review* 16.3 (1974), pp. 393–394.

- [118] Santtu Tikka, Antti Hyttinen, and Juha Karvanen. “Identifying Causal Effects via Context-specific Independence Relations”. In: *Advances in Neural Information Processing Systems*. Vol. 32. Curran Associates, Inc., 2019.
- [119] Caroline Uhler. “Gaussian graphical models: An algebraic and geometric perspective”. In: *Handbook on Graphical Models* (2018). Ed. by M. Drton, S. Lauritzen, M. Maathuis and M. Wainwright.
- [120] Georges Voronoi. “Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire. Recherches sur les paralléloèdres primitifs”. In: *J. Reine Angew. Math.* 134 (1908), pp. 198–287.
- [121] Madeleine Weinstein. “Metric Algebraic Geometry”. PhD thesis. University of California, Berkeley, 2021.
- [122] Nanny Wermuth. “On block-recursive linear regression equations”. In: *Rebrape* 6.1 (1992). With discussion and a reply by the author, pp. 1–56.
- [123] Takashi Yamaguchi, Mitsunori Ogawa, and Akimichi Takemura. “Markov degree of the Birkhoff model”. In: *Journal of Algebraic Combinatorics* 40 (2014), pp. 293–311.
- [124] Nico Zavallos. *Optimizing ‘The Sacrifice’*. Accessed: 2020-06-09. 2019. URL: <https://nicoz.net/offret/index.html>.
- [125] Yifan Zhang and Joe Kileel. “Moment estimation for nonparametric mixture models through implicit tensor decomposition”. In: *SIAM Journal on Mathematics of Data Science* (2023). To appear.
- [126] Günter M. Ziegler. *Lectures on polytopes*. Vol. 152. Graduate Texts in Mathematics. Springer-Verlag, New York, 1995.
- [127] Piotr Zwiernik, Caroline Uhler, and Donald Richards. “Maximum likelihood estimation for linear Gaussian covariance models”. In: *J. R. Stat. Soc. Ser. B. Stat. Methodol.* 79.4 (2017), pp. 1269–1292.