The Algebraic Statistics of Sampling, Likelihood, and Regression

Der Fakultät für Mathematik und Informatik
der Universität Leipzig eingereichte

Dissertation

zur Erlangung des akademischen Grades

Doctor Rerum Naturalium
(Dr. rer. nat.)

im Fachgebiet

Mathematik

vorgelegt von

Orlando Marigliano

geboren am 5. August 1994 in Rom

Leipzig, den 2. Juni 2020

Abstract

This thesis is about statistical models and algebraic varieties. Algebraic Statistics unites these two concepts, turning algebraic structure into statistical insight. Featured here are three types of models that have such an algebraic structure.

Linear Gaussian covariance models are continuous models which are simple to define but hard to analyze. We compute their maximum likelihood degree in dimension two and find it equal to $2n - 3$ generically if the model has n covariates.

Discrete models with rational MLE are those discrete models for which likelihood estimation is easiest. We characterize them geometrically by building on the work of Huh and Kapranov on Horn uniformization.

Algebraic manifolds are a more general kind of object which is used to encode continuous data. We introduce a new method for computing integrals and sampling from distributions on them, based on intersecting with random linear spaces.

A brief report on mathematics in the sciences featuring case studies from soil ecology and nonparametric statistics closes the thesis.

About the author

Orlando Marigliano studied mathematics at the University of Bonn. He completed his Master's degree in 2018. His Master's thesis in algebraic geometry was supervised by Daniel Huybrechts. Since April 2018 he is a member of the Nonlinear Algebra Group at Max Planck Institute for Mathematics in the Sciences, Leipzig. His doctoral studies in algebraic statistics are supervised by Bernd Sturmfels and Christiane Görgen.

Bibliographic data

Orlando Marigliano. *The Algebraic Statistics of Sampling, Likelihood, and Regression*. PhD thesis, Universität Leipzig, 2020.

```
@phdthesis{om,
  author = {Orlando Marigliano},
  title  = {The Algebraic Statistics of Sampling, Likelihood, and Regression},
  school = {Universit{\"a}t Leipzig},
  year   = {2020}
}
```

Authorship

This thesis is presented to the Faculty of Mathematics and Computer Science at the University of Leipzig in partial fulfillment of the requirements for a doctoral degree. The contents of this thesis are either my own work or represent joint work of my co-authors and myself.

Chapter 1 is written by myself.

Chapter 2 is based on the joint article [21] with Jane I. Coons and Michael Ruddy. My coauthors and I contributed equally to all mathematical exposition and results. Michael Ruddy executed the experiments in Example 2.1.6. The article was submitted to *Algebraic Statistics* and is in revision.

Chapter 3 is based on the joint article [29] with Eliana Duarte and Bernd Sturmfels. My coauthors and I contributed equally to all mathematical exposition and results. I executed the experiments in Examples 3.5.1, 3.5.2, and 3.5.3. The article was accepted for publication in *Bernoulli* and is to appear.

Chapter 4 is based on the joint article [11] with Paul Breiding. My coauthor and I contributed equally to all mathematical exposition and results. Paul Breiding executed the experiments in Sections 4.2 and 4.6. The article was accepted for publication in *SIAM Journal on Mathematics of Data Science* and is to appear.

Chapter 5 is written by myself. Section 5.1 is based on the joint article [43] with Carlos Guerra, Manuel Delgado Baquerizo, Eliana Duarte, Christiane Görgen, Fernando T. Maestre, and Nico Eisenhauer. The article was submitted to *Journal of Global Ecology and Biogeography* and is in revision. Section 5.2 is based on the joint article [76] with Leonie Selk and Charles Tillier. My

coauthors and I contributed equally to the proof of Lemma 5.2.1 in that section. It can be found in the cited article. The article was submitted to *Biometrika* and is in revision.

Appendix A is written by myself.

Erklärung

Hiermit erkläre ich, die vorliegende Dissertation selbständig und ohne unzulässige fremde Hilfe angefertigt zu haben. Ich habe keine anderen als die angeführten Quellen und Hilfsmittel benutzt und sämtliche Textstellen, die wörtlich oder sinngemäß aus veröffentlichten oder unveröffentlichten Schriften entnommen wurden, und alle Angaben, die auf mündlichen Auskünften beruhen, als solche kenntlich gemacht. Ebenfalls sind alle von anderen Personen bereitgestellten Materialen oder erbrachten Dienstleistungen als solche gekennzeichnet.

Leipzig, den 2. Juni 2020

Orlando Marigliano

to Gregorio

Contents

List of Figures

1. Introduction

A traditional algebraic geometer does not suspect that their discipline can be useful outside of mathematics. I was surprised to discover that in *Algebraic Statistics* this is the case. In this field of mathematics, researchers use recent techniques from algebra, geometry, and combinatorics to solve problems in theoretical statistics. For instance, computational commutative algebra can help address problems in sampling and experimental design, as the early works [26] and [69] demonstrated. Since then, algebraic statistics has expanded. It now connects to a range of topics including polytope theory [52], causality [85], toric geometry [65] and phylogenetics [68].

This thesis contributes to Algebraic Statistics in the realm of *maximum likelihood estimation*. This is an important problem in statistics related to model selection and parameter estimation. The algebraic statistician addresses this problem by focusing on *algebraic models*. These models are attractive because they lend themselves to algebraic techniques. This thesis features special discrete algebraic models in Chapter 3 and linear Gaussian models in Chapter 2.

More broadly, this thesis is about exploiting the geometric and algebraic structure of data. This promising approach to data analysis is taken for instance in the articles [10] and [30]. Continuous data can sometimes be modeled by a manifold. Hence, sampling from one in an efficient way is an important problem. Chapter 4 addresses this problem in the case of *algebraic manifolds*.

In Chapter 5, I include two case studies on mathematics in the sciences. There, I record the main mathematical insights and their significance to their respective applications. That chapter features the fields of soil ecology and nonparametric statistics.

In the next four sections, I introduce some common terminology and illustrate the main themes of this thesis with simple examples. I also set up problem statements to be solved by each of the three main chapters. The examples are well-known toy examples in algebraic statistics and can be found for instance in the textbook [83].

Notation

In the next sections, I unify the formalism for discrete and continuous statistical models. Hence, some of the notation might appear non-standard to the reader. In general, I denote statistical models (varieties, manifolds) by calligraphic capital letters such as $\mathcal{M}$. I denote the points of these models by x, y, z. When the points are interpreted as probability

distributions I write p, q, or Σ instead. I identify a Gaussian probability distribution (having zero mean) with its covariance matrix Σ. I use X, Y for random variables and e for random events.

Acknowledgements

Writing this thesis was only possible thanks to the many people who accompanied me in my life as a doctoral student in Leipzig. An acknowledgements section is too limited to do you all justice in a way that feels right to me. Nevertheless, thank you. In particular:

Thanks to my advisors Bernd Sturmfels and Christiane Görgen for their support, friendship and guidance throughout my doctoral studies. For encouraging me to write better, communicate better, think better. For collaborating with me on projects and for introducing me to a mathematical community that makes me enjoy my research immensely.

Thanks to my coauthors Paul Breiding, Jane Ivy Coons, Eliana Duarte, Carlos Guerra, Leonie Selk, Michael Ruddy, and Charles Tillier for being patient and uplifting collaborators. At all times, I was always more excited to do mathematics with you than to work alone.

Thanks to my family for their loving support throughout all of my studies.

Thanks to all the people who refined my thinking on the material in this thesis through conversation: Carlos Améndola, Benjamin Gess, János Kollár, Avinash Kulkarni, Tony Lelievre, Emre Sertöz, Sascha Timme, Caroline Uhler, Lorenzo Venturello, and Piotr Zwiernik.

1.1. Algebraic models

A *statistical model* $\mathcal{M}$ is a subset of the set of probability distributions on some sample space Ω. In statistics, this subset is often the image of some parameter map $\Theta \to \mathcal{M}$. The basic problem of maximum likelihood estimation is to find the distribution $x \in \mathcal{M}$ that best explains a set of observations $e \in \Omega^N$. If the distributions $x \in \mathcal{M}$ have densities f_x, this amounts to maximizing the *log-likelihood*

$$\sum_{i=1}^{N} \log f_x(e_i) \tag{1.1.1}$$

over all $x \in \mathcal{M}$. The implicit assumption here is that there is a true probability distribution on Ω that generated the samples $e_1, \ldots, e_N$ independently of each other. Maximizing the log-likelihood has thus the effect of maximizing the density of e in the product space Ω^N. For more background on likelihood estimation and statistics in general, see the textbook [16].

Example 1.1.1. Let $\Omega = \{0, 1, 2, 3\}$. The space of all probability distributions on Ω can be represented by the tetrahedron

$$\Delta_3 = \{(p_0, p_1, p_2, p_3) \in \mathbb{R}^4 \mid p_j \geq 0 \text{ for all } j,\ \textstyle\sum_j p_j = 1\}.$$

The *independence model* for two binary random variables is the set

$$\mathcal{M} = \{p \in \Delta_3 \mid p_0 p_2 - p_1 p_3 = 0\}.$$

Suppose we observe the sample $e = (0, 1, 2, 2, 3, 2, 0)$. Maximizing the log-likelihood (1.1.1) amounts to maximizing the sum

$$2 \log p_0 + \log p_1 + 3 \log p_2 + \log p_3$$

over $\mathcal{M}$. Observe that this sum does not depend on the ordering of the entries of e. Instead, it only depends on the vector of counts u defined by $u_j = \#\{i \mid e_i = j\}$ for $j = 0, \ldots, 3$. Here, $u = (2, 1, 3, 1)$.

Definition 1.1.2. A *discrete statistical model* on $n + 1$ outcomes is a subset $\mathcal{M}$ of the probability simplex

$$\Delta_n \coloneqq \{p \in \mathbb{R}^{n+1} \mid p_j \geq 0 \text{ for all } j,\ \textstyle\sum_j p_j = 1\},$$

which we interpret as the set of probability distributions on the finite set $\Omega = \{0, \ldots, n\}$. Example 1.1.1 illustrates that for the purpose of likelihood estimation it suffices to represent a sample $e \in \Omega^N$ by its vector of counts $u \in \mathbb{N}^{|\Omega|}$. The sample size is recovered by the quantity $|u| \coloneqq \sum_j u_j$. Another way to represent the sample e is to identify it with its *empirical distribution*, which is defined by $q \coloneqq u/|u|$. The maximum likelihood estimation problem for an empirical distribution $q \in \Delta_n$ is then to maximize the value of the *log-likelihood function*

$$\ell(p, q) \coloneqq \sum_{j=0}^{n} q_j \log p_j \tag{1.1.2}$$

over all $p \in \mathcal{M}$. Note that if e is a sample and u its vector of counts, then $\ell(p, u/|u|)$ differs from the expression in (1.1.1) by a multiplicative constant that does not depend on p.

Example 1.1.3. Let $\Omega = \mathbb{R}^3$. The space of all Gaussian probability distributions on Ω with zero mean can be represented by the set

$$\mathrm{PD}_3 = \{\Sigma \in \mathrm{Mat}(3 \times 3, \mathbb{R}) \mid \Sigma \text{ symmetric, positive definite}\}$$

of all 3×3 covariance matrices. The *conditional independence model* of two Gaussian random variables given a third is the set

$$\mathcal{M} = \{(\sigma_{ij})_{i,j} \in \mathrm{PD}_3 \mid \sigma_{12}\sigma_{33} - \sigma_{13}\sigma_{32} = 0\}.$$

Suppose we observe the sample $e = ((1,1,1),(1,2,3),(-1,0,0))$. Let

$$S = \frac{1}{3}\left((1,1,1)^T(1,1,1) + (1,2,3)^T(1,2,3) + (-1,0,0)^T(-1,0,0)\right) = \frac{1}{3}\begin{pmatrix} 3 & 3 & 4 \\ 3 & 5 & 7 \\ 4 & 7 & 10 \end{pmatrix}$$

be the *sample covariance matrix* of e. Up to additive and multiplicative constants that do not depend on Σ, the log-likelihood (1.1.1) reduces to the expression

$$-\log\det(\Sigma) - \operatorname{tr}(S\Sigma^{-1}).$$

Definition 1.1.4. A *Gaussian statistical model* on n random variables is a subset $\mathcal{M}$ of the cone of positive definite symmetric matrices

$$\mathrm{PD}_n \coloneqq \{\Sigma \in \mathrm{Mat}(n \times n, \mathbb{R}) \mid \Sigma \text{ symmetric, positive definite}\},$$

which we interpret as the set of Gaussian probability distributions on $\Omega = \mathbb{R}^n$. Example 1.1.3 illustrates that for the purpose of likelihood estimation, it is enough to represent a sample $e \in \Omega^N$ by its sample covariance matrix $S = (1/N)\sum_{i=1}^N e_i e_i^T$. The maximum likelihood estimation problem for a sample covariance matrix $S \in \mathrm{PD}_n$ is then to minimize the value of the *negated log-likelihood function*

$$\ell(\Sigma, S) \coloneqq \log\det(\Sigma) + \operatorname{tr}(S\Sigma^{-1}) \tag{1.1.3}$$

over all $\Sigma \in \mathcal{M}$.

This thesis focuses on algebraic models, defined below. Recall that a *semi-algebraic set* is a subset of $\mathbb{R}^n$ defined by polynomial equations and inequalities.

Definition 1.1.5. An *algebraic model* is a semi-algebraic subset $\mathcal{M}$ of a semi-algebraic set $\mathcal{N}$ of distributions on some space Ω. We call $\mathcal{N}$ the space of *empirical distributions*.

Examples 1.1.1 and 1.1.3 are algebraic models. Their space of empirical distributions is Δ_n and PD_n, respectively.

1.2. The maximum likelihood degree

The likelihood estimation problem for an algebraic model $\mathcal{M} \subseteq \mathcal{N}$ starts with the log-likelihood function $\ell\colon \mathcal{M} \times \mathcal{N} \to \mathbb{R}$ defined for the model and asks for the critical points of the restriction of ℓ to $\mathcal{M} \times \{y\}$ for a fixed $y \in \mathcal{N}$. One method for finding these points uses Lagrangian multipliers. The first step is to set up the Lagrange multiplier function

$$\mathcal{L}(x, y, \lambda) = \ell(x, y) - \sum_{k=1}^{r} \lambda_k g_k(x, y)$$

where the $g_1, \ldots, g_r$ are the polynomials specifying the equalities in the definition of $\mathcal{M}$, and $\lambda \in \mathbb{R}^r$. The next step is to solve the system of equations $\nabla_{(x,\lambda)}\mathcal{L}(x, y, \lambda) = 0$, called the *score equations*, which depend on the gradients of ℓ and g_k.

For discrete or Gaussian algebraic models, the score equations are *rational* functions of their arguments. This fact is not true in general. It is an important advantage of these algebraic models. It allows us to use algebraic techniques for the likelihood estimation problem. To illustrate this, we first need to pass to complex coefficients.

Every algebraic model $\mathcal{M} \subseteq \mathcal{N}$ has an associated embedding $\mathcal{M}_\mathbb{C} \subseteq \mathcal{N}_\mathbb{C}$ of complex algebraic varieties. It is obtained by changing the base field to $\mathbb{C}$ and only considering the defining equalities for $\mathcal{M}$ and $\mathcal{N}$, discarding their defining inequalities. Now we can view the score equations as rational functions of $x \in \mathcal{M}_\mathbb{C}$, $y \in \mathcal{N}_\mathbb{C}$, and $\lambda \in \mathbb{C}^r$.

An argument in algebraic geometry shows that for general $y \in \mathcal{N}_\mathbb{C}$, the number of solutions to the score equations does not depend on y. More precisely, there exists a number d and a Zariski-open subset U of $\mathcal{N}_\mathbb{C}$ such that for all $y \in U$, the number of complex solutions to the score equations, counted with multiplicties, is d. That number is called the *maximum likelihood degree*, or *ML degree*, of the model $\mathcal{M}$. For a detailed explanation of the terminology around "general points" see the end of Section 2.1.

Knowledge of the ML degree of a model is important when applying numerical algebraic geometry methods to solve the maximum likelihood estimation problem. In particular, it gives a stopping criterion for monodromy methods [82, Sec. 5]. More intuitively, the ML degree of a model is interesting because it gives a measure of its "likelihood estimation complexity". By analogy, the algebraic degree of an algebraic variety is a measure of its "intersection-theoretic complexity".

We can define the ML degree for any algebraic model with rational score equations. For the algebraic argument needed to make this notion well-defined, see Lemma A.1. For more background on ML degrees, see [17] and [27, Ch. 2].

When an algebraic model is parametrized by a polynomial map, the score equations become simpler since there is no need to introduce the Lagrange multipliers λ.

Example 1.2.1. The independence model for two binary random variables is parametrized by the map $\varphi\colon [0,1]^2 \to \mathcal{M}$ defined by $\varphi(x, y) = (xy, (1-x)y, x(1-y), (1-x)(1-y))$. With this parametrization, the log-likelihood function is the function obtained by composing the function ℓ from Equation (1.1.2) with φ. For fixed $q \in \Delta_3$, the resulting score equations $\ell_{1,2}\colon [0,1]^2 \times N \to \mathbb{R}$ are

$$\ell_1(x, y) = \frac{q_1}{x} - \frac{q_2}{1-x} + \frac{q_3}{x} - \frac{q_4}{1-x},$$
$$\ell_2(x, y) = \frac{q_1}{y} + \frac{q_2}{y} - \frac{q_3}{1-y} - \frac{q_4}{1-y}.$$

Example 1.2.2. Let A and B be two positive definite symmetric $n \times n$ matrices. The parametrized statistical model

$$\mathcal{M} = \{\Sigma = xA + yB \mid x, y \in \mathbb{R}^2,\ \Sigma \text{ positive definite}\}$$

is the *two-dimensional linear covariance model* on n variables. Its log-likelihood function is obtained by substituting $\Sigma = xA + yB$ in the expression in (1.1.3). The score equations are obtained by differentiating with respect to x and y. Setting $P = \det(xA + yB)$ and $T = \operatorname{tr}(S \operatorname{adj}(xA + yB))$, we thus have the score equations

$$\ell_x(x, y) = \frac{P_x}{P} + \frac{PT_x - TP_x}{P^2}$$
$$\ell_y(x, y) = \frac{P_y}{P} + \frac{PT_y - TP_y}{P^2},$$

where the subscripts indicate differentiation with respect to that variable.

The family of models in Example 1.2.2 is of special interest for this thesis:

Problem (Chapter 2)**.** Find the ML degree of a general two-dimensional linear Gaussian covariance model on n variables.

In Chapter 2, we will see how to solve this problem using the intersection theory of plane curves. The two score equations define two curves in the projective plane with some points removed. By analyzing carefully how they intersect, taking the removed points into account, we arrive at the figure of $2n - 3$ intersection points. This is the ML degree of the model.

Asking the ML degree of some given model has its inverse problem: given an ML degree, we could ask which models of a certain class have it. The case where the ML degree is *one* is special because in that case, the unique solution to the score equations can be expressed as a *rational* function of the data. This is implied by the proof of Lemma A.1. The next section focuses on discrete models with ML degree one.

1.3. ML degree one and rational MLE

Let $\mathcal{M} \subseteq \Delta_n$ be a discrete statistical model with ML degree one. Then there exists a rational map $\Phi_{\mathbb{C}} \colon \Delta_{n,\mathbb{C}} \to \mathcal{M}_{\mathbb{C}}$ sending a point $q \in \Delta_{n,\mathbb{C}}$ to the unique critical point of the score equations. Here, the "$\mathbb{C}$" in the index indicates complexification of algebraic sets as defined in Section 1.2. Restricting to the real numbers, we get a rational map $\Phi \colon \Delta_n \to \mathcal{M}$ that sends an empirical distribution q to the unique maximizer of the log-likelihood function.

Definition 1.3.1. The map Φ is called the *maximum likelihood estimator* of $\mathcal{M}$.

Example 1.3.2. The score equations from Example 1.2.1 for an empirical distribution $q = (q_1, \ldots, q_4)$ are solved by $\hat{x}, \hat{y} \neq 0, 1$ with

$$\hat{x} = \frac{q_1 + q_3}{q_1 + q_2 + q_3 + q_4}, \quad \hat{y} = \frac{q_2 + q_4}{q_1 + q_2 + q_3 + q_4},$$

which are rational expressions in q. Thus, the model has ML degree one with maximum likelihood estimator $\Phi(q) = \varphi(\hat{x}, \hat{y})$.

In Example 1.3.2, we see that the coordinates of Φ are not only rational functions but even alternating products of linear forms in $q = (q_1, \ldots, q_n)$. This is no coincidence. Huh [46] proved that each of the coordinates of $\Phi_{\mathbb{C}}$ is an alternating product of linear forms, with numerator and denominator of the same degree. Huh further showed that this alternating product must take a very specific shape. That shape was discovered by Kapranov [51] who named it the *Horn uniformization*.

Horn uniformizations need not come from statistical models, in general they just have the form $\Phi_{\mathbb{C}} \colon \mathbb{C}^n \to \mathbb{C}^n$. However, any Horn uniformization that restricts to a map $\Phi \colon \Delta_n \to \Delta_n$ gives rise to a statistical model by setting $\mathcal{M} = \mathrm{im}(\Phi)$. Thus, a strategy for characterizing discrete statistical models with ML degee one emerges: characterize precisely those cases where the restriction Φ of a Horn uniformization $\Phi_{\mathbb{C}}$ to the real numbers defines a statistical model $\mathrm{im}(\Phi)$.

Problem (Chapter 3)**.** Characterize all discrete statistical models with rational MLE.

We will solve this problem in Chapter 3 by a topological argument that exploits the real setting, coupled with a careful analysis of the Horn uniformization. Additionally, every Horn uniformization comes with a corresponding toric variety. We will see what conditions the toric variety must satify to define a model. This will give a recipe for creating many examples of models with ML degree one starting from toric varieties.

1.4. Algebraic manifolds

In statistics and applied mathematics, manifolds are useful models for many types of continuous data. For example, in *computational statistical physics* the state space of a collection of particles is a manifold. Each point on this manifold records the positions of all particles in space. Meanwhile, *Topological data analysis* studies the geometric properties of a point cloud in some Euclidean space. Learning the manifold that best explains the position of the points is a research topic in this field. In Section 4.2 we will see examples of both applications.

Definition 1.4.1. An *algebraic manifold* is an open submanifold $\mathcal{M}$ of the set of non-singular solutions to a system of polynomial equations.

Any algebraic statistical model with singularities removed is an example of an algebraic manifold.

Whenever we use a manifold $\mathcal{M}$ to model data, we are interested in performing two basic computations: to integrate a function on $\mathcal{M}$, and to sample from some probability distribution on $\mathcal{M}$. In this thesis, we are interested in the case where $\mathcal{M}$ is algebraic:

Problems (Chapter 4)**.** Let $\mathcal{M}$ be an algebraic manifold given as the solution set of a system of polynomial equations and inequalities.

(1) Approximate the Lebesgue integral $\int_{\mathcal{M}} f(x)\,\mathrm{d}x$ of a given function f on $\mathcal{M}$.

(2) Sample from a probability distribution with a given density on $\mathcal{M}$.

Chapter 4 proposes a new method to solve the above problems based on intersecting $\mathcal{M}$ with random hyperplanes. We achieve this in practice with computational methods from numerical algebraic geometry. The main tool to establish the theoretical results will be the coarea formula from geometric measure theory.

2. Two-dimensional linear covariance models

A linear Gaussian covariance model is a collection of multivariate Gaussian probability distributions whose covariance matrices are linear combinations of some fixed symmetric matrices. This chapter is about the *two-dimensional linear Gaussian covariance model* seen in Example 1.2.2. Here, all of the covariance matrices in the model lie in a two-dimensional linear space. Linear Gaussian covariance models were first studied by Anderson in [2] in the analysis of time series models. They continue to be studied towards this end, for example, in [87].

One common type of linear Gaussian covariance model is obtained by prescribing zero entries in the covariance matrix. These models often have a clear statistical interpretation. For instance, given a Gaussian random vector $(X_1, \ldots, X_n)$ with covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$, we can discern independence statements from the zeros in Σ. For example, the disjoint subvectors $(X_{i_1}, \ldots, X_{i_k})$ and $(X_{j_1}, \ldots, X_{j_l})$ are independent if and only if the submatrix of Σ that consists of rows $i_1, \ldots, i_k$ and columns $j_1, \ldots, j_l$ is the zero matrix [83, Prop. 2.4.4].

Maximum likelihood estimation for covariance matrices with a fixed independence structure was studied in [18]. These types of models find applications for example in the study of gene expression using relevance networks [15]. The nodes of these networks represent genes. They are connected with an edge if their expressions are sufficiently correlated. The edges and non-edges in the resulting graph dictate the sparsity structure of the covariance matrix. Problems related to estimation of sparse covariance matrices have been studied in [6] and [73].

Linear Gaussian covariance models also find applications in the field of phylogenetics. In particular, Brownian Motion tree models, which model the evolution of normally distributed traits along an evolutionary tree, are linear Gaussian covariance models [34]. This chapter focuses on linear combinations of two matrices, while the covariance matrices of Brownian Motion tree models usually require more. However, the results in this chapter could find applications to mixtures of such models. They apply, for example, to models of trait evolution that consider two genes instead of just one [49].

Algorithms for computing the maximum likelihood estimate for generic linear Gaussian covariance models have been extensively studied [1, 2, 6, 18]. Nevertheless, to fully understand these models we need many more theoretical results. Zwiernik, Uhler and Richards have shown that when the number of data points is sufficiently large, maximum

likelihood estimation for such models behaves like a convex optimization problem in a large convex region containing the maximum likelihood estimate [89]. This chapter contributes a result about the ML degree of two-dimensional linear covariance models for an arbitrary number of covariates.

Section 2.1 formally defines the models to consider, states our main result (Theorem 2.1.4), and offers a proof strategy. The remainder of the chapter carries out this strategy, starting with some important geometric preliminaries in Section 2.2. Afterwards, the problem reduces to computing two intersection multiplicities. The first one is done in Section 2.3. The second one and the proof of the main result are detailed in Section 2.4. Section 2.5 closes the chapter with a discussion on higher-dimensional linear covariance models.

2.1. The model and its score equations

Let n be a natural number and let $\mathrm{PD}_n \subset \mathbb{R}^{\binom{n+1}{2}}$ denote the cone of all $n \times n$ symmetric positive definite matrices. We view PD_n as the space of covariance matrices of all normal distributions $\mathcal{N}(0, \Sigma)$ with zero mean.

Definition 2.1.1. Let A and B be symmetric $n \times n$ matrices. The *two-dimensional linear Gaussian covariance model* with respect to A and B is the algebraic subset $\mathcal{M}_{A,B}$ of PD_n defined by

$$\mathcal{M}_{A,B} = \{xA + yB \mid x, y \in \mathbb{R}\} \cap \mathrm{PD}_n.$$

So, $\mathcal{M}_{A,B}$ is the intersection of the positive definite cone with the linear span of A and B. By convention, we assume this intersection to be non-empty.

Given independent, identically distributed (i.i.d.) samples $e_1, \ldots, e_N \in \mathbb{R}^n$ from some normal distribution, the maximum likelihood estimation problem for $\mathcal{M}_{A,B}$ is to find a covariance matrix $\hat{\Sigma} \in \mathcal{M}_{A,B}$, if it exists, that maximizes the value of the log-likelihood

$$\sum_{i=1}^{N} \log f_\Sigma(e_i) \tag{2.1.1}$$

from Section 1.1, where f_Σ is the density of $\mathcal{N}(0, \Sigma)$. Let

$$S \coloneqq \frac{1}{N} \sum_{i=1}^{N} e_i e_i^T.$$

The matrix S is called the *sample covariance matrix*. Since for all Σ the value of (2.1.1) only depends on S, we identify the observation represented by N i.i.d. samples from a normal distribution with the empirical distribution defined by their sample covariance matrix S. More precisely (2.1.1) can be written in terms of S as

$$-\frac{Nn}{2} \log(2\pi) - \frac{N}{2} \log \det(\Sigma) - \frac{N}{2} \mathrm{tr}(S\Sigma^{-1}). \tag{2.1.2}$$

The maximizer of (2.1.2) is equal to the minimizer of the following function.

Definition 2.1.2. The *negated log-likelihood function* ℓ of a positive definite symmetric matrix Σ and a symmetric matrix S is defined by

$$\ell(\Sigma, S) := \log\det(\Sigma) + \operatorname{tr}(S\Sigma^{-1}). \tag{2.1.3}$$

When we restrict to the model $\mathcal{M}_{A,B}$, we require that $\Sigma = xA + yB$ for some $x, y \in \mathbb{R}$ such that $xA + yB$ is positive definite. So the maximum likelihood estimation problem in this case is equivalent to

$$\begin{aligned} \operatorname*{argmin}_{x,y} \quad & \ell(xA + yB, S) \\ \text{subject to} \quad & xA + yB \in \mathrm{PD}_n. \end{aligned}$$

To find local extrema of (2.1.3), we set its gradient equal to 0 and solve for x and y. The two resulting equations are called the score equations.

Definition 2.1.3. The *score equations* for $\mathcal{M}_{A,B}$ are the partial derivatives of the function $\ell(xA + yB, S)$ with respect to x and y. The *maximum likelihood degree* or *ML degree* of $\mathcal{M}_{A,B}$ is the number of complex solutions to the score equations, counted with multiplicities, for a generic sample covariance matrix S.

Definition 2.1.3 refers to a *generic* sample covariance matrix. We give a detailed explanation of this term from algebraic geometry at the end of this section.

The score equations are rational functions of the empirical distribution S. This allows us to analyze them using algebraic methods as explained in Section 1.2. To see this, let $\Sigma = xA + yB$. For the sake of brevity, we will denote

$$P(x, y) = \det\Sigma \quad \text{and} \quad T(x, y) = \operatorname{tr}(S \operatorname{adj}\Sigma),$$

where $\operatorname{adj}\Sigma$ is the classical adjoint. With this notation, the function ℓ takes the form

$$\ell(\Sigma \mid S) = \log P + \frac{T}{P}.$$

Accordingly, the score equations are

$$\begin{aligned} \ell_x(x, y) &= \frac{P_x}{P} + \frac{PT_x - TP_x}{P^2}, \\ \ell_y(x, y) &= \frac{P_y}{P} + \frac{PT_y - TP_y}{P^2}. \end{aligned}$$

Here and throughout, the notation h_x is used for the derivative of a function h with respect to the variable x. We are concerned with values of $(x, y) \in \mathbb{C}^2$ where both of the score equations are zero. We clear denominators by multiplying ℓ_x and ℓ_y by P^2 to obtain two polynomials,

$$\begin{aligned} f(x, y) &:= PP_x + PT_x - TP_x, \\ g(x, y) &:= PP_y + PT_y - TP_y. \end{aligned} \tag{2.1.4}$$

Next, we record the *degrees* of each relevant term for generic A, B and S. Specifically, their total degree with respect to their variables x and y are:

$$\begin{aligned} \deg P &= n, \\ \deg P_x = \deg P_y = \deg T &= n-1, \\ \deg T_x = \deg T_y &= n-2. \end{aligned}$$

A polynomial h is called a *homogeneous form* if each of its terms has the same degree. The polynomials f and g in (2.1.4) can be written as a sum of a homogeneous degree $2n-1$ form with a homogeneous degree $2n-2$ form. For instance, we have $f = (PP_x + PT_x) - (TP_x)$.

The critical points of ℓ are in the variety $V(f,g)$. However, this variety also contains points at which ℓ and the score equations are not defined since we cleared denominators. The ideal whose variety contains precisely the critical points of ℓ is the *saturation*

$$\begin{aligned} J &= \mathcal{I}(f,g) : \langle P \rangle^{\infty} \\ &\coloneqq \{h \in \mathbb{C}[x,y] \mid hP^r \in \mathcal{I}(f,g) \text{ for some } r \geq 0\}. \end{aligned}$$

Here, saturating with $P = \det \Sigma$ removes all points in $V(f,g)$ where the determinant is zero and ℓ is undefined. For more details on the geometric content of saturation, see Chapter 7 of [83]. In Lemmas 2.2.2 and 2.2.3, we will show that $\mathcal{I}(f,g)$ and hence J is zero-dimensional. The ML degree of the model is hence the degree of J. This is equal to the number of isolated points in the variety of J counted with multiplicity. For more background on degrees of general varieties, see [44, Lec. 13] and [77, Ch. 4, Sec. 1.4].

We can now state the main result and offer an outline for its proof which we will follow in the remaining sections.

Theorem 2.1.4. *For generic $n \times n$ symmetric matrices A and B, the maximum likelihood degree of the two-dimensional linear Gaussian covariance model $\mathcal{M}_{A,B}$ is $2n-3$.*

This means that the expected number of complex solutions to the score equations for a model $\mathcal{M}_{A,B}$ is $2n-3$. For the proof strategy of Theorem 2.1.4, we use Bézout's Theorem, a proof of which can be found in Chapter 5.3 of [37]. Our plan to prove Theorem 2.1.4 is to first turn the score equations f and g into two projective plane curves and then to calculate their intersection points. Bézout's Theorem does exactly this, but it overcounts since we added some superfluous points. So, the second part of the proof strategy is to remove those superfluous points.

Theorem 2.1.5 (Bézout's Theorem)**.** *Let H and K be projective plane curves of degrees d_1 and d_2 respectively. Suppose further that H and K share no common component. Then the intersection of H and K is zero-dimensional and the number of intersection points of H and K, counted with multiplicities, is $d_1 d_2$.*

Next, we turn f and g into plane curves. Let $F(x,y,z)$ and $G(x,y,z)$ denote the homogenizations of f and g with respect to z. Then F and G both define projective plane curves

of degree $2n-1$ because this is the degree of the highest degree form appearing in f and g. Lemmas 2.2.2 and 2.2.3 will show that F and G do not share a common component. So, we can apply Bézout's Theorem to count their intersection points.

Let $q = [x : y : z]$ be a point in $\mathbb{CP}^2$. Then by Bézout's theorem,

$$(2n-1)^2 = \sum_{q \in V(F,G)} I_q(F,G), \tag{2.1.5}$$

where $I_q(F,G)$ denotes the *intersection multiplicity* of F and G at q. The definition and properties of the intersection multiplicity of a pair of algebraic curves at a point can be found in [37, Sec. 3, Thm. 3]. For affine points $(x,y) \in V(f,g)$ we sometimes denote the intersection multiplicity as $I_{(x,y)}(f,g) := I_{[x:y:1]}(F,G)$.

We just used Bézout's theorem to count the intersection points of the plane curves defined by F and G. However, we introducted an extraneous factor of $\det \Sigma$ to turn the score equations into polynomials. We also introduced extraneous points at infinity to make the curves projective. In the next few paragraphs, we will see how to exclude the extraneous points.

We will show in Proposition 2.2.4 that saturating the ideal $\mathcal{I}(f,g)$ with $\det \Sigma$ corresponds to removing only the origin from the affine variety of f and g. This in turn corresponds to removing the point $[0:0:1]$ from the projective variety $V(F,G)$. Since we are only interested in the affine intersection points of F and G outside of the origin, we split the sum on the right-hand side of the above equation as follows:

$$(2n-1)^2 = I_{[0:0:1]}(F,G) + \sum_{\substack{q \in V(F,G) \\ q \notin \{[0:0:1]\} \cup V(F,G,z)}} I_q(F,G) + \sum_{q \in V(F,G,z)} I_q(F,G). \tag{2.1.6}$$

The middle term of the right-hand side of (2.1.6) is exactly the degree of the saturated ideal $J = \mathcal{I}(f,g) : \langle \det \Sigma \rangle^\infty$. Thus we can find the degree of J by computing the intersection multiplicities of F and G at the origin and at their intersection points at infinity. We compute the former in Section 2.3 and the latter in Section 2.4 to obtain

$$I_{[0:0:1]}(F,G) = (2n-2)^2 \quad \text{and} \quad \sum_{q \in V(F,G,z)} I_q(F,G) = 2n$$

for generic A, B and S. Thus, by rearranging (2.1.6),

$$\sum_{\substack{q \in V(F,G) \\ q \notin \{[0:0:1]\} \cup V(F,G,z)}} I_q(F,G) = (2n-1)^2 - (2n-2)^2 - 2n = 2n-3,$$

which implies $\deg(J) = 2n-3$, completing the proof.

We dedicate the rest of the section to an example and an explanation of the term "generic". Then, Sections 2.2–2.4 will execute the above proof strategy, filling in the details.

Example 2.1.6. In this example we use two software packages written in the numerical computing programming language `Julia` [5]. The package `LinearCovarianceModels.jl` [84] provides functionality to compute with linear covariance models, in particular we can compute the ML degree of any given model. It is powered by the package `HomotopyContinuation.jl` [12], which enables the user to solve systems of polynomial equations numerically.

Let $n = 3$ and consider the model $\mathcal{M}_{A,B}$ defined by the positive definite matrices

$$A = \begin{pmatrix} 5 & 1 & 0 \\ 1 & 3 & -2 \\ 0 & -2 & 6 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & -1 & 0 \\ -1 & 6 & -2 \\ 0 & -2 & 1 \end{pmatrix}.$$

Using `LinearCovarianceModels.jl` we find that the maximum likelihood degree of $\mathcal{M}_{A,B}$ is $2 \cdot 3 - 3 = 3$. Thus, for a generic sample covariance matrix there will be three complex solutions to the score equations. If we take the sample covariance matrix

$$S = \begin{pmatrix} 1 & 2 & -2 \\ 2 & 6 & -7 \\ -2 & -7 & 9 \end{pmatrix}$$

then the equations f, g from (2.1.4) for $\mathcal{M}_{A,B}$ and S are

$$\begin{aligned} f = 12288x^5 + 57600x^4y + 74272x^3y^2 + 20172x^2y^3 + 1729xy^4 + 37y^5 \\ -10496x^4 - 33792x^3y - 45484x^2y^2 - 7232xy^3 - 513y^4 \end{aligned}$$

and

$$\begin{aligned} g = 11520x^5 + 37136x^4y + 20172x^3y^2 + 3458x^2y^3 + 185xy^4 + 3y^5 \\ -12624x^4 - 9448x^3y - 6480x^2y^2 - 528xy^3 - 21y^4. \end{aligned}$$

We can then use `HomotopyContinuation.jl` to find the solutions to the system of equations $f = 0$ and $g = 0$. The solution set consists of the origin (with multiplicity 16) and three points corresponding to the critical points of the negated log-likelihood function:

$$\begin{aligned} \{(0.6897, 0.1773), (0.2655 + 0.3071i, 0.9865 - 2.4601i), \\ (0.2655 - 0.3071i, 0.9865 + 2.4601i)\}. \end{aligned}$$

The number of critical points and the multiplicity at the origin are predicted by Theorem 2.1.4 and Corollary 2.3.2 respectively. This fits into Equation (2.1.6) which for $n = 3$ becomes $5^2 = 16 + 3 + 6$. The maximum likelihood estimate for $\mathcal{M}_{A,B}$ and S is the real point in the list above, which corresponds to the positive definite covariance matrix

$$\Sigma = \begin{pmatrix} 3.6257 & 0.5124 & 0 \\ 0.5124 & 3.1329 & -1.7340 \\ 0 & -1.7340 & 4.3154 \end{pmatrix}$$

that maximizes the log-likelihood (2.1.1).

Properties that hold generically

The ML degree formula of Theorem 2.1.4 only applies if the matrices A, B, S from which we generate the score equations are "generic enough". In computations, this means using a random number generator to generate the entries of A, B, and S, since the set of triples which are not "generic enough" has Lebesgue measure zero. We now explain the precise notion of genericity in classical algebraic geometry.

Let $\mathcal{M}$ be an algebraic variety and $\mathcal{P}$ a property of the points of $\mathcal{M}$. We say that $\mathcal{P}(x)$ *holds for generic* $x \in \mathcal{M}$, or *holds generically* on $\mathcal{M}$, if there exists a non-empty Zariski open set U of $\mathcal{M}$ such that $\mathcal{P}(x)$ holds for all $x \in U$.

Consider the case $\mathcal{M} = \mathbb{C}^N$. A Zariski open set in $\mathbb{C}^N$ is the complement of a set $V = V(f_1, \ldots, f_k)$ of common zeros of a collection of polynomials $f_1, \ldots, f_k$ in N variables.

Thus, to verify that some property $\mathcal{P}$ holds generically on $\mathbb{C}^N$, we first have to find such a set V with the property that for all x, if $\mathcal{P}(x)$ does not hold then $x \in V$. This verifies that $\mathcal{P}(x)$ holds for all $x \in U$, where $U = \mathbb{C}^N \setminus V$. We also have to verify that U is non-empty which amounts to finding a specific element x_0 such that $x_0 \notin V$.

Note that $\dim V$ is at most $N - 1$. In particular it is expected that a point $x \in \mathbb{C}^N$ taken at random (say, according to the multivariate normal distribution on $\mathbb{R}^{2N}$) will lie in U. This justifies the term "generic".

Suppose that $\mathcal{Q}$ is another property of the points of $\mathcal{M}$ and we want to show that both $\mathcal{P}(x)$ and $\mathcal{Q}(x)$ hold generically on $\mathcal{M}$. Then it is enough to show separately that $\mathcal{P}(x)$ holds for generic x and that $\mathcal{Q}(x)$ holds for generic x. This follows from the fact that the intersection of two non-empty Zariski open sets U_1, U_2 is always a nonempty Zariski open set. In practice, this means that after finding U_1 and U_2 it is enough to find two separate elements $x_1 \in U_1$ and $x_2 \in U_2$ which is often easier than finding an element $x_0 \in U_1 \cap U_2$.

For this chapter, the notion of a property holding generically is important for two reasons. First it is needed for the definition of the ML degree. Indeed, the number (with multiplicities) of solutions $(\hat{x}, \hat{y})$ to the score equations $\ell_x(x, y)$ and $\ell_y(x, y)$ given an empirical covariance matrix S could vary with S. Nevertheless it is constant for generic S, which justifies the use of a single number. Second we consider a family of models $\mathcal{M}_{A,B}$ parametrized by pairs of symmetric matrices (A, B) and compute its ML degree only for generic A, B. To perform the computation, we use several properties that hold for generic A, B, and S. The fact that these properties hold generically is first proved separately for each one of them. Then, we use that generically all of them hold at the same time, as explained above.

2.2. Geometry of the score equations

In this section, we use Bézout's Theorem to derive a formula for computing the maximum likelihood degree of a generic $\mathcal{M}_{A,B}$. We will use Euler's homogeneous function theorem [32, Part I, § 222], which says that if $H(x, y)$ is a homogeneous function of degree m, then $mH = xH_x + yH_y$. We also use the fact that if $H(x, y)$ is a homogeneous polynomial over $\mathbb{C}$, then H factors as a product of linear forms. This follows by applying the fundamental theorem of algebra to the polynomial $H(x, 1)$. We further note that a generic such H factors as a product of *distinct* linear forms, since having a double root is a closed condition on the space of coefficients of a polynomial [70, Sec. 0.12].

We will use the following lemma throughout the rest of the chapter to prove claims of the form "the property $\mathcal{P}$ holds generically".

Lemma 2.2.1. *For generic A, B and S, the following projective varieties are empty:*

(1) $V(P, P_x), V(P, T), V(P, P_y)$

(2) $V(P_x, P_y)$

(3) $V(P_x, T_x), V(P_y, T_y), V(T, T_x), V(T, T_y)$

Proof. The emptiness of the varieties in the statement is an open condition in the space of parameters

$$\mathbb{A} \coloneqq \{(A, B, S) \mid A, B, S \in \mathbb{C}^{\binom{n+1}{2}}\}.$$

For instance, the subset of $\mathbb{A}$ where $V(P, P_x)$ is non-empty is the image of the variety defined by P and P_x in the space $\mathbb{A} \times \mathbb{P}^1_{[x:y]}$ under the first projection. This is a Zariski-closed subset of the parameter space by the projective elimination theorem [24, Ch. 8.5].

To show that the projective varieties in the statement are empty, we show that the polynomials defining them have no common factors.

First, consider the case where A is the $n \times n$ identity matrix, B is the diagonal matrix with diagonal entries $1, \ldots, n$, and $S = uu^T$ where u is the vector of all ones. We have

$$P = \prod_{k=1}^{n}(x + ky) \quad \text{and} \quad P_x = \sum_{k=1}^{n}\prod_{j \neq k}(x + jy).$$

From this we deduce that if $p = x + ky$ is a linear form that divides P, then it does not divide P_x. This shows that $V(P, P_x)$ is empty. The variety $V(P, T)$ is empty as well since $P_x = T$ in this case. We can show that $V(P, P_y)$ is empty in a similar way.

Euler's homogeneous function theorem applied to P says that $nP = xP_x + yP_y$. Since $V(P, P_x)$ is generically empty, the same holds for $V(P_x, P_y)$.

To prove the rest of the statements, we switch to an element (A, B, S) that makes the form of T particularly simple. This is allowed when combining generic properties as explained at the end of Section 2.1.

Let A and B be as before and $u = (1, 0, \ldots, 0)$. In this case we have

$$T = \prod_{k \neq 1} (x + ky) \quad \text{and} \quad P_x = T + (x + y)T_x.$$

Assume p divides P_x and T_x. Then p divides T, hence we may assume $p = x + ky$ with $k \neq 1$. However, we have $p \nmid P_x$ as before. This contradiction shows that $V(P_x, T_x)$ is empty. Similarly, $V(P_y, T_y)$ is empty. This example also has T with no common roots, hence $V(T, T_x)$ and $V(T, T_y)$ are generically empty. □

Now we will show that the projective curves defined by F and G satisfy the hypothesis of Bézout's theorem; that is, that they do not share a common component. This justifies our application of Bézout's theorem and allows us to count the points in their variety. To prove this, we show that the polynomials f and g in (2.1.4) generically are irreducible and do not share a common factor.

Lemma 2.2.2. *The polynomials f and g in* (2.1.4) *are irreducible for generic A, B and S.*

Proof. We prove the statement for f. The proof for g is analogous. We start by writing $f = F_{2n-1} + F_{2n-2}$, where

$$F_{2n-1} = PP_x \quad \text{and} \quad F_{2n-2} = PT_x - TP_x.$$

If f decomposes into a product of two polynomials, then at least one of them is homogeneous and we call it h. Indeed, otherwise the degrees of F_{2n-1} and F_{2n-2} would be at least two apart, when in fact they differ by one. Since h is homogeneous and divides a nonzero sum of homogeneous polynomials, h divides each of the summands F_{2n-1} and F_{2n-2}. Let h_0 be a linear factor of h. Since h_0 divides F_{2n-1} and is irreducible, h_0 divides P or P_x. In the first case, since h_0 divides F_{2n-2}, it would have to divide either T or P_x. This would imply that one of the projective varieties $V(P, T)$ and $V(P, P_x)$ is nonempty. By Lemma 2.2.1 this does not happen generically. In the second case, it would have to divide either P or T_x, which for the same reason does not happen generically. □

Lemma 2.2.3. *For generic A, B and S, the polynomials f and g in* (2.1.4) *are not constant multiples of one another.*

Proof. If f and g are constant multiples of each other, then so are their highest degree terms PP_x and PP_y. This does not happen generically since by Lemma 2.2.1 the projective variety $V(P_x, P_y)$ is generically empty. □

Furthermore, we can describe exactly which points are removed from the affine variety $V(f,g)$ after we saturate with the determinant. For generic parameters, the only point that is removed after saturation is the origin.

Proposition 2.2.4. *For generic A, B and S, we have*

$$V(f,g) \setminus V(\det \Sigma) = V(f,g) \setminus \{(0,0)\}.$$

Proof. Let $q \in V(P,f,g)$. Then $f(q) = -T(q)P_x(q)$ and $g(q) = -T(q)P_y(q)$. In order to have $f(q) = g(q) = 0$, we must either have both $P_x(q) = P_y(q) = 0$ or $T(q) = 0$. By Lemma 2.2.1, for generic A, B and S, both of these imply $q = (0,0)$. □

Proposition 2.2.5. *For generic A and B, the* ML*-degree of the model $\mathcal{M}_{A,B}$ is*

$$(2n-1)^2 - I_{[0:0:1]}(F,G) - \sum_{q \in V(F,G,z)} I_q(F,G).$$

Proof. The ML degree of $\mathcal{M}_{A,B}$ is the degree of the ideal $J = \langle f,g \rangle : (\det \Sigma)^\infty$. The finite set $V(J)$ in the projective plane can be written as

$$V(F,G) \setminus (V(F,G,z) \cup V(\det \Sigma)).$$

By Bézout's Theorem (Theorem 2.1.5), Lemmas 2.2.2 and 2.2.3 imply that the variety $V(F,G)$ is zero-dimensional. Using Proposition 2.2.4 we have

$$\begin{aligned} \deg(J) &= \sum_{\substack{q \in V(F,G) \\ q \notin \{[0:0:1]\} \cup V(F,G,z)}} I_q(F,G) \\ &= \sum_{q \in V(F,G)} I_q(F,G) - I_{[0:0:1]}(F,G) - \sum_{q \in V(F,G,z)} I_q(F,G). \end{aligned}$$

Both F and G have degree $2n-1$. Applying Theorem 2.1.5 to F and G gives the desired equality. □

2.3. Multiplicity at the origin

In this section we compute the intersection multiplicity of the polynomials f, g in (2.1.4) at the origin, denoted by $I_{[0:0:1]}(F,G)$ and also $I_{(0,0)}(f,g)$.

For a polynomial in two variables h there is a notion of *multiplicity* of h at the origin, denoted $m_{(0,0)}(h)$. This is the degree of the lowest-degree summand in the decomposition of h as a sum of homogeneous polynomials (for details, see [37, Sec. 3.1]). Since the polynomials f, g can be written as the sum of a homogeneous degree $2n-2$ form with a

homogeneous degree $2n-1$ form, we have $m_{(0,0)}(f) = m_{(0,0)}(g) = 2n-2$. We have the identity

$$I_{(0,0)}(f,g) = m_{(0,0)}(f) \cdot m_{(0,0)}(g) \tag{2.3.1}$$

if the lowest-degree homogeneous forms of f and g share no common factors [37, Sec. 3.3]. The degree $2n-2$ parts of f and g are $Q = PT_x - TP_x$ and $R = PT_y - TP_y$, respectively.

Proposition 2.3.1. *For generic A, B and S, the polynomials Q and R share no common factor.*

Proof. By the definition of Q and R and two applications of Euler's homogeneous function theorem we have

$$\begin{aligned} xQ + yR &= (xT_x + yT_y)P - (xP_x + yP_y)T \\ &= (2n-2)TP - (2n-1)PT \\ &= -PT. \end{aligned}$$

Assume that Q and R share a common factor p, which we may assume is irreducible. Then p divides PT. So p divides P or p divides T, but not both by Lemma 2.2.1. If p divides P, then since $Q = PT_x - TP_x$ and p is a factor of Q, p also divides TP_x. Similarly if p divides T, then p also divides PT_x. But then either P and TP_x share a common factor, or T and PT_x do. Each of the resulting four further cases does not occur generically by Lemma 2.2.1. □

Corollary 2.3.2. *For generic A, B and S, the intersection multiplicity of f and g at the origin is $(2n-2)^2$.*

Proof. By Proposition 2.3.1, this follows from (2.3.1). □

2.4. Multiplicity at infinity and ML degree

In this section we compute the intersection multiplicity at a point at infinity for the curves $V(f)$ and $V(g)$ defined by the polynomials in (2.1.4) for generic A, B and S. To do this we use the connection between the intersection multiplicity of curves and their series expansions about an intersection point.

Consider an irreducible polynomial h in two variables with $h(0,0) = 0$ and $h_y(0,0) \neq 0$. By [35, Sec. 7.11, Cor. 2], there exists an infinite series $\alpha = \sum_{m=1}^{\infty} a_m t^m$ and an open neighborhood $U \subset \mathbb{C}$ containing $t = 0$ such that $h(t, \alpha(t)) = 0$ for all $t \in U$. The series α is called the *series expansion* of h at the origin. The *valuation* of a series is the number M such that $a_M \neq 0$ and $a_m = 0$ for all $m < M$.

Proposition 2.4.1. *Let h and k be irreducible polynomials in two variables such that h and k vanish at $(0,0)$ and h_y and k_y do not. Let α and β be infinite series expansions of h resp. k at $(0,0)$. The intersection multiplicity $I_{(0,0)}(h,k)$ is the valuation of the series $\alpha - \beta$.*

Proof. By [35, Sec. 8.7], the intersection multiplicity of h and k at $(0,0)$ is the valuation of the infinite series $h(t,\beta(t))$. We prove that this is the same as the valuation of $\alpha - \beta$. First, let $s(t) = \sum_{m=1}^{\infty} s_m t^m$ be any infinite series and write $h = \sum_{i,j} c_{i,j} x^i y^j$, where the sum ranges over the pairs (i,j) with $0 < i + j \leq \deg(h)$. We have

$$\begin{aligned} h(t,s(t)) = \sum_{i,j} c_{i,j} t^i \left(\sum_{m=1}^{\infty} s_m t^m \right)^j &= \sum_{i,j} c_{i,j} t^i \left(\sum_{\nu=0}^{\infty} \left(\sum_{|a|=\nu} s_{a_1} \cdots s_{a_j} \right) t^\nu \right) \\ &= \sum_{i,j} \sum_{\nu=0}^{\infty} \sum_{|a|=\nu} c_{i,j} s_{a_1} \cdots s_{a_j} t^{\nu+i}, \end{aligned}$$

The coefficient r_m of t^m in this infinite series is a finite sum of products of the form $c_{i,j} s_{a_1} \cdots s_{a_j}$ with $a_j \leq m$ and $|a| + i = m$. The term s_m only appears in r_m when $j = 1$ and $i = 0$. Hence, we have $r_m = c_{0,1} s_m + p(s_1, \ldots, s_{m-1})$ for some polynomial p, where $c_{0,1} \neq 0$ since $h_y(0,0) \neq 0$. For example, the coefficient r_0 is zero since h, k vanishing at the origin implies that $c_{0,0}$ and s_0 are zero, and the coefficient of the first non-zero term is given by $r_1 = c_{0,1} s_1 + c_{1,0}$.

Write $\alpha(t) = \sum_{m=1}^{\infty} a_m t^m$ and $\beta(t) = \sum_{m=1}^{\infty} b_m t^m$. Suppose that the valuation of the series $\alpha - \beta$ is M. Then $a_M \neq b_M$ and $a_m = b_m$ for all $m < M$. We show that this is equivalent to $h(t,\beta(t)) = \sum_{m=1}^{\infty} r_m t^m$ having valuation M. Suppose that $M = 1$; then $a_1 \neq b_1$. Since $h(t,\alpha(t))$ is identically zero in a neighborhood of $t = 0$, we have $r_m(a_1, \ldots, a_m) = 0$ for all m. In particular $r_1(a_1) = c_{0,1} a_1 + c_{1,0} = 0$. Since $a_1 \neq b_1$ this implies that $r_1(b_1) = c_{0,1} b_1 + c_{1,0} \neq 0$ and $h(t,\beta(t))$ has valuation one. Similarly if $h(t,\beta(t))$ has valuation one, then $r_1(a_1) \neq r_1(b_1)$ implying $a_1 \neq b_1$. Thus $\alpha - \beta$ has valuation one if and only if $h(t,\beta(t))$ has valuation one.

Now suppose $M > 1$. By the form of r_m it now follows from an inductive argument on m that a_m and b_m agree up to $m = M$ and differ at $m = M+1$ if and only if $r_m(a_1, \ldots, a_m)$ and $r_m(b_1, \ldots, b_m)$ agree up to $m = M$ and differ at $m = M+1$. Since $r_m(a_1, \ldots, a_m) = 0$ for all m, the latter is equivalent to $h(t,\beta(t))$ having valuation M. □

Remark 2.4.2. In the context of Proposition 2.4.1, consider instead polynomials h and k defining the curves $\mathcal{X}$ resp. $\mathcal{Y}$ such that $\mathcal{X}$ and $\mathcal{Y}$ meet at a non-singular point q. Also, let v be a vector such that the directional derivatives h_v and k_v do not vanish at q. Choose an affine-linear transformation $\varphi\colon \mathbb{C}^2 \to \mathbb{C}^2$ sending $(0,0)$ to q and $(0,1)$ to v. Then $I_q(h,k) = I_{(0,0)}(h \circ \varphi, k \circ \varphi)$ and the polynomials $h \circ \varphi, k \circ \varphi$ satisfy the hypotheses of Proposition 2.4.1. Thus we can compute the intersection multiplicity at any non-singular intersection point of h, k using Proposition 2.4.1.

Remark 2.4.3. When the series $\alpha - \beta$ has valuation M, we say that h and k have *contact order* or *order of tangency* $M-1$ at q. Therefore the contact order of two curves at an intersection point is always one less than the intersection multiplicity. For more on contact order of algebraic curves see [74, Ch. 5].

Remark 2.4.4. The fact that the curves $\mathcal{X}$ and $\mathcal{Y}$ have intersection multiplicity one at q if and only if the gradients of h and k at q are linearly independent (see e.g. [37, Sec. 3.3]) arises as a special case of Proposition 2.4.1 once we compute the first terms of the series α and β.

Returning to the expressions in (2.1.4), recall that F and G denote the homogenizations of f and g with respect to the new variable z. The intersection points of $V(f)$ and $V(g)$ at infinity are exactly the variety $V(F, G, z)$.

Lemma 2.4.5. *For generic A, B and S, the projective variety $V(F, G, z)$ consists of n points of the form $[q_1 : q_2 : 0]$ such that $P(q_1, q_2) = 0$.*

Proof. Let $q = [q_1 : q_2 : 0]$ be a projective point of $V(F, G)$. We have

$$F = PP_x + z(PT_x - TP_x)$$
$$G = PP_y + z(PT_y - TP_y),$$

and hence $V(F, G, z)$ consists of points q where $[q_1 : q_2] \in V(PP_x, PP_y)$. Clearly if $P(q_1, q_2) = 0$, then $q \in V(F, G, z)$. These are the only such points since, by Lemma 2.2.1, for generic A, B and S the variety $V(P_x, P_y)$ is empty. The homogeneous binary form $P(x, y)$ factors in n linear forms. These forms are distinct, since a repeated factor would divide both P and P_x, while $V(P, P_x)$ is empty by Lemma 2.2.1. Thus there are n distinct points in $V(F, G, z)$. □

Lemma 2.4.6. *For generic A, B and S, the projective variety $V(P, P_yT_x - P_xT_y)$ is empty.*

Proof. Let $H = P_yT_x - P_xT_y$. By applying Euler's homogeneous function theorem twice in the following chain of equalities, we have

$$nT_xP - yH = T_x(nP - yP_y) + yP_xT_y = P_x(yT_y + xT_x) = (n-1)P_xT.$$

If P and H have an irreducible common factor p, then $p \mid P_xT$. This does not happen generically by Lemma 2.2.1. □

Lemma 2.4.7. *For generic A, B and S, if $q \in V(F, G, z)$ then $I_q(F, G) = 2$.*

Proof. By Lemma 2.4.5, such points are of the form $q = [q_1 : q_2 : 0]$ where $P(q_1, q_2) = 0$. Fix such a point q and assume for simplicity that $q_1 \neq 0$. This is not a restriction since the conditions $q_1 = 0$ and $P(q) = 0$ imply $\det(B) = 0$ which is a closed condition on the parameter space. Thus we can assume q is of the form $[1 : q_2 : 0]$.

Since intersection multiplicity at a point is a local quantity, we may dehomogenize with respect to x and consider the intersection multiplicity of the affine curves $V(F(1,y,z))$ and $V(G(1,y,z))$ at q. We can compute the partial derivatives with respect to y and z:

$$F_y(x,y,z) = P_yP_x + PP_{xy} + z\left(\frac{d}{dy}(PT_x - TP_x)\right), \quad F_z(x,y,z) = PT_x - TP_x,$$
$$G_y(x,y,z) = P_y^2 + PP_{yy} + z\left(\frac{d}{dy}(PT_y - TP_y)\right), \qquad G_z(x,y,z) = PT_y - TP_y. \quad (2.4.1)$$

Consider the polynomials obtained by translating q to $[1:0:0]$ given by $\tilde{F} = F(1, y+q_2, z)$ and $\tilde{G} = G(1, y+q_2, z)$. Then $\tilde{F}_z(1:0:0), \tilde{G}_z(1:0:0) \neq 0$ if and only if $F_z(q), G_z(q) \neq 0$, and from (2.4.1), we have that

$$F_z(q) = (-TP_x)(1,q_2) \quad \text{and} \quad G_z(q) = (-TP_y)(1,q_2).$$

Since $P(1,q_2) = 0$, Lemma 2.2.1 implies that $F_z(q), G_z(q) \neq 0$. Thus there exist series expansions $\alpha = \sum_{m=1}^{\infty} a_m t^m$ and $\beta = \sum_{m=1}^{\infty} b_m t^m$ such that, for all t in a neighborhood of $t = 0$,

$$\tilde{F}\left(1, t, \sum_{m=1}^{\infty} a_m t^m\right) = 0 \quad \text{and} \quad \tilde{G}\left(1, t, \sum_{m=1}^{\infty} b_m t^m\right) = 0,$$

and hence

$$F\left(1, t+q_2, \sum_{m=1}^{\infty} a_m t^m\right) = 0 \quad \text{and} \quad G\left(1, t+q_2, \sum_{m=1}^{\infty} a_m t^m\right) = 0, \quad (2.4.2)$$

in the same neighborhood. Since $I_{[1:0:0]}(\tilde{F}, \tilde{G}) = I_q(F,G)$, by Proposition 2.4.1 we can compute the valuation of the series $\alpha - \beta$ to determine $I_q(F,G)$. Differentiating the expressions in (2.4.2) with respect to t, then substituting $t = 0$ yields

$$F_y(q) + F_z(q)a_1 = 0 \quad \text{and} \quad G_y(q) + G_z(q)b_1 = 0.$$

Thus $a_1 = \frac{-F_y(q)}{F_z(q)}$ and $b_1 = \frac{-G_y(q)}{G_z(q)}$, and $(F_yG_z - F_zG_y)(q) = 0$ implies that $a_1 - b_1 = 0$. By differentiating (2.4.2) twice with respect to t and substituting these values for a_1 and b_1, we can similarly show that

$$a_2 = \left(\frac{-F_{yy}F_z^2 + 2F_{yz}F_yF_z - F_{zz}F_y^2}{2F_z^3}\right)\Bigg|_q$$
$$b_2 = \left(\frac{-G_{yy}G_z^2 + 2G_{yz}G_yG_z - G_{zz}G_y^2}{2G_z^3}\right)\Bigg|_q. \quad (2.4.3)$$

Since we know that $a_1 - b_1 = 0$, the valuation is *at least* two. We now show that the valuation of $\alpha - \beta = \sum_{m=1}^{\infty}(a_m - b_m)t^m$ is exactly two for generic A, B and S. We can verify that $a_2 - b_2 \neq 0$ with the help of the computer algebra system *Maple* [64] by the

following steps. First, compute all second-order derivatives of F and G with respect to y and z in terms of partial derivatives of P and T, by differentiating the expressions in (2.4.1). Then, substitute $P = 0$ and $z = 0$ in these expressions, which corresponds to evaluation at q. Thus from (2.4.3) we obtain expressions for a_2 and b_2 evaluated at q in terms of partial derivatives of P and T. Next, clear denominators in the resulting expression for $a_2 - b_2$, which yields

$$(a_2 - b_2)(q) = (T^4 P_x^2 P_y^4 (P_y T_x - P_x T_y))(1, q_2).$$

Since $P(1, q_2) = 0$, this expression does not generically evaluate to 0 by Lemmas 2.2.1 and 2.4.6. □

Corollary 2.4.8. *For generic A, B and S, we have $\sum_{q \in V(F,G,z)} I_q(F, G) = 2n$.*

Proof. This follows from Lemmas 2.4.5 and 2.4.7. □

Now we can prove the main result that $\deg(J) = 2n - 3$:

Proof of Theorem 2.1.4. Combining Proposition 2.2.5 with Corollaries 2.3.2 and 2.4.8 shows that the ML-degree of $\mathcal{M}_{A,B}$ for generic A and B is

$$(2n - 1)^2 - (2n - 2)^2 - 2n = 2n - 3.$$ □

2.5. Higher dimensions

In [82], Sturmfels, Timme and Zwiernik use numerical algebraic geometry methods implemented in the Julia package `LinearGaussianCovariance.jl` to compute the ML degrees of linear Gaussian covariance models for several values of n and m, where n is the size of the covariance matrix and m is the dimension of model. In this chapter we proved that for $m = 2$ and arbitrary n, the ML degree is $2n - 3$, which agrees with the computations in Table 1 of [82].

For higher dimensional linear spaces, where $m > 2$, the score equations consist of the partial derivatives of $\tilde{\ell}$ with respect to the m parameters of the linear space. Again, in this case, these are rational functions of the data and the parameters. For instance when $m = 3$, we can consider the linear span of three $n \times n$ matrices A, B, and C. Then if $\Sigma = xA + yB + zC$, $P = \det \Sigma$ and $T = \operatorname{tr}(S \operatorname{adj} \Sigma)$, the score equations are

$$\begin{aligned}
\tilde{\ell}_x(x, y, z) &= \frac{P_x}{P} + \frac{PT_x - TP_x}{P^2} \\
\tilde{\ell}_y(x, y, z) &= \frac{P_y}{P} + \frac{PT_y - TP_y}{P^2} \\
\tilde{\ell}_z(x, y, z) &= \frac{P_z}{P} + \frac{PT_z - TP_z}{P^2},
\end{aligned}$$

and we can similarly define polynomials

$$f(x,y,z) := PP_x + PT_x - TP_x$$
$$g(x,y,z) := PP_y + PT_y - TP_y$$
$$h(x,y,z) := PP_z + PT_z - TP_z,$$

such that the ML-degree of the model is the degree of $J = \mathcal{I}(f,g,h) : \langle \det \Sigma \rangle^\infty$. The authors of [82] conjecture that the ML-degree in this case is $3n^2 - 9n + 7$. To prove this conjecture as we did for $m = 2$, one might turn to a higher dimensional generalization of Bézout's Theorem, which says that the number of solutions to $V(f,g,h)$ counted with multiplicity is the product $\deg(f)\deg(g)\deg(h)$ *provided that* $V(f,g,h)$ *is zero-dimensional* (see for example [23, Sec. 3, Ch. 3] or [77, Sec. 2.1, Ch. 3]). This zero-dimensionality restriction is necessary for equality, otherwise the product of the degrees in this case simply gives an upper bound for the number of zero-dimensional solutions counted with multiplicity [36, Thm. 12.3].

Indeed the variety $V(f,g,h)$ contains the one-dimensional affine variety $V(P,T)$ as well as a "curve at infinity" corresponding to the vanishing of P. When $m = 2$, the variety $V(P,T) \subset \mathbb{C}^2$ consisted of only the origin and the elements at infinity were points whose multiplicity we could calculate using properties of curves. This illustrates the added difficulties in counting solutions when moving from planar intersection theory to higher dimensional intersections.

The authors of [82] also consider the *generic diagonal model*, in which the linear space that comprises the model consists of diagonal matrices. Their computations show that for $m = 2$, the ML degree of the generic diagonal model for the first several values of n is also $2n-3$, see [82, Table 2]. It follows from the proof of our result that this ML-degree is indeed $2n-3$ for all n, as the witnesses for the non-emptiness of the open dense sets that we saw in the proof of Lemma 2.2.1 were all diagonal matrices. For $m > 2$ and $n > 3$, the ML-degree of the generic diagonal model is conjectured in [82] to be strictly less than the corresponding generic linear Gaussian covariance model. This suggests that the study of linear Gaussian covariance models of arbitrary dimension will require us to look beyond diagonal matrices as witnesses to the non-triviality of some open conditions.

Indeed, many of the projective varieties in Lemma 2.2.1 are nonempty for diagonal matrices when $m > 2$. For example, when $m \geq 3$, the determinant of P for a diagonal Σ has a nonempty singular locus. Let $m = 3$ and let

$$\Sigma = xA + yB + zC$$

where A, B, and C are the diagonal matrices with diagonal entries $(a_1, \ldots, a_n)$, $(b_1, \ldots, b_n)$ and $(c_1, \ldots, c_n)$, respectively. Then we have

$$P = \prod_{i=1}^{n} (a_i x + b_i y + c_i z).$$

The derivatives of P are of the form

$$P_x = \sum_{i=1}^{n} a_i \prod_{\substack{j=1 \\ j \neq i}}^{n} (a_j x + b_j y + c_j z),$$

and similarly for P_y and P_z. So any projective point in the intersection of linear spaces of the form

$$V(a_i x + b_i y + c_i z) \cap V(a_j x + b_j y + c_j z)$$

for $i \neq j$ is a singular point of P. When $m > 2$, these intersections are nonempty, so such singular points exist.

Thus, when Σ is not defined by diagonal matrices, the problem of finding witnesses to the emptiness of the varieties in Lemma 2.2.1 for arbitrary n is more difficult, which adds another layer of difficulty for establishing the ML degree when $m > 2$. Nevertheless we believe that examining the structure of the score equations for $m = 2$ provides a possible blueprint for approaching the problem for $m > 2$, although it will require different tools from intersection theory.

For the purposes of statistical inference, we are most interested in *real* solutions to the score equations, as these are the ones that may have statistical meaning. Furthermore, it would be nice to understand whether there are truly $2n - 3$ distinct (complex) solutions to the score equations, as opposed to some having higher multiplicity.

Open problems. Regarding the $m = 2$ case, we still have the following open questions.

(1) How many real solutions can the score equations of a generic two-dimensional linear Gaussian covariance model have?

(2) From experiments done so far, it seems that the score equations of a generic two-dimensional linear Gaussian covariance models always has $2n - 3$ *distinct* complex solutions. Is this true?

3. Discrete models with rational MLE

Likelihood inference is a fundamental problem to many applications. For instance parametric statistics makes inferences about model parameters of interest. In this context, the notion of the *maximum likelihood estimator (MLE)* for that set of parameters is central. The MLE is the best guess for the true value of the parameters given the available data. See for instance the textbook [16] for background and examples.

Let $\mathcal{M} \subseteq \Delta_n$ be a discrete statistical model. In this chapter, we take its maximum likelihood estimator, if it exists, to be a function $\Phi\colon \Delta_n \to \mathcal{M}$. The MLE Φ takes an empirical distribution p to a distribution $\hat{p}$ which is best in the sense of likelihood inference. That is, it maximizes the value of of the log-likelihood function (1.1.2).

This general notion of the MLE of a model makes it possible for two fields at the crossroads of mathematics and data science to study its properties. *Information Geometry* [3] views the MLE as the nearest point map with respect to the Kullback-Leibler divergence, whose second derivative defines a Riemannian metric on Δ_n called the Fisher metric. Algebraic Statistics as in [27, 83], on the other hand, is concerned with models $\mathcal{M}$ whose MLE Φ is an algebraic function of u.

The question that we address in this chapter is: *For which models $\mathcal{M}$ is the MLE Φ a rational function in the empirical distribution p?* Having a rational MLE means not only that we have a closed formula for computing maximum likelihood estimates. It is also the simplest possible such formula. Hence, it is desirable to know exactly for which kind of model such a formula exists.

Example 1.1.1 is one such model, well-known to both fields. It belongs to the class of *graphical models*, a broad class of models $\mathcal{M}_G$ associated to graphs G that may be directed or undirected. In fact, all directed graphical models and all undirected *decomposable* graphical models have rational MLE. The book [55] provides proofs of these facts and more background on graphical models. In these cases, the structure of the MLE formula can be read from the graph G used to define the model. As indicated in Section 1.3, such an MLE formula is not only a rational function of the data but decomposes into an alternating product of linear forms, in a way that can again be read from G.

This chapter characterizes all discrete models with rational MLE. To do this, we need to introduce a considerable amount of terminology. In order to help the reader follow along, we start with the punchline by stating the main result (Theorem 3.1.1) right at the beginning of Section 3.1. The remainder of that section then defines all terms appearing in the theorem statement and presents many examples to illustrate it. Section 3.2 discusses

models with rational MLE that are familiar to statisticians, such as decomposable graphical models and Bayesian networks. The focus lies on *staged tree models*, a far-reaching generalization of discrete Bayesian networks, described in the book by Collazo, Görgen and Smith [20]. We will see how the main result applies to these models. Then Section 3.3 presents the proof of Theorem 3.1.1. This is the technical heart of this chapter, building on the likelihood geometry of [47, §3]. Section 3.4 discusses the connection to toric geometry and geometric modeling developed by Clarke and Cox [19]. Section 3.5 presents an algorithm for constructing models with rational MLE, and discusses its implementation and some experiments. The input is an integer matrix representing a toric variety, and the output is a list of models derived from that matrix. Our results suggest that only a very small fraction of Huh's varieties in [46] are statistical models.

In the following sections we will use notation and concepts from toric geometry. The toric variety associated to an integer matrix A by a polynomial parametrization is denoted by Y_A. Its dual toric variety is Y_A^*. There will also be an associated discriminant-like polynomial, denoted Δ or Δ_A. In contrast, for a natural number n, the notation Δ_n still indicates the probability n-simplex. A background in toric geometry is not necessary to understand this chapter. The interested reader is nevertheless referred to the textbook [25].

3.1. How to be rational

Let $\mathcal{M}$ be a discrete statistical model in the open simplex Δ_n with a well-defined maximum likelihood estimator $\Phi : \Delta_n \to \mathcal{M}$. We also write $\Phi : \mathbb{R}^{n+1}_{>0} \to \mathcal{M}$ for the induced map $u \mapsto \Phi(u/|u|)$ on positive integer vectors u. In general, in this chapter we often deal with empirical distributions p and positive integer data vectors u interchangeably. If the $n+1$ coordinates of Φ are rational functions in u, then we say that $\mathcal{M}$ *has rational MLE*. The following is our main result.

Theorem 3.1.1. *The following are equivalent for the statistical model $\mathcal{M}$ with MLE Φ:*

*(1) The model $\mathcal{M}$ has **rational MLE**.*

*(2) There exists a **Horn pair** (H, λ) such that $\mathcal{M}$ is the image of the Horn map*

$$\varphi_{(H,\lambda)} : \mathbb{R}^{n+1}_{>0} \to \mathbb{R}^{n+1}_{>0}.$$

*(3) There exists a **discriminantal triple** $(A, \Delta, \mathbf{m})$ such that $\mathcal{M}$ is the image under the monomial map $\phi_{(\Delta,\mathbf{m})}$ of precisely one orthant (3.1.9) of the dual toric variety Y_A^*.*

If (1)–(3) are true then the MLE of the model satisfies the following relation on the open orthant $\mathbb{R}^{n+1}_{>0}$:

$$\Phi \;=\; \varphi_{(H,\lambda)} \;=\; \phi_{(\Delta,\mathbf{m})} \circ H. \tag{3.1.1}$$

This theorem is relevant to statistics because it reveals when a model has an MLE of the simplest possible closed form. Property (2) says that the polynomials appearing in the numerators and denominators of the rational formulas must factor into linear forms with positive coefficients. Property (3) offers a recipe, based on toric geometry, for explicitly constructing all such models. The advance over [46] is that Theorem 3.1.1 deals with positive real numbers. This adaptation is essential for using Huh's result in statistics.

The goal of this section is to define all the terms seen in parts (2) and (3) of Theorem 3.1.1.

Example 3.1.2. We first discuss Theorem 3.1.1 for a simple experiment: *Flip a biased coin. If it shows heads, flip it again.* This is the model with $n = 2$ given by the tree diagram

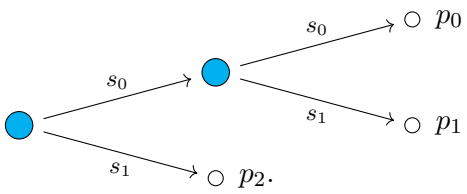

The model $\mathcal{M}$ is a curve in the probability triangle Δ_2. The tree shows its parametrization

$$\Delta_1 \to \Delta_2\,,\ (s_0, s_1) \mapsto (s_0^2, s_0 s_1, s_1) \qquad \text{where } s_0, s_1 > 0 \text{ and } s_0 + s_1 = 1.$$

The implicit representation of the curve $\mathcal{M}$ is the equation $p_0 p_2 - (p_0 + p_1)p_1 = 0$. Let (u_0, u_1, u_2) be the counts from repeated experiments. A total of $2u_0 + 2u_1 + u_2$ coin tosses were made. We estimate the parameters as the empirical frequency of heads resp. tails:

$$\hat{s}_0 \,=\, \frac{2u_0 + u_1}{2u_0 + 2u_1 + u_2} \quad \text{and} \quad \hat{s}_1 \,=\, \frac{u_1 + u_2}{2u_0 + 2u_1 + u_2}.$$

The MLE is the retraction from the triangle Δ_2 to the curve $\mathcal{M}$ given by the formula

$$\begin{aligned} \Phi(u_0, u_1, u_2) \;&=\; (\hat{s}_0^2, \hat{s}_0 \hat{s}_1, \hat{s}_1) \\ &=\; \left(\frac{(2u_0 + u_1)^2}{(2u_0 + 2u_1 + u_2)^2}\,,\, \frac{(2u_0 + u_1)(u_1 + u_2)}{(2u_0 + 2u_1 + u_2)^2}\,,\, \frac{u_1 + u_2}{2u_0 + 2u_1 + u_2} \right). \end{aligned}$$

Hence $\mathcal{M}$ has rational MLE. We see that the Horn pair from part (2) in Theorem 3.1.1 has

$$H \,=\, \begin{pmatrix} 2 & 1 & 0 \\ 0 & 1 & 1 \\ -2 & -2 & -1 \end{pmatrix} \quad \text{and} \quad \lambda \,=\, (1, 1, -1).$$

We next exhibit the discriminantal triple $(A, \Delta, \mathbf{m})$ in part (3) of Theorem 3.1.1. The matrix $A = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}$ gives a basis of the left kernel of H. The second entry is the polynomial

$$\Delta \;=\; x_3^2 - x_1^2 - x_1 x_2 + x_2 x_3 \;=\; (x_3 - x_1)(x_1 + x_2 + x_3). \tag{3.1.2}$$

The third entry marks the leading term $\mathbf{m} = x_3^2$. The discriminantal triple defines the monomial map

$$\phi_{(\Delta,\mathbf{m})} \ : \ (x_1, x_2, x_3) \mapsto \left(\frac{x_1^2}{x_3^2}, \frac{x_1 x_2}{x_3^2}, -\frac{x_2}{x_3} \right).$$

The toric variety of the matrix A is the point $Y_A = \{(1:1:1)\}$ in $\mathbb{P}^2$. Our polynomial Δ vanishes on the line $Y_A^* = \{x_1 + x_2 + x_3 = 0\}$ that is dual to Y_A. The relevant orthant is the open line segment $Y_{A,\sigma}^* := \{(x_1 : x_2 : x_3) \in Y_A^* \ : \ x_1, x_2 > 0 \text{ and } x_3 < 0\}$. Part (3) in Theorem 3.1.1 says that $\mathcal{M}$ is the image of $Y_{A,\sigma}^*$ under $\phi_{(\Delta,\mathbf{m})}$. The MLE is $\Phi = \phi_{(\Delta,\mathbf{m})} \circ H$.

We now come to the definitions needed for Theorem 3.1.1. Let $H = (h_{ij})$ be an $m \times (n+1)$ integer matrix whose columns sum to zero, i.e. $\sum_{i=1}^m h_{ij} = 0$ for $j = 0, \ldots, n$. We call such a matrix a *Horn matrix* and denote its columns by $h_0, h_1, \ldots, h_n$. The following alternating products of linear forms are rational functions of degree zero:

$$(Hu)^{h_j} \ := \ \prod_{i=1}^m \big(h_{i0}u_0 + h_{i1}u_1 + \cdots + h_{in}u_n\big)^{h_{ij}} \qquad \text{for } j = 0, 1, \ldots, n.$$

The Horn matrix H is *friendly* if there exists a real vector $\lambda = (\lambda_0, \ldots, \lambda_n)$ with $\lambda_i \neq 0$ for all i such that the following identity holds in the rational function field $\mathbb{R}(u_0, u_1, \ldots, u_n)$:

$$\lambda_0 (Hu)^{h_0} + \lambda_1 (Hu)^{h_1} + \cdots + \lambda_n (Hu)^{h_n} \ = \ 1. \tag{3.1.3}$$

If this holds, then we call (H, λ) a *friendly pair*, and we consider the rational function

$$\mathbb{R}^{n+1} \to \mathbb{R}^{n+1}, \ u \mapsto \big(\lambda_0 (Hu)^{h_0}, \lambda_1 (Hu)^{h_1}, \ldots, \lambda_n (Hu)^{h_n}\big). \tag{3.1.4}$$

The friendly pair (H, λ) is called a *Horn pair* if the function (3.1.4) is defined for all positive vectors and if it maps these to positive vectors. If these conditions hold then we write $\varphi_{(H,\lambda)} : \mathbb{R}^{n+1}_{>0} \to \mathbb{R}^{n+1}_{>0}$ for the restriction of (3.1.4) to the positive orthant. We call $\varphi_{(H,\lambda)}$ the *Horn map* associated to the Horn pair (H, λ).

The difference between our Horn pairs and the more general pairs considered by Huh in [46] is the positivity condition we just introduced, along with the "friendliness" condition. These conditions guarantee that the image of the Horn map lies in the probability simplex, which is necessary for its interpretation as a statistical model. They also imply special properties for the Horn pair, see Propositions 3.3.2 and 3.3.3 in Section 3.3. The examples in Section 3.5 show that only a fraction of Huh's pairs (H, λ) are Horn pairs.

Different Horn pairs may give rise to the same Horn map. For example, the Horn pair

$$H' \ = \ \begin{pmatrix} 0 & 2 & 2 \\ 2 & 1 & 0 \\ 0 & -1 & -1 \\ -2 & -2 & -1 \end{pmatrix} \quad \text{and} \quad \lambda' \ = \ \left(1, -\frac{1}{4}, \frac{1}{4}\right)$$

also gives the map in Example 3.1.2. This is because the first and third rows of H' are collinear, causing the cancellation of linear factors in the Horn map. Following [19], a Horn pair (H, λ) is *minimal* if the matrix H has no zero rows and no pair of collinear rows.

Lemma 3.1.3. *Let (H', λ') be a Horn pair arising from the Horn pair (H, λ) by replacing two collinear rows r_k and r_ℓ in H such that $r_\ell = \mu r_k$ with their sum $r_k + r_\ell$ and setting*

$$\lambda'_j \;=\; \frac{\lambda_j \mu^{\mu \cdot h_{kj}}}{(1+\mu)^{(1+\mu)h_{kj}}} \quad \text{for all } j = 0, \ldots, n.$$

Then the Horn maps $\varphi_{(H',\lambda')}$ and $\varphi_{(H,\lambda)}$ are equal.

Proof. Let w_k and w_ℓ be the linear forms associated to the rows r_k and r_ℓ respectively. Fix a column index j. We have $w_\ell = \mu w_k$ and $h_{\ell j} = \mu h_{kj}$. The factors of the j-th coordinates of the Horn maps $\varphi_{(H,\lambda)}$ and $\varphi_{(H',\lambda')}$ that have changed after the operation are

$$\begin{aligned} \lambda_j w_k^{h_{kj}} w_\ell^{h_{\ell j}} &= \lambda_j \mu^{\mu \cdot h_{kj}} w_k^{(1+\mu)h_{kj}} && \text{for } (H, \lambda) \text{ and} \\ \lambda'_j (w_k + w_\ell)^{(1+\mu)h_{kj}} &= \lambda'_j (1+\mu)^{(1+\mu)h_{kj}} w_k^{(1+\mu)h_{kj}} && \text{for } (H', \lambda'). \end{aligned}$$

Equating these two gives the desired formula. □

Every Horn map is represented by a unique minimal Horn pair. To make a Horn pair minimal, while retaining the Horn map, we can use Lemma 3.1.3 repeatedly, deleting zero rows as they appear. This follows by unique factorization, see also [19, Prop. 6.11].

Example 3.1.4. We illustrate the equivalence of (1) and (2) in Theorem 3.1.1 for the model described in [47, Expl. 3.11]. Here $n = 3$ and $m = 4$ and the Horn matrix equals

$$H \quad = \quad \begin{pmatrix} -1 & -1 & -2 & -2 \\ 1 & 0 & 3 & 2 \\ 1 & 3 & 0 & 2 \\ -1 & -2 & -1 & -2 \end{pmatrix}. \tag{3.1.5}$$

This Horn matrix is friendly because the following vector satisfies the identity (3.1.3):

$$\lambda \;=\; (\lambda_0, \lambda_1, \lambda_2, \lambda_3) \;=\; \left(\frac{2}{3}\,,\, -\frac{4}{27}\,,\, -\frac{4}{27}\,,\, \frac{1}{27} \right). \tag{3.1.6}$$

The pair (H, λ) is a Horn pair, with associated Horn map

$$\varphi_{(H,\lambda)} : \mathbb{R}^4_{>0} \,\to\, \mathbb{R}^4_{>0}\,,\; \begin{pmatrix} u_0 \\ u_1 \\ u_2 \\ u_3 \end{pmatrix} \mapsto \begin{pmatrix} \frac{2(u_0+3u_2+2u_3)(u_0+3u_1+2u_3)}{3(u_0+u_1+2u_2+2u_3)(u_0+2u_1+u_2+2u_3)} \\ \frac{4(u_0+3u_1+2u_3)^3}{27(u_0+u_1+2u_2+2u_3)(u_0+2u_1+u_2+2u_3)^2} \\ \frac{4(u_0+3u_2+2u_3)^3}{27(u_0+u_1+2u_2+2u_3)^2(u_0+2u_1+u_2+2u_3)} \\ \frac{(u_0+3u_2+2u_3)^2(u_0+3u_1+2u_3)^2}{27(u_0+u_1+2u_2+2u_3)^2(u_0+2u_1+u_2+2u_3)^2} \end{pmatrix}. \tag{3.1.7}$$

Indeed, this rational function takes positive vectors to positive vectors. The image of the map $\varphi_{(H,\lambda)}$ is a subset $\mathcal{M}$ of the tetrahedron $\Delta_3 = \{p \in \mathbb{R}^4_{>0} : p_0+p_1+p_2+p_3 = 1\}$. We regard the subset $\mathcal{M}$ as a discrete statistical model on the state space $\{0,1,2,3\}$. The model $\mathcal{M}$ is the curve of degree 4 inside Δ_3 defined by the two quadratic equations

$$9p_1p_2 - 8p_0p_3 \,=\, p_0^2 - 12p_3 \,=\, 0.$$

As in [47, Expl. 3.11], one verifies that $\mathcal{M}$ has rational MLE, namely $\Phi = \varphi_{(H,\lambda)}$.

We next define all the terms used in part (3) of Theorem 3.1.1. Fix a matrix $A \in \mathbb{Z}^{r\times m}$ of rank r with entries (a_{ij}) that has the vector $(1,\ldots,1)$ in its row span. The connection to part (2) of Theorem 3.1.1 will be that the rows of A span the left kernel of H. We identify the columns of A with Laurent monomials in r unknowns $t_1,\ldots,t_r$. The associated monomial map is

$$\gamma_A \,:\, (\mathbb{R}^*)^r \to \mathbb{RP}^{m-1}\,,\;\; (t_1,\ldots,t_r) \mapsto \Big(\prod_{i=1}^r t_i^{a_{i1}} : \prod_{i=1}^r t_i^{a_{i2}} : \,\cdots\, : \prod_{i=1}^r t_i^{a_{im}}\Big). \qquad (3.1.8)$$

Here $\mathbb{R}^* = \mathbb{R}\backslash\{0\}$ and $\mathbb{RP}^{m-1}$ denotes the real projective space of dimension $m-1$. Let Y_A be the closure of the image of γ_A. This is the projective toric variety given by A.

Every point $x = (x_1 : \cdots : x_m)$ in the dual projective space $(\mathbb{RP}^{m-1})^\vee$ corresponds to a hyperplane H_x in $\mathbb{RP}^{m-1}$. The *dual variety* Y_A^* to the toric variety Y_A is the closure of

$$\big\{\, x \in (\mathbb{RP}^{m-1})^\vee \;\mid\; \gamma_A^{-1}(H_x \cap Y_A) \;\text{ is singular}\,\big\}.$$

Here, the term *singular* means that the variety $\gamma_A^{-1}(H_x \cap Y_A)$ has a singular point in $(\mathbb{R}^*)^r$. A general point x in Y_A^* hence corresponds to a hyperplane H_x that is tangent to the toric variety Y_A at a point $\gamma_A(t)$ with nonzero coordinates. We identify sign vectors $\sigma \in \{-1,+1\}^m$ with orthants in $\mathbb{R}^m$. These map in a 2-to-1 manner to orthants in $\mathbb{RP}^{m-1}$. If we intersect them with Y_A^*, then we get the *orthants* of the dual toric variety:

$$Y^*_{A,\sigma} \;=\; \big\{\, x \in Y_A^* \,:\, \sigma_i \cdot x_i > 0 \text{ for } i = 1,2,\ldots,m \,\big\} \quad\subset\quad \mathbb{RP}^{m-1}. \qquad (3.1.9)$$

One of these is the distinguished orthant in Theorem 3.1.1, part (3).

Example 3.1.5. Fix $m = 4$ and $r = 2$. The following matrix has $(1,1,1,1)$ in its row span:

$$A \,=\, \begin{pmatrix} 3 & 2 & 1 & 0 \\ 0 & 1 & 2 & 3 \end{pmatrix}. \qquad (3.1.10)$$

As in [47, Expl. 3.9], the toric variety of A is the *twisted cubic curve* in 3-space:

$$Y_A \,=\, \overline{\big\{(t_1^3 : t_1^2t_2 : t_1t_2^2 : t_2^3) \in \mathbb{RP}^3 \,:\, t_1,t_2 \in \mathbb{R}^*\big\}}.$$

The dual toric variety Y_A^* is a surface in $(\mathbb{RP}^3)^\vee$. Its points x represent planes in $\mathbb{RP}^3$ that are tangent to the curve Y_A. Such a tangent plane corresponds to a univariate cubic

$x_1t^3+x_2t^2+x_3t+x_4$ with a double root. Just as we recognize quadrics with a double root by the vanishing of the quadratic discriminant, a cubic with coefficients (x_1, x_2, x_3, x_4) has a double root if and only if the following discriminant vanishes:

$$\Delta_A \;=\; \underline{27x_1^2x_4^2} - 18x_1x_2x_3x_4 + 4x_1x_3^3 + 4x_2^3x_4 - x_2^2x_3^2. \tag{3.1.11}$$

Hence, Y_A^* is the surface of degree 4 in $(\mathbb{RP}^3)^\vee$ defined by Δ_A. All eight orthants $Y_{A,\sigma}^*$ are non-empty. The coefficient vectors of the following eight cubics lie on different orthants:

$$\begin{matrix}(t+1)^2(t+3),\ (t+5)^2(t-1),\ (t-1)^2(t+3),\ (t+5)^2(t-8),\\ (t-3)^2(t+1),\ (t-1)^2(t-3),\ \underline{(t-2)^2(t+3)},\ (t+1)^2(t-3).\end{matrix}$$

For instance, the underlined cubic corresponds to the point $x = (1, -1, -8, 12)$ in the orthant $Y_{A,\sigma}^*$ associated with the sign vector $\sigma = (+1, -1, -1, +1)$.

Let Δ be a homogeneous polynomial in m variables with $n+2$ monomials and $\mathbf{m}$ one of these monomials. There is a one-to-one correspondence between such pairs $(\Delta, \mathbf{m})$ and pairs (H, λ) where H is a Horn matrix of size $m \times (n+1)$ and λ is a coefficient vector. Namely, for $k = 0, \ldots, n$ write h_k^+ resp. h_k^- for the positive resp. negative part of the column vector h_k, so that $h_k = h_k^+ - h_k^-$. In addition, let $\max_k(h_k^-)$ be the entrywise maximum of the h_k^-. We pass from pairs (H, λ) to pairs $(\Delta, \mathbf{m})$ as follows:

$$\mathbf{m} \;=\; x^{\max_k(h_k^-)} \quad \text{and} \quad \Delta \;=\; \mathbf{m} \cdot \left(1 - \sum_{k=0}^{n} \lambda_k x^{h_k}\right). \tag{3.1.12}$$

For the converse, from pairs $(\Delta, \mathbf{m})$ to pairs (H, λ), we divide Δ by $\mathbf{m}$ and use the same equations to determine the pair (H, λ). Note that the polynomial Δ being homogeneous and the matrix H being a Horn matrix are equivalent conditions using the equations (3.1.12). Given a pair $(\Delta, \mathbf{m})$ with associated pair (H, λ), we define the monomial map

$$\phi_{(\Delta,\mathbf{m})} \,:\, (\mathbb{R}^*)^m \to \mathbb{R}^{n+1},\ x \,\mapsto\, \big(\lambda_0 x^{h_0},\, \lambda_1 x^{h_1},\, \ldots,\, \lambda_n x^{h_n}\big).$$

We now come to the definition that is needed for part (3) of Theorem 3.1.1.

Definition 3.1.6. A *discriminantal triple* $(A, \Delta, \mathbf{m})$ consists of

(1) an $r \times m$ integer matrix A of rank r having $(1, 1, \ldots, 1)$ in its row span,

(2) an A-homogeneous polynomial Δ that vanishes on the dual toric variety Y_A^*,

(3) a distinguished term $\mathbf{m}$ among those that occur in the polynomial Δ,

such that the pair (H, λ) associated to $(\Delta, \mathbf{m})$ is a Horn pair. Here, the polynomial Δ being *A-homogeneous* means that $Av = Aw$ for any two exponent vectors v and w of Δ.

All definitions are now complete. Here is Definition 3.1.6 for our running example:

Example 3.1.7. Let A be the 2×4 matrix in (3.1.10), $\Delta = \Delta_A$ its discriminant in (3.1.11), and $\mathbf{m} = 27x_1^2x_4^2$ the special term. Then $(A, \Delta, \mathbf{m})$ is a discriminantal triple with associated sign vector $\sigma = (+1, -1, -1, +1)$. The orthant $Y^*_{A,\sigma}$, highlighted in Example 3.1.5, is a semialgebraic surface in $Y^*_A \subset \mathbb{RP}^3$. This surface is mapped into the tetrahedron Δ_3 by

$$\phi_{(\Delta,\mathbf{m})} : (x_1, x_2, x_3, x_4) \mapsto \left(\frac{2}{3}\frac{x_2x_3}{x_1x_4}, -\frac{4}{27}\frac{x_3^3}{x_1x_4^2}, -\frac{4}{27}\frac{x_2^3}{x_1^2x_4}, \frac{1}{27}\frac{x_2^2x_3^2}{x_1^2x_4^2}\right). \tag{3.1.13}$$

The image of this map is a curve in Δ_3, namely the model $\mathcal{M}$ in Example 3.1.4. We verify (3.1.1) by comparing (3.1.7) with (3.1.13). The former is obtained from the latter by setting $x = Hu$.

3.2. Staged trees

We consider contingency tables $u = (u_{i_1 i_2 \cdots i_m})$ of format $r_1 \times r_2 \times \cdots \times r_m$. Following [27,55], these represent joint distributions of discrete statistical models with $n + 1 = r_1 r_2 \cdots r_m$ states. Namely, the contingency table u represents the probability distribution $p := u/|u|$. For any subset $C \subset \{1, \ldots, m\}$, we consider the marginal table u_C that is obtained by summing out all indices not in C. The entries of the marginal table u_C are sums of entries in u. To obtain the entry $u_{I,C}$ of u_C for any state $I = (i_1, i_2, \ldots, i_m)$, we fix the indices of the states in C and sum over the indices not in C. For example, if $m = 4$, $C = \{1, 3\}$, $I = (i, j, k, l)$, then u_C is the $r_1 \times r_3$ matrix with entries

$$u_{I,C} \;=\; u_{i+k+} \;=\; \sum_{j=1}^{r_2}\sum_{l=1}^{r_4} u_{ijkl}.$$

Such linear forms are the basic building blocks for familiar models with rational MLE.

Consider an undirected graph G with vertex set $\{1, \ldots, m\}$ which is assumed to be *chordal*. This means that every induced cycle of G has exactly three vertices. The associated *decomposable graphical model* $\mathcal{M}_G$ in Δ_n has the rational MLE

$$\hat{p}_I \;=\; \frac{\prod_C u_{I,C}}{\prod_S u_{I,S}}, \tag{3.2.1}$$

where the product in the numerator is over all maximal cliques C of G, and the product in the denominator is over all separators S in a junction tree for G. See [55, §4.4.1]. We shall regard G as a directed graph, with edge directions given by a perfect elimination ordering on the vertex set $\{1, \ldots, m\}$. This turns $\mathcal{M}_G$ into a Bayesian network. More generally, a *Bayesian network* $\mathcal{M}_G$ is given by a directed acyclic graph G. We write $\mathrm{pa}(j)$ for the set of parents of the node j. The model $\mathcal{M}_G$ in Δ_n has the rational MLE

$$\hat{p}_I \;=\; \prod_{j=1}^{m} \frac{u_{I,\mathrm{pa}(j)\cup\{j\}}}{u_{I,\mathrm{pa}(j)}}. \tag{3.2.2}$$

If G comes from an undirected chordal graph then (3.2.1) arises from (3.2.2) by cancellations.

In the following examples we will specify small graphs by a complete list of their edges, each written as $[ij]$ where i and j represent two nodes of the graph.

Example 3.2.1 ($m = 4$)**.** We revisit two examples from page 36 in [27, §2.1]. The *star graph* $G = [14][24][34]$ is chordal. The MLE for $\mathcal{M}_G$ is the map Φ with coordinates

$$\hat{p}_{ijkl} \;=\; \frac{u_{i++l} \cdot u_{+j+l} \cdot u_{++kl}}{u_{++++} \cdot u_{+++l}^2} \;=\; \frac{u_{i+++}}{u_{++++}} \cdot \frac{u_{+j+l}}{u_{+++l}} \cdot \frac{u_{++kl}}{u_{+++l}} \cdot \frac{u_{i++l}}{u_{i+++}}.$$

The left expression is (3.2.1). The right is (3.2.2) for the directed graph $1 \to 4$, $4 \to 2$, $4 \to 3$.

The *chain graph* $G = [12][23][34]$ is chordal. Its MLE is the map Φ with coordinates

$$\hat{p}_{ijkl} \;=\; \frac{u_{ij++} \cdot u_{+jk+} \cdot u_{++kl}}{u_{+j++} \cdot u_{++k+} \cdot u_{++++}} \;=\; \varphi_{(H,\lambda)}(u)_{ijkl}.$$

This is the Horn map given by H as in Figure 3.2.1 and $\lambda = (1, \ldots, 1)$.

All Bayesian networks are special cases of staged tree models, which we introduce below. All staged tree models have rational MLE, and under this interpretation Equation (3.2.2) becomes a special case of a more general MLE formula for staged tree models, which is given in Proposition 3.2.4

Staged trees were introduced by Smith and Anderson [80] as a generalization of discrete Bayesian networks. They furnish an intuitive representation of many situations that the above graphs G cannot capture. In spite of their wide scope, staged tree models are appealing because of their intuitive formalism for encoding events. For an introduction see the textbook [20]. In what follows we study parts (1) and (2) in Theorem 3.1.1 for staged trees.

To define a *staged tree model*, we consider a directed rooted tree $\mathcal{T}$ with at least two edges emanating from each non-leaf vertex, a label set $S = \{s_i \mid i \in I\}$, and a labeling $\theta\colon \mathrm{E}(\mathcal{T}) \to S$ of the edges of the tree. Each vertex of $\mathcal{T}$ has a corresponding *floret*, which is the multiset of edge labels emanating from it. The labeled tree $\mathcal{T}$ is a *staged tree* if any two florets are either equal or disjoint. Two vertices in $\mathcal{T}$ are said to be in the same stage if their corresponding florets are the same. From now on, F denotes the set of florets of $\mathcal{T}$.

Definition 3.2.2. Let J be the set of root-to-leaf paths in the tree $\mathcal{T}$. We set $|J| = n + 1$. For $i \in I$ and $j \in J$, let μ_{ij} denote the number of times edge label s_i appears in the j-th root-to-leaf path. The *staged tree model* $\mathcal{M}_\mathcal{T}$ is the image of the parametrization

$$\phi_\mathcal{T} : \Theta \to \Delta_n\,,\ (s_i)_{i\in I} \mapsto (p_j)_{j\in J},$$

where the parameter space is $\Theta := \big\{(s_i)_{i\in I} \in (0,1)^{|I|} \mid \sum_{s_i \in f} s_i = 1 \text{ for all } f \in F\big\}$, and $p_j = \prod_{i\in I} s_i^{\mu_{ij}}$ is the product of the edge parameters on the j-th root-to-leaf path.

In the model $\mathcal{M}_\mathcal{T}$, the tree $\mathcal{T}$ represents possible sequences of events. The parameter s_i associated to an edge vv' is the transition probability from v to v'. All parameter labels in a floret sum to 1. The fact that distinct nodes in $\mathcal{T}$ can have the same floret of parameter labels enables staged tree models to encode conditional independence statements [80]. This allows us to represent any discrete Bayesian network or decomposable model as a staged tree model. Our first staged tree was seen in Example 3.1.2. Here is another specimen.

Example 3.2.3 ($n = 15$)**.** Consider the decomposable model for binary variables given by the 4-chain $G = [12][23][34]$ as in Example 3.2.1. Figure 3.2.1 shows a realization of $\mathcal{M}_G$ as a staged tree model $\mathcal{M}_\mathcal{T}$. The leaves of $\mathcal{T}$ represent the outcome space $\{0,1\}^4$. Nodes with the same color have the same associated floret. The blank nodes all have different florets. The seven florets of $\mathcal{T}$ are

$$f_1{=}\{s_0, s_1\}, f_2{=}\{s_2, s_3\}, f_3{=}\{s_4, s_5\}, f_4{=}\{s_6, s_7\},$$
$$f_5{=}\{s_8, s_9\}, f_6{=}\{s_{10}, s_{11}\}, f_7{=}\{s_{12}, s_{13}\}.$$

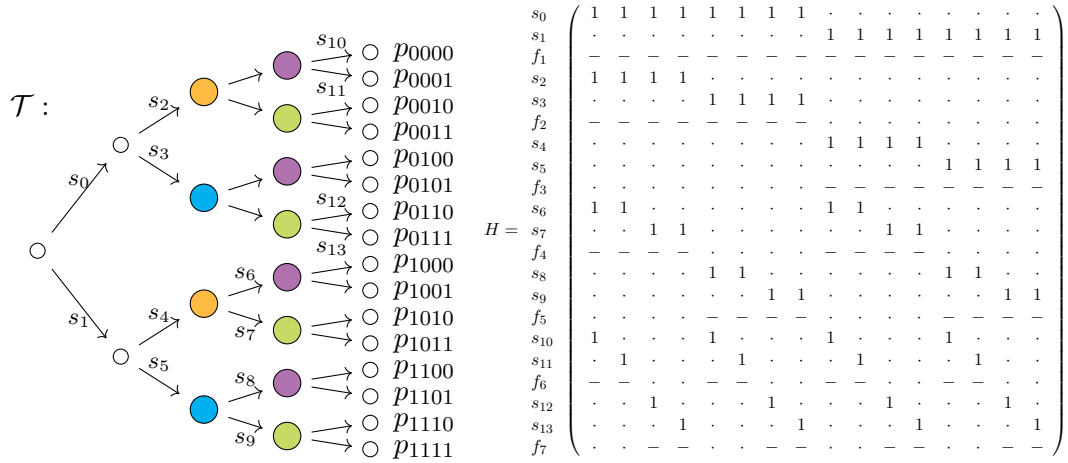

Figure 3.2.1: A staged tree $\mathcal{T}$ and its Horn matrix H in Proposition 3.2.4. A dot represents the entry 0 and a minus sign represents the entry -1.

Next we show that staged tree models have rational MLE, so they satisfy part (1) of Theorem 3.1.1. Our formula for Φ uses the notation for I, J and μ_{ij} introduced in Definition 3.2.2. This formula is known in the literature on chain event graphs (see e.g. [78]).

Proposition 3.2.4. *Let $\mathcal{M}_\mathcal{T}$ be a staged tree model, and let $u = (u_j)_{j\in J}$ be a vector of counts. For $i \in I$, let f be the floret containing the label s_i, and define the estimates*

$$\hat{s}_i \;:=\; \frac{\sum_j \mu_{ij} u_j}{\sum_{s_\ell \in f} \sum_j \mu_{\ell j} u_j} \quad \text{and} \quad \hat{p}_j \;:=\; \prod_{i\in I} (\hat{s}_i)^{\mu_{ij}}.$$

The rational function Φ that sends $(u_j)_{j\in J}$ to $(\hat{p}_j)_{j\in J}$ is the MLE of the model $\mathcal{M}_\mathcal{T}$.

Proof. We prove that the likelihood function $L(p,u)$ has a unique maximum at $p=(\hat{p}_j)_{j\in J}$. For a floret $f\in F$, we fix the vector of parameters $s_f=(s_i)_{s_i\in f}$, and we define the local likelihood function $L_f(s_f,u)=\prod_{s_i\in f}s_i^{\alpha_i}$, where $\alpha_i=\sum_j\mu_{ij}u_j$. We have

$$L(p,u)\;=\;\prod_j p_j^{u_j}\;=\;\prod_j\prod_i s_i^{u_j\mu_{ij}}\;=\;\prod_i s_i^{\alpha_i}\;=\;\prod_{f\in F}L_f(s_f,u).$$

Since the L_f depend on disjoint sets of unknowns, maximizing L is achieved by maximizing the factors L_f separately. But L_f is the likelihood function of the full model $\Delta_{|f|-1}$, given the data vector $(\alpha_i)_{s_i\in f}$. The MLE of that model is $\hat{s}_i=\alpha_i/\sum_{s_\ell\in f}\alpha_\ell$, where $s_i\in f$. Hence, $\operatorname{argmax}_{s_f}\big(L_f(s_f,u)\big)=(\hat{s}_i)_{s_i\in f}$ and $\operatorname{argmax}_p\big(L(p,u)\big)=(\hat{p}_j)_{j\in J}$. □

Remark 3.2.5. Here is a method for evaluating the MLE in Proposition 3.2.4. Let $[v]\subset J$ be the set of root-to-leaf paths through a node v in the tree $\mathcal{T}$ and define $u_{[v]}=\sum_{j\in[v]}u_j$. The ratio $u_{[v']}/u_{[v]}$ is the empirical transition probability from v to v' given arrival at v. To obtain $\hat{s}_i$ we first compute the quotients $u_{[v']}/u_{[v]}$ for all edges vv' with parameter label s_i. We aggregate them by adding their numerators and denominators separately. This gives $s_i=(\sum u_{[v']})/(\sum u_{[v]})$, where both sums range over all edges vv' with parameter label s_i.

Proposition 3.2.4 yields an explicit description of the Horn pair (H,λ) associated to the model $\mathcal{M}_\mathcal{T}$.

Corollary 3.2.6. *Fix a staged tree model $\mathcal{M}_\mathcal{T}$ as above. Let H be the $(|I|+|F|)\times|J|$ matrix whose rows are indexed by the set $I\sqcup F$ and entries are given by*

$$\begin{aligned} h_{ij}\;&=\;\mu_{ij}\ \text{for } i\in I,\ \text{and}\\ h_{fj}\;&=\;-\sum_{s_\ell\in f}\mu_{\ell j}\ \text{for } f\in F.\end{aligned}$$

Define $\lambda\in\{-1,+1\}^{|J|}$ by $\lambda_j=(-1)^{\sum_f h_{fj}}$. Then (H,λ) is a Horn pair for $\mathcal{M}_\mathcal{T}$.

Given a staged tree $\mathcal{T}$, we call the matrix H in Corollary 3.2.6 the *Horn matrix* of $\mathcal{T}$.

Remark 3.2.7. In Corollary 3.2.6, for a floret f, let H_f be the submatrix of H with row indices $\{i:s_i\in f\}\cup\{f\}$. Then H is the vertical concatenation of the matrices H_f for $f\in F$.

Example 3.2.8. For the tree $\mathcal{T}$ in Example 3.2.3, the Horn matrix H of $\mathcal{M}_\mathcal{T}$ is given in Figure 3.2.1. The vector λ of the Horn pair (H,λ) is the vector of ones $(1,\dots,1)\in\mathbb{R}^{16}$. The rows of H are indexed by the florets and labels

$$(s_0,s_1,f_1,s_2,s_3,f_2,s_4,s_5,f_3,s_6,s_7,f_4,s_8,s_9,f_5,s_{10},s_{11},f_6,s_{12},s_{13},f_7).$$

Note that (H,λ) is not minimal. Following the recipe in Lemma 3.1.3, we can delete the rows s_0,s_1,f_2,f_3 of the matrix H by summing the pairs (s_0,f_2) and (s_1,f_3) and deleting zero rows. The result is the minimal Horn pair (H',λ'), where $\lambda'=(-1,\dots,-1)$.

Two staged trees $\mathcal{T}$ and $\mathcal{T}'$ are called *statistically equivalent* in [40] if there exists a bijection between the sets of root-to-leaf paths of $\mathcal{T}$ and $\mathcal{T}'$ such that, after applying this bijection, $\mathcal{M}_\mathcal{T} = \mathcal{M}_{\mathcal{T}'}$ in the open simplex Δ_n. A staged tree model may have different but statistically equivalent tree representations. In [40, Theorem 1], it is shown that statistical equivalence of staged trees can be determined by a sequence of operations on the trees, named *swap* and *resize*. One of the advantages of describing a staged tree model via its Horn pair is that it gives a new criterion to decide whether two staged trees are statistically equivalent. This is simpler to implement than the criterion given in [40].

Corollary 3.2.9. *Two staged trees are statistically equivalent if and only if their associated Horn pairs reduce to the same minimal Horn pair.*

One natural operation on a staged tree $\mathcal{T}$ is identifying two florets of the same size. This gives a new staged tree $\mathcal{T}'$ whose Horn matrix is easy to get from that of $\mathcal{T}$.

Corollary 3.2.10. *Let $\mathcal{T}'$ be a staged tree arising from $\mathcal{T}$ by identifying two florets f and f', say by the bijection $(-)' \colon f \to f'$. The Horn matrix H' of $\mathcal{M}_{\mathcal{T}'}$ arises from the Horn matrix H of $\mathcal{M}_\mathcal{T}$ by replacing the blocks H_f and $H_{f'}$ in H by the block H'_f defined by*

$$\begin{aligned} h'_{ij} &= h_{ij} + h_{i'j} \quad \text{for } s_i \in f, \\ h'_{fj} &= h_{fj} + h_{f'j}. \end{aligned}$$

Proof. This follows from the definition of the Horn matrices for $\mathcal{M}_\mathcal{T}$ and $\mathcal{M}_{\mathcal{T}'}$. □

Example 3.2.11. Let $\mathcal{T}'$ be the tree obtained from Example 3.2.3 by identifying florets f_4 and f_5 in $\mathcal{T}$. Then $\mathcal{M}_{\mathcal{T}'}$ is the independence model of two random variables with four states.

Now we turn to part (3) of Theorem 3.1.1. We describe the triple $(A, \Delta, \mathbf{m})$ for a staged tree model $\mathcal{M}_\mathcal{T}$. The pair (H, λ) was given in Corollary 3.2.6. Let A be any matrix whose rows span the left kernel of H, set $m = |I| + |F|$, and write s for the m-tuple of parameters $(s_i, s_f)_{i \in I, f \in F}$. From the Horn matrix in Corollary 3.2.6 we see that

$$\Delta = \mathbf{m} \cdot \left(1 - \sum_j (-1)^{\varepsilon_j} \prod_i \left(\frac{s_i}{s_f} \right)^{\mu_{ij}} \right),$$

where f depends on i, $\mathbf{m} = \operatorname{lcm}(\prod_i s_f^{\mu_{ij}} : f \in F)$ and $\varepsilon_j = \sum_i \mu_{ij}$. The sign vector σ for the triple $(A, \Delta, \mathbf{m})$ is given by $\sigma_i = +1$ for $i \in I$ and $\sigma_f = -1$ for $f \in F$. Then $Y^*_{A,\sigma}$ gets mapped to $\mathcal{M}_\mathcal{T}$ via $\phi_{(\Delta,\mathbf{m})}$. Moreover, the map $\phi_\mathcal{T}$ from Definition 3.2.2 factors through $\phi_{(\Delta,\mathbf{m})}$. Indeed, if we define $\iota : \Theta \to Y^*_{A,\sigma}$ by $(s_i)_{i \in I} \mapsto (s_i, -1)_{i \in I, f \in F}$, then $\phi_\mathcal{T} = \phi_{(\Delta,\mathbf{m})} \circ \iota$. The following derivation is an extension of that in [47, Expl. 3.13].

Example 3.2.12. Let $\mathcal{M}_\mathcal{T}$ be the 4-chain model in Example 3.2.3. Its discriminant is

$$\begin{aligned}\Delta = f_1f_2f_3f_4f_5f_6f_7 - s_0s_2s_6s_{10}f_3f_5f_7 - s_0s_2s_6s_{11}f_3f_5f_7 - s_0s_2s_7s_{12}f_3f_5f_6 - s_0s_2s_7s_{13}f_3f_5f_6 \\ - s_0s_3s_8s_{10}f_3f_4f_7 - s_0s_3s_8s_{11}f_3f_4f_7 - s_0s_3s_9s_{12}f_3f_4f_6 - s_0s_3s_9s_{13}f_3f_4f_6 \\ - s_1s_4s_6s_{10}f_2f_5f_7 - s_1s_4s_6s_{11}f_2f_5f_7 - s_1s_4s_7s_{12}f_2f_5f_6 - s_1s_4s_7s_{13}f_2f_5f_6 \\ - s_1s_5s_8s_{10}f_2f_4f_7 - s_1s_5s_8s_{11}f_2f_4f_7 - s_1s_5s_9s_{12}f_2f_4f_6 - s_1s_5s_9s_{13}f_2f_4f_6.\end{aligned}$$

Our notation for the parameters matches the row labels of the Horn matrix H in Figure 3.2.1. This polynomial of degree 7 is irreducible, so it equals the A-discriminant: $\Delta = \Delta_A$. The underlying matrix A has format 13×21, and we represent it by its associated toric ideal

$$\begin{aligned}I_A = \langle &s_{10} - s_{11}, s_1s_5f_2 - s_0s_3f_3, s_1s_4f_2 - s_0s_2f_3, s_5s_9f_4 - s_4s_7f_5, s_3s_9f_4 - s_2s_7f_5, \\ &s_{12} - s_{13}, s_5s_8f_4 - s_4s_6f_5, s_3s_8f_4 - s_2s_6f_5, s_9s_{13}f_6 - s_8s_{11}f_7, s_7s_{13}f_6 - s_6s_{11}f_7, \\ &s_0s_2s_6s_{11} - f_1f_2f_4f_6, s_0s_2s_7s_{13} - f_1f_2f_4f_7, s_0s_3s_8s_{11} - f_1f_2f_5f_6, s_0s_3s_9s_{13} - f_1f_2f_5f_7, \\ &s_1s_4s_6s_{11} - f_1f_3f_4f_6, s_1s_4s_7s_{13} - f_1f_3f_4f_7, s_1s_5s_9s_{13} - f_1f_3f_5f_7, s_1s_5s_8s_{11} - f_1f_3f_5f_6\rangle.\end{aligned}$$

The toric variety $Y_A = V(I_A)$ has dimension 12 and degree 141. It lives in a linear space of codimension 2 in $\mathbb{P}^{20}$, where it is defined by eight cubics and eight quartics. The dual variety $Y_A^* = V(\Delta_A)$ is the above hypersurface of degree seven. We have $\mathbf{m} = f_1f_2f_3f_4f_5f_6f_7$, and σ is the vector in $\{-1,+1\}^{21}$ that has entry $+1$ at the indices corresponding to the s_i and entry -1 at the indices corresponding to the f_i.

It would be interesting to study the combinatorics of discriminantal triples for staged tree models. Our computations suggest that, for many such models, the polynomial Δ is irreducible and equals the A-discriminant Δ_A of the underlying configuration A. However, this is not true for all staged trees, as seen in Equation (3.1.2) of Example 3.1.2. We close this section with a familiar class of models with rational MLE whose associated polynomials Δ factor.

Example 3.2.13. The *multinomial distribution* encodes the experiment of rolling a k-sided die m times and recording the number of times we observed the j-th side, for $j = 1, \ldots, k$. The associated model $\mathcal{M}$ is the independence model for m identically distributed random variables on k states. We have $n+1 = \binom{k+m-1}{m}$. The Horn matrix H is the $(k+1) \times (n+1)$ matrix whose columns are the vectors $(-m, i_1, i_2, \ldots, i_k)^T$ where $i_1, i_2, \ldots, i_k$ are nonnegative integers whose sum equals m. Here, $A = (1\ 1\ \cdots\ 1)$, so the A-discriminant equals $\Delta_A = x_0 + x_1 + \cdots + x_k$. The following polynomial is a multiple of Δ_A:

$$\Delta \ = \ (-x_0)^m - (x_1 + x_2 + \cdots + x_k)^m.$$

This Δ, with its marked term $\mathbf{m} = (-x_0)^m$, encodes the MLE for the model $\mathcal{M}$:

$$\hat{p}_{(i_1,\ldots,i_k)} \ = \ \prod_{j=1}^{k} \left(\frac{\sum_{|I|=m} u_I \cdot I_j}{m \sum_{|I|=m} u_I} \right)^{i_j}$$

Here, I ranges over all vectors in $\mathbb{N}^k$ that sum to m, and I_j denotes the j-th entry of I.

3.3. Proof of Theorem 3.1.1

In this section we prove Theorem 3.1.1. For a pair (H, λ) consisting of a Horn matrix H and a coefficient vector λ, let φ be the rational map defined in (3.1.4). We use φ and $\varphi_{(H,\lambda)}$ interchangeably in this section, as well as ϕ and $\phi_{(\Delta,\mathbf{m})}$. Recall that its j-th coordinate is

$$\varphi_j(v) \;=\; \lambda_j \prod_{i=1}^{m} \Big(\sum_{k=0}^{n} h_{ik} v_k \Big)^{h_{ij}}. \tag{3.3.1}$$

For a fixed data vector $u \in \mathbb{N}^{n+1}$, we define the likelihood function for the image of φ:

$$L_u \,:\, \mathbb{R}^{n+1} \to \mathbb{R}\,, \;\; v \mapsto \prod_{j=0}^{n} \varphi_j(v)^{u_j}. \tag{3.3.2}$$

Lemma 3.3.1. *Let $H = (h_{ij})$ be a Horn matrix, λ a vector satisfying (3.1.3) and $u \in \mathbb{N}^{n+1}$. Then u is a critical point of its own likelihood function L_u. Furthermore, if u' is another critical point of L_u, then $\varphi(u) = \varphi(u')$.*

Proof. We compute the partial derivatives of L_u. For $\ell = 0, \ldots, n$ we find

$$\begin{aligned}
\frac{\partial}{\partial v_\ell} L_u(v) \;&=\; \sum_{j=0}^{n} u_j \, \frac{L_u(v)}{\varphi_j(v)} \, \frac{\partial}{\partial v_\ell} \varphi_j(v) \\
&=\; \sum_{j=0}^{n} u_j \, \frac{L_u(v)}{\varphi_j(v)} \, \sum_{i=1}^{m} h_{ij} \, \frac{\varphi_j(v)}{\sum_{k=0}^{n} h_{ik} v_k} \, h_{i\ell} \\
&=\; L_u(v) \sum_{i=1}^{m} \sum_{j=0}^{n} \frac{u_j \, h_{ij} \, h_{i\ell}}{\sum_{k=0}^{n} h_{ik} v_k} \quad = \quad L_u(v) \sum_{i=1}^{m} \frac{h_{i\ell} \, \sum_{j=0}^{n} h_{ij} u_j}{\sum_{k=0}^{n} h_{ik} v_k}.
\end{aligned}$$

For $v = u$, this evaluates to zero, since the sums in the fraction cancel and the ℓ-th column of H sums to zero. This shows that u is a critical point.

Next, let u' be another critical point of L_u. Using terminology from [46, Theorem 1], this means that $\varphi(u')$ is a critical point of the likelihood function $L(p, u)$ of the model $\mathcal{M}$ defined as the image of φ. The same holds for $\varphi(u)$. By the implication (ii) to (i) in [46, Thm. 1], the model $\mathcal{M}$ has ML degree one. This implies $\varphi(u) = \varphi(u')$. □

We use [46] to explain the relation between models with rational MLE and Horn pairs.

Proof of Theorem 3.1.1, Equivalence of (1) and (2). Let $\mathcal{M}$ be a model with rational MLE Φ. The Zariski closure of $\mathcal{M}$ is a variety whose likelihood function has a unique critical point. By [46, Thm. 1], there is a Horn matrix H and a coefficient vector λ such that $\varphi_{(H,\lambda)} = \Phi$. Now, the required sum-to-one and positivity conditions for $\varphi_{(H,\lambda)}$ are satisfied because they are satisfied by the MLE Φ. Indeed, the MLE of any discrete

statistical model maps positive vectors u in $\mathbb{R}^{n+1}_{>0}$ into the simplex Δ_n. Conversely, we claim that every Horn pair (H, λ) specifies a nonempty model $\mathcal{M}$ with rational MLE. Indeed, define $\mathcal{M}$ to be the image of $\varphi_{(H,\lambda)}$. By the defining properties of the Horn pair, we have $\mathcal{M} \subset \Delta_n$. Lemma 3.3.1 shows that $\varphi_{(H,\lambda)}$ is the MLE of $\mathcal{M}$. □

Next, we relate Horn pairs to discriminantal triples.

Proof of Theorem 3.1.1, Equivalence of (2) and (3). We already exhibited a bijection between pairs (H, λ) and pairs $(\Delta, \mathbf{m})$ given by Equation 3.1.12. The matrix A is the left kernel of H and forms the triple $(A, \Delta, \mathbf{m})$. It is a matrix of size $r \times m$ of rank r. When H is a Horn matrix, A contains $(1, \ldots, 1)$ in its row span. This implies that the polynomial Δ is homogeneous, which in turn implies that it is A-homogeneous by $AH = 0$.

Next, we show that the pair (H, λ) being friendly corresponds to the polynomial Δ vanishing on Y_A^*. This is part of the desired equivalence.

Claim. *The pair (H, λ) is friendly if and only if the A-homogeneous polynomial Δ vanishes on the dual toric variety Y_A^*.*

Proof of Claim. Let (H, λ) be friendly and A be a matrix as above. The Laurent polynomial $q := \Delta/\mathbf{m}$ is a rational function on $\mathbb{P}^{m-1}$ that vanishes on the dual toric variety Y_A^*. To see this, consider the exponentiation map $\varphi_2 : \mathbb{P}^{m-1} \to \mathbb{R}^{n+1}$, $x \mapsto \lambda * x^H$, where $*$ is the entrywise product and $x^H := (x^{h_0}, \ldots, x^{h_n})$. Let $f = 1 - (p_0 + \cdots + p_n)$. We have $q = f \circ \varphi_2$. By [46, Thms. 1 and 2], the function φ_2 maps an open dense subset of Y_A^* dominantly to the closure $\mathcal{M}$ of the image of $\varphi_{(H,\lambda)}$. Since $f = 0$ on $\mathcal{M}$, we have $f \circ \varphi_2 = 0$ on an open dense subset of Y_A^*, hence $q = 0$ on Y_A^*, so $\Delta = 0$ there as well.

Conversely, let Δ vanish on Y_A^*. We claim that $q(x)$ is zero for all $x = Hu$ in the image of the linear map H. We may assume $\mathbf{m}(x) \neq 0$. We only need to show that x is in the dual toric variety Y_A^*, since Δ vanishes on it. So, let $x_i = \sum_{j=0}^n h_{ij} u_j$ for $i = 1, \ldots m$. We claim that $t = (1, \ldots, 1)$ is a singular point of the hypersurface

$$\gamma_A^{-1}(H_x \cap Y_A) \;=\; \left\{ t \in \mathbb{C}^r \mid \sum_{i=1}^m x_i t^{a_i} = 0 \right\}.$$

First, the point t lies on that hypersurface since the columns of H sum to zero:

$$\sum_{i=1}^m x_i \;=\; \sum_{i=1}^m \sum_{j=0}^n h_{ij} u_j \;=\; \sum_{j=0}^n u_j \sum_{i=1}^m h_{ij} \;=\; 0.$$

For $s = 1, \ldots, r$ we have $\frac{\partial}{\partial t_s} t^{a_i} = a_{si} t^{a_i - e_s}$, with e_s the standard basis vector of $\mathbb{Z}^r$, and

$$\frac{\partial}{\partial t_s} \sum_{i=1}^m x_i t^{a_i} \;=\; \sum_{i=1}^m \sum_{j=0}^n h_{ij} u_j a_{si} t^{a_i - e_s} \;=\; \sum_{j=0}^n u_j \sum_{i=1}^m a_{si} h_{ij} t^{a_i - e_s}.$$

This is zero at $t = (1, \ldots, 1)$ because $AH = 0$. ■

We now prove the rest of the equivalence. Let (H, λ) be a Horn pair, let φ be its Horn map and let ϕ be the associated monomial map. Let $\mathcal{M}$ be the statistical model with MLE φ, so $\mathcal{M} = \varphi(\mathbb{R}^{n+1}_{>0})$. We have $\varphi = \phi \circ H$. By Proposition 3.3.3, there exists a unique sign vector σ such that $\operatorname{im} H|_{\mathbb{R}^{n+1}_{>0}} \subseteq \mathbb{R}^m_\sigma$. From the proof of the above claim we know that $\operatorname{im} H \subseteq Y_A^*$. Together, we have

$$\mathcal{M} = \varphi(\mathbb{R}^{n+1}_{>0}) = \phi(\operatorname{im} H|_{\mathbb{R}^{n+1}_{>0}}) \subseteq \phi(Y^*_{A,\sigma}).$$

By [46, Theorems 1 and 2] we have $\phi(Y_A^*) \subseteq \mathcal{M}'$, where $\mathcal{M}'$ is the real part of $\overline{\varphi(\mathbb{C}^{n+1})}$. We also have $\phi(Y^*_{A,\sigma}) \subseteq \mathbb{R}^{n+1}_{>0}$ by definition of the orthant. Thus $\phi(Y^*_{A,\sigma}) \subseteq \mathcal{M}' \cap \mathbb{R}^{n+1}_{>0}$. Every element in the latter set is a fixed point of the rational function φ, by a similar argument as in Lemma 3.3.1 for complex space. Hence $\mathcal{M}' \cap \mathbb{R}^{n+1}_{>0} = \mathcal{M}$, so $\phi(Y^*_{A,\sigma}) \subseteq \mathcal{M}$.

Finally, if $(A, \Delta, \mathbf{m})$ is a discriminantal triple then (H, λ) is a Horn pair by definition. This completes the proof of Theorem 3.1.1. □

In the next two propositions, we formulate simple criteria to decide whether the image of the map $\varphi_{(H,\lambda)}$ associated to a Horn matrix H and a coefficient vector λ is a statistical model. These are essential for constructing models with rational MLE in Algorithm 1.

Proposition 3.3.2. *Let (H, λ) be a friendly pair. If there exists a vector $u_0 \in \mathbb{R}^{n+1}$ such that $\varphi(u_0) > 0$, then we have $\varphi(u) > 0$ for all u in $\mathbb{R}^{n+1}_{>0}$ where it is defined.*

Proof. The function φ is homogeneous of degree zero. It suffices to prove that each coordinate of $\varphi(u)$ is a positive real number, for all vectors u with positive integer entries. Indeed, every positive u in $\mathbb{R}^{n+1}$ can be approximated by rational vectors, which can be scaled to be integral. The open subset $U = \varphi^{-1}(\Delta_n)$ of $\mathbb{R}^{n+1}$ contains u_0 by our assumptions. If $U = \mathbb{R}^{n+1}$, then we are done. Else, U has a nonempty boundary ∂U. By continuity, $\partial U \subseteq \varphi^{-1}(\partial \Delta_n)$. The likelihood function L_u for the data vector u vanishes on ∂U.

We claim that L_u has a critical point in U. The closed subset $\overline{U}$ is homogeneous. Seen in projective space $\mathbb{P}^n$, it becomes compact. The likelihood function L_u is well defined on this compact set in $\mathbb{P}^n$, since it is homogeneous of degree zero, and L_u vanishes on the boundary. Hence the restriction $L_u|_U$ is either identically zero or it has a critical point in U. But, since $u_0 \in U$ is a point with $L_u(u_0) \neq 0$, the second statement must be true.

Pick such a critical point u'. Since U is open in $\mathbb{R}^{n+1}$, the point u' is also critical point of L_u. By Lemma 3.3.1 and since $u' \in U$, we have $\varphi(u) = \varphi(u') > 0$. □

Proposition 3.3.3. *Let (H, λ) be a friendly pair, with no zero or collinear rows in H. Then (H, λ) is a Horn pair if and only if for every row r_i of H all nonzero entries of r_i have the same sign σ_i, and the sign vector $\sigma = (\sigma_i)$ satisfies $\lambda_j \sigma^{h_j} > 0$ for all columns j.*

Proof. Let (H, λ) be a Horn pair. Let $\ell_1, \ldots, \ell_k$ be the linear forms corresponding to the rows in H that have both positive and negative entries. Since ℓ_1 has positive and negative coefficients, there exists a positive vector u such that $\ell_1(u) = 0$. Since (H, λ) is minimal, we may choose $u > 0$ such that $\ell_1(u) = 0$ but $\ell_{k'}(u) \neq 0$ for all $k' \neq 1$. The form ℓ_1 appears in the numerator of some coordinate of φ, making this coordinate zero at u. But this contradicts the fact that (H, λ) is a Horn pair. Therefore we cannot have rows with both positive and negative entries. The inequalities $\lambda_j \sigma^{h_j} > 0$ then follow from the definition of a Horn pair by evaluating $\varphi(u)$ for some positive vector u.

Conversely, if the sign vector σ is well-defined, the inequalities $\lambda_j \sigma^{h_j} > 0$ imply $\varphi(u) > 0$ for all positive u. Hence (H, λ) is a Horn pair. □

Every model with rational MLE *arises from* a toric variety Y_A. In some cases, the model *is itself* a toric variety Y_C. It is crucial to distinguish the two matrices A and C. The two toric structures are very different. For instance, every undirected graphical model is toric [27, Prop. 3.3.3]. The toric varieties Y_C among staged tree models $\mathcal{M}_\mathcal{T}$ were classified in [28]. The 4-chain model $\mathcal{M}_\mathcal{T} = Y_C$ *is itself* a toric variety of dimension 7 in $\mathbb{P}^{15}$. But it *arises from* a toric variety Y_A of dimension 12 in $\mathbb{P}^{20}$, seen in Example 3.2.12.

3.4. Geometric modeling

Toric models with rational MLE play an important role in *geometric modeling* [19, 38]. Given a matrix $C \in \mathbb{Z}^{r \times (n+1)}$ and a vector of weights $w \in \mathbb{R}^{n+1}_{>0}$, one considers the *scaled projective toric variety* $Y_{C,w}$ in $\mathbb{RP}^n$. This is defined as the closure of the image of

$$\gamma_{C,w} \,:\, (\mathbb{R}^*)^r \to \mathbb{RP}^n, \;\; (t_1, \ldots, t_r) \mapsto \left(w_0 \prod_{i=1}^r t_i^{c_{i0}}, \, w_1 \prod_{i=1}^r t_i^{c_{i1}t}, \, \ldots \, , \, w_n \prod_{i=1}^r t_i^{c_{in}} \right).$$

The set $\mathcal{M}_{C,w}$ of positive points in $Y_{C,w}$ is a statistical model in Δ_n. There is a natural homeomorphism from the toric model $\mathcal{M}_{C,w}$ onto the polytope of C. This is known in geometry as the *moment map*. For a reference from algebraic statistics, see [27, Proposition 2.1.5]. In geometric modeling the pair (C, w) defines *toric blending functions* [53].

It is desirable for the toric blending functions to have *rational linear precision* [19, 53]. The property is rare and it depends in a subtle way on (C, w). Garcia-Puente and Sottile [38] established the connection to algebraic statistics. They showed that rational linear precision holds for (C, w) if and only if the statistical model $\mathcal{M}_{C,w}$ has rational MLE.

Example 3.4.1. The most classical blending functions with rational linear precision live on the triangle $\{x \in \mathbb{R}^3_{>0} : x_1 + x_2 + x_3 = 1\}$. They are the *Bernstein basis polynomials*

$$\frac{m!}{i!j!(m-i-j)!} x_1^i x_2^j x_3^{m-i-j} \quad \text{for} \quad i, j \geq 0, \, i + j \leq m. \tag{3.4.1}$$

Here C is the $3 \times \binom{m+1}{2}$ matrix whose columns are the vectors $(i, j, m-i-j)$. The weights are $w_{(i,j)} = \frac{m!}{i!j!(m-i-j)!}$. The toric model $\mathcal{M}_{C,w}$ is the multinomial family, where (3.4.1) is the probability of observing i times 1, j times 2 and $m-i-j$ times 3 in m trials. This model has rational MLE, as seen in Example 3.2.13. Again, notice the distinction between the two toric varieties. Here, Y_A is a point in $\mathbb{P}^m$, whereas Y_C is a surface in $\mathbb{P}^{\binom{m}{2}-1}$.

Clarke and Cox [19] raise the problem of characterizing all pairs (C, w) with rational linear precision. This was solved by Duarte and Görgen [28] for pairs arising from staged trees. While the problem remains open in general, the theory in this chapter offers new tools. We may ask for a characterization of discriminantal triples whose models are toric.

3.5. Constructing models with rational MLE

Part (3) in Theorem 3.1.1 allows us to construct models with rational MLE starting from a matrix A that defines a projective toric variety Y_A. To carry out this construction effectively we propose Algorithm 1. In most cases, the dual variety Y_A^* is a hypersurface, and we can compute its defining polynomial Δ_A, the *discriminant* [39]. The polynomial Δ in a discriminantal triple can be any homogeneous multiple of Δ_A, but we just take $\Delta = \Delta_A$. For all terms $\mathbf{m}$ in Δ_A, we check whether $(A, \Delta_A, \mathbf{m})$ is a discriminantal triple.

Lines 1 and 15 of Algorithm 1 are computations with Gröbner bases. Executing Line 15 can be very slow. It may be omitted if we are satisfied with obtaining the parametric description and MLE $\Phi^{(\ell)}$ of the model $\mathcal{M}_\ell$. For the check in Line 14, we rely on Proposition 3.3.2 for correctness. A check based on the criterion in Proposition 3.3.3 is also possible.

Example 3.5.1 ($r = 2, m = 4$)**.** For distinct integers $\alpha, \beta, \gamma > 0$ with $\gcd(\alpha, \beta, \gamma) = 1$ let

$$A_{\alpha,\beta,\gamma} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & \alpha & \beta & \gamma \end{pmatrix}.$$

We ran Algorithm 1 for all 613 such matrices with $0 < \alpha < \beta < \gamma \leq 17$. Line 1 computes the discriminant Δ_A of the univariate polynomial $f(t) = x_1 + x_2 t^\alpha + x_3 t^\beta + x_4 t^\gamma$. The number $n+2$ of terms of these discriminants equals 7927 in total, for an average of 12.93 terms per discriminant. Thus a total of 7927 candidate triples $(A, \Delta_A, \mathbf{m})$ were tested in Lines 12 to 21. Precisely 123 of these were found to be discriminantal triples. This is a fraction of 1.55 %. Hence, only 1.55 % of the resulting complex varieties permitted by [46] are actually statistical models.

Here is a typical model that was discovered. Take $\alpha = 1, \beta = 4, \gamma = 7$. The discriminant

$$\begin{aligned} \Delta_A \;=\; & 729x_2^4x_3^6 - 6912x_1^3x_3^7 - 8748x_2^5x_3^4x_4 + 84672x_1^3x_2x_3^5x_4 + 34992x_2^6x_3^2x_4^2 \\ & -351918x_1^3x_2^2x_3^3x_4^2 - 46656x_2^7x_4^3 + 518616x_1^3x_2^3x_3x_4^3 \underline{\,-\,823543x_1^6x_4^4} \end{aligned}$$

Algorithm 1: From toric varieties to statistical models

Input : An integer matrix A of size $r \times m$ with $(1, \ldots, 1)$ in its row span

Output : An integer n and a collection of statistical models $\mathcal{M}^{(\ell)} = (\Phi^{(\ell)}, I^{(\ell)})$, where $\Phi^{(\ell)} \colon \mathbb{R}^{n+1} \to \mathbb{R}^{n+1}$ is a rational MLE for $\mathcal{M}^{(\ell)}$, and $I^{(\ell)} \subseteq \mathbb{R}[p_0, \ldots, p_n]$ is the defining prime ideal of $\mathcal{M}^{(\ell)}$.

```
Compute the A-discriminant Δ_A ∈ ℤ[x_1, ..., x_m];
n ← #terms(Δ_A) − 2;
models ← {};
for 0 ≤ ℓ ≤ n + 1 do
    m ← terms(Δ_A)_ℓ;
    q ← 1 − Δ_A/m;
    for 0 ≤ i ≤ n do
        λ_i ← coefficients(q)_i;
        h_i ← exponent_vectors(q)_i;
        Φ_i^(ℓ) ← (u ↦ λ_i ∏_{j=1}^m (∑_{k=0}^n h_jk u_k)^{h_ji});
    end
    H ← (h_i)_i;
    Choose any positive vector v in ℝ_{>0}^{n+1};
    if Φ_i^(ℓ)(v) > 0 for i = 0, 1, ..., n then
        Compute the ideal I^(ℓ) of the image of Φ^(ℓ);
        models ← models ∪ {(Φ^(ℓ), I^(ℓ))};
    end
end
return models;
```

has 9 terms, so $n = 7$. The term $\mathbf{m}$ is underlined. The associated model is a curve of degree ten in Δ_7. Its prime ideal $I^{(\ell)}$ is generated by 18 quadrics. Among them are 15 binomials that define a toric surface of degree six: $49p_1p_2 - 48p_0p_3, 3p_0p_4 - p_2^2, \ldots, 361p_3p_7 - 128p_5^2$. Inside that surface, our curve is cut out by three quadrics, here is one of them:

$$26068p_2^2 + 73728p_0p_5 + 703836p_0p_6 + 234612p_2p_6 + 78204p_4p_6 + 612864p_0p_7 + 212268p_2p_7 + 78204p_4p_7 - 8379p_7^2.$$

Example 3.5.2 ($r = 3, m = 6$)**.** For any positive integers $\alpha, \beta, \gamma, \varepsilon$, we consider the matrix

$$A = \begin{pmatrix} 0 & \alpha & \beta & 0 & \gamma & \varepsilon \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}.$$

The discriminant Δ_A is the *resultant* of two trinomials $x_1 + x_2t^\alpha + x_3t^\beta$ and $x_4 + x_5t^\gamma + x_6t^\varepsilon$. We ran Algorithm 1 for all 138 such matrices with

$$0 < \alpha < \beta \leq 17, \quad 0 < \gamma < \varepsilon \leq 17, \quad \gcd(\alpha, \beta) = \gcd(\gamma, \varepsilon) = 1.$$

The number $n+2$ of terms of these discriminants equals 2665 in total, for an average of 19.31 terms per discriminant. Thus a total of 2665 candidate triples $(A, \Delta_A, \mathbf{m})$ were tested in Line 13. Precisely 93 of these are discriminantal triples. This is only 3.49 %.

We now shift gears by looking at polynomials Δ that are multiples of the A-discriminant.

Example 3.5.3 ($r = 1, m = 4$)**.** We saw in Examples 3.1.2 and 3.2.13 that interesting models arise from the matrix $A = (1\ 1\ \cdots\ 1)$ whose toric variety is just a point. Any homogeneous multiple Δ of the linear form $\Delta_A = x_1 + x_2 + \cdots + x_m$ can be the input in Line 1 of Algorithm 1. Here, taking $\Delta = \Delta_A$ results in the model given by the full simplex Δ_{m-2}.

Let $m = 4$ and abbreviate $x^a = x_1^{a_1}x_2^{a_2}x_3^{a_3}x_4^{a_4}$ and $|a| = a_1+a_2+a_3+a_4$ for $a \in \mathbb{N}^4$. We conducted experiments with two families of multiples. The first uses binomial multipliers:

$$\Delta \,=\, (x^a + x^b)\Delta_A \quad \text{or} \quad \Delta = (x^a - x^b)\Delta_A,$$

where $|a| = |b| \in \{1, 2, \ldots, 8\}$ and $\gcd(x^a, x^b) = 1$. This gives 1028 polynomials Δ. The numbers of polynomials of degree $2, 3, 4, 5, 6, 7, 8, 9$ is 6, 21, 46, 81, 126, 181, 246, 321. For the second family we use the trinomial multiples

$$\Delta = (x^a + x^b + x^c)\Delta_A \quad \text{or} \quad \Delta = (x^a + x^b - x^c)\Delta_A,$$

where $|a|=|b|=|c| \in \{1, 2, 3\}$ and $\gcd(x^a, x^b, x^c) = 1$. Each list contains 4 quadrics, 104 cubics and 684 quartics. We report our findings in a table:

Family	Pairs $(\Delta, \mathbf{m})$	Horn pairs	Percentage
$(x^a - x^b)\Delta_A$	8212	12	0.15%
$(x^a + x^b)\Delta_A$	8218	0	0%
$(x^a + x^b - x^c)\Delta_A$	8678	8	0.01%
$(x^a + x^b + x^c)\Delta_A$	8968	0	0%

All 12 Horn pairs in the first family represent the same model, up to permuting coordinates. All are coming from the six quadrics of the family. The model is the surface in Δ_4 defined by the 2×2 minors of the matrix

$$\begin{pmatrix} p_0 & p_1 & p_2 \\ p_0{+}p_1{+}p_2 & p_3 & p_4 \end{pmatrix}.$$

This is a staged tree model similar to Example 2, but now with three choices at each blue node instead of two. The eight Horn pairs in the third family represent two distinct models. Four of the eight Horn pairs represent a surface in Δ_5 and the rest represent a surface in Δ_6.

Our construction of models with rational MLE starts with families where r and m are fixed. However, as the entries of the matrix A go up, the number $n+1$ of states increases. This suggests the possibility of listing all models for fixed values of n. Is this list finite?

Problem. Suppose that n is fixed. Are there only finitely many models with rational MLE in the simplex Δ_n? Can we find absolute bounds, depending only on n, for the dimension, degree and number of ideal generators of the associated varieties in $\mathbb{P}^n$?

Algorithm 1 is a tool for studying these questions experimentally. At present, we do not have any clear answers, even for $n = 3$, where the models are curves in a triangle.

4. Algebraic manifolds

Let $\mathcal{M}$ be an algebraic manifold given as the solution set of a system of polynomial equations and inequalities. In this chapter we develop a numerical method to:

(1) approximate the Lebesgue integral $\int_{\mathcal{M}} f(x)\,\mathrm{d}x$ of a given function f on $\mathcal{M}$, and

(2) sample from a probability distribution with a given density on $\mathcal{M}$.

These problems are closely related in theory. In applications however, they may occur separately. For instance, in Section 4.2 we encounter an example from computational physics that involves only (1). On the other hand, the subsequent example from topological data analysis concerns (2).

When $\mathcal{M}$ is given as a set of solutions to implicit equations, the standard techniques to solve (2) involve moving randomly from one sample point on $\mathcal{M}$ to the next nearby, and fall under the umbrella term of *Markov Chain Monte Carlo (MCMC)*. The literature on MCMC methods is vast. The interested reader is referred to the recent articles [60, 88].

This chapter presents a new method that solves (1) and (2) when the functions F_i are *polynomial*. In simple terms, the method can be described as follows. First, we choose a random linear subspace of complementary dimension and calculate its intersection with $\mathcal{M}$. Since the implicit equations are polynomial, the intersection can be efficiently determined using numerical polynomial equation solvers. For example, we could pick the package `HomotopyContinuation.jl` [12] from Example 2.1.6 as our solver. The number of intersection points is finite and bounded by the *degree* of the algebraic manifold. Next, if we want to solve (1) we evaluate a modified function $\overline{f}$ at each intersection point, sum its values, and repeat the process to approximate the desired integral. Else if we want to solve (2), after a rejection step we pick one of the intersection points at random to be our sample point. We then repeat the process to obtain more samples of the desired density.

Compared to MCMC sampling, our method has two main advantages. First, we have the option to generate points that are independent of each other. Second, the method is global in the sense that it also works when the manifold has multiple distinct connected components, and does not require picking a starting point $x_0 \in \mathcal{M}$.

The main theoretical result supporting the method is Theorem 4.1.1, which we state in Section 4.1. It is in line with a series of classical results commonly known as *Crofton's formulae* or *kinematic formulae* [75] that relate the volume of a manifold to the expected

number of its intersection points with random linear spaces. Section 4.2 presents applications of the method to examples in topological data analysis and statistical physics. Section 4.3 introduces the tools that are needed for the proof of the main theorem. The proof is carried out in Section 4.4 where we also prove a rate of convergence bound for our method. We prove a variant of Theorem 4.1.1 for projective algebraic manifolds in Section 4.5. In Section 4.6 we review other methods that make use of a kinematic formula for sampling. We also discuss the advantages and limitations of our method and possible future work.

Before we continue, let us fix some notation. Throughout this chapter we fix an n-dimensional algebraic manifold $\mathcal{M} \subseteq \mathbb{R}^N$. The Euclidean inner product on $\mathbb{R}^N$ is defined by $\langle x, y\rangle := x^T y$ and the associated norm is $\|x\| := \sqrt{\langle x, x\rangle}$. The unit sphere in $\mathbb{R}^N$ is $\mathbb{S}^{N-1} := \{x \in \mathbb{R}^N : \|x\| = 1\}$. For a function $f : \mathcal{M} \to \mathcal{N}$ between manifolds we denote by $\mathrm{D}_x f$ the derivative of f at $x \in \mathcal{M}$. The tangent space of $\mathcal{M}$ at x is denoted $\mathrm{T}_x\mathcal{M}$ and the normal space is $\mathrm{N}_x\mathcal{M}$.

4.1. A method for sampling

To state our result, we fix a measurable function $f\colon \mathcal{M} \to \mathbb{R}_{\geq 0}$ with finite integral over $\mathcal{M}$. We define the auxiliary function $\overline{f}\colon \mathbb{R}^{n\times N} \times \mathbb{R}^n \to \mathbb{R}$ as follows:

$$\overline{f}(A, b) \coloneqq \sum_{x\in\mathcal{M}: Ax=b} \frac{f(x)}{\alpha(x)} \quad \text{where} \quad \alpha(x) \coloneqq \frac{\sqrt{1 + \langle x, \Pi_{\mathrm{N}_x\mathcal{M}}\, x\rangle}}{1 + \|x\|^2} \, \frac{\Gamma\left(\frac{n+1}{2}\right)}{\sqrt{\pi}^{n+1}}$$

and where $\Pi_{\mathrm{N}_x\mathcal{M}} : \mathbb{R}^N \to \mathrm{N}_x\mathcal{M}$ denotes the orthogonal projection onto the normal space of $\mathcal{M}$ at $x \in \mathcal{M}$. Note that this projection can be computed from the implicit equations for $\mathcal{M}$, because $\mathrm{N}_x\mathcal{M}$ is the row-span of the Jacobian matrix $J(x) = [\frac{\partial F_i}{\partial x_j}(x)]_{1\leq i\leq r, 1\leq j\leq N}$. Therefore, if $Q \in \mathbb{R}^{N\times r}$ is the Q-factor from the QR-decomposition of $J(x)^T$, then we have $\mathrm{N}_x\mathcal{M} = QQ^T$. This means that we can easily compute $\overline{f}(A, b)$ from $\mathcal{M} \cap \mathcal{L}_{A,b}$.

The operator $f \mapsto \overline{f}$ allow us to state the following main result, making precise the method described at the beginning of the chapter.

Theorem 4.1.1. *Let $\varphi(A, b)$ be the probability density for which the entries of $A \in \mathbb{R}^{n\times N}$ and $b \in \mathbb{R}^n$ are i.i.d. standard normal. In the notation introduced above:*

(1) The integral of f over $\mathcal{M}$ is the expected value of $\overline{f}$:

$$\int_{\mathcal{M}} f(x)\,\mathrm{d}x = \mathbb{E}_{(A,b)\sim\varphi}\overline{f}(A, b).$$

(2) Assume that $f : \mathcal{M} \to \mathbb{R}$ is nonnegative and that $\int_{\mathcal{M}} f(x)\,\mathrm{d}x$ is positive and finite. Let $X \in \mathcal{M}$ be the random variable obtained by choosing a pair $(A, b) \in \mathbb{R}^{n\times N} \times \mathbb{R}^n$ with probability

$$\psi(A, b) := \frac{\varphi(A, b)\,\overline{f}(A, b)}{\mathbb{E}_\varphi(\overline{f})}$$

and choosing one of the finitely many points X of the intersection $\mathcal{M} \cap \mathcal{L}_{A,b}$ with probability $f(x)\alpha(x)^{-1}\overline{f}(A,b)^{-1}$. Then X is distributed according to the scaled density $f(x)/(\int_{\mathcal{M}} f(x)\,\mathrm{d}x)$ associated to $f(x)$.

Using the formula for $\alpha(x)$ we can already evaluate $\overline{f}(A,b)$ for an integrable function f. Thus we can approximate the integral of f by computing the empirical mean

$$\mathrm{E}(f,k) = \tfrac{1}{k}(\overline{f}_1(A,b) + \cdots + \overline{f}_k(A,b))$$

of a sample drawn from $(A,b) \sim \varphi$. The next lemma yields a bound for the rate of convergence of this approach. It is an application of Chebyshev's inequality and proved in Section 4.4.

Lemma 4.1.2. *Assume that $|f(x)|$ and $\|x\|$ are bounded on $\mathcal{M}$. The variance $\sigma^2(\overline{f})$ of $\overline{f}(A,b)$ is finite and for $\varepsilon > 0$ we have* $\mathrm{Prob}\{|\mathrm{E}(f,k) - \int_{\mathcal{M}} f(x)\,\mathrm{d}x| \geq \varepsilon\} \leq \frac{\sigma^2(\overline{f})}{\varepsilon^2 k}$.

In (4.4.1) we will see a deterministic bound for $\sigma^2(\overline{f})$, which involves the degree of the ambient variety of $\mathcal{M}$ and upper bounds for $\|x\|$ and $|f(x)|$ on $\mathcal{M}$. In our experiments we also use the empirical variance $s^2(\overline{f})$ of a sample for estimating $\sigma^2(\overline{f})$.

For sampling $(A,b) \sim \psi$ in the second part of Theorem 4.1.1 we could use MCMC sampling. Note that this would employ MCMC sampling for the flat space $\mathbb{R}^{n\times N} \times \mathbb{R}^n$, which is easier than MCMC for nonlinear spaces like $\mathcal{M}$. Nevertheless, in this chapter we use the simplest method for sampling ψ, namely *rejection sampling*. This is used in the experiment section and explained in Section 4.3.

4.2. Experiments

In this section we apply our main result to examples. All experiments have been performed on macOS 10.14.2 on a computer with Intel Core i5 2,3GHz (two cores) and 8 GB RAM memory. For computing the intersections with linear spaces, we use the numerical polynomial equation solver `HomotopyContinuation.jl` [12]. For plotting we use `Matplotlib` [48]. For sampling from the distribution $\psi(A,b)$ we use rejection sampling as described in Section 4.3.

As a first simple example, consider the plane curve $\mathcal{M}$ given by the equation

$$x^4 + y^4 - 3x^2 - xy^2 - y + 1 = 0. \tag{4.2.1}$$

We have $\mathrm{vol}(\mathcal{M}) = \mathbb{E}_\varphi(\overline{1})$. We can therefore estimate the length of the curve $\mathcal{M}$ by taking a sample of i.i.d. pairs (A,b) and computing the empirical mean $\mathrm{E}(1,k)$ of $\overline{1}$ of the sample. A sample of $k = 10^5$ yields $\mathrm{E}(1,k) = 11.2$. In Lemma 4.1.2 we take $\varepsilon = 0.1$ and the variance of the sample s^2, and get an upper bound of $\frac{s^2}{\varepsilon^2 k} = 0.008$. Therefore, we expect that 11.2 is a good approximation of the true length. We can also take the deterministic

upper bound from (4.4.1) for the variance σ^2 of $\overline{1}$. Here, we take $\sup_{x\in\mathcal{M}} \|x\| = \sqrt{8}$. To get an estimate with accuracy at least $\varepsilon = 0.1$ with probability at least 0.9 we need a sample of size $k \geq \frac{\sigma^2}{\varepsilon^2 \cdot 0.9} \geq 1421300$. Taking such a sample size we get an estimated length of ≈ 11.217.

Next, we use the second part of Theorem 4.1.1 to generate random samples on $\mathcal{M}$. The left picture of Figure 4.2.1 shows a sample of 200 points drawn uniformly from the curve. The right picture shows 200 points drawn from the scaled density associated to $f(x, y) = e^{2y}$. As can be seen from the pictures the points drawn from the second distribution concentrate in the upper half of $\mathcal{M}$, whereas points from the first distribution spread equally around the curve. This experiment also shows how our method generates global samples. The curve has more than one connected component, which is not an obstacle for our method.

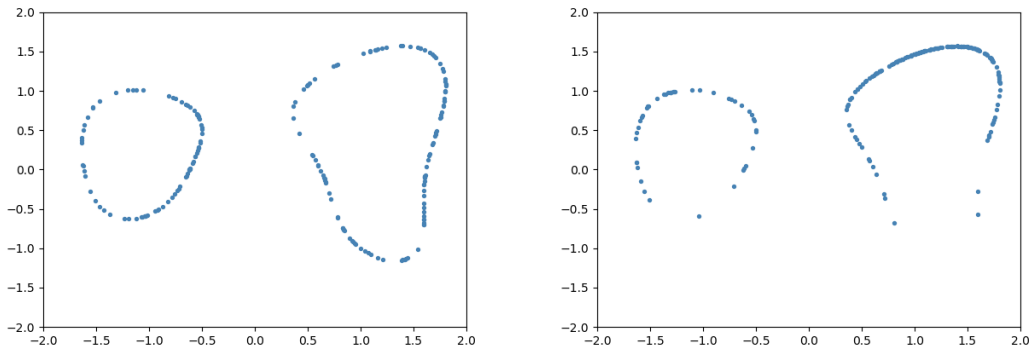

Figure 4.2.1: Left picture: a sample of 200 points from the uniform distribution on the curve (4.2.1). Right picture: a sample of 200 points from the same distribution scaled by e^{2y}.

Our method is particularly appealing for hypersurfaces like (4.2.1) because intersecting a hypersurface with a linear space of dimension one reduces to solving a single univariate polynomial equation. This can be done very efficiently, for instance using the algorithm from [79], and so for hypersurfaces we can easily generate large sample sets.

The pictures suggest to use sampling for visualization. For instance, we can visualize a semialgebraic piece of the complex and real part of the *Trott curve* T, defined by the equation

$$144(x_1^4 + x_2^4) - 225(x_1^2 + x_2^2) + 350x_1^2x_2^2 + 81 = 0. \tag{4.2.2}$$

The associated complex variety in $\mathbb{C}^2$ can be seen as a real variety $T_{\mathbb{C}}$ in $\mathbb{R}^4$. We sample from the real Trott curve T and the complex Trott curve $T_{\mathbb{C}}$ intersected with the box $-1.5 < \text{Real}(x_1), \text{Imag}(x_1), \text{Real}(x_2), \text{Imag}(x_2) < 1.5$. Then, we take a random projection $\mathbb{R}^4 \to \mathbb{R}^3$ to obtain a sample in $\mathbb{R}^3$ (the projected sample is *not* uniform on the projected semialgebraic variety). The outcome of this experiment is shown in Figure 4.2.2.

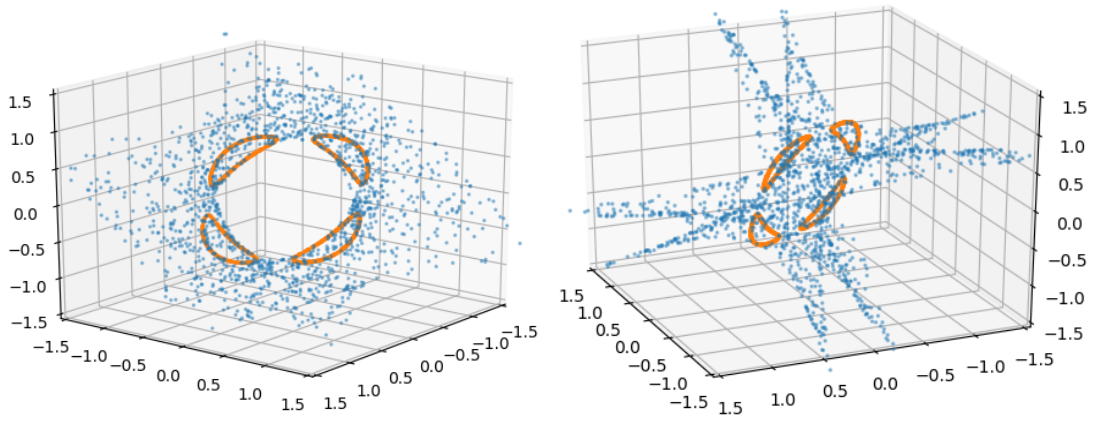

Figure 4.2.2: The blue points are a sample of 1569 points from the complex Trott curve (4.2.2) seen as a variety in $\mathbb{R}^4$ projected to $\mathbb{R}^3$. The orange points are a sample of 1259 points from the real part of the Trott curve. The two pictures show two different projections to $\mathbb{R}^3$.

Application to statistical physics

Next, we apply Theorem 4.1.1 to study a physical system of N particles $q = (q_1, \ldots, q_N)$ interpreted as elements of a manifold $\mathcal{M} \subseteq (\mathbb{R}^3)^N$ that models the spatial constraints of the q_i. In our example we have $N = 6$ and the q_i are the spatial positions of carbon atoms in a *cyclohexane molecule*. The constraints of this molecule are algebraic and define the manifold

$$\mathcal{M} = \{q = (q_1, \ldots, q_6) \in (\mathbb{R}^3)^6 \mid \|q_1 - q_2\|^2 = \cdots = \|q_5 - q_6\|^2 = \|q_6 - q_1\|^2 = c^2\}, \quad (4.2.3)$$

where c is the *bond length* between two neighboring atoms (the vectors $q_i - q_{i+1}$ are called *bonds*). In our example we take $c^2 = 5$ (unitless). Due to rotational and translational invariance of the equations we define q_1 to be the origin, $q_6 = (c, 0, 0)$ and q_5 to be rotated in such a way that its last entry is equal to zero. We thus have 11 variables.

Lelievre et. al. [59] write "In the framework of statistical physics, macroscopic quantities of interest are written as averages over [...] probability measures on all the admissible microscopic configurations." As the probability measure we take the canonical ensemble [59]. If $E(q)$ denotes the total energy of a configuation q, the density in the canonical ensemble is proportional to $f(q) = e^{-E(q)}$. That is, a configuration is most likely to appear when its energy is minimal. We model the energy of a molecule using an interaction potential, namely the *Lennard Jones potential* $V(r) = \frac{1}{4}\left(\frac{c}{r}\right)^{12} - \frac{1}{2}\left(\frac{c}{r}\right)^{6}$ see e.g. [59, Eq. (1.5)]. Then, the energy function of a system is

$$E(q) = \sum_{1 \leq i < j \leq N} V(\|q_i - q_j\|).$$

In this example we consider the average angle between neighboring bonds $q_{i-1} - q_i$ and $q_{i+1} - q_i$ given by the formula

$$\theta(q) = \frac{\angle(q_6 - q_1, q_2 - q_1) + \cdots + \angle(q_5 - q_6, q_1 - q_6)}{6},$$

where $\angle(b_1, b_2) := \arccos \frac{\langle b_1, b_2 \rangle}{\|b_1\| \|b_2\|}$. We compute the macroscopic state of $\theta(q)$ by determining its distribution $\text{Prob}\{\theta(q) = \theta_0\} = \frac{1}{H} \int_{\theta(q)=\theta_0} f(q)\text{d}q$, where $H = \int_V f(q)\text{d}q$ is the normalizing constant. For comparing the probabilities of different values for θ it suffices to compute

$$\rho(\theta_0) = \int_{\theta(q)=\theta_0} f(q)\text{d}q.$$

We approximate this integral as $\rho(\theta_0) \approx \frac{\mu_1(\theta_0)}{\mu_2(\theta_0)}$, where

$$\mu_1(\theta_0) = \int_{\substack{\theta(q)>\theta_0-\Delta\theta \\ \theta(q)<\theta_0+\Delta\theta}} f(q)\ \text{d}q \quad \text{and} \quad \mu_2(\theta_0) = \int_{\substack{\theta(q)>\theta_0-\Delta\theta \\ \theta(q)<\theta_0+\Delta\theta}} 1\ \text{d}q$$

for some $\Delta\theta > 0$ (in our experiment we take $\Delta\theta = 3°$), and we approximate both $\mu_1(\theta_0)$ and $\mu_2(\theta_0)$ for several values of θ by their empirical means $\text{E}(f, k)$ and $\text{E}(1, k)$, using Theorem 4.1.1. We take $k = 10^4$ samples in both cases.

Figure 4.2.3 shows both the values of the empirical means in the logarithmic scale, and the ratio of $\mu_1(\theta_0)$ and $\mu_2(\theta_0)$.

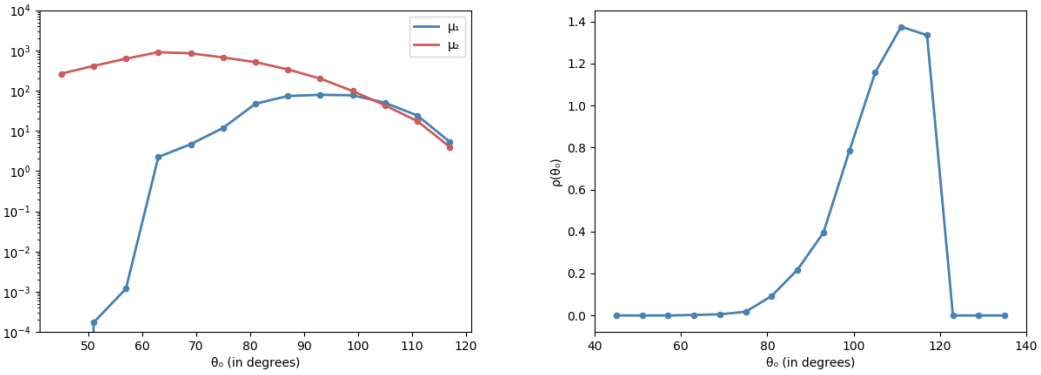

Figure 4.2.3: The left picture shows the approximations of $\mu_1(\theta_0)$ and $\mu_2(\theta_0)$ by the empirical means $\text{E}(f, k)$ and $\text{E}(1, k)$. Both integrals were approximated independently, each by an empirical mean obtained from 10^4 intersections with linear spaces. The right picture shows the ratio of the empirical means, which approximate $\rho(\theta_0)$.

How good is our estimate? From the plot above we can deduce that $\varepsilon = 2$ is a good accuracy for both $\mu_1(\theta_0)$ and $\mu_2(\theta_0)$. Using the variances s_1^2 and s_2^2 of the samples, respectively, we get $\frac{s_1^2}{\varepsilon^2 k} = 0.02$ and $\frac{s_2^2}{\varepsilon^2 k} = 0.04$. Hence, by Lemma 4.1.2 we expect that

the probability that the empirical mean $E(f,k)$ deviates from $\mu_1(\theta_0)$ by more than ε is at most 2%. And the probability that $E(1,k)$ deviates from $\mu_2(\theta_0)$ by more than ε is at most 4%. We conclude that our approximation of $\rho(\theta) = H\,\mathrm{Prob}\{\theta(q) = \theta_0\}$ is a good approximation.

In fact, Figure (4.2.3) shows a peak at around $\theta = 110°$. It is known that the total energy of the cyclohexane system is minimized when all angles between consecutive bonds achieve 110.9°; see [13, Ch. 2]. Therefore, our experiment gives a good approximation of the molecular geometry of cyclohexane. An example where all the angles between consecutive bonds are 110.9° is shown in Figure 4.2.4.

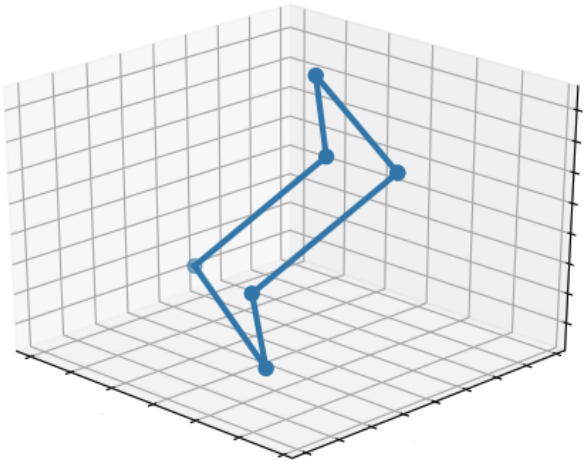

Figure 4.2.4: The picture shows a point from the variety (4.2.3), for which the angles between two consecutive bonds are all equal to 110.9° degrees. This configuration is also known as the "chair" [67].

Application to topological data analysis

Theorem 4.1.1 can be of interest for researchers working in topological data analysis using persistent homology (PH). Persistent homology is a tool to estimate the homology groups of a topological space from a finite point sample. The underlying idea is as follows: for varying t, put a ball of radius t around each point and compute the homology of the union of those balls. One then looks at topological features that persists for large intervals in t. It is intuitively clear that the point sample should be large enough to capture all of the topological information of its underlying space, and, on the other hand, the sample should be small enough to remain feasible for computations. Dufresne et al. [30] comment "Both the theoretical framework for the PH pipeline and its computational costs drive the requirements of a suitable sampling algorithm." (For an explanation of the PH pipeline

see [30, Sec. 2] and the references therein). They develop an algorithm that takes as input a denseness parameter ε and outputs a sample where each point has at most distance ε to its nearest neighbor. At the same time, their algorithm tries to keep the sample size as small as possible. In the context of topological data analysis our algorithm can be used as an alternative to [30].

In the following we use Theorem 4.1.1 for generating samples as input for the PH pipeline from [30]. The output of this pipeline is a *persistence diagram*. It shows the appearance and the vanishing of topological features in a 2-dimensional plot. Each point in the plot corresponds to an i-dimensional "hole", where the x-coordinate represents the time t when the hole appears, and the y-coordinate is the time when it vanishes. Points that are far from the line $x = y$ should be interpreted as signals coming from the underlying space. The number of those points is used as an estimator for the Betti number β_i. For computing persistence diagrams we use `Ripser` [4].

First, we consider two toy examples from [30, Sec. 5]: the surface $\mathcal{S}_1$ is given by

$$4x_1^4+7x_2^4+3x_3^4-3-8x_1^3+2x_1^2x_2-4x_1^2-8x_1x_2^2-5x_1x_2+8x_1-6x_2^3+8x_2^2+4x_2 = 0. \quad (4.2.4)$$

Figure 4.2.5 shows a sample of 386 points from the uniform distribution on $\mathcal{S}_1$. The associated persistence diagram suggests one connected component, two 1-dimensional and two 2-dimensional holes. The latter two come from the two sphere-like features of the variety. The outcome is similar to the diagram from [30, Fig. 6]. Considering that the diagram in this reference was computed using 1500 points [31], we conclude that the quality of our diagram is good.

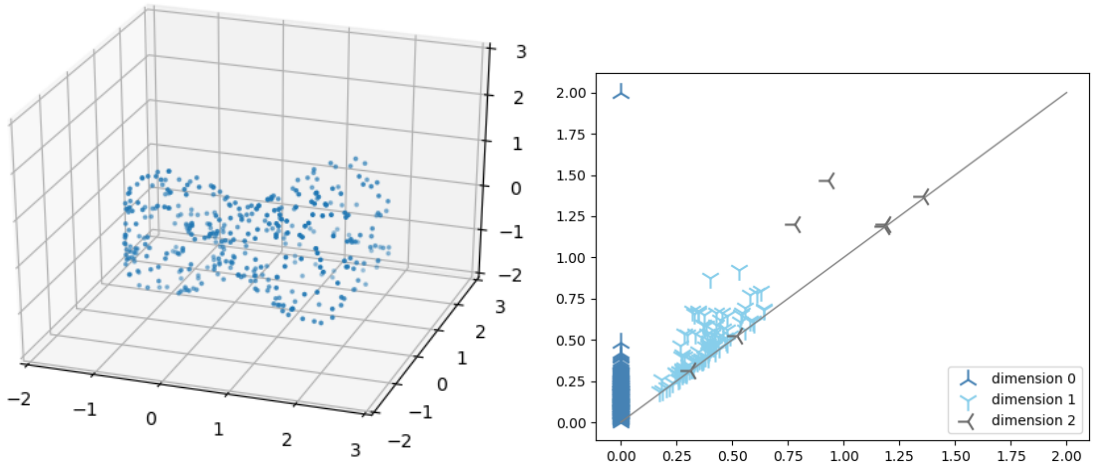

Figure 4.2.5: The left picture shows a sample of 386 points from the variety (4.2.4). The right picture shows the corresponding persistence diagram.

The second example is the surface $\mathcal{S}_2$ given by the equation

$$144(x_1^4+x_2^4)-225(x_1^2+x_2^2)x_3^2+350x_1^2x_2^2+81x_3^4+x_1^3+7x_1^2x_2+3(x_1^2+x_1x_2^2)-4x_1-5(x_2^3-x_2^2-x_2) = 0.$$

Figure 4.2.6 shows a sample of 651 points from the uniform distribution on $\mathcal{S}_2$. The persistence diagram on the right suggest one or five connected components. The true answer is five connected components. The diagram from [30, Fig. 6] captures the correct homology more clearly, but was generated from a sample of 10000 points [31].

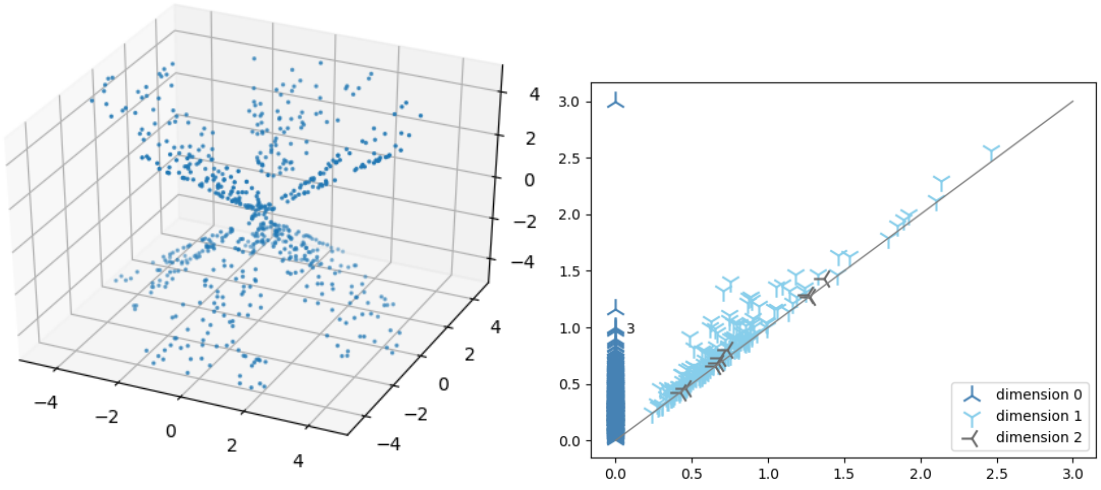

Figure 4.2.6: The left picture shows a sample of 651 points from the variety $\mathcal{S}_2$. The right picture shows the corresponding persistence diagram.

The next example is from a specific application in kinematics. Quoting [30, Sec. 5.3]: "Consider a regular pentagon in the plane consisting of links with unit length, and with one of the links fixed to lie along the x-axis with leftmost point at $(0, 0)$. The set of all possible configurations of this regular pentagon is a real algebraic variety." The equations of the configuration space are

$$(x_1+x_2+x_3)^2+(1+x_4+x_5+x_6)^2-1=0,\ x_1^2+x_4^2=1,\ x_2^2+x_5^2=1,\ x_3^2+x_6^2=1. \quad (4.2.5)$$

Here, the zeroth homology is of particular importance because if the variety is connected, "the mechanism has one assembly mode which can be continuously deformed to all possible configurations" [30]. Figure 4.2.7 shows the persistence diagram of a sample of 1400 points from the configuration space. It suggests that the variety indeed has only one connected component. We moreover observe eight holes of dimension 1 and one or three 2-dimensional holes. The correct Betti numbers are $\beta_0 = 1, \beta_1 = 8, \beta_2 = 1$; see [33].

4.3. Algebraic, geometric, and probabilistic tools

In this section we will first define the degree of a real algebraic variety and see why the number of intersection points of $\mathcal{M}$ with a linear space of the right codimension does not exceed the degree of its ambient variety. Then we recall the coarea formula of integration

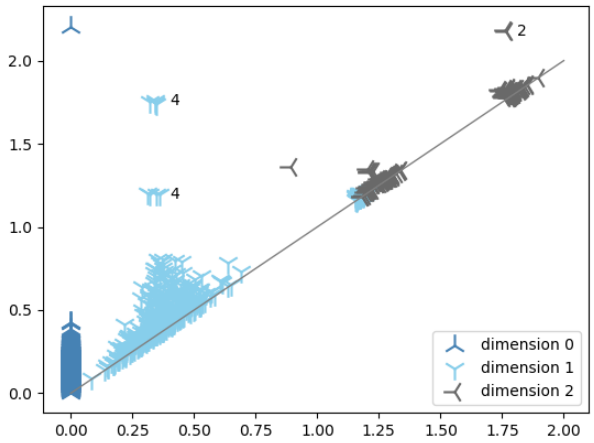

Figure 4.2.7: The picture shows the persistence diagram of a sample of 1400 points from the variety given by (4.2.5).

and discuss some consequences. Finally, we will see how to sample from $\psi(A, b)$ using rejection sampling and prove an algorithm for sampling the $\mathcal{L}_{A,b}$ in implicit form.

Real algebraic varieties

For the purpose of this chapter, a (real, affine) *algebraic variety* is a subset $\mathcal{V}$ of $\mathbb{R}^N$ such that there exists a set of polynomials $F_1, \ldots, F_k$ in N variables such that $\mathcal{V}$ is their set of common zeros. All varieties have a *dimension* and a *degree*. The dimension of $\mathcal{V}$ is defined as the dimension of its subspace of non-singular points $\mathcal{V}_0$, which is a manifold. For the degree we give a definition in the following steps.

An algebraic variety $\mathcal{V}$ in $\mathbb{R}^N$ is *homogeneous* if for all $t \in \mathbb{R} \setminus \{0\}$ and $x \in \mathcal{V}$ we have $tx \in \mathcal{V}$. Homogeneous varieties are precisely the ones where we can choose the F_i above to be homogenous polynomials. Homogeneous varieties live in the $(N-1)$-dimensional *real projective space* $\mathbb{P}^{N-1}$. This space is defined as the set $(\mathbb{R}^N \setminus \{0\})/\sim$, where $x \sim y$ if x and y are collinear. It comes with a canonical projection map $p\colon (\mathbb{R}^N \setminus 0) \to \mathbb{P}^{N-1}$. Then, a *projective variety* is defined as the image of a homogeneous variety $\mathcal{V}$ under p. Its *dimension* is $\dim \mathcal{V} - 1$.

Similarly, we define *complex affine, homogeneous, and projective varieties* by replacing $\mathbb{R}$ with $\mathbb{C}$ in the previous definitions. We can pass from real to complex varieties as follows. Let $\mathcal{V} \subset \mathbb{R}^N$ be a real affine variety. Its *complexification* $\mathcal{V}_\mathbb{C}$ is defined as the complex affine variety

$$\mathcal{V}_\mathbb{C} \coloneqq \{x \in \mathbb{C}^n : f(x) = 0 \text{ for all real polynomials } f \text{ vanishing on } \mathcal{V}\}.$$

The "all" is crucial here. Consider for instance the variety in $\mathbb{R}^2$ defined by $x_1^2 + x_2^2 = 0$. Obviously, this variety is a single point $\{(0,0)\}$, but the set $\{x \in \mathbb{C}^2 : x_1^2 + x_2^2 = 0\}$ equals the set $\{(t, \sqrt{-1}\, t) : t \in \mathbb{C}\}$ and is thus one-dimensional. Nevertheless, the polynomials $x_1 = 0, x_2 = 0$ also vanish on $\{(0,0)\}$ and so the complexification of $\mathcal{V} = \{(0,0)\}$ is $\mathcal{V}_{\mathbb{C}} = \{(0,0)\}$. The following lemma is important.

Lemma 4.3.1 (Lemma 8 in [86]). *The real dimension of $\mathcal{V}$ and the complex dimension of its complexification $\mathcal{V}_{\mathbb{C}}$ agree.*

The *Grassmannian* is a smooth algebraic variety $\mathrm{G}(k, \mathbb{C}^N)$ that parametrizes linear subspaces of $\mathbb{C}^N$ of dimension k. Furthermore, k-dimensional affine-linear subspaces of $\mathbb{C}^N$ can be seen as $(k+1)$-dimensional linear subspaces of $\mathbb{C}^{N+1}$ and are parametrized by the *affine Grassmannian* $\mathrm{G}_{\mathrm{Aff}}(k, \mathbb{C}^N)$. A *projective linear space* of dimension k is the image of a linear space $\mathcal{L} \in \mathrm{G}(k+1, \mathbb{C}^N)$ under the projection p. This motivates to define the projective Grassmannian as $\mathrm{G}(k, \mathbb{P}^{N-1}) \coloneqq \{p(\mathcal{L}) : \mathcal{L} \in \mathrm{G}(k+1, \mathbb{C}^N)\}$.

We have gathered all the material to give a precise definition of the degree: let $\mathcal{V} \subset \mathbb{P}_{\mathbb{C}}^{N-1}$ be a complex projective variety of dimension n. There exists a unique natural number d and a lower-dimensional subvariety $\mathcal{W}$ of $\mathrm{G}(N-n, \mathbb{P}^{N-1})$ with the property that for all linear spaces $\mathcal{L} \in \mathrm{G}(N-n, \mathbb{P}^{N-1}) \backslash \mathcal{W}$ the intersection $\mathcal{V} \cap \mathcal{L}$ consists of d distinct points [44, Sect. 18]. Furthermore, the number of such intersection points only decreases when $\mathcal{L} \in \mathcal{W}$. This number d is called the *degree* of the projective variety $\mathcal{V}$. The *degree* of a complex affine variety $\mathcal{V} \subset \mathbb{C}^N$ is defined as the degree of the smallest projective variety containing the image of $\mathcal{V}$ under the embedding $\mathbb{C}^N \hookrightarrow \mathbb{P}_{\mathbb{C}}^N$ sending x to $p([1, x])$.

The definition of degree of complex varieties is standard in algebraic geometry. In this chapter however we are solely dealing with real varieties. We therefore make the following definition, which is not standard in the literature, but which fits in our setting.

Definition 4.3.2. The *degree* of a real affine variety $\mathcal{V}$ is the degree of its complexification. The *degree* of a real projective variety $\mathcal{V}$ is the degree of the image of the complexification of $p^{-1}(\mathcal{V})$ under p.

Using Lemma 4.3.1 we make the following conclusions, after passing from the Grassmannians $\mathrm{G}_{\mathrm{Aff}}(N-n, \mathbb{R}^N)$ and $\mathrm{G}(N-n, \mathbb{R}^N)$ to the parameter spaces $\mathbb{R}^{n \times N} \times \mathbb{R}^n$ and $\mathbb{R}^{n \times N}$.

Lemma 4.3.3. *Let $\mathcal{V} \subset \mathbb{R}^N$ be an affine variety of dimension n and degree d. Except for a lower-dimensional subset of $\mathbb{R}^{n \times N} \times \mathbb{R}^N$, all affine linear subspaces $\mathcal{L}_{A,b} \subset \mathbb{R}^N$ defined by $\mathcal{L}_{A,b} = \{x \in \mathbb{R}^N : Ax = b\}$ intersect $\mathcal{V}$ in at most d many points.*

Let $\mathcal{V} \subset \mathbb{P}^{N-1}$ be a projective variety of dimension n and degree d. Except for a lower-dimensional subset of $\mathbb{R}^{n \times N}$, all linear subspaces $\mathcal{L}_A = \{x \in \mathbb{P}^{N-1} : Ax = 0\} \subset \mathbb{P}^{N-1}$ intersect $\mathcal{V}$ in at most d many points.

The coarea formula

The *coarea formula of integration* is a key ingredient in the proof of Theorem 4.1.1. This formula says how integrals transform under smooth maps. A well-known special case is integration by substitution. The coarea formula generalizes this from integrals defined on the real line to integrals defined on differentiable manifolds.

Let $\mathcal{M},\mathcal{N}$ be Riemannian manifolds and $\mathrm{d}v$, $\mathrm{d}w$ be the respective volume forms. Furthermore, let $h : \mathcal{M} \to \mathcal{N}$ be a smooth map. A point $v \in \mathcal{M}$ is called a *regular point of* h if $\mathrm{D}_v h$ is surjective. Note that a necessary condition for regular points to exist is $\dim \mathcal{M} \geq \dim \mathcal{N}$.

For any $v \in \mathcal{M}$ the Riemannian metric on $\mathcal{M}$ defines orthogonality on $\mathrm{T}_v\mathcal{M}$. For a regular point v of h this implies that the restriction of $\mathrm{D}_v h$ to the orthogonal complement of its kernel is a linear isomorphism. The absolute value of the determinant of that isomorphism is the *normal Jacobian of* h *at* v. Let us summarize this in a definition.

Definition 4.3.4. Let $h : \mathcal{M} \to \mathcal{N}$ be a smooth map and $v \in \mathcal{M}$ be a regular point of h. Let $(\,\cdot\,)^{\perp}$ denote the orthogonal complement. The normal Jacobian of h at v is defined as

$$\mathrm{NJ}(h, v) := \left|\det\left(\mathrm{D}_v h\,|_{(\ker \mathrm{D}_v h)^{\perp}}\right)\right|.$$

We also need the following theorem (see, e.g., [14, Theorem A.9]).

Theorem 4.3.5. *Let* $\mathcal{M},\mathcal{N}$ *be smooth manifolds with* $\dim \mathcal{M} \geq \dim \mathcal{N}$ *and let* $h : \mathcal{M} \to \mathcal{N}$ *be a smooth map. Let* $w \in \mathcal{N}$ *be such that all* $v \in h^{-1}(w)$ *are regular points of* h*. Then, the fiber* $h^{-1}(w)$ *over* w *is a smooth submanifold of* $\mathcal{M}$ *of dimension* $\dim \mathcal{M} - \dim \mathcal{N}$ *and the tangent space of* $h^{-1}(w)$ *at* v *is* $\mathrm{T}_v h^{-1}(w) = \ker \mathrm{D}_v h$*.*

A point $w \in \mathcal{N}$ satisfying the properties in the previous theorem is called *regular value of* h. By *Sard's lemma* the set of all $w \in \mathcal{N}$ that are not a regular value of h is a set of measure zero. We are now equipped with all we need to state the coarea formula. See [45, (A-2)] for a proof.

Theorem 4.3.6 (The coarea formula of integration)**.** *Suppose that* $\mathcal{M},\mathcal{N}$ *are Riemannian manifolds, and let* $h : \mathcal{M} \to \mathcal{N}$ *be a surjective smooth map. For every function* $a : \mathcal{M} \to \mathbb{R}$ *that is integrable with respect to the volume measure of* $\mathcal{M}$ *we have*

$$\int_{v\in\mathcal{M}} a(v)\,\mathrm{d}v = \int_{w\in\mathcal{N}} \left(\int_{u\in h^{-1}(w)} \frac{a(u)}{\mathrm{NJ}(h,u)}\,\mathrm{d}u\right)\mathrm{d}w,$$

where $\mathrm{d}u$ *is the volume form on the submanifold* $h^{-1}(w)$*.*

The following corollary from the coarea formula is important.

Corollary 4.3.7. *Let $h\colon \mathcal{M} \to \mathcal{N}$ be a smooth surjective map of Riemannian manifolds.*

(1) Let X be a random variable on $\mathcal{M}$ with density β. Then $h(X)$ is a random variable on $\mathcal{N}$ with density

$$\gamma(y) = \int_{x \in h^{-1}(y)} \frac{\beta(x)}{\mathrm{NJ}(h,x)}\,\mathrm{d}x.$$

(2) Let ψ be a density on $\mathcal{N}$ and for all $y \in \mathcal{N}$, let ρ_y be a density on $h^{-1}(y)$. The random variable X on $\mathcal{M}$ obtained by independently taking $Y \in \mathcal{N}$ with density ψ and $X \in h^{-1}(Y)$ with density ρ_Y has density

$$\beta(x) = \psi(h(x))\rho_{h(x)}(x)\mathrm{NJ}(h,x).$$

Proof. The first part follows directly from the coarea formula. For the second part, it suffices to note that for measurable $\mathcal{U} \subset \mathcal{M}$ we have

$$\begin{aligned}\int_{y \in h(\mathcal{U})} \int_{x \in h^{-1}(y) \cap \mathcal{U}} \psi(h(x))\rho_{h(x)}(x)\,\mathrm{d}x\,\mathrm{d}y &= \int_{y \in h(\mathcal{U})} \int_{x \in h^{-1}(y) \cap \mathcal{U}} \frac{\beta(x)}{\mathrm{NJ}(h,x)}\,\mathrm{d}x\,\mathrm{d}y \\ &= \int_{x \in \mathcal{U}} \beta(x)\,\mathrm{d}x,\end{aligned}$$

see also [14, Rem. 17.11]. □

Sampling from the density on affine-linear subspaces

Our method for sampling from an algebraic manifold involves taking a distribution φ on the parameter space $\mathbb{R}^{n\times N} \times \mathbb{R}^n$ of hyperplanes of the right dimension which is easy to sample, and turning it into another density ψ. Here, we explain how to sample from ψ with rejection sampling. In the following we denote elements of $\mathbb{R}^{n\times N} \times \mathbb{R}^n$ by (A, b).

Proposition 4.3.8. *Let κ be any number satisfying $0 < \kappa \cdot \sup_{(A,b)} \overline{f}(A,b) \leq 1$. Consider the binary random variable $Z \in \{0,1\}$ with $\mathrm{Prob}\{Z = 1 \mid (A,b)\} = \kappa\,\overline{f}(A,b)$. Then, ψ is the density of the conditional random variable $((A,b) \mid Z = 1)$.*

Proof. We denote the density of the conditional random variable $((A,b) \mid Z = 1)$ by λ. Bayes' Theorem implies $\lambda(A,b)\ \mathrm{Prob}\{Z = 1\} = \mathrm{Prob}\{Z = 1 \mid (A,b)\}\,\varphi(A,b)$, which by assumption is equivalent to

$$\lambda(A,b) = \frac{\kappa\,\overline{f}(A,b)\,\varphi(A,b)}{\mathrm{Prob}\{Z = 1\}}.$$

By the definition of Z we have $\mathrm{Prob}\{Z = 1\} = \kappa\ \mathbb{E}_{(A,b)\sim\varphi}\ \overline{f}(A,b)$. Hence,

$$\lambda(A,b) = \frac{\overline{f}(A,b)\,\varphi(A,b)}{\mathbb{E}_{(A,b)\sim\varphi}\ \overline{f}(A,b)} = \psi(A,b).$$

This finishes the proof. □

Proposition 4.3.8 shows that ψ is the density of a conditional distribution. A way to sample from such distributions is by *rejection sampling*: for sampling $((A,b) \mid Z=1)$ we may sample from the joint distribution (A,b,Z) and then keep only the points with $Z=1$. The strong law of large numbers implies the correctness of rejection sampling. Indeed, if (A_i,b_i,Z_i) is a sequence of i.i.d. copies of (A,b,Z) and $\mathcal{U}$ is a measurable set with respect to the Lebesgue measure on $\mathbb{R}^{n\times N}\times\mathbb{R}^n$, then we have the following convergence with probability one:

$$\begin{aligned}\frac{\#\{i \mid (A_i,b_i)\in\mathcal{U}, Z_i=z, i\le n\}}{\#\{i \mid Z_i=z, i\le n\}} &= \frac{\frac{1}{n}\#\{i \mid (A_i,b_i)\in\mathcal{U}, Z_i=z, i\le n\}}{\frac{1}{n}\#\{i \mid Z_i=z, i\le n\}}\\ &\xrightarrow{n\to\infty} \frac{\operatorname{Prob}\{(A,b)\in\mathcal{U}, Z=z\}}{\operatorname{Prob}\{Z=z\}}\\ &= \operatorname{Prob}_{(A,b)\mid Z=z}(\mathcal{U}).\end{aligned}$$

For sampling Z, however, we must compute a suitable κ. This can be done as follows. Let d be the degree of the ambient variety of $\mathcal{M}$. We assume we know upper bounds K for $f(x)$ and C for $\|x\|^2$, both as x ranges over $\mathcal{M}$, and set

$$\kappa = \frac{1}{dK}\frac{\Gamma(\frac{n+1}{2})}{\sqrt{\pi}^{n+1}}\frac{1}{1+C}.$$

Then we have $0<\kappa\overline{f}(A,b)\le 1$ for all (A,b) as needed. With κ, we have everything we need to carry out the sampling method.

How to obtain the upper bounds K and C? For K, we might just know the maximum of f. For example, if we want to sample from the uniform distribution, then we may use $f=1$. In more complicated cases, we could approximate $\max f$ by repeatedly sampling $(A,b)\sim\varphi$ and recording the highest value f takes on the points in the intersection $\mathcal{M}\cap\mathcal{L}_{(A,b)}$. Casella and Robert [72] call this approach *stochastic exploration*.

We might know C a priori, for example because we restrict the manifold $\mathcal{M}$ to a box in $\mathbb{R}^N$. We can also restrict the manifold to a box after determining by sampling what the size of the box should be. We could also estimate $\max\|x\|^2$ by sampling as for $\max f$. Sometimes we can also use *semidefinite programming* [7] to bound polynomial functions like $\|x\|^2$ on a variety. Note that the probability for rejection increases as C increases. We thus seek a C which is as small as possible. If our given function f is invariant under translation, we may translate $\mathcal{M}$ to decrease C. For instance, sampling from the uniform distribution on the circle $(x_1-100)^2+(x_2-100)^2=1$ is the same as sampling on $x_1^2+x_1^2=1$ and then translating by adding $(100,100)$ to each sample point. The difference between the two is that for the first variety we need $C=101$, whereas for the second we can use $C=1$.

Sampling linear spaces in explicit form

Sometimes it is useful to sample the linear space $\mathcal{L}_{A,b}$ in explicit form, and not in implicit form $Ax = b$. For instance, if $\mathcal{V}$ is a hypersurface given by an equation $F(x) = 0$, then intersecting $\mathcal{V}$ with a line $u + tv$ can be done by solving the univariate equation $F(u+tv) = 0$. The next lemma shows how to pass from implicit to explicit representation in the Gaussian case.

Lemma 4.3.9. *Let $\mathcal{L}_{A,b} = \{x \in \mathbb{R}^N \mid Ax = b\}$ be a random affine linear space given by i.i.d. standard normal entries for $A \in \mathbb{R}^{n\times N}$ and $b \in \mathbb{R}^n$. Consider another random linear space*

$$\mathcal{K}_{u,v_1,\ldots,v_{N-m}} = \{u + t_1v_1 + \cdots t_{N-n}v_{N-n} \mid t_1, \ldots, t_{N-n} \in \mathbb{R}\},$$

where $u, v_1, \ldots, v_{N-m}$ are obtained as follows. Sample a matrix $U \in \mathbb{R}^{(N-n+1)\times(N+1)}$ with i.i.d. standard normal entries, and let

$$\begin{pmatrix} u \\ 1 \end{pmatrix}, \begin{pmatrix} v_1 \\ 0 \end{pmatrix}, \ldots, \begin{pmatrix} v_{N-n} \\ 0 \end{pmatrix} \in \operatorname{rowspan}(U).$$

Then, we have $\mathcal{K}_{u,v_1,\ldots,v_{N-m}} \sim \mathcal{L}_{A,b}$.

Proof. Consider the linear space $\widetilde{\mathcal{L}}_{A,b} := \{z \in \mathbb{R}^{N+1} \mid [A, -b]z = 0\}$. This is a random linear space in the Grassmannian $\mathrm{G}(N+1-n, \mathbb{R}^{N+1})$. The affine linear space is given as $\mathcal{L}_{A,b} := \{u + t_1v_1 + \cdots t_{N-n}v_{N-n}\}$, where

$$\begin{pmatrix} u \\ 1 \end{pmatrix}, \begin{pmatrix} v_1 \\ 0 \end{pmatrix}, \ldots, \begin{pmatrix} v_{N-n} \\ 0 \end{pmatrix} \in \ker([A, -b]).$$

Now the kernel of $[A, -b]$ is a random linear space in $\mathrm{G}(n, \mathbb{R}^{N+1})$, which is invariant under orthogonal transformations. By [58] there is a unique orthogonally invariant probability distribution on the Grassmannian $\mathrm{G}(n, \mathbb{R}^{N+1})$. Since $\operatorname{rowspan}(U)$ is also orthogonally invariant, we find that $\operatorname{rowspan}(U) \sim \ker([A, -b])$, which concludes the proof. □

4.4. Proofs of Theorem 4.1.1 and Lemma 4.1.2

We begin by giving an alternate description of the function $\alpha(x)$.

Lemma 4.4.1. $\alpha(x) = \int_{A\in\mathbb{R}^{n\times N}} \varphi(A, Ax) |\det(A|_{\mathrm{T}_x\mathcal{M}})| \, \mathrm{d}A$.

Proof. Let $\alpha'(x)$ be the right hand side of the formula. Let $U \in O(N)$ be an orthogonal matrix such that $Ux = (0, \ldots, 0, x_N)^T$ and consider the manifold $\mathcal{N} = U \cdot \mathcal{M}$. We have $\mathrm{T}_{Ux}\mathcal{N} = U\mathrm{T}_x\mathcal{M}$ and $\det(A|_{\mathrm{T}_x\mathcal{M}}) = \det(AU^T\big|_{\mathrm{T}_{Ux}\mathcal{N}})$. After the change of variables $A \mapsto AU^T$ we get

$$\alpha'(x) = \int_{A\in\mathbb{R}^{n\times N}} \left|\det(A|_{\mathrm{T}_{Ux}\mathcal{N}})\right| \, \varphi(A, AUx) \, \mathrm{d}A.$$

By definition of the Gaussian density, we have $\varphi(A, AUx) = \frac{1}{(\sqrt{2\pi})^n}\phi(AR)$, where ϕ is the Gaussian density on $\mathbb{R}^{n\times N}$ and $R = \operatorname{diag}(1,\dots,1,\sqrt{1+x_N^2}) \in \mathbb{R}^{N\times N}$. Let us write $B = AR$. A change of variables from A to B yields

$$\alpha'(x) = \frac{1}{\sqrt{1+x_N^2}(\sqrt{2\pi})^n}\int_{B\in\mathbb{R}^{n\times N}} \left|\det(BR^{-1}\big|_{\mathrm{T}_{Ux}\mathcal{N}})\right| \ \phi(B)\,\mathrm{d}B.$$

Let $W \in \mathbb{R}^{N\times n}$ be a matrix whose columns form an orthonormal basis for $\mathrm{T}_{Ux}\mathcal{N}$ and write $M := R^{-1}W$. Then we have $\det(BR^{-1}\big|_{\mathrm{T}_{Ux}\mathcal{N}}) = \det(BM)$ and so

$$\alpha'(x) = \frac{1}{\sqrt{1+x_N^2}\,\sqrt{2\pi}^n}\ \mathop{\mathbb{E}}_{B\sim\phi}|\det(BM)|\,.$$

We now write $\mathbb{E}_{B\sim\phi}|\det(BM)| = \mathbb{E}_{B\sim\phi}\det(M^TB^TBM)^{\frac{1}{2}}$. By [66, Thm. 3.2.5] the matrix $C := M^TB^TBM \in \mathbb{R}^{n\times n}$ is a *Wishart matrix* with covariance matrix M^TM. By [66, Thm. 3.2.15], we have $\mathbb{E}\det(C)^{\frac{1}{2}} = \det(M^TM)^{\frac{1}{2}}\frac{1}{\sqrt{\pi}}\sqrt{2}^n\Gamma(\frac{n+1}{2})$. Altogether, this shows that

$$\alpha'(x) = \frac{\det(M^TM)^{\frac{1}{2}}}{\sqrt{1+x_N^2}}\,\frac{\Gamma\left(\frac{n+1}{2}\right)}{\sqrt{\pi}^{n+1}}.$$

Moreover, we have

$$\begin{aligned} M^TM &= W^TR^{-T}R^{-1}W \\ &= W^T\operatorname{diag}(1,\dots,1,\tfrac{1}{1+x_N^2})W \\ &= \mathbf{1} - \tfrac{1}{1+x_N^2}W\operatorname{diag}(0,\dots,0,x_N^2)W. \end{aligned}$$

The second summand in the last expression is a rank-one matrix with the single non-zero eigenvalue $-\frac{||W^TUx||^2}{1+||Ux||^2}$. Taking determinants we get

$$\det(M^TM) = 1 - \frac{||W^TUx||^2}{1+||Ux||^2} = \frac{1+||Ux||^2-||W^TUx||^2}{1+||Ux||^2} = \frac{1+||x||^2-||\Pi_{\mathrm{T}_x\mathcal{M}}x||^2}{1+||x||^2},$$

where $\Pi_{\mathrm{T}_x\mathcal{M}}$ denotes the orthogonal projection onto the tangent space. Since we have $||x||^2 - ||\Pi_{\mathrm{T}_x\mathcal{M}}x||^2 = \langle x, \Pi_{\mathrm{N}_x\mathcal{M}}\,x\rangle$, this implies

$$\alpha'(x) = \frac{\sqrt{1+\langle x, \Pi_{\mathrm{N}_x\mathcal{M}}\,x\rangle}}{1+\|x\|^2}\,\frac{\Gamma\left(\frac{n+1}{2}\right)}{\sqrt{\pi}^{n+1}} = \alpha(x).$$

This concludes the proof. □

We are now prepared to prove Theorem 4.1.1.

Proof of Theorem 4.1.1. We first prove the first part. The support of $\overline{f}(A,b)$ is a full dimensional subset of $\mathbb{R}^{n\times N}\times\mathbb{R}^n$ and it is contained in the complement of the set of all (A,b) for which $\mathcal{M}\cap\mathcal{L}_{A,b}=\emptyset$. We let $\mathcal{X}$ denote the interior of the support of $\overline{f}(A,b)$, so that $\mathbb{E}_{(A,b)\sim\varphi}\,\overline{f}(A,b)=\int_{\mathcal{X}}\overline{f}(A,b)\varphi(A,b)\mathrm{d}(A,b)$.

Let $\pi\colon\mathbb{R}^{n\times N}\times\mathcal{M}\to\mathcal{X}$ be the map sending a pair (A,x) to (A,Ax). We have $\mathrm{D}_{(A,x)}\pi\,(\dot{A},\dot{x})=(\dot{A},\dot{A}x+A\dot{x})$, so the derivative of π can be identified with the matrix $\left(\begin{smallmatrix}\mathbf{1} & 0\\ * & A\end{smallmatrix}\right)$. This shows that $\mathrm{NJ}(\pi,(A,x))=|\det(A|_{\mathrm{T}_x\mathcal{M}})|$. Therefore, by Theorem 4.3.6,

$$\mathop{\mathbb{E}}_{(A,b)\sim\varphi}\overline{f}(A,b)=\int_{\mathbb{R}^{n\times N}\times\mathcal{M}}\frac{f(x)}{\alpha(x)}\,|\det(A|_{\mathrm{T}_x\mathcal{M}})|\,\varphi(A,Ax)\;\mathrm{d}(A,x).$$

The projection $\mathbb{R}^{n\times N}\times\mathcal{M}\to\mathcal{M}$ on the second factor has normal Jacobian one everywhere. Applying Theorem 4.3.6 again yields

$$\mathop{\mathbb{E}}_{(A,b)\sim\varphi}\overline{f}(A,b)=\int_{\mathcal{M}}\frac{f(x)}{\alpha(x)}\left(\int_{\mathbb{R}^{n\times N}}|\det(A|_{\mathrm{T}_x\mathcal{M}})|\,\varphi(A,Ax)\;\mathrm{d}A\right)\mathrm{d}x=\int_{\mathcal{M}}f(x)\mathrm{d}x,$$

the second inequality by Lemma 4.4.1. This proves the first part.

Now, we prove the second part, where we assume that $f\colon\mathcal{M}\to\mathbb{R}_{>0}$ is nonnegative. Recall that $\psi(A,b)=\frac{\varphi(A,b)\overline{f}(A,b)}{\mathbb{E}_\varphi(\overline{f})}$. Since $\mathbb{E}_\varphi(\overline{f})=\int_{\mathcal{M}}f(x)\mathrm{d}x$ is positive and finite by the first part of the theorem, we find that ψ is a well defined probability density. The support of ψ is contained in the closure of $\mathcal{X}$ and therefore $\mathcal{M}\cap\mathcal{L}_{A,b}$ is almost surely non-empty and finite.

Let $Y=(A,x)\in\mathbb{R}^{n\times N}\times\mathcal{M}$ be the random variable defined by first choosing $(A,b)\sim\psi$ and then taking $x\in\mathcal{M}\cap\mathcal{L}_{A,b}$ with probability $f(x)\alpha(x)^{-1}\overline{f}(A,b)^{-1}$. By construction, $\pi(Y)\sim\psi$. We use Corollary 4.3.7 (2) and find that Y has density

$$\beta(A,x)=\frac{\psi(A,Ax)f(x)\mathrm{NJ}(\pi,(A,x))}{\alpha(x)\overline{f}(A,Ax)}.$$

Recall that $\mathrm{NJ}(\pi,(A,x))=|\det(A|_{\mathrm{T}_x\mathcal{M}})|$ and that the projection $\mathbb{R}^{n\times N}\times\mathcal{M}\to\mathcal{M}$ on the second factor has normal Jacobian one everywhere. Therefore, by Corollary 4.3.7 (1), the random point $x\in\mathcal{M}$ has density γ with

$$\begin{aligned}\gamma(x)&=\int_{A\in\mathbb{R}^{n\times N}}\beta(A,x)\,\mathrm{d}A\\&=\frac{f(x)}{\alpha(x)}\int_{A\in\mathbb{R}^{n\times N}}\frac{\psi(A,Ax)|\det(A|_{\mathrm{T}_x\mathcal{M}})|}{\overline{f}(A,Ax)}\,\mathrm{d}A\\&=\frac{f(x)}{\alpha(x)\mathbb{E}_\varphi(\overline{f})}\int_{A\in\mathbb{R}^{n\times N}}\varphi(A,Ax)|\det(A|_{\mathrm{T}_x\mathcal{M}})|\,\mathrm{d}A.\end{aligned}$$

Using Lemma 4.4.1 yields $\gamma(x)=\frac{f(x)}{\mathbb{E}_\varphi(\overline{f})}$. This finishes the proof of Theorem 4.1.1. □

Next, we prove the lemma about the rate of convergence of our estimator.

Proof of Lemma 4.1.2. First, we have the bound $\alpha(x) \geq \frac{1}{1+\sup_{x\in\mathcal{M}}\|x\|^2}\,\frac{\Gamma\left(\frac{n+1}{2}\right)}{\sqrt{\pi}^{n+1}}$. Let d be the degree of the ambient variety of $\mathcal{M}$. With probability one $\mathcal{M}\cap\mathcal{L}_{A,b}$ consists of at most d points and so we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{(A,b)\sim\varphi}\overline{f}(A,b)^2 &= \mathop{\mathbb{E}}_{(A,b)\sim\varphi}\left(\sum_{x\in\mathcal{M}\cap\mathcal{L}_{A,b}}\frac{f(x)}{\alpha(x)}\right)^2 \\
&\leq \mathop{\mathbb{E}}_{(A,b)\sim\varphi}\left(\sum_{x\in\mathcal{M}\cap\mathcal{L}_{A,b}}\frac{|f(x)|}{\alpha(x)}\right)^2 \\
&\leq \mathop{\mathbb{E}}_{(A,b)\sim\varphi}\left(\sum_{x\in\mathcal{M}\cap\mathcal{L}_{A,b}}\frac{\sup_{x\in\mathcal{M}}|f(x)|}{\inf_{x\in\mathcal{M}}\alpha(x)}\right)^2 \\
&\leq d^2(1+\sup_{x\in\mathcal{M}}\|x\|^2)^2\frac{\pi^{n+1}}{\Gamma\left(\frac{n+1}{2}\right)^2}\sup_{x\in\mathcal{M}}f(x)^2.
\end{aligned}$$

We also have $\sigma^2(\overline{f}) \leq \mathbb{E}_{(A,b)\sim\varphi}\,\overline{f}(A,b)^2$, and therefore

$$\sigma^2(\overline{f}) \leq d^2(1+\sup_{x\in\mathcal{M}}\|x\|^2)^2\frac{\pi^{n+1}}{\Gamma\left(\frac{n+1}{2}\right)^2}\sup_{x\in\mathcal{M}}f(x)^2 \tag{4.4.1}$$

is finite. We may therefore use Chebyshev's inequality to deduce that

$$\mathrm{Prob}\left\{\left|\mathrm{E}(f,k)-\int_{\mathcal{M}}f(x)\,\mathrm{d}x\right|\geq\varepsilon\right\}\leq\frac{\sigma^2(\overline{f})}{\varepsilon^2 k}. \tag{4.4.2}$$

This finishes the proof. □

4.5. Sampling from projective manifolds

In this section we prove a variation of Theorem 4.1.1 for projective algebraic manifolds.

Real projective space $\mathbb{P}^{N-1}$ from Section 4.3 is a compact Riemannian manifold with a canonical metric, the *Fubini-Study* metric. Namely, let $p:\mathbb{R}^N\backslash\{0\}\to\mathbb{P}^{N-1}$ be the canonical projection. Restricted to the unit sphere $\mathbb{S}^{N-1}$, the projection p identifies antipodal points. We define a subset $\mathcal{U}\subset\mathbb{P}^{N-1}$ to be open if and only if $p|_{\mathbb{S}^{N-1}}^{-1}(\mathcal{U})$ is open. This gives $\mathbb{P}^{N-1}$ the structure of a differential manifold. The Riemannian structure on $\mathbb{P}^{N-1}$ is defined as $\langle\dot{a},\dot{b}\rangle := \langle \mathrm{D}_x p^{-1}\dot{a}, \mathrm{D}_x p^{-1}\dot{b}\rangle$ for $\dot{a},\dot{b}\in\mathrm{T}_x\mathbb{P}^{N-1}$. This metric is called the Fubini-Study metric, and it induces the *standard measure* on $\mathbb{P}^{N-1}$.

We say that $\mathcal{M}$ is a *projective algebraic manifold* if it is an open submanifold of the smooth part of a real projective variety $\mathcal{V} \subset \mathbb{P}^{N-1}$. We assume $\mathcal{M}$ to be n-dimensional, and consider a function $f\colon \mathcal{M} \to \mathbb{R}_{\geq 0}$ with a well-defined scaled probability density $f(x)/\int_{\mathcal{M}} f(x)\mathrm{d}x$. For $A \in \mathbb{R}^{n\times N}$, define the linear space $\mathcal{L}_A = \{x \in \mathbb{P}^{N-1} \mid Ax = 0\}$ and write

$$\overline{f}(A) := \sum_{x \in \mathcal{M}\cap\mathcal{L}_A} f(x).$$

In this section, we denote by φ_ℓ the density of the multivariate standard normal distribution on $\mathbb{R}^\ell$.

Theorem 4.5.1. *In the notation introduced above:*

(1) Let f be an integrable function on $\mathcal{M}$. We have

$$\int_{\mathcal{M}} f(x)\,\mathrm{d}(x) = \mathrm{vol}(\mathbb{P}^n) \mathop{\mathbb{E}}_{A\sim\varphi_{n\times N}} \overline{f}(A).$$

(2) Let f be nonnegative and assume that the integral $\int_{\mathcal{M}} f(x)\,\mathrm{d}(x)$ is finite and nonzero. Let $X \in \mathcal{M}$ be the random variable obtained by choosing $A \in \mathbb{R}^{n\times N}$ with probability $\psi(A) := \frac{\varphi(A)\,\overline{f}(A)}{\mathbb{E}_{A\sim\varphi_{n\times N}}\,\overline{f}(A)}$ and one of the finitely many points $X \in \mathcal{M}\cap\mathcal{L}_A$ with probability $f(x)/\overline{f}(A)$. Then X is distributed according to the density $f(x)/\int_{\mathcal{M}} f(x)\mathrm{d}x$.

Remark 4.5.2. In [54, Sec. 2.4] Lairez proved a similar theorem for the uniform distribution on complex projective varieties.

Sampling $\mathcal{L}_A$ with $A \sim \varphi_{n\times N}$ yields a special distribution on $\mathrm{G}(N-n-1,\mathbb{P}^{N-1})$. By [58] there is a unique orthogonally invariant probability measure ν on the Grassmannian $\mathrm{G}(N-n-1,\mathbb{P}^{N-1})$. Since the distribution of the kernel of a Gaussian A is invariant under orthogonal transformations, the projective plane $\mathcal{L}_A = \{x \in \mathbb{P}^{N-1} : Ax = 0\}$ has distribution ν.

Furthermore, setting $f = 1$ in Theorem 4.5.1 gives the formula

$$\mathrm{vol}(\mathcal{M}) = \mathrm{vol}(\mathbb{P}^n) \mathop{\mathbb{E}}_{A\sim\varphi_{n\times N}} |\mathcal{M}\cap\mathcal{L}_A|.$$

This is the *kinematic formula for projective manifolds* from [45, Thm. 3.8] in disguise.

Before we can prove Theorem 4.5.1, we have to prove an auxiliary lemma, similar to Lemma 4.4.1.

Lemma 4.5.3. *For any $x \in \mathcal{M}$ we have*

$$\int_{A\in\mathbb{R}^{n\times N}:Ax=0} |\det(A|_{\mathrm{T}_x\mathcal{M}})|\,\varphi_{n\times N}(A)\,\mathrm{d}A = \frac{1}{\mathrm{vol}(\mathbb{P}^n)}.$$

In particular, the integral is independent of x.

Proof. Let $\mathcal{H}(x) := \{A \in \mathbb{R}^{n\times N} \mid Ax = 0\}$. It is a linear subspace of $\mathbb{R}^{n\times N}$ of codimension n. Let $U \in \mathbb{R}^{N\times n}$ be a matrix whose columns form an orthonormal basis for $\mathrm{T}_x\mathcal{M}$, so that $\det(A|_{\mathrm{T}_x\mathcal{M}}) = \det(AU)$. Furthermore, let $O \in \mathbb{R}^{N\times N}$ be an orthogonal matrix with $Ox = e_1$, where $e_1 = (1, 0, \ldots, 0)^T \in \mathbb{R}^N$. Then, $\mathcal{H}(e_1)O = \mathcal{H}(x)$. Making a change of variables $A \mapsto AO$ we get

$$\int_{A\in\mathcal{H}(x)} |\det(A|_{\mathrm{T}_x\mathcal{M}})|\, \varphi_{n\times N}(A)\,\mathrm{d}A = \int_{A\in\mathcal{H}(e_1)} |\det(AOU)|\, \varphi_{n\times N}(AO)\,\mathrm{d}A. \tag{4.5.1}$$

We have $\varphi_{n\times N}(AO) = \varphi_{n\times N}(A)$, because the Gaussian distribution is orthogonally invariant. Moreover, any $A \in \mathcal{H}(e_1)$ is of the form $A = [0, A']$ with $A' \in \mathbb{R}^{n\times(N-1)}$, and we have $\varphi_{n\times N}(A) = \frac{1}{\sqrt{2\pi}^n}\varphi_{n\times(N-1)}(A')$. Let us denote by O' the lower $(N-1)\times n$ part of OU, so that $AOU = A'O'$. It follows that (4.5.1) is equal to

$$\frac{1}{\sqrt{2\pi}^n}\int_{A'\in\mathbb{R}^{n\times(N-1)}} |\det(A'O')|\, \varphi_{n\times(N-1)}(A')\,\mathrm{d}A'.$$

We show that O' has orthonormal columns: since $\mathcal{M} \subset \mathbb{S}^{N-1}$, the tangent space $\mathrm{T}_x\mathcal{M}$ is orthogonal to x, which implies $U^Tx = 0$. Furthermore, $e_1^TOU = (U^TO^Te_1)^T = (U^Tx)^T$. It follows that the first row of OU contains only zeros and so the columns of O' must be pairwise orthogonal and of norm one. A standard normal matrix multiplied with a matrix with orthonormal columns is also standard normal, so we have

$$\int_{A'\in\mathbb{R}^{n\times(N-1)}} |\det(A'O')|\, \varphi_{n\times(N-1)}(A')\,\mathrm{d}A' = \int_{M\in\mathbb{R}^{n\times n}} |\det(M)|\varphi_{n\times n}(M)\,\mathrm{d}M.$$

This implies

$$\int_{A\in\mathcal{H}(x)} \varphi_{n\times N}(A)|\det(A|_{\mathrm{T}_x\mathcal{M}})|\,\mathrm{d}A = \frac{\mathbb{E}_{M\sim\varphi_{n\times n}}|\det(M)|}{\sqrt{2\pi}^n}.$$

Finally, we compute $\mathrm{vol}(\mathbb{P}^n) = \frac{1}{2}\mathrm{vol}(\mathbb{S}^n) = \frac{\sqrt{\pi}^{n+1}}{\Gamma(\frac{n+1}{2})}$, and by [66, Thm. 3.2.15], we have $\mathbb{E}_{M\sim\varphi_{n\times n}}\det(M^TM)^{\frac{1}{2}} = \frac{1}{\sqrt{\pi}}\sqrt{2}^n\Gamma(\frac{n+1}{2})$. This finishes the proof. □

Proof of Theorem 4.5.1. We define $\mathcal{I} := \{(A, x) \in \mathbb{R}^{n\times N}\times\mathcal{M} \mid Ax = 0\}$. It is an algebraic subvariety of $\mathbb{R}^{n\times N}\times\mathcal{M}$. One can show that $\mathcal{I}$ is smooth, but for our purposes it suffices to integrate over the dense subset of $\mathcal{I}$ that is obtained by removing potential singularities from $\mathcal{I}$.

Let π_1 and π_2 be the projections from $\mathcal{I}$ to $\mathbb{R}^{n\times N}$ and $\mathcal{M}$, respectively. Applying Theorem 4.3.6 first to π_1 and then to π_2 yields

$$\mathop{\mathbb{E}}_{A\sim\varphi_{n\times N}} \overline{f}(A) = \int_{\mathcal{M}} f(x)\left(\int_{A\in\pi_1(\pi_2^{-1}(x))} \varphi_{n\times N}(A)\frac{\mathrm{NJ}(\pi_2,(A,x))}{\mathrm{NJ}(\pi_1,(A,x))}\,\mathrm{d}A\right)\mathrm{d}x.$$

By [8, Sec. 13.2, Lemma 3], the ratio of normal Jacobians in the integrand equals $|\det(A|_{\mathrm{T}_x\mathcal{M}})|$. We get

$$\begin{aligned}\mathop{\mathbb{E}}_{A\sim\varphi_{n\times N}} \overline{f}(A) &= \int_{\mathcal{M}} f(x) \left(\int_{A\in\pi_1(\pi_2^1(x))} \varphi_{n\times N}(A) |\det(A|_{\mathrm{T}_x\mathcal{M}})| \,\mathrm{d}A \right) \mathrm{d}x \\ &= \frac{1}{\mathrm{vol}(\mathbb{P}^n)} \int_{\mathcal{M}} f(x)\mathrm{d}x;\end{aligned}$$

the second equality by Lemma 4.5.3. This proves the first part.

Let now $Y \in \mathcal{I}$ be the random variable obtained by choosing $A \in \mathbb{R}^{n\times N}$ with distribution $\psi(A) = \frac{\varphi(A)\,\overline{f}(A)}{\mathbb{E}_{A\sim\varphi_{n\times N}}\,\overline{f}(A)}$ and, independently of A, a point $x \in \mathcal{M}\cap\mathcal{L}_A$ with probability $f(x)\overline{f}(A)^{-1}$. Then, by construction, $X = \pi_2(Y)$. Let γ be the density of X. Applying the first part of Corollary 4.3.7 to π_1 and then the second part to π_2, we have

$$\begin{aligned}\gamma(x) &= \int_{(A,x)\in\pi_2^{-1}(x)} \frac{\psi(A) f(x) \mathrm{NJ}(\pi_2,(A,x))}{\overline{f}(A)\mathrm{NJ}(\pi_1,(A,x))} \,\mathrm{d}A \\ &= \frac{f(x)}{\mathbb{E}_\varphi(\overline{f})} \int_{(A,x)\in\pi_2^{-1}(x)} \varphi(A)\,|\det(A|_{\mathrm{T}_x\mathcal{M}})|\,\mathrm{d}A \\ &= \frac{f(x)}{\mathbb{E}_\varphi(\overline{f})} \frac{1}{\mathrm{vol}(\mathbb{P}^n)} \\ &= \frac{f(x)}{\int_{\mathcal{M}} f(x)\,\mathrm{d}x};\end{aligned}$$

the last penultimate equality again by Lemma 4.5.3, and the last equality by the first part of the theorem. This finishes the proof. □

4.6. Previous methods

We now briefly review the established use of kinematic formulae in applications and compare our method to them.

In [9], the authors use a Crofton-type formula for curves to establish a link between a discrete *cut metric* on a grid, which is an object in combinatorial optimization, and an Euclidean metric on $\mathbb{R}^2$. This is then applied to a problem in image segmentation.

In [56], the authors use the Cauchy formula and Crofton formula to compute Minkowski measures (e.g. surface area, perimeter) of discrete binary 2D or 3D pictures given as a grid of white or black pixels. Since the picture is discrete, the set of lines is also appropriately discretized, as well as Crofton's formula itself. So the difference to our method here lies in the discretization. In [57], a more efficient way to evaluate the discretized Crofton's formula is proposed, using run-length encoding.

In [61], the authors use Crofton's formula to approximate the volume of a body $\mathcal{M}$. They use a different sampling method, which goes as follows: (1) Find a compact body $\mathcal{E}$ containing $\mathcal{M}$, of known volume $\mathrm{vol}(\mathcal{E})$, such that the space of lines intersecting $\mathcal{E}$ is approximately the same as the space of lines intersecting $\mathcal{M}$. For example, $\mathcal{E}$ could be a sphere containing $\mathcal{M}$. (2) Sample uniformly from the set of lines that intersect $\mathcal{E}$. (3) Compute the total number of intersection points of all the sampled lines with $\mathcal{E}$ (call it g) and with $\mathcal{M}$ (call it h). (4) Approximate the volume of $\mathcal{M}$ as $\frac{h}{g}\,\mathrm{vol}(\mathcal{E})$.

As the authors write in [61], this method can only give an approximation for the volume of $\mathcal{M}$. Its accuracy depends on the choice of $\mathcal{E}$, i.e. on how well the uniform distribution on the set of lines intersecting $\mathcal{E}$ approximates the same with respect to $\mathcal{M}$. On the other hand, the method presented in this chapter is guaranteed to converge to the true volume given enough samples. We tested both methods on the curve $\mathcal{M}$ from (4.2.1) as well as on the ellipse $\mathcal{M}_1 = \{x \in \mathbb{R}^2 \mid (x/3)^2 + y^2 - 1 = 0\}$, choosing $\mathcal{E}$ to be the centered circle of radius 3. We plotted the results in Figure 4.6.1.

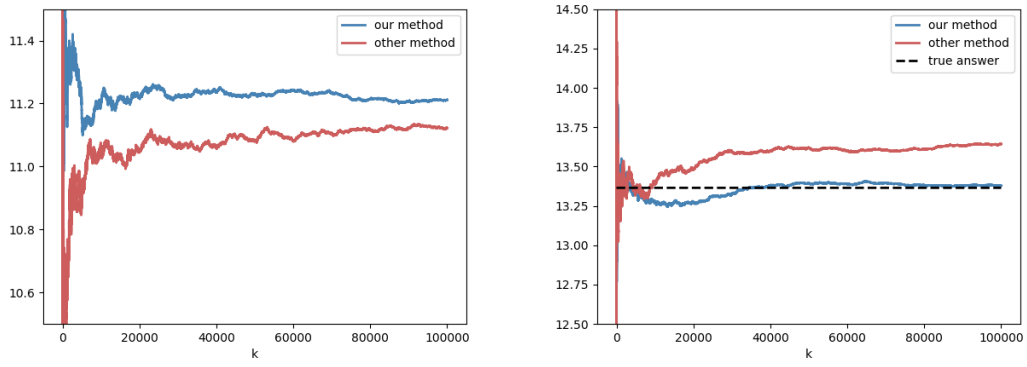

Figure 4.6.1: The plot shows estimates for the volumes of two curves $\mathcal{M}$ obtained from empirical estimates for $1 \leq k \leq 10^5$ samples. On the left, $\mathcal{M}$ is the curve from (4.2.1). On the right, $\mathcal{M}$ is the ellipse $\{x \in \mathbb{R}^2 \mid (x/3)^2 + y^2 - 1 = 0\}$. Its volume is known and shown by the black line. The blue curve shows our method, and the red curve shows the method from [61], where we have used the circle of radius 3 for $\mathcal{E}$. Our method is guaranteed to converge to the true volume of $\mathcal{M}$ for $k \to \infty$, while the other method is not, as exemplified by the plot. Our method seems to converge at least as quickly as the other method, if not slightly faster.

Finally, let us briefly review the references [62, 63]. In these works the authors derive MCMC methods for sampling $\mathcal{M}$ by intersecting it with random subspaces moving according to the *kinematic measure* in $\mathbb{R}^N$. This is related to our discussion from the introduction, where we proposed sampling from $\psi(A, b)$ using MCMC methods. Taking this approach and comparing it to [62, 63] is left for future work.

One feature of the implementation we described is that it generates independent samples from the density $\psi(A, b)$ by rejection sampling, hence independent points $x \in \mathcal{M}$. As we made experiments, we observed some downsides of our method. Namely, our method

becomes slow when the degree of the variety is large in which case it is not easy to find a good κ, and the rejection rate in the sampling process becomes infeasibly large.

But we could also sample from $\psi(A, B)$ using an MCMC method with the goal of improving the rejection rate, at the cost of introducing dependencies between samples. In contrast to the known MCMC methods for nonlinear manifolds, our method would employ MCMC on a flat space. We name using MCMC methods for sampling $\psi(A, B)$ as a possible direction for future research.

5. Mathematics in the sciences

In this expository chapter I describe two applied projects I worked on from the mathematical perspective. In both projects, I helped develop the mathematical theory. Usually, applications do not use contemporary theoretical mathematics. Nevertheless, it is interesting to see how mathematical thinking interfaces with research in other fields. My two example applications both deal with statistical modeling, the first in soil ecology and the second in the more abstract nonparametric statistics.

5.1. A Bayesian network in soil ecology

Soil ecologists are interested in the soil as an ecosystem. In each square centimeter of soil there exist a multitude of fungal and bacterial species. They respond to various environmental pressures and in turn influence the above-ground ecosystem by governing plant growth. The diversity of these species is an important characteristic of the ecosystem. In ecology, it is measured by two numbers that can be calculated from a sample of the soil called *alpha* and *beta diversity*.

The article [43] discusses the influence of various climatic factors on these diversity measures, among others. We try to predict how the soil ecosystem diversity will change in response to global climate change. For this, we consider a set of factors related to climate that could possibly have an influence on biodiversity. Our model sets up a web of causal relationships between the factors that influence alpha diversity. It is represented by the graph in Figure 5.1.1.

The statistical model associated to such a graph is called a *structural equation model*. It is the continuous version of the discrete Bayesian networks introduced in Section 3.2. Each node of the graph corresponds to a real random variable whose name is noted in parentheses in Figure 5.1.1. In the model, we assume that every random variable is a linear expression in the variables pointing to it. For the alpha diversity measure α_{bact}, the model postulates

$$\begin{aligned}\alpha_{\text{bact}} = \lambda_{\text{elev,alpha}}\, I_{\text{elev}} + \lambda_{\text{veg,alpha}}\, I_{\text{veg}} + \lambda_{\text{pV,alpha}}\, M_{\text{p.veg}} + \lambda_{\text{pH,alpha}}\, M_{\text{pH}} \\ + \lambda_{\text{carb,alpha}}\, I_{\text{carb}} + \lambda_{\text{text,alpha}}\, I_{\text{text}}\end{aligned}$$

where the λ are real model parameters.

In our project, we make predictions for α_{bact} based on the projected values for the variables with no incoming edges (the I). These are in turn based on the climate change

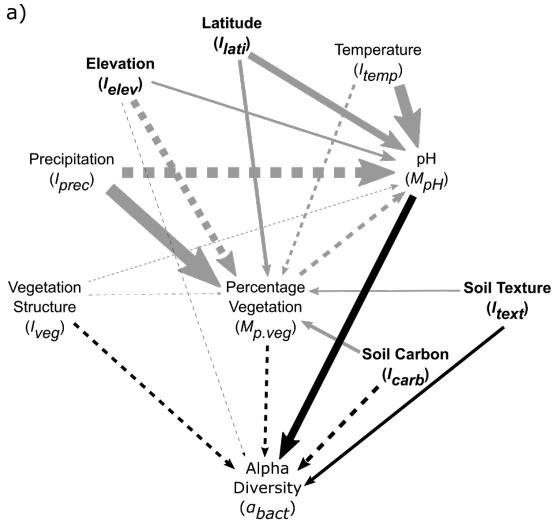

Figure 5.1.1: A Bayesian network that represents possible causal relationships between various environmental factors that influence bacterial alpha diversity. Thicker arrows indicate stronger correlations, dashed arrows indicate negative correlations.

scenarios described in [71]. The first task is to find the parameters λ. This can be accomplished by a chain of multivariate regressions following a topological ordering of the graph. In concrete terms, this means first estimating the parameters associated to the edges pointing to $M_{\text{p.veg}}$, then the ones pointing to M_{pH}, and finally the ones pointing to α_{bact}, respectively using the previously found estimators.

To make predictions based on these parameters, it is desirable to have a single equation with α_{bact} on the left-hand side and a linear expression in the I on the right-hand side. Of course, this is accomplished by substituting the linear expressions for $M_{\text{p.veg}}$ and M_{pH} into the formula for α_{bact} above. The predictions then are made by evaluating the expression for α_{bact} using the values of the I from the climate change scenarios.

An important decision when using structural equation models is what network to use. Which edges should there be and what should their direction be? In our project, we base our model selection on *expert knowledge*, which means using established knowledge about the causal relationships between each variable pair to draw the arrows of the graph. This approach first gives more arrows than the ones depicted. We then eliminate the arrows that after an initial parameter estimation we find to be (near) zero or not statistically significant. We arrive at the final graph depicted in Figure 5.1.1.

Another way for deciding the graph structure is learning it from the data alone. The standard algorithm to do so is called the *PC algorithm*. In general however, the PC algorithm returns multiple candidate graphs that it cannot distinguish from each other based only on the data. Such candidate graphs are said to lie in the same *Markov*

equivalence class. For example, the graphs

$$A \to B \qquad \text{and} \qquad A \leftarrow B$$

belong to the same Markov equivalence class even though they express opposite causal relationships. The PC algorithm also has more fundamental limitations, discussed in detail in [85].

In our project we use a dataset of 231 soil samples from a wide variety of ecosystems around the world. The samples were collected between 2003 and 2015. They contain information on DNA sequencing for bacteria and fungi that in particular determine the values of the diversity measures. The data incorporates many vegetation types, climatic ranges, and soil properties.

Using our structural equation model on the dataset, we predict an increase in alpha diversity and a decrease in beta diversity in response to man-made climate change. Of course, this preliminary prediction has to be corroborated by new data. Samples of soil ecosystems are difficult to collect and process. With our model and predictions about the relationship between the climate and the soil ecosystem, we hope to encourage this necessary field work.

5.2. Polynomials in nonparametric regression

Boundary regression models and their variants are used to estimate the frontier, or boundary, of a data set, beyond which no new data is expected to appear. They are used in the field of *data envelopment analysis* [22] to evaluate the technical efficiency of a production unit. This modeling technique has many applications, for instance in health care, supply chain management, energy systems, and the social sciences. The mathematical foundations required for applying these models are developed in large part by the field of nonparametric statistics.

The basic problem of regression in statistics is to estimate a function $g\colon \mathbb{R}^q \to \mathbb{R}$ based on observations $(X_i, Y_i) \in \mathbb{R}^q \times \mathbb{R}$ under the assumption

$$Y_i = g(X_i) + \varepsilon_i.$$

Here, X_i and ε_i are random variables that may be subjected to various constraints depending on the problem. In [76], we study the *nonparametric boundary regression model* and propose a method for estimating g in that setting.

In the model, the space $\mathbb{R}^q$ is replaced by the bounded $[0,1]^q$. We assume g to be in the Hölder space $C^{\beta^*,\delta}((0,1)^q)$, where $\beta^* \in \mathbb{N}$ and $\delta \in (0,1]$. We further assume the errors ε_i to be i.i.d. and to always satisfy $\varepsilon_i \leq 0$. More precisely, we assume that the cumulative density function F of the errors satisfies

$$(1 - F(y)) = c|y|^\alpha + r(y) \quad \text{for all} \quad y < 0,$$

where $\alpha, c > 0$ and $r(y) = o(|y|^\alpha)$ when $y \nearrow 0$. Figure 5.2.1 depicts an example of such a model where $q = 1$.

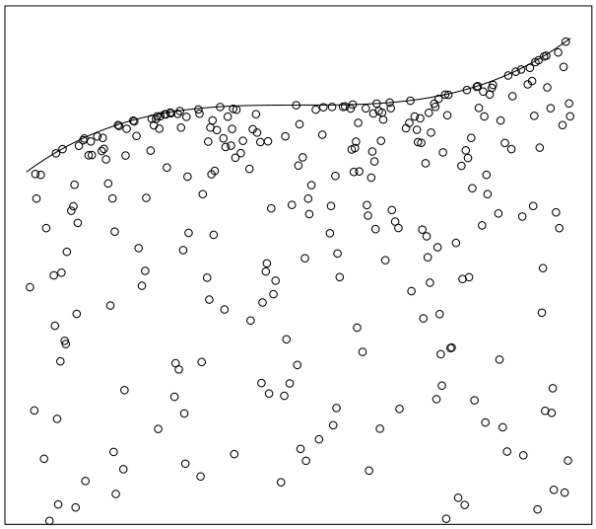

Figure 5.2.1: Example of the samples generated by a univariate boundary regression model. Here we have $g(x) = (x - 0.5)^3 + 2$ and the errors ε are distributed on according to the cumulative density function $y \mapsto \exp(y^{0.5})$ on $(-\infty, 0)$

The article [50] proposes a method to estimate g in the case where the inputs X_i are univariate, that is when $q = 1$. Our contribution is to solve the same problem in the multivariate case. Both solutions use the same general strategy. To define the approximated value $\hat{g}(x)$ of g at the point x, first approximate g by a polynomial P of degree β^* in a neighborhood of x, then set $\hat{g}(x) = P(x)$.

In our article, we define P as the solution to a linear optimization problem where the problem space is the finite dimensional vector space of all polynomials of degree β^*. The reader is referred to the article for more details.

To analyze the convergence rate of the estimator $\hat{g}$, it was important to prove a multivariate version of Lemma 6.1 in [50]. Intuitively, this lemma is about controlling the values of a polynomial function $P\colon [0,1] \to \mathbb{R}$ by considering the restriction of the polynomial to a well-chosen ball.

In [50], the authors prove the lemma by passing to the complex numbers and using the fundamental theorem of algebra, which is not available in multiple dimensions. In [76], we prove the following generalization without passing to the complex numbers.

Lemma 5.2.1. *Let β^* and q be natural numbers. There exist positive real numbers δ, c such that for all polynomial functions $P\colon [0,1]^q \to \mathbb{R}$ of degree β^* with non-negative integral over $[0,1]^q$ there exists a δ-ball $B_\delta \subseteq [0,1]^q$ with respect to the maximum norm*

such that $P \geq 0$ *on* B_δ *and*

$$\inf_{B_\delta} P \geq c \cdot \sup_{[0,1]^q} |P|.$$

Intuitively, Lemma 5.2.1 can be explained as follows. The zero set of a polynomial defined on $[0,1]^q$ has measure zero, so we can find a ball B_δ contained in its complement. With some effort, we can prove that we may fix δ at a value that works for all polynomials of a given degree β^*, so long as we are allowed to move the ball B_δ. Lemma 5.2.1 states the even stronger fact that when the integral of P is non-negative and $|P|$ has supremum one on $[0,1]^q$ we can make the smallest value of P on B_δ greater than a constant that does not depend on P, only on its degree.

The proof of Lemma 5.2.1 rests on the fact that bounded polynomials of supremum one are Lefschetz continuous with Lefschetz constant only depending on their degree. More precisely, for all polynomial functions $P\colon [0,1]^q \to \mathbb{R}$ of degree β^* and all $x, y \in [0,1]^q$ we have

$$|P(x) - P(y)| < L \cdot ||x - y|| \cdot \sup |P|.$$

with $L = 4(\beta^*)^2\sqrt{q}$. For the rest of the proof, the reader is referred to [76].

Using this lemma and the definition of $\hat{g}$, we establish the deterministic part $O(h_n^\beta)$ of the rate of convergence equality

$$\sup_{x \in [0,1]^q} |\hat{g}(x) - g(x)| = O(h_n^\beta) + O_{\mathcal{P}}\left(\left(\frac{\log(n)}{nh_n^q}\right)^{1/\alpha}\right),$$

while we establish the random part $O_{\mathcal{P}}$ by other methods.

In sum, our project is about generalizing a univariate regression technique to a multivariate one. As it turns out, proving that this is sound requires something other than a straight generalization of the proofs. The crucial Lemma 5.2.1 has a satisfyingly elementary proof that still uses theorems about polynomials in an essential way.

5.3. Conclusion

Where to go from here? The two applied projects presented in this chapter are part of much broader research stories. Algebraic geometry can contribute to both.

The structural equation models from Section 5.1 are part of the broader story of causality. The main characters of this story are tools such as graphical models, Bayesian networks, and the staged trees from Section 3.2. These graphical tools generate algebraic statistical models that can be probed for algebraic properties. In particular, staged trees are a versatile new tool for modeling discrete sequences of events. Here are three simple but important questions to advance the theory of staged trees.

(1) Are staged trees exponential families? An affirmative answer would unlock the well-understood theory of exponential families for the study of staged tree models.

(2) Can we reconstruct a Bayesian network from its staged tree? If so, we could more easily translate between the two types of tool.

(3) Can we use staged trees for causal discovery? This would greatly enhance staged trees' capacity for modeling causality.

The boundary regression model from Section 5.2 is part of the broader story of statistical learning. Nowadays, the protagonist of this story is the neural network. Often, the space of functions defined by a neural network is a semialgebraic set. In this setting, or in the tamer setting of nonparametric boundary regression, algebraic geometry can help answer the following questions.

(1) In [76], we use a linear optimization problem to define our estimator for the boundary regression model. When does a unique solution exist?

(2) Can we fully understand the critical locus of the loss function of a neural network? This would help explain why its optimization is not hindered by local minima.

Researchers in Algebraic Statistics will continue developing bridges between theory and application in the coming years. My participation in the field will in particular involve studying the ML degree of Gaussian statistical models. Here, a promising *Ansatz* is to consider the moduli space of complete quadrics. This classical object from algebraic geometry contains the set of Gaussian models as a subset. Here are two natural questions in this line of research.

(1) Can we classify all Gaussian models with ML degree one? This is desirable for the same reasons as the discrete version of this question raised in Section 3.5.

(2) Can we solve the ML degree problem for linear covariance models by turning it into an intersection theory problem in the space of complete quadrics?

In the two projects presented in this chapter, I did not use tools from algebraic geometry. Instead, my support role as a mathematician was to help build the mathematical theory and ensure it stands on solid ground. I believe this is as valuable to science as basic mathematical research. The future will bring many fruitful collaborations like these to the researcher who seeks them out.

A. Appendix: A lemma on the ML degree

The following lemma implies that the maximum likelihood degree of a discrete or continuous parametrized algebraic model is well-defined. We use the language of schemes for precision.

Lemma A.1. *Let $\mathcal{M} \subseteq \mathcal{N}$ be an algebraic model with rational score equations. Let $\mathcal{N}_{\mathbb{C}}$ be integral. For general $y \in \mathcal{N}_{\mathbb{C}}$, the number of solutions to the score equations is independent of y.*

Proof. Let x and y denote the coordinates of $\mathcal{M}_{\mathbb{C}}$ and $\mathcal{N}_{\mathbb{C}}$, respectively. Write the score equations as $\ell_j = p_j/q_j$ with $p_j, q_j \in \mathbb{R}[x, y, \lambda]$. Consider the quasi-affine variety $V \subseteq \mathcal{M}_{\mathbb{C}} \times \mathcal{N}_{\mathbb{C}} \times \mathbb{C}^r$ defined as the set of triples (x, y, λ) such that $p(x, y, \lambda) = 0$ and $q(x, y, \lambda) \neq 0$. Let $\alpha : V \to \mathcal{N}_{\mathbb{C}}$ denote the second projection. The number of solutions to the score equations for general y is the number of elements, counted with multiplicities, of the fiber $V_y = \alpha^{-1}(y)$. More precisely, it is the number $d_y \coloneqq \sum_{x \in \alpha^{-1}(y)} \dim_{\mathbb{C}}(\mathcal{O}_{V_y, x})$.

A general fiber V_y of α is of the same dimension as the generic fiber V_η [81, Lemma 05F7]. If the latter is positive, then a general fiber has infinitely many points and we are done. So, we assume that the generic fiber is zero-dimensional and thus, since α is of finite type, that it contains a finite number of points. By [81, Lemma 02NW], we may assume that α is finite. By generic flatness [81, Proposition 052A], we may assume that α is flat. Finally, since $\mathcal{N}_{\mathbb{C}}$ is connected the map α has constant degree and we may apply [41, Prop. 12.21] to obtain $\deg \alpha = d_y$ for all y. □

Bibliography

[1] Theodore W. Anderson. Asymptotically efficient estimation of covariance matrices with linear structure. *The Annals of Statistics* **1** (1973), pp. 135–141.

[2] Theodore W. Anderson. Estimation of covariance matrices which are linear combinations or whose inverses are linear combinations of given matrices. *Essays in Probability and Statistics*, pp. 1–24. UNC Press, 1970.

[3] Nihat Ay, Jürgen Jost, Hông Vân Lê, and Lorenz Schwachhöfer. *Information Geometry*. Springer-Verlag New York, 2017.

[4] Ulrich Bauer. Ripser: efficient computation of Vietoris-Rips peristence barcodes. `ArXiv:1908.02518`, 2019.

[5] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. Julia: a fresh approach to numerical computing. *SIAM Review* **59** (2017), pp. 65–98.

[6] Jacob Bien and Robert J. Tibshirani. Sparse estimation of a covariance matrix. *Biometrika* **98** (2011), pp. 807–820.

[7] Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas, editors. *Semidefinite Optimization and Convex Algebraic Geometry*. SIAM, 2013.

[8] Lenore Blum, Felipe Cucker, Michael Shub, and Steve Smale. *Complexity and Real Computation*. Springer-Verlag New York, 1998.

[9] Yuri Boykov and Vladimir Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. *Proceedings Ninth IEEE International Conference on Computer Vision*, pp. 26–33, 2003.

[10] Paul Breiding, Sara Kališnik, Bernd Sturmfels, and Madeleine Weinstein. Learning algebraic varieties from samples. *Revista Matemática Complutense* **31** (2018), pp. 545–593.

[11] Paul Breiding and Orlando Marigliano. Random points on an algebraic manifold. *SIAM Journal on Mathematics of Data Science* (2020, to appear).

[12] Paul Breiding and Sascha Timme. HomotopyContinuation.jl: a package for homotopy continuation in Julia. *International Congress on Mathematical Software*, pp. 458–465, 2018.

[13] William H. Brown and Lawrence S. Brown. *Organic Chemistry, Enhanced Edition*. Brooks Cole, 2010.

[14] Peter Bürgisser and Felipe Cucker. *Condition*. Springer-Verlag Berlin Heidelberg, 2013.

[15] Atul J. Butte, Pablo Tamayo, Donna Slonim, Todd R. Golub, and Isaac S. Kohane. Discovering functional relationships between RNA expression and chemotherapeutic susceptibility using relevance networks. *Proceedings of the National Academy of Sciences* **97** (2000), pp. 12182–12186.

[16] George Casella and Roger L. Berger. *Statistical Inference*. Thomson Learning, 2002.

[17] Fabrizio Catanese, Serkan Hoşten, Amit Khetan, and Bernd Sturmfels. The maximum likelihood degree. *American Journal of Mathematics* **128** (2006), pp. 671–697.

[18] Sanjay Chaudhuri, Mathias Drton, and Thomas S. Richardson. Estimation of a covariance matrix with zeros. *Biometrika* **94** (2007), pp. 199–216.

[19] Patrick Clarke and David A. Cox. Moment maps, strict linear precision, and maximum likelihood degree one. `ArXiv:1810.03672`, 2018.

[20] Rodrigo A. Collazo, Christiane Görgen, and Jim Q. Smith. *Chain Event Graphs*. CRC Press, 2018

[21] Jane Ivy Coons, Orlando Marigliano, and Michael Ruddy. Maximum likelihood degree of the small linear gaussian covariance model. `ArXiv:1909.04553`, 2019.

[22] William W. Cooper, Lawrence M. Seiford, and Joe Zhu. *Handbook on Data Envelopment Analysis*. Springer US, 2004.

[23] David A. Cox, John B. Little, and Donal O'Shea. *Using Algebraic Geometry*. Springer-Verlag New York, 1998.

[24] David A. Cox, John B. Little, and Donal O'Shea. *Ideals, Varieties, and Algorithms*. Springer-Verlag New York, 1992.

[25] David A. Cox, John B. Little, and Henry K. Schenk. *Toric Varieties*. American Mathematical Society, 2011.

[26] Persi Diaconis and Bernd Sturmfels. Algebraic algorithms for sampling from conditional distributions. *The Annals of Statistics* **26** (1998), pp. 363–397.

[27] Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on Algebraic Statistics*. Birkhäuser Basel, 2008.

[28] Eliana Duarte and Christiane Görgen. Equations defining probability tree models. *Journal of Symbolic Computation* **99** (2020), pp. 127–146.

[29] Eliana Duarte, Orlando Marigliano, and Bernd Sturmfels. Discrete statistical models with rational maximum likelihood estimator. *Bernoulli* (2020, to appear).

[30] Emilie Dufresne, Parker B. Edwards, Heather A. Harrington, and Jonathan D. Hauenstein. Sampling real algebraic varieties for topological data analysis. *18th IEEE International Conference on Machine Learning and Applications*, pp. 1531–1536, 2019.

[31] Parker Edwards. Personal communication.

[32] Leonhard Euler. *Institutiones Calculi Differentialis*. Academia Imperialis Scientiarum Petropolitanae, 1755

[33] Michael Farber and Dirk Schütz. Homology of planar polygon spaces. *Geometriae Dedicata* **125** (2007), pp. 75–92.

[34] Joseph Felsenstein. Maximum-likelihood estimation of evolutionary trees from continuous characters. *American Journal of Human Genetics* **25** (1973), pp. 471–492.

[35] Gerd Fischer. *Plane Algebraic Curves*. American Mathematical Society, 2001.

[36] William Fulton. *Intersection Theory*. Springer-Verlag New York, 1998.

[37] William Fulton. *Algebraic Curves*. Addison-Wesley, 1989.

[38] Luis David Garcia-Puente and Frank Sottile. Linear precision for parametric patches. *Advances in Computational Mathematics* **33** (2010), pp. 191–214.

[39] Izrail' M. Gel'fand, Mikhail M. Kapranov, and Andrey V. Zelevinsky. *Discriminants, Resultants, and Multidimensional Determinants*. Birkhäuser Boston, 1994

[40] Christiane Görgen and Jim Q. Smith. Equivalence classes of staged trees. *Bernoulli* **24** (2018), pp. 2676–2692.

[41] Ulrich Görtz and Torsten Wedhorn. *Algebraic Geometry I*. Vieweg+Teubner, 2010

[42] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. `faculty.math.illinois.edu/Macaulay2/`.

[43] Carlos Guerra, Manuel Delgado Baquerizo, Eliana Duarte, Fernando T. Maestre, Nico Eisenhauer, Orlando Marigliano, and Christiane Görgen. Global projections of the soil microbiome. 2019, submitted.

[44] Joe Harris. *Algebraic Geometry*. Springer-Verlag New York, 1995.

[45] Ralph Howard. *The Kinematic Formula in Riemannian Homogeneous Spaces*. American Mathematical Society, 1993.

[46] June Huh. Varieties with maximum likelihood degree one. *Journal of Algebraic Statistics* **5** (2014), pp. 1–17. `repository.iit.edu/islandora/object/islandora%3A1007817`.

[47] June Huh and Bernd Sturmfels. Likelihood geometry. *Combinatorial Algebraic Geometry*, pages 63–117. Springer Cham, 2014.

[48] John D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science Engineering* **9** (2007), pp. 90–95.

[49] Huan Jiang. *Modeling Trait Evolutionary Processes with More Than One Gene*. PhD thesis, University of New Mexico, 2017.

[50] Moritz Jirak, Alexander Meister, and Markus Reiß. Adaptive function estimation in nonparametric regression with one-sided errors. *The Annals of Statistics* **42** (2014), pp. 1970–2002.

[51] Mikhail M. Kapranov. A characterization of A-discriminantal hypersurfaces in terms of the logarithmic Gauss map. *Mathematische Annalen* **290** (1991), pp. 277–285.

[52] Kathlén Kohn, Boris Shapiro, and Bernd Sturmfels. Moment varieties of measures on polytopes. `ArXiv:1807.10258`, 2018.

[53] Rimvydas Krasauskas. Toric surface patches. *Advances in Computational Mathematics* **17** (2002), pp. 89–113.

[54] Pierre Lairez. Rigid continuation paths I. Quasilinear average complexity for solving polynomial systems. `ArXiv:1711.03420`, 2017.

[55] Steffen L. Lauritzen. *Graphical Models*. Oxford University Press, 1996.

[56] David Legland, Kiên Kiêu, and Marie-Françoise Devaux. Computation of Minkowski measures on 2D and 3D binary images. *Image Analysis & Stereology* **26** (2007), pp. 83–92.

[57] Gaëtan Lehmann and David Legland. Efficient N-dimensional surface estimation using Crofton formula and run-length encoding. `hdl.handle.net/10380/3342`, 2012.

[58] Kurt Leichtweiss. Zur Riemannschen Geometrie in Grassmannschen Mannigfaltigkeiten. *Mathematische Zeitschrift* **76** (1961), pp. 334–366.

[59] Tony Lelièvre, Mathias Rousset, and Gabriel Stoltz. *Free Energy Computations*. Imperial College Press London, 2010.

[60] Tony Lelièvre, Mathias Rousset, and Gabriel Stoltz. Hybrid Monte Carlo methods for sampling probability measures on submanifolds. *Numerische Mathematik* **143** (2019), pp. 379–421.

[61] Xueqing Li, Wenping Wang, Ralph R. Martin, and Adrian Bowyer. Using low-discrepancy sequences and the crofton formula to compute surface areas of geometric models. *Computer-Aided Design* **35** (2003), pp. 771–782.

[62] Oren Mangoubi. *Integral Geometry, Hamiltonian Dynamics, and Markov Chain Monte Carlo*. PhD thesis, Massachusetts Institute of Technology, 2016.

[63] Oren Mangoubi and Alan Edelman. Integral geometry for Markov chain Monte Carlo: overcoming the curse of search-subspace dimensionality. `ArXiv:1503.03626`, 2015.

[64] *Maple User Manual*. Maplesoft, 2019.

[65] Mateusz Michałek, Bernd Sturmfels, Caroline Uhler, and Piotr Zwiernik. Exponential varieties. *Proceedings of the London Mathematical Society* **112** (2016), pp. 27–56.

[66] Robb J. Muirhead. *Aspects of Multivariate Statistical Theory*. Wiley New Jersey, 1982.

[67] Donna J. Nelson and Christopher N. Brammer. Toward consistent terminology for cyclohexane conformers in introductory organic chemistry. *Journal of Chemical Education* **88** (2011), pp. 292–294.

[68] Lior Pachter and Bernd Sturmfels, editors. *Algebraic Statistics for Computational Biology*. Cambridge University Press, 2005.

[69] Giovanni Pistone, Eva Riccomagno, and Henry P. Wynn. *Algebraic Statistics*. Chapman & Hall/CRC, 2001.

[70] Vladimir L. Popov and Ernest B. Vinberg. Invariant Theory. *Algebraic Geometry IV*, pp. 123–278. Springer Berlin, 1994.

[71] Alexander Popp et al. Land-use futures in the shared socio-economic pathways. *Global Environmental Change* **42** (2017), pp. 331–345.

[72] Christian P. Robert and George Casella. *Monte Carlo Statistical Methods*. Springer-Verlag New York, 1999.

[73] Adam J. Rothman, Elizaveta Levina, and Ji Zhu. A new approach to Cholesky-based covariance regularization in high dimensions. *Biometrika* **97** (2010), pp. 539–550.

[74] John W. Rutter. *Geometry of Curves*. Chapman & Hall/CRC, 2000.

[75] Luis A. Santaló. *Integral Geometry and Geometric Probability*. Cambridge University Press, 2004.

[76] Leonie Selk, Charles Tillier, and Orlando Marigliano. Multivariate boundary regression models. `ArXiv:2004.09881`, 2020.

[77] Igor R. Shafarevich. *Basic Algebraic Geometry 1*. Springer Heidelberg, 2013.

[78] Tomi Silander and Tze-Yun Leong. A dynamic programming algorithm for learning chain event graphs. *Discovery Science*, pp. 201–216. Springer Berlin Heidelberg, 2013.

[79] Jan Skowron and Andrew Gould. General complex polynomial root solver and its further optimization for binary microlenses. `ArXiv:1203.1034`, 2012.

[80] Jim Q. Smith and Paul E. Anderson. Conditional independence and chain event graphs. *Artificial Intelligence* **172** (2008), pp. 42–68.

[81] The Stacks Project Authors. *Stacks Project*. `https://stacks.math.columbia.edu`, 2018.

[82] Bernd Sturmfels, Sascha Timme, and Piotr Zwiernik. Estimating linear covariance models with numerical nonlinear algebra. `ArXiv:1909.00566`, 2019.

[83] Seth Sullivant. *Algebraic Statistics*. American Mathematical Society, 2018.

[84] Sasha Timme. LinearCovarianceModels. `juliapackages.com/p/linearcovariancemodels`, 2019.

[85] Caroline Uhler, Garvesh Raskutti, Peter Bühlmann, and Bin Yu. Geometry of the faithfulness assumption in causal inference. *The Annals of Statistics* **41** (2013), pp. 436–463.

[86] Hassler Whitney. Elementary structure of real algebraic varieties. *Annals of Mathematics* **66** (1957), pp. 545–556.

[87] Wei Biao Wu and Han Xiao. Covariance matrix estimation in time series. *Time Series Analysis: Methods and Applications*, pp. 187–209. Elsevier, 2012.

[88] Emilio Zappa, Miranda Holmes-Cerfon, and Jonathan Goodman. Monte Carlo on manifolds: sampling densities and integrating functions. *Communications on Pure and Applied Mathematics* **71** (2018), pp. 2609–2647.

[89] Piotr Zwiernik, Caroline Uhler, and Donald Richards. Maximum likelihood estimation for linear Gaussian covariance models. *Journal of the Royal Statistical Society* **B 79** (2017), pp. 1269–1292.