

Quality Improvement for Chronic Cases in Pulmonary Disease

Joseph E. Williams

Dept. of Mathematics
American Mathematics Society
United States

Numerical nonlinear algebra is applied to maximum likelihood estimation for Gaussian models defined by linear constraints on the covariance matrix. We examine the generic case as well as special models (e.g., Toeplitz, sparse, trees) that are of interest in statistics. We study the maximum likelihood degree and its dual analogue, and we introduce a new software package `LinearCovarianceModels.jl` for solving the score equations. All local maxima can thus be computed reliably. In addition we identify several scenarios for which the estimator is a rational function.

1. Introduction

In many statistical applications, the covariance matrix Σ has a special structure. A natural setting is that one imposes linear constraints on Σ or its inverse Σ^{-1} . Here we study models for Gaussians whose covariance matrix Σ lies in a given linear space. Such linear Gaussian covariance models were introduced by [1]. He was motivated by the Toeplitz structure of Σ in time series analysis. Recent applications of such models include repeated time series, longitudinal data, and a range of engineering problems [26]. Other occurrences are Brownian motion tree models [29], as well as pairwise independence models, where some entries of Σ are set to zero.

The literature on estimating a covariance matrix is extremely rich. Its development has been particularly dynamic in high-dimensional statistics under a sparsity assumption on Σ or its inverse; see [12] for an overview. Although the sample covariance matrix is known to have poor statistical properties, for many Gaussian models the maximum likelihood estimator (MLE) remains an important reference point.

Maximum likelihood estimation for linear covariance models is a nonlinear algebraic optimization problem over a spectrahedral cone, namely the convex cone of positive definite matrices Σ that satisfy the linear constraints of interest. The objective function is not convex and can have multiple local maxima. Yet, if the sample size is large relative to the dimension, then the problem is essentially convex. This was shown in [32]. In general, however, the MLE problem is poorly understood, and there is a need for accurate methods that reliably identify all local maxima.

Nonlinear algebra furnishes such a method, namely solving the score equations [30, §7.1] using numerical homotopy continuation [27]. This is guaranteed to find all critical points of the likelihood function and hence all local maxima. A key step is the knowledge of the maximum likelihood degree

(ML degree). This is the number of complex critical points. The ML degree of a linear covariance model is an invariant of a linear space of symmetric matrices which is of interest in its own right.

Our presentation is organized as follows. In [Section 2](#) we introduce various models to be studied, ranging from generic linear equations to colored graph models. In [Section 3](#) we discuss the maximum likelihood estimator as well as the dual maximum likelihood estimator. Starting from [\[30, Proposition 7.1.10\]](#), we derive a convenient form of the score equations. The natural point of entry for an algebraic geometer is the study of generic linear constraints. This is our topic in [Section 4](#). We compute a range of ML degrees, and we compare them to the dual degrees in [\[28, §2.2\]](#).

In [Section 5](#) we present our software `LinearCovarianceModels.jl` [\[31\]](#). This is written in Julia and is easy to use. It computes the ML degree and the dual ML degree for a given subspace \mathcal{L} , and it determines all complex critical points for a given sample covariance matrix S . Among these, it identifies the real and positive definite solutions, and it then selects those that are local maxima. The package rests on the software `HomotopyContinuation.jl` of [\[2\]](#).

[Section 6](#) discusses instances where the likelihood function has multiple local maxima. This is meant to underscore the strength of our approach. We then turn to models where the maximum is unique and the MLE is a rational function. In [Section 7](#) we examine Brownian motion tree models. Here the linear constraints are determined by a rooted phylogenetic tree. We study the ML degree and dual ML degree. We show that the latter equals one for binary trees, and we derive the explicit rational formula for their MLE. A census of these degrees is found in [Table 5](#).

2. Models

Let \mathbb{S}^n be the $\binom{n+1}{2}$ -dimensional real vector space of $n \times n$ symmetric matrices $\Sigma = (\sigma_{ij})$. The subset \mathbb{S}_+^n of positive definite matrices is a full-dimensional open convex cone. Consider any linear subspace \mathcal{L} of \mathbb{S}^n whose intersection with \mathbb{S}_+^n is nonempty. Then $\mathbb{S}_+^n \cap \mathcal{L}$ is a relatively open convex cone. In optimization, where one uses the closure, this is known as a *spectrahedral cone*. In statistics, the intersection $\mathbb{S}_+^n \cap \mathcal{L}$ is a *linear covariance model*. These are the models we study in this paper. In what follows we discuss various families of linear spaces \mathcal{L} that are of interest to us.

Generic linear constraints. Fix a positive integer $m \leq \binom{n+1}{2}$, and suppose that \mathcal{L} is a generic linear subspace of \mathbb{S}^n . Here “generic” is meant in the sense of algebraic geometry; i.e., \mathcal{L} is a point in the Grassmannian that lies outside a certain algebraic hypersurface. This hypersurface has measure zero, so a random subspace will be generic with probability one. For a geometer, it is natural to begin with the generic case, since its complexity controls the complexity of any special family of linear spaces. In particular, the ML degree for a generic \mathcal{L} depends only on m and n , and this furnishes an upper bound for the ML degree of the special families below.

Diagonal covariance matrices. Here we take $m \leq n$, and we assume that \mathcal{L} is a linear space that consists of diagonal matrices. Restricting to covariance matrices that are diagonal is natural when modeling independent Gaussians. We use the term *generic diagonal model* when \mathcal{L} is a generic point in the $(n - m)m$ -dimensional Grassmannian of m -dimensional subspaces inside the diagonal $n \times n$ matrices.

Brownian motion tree models. A tree is a connected graph with no cycles. A rooted tree is obtained by fixing a vertex, called the root, and directing all edges away from the root. Fix a rooted tree T with n leaves. Every vertex v of T defines a *clade*, namely the set of leaves that are descendants of v . For the Brownian motion tree model on T , the space \mathcal{L} is spanned by the rank-one matrices $e_A e_A^T$, where $e_A \in \{0, 1\}^n$ is the indicator vector of A . Hence, if \mathcal{C} is the set of all clades of T , then

$$\Sigma = \sum_{A \in \mathcal{C}} \theta_A e_A e_A^T, \quad \text{where } \theta_A \text{ are model parameters.} \tag{1}$$

The linear equations for the subspace \mathcal{L} are $\sigma_{ij} = \sigma_{kl}$ whenever the least common ancestors $\text{lca}(i, j)$ and $\text{lca}(k, l)$ agree in the tree T . Assuming $\theta_A \geq 0$, the union of the models for all trees T is characterized by the ultrametric condition $\sigma_{ij} \geq \min\{\sigma_{ik}, \sigma_{jk}\} \geq 0$. Matrices of this form also play an important role in hierarchical clustering [15, §14.3.12], phylogenetics [13], and random walks on graphs [9].

Maximum likelihood estimation for this class of models is generally complicated, but recently there has been progress [16; 29] on exploiting the nice structure of the matrices Σ above. In Section 7 we study computational aspects of the MLE and, more importantly, provide a significant advance by considering the dual MLE.

Covariance graph models. We consider models \mathcal{L} that arise from imposing zero restrictions on entries of Σ . This was studied in [5; 10]. This is similar to Gaussian graphical models where zero restrictions are placed on the inverse Σ^{-1} . We encode the sparsity structure with a graph whose edges correspond to nonzero off-diagonal entries of Σ . Zero entries in Σ correspond to pairwise marginal independences. These arise in statistical modeling in the context of causal inference [8]. Models with zero restrictions on the covariance matrix are known as covariance graph models. Maximum likelihood in these Gaussian models can be carried out using iterative conditional fitting [5; 10], which is implemented in the `ggm` package in R [22].

Toeplitz matrices. Suppose $X = (X_1, \dots, X_n)$ follows the autoregressive model of order 1, that is, $X_t = \rho X_{t-1} + \epsilon_t$, where $\rho \in \mathbb{R}$ and $\epsilon_t \sim N(0, \sigma)$ for some σ . Assume that the ϵ_t are mutually uncorrelated. Then $\text{cov}(X_t, X_{t-k}) = \rho^k$, and hence Σ is a Toeplitz matrix. More generally, covariance matrices from stationary time series are Toeplitz. Multichannel and multidimensional processes have covariance matrices of block Toeplitz form [3; 24]. Similarly, if X follows the moving average process of order q , then $\text{cov}(X_t, X_{t-k}) = \gamma_k$ if $k \leq q$ and is zero otherwise; see, for example, [14, §3.3]. Thus, in time series analysis, we encounter matrices like

$$\begin{bmatrix} \gamma_0 & \gamma_1 & \gamma_2 & \gamma_3 & \gamma_4 \\ \gamma_1 & \gamma_0 & \gamma_1 & \gamma_2 & \gamma_3 \\ \gamma_2 & \gamma_1 & \gamma_0 & \gamma_1 & \gamma_2 \\ \gamma_3 & \gamma_2 & \gamma_1 & \gamma_0 & \gamma_1 \\ \gamma_4 & \gamma_3 & \gamma_2 & \gamma_1 & \gamma_0 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} \gamma_0 & \gamma_1 & 0 & 0 & 0 \\ \gamma_1 & \gamma_0 & \gamma_1 & 0 & 0 \\ 0 & \gamma_1 & \gamma_0 & \gamma_1 & 0 \\ 0 & 0 & \gamma_1 & \gamma_0 & \gamma_1 \\ 0 & 0 & 0 & \gamma_1 & \gamma_0 \end{bmatrix}. \tag{2}$$

We found that the ML degree for such models is surprisingly low. This means that nonlinear algebra can reliably estimate Toeplitz matrices that are fairly large.

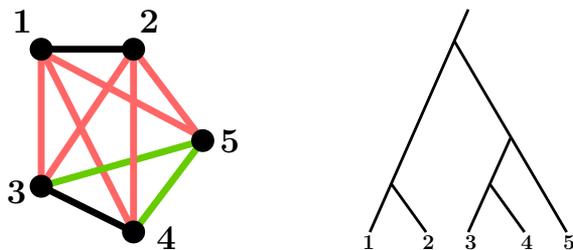


Figure 1. A covariance graph model with edge symmetries and the rooted tree for the corresponding Brownian motion tree model.

Colored covariance graph models. A generalization of covariance graph models is obtained by following [17], which introduces graphical models with vertex and edge symmetries. Models of this type also generalize the Toeplitz matrices and the Brownian motion tree models. Following the standard convention, we use the same colors for edges or vertices when the corresponding entries of Σ are equal. The black color is considered neutral and encodes no restrictions.

The Brownian motion tree model corresponds to a colored model over the complete graph, where edge symmetries are encoded by the tree; see Figure 1. Also, both matrices in (2) represent covariance graph models with edge and vertex symmetries.

3. Maximum likelihood estimator and its dual

Now that we have seen motivating examples, we formally define the MLE problem for a linear covariance model \mathcal{L} . Suppose we observe a random sample $X^{(1)}, \dots, X^{(N)}$ in \mathbb{R}^n from $N_n(0, \Sigma)$. The sample covariance matrix is $S = (1/N) \sum_{i=1}^N X^{(i)} X^{(i)T}$. The matrix S is positive semidefinite. Our aim is to maximize the function

$$\ell(\Sigma) = \log \det \Sigma^{-1} - \text{tr}(S\Sigma^{-1}) \quad \text{subject to } \Sigma \in \mathcal{L}. \tag{3}$$

Following [30, Proposition 7.1.10], this equals the log-likelihood function times $N/2$.

We fix the standard inner product $\langle A, B \rangle = \text{tr}(AB)$ on the space \mathbb{S}^n of symmetric matrices. The orthogonal complement \mathcal{L}^\perp to a subspace $\mathcal{L} \subset \mathbb{S}^n$ is defined as usual.

Proposition 3.1. *Finding all the critical points of the log-likelihood function $\ell(\Sigma)$ amounts to solving the following system of linear and quadratic equations in $2 \cdot \binom{n+1}{2}$ unknowns:*

$$\Sigma \in \mathcal{L}, \quad K\Sigma = I_n, \quad KSK - K \in \mathcal{L}^\perp. \tag{4}$$

Proof. The matrix Σ is a critical point ℓ if and only if, for every $U \in \mathcal{L}$, the derivative of ℓ at Σ in the direction U vanishes. This directional derivative equals

$$-\text{tr}(\Sigma^{-1}U) + \text{tr}(S\Sigma^{-1}U\Sigma^{-1}).$$

This formula follows by multivariate calculus from two facts: (i) the derivative of the matrix mapping $\Sigma \mapsto \Sigma^{-1}$ is the linear transformation $U \mapsto \Sigma^{-1}U\Sigma^{-1}$; (ii) the derivative of the function $\Sigma \mapsto \log \det \Sigma$ is the linear functional $U \mapsto \text{tr}(\Sigma^{-1}U)$.

Using the identity $K = \Sigma^{-1}$, vanishing of the directional derivative is equivalent to

$$-\langle K, U \rangle + \langle KSK, U \rangle = 0.$$

The condition $\langle KSK - K, U \rangle = 0$ for all $U \in \mathcal{L}$ is equivalent to $KSK - K \in \mathcal{L}^\perp$. □

Example 3.2 (3×3 Toeplitz matrices). Let \mathcal{L} be the space of Toeplitz matrices

$$\Sigma = \begin{bmatrix} \gamma_0 & \gamma_1 & \gamma_2 \\ \gamma_1 & \gamma_0 & \gamma_1 \\ \gamma_2 & \gamma_1 & \gamma_0 \end{bmatrix}.$$

This space has dimension 3 in $\mathbb{S}^3 \simeq \mathbb{R}^6$. Fix a sample covariance matrix $S = (s_{ij})$ with real entries. We need to solve the system (4). This consists of $3 + 9 + 3 = 15$ equations in $6 + 6 = 12$ unknowns, namely the entries of the covariance matrix $\Sigma = (\sigma_{ij})$ and its inverse $K = (k_{ij})$. The condition $\Sigma \in \mathcal{L}$ gives three linear polynomials

$$\sigma_{11} - \sigma_{33}, \quad \sigma_{12} - \sigma_{23}, \quad \sigma_{22} - \sigma_{33}.$$

The condition $K\Sigma = I_3$ translates into nine bilinear polynomials

$$\begin{array}{lll} \sigma_{11}k_{11} + \sigma_{12}k_{12} + \sigma_{13}k_{13} - 1, & \sigma_{12}k_{11} + \sigma_{22}k_{12} + \sigma_{23}k_{13}, & \sigma_{13}k_{11} + \sigma_{23}k_{12} + \sigma_{33}k_{13}, \\ \sigma_{11}k_{12} + \sigma_{12}k_{22} + \sigma_{13}k_{23}, & \sigma_{12}k_{12} + \sigma_{22}k_{22} + \sigma_{23}k_{23} - 1, & \sigma_{13}k_{12} + \sigma_{23}k_{22} + \sigma_{33}k_{23}, \\ \sigma_{11}k_{13} + \sigma_{12}k_{23} + \sigma_{13}k_{33}, & \sigma_{12}k_{13} + \sigma_{22}k_{23} + \sigma_{23}k_{33}, & \sigma_{13}k_{13} + \sigma_{23}k_{23} + \sigma_{33}k_{33} - 1. \end{array}$$

Finally, the condition $KSK - K \in \mathcal{L}^\perp$ translates into three quadratic polynomials

$$\begin{aligned} & k_{11}^2s_{11} + k_{12}^2s_{11} + k_{13}^2s_{11} + 2k_{11}k_{12}s_{12} + 2k_{12}k_{22}s_{12} + 2k_{13}k_{23}s_{12} + 2k_{11}k_{13}s_{13} \\ & + 2k_{12}k_{23}s_{13} + 2k_{13}k_{33}s_{13} + k_{12}^2s_{22} + k_{22}^2s_{22} + k_{23}^2s_{22} + 2k_{12}k_{13}s_{23} + 2k_{22}k_{23}s_{23} \\ & + 2k_{23}k_{33}s_{23} + k_{13}^2s_{33} + k_{23}^2s_{33} + k_{33}^2s_{33} - k_{11} - k_{22} - k_{33}, \\ & k_{23}s_{13} + k_{12}k_{33}s_{13} + k_{12}k_{22}s_{22} + k_{22}k_{23}s_{22} + k_{13}k_{22}s_{23} + k_{12}k_{23}s_{23} \\ & + k_{23}^2s_{23} + k_{22}k_{33}s_{23} + k_{13}k_{23}s_{33} + k_{23}k_{33}s_{33} - k_{12} - k_{23}, \\ & k_{11}k_{13}s_{11} + k_{12}k_{13}s_{12} + k_{11}k_{23}s_{12} + k_{13}^2s_{13} + k_{11}k_{33}s_{13} + k_{12}k_{23}s_{22} + k_{13}k_{23}s_{23} + k_{12}k_{33}s_{23} + k_{13}k_{33}s_{33} - k_{13}. \end{aligned}$$

The zero set of these 15 polynomials in 12 unknowns consists of three points $(\widehat{\Sigma}, \widehat{K})$. We present a concrete instance with multiple local solutions:

$$S = \begin{bmatrix} 4/5 & -9/5 & -1/25 \\ -9/5 & 79/16 & 25/24 \\ -1/25 & 25/24 & 17/16 \end{bmatrix} \approx \begin{bmatrix} 0.8000 & -1.8000 & -0.0400 \\ -1.8000 & 4.9375 & 1.0417 \\ -0.0400 & 1.0417 & 1.0625 \end{bmatrix}. \quad (5)$$

For this sample covariance matrix all three critical points are real and positive definite. The three Toeplitz matrices that solve the score equations for this S are:

$\hat{\gamma}_0$	$\hat{\gamma}_1$	$\hat{\gamma}_2$	log-likelihood value	
2.52783	-0.215929	-1.45229	-5.35	global maximum
2.39038	-0.286009	0.949965	-5.41	local maximum
2.28596	-0.256394	0.422321	-5.42	saddle point

So, even in this tiny example, our optimization problem has multiple local maxima in the cone \mathbb{S}_+^3 . A numerical study of this phenomenon will be presented in [Section 6](#).

In this paper we also consider the *dual* maximum likelihood estimator as a more computationally efficient alternative. Dual estimation is based on the maximization of a dual likelihood function. In the Gaussian case this is motivated by interchanging the role of the parameter matrix Σ and the empirical covariance matrix S . The *Kullback–Leibler divergence* of two Gaussian distributions $N(0, \Sigma_0)$ and $N(0, \Sigma_1)$ on \mathbb{R}^n is equal to

$$\text{KL}(\Sigma_0, \Sigma_1) = \frac{1}{2} \left(\text{tr}(\Sigma_1^{-1} \Sigma_0) - n + \log \frac{\det \Sigma_1}{\det \Sigma_0} \right).$$

Computing the MLE is equivalent to minimizing $\text{KL}(\Sigma_0, \Sigma_1)$ with respect to Σ_1 with $\Sigma_0 = S$. On the other hand, the dual MLE is obtained by minimizing $\text{KL}(\Sigma_0, \Sigma_1)$ with respect to Σ_0 with $\Sigma_1 = S$. Equivalently, we set $W = S^{-1}$ and maximize

$$\ell^\vee(\Sigma) = \log \det \Sigma - \text{tr}(W \Sigma).$$

The idea of utilizing the “wrong” Kullback–Leibler distance is ubiquitous in variational inference and is central for mean field approximation and related methods. The idea of using this estimation method for Gaussian linear covariance models is very natural. It results in a unique maximum, since $\Sigma \mapsto \ell^\vee(\Sigma)$ is a convex function on the positive definite cone \mathbb{S}_+^n . See [\[6\]](#) and also [\[5, §3.2; 18, §4\]](#).

The following algebraic formulation is the analogue to [Proposition 3.1](#).

Proposition 3.3. *Finding all the critical points of the dual log-likelihood function ℓ^\vee amounts to solving the following system of equations in $2 \cdot \binom{n+1}{2}$ unknowns:*

$$\Sigma \in \mathcal{L}, \quad K \Sigma = I_n, \quad K - W \in \mathcal{L}^\perp. \tag{6}$$

Proof. After switching the roles of K and Σ , and of W and S , our problem becomes MLE for linear concentration models. [Equation \(6\)](#) is found in [\[28, \(10\)\]](#). □

The next result lists properties of the dual MLE that are important for statistics.

Proposition 3.4. *The dual maximum likelihood estimator of a Gaussian linear covariance model is consistent, asymptotically normal, and first-order efficient.*

Proof. See Theorems 3.1 and 3.2 in [\[6\]](#). □

First-order efficiency means that the asymptotic variance of the properly normalized dual MLE is optimal, that is, it equals the asymptotic variance of the MLE.

In this paper, we focus on algebraic structures, and we note the following important distinction between our two estimators. The MLE requires the quadratic equations $KS K - K \in \mathcal{L}^\perp$ in (4), whereas the dual MLE requires the linear equations $K - W \in \mathcal{L}^\perp$ in (6). The latter are easier to solve than the former, and they give far fewer solutions. This is quantified by the tables for the ML degrees in the next sections.

We are particularly interested in models whose dual ML estimator $(\check{\Sigma}, \check{K})$ can be written as an explicit expression in the sample covariance matrix S . We identify such scenarios in Sections 6 and 7. Here is a first example to illustrate this point.

Example 3.5. We revisit the Toeplitz model in Example 3.2. For the dual MLE, the three quadratic polynomials in K are now replaced by three linear polynomials

$$k_{11} + k_{22} + k_{33} - w_{11} - w_{22} - w_{33}, \quad k_{12} + k_{23} - w_{12} - w_{23}, \quad k_{13} - w_{13}.$$

The w_{ij} are the entries of the inverse sample covariance matrix $W = S^{-1}$. The new system has two solutions, and we can write the $\check{\sigma}_{ij}$ and \check{k}_{ij} in terms of the w_{ij} (or the s_{ij}) using the familiar formula for solving quadratic equations in one variable. Specifically, for the covariance matrix S in (5) we find that the dual MLE is given by

$$\begin{aligned} [\check{\gamma}_0, \check{\gamma}_1, \check{\gamma}_2] &= [0.203557267562, -0.189349961613, 0.1963649733282] \\ &= \left[\frac{1284368265268038839512}{12363704694314904961417} + \frac{52\sqrt{561647777654592987689702150027364667081}}{12363704694314904961417}, \right. \\ &\quad \left. - \frac{5817390611804320873051}{61818523471574524807085} - \frac{655679934637\sqrt{561647777654592987689702150027364667081}}{163146905524715599705244729886305}, \right. \\ &\quad \left. \frac{1990451408446510673691859}{22254668449766828930550600} + \frac{264990063915733\sqrt{561647777654592987689702150027364667081}}{58732885988897615893888102759069800} \right]. \end{aligned}$$

Needless to say, nonlinear algebra goes much beyond the quadratic formula. In what follows we shall employ state-of-the-art methods for solving polynomial equations.

4. General linear constraints

The *maximum likelihood degree* of a linear covariance model \mathcal{L} is, by definition, the number of complex solutions to the likelihood equations (4) for generic data S . This is abbreviated *ML degree* [30, §7.1]. To compute the ML degree, take S to be a random symmetric $n \times n$ matrix and count all complex critical points of the likelihood function $\ell(\Sigma)$ for $\Sigma \in \mathcal{L}$. Equivalently, the ML degree of the model \mathcal{L} is the number of complex solutions (Σ, K) to the polynomial equations in (4).

We also consider the complex critical points of the dual likelihood function $\ell^\vee(\Sigma)$. Their number, for a generic matrix $S \in \mathbb{S}^n$, is the *dual ML degree* of \mathcal{L} . It coincides with the number of complex solutions (Σ, K) to the polynomial equations in (6).

Our ML degrees can be computed symbolically in a computer algebra system that rests on Gröbner bases. However, this approach is limited to small instances. To get further, we use the methods from numerical nonlinear algebra described in Section 5.

m	n				
	2	3	4	5	6
2	1	3	5	7	9
3	1	7	19	37	61
4		7	45	135	299
5		3	71	361	1121
6		1	81	753	3395
7			63	1245	8513
8			29	1625	17867
9			7	1661	31601
10			1	1323	47343
11				801	60177
12				347	64731
13				97	58561
14				15	44131
15				1	27329
16					13627
17					5341
18					1511
19					289
20					31
21					1

m	n				
	2	3	4	5	6
2	1	2	3	4	5
3	1	4	9	16	25
4		4	17	44	90
5		2	21	86	240
6		1	21	137	528
7			17	188	1016
8			9	212	1696
9			3	188	2396
10			1	137	2886
11				86	3054
12				44	2886
13				16	2396
14				4	1696
15				1	1016
16					528
17					240
18					90
19					25
20					5
21					1

Table 1. ML degrees and dual ML degrees for generic models.

Here we focus on a generic m -dimensional linear subspace \mathcal{L} of \mathbb{S}^n . In practice this means that a basis for \mathcal{L} is chosen by sampling m matrices at random from \mathbb{S}^n .

Proposition 4.1. *The ML degree and the dual ML degree of a generic subspace \mathcal{L} of dimension m in \mathbb{S}^n depends only on m and n . It is independent of the particular choice of \mathcal{L} . For small parameter values, these ML degrees are listed in Table 1.*

Proof. The independence rests on general results in algebraic geometry [27, Corollary A.14.2], to the effect that the system (4) (resp. (6)) can be considered as a system parametrized by the coordinates of \mathcal{L} and S (resp. W). The ML degree will be the same for all specializations to \mathbb{R} that remain outside a certain discriminant hypersurface. Table 1 and further values are computed rapidly using the software described in Section 5. □

The dual ML degree was already studied by [28, §2]. Our table on the right is in fact found in their paper. The symmetry along its columns is proved in [28, Theorem 2.3]. It states that the dual ML degree for dimension m coincides with the dual ML degree for codimension $m - 1$. This is derived from the equations (6) by an appropriate homogenization. Namely, the middle equation is clearly symmetric under switching the roles of K and Σ , and the linear equations on the left and on the right in (6) can also be interchanged under this switch.

It was conjectured in [28, §2] that, for fixed m , the dual ML degree is a polynomial of degree $m - 1$ in the matrix size n . This is easy to see for $m \leq 3$. The polynomials for $m = 4$ and $m = 5$ were also derived in [28, §2].

The situation is similar but more complicated for the ML degree. First of all, the symmetry along columns no longer holds as seen on the left in Table 1. This is explained by the fact that the linear equation $K - W \in \mathcal{L}^\perp$ is now replaced by the quadratic equation $KS K - K \in \mathcal{L}^\perp$. However, the polynomiality along the rows of Table 1 seems to persist. For $m = 2$ the ML degree equals $2n - 3$, as shown recently by Coons, Marigliano, and Ruddy [7]. For $m \geq 3$ we propose the following conjecture.

Conjecture 4.2. The ML degree of a linear covariance model of dimension m is a polynomial of degree $m - 1$ in the ambient dimension n . For $m = 3$ this ML degree equals $3n^2 - 9n + 7$, and for $m = 4$ it equals $\frac{11}{3}n^3 - 18n^2 + \frac{85}{3}n - 15$.

We now come to *diagonal linear covariance models*. For these models, \mathcal{L} is a linear subspace of dimension m inside the space \mathbb{R}^n of diagonal $n \times n$ matrices. We wish to determine the ML degree and dual ML degree when \mathcal{L} is generic in \mathbb{R}^n .

In the diagonal case, the score equations simplify as follows. Both the covariance matrix and the concentration matrix are diagonal. We eliminate the entries of Σ by setting $K = \text{diag}(k_1, \dots, k_n)$ and $\Sigma = \text{diag}(k_1^{-1}, \dots, k_n^{-1})$. We also write s_1, \dots, s_n for the diagonal entries of the sample covariance matrix S and $w_i = s_i^{-1}$ for their reciprocals. Finally, let \mathcal{L}^{-1} denote the *reciprocal linear space* of \mathcal{L} , i.e., the variety obtained as the closure of the set of coordinatewise reciprocals of vectors in $\mathcal{L} \cap (\mathbb{R}^*)^n$.

Proposition 4.3. Let $\mathcal{L} \subset \mathbb{R}^n$ be a linear space, viewed as a Gaussian covariance model of diagonal matrices. The score equations for the likelihood in (4) and the dual likelihood in (6) can be written as systems of n equations in n unknowns:

$$(k_1, \dots, k_n) \in \mathcal{L}^{-1} \quad \text{and} \quad (s_1 k_1^2 - k_1, s_2 k_2^2 - k_2, \dots, s_n k_n^2 - k_n) \in \mathcal{L}^\perp, \tag{4'}$$

$$(k_1, \dots, k_n) \in \mathcal{L}^{-1} \quad \text{and} \quad (k_1 - w_1, k_2 - w_2, \dots, k_n - w_n) \in \mathcal{L}^\perp. \tag{6'}$$

The number of complex solutions to (6') for generic \mathcal{L} of dimension m equals $\binom{n-1}{m-1}$.

Proof. The translation of (4) and (6) to (4') and (6') is straightforward. The equations (6') represent a general linear section of the reciprocal linear space \mathcal{L}^{-1} . Proudfoot and Speyer showed that the degree of \mathcal{L}^{-1} equals the Möbius invariant of the underlying matroid. We refer to [19] for a recent study. This Möbius invariant equals $\binom{n-1}{m-1}$ in the generic case, when the matroid is uniform. \square

It would be desirable to express the number of complex solutions to (4') as a matroid invariant, and thereby explain the entries on the left side of Table 2. As before, the m -th row gives the values of a polynomial of degree $m - 1$. For instance, for $m = 3$ we find $2n^2 - 8n + 7$, and for $m = 4$ we find $\frac{4}{3}n^3 - 10n^2 + \frac{68}{3}n - 15$.

m	n				
	3	4	5	6	7
2	3	5	7	9	11
3	1	7	17	31	49
4		1	15	49	111
5			1	31	129
6				1	63
7					1

m	n				
	3	4	5	6	7
2	2	3	4	5	6
3	1	3	6	10	15
4		1	4	10	21
5			1	5	21
6				1	15
7					1

Table 2. ML degrees and dual ML degrees for generic diagonal models.

5. Numerical nonlinear algebra

Linear algebra is the foundation of scientific computing and applied mathematics. *Nonlinear algebra* [23] is a generalization where linear systems are replaced by nonlinear equations and inequalities. At the heart of this lies algebraic geometry, but there are links to many other branches, such as combinatorics, algebraic topology, commutative algebra, convex and discrete geometry, tensors and multilinear algebra, number theory, and representation theory. Nonlinear algebra is not simply a rebranding of algebraic geometry. It highlights that the focus is on computation and applications, and the theoretical needs that this requires results in a new perspective.

We refer to *numerical nonlinear algebra* as the branch of nonlinear algebra which is concerned with the efficient numerical solution of polynomial equations and inequalities. In the existing literature, this is referred to as numerical algebraic geometry. In the following we discuss the numerical solution of polynomial equations, and we describe the techniques used for deriving the computational results in this paper.

One of our main contributions is the Julia package `LinearCovarianceModels.jl` for estimating linear covariance models [31]. Given \mathcal{L} , our package computes the ML degree and the dual ML degree. For any S , it finds all critical points and selects those that are local maxima. The following example explains how this is done.

Example 5.1. We use the package to verify [Example 3.2](#):

```
julia> using LinearCovarianceModels
julia> Σ = toeplitz(3)
3-dimensional LCMoel:
  θ1 θ2 θ3
  θ2 θ1 θ2
  θ3 θ2 θ1
```

We compute the ML degree of the family Σ by computing all solutions for a generic instance. The pair of solutions and generic instance is called an *ML degree witness*:

```
julia> W = ml_degree_witness(Σ)
MLDegreeWitness:
```

- ML degree $\rightarrow 3$
- model dimension $\rightarrow 3$
- dual \rightarrow false

By default, the computation of the ML degree witness relies on a heuristic stopping criterion. We can numerically verify the correctness by using a trace test [21]:

```
julia> verify(W)
Compute additional witnesses for completeness...
Found 10 additional witnesses
Compute trace...
Norm of trace: 2.6521474798326718e-12
true
```

We now input the specific sample covariance matrix in (5), and we compute all critical points of this MLE problem using the ML degree witness from the previous step:

```
julia> S = [4/5 -9/5 -1/25; -9/5 79/16 25/24; -1/25 25/24 17/16];
julia> critical_points(W, S)
3-element Array{Tuple{Array{Float64,1},Float64,Symbol},1}:
 ([2.39038, -0.286009, 0.949965], -5.421751313919751, :local_maximum)
 ([2.52783, -0.215929, -1.45229], -5.346601549034418, :global_maximum)
 ([2.28596, -0.256394, 0.422321], -5.424161999175718, :saddle_point)
```

If only the global maximum is of interest then this can also be computed directly:

```
julia> mle(W, S)
3-element Array{Float64,1}:
 2.527832268219689
-0.21592947057775033
-1.4522862659134732
```

By default only positive definite solutions are reported. To list all critical points we run the command with an additional option:

```
julia> critical_points(W, S, only_positive_definite=false)
```

In this case, since the ML degree is 3, we are not getting more solutions.

In the rest of this section we explain the mathematics behind our software, and how it applies to our MLE problems. A textbook introduction to the numerical solution of polynomial systems by homotopy continuation methods is [27].

Suppose we are given m polynomials f_1, \dots, f_m in n unknowns x_1, \dots, x_n with complex coefficients, where $m \geq n$. We are interested in computing all isolated complex solutions of the system $f_1(x) = \dots = f_m(x) = 0$. These solutions comprise the zero-dimensional components of the variety $V(F)$ where $F = (f_1, \dots, f_m)$.

The general idea of homotopy continuation is as follows. Assume we have another system $G = (g_1, \dots, g_m)$ of polynomials for which we know some or all of its solutions. Suppose there is a *homotopy* $H(x, t)$ with $H(x, 0) = G(x)$ and $H(x, 1) = F(x)$ with the property that, for every $x^* \in V(G)$, there exists a smooth path $x : [0, 1) \rightarrow \mathbb{C}^n$ with $x(0) = x^*$ and $H(x(t), t) = 0$ for all $t \in [0, 1)$. Then we can track each point in $V(G)$ to a point in $V(F)$. This is done by solving the *Dauidenko differential equation*

$$\frac{\partial H}{\partial x}(x(t), t) \cdot \dot{x}(t) + \frac{\partial H}{\partial t}(x(t), t) = 0$$

with initial condition $x(0) = x^*$. Using a *predictor-corrector* scheme for numerical path tracking, both the local and global error can be controlled. Methods called *endgames* are used to handle divergent paths and singular solutions [27, Chapter 10].

Here is a general framework for start systems and homotopies. Embed F in a family of polynomial systems \mathcal{F}_Q , continuously parametrized by a convex open set $Q \subset \mathbb{C}^k$. We have $F = F_q \in \mathcal{F}_Q$ for some $q \in Q$. Outside a Zariski closed set $\Delta \subset Q$, every system in \mathcal{F}_Q has the same number of solutions. If $p \in Q \setminus \Delta$, then F_p is such a *generic instance* of the family \mathcal{F}_Q , and the following is a suitable homotopy [25]:

$$H(x, t) = F_{(1-t)p+ tq}(x). \tag{7}$$

Now, to compute $V(F_q)$, it suffices to find all solutions of a generic instance F_p and then track these along the homotopy (7). Obtaining all solutions of a generic instance can be a challenge, but this has to be done *only once!* That is the *offline phase*. Tracking from a generic to a specific instance of interest is the *online phase*.

A key point in applying this method is the choice of the family \mathcal{F}_Q . For MLE problems in statistics, it is natural to choose Q as the space of data or instances. In our scenario, Q is \mathbb{S}^n , or a complex version thereof. We shall discuss this below.

First, we explain the *monodromy method* for an arbitrary family \mathcal{F}_Q . Suppose the general instance has d solutions, and that we are given one *start pair* (x_0, p_0) . This means that x_0 is a solution to the instance F_{p_0} . Consider the incidence variety

$$Y := \{(x, p) \in \mathbb{C}^n \times Q \mid F_p(x) = 0\}.$$

Let π be the projection from $\mathbb{C}^n \times Q$ onto the second factor. For $q \in Q \setminus \Delta$, the fiber $\pi^{-1}(q)$ has exactly d points. A loop in $Q \setminus \Delta$ based at q has d lifts to Y . Associating a point in the fiber to the endpoint of the corresponding lift gives a permutation in S_d . This defines an action of the fundamental group of $Q \setminus \Delta$ on the fiber $\pi^{-1}(q)$. The *monodromy group* of our family is the image of the fundamental group in S_d .

The monodromy method fills the fiber $\pi^{-1}(p_0)$ by exploiting the monodromy group. For this, the start solution x_0 is numerically tracked along a loop in $Q \setminus \Delta$, yielding a solution x_1 at the end. If $x_1 \neq x_0$, then x_1 is also tracked along the *same* loop, possibly again yielding a new solution. This is done until no more solutions are found. Then, all solutions are tracked along a new loop, where the process is repeated. This process is stopped by use of a *trace test*. For a detailed description of the monodromy method and the trace test, see [4; 21]. To get this off the ground, one needs a start pair (x_0, p_0) . This can often be

found by inverting the problem. Instead of finding a solution x_0 to a given p_0 , we start with x_0 and look for p_0 such that $F_{p_0}(x_0) = 0$.

We now explain how this works for the score equations (4) of our MLE problem. First pick a random matrix Σ_0 in the subspace \mathcal{L} . We next compute K_0 by inverting Σ_0 . Finally we need to find a symmetric matrix S_0 such that $K_0 S_0 K_0 - K_0 \in \mathcal{L}^\perp$. Note that this is a linear system of equations and hence directly solvable. In this manner, we easily find a start pair (x_0, p_0) by setting $p_0 = S_0$ and $x_0 = (\Sigma_0, K_0)$.

The number d of solutions to a generic instance is the ML degree of our model. A priori knowledge of d is useful because it serves as a stopping criterion in the monodromy method. This is one reason for focusing on the ML degree in this paper.

6. Local maxima versus rational MLE

The theme of this paper is maximum likelihood inference for linear covariance models. We developed some numerical nonlinear algebra for this problem, and we offer a software package [31]. From the applications perspective, this is motivated by the fact that the likelihood function is nonconvex. It can have multiple local maxima. A concrete instance for 3×3 Toeplitz matrices was shown in Example 3.2.

In this section we undertake a more systematic experimental study of local maxima. Our aim is to answer the following question: there is the theoretical possibility that $\ell(\Sigma)$ has many local maxima, but can we also observe this in practice?

To address this question, we explored a range of linear covariance models \mathcal{L} . For each model, we conducted the following experiment. We repeatedly generated sample covariance matrices $S \in \mathbb{S}_+^n$. This was done as follows. We first sample a matrix $X \in \mathbb{R}^{n \times n}$ by picking each entry independently from a normal distribution with mean zero and variance one. And then we set $S := XX^T/n$. This is equivalent to sampling $nS \in \mathbb{S}_+^n$ from the standard Wishart distribution with n degrees of freedom.

For each of the generated sample covariance matrices S , we computed the real solutions of the likelihood equations (4). From these, we identified the set of all local maxima in \mathbb{S}^n , and we extracted its subset of local maxima in the positive definite cone \mathbb{S}_+^n . We recorded the numbers of these local maxima. Moreover, we kept track of the fraction of instances S for which there were multiple (positive definite) local maxima. In Table 3 we present our results for $n = 5$ and generic linear subspaces \mathcal{L} .

	m												
	2	3	4	5	6	7	8	9	10	11	12	13	14
ML degree	7	37	135	361	753	1245	1625	1661	1323	801	347	97	15
max	2	3	3	5	5	5	5	6	7	5	4	2	1
max pd	1	2	3	3	4	4	4	4	5	5	4	2	1
multiple	0.4%	5.8	13.8	31.2	37.2	39.0	40.6	37.4	32.0	20.4	13.8	3.0	0.0
multiple pd	0.0%	4.6	11.2	22.4	25.2	31.6	33.0	34.8	29.6	19.4	13.0	3.0	0.0

Table 3. Experiments for generic m -dimensional linear subspaces of \mathbb{S}^5 .

	tree number										
	1	2	3	4	5	6	7	8	9	10	11
ML degree	37	37	81	31	27	31	31	27	13	17	17
max	3	3	4	3	3	3	4	3	3	3	3
max pd	3	2	3	3	3	3	2	2	3	3	3
multiple	21.2%	22.8	24.2	15.6	23.0	21.2	21.2	15.4	13.8	16.2	12.4
multiple pd	8.2%	9.4	14.0	10.0	15.8	13.0	12.2	8.8	13.8	16.2	12.4

Table 4. Experiments for eleven Brownian motion tree models with 5 leaves.

For each m between 2 and 14, we selected five generic linear subspaces \mathcal{L} in the 15-dimensional space \mathbb{S}^5 . Each linear subspace \mathcal{L} was constructed by choosing a basis of positive definite matrices. The basis elements were constructed with the same sampling method as the sample covariance matrices. The ML degree of this linear covariance model is the corresponding entry in the $n = 5$ column on the left in Table 1. These degrees are repeated in the row named *ML degree* in Table 3.

For each model \mathcal{L} , we generated 100 sample covariance matrices S , and we solved the likelihood equations (4) using our software `LinearCovarianceModels.jl`. The row *max* denotes the largest number of local maxima that was observed in these 100 experiments. The row *multiple* gives the fraction of instances which resulted in two or more local maxima. These two numbers pertain to local maxima in \mathbb{S}^5 . The rows *max pd* and *multiple pd* are the analogues restricted to the positive definite cone \mathbb{S}_+^5 .

For an illustration, let us discuss the models of dimension $m = 7$. These equations (4) have 1245 complex solutions, but the number of real solutions is much smaller. Nevertheless, in two fifths of the instances (39.0%) there were two or more local maxima in \mathbb{S}^5 . In one third of the instances (31.6%) the same happened \mathbb{S}_+^5 . The latter is the case of interest in statistics. One instance had four local maxima in \mathbb{S}_+^5 .

The second experiment we report concerns a combinatorially defined class of linear covariance models, namely the Brownian motion tree models in (1). We consider eleven combinatorial types of trees with 5 leaves. For each model we perform the experiment described above, but we now used 500 sample covariance matrices per model. Our results are presented in Table 4, in the same format as in Table 3.

The eleven trees are numbered by the order in which they appear in Table 5. For instance, tree 1 gives the 7-dimensional model in \mathbb{S}_+^5 whose covariance matrices are

$$\Sigma = \begin{bmatrix} \gamma_1 & \gamma_6 & \gamma_6 & \gamma_6 & \gamma_7 \\ \gamma_6 & \gamma_2 & \gamma_6 & \gamma_6 & \gamma_7 \\ \gamma_6 & \gamma_6 & \gamma_3 & \gamma_6 & \gamma_7 \\ \gamma_6 & \gamma_6 & \gamma_6 & \gamma_4 & \gamma_7 \\ \gamma_7 & \gamma_7 & \gamma_7 & \gamma_7 & \gamma_5 \end{bmatrix}.$$

This model has ML degree 37. Around eight percent of the instances led to multiple maxima among positive definite matrices. Up to three such maxima were observed.

The results reported in Tables 3 and 4 show that the maximal number of local maxima increases with the ML degree. But they do not increase as fast as one would expect from the growth of the ML degree. On the other hand, the frequency of observing multiple local maxima seems to be roughly related to the ML degree.

n	clades	ML degree	dual ML degree
5	{1, 2, 3, 4}	37	11
5	{1, 2}	37	11
5	{1, 2, 3}	81	16
5	{1, 2}, {3, 4, 5}	31	4
5	{1, 2}, {3, 4}	27	4
5	{1, 2, 3}, {1, 2, 3, 4}	31	4
5	{1, 2}, {1, 2, 3}	31	4
5	{1, 2}, {1, 2, 3, 4}	27	4
5	{1, 2}, {3, 4}, {1, 2, 3, 4}	13	1
5	{1, 2}, {3, 4}, {1, 2, 5}	17	1
5	{1, 2}, {1, 2, 3}, {1, 2, 3, 4}	17	1
6	{1, 2, 3, 4, 5}	95	26
6	{1, 2}	95	26
6	{1, 2, 3, 4}	259	44
6	{1, 2, 3}	259	44
6	{1, 2, 3}, {4, 5, 6}	221	16
6	{1, 2}, {3, 4, 5, 6}	101	11
6	{1, 2, 3, 4}, {1, 2, 3, 4, 5}	101	11
6	{1, 2}, {3, 4}	81	11
6	{1, 2}, {1, 2, 3}	101	11
6	{1, 2}, {3, 4, 5}	181	16
6	{1, 2}, {1, 2, 3, 4, 5}	81	11
6	{1, 2, 3}, {1, 2, 3, 4}	221	16
6	{1, 2, 3}, {1, 2, 3, 4, 5}	181	16
6	{1, 2}, {1, 2, 3, 4}	181	16
6	{1, 2}, {3, 4}, {5, 6}	63	4
6	{1, 2}, {3, 4}, {1, 2, 3, 4}	99	4
6	{1, 2}, {1, 2, 3}, {4, 5, 6}	115	4
6	{1, 2}, {3, 4, 5}, {3, 4, 5, 6}	115	4
6	{1, 2}, {3, 4, 5}, {1, 2, 3, 4, 5}	99	4
6	{1, 2}, {3, 4}, {1, 2, 5, 6}	83	4
6	{1, 2}, {3, 4}, {1, 2, 3, 4, 5}	63	4
6	{1, 2, 3}, {1, 2, 3, 4}, {1, 2, 3, 4, 5}	115	4
6	{1, 2}, {3, 4}, {1, 2, 5}	83	4
6	{1, 2}, {1, 2, 3}, {1, 2, 3, 4}	115	4
6	{1, 2}, {1, 2, 3}, {1, 2, 3, 4, 5}	83	4
6	{1, 2}, {1, 2, 3, 4}, {1, 2, 3, 4, 5}	83	4
6	{1, 2}, {3, 4}, {5, 6}, {1, 2, 3, 4}	53	1
6	{1, 2}, {3, 4}, {1, 2, 5}, {3, 4, 6}	61	1
6	{1, 2}, {3, 4}, {1, 2, 3, 4}, {1, 2, 3, 4, 5}	53	1
6	{1, 2}, {3, 4}, {1, 2, 5}, {1, 2, 5, 6}	61	1
6	{1, 2}, {3, 4}, {1, 2, 5}, {1, 2, 3, 4, 5}	53	1
6	{1, 2}, {1, 2, 3}, {1, 2, 3, 4}, {1, 2, 3, 4, 5}	61	1

Table 5. ML degrees and dual ML degrees for Brownian motion tree models with five and six leaves. Binary trees are in bold.

Here is an interesting observation to be made in [Table 4](#). The last three trees, labeled 9, 10, and 11, are the binary trees. These have the maximum dimension $2n - 2$. For these models, every local maximum in \mathbb{S}^n is also in the positive definite cone \mathbb{S}_+^n . We also verified this for all binary trees with $n = 6$ leaves. This is interesting since the positive-definiteness constraint is the hardest to respect in an optimization routine. It is tempting to conjecture that this persists for all binary trees with $n \geq 7$.

There is another striking observation in [Table 5](#). The dual ML degree for binary trees is always equal to one. We shall prove in [Theorem 7.3](#) that this holds for any n . This means that the dual MLE can be expressed as a rational function in the data S . Hence there is only one local maximum, which is therefore the global maximum.

We close this section with a few remarks on the important special case when the ML degree or the dual ML degree is equal to one. This holds if and only if each entry of the estimated matrix $\widehat{\Sigma}$ or $\check{\Sigma}$ is a rational function in the $\binom{n+1}{2}$ quantities s_{ij} .

Rationality of the MLE has received a lot of attention in the case of *discrete random variables*. See [\[30, §7.1\]](#) for a textbook reference. If the MLE of a discrete model is rational, then its coordinates are alternating products of linear forms in the data [\[30, Theorem 7.3.4\]](#). This result due to Huh was refined in [\[11, Theorem 1\]](#). At present we have no idea what the analogue in the Gaussian case might look like.

Problem 6.1. Characterize all Gaussian models whose MLE is a rational function.

In addition to the binary trees in [Theorem 7.3](#), statisticians are familiar with a number of situations when the dual MLE is rational. The dual MLE is the MLE of a linear concentration model with the sample covariance matrix S replaced by its inverse W . This is studied in [\[28\]](#) and in many other sources on Gaussian graphical models and exponential families. The following result paraphrases [\[28, Theorem 4.3\]](#).

Proposition 6.2. *If a linear covariance model \mathcal{L} is given by zero restrictions on Σ , then the dual ML degree is equal to one if and only if the associated graph is chordal.*

It would be interesting to extend this result to other combinatorial families, such as colored covariance graph models [\[17\]](#), including structured Toeplitz matrices.

The following example illustrates [Problem 6.1](#) and raises some further questions.

Example 6.3. We present a linear covariance model such that both the MLE and the dual MLE are rational functions. Fix $n \geq 2$ and let \mathcal{L} be the hyperplane with equation $\sigma_{12} = 0$. By [Proposition 6.2](#), the dual ML degree of \mathcal{L} is one. The model is dual to the decomposable undirected graphical model with missing edge $\{1, 2\}$.

Following [\[20; 28\]](#), we obtain the rational formula for its dual MLE:

$$\check{k}_{12} = W_{1,R} W_{R,R}^{-1} W_{R,2} \quad \text{and} \quad \check{k}_{ij} = w_{ij} \quad \text{for } (i, j) \neq (1, 2). \tag{8}$$

Here $R = \{3, \dots, n\}$ and $W_{\cdot, \cdot}$ is our notation for submatrices of $W = (w_{ij}) = S^{-1}$.

The ML degree of the model \mathcal{L} is also one. To see this, we note that \mathcal{L} is the DAG model with edges $i \rightarrow j$ whenever $i < j$ unless $(i, j) = (1, 2)$. By [\[20, §5.4.1\]](#), the MLE of any Gaussian DAG model is

rational. In our case, we find $\widehat{K} = W + A$, where A is the $n \times n$ matrix which is zero apart from the upper left 2×2 block

$$A_{12,12} = \begin{bmatrix} s_{11}^{-1} & 0 \\ 0 & s_{22}^{-1} \end{bmatrix} - \frac{1}{s_{11}s_{22} - s_{12}^2} \begin{bmatrix} s_{22} & -s_{12} \\ -s_{12} & s_{11} \end{bmatrix}.$$

The entries in $\check{\Sigma} = (\check{K})^{-1}$ and $\widehat{\Sigma} = (\widehat{K})^{-1}$ are rational functions in the data s_{ij} . But, unlike in the discrete case of [11], here the rational functions are not products of linear forms. Problem 6.1 asks for an understanding of its irreducible factors.

Example 6.3 raises many questions. First of all, can we characterize all linear spaces \mathcal{L} with rational formulas for their MLE, or their dual MLE, or both of them? Second, it would be interesting to study arbitrary models \mathcal{L} that are hyperplanes. Consider the entries for $m = \binom{n+1}{2} - 1$ in Tables 1 and 3. We know from [28, §2.2] that the dual ML degree equals $n - 1$. The ML degree seems to be $2^{n-1} - 1$. In all cases there seems to be only one local (and hence global) maximum. How could one prove these observations? Finally, it is worthwhile to study the MLE when \mathcal{L}^\perp is a generic symmetric matrix of rank r . What is the ML degree in terms of r and n ?

7. Brownian motion tree models

We now study the linear space \mathcal{L}_T associated with a rooted tree T with n leaves. The equations of \mathcal{L}_T are $\sigma_{ij} = \sigma_{kl}$ whenever $\text{lca}(i, j) = \text{lca}(k, l)$. In the literature [13; 29] one assumes that the parameters θ_A in (1) are nonnegative. Here, we relax this hypothesis: we allow all covariance matrices in the spectrahedron $\mathcal{L}_T \cap \mathbb{S}_+^n$.

The ML degree and its dual do not depend on how the leaves of a tree are labeled but only on the tree topology. For fixed n each tree topology is uniquely identified by the set of clades. Since the root clade $\{1, \dots, n\}$ and the leaf-clades $\{1\}, \dots, \{n\}$ are part of every tree, they are omitted in our notation. For example, if $n = 5$, then the tree $\{\{1, 2\}, \{3, 4\}, \{3, 4, 5\}\}$ is the binary tree with four inner vertices corresponding to the three nontrivial clades mentioned explicitly. This tree is depicted in Figure 1.

We computed the ML degree and the dual ML degree of \mathcal{L}_T for many trees T . In Table 5 we report results for five and six leaves. We notice that the dual ML degree is exactly one for all binary trees. This suggests that the dual MLE is a rational function. Our main result in this section (Theorem 7.3) says that this is indeed true.

The equations (6) for the dual ML degree can be written as $e_A^T(K - W)e_A = 0$ for all clades A . Here $W = (w_{ij})$ is given and $K^{-1} \in \mathcal{L}_T$ is unknown. We abbreviate

$$w_{A,B} = \sum_{i \in A} \sum_{j \in B} w_{ij} = e_A^T W e_B. \tag{9}$$

The same notation is used for general matrices. We present two examples with $n = 4$.

Example 7.1. Consider the tree with clades $\{1, 2\}, \{3, 4\}$, shown in [29, Figure 1]. The dual MLE \check{K} satisfies $\check{k}_{ii} = w_{ii}$ for $i = 1, 2, 3, 4$, and $\check{k}_{12} = w_{12}$, $\check{k}_{34} = w_{34}$, and

$$\check{k}_{ij} = w_{12,34} \frac{w_{i,12} w_{j,34}}{w_{12,12} w_{34,34}} \quad \text{for } i \in \{1, 2\} \text{ and } j \in \{3, 4\}.$$

Example 7.2. The tree with clades $\{1, 2\}, \{1, 2, 3\}$ has $\check{k}_{ii} = w_{ii}, \check{k}_{12} = w_{12}$, and

$$\check{k}_{13} = w_{12,3} \frac{w_{1,12}}{w_{12,12}}, \quad \check{k}_{14} = w_{123,4} \frac{w_{1,12} w_{12,123}}{w_{12,12} w_{123,123}}, \quad \check{k}_{23} = w_{12,3} \frac{w_{2,12}}{w_{12,12}},$$

$$\check{k}_{24} = w_{123,4} \frac{w_{2,12} w_{12,123}}{w_{12,12} w_{123,123}}, \quad \check{k}_{34} = w_{123,4} \frac{w_{123,3}}{w_{123,123}}.$$

Both examples were computed in Mathematica using the description of the Brownian motion tree model in terms of the inverse covariance matrix given in [29].

Recall that for $v \in V$ we write $\text{de}(v)$ for the set of leaves of T that are descendants of v . The following theorem generalizes formulas in the above two examples. It is our main result in Section 7.

Theorem 7.3. Consider the model \mathcal{L}_T given by a rooted binary tree T with n leaves. The dual MLE $\check{K} = (\check{k}_{ij})$ satisfies $\check{k}_{ii} = w_{i,i}$ for all i , and its off-diagonal entries are

$$\check{k}_{ij} = w_{A,B} \prod_{u \rightarrow v} \frac{w_{\text{de}(v), \text{de}(u)}}{w_{\text{de}(u), \text{de}(u)}} \quad \text{for } 1 \leq i < j \leq n. \quad (10)$$

Here A, B are the clades of the two children of $\text{lca}(i, j)$. The product is over all edges $u \rightarrow v$ of T , except for the two edges with $u = \text{lca}(i, j)$, on the path from i to j in T .

Proof. Define $p_{ij} = -k_{ij}$ for $1 \leq i < j \leq n$ and $p_{0i} = \sum_{j=1}^n k_{ij}$ for $1 \leq i \leq n$. By [29, Theorem 1.2], in the new coordinates \mathcal{L}_T is a toric variety with a monomial parametrization $p_{0i} = 1/t_i$ and $p_{ij} = t_{\text{lca}(i,j)}/(t_i t_j)$. The condition $\Sigma \in \mathcal{L}_T, K \Sigma = I_n$ in (6) is therefore equivalent to requiring that the coordinates p_{ij}, p_{0i} admit such a monomial parametrization.

The last condition, $K - W \in \mathcal{L}_T^\perp$, in (6) means that $k_{A,A} = w_{A,A}$ for every clade A of T . This can be rewritten in the new coordinates as

- (i) $\sum_{j \neq i} p_{ij} = w_{i,i}$ for all $1 \leq i \leq n$, and
- (ii) $p_{A,B} = -w_{A,B}$ for all inner vertices u of T , where $A | B$ is the partition of $\text{clade}(u) = A \cup B$ given by the two children of u .

Now fix u with clade partition $A | B$ as above, so $u = \text{lca}(i, j)$ for all $i \in A$ and $j \in B$. The parametrization $p_{0i} = 1/t_i$ and $p_{ij} = t_{\text{lca}(i,j)}/(t_i t_j)$ yields

$$p_{A,B} = t_u \cdot \sum_{i \in A} \frac{1}{t_i} \cdot \sum_{j \in B} \frac{1}{t_j} = t_u \cdot p_{0,A} \cdot p_{0,B}.$$

Using the equations in (ii), we obtain

$$t_u = -\frac{w_{A,B}}{p_{0,A} p_{0,B}} \quad \text{and hence} \quad p_{ij} = -w_{A,B} \frac{p_{0i}}{p_{0,A}} \frac{p_{0j}}{p_{0,B}}. \quad (11)$$

We claim that the following identity holds for any clade $A \subseteq [n]$:

$$p_{0,A} = w_{[n],[n]} \prod_{u \rightarrow v} \frac{w_{\text{de}(v), \text{de}(u)}}{w_{\text{de}(u), \text{de}(u)}}, \quad (12)$$

where the product is over all edges $u \rightarrow v$ of T in the path from the root to the node with clade A . Note that (11) and (12) imply (10) and so the theorem.

We now prove (12). Since $p_{0,[n]} = w_{[n],[n]}$, the claim holds for $A = [n]$. Fix a clade $A \subset [n]$ and assume (12) for all clades $A_1 \subset \dots \subset A_k \subset A_{k+1} = [n]$ strictly containing $A_0 = A$. For each $i = 0, \dots, k$ denote

$$\alpha_i := w_{A_{k+1}, A_{k+1}} \frac{w_{A_k, A_{k+1}}}{w_{A_{k+1}, A_{k+1}}} \dots \frac{w_{A_i, A_{i+1}}}{w_{A_{i+1}, A_{i+1}}}.$$

By the induction hypothesis $p_{0, A_i} = \alpha_i$ for all $i = 1, \dots, k$. Our goal is to prove that $p_{0, A} = \alpha_0$. The clades $A_1 \setminus A, \dots, A_{k+1} \setminus A_k$ form a partition of $\bar{A} = [n] \setminus A$. We have

$$\begin{aligned} p_{0, A} &= w_{A, A} + k_{A, \bar{A}} = w_{AA} - p_{A, A_1 \setminus A} - p_{A, A_2 \setminus A_1} - \dots - p_{A, A_{k+1} \setminus A_k} \\ &= w_{A, A} + w_{A, A_1 \setminus A} + \sum_{i=1}^k w_{A_i, A_{i+1} \setminus A_i} \frac{p_{0A}}{p_{0, A_i}}. \end{aligned} \tag{13}$$

Here the last equality follows because for every $i \in A$ and every $j \in A_{l+1} \setminus A_l$ the vertex $u = \text{lca}(i, j)$ is the same. The clades of the children of u are A_l and $A_{l+1} \setminus A_l$. Therefore, using (11), we get $p_{A, A_{l+1} \setminus A_l} = w_{A_l, A_{l+1} \setminus A_l} p_{0, A} / p_{0, A_l}$. We rewrite (13) as

$$w_{A, A_1} = p_{0, A} \left(1 - \sum_{i=1}^k \frac{w_{A_i, A_{i+1}} - w_{A_i, A_i}}{\alpha_i} \right). \tag{14}$$

To simplify the bracketed expression, note that $w_{A_i, A_{i+1}} / \alpha_i = w_{A_{i+1}, A_{i+1}} / \alpha_{i+1}$, and so

$$1 - \sum_{i=1}^k \frac{w_{A_i, A_{i+1}} - w_{A_i, A_i}}{\alpha_i} = 1 + \frac{w_{A_1, A_1}}{\alpha_1} - \frac{w_{A_k, A_{k+1}}}{\alpha_k} = \frac{w_{A_1, A_1}}{\alpha_1}.$$

Plugging this back into (14) gives

$$p_{0, A} = \alpha_1 \frac{w_{A, A_1}}{w_{A_1, A_1}} = \alpha_0.$$

This proves the correctness of (12).

We have shown that (6) implies the rational formula (10) for \check{K} in terms of W . To argue that this is the MLE, we need that $W \in \mathbb{S}_+^n$ implies $\check{K} \in \mathbb{S}_+^n$. For this, we use an analytic argument. Since W is positive definite, the dual likelihood function has a unique maximum $K = W$ over the whole cone \mathbb{S}_+^n . The model $\mathcal{L}_T \cap \mathbb{S}_+^n$ is a relatively closed subset of \mathbb{S}_+^n and so the dual likelihood restricted to this set attains its maximum. The ML degree is equal to one and so there is at most one optimum in \mathcal{L}_T . We conclude there is exactly one optimum in $\mathcal{L}_T \cap \mathbb{S}_+^n$ and it is equal to \check{K} . \square

In our concluding example we compare the MLE and its dual in a special case.

Example 7.4. Fix the five-leaf tree in Figure 1, with clades $\{1, 2\}, \{3, 4\}, \{3, 4, 5\}$. For simplicity assume that the data-generating distribution has all parameters θ_A in (1) equal to one. For each sample size $n = 50, 200, 500, 5000$, we run 1000 iterations to obtain a simple Monte Carlo estimator of the mean squared errors as measured by $\mathbb{E} \|\hat{\Sigma} - \Sigma^*\|^2$ and $\mathbb{E} \|\check{\Sigma} - \Sigma^*\|^2$, where Σ^* is the true covariance matrix and

$\|\cdot\|$ is a given matrix norm. To have a direct comparison between both estimators we also approximate $\mathbb{E}\|\widehat{\Sigma} - \check{\Sigma}\|^2$. We obtain the following numbers for the operator norm:

	50	200	500	5000
approx. $\mathbb{E}\ \widehat{\Sigma} - \Sigma^*\ ^2$	5.44	1.30	0.55	0.05
approx. $\mathbb{E}\ \check{\Sigma} - \Sigma^*\ ^2$	5.28	1.31	0.55	0.05
approx. $\mathbb{E}\ \widehat{\Sigma} - \check{\Sigma}\ ^2$	0.38	0.02	0.00	0.00

We see that the two estimators have essentially the same statistical performance. On average, they lie very close to each other. The dual MLE, which is available in closed form, thus offers a very attractive alternative to the MLE. Similar results were obtained for the Frobenius norm and the ℓ_∞ -norm but they are not reported here.

The estimates in the previous example were computed by evaluating the function in Theorem 7.3. We stress that nonlinear algebra and our software [31] played an essential role in getting to this point. Namely, with computations as described in Section 5, we created Table 5. After seeing that table, we conjectured that the dual MLE for binary trees is one. This led us to find the rational formula. The expression (10) is an alternating product of linear forms, reminiscent of [11, Theorem 1]. However, this structure does not generalize, by Example 6.3, thus underscoring Problem 6.1.

Acknowledgements

We thank Steffen Lauritzen for useful discussions and for providing references. We also thank Jane Coons, Orlando Marigliano, and Michael Ruddy for their comments. Zwiernik was supported by Spanish government grants RYC-2017-22544, PGC2018-101643-B-I00, and SEV-2015-0563 and by the Ayudas Fundación BBVA a Equipos de Investigación Científica 2017. Timme was supported by the Deutsche Forschungsgemeinschaft Graduiertenkolleg *Facets of Complexity* (GRK 2434).

References

- [1] T. W. Anderson, “Estimation of covariance matrices which are linear combinations or whose inverses are linear combinations of given matrices”, pp. 1–24 in *Essays in probability and statistics*, edited by R. C. Bose et al., University of North Carolina, Chapel Hill, NC, 1970.
- [2] P. Breiding and S. Timme, “HomotopyContinuation.jl: a package for homotopy continuation in Julia”, pp. 458–465 in *Mathematical software — ICMS 2018* (South Bend, IN, 2018), edited by J. H. Davenport et al., Lecture Notes in Computer Science **10931**, Springer, 2018.
- [3] J. P. Burg, D. G. Luenberger, and D. L. Wenger, “Estimation of structured covariance matrices”, *Proceedings of the IEEE* **70:9** (1982), 963–974.
- [4] A. Martín del Campo and J. I. Rodríguez, “Critical points via monodromy and local methods”, *J. Symbolic Comput.* **79:3** (2017), 559–574.
- [5] S. Chaudhuri, M. Drton, and T. S. Richardson, “Estimation of a covariance matrix with zeros”, *Biometrika* **94:1** (2007), 199–216.
- [6] E. S. Christensen, “Statistical properties of I -projections within exponential families”, *Scand. J. Statist.* **16:4** (1989), 307–318.

- [7] J. I. Coons, O. Marigliano, and M. Ruddy, “Maximum likelihood degree of the two-dimensional linear Gaussian covariance model”, *Alg. Stat.* **11**:2 (2020). Preprint available at arXiv 1909.04553v2.
- [8] D. R. Cox and N. Wermuth, “Linear dependencies represented by chain graphs”, *Statist. Sci.* **8**:3 (1993), 204–218.
- [9] C. Dellacherie, S. Martinez, and J. San Martín, *Inverse M-matrices and ultrametric matrices*, Lecture Notes in Mathematics **2118**, Springer, 2014.
- [10] M. Drton and T. Richardson, “A new algorithm for maximum likelihood estimation in Gaussian graphical models for marginal independence”, pp. 184–191 in *Proceedings of the Nineteenth Annual Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, 2003.
- [11] E. Duarte, O. Marigliano, and B. Sturmfels, “Discrete statistical models with rational maximum likelihood estimator”, preprint, 2019. arXiv 1903.06110v1
- [12] J. Fan, Y. Liao, and H. Liu, “An overview of the estimation of large covariance and precision matrices”, *Econom. J.* **19**:1 (2016), C1–C32.
- [13] J. Felsenstein, “Maximum-likelihood estimation of evolutionary trees from continuous characters”, *Am. J. Hum. Genet.* **25**:5 (1973), 471–492.
- [14] J. D. Hamilton, *Time series analysis*, Princeton University, 1994.
- [15] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, Springer, 2001.
- [16] L. Ho and C. Ané, “A linear-time algorithm for Gaussian and non-Gaussian trait evolution models”, *Syst. Biol.* **63**:3 (2014), 397–408.
- [17] S. Højsgaard and S. L. Lauritzen, “Graphical Gaussian models with edge and vertex symmetries”, *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70**:5 (2008), 1005–1027.
- [18] G. Kauermann, “On a dualization of graphical Gaussian models”, *Scand. J. Statist.* **23**:1 (1996), 105–116.
- [19] M. Kummer and C. Vinzant, “The Chow form of a reciprocal linear space”, *Michigan Math. J.* **68**:4 (2019), 831–858.
- [20] S. L. Lauritzen, *Graphical models*, Oxford Statistical Science Series **17**, Clarendon, New York, 1996.
- [21] A. Leykin, J. I. Rodriguez, and F. Sottile, “Trace test”, *Arnold Math. J.* **4**:1 (2018), 113–125.
- [22] G. M. Marchetti, “Independencies induced from a graphical Markov model after marginalization and conditioning: the R package **ggm**”, *J. Stat. Softw.* **15**:6 (2006).
- [23] M. Michałek and B. Sturmfels, “Invitation to nonlinear algebra”, book in progress, <https://personal-homepages.mis.mpg.de/michalek/NonLinearAlgebra.pdf>.
- [24] M. I. Miller and D. L. Snyder, “The role of likelihood and entropy in incomplete-data problems: applications to estimating point-process intensities and Toeplitz constrained covariances”, *Proceedings of the IEEE* **75**:7 (1987), 892–907.
- [25] A. P. Morgan and A. J. Sommese, “Coefficient-parameter polynomial continuation”, *Appl. Math. Comput.* **29**:2 (1989), 123–160.
- [26] M. Pourahmadi, “Joint mean-covariance models with applications to longitudinal data: unconstrained parameterisation”, *Biometrika* **86**:3 (1999), 677–690.
- [27] A. J. Sommese and C. W. Wampler, II, *The numerical solution of systems of polynomials arising in engineering and science*, World Scientific, Hackensack, NJ, 2005.
- [28] B. Sturmfels and C. Uhler, “Multivariate Gaussian, semidefinite matrix completion, and convex algebraic geometry”, *Ann. Inst. Statist. Math.* **62**:4 (2010), 603–638.
- [29] B. Sturmfels, C. Uhler, and P. Zwiernik, “Brownian motion tree models are toric”, preprint, 2019. arXiv 1902.09905v1
- [30] S. Sullivant, *Algebraic statistics*, Graduate Studies in Mathematics **194**, American Mathematical Society, Providence, RI, 2018.
- [31] S. Timme, `LinearCovarianceModels.jl`, 2019, <https://github.com/saschatimme/LinearCovarianceModels.jl>. Julia package, version 0.1.2.
- [32] P. Zwiernik, C. Uhler, and D. Richards, “Maximum likelihood estimation for linear Gaussian covariance models”, *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79**:4 (2017), 1269–1292.