

LOCALLY ASSOCIATED GRAPHICAL MODELS AND MIXED CONVEX EXPONENTIAL FAMILIES

BY STEFFEN LAURITZEN^{1,a} AND PIOTR ZWIERNIK^{2,b}

¹*Department of Mathematical Sciences, University of Copenhagen, a.lauritzen@math.ku.dk*

²*Department of Statistical Sciences, University of Toronto, piotr.zwiernik@utoronto.ca*

The notion of multivariate total positivity has proved to be useful in finance and psychology but may be too restrictive in other applications. In this paper, we propose a concept of local association, where highly connected components in a graphical model are positively associated and study its properties. Our main motivation comes from gene expression data, where graphical models have become a popular exploratory tool. The models are instances of what we term *mixed convex exponential families* and we show that a *mixed dual likelihood estimator* has simple exact properties for such families as well as asymptotic properties similar to the maximum likelihood estimator. We further relax the positivity assumption by penalizing negative partial correlations in what we term the *positive graphical lasso*. Finally, we develop a GOLAZO algorithm based on block-coordinate descent that applies to a number of optimization procedures that arise in the context of graphical models, including the estimation problems described above. We derive results on existence of the optimum for such problems.

1. Introduction and summary. It has been illustrated recently in a number of publications that explicitly incorporating positive dependence constraints can be useful for modelling in various contexts where components are naturally positively associated (e.g., finance or psychology) [1, 17, 33, 34, 46]. The main distinctive feature of this line of work as opposed to more classical literature on positive dependence is that they link to techniques used in high-dimensional statistics and graphical models using the positivity constraint as an implicit regularizer.

In the Gaussian case, a natural positivity constraint is that all partial correlations are non-negative or, equivalently, the inverse covariance matrix is an M-matrix, $(\Sigma^{-1})_{ij} \leq 0$ for all $i \neq j$. Optimizing a loss function under this restriction typically results in a sparse estimate, which was the driving idea in [46]. For standard stock market datasets, this may lead to an estimate that gives both a sparser graph and a higher value of the likelihood function than estimates from the graphical lasso approach [45].

Although useful, this global positivity constraint is often too restrictive. In this paper, we propose and study natural relaxations of the condition. With an underlying graph representing the dependence structure between the variables, we will require that highly connected components are positively dependent, in the precise sense that variables in the same clique are *associated* [16], which in the Gaussian case is equivalent to having a covariance matrix with positive entries [42]. Unfortunately, maximum likelihood estimation (MLE) in this type of model is problematic as the likelihood function may get complicated. The model is an instance of what we term a *mixed convex exponential family*. We develop an associated *mixed dual estimator* (MDE) which overcomes the problems faced by the MLE. The MDE can be

Received June 2021; revised July 2022.

MSC2020 subject classifications. Primary 62H05, 62H12; secondary 62H22.

Key words and phrases. Association, convex optimization, dual likelihood, exponential families, Gaussian distribution, graphical lasso, Kullback–Leibler divergence, mixed parametrization, positive correlations, structure learning.

found by solving two convex optimization problems, and has asymptotic properties similar to the MLE. We note that Maliutov *et al.* [38] suggest a different relaxation of the MTP₂ condition, unrelated to ours.

In this situation, when the underlying graph is not known, we also consider a further relaxation of local association through what we name the *positive graphical lasso* which penalizes large negative partial correlations. As for the now classical graphical lasso, this will typically identify a sparse structure.

1.1. A motivating example. A motivating problem is the exploratory analysis of gene expression data. The graphical lasso has become a standard technique for estimating gene expression networks. While constructing and interpreting the network, researchers often focus on positive co-expression (e.g., [40]), where pairs of genes show a proportional expression pattern across samples. Also, in various scenarios it has been observed that positively co-expressed genes within the same pathway tend to cluster close together in the pathway structure, while negatively correlated genes typically occupy more distant positions; see, for example, [29, 52]. We shall later, in Section 9.2, analyse publicly available microarray expression data profiling umbilical cord tissue in a study of fetal inflammatory response (FIR); cf. [11, 12]. From an initial set of 12,093 genes with reliable expression, Costa and Castelo [12] identified 1097 as differentially expressed between FIR-affected and unaffected infants, from which 592 were upregulated in FIR. This subset of 592 upregulated genes was significantly enriched by 136 genes involved in the innate immune response ([7]) and we shall focus our analysis on this subset of 136 genes.

A typical approach to explicitly model positive co-expression is by building weighted gene co-expression networks where correlations are mapped monotonely from $[-1, 1]$ to $[0, 1]$ and then thresholded. This approach is subject to standard problems with building co-expression networks based on correlations alone not taking the effect of other genes into account. Our approach is based on partial correlations and so may provide more meaningful estimates of the underlying network.

1.2. An optimization algorithm. An important first step in modelling large systems that satisfy some positive dependence constraint is to reduce to a sparser representation without loosing the positive dependence information. In this respect, ℓ_1 -regularized approaches do not work well, since they treat positive and negative partial correlations in an equal manner. In this paper, we propose a version where only negative partial correlations are penalized.

This approach is developed further in a general GOLAZO¹ algorithm (Graphical Oriented LAZY Optimization) where a penalty of the form

$$\sum_{i \neq j} \max\{L_{ij}K_{ij}, U_{ij}K_{ij}\}$$

is employed to obtain sparse estimation of $K = \Sigma^{-1}$. Here the penalty parameters L, U satisfy $-\infty \leq L_{ij} \leq 0 \leq U_{ij} \leq +\infty$; including zero and infinite values solves several optimization problems proposed in this paper and a number of related problems in graphical models. The advantage of our general approach is that it provides a detailed analysis of convergence and existence of the optimum.

1.3. Main contributions and structure of paper. The main contributions in this paper are the following.

¹Pronounced *goh-lah-soh*, like the Spanish word *golazo* but without the Castilian lisp.

- (i) A class of Gaussian graphical models with relaxation of positivity restrictions, either via a graph or via a positive lasso-type penalty, or both;
- (ii) A general framework for mixed convex exponential families with an associated method of estimation that has asymptotic properties similar to maximum likelihood but more favourable computational properties;
- (iii) The GOLAZO algorithm for solving a range of estimation problems associated with graphical models, including positivity restrictions. The associated R package `golazo` is available on GitHub.

The structure of the paper is as follows. In Section 2, we study the basic properties of locally associated distributions and discuss their relation to other models involving positivity. Section 3 introduces and studies mixed convex exponential families. The mixed dual estimator is introduced in Section 4 and applied to locally associated Gaussian graphical models (laGGMs) in Section 5. The asymptotic properties of the estimator are established in Section 6. In Section 7, we focus on learning the structure of a laGGM and introduce the positive graphical lasso for this purpose. We derive the GOLAZO algorithm in Section 8 and argue that a number of optimization problems in graphical models can be seen as special instances and hence solved by this algorithm. Section 9 concludes the paper by applying the methodology to two datasets.

2. Locally associated distributions. In this section, we define local association and locally associated Gaussian graphical models linking to other relevant statistical models.

2.1. *Definition and basic properties.* We say that a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is increasing if $x \leq x'$ (coordinatewise) implies that $f(x) \leq f(x')$. A d -dimensional random vector X is (positively) *associated* if for any two increasing functions $f, g : \mathbb{R}^d \rightarrow \mathbb{R}$, the covariance $\mathbb{V}(f(X), g(X))$ is nonnegative; for basic properties of this notion see [16].

In general, association is hard to check and [26] introduced the strictly stronger notion of multivariate total positivity which in the Gaussian case is equivalent to the covariance Σ being an inverse M-matrix [27, 33]. We recall that $K = \Sigma^{-1}$ is an M-matrix if it is positive definite and $K_{ij} \leq 0$ for all $i \neq j$; so this condition corresponds to assuming that all partial correlations are nonnegative.

In the Gaussian case, there is a simple condition for association, as stated in the theorem below.

THEOREM 2.1 (Pitt [42]). *Suppose X is a Gaussian vector with covariance matrix Σ . It then holds that X is associated if and only if $\Sigma_{ij} \geq 0$ for all i, j .*

Positive semi-definite matrices with nonnegative entries are also called doubly nonnegative matrices. If Σ is an inverse M-matrix, then it is doubly nonnegative, but the reverse implication does not hold.

The requirement that a distribution is associated is strong and we wish to relax this in relation to a local structure given by a graph. We build on the standard terminology for undirected and directed graphs as given, for example, in [35]. If $G = (V, E)$ is an undirected graph with vertex set $V = \{1, \dots, d\}$ and edge set E , a *complete subset* of G is any subset C of V such that any two vertices $i, j \in C$ are connected by an edge, that is, $ij \in E$. A *clique* is a complete subset that is maximal with respect to inclusion. Let $X = (X_1, \dots, X_d)$ be a random vector and fix a graph $G = (V, E)$. For any $C \subseteq V$, by X_C denote the subvector of X with entries $X_i, i \in C$.

DEFINITION 2.2. The random vector X is said to be *locally associated* w.r.t. G if it holds for any clique C of G that the subvector X_C is associated.

Denote by $A(G)$ the set of covariance matrices of Gaussian vectors that are locally associated with respect to G . By Theorem 2.1, these are the positive definite matrices Σ such that $\Sigma_{ij} \geq 0$ for all edges $ij \in E$.

2.2. *Locally associated Gaussian graphical models.* Our main interest lies in locally associated distributions for Gaussian graphical models. We say that a distribution of a random vector X is *Markov* with respect to G , or $M(G)$, if X satisfies global Markov properties over the graph G ; for more on graphical models see [35]. A Gaussian vector X with covariance matrix Σ is Markov with respect to G if $(\Sigma^{-1})_{ij} = 0$ for all $ij \notin E(G)$. The distributions that are Markov and locally associated with respect to G are denoted by $M_+(G)$. We refer to $M_+(G)$ as a *locally associated Gaussian graphical model (locally associated GGM)*. By definition, we have

$$M_+(G) = A(G) \cap M(G).$$

The set of locally associated Gaussian distributions that are Markov with respect to a graph is given as the intersection of a set that is convex in K with a set that is convex in Σ . The intersection is typically neither convex in K nor in Σ but a locally associated GGM is an instance of what we shall term a *mixed convex exponential family*; see Section 3 below.

2.3. *Positive linear systems and factor models.* In this section, we link locally associated Gaussian graphical models to a broad class of models that includes, for example, factor analysis models with nonnegative loadings.

Recall that a Gaussian model over a directed acyclic graph (DAG) D has the linear structural representation

$$Y_i = \sum_{i \rightarrow j} \lambda_{ij} Y_j + \epsilon_i \quad \text{for all } i \in V,$$

where $(\epsilon_1, \dots, \epsilon_d)$ is a mean-zero vector of Gaussian independent noise terms and $\lambda_{ij} \in \mathbb{R}$. Write $M(D)$ for the class of all such distributions parameterized by $\Lambda = [\lambda_{ij}]$ and the variances of ϵ_i . Moreover, $M_+(D)$ denotes the subset of $M(D)$ where $\Lambda \geq 0$, that is, where all regression coefficients are nonnegative.

PROPOSITION 2.3. *Suppose that the distribution of a zero-mean Gaussian Y lies in $M_+(D)$. Then Y is associated and so is each margin of Y .*

PROOF. Since $(I - \Lambda)Y = \epsilon$ then $\mathbb{V}(Y) = (I - \Lambda)^{-1} \Omega (I - \Lambda)^{-T}$, where Ω is a diagonal matrix with the variances of ϵ on the diagonal and $(I - \Lambda)$ being invertible by acyclicity of D . Since $(I - \Lambda)^{-1} = I + \Lambda + \Lambda^2 + \dots$ and $\Lambda \geq 0$ we conclude that $\mathbb{V}(Y)$ has only nonnegative entries. This concludes the proof. \square

Proposition 2.3 shows that associated distributions contain the interesting family of positive DAG models $M_+(D)$, potentially with some nodes unobserved. Factor analysis models with nonnegative loadings form a particular instance of margins of DAG models of the form $M_+(D)$. Recall that the factor analysis model $\mathbb{F}_{d,s}$ is the family of multivariate Gaussian distributions with an arbitrary mean and whose covariance matrix Σ is of the form $\Sigma = \Delta + \Lambda \Lambda^T$ with a positive diagonal matrix Δ and $\Lambda \in \mathbb{R}^{d \times s}$. We write $\mathbb{F}_{d,s}^+$ if the loading matrix Λ is restricted to have nonnegative entries.

One of the standard arguments for why MTP_2 distributions may be useful in statistical modelling is that they contain the one factor model with nonnegative loadings, $\mathbb{F}_{d,1}^+$. This observation and the corresponding link to the Capital Asset Pricing Model was used in [1]

to argue why MTP_2 distributions are particularly suitable for modelling financial data. However, it is easy to show by explicit examples that distributions in $\mathbb{F}_{d,s}^+$ for $s > 1$ need not be MTP_2 . In this context, the fact that all of them are (globally) associated may provide useful regularization procedures in applications where factor analysis models with nonnegative loadings are expected to perform well; see Section 9 for some evidence.

2.4. Gaussian distributions and Gaussian copulas. Since association is preserved after applying a strictly increasing function $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ to each X_i , our definition of local association naturally extends to Gaussian copula models, as in Proposition 2.4 below. A d -dimensional random vector Y has a nonparanormal distribution if there exist strictly increasing functions $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ for $i = 1, \dots, d$ such that the vector $\phi(Y) := (\phi_1(Y_1), \dots, \phi_d(Y_d))$ has a Gaussian distribution.

PROPOSITION 2.4. *If $G = (V, E)$ is an undirected graph and Y has a nonparanormal distribution, then Y is locally associated with respect to G if and only if $\phi(Y)$ is in $A(G)$. Moreover, Y is Markov with respect to G if and only if $\phi(Y)$ is.*

3. Mixed convex exponential families. It is useful to see locally associated Gaussian graphical models as a special case of a more general type of models. Consider a random variable X with values in a general state-space \mathcal{X} . Suppose that the distribution of X is in a minimally represented steep exponential family $\mathcal{E} = \{P_\theta \mid \theta \in \Theta\}$ with canonical statistic $t : \mathcal{X} \mapsto \mathbb{R}^k$ and canonical parameter θ . This means that the density function $p(x; \theta)$ of the distribution P_θ with respect to some underlying measure ν on \mathcal{X} takes the form

$$(1) \quad p(x; \theta) = \exp\{\langle \theta, t(x) \rangle - A(\theta)\} \quad \text{for } \theta \in \Theta,$$

where $\nu\{x : \langle \lambda, t(x) \rangle = c\} = 0$ if $\lambda \neq 0$. The space of canonical parameters

$$(2) \quad \Theta := \text{int} \left\{ \theta \in \mathbb{R}^k : \int_{\mathcal{X}} \exp\{\langle \theta, t(x) \rangle\} \nu(dx) < \infty \right\}$$

is an open set in \mathbb{R}^k and the cumulant function $A : \Theta \rightarrow \mathbb{R}$ is strictly convex and smooth. The map μ between canonical parameter $\theta \in \Theta$ and the mean parameter $\mu \in M$ satisfies

$$\mu(\theta) = \nabla A(\theta)$$

and establishes a bijection between Θ and M . Moreover, M is the interior of the convex hull of $t(\mathcal{X})$; see any of the references [5, 8, 49] for more details. The inverse map is denoted by θ , that is, $\theta = \theta(\mu)$.

REMARK 3.1. The class of steep exponential families contains a more popular class of regular exponential families, for which the set $\{\theta \in \mathbb{R}^k : \int_{\mathcal{X}} \exp\{\langle \theta, t(x) \rangle\} \nu(dx) < \infty\}$ is already open, and so the interior operation in (2) is redundant.

Suppose we split the sufficient statistics into two subvectors $t(x) = (u, v)$ of dimension r, s where $r + s = k$. Let $\theta = (\theta_u, \theta_v), \mu = (\mu_u, \mu_v)$ be the corresponding splits in the canonical and in the mean parameter. In analogy with $\mu = \mu(\theta)$, we also use the notation $\mu_u(\theta) = \mu_u, \mu_v(\theta) = \mu_v, \theta_u(\mu) = \theta_u,$ and $\theta_v(\mu) = \theta_v$. For example, $\mu_u(\theta)$ is the composition of $\mu : \Theta \rightarrow M$ with the projection $(\mu_u, \mu_v) \mapsto \mu_u$.

By [8], Theorem 3.9 (see also the discussion preceding this theorem and Exercise 3.9.1), the pair (μ_u, θ_v) forms an alternative parametrization for the exponential family \mathcal{E} called the *mixed parametrization*. The parameters μ_u and θ_v are *variation independent*, that is, the parameter space for (μ_u, θ_v) is the Cartesian product space $M_u \times \Theta_v$ where Θ_v is the

projection of Θ on θ_v and M_u is the projection of the space of mean parameters M onto μ_u . So we may without ambiguity write

$$\mathcal{E} = \{P_\theta \mid \theta \in \Theta\} = \{P_\mu \mid \mu \in M\} = \{P_{(\mu_u, \theta_v)} \mid \mu_u \in M_u, \theta_v \in \Theta_v\}.$$

We may thus consider the model \mathcal{E}' obtained from \mathcal{E} by a convex restriction on μ_u and a convex restriction on θ_v . More precisely, we define the following definition.

DEFINITION 3.2. Fix a mixed parametrization $(\mu_u, \theta_v) \in M_u \times \Theta_v$ of the exponential family \mathcal{E} . The model \mathcal{E}' is called a *mixed convex* submodel of \mathcal{E} and a *mixed convex exponential family* if it consists of all mixed parameters $(\mu_u, \theta_v) \in M'_u \times \Theta'_v$, where M'_u and Θ'_v are convex and relatively closed subsets of M_u and Θ_v respectively. We say that \mathcal{E}' is a *mixed linear exponential family*, if both of M'_u and Θ'_v are given by affine restrictions.

It is useful to introduce the following notation:

$$(3) \quad C_u = \{\mu \in M : \mu_u \in M'_u\}, \quad C_v = \{\theta \in \Theta : \theta_v \in \Theta'_v\}.$$

The mixed convex exponential family is then given as an intersection $\mathcal{E}' = \mathcal{E}_u \cap \mathcal{E}_v \subseteq \mathcal{E}$, where $\mathcal{E}_u = \{P_\mu \mid \mu \in C_u\}$ and $\mathcal{E}_v = \{P_\theta \mid \theta \in C_v\}$ (one model is convex in the mean parametrization and the other is convex in the canonical parametrization). Note that these restrictions are also variation independent so that

$$\mathcal{E}' = \{P_{(\mu_u, \theta_v)} \mid \mu \in C_u, \theta \in C_v\} = \{P_{(\mu_u, \theta_v)} \mid (\mu_u, \theta_v) \in M'_u \times \Theta'_v\}.$$

We now discuss a few examples of this. Recall that the family of multivariate Gaussian distributions with zero mean and unknown covariance matrix Σ is a regular exponential family with inner product $\langle A, B \rangle = \text{tr}(AB)$ and

$$t(x) = -xx^T/2, \quad \theta = K, \quad \mu = -\Sigma/2, \quad A(K) = -\frac{1}{2} \log \det K.$$

The space of canonical parameters is the cone of positive definite matrices and the space of mean parameters is the cone of negative definite matrices.

EXAMPLE 3.3. Fix a graph $G = (V, E)$ on d vertices $V = \{1, \dots, d\}$ and consider the family of d -variate mean zero Gaussian distributions. We split the sufficient statistics into $u = (-x_i x_j / 2)_{ij \in E}$ and $v = (-x_i x_j / 2)_{ij \notin E}$. The diagonal entries $-x_i^2 / 2$ are included in u . Then $\mu_u = (-\Sigma_{ij} / 2)_{ij \in E}$ and $\theta_v = (K_{ij})_{ij \notin E}$. We may consider a mixed convex family given by $\mu_u \leq 0$ and $\theta_v = 0$. These are precisely the locally associated Gaussian graphical models discussed in Section 2.2.

Mixed convex exponential families enable easy formulation of other relevant models encoding positive dependence in Gaussian distributions.

EXAMPLE 3.4. With the set-up as in Example 3.3, we alternatively split the sufficient statistics into $u = (-x_i x_j / 2)_{ij \notin E}$ and $v = (-x_i x_j / 2)_{ij \in E}$. The diagonal entries are now included in v . Then $\mu_u = (-\Sigma_{ij} / 2)_{ij \notin E}$ and $\theta_v = (K_{ij})_{ij \in E}$. We may consider a mixed convex family given by $\mu_u \leq 0$ and $\theta_v \leq 0$. Here cliques in the graph correspond to subsystems characterized by a strong notion of positive dependence (all partial correlations nonnegative) or, in other words, the conditional distribution of variables in a clique given the remaining variables is MTP₂. Otherwise the system is weakly positively dependent (positive correlations).

An example of a mixed linear model can be easily motivated by causal analysis [41], where zero restrictions on some entries of Σ correspond to marginal independence and zero restrictions on K correspond to conditional independence. Models of this form fit our general set-up.

EXAMPLE 3.5. Given a graph G_{COV} , called the covariance graph, over $V = \{1, \dots, d\}$, we define the corresponding covariance graph model $B(G_{\text{COV}})$ given by all covariance matrices that satisfy $\Sigma_{ij} = 0$ if $ij \notin E(G_{\text{COV}})$. Given the covariance graph G_{COV} and the concentration graph G_{CON} we want to understand the intersection $B(G_{\text{COV}}) \cap M(G_{\text{CON}})$. In the special case when $ij \notin E(G_{\text{CON}})$ implies $ij \in E(G_{\text{COV}})$, the corresponding model is a mixed linear model. Models of this type were studied in detail in [28], see also [6, 14, 41].

EXAMPLE 3.6. Recently [13] discussed the problem of testing equality of mean zero d -variate Gaussian distributions under the assumption that they lie in a fixed graphical model over the graph G . So suppose that X, Y are independent Gaussian with covariance matrices $\Sigma^{(1)}, \Sigma^{(2)}$. Here the mixed parameters are $(\Sigma_{ij}^{(1)}, \Sigma_{ij}^{(2)})_{ij \in E}$ for the mean part and $(K_{ij}^{(1)}, K_{ij}^{(2)})_{ij \notin E}$ for the canonical part. The mixed linear model which assures that both distributions are equal is given by $\Sigma_{ij}^{(1)} = \Sigma_{ij}^{(2)}$ for all $ij \in E$ and $K_{ij}^{(1)} = K_{ij}^{(2)} = 0$ for all $ij \notin E$.

Although the Gaussian case is our focus here, there are interesting examples beyond this case.

EXAMPLE 3.7. Consider discrete random variables X and Y taking values in $\mathcal{S} = \{0, 1, \dots, k\}$ and let $p_{xy} = P(X = x, Y = y)$ with $p_{xy} > 0$ for all $x, y \in \mathcal{S}$. This specifies a regular exponential family with canonical parameters

$$\lambda_{xy}, \quad x, y \in \mathcal{S} \text{ and } (x, y) \neq (0, 0),$$

where

$$\lambda_{xy} = \log p_{xy} - \log p_{00}$$

and so $\lambda_{00} = 0$. The corresponding sufficient statistics are $t = (t_{xy})$ where

$$t_{xy}(x', y') = \mathbf{1}_{\{(x,y)\}}(x', y')$$

and the corresponding mean parameters are p_{xy} for $(x, y) \neq (0, 0)$. We will consider a linear transformation of the canonical parameters and sufficient statistics into

$$\theta_{xy} = \lambda_{xy} - \lambda_{x0} - \lambda_{0y}, \quad x, y \in \mathcal{S} \setminus \{0\}, \quad \theta_{x0} = \lambda_{x0}, \quad \theta_{0y} = \lambda_{0y}, \quad x, y \in \mathcal{S} \setminus \{0\}$$

with corresponding sufficient statistics

$$t_{xy}, \quad x, y \in \mathcal{S} \setminus \{0\}; \quad t_{x+} = \sum_{y \in \mathcal{S}} t_{xy}, \quad x \in \mathcal{S} \setminus \{0\}, \quad t_{+y} = \sum_{x \in \mathcal{S}} t_{xy}, \quad y \in \mathcal{S} \setminus \{0\}.$$

This exponential family may thus be mixed parametrized with the *marginals*

$$\begin{aligned} \mu_{x+} &= p_{x+} = \mathbb{E}\{t_{x+}(X, Y)\} = \sum_y p_{xy}, \quad x \in \mathcal{S} \setminus \{0\}, \\ \mu_{+y} &= p_{+y} = \mathbb{E}\{t_{+y}(X, Y)\} = \sum_x p_{xy}, \quad y \in \mathcal{S} \setminus \{0\} \end{aligned}$$

and the *interactions*

$$\theta_{xy}, \quad x, y \in \mathcal{S} \setminus \{0\}.$$

We may then consider the hypothesis of *marginal homogeneity* [31], that is,

$$(4) \quad p_{x+} = p_{+x} \quad \text{for all } x \in \mathcal{S}$$

in combination with the distribution being MTP_2 ; the latter is equivalent to the condition

$$(5) \quad \theta_{xy} + \theta_{x'y'} - \theta_{xy'} - \theta_{x'y} \geq 0 \quad \text{for all } x \geq x' \text{ and } y \geq y'.$$

The restriction (4) is convex (in fact linear) in the mean parameters and (5) is convex in θ_{xy} so these restrictions jointly specify a mixed convex exponential family.

Another alternative would exploit that categories are ordered and for example specify that p_{i+} is stochastically smaller than p_{+i} that is,

$$\sum_{x=0}^j p_{x+} \leq \sum_{y=0}^j p_{+y} \quad \text{for all } j \in \mathcal{S},$$

yielding a convex restriction also on the mean parameters; see [2] and [3] for further details of this model.

The mixed parametrization can be naturally used in discrete exponential families when the mean vector is regressed via a link function on some covariates and the remaining canonical parameters are used to handle higher order interactions. Similar ideas emerge for models with restrictions on marginal and conditional distributions; see, for example, [19] and [22].

4. Estimation in mixed convex exponential families. Since mixed convex models are not necessarily convex exponential families (given by convex constraints on θ only), maximum likelihood estimation leads in general to nonconvex optimization problems that may have many local optima. In this section, we propose a simple alternative approach leading to *two* convex optimization problems. In Section 6, we will show that, asymptotically, the resulting estimator has the same asymptotic distribution as the MLE up to the first order in the sense that the difference between the estimators converges in probability to zero, even after multiplying with \sqrt{n} .

4.1. *Likelihood and its dual.* Before we present our optimization procedure, we quickly recall the definition of the dual likelihood function; cf. Chapter 6 in [8]. Given a random sample $X^{(1)}, \dots, X^{(n)}$ of size n from the exponential family \mathcal{E} in (1), denote $\mathbf{t} = \mathbf{t}_n = \sum_{i=1}^n \mathbf{t}(X^{(i)})/n$. The log-likelihood function is a strictly concave function of θ given as

$$(6) \quad \ell(\theta; \mathbf{t}) = \langle \theta, \mathbf{t} \rangle - A(\theta),$$

where we here and in the following have suppressed the index n as we are not yet considering the asymptotic behaviour. Since

$$\nabla \ell(\theta; \mathbf{t}) = \mathbf{t} - \nabla A(\theta) = \mathbf{t} - \mu(\theta),$$

the unconstrained optimum based on data \mathbf{t} is the parameter θ for which the mean parameter $\mu(\theta)$ is equal to \mathbf{t} ; in other words, this is $\theta(\mathbf{t})$ and is well defined if $\mathbf{t} \in M$. In what follows we ignore that \mathbf{t} comes from data and write $\ell(\theta; \mu)$, where μ is a general point in the topological closure \bar{M} of the space of mean parameters.

The *Fenchel conjugate* of the cumulant function A is the function

$$A^*(\mu) = \sup\{\ell(\bar{\theta}; \mu) : \bar{\theta} \in \mathbb{R}^k\}.$$

The function A^* is convex as a supremum of linear functions and indeed strictly convex. The unique optimizer of the log-likelihood is $\theta(\boldsymbol{\mu})$, so

$$(7) \quad A^*(\boldsymbol{\mu}) = \ell(\theta(\boldsymbol{\mu}); \boldsymbol{\mu}) = \langle \theta(\boldsymbol{\mu}), \boldsymbol{\mu} \rangle - A(\theta(\boldsymbol{\mu}))$$

implying in particular that A^* is smooth, that is, infinitely often differentiable, since θ and A are both smooth. For any fixed $\bar{\boldsymbol{\theta}}$, the function

$$(8) \quad \tilde{\ell}(\boldsymbol{\mu}; \bar{\boldsymbol{\theta}}) := \langle \bar{\boldsymbol{\theta}}, \boldsymbol{\mu} \rangle - A^*(\boldsymbol{\mu})$$

is a strictly concave function of $\boldsymbol{\mu} \in M$ called the *dual log-likelihood function*. Analogously to the log-likelihood function, $\tilde{\ell}$ satisfies

$$\nabla_{\boldsymbol{\mu}} \tilde{\ell}(\boldsymbol{\mu}; \bar{\boldsymbol{\theta}}) = \bar{\boldsymbol{\theta}} - \nabla_{\boldsymbol{\mu}} A^*(\boldsymbol{\mu}) = \bar{\boldsymbol{\theta}} - \theta(\boldsymbol{\mu}),$$

which follows by composite differentiation in (7) since $\boldsymbol{\mu}(\theta(\boldsymbol{\mu})) = \boldsymbol{\mu}$ and thus

$$\nabla_{\boldsymbol{\mu}} A^*(\boldsymbol{\mu}) = \theta(\boldsymbol{\mu}) + \nabla_{\boldsymbol{\mu}} \theta(\boldsymbol{\mu}) \cdot \boldsymbol{\mu} - \nabla_{\boldsymbol{\mu}} \theta(\boldsymbol{\mu}) \cdot \boldsymbol{\mu}(\theta(\boldsymbol{\mu})) = \theta(\boldsymbol{\mu}).$$

Consider two distributions in \mathcal{E} , one with the mean parameter $\boldsymbol{\mu}_1$ and the other with canonical parameter $\boldsymbol{\theta}_2$. The *Kullback–Leibler divergence* ([32]) between these two distributions is

$$(9) \quad \mathbf{K}(\boldsymbol{\mu}_1, \boldsymbol{\theta}_2) = -\langle \boldsymbol{\mu}_1, \boldsymbol{\theta}_2 \rangle + A^*(\boldsymbol{\mu}_1) + A(\boldsymbol{\theta}_2),$$

which follows directly from Proposition 6.3 in [8] and the definition of $A^*(\boldsymbol{\mu}_1)$. The Kullback–Leibler divergence is well defined and nonnegative over $M \times \Theta$. Moreover, $\mathbf{K}(\boldsymbol{\mu}_1, \boldsymbol{\theta}_2) = 0$ if and only if $\boldsymbol{\mu}_1 = \boldsymbol{\mu}(\boldsymbol{\theta}_2)$. We will extend the definition of $\mathbf{K}(\cdot, \boldsymbol{\theta})$ to all \mathbb{R}^k by semicontinuity; cf. [8], page 175. Then $\mathbf{K}(\boldsymbol{\mu}, \boldsymbol{\theta})$ may be well defined even if $\boldsymbol{\mu}$ does not lie in the space of mean parameters but in its closure.

The reason to express the Kullback–Leibler distance in terms of $\boldsymbol{\mu}_1$ and $\boldsymbol{\theta}_2$ rather than $\boldsymbol{\theta}_1$, $\boldsymbol{\theta}_2$ (as usually done in the literature) is that we wish to exploit the following basic result.

PROPOSITION 4.1. *The Kullback–Leibler divergence $\mathbf{K}(\boldsymbol{\mu}_1, \boldsymbol{\theta}_2)$ is strictly convex both in $\boldsymbol{\mu}_1$ and in $\boldsymbol{\theta}_2$.*

PROOF. This follows directly from (9) and the fact that both $A(\boldsymbol{\theta})$ and $A^*(\boldsymbol{\mu})$ are strictly convex functions. \square

REMARK 4.2. Minimizing $\mathbf{K}(\boldsymbol{t}, \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ with \boldsymbol{t} fixed is equivalent to maximizing $\ell(\boldsymbol{\theta}; \boldsymbol{t})$ in (6). Similarly, minimizing $\mathbf{K}(\boldsymbol{\mu}, \bar{\boldsymbol{\theta}})$ with respect to $\boldsymbol{\mu}$ with $\bar{\boldsymbol{\theta}}$ fixed is equivalent to maximizing the dual log-likelihood $\tilde{\ell}(\boldsymbol{\mu}; \bar{\boldsymbol{\theta}})$ in (8).

4.2. The mixed dual estimator. Recall the definition of the sets C_u, C_v in (3) and that $\boldsymbol{t} = \sum_{i=1}^n \boldsymbol{t}(X^{(i)})/n$. We propose the following two-step procedure to estimate the mixed parameter $(\boldsymbol{\mu}_u, \boldsymbol{\theta}_v)$ in the mixed convex family \mathcal{E}' from data \boldsymbol{t} :

(S1) First minimize $\mathbf{K}(\boldsymbol{t}, \boldsymbol{\theta})$ over $\boldsymbol{\theta} \in C_v \subseteq \Theta$. Denote the *unique* optimum, assuming it exists, by $\hat{\boldsymbol{\theta}}$.

(S2) Then minimize $\mathbf{K}(\boldsymbol{\mu}, \hat{\boldsymbol{\theta}})$ subject to $\boldsymbol{\mu} \in C_u \subseteq M$. Denote the unique optimum by $\tilde{\boldsymbol{\mu}}$.

We shall term our estimator $\tilde{\boldsymbol{\mu}}$ the *mixed dual estimator* (MDE) and show below in Theorem 4.5 that indeed $P_{\tilde{\boldsymbol{\mu}}} \in \mathcal{E}'$. By Proposition 4.1, both steps (S1) and (S2) rely on solving a convex optimization problem. Note that the optimum in (S1) is the MLE under the convex exponential family given by $\boldsymbol{\theta} \in C_v$ (cf. Remark 4.2). This MLE may not exist (if \boldsymbol{t} lies on the boundary of the space of mean parameters) but we have the following proposition.

PROPOSITION 4.3. *If the optimum $\widehat{\theta}$ in (S1) exists, then it is unique and the optimum $\widetilde{\mu}$ in (S2) exists and is unique too.*

PROOF. Suppose that the optimum in (S1) exists. Uniqueness follows by strict convexity. Note that (S2) is equivalent to the maximization of the dual likelihood $\check{\ell}(\mu, \widehat{\theta})$ over all $\mu \in C_u$. Let $\widetilde{M} = \{\mu : A^*(\mu) < \infty\}$, which contains M but is typically bigger, $M \subseteq \widetilde{M} \subseteq \overline{M}$ where \overline{M} is the topological closure of M . Let S be the closure (in \widetilde{M}) of C_u . Since C_u is relatively closed in M , the only extra points are those in $\widetilde{M} \setminus M$. Now, by Theorem 6.13 in [8], the fact that $\widehat{\theta} \in \Theta$ implies that $\check{\ell}(\mu, \widehat{\theta})$ attains an optimum over S and the optimum is uniquely attained in $S \cap M = C_u$. In particular, the optimum in (S2) exists. (Note that the extra points we added to the model played only an auxiliary role in this proof.) \square

REMARK 4.4. Let $t = (u, v)$. Theorem 4.5 below implies that $u \in M_u$ is a necessary condition for $\widehat{\theta}$ to exist.

Our main result of this section shows that after the steps (S1) and (S2), we indeed obtain a point in the mixed convex family \mathcal{E}' .

THEOREM 4.5. *Let $t = (u, v)$ and suppose that $\widehat{\theta}$ in step (S1) exists. Then, $\mu_u(\widehat{\theta}) = u \in M_u$ and in step (S2) we get that $\theta_v(\widetilde{\mu}) = \widehat{\theta}_v$. In particular, after steps (S1) and (S2), the optimum $\widetilde{\mu}$ represents an element in the mixed convex family \mathcal{E}' .*

PROOF. If the optimum $\widehat{\theta} \in C_v$ in step (S1) exists then, by convexity of $\mathbf{K}(t, \theta)$ and of C_v , it must satisfy $\langle \nabla_{\theta} \mathbf{K}(t, \widehat{\theta}), \theta - \widehat{\theta} \rangle \geq 0$ for all $\theta \in C_u$. Since Θ is an open set, a small perturbation $\widehat{\theta} + \tau$ also lies in Θ . If, in addition, $\tau_v = \mathbf{0}$ then this perturbation lies in C_v . Thus, for any sufficiently small vector $\tau \in \mathbb{R}^k$ such that $\tau_v = \mathbf{0}$, we must have

$$\langle \nabla_{\theta} \mathbf{K}(t, \widehat{\theta}), \tau \rangle = \langle \mu_u(\widehat{\theta}) - u, \tau_u \rangle \geq 0.$$

Since τ_u is small but otherwise arbitrary, this is only possible if $\mu_u(\widehat{\theta}) = u$, which proves the first part of the theorem.

The second part is proved in the same way: In step (S2), the optimum $\widetilde{\mu}$ exists and is unique by Proposition 4.3. By convexity of $\mathbf{K}(\mu, \widehat{\theta})$, the optimum satisfies $\langle \nabla_{\mu} \mathbf{K}(\widetilde{\mu}, \widehat{\theta}), \mu - \widetilde{\mu} \rangle \geq 0$ for all $\mu \in C_u$. Since $\widetilde{\mu}$ is an interior point of M , a small perturbation $\widetilde{\mu} + \tau$ also lies in M and, if $\tau_u = \mathbf{0}$, it also lies in C_u . For any such perturbation, we necessarily have

$$\langle \nabla_{\mu} \mathbf{K}(\widetilde{\mu}, \widehat{\theta}), \tau \rangle = \langle \theta_v(\widetilde{\mu}) - \widehat{\theta}_v, \tau_v \rangle \geq 0.$$

Since τ_v is small but otherwise arbitrary, this is only possible if $\theta_v(\widetilde{\mu}) = \widehat{\theta}_v$. \square

REMARK 4.6. In principle, we may interchange the order of optimization with respect to θ and μ by starting from $\theta(t)$ and running Step (S2) first. However, if t does not lie in the space of mean parameters (but in its closure) $\theta(t)$ is not well defined. Here we exploit that $\widehat{\theta}$ might exist even if t is on the boundary of M due to the additional convex restriction $\theta \in C_v$.

REMARK 4.7. Note that when C_v is given by affine constraints, the first step just corresponds to reducing data by sufficiency; then the MDE is simply the *dual likelihood estimate* (DLE), studied by [8, 10, 15] and used extensively by [30]. This provides, for example, a straight-forward way to test equality of distributions in Example 3.6 because the likelihood ratio statistics based on the dual likelihood has asymptotically the same distribution as the standard likelihood ratio statistics, cf. [10], Theorem 3.3.

5. Estimating Gaussian locally associated distributions. As mentioned in Section 2.2, a locally associated Gaussian graphical model $M_+(G)$ is determined as

$$M_+(G) = A(G) \cap M(G)$$

and thus it forms an instance of a mixed convex exponential family since $A(G)$ is convex in Σ and $M(G)$ is convex in K and the restrictions refer to distinct parts of the canonical statistic and parameters; see Example 3.3 for details.

5.1. *The Gaussian log-likelihood.* Given the data $X \in \mathbb{R}^{n \times d}$ with independent rows distributed as $N(0, \Sigma^*)$ our goal is to estimate Σ^* . Let $S = X^T X/n$ be the sample covariance matrix. The Gaussian log-likelihood (ignoring a multiplicative constant n) is

$$(10) \quad \ell(K) = \frac{1}{2} \log \det(K) - \frac{1}{2} \text{tr}(SK).$$

Note that here and elsewhere we ignore the multiplicative constant n in the log-likelihood function. Recall that for the Gaussian family

$$t(x) = -xx^T/2, \quad \theta = K, \quad \mu = -\Sigma/2, \quad A(K) = -\frac{1}{2} \log \det K.$$

The unique optimizer of $\ell(K)$ is S^{-1} and so

$$A^*(-\Sigma/2) = -\frac{1}{2} \log \det(\Sigma) - \frac{d}{2}$$

and thus the Kullback–Leibler divergence as above becomes

$$\mathbf{K}(-\Sigma_1/2, K_2) = \frac{1}{2} \text{tr}(\Sigma_1 K_2 - I) - \frac{1}{2} \log \det(\Sigma_1 K_2).$$

5.2. *The mixed dual estimator.* We consider now the mixed dual estimator MDE as developed above—in this case equivalent to the dual likelihood estimator, since the first step is just estimation in the standard graphical model given by edge restrictions on K (see, e.g., [35], Section 5.2.1).

The second step corresponds to the following convex optimization problem:

$$(11) \quad \begin{aligned} & \underset{\Sigma \geq 0}{\text{minimize}} \quad -\log \det(\Sigma) + \text{tr}(\Sigma \widehat{K}) \\ & \text{subject to} \quad \Sigma_{ij} \geq 0 \text{ for all } ij \in E(G), \end{aligned}$$

where, as before, \widehat{K} is the MLE of K in the linear exponential family $M(G)$. Note that we ignore the Markov constraint on Σ that would destroy convexity of this problem. As shown in Theorem 4.5 above, the Markov constraint becomes automatically fulfilled in the optimization. The Lagrangian for this problem is

$$L(\Sigma, \Lambda) = -\log \det \Sigma + \text{tr}(\Sigma \widehat{K}) - \text{tr}(\Sigma \Lambda) = -\log \det \Sigma + \text{tr}(\Sigma(\widehat{K} - \Lambda)),$$

where Λ is a symmetric matrix with diagonal equal to 0, $\Lambda_{ij} = 0$ unless $ij \in E(G)$, and $\Lambda_{ij} \geq 0$. The Lagrangian is minimized in Σ for fixed Λ by the matrix Σ^* given as

$$(12) \quad \Sigma^* = (\widehat{K} - \Lambda)^{-1}$$

and complementary slackness implies for the optimal $\Lambda = \Lambda^*$ that

$$\Sigma_{ij}^* \Lambda_{ij}^* = 0, \quad ij \in E$$

or, equivalently,

$$\Sigma_{ij}^*(\widehat{K}_{ij} - K_{ij}^*) = 0, \quad ij \in E.$$

Note that it follows directly from (12) that $\Sigma^* \in M(G)$ if $\widehat{\Sigma} \in M(G)$, conforming with Theorem 4.5 and hence $\Sigma^* = \widetilde{\Sigma}$ is in fact the MDE we are looking for. In summary, we have the following theorem.

THEOREM 5.1. *If the MDE $\widetilde{\Sigma}$ of Σ under $M_+(G)$ exists, it is given as the unique positive definite solution to the following system of equations and inequalities, where $\widetilde{K} = \widetilde{\Sigma}^{-1}$ and $\widehat{K} = \widehat{\Sigma}^{-1}$:*

- (i) $\widetilde{\Sigma}_{ij} \geq 0, ij \in E(G),$
- (ii) $\widetilde{\Sigma}_{ij} = S_{ij}, ij \in E(G),$
- (iii) $\widetilde{\Sigma}_{ii} = S_{ii}, i \in V(G),$
- (iv) $\widetilde{K}_{ij} = \widehat{K}_{ij} = 0, ij \notin E(G),$
- (v) $\widetilde{K}_{ij} \leq \widehat{K}_{ij}, ij \in E(G),$
- (vi) $\widetilde{K}_{ii} = \widehat{K}_{ii}, i \in V(G),$
- (vii) $\widetilde{\Sigma}_{ij}(\widehat{K}_{ij} - \widetilde{K}_{ij}) = 0, ij \in E(G).$

Note that the equations (ii), (iii), and the last part of (iv) are equations determining the MLE \widehat{K} under $M(G)$. The condition (vii) naturally induces sparsity in $\widetilde{\Sigma}$.

Further, denoting by \widetilde{G} the graph whose edges correspond to nonzero entries of $\widetilde{\Sigma}$, we note that the zero entries in Σ obtained in this way are complementary to the zero entries in K ; cf. also Example 3.5. We also have the following.

COROLLARY 5.2. *The mixed dual estimate $\widetilde{\Sigma}$ in the model $B(\widetilde{G}) \cap M(G)$ is identical to the mixed dual estimate in the model $M_+(G) = A(G) \cap M(G)$ as determined in Theorem 5.1 above.*

PROOF. By construction, if $ij \notin E(\widetilde{G})$ then $ij \in E(G)$. Denoting by E_0 the set of pairs that do not lie in $E(\widetilde{G})$ we get that $E_0 \subseteq E(G)$. The optimality conditions for \widehat{K} over $M(G)$ are standard and given by:

- (a) $\widehat{\Sigma}_{ii} = S_{ii}, i \in V(G),$
- (b) $\widehat{\Sigma}_{ij} = S_{ij}, ij \in E(G),$
- (c) $\widehat{K}_{ij} = 0, ij \notin E(G).$

By an analogous argument, the dual likelihood estimate in $B(\widetilde{G})$ based on \widehat{K} is the unique positive definite matrix $\widetilde{\Sigma} = \widetilde{K}^{-1}$ satisfying:

- (d) $\widetilde{K}_{ii} = \widehat{K}_{ii}, i \in V(G),$
- (e) $\widetilde{K}_{ij} = \widehat{K}_{ij}, ij \in E(\widetilde{G}),$
- (f) $\widetilde{\Sigma}_{ij} = 0, ij \in E_0.$

Our aim is to show that $\widetilde{\Sigma} = \widehat{\Sigma}$. First note that, (d) together with condition (vi) in Theorem 5.1, implies that $\widetilde{K}_{ii} = \widehat{K}_{ii}$ for all $i \in V$. Similarly, (iv) and (e) imply that $\widetilde{K}_{ij} = \widehat{K}_{ij} = \widetilde{K}_{ij}$ for $ij \notin E(G)$. This equality extends to all $ij \in E(G) \setminus E_0$ by (vii) and (e). Moreover, $\widetilde{\Sigma}_{ij} = \widehat{\Sigma}_{ij} = 0$ for all $ij \in E_0$. It is a standard result that there is a unique completion of a partially specified positive definite matrix K to a matrix such that $\Sigma = K^{-1}$ has zero entries on the unspecified entries of K . It follows that $\widetilde{K} = \widehat{K}$. \square

Note that the statement in Corollary 5.2 is not trivial since $B(\tilde{G}) \cap M(G)$ is not a subset of $M_+(G)$.

The conditions for existence of $\tilde{\Sigma}$ do not seem to simplify beyond the conditions for existence of \tilde{K} ; see also Theorem 4.5. Conditions for existence of \tilde{K} may be rather involved; see, for example, [23, 50].

6. Asymptotic behaviour of the mixed dual estimator. We now return to the the general mixed dual estimator. In this section, we work entirely in the corresponding mixed parametrization $\psi = (\mu_u, \theta_v)$. The MLE and the MDE for \mathcal{E}' in this parametrization are denoted simply by $\tilde{\psi}$ and $\tilde{\psi}$. The maximizer of the log-likelihood function obtained in step (S1) is denoted $\tilde{\psi}$. Suppose that the true data generating distribution with the mixed parameter ψ_0 lies in \mathcal{E}' , that is, $\psi_0 \in M'_u \times \Theta'_v$. We study the asymptotic distribution of the MDE $\tilde{\psi}_n$, where n is the sample size and show that the MDE is \sqrt{n} -consistent and has exactly the same asymptotic distribution as the maximum likelihood estimator $\tilde{\psi}_n$.

THEOREM 6.1. *The MDE and MLE are asymptotically equivalent, that is, it holds that $\sqrt{n}(\tilde{\psi}_n - \tilde{\psi}_n) \rightarrow 0$ in probability, implying that $\sqrt{n}(\tilde{\psi}_n - \psi_0)$ converges in distribution and the limiting distribution of $\sqrt{n}(\tilde{\psi}_n - \psi_0)$ equals the limiting distribution of $\sqrt{n}(\tilde{\psi}_n - \psi_0)$.*

PROOF. The proof is provided in Appendix A. \square

Let $\bar{\psi}_n = (\mathbf{u}_n, \theta_v(\mathbf{t}_n))$ be the MLE in the unrestricted exponential family \mathcal{E} expressed in the mixed parametrization. The limiting distribution in Theorem 6.1 is obtained by projecting the Gaussian limiting distribution of $\sqrt{n}(\bar{\psi}_n - \psi_0)$ onto the tangent cone of the mixed exponential family at the true parameter ψ_0 ; cf. [21]. If the constraints defining Θ'_v are affine, it is useful to equivalently describe this distribution as the limiting distribution of $\sqrt{n}(\hat{\psi}_n - \psi_0)$ (which is also Gaussian), onto the tangent cone $T_{M'_u \times \Theta'_v}(\psi_0)$. In the case of locally associated Gaussian graphical models, this results with the Gaussian distribution (the asymptotic distribution of the MLE in a Gaussian graphical model; cf. [35], Proposition 5.8) projected onto the cone given by the edge covariances being nonnegative. The problem is that even if we can describe this distribution exactly, it depends in a very complicated way on the true covariance matrix. It will be given as a mixture of normal distributions that are ‘truncated’ to regions projecting onto the various facets with weights that are generally impossible to compute exactly. As an alternative to using asymptotic results, the distribution of the MDE may be simulated using bootstrap methods, as the estimation algorithm is fast and guaranteed to converge, whereas simulating the distribution of the MLE is difficult as the MLE may not even be well defined for all bootstrap samples.

7. Learning the local structure. In this section, we consider the situation where the graph G defining the local structure in locally associated Gaussian graphical models is unknown. We aim at obtaining a sparse structure in K through a lasso-type penalty.

7.1. The positive graphical lasso. To avoid losing any positive dependence information we only penalize positive values in the inverse covariance matrix, corresponding to large negative partial correlations. More precisely, we want to solve the following optimization problem

$$(13) \quad \text{minimize } -2\ell(K) + \rho \|K^+\|_1 \quad \text{subject to } K \succ 0,$$

where $\ell_n(K)$ is the Gaussian log-likelihood in (10) and

$$\|K^+\|_1 = \sum_{i \neq j} \max\{0, K_{ij}\}.$$

We shall refer to this procedure as the *positive graphical lasso*. Note that for $\rho = \infty$, the penalty forces the solution \widehat{K}^ρ to be an M-matrix and hence the positive graphical lasso can be seen as a direct relaxation of the estimation under the assumption that the distribution is MTP_2 ([33, 46]); cf. Section 8.4.

By computing subgradients, we easily check that \widehat{K}^ρ is the unique optimal point of (13) and only if

$$(14) \quad \widehat{\Sigma}_{ij}^\rho - S_{ij} \in \begin{cases} \{0\} & \text{if } \widehat{K}_{ij}^\rho < 0, \\ [0, \rho] & \text{if } \widehat{K}_{ij}^\rho = 0, \\ \{\rho\} & \text{if } \widehat{K}_{ij}^\rho > 0 \end{cases} \quad \text{for all } i \neq j \in V.$$

As a corollary we get an alternative characterization of the optimal solution that links it to the MLE in the Gaussian graphical model.

COROLLARY 7.1. *Let $\widehat{G}^\rho = (V, \widehat{E}^\rho)$ be the graph with edges determined by \widehat{K}^ρ and define the modified sample covariance as*

$$S_{ij}^\rho = \begin{cases} S_{ij} & \text{if } \widehat{K}_{ij}^\rho \leq 0, \\ S_{ij} + \rho & \text{if } \widehat{K}_{ij}^\rho > 0 \end{cases} \quad \text{for all } i \neq j \in V.$$

Then \widehat{K}^ρ is the MLE under the Gaussian graphical model determined by \widehat{G}^ρ , based on the modified sample covariance S^ρ .

PROOF. The MLE is uniquely determined by fitting covariances along edges in \widehat{E}^ρ and satisfying $\widehat{K}_{ij} = 0$ for nonedges. \square

See also Proposition 8.2 and Corollary 8.10 below for generalizations of this result.

The positive graphical lasso estimate, as described and calculated above, will avoid large negative partial correlations and as such it may often directly result in a locally associated covariance matrix, in particular for large penalty parameters, as shown in Theorem 8.9. If this is not the case, we may wish to take the additional restriction of edge positivity into account using the estimate $\widetilde{\Sigma}^\rho$ given as

$$\widetilde{\Sigma}^\rho = \underset{\Sigma \in M_+(\widehat{G}^\rho)}{\operatorname{arg\,min}} \mathbf{K}(-\Sigma/2, \widehat{K}^\rho) = \underset{\Sigma \in A(\widehat{G}^\rho)}{\operatorname{arg\,min}} \mathbf{K}(-\Sigma/2, \widehat{K}^\rho).$$

This is exactly the dual likelihood estimate in (11) calculated with \widehat{K}^ρ as starting point, rather than \widehat{K} . We may then again think of the two-step procedure as first obtaining a compact representation \widehat{K}^ρ of the data matrix S , adapting and taking into account the penalty for negative partial correlations, and subsequently approximating this compact representation with a locally associated, and hence locally associated Gaussian distribution $\widetilde{\Sigma}^\rho$. We refer to this procedure as the *dual penalized likelihood estimate* (DPLE).

7.2. A comment on high-dimensional analysis. A careful analysis of the high-dimensional properties of the positive graphical lasso estimator is beyond the scope of this paper. Here we share some cautionary remarks. In [47] the authors analysed the convergence rates for the operator norm $\|\widehat{K} - K^*\|$ for the problem of M-matrix estimation. As we argue in Section 8.4 below, M-matrix estimation is a special case of our positive graphical LASSO

set-up. In particular, Section 3 in [47] suggests that we cannot expect good rates for the operator norm $\|\widehat{K}^\rho - K^*\|$ if K^* is sparse. So in the high-dimensional regime, the spectral properties $\widehat{\Sigma}^\rho$ should be interpreted with caution. Similarly, \widehat{K} will not have good support recovery. For example, if $K^* = I_d$, \widehat{K} will typically not even be sparse. A natural way to obtain an estimator with better statistical properties is by replacing the sample covariance matrix S in (13) with a better estimator of Σ^* (e.g., shrinkage estimator). Our motivating example is high-dimensional with $d = 136$ and $n = 43$. In this case, the positive graphical lasso estimator outperforms the graphical lasso estimator by a large margin also producing a much sparser graph; cf. Section 9.2.

8. The GOLAZO algorithm. In this section, we formulate a general optimization problem and algorithm that unifies the positive graphical lasso, estimation under MTP_2 constraints, under local association, and a number of other forms of likelihood-based estimation for graphical models. It allows us to flexibly introduce sign constraints, zero restrictions, and to penalize different signs of K_{ij} at different rates.

Let L, U be two symmetric matrices with entries in $\mathbb{R} \cup \{-\infty, +\infty\}$ with the restriction that $L_{ij} \leq 0 \leq U_{ij}$ for all $i \neq j$, and $L_{ii} = U_{ii} = 0$ for all i . Denote

$$\|K\|_{LU} := \sum_{i \neq j} \max\{L_{ij}K_{ij}, U_{ij}K_{ij}\}.$$

The function $\|K\|_{LU}$ is convex, positively homogeneous, continuous, and nonnegative. Although it is sublinear, that is $\|K + K'\|_{LU} \leq \|K\|_{LU} + \|K'\|_{LU}$, it does not define a norm unless $|L_{ij}| = |U_{ij}|$ for all $i \neq j$.

The penalty function $\|K\|_{LU}$ enables us to obtain sparse estimates of K in a way that takes into account the signs of K or, equivalently, the signs of the partial correlations. We aim at solving the following problem

$$(15) \quad \text{minimize } -2\ell(K) + \|K\|_{LU},$$

and refer to this as *Graphical Oriented LAZY Optimization (GOLAZO)*. To get a procedure that is invariant under diagonal rescaling, we also typically replace the sample covariance matrix S in $\ell_n(K)$ with the sample correlation matrix R .

REMARK 8.1. For nonparanormal distributions, we replace the sample correlation matrix R with another estimate of the correlation matrix. Following the SKEPTIC approach of [37], we first compute Kendall’s tau $\widehat{\tau}_{ij}$, which can be estimated without knowledge of the underlying monotone transformations f_i . Then we compute $\widehat{\rho}_{ij} = \sin(\frac{\pi}{2}\widehat{\tau}_{ij})$, which is a natural plug-in estimator of the correlation based on the main result of [36].

To illustrate usefulness of this general approach, we list a number of situation that are included in this set-up.

Graphical lasso and SCAD penalties: If $|L_{ij}| = |U_{ij}| = \rho > 0$ for all $i \neq j$, (15) corresponds to the standard graphical lasso. More generally, if $|L_{ij}| = |U_{ij}| = \rho_{ij}$, that is when $\|K\|_{LU}$ is a norm, we obtain a version of the graphical lasso that takes into account different scalings of the variables. This general version is used in the adaptive GLASSO procedure and the local linear approximation algorithm used for general concave penalties rely on solving one or more problems of this form; see [18] for details.

Asymmetric graphical lasso: If $L_{ij} = -\rho_-$ and $U_{ij} = \rho_+$ for all $i \neq j$ where $0 < \rho_-, \rho_+ < +\infty$, we obtain a version of the graphical lasso, where positive entries of K are penalized at a different rate than the negative entries.

Positive graphical lasso: If $L = 0$ and $U_{ij} = \rho$ for all $i \neq j$ then (15) corresponds to the positive graphical lasso problem in (13). This looks like the asymmetric graphical lasso

problem with $\rho_- = 0$ but as we will see, a zero penalty introduces additional complications concerning existence of the optimum.

MTP₂ distributions: If $L = 0$ and $U_{ij} = +\infty$ for all $i \neq j$, then (15) gives the maximum likelihood estimator for constrained M-matrix estimation. In Remark 8.3 we show that, rather than infinite, U_{ij} must be sufficiently large.

Gaussian graphical models: In certain situations, we may in advance wish to specify that some entries of K should be zero. If $K_{ij} = 0$, this can be imposed by setting $L_{ij} = -\infty$, $U_{ij} = +\infty$ (by the standard convention $\pm\infty \cdot 0 = 0$). Thus, the optimization algorithm detailed in Section 8.1 also yields an interesting alternative to the IPS algorithm and other edge based algorithms in [48] which may occasionally be slow.

Dual likelihood estimate: The optimization problem in (11) is equivalent to (15) just replacing K with Σ , S with \widehat{K} , setting $L_{ij} = -\infty$ and $U_{ij} = 0$.

PROPOSITION 8.2. *If the optimum in (15) exists, it is the unique positive definite matrix \widehat{K} (with $\widehat{\Sigma} = \widehat{K}^{-1}$) satisfying*

$$\widehat{\Sigma}_{ij} - S_{ij} \in \begin{cases} \{L_{ij}\} & \text{if } \widehat{K}_{ij} < 0, \\ [L_{ij}, U_{ij}] & \text{if } \widehat{K}_{ij} = 0, \\ \{U_{ij}\} & \text{if } \widehat{K}_{ij} > 0 \end{cases} \quad \text{for all } i, j \in V.$$

PROOF. The subgradient of the function $\max\{L_{ij}K_{ij}, U_{ij}K_{ij}\}$ at $K_{ij} = 0$ is the interval $[L_{ij}, U_{ij}]$. This subgradient is $\{L_{ij}\}$, $\{U_{ij}\}$ if $K_{ij} < 0$, $K_{ij} > 0$ respectively. Now the conclusion follows from the standard theory for nondifferentiable convex functions; see, for example, [43]. \square

The problem (15) is a convex optimization problem. Its dual problem is particularly simple and admits a straight-forward block-coordinate descent procedure. To obtain the dual, note that

$$\max\{L_{ij}K_{ij}, U_{ij}K_{ij}\} = \sup_{L_{ij} \leq \Gamma_{ij} \leq U_{ij}} \Gamma_{ij}K_{ij}$$

and so

$$\|K\|_{LU} = \sup_{L \leq \Gamma \leq U} \text{tr}(\Gamma K),$$

whereby (15) becomes

$$\inf_{K > 0} \sup_{L \leq \Gamma \leq U} \{-\log \det K + \text{tr}((S + \Gamma)K)\}.$$

Swapping inf with sup and using the fact that the unique optimizer with respect to K of $-\log \det K + \text{tr}((S + \Gamma)K)$ (if exists) is $(S + \Gamma)^{-1}$, we obtain the dual problem by letting $\Sigma = S + \Gamma$:

$$(16) \quad \text{maximize } \log \det \Sigma + d \quad \text{subject to } S + L \leq \Sigma \leq S + U.$$

In particular, every feasible point of the dual problem (16) has the same diagonal as S .² In particular, if the correlation matrix R is used as the data input, the optimum is a correlation matrix too.

²Our setting can be easily extended to the case when the diagonal entries of K are also penalized. In this case the optimal point of the dual problem has the same diagonal as $S + U$.

REMARK 8.3. Since Σ is positive definite we have in particular that $-\sqrt{S_{ii}S_{jj}} < \Sigma_{ij} < \sqrt{S_{ii}S_{jj}}$ for all $i \neq j$. It follows that every dually feasible Σ satisfies

$$\max\{S_{ij} + L_{ij}, -\sqrt{S_{ii}S_{jj}}\} \leq \Sigma_{ij} \leq \min\{S_{ij} + U_{ij}, \sqrt{S_{ii}S_{jj}}\}.$$

This allows us to replace U_{ij} with $U_{ij} \wedge (\sqrt{S_{ii}S_{jj}} - S_{ij})$ and L_{ij} with $L_{ij} \vee (-S_{ij} - \sqrt{S_{ii}S_{jj}})$, which is particularly useful if $L_{ij} = -\infty$ or $U_{ij} = +\infty$.

COROLLARY 8.4. Let \widehat{K} be the optimal solution to (15). If $L_{ij} \leq -S_{ij} - \sqrt{S_{ii}S_{jj}}$ then $\widehat{K}_{ij} \geq 0$. If $U_{ij} \geq -S_{ij} + \sqrt{S_{ii}S_{jj}}$ then $\widehat{K}_{ij} \leq 0$. In particular, $\widehat{K}_{ij} = 0$ if both conditions hold.

Since (15) is always feasible, feasibility of the dual problem (16) assures that the optimum of (15) exists and is unique. We show below that it always holds if L and U have no zeros outside of the diagonal. Under minor conditions this also holds for the positive graphical lasso in which case L is a zero matrix. We provide a more detailed treatment of this problem in Section 8.2. But first we introduce our optimization algorithm.

8.1. *Optimization.* To solve (16), we use a straightforward block-coordinate descent approach that is a direct modification of the algorithm used for the dual graphical lasso problem in [4]. An important difference is that, by default, we do not penalize the diagonal, which leads to additional issues that may arise. We optimize the determinant of Σ updating row by row, but keep the diagonal entries fixed to be equal to the diagonal of S .

For the j th row we consider $\log \det \Sigma$ as the function of $\Sigma_{j, \setminus j}$ keeping the other entries of Σ fixed. By standard matrix algebra

$$\log \det \Sigma = \log \det \Sigma_{\setminus j} + \log \det(\Sigma_{jj} - \Sigma_{j, \setminus j}(\Sigma_{\setminus j})^{-1}\Sigma_{\setminus j, j}).$$

Thus, maximizing $\log \det \Sigma$ with respect to $y := \Sigma_{j, \setminus j}$ is equivalent to minimizing $y^T \times (\Sigma_{\setminus j})^{-1}y$, where we need to impose linear conditions that $S_{ij} + L_{ij} \leq y_i \leq S_{ij} + U_{ij}$ for every $i \in V \setminus \{j\}$. This is an instance of a quadratic program. The following result is a straightforward generalization of [4], Theorem 4. It allows to quickly identify disconnected nodes in the underlying graph.

LEMMA 8.5. If $S_{j, \setminus j} + L_{j, \setminus j} \leq 0 \leq S_{j, \setminus j} + U_{j, \setminus j}$ for some $j \in V$, then $\widehat{\Sigma}_{j, \setminus j} = \widehat{K}_{j, \setminus j} = 0$.

The starting point Σ^0 of the algorithm needs to be chosen carefully so that Σ^0 is dually feasible. In this case, each iterate of the algorithm is guaranteed to be dually feasible. In Section 8.2, we show how to find such a starting point. Given the starting point, our procedure is straightforward and described in Algorithm 1. Corollary 8.4 and Lemma 8.5 give an obvious way to speed up the computations but reducing the number of nodes that have to be visited at each step. To solve the quadratic problem in each iteration, we use the `quadprog` package in R.

To establish convergence in Algorithm 1, we track the duality gap

$$\text{tr}(SK) - d + \|K\|_{LU},$$

which is guaranteed to be nonnegative for each step of the algorithm, decrease at each iteration, and to be zero at the optimum. We stop the algorithm once this positive gap becomes sufficiently close to zero.

In the actual implementation, it is important to compute the dual gap carefully in case L, U contain infinite entries. We simply make use of Remark 8.3 and replace $\pm\infty$ with appropriate finite bounds. The only remaining issue to complete the description of the algorithm is a procedure to obtain a dually feasible starting point. We address this issue in the next section.

Data: A positive semidefinite matrix S , penalty matrices $L \leq 0 \leq U$.

Result: A maximizer of (13).

Initialize: $\Sigma = \Sigma^0$ (a dually feasible point);

while there is no convergence **do**

for $j = 1, \dots, d$ **do**

 Update $\Sigma_{j,\setminus j} \leftarrow \hat{y}$, where

$$\hat{y} = \arg \min_y \{y^T (\Sigma_{\setminus j})^{-1} y : S_{j,\setminus j} + L_{j,\setminus j} \leq y \leq S_{j,\setminus j} + U_{j,\setminus j}\}.$$

end

end

Algorithm 1: The Graphical Oriented LAZY Optimization (GOLAZO) algorithm

8.2. *Dually feasible starting point.* Recall that Σ is dually feasible if Σ is positive definite and $L \leq \Sigma - S \leq U$. If S is positive definite, it is dually feasible, and we take $\Sigma^0 = S$. We then focus on the case when S is rank deficient. Since S is a sample covariance matrix it is always positive semidefinite and it has rank at least one. This implies that (with probability one) the diagonal entries are strictly positive. Clearly, if some $S_{ii} = 0$ then no feasible point exists. We then always assume that $S_{ii} > 0$ for all $i \in V$.

We first show how to construct a starting point in the case when for each $i \neq j$ both the negative and the positive values of K_{ij} are penalized. Denote by $\text{diag}(S)$ the diagonal matrix whose diagonal is equal to the diagonal of S .

LEMMA 8.6. *If $L_{ij} < 0 < U_{ij}$ for all $i \neq j$, then there exists $t \in (0, 1)$ such that $\Sigma^0 = (1 - t)S + t \text{diag}(S)$ is dually feasible. The condition for dual feasibility becomes that $L_{ij} \leq -tS_{ij} \leq U_{ij}$ for all $i \neq j$.*

PROOF. Since $\text{diag}(S)$ is positive definite, $\Sigma^0 = (1 - t)S + t \text{diag}(S)$ is positive definite for every $t \in (0, 1]$. We have $\Sigma^0 - S = t(\text{diag}(S) - S)$ and so $(\Sigma^0 - S)_{ij} = -tS_{ij}$ for all $i \neq j$. Since $L_{ij} < 0$ and $U_{ij} > 0$, it holds that if t is sufficiently small, then $L_{ij} \leq t(\text{diag}(S) - S)_{ij} \leq U_{ij}$ for all $i \neq j$. \square

The conditions of Lemma 8.6 are satisfied for the graphical lasso and this result can be turned into an explicit procedure for computing a starting point. Note that using `glasso` or `EBICglasso` in R outputs either a warning or an error when S is not positive definite. As a consequence of Lemma 8.6, we have the following theorem.

THEOREM 8.7. *The optimum in the graphical lasso problem always exists and is unique. More generally, this also holds for the GOLAZO if $L_{ij} < 0 < U_{ij}$ for all $i \neq j$.*

The conditions of Lemma 8.6 are not satisfied for the positive graphical lasso and, in general, for any case where S lies on the boundary of the rectangle $\{\Sigma : L \leq \Sigma - S \leq U\}$. If L has potentially zero entries but $U_{ij} > 0$ for all $i \neq j$, then we still can prove that a dually feasible point exists under very mild additional assumptions on S . This relies on an explicit construction of a dually feasible point based on the definition of a single-linkage matrix of S given in [33]. We provide the details of this construction in Appendix B, which also contains the proof of the following result.

THEOREM 8.8. *The optimum in the positive graphical lasso problem exists and is unique as long as $S_{ij} < \sqrt{S_{ii}S_{jj}}$ for all $i \neq j$, which happens with probability one if the sample size is at least two. The same holds in general whenever $U_{ij} > 0$ for all $i \neq j$.*

8.3. *Choice of the penalty parameter.* In this section, we propose a simple method to choose the penalty parameters in L and U . Our method is based on the Extended Bayesian Information Criterion (EBIC) proposed in [9] and adapted to problems of graphical lasso type in [20]. Given a sample of n independent and identically distributed observations, let \widehat{K} denote an estimate and \widehat{E} the set of edges of the underlying graph of \widehat{K} . The EBIC criterion takes the form

$$\text{BIC}_\gamma(\widehat{E}) = -n\ell_n(\widehat{K}) + |\widehat{E}|(\log(n) + 4\gamma \log(d)).$$

The criterion is indexed by a parameter $\gamma \in [0, 1]$. If $\gamma = 0$, then the classical BIC is recovered, which is known to be asymptotically consistent for model selection in case d is fixed and n goes to infinity. Positive values of γ lead to better graph estimates in the case when d and n are comparable. This observation can be formalized in certain scenarios but otherwise relies on numerical experiments; cf. [20].

The model selection procedure based on EBIC relies on choosing a sequence of potential penalty parameters ρ_1, \dots, ρ_l . Then for fixed L, U we then compute l optima for the GOLAZO problem with parameters $\rho_i L, \rho_i U$. For each of these we compute EBIC and choose ρ_i that minimizes this criterion. For positive glasso the canonical choice is $L_{ij} = 0$ and $U_{ij} = 1$ for all $i \neq j$. Finally, note that this procedure is trivially parallelizable, which we exploit in our code.

8.4. *Sign-constrained likelihood optimization.* The positive graphical lasso problem links to the problem of maximum likelihood estimation under M-matrix constraints [33, 46], that is to the problem

$$\text{minimize } -\ell(K) \quad \text{subject to } K_{ij} \leq 0 \text{ for all } i \neq j.$$

More generally, if $L_{ij} \in \{-\infty, 0\}$ and $U_{ij} \in \{0, +\infty\}$ then (15) is equivalent to optimizing the Gaussian likelihood under sign-constraints. If E_L is the set of all $i \neq j$ such that $L_{ij} = -\infty$ and E_U is the set of all $i \neq j$ such that $U_{ij} = +\infty$, then (15) amounts to maximizing the Gaussian likelihood over the set of all inverse covariance matrices K such that $K_{ij} \geq 0$ for all $ij \in E_L$ and $K_{ij} \leq 0$ for all $ij \in E_U$. If $ij \in E_L \cap E_U$ then $K_{ij} = 0$.

One of the important reasons why we choose not to penalize the diagonal of K is the following result.

THEOREM 8.9. *Let L, U be such that $L_{ij} \in \{-\infty, 0\}$ and $U_{ij} \in \{0, +\infty\}$ for all $i \neq j$. The GOLAZO estimator is equal to the maximum likelihood estimator under the constraints $K_{ij} \geq 0$ for all $ij \in E_L$ and $K_{ij} \leq 0$ for all $ij \in E_U$.*

Recall from Remark 8.3 that infinity can be replaced by a sufficiently large number.

8.5. *A link to graphical models.* Let \widehat{G} denote the graph determined by the support of \widehat{K} . Proposition 8.2 implies the following result.

COROLLARY 8.10. *The optimum \widehat{K} of (15) is equal to the maximum likelihood estimator in the Gaussian graphical model $M(\widehat{G})$ with the sufficient statistics $(\widehat{S}_{ij})_{ij \in \widehat{G}}$, where $\widehat{S}_{ij} = S_{ij} + L_{ij}$ if $\widehat{K}_{ij} < 0$ and $\widehat{S}_{ij} = S_{ij} + U_{ij}$ if $\widehat{K}_{ij} > 0$.*

PROOF. We have that $\widehat{\Sigma} \in M(\widehat{G})$ and it coincides with \widehat{S} on the edges of \widehat{G} . It is then the unique optimum of the MLE problem. \square

This result implies that as soon as we have access to the matrix of signs of \widehat{K} with $\text{sgn}(\widehat{K}_{ij}) \in \{-1, 0, 1\}$ we could alternatively find \widehat{K} by augmenting the sufficient statistics and fitting the corresponding Gaussian graphical model. If \widehat{G} is decomposable, the optimum is then given in a closed form.

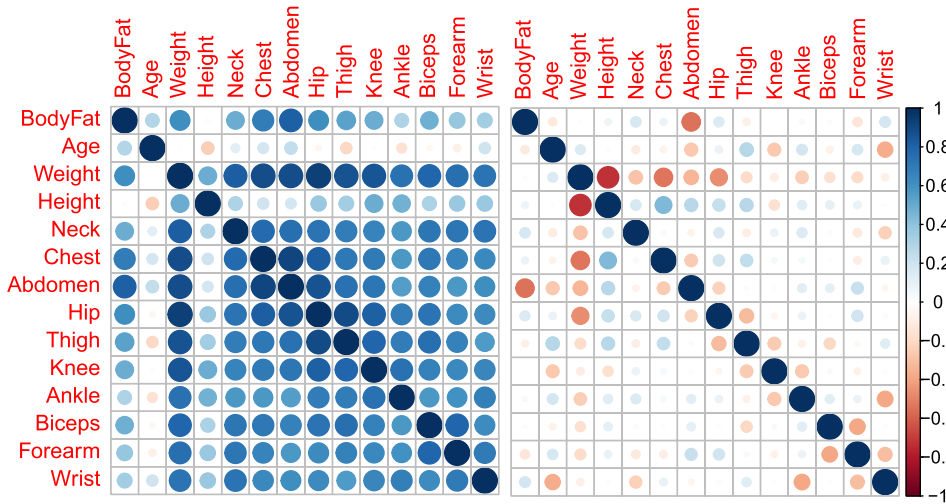


FIG. 1. The sample correlation matrix for the body fat data (left) and its normalized inverse (right).

9. Applications and simulations. In this section, we illustrate our methods with two applications and a small simulation study. The computations are based on the GOLAZO algorithm described in Section 8 as implemented in the R package `golazo` available on GitHub. The corresponding R Markdown files can be downloaded from <https://pzwiernik.github.io/supps/2020-LZ-golazo.zip>.

9.1. Body fat data. As a simple illustration for our method, we analyse the body fat data first studied in [25] and available in the R package `gRim`.

In the data, percentage of body fat, age, weight, height, and ten body circumference measurements are recorded for 252 men. We remove 11 individuals from the study for various problems reported for their respective observations. As shown in Figure 1 the variables in this dataset are strongly positively correlated. The only exception is Age. Its negative correlation with height reflects the fact that the growth in wealth, especially after the second world war, has made the new generations taller. Also, the negative correlation with thigh and ankle could reflect that muscle mass tends to be reduced with age. In other words, the relation between Age and the other variables is complex and certainly not approximately linear. For these reasons and for simplicity, we remove Age from our analysis.

As the inverse of the sample covariance matrix has a significant number of positive entries, MTP_2 seems to be too strong hypothesis for this dataset. After normalizing the data we run the positive graphical lasso procedure with $\rho = 0.11$. This choice was based on the EBIC criterion with parameter $\gamma = 0.5$ as described in Section 8.3. Figure 2 shows the graph (together with signs of the partial correlations) of the optimum \hat{K} (left) compared with the graphical lasso estimate. The penalty in the graphical lasso estimate is chosen close to zero both by cross-validation and by the EBIC criterion. Although the positive glasso estimate gives lower likelihood than the glasso estimate, it is much sparser and beats glasso in the EBIC criterion with $\gamma = 0.5$: -364.9 for the positive glasso and -237.65 for the graphical lasso.

In the second step of our procedure we take the resulting estimate \hat{K} of the positive glasso procedure and compute the dual MLE $\check{\Sigma}$ under edge positivity to further regularize the positive glasso estimate. However, although \hat{K} is not an M-matrix, it corresponds to a locally associated distribution and so the second step of the algorithm becomes redundant.

9.2. Positive co-expression gene network. As we argued in the [Introduction](#), our main motivation was to develop statistically sound methods for building gene expression networks

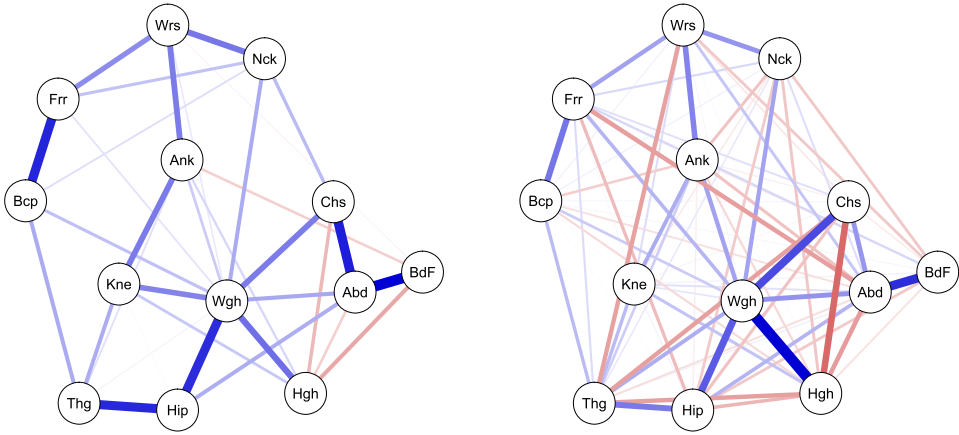


FIG. 2. Partial correlations in the estimated positive glasso graph (left) and glasso graph (right).

that focus on positive co-expression. For illustrative purposes we focus on a relatively small subsample of genes. We analyse a publicly available microarray expression data profiling umbilical cord tissue; cf. [11, 12]. From <https://functionalgenomics.upf.edu/supplements/FIRinELGANs> we downloaded the normalized and filtered gene expression data, as well as its corresponding phenotype data, including a batch indicator variable that specifies the groups in which samples were processed, birth weight, gestational age, sex and fetal inflammatory response (FIR) status.

As described in the Introduction, we focus our analysis on 136 genes which were coordinately upregulated in FIR-affected infants to trigger an innate immune response, and therefore, we can assume their positive co-expression.

We first run the GOLAZO procedure penalizing negative partial correlations with an EBIC optimal penalty parameter. This computation takes less than a minute on a standard laptop. The penalty $\rho = 0.5$ was chosen as optimal with respect to the EBIC criterion with $\gamma = 0.5$. The resulting estimate \hat{K}^ρ is very sparse with the edge density 0.067 (an entry of \hat{K}^ρ is always treated nonzero if its absolute value exceeded 10^{-6}). Still the diameter of this graph, displayed in Figure 3, is very small, just 5.

Like in the Body Fat example in the previous section, also here the EBIC criterion chooses much smaller penalty parameter for the standard graphical lasso, $\rho = 0.1$. The EBIC criterion for the optimal positive graphical lasso model penalty is -8784 on the other hand, the graphical lasso gives a much denser graph with edge density 0.21 and the EBIC is also much higher, 13,478. The optimum \hat{K}^ρ is not an M-matrix and there are several significant negative partial correlations, displayed in red in Figure 3. However, the distribution is already locally associated so the second step of our procedure is again redundant. This also confirms that local association is a reasonable assumption for this dataset.

APPENDIX A: PROOF OF THEOREM 6.1

A.1. Chernoff regularity and convexity. The asymptotic analysis of statistical procedures under constraints typically involves technical assumptions on the local geometry of the constrained space around the true parameter ψ_0 . Conditions of this form are called Chernoff regularity conditions; cf. [21] and references therein. In our case, convexity ensures these conditions to hold, but we provide the relevant definitions for completeness.

DEFINITION A.1. The *tangent cone* $T_C(\psi_0)$ of the set $C \subseteq \mathbb{R}^k$ at the point ψ_0 is the set of vectors in \mathbb{R}^k that are limits of sequences $\alpha_n(\psi_n - \psi_0)$, where α_n are positive reals and $\psi_n \in C$ converge to μ_0 .

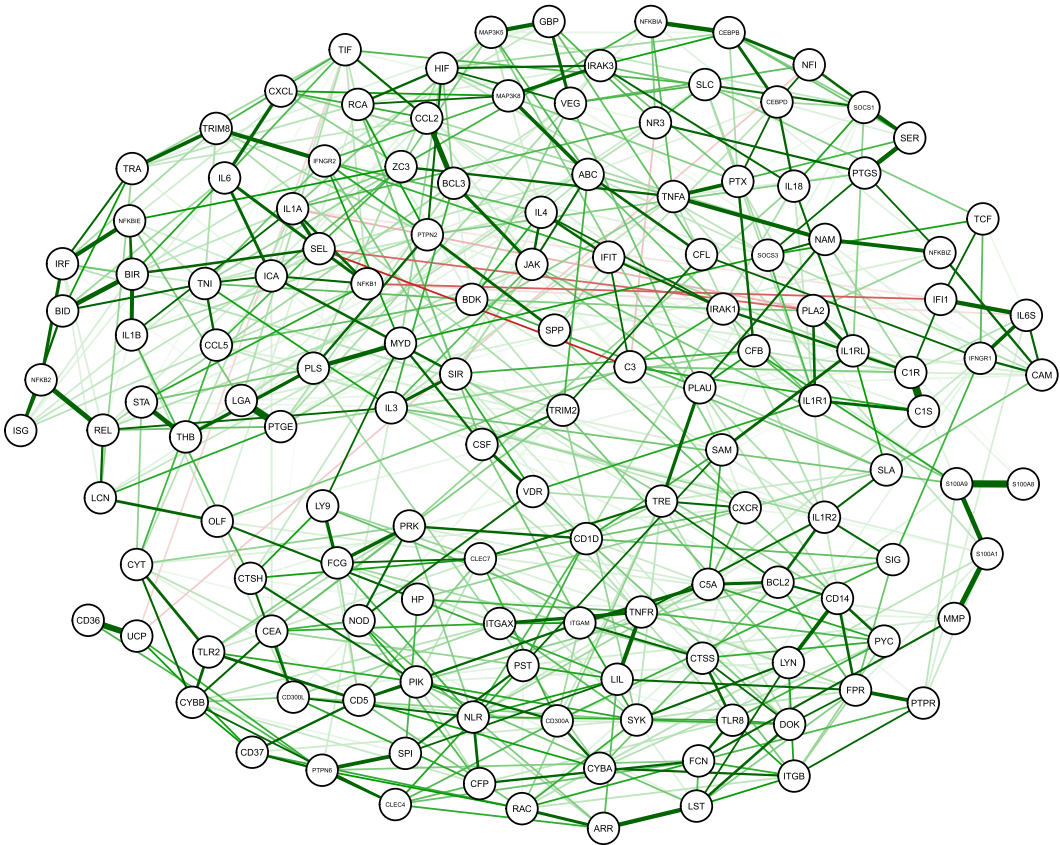


FIG. 3. The concentration graph for the co-expression gene network. Positive partial correlations are indicated with green color and negative partial correlations with red color. The thickness of edges is proportional to their absolute size.

DEFINITION A.2. The set $C \subseteq \mathbb{R}^k$ is *Chernoff regular* at ψ_0 if for every vector τ in the tangent cone $T_C(\psi_0)$ there exists $\varepsilon > 0$ and a map $\alpha : [0, \varepsilon) \rightarrow C$ with $\alpha(0) = \psi_0$ such that $\tau = \lim_{t \rightarrow 0^+} [\alpha(t) - \alpha(0)]/t$. In this case, we say that $T_C(\psi_0)$ is *derivable*; cf. Definition 6.1 in [44].

The standard asymptotic results typically assume Chernoff regularity. We will use the following result.

THEOREM A.3 (Theorem 6.9, [44]). A convex set $C \subseteq \mathbb{R}^k$ is Chernoff regular at any $\psi_0 \in C$.

It is clear from the definition that Chernoff regularity is preserved under a smooth and regular change of variables $G : \mathbb{R}^k \rightarrow \mathbb{R}^k$. Indeed, the tangent cone $T_{G(C)}(G(\psi_0))$ is equal to $\nabla G(\psi_0) \cdot T_C(\psi_0)$; cf. Section 6.C in [44]. If $\tau \in T_{G(C)}(G(\psi_0))$ and $T_C(G(\psi_0))$ is derivable then

$$(\nabla G(\psi_0))^{-1} \cdot \tau = \lim_{t \rightarrow 0^+} [\alpha(t) - \alpha(0)]/t$$

for some $\alpha : [0, \varepsilon) \rightarrow C$. Then $\alpha_G = (\nabla G(\psi_0)) \cdot \alpha : [0, \varepsilon) \rightarrow G(C)$ can be used to show that τ is derivable.

A.2. Asymptotics of the maximum likelihood and mixed dual estimator. Recall our notation $\mathbf{t}_n = \sum_{i=1}^n \mathbf{t}(X^{(i)})/n$. It is a standard result that $\sqrt{n}(\mathbf{t}_n - \boldsymbol{\mu}_0)$ is asymptotically normally distributed; see Proposition 4.3 in [49]. In this section we show that the maximum likelihood estimator under a convex restriction has a similar rate of convergence, with the limiting distribution not necessarily being normal. Using equivariance of the MLE and the delta method, we can show this also holds for the MLE and MDE in a mixed convex exponential family.

Let $\boldsymbol{\psi} = \boldsymbol{\psi}(\boldsymbol{\theta})$ be an alternative smooth, regular, and bijective parametrization of the exponential family \mathcal{E} , so that, $\boldsymbol{\theta} = \boldsymbol{\psi}^{-1}(\boldsymbol{\psi})$. The log-likelihood function expressed in this new parametrization is denoted by $\ell(\boldsymbol{\psi}; \mathbf{t}_n) = \langle \boldsymbol{\psi}^{-1}(\boldsymbol{\psi}), \mathbf{t}_n \rangle - A(\boldsymbol{\psi}^{-1}(\boldsymbol{\psi}))$.

PROPOSITION A.4. *Consider an alternative smooth and regular parametrization $\boldsymbol{\psi} = \boldsymbol{\psi}(\boldsymbol{\theta})$ of the exponential family \mathcal{E} . Let C be a closed and convex subset of the parameter space $\boldsymbol{\psi}(\Theta)$. Let $\hat{\boldsymbol{\psi}}_n = \arg \max_{\boldsymbol{\psi} \in C} \ell(\boldsymbol{\psi}; \mathbf{t}_n)$ be the maximum likelihood estimator over C . If the data are generated from the distribution with parameter $\boldsymbol{\psi}_0 \in C$ then $\sqrt{n}(\hat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$ converges in distribution.*

PROOF. Since the maximum likelihood estimator is equivariant, we have that $\hat{\boldsymbol{\psi}}_n = \boldsymbol{\psi}(\hat{\boldsymbol{\theta}}_n)$, where $\hat{\boldsymbol{\theta}}_n = \arg \max_{\boldsymbol{\theta} \in \boldsymbol{\psi}^{-1}(C)} \ell(\boldsymbol{\theta}; \mathbf{t}_n)$. By the delta method (see Theorem 3.1 in [51]), it is enough to show that $\sqrt{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ converges in distribution. Since the MLE is an M-estimator, this follows from Theorem 4.4 in [21]. This theorem uses a number of assumptions that we verify one by one: Assumption A holds because the function $F(\boldsymbol{\theta}) = A(\boldsymbol{\theta}) - \langle \boldsymbol{\theta}, \boldsymbol{\mu}_0 \rangle$ admits a quadratic approximation around $\boldsymbol{\theta}_0$ with a positive definite Hessian $\nabla^2 A(\boldsymbol{\theta}_0)$ and $\nabla F(\boldsymbol{\theta}_0) = 0$. Assumption B is satisfied simply because the second derivative of the likelihood does not depend on the data at all and so condition (4.3) in [21] trivially holds. Assumption C requires that the standard central limit theorem for the gradient of the likelihood function holds, which again is automatic for exponential families. Assumption D holds simply because $\hat{\boldsymbol{\theta}}_n$ is the exact minimizer of $F_n(\boldsymbol{\theta}) = -\ell(\boldsymbol{\theta}; \mathbf{t}_n)$ over $\boldsymbol{\psi}^{-1}(C)$. Finally, Chernoff regularity of $\boldsymbol{\psi}^{-1}(C)$ at $\boldsymbol{\theta}_0$ follows by convexity of C and the fact that the property is invariant under smooth transformations; cf. Theorem A.3 and the discussion below it. \square

As in Section 6, we now reserve the notation $\boldsymbol{\psi}$ for the underlying mixed parametrization, $\boldsymbol{\psi} = (\boldsymbol{\mu}_u, \boldsymbol{\theta}_v)$. The maximum likelihood estimator in our problem is obtained by maximizing $\ell(\boldsymbol{\psi}; \mathbf{t}_n)$ over all parameters $\boldsymbol{\psi}$ in the mixed convex exponential family \mathcal{E}' , $\boldsymbol{\psi} \in M'_u \times \Theta'_v$. This MLE is denoted $\tilde{\boldsymbol{\psi}}_n$ ($\tilde{\boldsymbol{\theta}}_n, \tilde{\boldsymbol{\mu}}_n$ resp.) to distinguish from the estimator $\hat{\boldsymbol{\theta}}_n$ obtained in step (S1) of the procedure for finding the MDE. Recall that $\bar{\boldsymbol{\psi}}_n = (\mathbf{u}_n, \theta_v(\mathbf{t}_n))$ and denote $\hat{\boldsymbol{\psi}}_n = (\mathbf{u}_n, \hat{\boldsymbol{\theta}}_v)$, that is, $\hat{\boldsymbol{\psi}}_n$ is $\hat{\boldsymbol{\theta}}_n = (\hat{\boldsymbol{\theta}}_u, \hat{\boldsymbol{\theta}}_v)$ expressed in the mixed parametrization, where the fact that the first component of $\hat{\boldsymbol{\psi}}_n$ is equal to \mathbf{u}_n follows from Theorem 4.5.

COROLLARY A.5. *The sequences $\sqrt{n}(\bar{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$, $\sqrt{n}(\hat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$, and $\sqrt{n}(\tilde{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$ all converge in distribution.*

PROOF. The estimators $\bar{\boldsymbol{\psi}}_n$, $\hat{\boldsymbol{\psi}}_n$, and $\tilde{\boldsymbol{\psi}}_n$ are all maximum likelihood estimators in families that satisfy the conditions in Proposition A.4. \square

The proof of our main result relies on the fact that the log-likelihood and dual log-likelihood have locally a similar shape around their global maximum. Moreover, the Hessian at this point is block diagonal as shown, for example, in Proposition 3.20 of [49]. For convenience of the reader, we have included this result reformulated in Proposition A.6 to conform with our notation and extended to the case of steep families.

PROPOSITION A.6. *If a steep exponential family is given a mixed parameterization by $(\boldsymbol{\mu}_u, \boldsymbol{\theta}_v)$, the two parameter components are variation independent and information orthogonal, and the Fisher information for $(\boldsymbol{\mu}_u, \boldsymbol{\theta}_v)$ is*

$$I(\boldsymbol{\mu}_u, \boldsymbol{\theta}_v) = \begin{bmatrix} \text{var}(\mathbf{u})^{-1} & \mathbf{0} \\ \mathbf{0} & (\text{var}(\mathbf{t})^{vv})^{-1} \end{bmatrix},$$

where $\text{var}(\mathbf{t})^{vv}$ stands for the $\mathbf{v}\mathbf{v}$ -block of $\text{var}(\mathbf{t})^{-1}$. The same formula holds for the observed information $J(\mathbf{u}_n, \boldsymbol{\theta}_v(\mathbf{t}_n))$ in the unrestricted MLE $\bar{\boldsymbol{\psi}}_n$.

LEMMA A.7. *Let $\mathbf{t} = (\mathbf{u}, \mathbf{v}) \in M$, $\bar{\boldsymbol{\psi}} = (\mathbf{u}, \boldsymbol{\theta}_v(\mathbf{t}))$, and $\bar{\boldsymbol{\theta}} = \boldsymbol{\theta}(\mathbf{t})$. Then $\nabla_{\boldsymbol{\psi}} \ell(\bar{\boldsymbol{\psi}}; \mathbf{t}) = \nabla_{\boldsymbol{\psi}} \tilde{\ell}(\bar{\boldsymbol{\psi}}; \mathbf{t}) = \mathbf{0}$ and*

$$\nabla_{\boldsymbol{\psi}}^2 \ell(\bar{\boldsymbol{\psi}}; \mathbf{t}) = \nabla_{\boldsymbol{\psi}}^2 \tilde{\ell}(\bar{\boldsymbol{\psi}}; \mathbf{t}) = - \begin{bmatrix} \text{var}(\mathbf{u})^{-1} & \mathbf{0} \\ \mathbf{0} & (\text{var}(\mathbf{t})^{vv})^{-1} \end{bmatrix},$$

where the variance is computed with respect to the distribution corresponding to the mixed parameter $\bar{\boldsymbol{\psi}}$.

PROOF. Slightly abusing notation, we write $\boldsymbol{\mu}(\boldsymbol{\psi})$ ($\boldsymbol{\theta}(\boldsymbol{\psi})$) for the map that maps the mixed parameter $\boldsymbol{\psi}$ to the corresponding mean (canonical) parameter. We have

$$\begin{aligned} \nabla_{\boldsymbol{\psi}} \ell(\boldsymbol{\psi}, \mathbf{t}) &= \left(\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\psi}} \right)^T \cdot (\mathbf{t} - \boldsymbol{\mu}(\boldsymbol{\psi})), \\ \nabla_{\boldsymbol{\psi}} \tilde{\ell}(\boldsymbol{\psi}, \mathbf{t}) &= \left(\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\psi}} \right)^T \cdot (\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}(\boldsymbol{\psi})) \end{aligned}$$

from which the statement about the gradients easily follows by plugging $\boldsymbol{\psi} = \bar{\boldsymbol{\psi}}$ as $\boldsymbol{\mu}(\bar{\boldsymbol{\psi}}) = \mathbf{t}$ and $\boldsymbol{\theta}(\bar{\boldsymbol{\psi}}) = \bar{\boldsymbol{\theta}}$. The particular block-diagonal form of the Hessian $\nabla_{\boldsymbol{\psi}}^2 \ell(\bar{\boldsymbol{\psi}}; \mathbf{t})$ follows by the last part of Proposition A.6. It remains to show that $\nabla_{\boldsymbol{\psi}}^2 \ell(\bar{\boldsymbol{\psi}}; \mathbf{t}) = \nabla_{\boldsymbol{\psi}}^2 \tilde{\ell}(\bar{\boldsymbol{\psi}}; \mathbf{t})$. Using the Leibniz rule, we observe that differentiating $\nabla_{\boldsymbol{\psi}} \ell(\boldsymbol{\psi}, \mathbf{t})$ with respect to $\boldsymbol{\psi}$, we get one term that vanishes at $\boldsymbol{\psi} = \bar{\boldsymbol{\psi}}$ and so

$$\nabla_{\boldsymbol{\psi}}^2 \ell(\bar{\boldsymbol{\psi}}; \mathbf{t}) = - \left(\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\psi}} \right)^T \frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\psi}}.$$

Using the same argument for $\nabla_{\boldsymbol{\psi}} \tilde{\ell}(\boldsymbol{\psi}, \mathbf{t})$ we get

$$\nabla_{\boldsymbol{\psi}}^2 \tilde{\ell}(\bar{\boldsymbol{\psi}}; \mathbf{t}) = - \left(\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\psi}} \right)^T \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\psi}}.$$

Since both matrices are symmetric, the equality follows. \square

Using Lemma A.7 with $\mathbf{t} = \mathbf{t}_n$, we see that the observed information satisfies

$$J_n := -\nabla_{\boldsymbol{\psi}}^2 \ell(\bar{\boldsymbol{\psi}}_n; \mathbf{t}_n) = -\nabla_{\boldsymbol{\psi}}^2 \tilde{\ell}(\bar{\boldsymbol{\psi}}_n; \mathbf{t}_n).$$

In particular, J_n is always positive definite. Taking $\mathbf{t} = \hat{\boldsymbol{\mu}}_n$ we get

$$\hat{J}_n := -\nabla_{\boldsymbol{\psi}}^2 \ell(\hat{\boldsymbol{\psi}}_n; \hat{\boldsymbol{\mu}}_n) = -\nabla_{\boldsymbol{\psi}}^2 \tilde{\ell}(\hat{\boldsymbol{\psi}}_n; \hat{\boldsymbol{\mu}}_n).$$

Both of J_n and \hat{J}_n converge in probability to the Fisher information matrix since the mapping $\mathbf{t} \mapsto -\nabla_{\boldsymbol{\psi}}^2 \ell(\bar{\boldsymbol{\psi}}; \mathbf{t})$ is continuous and both of \mathbf{t}_n and $\hat{\boldsymbol{\mu}}_n$ converge in probability to $\boldsymbol{\mu}_0$. We thus conclude

$$(17) \quad \hat{J}_n = J_n + o_P(1).$$

Denote by $\|\cdot\|_{J_n}$ ($\|\cdot\|_{\widehat{J}_n}$) the norm induced by the matrix J_n (\widehat{J}_n respectively), that is,

$$\|\mathbf{x}\|_{J_n} := \sqrt{\mathbf{x}^T J_n \mathbf{x}}, \quad \|\mathbf{x}\|_{\widehat{J}_n} := \sqrt{\mathbf{x}^T \widehat{J}_n \mathbf{x}}.$$

Note that for every $\mathbf{x} \in \mathbb{R}^d$

$$\lambda_{\min}(J_n)\|\mathbf{x}\|^2 \leq \|\mathbf{x}\|_{J_n}^2 \leq \lambda_{\max}(J_n)\|\mathbf{x}\|^2,$$

where $\lambda_{\min}(J_n)$ and $\lambda_{\max}(J_n)$ are the minimal and the maximal eigenvalue of J_n . Since the eigenvalues are continuous functions of the matrix (use [24], Theorem 2.6.4) and the fact that for a positive definite matrix singular values are equal to the eigenvalues), the sequences $\lambda_{\min}(J_n)$ and $\lambda_{\max}(J_n)$ converge in probability to the corresponding eigenvalues of the Fisher information matrix, which we denote by λ_{\min}^* , λ_{\max}^* . The same argument applies to $\lambda_{\min}(\widehat{J}_n)$ and $\lambda_{\max}(\widehat{J}_n)$. We conclude that for every $\mathbf{x} \in \mathbb{R}^d$

$$(18) \quad (\lambda_{\min}^* + o_P(1))\|\mathbf{x}\|^2 \leq \|\mathbf{x}\|_{J_n}^2 \leq (\lambda_{\max}^* + o_P(1))\|\mathbf{x}\|^2,$$

$$(19) \quad (\lambda_{\min}^* + o_P(1))\|\mathbf{x}\|^2 \leq \|\mathbf{x}\|_{\widehat{J}_n}^2 \leq (\lambda_{\max}^* + o_P(1))\|\mathbf{x}\|^2.$$

Now let

$$F_n(\boldsymbol{\psi}) = -\ell(\boldsymbol{\psi}; \mathbf{t}_n) \quad \text{and} \quad \widetilde{F}_n(\boldsymbol{\psi}) = -\widetilde{\ell}(\boldsymbol{\psi}; \widehat{\boldsymbol{\mu}}_n).$$

The MLE $\widetilde{\boldsymbol{\psi}}_n$ is the minimizer of F_n over $M'_u \times \Theta'_v$ and $\widehat{\boldsymbol{\psi}}_n$ is the minimizer of \widetilde{F}_n over $M'_u \times \Theta_v$ (equivalently over $M'_u \times \Theta'_v$, as argued in Theorem 4.5, because $\widehat{\boldsymbol{\psi}}_n \in M_u \times \Theta'_v$).

Recall that $U_n = O_P(1)$ denotes that U_n is bounded in probability, that is, for every $\epsilon > 0$ there exists $M \in \mathbb{R}$ such that $\mathbb{P}(\|U_n\| > M) < \epsilon$.

LEMMA A.8. *It holds that $\sqrt{n}(\overline{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0) = O_P(1)$, $\sqrt{n}(\widetilde{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0) = O_P(1)$, and $\sqrt{n}(\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0) = O_P(1)$.*

PROOF. By Corollary A.5, $\sqrt{n}(\overline{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$ and $\sqrt{n}(\widetilde{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$ converge in distribution and thus they are bounded in probability. To show the same for the other estimator, we use the fact that locally it is obtained by suitably projecting $\widehat{\boldsymbol{\psi}}_n$ on $M'_u \times \Theta_v$ and $\sqrt{n}(\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$ is bounded in probability because it converges in distribution by Corollary A.5. More formally, since $\widehat{\boldsymbol{\psi}}_n$ is the global minimizer of \widetilde{F}_n , $\widetilde{\boldsymbol{\psi}}_n$ is the minimizer over $M'_u \times \Theta_v$, and $\boldsymbol{\psi}_0 \in M'_u \times \Theta_v$, we get

$$0 \leq \widetilde{F}_n(\widetilde{\boldsymbol{\psi}}_n) - \widetilde{F}_n(\widehat{\boldsymbol{\psi}}_n) \leq \widetilde{F}_n(\boldsymbol{\psi}_0) - \widetilde{F}_n(\widehat{\boldsymbol{\psi}}_n).$$

Using the second-order expansion of $\widetilde{F}_n(\boldsymbol{\psi})$ around $\boldsymbol{\psi} = \widehat{\boldsymbol{\psi}}_n$ and Lemma A.7, we get

$$0 \leq \frac{1}{2}\|\widetilde{\boldsymbol{\psi}}_n - \widehat{\boldsymbol{\psi}}_n\|_{\widehat{J}_n}^2 + o_P(\|\widetilde{\boldsymbol{\psi}}_n - \widehat{\boldsymbol{\psi}}_n\|^2) \leq \frac{1}{2}\|\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0\|_{\widehat{J}_n}^2 + o_P(\|\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0\|^2).$$

Here we have also used the basic fact that $R(h) = o(\|h\|^2)$ implies $R(U_n) = o_P(\|U_n\|^2)$; cf. Lemma 2.12 in [51]. Multiply this inequality by n so that the right-hand side becomes

$$\frac{1}{2}\|\sqrt{n}(\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)\|_{\widehat{J}_n}^2 + o_P(\|\sqrt{n}(\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)\|^2).$$

Since $\sqrt{n}(\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)$ is bounded in probability, so is $\|\sqrt{n}(\widehat{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0)\|^2$. Using (19) and the fact that $o_P(O_P(1)) = o_P(1)$ we conclude that the right-hand side is bounded in probability and so the left-hand side is bounded too, that is, $\sqrt{n}(\widetilde{\boldsymbol{\psi}}_n - \widehat{\boldsymbol{\psi}}_n) = O_P(1)$. The triangle inequality now implies $\sqrt{n}(\widetilde{\boldsymbol{\psi}}_n - \boldsymbol{\psi}_0) = O_P(1)$, which concludes the proof. \square

We are now ready to prove the theorem.

PROOF OF THEOREM 6.1. Since, by Corollary A.5, $\sqrt{n}(\tilde{\psi}_n - \psi_0)$ converges in distribution, it remains to show that $\sqrt{n}(\tilde{\psi}_n - \bar{\psi}_n) = o_P(1)$; see [39], Theorem 3. The standard local first order conditions for optimality of $\tilde{\psi}_n$ are

$$(20) \quad \nabla F_n(\tilde{\psi}_n)^T \cdot (\psi - \tilde{\psi}_n) \geq 0 \quad \text{for all } \psi \in M'_u \times \Theta'_v,$$

expressing that the directional derivative in any feasible direction must be nonnegative. First-order Taylor expansion at $\bar{\psi}_n$ gives

$$\nabla F_n(\tilde{\psi}_n) = J_n \cdot (\tilde{\psi}_n - \bar{\psi}_n) + o_P(\|\tilde{\psi}_n - \bar{\psi}_n\|),$$

where we used Lemma A.7 and the fact that $\nabla F_n(\bar{\psi}_n) = 0$. After multiplying by \sqrt{n} the last term becomes $o_P(\sqrt{n}\|\tilde{\psi}_n - \bar{\psi}_n\|)$. By the triangle inequality and Lemma A.8, we have $\sqrt{n}\|\tilde{\psi}_n - \bar{\psi}_n\| \leq \sqrt{n}\|\tilde{\psi}_n - \psi_0\| + \sqrt{n}\|\psi_0 - \bar{\psi}_n\| = O_P(1)$. Using the fact that $o_P(O_P(1)) = o_P(1)$, we get

$$(21) \quad \sqrt{n}\nabla F_n(\tilde{\psi}_n) = J_n \cdot \sqrt{n}(\tilde{\psi}_n - \bar{\psi}_n) + o_P(1).$$

Multiply (20) by n and insert the expression in (21) for $\sqrt{n}\nabla F_n(\tilde{\psi}_n)$ to conclude that for all $\psi \in M'_u \times \Theta'_v$

$$(22) \quad \sqrt{n}(\tilde{\psi}_n - \bar{\psi}_n)^T \cdot J_n \cdot \sqrt{n}(\psi - \tilde{\psi}_n) + \|\sqrt{n}(\psi - \tilde{\psi}_n)\| o_P(1) \geq 0.$$

Here we used a basic fact that if r_n is a sequence of random vectors in a finite dimensional vector space such that each coordinate is $o_P(1)$ and q_n is another sequence of random vectors then $r_n^T q_n \leq \|q_n\| o_P(1)$.

Similarly, local optimality conditions for $\tilde{\psi}_n$ are

$$\nabla \tilde{F}_n(\tilde{\psi}_n)^T \cdot (\psi - \tilde{\psi}_n) \geq 0 \quad \text{for all } \psi \in M'_u \times \Theta_v.$$

By Lemma A.7, we have $\nabla \tilde{F}_n(\hat{\psi}_n) = 0$ and $\nabla^2 \tilde{F}_n(\hat{\psi}_n) = \hat{J}_n = J_n + o_P(1)$ (cf. (17)), which gives that

$$\nabla \tilde{F}_n(\tilde{\psi}_n) = J_n \cdot (\tilde{\psi}_n - \hat{\psi}_n) + o_P(\|\tilde{\psi}_n - \hat{\psi}_n\|),$$

and so, again using Lemma A.8, the first order optimality conditions for $\tilde{\psi}_n$, we get that for all $\psi \in M'_u \times \Theta_v$

$$(23) \quad \sqrt{n}(\tilde{\psi}_n - \hat{\psi}_n)^T \cdot J_n \cdot \sqrt{n}(\psi - \tilde{\psi}_n) + \|\sqrt{n}(\tilde{\psi}_n - \psi)\| o_P(1) \geq 0.$$

Note also that the optimality conditions for $\hat{\psi}_n$ are

$$\nabla F_n(\hat{\psi}_n)^T \cdot (\psi - \hat{\psi}_n) \geq 0 \quad \text{for all } \psi \in M_u \times \Theta'_v.$$

Like in the previous two cases, we argue that for all $\psi \in M_u \times \Theta'_v$

$$(24) \quad \sqrt{n}(\hat{\psi}_n - \bar{\psi}_n)^T \cdot J_n \cdot \sqrt{n}(\psi - \hat{\psi}_n) + \|\sqrt{n}(\psi - \hat{\psi}_n)\| o_P(1) \geq 0.$$

The rest of the proof relies on using (22), (23), and (24) evaluated at various points ψ in order to show that $\sqrt{n}(\tilde{\psi}_n - \bar{\psi}_n) = o_P(1)$. We now drop dependence on n from our notation to keep it simple. We will exploit the fact that J is positive definite and has a block-diagonal form with blocks, which we denote by J_{uu} and J_{vv} but the exact form of these blocks given in Lemma A.7 is irrelevant here.

Insert $(\tilde{\mu}_u, \tilde{\theta}_v)$ for ψ in (22) and $(\tilde{\mu}_u, \tilde{\theta}_v)$ for ψ in (23) noting that both are valid points in $M'_u \times \Theta'_v$; then add both expressions and use that the u -block of $\tilde{\psi}_n$ and $\hat{\psi}_n$ are both equal to u_n (cf. Theorem 4.5). From this calculation, we get

$$-\|\sqrt{n}(\tilde{\mu}_u - \hat{\mu}_u)\|_{J_{uu}}^2 + \|\sqrt{n}(\tilde{\mu}_u - \hat{\mu}_u)\| \cdot o_P(1) \geq 0.$$

This can be rewritten as

$$0 \leq \|\sqrt{n}(\tilde{\mu}_u - \hat{\mu}_u)\|_{J_{uu}}^2 \leq \|\sqrt{n}(\tilde{\mu}_u - \hat{\mu}_u)\| \cdot o_P(1).$$

Using a version of (18) for J_{uu} , we conclude that $\sqrt{n}(\tilde{\mu}_u - \hat{\mu}_u) = o_P(1)$.

Equation (22) evaluated at $\psi = (\tilde{\mu}_u, \tilde{\theta}_v) \in M'_u \times \Theta'_v$ (note that $\tilde{\theta}_v = \hat{\theta}_v$ by Theorem 4.5) yields

$$\sqrt{n}(\tilde{\theta}_v - \bar{\theta}_v)^T \cdot J_{vv} \cdot \sqrt{n}(\hat{\theta}_v - \tilde{\theta}_v) + \|\sqrt{n}(\hat{\theta}_v - \tilde{\theta}_v)\| o_P(1) \geq 0.$$

Similarly, insert $\psi = (u, \tilde{\theta}_v) = (\hat{\mu}_u, \tilde{\theta}_v) \in M_u \times \Theta'_v$ into (24) to get

$$\sqrt{n}(\hat{\theta}_v - \bar{\theta}_v)^T \cdot J_{vv} \cdot \sqrt{n}(\tilde{\theta}_v - \hat{\theta}_v) + \|\sqrt{n}(\tilde{\theta}_v - \hat{\theta}_v)\| o_P(1) \geq 0.$$

Adding these two inequalities yields

$$0 \leq \|\sqrt{n}(\tilde{\theta}_v - \hat{\theta}_v)\|_{J_{vv}}^2 \leq \|\sqrt{n}(\tilde{\theta}_v - \hat{\theta}_v)\| \cdot o_P(1).$$

Using a version of (18) for J_{vv} , we conclude that $\sqrt{n}(\tilde{\theta}_v - \hat{\theta}_v) = o_P(1)$. This finally gives that $\sqrt{n}(\tilde{\psi}_n - \hat{\psi}_n) = o_P(1)$, as desired. \square

APPENDIX B: THE SINGLE-LINKAGE MATRIX

We first define the single-linkage matrix of a covariance matrix S . Let R be a symmetric $p \times p$ positive semidefinite matrix such that $R_{ii} = 1$ for all $i = 1, \dots, p$. In our case R will be the corresponding correlation matrix of S . Consider the graph G^+ over $V = \{1, \dots, d\}$ with an edge between i and j whenever $R_{ij} > 0$. Assign to each edge the corresponding positive weight R_{ij} and note that G^+ in general does not have to be connected. Define a $d \times d$ matrix Z by setting $Z_{ii} = 1$ for all i and

$$(25) \quad Z_{ij} := \max_P \min_{uv \in P} R_{uv},$$

where the maximum is taken over all paths P in G^+ between i and j and is set to zero if no such path exists. We call Z the *single-linkage matrix* of R .

PROPOSITION B.1 ([33]). *Let R be a symmetric $d \times d$ positive semidefinite matrix satisfying $R_{ii} = 1$ for all $i = 1, \dots, p$. Then the single-linkage matrix Z of R has ones on the diagonal and satisfies $Z \geq R$. If, in addition, $R_{ij} < 1$ for all $i \neq j$, then Z is an inverse M -matrix.*

Now if S is a symmetric positive semidefinite matrix with strictly positive entries on the diagonal and such that $S_{ij} < \sqrt{S_{ii}S_{jj}}$. Then, by Proposition B.1, there exists an inverse M -matrix Z such that $Z \geq S$ and $Z_{ii} = S_{ii}$ for all $i = 1, \dots, d$, obtained by appropriate scaling of the correlation matrix R of S . This matrix is called the single-linkage matrix of S . If S is a sample covariance matrix based on at least two observations, the single-linkage matrix is always positive definite with probability one. We are now ready to show that the positive graphical lasso estimate exists for $n \geq 2$ observations.

PROOF OF THEOREM 8.8. We must construct a feasible point Σ^0 of the dual problem (16) in the case when L is the zero matrix. If S is positive definite, we can take $\Sigma^0 = S$ so assume that S is rank deficient. The single-linkage matrix Z of S by construction satisfies $Z \geq S$. If the entries of U are sufficiently large then $S + L \leq Z \leq S + U$ and so Z is dually feasible. If Z is not upper bounded by $S + U$ we proceed as follows. Let $\rho = \min_{i \neq j} U_{ij}$. Since $\|Z - S\|_\infty > 0$, we can define

$$(26) \quad \Sigma^0 := (1 - t^*)S + t^*Z, \quad \text{where } t^* = \min \left\{ 1, \frac{\rho}{\|Z - S\|_\infty} \right\}$$

which lets $\Sigma^0 = Z$ if $\rho = \infty$. Then $\Sigma^0 \geq S$ and it is equal to S on the diagonal. Moreover,

$$\|\Sigma^0 - S\| = t^* \|Z - S\|_\infty \leq \rho$$

and hence Σ^0 is dually feasible. Since Σ_0 is dually feasible, the optimum exists. This concludes the proof. \square

Finally, we comment briefly on computational issues. Computing Z can be done efficiently using the link of this construction to single-linkage clustering; cf. [33], Proposition 3.7. More precisely, we first take the corresponding correlation matrix R and form a dissimilarity matrix D , where

$$D_{ij} = \begin{cases} -\log R_{ij} & \text{if } R_{ij} > 0, \\ \infty & \text{otherwise.} \end{cases}$$

By construction $D_{ii} = 0$ for all $i \in V$. We then run the single linkage clustering algorithm on D . The time complexity of this step is $O(d^2)$. The R function `hclust` by default does not return the underlying ultrametric matrix \widehat{D} of distances but this information can be recovered from the standard output with a bit of work. Now the single-linkage matrix of R is simply

$$Z_{ij} = \exp(-\widehat{D}_{ij}).$$

The single linkage matrix of S is obtained by rescaling the matrix Z with the diagonal entries of S .

Acknowledgments. We are grateful to Robert Castelo for providing us with an interesting dataset and for patiently explaining the underlying biology to us. We would also like the referees for helpful comments and for pointing out a mistake in an earlier version of the paper.

Funding. The second author was supported in part by Ayudas Fundación BBVA a Proyectos de Investigación Científica en Matemáticas 2022.

REFERENCES

- [1] AGRAWAL, R., ROY, U. and UHLER, C. (2019). Covariance matrix estimation under total positivity for portfolio selection. arXiv preprint. Available at [arXiv:1909.04222](https://arxiv.org/abs/1909.04222).
- [2] AGRESTI, A. (1983). Testing marginal homogeneity for ordinal categorical variables. *Biometrics* **39** 505–510.
- [3] AGRESTI, A. (2002). *Categorical Data Analysis*, 2nd ed. *Wiley Series in Probability and Statistics*. Wiley-Interscience, New York. MR1914507 <https://doi.org/10.1002/0471249688>
- [4] BANERJEE, O., EL GHAOU, L. and D'ASPROMONT, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *J. Mach. Learn. Res.* **9** 485–516. MR2417243
- [5] BARNDORFF-NIELSEN, O. E. (1978). *Information and Exponential Families in Statistical Theory*. *Wiley Series in Probability and Mathematical Statistics*. Wiley, Chichester. MR0489333

- [6] BOEGE, T., KAHLE, T., KRETSCHMER, A. and RÖTTGER, F. (2021). The geometry of Gaussian double Markovian distributions. arXiv preprint. Available at [arXiv:2107.00134](https://arxiv.org/abs/2107.00134).
- [7] BREUER, K., FOROUSHANI, A. K., LAIRD, M. R., CHEN, C., SRIBNAIA, A., LO, R., WINSOR, G. L., HANCOCK, R. E. W., BRINKMAN, F. S. L. and LYNN, D. J. (2013). InnateDB: Systems biology of innate immunity and beyond—recent updates and continuing curation. *Nucleic Acids Res.* **41** D1228–D1233.
- [8] BROWN, L. D. (1986). *Fundamentals of Statistical Exponential Families. IMS-Monographs IX*. IMS, Hayward, CA.
- [9] CHEN, J. and CHEN, Z. (2008). Extended Bayesian information criteria for model selection with large model spaces. *Biometrika* **95** 759–771. [MR2443189 https://doi.org/10.1093/biomet/asn034](https://doi.org/10.1093/biomet/asn034)
- [10] CHRISTENSEN, E. S. (1989). Statistical properties of I -projections within exponential families. *Scand. J. Stat.* **16** 307–318. [MR1039285](https://doi.org/10.1093/biomet/16.3.307)
- [11] COHEN, J., VAN MARTER, L. J., SUN, Y., ALLRED, E., LEVITON, A. and KOHANE, I. S. (2007). Perturbation of gene expression of the chromatin remodeling pathway in premature newborns at risk for bronchopulmonary dysplasia. *Genome Biol.* **8** R210. <https://doi.org/10.1186/gb-2007-8-10-r210>
- [12] COSTA, D. and CASTELO, R. (2016). Umbilical cord gene expression reveals the molecular architecture of the fetal inflammatory response in extremely preterm newborns. *Pediatr. Res.* **79** 473–481.
- [13] DJORDJILović, V. and CHIOGNA, M. (2018). Searching for a source of difference in Gaussian graphical models. arXiv preprint. Available at [arXiv:1811.02503](https://arxiv.org/abs/1811.02503).
- [14] DRTON, M. and RICHARDSON, T. S. (2002). A new algorithm for maximum likelihood estimation in Gaussian graphical models for marginal independence. In *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence* 184–191. Morgan Kaufmann, San Mateo.
- [15] EFRON, B. (1978). The geometry of exponential families. *Ann. Statist.* **6** 362–376. [MR0471152](https://doi.org/10.1214/aoms/1177698701)
- [16] ESARY, J. D., PROSCHAN, F. and WALKUP, D. W. (1967). Association of random variables, with applications. *Ann. Math. Stat.* **38** 1466–1474. [MR0217826 https://doi.org/10.1214/aoms/1177698701](https://doi.org/10.1214/aoms/1177698701)
- [17] FALLAT, S., LAURITZEN, S., SADEGHI, K., UHLER, C., WERMUTH, N. and ZWIERNIK, P. (2017). Total positivity in Markov structures. *Ann. Statist.* **45** 1152–1184. [MR3662451 https://doi.org/10.1214/16-AOS1478](https://doi.org/10.1214/16-AOS1478)
- [18] FAN, J., FENG, Y. and WU, Y. (2009). Network exploration via the adaptive lasso and SCAD penalties. *Ann. Appl. Stat.* **3** 521–541. [MR2750671 https://doi.org/10.1214/08-AOAS215](https://doi.org/10.1214/08-AOAS215)
- [19] FITZMAURICE, G. M. and LAIRD, N. M. (1993). A likelihood-based method for analysing longitudinal binary responses. *Biometrika* **80** 141–151.
- [20] FOYGEL, R. and DRTON, M. (2010). Extended Bayesian information criteria for Gaussian graphical models. In *Advances in Neural Information Processing Systems* 604–612.
- [21] GEYER, C. J. (1994). On the asymptotics of constrained M -estimation. *Ann. Statist.* **22** 1993–2010. [MR1329179 https://doi.org/10.1214/aos/1176325768](https://doi.org/10.1214/aos/1176325768)
- [22] GLONEK, G. F. V. (1996). A class of regression models for multivariate categorical responses. *Biometrika* **83** 15–28. [MR1399152 https://doi.org/10.1093/biomet/83.1.15](https://doi.org/10.1093/biomet/83.1.15)
- [23] GROSS, E. and SULLIVANT, S. (2018). The maximum likelihood threshold of a graph. *Bernoulli* **24** 386–407. [MR3706762 https://doi.org/10.3150/16-BEJ881](https://doi.org/10.3150/16-BEJ881)
- [24] HORN, R. A. and JOHNSON, C. R. (2013). *Matrix Analysis*, 2nd ed. Cambridge Univ. Press, Cambridge. [MR2978290](https://doi.org/10.1017/9781107057222)
- [25] JOHNSON, R. W. (1996). Fitting percentage of body fat to simple body measurements. *J. Stat. Educ.* **4**.
- [26] KARLIN, S. and RINOTT, Y. (1980). Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions. *J. Multivariate Anal.* **10** 467–498. [MR0599685 https://doi.org/10.1016/0047-259X\(80\)90065-2](https://doi.org/10.1016/0047-259X(80)90065-2)
- [27] KARLIN, S. and RINOTT, Y. (1983). M -matrices as covariance matrices of multivariate normal distributions. *Linear Algebra Appl.* **52/53** 419–438. [MR0709364 https://doi.org/10.1016/0024-3795\(83\)80027-5](https://doi.org/10.1016/0024-3795(83)80027-5)
- [28] KAUFMANN, G. (1996). On a dualization of graphical Gaussian models. *Scand. J. Stat.* **23** 105–116. [MR1380485](https://doi.org/10.1080/00141801.1996.10555845)
- [29] KHARCHENKO, P., CHURCH, G. M. and VITKUP, D. (2005). Expression dynamics of a cellular metabolic network. *Mol. Syst. Biol.* **1** 2005-0016.
- [30] KULLBACK, S. (1959). *Information Theory and Statistics*. Wiley, New York. [MR0103557](https://doi.org/10.1017/CBO9780511528364)
- [31] KULLBACK, S. (1971). Marginal homogeneity of multidimensional contingency tables. *Ann. Math. Stat.* **42** 594–606. [MR0297084 https://doi.org/10.1214/aoms/1177693409](https://doi.org/10.1214/aoms/1177693409)
- [32] KULLBACK, S. and LEIBLER, R. A. (1951). On information and sufficiency. *Ann. Math. Stat.* **22** 79–86. [MR0039968 https://doi.org/10.1214/aoms/1177729694](https://doi.org/10.1214/aoms/1177729694)
- [33] LAURITZEN, S., UHLER, C. and ZWIERNIK, P. (2019). Maximum likelihood estimation in Gaussian models under total positivity. *Ann. Statist.* **47** 1835–1863. [MR3953437 https://doi.org/10.1214/17-AOS1668](https://doi.org/10.1214/17-AOS1668)

- [34] LAURITZEN, S., UHLER, C. and ZWIERNIK, P. (2021). Total positivity in exponential families with application to binary variables. *Ann. Statist.* **49** 1436–1459. [MR4298870](#) <https://doi.org/10.1214/20-aos2007>
- [35] LAURITZEN, S. L. (1996). *Graphical Models. Oxford Statistical Science Series 17*. Clarendon, Oxford. [MR1419991](#)
- [36] LINDSKOG, F., MCNEIL, A. and SCHMOCK, U. (2003). Kendall's tau for elliptical distributions. In *Credit Risk* 149–156. Springer, Berlin.
- [37] LIU, H., HAN, F., YUAN, M., LAFFERTY, J. and WASSERMAN, L. (2012). High-dimensional semiparametric Gaussian copula graphical models. *Ann. Statist.* **40** 2293–2326. [MR3059084](#) <https://doi.org/10.1214/12-AOS1037>
- [38] MALIOUTOV, D. M., JOHNSON, J. K. and WILLSKY, A. S. (2006). Walk-sums and belief propagation in Gaussian graphical models. *J. Mach. Learn. Res.* **7** 2031–2064. [MR2274432](#)
- [39] MANN, H. B. and WALD, A. (1943). On stochastic limit and order relationships. *Ann. Math. Stat.* **14** 217–226. [MR0009261](#) <https://doi.org/10.1214/aoms/1177731415>
- [40] MASON, M. J., FAN, G., PLATH, K., ZHOU, Q. and HORVATH, S. (2009). Signed weighted gene co-expression network analysis of transcriptional regulation in murine embryonic stem cells. *BMC Genomics* **10** 327. <https://doi.org/10.1186/1471-2164-10-327>
- [41] PEARL, J. and WERMUTH, N. (1994). When can association graphs admit a causal interpretation? In *Selecting Models from Data* 205–214 Springer, New York.
- [42] PITT, L. D. (1982). Positively correlated normal variables are associated. *Ann. Probab.* **10** 496–499. [MR0665603](#)
- [43] ROCKAFELLAR, R. T. (1970). *Convex Analysis. Princeton Mathematical Series 28*. Princeton Univ. Press, Princeton. [MR0274683](#)
- [44] ROCKAFELLAR, R. T. and WETS, R. J.-B. (2009). *Variational Analysis 317*. Springer, Berlin.
- [45] ROSSELL, D. and ZWIERNIK, P. (2021). Dependence in elliptical partial correlation graphs. *Electron. J. Stat.* **15** 4236–4263. [MR4312203](#) <https://doi.org/10.1214/21-ejs1891>
- [46] SLAWSKI, M. and HEIN, M. (2015). Estimation of positive definite M -matrices and structure learning for attractive Gaussian Markov random fields. *Linear Algebra Appl.* **473** 145–179. [MR3338330](#) <https://doi.org/10.1016/j.laa.2014.04.020>
- [47] SOLOFF, J. A., GUNTUBOYINA, A. and JORDAN, M. I. (2020). Covariance estimation with nonnegative partial correlations. arXiv preprint. Available at [arXiv:2007.15252](https://arxiv.org/abs/2007.15252).
- [48] SPEED, T. P. and KIIVERI, H. T. (1986). Gaussian Markov distributions over finite graphs. *Ann. Statist.* **14** 138–150. [MR0829559](#) <https://doi.org/10.1214/aos/1176349846>
- [49] SUNDBERG, R. (2019). *Statistical Modelling by Exponential Families. Institute of Mathematical Statistics Textbooks 12*. Cambridge Univ. Press, Cambridge. [MR3969949](#) <https://doi.org/10.1017/9781108604574>
- [50] UHLER, C. (2012). Geometry of maximum likelihood estimation in Gaussian graphical models. *Ann. Statist.* **40** 238–261. [MR3014306](#) <https://doi.org/10.1214/11-AOS957>
- [51] VAN DER VAART, A. W. (1998). *Asymptotic Statistics. Cambridge Series in Statistical and Probabilistic Mathematics 3*. Cambridge Univ. Press, Cambridge. [MR1652247](#) <https://doi.org/10.1017/CBO9780511802256>
- [52] WEI, H., PERSSON, S., MEHTA, T., SRINIVASASAINAGENDRA, V., CHEN, L., PAGE, G. P., SOMERVILLE, C. and LORAINÉ, A. (2006). Transcriptional coordination of the metabolic network in Arabidopsis. *Plant Physiol.* **142** 762–774.