

LATENT VARIABLE GRAPHICAL MODEL SELECTION VIA CONVEX OPTIMIZATION*

BY VENKAT CHANDRASEKARAN [†], PABLO A. PARRILO , AND ALAN S.
WILLSKY

Massachusetts Institute of Technology

Suppose we observe samples of a *subset* of a collection of random variables. No additional information is provided about the number of latent variables, nor of the relationship between the latent and observed variables. Is it possible to discover the number of latent components, and to learn a statistical model over the entire collection of variables? We address this question in the setting in which the latent and observed variables are jointly Gaussian, with the conditional statistics of the observed variables conditioned on the latent variables being specified by a graphical model. As a first step we give natural conditions under which such latent-variable Gaussian graphical models are identifiable given marginal statistics of only the observed variables. Essentially these conditions require that the conditional graphical model among the observed variables is sparse, while the effect of the latent variables is “spread out” over most of the observed variables. Next we propose a tractable convex program based on regularized maximum-likelihood for model selection in this latent-variable setting; the regularizer uses both the ℓ_1 norm and the nuclear norm. Our modeling framework can be viewed as a combination of dimensionality reduction (to identify latent variables) and graphical modeling (to capture remaining statistical structure not attributable to the latent variables), and it consistently estimates both the number of latent components and the conditional graphical model structure among the observed variables. These results are applicable in the high-dimensional setting in which the number of latent/observed variables grows with the number of samples of the observed variables. The geometric properties of the algebraic varieties of sparse matrices and of low-rank matrices play an important role in our analysis.

1. Introduction and Setup. Statistical model selection in the high-dimensional regime arises in a number of applications. In many data analysis

*This work was supported in part by AFOSR grant FA9550-08-1-0180, in part under a MURI through AFOSR grant FA9550-06-1-0324, in part under a MURI through AFOSR grant FA9550-06-1-0303, and in part by NSF FRG 0757207.

[†]Corresponding author.

AMS 2000 subject classifications: 62F12, 62H12

Keywords and phrases: Gaussian graphical models, covariance selection, latent variables, regularization, sparsity, low-rank, algebraic statistics, high-dimensional asymptotics

problems in geophysics, radiology, genetics, climate studies, and image processing, the number of samples available is comparable to or even smaller than the number of variables. As empirical statistics in these settings may not be well-behaved (see [19, 24]), high-dimensional model selection is therefore both challenging and of great interest. A model selection problem that has received considerable attention recently is the estimation of covariance matrices in the high-dimensional setting. As the sample covariance matrix is poorly behaved in such a regime, some form of *regularization* of the sample covariance is adopted based on assumptions about the true underlying covariance matrix [1, 2, 14, 15, 22, 39].

Graphical models. A number of papers have studied covariance estimation in the context of *Gaussian graphical model selection*. A *Gaussian graphical model* [21, 32] (also commonly referred to as a Gauss-Markov random field) is a statistical model defined with respect to a graph, in which the nodes index a collection of jointly Gaussian random variables and the edges represent the conditional independence relations (Markov structure) among the variables. In such models the sparsity pattern of the inverse of the covariance matrix, or the *concentration* matrix, directly corresponds to the graphical model structure. Specifically, consider a Gaussian graphical model in which the covariance matrix is given by a positive-definite Σ^* and the concentration matrix is given by $K^* = (\Sigma^*)^{-1}$. Then an edge $\{i, j\}$ is present in the underlying graphical model if and only if $K_{i,j}^* \neq 0$. In particular the absence of an edge between two nodes implies that the corresponding variables are independent conditioned on all the other variables. The model selection method usually studied in such a Gaussian graphical model setting is ℓ_1 -regularized maximum-likelihood, with the ℓ_1 penalty applied to the entries of the concentration matrix to induce sparsity. The consistency properties of such an estimator have been studied [20, 28, 31], and under suitable conditions [20, 28] this estimator is also “sparsistent”, i.e., the estimated concentration matrix has the same sparsity pattern as the true model from which the samples are generated. An alternative approach to ℓ_1 -regularized maximum-likelihood is to estimate the sparsity pattern of the concentration matrix by performing regression separately on each variable [25]; while such a method consistently estimates the sparsity pattern, it does not directly provide estimates of the covariance or concentration matrix.

In many applications throughout science and engineering (e.g., psychology, computational biology, and economics), a challenge is that one may not have access to observations of all the relevant phenomena, i.e., some of the relevant variables may be latent or unobserved. In general latent variables pose a significant difficulty for model selection because one may not know

the number of relevant latent variables, nor the relationship between these variables and the observed variables. Typical algorithmic methods that try to get around this difficulty usually fix the number of latent variables as well as the structural relationship between latent and observed variables (e.g., the graphical model structure between latent and observed variables), and use the EM algorithm to fit parameters [10]. This approach suffers from the problem that one optimizes non-convex functions, and thus one may get stuck in sub-optimal local minima. An alternative suggestion [13] is one based on a greedy, local, combinatorial heuristic that assigns latent variables to groups of observed variables, via some form of clustering of the observed variables; however, this approach has no consistency guarantees.

Our setup. In this paper we study the problem of latent-variable graphical model selection in the setting where all the variables, both observed and latent, are jointly Gaussian. More concretely X is a Gaussian random vector in \mathbb{R}^{p+h} , O and H are disjoint subsets of indices in $\{1, \dots, p+h\}$ of cardinalities $|O| = p$ and $|H| = h$, and the corresponding subvectors of X are denoted by X_O and X_H respectively. Let the covariance matrix underlying X be denoted by $\Sigma_{(O\ H)}^*$. The marginal statistics corresponding to the observed variables X_O are given by the marginal covariance matrix Σ_O^* , which is simply a submatrix of the full covariance matrix $\Sigma_{(O\ H)}^*$. However suppose that we parameterize our model by the concentration matrix $K_{(O\ H)}^* = (\Sigma_{(O\ H)}^*)^{-1}$, which as discussed above reveals the connection to graphical models. Here the submatrices $K_O^*, K_{O,H}^*, K_H^*$ specify (in the full model) the dependencies among the observed variables, between the observed and latent variables, and among the latent variables, respectively. In such a parametrization, the *marginal concentration matrix* $(\Sigma_O^*)^{-1}$ corresponding to the observed variables X_O is given by the Schur complement [18] with respect to the block K_H^* :

$$(1.1) \quad \tilde{K}_O^* = (\Sigma_O^*)^{-1} = K_O^* - K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*.$$

Thus if we only observe the variables X_O , we only have access to Σ_O^* (or \tilde{K}_O^*). The two terms that compose \tilde{K}_O^* above have interesting properties. The matrix K_O^* specifies the concentration matrix of the *conditional statistics* of the observed variables given the latent variables. If these conditional statistics are given by a sparse graphical model then K_O^* is *sparse*. On the other hand the matrix $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$ serves as a *summary* of the effect of marginalization over the latent variables X_H . This matrix has small rank if the number of latent, unobserved variables X_H is small relative to the number of observed variables X_O . Therefore the marginal concentration matrix \tilde{K}_O^* is generally *not sparse* due to the additional low-rank term

$K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$. Hence standard graphical model selection techniques applied directly to the observed variables X_O are not useful.

A modeling paradigm that infers the effect of the latent variables X_H would be more suitable in order to provide a concise explanation of the underlying statistical structure. Hence we approximate the sample covariance by a model in which the concentration matrix *decomposes* into the sum of a sparse matrix and a low-rank matrix, which reveals the conditional graphical model structure in the observed variables as well as the number of and effect due to the unobserved latent variables. Such a method can be viewed as a blend of principal component analysis and graphical modeling. In standard graphical modeling one would directly approximate a concentration matrix by a sparse matrix to learn a sparse graphical model, while in principal component analysis the goal is to explain the statistical structure underlying a set of observations using a small number of latent variables (i.e., approximate a covariance matrix as a low-rank matrix). In our framework we learn a sparse graphical model among the observed variables *conditioned* on a few (additional) latent variables. These latent variables are *not* principal components, as the conditional statistics (conditioned on these latent variables) are given by a graphical model. Therefore we refer to these latent variables informally as *latent components*.

Contributions. Our first contribution in Section 3 is to address the fundamental question of *identifiability* of such latent-variable graphical models given the marginal statistics of only the observed variables. The critical point is that we need to tease apart the correlations induced due to marginalization over the latent variables from the conditional graphical model structure among the observed variables. As the identifiability problem is one of *uniquely* decomposing the sum of a sparse matrix and a low-rank matrix into the individual components, we study the algebraic varieties of sparse matrices and low-rank matrices. An important theme in this paper is the connection between the tangent spaces to these algebraic varieties and the question of identifiability. Specifically let $\Omega(K_O^*)$ denote the tangent space at K_O^* to the algebraic variety of sparse matrices, and let $T(K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*)$ denote the tangent space at $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$ to the algebraic variety of low-rank matrices. Then the *statistical* question of identifiability of K_O^* and $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$ given \tilde{K}_O^* is determined by the *geometric* notion of *transversality* of the tangent spaces $\Omega(K_O^*)$ and $T(K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*)$. The study of the transversality of these tangent spaces leads to natural conditions for identifiability. In particular we show that latent-variable models in which (1) the sparse matrix K_O^* has a small number of nonzeros per row/column, and (2) the low-rank matrix $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$ has row/column spaces that

are not closely aligned with the coordinate axes, are identifiable. These conditions have natural statistical interpretations. The first condition ensures that there are no densely-connected subgraphs in the conditional graphical model structure among the observed variables, i.e., that these conditional statistics are indeed specified by a sparse graphical model. Such statistical relationships may otherwise be mistakenly attributed to the effect of marginalization over some latent variable. The second condition ensures that the effect of marginalization over the latent variables is “spread out” over many observed variables; thus, the effect of marginalization over a latent variable is not confused with the conditional graphical model structure among the observed variables. In fact the first condition is often assumed in standard graphical model selection without latent variables (e.g., [28]).

As our next contribution we propose a *regularized maximum-likelihood decomposition* framework to approximate a given sample covariance matrix by a model in which the concentration matrix decomposes into a sparse matrix and a low-rank matrix. Based on the effectiveness of the ℓ_1 norm as a tractable convex relaxation for recovering sparse models [4, 11, 12] and the nuclear norm for low-rank matrices [5, 16, 29], we propose the following penalized likelihood method given a sample covariance matrix $\Sigma_{\mathcal{O}}^n$ formed from n samples of the observed variables:

$$(1.2) \quad \begin{aligned} (\hat{S}_n, \hat{L}_n) = \arg \min_{S, L} & \quad -\ell(S - L; \Sigma_{\mathcal{O}}^n) + \lambda_n (\gamma \|S\|_1 + \text{tr}(L)) \\ \text{s.t.} & \quad S - L \succ 0, \quad L \succeq 0. \end{aligned}$$

The constraints $\succ 0$ and $\succeq 0$ impose positive-definiteness and positive-semidefiniteness. The function ℓ represents the Gaussian log-likelihood $\ell(K; \Sigma) = \log \det(K) - \text{tr}(K\Sigma)$ for $K \succ 0$, where tr is the trace of a matrix and \det is the determinant. Here \hat{S}_n provides an estimate of $K_{\mathcal{O}}^*$, which represents the conditional concentration matrix of the observed variables; \hat{L}_n provides an estimate of $K_{\mathcal{O}, H}^* (K_H^*)^{-1} K_{H, \mathcal{O}}^*$, which represents the effect of marginalization over the latent variables. The regularizer is a combination of the ℓ_1 norm applied to S and the nuclear norm applied to L (the nuclear norm reduces to the trace over the cone of symmetric, positive-semidefinite matrices), with γ providing a tradeoff between the two terms. This variational formulation is a *convex optimization* problem, and it is a regularized max-det program that can be solved in polynomial time using general-purpose solvers [36].

Our main result in Section 4 is a proof of the consistency of the estimator (1.2) in the high-dimensional regime in which both the number of observed variables and the number of latent components are allowed to grow with the number of samples (of the observed variables). We show that for

a suitable choice of the regularization parameter λ_n , there exists a range of values of γ for which the estimates (\hat{S}_n, \hat{L}_n) have the same sparsity (and sign) pattern and rank as $(K_O^*, K_{O,H}^* (K_H^*)^{-1} K_{H,O}^*)$ with high probability (see Theorem 4.1). The key technical requirement is an identifiability condition for the two components of the marginal concentration matrix \tilde{K}_O^* with respect to the Fisher information (see Section 3.4). We make connections between our condition and the irrepresentability conditions required for support/graphical-model recovery using ℓ_1 regularization [28, 35, 40]. Our results provide numerous scaling regimes under which consistency holds in latent-variable graphical model selection. For example we show that *under suitable identifiability conditions consistent model selection is possible even when the number of samples and the number of latent variables are on the same order as the number of observed variables (see Section 4.2)*.

Related previous work. The problem of decomposing the sum of a sparse matrix and a low-rank matrix via convex optimization into the individual components was initially studied in [8] by a superset of the authors of the present paper, with conditions derived under which the convex program exactly recovers the underlying components. In subsequent work Candès et al. [6] also studied this sparse-plus-low-rank decomposition problem, and provided guarantees for exact recovery using the convex program proposed in [8]. The problem setup considered in the present paper is quite different and is more challenging because we are only given access to an inexact sample covariance matrix, and we wish to produce an *inverse* covariance matrix that can be decomposed as the sum of sparse and low-rank components (preserving the sparsity pattern and rank of the components in the true underlying model). In addition to proving the consistency of the estimator (1.2), we also provide a statistical interpretation of our identifiability conditions and describe natural classes of latent-variable Gaussian graphical models that satisfy these conditions. As such our paper is closer in spirit to the many recent papers on covariance selection, but with the important difference that some of the variables are not observed.

Outline. Section 2 gives some background and a formal problem statement. Section 3 discusses the identifiability question, Section 4 states the main results of this paper, and Section 5 gives some proofs. We provide experimental demonstration of the effectiveness of our estimator on synthetic and real data in Section 6, and conclude with a brief discussion in Section 7. Some of our technical results are deferred to supplementary material [7].

2. Problem Statement and Background. We give a formal statement of the latent-variable model selection problem. We also briefly describe

various properties of the algebraic varieties of sparse matrices and of low-rank matrices, and the properties of the Gaussian likelihood function.

The following matrix norms are employed throughout this paper. $\|M\|_2$ denotes the spectral norm, or the largest singular value of M . $\|M\|_\infty$ denotes the largest entry in magnitude of M . $\|M\|_F$ denotes the Frobenius norm, or the square-root of the sum of the squares of the entries of M . $\|M\|_*$ denotes the nuclear norm, or the sum of the singular values of M (this reduces to the trace for positive-semidefinite matrices). $\|M\|_1$ denotes the sum of the absolute values of the entries of M . A number of *matrix operator* norms are also used. For example, let $\mathcal{Z} : \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p \times p}$ be a linear operator acting on matrices. Then the induced operator norm is defined as $\|\mathcal{Z}\|_{q \rightarrow q} \triangleq \max_{N \in \mathbb{R}^{p \times p}, \|N\|_q \leq 1} \|\mathcal{Z}(N)\|_q$. Therefore, $\|\mathcal{Z}\|_{F \rightarrow F}$ denotes the spectral norm of the operator \mathcal{Z} . The only vector norm used is the Euclidean norm, which is denoted by $\|\cdot\|$. Given any norm $\|\cdot\|_q$ (either a vector norm, a matrix norm, or a matrix operator norm) the dual norm is given by $\|M\|_q^* \triangleq \sup\{\langle M, N \rangle \mid \|N\|_q \leq 1\}$.

2.1. Problem statement. In order to analyze latent-variable model selection methods, we need to define an appropriate notion of model selection consistency for latent-variable graphical models. Given the two components K_O^* and $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$ of the concentration matrix of the marginal distribution (1.1), there are *infinitely* many configurations of the latent variables (i.e., matrices $K_H^* \succ 0, K_{O,H}^* = (K_{H,O}^*)^T$) that give rise to the *same* low-rank matrix $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$. Specifically for any non-singular matrix $B \in \mathbb{R}^{|H| \times |H|}$, one can apply the transformations $K_H^* \rightarrow BK_H^*B^T, K_{O,H}^* \rightarrow K_{O,H}^*B^T$ and still preserve the low-rank matrix $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$. In *all* of these models the marginal statistics of the observed variables X_O remain the same upon marginalization over the latent variables X_H . The key *invariant* is the low-rank matrix $K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$, which summarizes the effect of marginalization over the latent variables. Consequently from here on we use the notation $S^* = K_O^*$ and $L^* = K_{O,H}^*(K_H^*)^{-1}K_{H,O}^*$. These observations give rise to the following notion of structure recovery:

DEFINITION 2.1. *An pair of $|O| \times |O|$ symmetric matrices (\hat{S}, \hat{L}) is an algebraically correct estimate of a latent-variable Gaussian graphical model given by the concentration matrix $K_{(O \ H)}^*$ if the following conditions hold:*

1. *The sign-pattern of \hat{S} is the same as that of S^* (here $\text{sign}(0) = 0$):*

$$\text{sign}(\hat{S}_{i,j}) = \text{sign}(S_{i,j}^*), \quad \forall i, j.$$

2. The rank of \hat{L} is the same as the rank of L^* :

$$\text{rank}(\hat{L}) = \text{rank}(L^*).$$

3. The concentration matrix $\hat{S} - \hat{L}$ can be realized as the marginal concentration matrix of an appropriate latent-variable model:

$$\hat{S} - \hat{L} \succ 0, \quad \hat{L} \succeq 0.$$

When a sequence of estimators is algebraically correct with probability approaching one in a suitable high-dimensional scaling regime, then we say that the estimators are *algebraically consistent*. The first condition ensures that \hat{S} provides the correct structural estimate of the conditional graphical model of the observed variables conditioned on the latent components. This property is the same as the “sparsistency” property studied in standard graphical model selection [20, 28]. The second condition ensures that the number of latent components is properly estimated. Finally, the third condition ensures that the pair of matrices (\hat{S}, \hat{L}) leads to a realizable latent-variable model. In particular this condition implies that there exists a valid latent-variable model in which (a) the conditional graphical model structure among the observed variables is given by \hat{S} , (b) the number of latent variables is equal to the rank of \hat{L} , and (c) the extra correlations induced due to marginalization over the latent variables is equal to \hat{L} . Any method for matrix factorization (e.g., [38]) can be used to further factorize \hat{L} , depending on the property that one desires in the factors (e.g., sparsity).

We also study estimation error rates in the usual sense, i.e., we show that one can produce estimates (\hat{S}, \hat{L}) that are close in various norms to the matrices (S^*, L^*) . Notice that bounding the estimation error in some norm does not in general imply that the support/sign-pattern and rank of (\hat{S}, \hat{L}) are the same as those of (S^*, L^*) . Therefore bounded estimation error is different from algebraic correctness, which requires that (\hat{S}, \hat{L}) have the same support/sign-pattern and rank as (S^*, L^*) .

Goal. Let $K_{(O\ H)}^*$ denote the concentration matrix of a Gaussian model. Suppose that we have n samples $\{X_O^i\}_{i=1}^n$ of the observed variables X_O . We would like to produce estimates (\hat{S}_n, \hat{L}_n) that, with high-probability, are algebraically correct and have bounded estimation error (in some norm).

Our approach. We propose the regularized likelihood convex program (1.2) to produce estimates (\hat{S}_n, \hat{L}_n) . Specifically, the sample covariance matrix Σ_O^n in (1.2) is defined as:

$$\Sigma_O^n \triangleq \frac{1}{n} \sum_{i=1}^n X_O^i X_O^{i\ T}.$$

We give conditions on the underlying model $K_{(O\ H)}^*$ and suitable choices for the parameters λ_n, γ under which the estimates (\hat{S}_n, \hat{L}_n) are consistent (see Theorem 4.1).

2.2. Likelihood function and Fisher information. Given n samples $\{X^i\}_{i=1}^n$ of a finite collection of jointly Gaussian zero-mean random variables with concentration matrix K^* , it is easily seen that the log-likelihood function is given by:

$$(2.1) \quad \ell(K; \Sigma^n) = \log \det(K) - \text{tr}(K \Sigma^n),$$

where $\ell(K; \Sigma^n)$ is a function of K . Notice that this function is strictly concave for $K \succ 0$. In the latent-variable modeling problem with sample covariance Σ_O^n , the likelihood function with respect to the parametrization (S, L) is given by $\ell(S - L; \Sigma_O^n)$. This function is *jointly concave* with respect to the parameters (S, L) whenever $S - L \succ 0$, and it is employed in our variational formulation (1.2) to learn a latent-variable model.

In the analysis of a convex program involving the likelihood function, the Fisher information plays an important role as it is the negative of the Hessian of the likelihood function and thus controls the curvature. As the first term in the likelihood function is linear, we need only study higher-order derivatives of the log-determinant function in order to compute the Hessian. In the latent-variable setting with the marginal concentration matrix of the observed variables given by $\tilde{K}_O^* = (\Sigma_O^*)^{-1}$ (see (1.1)), the corresponding Fisher information matrix is

$$(2.2) \quad \mathcal{I}(\tilde{K}_O^*) = (\tilde{K}_O^*)^{-1} \otimes (\tilde{K}_O^*)^{-1} = \Sigma_O^* \otimes \Sigma_O^*.$$

Here \otimes denotes the tensor product between matrices. Notice that this is precisely the $|O|^2 \times |O|^2$ submatrix of the full Fisher information matrix $\mathcal{I}(K_{(O\ H)}^*) = \Sigma_{(O\ H)}^* \otimes \Sigma_{(O\ H)}^*$ with respect to all the parameters $K_{(O\ H)}^* = (\Sigma_{(O\ H)}^*)^{-1}$ (corresponding to the situation in which *all* the variables $X_{O \cup H}$ are observed). In Section 3.4 we impose various conditions on the Fisher information matrix $\mathcal{I}(\tilde{K}_O^*)$ under which our regularized maximum-likelihood formulation provides consistent estimates.

2.3. Algebraic varieties of sparse and low-rank matrices. The set of sparse matrices and the set of low-rank matrices can be naturally viewed as algebraic varieties (solution sets of systems of polynomial equations). Here we describe these varieties, and discuss some of their geometric properties such as the tangent space and local curvature at a (smooth) point.

Let $\mathcal{S}(k)$ denote the set of matrices with at most k nonzeros:

$$(2.3) \quad \mathcal{S}(k) \triangleq \{M \in \mathbb{R}^{p \times p} \mid |\text{support}(M)| \leq k\}.$$

Here support denotes the locations of nonzero entries. The set $\mathcal{S}(k)$ is an algebraic variety, and can in fact be viewed as a union of $\binom{p^2}{k}$ subspaces in $\mathbb{R}^{p \times p}$. This variety has dimension k , and it is smooth everywhere except at those matrices that have support size strictly smaller than k . For any matrix $M \in \mathbb{R}^{p \times p}$, consider the variety $\mathcal{S}(|\text{support}(M)|)$; M is a smooth point of this variety, and the tangent space at M is given by

$$(2.4) \quad \Omega(M) = \{N \in \mathbb{R}^{p \times p} \mid \text{support}(N) \subseteq \text{support}(M)\}.$$

Next let $\mathcal{L}(r)$ denote the algebraic variety of matrices with rank at most r :

$$(2.5) \quad \mathcal{L}(r) \triangleq \{M \in \mathbb{R}^{p \times p} \mid \text{rank}(M) \leq r\}.$$

It is easily seen that $\mathcal{L}(r)$ is an algebraic variety because it can be defined through the vanishing of all $(r+1) \times (r+1)$ minors. This variety has dimension equal to $r(2p-r)$, and it is smooth everywhere except at those matrices that have rank strictly smaller than r . Consider a rank- r matrix M with singular value decomposition (SVD) given by $M = UDV^T$, where $U, V \in \mathbb{R}^{p \times r}$ and $D \in \mathbb{R}^{r \times r}$. The matrix M is a smooth point of the variety $\mathcal{L}(\text{rank}(M))$, and the tangent space at M with respect to this variety is given by

$$(2.6) \quad T(M) = \{UY_1^T + Y_2V^T \mid Y_1, Y_2 \in \mathbb{R}^{p \times r}\}.$$

We view both $\Omega(M)$ and $T(M)$ as subspaces in $\mathbb{R}^{p \times p}$. In Section 3 we explore the connection between geometric properties of these tangent spaces and the identifiability problem in latent-variable graphical models.

Curvature of rank variety. The sparse matrix variety $\mathcal{S}(k)$ has the property that it has *zero* curvature at any smooth point. The situation is more complicated for the low-rank matrix variety $\mathcal{L}(r)$, because the curvature at any smooth point is nonzero. We analyze how this variety curves locally, by studying how the tangent space changes from one point to a neighboring point. Indeed the amount of curvature at a point is directly related to the “angle” between the tangent space at that point and the tangent space at a neighboring point. For any linear subspace T of matrices, let \mathcal{P}_T denote the projection onto T . Given two subspaces T_1, T_2 of the same dimension, we measure the “twisting” between these subspaces by considering the following quantity.

$$(2.7) \quad \rho(T_1, T_2) \triangleq \|\mathcal{P}_{T_1} - \mathcal{P}_{T_2}\|_{2 \rightarrow 2} = \max_{\|N\|_2 \leq 1} \|[\mathcal{P}_{T_1} - \mathcal{P}_{T_2}](N)\|_2.$$

In the supplement [7] we review relevant results from matrix perturbation theory, which suggest that the magnitude of the smallest nonzero singular value is closely tied to the local curvature of the variety. Therefore we control the twisting between tangent spaces at nearby points by bounding the smallest nonzero singular value away from zero.

3. Identifiability. In the absence of additional conditions, the latent-variable model selection problem is ill-posed. In this section we discuss a set of conditions on latent-variable models that ensure that these models are identifiable given marginal statistics for a subset of the variables. Some of the discussion in Section 3.1 and Section 3.2 is presented in greater detail in [8].

3.1. *Structure between latent and observed variables.* Suppose that the low-rank matrix that summarizes the effect of the latent components is itself sparse. This leads to identifiability issues in the sparse-plus-low-rank decomposition problem. Statistically the additional correlations induced due to marginalization over the latent variables could be mistaken for the conditional graphical model structure of the observed variables. In order to avoid such identifiability problems the effect of the latent variables must be “diffuse” across the observed variables. To address this point the following quantity was introduced in [8] for any matrix M , defined with respect to the tangent space $T(M)$:

$$(3.1) \quad \xi(T(M)) \triangleq \max_{N \in T(M), \|N\|_2 \leq 1} \|N\|_\infty.$$

Thus $\xi(T(M))$ being small implies that elements of the tangent space $T(M)$ cannot have their support concentrated in a few locations; as a result M cannot be too sparse. This idea is formalized in [8] by relating $\xi(T(M))$ to a notion of “incoherence” of the row/column spaces, where the row/column spaces are said to be incoherent with respect to the standard basis if these spaces are not aligned closely with any of the coordinate axes. Typically a matrix M with incoherent row/column spaces would have $\xi(T(M)) \ll 1$. This point is quantified precisely in [8]. Specifically, we note that $\xi(T(M))$ can be as small as $\sim \sqrt{\frac{r}{p}}$ for a rank- r matrix $M \in \mathbb{R}^{p \times p}$ with row/column spaces that are almost maximally incoherent (e.g., if the row/column spaces span any r columns of a $p \times p$ orthonormal Hadamard matrix). On the other hand $\xi(T(M)) = 1$ if the row/column spaces of M contain a standard basis vector.

Based on these concepts we roughly require that the low-rank matrix that summarizes the effect of the latent variables be *incoherent*, thereby ensuring

that the extra correlations due to marginalization over the latent components cannot be confused with the conditional graphical model structure of the observed variables. Notice that the quantity ξ is not just a measure of the number of latent variables, but also of the overall effect of the correlations induced by marginalization over these variables.

Curvature and change in ξ : As noted previously an important technical point is that the algebraic variety of low-rank matrices is locally curved at any smooth point. Consequently the quantity ξ changes as we move along the low-rank matrix variety smoothly. The quantity $\rho(T_1, T_2)$ introduced in (2.7) allows us to bound the variation in ξ as follows (proof in Section 5):

LEMMA 3.1. *Let T_1, T_2 be two linear subspaces of matrices of the same dimension with the property that $\rho(T_1, T_2) < 1$, where ρ is defined in (2.7). Then we have that*

$$\xi(T_2) \leq \frac{1}{1 - \rho(T_1, T_2)} [\xi(T_1) + \rho(T_1, T_2)].$$

3.2. *Structure among observed variables.* An identifiability problem also arises if the conditional graphical model among the observed variables contains a densely connected subgraph. These statistical relationships might be mistaken as correlations induced by marginalization over latent variables. Therefore we need to ensure that the conditional graphical model among the observed variables is sparse. We impose the condition that this conditional graphical model must have small “degree”, i.e., no observed variable is directly connected to too many other observed variables conditioned on the latent components. Notice that bounding the degree is a more refined condition than simply bounding the total number of nonzeros as the *sparsity pattern* also plays a role. In [8] the authors introduced the following quantity in order to provide an appropriate measure of the sparsity pattern of a matrix:

$$(3.2) \quad \mu(\Omega(M)) \triangleq \max_{N \in \Omega(M), \|N\|_\infty \leq 1} \|N\|_2.$$

The quantity $\mu(\Omega(M))$ being small for a matrix implies that the spectrum of any element of the tangent space $\Omega(M)$ is not too “concentrated”, i.e., the singular values of the elements of the tangent space are not too large. In [8] it is shown that a sparse matrix M with “bounded degree” (a small number of nonzeros per row/column) has small $\mu(M)$. Specifically, if $M \in \mathbb{R}^{p \times p}$ is any matrix with at most $\deg(M)$ nonzero entries per row/column, then we have that

$$\mu(\Omega(M)) \leq \deg(M).$$

3.3. *Transversality of tangent spaces.* Suppose that we have the sum of two vectors, each from two known subspaces. It is possible to uniquely recover the individual vectors from the sum if and only if the subspaces have a transverse intersection, i.e., they only intersect at the origin. This simple observation leads to an appealing geometric notion of identifiability. Suppose now that we have the sum of a sparse matrix and a low-rank matrix, and that we are also given the tangent spaces at these matrices with respect to the algebraic varieties of sparse and low-rank matrices respectively. Then a necessary and sufficient condition for identifiability with respect to the tangent spaces is that these spaces have a transverse intersection. This transverse intersection condition is also sufficient for local identifiability in a neighborhood around the sparse matrix and low-rank matrix with respect to the varieties of sparse and low-rank matrices (due to the inverse function theorem). It turns out that these tangent space transversality conditions are also sufficient for the convex program (1.2) to provide consistent estimates of a latent-variable graphical model (without any side information about the tangent spaces).

In order to quantify the level of transversality between the tangent spaces Ω and T we study the *minimum gain* with respect to some norm of the addition operator (which adds two matrices) $\mathcal{A} : \mathbb{R}^{p \times p} \times \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p \times p}$ restricted to the cartesian product $\mathcal{Y} = \Omega \times T$. Then given any matrix norm $\|\cdot\|_q$ on $\mathbb{R}^{p \times p} \times \mathbb{R}^{p \times p}$, the minimum gain of \mathcal{A} restricted to \mathcal{Y} is defined as:

$$\epsilon(\Omega, T, \|\cdot\|_q) \triangleq \min_{(S,L) \in \Omega \times T, \|(S,L)\|_q=1} \|\mathcal{P}_{\mathcal{Y}} \mathcal{A}^\dagger \mathcal{A} \mathcal{P}_{\mathcal{Y}}(S, L)\|_q,$$

where $\mathcal{P}_{\mathcal{Y}}$ denotes the projection onto \mathcal{Y} , and \mathcal{A}^\dagger denotes the adjoint of the addition operator (with respect to the standard Euclidean inner-product). The “level” of transversality of Ω and T is measured by the magnitude of $\epsilon(\Omega, T, \|\cdot\|_q)$, with transverse intersection being equivalent to $\epsilon(\Omega, T, \|\cdot\|_q) > 0$. Note that $\epsilon(\Omega, T, \|\cdot\|_F)$ is the square of the *minimum singular value* of the addition operator \mathcal{A} restricted to $\Omega \times T$.

A natural norm with which to measure transversality is the dual norm of the regularization function in (1.2), as the subdifferential of the regularization function is specified in terms of its dual. The reasons for this will become clearer as we proceed through this paper. Recall that the regularization function used in the variational formulation (1.2) is given by:

$$f_\gamma(S, L) = \gamma \|S\|_1 + \|L\|_*,$$

where the nuclear norm $\|\cdot\|_*$ reduces to the trace function over the cone of positive-semidefinite matrices. This function is a norm for all $\gamma > 0$. The

dual norm of f_γ is given by

$$g_\gamma(S, L) = \max \left\{ \frac{\|S\|_\infty}{\gamma}, \|L\|_2 \right\}.$$

Next we define the quantity $\chi(\Omega, T, \gamma)$ as follows in order to study the transversality of the spaces Ω and T with respect to the g_γ norm:

$$(3.3) \quad \chi(\Omega, T, \gamma) \triangleq \max \left\{ \frac{\xi(T)}{\gamma}, 2\mu(\Omega)\gamma \right\}$$

Here μ and ξ are defined in (3.2) and (3.1). We then have the following result (proved in Section 5):

LEMMA 3.2. *Let $S \in \Omega, L \in T$ be matrices such that $\|S\|_\infty = \gamma$ and let $\|L\|_2 = 1$. Then we have that $g_\gamma(\mathcal{P}_\mathcal{Y} \mathcal{A}^\dagger \mathcal{A} \mathcal{P}_\mathcal{Y}(S, L)) \in [1 - \chi(\Omega, T, \gamma), 1 + \chi(\Omega, T, \gamma)]$, where $\mathcal{Y} = \Omega \times T$ and $\chi(\Omega, T, \gamma)$ is defined in (3.3). In particular we have that $1 - \chi(\Omega, T, \gamma) \leq \epsilon(\Omega, T, g_\gamma)$.*

The quantity $\chi(\Omega, T, \gamma)$ being small implies that the addition operator is essentially isometric when restricted to $\mathcal{Y} = \Omega \times T$. Stated differently the magnitude of $\chi(\Omega, T, \gamma)$ is a measure of the level of transversality of the spaces Ω and T . If $\mu(\Omega)\xi(T) < \frac{1}{2}$ then $\gamma \in (\xi(T), \frac{1}{2\mu(\Omega)})$ ensures that $\chi(\Omega, T, \gamma) < 1$, which in turn implies that the tangent spaces Ω and T have a transverse intersection.

Observation: Thus we have that the smaller the quantities $\mu(\Omega)$ and $\xi(T)$, the more transverse the intersection of the spaces Ω and T as measured by $\epsilon(\Omega, T, g_\gamma)$.

3.4. Conditions on Fisher information. The main focus of Section 4 is to analyze the regularized maximum-likelihood convex program (1.2) by studying its optimality conditions. The log-likelihood function is well-approximated in a neighborhood by a quadratic form given by the Fisher information (which measures the curvature, as discussed in Section 2.2). Let $\mathcal{I}^* = \mathcal{I}(\tilde{K}_\mathcal{O}^*)$ denote the Fisher information evaluated at the true marginal concentration matrix $\tilde{K}_\mathcal{O}^*$ (see (1.1)). The appropriate measure of transversality between the tangent spaces¹ $\Omega = \Omega(S^*)$ and $T = T(L^*)$ is then in a space in which the inner-product is given by \mathcal{I}^* . Specifically, we need to analyze the minimum gain of the operator $\mathcal{P}_\mathcal{Y} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_\mathcal{Y}$ restricted to the

¹We implicitly assume that these tangent spaces are subspaces of the space of *symmetric* matrices.

space $\mathcal{Y} = \Omega \times T$. Therefore we impose several conditions on the Fisher information \mathcal{I}^* . We define quantities that control the gains of \mathcal{I}^* restricted to Ω and T separately; these ensure that elements of Ω and elements of T are individually identifiable under the map \mathcal{I}^* . In addition we define quantities that, in conjunction with bounds on $\mu(\Omega)$ and $\xi(T)$, allow us to control the gain of \mathcal{I}^* restricted to the direct-sum $\Omega \oplus T$.

\mathcal{I}^* restricted to Ω : The minimum gain of the operator $\mathcal{P}_\Omega \mathcal{I}^* \mathcal{P}_\Omega$ restricted to Ω is given by

$$\alpha_\Omega \triangleq \min_{M \in \Omega, \|M\|_\infty=1} \|\mathcal{P}_\Omega \mathcal{I}^* \mathcal{P}_\Omega(M)\|_\infty.$$

The maximum effect of elements in Ω in the orthogonal direction Ω^\perp is given by

$$\delta_\Omega \triangleq \max_{M \in \Omega, \|M\|_\infty=1} \|\mathcal{P}_{\Omega^\perp} \mathcal{I}^* \mathcal{P}_\Omega(M)\|_\infty.$$

The operator \mathcal{I}^* is injective on Ω if $\alpha_\Omega > 0$. The ratio $\frac{\delta_\Omega}{\alpha_\Omega} \leq 1 - \nu$ implies the irrepresentability condition imposed in [28], which gives a sufficient condition for consistent recovery of graphical model structure using ℓ_1 -regularized maximum-likelihood. Notice that this condition is a generalization of the usual Lasso irrepresentability conditions [35, 40], which are typically imposed on the covariance matrix. Finally we also consider the following quantity, which controls the behavior of \mathcal{I}^* restricted to Ω in the spectral norm:

$$\beta_\Omega \triangleq \max_{M \in \Omega, \|M\|_2=1} \|\mathcal{I}^*(M)\|_2.$$

\mathcal{I}^* restricted to T : Analogous to the case of Ω one could control the gains of the operators $\mathcal{P}_{T^\perp} \mathcal{I}^* \mathcal{P}_T$ and $\mathcal{P}_T \mathcal{I}^* \mathcal{P}_T$. However as discussed previously one complication is that the tangent spaces at nearby smooth points on the rank variety are in general different, and the amount of twisting between these spaces is governed by the local curvature. Therefore we control the gains of the operators $\mathcal{P}_{T'^\perp} \mathcal{I}^* \mathcal{P}_{T'}$ and $\mathcal{P}_{T'} \mathcal{I}^* \mathcal{P}_{T'}$ for all tangent spaces T' that are “close to” the nominal T (at the true underlying low-rank matrix), measured by $\rho(T, T')$ (2.7) being small. The minimum gain of the operator $\mathcal{P}_{T'} \mathcal{I}^* \mathcal{P}_{T'}$ restricted to T' (close to T) is given by

$$\alpha_{T'} \triangleq \min_{\rho(T', T) \leq \frac{\xi(T)}{2}} \min_{M \in T', \|M\|_2=1} \|\mathcal{P}_{T'} \mathcal{I}^* \mathcal{P}_{T'}(M)\|_2.$$

Similarly the maximum effect of elements in T' in the orthogonal direction T'^\perp (for T' close to T) is given by

$$\delta_{T'} \triangleq \max_{\rho(T', T) \leq \frac{\xi(T)}{2}} \max_{M \in T', \|M\|_2=1} \|\mathcal{P}_{T'^\perp} \mathcal{I}^* \mathcal{P}_{T'}(M)\|_2.$$

Implicit in the definition of α_T and δ_T is the fact that the outer minimum and maximum are only taken over spaces T' that are tangent spaces to the rank-variety. The operator \mathcal{I}^* is injective on all tangent spaces T' such that $\rho(T', T) \leq \frac{\xi(T)}{2}$ if $\alpha_T > 0$. An irrepresentability condition (analogous to those developed for the sparse case) for tangent spaces near T to the rank variety would be that $\frac{\delta_T}{\alpha_T} \leq 1 - \nu$. Finally we also control the behavior of \mathcal{I}^* restricted to T' close to T in the ℓ_∞ norm:

$$\beta_T \triangleq \max_{\rho(T', T) \leq \frac{\xi(T)}{2}} \max_{M \in T', \|M\|_\infty = 1} \|\mathcal{I}^*(M)\|_\infty.$$

The two sets of quantities $(\alpha_\Omega, \delta_\Omega)$ and (α_T, δ_T) essentially control how \mathcal{I}^* behaves when restricted to the spaces Ω and T *separately* (in the natural norms). The quantities β_Ω and β_T are useful in order to control the gains of the operator \mathcal{I}^* restricted to the *direct sum* $\Omega \oplus T$. Notice that although the magnitudes of elements in Ω are measured most naturally in the ℓ_∞ norm, the quantity β_Ω is specified with respect to the spectral norm. Similarly elements of the tangent spaces T' to the rank variety are most naturally measured in the spectral norm, but β_T provides control in the ℓ_∞ norm. These quantities, combined with $\mu(\Omega)$ and $\xi(T)$ (defined in (3.2) and (3.1)), provide the “coupling” necessary to control the behavior of \mathcal{I}^* restricted to elements in the direct sum $\Omega \oplus T$. In order to keep track of fewer quantities, we summarize the six quantities as follows:

$$\alpha \triangleq \min(\alpha_\Omega, \alpha_T); \quad \delta \triangleq \max(\delta_\Omega, \delta_T); \quad \beta \triangleq \max(\beta_\Omega, \beta_T).$$

Main assumption There exists a $\nu \in (0, \frac{1}{2}]$ such that:

$$\frac{\delta}{\alpha} \leq 1 - 2\nu.$$

This assumption is to be viewed as a generalization of the irrepresentability conditions imposed on the covariance matrix [35, 40] or the Fisher information matrix [28] in order to provide consistency guarantees for sparse model selection using the ℓ_1 norm. With this assumption we have the following proposition, proved in Section 5, about the gains of the operator \mathcal{I}^* restricted to $\Omega \oplus T$. This proposition plays a fundamental role in the analysis of the performance of the regularized maximum-likelihood procedure (1.2). Specifically, it gives conditions under which a suitable primal-dual pair can be specified to certify optimality with respect to (1.2) (see Section 5.2 for more details).

PROPOSITION 3.3. *Let Ω and T be the tangent spaces defined in this section, and let \mathcal{I}^* be the Fisher information evaluated at the true marginal concentration matrix. Further let α, β, ν be as defined above. Suppose that*

$$\mu(\Omega)\xi(T) \leq \frac{1}{6} \left(\frac{\nu\alpha}{\beta(2-\nu)} \right)^2,$$

and that γ is in the following range:

$$\gamma \in \left[\frac{3\xi(T)\beta(2-\nu)}{\nu\alpha}, \frac{\nu\alpha}{2\mu(\Omega)\beta(2-\nu)} \right].$$

Then we have the following two conclusions for $\mathcal{Y} = \Omega \times T'$ with $\rho(T', T) \leq \frac{\xi(T)}{2}$:

1. The minimum gain of \mathcal{I}^* restricted to $\Omega \oplus T'$ is bounded below:

$$\min_{(S,L) \in \mathcal{Y}, \|S\|_\infty = \gamma, \|L\|_2 = 1} g_\gamma(\mathcal{P}_\mathcal{Y} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_\mathcal{Y}(S, L)) \geq \frac{\alpha}{2}.$$

Specifically this implies that for all $(S, L) \in \mathcal{Y}$

$$g_\gamma(\mathcal{P}_\mathcal{Y} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_\mathcal{Y}(S, L)) \geq \frac{\alpha}{2} g_\gamma(S, L).$$

2. The effect of elements in $\mathcal{Y} = \Omega \times T'$ on the orthogonal complement $\mathcal{Y}^\perp = \Omega^\perp \times T'^\perp$ is bounded above:

$$\left\| \mathcal{P}_{\mathcal{Y}^\perp} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_\mathcal{Y} \left(\mathcal{P}_\mathcal{Y} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_\mathcal{Y} \right)^{-1} \right\|_{g_\gamma \rightarrow g_\gamma} \leq 1 - \nu.$$

Specifically this implies that for all $(S, L) \in \mathcal{Y}$

$$g_\gamma(\mathcal{P}_{\mathcal{Y}^\perp} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_\mathcal{Y}(S, L)) \leq (1 - \nu) g_\gamma(\mathcal{P}_\mathcal{Y} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_\mathcal{Y}(S, L)).$$

The last quantity we consider is the spectral norm of the marginal covariance matrix $\Sigma_O^* = (\tilde{K}_O^*)^{-1}$:

$$(3.4) \quad \psi \triangleq \|\Sigma_O^*\|_2 = \|(\tilde{K}_O^*)^{-1}\|_2.$$

A bound on ψ is useful in the probabilistic component of our analysis, in order to derive convergence rates of the sample covariance matrix to the true covariance matrix. We also observe that

$$\|\mathcal{I}^*\|_{2 \rightarrow 2} = \|(\tilde{K}_O^*)^{-1} \otimes (\tilde{K}_O^*)^{-1}\|_{2 \rightarrow 2} = \psi^2.$$

Remarks. The quantities α, β, δ bound the gains of the Fisher information \mathcal{I}^* restricted to the spaces Ω and T (and tangent spaces near T). One can make stronger assumptions on \mathcal{I}^* that are more easily interpretable. For example $\alpha_\Omega, \beta_\Omega$ could bound the minimum/maximum gains of \mathcal{I}^* for *all* matrices (rather than just those in Ω), and δ_Ω the \mathcal{I}^* -inner-product for *all* pairs of orthogonal matrices (rather than just those in Ω and Ω^\perp). Similarly α_T, β_T could bound the minimum/maximum gains of \mathcal{I}^* for all matrices (rather than just those near T), and δ_T the \mathcal{I}^* -inner-product for all pairs of orthogonal matrices (rather than just those near T and T^\perp). Such bounds would apply in either the $\|\cdot\|_{2 \rightarrow 2}$ norm (for $\alpha_T, \delta_T, \beta_\Omega$) or the $\|\cdot\|_{\infty \rightarrow \infty}$ norm (for $\alpha_\Omega, \delta_\Omega, \beta_T$). These modified assumptions are global in nature (not restricted just to Ω or near T) and are consequently stronger (they lower bound the original α_Ω, α_T and they upper bound the original $\beta_\Omega, \beta_T, \delta_\Omega, \delta_T$), and they essentially control the gains of the operator \mathcal{I}^* in the $\|\cdot\|_{2 \rightarrow 2}$ norm and the $\|\cdot\|_{\infty \rightarrow \infty}$ norm. In contrast previous works on covariance selection [1, 2, 31] consider *well-conditioned* families of covariance matrices by bounding the minimum/maximum eigenvalues (i.e., gain with respect to the spectral norm).

4. Consistency of Regularized Maximum-Likelihood Program.

4.1. *Main results.* Recall that $K_{(O \ H)}^*$ denotes the full concentration matrix of a collection of zero-mean jointly-Gaussian observed and latent variables. Let $p = |O|$ denote the number of observed variables, and let $h = |H|$ denote the number of latent variables. We are given n samples $\{X_O^i\}_{i=1}^n$ of the observed variables X_O . We consider the high-dimensional setting in which (p, h, n) are all allowed to grow simultaneously. We present our main result next demonstrating the consistency of the estimator (1.2), and then discuss classes of latent-variable graphical models and various scaling regimes in which our estimator is consistent. Recall from (1.2) that λ_n is a regularization parameter, and γ is a tradeoff parameter between the rank and sparsity terms. Notice from Proposition 3.3 that the choice of γ depends on the values of $\mu(\Omega(S^*))$ and $\xi(T(L^*))$. While these quantities may not be known *a priori*, we discuss a method to choose γ numerically in our experimental results (see Section 6). The following theorem shows that the estimates (\hat{S}_n, \hat{L}_n) provided by the convex program (1.2) are consistent for a suitable choice of λ_n . In addition to the appropriate identifiability conditions (as specified by Proposition 3.3), we also impose lower bounds on the minimum magnitude nonzero entry θ of the sparse conditional graphical model matrix S^* and on the minimum nonzero singular value σ of the low-

rank matrix L^* summarizing the effect of the latent variables. The theorem is stated in terms of the quantities α, β, ν, ψ , and we particularly emphasize the dependence on $\mu(\Omega(S^*))$ and $\xi(T(L^*))$ because these control the complexity of the underlying latent-variable graphical model given by $K_{(O \ H)}^*$. A number quantities play a role in our theorem: Let $D = \max\{1, \frac{\nu\alpha}{3\beta(2-\nu)}\}$, $C_1 = \psi(1 + \frac{\alpha}{6\beta})$, $C_2 = \frac{48}{\alpha} + \frac{1}{\psi^2}$, $C_{\text{samp}} = \frac{\alpha\nu}{32(3-\nu)D} \min\left\{\frac{1}{4C_1}, \frac{\alpha\nu}{256D(3-\nu)\psi C_1^2}\right\}$, $C_\lambda = \frac{48\sqrt{2}D\psi(2-\nu)}{\xi(T)^\nu}$, $C_S = \max\left\{\left(\frac{6(2-\nu)}{\nu} + 1\right) C_2^2\psi^2D, C_2 + \frac{3\alpha C_2^2(2-\nu)}{16(3-\nu)}\right\}$, and $C_L = \frac{C_2\nu\alpha}{\beta(2-\nu)}$.

THEOREM 4.1. *Let $K_{(O \ H)}^*$ denote the concentration matrix of a Gaussian model. We have n samples $\{X_O^i\}_{i=1}^n$ of the p observed variables denoted by O . Let $\Omega = \Omega(S^*)$ and $T = T(L^*)$ denote the tangent spaces at S^* and at L^* with respect to the sparse and low-rank matrix varieties respectively.*

Assumptions: *Suppose that the quantities $\mu(\Omega)$ and $\xi(T)$ satisfy the assumption of Proposition 3.3 for identifiability, and γ is chosen in the range specified by Proposition 3.3. Further suppose that the following conditions hold:*

1. *Let $n \geq \frac{p}{\xi(T)^4} \max\left\{\frac{128\psi^2}{C_{\text{samp}}^2}, 2\right\}$, i.e., we require that $n \gtrsim \frac{p}{\xi(T)^4}$.*
2. *Set $\lambda_n = \frac{48\sqrt{2}D\psi(2-\nu)}{\xi(T)^\nu} \sqrt{\frac{p}{n}}$, i.e., we require that $\lambda_n \asymp \frac{1}{\xi(T)} \sqrt{\frac{p}{n}}$.*
3. *Let $\sigma \geq \frac{C_L\lambda_n}{\xi(T)^2}$, i.e., we require that $\sigma \gtrsim \frac{1}{\xi(T)^3} \sqrt{\frac{p}{n}}$.*
4. *Let $\theta \geq \frac{C_S\lambda_n}{\mu(\Omega)}$, i.e., we require that $\theta \gtrsim \frac{1}{\xi(T)\mu(\Omega)} \sqrt{\frac{p}{n}}$.*

Conclusions: *Then with probability greater than $1 - 2\exp\{-p\}$ we have algebraic correctness and estimation error given by:*

1. $\text{sign}(\hat{S}_n) = \text{sign}(S^*)$ and $\text{rank}(\hat{L}_n) = \text{rank}(L^*)$;
2. $g_\gamma(\hat{S}_n - S^*, \hat{L}_n - L^*) \leq \frac{512\sqrt{2}(3-\nu)D\psi}{\nu\alpha\xi(T)} \sqrt{\frac{p}{n}} \lesssim \frac{1}{\xi(T)} \sqrt{\frac{p}{n}}$.

The proof of this theorem is given in Section 5. The theorem essentially states that if the minimum nonzero singular value of the low-rank piece L^* and minimum nonzero entry of the sparse piece S^* are bounded away from zero, then the convex program (1.2) provides estimates that are both algebraically correct and have bounded estimation error (in the ℓ_∞ and spectral norms).

Notice that the condition on the minimum singular value of L^* is more stringent than the one on the minimum nonzero entry of S^* . One role played by these conditions is to ensure that the estimates (\hat{S}_n, \hat{L}_n) do not have

smaller support size/rank than (S^*, L^*) . However the minimum singular value bound plays the additional role of bounding the curvature of the low-rank matrix variety around the point L^* , which is the reason for this condition being more stringent. Notice also that the number of latent variables h does not explicitly appear in the bounds in Theorem 4.1, which only depend on $p, \mu(\Omega(S^*)), \xi(T(L^*))$. However the dependence on h is implicit in the dependence on $\xi(T(L^*))$, and we discuss this point in greater detail in the following section.

Finally we note that consistency holds in Theorem 4.1 for a *range* of values of $\gamma \in \left[\frac{3\beta(2-\nu)\xi(T)}{\nu\alpha}, \frac{\nu\alpha}{2\beta(2-\nu)\mu(\Omega)} \right]$. In particular the assumptions on the sample complexity, the minimum nonzero singular value of L^* , and the minimum magnitude nonzero entry of S^* are governed by the lower end of this range for γ . These assumptions can be weakened if we only require consistency for a smaller range of values of γ . The next result conveys this point with a specific example.

COROLLARY 4.2. *Consider the same setup and notation as in Theorem 4.1. Suppose that the quantities $\mu(\Omega)$ and $\xi(T)$ satisfy the assumption of Proposition 3.3 for identifiability, and that $\gamma = \frac{\nu\alpha}{2\beta(2-\nu)\mu(\Omega)}$ (the upper end of the range specified in Proposition 3.3), i.e., $\gamma \asymp \frac{1}{\mu(\Omega)}$. Further suppose that: (1) $n \gtrsim \mu(\Omega)^4 p$; (2) $\lambda_n \asymp \mu(\Omega)\sqrt{\frac{p}{n}}$; (3) $\sigma \gtrsim \frac{\mu(\Omega)^2}{\xi(T)}\sqrt{\frac{p}{n}}$; (4) $\theta \gtrsim \sqrt{\frac{p}{n}}$. Then with probability greater than $1 - 2\exp\{-p\}$ we have estimates (\hat{S}_n, \hat{L}_n) that are algebraically correct, and with the error bounded as $g_\gamma(\hat{S}_n - S^*, \hat{L}_n - L^*) \lesssim \mu(\Omega)\sqrt{\frac{p}{n}}$.*

The proof of this corollary² is analogous to that of Theorem 4.1. We emphasize that in practice it is often beneficial to have consistent estimates for a range of values of γ (as in Theorem 4.1). Specifically the stability of the sparsity pattern and rank of the estimates (\hat{S}_n, \hat{L}_n) for a range of tradeoff parameters is useful in order to choose a suitable value of γ , as prior information about the quantities $\mu(\Omega(S^*))$ and $\xi(T(L^*))$ is not typically available (see Section 6).

We remark here that the identifiability conditions of Proposition 3.3 are the main sufficient conditions required for Theorem 4.1 and Corollary 4.2 to

²By making stronger assumptions on the Fisher information matrix \mathcal{I}^* , one can further remove the factor of $\xi(T)$ in the lower bound for σ . Specifically the lower bound $\sigma \gtrsim \mu(\Omega)^3\sqrt{\frac{p}{n}}$ suffices for consistent estimation if the bounds defined by the quantities $\alpha_T, \beta_T, \delta_T$ can be strengthened as described in the remarks at the end of Section 3.4.

hold. It would be interesting to obtain necessary conditions as well for these results, analogous to the necessity and sufficiency of the irrepresentability conditions for the Lasso [35, 40].

4.2. Scaling regimes. Next we consider classes of latent-variable models that satisfy the conditions of Theorem 4.1. Recall from Section 3.2 that $\mu(\Omega(S^*)) \leq \deg(S^*)$. Throughout this section, we consider latent-variable models in which the low-rank matrix L^* is almost maximally incoherent, i.e., $\xi(T(L^*)) \sim \sqrt{\frac{h}{p}}$ so the effect of marginalization over the latent variables is diffuse across almost all the observed variables. We suppress the dependence on the quantities α, β, ν, ψ defined in Section 3.4 in our scaling results, and specifically focus on the tradeoff between $\xi(T(L^*))$ and $\mu(\Omega(S^*))$ for consistent estimation (we also suppress the dependence of these quantities on n). Thus, based on Proposition 3.3 we study latent-variable models in which

$$\xi(T(L^*)) \mu(\Omega(S^*)) = \mathcal{O}\left(\sqrt{\frac{h}{p}} \deg(S^*)\right) = \mathcal{O}(1).$$

As we describe next, there are non-trivial classes of latent-variable graphical models in which this condition holds.

Bounded degree: The first class of latent-variable models that we consider are those in which the conditional graphical model among the observed variables (given by K_O^*) has constant degree:

$$\deg(S^*) = \mathcal{O}(1), \quad h \sim p.$$

Such models can be estimated consistently from $n \sim p$ samples. Thus consistent latent-variable model selection is possible even when the number of samples and the number of latent variables are on the same order as the number of observed variables.

Polylogarithmic degree: The next class of models that we consider are those in which the degree of the conditional graphical model of the observed variables grows polylogarithmically with p :

$$\deg(S^*) \sim \log(p)^q, \quad h \sim \frac{p}{\log(p)^{2q}},$$

Such latent-variable graphical models can be consistently estimated as long as $n \sim p \text{ polylog}(p)$.

For standard graphical model selection with no latent variables ℓ_1 -regularized maximum likelihood is shown to be consistent with $n = \mathcal{O}(\log p)$ samples [28]. On the other hand our results prove consistency in the setting with latent variables when $n = \mathcal{O}(p)$ samples. It would be interesting to study whether these rates are inherent to latent-variable model selection.

4.3. *Rates for covariance matrix estimation.* Theorem 4.1 gives conditions under which we can consistently estimate the sparse and low-rank parts that compose the marginal concentration matrix \tilde{K}_O^* . Here we state a corollary that gives rates for covariance matrix estimation, i.e., the quality of the estimate $(\hat{S}_n - \hat{L}_n)^{-1}$ with respect to the “true” marginal covariance matrix Σ_O^* .

COROLLARY 4.3. *Under the same conditions as in Theorem 4.1, we have with probability greater than $1 - 2 \exp\{-p\}$ that*

$$g_\gamma(\mathcal{A}^\dagger[(\hat{S}_n - \hat{L}_n)^{-1} - \Sigma_O^*]) \leq \lambda_n \left[1 + \frac{\nu}{6(2-\nu)} \right].$$

This corollary implies that $\|(\hat{S}_n - \hat{L}_n)^{-1} - \Sigma_O^*\|_2 \lesssim \frac{1}{\xi(T)} \sqrt{\frac{p}{n}}$ based on the choice of λ_n in Theorem 4.1, and that $\|(\hat{S}_n - \hat{L}_n)^{-1} - \Sigma_O^*\|_2 \lesssim \mu(\Omega) \sqrt{\frac{p}{n}}$ based on the choice of λ_n in Corollary 4.2.

5. Proofs.

5.1. *Proofs of Section 3.* Here we give proofs of the results stated in Section 3.

Proof of Lemma 3.1: Since $\rho(T_1, T_2) < 1$ the largest principal angle between T_1 and T_2 is strictly less than $\frac{\pi}{2}$. Consequently, the mapping $\mathcal{P}_{T_2} : T_1 \rightarrow T_2$ restricted to T_1 is bijective (as it is injective, and the spaces T_1, T_2 have the same dimension). Consider the maximum and minimum gains of \mathcal{P}_{T_2} restricted to T_1 ; for any $M \in T_1$, $\|M\|_2 = 1$:

$$\|\mathcal{P}_{T_2}(M)\|_2 = \|M + [\mathcal{P}_{T_2} - \mathcal{P}_{T_1}](M)\|_2 \in [1 - \rho(T_1, T_2), 1 + \rho(T_1, T_2)].$$

Therefore, we can rewrite $\xi(T_2)$ as follows:

$$\begin{aligned} \xi(T_2) &= \max_{N \in T_2, \|N\|_2 \leq 1} \|N\|_\infty = \max_{N \in T_2, \|N\|_2 \leq 1} \|\mathcal{P}_{T_2}(N)\|_\infty \\ &\leq \max_{N \in T_1, \|N\|_2 \leq \frac{1}{1 - \rho(T_1, T_2)}} \|\mathcal{P}_{T_2}(N)\|_\infty \\ &\leq \max_{N \in T_1, \|N\|_2 \leq \frac{1}{1 - \rho(T_1, T_2)}} [\|N\|_\infty + \|[\mathcal{P}_{T_1} - \mathcal{P}_{T_2}](N)\|_\infty] \\ &\leq \frac{1}{1 - \rho(T_1, T_2)} \left[\xi(T_1) + \max_{N \in T_1, \|N\|_2 \leq 1} \|[\mathcal{P}_{T_1} - \mathcal{P}_{T_2}](N)\|_\infty \right] \\ &\leq \frac{1}{1 - \rho(T_1, T_2)} \left[\xi(T_1) + \max_{\|N\|_2 \leq 1} \|[\mathcal{P}_{T_1} - \mathcal{P}_{T_2}](N)\|_2 \right] \\ &\leq \frac{1}{1 - \rho(T_1, T_2)} [\xi(T_1) + \rho(T_1, T_2)]. \end{aligned}$$

This concludes the proof of the lemma. \square

Proof of Lemma 3.2: We have that $\mathcal{A}^\dagger \mathcal{A}(S, L) = (S + L, S + L)$; therefore, $\mathcal{P}_Y \mathcal{A}^\dagger \mathcal{A} \mathcal{P}_Y(S, L) = (S + \mathcal{P}_\Omega(L), \mathcal{P}_T(S) + L)$. We need to bound $\|S + \mathcal{P}_\Omega(L)\|_\infty$ and $\|\mathcal{P}_T(S) + L\|_2$. First, we have

$$\begin{aligned} \|S + \mathcal{P}_\Omega(L)\|_\infty &\in [\|S\|_\infty - \|\mathcal{P}_\Omega(L)\|_\infty, \|S\|_\infty + \|\mathcal{P}_\Omega(L)\|_\infty] \\ &\subseteq [\|S\|_\infty - \|L\|_\infty, \|S\|_\infty + \|L\|_\infty] \\ &\subseteq [\gamma - \xi(T), \gamma + \xi(T)]. \end{aligned}$$

Similarly, one can check that

$$\begin{aligned} \|\mathcal{P}_T(S) + L\|_2 &\in [-\|\mathcal{P}_T(S)\|_2 + \|L\|_2, \|\mathcal{P}_T(S)\|_2 + \|L\|_2] \\ &\subseteq [1 - 2\|S\|_2, 1 + 2\|S\|_2] \\ &\subseteq [1 - 2\gamma\mu(\Omega), 1 + 2\gamma\mu(\Omega)]. \end{aligned}$$

These two bounds give us the desired result. \square

Proof of Proposition 3.3: Before proving the two parts of this proposition we make a simple observation about $\xi(T')$ using the condition that $\rho(T, T') \leq \frac{\xi(T)}{2}$ by applying Lemma 3.1:

$$\xi(T') \leq \frac{\xi(T) + \rho(T, T')}{1 - \rho(T, T')} \leq \frac{\frac{3\xi(T)}{2}}{1 - \frac{\xi(T)}{2}} \leq 3\xi(T).$$

Here we used the property that $\xi(T) \leq 1$ in obtaining the final inequality. Consequently, noting that $\gamma \in \left[\frac{3\beta(2-\nu)\xi(T)}{\nu\alpha}, \frac{\nu\alpha}{2\beta(2-\nu)\mu(\Omega)} \right]$ implies that

$$(5.1) \quad \chi(\Omega, T', \gamma) = \max \left\{ \frac{\xi(T')}{\gamma}, 2\mu(\Omega)\gamma \right\} \leq \frac{\nu\alpha}{\beta(2-\nu)}.$$

Part 1: The proof of this step proceeds in a similar manner to that of Lemma 3.2. First we have for $S \in \Omega, L \in T'$ with $\|S\|_\infty = \gamma, \|L\|_2 = 1$:

$$\|\mathcal{P}_\Omega \mathcal{I}^*(S + L)\|_\infty \geq \|\mathcal{P}_\Omega \mathcal{I}^* S\|_\infty - \|\mathcal{P}_\Omega \mathcal{I}^* L\|_\infty \geq \alpha\gamma - \|\mathcal{I}^* L\|_\infty \geq \alpha\gamma - \beta\xi(T').$$

Next under the same conditions on S, L ,

$$\|\mathcal{P}_{T'} \mathcal{I}^*(S + L)\|_2 \geq \|\mathcal{P}_{T'} \mathcal{I}^* L\|_2 - \|\mathcal{P}_{T'} \mathcal{I}^* S\|_2 \geq \alpha - 2\|\mathcal{I}^* S\|_2 \geq \alpha - 2\beta\mu(\Omega)\gamma.$$

Combining these last two bounds with (5.1), we conclude that

$$\begin{aligned} &\min_{(S, L) \in \mathcal{Y}, \|S\|_\infty = \gamma, \|L\|_2 = 1} g_\gamma(\mathcal{P}_Y \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_Y(S, L)) \\ &\geq \alpha - \beta \max \left\{ \frac{\xi(T')}{\gamma}, 2\mu(\Omega)\gamma \right\} \geq \alpha - \frac{\nu\alpha}{2-\nu} = \frac{2\alpha(1-\nu)}{2-\nu} \geq \frac{\alpha}{2}, \end{aligned}$$

where the final inequality follows from the assumption that $\nu \in (0, \frac{1}{2}]$.

Part 2: Note that for $S \in \Omega, L \in T'$ with $\|S\|_\infty \leq \gamma, \|L\|_2 \leq 1$

$$\|\mathcal{P}_{\Omega^\perp} \mathcal{I}^*(S + L)\|_\infty \leq \|\mathcal{P}_{\Omega^\perp} \mathcal{I}^* S\|_\infty + \|\mathcal{P}_{\Omega^\perp} \mathcal{I}^* L\|_\infty \leq \delta\gamma + \beta\xi(T').$$

Similarly

$$\|\mathcal{P}_{T'^\perp} \mathcal{I}^*(S + L)\|_2 \leq \|\mathcal{P}_{T'^\perp} \mathcal{I}^* S\|_2 + \|\mathcal{P}_{T'^\perp} \mathcal{I}^* L\|_2 \leq \beta\gamma\mu(\Omega) + \delta.$$

Combining these last two bounds with the bounds from the first part, we have that

$$\begin{aligned} & \left\| \mathcal{P}_{\mathcal{Y}^\perp} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_{\mathcal{Y}} (\mathcal{P}_{\mathcal{Y}} \mathcal{A}^\dagger \mathcal{I}^* \mathcal{A} \mathcal{P}_{\mathcal{Y}})^{-1} \right\|_{g_\gamma \rightarrow g_\gamma} \\ & \leq \frac{\delta + \beta \max\left\{\frac{\xi(T')}{\gamma}, 2\mu(\Omega)\gamma\right\}}{\alpha - \beta \max\left\{\frac{\xi(T')}{\gamma}, 2\mu(\Omega)\gamma\right\}} \leq \frac{\delta + \frac{\nu\alpha}{2-\nu}}{\alpha - \frac{\nu\alpha}{2-\nu}} \leq \frac{(1-2\nu)\alpha + \frac{\nu\alpha}{2-\nu}}{\alpha - \frac{\nu\alpha}{2-\nu}} = 1 - \nu. \end{aligned}$$

This concludes the proof of the proposition. \square

5.2. Proof strategy for Theorem 4.1. Standard results from convex analysis [30] state that (\hat{S}_n, \hat{L}_n) is a minimum of the convex program (1.2) if the zero matrix belongs to the subdifferential of the objective function evaluated at (\hat{S}_n, \hat{L}_n) (in addition to (\hat{S}_n, \hat{L}_n) satisfying the constraints). Elements of the subdifferentials with respect to the ℓ_1 norm and the nuclear norm at a matrix M have the key property that they decompose with respect to the tangent spaces $\Omega(M)$ and $T(M)$ [37]. This decomposition property plays a critical role in our analysis. In particular it states that the optimality conditions consist of two parts, one part corresponding to the tangent spaces Ω and T and another corresponding to the normal spaces Ω^\perp and T^\perp .

Our analysis proceeds by constructing a primal-dual pair of variables that certify optimality with respect to (1.2). Consider the optimization problem (1.2) with the additional (non-convex) constraints that the variable S belongs to the algebraic variety of sparse matrices and that the variable L belongs to the algebraic variety of low-rank matrices. While this new optimization problem is non-convex, it has a very interesting property. At a globally optimal solution (and indeed at any locally optimal solution) (\tilde{S}, \tilde{L}) such that \tilde{S} and \tilde{L} are smooth points of the algebraic varieties of sparse and low-rank matrices, the first-order optimality conditions state that the Lagrange multipliers corresponding to the additional variety constraints must lie in the *normal spaces* $\Omega(\tilde{S})^\perp$ and $T(\tilde{L})^\perp$. This basic observation, combined with the decomposition property of the subdifferentials of the ℓ_1 and nuclear norms, suggests the following high-level proof strategy: Considering the solution (\tilde{S}, \tilde{L}) of the variety-constrained problem, we show under suitable

conditions that the second part of the subgradient optimality conditions of (1.2) (without any variety constraints) corresponding to components in the normal spaces $\Omega(\tilde{S})^\perp$ and $T(\tilde{L})^\perp$ is also satisfied by (\tilde{S}, \tilde{L}) . Thus, we show that (\tilde{S}, \tilde{L}) satisfy the optimality conditions of the *original convex program* (1.2). Consequently (\tilde{S}, \tilde{L}) is also the optimum of the convex program (1.2). As this estimate is obtained as the solution to the problem with the variety constraints, the algebraic correctness of (\tilde{S}, \tilde{L}) can be directly concluded. We emphasize here that the variety-constrained optimization problem is used solely as an analysis tool in order to prove consistency of the estimates provided by the convex program (1.2). The key technical complication is that the tangent spaces at \tilde{L} and L^* are in general different. We bound the twisting between these tangent spaces by using the fact that the minimum nonzero singular value of L^* is bounded away from zero (as assumed in Theorem 4.1; see also the supplement [7]).

5.3. Results proved in supplement. In this section we give the statements of some results that are proved in a separate supplement [7]. These results are critical to the proof of our main theorem, but they deal mainly with non-statistical aspects such as the curvature of the algebraic variety of low-rank matrices. Recall that $\Omega = \Omega(S^*)$ and $T = T(L^*)$. We also refer frequently to the constants defined in Theorem 4.1.

As the gradient of the log-determinant function is given by a matrix inverse, a key step in analyzing the properties of the convex program (1.2) is to show that the change in the inverse of a matrix due to small perturbations is well-approximated by the first-order term in the Taylor series expansion. Consider the Taylor series of the inverse of a matrix:

$$(M + \Delta)^{-1} = M^{-1} - M^{-1}\Delta M^{-1} + R_{M^{-1}}(\Delta),$$

where

$$R_{M^{-1}}(\Delta) = M^{-1} \left[\sum_{k=2}^{\infty} (-\Delta M^{-1})^k \right].$$

This infinite sum converges for Δ sufficiently small. The following proposition provides a bound on the second-order term specialized to our setting:

PROPOSITION 5.1. *Suppose that γ is in the range given by Proposition 3.3. Further suppose $\Delta_S \in \Omega$, and let $g_\gamma(\Delta_S, \Delta_L) \leq \frac{1}{2C_1}$. Then we have that*

$$g_\gamma(\mathcal{A}^\dagger R_{\Sigma_O^*}(\mathcal{A}(\Delta_S, \Delta_L))) \leq \frac{2D\psi C_1^2 g_\gamma(\Delta_S, \Delta_L)^2}{\xi(T)}.$$

Next we analyze the following convex program subject to certain additional constraints:

$$(5.2) \quad \begin{aligned} (\hat{S}_\Omega, \hat{L}_{\tilde{T}}) &= \arg \min_{S, L} \text{tr}[(S - L) \Sigma_O^n] - \log \det(S - L) + \lambda_n [\gamma \|S\|_1 + \|L\|_*] \\ \text{s.t. } & S - L \succ 0, \quad S \in \Omega, \quad L \in \tilde{T}, \end{aligned}$$

for some subspace \tilde{T} . Comparing (5.2) with the convex program (1.2), we also do not constrain the variable L to be positive semidefinite in (5.2) for ease of proof of the next result (see the supplement [7] for more details; recall that the nuclear norm of a positive semidefinite matrix is equal to its trace). We show that if \tilde{T} is any tangent space to the low-rank matrix variety such that $\rho(T, \tilde{T}) \leq \frac{\xi(T)}{2}$, then we can bound the error $(\Delta_S, \Delta_L) = (\hat{S}_\Omega - S^*, L^* - \hat{L}_{\tilde{T}})$. Let $\mathcal{C}_{\tilde{T}} = \mathcal{P}_{\tilde{T}^\perp}(L^*)$ denote the normal component of the true low-rank matrix at \tilde{T} , and let $E_n = \Sigma_O^n - \Sigma_O^*$ denote the difference between the true marginal covariance and the sample covariance. The proof of the following result uses Brouwer's fixed-point theorem [27], and is inspired by the proof of a similar result in [28] for standard sparse graphical model recovery without latent variables.

PROPOSITION 5.2. *Let the error (Δ_S, Δ_L) in the solution of the convex program (5.2) (with \tilde{T} such that $\rho(\tilde{T}, T) \leq \frac{\xi(T)}{2}$) be as defined above, and define*

$$r = \max \left\{ \frac{8}{\alpha} \left[g_\gamma(\mathcal{A}^\dagger E_n) + g_\gamma(\mathcal{A}^\dagger T^* \mathcal{C}_{\tilde{T}}) + \lambda_n \right], \|\mathcal{C}_{T'}\|_2 \right\}.$$

If $r \leq \min \left\{ \frac{1}{4C_1}, \frac{\alpha\xi(T)}{64D\psi C_1^2} \right\}$ for γ as in Proposition 3.3, then $g_\gamma(\Delta_S, \Delta_L) \leq 2r$.

Finally we give a proposition that summarizes the algebraic component of our proof.

PROPOSITION 5.3. *Assume that γ is in the range specified by Proposition 3.3, $\sigma \geq \frac{C_L \lambda_n}{\xi(T)^2}$, $\theta \geq \frac{C_S \lambda_n}{\mu(\Omega)}$, $g_\gamma(\mathcal{A}^\dagger E_n) \leq \frac{\lambda_n \nu}{6(2-\nu)}$, and that $\lambda_n \leq \frac{3\alpha(2-\nu)}{16(3-\nu)} \min \left\{ \frac{1}{4C_1}, \frac{\alpha\xi(T)}{64D\psi C_1^2} \right\}$. Then there exists a T' and a corresponding unique solution $(\hat{S}_\Omega, \hat{L}_{T'})$ of (5.2) with $\tilde{T} = T'$ with the following properties:*

1. $\text{sign}(\hat{S}_\Omega) = \text{sign}(S^*)$ and $\text{rank}(\hat{L}_{T'}) = \text{rank}(L^*)$, with $\hat{L}_{T'} \succeq 0$. Further $T(\hat{L}_{T'}) = T'$ and $\rho(T, T') \leq \frac{\xi(T)}{4}$.
2. Letting $\mathcal{C}_{T'} = \mathcal{P}_{T'^\perp}(L^*)$ we have that $g_\gamma(\mathcal{A}^\dagger T^* \mathcal{C}_{T'}) \leq \frac{\lambda_n \nu}{6(2-\nu)}$, and that $\|\mathcal{C}_{T'}\|_2 \leq \frac{16(3-\nu)\lambda_n}{3\alpha(2-\nu)}$.

Further if $g_\gamma(\mathcal{A}^\dagger R_{\Sigma_O^*}(\mathcal{A}(\hat{S}_\Omega - S^*, L^* - \hat{L}_{T'}))) \leq \frac{\lambda_n \nu}{6(2-\nu)}$ then the tangent space constraints $S \in \Omega, L \in T'$ are inactive in (5.2). Consequently the unique solution of (1.2) is $(\hat{S}_n, \hat{L}_n) = (\hat{S}_\Omega, \hat{L}_{T'})$.

5.4. *Probabilistic analysis.* The results given thus far in this section have been completely deterministic in nature. Here we present the probabilistic component of our proof by studying the rate at which the sample covariance matrix Σ_O^n converges to the true covariance matrix Σ_O^* in spectral norm. This result is well-known and follows directly from Theorem II.13 in [9]; we mainly discuss it here for completeness and also to show explicitly the dependence on $\psi = \|\Sigma_O^*\|_2$ defined in (3.4). See the supplement [7] for a proof.

LEMMA 5.4. *Let $\psi = \|\Sigma_O^*\|_2$. Given any $\delta > 0$ with $\delta \leq 8\psi$, let the number of samples n be such that $n \geq \frac{64p\psi^2}{\delta^2}$. Then we have that*

$$\Pr[\|\Sigma_O^n - \Sigma_O^*\|_2 \geq \delta] \leq 2 \exp\left\{-\frac{n\delta^2}{128\psi^2}\right\}.$$

The following corollary relates the number of samples required for an error bound to hold with probability $1 - 2 \exp\{-p\}$.

COROLLARY 5.5. *Let Σ_O^n be the sample covariance formed from n samples of the observed variables. Set $\delta_n = \sqrt{\frac{128p\psi^2}{n}}$. If $n \geq 2p$ then*

$$\Pr[\|\Sigma_O^n - \Sigma_O^*\|_2 \leq \delta_n] \geq 1 - 2 \exp\{-p\}.$$

Proof: Note that $n \geq 2p$ implies that $\delta_n \leq 8\psi$, and apply Lemma 5.4. \square

5.5. *Proof of Theorem 4.1 and Corollary 4.3.* We first combine the results obtained thus far to prove Theorem 4.1. Set $E_n = \Sigma_O^n - \Sigma_O^*$, set $\delta_n = \sqrt{\frac{128p\psi^2}{n}}$, and then set $\lambda_n = \frac{6D\delta_n(2-\nu)}{\xi(T)\nu}$. This setting of λ_n is equivalent to the specification in the statement of Theorem 4.1.

Proof of Theorem 4.1: We mainly need to show that the various sufficient conditions of Proposition 5.3 are satisfied. We condition on the event that $\|E_n\|_2 \leq \delta_n$, which holds with probability greater than $1 - 2 \exp\{-p\}$ from Corollary 5.5 as $n \geq 2p$ by assumption. Based on the bound on n , we also have that

$$\delta_n \leq \xi(T)^2 \left[\frac{\alpha\nu}{32(3-\nu)D} \min\left\{\frac{1}{4C_1}, \frac{\alpha\nu}{256D(3-\nu)\psi C_1^2}\right\} \right].$$

In particular, these bounds imply that

$$(5.3) \quad \delta_n \leq \frac{\alpha\xi(T)\nu}{32(3-\nu)D} \min\left\{\frac{1}{4C_1}, \frac{\alpha\xi(T)}{64D\psi C_1^2}\right\} \quad ; \quad \delta_n \leq \frac{\alpha^2\xi(T)^2\nu^2}{8192\psi C_1^2(3-\nu)^2 D^2}.$$

Both these weaker bounds are used later.

Based on the assumptions of Theorem 4.1, the requirements of Proposition 5.3 on σ and θ are satisfied. Next we verify the bounds on λ_n and $g_\gamma(\mathcal{A}^\dagger E_n)$. Based on the setting of λ_n above and the bound on δ_n from (5.3), we have that

$$\lambda_n = \frac{6D(2-\nu)\delta_n}{\xi(T)^\nu} \leq \frac{3\alpha(2-\nu)}{16(3-\nu)} \min \left\{ \frac{1}{4C_1}, \frac{\alpha\xi(T)}{64D\psi C_1^2} \right\}.$$

Next we combine the facts that $\lambda_n = \frac{6D\delta_n(2-\nu)}{\xi(T)^\nu}$ and that $\|E_n\|_2 \leq \delta_n$ to conclude that

$$(5.4) \quad g_\gamma(\mathcal{A}^\dagger E_n) \leq \frac{D\delta_n}{\xi(T)} = \frac{\lambda_n\nu}{6(2-\nu)}.$$

Thus, we have from Proposition 5.3 that there exists a T' and corresponding solution $(\hat{S}_\Omega, \hat{L}_{T'})$ of (5.2) with the prescribed properties. Next we apply Proposition 5.2 with $\tilde{T} = T'$ to bound the error $(\hat{S}_\Omega - S^*, L^* - \hat{L}_{T'})$. Noting that $\rho(T, T') \leq \frac{\xi(T)}{4}$, we have that

$$(5.5) \quad \frac{8}{\alpha} \left[g_\gamma(\mathcal{A}^\dagger E_n) + g_\gamma(\mathcal{A}^\dagger \mathcal{I}^* \mathcal{C}_{T'}) + \lambda_n \right] \leq \frac{8}{\alpha} \left[\frac{\nu}{3(2-\nu)} + 1 \right] \lambda_n$$

$$(5.6) \quad = \frac{16(3-\nu)\lambda_n}{3\alpha(2-\nu)}$$

$$(5.7) \quad = \frac{32(3-\nu)D}{\alpha\xi(T)^\nu} \delta_n$$

$$(5.7) \quad \leq \min \left\{ \frac{1}{4C_1}, \frac{\alpha\xi(T)}{64D\psi C_1^2} \right\}.$$

In the first inequality we used the fact that $g_\gamma(\mathcal{A}^\dagger E_n) \leq \frac{\lambda_n\nu}{6(2-\nu)}$ (from above) and that $g_\gamma(\mathcal{A}^\dagger \mathcal{I}^* \mathcal{C}_{T'})$ is similarly bounded (from Proposition 5.3). In the second equality we used the relation $\lambda_n = \frac{6D\delta_n(2-\nu)}{\xi(T)^\nu}$. In the final inequality we used the bound on δ_n from (5.3). This satisfies one of the requirements of Proposition 5.2. The second requirement of Proposition 5.2 on $\|\mathcal{C}_{T'}\|_2$ is also similarly satisfied as we have that $\|\mathcal{C}_{T'}\|_2 \leq \frac{16(3-\nu)\lambda_n}{3\alpha(2-\nu)}$ from Proposition 5.3, and we use the same sequence of inequalities as above. Thus we conclude from Proposition 5.2 and from (5.5) that

$$(5.8) \quad g_\gamma(\hat{S}_\Omega - S^*, L^* - \hat{L}_{T'}) \leq \frac{32(3-\nu)\lambda_n}{3\alpha(2-\nu)} \lesssim \frac{1}{\xi(T)} \sqrt{\frac{p}{n}}.$$

Here the last inequality follows from the bound on λ_n .

If we show that $(\hat{S}_n, \hat{L}_n) = (\hat{S}_\Omega, \hat{L}_{T'})$ we can conclude the proof of Theorem 4.1 since algebraic correctness of $(\hat{S}_\Omega, \hat{L}_{T'})$ holds from Proposition 5.3 and the estimation error bound follows from (5.8). In order to complete this

final step, we again revert to Proposition 5.3 and prove the requisite bound on $g_\gamma(\mathcal{A}^\dagger R_{\Sigma_O^*}(\mathcal{A}(\hat{S}_\Omega - S^*, L^* - \hat{L}_{T'})$)).

Since the bound (5.8) combined with the inequality (5.7) satisfies the condition of Proposition 5.1 (i.e., we have that $g_\gamma(\hat{S}_\Omega - S^*, L^* - \hat{L}_{T'}) \leq \frac{1}{2C_1}$):

$$\begin{aligned}
g_\gamma(\mathcal{A}^\dagger R_{\Sigma_O^*}(\mathcal{A}(\hat{S}_\Omega - S^*, L^* - \hat{L}_{T'}))) &\leq \frac{2D\psi C_1^2}{\xi(T)} g_\gamma(\hat{S}_\Omega - S^*, L^* - \hat{L}_{T'})^2 \\
&\leq \frac{2D\psi C_1^2}{\xi(T)} \left(\frac{64(3-\nu)D}{\alpha\xi(T)\nu} \right)^2 \delta_n^2 \\
&= \left[\frac{8192\psi C_1^2 (3-\nu)^2 D^2}{\alpha^2 \xi(T)^2 \nu^2} \delta_n \right] \frac{D\delta_n}{\xi(T)} \\
&\leq \frac{D\delta_n}{\xi(T)} \\
&= \frac{\lambda_n \nu}{6(2-\nu)}.
\end{aligned}$$

In the second inequality we used (5.6) and (5.8), in the final inequality we used the bound (5.3) on δ_n , and in the final equality we used the relation $\lambda_n = \frac{6D\delta_n(2-\nu)}{\xi(T)\nu}$. \square

Proof of Corollary 4.3: Based on the optimality conditions of the modified convex program (5.2), we have that

$$g_\gamma(\mathcal{A}^\dagger[(\hat{S}_n - \hat{L}_n)^{-1} - \Sigma_O^n]) \leq \lambda_n.$$

Combining this with the bound (5.4) yields the desired result. \square

6. Simulation Results. In this section we give experimental demonstration of the consistency of our estimator (1.2) on synthetic examples, and its effectiveness in modeling real-world stock return data. Our choices of λ_n and γ are guided by Theorem 4.1. Specifically, we choose λ_n to be proportional to $\sqrt{\frac{p}{n}}$. For γ we observe that the support/sign-pattern and the rank of the solution (\hat{S}_n, \hat{L}_n) are the same for a *range* of values of γ . Therefore one could solve the convex program (1.2) for several values of γ , and choose a solution in a suitable range in which the sign-pattern and rank of the solution are stable (see [8] for details). In practical problems with real-world data these parameters may be chosen via cross-validation (it would be of interest to consider methods such as those developed in [26]). For small problem instances we solve the convex program (1.2) using a combination of YALMIP [23] and SDPT3 [33]. For larger problem instances we use the special purpose solver LogdetPPA [36] developed for log-determinant semidefinite programs.

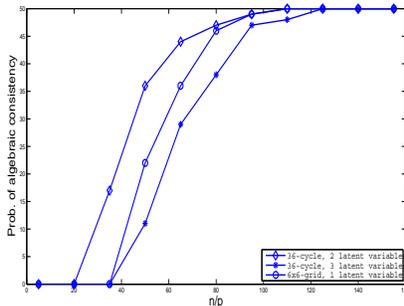


FIG 1. *Synthetic data: Plot showing probability of algebraically correct estimation. The three models studied are (a) 36-node conditional graphical model given by a cycle with $h = 2$ latent variables, (b) 36-node conditional graphical model given by a cycle with $h = 3$ latent variables, and (c) 36-node conditional graphical model given by a 6×6 grid with $h = 1$ latent variable. For each plotted point, the probability of algebraically correct estimation is obtained over 50 random trials.*

6.1. *Synthetic data.* In the first set of experiments we consider a setting in which we have access to samples of the observed variables of a latent-variable graphical model. We consider several latent-variable Gaussian graphical models. The first model consists of $p = 36$ observed variables and $h = 2$ latent variables. The conditional graphical model structure of the observed variables is a cycle with the edge partial correlation coefficients equal to 0.25; thus, this conditional model is specified by a sparse graphical model with degree 2. The second model is the same as the first one, but with $h = 3$ latent variables. The third model consists of $h = 1$ latent variable, and the conditional graphical model structure of the observed variables is given by a 6×6 nearest-neighbor grid (i.e., $p = 36$ and degree 4) with the partial correlation coefficients of the edges equal to 0.15. In all three of these models each latent variable is connected to a random subset of 80% of the observed variables (and the partial correlation coefficients corresponding to these edges are also random). Therefore the effect of the latent variables is “spread out” over most of the observed variables, i.e., the low-rank matrix summarizing the effect of the latent variables is incoherent.

For each model we generate n samples of the observed variables, and use the resulting sample covariance $\Sigma_{\mathcal{O}}^n$ as input to our convex program (1.2). Figure 1 shows the probability of obtaining algebraically correct estimates as a function of n . This probability is evaluated over 50 experiments for each value of n . In all of these cases standard graphical model selection applied directly to the observed variables is not useful as the marginal concentration

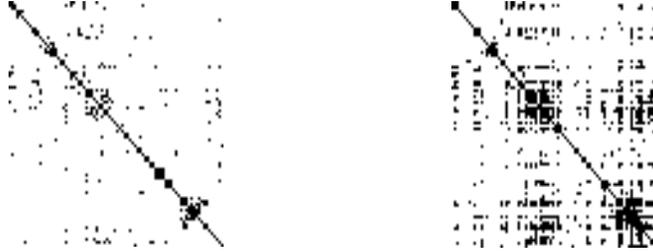


FIG 2. *Stock returns: The figure on the left shows the sparsity pattern (black denotes an edge, and white denotes no edge) of the concentration matrix of the conditional graphical model (135 edges) of the stock returns, conditioned on 5 latent variables, in a latent-variable graphical model (total number of parameters equals 639). This model is learned using (1.2), and the KL divergence with respect to a Gaussian distribution specified by the sample covariance is 17.7. The figure on the right shows the concentration matrix of the graphical model (646 edges) of the stock returns, learned using standard sparse graphical model selection based on solving an ℓ_1 -regularized maximum-likelihood program (total number of parameters equals 730). The KL divergence between this distribution and a Gaussian distribution specified by the sample covariance is 44.4.*

matrix of the observed variables is not well-approximated by a sparse matrix. These experiments agree with our theoretical results that the convex program (1.2) is an algebraically consistent estimator of a latent-variable model given (sufficiently many) samples of only the observed variables.

6.2. *Stock return data.* In the next experiment we model the statistical structure of monthly stock returns of 84 companies in the S&P 100 index from 1990 to 2007; we disregard 16 companies that were listed after 1990. The number of samples n is equal to 216. We compute the sample covariance based on these returns and use this as input to (1.2).

The model learned using (1.2) for suitable values of λ_n, γ consists of $h = 5$ latent variables, and the conditional graphical model structure of the stock returns conditioned on these latent components consists of 135 edges. Therefore the number of parameters in the model is $84 + 135 + (5 \times 84) = 639$. The resulting KL divergence between the distribution specified by this model and a Gaussian distribution specified by the sample covariance is 17.7. Figure 2 (left) shows the *conditional* graphical model structure. The strongest edges in this conditional graphical model, as measured by partial correlation, are between Baker Hughes - Schlumberger, A.T.&T. - Verizon, Merrill Lynch - Morgan Stanley, Halliburton - Baker Hughes, Intel - Texas Instruments, Apple - Dell, and Microsoft - Dell. It is of interest to note that in the Standard

Industrial Classification³ system for grouping these companies, several of these pairs are in different classes. As mentioned in Section 2.1 our method estimates a low-rank matrix that summarizes the effect of the latent variables; in order to factorize this low-rank matrix, for example into sparse factors, one could use methods such as those described in [38].

We compare these results to those obtained using a sparse graphical model learned using ℓ_1 -regularized maximum-likelihood (see for example [28]), without introducing any latent variables. Figure 2 (right) shows this graphical model structure. The number of edges in this model is 646 (the total number of parameters is equal to $646 + 84 = 730$), and the resulting KL divergence between this distribution and a Gaussian distribution specified by the sample covariance is 44.4.

These results suggest that a latent-variable graphical model is better suited than a standard sparse graphical model for modeling stock returns. This is likely due to the presence of global, long-range correlations in stock return data that are better modeled via latent variables.

7. Discussion. We have studied the problem of modeling the statistical structure of a collection of random variables as a sparse graphical model conditioned on a few additional latent components. As a first contribution we described conditions under which such latent-variable graphical models are identifiable given samples of only the observed variables. We also proposed a convex program based on ℓ_1 and nuclear norm regularized maximum-likelihood for latent-variable graphical model selection. Given samples of the observed variables of a latent-variable Gaussian model we proved that this convex program provides consistent estimates of the number of latent components as well as the conditional graphical model structure among the observed variables conditioned on the latent components. Our analysis holds in the high-dimensional regime in which the number of observed/latent variables are allowed to grow with the number of samples of the observed variables. These theoretical predictions are verified via a set of experiments on synthetic data. We also demonstrate the effectiveness of our approach in modeling real-world stock return data.

Several questions arise that are worthy of further investigation. While (1.2) can be solved in polynomial time using off-the-shelf solvers, it is preferable to develop more efficient special-purpose solvers to scale to massive datasets by taking advantage of the structure of (1.2). It is also of interest to develop statistically consistent convex optimization methods for latent-variable modeling with non-Gaussian variables, e.g., for categorical data.

³See the U.S. SEC website at <http://www.sec.gov/info/edgar/siccodes.htm>

Acknowledgements. We would like to thank James Saunderson and Myung Jin Choi for helpful discussions, and Kim-Chuan Toh for kindly providing us specialized code to solve larger instances of our convex program.

References.

- [1] BICKEL, P. J. AND LEVINA, E. (2008). Regularized estimation of large covariance matrices. *Ann. Statistics*. **36** 199–227.
- [2] BICKEL, P. J. AND LEVINA, E. (2008). Covariance regularization by thresholding. *Ann. Statistics*. **36** 2577–2604.
- [3] BOYD, S. P. AND VANDENBERGHE, L. (2004). *Convex optimization*. Cambridge University Press.
- [4] CANDÈS, E. J., ROMBERG, J., AND TAO, T. (2006). Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Info. Theory*. **52** 489–509.
- [5] CANDÈS, E. J. AND RECHT, B. (2009). Exact matrix completion via convex optimization. *Found. of Comput. Math.* **9** 717–772.
- [6] CANDÈS, E. J., LI, X., MA, Y. AND WRIGHT, J. (2009). Robust principal component analysis? *Preprint*.
- [7] CHANDRASEKARAN, V., PARRILO, P. A. AND WILLSKY, A. S. (2011). Latent Variable Graphical Model Selection via Convex Optimization – Supplementary Material.
- [8] CHANDRASEKARAN, V., SANGHAVI, S., PARRILO, P. A. AND WILLSKY, A. S. (2009). Rank-sparsity incoherence for matrix decomposition. *SIAM Jour. Opt.* **21** 572–596.
- [9] DAVIDSON, K. R. AND SZAREK, S.J. (2001). Local operator theory, random matrices and Banach spaces. *Handbook of the Geometry of Banach Spaces*. **I** 317–366.
- [10] DEMPSTER, A. P., LAIRD, N. M., AND RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc. B*. **39** 1–38.
- [11] DONOHO, D. L. (2006). For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Comm. on Pure and Applied Math.* **59** 797–829.
- [12] DONOHO, D. L. (2006). Compressed sensing. *IEEE Trans. Info. Theory*. **52** 1289–1306.
- [13] ELIDAN, G., NACHMAN, I., AND FRIEDMAN, N. (2007). “Ideal Parent” structure learning for continuous variable Bayesian networks. *J. Mach. Lear. Res.* **8** 1799–1833.
- [14] EL KAROUI, N. (2008). Operator norm consistent estimation of large-dimensional sparse covariance matrices. *Ann. Statistics*. **36** 2717–2756.
- [15] FAN, J., FAN, Y., AND LV, J. (2008). High dimensional covariance matrix estimation using a factor model. *J. Econometrics*. **147** 186–197.
- [16] FAZEL, M. (2002). Matrix rank minimization with applications. *PhD thesis, Dep. Elec. Eng., Stanford University*. 2002.
- [17] GOLUB, G. H. AND VAN LOAN, C. H. (1990). *Matrix computations*. The Johns Hopkins Univ. Press.
- [18] HORN, R. A. AND JOHNSON, C. R. (1990). *Matrix analysis*. Cambridge University Press.
- [19] JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statistics*. **29** 295–327.
- [20] LAM, C. AND FAN, J. (2009). Sparsistency and rates of convergence in large covariance matrix estimation. *Ann. Statistics*. **37** 4254–4278.
- [21] LAURITZEN, S. L. (1996). *Graphical models*. Oxford University Press.

- [22] LEDOIT, O. AND WOLF, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivar. Analysis.* **88** 365–411.
- [23] LÖFBERG, J. (2004). YALMIP: A Toolbox for Modeling and Optimization in MATLAB. *Proceedings of the CACSD Conference, Taiwan*. Available from <http://control.ee.ethz.ch/~joloef/yalmip.php>.
- [24] MARCENKO, V. A. AND PASTUR, L. A. (1967). Distributions of eigenvalues of some sets of random matrices. *Math. USSR-Sb.* **1** 507–536.
- [25] MEINSHAUSEN, N. AND BÜHLMANN, P. (2006). High dimensional graphs and variable selection with the Lasso. *Ann. Statistics.* **34** 1436–1462.
- [26] MEINSHAUSEN, N. AND BÜHLMANN, P. (2010). Stability selection. *Jour. Roy. Stat. Soc., Ser. B.* **72** 417–473.
- [27] ORTEGA, J. M. AND RHEINBOLDT, W. G. (1970). *Iterative solution of nonlinear equations in several variables*. Academic Press.
- [28] RAVIKUMAR, P., WAINWRIGHT, M. J., RASKUTTI, G., AND YU, B. (2011). High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Elec. Jour. of Stat.* **4** 935–980.
- [29] RECHT, B., FAZEL, M., AND PARRILO, P. A. (2010). Guaranteed minimum rank solutions to linear matrix equations via nuclear norm minimization. *SIAM Rev.* **52** 471–501.
- [30] ROCKAFELLAR, R. T. (1996). *Convex Analysis*. Princeton University Press.
- [31] ROTHMAN, A. J., BICKEL, P. J., LEVINA, E., AND ZHU, J. (2008). Sparse permutation invariant covariance estimation. *Elec. J. Statistics.* **2** 494–515.
- [32] SPEED, T. P. AND KIIVERI, H. T. (1986). Gaussian Markov distributions over finite graphs. *Ann. Statistics.* **14** 138–150.
- [33] K. C. TOH, M. J. TODD, AND R. H. TUTUNCU. *SDPT3 - a MATLAB software package for semidefinite-quadratic-linear programming*. Available from <http://www.math.nus.edu.sg/~mattohkc/sdpt3.html>.
- [34] VARGA, B. S. (2000). *Matrix iterative analysis*. Springer-Verlag.
- [35] WAINWRIGHT, M. J. (2009). Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans. Info. Theory.* **55** 2183–2202.
- [36] WANG, C., SUN, D. AND TOH, K. C. (2009). Solving log-determinant optimization problems by a Newton-CG primal proximal point algorithm. *Preprint*.
- [37] WATSON, G. A. (1992). Characterization of the subdifferential of some matrix norms. *Linear Algebra and Applications.* **170** 1039–1053.
- [38] WITTEN, D. M., TIBSHIRANI, R. AND HASTIE, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostat.* **10** 515–534.
- [39] WU, W. B. AND POURAHMADI, M. (2003). Nonparametric estimation of large covariance matrices of longitudinal data. *Biometrika.* **90** 831–844.
- [40] ZHAO, P. AND YU, B. (2006). On model selection consistency of lasso. *J. Mach. Lear. Res.* **7** 2541–2567.

LABORATORY FOR INFORMATION AND DECISION SYSTEMS
 DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
 MASSACHUSETTS INSTITUTE OF TECHNOLOGY
 CAMBRIDGE, MA 02139
 USA
 E-MAIL: venkatc@mit.edu
 parrilo@mit.edu
 willsky@mit.edu