

---

# Learning Latent Variable Gaussian Graphical Models

---

**Zhaoshi Meng**

MENGZS@UMICH.EDU

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

**Brian Eriksson**

BRIAN.ERIKSSON@TECHNICOLOR.COM

Technicolor Research Center, 735 Emerson Street, Palo Alto, CA 94301, USA

**Alfred O. Hero III**

HERO@EECS.UMICH.EDU

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

## Abstract

Gaussian graphical models (GGM) have been widely used in many high-dimensional applications ranging from biological and financial data to recommender systems. Sparsity in GGM plays a central role both statistically and computationally. Unfortunately, real-world data often does not fit well to sparse graphical models. In this paper, we focus on a family of latent variable Gaussian graphical models (LVGGM), where the model is conditionally sparse given latent variables, but marginally non-sparse. In LVGGM, the inverse covariance matrix has a low-rank plus sparse structure, and can be learned in a regularized maximum likelihood framework. We derive novel parameter estimation error bounds for LVGGM under mild conditions in the high-dimensional setting. These results complement the existing theory on the structural learning, and open up new possibilities of using LVGGM for statistical inference.

## 1. Introduction

Critical to many statistical inference tasks in complex real-world systems, such as prediction and detection, is the ability to extract and estimate distributional characteristics from the observations. Unfortunately, in the high-dimensional regime such model estimation often leads to ill-posed problems, particularly when the number of observations  $n$  (or sample size) is comparable to or fewer than the ambient dimensionality  $p$  of the model (*i.e.*, the “large  $p$ , small  $n$ ” problem). This challenge arises in

many modern real-world applications ranging from recommender systems, gene microarray data, and financial data, to name a few. To perform accurate model parameter estimation and subsequent statistical inference, low dimensional structure is often imposed for regularization (Negahban et al., 2012).

For Gaussian-distributed data, the central problem is often to estimate the inverse covariance matrix (alternatively known as the precision, concentration or information matrix). Gaussian graphical models (GGM) provide an efficient representation of the precision matrix through a graph that represents non-zeros in the matrix (Lauritzen, 1996). In high-dimensional regimes, this graph can be forced to be sparse, imposing a low-dimensional structure on the GGM. For sufficiently sparse GGM, statistically consistent estimates of the model structure (*i.e.*, sparsistency) can be achieved (*e.g.*, Ravikumar et al. (2011)). On the computational side, sparsity also leads to reduced complexity of the estimator (Hsieh et al., 2013). However, when the true distribution can not be well-approximated by a sparse GGM, the standard learning paradigm suffers from either large estimation bias due to enforcing a overly sparse model, or degraded computation time for a dense model. Both result in suboptimal performance in the subsequent inference tasks.

In this paper, we consider a new class of high-dimensional GGM for extending the standard sparse GGM. The proposed model is motivated by many real-world applications, where there exist certain exogenous and often latent factors affecting a large portion of the variables. Examples are the price of oil on the airlines’ stock price variables (Choi et al., 2010), and the genres on movie rating variables. Conditioning on these *global* effects, the variables are assumed to have highly *localized* interactions, which can be well-fitted by a sparse GGM. However, due to the marginalization over global effects, the observed (marginal) GGM, and its corresponding precision matrix, is not sparse. Unfortunately, in this regime, existing theoretical results and com-

putational tools for sparse GGM are not applicable.

To address this problem, we propose to use latent variable Gaussian graphical models (LVGGM) for modeling and statistical inference. LVGGM introduce latent variables to capture the correlations due to the global effects, and the remaining effects are captured by a conditionally sparse graphical model. The resulting marginal precision matrix of the LVGGM has a sparse plus low-rank structure, therefore we consider a regularized maximum likelihood (ML) approach for parameter estimation (previously considered by Chandrasekaran et al. (2012)). By utilizing the *almost strong convexity* (Kakade et al., 2010) of the log-likelihood, we derive a non-asymptotic parameter error bound for the regularized ML estimator. Our derived bounds apply to the high-dimensional setting of  $p \gg n$  due to restricted strong convexity (Negahban et al., 2012) and certain structural incoherence between the sparse and low-rank components of the precision matrix (Yang & Ravikumar, 2013).

We show that for sufficiently large  $n$ , the Frobenius norm error of the precision matrix of LVGGM converges at the rate  $\mathcal{O}(\sqrt{\frac{(s+r_{\text{eff}} \cdot r) \log p}{n}})$ , where  $s$  is the number of non-zeros in the conditionally sparse precision matrix,  $r_{\text{eff}}$  is the effective rank of the covariance matrix and  $r$  is the number of latent variables. This rate is in general significantly faster than the standard convergence rate of  $\mathcal{O}(\sqrt{\frac{p^2 \log p}{n}})$  for an unstructured dense GGM. This result offers a compelling argument for using LVGGM over sparse GGM for many inference problems.

The paper is structured as follows. In Section 2 we review the relevant prior literature. In Section 3 we formulate the LVGGM estimation problem. In Section 4 the main theoretical results are presented. Experimental results are shown in Section 5 and we conclude in Section 6. We use boldface letters to denote vectors and matrices.  $\|\cdot\|_1$ ,  $\|\cdot\|_2$ ,  $\|\cdot\|_F$ ,  $\|\cdot\|_*$  denote the elementwise  $\ell_1$ , spectral, Frobenius, and nuclear matrix norms, respectively.

## 2. Background and Related Work

The problem of learning GGM with sparse inverse covariance matrices using  $\ell_1$ -regularized maximum likelihood estimation, often referred to as the graphical lasso (Glasso) problem, has been studied in Friedman et al. (2008); Ravikumar et al. (2011); Rothman et al. (2008). In particular, the authors of Ravikumar et al. (2011) study the model selection consistency (*i.e.*, “sparsistency”) under certain incoherence condition. Beyond sparse GGM, Choi et al. (2010) propose a multi-resolution extension of a GGM augmented with sparse inter-level correlations, while in Choi et al. (2011) the authors consider latent tree-structured graphical models. Both models lead to computa-

tionally efficient inference and learning algorithms but restrict the latent structure to trees. Recently, Liu & Willsky (2013) consider a computationally efficient learning algorithm for a class of conditionally tree-structured LVGGM.

The work that is most relevant to ours is by Chandrasekaran et al. (2012), who study the LVGGM learning problem, but focus on the simultaneous model selection consistency of both the sparse and low-rank components. In contrast, in this paper we focus on the Frobenius norm error bounds for estimating the precision matrix of LVGGM. Although structural consistency can be useful for deriving insights, parameter estimation error analysis is of equal or greater importance in practice. Since it provides additional, and usually more direct, insights into factors influencing the performance of the subsequent statistical inference tasks, such as prediction and detection. Also, compared with Chandrasekaran et al. (2012), our Frobenius norm error bounds are derived under mild condition on the Fisher information of the distribution.

We note that there is a fundamentally different line of work on estimating models with a similar structural composition, known as *robust PCA* (Candès et al., 2011). In robust PCA, the data matrix is modeled as “low-rank plus sparse”. This model has been applied to extracting the salient foreground from background in videos, and detecting malicious user ratings in recommender system data (Xu et al., 2012). In contrast, the equivalent covariance model of our LVGGM can be decomposed into a low-rank plus a dense matrix whose *inverse* is sparse. A similar covariance model has recently been studied by Kalaitzis & Lawrence (2012), in which an EM algorithm is proposed for estimation but no theoretical error bounds are derived. In this paper, we instead focus on the precision matrix parameterization, which enables model estimation through a convex optimization. This formulation is of both theoretical and computational importance.

## 3. Problem Setup

In this section, we review Gaussian graphical models and formulate the problem of latent variable Gaussian graphical model estimation via a regularized maximum likelihood optimization.

### 3.1. Gaussian Graphical Models

Consider a  $p$ -dimensional random vector  $\mathbf{x}$  associated with an undirected graph  $\mathcal{G} = (V_G, E_G)$ , where  $V_G$  is a set of nodes corresponding to elements of  $\mathbf{x}$  and  $E_G$  is a set of edges connecting nodes (including self-edges for each node). Then  $\mathbf{x}$  follows a graphical model distribution if it satisfies the Markov property with respect to  $\mathcal{G}$ : for any pair of nonadjacent nodes in  $\mathcal{G}$ , the corresponding pair of

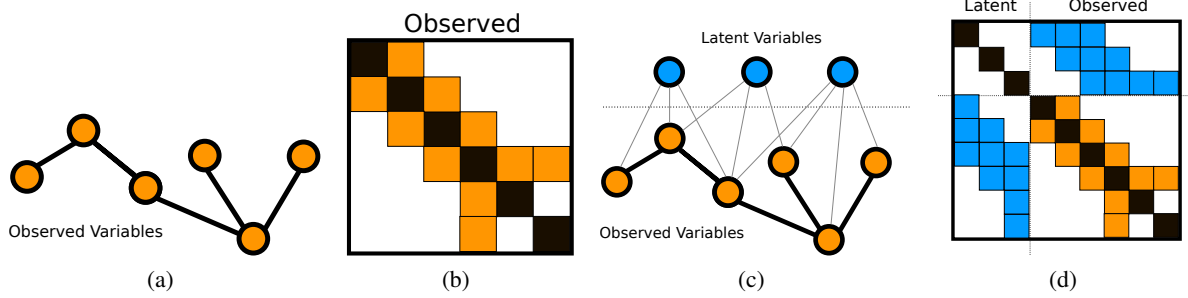


Figure 1. Illustrations of a sparse Gaussian graphical model (GGM) (left) and a latent variable Gaussian graphical model (LVGGM) (right). (A) Example of a sparse GGM with only observed variables, (B) Sparsity pattern of example sparse GGM’s precision matrix, (C) Example of a LVGGM with both observed and latent variables, (D) Sparsity pattern of example LVGGM’s precision matrix.

variables in  $\mathbf{x}$  are conditionally independent given the remaining variables, *i.e.*,  $x_i \perp x_j \mid \mathbf{x}_{\setminus\{i,j\}}$ , for all  $(i, j) \notin E_G$ .

If  $\mathbf{x}$  follows a multivariate Gaussian distribution, the corresponding graphical model is called a Gaussian graphical model (GGM). We assume without loss of generality that  $\mathbf{x}$  has zero mean. The Markov property in GGM is manifested in the sparsity pattern of the inverse covariance matrix  $\mathbf{J}$ :  $\mathbf{J}_{i,j} = 0$  for all  $i \neq j, (i, j) \notin E$ . An example of this property for sparse GGM is shown in Figure 1(a) and 1(b).

The precision matrix parameterization arises in many statistical inference problems for Gaussian distributions, in areas such as belief propagation, linear prediction, portfolio selection in financial data, and anomaly detection. Estimation of the precision matrix in GGM is the first step in these inference problems.

### 3.2. Latent Variable Gaussian Graphical Models

Unfortunately, due to the presence of global factors that destroy sparsity, real-world observations often do not conform exactly to a sparse GGM (Choi et al., 2010; 2011). By introducing latent variables (denoted as a  $r$ -dimensional random vector  $\mathbf{x}_L$ ) to capture global factors, we can generalize the GGM. Specifically, we construct a model that is conditionally a GGM, *i.e.*, one that has a sparse precision matrix given knowledge of latent variables,  $\mathbf{x}_L$ .

Defining the  $p$  observed variables as  $\mathbf{x}_O$ , we assume the joint distribution of the  $(p + r)$ -dimensional concatenated random vector  $\mathbf{x} = (\mathbf{x}_O, \mathbf{x}_L)$  follows a Gaussian distribution with covariance matrix  $\Omega$  and precision matrix  $\mathbf{J} = \Omega^{-1}$ . An example of this structure can be seen in Figure 1(c) and 1(d). Marginalizing over the latent variables  $\mathbf{x}_L$ , the distribution of the observed variables  $\mathbf{x}_O$  remains Gaussian with observed covariance matrix,  $\Sigma = \Omega_{O,O}$ .

The observed precision matrix  $\Theta \in \mathbb{R}^{p \times p}$  satisfies:

$$\Theta = \Sigma^{-1} = \underbrace{\mathbf{J}_{O,O}}_{\mathbf{S}} - \underbrace{\mathbf{J}_{O,L} \mathbf{J}_{L,L}^{-1} \mathbf{J}_{L,O}}_{\mathbf{L}}, \quad (1)$$

where we have defined  $\mathbf{S} := \mathbf{J}_{O,O}$  and  $\mathbf{L} := -\mathbf{J}_{O,L} \mathbf{J}_{L,L}^{-1} \mathbf{J}_{L,O}$ . Thus, the marginal precision matrix can be written as  $\Theta = \mathbf{S} + \mathbf{L}$ , the sum of a sparse and a low-rank matrix. Similar to standard GGM, we parameterize the marginal distribution through the precision matrix. We refer to this model as the latent variable GGM, or LVGGM.

The LVGGM is a hierarchical model that generalizes the (sparse) GGM. Note that  $\mathbf{S}^{-1} = \mathbf{J}_{O,O}^{-1} = \Omega_{O,O} - \Omega_{O,L} \Omega_{L,L}^{-1} \Omega_{L,O}$  is the covariance matrix of the conditional distribution of the observed variables. The matrix is not generally sparse, even though  $\mathbf{S}$  is assumed to be sparse. We will also assume that the number of latent variables is much smaller than the number of observed variables, *i.e.*,  $r \ll p$ . We place no sparsity restrictions on the dependencies between the observed and latent variables – the submatrices  $\mathbf{J}_{O,L}$  and  $\mathbf{J}_{L,O}$  could be dense. As a result, the  $p \times p$  matrix  $\mathbf{L} = -\mathbf{J}_{O,L} \mathbf{J}_{L,L}^{-1} \mathbf{J}_{L,O}$  is low-rank and potentially dense. The sparse plus low-rank structure of the marginal precision matrix  $\Theta$  is the key property of the precision matrix that will be exploited for model estimation.

The structural assumptions on the precision matrix of the LVGGM can be further motivated and validated on real-world recommender system data and stock return data. Due to the space limits, we defer these two motivating examples to Section A in the Appendix.

### 3.3. Effective Rank of Covariance Matrix

We introduce the *effective rank* of a matrix, which will be useful to derived high-dimensional error bounds. The effective rank of a matrix  $\Sigma$  is defined as (Vershynin, 2010):

$$r_{\text{eff}}(\Sigma) := \text{tr}(\Sigma) / \|\Sigma\|_2. \quad (2)$$

The effective rank can be considered a measure of the concentration level of the spectrum of  $\Sigma$ . As we will show in Section 5.1, in many situations the effective rank of the covariance matrix corresponding to a LVGGM is much smaller than  $p$ . Under this condition, our theoretical results in the sequel provide a tight Frobenius norm estimation error bound, which is significantly improved upon the error bound derived without the effective rank assumption.

### 3.4. Regularized ML Estimation of LVGGM

Available are  $n$  samples  $x_1, x_2, \dots, x_n$  from a LVGGM model  $\mathbf{x}_O$ , concatenated into a data matrix  $\mathbf{X} \in \mathbb{R}^{p \times n}$ . The negative log-likelihood function is

$$\mathcal{L}(\Theta; \mathbf{X}) = \langle \widehat{\Sigma}, \Theta \rangle - \log \det(\Theta), \quad (3)$$

where  $\widehat{\Sigma} := \frac{1}{n} \mathbf{X}^T \mathbf{X}$  is the sample covariance matrix. The regularized ML estimate minimizes the objective function  $\mathcal{L}(\Theta; \mathbf{X}) + \lambda \mathcal{R}(\Theta)$ , where the regularization parameter  $\lambda > 0$ , and the regularization function  $\mathcal{R}(\Theta)$  is designed to enforce the sparse plus low-rank structure on  $\Theta$ .

Similar to Chandrasekaran et al. (2012), we consider the following regularized ML estimation problem:

$$\begin{aligned} \min_{\mathbf{S}, \mathbf{L}} \quad & \mathcal{L}(\mathbf{S} + \mathbf{L}; \mathbf{X}) + \lambda \|\mathbf{S}\|_1 + \mu \|\mathbf{L}\|_* \\ \text{s.t.} \quad & -\mathbf{L} \succeq \mathbf{0}, \mathbf{S} + \mathbf{L} \succeq \mathbf{0}, \end{aligned} \quad (4)$$

where the corresponding regularization function is the sum of two regularizers:  $\mathcal{R}(\Theta) = \|\mathbf{S}\|_1 + \frac{\mu}{\lambda} \|\mathbf{L}\|_*$ , each of which has been shown to promote sparse (low-rank) structure in  $\mathbf{S}$  ( $\mathbf{L}$ , respectively) (Negahban et al., 2012). Constants  $\lambda, \mu > 0$  are regularization parameters corresponding to the two functions, respectively. The LVGGM estimator is defined as a solution to the above convex optimization problem (4). Efficient convex solver, such as Ma et al. (2013), can be used to solve.

## 4. Error Bounds on ML LVGGM Estimation

We analyze the regularized ML estimation problem (4) and provide Frobenius norm error bounds for estimating the precision matrix in high-dimensional setting. We adopt the decomposable regularization framework of (Negahban et al., 2012; Agarwal et al., 2012; Yang & Ravikumar, 2013) to derive these bounds. In contrast to this prior work, here we focus on multiple decomposable regularizers interacting with the non-quadratic log-likelihood loss function encountered in the LVGGM. Two important ingredients in the derivations are the restricted strong convexity of the loss function, and an incoherence condition between the two structured subspaces containing the sparse and low-rank components ( $\mathbf{S}$  and  $\mathbf{L}$ ). We show that under assumptions on the Fisher information these two conditions are verified.

In the following subsections, first we define some necessary notation, then we introduce the assumptions and place them in the context of prior literature, and finally we state the main results in Theorem 1 and Theorem 2.

### 4.1. Decomposable Regularizers and Subspace Notation

In this subsection we introduce the notion of decomposable regularizers and the corresponding subspace pairs. We refer the reader to Negahban et al. (2012) for more details.

Consider a pair of subspaces  $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ , where  $\mathcal{M} \subset \overline{\mathcal{M}} \subset \mathbb{R}^{p \times p}$ .  $\mathcal{R}(\cdot)$  is called a decomposable regularization function with respect to the subspace pair if, for any  $u \in \mathcal{M}, v \in \overline{\mathcal{M}}^\perp$ , we have  $\mathcal{R}(u + v) = \mathcal{R}(u) + \mathcal{R}(v)$ .

For the sparse and low-rank matrix-valued parameters, the following two subspace pairs and their corresponding decomposable regularizers are considered:

- *Sparse matrices.* Let  $E \subseteq \{1, \dots, p\} \times \{1, \dots, p\}$  be a subset of index pairs (edges). Define  $\mathcal{M}(E) = \overline{\mathcal{M}}(E)$  as the subspace of all sparse matrices in  $\mathbb{R}^{p \times p}$  that are supported in subsets of  $E$ , i.e.,  $\mathcal{P}_{\mathcal{M}(E)}(\mathbf{A}) = \mathbf{A}_E$ . A decomposable regularizer is the  $\ell_1$  norm, since  $\|\mathbf{A}\|_1 = \|\mathbf{A}_E\|_1 + \|\mathbf{A}_{E^c}\|_1$ .
- *Low-rank PSD matrices.* Consider a class of low-rank and positive semi-definite matrices  $\mathcal{A} \subset \mathbb{S}_+^{p \times p}$  which have rank  $r \leq p$ . For any given matrix  $\mathbf{A} \in \mathcal{A}$ , let  $\text{col}(\mathbf{A})$  denote its column space. Let  $U \subset \mathbb{R}^n$  be a  $r$ -dimensional subspace and define the subspace  $\mathcal{M}(U)$  and the perturbation subspace  $\overline{\mathcal{M}}^\perp(U)$  as

$$\begin{aligned} \mathcal{M}(U) &:= \{\mathbf{A} \in \mathbb{R}^{n \times p} \mid \text{col}(\mathbf{A}) \subseteq U\}, \\ \overline{\mathcal{M}}^\perp(U) &:= \{\mathbf{A} \in \mathbb{R}^{n \times p} \mid \text{col}(\mathbf{A}) \subseteq U^\perp\}. \end{aligned}$$

Then the nuclear norm  $\mathcal{R}_L(\cdot) = \|\cdot\|_*$  is a decomposable regularization function with respect to the subspace pair  $(\mathcal{M}(U), \overline{\mathcal{M}}^\perp(U))$ .

For the true model parameter  $\Theta^*$ , we define its associated *structural error set* with respect to a subspace  $\mathcal{M}$  as (Negahban et al., 2012):

$$\begin{aligned} \mathbb{C}(\mathcal{M}, \overline{\mathcal{M}}^\perp; \Theta^*) &:= \\ \left\{ \Delta \in \mathbb{R}^{n \times p} \mid \mathcal{R}(\Delta_{\overline{\mathcal{M}}^\perp}) \leq 3\mathcal{R}(\Delta_{\overline{\mathcal{M}}}) + 4\mathcal{R}(\Theta_{\overline{\mathcal{M}}^\perp}^*) \right\}. \end{aligned} \quad (5)$$

By construction, if the norm of the projection of the true parameter  $\Theta^*$  into  $\overline{\mathcal{M}}^\perp$  is small, then elements  $\Delta$  in this structural error set also have limited projection onto the perturbation subspace  $\overline{\mathcal{M}}^\perp$ .

Now let  $\Theta^*$  be the true (marginal) precision matrix of the LVGGM, and let the sparse and low-rank components

be  $\mathbf{S}^*$  and  $\mathbf{L}^*$ , respectively. For the defined subspace pairs  $(\mathcal{M}(E), \overline{\mathcal{M}}(E)^\perp)$  and  $(\mathcal{M}(U), \overline{\mathcal{M}}(U)^\perp)$ , we use  $\mathbb{C}(E)$  and  $\mathbb{C}(U)$  as the shorthand notations for the corresponding structural error sets centered at  $\mathbf{S}^*$  and  $\mathbf{L}^*$ , *i.e.*,  $\mathbb{C}(\mathcal{M}(E), \overline{\mathcal{M}}(E)^\perp; \mathbf{S}^*)$  and  $\mathbb{C}(\mathcal{M}(U), \overline{\mathcal{M}}(U)^\perp; \mathbf{L}^*)$ , respectively. Later, we will consider the perturbation of  $\Theta^*$  along restricted directions in these two sets.

## 4.2. Assumptions on Fisher Information

We characterize the interaction between the elements in the two subspaces through their inner products using the Hessian of the loss function, also known as the *Fisher information* of the distribution. Denoting the Fisher information matrix of a Gaussian distribution as  $\mathcal{F}^*$  (evaluated at  $\Theta^*$ ), we find that  $\mathcal{F}^* = \Theta^{*-1} \otimes \Theta^{*-1}$ , where  $\otimes$  is the Kronecker product. We define the *Fisher inner product* between two matrices  $\Delta_A$  and  $\Delta_B$  as

$$\langle \Delta_A, \Delta_B \rangle_{\mathcal{F}^*} := \text{vec}(\Delta_A)^T \mathcal{F}^* \text{vec}(\Delta_B) \quad (6)$$

$$= \text{Tr}(\Theta^{*-1} \Delta_A \Theta^{*-1} \Delta_B), \quad (7)$$

where  $\text{vec}(\cdot)$  denotes the vectorization of a matrix.

Similar to prior work of (Kakade et al., 2010), we define the induced *Fisher norm* of a matrix  $\Delta$  as

$$\|\Delta\|_{\mathcal{F}^*}^2 := \text{vec}(\Delta)^T \mathcal{F}^* \text{vec}(\Delta) \quad (8)$$

$$= \text{Tr}(\Theta^{*-1} \Delta \Theta^{*-1} \Delta). \quad (9)$$

The first assumption we make is the following *Restricted Fisher Eigenvalue* (RFE) condition on the true precision model with respect to the sparse and low-rank structural error sets.

**Assumption 1 (Restricted Fisher Eigenvalue).** *There exists some constant  $\kappa_{\min}^* > 0$ , such that for all  $\Delta \in \mathbb{C}(E) \cup \mathbb{C}(U)$ , the following holds:*

$$\|\Delta\|_{\mathcal{F}^*}^2 \geq \kappa_{\min}^* \|\Delta\|_F^2. \quad (10)$$

This RFE condition generalizes the restricted eigenvalue (RE) condition for sparsity-promoting linear regression problems (Bickel et al., 2009). It assumes that the minimum eigenvalue of the Fisher information is bounded away from zero along the directions  $\mathbb{C}(E)$  and  $\mathbb{C}(U)$ . Due to the identity (8) and properties of the Kronecker product, a trivial lower bound for  $\kappa_{\min}^*$  is  $\lambda_{\min}^2(\Theta^*)$ , where  $\lambda_{\min}(\cdot)$  denotes the minimum eigenvalue. In the high-dimensional setting, the RFE parameter  $\kappa_{\min}^*$ , which is defined only with respect to the above restricted set of directions, can be substantially larger than  $\lambda_{\min}^2(\Theta^*)$ . As a result, the derived error bounds, which depend on  $\kappa_{\min}^*$ , are generally tighter than the bounds depending on  $\lambda_{\min}^2(\Theta^*)$  (cf. Theorem 1).

Due to the sparse plus low-rank superpositioned structure, we impose a type of incoherence between the two structural

error sets to ensure consistent estimation of the combined model. The incoherence condition will limit the interaction between elements from the two sets. For our problem, such interaction occurs through their inner products with the Fisher information, which motivates the following *Structural Fisher Incoherence* (SFI) assumption (which generalizes the *C-Linear* assumption proposed in Yang & Ravikumar (2013)).

Let  $\mathcal{P}_E := \mathcal{P}_{\overline{\mathcal{M}}(E)}$  denote the projection operator corresponding to the subspace  $\overline{\mathcal{M}}(E)$ . Similarly define  $\mathcal{P}_U := \mathcal{P}_{\overline{\mathcal{M}}(U)}$ ,  $\mathcal{P}_{E^\perp} := \mathcal{P}_{\overline{\mathcal{M}}(E)^\perp}$ , and  $\mathcal{P}_{U^\perp} := \mathcal{P}_{\overline{\mathcal{M}}(U)^\perp}$ . We assume the following condition on the Fisher information.

**Assumption 2 (Structural Fisher Incoherence).** *Given a constant  $M > 6$ , a set of regularization parameters  $(\lambda, \mu)$ , and the subspace pairs  $(\mathcal{M}(E), \overline{\mathcal{M}}(E)^\perp)$  and  $(\mathcal{M}(U), \overline{\mathcal{M}}(U)^\perp)$  as defined above, let  $\Lambda = 2 + 3 \max \left\{ \frac{\lambda\sqrt{s}}{\mu\sqrt{r}}, \frac{\mu\sqrt{r}}{\lambda\sqrt{s}} \right\}$ , where  $s = |E|$  and  $r = \text{rank}(U)$ . Then the Fisher information  $\mathcal{F}^*$  satisfies:*

$$\max \{ \bar{\sigma}(\mathcal{P}_E \mathcal{F}^* \mathcal{P}_U), \bar{\sigma}(\mathcal{P}_{E^\perp} \mathcal{F}^* \mathcal{P}_U) \},$$

$$\bar{\sigma}(\mathcal{P}_E \mathcal{F}^* \mathcal{P}_{U^\perp}), \bar{\sigma}(\mathcal{P}_{E^\perp} \mathcal{F}^* \mathcal{P}_{U^\perp}) \} \leq \frac{\kappa_{\min}^*}{c_1 \Lambda^2},$$

where  $\bar{\sigma}(\cdot)$  denotes the maximum singular value, and constant  $c_1$  is defined as  $c_1 = \frac{16M}{M-6}$ .

The constant  $M$  is related to a “burn-in” period after which the likelihood loss function has desirable properties in a small neighborhood of the true parameter. In particular, when  $M = 7$ , the constant  $c_1 = 112$  suffices for our theory to hold. See the main theorem and its proof for more discussion on this quantity.

It is interesting to compare our SFI assumption to other similar assumptions in the literature of GGM estimation. In Ravikumar et al. (2011), a form of irrepresentability condition is assumed, which limits the induced  $\ell_1$  norm of a matrix that is similar to the projected Fisher information onto the sparse matrix subspace pair. In Chandrasekaran et al. (2012), the notion of irrepresentability is extended to two subspace pairs (*i.e.*, sparse and low-rank), but detailed behaviors of the projected Fisher information are controlled (see the main assumption on page 1949 of Chandrasekaran et al. (2012)). For model selection consistency, a more general form of irrepresentability has been shown to be necessary for model selection consistency, see Lee et al. (2013) for a recent discussion. In contrast to the above line of work, the SFI assumption we make only controls the maximum singular values of the projected Fisher information. This can be explained as we are interested in bounding a weaker quantity, the Frobenius norm of the parameter estimation error, instead of establishing the stronger model selection consistency of Ravikumar et al. (2011) or the algebraic consistency as in Chandrasekaran et al. (2012).

### 4.3. Error Bounds for LVGGM Estimation

We have the following bound on the parameter error of the estimated precision matrix of LVGGM,  $\hat{\Theta} = \hat{\mathbf{S}} + \hat{\mathbf{L}}$ , obtained by solving the regularized ML problem (4).

**Theorem 1.** *Suppose Assumption 1 and 2 hold for the true marginal precision matrix  $\Theta^*$ , and the regularization parameters are chosen such that*

$$\lambda \geq 2\|\Sigma^* - \hat{\Sigma}\|_\infty \text{ and } \mu \geq 2\|\Sigma^* - \hat{\Sigma}\|_2. \quad (11)$$

Given a constant  $M > 6$ , if an optimal solution pair  $(\hat{\mathbf{S}}, \hat{\mathbf{L}})$  to the convex program (4) satisfies

$$\max\{\|\hat{\mathbf{S}} - \mathbf{S}^*\|_{\mathcal{F}^*}, \|\hat{\mathbf{L}} - \mathbf{L}^*\|_{\mathcal{F}^*}\} \leq \frac{1}{6M^2}, \quad (12)$$

then we have the following error bound for the estimated precision matrix  $\hat{\Theta} = \hat{\mathbf{S}} + \hat{\mathbf{L}}$ :

$$\|\hat{\Theta} - \Theta^*\|_F \leq \frac{6}{\kappa_{\mathcal{L}}} \max\{\lambda\sqrt{s}, \mu\sqrt{r}\} + \sqrt{\frac{8r_{\perp}^*}{\kappa_{\mathcal{L}}}}, \quad (13)$$

where  $s = |E|$ ,  $r = \text{rank}(U)$ , and

$$\kappa_{\mathcal{L}} := \frac{M-2}{2(M-1)}\kappa_{\min}^*, \quad (14)$$

$$r_{\perp}^* := \lambda \sum_{(j,k) \notin E} |\mathbf{S}_{jk}^*| + \mu \sum_{j=r+1}^p \sigma_j(\mathbf{L}^*). \quad (15)$$

*Proof sketch.* The proof is inspired by Yang & Ravikumar (2013), in which a parameter estimation error bound is proven for estimating a class of superposition-structured parameters, such as sparse plus low-rank, through M-estimation with decomposable regularizers. Critical to specializing this framework to our LVGGM estimation problem is to verify two conditions on the log-likelihood loss function (3): the restricted strong convexity (RSC) and structural incoherence (SI). The RSC condition (which originally proposed in Negahban et al. (2012)) specifies the loss function to be sufficiently curved (*i.e.* lower bounded by a quadratic function) along a restricted set of directions (defined by  $\mathbb{C}(E)$  and  $\mathbb{C}(U)$ ). On the other hand, the SI condition effectively limits certain interaction between elements from the above two structural error sets. In Yang & Ravikumar (2013), under certain *C-linear* assumptions, the RSC and SI conditions are verified for several problems with quadratic loss functions. For the LVGGM estimation problem, however, the technical difficulty lies in the non-quadratic log-likelihood loss (3), for which the previously established RSC and SI conditions do not hold.

To deal with this difficulty, we leverage the *almost strong convexity* properties (Kakade et al., 2010) to characterize the convergence behavior of the sum of higher-order terms in the Taylor series of the log-likelihood loss function. We

show that in the regime specified by condition (12), the loss function can be well-approximated by the sum of a quadratic function and a residual term. Under this condition, the RFE assumption (Assumption 1) guarantees the RSC condition (cf. Lemma 2), and the SFI assumption (Assumption 2) leads to SI condition to hold (cf. Lemma 4). Theorem 1 can then be proven by the general theorem in Yang & Ravikumar (2013). A detailed proof of Theorem 1 can be found in Appendix B.  $\square$

We make the following remarks:

- The error bound (13) is a family of upper bounds defined by different sets of subspace pairs  $(\mathcal{M}(E), \overline{\mathcal{M}}(E)^\perp)$  and  $(\mathcal{M}(U), \overline{\mathcal{M}}(U)^\perp)$ . The tightest bound can be achieved by appropriately choosing  $E$  and  $U$ . The first additive term in (13) captures effect of the estimation error, while the second term captures the approximation error. In many cases it is reasonable to assume the approximation error is zero, then the error bound reduces to the first additive term.
- We note that similar derivations also apply to  $\ell_1$ -regularized estimation of sparse GGM. For the sparse GGM, only Assumption 1 is required, and the derivations largely simplify. The final error bound also contains estimation and approximation errors, depending only on the sparse matrix subspace pair. However, when the true precision matrix  $\Theta^*$  cannot be well-approximated as a sparse matrix (such as the LVGGM case), the approximation error would be much worse, leading to an inefficient learning rate.
- We finally remark that the SFI assumption can be relaxed to an even milder incoherence condition,  $\|\mathbf{L}\|_\infty \leq \alpha$ , as considered in Agarwal et al. (2012). Following similar derivations as in the proof of Theorem 1, the corresponding error bound can be obtained. However, as a result of this incoherence assumption, the error bound would contain an additional incoherence term which does not vanish to zero even with infinite samples. This disadvantage is overcome under the structural incoherence condition.

The statement of Theorem 1 is deterministic in nature and applies to any optimum of the convex program. However, the condition on the regularization parameters (11) and the error bound depend on the sampled data (in particular the sample covariance matrix  $\hat{\Sigma}$ ), which is random. Therefore the key to specifying the regularization parameters, and hence obtaining error bounds independent of data, is to derive tight deviation bounds of the sample covariance matrix in terms of the  $\ell_\infty$  and  $\ell_2$  norms, such that condition (11) holds with high probability. These bounds can be obtained by using concentration inequalities for Gaussian distributions, which leads to the following corollary.

**Corollary 1.** *Let the same assumptions in Theorem 1 hold.*

Given constants  $C_1 > 1$  and  $C_2 \geq 1$ , assume that the number of samples  $n$  satisfies  $n \geq \max\{4C_1^2 \log p, C_2^2 p\}$ , and that the regularization parameters satisfy

$$\lambda = 160C_1\bar{\sigma}^* \sqrt{\frac{\log p}{n}} \text{ and } \mu = 16C_2\rho^* \sqrt{\frac{p}{n}}, \quad (16)$$

where  $\bar{\sigma}^* = \max_i \Sigma_{i,i}^*$  and  $\rho^* = \|\Sigma^*\|_2$ . Then with probability at least  $1 - 4p^{-2(C_1-1)} - 2\exp(-\frac{C_2^2 p}{2})$ , we have

$$\|\hat{\Theta} - \Theta^*\|_F \leq c_1 \sqrt{\frac{s \log p}{n}} + c_2 \sqrt{\frac{rp}{n}}, \quad (17)$$

where  $c_1 = \frac{960}{\kappa_{\mathcal{L}}}\bar{\sigma}^*$  and  $c_2 = \frac{96}{\kappa_{\mathcal{L}}}\rho^*$ .

*Remark:* The estimation error (17) consists of two terms corresponding to the sparse and low-rank components, respectively. Note its resemblance to the error bounds of robust PCA (e.g., Agarwal et al. (2012); Yang & Ravikumar (2013)) and the derived bound in Chandrasekaran et al. (2012). In particular, the first term in (17) was on the same order as the estimation error of a sparse GGM (Ravikumar et al., 2011). However, due to the presence of latent variables, both the sample requirement (i.e.,  $n \gtrsim p$ ) and the combined error bound are worse than those for learning the sparse conditional GGM.

Next we consider a scenario under which this additional disadvantage is largely removed. Assume that the true marginal covariance matrix  $\Sigma^*$  has an effective rank  $r_{\text{eff}} := r_{\text{eff}}(\Sigma^*)$  (recall  $r_{\text{eff}}(\Sigma^*) := \text{tr}(\Sigma^*)/\|\Sigma^*\|_2$ ) that is much smaller than  $p$ . Then, by using recent advances on the asymptotic behavior of the sample covariance matrix (Lounici, 2012), we can obtain a much tighter bound which only depends on  $p$  logarithmically, as stated in the following theorem.

**Theorem 2.** *Let the same assumptions in Theorem 1 hold. Given a constant  $C_1 > 1$ , assume that the number of observations  $n$  satisfies  $n \geq \max\{4C_1 \log p, C_3 r_{\text{eff}} \log^2(2p)\}$ , and the regularization parameters satisfy*

$$\lambda = 160C_1\bar{\sigma}^* \sqrt{\frac{\log p}{n}} \text{ and } \mu = C_4\rho^* \sqrt{\frac{r_{\text{eff}} \log p}{n}}, \quad (18)$$

where  $\bar{\sigma}^* = \max_i \Sigma_{i,i}^*$ ,  $\rho^* = \|\Sigma^*\|_2$ , and  $C_3, C_4 > 0$  are sufficiently large constants. Then with probability at least  $1 - 2p^{-2(C_1-1)} - (2p)^{-1}$ , we have

$$\|\hat{\Theta} - \Theta^*\|_F \leq \tilde{c}_1 \sqrt{\frac{s \log p}{n}} + \tilde{c}_2 \sqrt{\frac{r_{\text{eff}} \cdot r \log(2p)}{n}}, \quad (19)$$

where  $\tilde{c}_1 = \frac{960}{\kappa_{\mathcal{L}}}\bar{\sigma}^*$ ,  $\tilde{c}_2 = \frac{8C_4}{3\kappa_{\mathcal{L}}}\rho^*$ .

*Proof sketch.* Same as Corollary 1, we need to verify that the choices of regularization parameters (18) satisfy the

condition (11) with high probability. Since the choice of  $\lambda$  has been verified in Corollary 1, it only remains to verify the condition on  $\mu$ . To this end, we make use of the following sharp bound on the spectral norm deviation of the sample covariance matrix:

**Lemma 1 (Lounici (2012)).** *Let  $\hat{\Sigma}$  be a sample covariance matrix constructed from  $n$  i.i.d. samples from a  $p$ -dimensional Gaussian distribution  $\mathcal{N}(0, \Sigma^*)$ . Then with probability at least  $1 - (2p)^{-1}$ ,*

$$\|\hat{\Sigma} - \Sigma^*\|_2 \leq C\|\Sigma^*\|_2 \cdot \max \left\{ \sqrt{\frac{2r_{\text{eff}} \log(2p)}{n}}, \frac{2r_{\text{eff}} \log(2p)(3/8 + \log(2pm))}{n} \right\},$$

where  $C > 0$  is an absolute constant.

Then as commented in Lounici (2012) (Prop. 3), when the sample size  $n$  is sufficiently large such that  $n \geq C_3 r_{\text{eff}} \log^2 \max\{2p, n\}$ , where  $C_3 > 0$  is a large constant, the choice of regularization parameter  $\mu$  as in (18) suffices for the condition (11) to hold with high probability.  $\square$

Notice that when  $r_{\text{eff}} \ll p$ , the error bound (19) is significantly tighter than the bound (17). Also the sample requirement  $n \gtrsim r_{\text{eff}} \log(p)$  is much milder. This result implies the efficiency of LVGGM learning when the true covariance model has a low effective rank.

## 5. Experiments

We use a set of simulations on synthetic data to verify our reduced effective rank assumption on the covariance matrix of LVGGM, and the derived error bounds in Theorem 2.

### 5.1. Effective Rank of Covariance of LVGGM

To better understand the effective rank of the covariance matrix of LVGGM, it is convenient to consider a hierarchical generating process for the observed variables:  $x_O \sim \mathbf{A}x_L + z$ , where  $x_L \sim \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_{L,L})$  are the latent variables,  $\mathbf{A} := \mathbf{J}_{O,O}^{-1} \mathbf{J}_{O,L} \in \mathbb{R}^{p \times r}$ , and  $z \sim \mathcal{N}(\mathbf{0}, \mathbf{S}^{-1})$  captures the conditional effects. The marginal covariance matrix of the observed variables can be represented as

$$\Sigma = \underbrace{\mathbf{A} \mathbf{\Omega}_{L,L} \mathbf{A}^T}_{\mathbf{G}} + \mathbf{S}^{-1}, \quad (20)$$

where  $\mathbf{G}$  is a low-rank covariance matrix (global effects), and  $\mathbf{S}^{-1}$  is a non-sparse covariance matrix (conditionally local effects) whose inverse is sparse. While the low-rank global effects naturally result in a concentrated spectrum, the sparse-inverse local effects generally contribute to a diffuse spectrum. The effective rank, which is the sum of all




Figure 2. Effective ranks of covariance matrices of LVGGM with various global/local energy ratios.




Figure 3. Simulations for chain graphical models with latent variables. Plots of Frobenius norm error  $\|\hat{\Theta} - \Theta^*\|_F$  versus the rescaled sample size  $n/(s \log(p) + r \log(2p))$ .

eigenvalues divided by the magnitude of the largest one, depends on the relative energy ratio between  $\mathbf{G}$  and  $\mathbf{S}^{-1}$ .

Since an exact characterization of the effective rank in terms of  $\mathbf{A}$ ,  $\mathbf{\Omega}_{L,L}$ , and  $\mathbf{S}$  tends to be difficult, we use Monte Carlo simulations to investigate synthetic LVGGM that conform to our assumptions. We generate LVGGM with independent latent variables (*i.e.*, diagonal  $\mathbf{J}_{L,L}$ ), dense latent-observed submatrix  $\mathbf{J}_{L,O}$ , and a sparse conditional GGM  $\mathbf{J}_{O,O}$  for observed variable with a random sparsity pattern (sparsity level  $\approx 5\%$ ). We fix the number of latent variables to be 10, and vary the number of observed variables  $p = \{80, 120, 200, 500\}$ . By scaling the magnitudes of the elements in the latent variable submatrix, we sweep through the relative energy ratio between the global and local factors, *i.e.*,  $\text{Tr}(\mathbf{G})/\text{Tr}(\mathbf{S}^{-1})$  from 0.1 to 10. After 550 realizations for each value of  $p$ , we plot the empirical effective ranks of observed covariance matrices in Figure 2.

As seen in the figure, when the global factor dominates (*i.e.*, the ratio is large), the effective rank of the covariance matrix is very small, as expected. On the other hand, when the local effects become stronger (*e.g.*, when the number of observed variables  $p$  increases) the effective rank increases, but at a very mild rate. In particular, when  $p$  increases from

80 to 500, the maximum empirical effective rank in our simulation only increases from 4 to 26. For all of our simulated LVGGM, the empirical effective ranks are observed as at least an order of magnitude smaller than  $p$ . This mild growing rate of the effective rank (compared to  $p$ ) will lead to our improved error bound in Theorem 2 to hold.

## 5.2. Frobenius Norm Error of LVGGM Estimation

We simulate LVGGM data with number of observed variables  $p = \{160, 200, 320, 400\}$  and number of latent variables in the set  $r = \{0.1, 0.15, 0.2, 0.3\}p$ . The sparse conditional GGM is a chain graph whose associated precision matrix is tridiagonal with off-diagonal elements  $\mathbf{S}_{i,i-1} = \mathbf{S}_{i,i+1} = 0.4\mathbf{S}_{i,i}$  for  $i = \{2, \dots, p-1\}$ . For each configuration of  $p$  and  $r$ , we draw  $n$  samples from the LVGGM, where  $n$  ranges from 200 to 1000. Using these samples, the precision matrix  $\hat{\Theta}$  is learned by solving the regularized ML estimation problem (4). As shown in Section 5.1, the effective rank of the covariance matrix grows mildly. Then Theorem 2 predicts that the Frobenius error of the estimated precision matrix of LVGGM should scale as  $\|\hat{\Theta} - \Theta^*\|_F \asymp \sqrt{(s \log(p) + r \log(2p))/n}$ , when the regularization parameters are chosen such that  $\lambda \asymp \bar{\sigma}^* \sqrt{\frac{\log(p)}{n}}$  and  $\mu \asymp \rho^* \sqrt{\frac{r_{\text{eff}} \log(p)}{n}}$ . Guided by this theoretical result, we set the regularization parameters as  $\lambda = C_a \bar{\sigma}^* \sqrt{\frac{\log(p)}{n}}$  and  $\mu = C_b \rho^* \sqrt{\frac{r_{\text{eff}} \log(p)}{n}}$ , where constants  $C_a$  and  $C_b$  are cross-validated and then fixed for all test data sets with different configurations. We plot the Frobenius estimation errors against the rescaled sample size  $n/(s \log(p) + r \log(2p))$  in Figure 3. With a wide range of configurations, almost all the empirical error curves for models align and have the form of  $f(t) \propto t^{-1/2}$  when the sample size is rescaled, as predicted by Theorem 2. In practice when the true model is unknown, one could set the regularization parameters according to the sample versions of the quantities  $\bar{\sigma}^*$  and  $\rho^*$ , as discussed in Lounici (2012).

## 6. Conclusions

We consider a family of latent variable Gaussian graphical model whose precision matrix has a sparse plus low-rank structure. We derive parameter error bounds for regularized maximum likelihood estimation. Future work includes extending the framework to other distributions and the application to tasks such as prediction and detection.

## Acknowledgement

This work is in part supported by ARO grant W911NF-11-1-0391. The authors thank the anonymous reviewers for their valuable comments, along with Jason Lee and Yuekai Sun for helpful discussions.



## References

- Agarwal, Alekh, Negahban, Sahand, and Wainwright, Martin J. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics*, 40(2):1171–1197, 2012.
- Bickel, Peter J, Ritov, Yaacov, and Tsybakov, Alexandre B. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- Candès, Emmanuel J, Li, Xiaodong, Ma, Yi, and Wright, John. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.
- Chandrasekaran, Venkat, Parrilo, Pablo A, and Willsky, Alan S. Latent variable graphical model selection via convex optimization. *Annals of Statistics*, 40(4):1935–1967, 2012.
- Choi, Myung Jin, Chandrasekaran, Venkat, and Willsky, Alan S. Gaussian multiresolution models: Exploiting sparse Markov and covariance structure. *Signal Processing, IEEE Transactions on*, 58(3):1012–1024, 2010.
- Choi, Myung Jin, Tan, Vincent YF, Anandkumar, Animesh, and Willsky, Alan S. Learning latent tree graphical models. *Journal of Machine Learning Research*, 12:1729–1770, 2011.
- Friedman, Jerome, Hastie, Trevor, and Tibshirani, Robert. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- Hsieh, Cho-Jui, Sustik, Mátyás A, Dhillon, Inderjit, Ravikumar, Pradeep, and Poldrack, Russell. BIG & QUIC: Sparse inverse covariance estimation for a million variables. In *Advances in Neural Information Processing Systems*, pp. 3165–3173, 2013.
- Kakade, Sham, Shamir, Ohad, Sindharan, Karthik, and Tewari, Ambuj. Learning exponential families in high-dimensions: Strong convexity and sparsity. In *International Conference on Artificial Intelligence and Statistics*, pp. 381–388, 2010.
- Kalaitzis, Alfredo and Lawrence, Neil D. Residual component analysis: Generalising PCA for more flexible inference in linear-Gaussian models. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pp. 209–216, 2012.
- Lauritzen, S.L. *Graphical models*, volume 17. Oxford University Press, USA, 1996.
- Lee, Jason, Sun, Yuekai, and Taylor, Jonathan E. On model selection consistency of penalized M-estimators: a geometric theory. In *Advances in Neural Information Processing Systems*, pp. 342–350, 2013.
- Liu, Ying and Willsky, Alan. Learning Gaussian graphical models with observed or latent FVSs. In *Advances in Neural Information Processing Systems*, pp. 1833–1841, 2013.
- Lounici, Karim. High-dimensional covariance matrix estimation with missing observations. *arXiv preprint arXiv:1201.2577*, 2012.
- Ma, Shiqian, Xue, Lingzhou, and Zou, Hui. Alternating direction methods for latent variable Gaussian graphical model selection. *Neural computation*, 25(8):2172–2198, 2013.
- Negahban, Sahand N, Ravikumar, Pradeep, Wainwright, Martin J, and Yu, Bin. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- Ravikumar, Pradeep, Wainwright, Martin J, Raskutti, Garvesh, and Yu, Bin. High-dimensional covariance estimation by minimizing  $\ell_1$ -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- Rothman, Adam J, Bickel, Peter J, Levina, Elizaveta, and Zhu, Ji. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.
- Vershynin, Roman. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Xu, Huan, Caramanis, Constantine, and Sanghavi, Sujay. Robust PCA via outlier pursuit. *Information Theory, IEEE Transactions on*, 58(5):3047–3064, 2012.
- Yang, Eunho and Ravikumar, Pradeep. Dirty statistical models. In *Advances in Neural Information Processing Systems 26*, pp. 611–619, 2013.