Efficient Statistics for Sparse Graphical Models from Truncated Samples

Arnab Bhattacharyya NUS
Rathin Desai BU
Sai Ganesh Nagarajan SUTD
Ioannis Panageas UC Irvine

Abstract

In this paper, we study high-dimensional estimation from truncated samples. We focus on two fundamental and classical problems: (i) inference of sparse graphical Gaussian models and (ii) support recovery of sparse linear models.

(i) For Gaussian graphical models, suppose d-dimensional samples $\mathbf{x}$ are generated from a Gaussian $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and observed only if they belong to a subset $S \subseteq \mathbb{R}^d$. We show that $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ can be estimated with error ε in the Frobenius norm, using[1] $\tilde{O}\left(\frac{\mathrm{nz}(\boldsymbol{\Sigma}^{-1})}{\varepsilon^2}\right)$ samples from a truncated $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and having access to a membership oracle for S. The set S is assumed to have non-trivial measure under the unknown distribution but is otherwise arbitrary.

(ii) For sparse linear regression, suppose samples $(\mathbf{x}, y)$ are generated where $y = {\boldsymbol{\Omega}^*}^\top \mathbf{x} + \mathcal{N}(0,1)$ and $(\mathbf{x}, y)$ is seen only if y belongs to a truncation set $S \subseteq \mathbb{R}$. We consider the case that $\boldsymbol{\Omega}^*$ is sparse with a support set of size k. Our main result is to establish precise conditions on the problem dimension d, the support size k, the number of observations n, and properties of the samples and the truncation that are sufficient to recover the support of $\boldsymbol{\Omega}^*$. Specifically, we show that under some natural assumptions, only $O(k^2 \log d)$ samples are needed to estimate $\boldsymbol{\Omega}^*$ in the ℓ_∞-norm up to a bounded error. Similar results are also estabilished for estimating $\boldsymbol{\Omega}^*$ in the Euclidean norm up to arbitrary error.

For both problems, the estimator is obtained by minimizing the sum of the empirical negative log-likelihood function and an ℓ_1-regularization term.

[1] $\mathrm{nz}(A)$ denotes the number of non-zero entries of A.

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130.

1 Introduction

Sparse high-dimensional models are a mainstay of modern statistics and machine learning. In this work, we consider two different sparse linear models that have been the subject of intensive study.

– **Sparse Gaussian Graphical Models.** Graphical models are used to represent the probabilistic relationships between a collection of variables. These models are used in a huge number of different domains, such as statistical physics, computational biology, finance, and machine learning; the books [24, 28, 38, 22] give an indication of the breadth of this area. We focus on Gaussian graphical models in which the d variables $X_1, \ldots, X_d$ are distributed according to a d-dimensional Gaussian. Specifically, the distribution is described in terms of a density function $p(\mathbf{X})$ where $\mathbf{X} = (X_1, \ldots, X_d)$ and $p(\mathbf{X}) = (2\pi)^{-d/2} \cdot (\det \boldsymbol{\Sigma})^{-1/2}$

$$\exp\left(-\frac{1}{2}(\mathbf{X}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{X}-\boldsymbol{\mu})\right).$$

Here, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ correspond to the mean and variance of the distribution respectively.

It is convenient to reparametrize the density function in terms of the inverse covariance matrix or the *precision matrix*, $\boldsymbol{\Theta} = \boldsymbol{\Sigma}^{-1}$: Then,

$$p(\mathbf{X}) = (2\pi)^{-d/2} \cdot \exp\big(\boldsymbol{\mu}^T \boldsymbol{\Theta}\mathbf{X} - \frac{1}{2}\mathbf{X}^\top \boldsymbol{\Theta}\mathbf{X} - \frac{1}{2}\boldsymbol{\mu}^\top \boldsymbol{\Theta}\boldsymbol{\mu} + \frac{1}{2}\log\det(\boldsymbol{\Theta})\big)$$

Note that the exponent is a quadratic polynomial in which the coefficient of X_iX_j is $\Theta_{i,j}$. The symmetric matrix $\mathbf{\Theta}$ naturally defines an undirected graph G on d vertices in which $(i,j) \in E(G)$ iff $\Theta_{i,j} \neq 0$. The graph G also admits a very nice probabilistic interpretation: X_i and X_j are independent conditioned on all other variables if and only if $\Theta_{i,j} = 0$. Thus, for natural systems, it is quite reasonable to assume that the degree of each node in G is small, as this corresponds to assuming that each variable is "directly" dependent on a small number of variables. Note that even if $\mathbf{\Theta}$ is sparse, $\mathbf{\Sigma}$ could be dense; in fact, in many typical systems, any pair of variables is correlated even though they are not directly dependent.

The problem of learning sparse high-dimensional Gaussian graphical models (in terms of the precision matrix) has a rich history. Popular approaches include the graphical Lasso [18, 41, 3, 13, 33, 32], neighborhood-based methods [4, 27, 36], and CLIME [6] which have been proved to work under different sets of assumptions.

– **Sparse Linear Regression.** A fundamental problem in data science is to solve the following inverse problem. Given pairs $(\mathbf{X}_1, Y_1), \ldots, (\mathbf{X}_n, Y_n) \in \mathbb{R}^d \times \mathbb{R}$, find the "best" choice of $\mathbf{\Omega} \in \mathbb{R}^d$ so that $Y_i - \mathbf{\Omega}^\top \mathbf{X_i}$ is small in some norm. It is natural to want $\mathbf{\Omega}$ to be sparse so that the prediction can be made based on a small number of variables.

Consider the model $Y = \mathbf{\Omega^*}^\top \mathbf{X} + \varepsilon$ where ε is a Gaussian random variable and $\mathbf{\Omega}^*$ is a sparse vector. There has been a huge amount of work on this problem. In the high-dimensional setting, a very popular approach is using ℓ_1-regularization, leading to the Lasso algorithm [34]. By now, we have an almost complete understanding of the necessary and sufficient conditions needed for Lasso to recover $\mathbf{\Omega}^*$; see the discussion and references in Chapter 7 of [37].

In our work, we study the above two problems in the setting where the samples are subject to *truncation*. Truncation is also a classic challenge in statistics, occurring whenever the observation process is dependent on the drawn sample. Following early work by Galton [19], there has been a sustained history of research on truncated distributions, in particular, truncated Gaussians (see the citations in [9]) and truncated linear regression [35, 1, 20, 5]. We pick up the thread at [10] who developed a computationally and statistically efficient algorithm to learn a multivariate Gaussian given truncated samples and assuming that the truncation set is known. A follow-up work, [11], extended the analysis to the linear regression problem where only those samples $(\mathbf{X}_i, Y_i)$ are seen in which $Y_i \in S$, the truncation set.

To the best of our knowledge, ours is the first work that examines the problems of learning sparse Gaussian graphical models and linear models with truncated samples. We state our results next.

Statement of the results The first contribution of the paper is the following theorem on learning Gaussian graphical models up to small Frobenius norm error. The sampling process is that samples from an unknown d-variate Gaussian are only revealed if they belong to a subset $S \subseteq \mathbb{R}^d$; otherwise, the samples are completely hidden.

Theorem 1.1 (Frobenius norm)**.** *Suppose that we are given oracle access to a measurable set S, so that $\int_S \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}^*, \mathbf{\Sigma}^*) d\mathbf{x} = \alpha > 0$ for some d-variate $\mathcal{N}(\boldsymbol{\mu}^*, \mathbf{\Sigma}^*)$ and constant[2] $\alpha > 0$. There exists an estimator $\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{\Sigma}}$ that uses $\tilde{O}\left(\frac{nz(\mathbf{\Sigma}^{*-1})}{\varepsilon^2}\right)$ samples from the truncated distribution $\mathcal{N}(\boldsymbol{\mu}^*, \mathbf{\Sigma}^*, S)$ so that with probability at least 99%*

$$\left\| I - \mathbf{\Sigma}^{*-1/2} \tilde{\mathbf{\Sigma}} \mathbf{\Sigma}^{*-1/2} \right\|_F \leqslant \varepsilon, \left\| \mathbf{\Sigma}^{*-1/2} (\boldsymbol{\mu}^* - \tilde{\boldsymbol{\mu}}) \right\|_2 \leqslant \varepsilon.$$

Remark 1.2. *We would like to note that the number of non-zero entries of $\mathbf{\Sigma}^{*-1}$ should be at least d (since $\alpha > 0$, $\mathbf{\Sigma}^*$ is invertible and so is $\mathbf{\Sigma}^{*-1}$). Therefore $nz(\mathbf{\Sigma}^{*-1})$ absorbs any term that is $\Theta(d)$.*

The second contribution of the paper solves the variable selection problem for linear models, under certain assumptions. The sampling process is as follows: each covariate $\mathbf{x}^{(i)} \in \mathbb{R}^d$ is picked arbitrarily, and the value $y_i = \mathbf{\Omega^*}^\top \mathbf{X}_i + \varepsilon_i$ is revealed only if $y_i \in S$. Here, $\varepsilon_i \sim \mathcal{N}(0,1)$, the standard normal distribution.

Theorem 1.3 (Linear regression, informal)**.** *Suppose that we are given oracle access to a measurable set S. Let $\mathbf{X}$ denote a design matrix consisting of n samples $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(n)} \in \mathbb{R}^d$. Let K denote the unknown support of $\mathbf{\Omega}^*$, and let $k = |K|$. Assume that:*

*(i) (**Survival probability**) For each observed $\mathbf{x}^{(i)}$, the probability that $\mathbf{\Omega^*}^\top \mathbf{x}^{(i)} + w_i$ survives the truncation is not too small.*

*(ii) (**Minimum eigenvalue**) The vector $\mathbf{\Omega}^*$ is identifiable if its support K was known a priori.*

*(iii) (**Mutual incoherence**) Covariates not in the support set K form columns in $\mathbf{X}$ that are approximately orthogonal to the space spanned by the columns corresponding to K.*

[2] Think of α like 1%.

(iv) ***(Normalization)*** *Each entry of* $\mathbf{X}$ *is small in magnitude.*

Then, with only $n = O(k^2 \log d)$ *samples* $(\mathbf{x}^{(i)}, y_i)$ *from the truncated distribution, one can recover a vector* $\hat{\mathbf{\Omega}}$ *such that with high probability:*

(a) The support of $\hat{\mathbf{\Omega}}$ *is contained in* K*.*

(b) If for some $j \in K$*,* Ω^*_j *is larger than a threshold* τ *(which depends on the problem parameters but not* d*), then* $\hat{\Omega}_j \neq 0$*.*

Finally, under standard assumptions on the covariates, we can get ℓ_2 *error* $\left\|\hat{\mathbf{\Omega}} - \mathbf{\Omega}^*\right\|_2 \leqslant \varepsilon$ *using* $O(\frac{k^2 \log d}{\varepsilon^2})$ *samples (see Section 5).*

Our techniques We first discuss the ideas behind Theorem 1.1. In [10], it was shown that using $n = \tilde{O}\left(\frac{d^2}{\varepsilon^2}\right)$ samples from a d-variate truncated Gaussian distribution with truncation set S of measure some constant $\alpha > 0$, the mean $\boldsymbol{\mu}^*$ and the covariance $\mathbf{\Sigma}^*$ of the untruncated distribution can be estimated with ε error in ℓ_2 and Frobenius norm respectively. The crux of their proof involves proving that the infinite population negative log-likelihood is κ-strongly convex in a neighborhood U ($U \subseteq \mathcal{S}_{d\times d} \times \mathbb{R}^d$) of the true parameters where the radius of U and κ are functions of α[3]. Moreover, they run projected SGD with an efficient projection procedure in that neighborhood U. SGD requires a sample from the true truncated distribution in every iteration, so the sample complexity of this approach is at least as much as the number of iterations of SGD. Due to variance reasons, for SGD to converge, the number of samples needed is $\Omega\left(\frac{d^2}{\varepsilon^2}\right)$.

To improve up on their sample complexity, our estimator is the minimizer of a different function - denoted by L_n - which is the *finite* population negative log-likelihood plus a regularization term (see Equation (3.4)). The regularization term is the sum of the absolute values of the entries of the precision matrix (excluding the diagonal entries). This approach is the well-known Graphical Lasso.

One first easy observation is that the finite population negative log-likelihood and the infinite population negative log-likelihood have the same Hessian (thus same convexity properties, see Equation (3.6)). Moreover, since the extra regularization term does not change the convexity properties of the finite population negative log-likelihood, we get for free from [10] that the function L_n is κ-strongly convex in a neighborhood U of the true parameters (same κ and U as before). The crucial part now is that for the Lasso approach to work, we need that the empirical mean and the empirical covariance (from the truncated distribution) is close in ℓ_∞ and max-norm respectively (and not in ℓ_2 and Frobenius norm). The only requirement for the proof to go through is that the number of samples gives the statistical guarantee for Lasso to work (see Lemma 3.2).

For the support recovery problem in the sparse linear model with truncated samples, we again consider the Lasso objective, i.e., the sum of the finite population negative log-likelihood plus $\lambda\|\mathbf{\Omega}\|_1$. This objective function is globally convex. Suppose we already know the support K of $\mathbf{\Omega}^*$, the true k-sparse coefficient vector. In this case, we can solve the Lasso objective restricted to the variables in K and hope that it is strongly convex so that the minimum is unique. For the untruncated case, the minimum eigen value assumption (Assumption (ii) in Theorem 1.3) implies global strong convexity. In the truncated case, we can only guarantee strong convexity in a neighborhood around $\mathbf{\Omega}^*$. By tuning the regularization parameter λ, we can ensure that the minimum of the restricted Lasso objective will be in this neighborhood, and hence, is uniquely defined.

The main challenge in proving Theorem 1.3 is to extend the above ideas to when K is not known. To this end, we use the *primal-dual witness method* that has proven very useful for studying many Lasso-type algorithms [36, 30, 21, 7, 29, 31, 25, 39, 40]. We identify a strict dual feasibility condition that implies uniqueness of the Lasso solution and then demonstrate for a set of parameters that the condition holds. In contrast to the untruncated case, we are not able to drive the ℓ_∞-error to zero as n grows to infinity. Also, we require a stronger normalization on the entries of the design matrix. We leave as an interesting open problem the question of overcoming these deficiencies in our analysis.

Another independent work: We would like to note that we became familiar with some concurrent and independent work for the problem of sparse truncated linear regression that is very similar to ours [12] (their setting and approach is very close to ours). The main difference is that their work focuses on ℓ_2 recovery (using Stochastic Gradient Descent, they give algorithmic -polynomial time- guarantees for truncation sets with specific structure, with the extra assumption that the covariates are i.i.d Gaussians) and ours

[3] Think of the radius r as $O\left(\frac{\log(1/\alpha)}{\alpha^2}\right)$ and κ to be $O(\alpha^{cr^5})$ where c some constant. U is a subset of $\mathcal{S}_{d\times d} \times \mathbb{R}^d$ where $\mathcal{S}_{d\times d}$ denotes the symmetric matrices of size $d \times d$.

focuses on both support and ℓ_2 recovery (statistical guarantees). Moreover, in this work we also provide results on recovery of Gaussian graphical models which is not provided in [12].

Other related works Our work comes under the purview of robust statistics where the body of work relating to [15, 16, 14, 23, 8] provided guarantees for computationally efficient robust estimators in the presence of corruptions of an ε fraction of the data, when the samples are drawn from a multivariate Gaussian distribution. In addition, [17] provide statistical query lower bounds on estimation problems related to multivariate Gaussians such as learning mixtures of high dimensional Gaussians. These works generally talk about the seemingly inherent trade-off between increasing the sample complexity for computational tractability. As a result, an important assumption about the underlying problem or the statistical model is that of sparsity. Aside from the works related to estimation in sparse models in classical statistics such as sparse linear regression (LASSO) and sparse PCA [42] to mention a few, there is a line of work related to robust estimation in sparse models, such as robust sparse mean estimation when the covariance matrix is identity and then detection of rank 1 sparse shifts of high dimensional covariances of Gaussian distributions when the mean is zero, using the spiked covariance model as studied in [26, 2].

2 Preliminaries

2.1 Definitions and Notations

Notation We use bold faces to denote vectors and matrices. By $\mathbf{x}_{-j}$ we denote the vector $\mathbf{x}$ that involves all coordinates but j. We use $\mathrm{vec}(A)$ to denote the standard vectorization of matrix A. Moreover, we use $\|\mathrm{vec}(A)\|_{1,\mathrm{off}}$ to denote the ℓ_1 norm of $\mathrm{vec}(A)$ by excluding the diagonal entries of matrix A and $\mathrm{nz}(A)$ for the number of non-zero entries of matrix A. We denote by $\mathcal{S}_{d\times d}$ the set of symmetric matrices.

Norms For a $d \times d$ matrix A,

$$\|A\|_2 = \max_{\|x\|_2=1} \|Ax\|_2\,, \quad \|A\|_\infty = \max_{j\in[n]} \sum_{i=1}^{n} |A_{ij}|,$$
$$\|A\|_F = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n} A_{ij}^2}.$$

When A is a symmetric matrix we have that $\|A\|_2 \leqslant \|A\|_\infty \leqslant \|A\|_F \leqslant \sqrt{n}\,\|A\|_2 \leqslant \sqrt{n}\,\|A\|_\infty$. For a vector $\mathbf{x} \in \mathbb{R}^d$ we also have,

$$\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^{d} \mathbf{x}_i^2},\ \|\mathbf{x}\|_\infty = \max_{j\in[d]} |\mathbf{x}_j|,\ \|\mathbf{x}\|_1 = \sum_{i=1}^{d} |\mathbf{x}_i|.$$

It holds that ℓ_1 is the dual of ℓ_∞ and for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ one can have $\mathbf{x}^T\mathbf{y} \leqslant \|\mathbf{x}\|_1 \|\mathbf{y}\|_\infty$ (Holder's inequality).

Truncated Gaussian Distribution A truncated Gaussian distribution for a measurable set S with parameters $\boldsymbol{\mu}, \boldsymbol{\Sigma}$ is defined as follows

$$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, S; \mathbf{x}) \stackrel{\text{def}}{=} \begin{cases} \dfrac{\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}; \mathbf{x})}{\int_S \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}; \mathbf{x}) d\mathbf{x}}\ , & \mathbf{x} \in S \\ 0\ , & \mathbf{x} \notin S \end{cases} \tag{2.1}$$

where

$$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}; \mathbf{x}) \stackrel{\text{def}}{=} \frac{1}{\sqrt{2\pi \det(\boldsymbol{\Sigma})}} \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}-\boldsymbol{\mu})\right).$$

Definition 2.1 (Membership oracle)**.** *Let $S \subset \mathbb{R}^d$ be a measurable set. A membership oracle of S is a function that given an arbitrary $\mathbf{x} \in \mathbb{R}^d$, it returns yes if it belongs to the set, otherwise no (i.e., it implements the indicator function of S). We assume oracle access to the indicator of S.*

Precision matrix and sparsity Let $G = (V, E)$ be an undirected graph with $V = [d]$. A random vector $\mathbf{X} \in \mathbb{R}^d$ is said to be distributed according to (undirected) Gaussian Graphical model with graph G if $\mathbf{X}$ has a multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with

$$\left(\boldsymbol{\Sigma}^{-1}\right)_{ij} = 0 \ \ \forall\, (i,j) \notin E, \tag{2.2}$$

$\boldsymbol{\Sigma}^{-1}$ which we denote by $\boldsymbol{\Theta}$ is known as the precision matrix. In our results, the sample complexity depends on the number of non-zero entries of $\boldsymbol{\Sigma}^{-1}$, i.e., $\mathrm{nz}(\boldsymbol{\Sigma}^{-1})$.

3 Statistics for Frobenius norm

3.1 Graphical Lasso and finite population Likelihood

The infinite population negative log-likelihood for a truncated Gaussian $\mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$ with variables $(\boldsymbol{\Theta}, \mathbf{v})$ where $\boldsymbol{\Theta}$ captures $\boldsymbol{\Sigma}^{-1}$ and $\mathbf{v} = \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$ is given by (see [10] for calculations)

$$\bar{l}(\boldsymbol{\Theta}, \mathbf{v}) := \mathbb{E}_{\mathbf{x}\sim\mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*, S)} \left[\frac{1}{2}\mathbf{x}^T\boldsymbol{\Theta}\mathbf{x} - \mathbf{x}^T\mathbf{v}\right] - \log\left(\int_S \exp(-\frac{1}{2}\mathbf{z}^T\boldsymbol{\Theta}\mathbf{z} + \mathbf{z}^T\mathbf{v}) d\mathbf{z}\right). \tag{3.1}$$

Moreover, the gradient of the function above $\bar{l}(\mathbf{\Theta}, \mathbf{v})$ is given by

$$\nabla \bar{l}(\mathbf{\Theta}, \mathbf{v}) := -\mathbb{E}_{\mathbf{x}\sim\mathcal{N}(\boldsymbol{\mu}^*, \mathbf{\Sigma}^*, S)} \left[\begin{pmatrix} \mathrm{vec}(-\frac{1}{2}\mathbf{x}\mathbf{x}^T) \\ \mathbf{x} \end{pmatrix} \right] + \mathbb{E}_{\mathbf{z}\sim\mathcal{N}(\mathbf{\Theta}^{-1}\mathbf{v}, \mathbf{\Theta}^{-1}, S)} \left[\begin{pmatrix} \mathrm{vec}(-\frac{1}{2}\mathbf{z}\mathbf{z}^T) \\ \mathbf{z} \end{pmatrix} \right] \tag{3.2}$$

and its Hessian is

$$\nabla^2 \bar{l}(\mathbf{\Theta}, \mathbf{v}) := \mathrm{Cov}_{\mathbf{z}\sim\mathcal{N}(\mathbf{\Theta}^{-1}\mathbf{v}, \mathbf{\Theta}^{-1}, S)} \left[\begin{pmatrix} \mathrm{vec}(-\frac{1}{2}\mathbf{z}\mathbf{z}^T) \\ \mathbf{z} \end{pmatrix}, \begin{pmatrix} \mathrm{vec}(-\frac{1}{2}\mathbf{z}\mathbf{z}^T) \\ \mathbf{z} \end{pmatrix} \right]. \tag{3.3}$$

We define the following score objective with parameter $\lambda > 0$ to be chosen later

$$L_n(\mathbf{\Theta}, \mathbf{v}) := l_n(\mathbf{\Theta}, \mathbf{v}) + \lambda \left\| \mathrm{vec}(\mathbf{\Theta}) \right\|_{1,\mathrm{off}}, \tag{3.4}$$

where

$$l_n(\mathbf{\Theta}, \mathbf{v}) = \frac{1}{n}\sum_{i=1}^{n} \frac{1}{2}\mathbf{x}_i^T \mathbf{\Theta} \mathbf{x}_i - \frac{1}{n}\sum_{i=1}^{n} \mathbf{x}_i^T \mathbf{v} - \log\left(\int_S \exp(-\frac{1}{2}\mathbf{z}^T \mathbf{\Theta}\mathbf{z} + \mathbf{z}^T\mathbf{v}) d\mathbf{z} \right), \tag{3.5}$$

given i.i.d samples $\mathbf{x}_1, ..., \mathbf{x}_n$ from the true truncated distribution i.e., it is the finite population negative log-likelihood.

From Lemma A.4 (one of the main Lemmas of [10]), we know that $\bar{l}(\mathbf{v}, \mathbf{\Theta})$ is strongly convex in some neighborhood $U \subseteq \mathcal{S}_{d\times d} \times \mathbb{R}^d$ of the true parameters. We can conclude that $L_n(\mathbf{\Theta}, \mathbf{v})$ is also strongly convex in the same neighborhood because the term $\lambda \left\| \mathrm{vec}(\mathbf{\Theta}) \right\|_{1,\mathrm{off}}$ is a convex function and

$$\nabla^2 l_n(\mathbf{\Theta}, \mathbf{v}) = \nabla^2 \bar{l}(\mathbf{\Theta}, \mathbf{v}) \tag{3.6}$$

i.e., they have same strong-convexity properties The following lemma indicates that the minimizer $\tilde{\mathbf{\Theta}}, \tilde{\mathbf{v}}$ of function L_n does not put too much weight on the coordinates ij of $\tilde{\mathbf{\Theta}}$ for which $\mathbf{\Theta}^*_{ij} = 0$, where $(\mathbf{\Theta}^*, \mathbf{v}^*)$ denote the true parameters.

Lemma 3.1 (Lasso guarantee)**.** *Let $(\tilde{\mathbf{\Theta}}, \tilde{\mathbf{v}})$ be the minimum of L_n and $(\mathbf{\Theta}^*, \mathbf{v}^*)$ be the true parameters. Assume that $\lambda \geqslant 2 \left\| \nabla_{\mathbf{\Theta}} l_n(\mathbf{\Theta}^*, \mathbf{v}^*) \right\|_{\infty, off}$ and $\Delta = \tilde{\mathbf{\Theta}} - \mathbf{\Theta}^*, \delta = \tilde{\mathbf{v}} - \mathbf{v}^*$ then it holds*

$$\frac{1}{3}\left\| vec(\Delta_{\tilde{T}}) \right\|_1 - \frac{1}{3}\left\| \delta \right\|_1 \leqslant \left\| vec(\Delta_T) \right\|_1,$$

where T denotes the support of $\mathbf{\Theta}^*$ and $\tilde{T}$ denotes the complement. Moreover, we may assume that $\tilde{\mathbf{\Theta}}$ is symmetric.

From Lemma 3.1 and Cauchy-Schwarz inequality we conclude that

$$\begin{aligned} \left\| \mathrm{vec}(\Delta) \right\|_1 + \left\| \delta \right\|_1 &\leqslant \left\| \mathrm{vec}(\Delta_T) \right\|_1 + 3\left\| \mathrm{vec}(\Delta_T) \right\|_1 + 2\left\| \delta \right\|_1 \\ &\leqslant 4\sqrt{\mathrm{nz}(\mathbf{\Theta}^*) + d}(\left\| \Delta \right\|_F + \left\| \delta \right\|_2), \end{aligned} \tag{3.7}$$

We can now prove using Lemma 3.1 that for an appropriate choice of λ, the minimizer $(\tilde{\mathbf{\Theta}}, \tilde{\mathbf{v}})$ of L_n will be close to the true parameters $(\mathbf{\Theta}^*, \mathbf{v}^*)$.

Lemma 3.2 ($(\tilde{\mathbf{\Theta}}, \tilde{\mathbf{v}})$ are close to the true parameters)**.** *Let L_n be κ-strong convex in a neighborhood of the true parameters. By choosing λ to be $O\left(\frac{\kappa\cdot\varepsilon}{\sqrt{nz(\mathbf{\Theta}^*)}}\right)$ and moreover $\lambda \geqslant 2\left\| \nabla l_n(\mathbf{\Theta}^*, \mathbf{v}^*) \right\|_\infty$ then $\left\| \tilde{\mathbf{\Theta}} - \mathbf{\Theta}^* \right\|_F + \left\| \tilde{\mathbf{v}} - \mathbf{v}^* \right\|_2 \leqslant \varepsilon$.*

We finish this section with a concentration lemma about how close the empirical mean and covariance is from the truncated mean and covariance in terms of ℓ_∞ and max norm respectively.

Lemma 3.3 (Concentration of gradient)**.** *Assume that n is $\Omega\left(\frac{\log d \log(1/\delta)}{t^2}\right)$ It holds that*

$$\mathbb{P}\left[\left\| \nabla l_n(\mathbf{\Theta}^*, \mathbf{v}^*) \right\|_\infty \geqslant \frac{t}{2} \right] \leqslant \delta.$$

3.2 Proof of Theorem 1.1

We choose λ to be $\tilde{O}\left(\frac{\varepsilon}{12\sqrt{\mathrm{nz}(\mathbf{\Theta}^*)+d}}\right)$ and consider the estimator $(\tilde{\mathbf{\Theta}}, \tilde{\mathbf{v}}) := \arg\min_{\mathbf{\Theta}, \mathbf{v}} L_n(\mathbf{\Theta}, \mathbf{v})$. Notice that we need the oracle access to the truncation set S so that we can compute the term $\mathbb{E}_{\mathbf{z}\sim\mathcal{N}(\mathbf{\Theta}^{-1}\mathbf{v}, \mathbf{\Theta}^{-1}, S)} \left[\begin{pmatrix} \mathrm{vec}(-\frac{1}{2}\mathbf{z}\mathbf{z}^T) \\ \mathbf{z} \end{pmatrix} \right]$ which *does not involve* the true parameters $(\mathbf{\Theta}^*, \mathbf{v}^*)$.

We will prove that $(\tilde{\mathbf{\Theta}}, \tilde{\mathbf{v}})$ satisfies the statement of Theorem 1.1.

From Lemma 3.3 we conclude that if n is $\tilde{O}\left(\frac{(\mathrm{nz}(\mathbf{\Theta}^*)+d)\log(1/\delta)}{\varepsilon^2}\right)$ we get that $\lambda \geqslant 2\left\| \nabla l_n(\mathbf{\Theta}^*, \mathbf{v}^*) \right\|_\infty$ with probability $1-\delta$. Therefore the assumptions of Lemma 3.2 hold and is guaranteed that the minimizer $(\tilde{\mathbf{\Theta}}, \tilde{\mathbf{v}})$ of L_n satisfies

$$\left\| \tilde{\mathbf{\Theta}} - \mathbf{\Theta}^* \right\|_F \leqslant \varepsilon \text{ and } \left\| \tilde{\mathbf{v}} - \mathbf{v}^* \right\|_2 \leqslant \varepsilon. \tag{3.8}$$

4 Sparse Linear Regression

Recall the model described in the Introduction for the linear regression problem. The probability of obtain-

ing a sample $(\mathbf{x}, y) \in \mathbb{R}^d \times \mathbb{R}$ is:

$$\frac{\exp\left(-\frac{1}{2}(y - \mathbf{\Omega}^{*\top}\mathbf{x})^2\right)}{\int \exp\left(-\frac{1}{2}(z - \mathbf{\Omega}^{*\top}\mathbf{x})^2\right) S(z)dz}$$

The infinite population negative log-likelihood function with n samples is then:

$$\bar{\ell}(\mathbf{\Omega}) = \frac{1}{n}\sum_{i=1}^{n} \mathop{\mathbb{E}}_{y\sim\mathcal{N}(\mathbf{\Omega}^{*\top}\mathbf{x}^{(i)},1,S)} \left[\frac{1}{2}y^2 - y\cdot\mathbf{\Omega}^\top\mathbf{x}^{(i)} - \log\int \exp\left(-\frac{1}{2}z^2 + z\cdot\mathbf{\Omega}^\top\mathbf{x}^{(i)}\right)dz\right] \quad (4.1)$$

As in the last section, we instead work with the finite sample negative log-likelihood, which is based on n samples $(\mathbf{x}^{(1)}, y^{(i)}), \ldots, (\mathbf{x}^{(n)}, y^{(i)})$ with each $y^{(i)}$ being drawn from the distribution $\mathcal{N}(\mathbf{\Omega}^{*\top}\mathbf{x}^{(i)}, 1, S)$:

$$\ell_n(\mathbf{\Omega}) = \frac{1}{n}\sum_{i=1}^{n}\left(\frac{1}{2}{y^{(i)}}^2 - y^{(i)}\mathbf{\Omega}^\top\mathbf{x}^{(i)} + \log\int \exp\left(-\frac{1}{2}z^2 + z\mathbf{\Omega}^\top\mathbf{x}^{(i)}\right)S(z)dz\right).$$

Note that ℓ_n is a random variable. We add a regularizer to the sample negative log-likelihood to obtain the *truncated Lasso estimator*:

$$\hat{\mathbf{\Omega}} \in \arg\min_{\mathbf{\Omega}\in\mathbb{R}^d}\left\{\ell_n(\mathbf{\Omega}) + \lambda\|\mathbf{\Omega}\|_1\right\}. \quad (4.2)$$

In the following, let $\mathbf{X}$ denote the n-by-d design matrix whose i'th row corresponds to the i'th sample $\mathbf{x}^{(i)}$. Also, we let $\mathbf{x}_j \in \mathbb{R}^n$ denote the j'th column of $\mathbf{X}$.

4.1 Assumptions

We now formally state the assumptions under which our result holds. For vectors $\mathbf{\Omega}$ and $\mathbf{x}$, let $\alpha(\mathbf{\Omega}, \mathbf{x}) \stackrel{\text{def}}{=} \mathbb{E}_{y\sim\mathcal{N}(\mathbf{\Omega}^\top\mathbf{x},1)}[S(y)]$. Also, in the following, let $K \subseteq [d]$ denote the support of $\mathbf{\Omega}^*$, and let $k = |K|$.

Our first assumption states that for every observed $\mathbf{x}^{(i)}$, there is a significant probability that the corresponding response variable $y^{(i)}$ is not truncated.

Assumption 4.1 (Survival Probability)**.** *There exists a constant $\alpha > 0$ such that for every $i \in [n]$, $\alpha(\mathbf{\Omega}^*, \mathbf{x}^{(i)}) \geqslant \alpha$.*

Our second assumption is quite mild. It ensures that the model is identifiable when the support set S is known in advance.

Assumption 4.2 (Lower Eigenvalue)**.** *There exists a constant $\sigma_{\min} > 0$ such that*

$$\frac{1}{n}\mathbf{X}_K^\top\mathbf{X}_K \succeq \sigma_{\min}\cdot\mathbf{I}.$$

Our third assumption ensures that the covariates corresponding to the support set are sufficiently prominent. More precisely, the mutual incoherence assumption below requires that if $j \notin K$, then $\mathbf{x}_j$ is approximately orthogonal to the span of the submatrix $\mathbf{X}_K$ corresponding to the covariates in K.

Assumption 4.3 (Mutual incoherence)**.** *There exists a constant $\beta \in (0, 1)$ such that:*

$$\max_{j\notin K}\|\mathbf{x}_j^\top\mathbf{X}_K(\mathbf{X}_K^\top\mathbf{X}_K)^{-1}\|_1 \leqslant \beta.$$

Mutual incoherence is known to hold, for example, with high probability when $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(n)}$ are drawn i.i.d. from $N(0, \mathbf{I}_{d\times d})$ as long as $n \geqslant \Omega(k\log d)$.

Our last assumption puts a bound on each entry of $\mathbf{X}$.

Assumption 4.4 (Normalization)**.** *There exists a parameter C such that $\max_{i\in[n]}\|\mathbf{x}^{(i)}\|_\infty = \max_{j\in[d]}\|\mathbf{x}_j\|_\infty \leqslant C$.*

4.2 Support Recovery

We formally state the main theorem of this section.

Theorem 4.5. *Consider a k-sparse linear regression model for which Assumptions 4.1, 4.2, 4.3, and 4.4 are all satisfied. Moreover, suppose that $\frac{C^2k}{\alpha^5\sigma_{\min}(1-\beta)}$ is a sufficiently small constant. Then, if*

$$n \geqslant \Omega\left(\frac{C^4k^2\log d}{(1-\beta)^2\sigma_{\min}^2\alpha^9}\right) \quad \text{and} \quad \lambda = \Theta\left(\frac{\alpha^4\sigma_{\min}}{Ck}\right),$$

any solution $\hat{\mathbf{\Omega}}$ to the objective (4.2) satisfies the following properties with high probability.

(a) ***Uniqueness****: There is a unique solution $\hat{\mathbf{\Omega}}$.*

(b) ***No false inclusion****: $\mathrm{supp}(\hat{\mathbf{\Omega}}) \subseteq \mathrm{supp}(\mathbf{\Omega}^*)$.*

(c) ***ℓ_∞-bounds****: The error $\hat{\mathbf{\Omega}} - \mathbf{\Omega}^*$ satisfies*

$$\|\hat{\mathbf{\Omega}} - \mathbf{\Omega}^*\|_\infty \leqslant O\left(\sqrt{\frac{\log(1/\alpha)}{\sigma_{\min}}} + \frac{\alpha^4}{C\sqrt{k}}\right).$$

In other words, if the non-zero entries of $\mathbf{\Omega}^*$ are greater than a particular threshold τ (which is independent of d), then the support of $\hat{\mathbf{\Omega}}$ exactly matches with the support of $\mathbf{\Omega}^*$.

In the untruncated setting, it is known (see Chapter 7 of [37]) that λ can be made to scale as $\sim \frac{1}{\sqrt{n}}$, and the ℓ_∞ error is the sum of two terms, one proportional to λ and the other to $\frac{1}{\sqrt{n}}$. Hence, by making n large, the ℓ_∞ error can be made arbitrarily small. In contrast, in our analysis, we cannot make λ arbitrarily small; so,

above, we fix it in terms of the other problem parameters.

The other notable aspect of Theorem 4.5 is the hypothesis that $\frac{C^2 k}{\alpha^5 \sigma_{\min}(1-\beta)}$ is small, which is also absent from the untruncated setting. The hypothesis can be satisfied if C is mildly decreasing in d (e.g., $\sim 1/\log(d)$), and d is very large.

5 Bounds on ℓ_2-error for truncated linear regression

In this section, we give bounds on the ℓ_2-error for the truncated sparse linear regression problem. We use the same notation as Section 4, with:

$$\hat{\mathbf{\Omega}} \in \arg\min_{\mathbf{\Omega}\in\mathbb{R}^d} \{l_n(\mathbf{\Omega}) + \lambda\|\mathbf{\Omega}\|_1\} \tag{5.1}$$

where l_n is the negative log-likelihood function:

$$l_n(\mathbf{\Omega}) = \frac{1}{n}\sum_{i=1}^{n}\left(\frac{1}{2}{y^{(i)}}^2 - y^{(i)}\mathbf{\Omega}^\top \mathbf{x}^{(i)} + \log\int \exp\left(-\frac{1}{2}z^2 + z\mathbf{\Omega}^\top \mathbf{x}^{(i)}\right) S(z)dz\right),$$

which is convex everywhere. Recall that each $y^{(i)}$ is distributed as $\mathcal{N}(\mathbf{x}^{(i)}\mathbf{\Omega}^*, 1, S)$; we denote K as the support of $\mathbf{\Omega}^*$ and let $k = |K|$. The matrix $\mathbf{X}$ is defined to have n rows $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(n)}$.

For the theorem in this section, we need the survival probability assumption (Assumption 4.1) and the normalization assumption (Assumption 4.4), but we replace the minimum eigenvalue and mutual incoherence assumptions by the following.

Assumption 5.1 (Restricted Eigenvalue)**.** *There exists $\beta > 0$ such that the matrix $\mathbf{X}$ satisfies:*

$$\frac{1}{n}\|\mathbf{X}\mathbf{\Delta}\|_2^2 \geqslant \beta\|\mathbf{\Delta}\|_2^2 \qquad \text{for all } \mathbf{\Delta} \text{ such that } \|\mathbf{\Delta}_{\overline{K}}\|_1 \leqslant 3\|\mathbf{\Delta}_K\|_1.$$

The restricted eigenvalue assumption is a very common one in the study of Lasso-type algorithms, and it is known that many families of random design matrices satisfy it (see Chapter 7 of [37]). We can now state the main result of this section:

Theorem 5.2. *Suppose that Assumptions 4.1, 4.4 and 5.1 hold. Then, for any $\varepsilon > 0$, if $\lambda = O\left(\frac{\beta\alpha^4\varepsilon}{Ck}\right)$ and $n = \tilde{O}\left(\frac{C^4k^2\log d}{\beta^2\alpha^8\varepsilon^2}\right)$, then $\|\hat{\mathbf{\Omega}} - \mathbf{\Omega}^*\|_2 \leqslant \varepsilon$.*

6 Experimental Evaluation and Conclusion

We studied the problem of parameter estimation for sparse Gaussian Graphical models and the problem of sparse linear regression, given samples that are subject to truncation. We provided sample efficient estimators for both aforementioned problems under suitable assumptions.

We conducted a few experiments to empirically investigate the problem of inferring Gaussian graphical models. The algorithm we used was a projected stochastic gradient descent algorithm. In each iteration of this algorithm, the current estimates $\mathbf{v}$ and $\mathbf{\Theta}$ are updated by adding a subgradient of the graphical Lasso objective (3.4), scaled by a regularization parameter (that is set in accordance with Lemma 3.2). The updated $\mathbf{\Theta}$ is projected so as to ensure that it is symmetric with minimum eigenvalue at least 10^{-5}. The distribution generating the original samples is a 10-dimensional Gaussian distribution with each co-ordinate truncated on a support $(-2, 2)$. The mean of the distribution $\mathbf{v}^*$ is set to be $(0, 0, \cdots, 0)$. Moreover, we set the precision matrix $\mathbf{\Theta}^*$ to be the identity matrix plus 0.2's entries in the the upper and lower diagonal, thus making the number of nonzero entries in the precision matrix to be 30 (out of 100).

The first experiment studies how the Frobenius norm error ε (between the true parameters and the estimates) varies with the number of samples (fixing the number of iterations of SGD to be 10^6). The second experiment focuses on how the error varies with number of iterations of projected SGD for a fixed number of samples from the true distribution. In Figure 1a we see that the number of samples scales, as expected, like $1/\varepsilon^2$ w.r.t the error ε. Figure 1b shows that projected SGD performs rather poorly computationally; this is also expected as the function we optimize L_n is locally strong convex and the initialization is not necessarily close enough to the true parameters.

We also performed a couple of experiments to understand how well we can recover the support of the model. One experiment studies how the sparsity of the estimate varies as a function of the number of iterations of the proposed algorithms for a fixed number of samples from the true distribution. Another focuses on understanding how the sparsity of the estimate varies with the number of initial samples (fixing the number of SGD iterations to 10^6). To quantify the sparsity in our solution, we "binarized" the estimated precision matrix by thresholding all elements lying in the interval $[-0.1, 0.1]$ to zero and the rest to one.

Now to compare the "closeness" in sparsity between

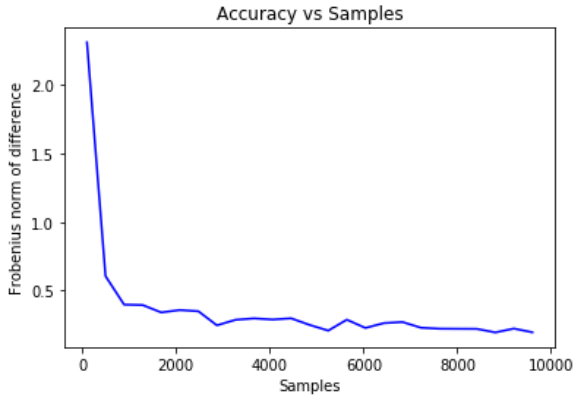

(a) Error vs Number of Samples

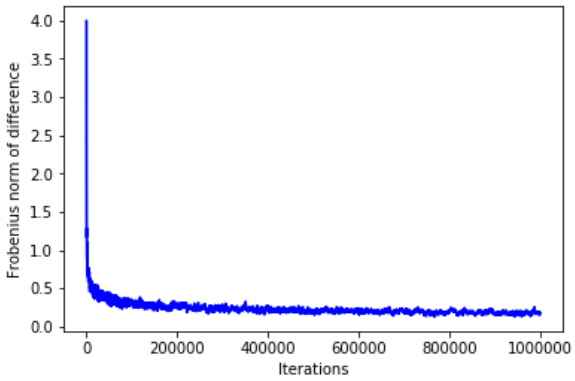

(b) Error vs Number of Iterations of SGD

the two binary matrices (the original and the estimated), we evaluate element wise Hamming distance between them. This is defined as follows:

$$h(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{k} \mathbf{1}\{x_i \neq y_i\}, \tag{6.1}$$

where $\mathbf{x}$ and $\mathbf{y}$ are two k-dimensional binary vectors.

Evaluation. The experiments suggest that our proposed algorithms recovers the true sparsity of the precision matrix. We observe from Figure 2a that the sparsity of the estimate goes reaches the required number over $2 * 10^5$ number of iterations. Although the estimate is sparse, we observe from our previous experiments (Figure 1b), that the desired accuracy is achieved with ($\approx 10^6$) number of iterations. In addition, our experiment for $n = 5000$ (see Figure 2b) suggests that even if we obtain a sparse estimate with less than 50000 initial samples, it takes a larger number of iterations to converge. This is evidenced in Figure 2b, where, the graph is noisy before ($\approx 10^6$) iterations and convergence happens beyond it.

In our second experiment we try to understand how the hamming distance varies with the initial number of samples. From Figure 3, we observe that $n = 2500$ are sufficient to achieve the required sparsity with 10^6 iterations.

Thus, together with the experiments on the convergence with respect to the Frobenius norm of the precision matrix, we can see that our proposed algorithm is able to recover the sparse precision matrix.

Acknowledgements

Arnab Bhattacharyya would like to acknowledge the National Research Foundation, Singapore Fellowship for AI (WBS R-252-000-B13-281) and an Amazon Research Award (WBS R-252-000-A61-720). Rathin Desai contributed to the work primarly while he was a research attachment at NUS and was supported by the Amazon Research Award (WBS R-252-000-A61-720). Sai Ganesh Nagarajan would like to acknowledge SUTD President's Graduate Fellowship (SUTD-PGF). Ioannis Panageas would like to acknowledge a UC Irvine start-up grant.

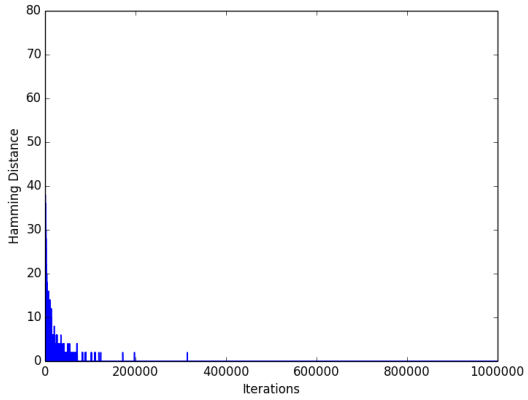

(a) Hamming distance vs Iterations($n = 50000$)

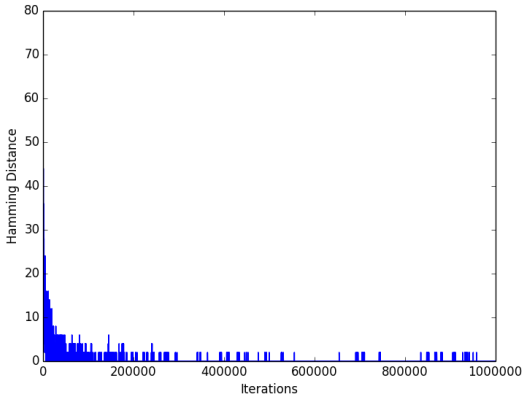

(b) Hamming distance vs Iterations($n = 5000$)

Figure 2: Hamming distance vs Iterations in low and high sample regimes.

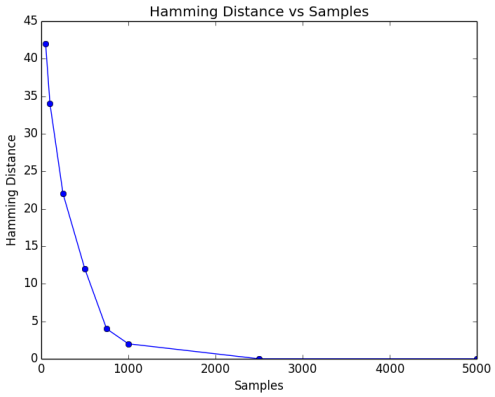

Figure 3: Hamming distance vs Number of Samples when keeping the number of iterations to be 10^6

References

[1] Takeshi Amemiya. Regression analysis when the dependent variable is truncated normal. *Econometrica*, pages 997–1016, 1973.

[2] Sivaraman Balakrishnan, Simon S. Du, Jerry Li, and Aarti Singh. Computationally efficient robust sparse estimation in high dimensions. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 169–212, Amsterdam, Netherlands, 07–10 Jul 2017. PMLR.

[3] Onureena Banerjee, Laurent El Ghaoui, and Alexandre d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine learning research*, 9(Mar):485–516, 2008.

[4] Julian Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):192–225, 1974.

[5] Richard Breen. *Regression Models: Censored, Sample Selected, or Truncated Data*. SAGE Publications, Inc, 1 edition, 1996.

[6] Tony Cai, Weidong Liu, and Xi Luo. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106(494):594–607, 2011.

[7] Venkat Chandrasekaran, Pablo A Parrilo, and Alan S Willsky. Latent variable graphical model selection via convex optimization. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1610–1613. IEEE, 2010.

[8] Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data, 2016.

[9] A Clifford Cohen. *Truncated and censored samples: theory and applications*. CRC press, 2016.

[10] Constantinos Daskalakis, Themis Gouleakis, Christos Tzamos, and Manolis Zampetakis. Efficient statistics, in high dimensions, from truncated samples. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 639–649, 2018.

[11] Constantinos Daskalakis, Themis Gouleakis, Christos Tzamos, and Manolis Zampetakis. Computationally and statistically efficient truncated regression. In *Conference on Learning Theory,*

COLT 2019, 25-28 June 2019, Phoenix, AZ, USA, pages 955–960, 2019.

[12] Constantinos Daskalakis, Dhruv Rohatgi, and Manolis Zampetakis. Truncated linear regression in high dimensions. *CoRR*, abs/2007.14539, 2020.

[13] Alexandre d'Aspremont, Onureena Banerjee, and Laurent El Ghaoui. First-order methods for sparse covariance selection. *SIAM Journal on Matrix Analysis and Applications*, 30(1):56–66, 2008.

[14] Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability, 2016.

[15] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Being robust (in high dimensions) can be practical, 2017.

[16] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robustly learning a gaussian: Getting optimal error, efficiently, 2017.

[17] Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 73–84. IEEE, 2017.

[18] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.

[19] Francis Galton. An examination into the registered speeds of american trotting horses, with remarks on their value as hereditary data. *Proceedings of the Royal Society of London*, 62(379-387):310–315, 1898.

[20] Jerry A Hauseman and David A Wise. Social experimentation, truncated distributions, and efficient estimation. *Econometrica*, pages 919–938, 1977.

[21] Ali Jalali, Sujay Sanghavi, Chao Ruan, and Pradeep K Ravikumar. A dirty model for multi-task learning. In *Advances in neural information processing systems*, pages 964–972, 2010.

[22] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[23] Kevin A. Lai, Anup B. Rao, and Santosh Vempala. Agnostic estimation of mean and covariance, 2016.

[24] Steffen L Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.

[25] Jason D Lee, Yuekai Sun, and Jonathan E Taylor. On model selection consistency of m-estimators with geometrically decomposable penalties. *arXiv preprint arXiv:1305.7477*, 241, 2013.

[26] Jerry Li. Robust sparse estimation tasks in high dimensions, 2017.

[27] Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, 34(3):1436–1462, 2006.

[28] Marc Mezard, Marc Mezard, and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.

[29] Sahand N Negahban and Martin J Wainwright. Simultaneous support recovery in high dimensions: Benefits and perils of block ℓ_1/ℓ_∞-regularization. *IEEE Transactions on Information Theory*, 57(6):3841–3863, 2011.

[30] Guillaume Obozinski, Martin J Wainwright, and Michael I Jordan. Union support recovery in high-dimensional multivariate regression. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 21–26. IEEE, 2008.

[31] Pradeep Ravikumar, Martin J Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing l1-penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.

[32] Pradeep Ravikumar, Martin J Wainwright, Garvesh Raskutti, Bin Yu, et al. High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.

[33] Adam J Rothman, Peter J Bickel, Elizaveta Levina, Ji Zhu, et al. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.

[34] Robert Tibshirani. Regression shrinkage and selection via the lasso. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, 58:267–288, 1994.

[35] James Tobin. Estimation of relationships for limited dependent variables. *Econometrica*, pages 24–36, 1958.

[36] M. J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, May 2009.

[37] Martin J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019.

[38] Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.

[39] Weiguang Wang, Yingbin Liang, and Eric P Xing. Collective support recovery for multi-design multi-response linear regression. *IEEE Transactions on Information Theory*, 61(1):513–534, 2014.

[40] Min Xu, Minhua Chen, and John Lafferty. Faithful variable screening for high-dimensional convex regression. *The Annals of Statistics*, 44(6):2624–2660, 2016.

[41] Ming Yuan and Yi Lin. Model selection and estimation in the gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.

[42] Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.