

ACTIVE COVARIANCE ESTIMATION BY RANDOM SUB-SAMPLING OF VARIABLES

Eduardo Pavez and Antonio Ortega

University of Southern California

ABSTRACT

We study covariance matrix estimation for the case of partially observed random vectors, where different samples contain different subsets of vector coordinates. Each observation is the product of the variable of interest with a 0 – 1 Bernoulli random variable. We analyze an unbiased covariance estimator under this model, and derive an error bound that reveals relations between the sub-sampling probabilities and the entries of the covariance matrix. We apply our analysis in an active learning framework, where the expected number of observed variables is small compared to the dimension of the vector of interest, and propose a design of optimal sub-sampling probabilities and an active covariance matrix estimation algorithm.

Index Terms— active covariance estimation, random sampling, missing data, graphical models, covariance matrix

1. INTRODUCTION

We study estimation of the covariance matrix of a random vector \mathbf{x} from observations of the form $\mathbf{y} = [\delta_1 x_1, \delta_2 x_2, \dots, \delta_n x_n]^\top$, where each component x_i is multiplied by a 0 – 1 Bernoulli random variable δ_i . We assume the probabilities $\mathbb{P}(\delta_i = 1) = p_i$ are known, independent of each other, and of \mathbf{x} . Partial observations consistent with this model may arise when there are some physical limitations to the observation process (and thus there is missing data) or when the cost of observation needs to be reduced by selecting what to sample (i.e., active learning).

Applications where probabilistic models need to be constructed from observations with missing data include transportation networks [1] and sensor networks [2]. For example, different sensors in the network might have different capabilities, e.g. different reliability or number of samples per hour that can be acquired, and these can be captured by associating different sub-sampling probabilities p_i to each sensor.

For covariance estimation, or more generally graphical model estimation, active learning approaches have been considered where variables are observed in a sequential and adaptive manner to identify the graphical model with minimal number of observations. Active learning approaches for covariance estimation can be useful in distributed computation environments like sensor networks, where there are acquisition, processing or communication constraints, and there is a need for optimal resource allocation.

One critical difference between missing data and active covariance estimation scenarios is the control over the observation model. In the missing data case, the sub-sampling probabilities are a problem feature that is given or has to be estimated from data, while in active learning problems, the sampling probabilities are a design parameter. The results from this paper can be useful for analysis in the

missing data case, as well as for designing active learning algorithms as we will discuss in more detail in Section 5.

In this work we analyze an unbiased covariance matrix estimator under sub-Gaussian assumptions on \mathbf{x} . Our main result is an error bound on the Frobenius norm that reveals the relation between number of observations, sub-sampling probabilities and entries of the true covariance matrix. We apply this error bound to the design of sub-sampling probabilities in an active covariance estimation scenario. An interesting conclusion from this work is that when the covariance matrix is approximately low rank, an active covariance estimation approach can perform almost as well as an estimator with complete observations. The paper is organized as follows, in Section 2 we review related work. Sections 3 and 4 introduce the problem and the main results respectively. In Section 5 we show an application of our bounds for optimal design of sampling probabilities for a batch based active covariance estimation algorithm. Proofs are presented in Section 6 and conclusion in Sections 7.

2. RELATED WORK

Lounici [3] studies covariance matrix estimation from missing data when all variables can be observed with the same probability. The focus of [3] is on estimation of approximately low rank covariance matrices with a matrix version of the lasso algorithm. In [4] the empirical covariance matrix under missing data is shown to be indefinite. Most recently Cai and Zhang [5] study the missing completely at random model (MCAR), which assumes arbitrary and unknown missing data probabilities. They show that a modified sample covariance matrix estimator, similar to ours but using the empirical estimation of probabilities p_i , achieves optimal minimax rates in spectral norm for bandable and sparse covariance matrices. Compared to [3, 5, 4], we allow all probabilities to be different and known, and we only study the unbiased sample covariance estimator. Moreover, all of the aforementioned papers consider estimation errors in spectral norm, while we consider errors in Frobenius norm, allowing us to derive error bounds with precise dependences between the entries of the true covariance matrix and the sampling probabilities p_i . These bounds can then be applied to design the sampling distribution in an adaptive manner.

For distributed computing applications, the work of [6] studies covariance estimation from random subspace projections using dense matrices, which are generalized in [7] to sparse projection matrices. The same approach from [7] is used in [8] for memory and complexity reduced PCA. In all these works [6, 8, 7], each measurement is a (possibly sparse) linear combination of a few variables, and even though the designs are random, they have a fixed distribution for all observations.

The work of [9] considers the case when all p_i are unknown, and uses an unbiased covariance matrix estimator as an input to the graphical lasso algorithm [10] for inverse covariance matrix estimation. Other interesting active learning approaches for graphi-

This work was supported in part by NSF under grant CCF-1410009. Contact author e-mail: pavezcar@usc.edu, ortega@sipi.usc.edu.

cal model selection are [11], [12] and [13]. Vats et. al. [12] uses a method that combines sampling marginals with conditional independence testing to learn the graphical model structure. [11], [13] consider a more general family of algorithms based on sampling high degree vertices, and prove upper and lower performance bounds.

Random sub-sampling and reconstruction of signals has been studied within graph signal processing [14, 15] and statistics [16]. The Bernoulli observation model we use has been studied by [14, 15] as a sampling strategy for graph signals, where sampling probability designs are proposed for reconstruction of deterministic band-limited signals. Also, Romero et. al. [17] derives covariance matrix estimators assuming the target covariance matrix is a linear combination of known covariance matrices, which leads to algorithms and theoretical analysis that are fundamentally different from ours.

3. PROBLEM FORMULATION

3.1. Notation

We denote scalars using regular font, while we use bold for vectors and matrices, e.g., $\mathbf{a} = (a_i)$, and $\mathbf{A} = (a_{ij})$. The Hadamard product between matrices is defined as $(\mathbf{A} \odot \mathbf{B})_{ij} = a_{ij}b_{ij}$. We use $\|\cdot\|_q$ for entry-wise matrix norms, with $q = 2$ corresponding to the Frobenius norm. $\|\cdot\|$ denotes ℓ_2 norm or spectral norm when applied to vectors or matrices respectively.

3.2. Unbiased estimation

Consider a random vector \mathbf{x} taking values in \mathbb{R}^n . We observe

$$\mathbf{y} = \boldsymbol{\delta} \odot \mathbf{x}, \quad (1)$$

where $\boldsymbol{\delta} = (\delta_i)$ is a vector of Bernoulli 0 – 1 random variables. The probability of observing the i -th variable is given by $\mathbb{P}(\delta_i = 1) = p_i$. The vector of probabilities is denoted by $\mathbf{p} = [p_1, \dots, p_n]^\top$, and $\mathbf{P} = \text{diag}(\mathbf{p})$ is a diagonal matrix. The average number of samples corresponds to $\mathbb{E}(\sum_{i=1}^n \delta_i) = \sum_{i=1}^n p_i = m$. Given $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(T)}$, i.i.d. realizations of \mathbf{y} , the i -th variable will be sampled in average $p_i T$ times, if all $p_i = 1$, then $\mathbf{y} = \mathbf{x}$, and we have perfect observation of \mathbf{x} . We are interested in studying covariance estimation for \mathbf{x} when $m < n$ and $0 < p_i$ for all i . Let $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ be the mean and covariance of \mathbf{x} , then

$$\mathbb{E}(\mathbf{y}) = \mathbf{P}\boldsymbol{\mu}, \text{ and } \text{Cov}(\mathbf{y}) = \boldsymbol{\Sigma} \odot \boldsymbol{\Xi} + (\mathbf{P} - \mathbf{P}^2) \text{diag}(\boldsymbol{\mu}\boldsymbol{\mu}^\top),$$

where $\boldsymbol{\Xi} = (\xi_{ij})$ is defined as $\xi_{ii} = p_i$ and $\xi_{ij} = p_i p_j$ when $i \neq j$.

For the rest of the paper we will assume $\boldsymbol{\mu} = \mathbf{0}$. Given a set of i.i.d. samples $\{\mathbf{y}^{(k)}\}_{k=1}^T$ of the random vector \mathbf{y} , define

$$\hat{\boldsymbol{\Sigma}} = \frac{1}{T} \sum_{k=1}^T \mathbf{y}^{(k)} \mathbf{y}^{(k)\top} \odot \boldsymbol{\Xi}^\dagger, \quad (2)$$

where $\boldsymbol{\Xi}^\dagger$ is the Hadamard (entry-wise) inverse of $\boldsymbol{\Xi}$. A simple calculation shows that $\hat{\boldsymbol{\Sigma}}$ is an unbiased estimator for $\boldsymbol{\Sigma}$. Indeed, $\mathbb{E}(\hat{\boldsymbol{\Sigma}}) = \frac{1}{T} \sum_{k=1}^T \mathbb{E}(\mathbf{y}^{(k)} \mathbf{y}^{(k)\top}) \odot \boldsymbol{\Xi}^\dagger = \boldsymbol{\Sigma} \odot \boldsymbol{\Xi} \odot \boldsymbol{\Xi}^\dagger = \boldsymbol{\Sigma}$. Because $\boldsymbol{\Xi}^\dagger \not\geq 0$, the matrix $\hat{\boldsymbol{\Sigma}}$ might not be positive semi-definite (conditions for $\hat{\boldsymbol{\Sigma}}$ to be positive semi-definite are given in [4]).

4. ESTIMATION ERROR

In this section we present an error analysis of the covariance matrix estimator from (2) when \mathbf{x} has sub-Gaussian entries. Sub-Gaussian

random variables include Gaussian, Bernoulli and Bounded random variables. For more information see [18] and references therein.

Definition 1 ([18]). If $\mathbb{E}[\exp(z^2/K^2)] \leq 2$ holds for some $K > 0$, we say z is sub-Gaussian. If $\mathbb{E}[\exp(|z|/K)] \leq 2$ holds for some $K > 0$, we say z is sub-exponential.

Definition 2 ([18]). The sub-Gaussian and sub-exponential norms are defined as

$$\|z\|_{\psi_\alpha} = \inf\{u > 0 : \mathbb{E}[\exp(|z|^\alpha/u^\alpha)] \leq 2\}.$$

for $\alpha = 2$ and $\alpha = 1$ respectively.

Sub-Gaussian and sub-exponential random variables, and their norms, are related as follows.

Proposition 1 ([18]). If z and w are sub-Gaussian, then z^2 and zw are sub-exponential with norms satisfying $\|z^2\|_{\psi_1} = \|z\|_{\psi_2}^2$, and $\|zw\|_{\psi_1} \leq \|z\|_{\psi_2} \|w\|_{\psi_2}$.

We have that the following characterization of the product of sub-Gaussian and Bernoulli random variables.

Lemma 1. Let $y_1 = \delta_1 x_1$ and $y_2 = \delta_2 x_2$ be a product of Bernoulli δ_1, δ_2 and sub-Gaussian x_1, x_2 random variables with Bernoulli probabilities p_1 and p_2 respectively, the only dependent variables are x_1 and x_2 , then

- 1) y_i is sub-Gaussian and $\|y_i\|_{\psi_2} \leq \|x_i\|_{\psi_2}$.
- 2) y_1^2, y_2^2 , and $y_1 y_2$ are sub-exponential with norms satisfying $\|y_i y_j\|_{\psi_1} \leq \|x_i x_j\|_{\psi_1}$ for $i, j = 1, 2$.

The proof of Lemma 1 can be easily obtained from the definition of sub-Gaussian and sub-exponential norms. We omit it for space considerations. We also define the matrix $\mathbf{H} = (h_{ij})$, with entries given by

$$h_{ij} = \begin{cases} \frac{\|x_i x_j\|_{\psi_1}}{p_i p_j} & i \neq j \\ \frac{\|x_i^2\|_{\psi_1}}{p_i} & i = j. \end{cases}$$

Now we state our main result, whose proof appears in Section 6.

Theorem 1. Let \mathbf{x} be zero mean random vector in \mathbb{R}^n with sub-Gaussian entries and norm $\|x_i\|_{\psi_2}$. Let $\mathbf{y} = \boldsymbol{\delta} \odot \mathbf{x}$, where each δ_i is a Bernoulli random variable with parameter $0 < p_i \leq 1$, independent of each other and of \mathbf{x} . Given i.i.d. realizations $\{\mathbf{y}^{(k)}\}_{k=1}^T$, the estimator from (2) satisfies

$$\|\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_q \leq \|\mathbf{H}\|_q \left\{ \sqrt{\gamma \frac{2 \log(n) + \log(\eta)}{T}} \vee \gamma \frac{2 \log(n) + \log(\eta)}{T} \right\}$$

with probability at least $1 - \frac{2}{\eta}$, where γ is an universal constant.

Moreover if $\|x_i\|_{\psi_2} = \sigma \sqrt{\sum_{i=1}^n p_i}$ and $q \geq 2$, then

$$\|\mathbf{H}\|_q \leq \frac{2\sigma^2}{\hat{p}^2} \mathbf{r}(\boldsymbol{\Sigma}) \|\boldsymbol{\Sigma}\|, \quad (3)$$

where $\hat{p} = \min p_i$, and $\mathbf{r}(\boldsymbol{\Sigma}) = \text{tr}(\boldsymbol{\Sigma})/\|\boldsymbol{\Sigma}\|$ is the effective rank.

Theorem 1 shows that the estimation error $\|\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_q \rightarrow 0$ in probability as the number of samples increases. More importantly, our result reveals that the sampling probabilities p_i are closely related to the sub-exponential norms of the variables $x_i x_j$ through the matrix $\|\mathbf{H}\|_q$. The bound from (3) suggests that distributions with smaller effective rank [3, 18] can tolerate a more aggressive sub-sampling factor (smaller m). This is not surprising since the effective rank $\mathbf{r}(\boldsymbol{\Sigma})$ is upper bounded by the actual rank, and can be significantly smaller for distributions whose energy concentrates in few principal components. We also note that the ratio $\mathbf{r}(\boldsymbol{\Sigma})/\hat{p}^2$ appears in the bounds of [3] as well.

Algorithm 1 Covariance estimation with adaptive sampling

Require: initial distribution $\mathbf{p}^{(0)}$, covariance $\widehat{\Sigma}^{(0)} = \mathbf{0}$, budget $\mathbf{1}^\top \mathbf{p}^{(0)} = m$, and batch size B .

- 1: **for** $t = 0$ to $N - 1$ **do**
- 2: Sample B i.i.d. realizations of $\mathbf{y} = \delta(\mathbf{p}^{(t)}) \odot \mathbf{x}$
- 3: Estimate $\widehat{\Sigma}$
- 4: Update $\widehat{\Sigma}^{(t+1)} \leftarrow \frac{1}{t+1} \widehat{\Sigma} + \frac{t}{t+1} \widehat{\Sigma}^{(t)}$
- 5: Update $\mathbf{p}^{(t+1)}$ by solving (5) with $\widehat{\Sigma}^{(t+1)}$ and budget m .
- 6: **end for**

5. ACTIVE COVARIANCE ESTIMATION

In this section we consider a scenario where we cannot observe (on average) more than m variables at a time, but we have the freedom to choose the sub-sampling probability distribution.

5.1. Sub-sampling distribution for true covariance matrix

Based on the error bound from Theorem 1, we propose designing the sub-sampling distribution by approximately minimizing $\|\mathbf{H}\|_q$. Theorem 1 suggests that p_i should be larger whenever the sub-Gaussian norm of x_i is large, but also the product $p_i p_j$ should be large when the sub-exponential norm of $x_i x_j$ is large. We assume that $\|x_i\|_{\psi_2} = \sigma \sqrt{\Sigma_{ii}}$. The bounds for h_{ii} and h_{ij} from (6) and (7) respectively suggest the approximation $p_i^2 \sim \Sigma_{ii}$. Given a sampling budget m , we estimate the sub-sampling probability vector \mathbf{p} by solving the following scaled projection problem

$$\min_{\mathbf{p}, \rho} \frac{1}{2} \|\mathbf{p} - \rho \text{diag}(\Sigma)^{\frac{1}{2}}\|_2^2, \text{ s.t. } \mathbf{1}^\top \mathbf{p} = m, \mathbf{0} \leq \mathbf{p} \leq \mathbf{1}. \quad (4)$$

5.2. Sub-sampling distribution for empirical covariance matrix

Since the true covariance matrix is unknown, (4) does not lead to a practical estimator. Instead, we consider a batch based algorithm that for a given budget m , and a starting sub-sampling distribution, it iteratively refines the sub-sampling probability distribution as a function of previous observations. We show the pseudo code for such procedure in Algorithm 1, which can be summarized in the following steps: observation with variable sub-sampling, covariance estimation, and sub-sampling distribution update. At the t -th iteration, B i.i.d. realizations are observed according to (1) with sub-sampling probabilities $\mathbf{p}^{(t)}$. The covariance estimator is a convex combination of the estimator at the previous iteration, and the estimator for the current batch. Finally, the new covariance matrix estimator is used to update the sub-sampling probabilities as

$$\begin{aligned} \mathbf{p}^{(t)} &= \arg \min_{\mathbf{p}, \rho} \frac{1}{2} \|\mathbf{p} - \rho \text{diag}(\widehat{\Sigma}^{(t)})^{\frac{1}{2}}\|_2^2, \\ \text{s.t. } \mathbf{1}^\top \mathbf{p} &= m, \mathbf{0} \leq \mathbf{p} \leq \mathbf{1}. \end{aligned} \quad (5)$$

Proposition 2. *Algorithm 1 produces an unbiased estimator for Σ .*

Proof. We will proceed by induction. For the first iteration we have used uniform sampling, thus we have that $\mathbb{E}(\widehat{\Sigma}^{(1)}) = \Sigma$. Now assume $\mathbb{E}(\widehat{\Sigma}^{(t)}) = \Sigma$, then at the $(t+1)$ -th iteration, the covariance of the new data satisfies $\mathbb{E}(\widehat{\Sigma}) = \Sigma$, which implies $\mathbb{E}(\widehat{\Sigma}^{(t+1)}) = \frac{1}{t+1} \mathbb{E}(\widehat{\Sigma}) + \frac{t}{t+1} \mathbb{E}(\widehat{\Sigma}^{(t)}) = \Sigma$. \square

5.3. Numerical evaluation

In this section we evaluate our proposed method using the MNIST [19] dataset, which consists of 28×28 images of scanned digits from 0 to 9. In our experiments we consider $N = 5851$ images of the digit 8. The vectorized, and mean removed images are denoted by $\{\mathbf{z}_i\}_{i=1}^N$ with covariance matrix \mathbf{C} . We consider estimation of the the covariance matrix of $\mathbf{x}_i = \mathbf{z}_i + \sqrt{\theta} \|\mathbf{C}\| \mathbf{e}$, where \mathbf{e} is zero mean Gaussian noise with unit variance, therefore $\Sigma = \mathbf{C} + \theta \|\mathbf{C}\| \mathbf{I}$.

To draw i.i.d. realizations of \mathbf{x} , we sample images \mathbf{z}_i without replacement and add Gaussian noise with variance $\theta \|\mathbf{C}\|$. The effective rank of Σ is controlled by the parameter θ and satisfies

$$r(\Sigma) = \frac{r(\mathbf{C}) + n\theta}{1 + \theta}.$$

We first compare uniform sub-sampling with non uniform sub-sampling for estimation of Σ with $\theta = 1/n$. We designed the non uniform sub-sampling distribution using (4) with the true covariance matrix. We report relative errors in Frobenius norm as a function of T/n in Figure 1a. Each point in the plot is an average over 50 independent trials. We observe that when $m = 0.75n$ the non uniform sampling distribution matches closely the performance of the estimator with full data. It is clear that when $m = 0.50n$ and $m = 0.25n$ performance decreases (for uniform and non uniform sampling) as m decreases, and non uniform sampling always outperforms uniform sampling. In Figure 1b we evaluate our active method from Algorithm 1 with batch size $B = 300$. We consider the same scenario as in Figure 1a for $m = 0.5n$ and $m = 0.25n$. The proposed active covariance estimation method quickly learns the optimal sub-sampling distribution, is always better than uniform sampling, and matches the performance of the non uniform sub-sampling method obtained from the true covariance matrix. Finally we show in Figure 1c the same experiment shown in Figure 1b, but now with covariance matrix with parameter $\theta = 10/n$, which changes the effective rank from $r(\Sigma) = 9.08$ with $\theta = 1/n$ to $r(\Sigma) = 17.86$. We observe that the problem becomes more difficult since the estimation errors are larger. There is no major difference between uniform, and non uniform sampling, thus the advantages of the proposed active covariance estimation method are limited. This might be due to various effects including, relative magnitudes of diagonal entries of Σ , sampling budget m , effective rank, and probability update algorithm. Moreover, since the effective rank did not change much (compared with n), this experiment suggests the effective rank does not quantify effectively the problem difficulty. Also, a more precise method to update the probabilities \mathbf{p} might help improving the performance of the active covariance estimation algorithm.

6. PROOF OF THEOREM 1

We first need the following concentration bounds

Lemma 2. *Under the same assumptions of Theorem 1, for any $\nu > 0$ we have*

$$\begin{aligned} \mathbb{P}(|\widehat{\Sigma}_{ij} - \Sigma_{ij}| > \nu) &\leq 2 \exp \left\{ -c_1 T \min \left(\nu^2 / h_{ij}^2, \nu / h_{ij} \right) \right\}, \\ \mathbb{P}(|\widehat{\Sigma}_{ii} - \Sigma_{ii}| > \nu) &\leq 2 \exp \left\{ -c_2 T \min \left(\nu^2 / h_{ii}^2, \nu / h_{ii} \right) \right\}, \end{aligned}$$

for off-diagonal and diagonal entries respectively.

Proof. The error events for off-diagonal entries satisfy

$$|\widehat{\Sigma}_{ij} - \Sigma_{ij}| > \nu \Leftrightarrow \left| \sum_{k=1}^T (y_i^{(k)} y_j^{(k)} - p_i p_j \Sigma_{ij}) \right| > \nu T p_i p_j.$$




Fig. 1: Performance of different covariance estimators in relative Frobenius norm. (U) corresponds to uniform sub-sampling $p_i = m/n$, (P) denotes non-uniform sub-sampling found using (4), (A) denotes active covariance estimation using Algorithm 1

We apply Bernstein's inequality [18] for sums of independent zero mean sub-exponential random variables, which combined with the bound from Lemma 1 leads to the desired bound. The proof for diagonal terms follows almost the same procedure. \square

We also need the following geometric result which we state without proof.

Lemma 3. Let $\mathcal{A} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_1 > \epsilon\}$, and $\mathcal{B}_i = \{\mathbf{x} \in \mathbb{R}^n : |x_i| > \alpha_i \epsilon\}$, then for all $\epsilon > 0$, and $\alpha_i \in (0, 1]$ that satisfy $\sum_{i=1}^n \alpha_i = 1$, we have that $\mathcal{A} \subset \bigcup_{i=1}^n \mathcal{B}_i$.

The proof of Theorem 1 starts by bounding the probability of the event $\|\hat{\Sigma} - \Sigma\|_q > \epsilon \|\Sigma\|_q$. Pick a set of $\alpha_{ij} \in (0, 1]$ such that $\sum_{i,j} \alpha_{ij} = 1$, and apply Lemma 3 and the union bound to get

$$\begin{aligned} \mathbb{P}(\|\hat{\Sigma} - \Sigma\|_q > \epsilon \|\Sigma\|_q) &= \mathbb{P}\left(\sum_{i,j} |\hat{\Sigma}_{ij} - \Sigma_{ij}|^q > \epsilon^q \|\Sigma\|_q^q\right) \\ &\leq \sum_{i,j} \mathbb{P}(|\hat{\Sigma}_{ij} - \Sigma_{ij}|^q > \alpha_{ij} \epsilon^q \|\Sigma\|_q^q) \\ &= \sum_{i,j} \mathbb{P}(|\hat{\Sigma}_{ij} - \Sigma_{ij}| > \alpha_{ij}^{1/q} \epsilon \|\Sigma\|_q) \\ &\leq \sum_{i=1}^n 2 \exp \left\{ -c_2 T \min \left(\frac{\alpha_{ii}^{2/q} \epsilon^2 \|\Sigma\|_q^2}{h_{ii}^2}, \frac{\alpha_{ii}^{1/q} \epsilon \|\Sigma\|_q}{h_{ii}} \right) \right\} + \\ &\quad \sum_{i \neq j=1}^n 2 \exp \left\{ -c_1 T \min \left(\frac{\alpha_{ij}^{2/q} \epsilon^2 \|\Sigma\|_q^2}{h_{ij}^2}, \frac{\alpha_{ij}^{1/q} \epsilon \|\Sigma\|_q}{h_{ij}} \right) \right\}. \end{aligned}$$

The last inequality follows from Lemma 2 with constants c_1, c_2 appropriately chosen so the inequalities hold for all pairs i, j . We can further simplify by choosing $\alpha_{ij} = h_{ij}^q / \|\mathbf{H}\|_q^q$

$$\begin{aligned} \mathbb{P}(\|\hat{\Sigma} - \Sigma\|_q > \epsilon \|\Sigma\|_q) &\leq 2n \exp \left\{ -c_2 T \min \left(\epsilon^2 \frac{\|\Sigma\|_q^2}{\|\mathbf{H}\|_q^2}, \epsilon \frac{\|\Sigma\|_q}{\|\mathbf{H}\|_q} \right) \right\} \\ &\quad + 2(n^2 - n) \exp \left\{ -c_1 T \min \left(\epsilon^2 \frac{\|\Sigma\|_q^2}{\|\mathbf{H}\|_q^2}, \epsilon \frac{\|\Sigma\|_q}{\|\mathbf{H}\|_q} \right) \right\} \\ &\leq 2n^2 \exp \left\{ -\frac{T}{\gamma} \min \left(\epsilon^2 \frac{\|\Sigma\|_q^2}{\|\mathbf{H}\|_q^2}, \epsilon \frac{\|\Sigma\|_q}{\|\mathbf{H}\|_q} \right) \right\}, \end{aligned}$$

where $1/\gamma = \min(c_1, c_2)$. The proof can be finished by equating to $2/\eta$, solving for ϵ , and doing some min / max manipulations. To derive the bound from (3) we bound the entries of \mathbf{H} obtaining

$$h_{ii} = \frac{\sigma^2 \Sigma_{ii}}{p_i} \leq \frac{\sigma^2 \Sigma_{ii}}{\hat{p}}, \quad (6)$$

$$h_{ij} \leq \frac{\sigma^2 \sqrt{\Sigma_{ii} \Sigma_{jj}}}{p_i p_j} \leq \frac{\sigma^2 \sqrt{\Sigma_{ii} \Sigma_{jj}}}{\hat{p}^2}. \quad (7)$$

Then, applying (6) and (7) followed by triangle inequality of the ℓ_q norm we have

$$\begin{aligned} \|\mathbf{H}\|_q &\leq \frac{\sigma^2}{\hat{p}} \left[\left(1 - \frac{1}{\hat{p}^q}\right) \sum_{i=1}^n \Sigma_{ii}^q + \frac{1}{\hat{p}^q} \left(\sum_{i=1}^n \Sigma_{ii}^{q/2} \right)^2 \right]^{\frac{1}{q}} \\ &\leq \frac{\sigma^2}{\hat{p}} \left[\left(\frac{1}{\hat{p}^q} - 1 \right)^{\frac{1}{q}} \|\text{diag}(\Sigma)\|_q + \frac{1}{\hat{p}} \|\text{diag}(\Sigma)\|_{\frac{q}{2}} \right] \\ &\leq \frac{\sigma^2 \text{tr}(\Sigma)}{\hat{p}^2} \left[(1 - \hat{p}^q)^{\frac{1}{q}} + 1 \right] \leq \frac{2\sigma^2}{\hat{p}^2} r(\Sigma) \|\Sigma\|. \end{aligned}$$

The last step uses the fact that $\|\mathbf{a}\|_q \leq \|\mathbf{a}\|_1$ for all $q \geq 1$, and the definition of effective rank.

7. CONCLUSION

We studied covariance matrix estimation when the variables are sub-sampled by a product with Bernoulli 0 – 1 variables. Variations of this model have been traditionally considered in the analysis of missing data. We study an unbiased estimator for the covariance matrix and derive a novel estimation error bound in entry-wise ℓ_q norm. Our bound illustrates the subtle relations between covariance matrix parameters and sub-sampling distribution. Using this bound, we propose an active covariance matrix estimation algorithm that also produces an unbiased estimator. We show with numerical experiments that the proposed active covariance estimation algorithm outperforms uniform sub-sampling, and closely matches non-uniform sub-sampling with complete knowledge of the true covariance matrix.

8. REFERENCES

- [1] Muhammad Tayyab Asif, Nikola Mitrovic, Justin Dauwels, and Patrick Jaillet, “Matrix and tensor based methods for missing data estimation in large traffic networks,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 7, pp. 1816–1825, 2016.
- [2] Amol Deshpande, Carlos Guestrin, Samuel R Madden, Joseph M Hellerstein, and Wei Hong, “Model-driven data acquisition in sensor networks,” in *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*. VLDB Endowment, 2004, pp. 588–599.
- [3] Karim Lounici, “High-dimensional covariance matrix estimation with missing observations,” *Bernoulli*, vol. 20, no. 3, pp. 1029–1058, 2014.
- [4] Kamil Jurczak and Angelika Rohde, “Spectral analysis of high-dimensional sample covariance matrices with missing observations,” *Bernoulli*, vol. 23, no. 4A, pp. 2466–2532, 2017.
- [5] T Tony Cai and Anru Zhang, “Minimax rate-optimal estimation of high-dimensional covariance matrices with incomplete data,” *Journal of multivariate analysis*, vol. 150, pp. 55–74, 2016.
- [6] Martin Azizyan, Akshay Krishnamurthy, and Aarti Singh, “Extreme compressive sampling for covariance estimation,” *arXiv preprint arXiv:1506.00898*, 2015.
- [7] Farhad Pourkamali-Anaraki, “Estimation of the sample covariance matrix from compressive measurements,” *IET Signal Processing*, vol. 10, no. 9, pp. 1089–1095, 2016.
- [8] Farhad P Anaraki and Shannon Hughes, “Memory and computation efficient pca via very sparse random projections,” in *Proceedings of the 31st International Conference on Machine Learning (ICML)*, 2014, pp. 1341–1349.
- [9] Mladen Kolar and Eric P Xing, “Consistent covariance selection from data with missing values,” in *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012, pp. 551–558.
- [10] Jerome Friedman, Trevor Hastie, and Robert Tibshirani, “Sparse inverse covariance estimation with the graphical lasso,” *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.
- [11] Gautamd Dasarathy, Aarti Singh, Maria-Florina Balcan, and Jong H Park, “Active learning algorithms for graphical model selection,” in *Artificial Intelligence and Statistics (AISTATS)*, 2016, pp. 1356–1364.
- [12] Divyanshu Vats, Robert Nowak, and Richard Baraniuk, “Active learning for undirected graphical model selection,” in *Artificial Intelligence and Statistics (AISTATS)*, 2014, pp. 958–967.
- [13] Jonathan Scarlett and Volkan Cevher, “Lower bounds on active learning for graphical model selection,” in *The 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- [14] Siheng Chen, Rohan Varma, Aarti Singh, and Jelena Kovačević, “Signal recovery on graphs: Fundamental limits of sampling strategies,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 4, pp. 539–554, 2016.
- [15] Gilles Puy, Nicolas Tremblay, Rémi Gribonval, and Pierre Vandergheynst, “Random sampling of bandlimited signals on graphs,” *Applied and Computational Harmonic Analysis*, 2016.
- [16] Po-Ling Loh and Martin J Wainwright, “High dimensional regression with noisy and missing data: provable guarantees with non-convexity,” *The Annals of Statistics*, vol. 40, no. 3, pp. 1637–1664, 2012.
- [17] Daniel Romero, Dyonisius Dony Ariananda, Zhi Tian, and Geert Leus, “Compressive covariance sensing: Structure-based compressive sensing beyond sparsity,” *IEEE signal processing magazine*, vol. 33, no. 1, pp. 78–93, 2016.
- [18] Roman Vershynin, *High Dimensional Probability: An Introduction with Applications in Data Science*, Cambridge University Press, 2018.
- [19] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278 – 2324, 1998.