

PARTIAL ESTIMATION OF COVARIANCE MATRICES

ELIZAVETA LEVINA AND ROMAN VERSHYNIN

ABSTRACT. A classical approach to accurately estimating the covariance matrix Σ of a p -variate normal distribution is to draw a sample of size $n > p$ and form a sample covariance matrix. However, many modern applications operate with much smaller sample sizes, thus calling for estimation guarantees in the regime $n \ll p$. We show that a sample of size $n = O(m \log^6 p)$ is sufficient to accurately estimate in operator norm an arbitrary symmetric part of Σ consisting of $m \leq n$ nonzero entries per row. This follows from a general result on estimating Hadamard products $M \cdot \Sigma$, where M is an arbitrary symmetric matrix.

1. INTRODUCTION

The problem of estimating the population covariance matrix from a sample of n i.i.d. observations X_1, \dots, X_n in \mathbb{R}^p has attracted a lot of recent interest in statistics, with the focus on small sample sizes, $n \ll p$. Estimation of covariance matrices plays a key role in many data analysis techniques (e.g. in principal component analysis, discriminant analysis, graphical models). Applications where $n \ll p$ have become very common in gene expression data, climate studies, spectroscopy, and so on.

Suppose X_k 's are drawn from the multivariate normal distribution $\mathcal{N}_p(0, \Sigma)$, where $\Sigma = \mathbb{E}X X^T$ is a symmetric positive-semidefinite $p \times p$ matrix. Thus Σ is the covariance matrix of that distribution. The usual estimator for Σ is the sample covariance matrix defined by¹

$$(1.1) \quad \hat{\Sigma}_n = \frac{1}{n} \sum_{k=1}^n X_k X_k^T .$$

The properties of $\hat{\Sigma}_n$ have been extensively studied [14, 11, 12]. Let us first consider the simpler case where $\Sigma = I$. The “Bai-Yin law” of random matrix theory determines the magnitude of the extreme eigenvalues of the Wishart random matrix $\hat{\Sigma}_n$ in the limit as $n, p \rightarrow \infty$ and $n/p \rightarrow \text{const}$. It says that the

Date: August 9, 2010, revised February 11, 2011.

Partially supported by NSF grants DMS 0805798 (E.L.) and FRG DMS 0918623, DMS 1001829 (R.V.).

¹Centering the data with zero mean rather than the sample mean is essentially without loss of generality, since the mean term is of higher order (see Remark 4 after Corollary 2.2).

spectrum of $\hat{\Sigma}_n$ is almost surely contained in the interval $[\frac{a^2}{n} + o(1), \frac{b^2}{n} + o(1)]$ where $a = (\sqrt{n} - \sqrt{p})_+$ and $b = \sqrt{n} + \sqrt{p}$. (The upper bound is due to Geman [10], the lower bound is due to Silverstein [17]; see [1] for a unified exposition). It follows that, with high probability,

$$\|\hat{\Sigma}_n - I\| \leq 2\sqrt{\frac{p}{n}} + \frac{p}{n} + o(1)$$

where $\|\cdot\|$ is the spectral norm, also known as the operator or ℓ_2 matrix norm.

A similar approximation holds for random vectors X_k drawn from a general multivariate normal distribution $\mathcal{N}_p(0, \Sigma)$ (where Σ depends on p). Indeed, arguing as above for the random vectors $\Sigma^{-1/2}X_k$ whose covariance matrix is identity, we easily conclude that with high probability,

$$(1.2) \quad \|\hat{\Sigma}_n - \Sigma\| \leq \left(2\sqrt{\frac{p}{n}} + \frac{p}{n} + o(1)\right) \|\Sigma\|.$$

This implies that for an arbitrary precision $\varepsilon \in (0, 1)$ and for p sufficiently large, the sample size

$$(1.3) \quad n \gtrsim \varepsilon^{-2}p \quad \text{suffices for} \quad \|\hat{\Sigma}_n - \Sigma\| \leq \varepsilon\|\Sigma\|.$$

More accurately, the following non-asymptotic result holds for arbitrary $p \in \mathbb{N}$, $\varepsilon \in (0, 1)$ and $t \geq 1$. If the sample size satisfies $n \geq C(t/\varepsilon)^2p$, then inequality (1.3) holds with probability at least $1 - 2\exp(-t^2n)$ [18, Remark 51]. Here C is an absolute constant. This provides a satisfactory answer to the covariance estimation problem in the regime $n \geq p$ for general p -variate normal distributions.

We will now focus on the more difficult regime $n < p$, where covariance estimation is impossible in general (indeed, if $\Sigma = I$ then clearly $\|\hat{\Sigma}_n - I\| \geq 1$ because of the rank deficiency of $\hat{\Sigma}_n$). Still, under natural structural assumptions on the covariance matrix Σ , a number of alternative estimation schemes have been proposed.

A common assumption is that most entries of Σ are zero or close to 0, and thus can be safely estimated as 0, which reduces the effective size of the problem. This assumption underlies banding or tapering the covariance matrix in the case of ordered variables [3, 9, 16, 5], and thresholding of the covariance matrix in the case of unordered variables [2, 7, 15, 4]. In both cases, the regularized estimators have the form $M \cdot \hat{\Sigma}_n$ where M is a regularizing ‘‘mask’’ matrix, and ‘‘ \cdot ’’ denotes the Hadamard (entrywise) product of matrices. For banding and hard thresholding, M is a sparse matrix containing only 0s and 1s; for tapering, its entries can be anywhere in the interval $[0, 1]$, depending on the taper used. For thresholding, the mask M has to be estimated from $\hat{\Sigma}$, whereas for banding and tapering it is determined by a tuning parameter.

The reason regularized estimators work is best understood from the following decomposition:

$$\|M \cdot \hat{\Sigma}_n - \Sigma\| \leq \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| + \|M \cdot \Sigma - \Sigma\| .$$

The first term (variance) is well-behaved because M is sparse or close to sparse, and the second term (bias) is well-behaved because of the assumptions made on the true Σ . The overall rate for covariance estimation is normally computed as the sum of these two terms, but they come from two quite different problems: the first term has to do with how much of the covariance matrix we can reliably estimate with the corresponding part of the sample covariance matrix; and the second term has to do with our model assumptions.

In this paper, we focus on the first question: *how large does the sample size n have to be so that $M \cdot \hat{\Sigma}_n$ is a good approximation to $M \cdot \Sigma$?* Of course, the answer will depend on the size of the part we would like to estimate, but it almost does not depend on the ambient dimension p as we will show.

As a simple example, consider the case where $M = (m_{ij})$ is a fixed minor, i.e. $m_{ij} = \mathbf{1}_{\{i,j \in S\}}$ for some a priori given index set $S \subset \{1, \dots, p\}$. Denoting $|S| = m$, we can easily deduce from the results above that

$$(1.4) \quad \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| \leq \left(2\sqrt{\frac{m}{n}} + \frac{m}{n} + o(1)\right) \|\Sigma\|.$$

Note that the quality of approximation does not depend on the ambient dimension p but only on the number of selected variables m , as one would expect.

The purpose of this paper is to reveal a general phenomenon behind (1.4). We shall show that, up to a necessary logarithmic factor in p , the upper bound in (1.4) *holds for arbitrary symmetric 0/1 matrices M with at most m nonzero entries per column*. This result is independent of any structural hypotheses on the locations of the nonzero entries.

Even more generally, a stable version of (1.4) holds. It applies to completely *arbitrary symmetric matrices M* , and therefore covers all forms of tapering previously considered in the literature. Instead of the sparsity parameter m , the bound is governed by two different operator norms of M that correspond to the different dependencies on m in the two terms in (1.4). These norms are the $\ell_1 \rightarrow \ell_2$ operator norm $\|M\|_{1,2} = \max_j (\sum_i m_{ij}^2)^{1/2}$, and the usual $\ell_2 \rightarrow \ell_2$ operator norm $\|M\|$.

Note that since the right hand side of (1.3) and (1.4) depends on $\|\Sigma\|$, these are relative rather than the absolute error bounds on the difference between $\hat{\Sigma}$ and Σ . The latter are more commonly analyzed in the literature, even though numerical results are commonly reported as relative errors to enable comparisons across models. The advantage of using the relative error is disentangling

the convergence rates of the estimator from assumptions on the true covariance Σ , which is one of our goals in this paper.

2. MAIN RESULT

Throughout this paper, we consider an arbitrary mean 0 normal distribution in \mathbb{R}^p . Its covariance matrix is denoted by Σ , and the sample covariance matrix (1.1) obtained from an i.i.d. sample of n observations by $\hat{\Sigma}_n$. Positive absolute constants will be denoted C, C_1, C_2 .

Theorem 2.1 (Estimation of Hadamard products). *Let M be an arbitrary fixed symmetric $p \times p$ matrix. Then*

$$(2.1) \quad \mathbb{E} \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| \leq C \log^3(2p) \left(\frac{\|M\|_{1,2}}{\sqrt{n}} + \frac{\|M\|}{n} \right) \|\Sigma\|.$$

Note that M does not depend on $\hat{\Sigma}_n$ or Σ . We will prove this result in Section 4.

Corollary 2.2 (Partial estimation). *Let M be an arbitrary fixed symmetric $p \times p$ matrix such that all of its entries are equal to 0 or 1, and there are at most m nonzero entries in each column. Then*

$$\mathbb{E} \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| \leq C \log^3(2p) \left(\sqrt{\frac{m}{n}} + \frac{m}{n} \right) \|\Sigma\|.$$

Proof. We note that $\|M\|_{1,2} \leq \sqrt{m}$ and $\|M\| \leq m$ and apply Theorem 2.1. \square

Remark 1 (Sample size). Corollary 2.2 implies that for every $\varepsilon \in (0, 1)$, the sample size

$$(2.2) \quad n \geq 4C^2 \varepsilon^{-2} m \log^6(2p) \quad \text{suffices for} \quad \mathbb{E} \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| \leq \varepsilon \|\Sigma\|.$$

For sparse matrices M with $m \ll p$, this makes partial estimation possible with $n \ll p$ observations. Therefore we regard (2.2) as a satisfactory ‘‘sparse’’ version of the classical bound (1.3). A logarithmic term necessarily has to appear in (2.2), see Remark 3 below.

Remark 2 (Compressed sensing). Results of a similar nature arise in *compressed sensing* (see e.g. [8]). There one tries to reconstruct a signal $x \in \mathbb{R}^p$ from its n linear measurements $y = Ax \in \mathbb{R}^n$. In the range $n \geq p$, where the problem is overdetermined, the reconstruction is achieved by inverting the $n \times p$ matrix A (assuming it has full rank). In the range $n \ll p$, the problem is underdetermined, and a reconstruction is only possible under certain structural assumptions on the signal x such as sparsity. If x has m non-zero coordinates, then the typical results of compressed sensing guarantee an exact and algorithmically effective reconstruction provided that $n \gtrsim m \log^C p$.

Our result (2.2) is of a similar nature: $n \sim m \log^6 p$ observations suffice to estimate an m -sparse part of a covariance matrix in dimension p . A significant difference is that the compressed sensing problem becomes easy if one knows the locations of the non-zero coordinates of the signal x . In contrast, the covariance estimation problem remains non-trivial even if one knows the location of the m -sparse part of the covariance matrix (which is the situation in Corollary 2.2).

Remark 3 (Logarithmic factor). In Corollary 2.2 note the mild, logarithmic dependence on the dimension p and the optimal dependence on the sparsity m , which is of the same type as in (1.4). The necessity of the logarithmic term in our results is evident from the example where $\Sigma = M = I$. Indeed, writing in coordinates $X_k = (X_{k1}, \dots, X_{kp}) \in \mathbb{R}^p$, we obtain

$$\mathbb{E} \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| = \mathbb{E} \max_{i \in [p]} \left| \frac{1}{n} \sum_{k=1}^n X_{ki}^2 - 1 \right| \sim \sqrt{\frac{\log(2p)}{n}}.$$

We may compare this with the conclusion of Theorem 2.1 and Corollary 2.2 for this example, which is

$$\mathbb{E} \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| \leq C \frac{\log^3(2p)}{\sqrt{n}}.$$

We see that the logarithmic terms in these results are unavoidable. However, the present exponent 3 of the logarithm is certainly not optimal, and it can probably be reduced to the optimal value $1/2$.

Remark 4 (Centering). If we center the data with the sample mean $\bar{X} = \frac{1}{n} \sum_{k=1}^n X_k$ rather than the true mean, as one would in practice, we have

$$\mathbb{E} \|M \cdot (\Sigma_n - \bar{X} \bar{X}^T) - M \cdot \Sigma\| \leq \mathbb{E} \|M \cdot \Sigma_n - M \cdot \Sigma\| + \mathbb{E} \|M \cdot (\bar{X} \bar{X}^T)\|$$

Then, denoting by $\|x\|_\infty = \max_i |x_i|$ the ℓ_∞ norm in \mathbb{R}^p , we can check that

$$\mathbb{E} \|M \cdot (\bar{X} \bar{X}^T)\| \leq \mathbb{E} \|M\| \|\bar{X}\|_\infty^2 \leq C \|M\| \|\Sigma\| \frac{\log(2p)}{n},$$

so this term can be absorbed into the second term in (2.1).

Remark 5 (Identifying the non-zero entries of Σ by thresholding). For thresholding covariance matrices, our result is not directly applicable since the matrix M is random and estimated from data. However, we note that if the assumption is made that all non-zero entries in Σ are bounded away from zero by a margin of $h > 0$, then a sample size of $n \gtrsim h^{-2} \log(2p)$ would assure that all their locations are estimated correctly with probability approaching 1. With this assumption, we could derive a bound for the thresholded estimator. Since our focus here is not on any one particular estimator, we do not pursue this

extension, particularly since there are already bounds available specifically for thresholding [2, 4].

Remark 6 (Previous results). Results similar to Corollary 2.2 have been known for *banding* covariance matrices. They apply to the mask matrices M with m diagonals consisting of 1's (while all other entries are zero). A result of [3] guarantees in this case an error bound of order $m\sqrt{\log p/n}$. While this bound has a slightly better dependence on p , it implies that the sample size n required for estimation grows as m^2 , whereas our result shows n growing linearly with m . Furthermore, for some tapered version of the m -diagonal matrix M , an error of order $\sqrt{(m + \log p)/n}$ was obtained in [5], which agrees with ours up to $\log p$. The latter result heavily exploits the structure of the tapered banded matrix M , which allows one to decompose it into a weighted sum of small minors and reduce the problem to using (1.4) for each minor. In contrast, our result does not require any structure of the mask matrix M .

Acknowledgement. The authors are grateful to the anonymous referee for careful reading of the manuscript and useful suggestions.

3. PRELIMINARIES

3.1. Proof outline. In the rest of this paper we prove Theorem 2.1. We shall observe that the quadratic form $\langle (M \cdot \hat{\Sigma}_n)x, y \rangle$ is a Gaussian chaos (defined in (3.2) below) for fixed unit vectors x, y on the sphere S^{p-1} . The main difficulty is to control the chaos uniformly for all x, y . This will be done by establishing concentration inequalities of varying power depending on the “sparsity” of x, y (which amounts to decoupling, conditioning, and Gaussian concentration), and combining this with covering arguments to “count” the number of sparse vectors x, y on the sphere.

3.2. Decoupling. The following observation does not seem to be found in literature on decoupling (such as [6]).

Proposition 3.1 (Decoupling of sample covariance matrices). *Let X be a centered random normal vector in \mathbb{R}^p with covariance matrix Σ , and let X_1, \dots, X_n and X'_1, \dots, X'_n be independent copies of X . Consider the sample covariance matrix and the decoupled sample covariance matrix of X defined as*

$$(3.1) \quad \hat{\Sigma}_n = \frac{1}{n} \sum_{k=1}^n X_k \otimes X_k, \quad \Sigma'_n = \frac{1}{n} \sum_{k=1}^n X'_k \otimes X_k.$$

Then for every symmetric $p \times p$ matrix M we have

$$\mathbb{E} \|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| \leq 2\mathbb{E} \|M \cdot \Sigma'_n\|.$$

Proof. Computing the operator norm as the maximal value of the associated quadratic form, we have

$$\begin{aligned} \mathbb{E}\|M \cdot \hat{\Sigma}_n - M \cdot \Sigma\| &= \mathbb{E}\|M \cdot \hat{\Sigma}_n - \mathbb{E}(M \cdot \hat{\Sigma}_n)\| \\ &= \mathbb{E} \sup_{a \in S^{p-1}} |\langle (M \cdot \hat{\Sigma}_n)a, a \rangle - \mathbb{E}\langle (M \cdot \hat{\Sigma}_n)a, a \rangle|. \end{aligned}$$

Writing the inner product in coordinates

$$\langle (M \cdot \hat{\Sigma}_n)a, a \rangle = \frac{1}{n} \sum_{k=1}^n \sum_{i,j=1}^p m_{ij} a_i a_j X_{ki} X_{kj}$$

we recognize in the right hand side a *quadratic Gaussian chaos*. Recall that generally, a quadratic Gaussian chaos in dimension p is a quadratic form

$$(3.2) \quad \sum_{i,j=1}^p a_{ij} Z_i Z_j = \langle AZ, Z \rangle$$

where $Z = (Z_1, \dots, Z_p)$ is a centered normal random vector in \mathbb{R}^p , and the coefficients $A = (a_{ij})_{i,j=1}^p$ are taken from some subset \mathcal{A} of $p \times p$ matrices. The conclusion of the proposition then follows from the next lemma. \square

Lemma 3.2 (Decoupling of a Gaussian chaos). *Let Z be a centered normal random vector in \mathbb{R}^d , and let Z' be an independent copy of Z . Let \mathcal{A} be a subset of symmetric $d \times d$ matrices. Then*

$$\mathbb{E} \sup_{A \in \mathcal{A}} |\langle AZ, Z \rangle - \mathbb{E}\langle AZ, Z \rangle| \leq 2 \mathbb{E} \sup_{A \in \mathcal{A}} |\langle AZ, Z' \rangle|.$$

Lemma 3.2 allows one to replace a Gaussian chaos $\sum_{i,j} a_{ij} Z_i Z_j$ by its decoupled version $\sum_{i,j} a_{ij} Z'_i Z_j$. While it can be derived from observations in [13, Section 3.2], it is simpler to give a complete proof.

Proof. Without loss of generality, we may assume that Z is a standard normal random vector. Indeed, there exists a symmetric $d \times d$ matrix T such that $Z = Tg$ and $Z' = Tg'$ for some independent standard Gaussian vectors g, g' in \mathbb{R}^d . Thus $\langle AZ, Z \rangle = \langle TATg, g \rangle$ and $\langle AZ, Z' \rangle = \langle TATg, g' \rangle$. Replacing A by TAT in the statement of the lemma, we see that we can assume that Z is a standard normal random vector.

Using the identical distribution of Z and Z' and Jensen's inequality, we obtain

$$\begin{aligned} E &:= \mathbb{E} \sup_{A \in \mathcal{A}} |\langle AZ, Z \rangle - \mathbb{E}\langle AZ, Z \rangle| = \mathbb{E} \sup_{A \in \mathcal{A}} |\langle AZ, Z \rangle - \mathbb{E}\langle AZ', Z' \rangle| \\ &\leq \mathbb{E} \sup_{A \in \mathcal{A}} |\langle AZ, Z \rangle - \langle AZ', Z' \rangle|. \end{aligned}$$

Note the identity $\langle AZ, Z \rangle - \langle AZ', Z' \rangle = 2\langle A(\frac{Z+Z'}{\sqrt{2}}, \frac{Z-Z'}{\sqrt{2}}) \rangle$. By rotation invariance of Gaussian measure, the pair of vectors $(\frac{Z+Z'}{\sqrt{2}}, \frac{Z-Z'}{\sqrt{2}})$ is distributed identically with (Z, Z') . Hence we conclude that

$$E \leq 2\mathbb{E} \sup_{A \in \mathcal{A}} |\langle AZ, Z' \rangle|.$$

This completes the proof. \square

Remark 7. It is not clear if a version of Decoupling Proposition 3.1 holds for general distributions. The argument we gave relied heavily on the rotation invariance of the normal distribution.

3.3. Concentration. The following observation is a form of the rotation invariance property of normal distributions.

Lemma 3.3. *Let $Z = (Z_1, \dots, Z_d)$ be a centered normal random vector in \mathbb{R}^d with covariance matrix Σ . Let $a = (a_1, \dots, a_d) \in \mathbb{R}^d$. Then $\sum_{i=1}^d a_i Z_i$ is a centered normal random variable with standard deviation $\|\Sigma^{1/2}a\|_2 \leq \|\Sigma\|^{1/2}\|a\|_2$. \square*

Our proof of Theorem 2.1 will use the concentration of measure in the Gauss space. Such concentration results are usually stated in the literature for the standard normal distribution, but it is straightforward to make an adjustment for general normal distributions.

Proposition 3.4 (Concentration in the Gauss space, see [13] Section 1.1). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a Lipschitz function, and denote $L = \|f\|_{\text{Lip}}$. Let Z be a centered normal vector in \mathbb{R}^d with covariance matrix Σ . Then*

$$\mathbb{P}\{f(Z) - \mathbb{E}f(Z) \geq t\} \leq \frac{1}{2} \exp\left(-\frac{t^2}{2L^2\|\Sigma\|}\right) \quad \text{for all } t \geq 0.$$

Remark 8. If f is a norm on \mathbb{R}^d , then by translation invariance, $\|f\|_{\text{Lip}}$ equals the minimal number L such that

$$f(x) \leq \|x\|_2 \quad \text{for all } x \in \mathbb{R}^d.$$

3.4. Discretization. The operator norm of a $p \times p$ matrix A can be computed via the associated bilinear form as

$$\|A\| = \max_{x, y \in S^{p-1}} \langle Ax, y \rangle$$

By an approximation argument, one can replace the sphere S^{p-1} by an ε -net:

Lemma 3.5 (Computing operator norm on a net, see [18]). *Let A be a $p \times p$ matrix. Let \mathcal{N} be a δ -net of S^{p-1} in the Euclidean metric for some $\delta \in [0, 1)$. Then*

$$\|A\| \leq (1 - \delta)^{-2} \max_{x, y \in \mathcal{N}} \langle Ax, y \rangle.$$

We will construct a net \mathcal{N} by rounding off the coefficients of a vector x .

Definition 3.6 (Regular vectors). Define the subset of *regular vectors* of the sphere S^{p-1} as follows:

$$\begin{aligned} \text{Reg}_p(s) &= \{x \in S^{p-1} : \text{all coordinates satisfy } x_i^2 \in \{0, 1/s\}\}, \quad s \in [p]; \\ \text{Reg}_p &= \bigcup_{s \in [p]} \text{Reg}_p(s). \end{aligned}$$

Note that $|\text{supp}(x)| = s$ for all $x \in \text{Reg}_p(s)$.

Lemma 3.7 (Computing operator norm on regular vectors). *Let A be a $p \times p$ matrix. Then*

$$\|A\| \leq 12 \lceil \ln(2p) \rceil^2 \max_{x, y \in \text{Reg}_p} \langle Ax, y \rangle.$$

Proof. Define $k_0 = \lceil \ln(2p) \rceil$ and let \mathcal{N} consist of all vectors $x \in S^{p-1}$ such that for each coordinate x_i there is $k \in \{0, 1, \dots, k_0\}$ such that $|x_i| = 2^{-k}$. It is easy to check that \mathcal{N} is a 0.71-net of S^{p-1} .

By considering the level sets of a vector $x \in \mathcal{N}$, one can represent $x = \lambda_1 x^{(1)} + \dots + \lambda_{k_0} x^{(k_0)}$ where all $\lambda_k \in [0, 1]$ and $x^{(k)}$ are regular vectors. Using this representation for x and y in Lemma 3.5 with $\delta = 0.71$, we conclude by the linearity of the inner product that the estimate in the lemma holds. \square

Lemma 3.8 (Computing operator norm: symmetric distributions). *Let A be a random $p \times p$ matrix such that A and A^T are identically distributed. Then, for every $t \geq 0$, we have*

$$\mathbb{P}\{\|A\| \geq t\} \leq 2\mathbb{P}\left\{12 \lceil \ln(2p) \rceil^2 \max_{\substack{r, s \in [p] \\ r \leq s}} \max_{\substack{x \in \text{Reg}_p(r) \\ y \in \text{Reg}_p(s)}} \langle Ax, y \rangle \geq t\right\}.$$

Proof. We can write the conclusion of Lemma 3.7 as

$$\|A\| \leq 12 \lceil \ln(2p) \rceil^2 \max_{r, s \in [p]} \max_{\substack{x \in \text{Reg}_p(r) \\ y \in \text{Reg}_p(s)}} \langle Ax, y \rangle.$$

Taking the maximum separately over $r \leq s$ and $r > s$ and replacing the maximum by the sum, we have

$$(3.3) \quad \|A\| \leq 12 \lceil \ln(2p) \rceil^2 \max \left(\max_{\substack{r, s \in [p] \\ r \leq s}} \max_{\substack{x \in \text{Reg}_p(r) \\ y \in \text{Reg}_p(s)}} \langle Ax, y \rangle, \max_{\substack{r, s \in [p] \\ s \leq r}} \max_{\substack{x \in \text{Reg}_p(r) \\ y \in \text{Reg}_p(s)}} \langle Ax, y \rangle \right).$$

By the symmetry assumption on the distribution of A , the random variables $\langle Ax, y \rangle$ and $\langle Ay, x \rangle$ are identically distributed for all x, y . Therefore, the two double maxima in (3.3) are identically distributed. This yields the conclusion of the lemma. \square

4. PROOF OF THEOREM 2.1

4.1. Decoupling and conditioning. By rescaling, we can assume without loss of generality that $\|\Sigma\| = 1$. We shall denote the entries of M by m_{ij} .

By Proposition 3.1, it suffices to estimate $\mathbb{E}\|M \cdot \Sigma'_n\|$, where Σ'_n denotes the decoupled sample covariance matrix defined in (3.1). By the definition of Σ'_n and the symmetry of M , the random matrices $M \cdot \Sigma'_n$ and $(M \cdot \Sigma'_n)^T$ are identically distributed. Therefore, Lemma 3.8 applies and we get for every $t \geq 0$ that

$$(4.1) \quad \mathbb{P}\{\|M \cdot \Sigma'_n\| \geq t\} \leq 2\mathbb{P}\left\{12\lceil\ln(2p)\rceil^2 \max_{\substack{r,s \in [p] \\ r \leq s}} \max_{\substack{x \in \text{Reg}_p(r) \\ y \in \text{Reg}_p(s)}} \langle (M \cdot \Sigma'_n)x, y \rangle \geq t\right\}.$$

Writing the inner product in coordinates and rearranging the terms, we have

$$(4.2) \quad \langle (M \cdot \Sigma'_n)x, y \rangle = \frac{1}{n} \sum_{k=1}^n \sum_{i=1}^p \left(\sum_{j=1}^p m_{ij} x_j X_{kj} \right) y_i X'_{ki}.$$

Let us fix x and y and condition on the random vectors (X_1, \dots, X_n) . Then the expression in (4.2) defines a centered normal random variable. We shall estimate its standard deviation by Lemma 3.3. Since the covariance matrix Σ of each vector $X'_k \in \mathbb{R}^p$ has operator norm at most 1, the covariance matrix of the concatenated vector $(X'_k)_{k=1}^n \in \mathbb{R}^{pn}$ also has operator norm at most 1. Then Lemma 3.3 yields that the expression in (4.2) is a centered normal random variable with standard deviation at most $\sigma_x(X_1, \dots, X_n) \|y\|_\infty$ where

$$\sigma_x(X_1, \dots, X_n) = \frac{1}{n} \left[\sum_{k=1}^n \sum_{i=1}^p \left(\sum_{j=1}^p m_{ij} x_j X_{kj} \right)^2 \right]^{1/2}.$$

We will need to bound this quantity uniformly for all x .

4.2. Concentration. Let $x \in \text{Reg}_p(r)$. We will estimate $\sigma_x(X_1, \dots, X_n)$ using concentration in Gauss space, Proposition 3.4. First, we have

$$\begin{aligned}
\mathbb{E}\sigma_x(X_1, \dots, X_n) &\leq (\mathbb{E}\sigma_x(X_1, \dots, X_n)^2)^{1/2} \\
&= \frac{1}{n} \left[\sum_{k=1}^n \sum_{i=1}^p \mathbb{E} \left(\sum_{j=1}^p m_{ij} x_j X_{kj} \right)^2 \right]^{1/2} \\
&\leq \frac{1}{n} \left(\sum_{k=1}^n \sum_{i=1}^p \sum_{j=1}^p m_{ij}^2 x_j^2 \right)^{1/2} \quad (\text{by Lemma 3.3 and } \|\Sigma\| = 1) \\
&\leq \frac{1}{\sqrt{n}} \max_{j \in [p]} \left(\sum_{i=1}^p m_{ij}^2 \right)^{1/2} \quad (\text{because } \|x\|_2 = 1) \\
(4.3) \quad &\leq \frac{\|M\|_{1,2}}{\sqrt{n}} \quad (\text{by definition}).
\end{aligned}$$

Next, we consider $\sigma_x : \mathbb{R}^{pn} \rightarrow \mathbb{R}$ as a function of the concatenated Gaussian vector $(X_1, \dots, X_n) \in \mathbb{R}^{pn}$. Computing the Lipschitz norm of σ_x becomes easy once we write this function as

$$\sigma_x(X_1, \dots, X_n) = \frac{1}{n} \left(\sum_{k=1}^n \|M(x \cdot X_k)\|_2^2 \right)^{1/2}$$

where as usual $x \cdot X_k$ denotes the Hadamard (coordinate-wise) product of vectors, and the multiplication by M is the ordinary (matrix) multiplication. Separating M and x , we obtain the bound

$$\sigma_x(X_1, \dots, X_n) \leq \frac{1}{n} \|M\| \|x\|_\infty \left(\sum_{k=1}^n \|X_k\|_2^2 \right)^{1/2}.$$

Using that $\|x\|_\infty = 1/\sqrt{r}$ since $x \in \text{Reg}_p(r)$, we conclude that

$$\sigma_x(X_1, \dots, X_n) \leq \frac{\|M\|}{\sqrt{r}n} \cdot \|(X_1, \dots, X_n)\|_2.$$

By the remark below Proposition 3.4, we have proved that

$$(4.4) \quad \|\sigma_x\|_{\text{Lip}} \leq \frac{\|M\|}{\sqrt{r}n}.$$

In addition, since the covariance matrix Σ of each vector $X_k \in \mathbb{R}^p$ has operator norm at most 1, the covariance matrix of the concatenated vector $(X_1, \dots, X_n) \in \mathbb{R}^{pn}$ also has operator norm at most 1. From this and the bounds on the expectation (4.3) and on the Lipschitz norm (4.4), we conclude by Proposition 3.4 that for all $x \in \text{Reg}_p(r)$ and $t \geq 0$,

$$(4.5) \quad \mathbb{P} \left\{ \sigma_x(X_1, \dots, X_n) > \frac{\|M\|_{1,2}}{\sqrt{n}} + t \right\} \leq \frac{1}{2} \exp \left(- \frac{t^2 r n^2}{2 \|M\|^2} \right).$$

4.3. Union bounds. We return to estimating the random variable $\langle (M \cdot \Sigma'_n)x, y \rangle$ which we initiated in Section 4.1. Let us fix $u \geq 1$. For each $x \in \text{Reg}_p(r)$, we consider the events

$$\mathcal{E}_x := \left\{ \sigma_x(X_1, \dots, X_n) \leq \frac{\|M\|_{1,2}}{\sqrt{n}} + u \frac{\|M\|}{n} \sqrt{10 \ln(2ep)} \right\}$$

By (4.5), we have

$$(4.6) \quad \mathbb{P}(\mathcal{E}_x) \geq 1 - \frac{1}{2} \exp(-5u^2 r \ln(2ep)).$$

Note that the function $\sigma_x(X_1, \dots, X_n)$ and thus also the events \mathcal{E}_x are independent of the random variables (X'_1, \dots, X'_n) .

Let $x \in \text{Reg}_p(r)$ and $y \in \text{Reg}_p(s)$. As we noted in Section 4.1, conditioned on a realization of random variables (X_1, \dots, X_n) satisfying \mathcal{E}_x , the random variable $\langle (M \cdot \Sigma'_n)x, y \rangle$ is distributed identically with a centered normal random variable h whose standard deviation is bounded by

$$\sigma_x(X_1, \dots, X_n) \|y\|_\infty \leq \frac{\|M\|_{1,2}}{\sqrt{sn}} + u \frac{\|M\|}{\sqrt{sn}} \sqrt{10 \ln(2ep)} =: \sigma.$$

Then by the usual tail estimate for Gaussian random variables, we have

$$\mathbb{P}\{\langle (M \cdot \Sigma'_n)x, y \rangle \geq \varepsilon \mid \mathcal{E}_x\} \leq \frac{1}{2} \exp(-\varepsilon^2/2\sigma^2) \quad \text{for } \varepsilon \geq 0.$$

Choosing

$$(4.7) \quad \varepsilon = \varepsilon(u) := 2\sqrt{3} u \frac{\|M\|_{1,2}}{\sqrt{n}} \sqrt{\ln(2ep)} + 2\sqrt{30} u^2 \frac{\|M\|}{n} \ln(2ep),$$

we obtain

$$\mathbb{P}\{\langle (M \cdot \Sigma'_n)x, y \rangle \geq \varepsilon \mid \mathcal{E}_x\} \leq \frac{1}{2} \exp(-6u^2 s \ln(2ep))$$

for all $x \in \text{Reg}_p(r)$, $y \in \text{Reg}_p(s)$. We would like to take the union bound in this estimate over all $y \in \text{Reg}_p(s)$. Note that

$$|\text{Reg}_p(s)| = \binom{p}{s} 2^s \leq \exp(s \ln(2ep/s))$$

as there are $\binom{p}{s}$ ways to choose the support and 2^s ways to choose the signs of the coefficients of a vector in $\text{Reg}_p(s)$. Then

$$\begin{aligned} \mathbb{P}\left\{ \max_{y \in \text{Reg}_p(s)} \langle (M \cdot \Sigma'_n)x, y \rangle \geq \varepsilon \mid \mathcal{E}_x \right\} &\leq \frac{1}{2} \exp(s \ln(2ep/s)) \exp(-6u^2 s \ln(2ep)) \\ &\leq \frac{1}{2} \exp(-5u^2 s \ln(2ep)) \end{aligned}$$

as $u \geq 1$. Therefore, using (4.6), we have

$$(4.8) \quad \mathbb{P}\left\{ \max_{y \in \text{Reg}_p(s)} \langle (M \cdot \Sigma'_n)x, y \rangle \geq \varepsilon \right\} \leq \mathbb{P}\left\{ \max_{y \in \text{Reg}_p(s)} \langle (M \cdot \Sigma'_n)x, y \rangle \geq \varepsilon \mid \mathcal{E}_x \right\} + \mathbb{P}\{\mathcal{E}_x^c\}$$

$$\leq \frac{1}{2} \exp(-5u^2 s \ln(2ep)) + \frac{1}{2} \exp(-5u^2 r \ln(2ep)).$$

Now we take a further union bound over $x \in \text{Reg}_p(r)$ for fixed $r, p \in [p]$, $r \leq s$. In this range, the second term in (4.8) dominates. Estimating as before $|\text{Reg}_p(r)| \leq \exp(r \ln(2ep/r))$, we obtain that

$$\mathbb{P}\left\{ \max_{\substack{x \in \text{Reg}_p(r) \\ y \in \text{Reg}_p(s)}} \langle (M \cdot \Sigma'_n)x, y \rangle \geq \varepsilon \right\} \leq \exp(r \ln(2ep/r)) \exp(-5u^2 r \ln(2ep))$$

$$\leq \exp(-4u^2 r \ln(2ep)).$$

Finally, we take the union bound over all allowed pairs r, s . Since $r \geq 1$ and $u \geq 1$, we have

$$\mathbb{P}\left\{ \max_{\substack{r, s \in [p] \\ r \leq s}} \max_{\substack{x \in \text{Reg}_p(r) \\ y \in \text{Reg}_p(s)}} \langle (M \cdot \Sigma'_n)x, y \rangle \geq \varepsilon \right\} \leq p^2 \exp(-4u^2 \ln(2ep))$$

$$\leq \exp(-2u^2 \ln(2ep)).$$

Using (4.1), we have shown that

$$\mathbb{P}\{\|M \cdot \Sigma'_n\| \geq 12 \lceil \ln(2p) \rceil^2 \varepsilon\} \leq 2 \exp(-2u^2 \ln(2ep))$$

for all $u \geq 1$ and for $\varepsilon = \varepsilon(u)$ defined in (4.7). Integration yields

$$(4.9) \quad \mathbb{E}\|M \cdot \Sigma'_n\| \leq 84 \frac{\|M\|_{1,2}}{\sqrt{n}} \lceil \ln(2ep) \rceil^{5/2} + 263 \frac{\|M\|}{n} \lceil \ln(2ep) \rceil^3.$$

Decoupling Proposition 3.1 completes the proof of Theorem 2.1, giving also an explicit bound on the absolute constant C and a slightly better dependence on p . \square

Remark 9 (Arbitrary distributions). It is likely that the results of this paper generalize from Gaussian to arbitrary distributions in \mathbb{R}^p with enough moments. However, it is not clear whether a version of Decoupling Proposition 3.1 holds for general distributions.

REFERENCES

- [1] Z. Bai and Y. Yin. Limit of the smallest eigenvalue of a large-dimensional sample covariance matrix. *Annals of Probability*, 21:1275–1294, 1993.
- [2] P. J. Bickel and E. Levina. Covariance regularization by thresholding. *Ann. Statist.*, 36(6):2577–2604, 2008.
- [3] P. J. Bickel and E. Levina. Regularized estimation of large covariance matrices. *Ann. Statist.*, 36(1):199–227, 2008.

- [4] T. T. Cai and H. H. Zhou. Minimax estimation of large covariance matrices under ℓ_1 -norm. 2010. Manuscript.
- [5] T. Tony Cai, C.-H. Zhang, and Harrison H. Zhou. Optimal rates of convergence for covariance matrix estimation. *Ann. Statist.*, 2010. To appear.
- [6] V. de la Peña and E. Giné. *Decoupling. From dependence to independence. Randomly stopped processes. U-statistics and processes. Martingales and beyond*. New York: Springer-Verlag, 1999.
- [7] Noureddine El Karoui. Operator norm consistent estimation of large dimensional sparse covariance matrices. *Ann. Statist.*, 36(6):2717–2756, 2008.
- [8] M. Fornasier and H. Rauhut. Compressive sensing. In O. Scherzer, editor, *Handbook of Mathematical Methods in Imaging*. Springer, 2011.
- [9] R. Furrer and T. Bengtsson. Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *Journal of Multivariate Analysis*, 98(2):227–255, 2007.
- [10] S. Geman. A limit theorem for the norm of random matrices. *Annals of Probability*, 8:252–261, 1980.
- [11] I. M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.*, 29(2):295–327, 2001.
- [12] I. M. Johnstone and A. Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, June 2009.
- [13] M. Ledoux and M. Talagrand. *Probability in Banach spaces: Isoperimetry and processes*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*. Springer-Verlag, Berlin, 1991.
- [14] V. A. Marčenko and L. A. Pastur. Distributions of eigenvalues of some sets of random matrices. *Math. USSR-Sb*, 1:507–536, 1967.
- [15] A. J. Rothman, E. Levina, and J. Zhu. Generalized thresholding of large covariance matrices. *J. Amer. Statist. Assoc. (Theory and Methods)*, 104(485):177–186, 2009.
- [16] A. J. Rothman, E. Levina, and J. Zhu. A new approach to Cholesky-based estimation of high-dimensional covariance matrices. *Biometrika*, 97(3):539–550, 2010.
- [17] J. Silverstein. The smallest eigenvalue of a large-dimensional wishart matrix. *Annals of Probability*, 13:1364–1368, 1985.
- [18] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y. Eldar and G. Kutyniok, editors, *Compressed sensing: theory and applications*. Cambridge University Press. Submitted.

DEPARTMENT OF STATISTICS, UNIVERSITY OF MICHIGAN, 1085 S. UNIVERSITY, ANN ARBOR, MI 48109, U.S.A.

E-mail address: `elevina@umich.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, 530 CHURCH ST., ANN ARBOR, MI 48109, U.S.A.

E-mail address: `romanv@umich.edu`