Journal of Machine Learning Research 17 (2016) 1-41 Submitted 4/15; Revised 4/16; Published 8/16

Sparse PCA via Covariance Thresholding

Yash Deshpande YASH.DESHPANDE@STANFORD.EDU
Department of Electrical Engineering
Stanford University
Stanford, CA 94305, USA

Andrea Montanari MONTANARI@STANFORD.EDU
Departments of Electrical Engineering and Statistics
Stanford University
Stanford, CA 94305, USA

Editor: Alexander Rakhlin

Abstract

In sparse principal component analysis we are given noisy observations of a low-rank matrix of dimension $n \times p$ and seek to reconstruct it under additional sparsity assumptions. In particular, we assume here each of the principal components $\mathbf{v}_1, \dots, \mathbf{v}_r$ has at most s_0 non-zero entries. We are particularly interested in the high dimensional regime wherein p is comparable to, or even much larger than n.

In an influential paper, Johnstone and Lu (2004) introduced a simple algorithm that estimates the support of the principal vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ by the largest entries in the diagonal of the empirical covariance. This method can be shown to identify the correct support with high probability if $s_0 \le K_1\sqrt{n/\log p}$, and to fail with high probability if $s_0 \ge K_2\sqrt{n/\log p}$ for two constants $0 < K_1, K_2 < \infty$. Despite a considerable amount of work over the last ten years, no practical algorithm exists with provably better support recovery guarantees.

Here we analyze a covariance thresholding algorithm that was recently proposed by Krauthgamer, Nadler, Vilenchik, et al. (2015). On the basis of numerical simulations (for the rank-one case), these authors conjectured that covariance thresholding correctly recover the support with high probability for $s_0 \le K\sqrt{n}$ (assuming n of the same order as p). We prove this conjecture, and in fact establish a more general guarantee including higher-rank as well as n much smaller than p. Recent lower bounds (Berthet and Rigollet, 2013; Ma and Wigderson, 2015) suggest that no polynomial time algorithm can do significantly better.

The key technical component of our analysis develops new bounds on the norm of kernel random matrices, in regimes that were not considered before. Using these, we also derive sharp bounds for estimating the population covariance, and the principal component (with ℓ_2-loss).

©2016 Yash Deshpande and Andrea Montanari.

1. Introduction

In the spiked covariance model proposed by Johnstone and Lu (2004), we are given data $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ with $\mathbf{x}_i \in \mathbb{R}^p$ of the form[1]:

$$\mathbf{x}_i = \sum_{q=1}^{r} \sqrt{\beta_q}\, u_{q,i}\, \mathbf{v}_q + \mathbf{z}_i\,, \tag{1}$$

Here $\mathbf{v}_1, \dots, \mathbf{v}_r \in \mathbb{R}^p$ is a set of orthonormal vectors, that we want to estimate, while $u_{q,i} \sim \mathsf{N}(0,1)$ and $\mathbf{z}_i \sim \mathsf{N}(0, \mathrm{I}_p)$ are independent and identically distributed. The quantity $\beta_q \in \mathbb{R}_{>0}$ is a measure of signal-to-noise ratio. In the rest of this introduction, in order to simplify the exposition, we will refer to the rank one case and drop the subscript $q \in \{1, 2, \dots, r\}$. Further, we will assume n to be of the same order as p. Our results and proofs hold for a broad range of scalings of r, p, n, and will be stated in general form.

The standard method of principal component analysis involves computing the sample covariance matrix $\mathbf{G} = n^{-1} \sum_{i=1}^{n} \mathbf{x}_i \mathbf{x}_i^{\mathsf{T}}$ and estimates $\mathbf{v} = \mathbf{v}_1$ by its principal eigenvector $\mathbf{v}_{\rm PC}(\mathbf{G})$. It is a well-known fact that, in the high dimensional regime, this yields an inconsistent estimate (see Johnstone and Lu (2009)). Namely $\|\mathbf{v}_{\rm PC} - \mathbf{v}\| \not\to 0$ unless $p/n \to 0$. Even worse, Baik, Ben Arous, and Péché (2005) and Paul (2007) demonstrate the following phase transition phenomenon. Assuming that $p/n \to \alpha \in (0, \infty)$, if $\beta < \sqrt{\alpha}$ the estimate is *asymptotically orthogonal* to the signal, i.e. $\langle \mathbf{v}_{\rm PC}, \mathbf{v} \rangle \to 0$. On the other hand, for $\beta > \sqrt{\alpha}$, $|\langle \mathbf{v}_{\rm PC}, \mathbf{v} \rangle|$ remains bounded away from zero as $n, p \to \infty$. This phase transition phenomenon has attracted considerable attention recently within random matrix theory (see, e.g. Féral and Péché, 2007; Capitaine et al., 2009; Benaych-Georges and Nadakuditi, 2011; Knowles and Yin, 2013).

These inconsistency results motivated several efforts to exploit additional structural information on the signal $\mathbf{v}$. In two influential papers, Johnstone and Lu (2004, 2009) considered the case of a signal $\mathbf{v}$ that is sparse in a suitable basis, e.g. in the wavelet domain. Without loss of generality, we will assume here that $\mathbf{v}$ is sparse in the canonical basis $\mathbf{e}_1, \dots \mathbf{e}_p$. In a nutshell, Johnstone and Lu (2009) propose the following:

1. Order the diagonal entries of the Gram matrix $\mathbf{G}_{i(1),i(1)} \ge \mathbf{G}_{i(2),i(2)} \ge \dots \ge \mathbf{G}_{i(p),i(p)}$, and let $J \equiv \{i(1), i(2), \dots, i(k)\}$ be the set of indices corresponding to the s_0 largest entries.

2. Set to zero all the entries $\mathbf{G}_{i,j}$ of $\mathbf{G}$ unless $i, j \in J$, and estimate $\mathbf{v}$ with the principal eigenvector of the resulting matrix.

Johnstone and Lu formalized the sparsity assumption by requiring that $\mathbf{v}$ belongs to a weak ℓ_q-ball with $q \in (0,1)$. Instead, here we consider a strict sparsity constraint where $\mathbf{v}$ has exactly s_0 non-zero entries, with magnitudes bounded below by $\theta/\sqrt{s_0}$ for some constant $\theta > 0$. Amini and Wainwright (2009) studied the more restricted case when every entry of $\mathbf{v}$ has equal magnitude of $1/\sqrt{s_0}$. Within this restricted model, they proved diagonal thresholding successfully recovers the support of $\mathbf{v}$ provided the sample size n satisfies[2]

1. Throughout the paper, we follow the convention of denoting scalars by lowercase, vectors by lowercase boldface, and matrices by uppercase boldface letters.
2. Throughout the introduction, we write $f(n) \gtrsim g(n)$ as a shorthand of '$f(n) \ge K\, g(n)$ *for a some constant* $K = K(r, \beta)$'.

$n \gtrsim s_0^2 \log p$ (see Amini and Wainwright, 2009). This result is a striking improvement over vanilla PCA. While the latter requires a number of samples scaling with the number of parameters $n \gtrsim p$, sparse PCA via diagonal thresholding achieves the same objective with a number of samples that scales with the number of *non-zero* parameters, $n \gtrsim s_0^2 \log p$.

At the same time, this result is not as strong as might have been expected. By searching exhaustively over all possible supports of size s_0 (a method that has complexity of order p^{s_0}) the correct support can be identified with high probability as soon as $n \gtrsim s_0 \log p$. No method can succeed for much smaller n, because of information theoretic obstructions. We refer the reader to Amini and Wainwright (2009) for more details.

Over the last ten years, a significant effort has been devoted to developing practical algorithms that outperform diagonal thresholding, see e.g. Moghaddam et al. (2005); Zou et al. (2006); d'Aspremont et al. (2007, 2008); Witten et al. (2009). In particular, d'Aspremont et al. (2007) developed a promising M-estimator based on a semidefinite programming (SDP) relaxation. Amini and Wainwright (2009) also carried out an analysis of this method and proved that, if[3] *(i)* $n \geq K(\beta)\, s_0 \log(p - s_0)p$, and *(ii)* the SDP solution has rank one, then the SDP relaxation provides a consistent estimator of the support of $\mathbf{v}$.

At first sight, this appears as a satisfactory solution of the original problem. No procedure can estimate the support of $\mathbf{v}$ from less than $s_0 \log p$ samples, and the SDP relaxation succeeds in doing it from –at most– a constant factor more samples. This picture was upset by a recent, remarkable result by Krauthgamer et al. (2015) who showed that the rank-one condition assumed by Amini and Wainwright does not hold for $\sqrt{n} \lesssim s_0 \lesssim (n/\log p)$. This result is consistent with recent work of Berthet and Rigollet (2013) demonstrating that sparse PCA cannot be performed in polynomial time in the regime $s_0 \gtrsim \sqrt{n}$, under a certain computational complexity conjecture for the so-called planted clique problem.

In summary, the sparse PCA problem demonstrates a fascinating interplay between computational and statistical barriers.

From a statistical perspective, and disregarding computational considerations, the support of $\mathbf{v}$ can be estimated consistently if and only if $s_0 \lesssim n/\log p$. This can be done, for instance, by exhaustive search over all the $\binom{p}{s_0}$ possible supports of $\mathbf{v}$. We refer to Vu and Lei (2012); Cai et al. (2013) for a minimax analysis.

From a computational perspective, the problem appears to be much more difficult. There is rigorous evidence (Berthet and Rigollet, 2013; **?**; Ma and Wigderson, 2015; Wang et al., 2014) that no polynomial algorithm can reconstruct the support unless $s_0 \lesssim \sqrt{n}$. On the positive side, a very simple algorithm (Johnstone and Lu's diagonal thresholding) succeeds for $s_0 \lesssim \sqrt{n/\log p}$.

Of course, several elements are still missing in this emerging picture. In the present paper we address one of them, providing an answer to the following question:

> *Is there a polynomial time algorithm that is guaranteed to solve the sparse PCA problem with high probability for $\sqrt{n/\log p} \lesssim s_0 \lesssim \sqrt{n}$?*

3. Throughout the paper, we denote by K constants that can depend on problem parameters r and β. We denote by upper case C (lower case c) generic absolute constants that are bigger (resp. smaller) than 1, but which might change from line to line.

We answer this question positively by analyzing a covariance thresholding algorithm that proceeds, briefly, as follows. (A precise, general definition, with some technical changes is given in the next section.)

1. Form the empirical covariance matrix $\mathbf{G}$ and set to zero all its entries that are in modulus smaller than $\tau/\sqrt{n}$, for τ a suitably chosen constant.
2. Compute the principal eigenvector $\widehat{\mathbf{v}}_1$ of this thresholded matrix.
3. Denote by $\mathsf{B} \subseteq \{1, \dots, p\}$ be the set of indices corresponding to the s_0 largest entries of $\widehat{\mathbf{v}}_1$.
4. Estimate the support of $\mathbf{v}$ by 'cleaning' the set B. (Briefly, $\mathbf{v}$ is estimated by thresholding $\mathbf{G}\widehat{\mathbf{v}}_{\mathsf{B}}$ with $\widehat{\mathbf{v}}_{\mathsf{B}}$ obtained by zeroing the entries outside B.)

Such a covariance thresholding approach was proposed in Krauthgamer et al. (2015), and is in turn related to earlier work by Bickel and Levina (2008b); Cai et al. (2010). The formulation discussed in the next section presents some technical differences that have been introduced to simplify the analysis. Notice that, to simplify proofs, we assume s_0 to be known: this issue is discussed in the next two sections.

The rest of the paper is organized as follows. In the next section we provide a detailed description of the algorithm and state our main results. The proof strategy for our results is explained in Section 3. Our theoretical results assume full knowledge of problem parameters for ease of proof. In light of this, in Section 4 we discuss a practical implementation of the same idea that does not require prior knowledge of problem parameters, and is data-driven. We also illustrate the method through simulations. The complete proofs are in Sections 5, 7 and 6 respectively.

A preliminary version of this paper appeared in (Deshpande and Montanari, 2014). This paper extends significantly the results in Deshpande and Montanari (2014). In particular, by following an analogous strategy, we improve greatly the bounds obtained by Deshpande and Montanari (2014). This signifantly improves the regimes of (s_0, p, n) on which we can obtain non-trivial results. The proofs follow a similar strategy but are, correspondingly, more careful.

2. Algorithm and main results

We provide a detailed description of the covariance thresholding algorithm for the general model (1) in Table 1. For notational convenience, we shall assume that $2n$ sample vectors are given (instead of n): $\{\mathbf{x}_i\}_{1\le i\le 2n}$.

We start by splitting the data into two halves: $(\mathbf{x}_i)_{1\le i\le n}$ and $(\mathbf{x}_i)_{n<i\le 2n}$ and compute the respective sample covariance matrices $\mathbf{G}$ and $\mathbf{G}'$ respectively. Define $\mathbf{\Sigma}$ to be the population covariance minus identity. i.e.

$$\mathbf{\Sigma} \equiv \sum_{q=1}^{r} \beta_q \mathbf{v}_q \mathbf{v}_q^{\mathsf{T}} \, . \tag{2}$$

Throughout, we let Q_q and s_q denote the support of $\mathbf{v}_q$ and its size respectively, for $q \in \{1, 2, \dots, r\}$. We further let $\mathsf{Q} = \cup_{q=1}^{r} \mathsf{Q}_q$ and $s_0 = |\mathsf{Q}|$. The matrix $\mathbf{G}$ is used, in steps 1 to

Algorithm 1 Covariance Thresholding

1: **Input:** Data $(\mathbf{x}_i)_{1\le i\le 2n}$, parameter $s_0 \in \mathbb{N}$, $\tau, \rho \in \mathbb{R}_{\ge 0}$;
2: Compute the empirical covariance matrices $\mathbf{G} \equiv \sum_{i=1}^{n} \mathbf{x}_i\mathbf{x}_i^{\mathsf{T}}/n$, $\mathbf{G}' \equiv \sum_{i=n+1}^{2n} \mathbf{x}_i\mathbf{x}_i^{\mathsf{T}}/n$;
3: Compute $\widehat{\mathbf{\Sigma}} = \mathbf{G} - \mathrm{I}_p$ (resp. $\widehat{\mathbf{\Sigma}}' = \mathbf{G}' - \mathrm{I}_p$);
4: Compute the matrix $\eta(\widehat{\mathbf{\Sigma}})$ by soft-thresholding the entries of $\widehat{\mathbf{\Sigma}}$:

$$\eta(\widehat{\mathbf{\Sigma}})_{ij} = \begin{cases} \widehat{\mathbf{\Sigma}}_{ij} - \frac{\tau}{\sqrt{n}} & \text{if } \widehat{\mathbf{\Sigma}}_{ij} \ge \tau/\sqrt{n}, \\ 0 & \text{if } -\tau/\sqrt{n} < \widehat{\mathbf{\Sigma}}_{ij} < \tau/\sqrt{n}, \\ \widehat{\mathbf{\Sigma}}_{ij} + \frac{\tau}{\sqrt{n}} & \text{if } \widehat{\mathbf{\Sigma}}_{ij} \le -\tau/\sqrt{n}, \end{cases}$$

5: Let $(\widehat{\mathbf{v}}_q)_{q\le r}$ be the first r eigenvectors of $\eta(\widehat{\mathbf{\Sigma}})$;
6: **Output:** $\widehat{\mathrm{Q}} = \{i \in [p] :\ \exists\, q \text{ s.t. } |(\widehat{\mathbf{\Sigma}}'\widehat{\mathbf{v}}_q)_i| \ge \rho\}$.

4 to obtain a good estimate $\eta(\widehat{\mathbf{\Sigma}})$ for the low rank part of the population covariance $\mathbf{\Sigma}$. The algorithm first computes $\widehat{\mathbf{\Sigma}}$, a centered version of the empirical covariance of the samples as follows:

$$\widehat{\mathbf{\Sigma}} \equiv \mathbf{G} - \mathrm{I}_p, \tag{3}$$

where $\mathbf{G} = n^{-1}\sum_{i\le n} \mathbf{x}_i\mathbf{x}_i^{\mathsf{T}}$ is the sample covariance matrix.

It then obtains the estimate $\eta(\widehat{\mathbf{\Sigma}}) \in \mathbb{R}^{p\times p}$ by *soft thresholding* each entry of $\widehat{\mathbf{\Sigma}}$ at a threshold $\tau/\sqrt{n}$. Explicitly:

$$\big(\eta(\widehat{\mathbf{\Sigma}})\big)_{ij} \equiv \eta\left(\widehat{\mathbf{\Sigma}}_{ij}; \frac{\tau}{\sqrt{n}}\right). \tag{4}$$

Here $\eta : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$ is the soft thresholding function

$$\eta(z;\lambda) = \begin{cases} z - \lambda & \text{if } z \ge \lambda \\ -z + \lambda & \text{if } z \le -\lambda \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

In step 5 of the algorithm, this estimate is used to construct good estimates $\widehat{\mathbf{v}}_q$ of the eigenvectors $\mathbf{v}_q$. Finally, in step 6, these estimates are combined with the (independent) second half of the data $\mathbf{G}'$ to construct estimators $\widehat{\mathrm{Q}}_q$ for the support of the individual eigenvectors $\mathbf{v}_q$. In the first two subsections we will focus on the estimation of $\mathbf{\Sigma}$ and the individual principal components. Our results on support recovery are provided in the final subsection.

2.1 Estimating the population covariance

Our first result bounds the estimation error of the soft thresholding procedure in operator norm.

Theorem 1 *There exist numerical constants $C_1, C_2, C > 0$ such that the following happens. Assume $n > C\log p$, $n > s_0^2$ and let $\tau_* = C_1(\beta \vee 1)\sqrt{\log(p/s_0^2)}$. We keep the thresholding level τ according to*

$$\tau = \begin{cases} \tau_* & \text{when } \tau_* \leq \sqrt{\log p}/2,\ s_0^2 \leq p/e \\ C_2\tau_* & \text{when } \tau_* \geq \sqrt{\log p}/2,\ s_0 \leq p/e \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

. Then with probability $1 - o(1)$:

$$\big\|\eta(\widehat{\mathbf{\Sigma}}) - \mathbf{\Sigma}\big\|_{op} \leq C\sqrt{\frac{s_0^2(\beta^2 \vee 1)}{n}\Big(\log\frac{p}{s_0^2} \vee 1\Big)}\,. \tag{7}$$

At this point, it is useful to compare Theorem 1 with available results in the literature. Classical denoising theory (Donoho and Johnstone, 1994; Johnstone, 2015) provides upper bounds on the estimation error of soft-thresholding. However, estimation error is measured by (element-wise) ℓ_p norm, while here we are interested in operator norm.

Bickel and Levina (2008a,b); Karoui (2008); Cai, Zhang, Zhou, et al. (2010); Cai and Liu (2011) considered the operator norm error of thresholding estimators for structured covariance matrices. Specializing to our case of exact sparsity, the result of Bickel and Levina (2008a) implies that, with high probability:

$$\big\|\eta_H(\widehat{\mathbf{\Sigma}}) - \mathbf{\Sigma}\big\|_{op} \leq C_0\sqrt{\frac{s_0^2\log p}{n}}\,. \tag{8}$$

Here $\eta_H(\cdot,\cdot)$ is the hard-thresholding function: $\eta_H(z) = z\mathbb{I}(|z| \geq \tau/\sqrt{n})$, and the threshold is chosen to be $\tau = C_1\sqrt{\log p}$. Also, $\eta_H(\mathbf{M})$ is the matrix obtained by thresholding the entries of $\mathbf{M}$. In fact, Cai et al. (2012) showed that the rate in (8) is minimax optimal over the class of sparse population covariance matrices, with at most s_0 non-zero entries per row, under the assumption $s_0^2/n \leq C(\log p)^{-3}$.

Theorem 1 ensures consistency under a weaker sparsity condition, viz. $s_0^2/n \to 0$ is sufficient. Also, the resulting rate depends on $\log(p/s_0^2)$ instead of $\log p$. In other words, in order to achieve $\|\eta(\widehat{\mathbf{\Sigma}}) - \mathbf{\Sigma}\|_{op} < \varepsilon$ for a fixed ε, it is sufficient $s_0 \lesssim \varepsilon\sqrt{n}$ as opposed to $s_0 \lesssim \sqrt{n/\log p}$.

Crucially, in this regime for $s_0 = \Theta(\varepsilon\sqrt{n})$, Theorem 1 suggests a threshold of order $\tau = \Theta(\sqrt{\log(1/\varepsilon)})$ as opposed to $\tau = C_1\sqrt{\log p}$ which is used in Bickel and Levina (2008a); Cai et al. (2012). As we will see in Section 3, this regime mathematically more challenging than the one of Bickel and Levina (2008a); Cai et al. (2012). By setting $\tau = C_1\sqrt{\log p}$ for a large enough constant C_1, all the entries of $\widehat{\mathbf{\Sigma}}$ outside the support of $\mathbf{\Sigma}$ are set to 0. In contrast, a large part of our proof is devoted to control the operator norm of the noise part of $\widehat{\mathbf{\Sigma}}$.

2.2 Estimating the principal components

We next turn to the question of estimating the principal components $\mathbf{v}_1, \dots \mathbf{v}_r$. Of course, these are not identifiable if there are degeneracies in the population eigenvalues $\beta_1, \beta_2, \dots, \beta_r$. We thus introduce the following identifiability condition.

A1 The spike strengths $\beta_1 > \beta_2 > \dots \beta_r$ are all *distinct*. We denote by $\beta \equiv \max(\beta_1, \dots, \beta_r)$ and $\beta_{\min} \equiv \min_{q \neq q'}(\beta_1 - \beta_2, \beta_2 - \beta_3, \dots, \beta_r)$. Namely, β is the largest signal strength and $\beta_{\min}$ is the minimum gap.

We measure estimation error through the following loss, defined for $\mathbf{x}, \mathbf{y} \in S^{p-1} \equiv \{\mathbf{v} \in \mathbb{R}^p : \|\mathbf{v}\| = 1\}$:

$$L(\mathbf{x}, \mathbf{y}) \equiv \frac{1}{2} \min_{s \in \{+1,-1\}} \big\|\mathbf{x} - s\,\mathbf{y}\big\|^2 \tag{9}$$

$$= 1 - |\langle \mathbf{x}, \mathbf{y} \rangle| \,. \tag{10}$$

Notice the minimization over the sign $s \in \{+1, -1\}$. This is required because the principal components $\mathbf{v}_1, \dots, \mathbf{v}_r$ are only identifiable up to a sign. Analogous results can obtained for alternate loss functions such as the projection distance:

$$L_p(\mathbf{x}, \mathbf{y}) \equiv \frac{1}{\sqrt{2}} \|\mathbf{x}\mathbf{x}^{\mathsf{T}} - \mathbf{y}\mathbf{y}^{\mathsf{T}}\|_F = \sqrt{1 - \langle \mathbf{x}, \mathbf{y} \rangle^2}. \tag{11}$$

The theorem below is an immediate consequence of Theorem 1. In particular, it uses the guarantee of Theorem 1 to show that the corresponding principal components of $\eta(\widehat{\mathbf{\Sigma}})$ provide good estimates of the principal components $\mathbf{v}_q, 1 \le q \le r$.

Theorem 2 *There exists a numerical constant C such that the following holds. Suppose that Assumption* A1 *holds in addition to the conditions $n > C \log p$, $s_0^2 < n$, and $s_0^2 < p/e$. Set τ as according to Theorem 1, and let $\widehat{\mathbf{v}}_1, \dots, \widehat{\mathbf{v}}_r$ denote the r principal eigenvectors of $\eta(\widehat{\mathbf{\Sigma}}; \tau/\sqrt{n})$. Then, with probability $1 - o(1)$*

$$\max_{q \in [r]} L(\widehat{\mathbf{v}}_q, \mathbf{v}_q) \le \frac{C}{\beta_{\min}^2} \, \frac{s_0^2(\beta^2 \vee 1)}{n} \, \log \frac{p}{s_0^2} \,. \tag{12}$$

Proof Let $\mathbf{\Delta} \equiv \eta(\widehat{\mathbf{\Sigma}}; \tau/\sqrt{n}) - \mathbf{\Sigma}$. By Davis-Kahn sin-theta theorem (Davis and Kahan, 1970), we have, for $\beta_{\min} > \|\mathbf{\Delta}\|_{op}$,

$$L(\widehat{\mathbf{v}}_q, \mathbf{v}_q) \le \frac{1}{2} \left(\frac{\|\mathbf{\Delta}\|_{op}}{\beta_{\min} - \|\mathbf{\Delta}\|_{op}} \right)^2 \,. \tag{13}$$

For $\beta_{\min}^2 > 2C(s_0^2(\beta^2 \vee 1)/n) \, \log(p/s_0^2)$, the claim follows by using Theorem 1. If $\beta_{\min}^2 \le 2C(s_0^2(\beta^2 \vee 1)/n) \, \log(p/s_0^2)$, the claim is obviously true since $L(\widehat{\mathbf{v}}_q, \mathbf{v}_q) \le 1$ always. ■

2.3 Support recovery

Finally, we consider the question of support recovery of the principal components $\mathbf{v}_q$. The second phase of our algorithm aims at estimating union of the supports $\mathsf{Q} = \mathsf{Q}_1 \cup \dots \cup \mathsf{Q}_r$ from the estimated principal components $\widehat{\mathbf{v}}_q$. Note that, although $\widehat{\mathbf{v}}_q$ is not even expected to be sparse, it is easy to see that the largest entries of $\widehat{\mathbf{v}}_q$ should have significant overlap with $\mathrm{supp}(\mathbf{v}_q)$. Step 6 of the algorithm exploit this property to construct a consistent estimator $\widehat{\mathsf{Q}}_q$ of the support of the spike $\mathbf{v}_q$.

We will require the following assumption to ensure support recovery.

A2 There exist constants $\theta, \gamma > 0$ such that the following holds. The non-zero entries of the spikes satisfy $|v_{q,i}| \geq \theta/\sqrt{s_0}$ for all $i \in \mathsf{Q}_q$. Further, for any q, q' $|v_{q,i}/v_{q',i}| \leq \gamma$ for every $i \in \mathsf{Q}_q \cap \mathsf{Q}_{q'}$. Without loss of generality, we will assume $\gamma \geq 1$.

Theorem 3 *Assume the spiked covariance model of Eq.* (1) *satisfying assumptions* A1 *and* A2*, and further $n > C \log p$, $s_0^2 < n$, and $s_0^2 < p/e$ for C a large enough numerical constant. Consider the Covariance Thresholding algorithm of Table 1, with τ as in Theorem 1 $\rho = \beta_{\min}\theta/(2\sqrt{s_0})$.*

Then there exists $K_0 = K_0(\theta, \gamma, \beta, \beta_{\min})$ such that, if

$$n \geq K_0 s_0^2 r \log \frac{p}{s_0^2} \tag{14}$$

then the algorithm recovers the union of supports of $\mathbf{v}_q$ with probability $1 - o(1)$ (i.e. we have $\widehat{\mathsf{Q}} = \mathsf{Q}$).

The proof in Section 7 also provides an explicit expression for the constant K_0.

Remark 4 *In Assumption* A2*, the requirement on the minimum size of $|v_{q,i}|$ is standard in support recovery literature (see, e.g. Wainwright, 2009; Meinshausen and Bühlmann, 2006). Additionally, however, we require that when the supports of $\mathbf{v}_q, \mathbf{v}_{q'}$ overlap, they are of the same order, quantified by the parameter γ. Relaxing this condition is a potential direction for future work.*

Remark 5 *Recovering the signed supports $\mathsf{Q}_{q,+} = \{i \in [p] : v_{q,i} > 0\}$ and $\mathsf{Q}_{q,-} = \{i \in [p] : v_{q,i} < 0\}$, up to a sign flip, is possible using the same technique as recovering the supports* $\mathrm{supp}(\mathbf{v}_q)$ *above, and poses no additional difficulty.*

3. Algorithm intuition and proof strategy

For the purposes of exposition, throughout this section, we will assume that $r = 1$ and drop the corresponding subscript q.

Denoting by $\mathbf{X} \in \mathbb{R}^{n \times p}$ the matrix with rows $\mathbf{x}_1, \dots \mathbf{x}_n$, by $\mathbf{Z} \in \mathbb{R}^{n \times p}$ the matrix with rows $\mathbf{z}_1, \dots \mathbf{z}_n$, and letting $\mathbf{u} = (u_1, u_2, \dots, u_n)$, the model (1) can be rewritten as

$$\mathbf{X} = \sqrt{\beta}\, \mathbf{u}\, \mathbf{v}^{\mathsf{T}} + \mathbf{Z}\,. \tag{15}$$

Recall that $\widehat{\mathbf{\Sigma}} = n^{-1}\mathbf{X}^{\mathsf{T}}\mathbf{X} - \mathrm{I}_p = \mathbf{G} - \mathrm{I}_p$. For $\beta > \sqrt{p/n}$, the principal eigenvector of $\mathbf{G}$, and hence of $\widehat{\mathbf{\Sigma}}$ is positively correlated with $\mathbf{v}$, i.e. $|\langle \widehat{\mathbf{v}}_1(\widehat{\mathbf{\Sigma}}), \mathbf{v}\rangle|$ is bounded away from zero. However, for $\beta < \sqrt{p/n}$, the noise component in $\widehat{\mathbf{\Sigma}}$ dominates and the two vectors become asymptotically orthogonal, i.e. for instance $\lim_{n\to\infty} |\langle \widehat{\mathbf{v}}_1(\widehat{\mathbf{\Sigma}}), \mathbf{v}\rangle| = 0$. In order to reduce the noise level, we must exploit the sparsity of the spike $\mathbf{v}$.

Now, letting $\beta' \equiv \beta \|\mathbf{u}\|^2/n \approx \beta$, and $\mathbf{w} \equiv \sqrt{\beta}\mathbf{Z}^{\mathsf{T}}\mathbf{u}/n$, we can rewrite $\widehat{\mathbf{\Sigma}}$ as

$$\widehat{\mathbf{\Sigma}} = \beta'\, \mathbf{v}\mathbf{v}^{\mathsf{T}} + \mathbf{v}\,\mathbf{w}^{\mathsf{T}} + \mathbf{w}\,\mathbf{v}^{\mathsf{T}} + \frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} \;\; - \mathrm{I}_p,. \tag{16}$$

For a moment, let us neglect the cross terms $(\mathbf{v}\mathbf{w}^{\mathsf{T}} + \mathbf{w}\mathbf{v}^{\mathsf{T}})$. The 'signal' component $\beta' \, \mathbf{v}\mathbf{v}^{\mathsf{T}}$ is sparse with s_0^2 entries of magnitude $\beta'\theta^2/s_0$, which (in the regime of interest $s_0 = \sqrt{n}/C$) is equivalent to $C\theta^2\beta/\sqrt{n}$. The 'noise' component $\mathbf{Z}^{\mathsf{T}}\mathbf{Z}/n - \mathrm{I}_p$ is dense with entries of order $1/\sqrt{n}$. Assuming $s_0/\sqrt{n} < c$ for some small constant c, it should be possible to remove most of the noise by thresholding the entries at level of order $1/\sqrt{n}$. For technical reasons, we use the soft thresholding function $\eta : \mathbb{R} \times \mathbb{R}_{\geq 0} \to \mathbb{R}$, $\eta(z;\tau) = \mathrm{sgn}(z)(|z| - \tau)_+$. We will omit the second argument from $\eta(\cdot;\cdot)$ wherever it is clear from context.

Consider again the decomposition (16). Since the soft thresholding function $\eta(z;\tau/\sqrt{n})$ is affine when $z \gg \tau/\sqrt{n}$, we would expect that the following decomposition holds approximately (for instance, in operator norm):

$$\eta(\widehat{\mathbf{\Sigma}}) \approx \eta\left(\beta'\mathbf{v}\mathbf{v}^{\mathsf{T}}\right) + \eta\left(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}_p\right). \tag{17}$$

Since $\beta' \approx \beta$ and each entry of $\mathbf{v}\mathbf{v}^{\mathsf{T}}$ has magnitude at least θ^2/s_0, the first term is still approximately rank one, with

$$\left\|\eta\left(\beta'\mathbf{v}\mathbf{v}^{\mathsf{T}}\right) - \beta\mathbf{v}\mathbf{v}^{\mathsf{T}}\right\|_{op} \leq \frac{s_0\tau}{\sqrt{n}}. \tag{18}$$

This is straightforward to see since soft thresholding introduces a maximum bias of $\tau/\sqrt{n}$ per entry of the matrix, while the factor s_0 comes due to the support size of $\mathbf{v}\mathbf{v}^{\mathsf{T}}$ (see Proposition 14 below for a rigorous argument).

The main technical challenge now is to control the operator norm of the perturbation $\eta(\mathbf{Z}^{\mathsf{T}}\mathbf{Z}/n - \mathrm{I}_p)$. We know that $\eta(\mathbf{Z}^{\mathsf{T}}\mathbf{Z}/n - \mathrm{I}_p)$ has entries of variance $\delta(\tau)/n$, for $\delta(\tau) \approx \exp(-c\tau^2)$. If entries were independent with mild tail conditions, this would imply –with high probability–

$$\left\|\eta\left(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}_p\right)\right\|_{op} \lesssim C\delta(\tau)\sqrt{\frac{p}{n}} = C\exp(-c\tau^2)\sqrt{\frac{p}{n}}, \tag{19}$$

for some constant C. Combining the bias bound from Eq. (18) and the heuristic decomposition of Eq. (19) with the decomposition (17) results in the bound

$$\left\|\eta(\widehat{\mathbf{\Sigma}}) - \beta\mathbf{v}\mathbf{v}^{\mathsf{T}}\right\|_{op} \leq \frac{s_0\tau}{\sqrt{n}} + C\exp(-c\tau^2)\sqrt{\frac{p}{n}}. \tag{20}$$

Our analysis formalizes this argument and shows that such a bound is correct when $p < n$.

The matrix $\eta\big(\mathbf{Z}^{\mathsf{T}}\mathbf{Z}/n - \mathrm{I}_p\big)$ is a special case of so-called inner-product kernel random matrices, which have attracted recent interest within probability theory (see El Karoui, 2010a,b; Cheng and Singer, 2013; Fan and Montanari, 2015). The basic object of study in this line of work is a matrix $\mathbf{M} \in \mathbb{R}^{p\times p}$ of the type:

$$M_{ij} = f_n\left(\frac{\langle \tilde{\mathbf{z}}_i, \tilde{\mathbf{z}}_j\rangle}{n} - \mathbb{I}(i=j)\right). \tag{21}$$

In other words, $f_n : \mathbb{R} \to \mathbb{R}$ is a kernel function and is applied entry-wise to the matrix $\mathbf{Z}^{\mathsf{T}}\mathbf{Z}/n - \mathrm{I}_p$, with $\mathbf{Z}$ a matrix with independent standard normal entries as above and $\tilde{\mathbf{z}}_i \in \mathbb{R}^n$ are the columns of $\mathbf{Z}$.

The key technical challenge in our proof is the analysis of the operator norm of such matrices, when f_n is the soft-thresholding function, with threshold of order $1/\sqrt{n}$. Earlier results are not general enough to cover this case. El Karoui (2010a,b) provide conditions under which the spectrum of $f_n(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)$ is close to a rescaling of the spectrum of $(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)$. We are interested instead in a different regime in which the spectrum of $f_n(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)$ is very different from the one of $(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)$. Cheng and Singer (2013) consider n-dependent kernels, but their results are asymptotic and concern the weak limit of the empirical spectral distribution of $f_n(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)$. This does not yield an upper bound on the spectral norm of $f_n(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)$. Finally, Fan and Montanari (2015) consider the spectral norm of kernel random matrices for smooth kernels f, only in the proportional regime $n/p \to c \in (0, \infty)$.

Our approach to proving Theorem 1 follows instead the ε-net method: we develop high probability bounds on the maximum Rayleigh quotient:

$$\max_{\mathbf{y}\in\mathbb{S}^{p-1}} \langle \mathbf{y}, \eta(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)\mathbf{y}\rangle = \max_{\mathbf{y}\in\mathbb{S}^{p-1}} \sum_{i,j} \eta\left(\frac{\langle \tilde{\mathbf{z}}_i, \tilde{\mathbf{z}}_j\rangle}{n}; \frac{\tau}{\sqrt{n}}\right) y_i y_j, \tag{22}$$

by discretizing $\mathbb{S}^{p-1} = \{\mathbf{y} \in \mathbb{R}^p : \|\mathbf{y}\| = 1\}$, the unit sphere in p dimensions. For a fixed $\mathbf{y}$, the Rayleigh quotient $\langle \mathbf{y}, \eta(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)\mathbf{y}\rangle$ is a (complicated) function of the underlying Gaussian random variables $\mathbf{Z}$. One might hope that it is Lipschitz continuous with some Lipschitz constant $B = B(n, p, \tau, \mathbf{y})$, thereby implying, by Gaussian isoperimetry (Ledoux, 2005), that it concentrates to the scale B around its expectation (i.e. 0). Then, by a standard union bound argument over a discretization of the sphere, one would obtain that the operator norm of $\eta\big(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p\big)$ is typically no more than $\sqrt{p}\,\sup_{\mathbf{y}\in\mathbb{S}^{p-1}} B(n, p, \tau, \mathbf{y})$.

Unfortunately, this turns out not to be true over the whole space of $\mathbf{Z}$, i.e. the Rayleigh quotient is not Lipschitz continuous in the underlying Gaussian variables $\mathbf{Z}$. Our approach, instead, shows that for *typical* values of $\mathbf{Z}$, we can control the gradient of $\langle \mathbf{y}, \eta(\mathbf{Z}^\mathsf{T}\mathbf{Z}/n - \mathrm{I}_p)\mathbf{y}\rangle$ with respect to $\mathbf{Z}$, and extract the required concentration only from such local information of the function. This is formalized in our concentration lemma 9, which we apply extensively while proving Theorem 1. This lemma is a signficantly improved version of the analogous result in Deshpande and Montanari (2014).

4. Practical aspects and empirical results

Specializing to the rank one case, Theorems 2 and 3 show that Covariance Thresholding succeeds with high probability for a number of samples $n \gtrsim s_0^2$, while Diagonal Thresholding requires $n \gtrsim s_0^2 \log p$. The reader might wonder whether eliminating the $\log p$ factor has any practical relevance or is a purely conceptual improvement. Figure 1 presents simulations on synthetic data under the strictly sparse model, and the Covariance Thresholding algorithm of Table 1, used in the proof of Theorem 3. The objective is to check whether the $\log p$ factor has an impact at moderate p. We compare this with Diagonal Thresholding.

We plot the empirical success probability as a function of $s_0/\sqrt{n}$ for several values of p, with $p = n$. The empirical success probability was computed by using 100 independent instances of the problem. A few observations are of interest: (*i*) Covariance Thresholding appears to have a significantly larger success probability in the 'difficult' regime where Diagonal Thresholding starts to fail; (*ii*) The curves for Diagonal Thresholding appear to

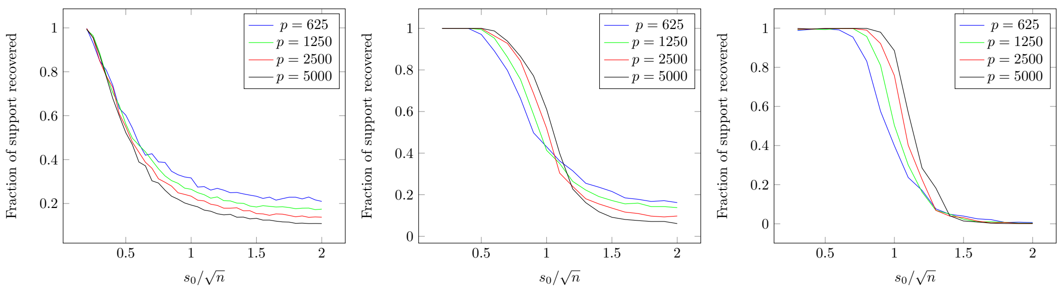

Figure 1: The support recovery phase transitions for Diagonal Thresholding (left) and Covariance Thresholding (center) and the data-driven version of Section 4 (right). For Covariance Thresholding, the fraction of support recovered correctly *increases* monotonically with p, as long as $s_0 \le c\sqrt{n}$ with $c \approx 1.1$. Further, it appears to converge to one throughout this region. For Diagonal Thresholding, the fraction of support recovered correctly *decreases* monotonically with p for all s_0 of order $\sqrt{n}$. This confirms that Covariance Thresholding (with or without knowledge of the support size s_0) succeeds with high probability for $s_0 \le c\sqrt{n}$, while Diagonal Thresholding requires a significantly sparser principal component.

decrease monotonically with p indicating that s_0 proportional to $\sqrt{n}$ is not the right scaling for this algorithm (as is known from theory); (*iii*) In contrast, the curves for Covariance Thresholding become steeper for larger p, and, in particular, the success probability increases with p for $s_0 \le 1.1\sqrt{n}$. This indicates a sharp threshold for $s_0 = \text{const} \cdot \sqrt{n}$, as suggested by our theory.

In terms of practical applicability, our algorithm in Table 1 has the shortcomings of requiring knowledge of problem parameters s_0, β, θ. Furthermore, the thresholds ρ, τ suggested by theory need not be optimal. We next describe a principled approach to estimating (where possible) the parameters of interest and running the algorithm in a purely data-dependent manner. Assume the following model, for $i \in [n]$

$$\mathbf{x}_i = \boldsymbol{\mu} + \sum_q \sqrt{\beta_q} u_{q,i} \mathbf{v}_q + \sigma \mathbf{z}_i,$$

where $\boldsymbol{\mu} \in \mathbb{R}^p$ is a fixed mean vector, $u_{q,i}$ have mean 0 and variance 1, and $\mathbf{z}_i$ have mean 0 and covariance I_p. Note that our focus in this section is not on rigorous analysis, but instead to demonstrate a principled approach to applying covariance thresholding in practice. We proceed as follows:

Estimating $\boldsymbol{\mu}$, σ: We let $\widehat{\boldsymbol{\mu}} = \sum_{i=1}^n \mathbf{x}_i/n$ be the empirical mean estimate for μ. Further letting $\overline{\mathbf{X}} = \mathbf{X} - \mathbf{1}\widehat{\boldsymbol{\mu}}^{\mathsf{T}}$ we see that $pn - (\sum_q k_q)n \approx pn$ entries of $\overline{\mathbf{X}}$ are mean 0 and variance σ^2. We let $\widehat{\sigma} = \mathrm{MAD}(\overline{\mathbf{X}})/\nu$ where $\mathrm{MAD}(\,\cdot\,)$ denotes the median absolute deviation of the entries of the matrix in the argument, and ν is a constant scale factor. Guided by the Gaussian case, we take $\nu = \Phi^{-1}(3/4) \approx 0.6745$.

Choosing τ: Although in the statement of the theorem, our choice of τ depends on the SNR β/σ^2, it is reasonable to instead threshold 'at the noise level', as follows. The noise component of entry i, j of the sample covariance (ignoring lower order terms) is given by $\sigma^2\langle \mathbf{z}_i, \mathbf{z}_j\rangle/n$. By the central limit theorem, $\langle \mathbf{z}_i, \mathbf{z}_j\rangle/\sqrt{n} \overset{\mathrm{d}}{\Rightarrow} \mathsf{N}(0, 1)$. Consequently, $\sigma^2\langle \mathbf{z}_i, \mathbf{z}_j\rangle/n \approx \mathsf{N}(0, \sigma^4/n)$, and we need to choose the (rescaled) threshold proportional to $\sqrt{\sigma^4} = \sigma^2$. Using previous estimates, we let $\tau = \nu' \cdot \widehat{\sigma}^2$ for a constant ν'. In simulations, a choice $3 \lesssim \nu' \lesssim 4$ appears to work well.

Estimating r: We define $\widehat{\mathbf{\Sigma}} = \overline{\mathbf{X}}^{\mathsf{T}}\overline{\mathbf{X}}/n - \widehat{\sigma}^2\mathrm{I}_p$ and soft threshold it to get $\eta(\widehat{\mathbf{\Sigma}})$ using τ as above. Our proof of Theorem 2 relies on the fact that $\eta(\widehat{\mathbf{\Sigma}})$ has r eigenvalues that are separated from the bulk of the spectrum. Hence, we estimate r using $\widehat{r}$: the number of eigenvalues separated from the bulk in $\eta(\widehat{\mathbf{\Sigma}})$. The edge of the spectrum can be computed numerically using the Stieltjes transform method as in Cheng and Singer (2013).

Estimating $\mathbf{v}_q$: Let $\widehat{\mathbf{v}}_q$ denote the q^{th} eigenvector of $\eta(\widehat{\mathbf{\Sigma}})$. Our theoretical analysis indicates that $\widehat{\mathbf{v}}_q$ is expected to be close to $\mathbf{v}_q$. In order to denoise $\widehat{\mathbf{v}}_q$, we assume $\widehat{\mathbf{\Sigma}}\widehat{\mathbf{v}}_q \approx (1-\delta)\mathbf{v}_q + \boldsymbol{\varepsilon}_q$, where $\boldsymbol{\varepsilon}_q$ is additive random noise (perhaps with some sparse corruptions). We then threshold $\widehat{\mathbf{\Sigma}}\mathbf{v}_q$ 'at the noise level' to recover a better estimate of $\mathbf{v}_q$. To do this, we estimate the standard deviation of the "noise" $\boldsymbol{\varepsilon}$ by $\widehat{\sigma_{\boldsymbol{\varepsilon}}} = \mathrm{MAD}(\widehat{\mathbf{v}}_q)/\nu$. Here we set –again guided by the Gaussian heuristic– $\nu \approx 0.6745$. Since $\mathbf{v}_q$ is sparse, this procedure returns a good estimate for the size of the noise deviation. We let $\widehat{\mathbf{v}}'_q$ denote the vector obtained by hard thresholding $\widehat{\mathbf{v}}_q$: set $\widehat{\mathbf{v}}'_i = \widehat{\mathbf{v}}_{q,i}$ if $|\widehat{v}_{q,i}| \geq \nu'\widehat{\sigma}_{\boldsymbol{\varepsilon}_q}$ and 0 otherwise. We then let $\widehat{\mathbf{v}}^*_q = \widehat{\mathbf{v}}'_q/\|\widehat{\mathbf{v}}'_q\|$ and return $\widehat{\mathbf{v}}^*_q$ as our estimate for $\mathbf{v}_q$.

Note that –while different in several respects– this empirical approach shares the same philosophy of the algorithm in Table 1. On the other hand, the data-driven algorithm presented in this section is less straightforward to analyze, a task that we defer to future work.

Figure 1 also shows results of a support recovery experiment using the 'data-driven' version of this section. Covariance thresholding in this form also appears to work for supports of size $s_0 \leq \mathrm{const}\sqrt{n}$. Figure 2 shows the performance of vanilla PCA, Diagonal Thresholding and Covariance Thresholding on the "Three Peak" example of Johnstone and Lu (2004). This signal is sparse in the wavelet domain and the simulations employ the data-driven version of covariance thresholding. A similar experiment with the "box" example of Johnstone and Lu is provided in Figure 3. These experiments demonstrate that, while for large values of n both Diagonal Thresholding and Covariance Thresholding perform well, the latter appears superior for smaller values of n.

5. Proof preliminaries

In this section we review some notation and preliminary facts that we will use throughout the paper.

5.1 Notation

We let $[m] = \{1, 2, \dots, m\}$ denote the set of first m integers. We will represent vectors using boldface lower case letters, e.g. $\mathbf{u}, \mathbf{v}, \mathbf{x}$. The entries of a vector $\mathbf{u} \in \mathbb{R}^n$ will be represented

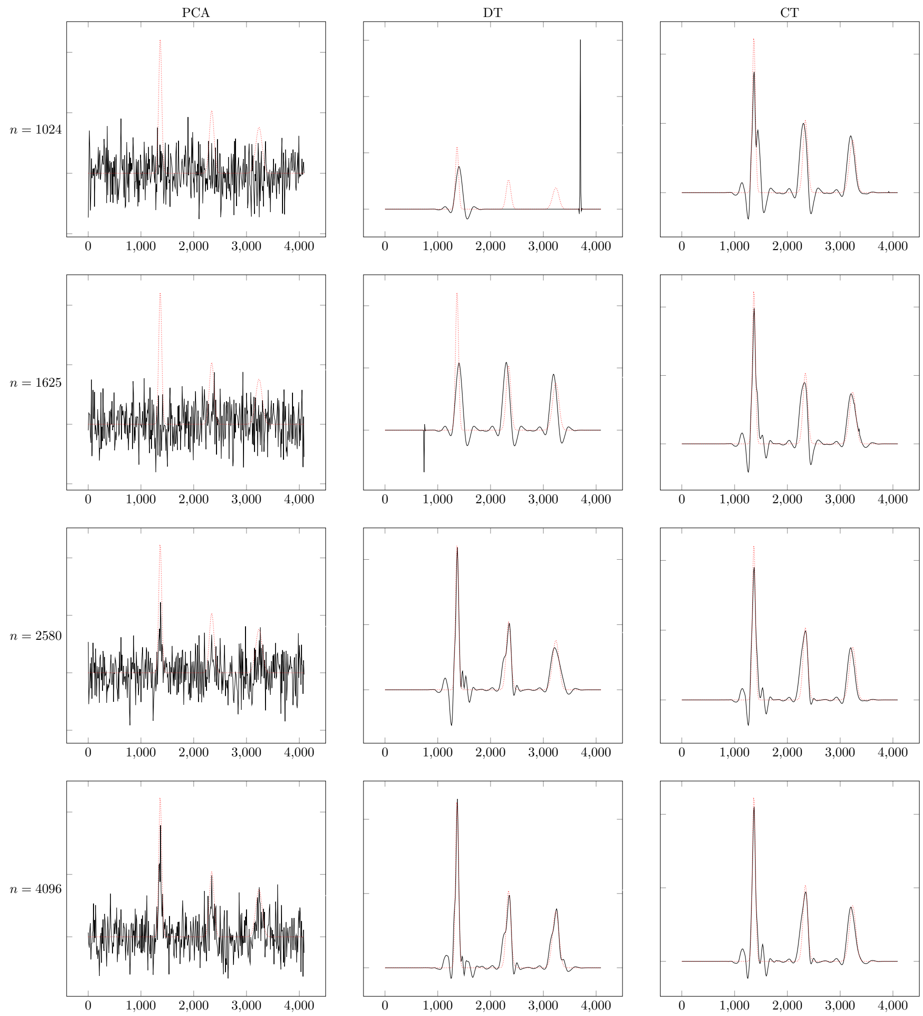

Figure 2: The results of Simple PCA, Diagonal Thresholding and Covariance Thresholding (respectively) for the "Three Peak" example of Johnstone and Lu (2009) (see Figure 1 of the paper). The signal is sparse in the 'Symmlet 8' basis. We use $\beta = 1.4, p = 4096$, and the rows correspond to sample sizes $n = 1024, 1625, 2580, 4096$ respectively. Parameters for Covariance Thresholding are chosen as in Section 4, with $\nu' = 4.5$. Parameters for Diagonal Thresholding are from Johnstone and Lu (2009). On each curve, we superpose the clean signal (dotted).

by $u_i, i \in [n]$. Matrices are represented using boldface upper case letters e.g. $\mathbf{A}, \mathbf{X}$. The entries of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ are represented by $\mathbf{A}_{ij}$ for $i \in [m], j \in [n]$. Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we generically let $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m$ denote its rows, and $\tilde{\mathbf{a}}_1, \tilde{\mathbf{a}}_2, \dots, \tilde{\mathbf{a}}_n$ its columns.

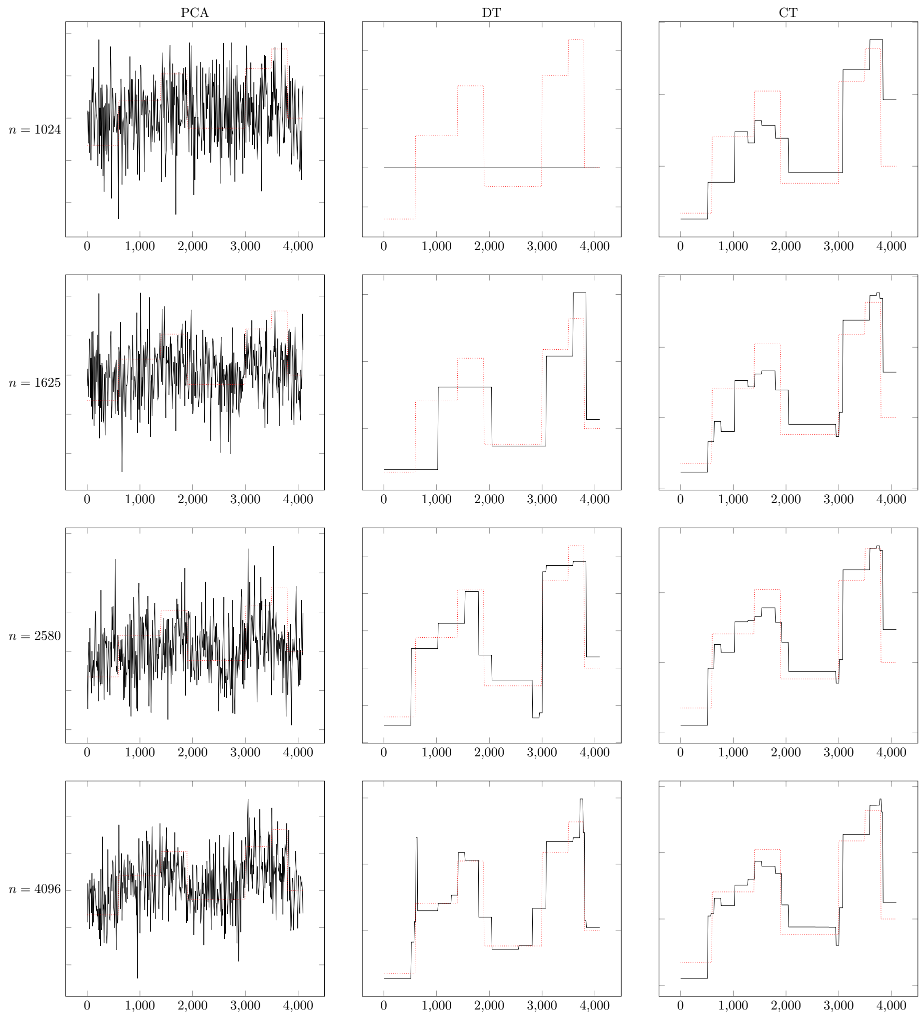

Figure 3: The results of Simple PCA, Diagonal Thresholding and Covariance Thresholding (respectively) for a synthetic block-constant function (which is sparse in the Haar wavelet basis). We use $\beta = 1.4, p = 4096$, and the rows correspond to sample sizes $n = 1024, 1625, 2580, 4096$ respectively. Parameters for Covariance Thresholding are chosen as in Section 4, with $\nu' = 4.5$. Parameters for Diagonal Thresholding are from Johnstone and Lu (2009). On each curve, we superpose the clean signal (dotted).

For $E \subseteq [m] \times [n]$, we define the projector operator $\mathcal{P}_E : \mathbb{R}^{m\times n} \to \mathbb{R}^{m\times n}$ by letting $\mathcal{P}_E(\mathbf{A})$ be the matrix with entries

$$\mathcal{P}_E(\mathbf{A})_{ij} = \begin{cases} \mathbf{A}_{ij} & \text{if } (i,j) \in E, \\ 0 & \text{otherwise.} \end{cases} \tag{23}$$

For a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, and a set $E \subseteq [n]$, we define its column restriction $\mathbf{A}_E \in \mathbb{R}^{m\times n}$ to be the matrix obtained by setting to 0 columns outside E:

$$(\mathbf{A}_E)_{ij} = \begin{cases} \mathbf{A}_{ij} & \text{if } j \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly $\mathbf{y}_E$ is obtained from $\mathbf{y}$ by setting to zero all indices outside E. The operator norm of a matrix $\mathbf{A}$ is denoted by $\|\mathbf{A}\|$ (or $\|\mathbf{A}\|_{op}$) and its Frobenius norm by $\|\mathbf{A}\|_F$. We write $\|\mathbf{x}\|$ for the standard ℓ_2 norm of a vector $\mathbf{x}$. Other vector norms such as ℓ_1 or ℓ_∞ are denoted with appropriate subscripts.

We let Q_q denotes the support of the q^{th} spike $\mathbf{v}_q$. Also, we denote the union of the supports of $\mathbf{v}_q$ by $\mathsf{Q} = \cup_q \mathsf{Q}_q$. The complement of a set $E \in [n]$ is denoted by E^c.

We write $\eta(\cdot;\cdot)$ for the soft-thresholding function. By $\partial\eta(\cdot;\tau)$ we denote the derivative of $\eta(\cdot;\tau)$ with respect to the *first* argument, which exists Lebesgue almost everywhere. To simplify the notation, we omit the second argument when it is understood from context.

For a random variable Z and a measurable set $\mathcal{A}$ we write $\mathbb{E}\{Z;\mathcal{A}\}$ to denote $\mathbb{E}\{Z\mathbb{I}(Z \in \mathcal{A})\}$, the expectation of Z constrained to the event $\mathcal{A}$.

In the statements of our results, consider the limit of large p and large n with certain conditions on p, n (as in Theorem 2). This limit will be referred to either as "n large enough" or "p large enough" where the phrase "large enough" indicates dependence of p (and thereby n) on specific problem parameters.

The Gaussian distribution function will be denoted by $\Phi(x) = \int_{-\infty}^{x} e^{-t^2/2}\,\mathrm{d}t/\sqrt{2\pi}$.

5.2 Preliminary facts

Let $\mathbb{S}^{N-1}$ denote the unit sphere in N dimensions, i.e. $\mathbb{S}^{N-1} = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\| = 1\}$. We use the following definition (see Vershynin, 2012, Definition 5.2) of the ε-net of a set $X \subseteq \mathbb{R}^n$:

Definition 6 (Nets, Covering numbers) *A subset $T^\varepsilon(X) \subseteq X$ is called an ε-net of X if every point in X may be approximated by one in $T^\varepsilon(X)$ with error at most ε. More precisely:*

$$\forall x \in X, \qquad \inf_{y\in T^\varepsilon(X)} \|x - y\| \le \varepsilon.$$

The minimum cardinality of an ε-net of X, if finite, is called its covering number.

The following two facts are useful while using ε-nets to bound the spectral norm of a matrix. For proofs, we refer the reader to (see Vershynin, 2012, Lemmas 5.2, 5.4).

Lemma 7 *Let S^{n-1} be the unit sphere in n dimensions. Then there exists an ε-net of S^{n-1}, $T^\varepsilon(S^{n-1})$ satisfying:*

$$|T^\varepsilon(S^{n-1})| \le \left(1 + \frac{2}{\varepsilon}\right)^n.$$

Lemma 8 *Let $\mathbf{A} \in \mathbb{R}^{n\times n}$ be a symmetric matrix. Then, there exists $\mathbf{x} \in T^{\varepsilon}(S^{n-1})$ such that*

$$|\langle \mathbf{x}, \mathbf{A}\mathbf{x}\rangle| \geq (1-2\varepsilon)\|\mathbf{A}\|. \tag{24}$$

Proof Firstly, we have $\|A\| = \max_{\mathbf{x}\in\mathbb{S}^{n-1}}|\langle \mathbf{x}, \mathbf{A}\mathbf{x}\rangle| = \max_{\mathbf{x}\in\mathbb{S}^{n-1}}\|\mathbf{A}\mathbf{x}\|$. Let $\mathbf{x}_*$ be the maximizer (which exists as $\mathbb{S}^{n-1}$ is compact and $|\langle \mathbf{x}, \mathbf{A}\mathbf{x}\rangle|$ is continuous in $\mathbf{x}$). Choose $\mathbf{x} \in T_n^{\varepsilon}$ so that $\|\mathbf{x}-\mathbf{x}_*\| \leq \varepsilon$. Then:

$$\langle \mathbf{x}, \mathbf{A}\mathbf{x}\rangle = \langle \mathbf{x}-\mathbf{x}_*, \mathbf{A}(\mathbf{x}+\mathbf{x}_*)\rangle + \langle \mathbf{x}_*, \mathbf{A}\mathbf{x}_*\rangle\,. \tag{25}$$

The lemma then follows as $|\langle \mathbf{x}, \mathbf{A}(\mathbf{x}-\mathbf{x}_*)\rangle| \leq \|\mathbf{x}+\mathbf{x}_*\|\|\mathbf{A}\|\|\mathbf{x}-\mathbf{x}_*\| \leq 2\varepsilon\|\mathbf{A}\|$. ∎

Throughout the paper we will denote by T_N^{ε} an ε-net on the unit sphere $\mathbb{S}^{N-1}$ that satisfies Lemma 7. For a subset of indices $\mathsf{S} \subset [N]$ we denote by $T_N^{\varepsilon}(\mathsf{S})$ the natural isometric embedding of $T_{\mathsf{S}}^{\varepsilon}$ in $\mathbb{S}^{N-1}$.

We now state a general concentration lemma. This will be our basic tool to establish Theorem 2, and thereby Theorem 3.

Lemma 9 *Let $\mathbf{z} \sim \mathsf{N}(0, \mathrm{I}_N)$ be vector of N i.i.d. standard normal variables. Suppose S is a finite set and we have functions $F_s : \mathbb{R}^N \to \mathbb{R}$ for every $s \in S$. Assume $\mathcal{G} \in \mathbb{R}^N \times \mathbb{R}^N$ is a Borel set such that for Lebesgue-almost every $(\mathbf{x}, \mathbf{y}) \in \mathcal{G}$:*

$$\max_{s\in S}\max_{t\in[0,1]}\|\nabla F_s(\sqrt{t}\mathbf{x} + \sqrt{1-t}\mathbf{y})\| \leq L\,. \tag{26}$$

Then, for any $\Delta > 0$:

$$\mathbb{P}\Big\{\max_{s\in S}|F_s(\mathbf{z}) - \mathbb{E}F_s(\mathbf{z})| \geq \Delta\Big\} \leq C|S|\,\exp\Big(-\frac{\Delta^2}{CL^2}\Big) + \frac{C}{\Delta^2}\mathbb{E}\Big\{\max_{s\in S}\big[(F_s(\mathbf{z}) - F_s(\mathbf{z}'))^2\big];\mathcal{G}^c\Big\}. \tag{27}$$

Here $\mathbf{z}'$ is an independent copy of $\mathbf{z}$.

Proof We use the Maurey-Pisier method along with symmetrization. By centering, assume that $\mathbb{E}F_s(\mathbf{z}) = 0$ for all $s \in S$. Further, by including the functions $-F_s$ in the set S (at most doubling its size), it suffices to prove the one-sided version of the inequality:

$$\mathbb{P}\{\max_{s\in S}F_s(\mathbf{z}) \geq \Delta\} \leq C|S|\,\exp\Big(-\frac{\Delta^2}{CL^2}\Big) + \frac{C}{\Delta^2}\mathbb{E}\{\max_s(F_s(\mathbf{z}) - F_s(\mathbf{z}'))^2;\mathcal{G}^c\}\,. \tag{28}$$

We first implement the symmetrization. Note that:

$$\{\mathbf{x} : \max_s F_s(\mathbf{x}) \geq \Delta\} \subseteq \{\mathbf{x} : \max_{x\in\mathbb{R}, s\in S}[2xF_s(\mathbf{x}) - x^2] \geq \Delta^2\} \tag{29}$$

$$\{\mathbf{x}, \mathbf{y} : \max_s[F_s(\mathbf{x}) - F_s(\mathbf{y})] \geq \Delta\} \subseteq \{\mathbf{x}, \mathbf{y} : \max_{x\in\mathbb{R}, s\in S}[2x(F_s(\mathbf{x}) - F_s(\mathbf{y})) - x^2] \geq \Delta^2\}. \tag{30}$$

Furthermore, by centering, $F_s(\mathbf{z}) = \mathbb{E}\{F_s(\mathbf{z}) - F_s(\mathbf{z}')|\mathbf{z}\}$. Hence for any non-decreasing convex function $\phi(z)$:

$$\mathbb{E}\Big\{\phi\big(\max_{x,s}[2xF_s(\mathbf{z}) - x^2]\big)\Big\} \le \mathbb{E}\Big\{\phi\Big(\max_{x,s}\big[\mathbb{E}\{2xF_s(\mathbf{z}) - 2xF_s(\mathbf{z}') - x^2|\mathbf{z}\}\big]\Big)\Big\} \tag{31}$$

$$\stackrel{(a)}{\le} \mathbb{E}\Big\{\phi\Big(\mathbb{E}\{\max_{x,s}[2x(F_s(\mathbf{z}) - F_s(\mathbf{z}')) - x^2]|\mathbf{z}\}\Big)\Big\} \tag{32}$$

$$\stackrel{(b)}{\le} \mathbb{E}\Big\{\phi\big(\max_{x,s}[2x(F_s(\mathbf{z}) - F_s(\mathbf{z}')) - x^2]\big)\Big\}. \tag{33}$$

Here we use Jensen's inequality with the monotonicity of $\phi(\cdot)$ to obtain (a) and with the convexity of $\phi(\cdot)$ to obtain (b).

Now we choose $\phi(z) = (z-a)_+$, for $a = \Delta^2/2$.

$$\mathbb{P}\{\max_s F_s(\mathbf{z}) \ge \Delta\} \le \mathbb{P}\big\{\max_{x,s}[2xF_s(\mathbf{z}) - x^2] \ge \Delta^2\big\} \tag{34}$$

$$\stackrel{(a)}{\le} \phi(\Delta^2)^{-1}\mathbb{E}\Big\{\phi\big(\max_{x,s}[2xF_s(\mathbf{z}) - x^2]\big)\Big\} \tag{35}$$

$$\stackrel{(b)}{\le} \phi(\Delta^2)^{-1}\mathbb{E}\Big\{\phi\big(\max_{x,s}[2x(F_s(\mathbf{z}) - F_s(\mathbf{z}')) - x^2]\big)\Big\} \tag{36}$$

$$= \phi(\Delta^2)^{-1}\mathbb{E}\Big\{\phi\big(\max_s[(F_s(\mathbf{z}) - F_s(\mathbf{z}'))^2]\big)\Big\} \tag{37}$$

$$\begin{aligned} &= \phi(\Delta^2)^{-1}\Bigg(\mathbb{E}\Big\{\phi\big(\max_s[(F_s(\mathbf{z}) - F_s(\mathbf{z}'))^2]\big); \mathcal{G}\Big\} \\ &\qquad + \mathbb{E}\Big(\big\{\phi(\max_s[(F_s(\mathbf{z}) - F_s(\mathbf{z}'))^2]; \mathcal{G}^c\big\}\Big)\Bigg). \end{aligned} \tag{38}$$

Here (a) is Markov's inequality, and (b) is the symmetrization bound Eq. (33), where we use the fact that $\phi(z) = (z-a)_+$ is non-decreasing and convex in z.

At this point, it is easy to see that the lemma follows if we are able to control the first term in Eq. (38). We establish this via the Maurey-Pisier method. Define the path $\mathbf{z}(\theta) \equiv \mathbf{z}\sin\theta + \mathbf{z}'\cos\theta$, the velocity $\dot{\mathbf{z}} \equiv \mathrm{d}\mathbf{z}/\mathrm{d}\theta = \mathbf{z}\cos\theta - \mathbf{z}'\sin\theta$.

$$\mathbb{E}\Big\{\phi\big(\max_s[(F_s(\mathbf{z}) - F_s(\mathbf{z}'))^2]\big); \mathcal{G}\Big\} = \int_0^\infty \mathbb{P}\Big\{\big(\max_s[(F_s(\mathbf{z}) - F_s(\mathbf{z}'))^2] - a\big)_+ \mathbb{I}(\mathcal{G}) \ge x\Big\}\mathrm{d}x \tag{39}$$

$$= \int_0^\infty \mathbb{P}\big\{\max_s[|F_s(\mathbf{z}) - F_s(\mathbf{z}')|] \ge \sqrt{x+a}; \mathcal{G}\big\}\mathrm{d}x \tag{40}$$

$$\le 2|S| \int_a^\infty e^{-\lambda\sqrt{x}} \max_s \Big[\mathbb{E}\big\{\exp\{\lambda(F_s(\mathbf{z}) - F_s(\mathbf{z}'))\}; \mathcal{G}\big\}\Big]\mathrm{d}x\,, \tag{41}$$

where, in the last inequality we use the union bound followed by Markov's inequality. To control the exponential moment, note that $F_s(\mathbf{z}) - F_s(\mathbf{z}') = \int_0^{\pi/2}\langle\nabla F(\mathbf{z}(\theta)), \dot{\mathbf{z}}(\theta)\rangle\mathrm{d}\theta$ whence, using Jensen's inequality:

$$\mathbb{E}\Big\{\exp\big\{\lambda(F_s(\mathbf{z})-F_s(\mathbf{z}'))\big\};\mathcal{G}\Big\} = \mathbb{E}\Big\{\exp\Big(\int_0^{\pi/2}\lambda\langle\nabla F_s(\mathbf{z}(\theta)),\dot{\mathbf{z}}(\theta)\rangle\mathrm{d}\theta\Big);\mathcal{G}\Big\} \tag{42}$$

$$\leq \frac{2}{\pi}\int_0^{\pi/2}\mathbb{E}\Big\{\exp\big(\lambda\pi\langle\nabla F_s(\mathbf{z}(\theta)),\dot{\mathbf{z}}(\theta)\rangle/2\big);\mathcal{G}\Big\}\mathrm{d}\theta. \tag{43}$$

Define the set $\mathcal{G}_\theta = \{(\mathbf{z},\mathbf{z}') : \max_s\|\nabla F_s(\mathbf{z}(\theta))\| \leq L\}$. Then:

$$\mathbb{E}\Big\{\exp\big\{\lambda(F_s(\mathbf{z})-F_s(\mathbf{z}'))\big\};\mathcal{G}\Big\} \stackrel{(a)}{\leq} \frac{2}{\pi}\int_0^{\pi/2}\mathbb{E}\Big\{\exp\big(\lambda\pi\langle\nabla F_s(\mathbf{z}(\theta)),\dot{\mathbf{z}}(\theta)\rangle/2\big);\mathcal{G}_\theta\Big\}\mathrm{d}\theta \tag{44}$$

$$\stackrel{(b)}{=} \frac{2}{\pi}\int_0^{\pi/2}\mathbb{E}\Big\{\exp\Big(\frac{\lambda^2\pi^2\|\nabla F_s(\mathbf{z}(\theta))\|^2}{8};\mathcal{G}_\theta\Big)\Big\}\mathrm{d}\theta \tag{45}$$

$$\stackrel{(c)}{\leq} \exp\Big(\frac{\lambda^2\pi^2L^2}{8}\Big). \tag{46}$$

Here (a) follows as $\mathcal{G}_\theta \supseteq \mathcal{G}$. Equality (b) follows from noting that $\mathcal{G}_\theta$ is measurable with respect to $\mathbf{z}(\theta)$ and, hence, first integrating with respect to $\dot{\mathbf{z}}(\theta) = \mathbf{z}\cos\theta - \mathbf{z}'\sin\theta$, a Gaussian random variable that is independent of $\mathbf{z}(\theta)$. The final inequality (c) follows by using the fact that $\|\nabla F_s(\mathbf{z}(\theta))\| \leq L$ on the set $\mathcal{G}_\theta$.

Since this bound is uniform over $s \in S$, we can use it in (41):

$$\mathbb{E}\Big\{\phi(\max_s(F_s(\mathbf{z})-F_s(\mathbf{z}'))^2);\mathcal{G}\Big\} \leq 2|S|\int_a^\infty \exp\Big(-\lambda\sqrt{x}+\frac{\lambda^2\pi^2L^2}{8}\Big)\mathrm{d}x \tag{47}$$

$$\leq \frac{4|S|}{\lambda^2}(1+\lambda\sqrt{a})\exp\Big(-\lambda\sqrt{a}+\frac{\lambda^2\pi^2L^2}{8}\Big) \tag{48}$$

We can now set $\lambda = 4\sqrt{a}/\pi^2L^2$, to obtain the exponent above as $-2a/\pi^2L^2 = -\Delta^2/\pi^2L^2$. The prefactor $(1+\lambda\sqrt{a})\lambda^{-2}$ is bounded by $CL^2\max(,L^2/\Delta^2)$ when $a = \Delta^2/2$. Therefore, as required, we obtain:

$$\mathbb{E}\Big\{\phi(\max_s(F_s(\mathbf{z})-F_s(\mathbf{z}'))^2);\mathcal{G}\Big\} \leq C\max(1,L^4/\Delta^4)\exp\Big(-\frac{\Delta^2}{CL^2}\Big) \tag{49}$$

Combining this with Eq. (38) and the fact that $\phi(\Delta^2)^{-1} \leq C\Delta^{-2}$ gives Eq. (28) and, consequently, the lemma. ■

By a simple application of Cauchy-Schwarz, this lemma implies the following.

Corollary 10 *Under the same conditions as Lemma 9,*

$$\mathbb{P}\Big\{\max_{s\in S}|F_s(\mathbf{z})-\mathbb{E}F_s(\mathbf{z})| \geq \Delta\Big\} \leq C|S|\exp\Big(-\frac{\Delta^2}{CL^2}\Big) + \frac{C}{\Delta^2}\mathbb{E}\Big\{\max_{s\in S}\big[(F_s(\mathbf{z})-F_s(\mathbf{z}'))^4\big]\Big\}^{1/2}\mathbb{P}\{\mathcal{G}^c\}^{1/2}. \tag{50}$$

The following two lemmas are well-known concentration of measure results. The forms below can be found in (Vershynin, 2012, Corollary 5.35), (Laurent and Massart, 2000, Lemma 1) respectively.

Lemma 11 *Let $\mathbf{A} \in \mathbb{R}^{M\times N}$ be a matrix with i.i.d. standard normal entries, i.e. $\mathbf{A}_{ij} \sim \mathsf{N}(0,1)$. Then, for every $t \geq 0$:*

$$\mathbb{P}\left\{\|\mathbf{A}\|_{op} \geq \sqrt{M} + \sqrt{N} + t\right\} \leq \exp\left(-\frac{t^2}{2}\right). \tag{51}$$

Lemma 12 *Let $\mathbf{z} \sim \mathsf{N}(0, \mathrm{I}_N)$. Then*

$$\mathbb{P}\{\|\mathbf{z}\|^2 \geq N + 2\sqrt{Nt} + 2t\} \leq \exp(-t). \tag{52}$$

6. Proof of Theorem 1

Since $\widehat{\mathbf{\Sigma}} = \mathbf{X}^{\mathsf{T}}\mathbf{X}/n - \mathrm{I}_p$, we have:

$$\begin{aligned}\widehat{\mathbf{\Sigma}} = \sum_{q=1}^{r}\left\{\frac{\beta_q\|\mathbf{u}_q\|^2}{n}\mathbf{v}_q(\mathbf{v}_q)^{\mathsf{T}} + \frac{\sqrt{\beta_q}}{n}\big(\mathbf{v}_q(\mathbf{Z}^{\mathsf{T}}\mathbf{u}_q)^{\mathsf{T}} + (\mathbf{Z}^{\mathsf{T}}\mathbf{u}_q)\mathbf{v}_q^{\mathsf{T}}\big)\right\} \\ + \sum_{q\neq q'}\left\{\frac{\sqrt{\beta_q\beta_{q'}}\langle\mathbf{u}_q, \mathbf{u}_{q'}\rangle}{n}\mathbf{v}_q(\mathbf{v}_{q'})^{\mathsf{T}}\right\} + \frac{\mathbf{Z}^{\mathsf{T}}\mathbf{Z}}{n} - \mathrm{I}_p\,.\end{aligned} \tag{53}$$

We let $\mathsf{D} = \{(i,i) : i \in [p] \setminus \mathsf{Q}\}$ be the diagonal entries not included in any support. (Recall that $\mathsf{Q} = \cup_q \mathsf{Q}_q$ denote the union of the supports.) Further let $\mathsf{E} = \mathsf{Q} \times \mathsf{Q}$, $\mathsf{F} = (\mathsf{Q}^c \times \mathsf{Q}^c)\backslash\mathsf{D}$, and $\mathsf{G} = [p] \times [p]\backslash(\mathsf{D} \cup \mathsf{E} \cup \mathsf{F})$, or, equivalently $\mathsf{G} = (\mathsf{Q} \times \mathsf{Q}^c) \cup (\mathsf{Q}^c \times \mathsf{Q})$. Since these are disjoint we have:

$$\eta(\widehat{\mathbf{\Sigma}}) = \underbrace{\mathcal{P}_{\mathsf{E}}\left\{\eta(\widehat{\mathbf{\Sigma}})\right\}}_{\mathbf{S}} + \underbrace{\mathcal{P}_{\mathsf{F}}\left\{\eta\left(\widehat{\mathbf{\Sigma}}\right)\right\}}_{\mathbf{N}} + \underbrace{\mathcal{P}_{\mathsf{G}}\left\{\eta(\widehat{\mathbf{\Sigma}})\right\}}_{\mathbf{C}} + \underbrace{\mathcal{P}_{\mathsf{D}}\left\{\eta(\widehat{\mathbf{\Sigma}})\right\}}_{\mathbf{D}}. \tag{54}$$

The first term corresponds to the 'signal' component, while the last three terms correspond to the 'noise' component.

Theorem 1 is a direct consequence of the next five propositions. The first demonstrates that, even for a low level of thresholding, viz. $\tau < \sqrt{\log p}/2$, the term $\mathbf{N}$ has small operator norm. The second demonstrates that the soft thresholding operation preserves the signal in the term $\mathbf{S}$. The next two propositions show that the cross and diagonal terms $\mathbf{C}$ and $\mathbf{D}$ are negligible as well. Finally, in the last proposition, we demonstrate that, for the regime of thresholding far above the noise level, i.e. $\tau > C\sqrt{\log p}$, the noise terms $\mathbf{N}$ and $\mathbf{C}$ vanish entirely.

Proposition 13 *Let $\mathbf{N}$ denote the second term of Eq. (54). Since $\mathsf{F} = \mathsf{Q}^c \times \mathsf{Q}^c\backslash\mathsf{D}$,*

$$\mathbf{N} = \mathcal{P}_{\mathsf{F}}\left(\eta(\widehat{\mathbf{\Sigma}})\right) = \mathcal{P}_{\mathsf{F}}\left\{\eta\left(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z}\right)\right\}. \tag{55}$$

Then, there exists an absolute constant C such that the following happens. Assuming that (i) $\tau < \sqrt{\log p}/2$ and (ii) $n > C\log p$, then with probability $1 - o(1)$

$$\|\mathbf{N}\|_{op} \leq C\left(\sqrt{\frac{p}{n}} \vee \frac{p}{n}\right)e^{-\tau^2/C}. \tag{56}$$

Proposition 14 *Let* $\mathbf{S}$ *denote the first term in Eq.* (54)*:*

$$\mathbf{S} = \mathcal{P}_{\mathsf{E}}\left\{\eta(\widehat{\mathbf{\Sigma}})\right\}. \tag{57}$$

Assume that (i) $s_0/n < 1$ *and (ii)*$n > C\log p$*: Then with probability* $1-o(1)$*:*

$$\left\|\mathbf{S} - \mathbf{\Sigma}\right\|_{op} \leq \frac{2\tau s_0}{\sqrt{n}} + C(\beta \vee 1)\sqrt{\frac{s_0}{n}}. \tag{58}$$

Proposition 15 *Let* $\mathbf{C}$ *denote the matrix corresponding to the third term of Eq.* (54)*:*

$$\mathbf{C} = \mathcal{P}_{\mathsf{G}}\left\{\eta(\widehat{\mathbf{\Sigma}})\right\}.$$

Assuming the conditions of Proposition 13 and, additionally, that $s_0^2 \leq p$*, there exist constants* C, c *such that with probability* $1-o(1)$

$$\|\mathbf{C}\|_{op} \leq C\,\tau e^{-c\tau^2/(\beta\vee 1)}\sqrt{\frac{p}{n}} \vee \frac{p}{n}. \tag{59}$$

Proposition 16 *Let* $\mathbf{D}$ *denote the matrix corresponding to the third term of Eq.* (54)*:*

$$\mathbf{D} = \mathcal{P}_{\mathsf{D}}\left\{\eta(\widehat{\mathbf{\Sigma}})\right\}.$$

With probability $1-o(1)$ *we have that* $\|\mathbf{D}\|_{op} \leq C\sqrt{n^{-1}\log p}$*.*

Proposition 17 *For some absolute constant* C_0*, we have for* $\tau \geq C_0(\beta \vee 1)\sqrt{\log p}$ *that, with probability* $1-o(1)$*:*

$$\forall i,j \quad N_{ij} = C_{ij} = 0. \tag{60}$$

Therefore, $\|\mathbf{N}\|_{op} = 0$ *and* $\|\mathbf{C}\|_{op} = 0$*.*

Remark 18 *At this point we remark that the probability* $1-o(1)$ *can be made quantitative, for e.g. of the form* $1-\exp(-\min(\sqrt{p}, n)/C_1)$*, for every* n *large enough. For simplicity of exposition we do not pursue this in the paper.*

We defer the proofs of Propositions 13, 14, 15, 16 and 17 to Sections 6.1, 6.2, 6.3, 6.4 and 6.5 respectively. By combining them for $\beta = O(1)$, we immediately obtain the following bound.

Theorem 19 *There exist numerical constants* C_0, C_1 *such that the following happens. Assume* $\beta \leq C_0$*,* $n > C_1 \log p$ *and* $\tau \leq \sqrt{\log p}/2$*. Then with probability* $1-o(1)$*:*

$$\left\|\eta(\widehat{\mathbf{\Sigma}}) - \mathbf{\Sigma}\right\|_{op} \leq \frac{2\tau s_0}{\sqrt{n}} + C\Big(\sqrt{\frac{p}{n}} \vee \frac{p}{n}\Big)\, e^{-\tau^2/C} + C\sqrt{\frac{s_0 \vee \log p}{n}}. \tag{61}$$

Proof The proof is obtained by adding the error terms from Propositions 13, 14, 15 and 16, and noting that β is bounded. ■

Using Propositions 13, 14, 15 and 16, together with a suitable choice of τ, we obtain the proof of Theorem 1.

Proof [Proof of Theorem 1] Note that in the case $s_0^2 > p/e$ there is no thresholding and hence the result follows from the fact that $\|\widehat{\mathbf{\Sigma}} - \mathbf{\Sigma}\|_{op} \le C\sqrt{p/n}$ (Vershynin, 2012, Remark 5.40).

We assume now that $s_0^2 \le p/e$ and the case that $\tau_* = C_1(\beta \vee 1)\sqrt{\log(p/s_0^2)} \le \sqrt{\log p}/2$. In that case we set $\tau = \tau_* \le \sqrt{\log p}/2$. Below we will keep C_1 a large enough constant, and check that each of the error terms in Propositions 13, 14, 15 and 16 is upper bounded by (a constant times) the right-hand side of Eq. (7). Throughout C will denote a generic constant that can be made as large as we want, and can change from line to line.

We start from Proposition 13:

$$\begin{aligned}
\|\mathbf{N}\|_{op} &\le C\left(\sqrt{\frac{p}{n}} \vee \frac{p}{n}\right)\left(\frac{s_0^2}{p}\right)^C && (62)\\
&\le C\sqrt{\frac{p}{n}\left(\frac{p}{s_0^2}\right)^{-C-1}} \vee C\sqrt{\left(\frac{p}{n}\right)^2\left(\frac{p}{s_0^2}\right)^{-C-2}} && (63)\\
&\le C\sqrt{\frac{s_0^2}{n}\left(\frac{p}{s_0^2}\right)^{-C}} \vee C\sqrt{\left(\frac{s_0^2}{n}\right)^2\left(\frac{p}{s_0^2}\right)^{-C}} && (64)\\
&\le C\sqrt{\frac{s_0^2}{n}\log\frac{p}{s_0^2}}\,, && (65)
\end{aligned}$$

where in the last step we used $(e\, s_0^2/p), (s_0^2/n) \le 1$.

Next consider Proposition 14:

$$\begin{aligned}
\big\|\mathbf{S} - \mathbf{\Sigma}\big\|_{op} &\le C\sqrt{\frac{s_0^2\tau^2}{n}} + C\sqrt{\frac{s_0(\beta \vee 1)^2}{n}} && (66)\\
&\le C\sqrt{\frac{s_0^2(\beta^2 \vee 1)}{n}\log\frac{p}{s_0^2}}\,. && (67)
\end{aligned}$$

From Proposition 15, we get, using the same argument as in Eq. (65)

$$\begin{aligned}
\|\mathbf{C}\|_{op} &\le C\sqrt{\beta \vee 1}\left(\sqrt{\frac{p}{n}} \vee \frac{p}{n}\right)\left(\frac{s_0^2}{p}\right)^C && (68)\\
&\le C(\beta \vee 1)\sqrt{\frac{s_0^2}{n}\log\frac{p}{s_0^2}}. && (69)
\end{aligned}$$

Finally, the term of Proposition 16 is also bounded as desired using $\log p \le s_0^2\log(p/s_0^2)$ (dividing both sides by p and using the fact that $x \mapsto x\log(1/x)$ is increasing).

The case of $\tau_* \ge \sqrt{\log p}/2$ is easier. In that case, we can keep $\tau = C_2\tau_*$ with C_2 large enough so that $\tau \ge C_0(\beta \vee 1)\sqrt{\log p}$ for C_0 of Proposition 17. Then, by Proposition 17, we know that $\mathbf{N} = 0$ and $\mathbf{C} = 0$. Therefore we only need consider the terms $\mathbf{S} - \mathbf{\Sigma}$ and $\mathbf{D}$. For these terms we can use Propositions 14 and 16 respectively and, arguing as in the earlier case $\tau_* \leq \sqrt{\log p}$, we obtain the desired result. ■

6.1 Proof of Proposition 13

Define $\widetilde{\mathbf{N}}$ as

$$\widetilde{\mathbf{N}} = \mathcal{P}_{\mathsf{nd}}\left\{\eta\left(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z}\right)\right\}.$$

Since $\mathbf{N}$ is a principal submatrix of $\widetilde{\mathbf{N}}$, it suffices to prove the same bound for $\widetilde{\mathbf{N}}$. Our main tool in the proof will be the concentration lemma 9 which we use on multiple occasions. With a view to using the lemma, we let let $\mathbf{Z}' \in \mathbb{R}^{n\times p}$ denote an independent copy of $\mathbf{Z}$, and $\tilde{\mathbf{z}}'_i$ it's i^{th} column. The proof relies on two preliminary lemmas. For some $A \geq 1$ (to be chosen later), we first state and prove the following lemma that controls the norm of *any principal submatrix of* $\tilde{\mathbf{N}}$ of size at most p/A.

Lemma 20 *Fix any $A \geq 1$. There exists an absolute constants C, c such that:*

$$\mathbb{P}\Big\{\max_{\mathsf{S}\subseteq[p],|\mathsf{S}|\leq p/A}\|\mathcal{P}_{\mathsf{S}\times\mathsf{S}}(\tilde{\mathbf{N}})\|_{op} \geq \Delta\Big\} \leq C\,\exp\left(p\frac{\log CA}{A} - \frac{n^2\Delta^2}{C(n+p)}\right) + C\frac{(np)^C}{\Delta^2}\exp(-cn). \tag{70}$$

Proof For any subset $\mathsf{S} \subset [p]$ recall that $T_p^{\varepsilon}(\mathsf{S})$ denotes an ε-net of unit vectors in $\mathbb{S}^{p-1}$ supported on the subset S. For simplicity let $T(A) = \cup_{\mathsf{S}:|\mathsf{S}|\leq p/A} T_p^{\varepsilon}(\mathsf{S})$. It suffices, by Lemma 8, to control $\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle$ on the set $T(A)$. In particular:

$$\mathbb{P}\Big\{\max_{\mathsf{S}\subseteq[p],|\mathsf{S}|\leq p/A}\|\mathcal{P}_{\mathsf{S}\times\mathsf{S}}(\tilde{\mathbf{N}})\|_{op} \geq \Delta\Big\} \leq \mathbb{P}\Big\{\max_{\mathbf{y}\in T(A)}|\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle| \geq \Delta(1-2\varepsilon)\Big\}. \tag{71}$$

Consider the good set $\mathcal{G}_1$ given by:

$$\mathcal{G}_1 = \{(\mathbf{Z}, \mathbf{Z}') : \max(\|\mathbf{Z}\|, \|\mathbf{Z}'\|) \leq \sqrt{2}(\sqrt{n} + \sqrt{p}))\}. \tag{72}$$

To use Lemma 9, we need to compute $\mathbb{E}\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle$ and the gradient of $\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle$ with respect to the underlying random variables $\mathbf{Z}$. Since $\eta(\cdot)$ is an odd function the expectation vanishes. To compute the gradient, we let $t \in [0,1]$ and $\mathbf{W} = \sqrt{t}\mathbf{Z} + \sqrt{1-t}\mathbf{Z}'$, and consider $\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle = \langle \mathbf{y}, \eta(\mathbf{W}^{\mathsf{T}}\mathbf{W}/n)\mathbf{y}\rangle$ as a function of the $\mathbf{W}$. Taking the gradient with respect to a column $\tilde{\mathbf{w}}_\ell$ for $\ell \in \mathsf{S}$:

$$\begin{aligned}
\nabla_{\tilde{\mathbf{w}}_\ell}\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle &= \frac{y_\ell}{n}\sum_{i\neq\ell, i\in\mathsf{S}} \tilde{\mathbf{w}}_i y_i \partial\eta(\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}_\ell\rangle/n) &(73)\\
&= \frac{y_\ell}{n}\mathbf{W}\boldsymbol{\sigma}, &(74)
\end{aligned}$$

where

$$\sigma_i = \begin{cases} y_i \partial\eta(\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}_\ell\rangle/n) & \text{if } i \neq \ell, i \in \mathsf{S} \\ 0 & \text{otherwise.} \end{cases} \tag{75}$$

Since $\|\boldsymbol{\sigma}\| \le \|\mathbf{y}\| = 1$, we have that $\|\nabla_{\tilde{\mathbf{w}}_\ell}\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle\|^2 \le y_\ell^2 \|\mathbf{W}\|^2/n^2$. Summing over $\ell \in \mathsf{S}$ we obtain the gradient bound, holding on the good set $\mathcal{G}_1$:

$$\begin{aligned} \|\nabla_{\mathbf{W}}\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle\|^2 &\le \frac{\sum_\ell y_\ell^2}{n^2}\|\mathbf{W}\|^2 \qquad (76)\\ &\le \frac{C(n+p)}{n^2}\,, \qquad (77)\end{aligned}$$

which holds because of triangle inequality and the fact that $\sqrt{t} + \sqrt{1-t} \le \sqrt{2}$. We can now apply Lemma 9 to bound the RHS of Eq. (71) and get:

$$\begin{aligned}\mathbb{P}\Big\{\max_{\mathsf{S}\subseteq[p], |\mathsf{S}|\le p/A} \mathcal{P}_{\mathsf{S}\times\mathsf{S}}(\tilde{\mathbf{N}}) \ge \Delta\Big\} &\le C|T(A)|\,\exp\Big(-\frac{n^2\Delta^2}{C(n+p)}\Big) \\ &\quad + \frac{C}{\Delta^2}\mathbb{E}\big\{\max_{\mathbf{y}\in T}\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle^2; \mathcal{G}_1^c\big\}. \qquad (78)\end{aligned}$$

We can simplify the terms on the right-hand side to obtain the result of the lemma. With $\varepsilon = 1/4$, Stirling's approximation and Lemma 7 we have:

$$|T(A)| \le \exp\Big(p\frac{\log CA}{A}\Big). \tag{79}$$

We use a crude bound on the complement of the good set $\mathcal{G}_1$. It is easy to see that, for any unit vector $\mathbf{y}$, $\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle^2 \le \|\tilde{\mathbf{N}}\|_F^2 \le \|\mathbf{Z}^{\mathsf{T}}\mathbf{Z}\|_F^2/n^2$. Cauchy-Schwarz then implies that

$$\begin{aligned}\mathbb{E}\{\max\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle^2; \mathcal{G}_1^c\} &\le n^{-2}\big(\mathbb{E}\{\|\mathbf{Z}^{\mathsf{T}}\mathbf{Z}\|_F^4\}\big)^{1/2}\mathbb{P}\{\mathcal{G}_1^c\}^{1/2} \qquad (80)\\ &\le (np)^C\exp(-c(n+p)), \qquad (81)\end{aligned}$$

where the bound on $\mathbb{P}\{\mathcal{G}_1^c\}$ follows from Lemma 11. This concludes the lemma. ∎

Note that Lemma 20, with $A = 1$, tells us that $\|\tilde{\mathbf{N}}\|_{op}$ is of order $\sqrt{p/n + (p/n)^2}$ (uniformly in τ) with high probability. Already this non-asymptotic bound is non-trivial, since the previous results of Cheng and Singer (2013) and Fan and Montanari (2015) do not extend to this case. However, Proposition 13 is stronger, and establishes a rate of decay with the thresholding level τ.

The second lemma we require controls the Rayleigh quotient $\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle$ when the entries of $\mathbf{y}$ are "spread out".

Lemma 21 *Assume that $\tau \le \sqrt{\log p}/2$. Given $A \ge 1$ and a unit vector $\mathbf{y}$, let $\mathsf{S} = \{i : |y_i| \le \sqrt{A/p}\}$ and $\mathbf{y}_{\mathsf{S}}, \mathbf{y}_{\mathsf{S}^c}$ denote the projections of $\mathbf{y}$ onto supports S, S^c respectively. We have:*

$$\mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^{1/4}}|\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle| \ge \Delta\Big\} \le C\exp\Big(-\frac{n^2\Delta^2}{L_1^2} + Cp\Big) + (np)^C\exp\big(-c\min(\sqrt{p}, n)\big), \tag{82}$$

for any $\Delta \geq L_1$ *where* $L_1 = C_1\sqrt{A\exp(-\tau^2/16)(n+p)/n^2}$. *The same bound holds for* $\mathbb{P}\big\{\max_{\mathbf{y}\in T_p^{1/4}} |\langle \mathbf{y}_{\mathsf{S}^c}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle| \geq \Delta\big\}$.

Proof We first prove the claim for $\langle \mathbf{y}_{\mathsf{S}}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle$. Firstly, we have $\mathbb{E}\langle \mathbf{y}_{\mathsf{S}}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle = 0$. Consider the "good set" $\mathcal{G}_2$ of pairs $(\mathbf{W}, \mathbf{W}') \in \mathbb{R}^{n\times p}\times\mathbb{R}^{n\times p}$ satisfying the conditions:

$$\|\mathbf{W}\|, \|\mathbf{W}'\| \leq \sqrt{2}(\sqrt{n}+\sqrt{p})\,, \tag{83}$$

$$\forall i\in[p], \quad \frac{1}{p}\sum_{j\in[p]\backslash i} |\mathbb{I}(\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}_j\rangle| \geq \tau\sqrt{n}/2) \leq 2\exp(-\tau^2/16)\,, \tag{84}$$

$$\forall i\in[p], \quad \frac{1}{p}\sum_{j\in[p]\backslash i} |\mathbb{I}(\langle \tilde{\mathbf{w}}'_i, \tilde{\mathbf{w}}'_j\rangle| \geq \tau\sqrt{n}/2) \leq 2\exp(-\tau^2/16)\,, \tag{85}$$

$$\forall i\in[p], \quad \frac{1}{p}\sum_{j\in[p]} \mathbb{I}(|\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}'_j\rangle| \geq \tau\sqrt{n}/2) \leq 2\exp(-\tau^2/16).\,. \tag{86}$$

Also, for any pair $\mathbf{W}, \mathbf{W}' \in \mathcal{G}_2$, for $\mathbf{W}(t) = \sqrt{t}\mathbf{W} + \sqrt{1-t}\mathbf{W}'$ (and its columns $\tilde{\mathbf{w}}(t)_i$ defined appropriately) we have:

$$\|\mathbf{W}(t)\| \leq \max_t(\sqrt{t}+\sqrt{1-t})(\sqrt{2n}+\sqrt{2p}) = 2(\sqrt{n}+\sqrt{p}), \tag{87}$$

$$\forall i\in[p] \quad \frac{1}{p}\sum_{j\in[p]\backslash i} \mathbb{I}(\langle \tilde{\mathbf{w}}(t)_i, \tilde{\mathbf{w}}(t)_j\rangle \geq \tau\sqrt{n}) \leq 6\exp(-\tau^2/16). \tag{88}$$

Equation (87) follows by a simple application of triangle inequality and condition (83) defining $\mathcal{G}_2$. For inequality (88), expanding the product $\langle \tilde{\mathbf{w}}(t)_i, \tilde{\mathbf{w}}(t)_j\rangle$:

$$\langle \tilde{\mathbf{w}}(t)_i, \tilde{\mathbf{w}}(t)_j\rangle = t\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}_j\rangle + (1-t)\langle \tilde{\mathbf{w}}'_i, \tilde{\mathbf{w}}'_j\rangle + \sqrt{t(1-t)}\langle \tilde{\mathbf{w}}_{\mathbf{i}}, \tilde{\mathbf{w}}'_j\rangle, \tag{89}$$

whence, by triangle inequality and $\sqrt{t(1-t)} < 1$

$$\begin{aligned}\mathbb{I}(|\langle \tilde{\mathbf{w}}(t)_i, \tilde{\mathbf{w}}(t)_j\rangle| \geq \tau\sqrt{n}) \leq \mathbb{I}(|\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}_j\rangle| \geq \tau\sqrt{n}/2) + \mathbb{I}(|\langle \tilde{\mathbf{w}}'_i, \tilde{\mathbf{w}}'_j\rangle| \geq \tau\sqrt{n}/2)\\ + \mathbb{I}(|\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}'_j\rangle| \geq \tau\sqrt{n}/2).\end{aligned} \tag{90}$$

The gradient of $\langle \mathbf{y}_{\mathsf{S}}, \eta(\mathbf{W}^{\mathsf{T}}\mathbf{W}/n)\mathbf{y}_{\mathsf{S}}\rangle$ with respect to a column $\tilde{\mathbf{w}}_\ell$ of $\mathbf{W}$ is given by:

$$\nabla_{\tilde{\mathbf{w}}_\ell}\langle \mathbf{y}_{\mathsf{S}}, \eta(\mathbf{W}^{\mathsf{T}}\mathbf{W}/n)\mathbf{y}_{\mathsf{S}}\rangle = \frac{y_\ell}{n}\sum_{j\in\mathsf{S}\backslash\ell} y_j\partial\eta\Big(\frac{\langle \tilde{\mathbf{w}}_j, \tilde{\mathbf{w}}_\ell\rangle}{n}; \frac{\tau}{\sqrt{n}}\Big)\tilde{\mathbf{w}}_j \tag{91}$$

$$= \frac{y_\ell}{n}\mathbf{W}\boldsymbol{\sigma}, \tag{92}$$

$$\text{where} \quad \sigma_i = \begin{cases}\partial\eta(\langle \tilde{\mathbf{w}}_i, \tilde{\mathbf{w}}_\ell\rangle/n; \tau/\sqrt{n})y_i & \text{when } i\in\mathsf{S}\backslash\ell\\ 0 & \text{otherwise.}\end{cases} \tag{93}$$

Therefore

$$\|\nabla_{\tilde{\mathbf{w}}_\ell}\langle \mathbf{y}_{\mathsf{S}}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle\|^2 \le \frac{y_\ell^2}{n^2}\|\mathbf{W}\|^2\|\boldsymbol{\sigma}\|^2 \tag{94}$$

$$\le \frac{y_\ell^2\|\mathbf{W}\|^2}{n^2}\sum_{i\neq\ell}(y_i\partial\eta(\langle\tilde{\mathbf{w}}_i,\tilde{\mathbf{w}}_\ell\rangle/n))^2 \tag{95}$$

$$\stackrel{(a)}{\le} \frac{y_\ell^2\|\mathbf{W}\|^2}{n^2}\sum_{i\neq\ell}\frac{A}{p}\mathbb{I}(|\langle\tilde{\mathbf{w}}_i,\tilde{\mathbf{w}}_\ell\rangle| \ge \tau\sqrt{n}) \tag{96}$$

$$\stackrel{(b)}{\le} \frac{y_\ell^2}{n^2}C(n+p)A\exp(-\tau^2/16) \tag{97}$$

Here (a) follows from fact that the entries of $\mathbf{y}_{\mathsf{S}}$ are bounded by $\sqrt{A/p}$ and the definition of the soft thresholding function. Inequality (b) follows follows when we set $\mathbf{W} = \mathbf{Z}(t) = \sqrt{t}\mathbf{Z} + \sqrt{1-t}\mathbf{Z}'$ and $(\mathbf{Z},\mathbf{Z}') \in \mathcal{G}_2$. Therefore, summing over ℓ we obtain the following bound for the gradient of $\langle \mathbf{y}_{\mathsf{S}}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle$

$$\|\nabla_{\mathbf{Z}(t)}\langle \mathbf{y}_{\mathsf{S}}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle\|^2 \le C_1\frac{A\exp(-\tau^2/16)(n+p)}{n^2} \equiv L_1^2. \tag{98}$$

We can use now Lemma 9, to get, for $L_1 > 0$ as defined above and any $\Delta \ge L_1$:

$$\mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^{1/4}}\langle \mathbf{y}_{\mathsf{S}}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle \ge \Delta\Big\} \le C\exp\Big(-\frac{\Delta^2}{CL_1^2} + Cp\Big) + CL_1^{-2}\mathbb{E}\{\max_{\mathbf{y}\in T_p^{\varepsilon}}\langle \mathbf{y}_{\mathsf{S}}, \widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle^2; \mathcal{G}_2\} \tag{99}$$

$$\le C\exp\Big(-\frac{\Delta^2}{CL_1^2} + Cp\Big) + C(np)^C\mathbb{P}\{\mathcal{G}_2^c\}^{1/2}, \tag{100}$$

where the last line follows by Cauchy-Schwarz, as in the proof of Lemma 20, and the fact that $L_1 \ge (np)^{-C_2}$ using the upper bound $\tau \le \sqrt{\log p}/2$.

To obtain the thesis, we need to now bound $\mathbb{P}\{\mathcal{G}_2^c\}$. It suffices to control the failure probability of conditions (83), (84), (85), (86) of the good set $\mathcal{G}_2$ individually, and apply the union bound. For $\mathbf{Z},\mathbf{Z}'$ independent, $\max(\|\mathbf{Z}\|,\|\mathbf{Z}'\|) \ge \sqrt{2}(\sqrt{n}+\sqrt{p})$ with probability at most $2\exp(-c(n+p))$ by Lemma 11. Now consider condition (84) with $i=1$, without loss of generality. First, for any $h > 0$ we have:

$$\mathbb{P}\Big\{\frac{1}{p}\sum_{j\neq 1}\mathbb{I}(|\langle\tilde{\mathbf{z}}_1,\tilde{\mathbf{z}}_j\rangle| \ge \tau\sqrt{n}/2) \ge h\Big\} \le \mathbb{P}\Big\{\frac{1}{p}\sum_{j\neq 1}\mathbb{I}(|\langle\tilde{\mathbf{z}}_1,\tilde{\mathbf{z}}_j\rangle| \ge \tau\sqrt{n}/2) \ge 2h; \|\tilde{\mathbf{z}}_1\| \le 2\sqrt{n}\Big\} + \mathbb{P}\big\{\|\tilde{\mathbf{z}}_1\| \ge \sqrt{2n}\big\}. \tag{101}$$

Lemma 12 guarantees that the second term is at most $\exp(-cn)$. To control the first term, we note that, conditional on $\tilde{\mathbf{z}}_1$, $\langle\tilde{\mathbf{z}}_j,\tilde{\mathbf{z}}_1\rangle, j\neq 1$ are independent Gaussian random variables with variance $\|\tilde{\mathbf{z}}_1\|^2$. Therefore, conditional on $\tilde{\mathbf{z}}_1$, $\mathbb{I}(|\langle\tilde{\mathbf{z}}_1,\tilde{\mathbf{z}}_j\rangle| \ge \tau\sqrt{n}/2)$ are independent Bernoulli random variables with success probability $h_0 = 2\Phi\big(-\tau\sqrt{n}/(2\|\tilde{\mathbf{z}}_1\|)\big)$, where $\Phi(\cdot)$ is the Gaussian cumulative distribution function. It follows, by the Chernoff-Hoeffding bound for Bernoulli random variables that

$$\mathbb{P}\Big\{\frac{1}{p}\sum_{j\neq 1}\mathbb{I}(|\langle\tilde{\mathbf{z}}_1,\tilde{\mathbf{z}}_j\rangle|\geq\tau\sqrt{n}/2)\geq h\big|\tilde{\mathbf{z}}_1\Big\}\leq\exp\big(-p\,D(h\|h_0)\big), \tag{102}$$

where $D(a\|b)=a\log(a/b)+(1-a)\log[(1-a)/(1-b)]$. Choosing $h=4\Phi(-\tau/(2\sqrt{2}))$, and conditional on $\|\tilde{\mathbf{z}}_1\|\leq\sqrt{2n}$, $D(h\|h_0)\geq ch$ for a constant c, implying that

$$\mathbb{P}\Big\{\frac{1}{p}\sum_{j\neq 1}\mathbb{I}(|\langle\tilde{\mathbf{z}}_1,\tilde{\mathbf{z}}_j\rangle|\geq\tau\sqrt{n}/2)\geq h;\|\tilde{\mathbf{z}}_1\|\leq\sqrt{2n}\Big\}\leq\exp(-cph). \tag{103}$$

By standard bounds $h=4\Phi(-\tau/2\sqrt{2})\leq 2\exp(-\tau^2/16)$ and, as $\tau\leq\sqrt{\log p}/2$, $h\geq 1/\sqrt{p}$, we have

$$\mathbb{P}\Big\{\frac{1}{p}\sum_{j\neq 1}\mathbb{I}(|\langle\tilde{\mathbf{z}}_1,\tilde{\mathbf{z}}_j\rangle|\geq\tau\sqrt{n}/2)\geq h;\|\tilde{\mathbf{z}}_1\|\leq\sqrt{2n}\Big\}\leq\exp(-c\sqrt{p}). \tag{104}$$

Combining this with Eq. (101) we now get:

$$\mathbb{P}\Big\{\frac{1}{p}\sum_{j\neq 1}\mathbb{I}(|\langle\tilde{\mathbf{z}}_1,\tilde{\mathbf{z}}_j\rangle|\geq\tau\sqrt{n}/2)\geq h\Big\}\leq 2\exp(-c\min(n,\sqrt{p})). \tag{105}$$

A similar bound holds for $i\neq 1$ and the other conditions (85) and (86), whence we have by the union bound that $\mathbb{P}\{\mathcal{G}_2^c\}\leq p^2\exp(-c\min(\sqrt{p},n))$. This completes the proof of the claim (82).

The proof of the claim for $\langle\mathbf{y}_{\mathsf{S}},\widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle$ is analogous, so we only sketch the points at which it differs from that of Eq. (82). We use the same good set $\mathcal{G}_2$, as defined earlier. Computing the gradient as for $\langle\mathbf{y}_{\mathsf{S}},\widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle$ we obtain:

$$\nabla_{\tilde{\mathbf{w}}_\ell}\langle\mathbf{y}_{\mathsf{S}},\widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle=\frac{y_\ell}{n}\sum_{j\in\mathsf{S}(\ell)}y_j\tilde{\mathbf{w}}_j\partial\eta\Big(\frac{\langle\tilde{\mathbf{w}}_j,\tilde{\mathbf{w}}_\ell\rangle}{n};\frac{\tau}{\sqrt{n}}\Big). \tag{106}$$

Here $\mathsf{S}(\ell)=\mathsf{S}^c$ if $\ell\in\mathsf{S}$ and S otherwise. Define the vector $\boldsymbol{\sigma}(\ell)\in\mathbb{R}^p$ as

$$(\boldsymbol{\sigma}(\ell))_j=\begin{cases}y_\ell y_j\partial\eta\Big(\frac{\langle\tilde{\mathbf{w}}_j,\tilde{\mathbf{w}}_\ell\rangle}{n};\frac{\tau}{\sqrt{n}}\Big) & \text{if } j\in\mathsf{S}(\ell)\\ 0 & \text{otherwise.}\end{cases} \tag{107}$$

As before, we have that $\|\nabla_{\tilde{\mathbf{w}}_\ell}\langle\mathbf{y}_{\mathsf{S}},\widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle\|=n^{-1}\|\mathbf{W}\boldsymbol{\sigma}(\ell)\|\leq n^{-1}\|\mathbf{W}\|\|\boldsymbol{\sigma}(\ell)\|$. Therefore, summing over $\ell\in[p]$:

$$\begin{aligned}\|\nabla_{\mathbf{W}}\langle\mathbf{y}_{\mathsf{S}},\widetilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle\|^2&\leq\frac{\|\mathbf{W}\|}{n^2}\sum_{\ell\in[p]}\|\boldsymbol{\sigma}(\ell)\|^2 &(108)\\&\leq\frac{\|\mathbf{W}\|^2}{n^2}\sum_{\ell\in[p]}\sum_{j\in\mathsf{S}(\ell)}y_\ell^2y_j^2\partial\eta\Big(\frac{\langle\tilde{\mathbf{w}}_j,\tilde{\mathbf{w}}_\ell\rangle}{n};\frac{\tau}{\sqrt{n}}\Big) &(109)\\&=\frac{2\|\mathbf{W}\|^2}{n^2}\sum_{\ell\in\mathsf{S}}\sum_{j\in\mathsf{S}^c}y_j^2y_\ell^2\partial\eta\Big(\frac{\langle\tilde{\mathbf{w}}_j,\tilde{\mathbf{w}}_\ell\rangle}{n};\frac{\tau}{\sqrt{n}}\Big) &(110)\\&\leq\frac{2\|\mathbf{W}\|^2}{n^2}\frac{A}{p}\max_{\ell\in[p]}\sum_{j\neq p}\partial\eta\Big(\frac{\langle\tilde{\mathbf{w}}_j,\tilde{\mathbf{w}}_\ell\rangle}{n};\frac{\tau}{\sqrt{n}}\Big). &(111)\end{aligned}$$

Under the condition of $\mathcal{G}_2$, the gradient also satisfies, when evaluated at $\mathbf{W} = \mathbf{Z}(t) = \sqrt{t}\mathbf{Z} + \sqrt{1-t}\mathbf{Z}'$:

$$\|\nabla_{\mathbf{Z}(t)}\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle\|^2 \le \frac{CA\exp(-\tau^2/16)(n+p)}{n^2}. \tag{112}$$

The rest of the proof is then the same as before. ■

Given these lemmas, we can now establish Proposition 13.

Proof [Proof of Proposition 13] We use a variant of the ε-net argument of Lemma 20. To bound the probability that $\|\tilde{\mathbf{N}}\|_{op}$ is large, with Lemma 8, we obtain:

$$\mathbb{P}\Big\{\|\tilde{\mathbf{N}}\|_{op} \ge \Delta\Big\} \le \mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^\varepsilon}|\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle| \ge \Delta(1-2\varepsilon)\Big\}. \tag{113}$$

Let $\mathsf{S} = \{i : |y_i| \le \sqrt{A/p}\}$ for some $A \ge 1$ to be chosen later. Then let $\mathbf{y} = \mathbf{y}_{\mathsf{S}} + \mathbf{y}_{\mathsf{S}^c}$ denote the projections of $\mathbf{y}$ onto supports S, S^c respectively. Since $\langle \mathbf{y}, \tilde{\mathbf{N}}\mathbf{y}\rangle = \langle \mathbf{y}_{\mathsf{S}^c}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle + \langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle + 2\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle$ by triangle inequality and union bound:

$$\mathbb{P}\Big\{\|\tilde{\mathbf{N}}\|_{op} \ge \Delta\Big\} \le \mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^\varepsilon}|\langle \mathbf{y}_{\mathsf{S}^c}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle| + |\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle| + 2|\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{sS^c}\rangle| \ge \Delta(1-2\varepsilon)\Big\} \tag{114}$$

$$\le \mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^\varepsilon}|\langle \mathbf{y}_{\mathsf{S}^c}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle| \ge \Delta(1-2\varepsilon)/4\Big\} + \mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^\varepsilon}|\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle| \ge \Delta(1-2\varepsilon)/4\Big\}$$
$$+ \mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^\varepsilon}|\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle| \ge \Delta(1-2\varepsilon)/4\Big\} \tag{115}$$

$$\le \mathbb{P}\Big\{\max_{\mathsf{S}':|\mathsf{S}'|\le p/A}\|\mathcal{P}_{\mathsf{S}'\times\mathsf{S}'}(\tilde{\mathbf{N}})\| \ge \Delta(1-2\varepsilon)/4\Big\} + \mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^\varepsilon}|\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}}\rangle| \ge \Delta(1-2\varepsilon)/4\Big\}$$
$$+ \mathbb{P}\Big\{\max_{\mathbf{y}\in T_p^\varepsilon}|\langle \mathbf{y}_{\mathsf{S}}, \tilde{\mathbf{N}}\mathbf{y}_{\mathsf{S}^c}\rangle| \ge \Delta(1-2\varepsilon)/4\Big\}. \tag{116}$$

With $\varepsilon = 1/4$, the first term is controlled by Lemma 20 while the final two are controlled by Lemma 21. We choose $\varepsilon = 1/4$ in Eq. (116), and

$$\Delta = \Delta_* \equiv C\sqrt{\frac{p}{n}\Big(1+\frac{p}{n}\Big)\Big(\frac{\log A}{A} + A\exp\Big(-\frac{\tau^2}{16}\Big)\Big)}, \tag{117}$$

for large enough C so that, using the bounds of Lemmas 20 and 21, we have:

$$\mathbb{P}\Big\{\tilde{\mathbf{N}} \ge \Delta_*\Big\} \le C(np)^C \exp\Big[-c\min\Big(\sqrt{p}, n, p\frac{\log A}{A}\Big)\Big]. \tag{118}$$

This probability bound is $o(1)$ provided A is not too large: we choose $A = 0.25\sqrt{\tau\exp(\tau^2/16)} \ll \sqrt{p}$ which guarantees that the bound above is $o(1)$ when $n > C\log p$ for some C large enough. This concludes the proposition. ■

6.2 Proof of Proposition 14

We decompose the empirical covariance matrix (53) as

$$\mathcal{P}_{\mathsf{E}}(\widehat{\mathbf{\Sigma}}) = \mathbf{\Sigma} + \mathbf{\Delta}_1 + \mathbf{\Delta}_2 + \mathbf{\Delta}_2^{\mathsf{T}} + \mathcal{P}_{\mathsf{E}}\Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}_p\Big)\,, \tag{119}$$

$$\mathbf{\Delta}_1 \equiv \sum_{q,q'=1}^{r} \sqrt{\beta_q\beta_q'}\Big(\frac{1}{n}\langle \mathbf{u}_q, \mathbf{u}_q'\rangle - \mathbf{1}_{q=q'}\Big)\mathbf{v}_q\mathbf{v}_q'^{\mathsf{T}}\,, \tag{120}$$

$$\mathbf{\Delta}_2 \equiv \sum_{q=1}^{r} \frac{\sqrt{\beta_q}}{n}\mathbf{v}_q(\mathbf{Z}^{\mathsf{T}}\mathbf{u}_q)_{\mathsf{Q}}^{\mathsf{T}}\,. \tag{121}$$

Next notice that, for any $x \in \mathbb{R}$,

$$\big|\eta(x) - x\big| \le \frac{\tau}{\sqrt{n}}\,. \tag{122}$$

With a view to employing this inequality, we use Eq. (119) and the triangle inequality:

$$\big\|\mathcal{P}_{\mathsf{E}}(\eta(\widehat{\mathbf{\Sigma}})) - \mathbf{\Sigma}\big\|_{op} = \Big\|\mathcal{P}_{\mathsf{E}}\big(\eta(\widehat{\mathbf{\Sigma}})\big) - \mathcal{P}_{\mathsf{E}}\big(\widehat{\mathbf{\Sigma}}\big) - \mathbf{\Delta}_1 - \mathbf{\Delta}_2 - \mathbf{\Delta}_2^{\mathsf{T}} - \mathcal{P}_{\mathsf{E}}\Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}_p\Big)\Big\|_{op} \tag{123}$$

$$\le \big\|\mathcal{P}_{\mathsf{E}}\big(\eta(\widehat{\mathbf{\Sigma}}) - \widehat{\mathbf{\Sigma}}\big)\big\|_{op} + \|\mathbf{\Delta}_1\|_{op} + 2\|\mathbf{\Delta}_2\|_{op} + \Big\|\mathcal{P}_{\mathsf{E}}\Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}_p\Big)\Big\|_{op} \tag{124}$$

$$\le \frac{s_0\tau}{\sqrt{n}} + \|\mathbf{\Delta}_1\|_{op} + 2\|\mathbf{\Delta}_2\|_{op} + \Big\|\mathcal{P}_{\mathsf{E}}\Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}_p\Big)\Big\|_{op}, \tag{125}$$

where the last line follows by noticing that the first term is supported on E of size $s_0 \times s_0$ and then using bias bound Eq. (122) entry-wise. We next bound each of the three terns on the right hand side.

For the first term in Eq (125), note that with a change of basis to the orthonormal set $\mathbf{v}_1, \dots \mathbf{v}_r$ $\mathbf{\Delta}_1$ is equivalent to an $r \times r$ matrix with entries $M_{qq'}\sqrt{\beta_q\beta_{q'}}$, where $M_{qq'} = \big(\langle \mathbf{u}_q, \mathbf{u}_q'\rangle/n - \mathbf{1}_{q=q'}\big)$. Denote by $\mathbf{B} \in \mathbb{R}^{r\times r}$ the diagonal matrix with $B_{qq} = \sqrt{\beta_q}$ and by $\mathbf{U} \in \mathbb{R}^{r\times n}$, the matrix with columns $\mathbf{u}_1, \dots \mathbf{u}_r$. Then, we have, with high probability

$$\|\mathbf{\Delta}_1\|_{op} = \|\mathbf{B}\mathbf{M}\mathbf{B}\|_{op} \tag{126}$$

$$\le \|\mathbf{B}\|_{op}^2\|\mathbf{M}\|_{op} = \beta\big\|\frac{1}{n}\mathbf{U}^{\mathsf{T}}\mathbf{U} - \mathrm{I}_{r\times r}\big\|_{op} \tag{127}$$

$$\le C\beta\sqrt{\frac{r}{n}}\,. \tag{128}$$

The last inequality follows from the Bai-Yin law on eigenvalues of Wishart matrices (see Vershynin, 2012, Corollary 5.35).

Consider the second term in Eq (125). By orthonormality of $\mathbf{v}_1, \dots, \mathbf{v}_r$, the matrix $\mathbf{\Delta}_2$ is orthogonally equivalent to $\mathbf{B}\mathbf{Z}_{\mathsf{Q}}^{\mathsf{T}}\mathbf{U}/n$, where we recall that $\mathbf{Z}_{\mathsf{Q}}$ denotes the submatrix of $\mathbf{Z}$ formed by the columns in Q. Denoting by $\mathbf{P}_U$ the orthogonal projector onto the column space of $\mathbf{U}$, we then have, with high probability,

$$\|\mathbf{\Delta}_2\|_{op} \leq \frac{1}{n}\,\|\mathbf{B}\|_{op}\|\mathbf{Z}_{\mathsf{Q}}^{\mathsf{T}}\mathbf{P}_U\mathbf{U}\|_{op} \tag{129}$$

$$\leq \frac{\beta}{n}\,\|\mathbf{P}_U\mathbf{Z}_{\mathsf{Q}}\|_{op}\|\mathbf{U}\|_{op} \tag{130}$$

$$\leq \frac{C\beta}{n}\,\big(\sqrt{s_0}+\sqrt{r}\big)\big(\sqrt{n}+\sqrt{r}\big) \leq C\beta\sqrt{\frac{s_0}{n}}\,. \tag{131}$$

Here the penultimate inequality follows by Lemma 11 noting that, by invariance under rotations (and since $\mathbf{P}_U$ project onto a random subspace of r dimensions independent of $\mathbf{Z}$), $\|\mathbf{P}_U\mathbf{Z}_{\mathsf{Q}}\|_{op}$ is distributed as the norm of a matrix with i.i.d. standard normal entries, with dimensions $|\mathsf{Q}|\times r$, $|\mathsf{Q}|\leq s_0$.

Finally, for the third term of Eq. (125) we use the Bai-Yin law of Wishart matrices (see Vershynin, 2012, Corollary 5.35) to obtain, with high probability:

$$\Big\|\mathcal{P}_{\mathsf{E}}\Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z}-\mathrm{I}_p\Big)\Big\|_{op} = \Big\|\frac{1}{n}\mathbf{Z}_{\mathsf{Q}}^{\mathsf{T}}\mathbf{Z}_{\mathsf{Q}}-\mathrm{I}_{s_0}\Big\|_{op} \tag{132}$$

$$\leq C\sqrt{\frac{s_0}{n}}, \tag{133}$$

Finally, substituting the above bounds in Eq. (125), we get

$$\big\|\mathcal{P}_{\mathsf{E}}(\eta(\widehat{\mathbf{\Sigma}}))-\mathbf{\Sigma}\big\|_{op} = \frac{\tau s_0}{\sqrt{n}}+C(1+\beta)\sqrt{\frac{s_0}{n}}, \tag{134}$$

which implies the proposition.

6.3 Proof of Proposition 15

Note that $\mathbf{C}=\bar{\mathbf{C}}+\bar{\mathbf{C}}^{\mathsf{T}}$ where $\bar{\mathbf{C}}=\mathcal{P}_{\mathsf{Q}\times\mathsf{Q}^c}\big(\eta(\widehat{\mathbf{\Sigma}})\big)$. It is therefore sufficient to control $\bar{\mathbf{C}}$, and then use triangle inequality. The proof is similar to that of Proposition 13. We let $\mathbf{U}\in\mathbb{R}^{n\times r}$ denote the matrix with columns $\mathbf{u}_1$, $\mathbf{u}_2$,… $\mathbf{u}_r$, and introduce the set

$$\mathcal{U}\equiv\Big\{\mathbf{U}\in\mathbb{R}^{n\times r}:\ \ \Big\|\frac{1}{n}\mathbf{U}^{\mathsf{T}}\mathbf{U}-\mathrm{I}_{r\times r}\Big\|_{op}\leq 5\sqrt{\frac{r}{n}}\Big\}\,. \tag{135}$$

We then have

$$\mathbb{P}\big(\|\bar{\mathbf{C}}\|_{op}\geq\Delta\big)\leq\sup_{\mathbf{U}\in\mathcal{U}}\mathbb{P}\big(\|\bar{\mathbf{C}}\|_{op}\geq\Delta\,\big|\,\mathbf{U}\big)+\mathbb{P}\big(\mathbf{U}\not\in\mathcal{U}\big)\,. \tag{136}$$

Notice that, by the Bai-Yin law on eigenvalues of Wishart matrices (see Vershynin, 2012, Corollary 5.35), $\lim_{n\to\infty}\mathbb{P}(\mathbf{U}\in\mathcal{U})=1$ (throughout $r<c\,n$ for c a small constant). It is therefore sufficient to show $\sup_{\mathbf{U}\in\mathcal{U}}\mathbb{P}\big(\|\bar{\mathbf{C}}\|_{op}\geq\Delta\,\big|\,\mathbf{U}\big)\to 0$ for Δ as in the statement of the theorem.

In order to lighten the notation, we will write $\widetilde{\mathbb{P}}(\,\cdot\,)\equiv\mathbb{P}(\,\cdot\,|\mathbf{U})$ and bound the above probability uniformly over $\mathbf{U}\in\mathcal{U}$. (In other words $\widetilde{\mathbb{P}}$ denotes expectation over $\mathbf{Z}$ with $\mathbf{U}$ fixed). We first control the norms of small submatrices of $\bar{\mathbf{C}}$, following which we control the full matrix.

Lemma 22 *Fix an $A \in [1, p^{1/3}]$, and let $L = \sqrt{((\beta \vee 1)n + p)/n^2}$. Then, there exists an absolute constant $C > 0$ such that, for any $\Delta > 0$:*

$$\widetilde{\mathbb{P}}\Big\{ \max_{\mathsf{Q}^c \supseteq \mathsf{S}: |\mathsf{S}| \le p/A} \|\mathcal{P}_{\mathsf{Q}\times\mathsf{S}}\big(\eta(\widehat{\mathbf{\Sigma}})\big)\|_{op} \ge \Delta \Big\} \le C \exp\Big(C s_0 + \frac{p\log(CA)}{A} - \frac{\Delta^2}{CL^2}\Big) \\ + L^{-2}(np)^C \exp(-n/C). \tag{137}$$

Proof Let, as before, $T_p^\varepsilon(\mathsf{S})$ denote the ε-net of unit vectors supported on $\mathsf{S} \subset \mathsf{Q}^c$ of size at most p/A and let $T = \cup_{\mathsf{S}} T_p^\varepsilon(\mathsf{S})$. Then, by Lemma 8, with $\varepsilon = 1/4$:

$$\widetilde{\mathbb{P}}\Big\{ \max_{\mathsf{S}\subseteq \mathsf{Q}^c |\mathsf{S}| \le p/A} \big\|\mathcal{P}_{\mathsf{Q}\times\mathsf{S}}\big(\eta(\widehat{\mathbf{\Sigma}})\big)\big\|_{op} \ge \Delta \Big\} \le \widetilde{\mathbb{P}}\Big\{ \max_{\mathbf{y}\in T, \mathbf{w}\in T_{s_0}^\varepsilon} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle \ge \Delta(1-2\varepsilon)/2 \Big\}. \tag{138}$$

It now suffices to control the right hand side via Lemma 9. We first compute the gradients with respect to $\tilde{\mathbf{z}}_\ell$ as before:

$$\nabla_{\tilde{\mathbf{z}}_\ell} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle = \begin{cases} \frac{w_\ell}{n} \sum_{i\in\mathsf{Q}^c} y_i \partial\eta(\langle \tilde{\mathbf{x}}_\ell, \tilde{\mathbf{z}}_i\rangle/n)\tilde{\mathbf{z}}_j & \text{when } \ell \in \mathsf{Q}, . \\ \frac{y_\ell}{n} \sum_{i\in\mathsf{Q}} w_i \partial\eta(\langle \tilde{\mathbf{z}}_\ell, \tilde{\mathbf{x}}_i\rangle/n)\tilde{\mathbf{x}}_i & \text{when } \ell \in \mathsf{Q}^c, \end{cases} \tag{139}$$

Therefore, arguing as in proof of Proposition 13 (see Lemma 20):

$$\|\nabla_{\mathbf{Z}} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle\|_F^2 = \sum_\ell \|\nabla_{\tilde{\mathbf{z}}_\ell}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle\|^2 \le \frac{\|\mathbf{Z}\|^2 + \|\mathbf{X}_{\mathsf{Q}}\|^2}{n^2}\,. \tag{140}$$

Let $\mathbf{B} \in \mathbb{R}^{r\times r}$ be the diagonal matrix with entries $B_{q,q} = \sqrt{\beta_q}$, and $\mathbf{V} \in \mathbb{R}^{p\times r}$ be the matrix with columns $\mathbf{v}_1, \dots, \mathbf{v}_r$. We then have $\mathbf{X} = \mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}} + \mathbf{Z}$, whence, recalling $\mathbf{U} \in \mathcal{U}$, and $r \le c\,n$ with c small enough

$$\|\mathbf{X}_{\mathsf{Q}}\| \le \|\mathbf{X}\| \le \|\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\| + \|\mathbf{Z}\| \tag{141}$$
$$\le \sqrt{\beta}\|\mathbf{U}\| + \|\mathbf{Z}\| \le 5\sqrt{\beta n} + \|\mathbf{Z}\|\,. \tag{142}$$

Consider the good set $\mathcal{G}_4$ of pairs $\big(\mathbf{Z}, \mathbf{Z}'\big)$ satisfying:

$$\max(\|\mathbf{Z}\|, \|\mathbf{Z}'\|) \le \sqrt{2n} + \sqrt{2p}\,, \tag{143}$$
$$\max(\|\mathbf{Z}_{\mathsf{Q}}\|, \|\mathbf{Z}'_{\mathsf{Q}}\|) \le \sqrt{2n} + \sqrt{2k}\,. \tag{144}$$

For $\big((\mathbf{Z}, \mathbf{Z}'\big) \in \mathcal{G}_4$, and $t \in [0,1]$, define $\mathbf{Z}(t) = \sqrt{t}\mathbf{Z} + \sqrt{1-t}\mathbf{Z}'$. Now Using Eqs. (140) and (142, the gradient $\nabla\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle$ evaluated at $\mathbf{Z}(t)$ satisfies:

$$\|\nabla\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle\|^2 \le \frac{3\|\mathbf{Z}(t)\|^2 + 10\beta n}{n^2} \tag{145}$$
$$\le C\frac{(n+p) + \beta n}{n^2} \tag{146}$$
$$\le C\frac{(\beta\vee 1)n + p}{n^2}. \tag{147}$$

Now applying Corollary 10, for $L = C\sqrt{((\beta \vee 1)n + p)/n^2}$:

$$\widetilde{\mathbb{P}}\Big\{ \max_{\mathsf{S}\subseteq \mathsf{Q}^c |\mathsf{S}|\le p/A} \big\|\mathcal{P}_{\mathsf{Q}\times\mathsf{S}}\big(\eta(\widehat{\mathbf{\Sigma}})\big)\big\|_{op} \ge \Delta \Big\} \le C|T| \exp\Big(- \frac{\Delta^2}{CL^2}\Big) + CL^{-2}\widetilde{\mathbb{E}}\{\max_{\mathbf{w},\mathbf{y}} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle^4\}^{1/4}\mathbb{P}\{\mathcal{G}_4\}^{1/2}. \tag{148}$$

Let $\varepsilon = 1/4$, observing that $T \subseteq \cup_{\mathsf{S}:|\mathsf{S}|\le p/A} T_p^{\varepsilon}(\mathsf{S})$, we have the bound (using Lemma 7 and Stirling's approximation):

$$|T| \le \exp(Cs_0 + A^{-1}p\log CA), \tag{149}$$

for some absolute C. Now, as in the proof of Proposition 13, $|\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle| \le \|\mathbf{C}\| \le \|\mathbf{C}\|_F \le \|\widehat{\mathbf{\Sigma}}\|_F$. From this it follows that $\widetilde{\mathbb{E}}\big\{ \max_{\mathbf{w},\mathbf{y}} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle^4 \big\} \le (np)^C$ for some C. Finally $\mathbb{P}\{\mathcal{G}_4^c\} \le \exp(-cn)$ using Lemmas 11, 12 and the union bound. Combining these bounds in Eq. (148) yields the lemma. ■

Now we prove a similar lemma when $\mathbf{y}$ has entries that are "spread out".

Lemma 23 *Fix an $A \in [1, p^{1/3}]$, and a unit vector $\mathbf{y} \in \mathbb{R}^{\mathsf{Q}^c}$ let $\mathsf{S} = \{i : |y_i| \le \sqrt{A/p}$, and $\mathbf{y}_{\mathsf{S}}$ denote the projection of $\mathbf{y}$ on the set of indices S. Then there exists a numerical constant C such that, assuming $\tau \le \sqrt{\log p}/2$, we have*

$$\widetilde{\mathbb{P}}\Big\{ \max_{\mathbf{w}\in T^{\varepsilon}_{\mathsf{Q}}, \mathbf{y}\in T^{\varepsilon}_{\mathsf{Q}^c}} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}_{\mathsf{S}}\rangle \ge \Delta \Big\} \quad \le C\exp\Big(- \frac{\Delta^2}{CL_*^2} + Cp\Big) + (np)^C \exp\big(- c\min(\sqrt{p}, n)\big), \tag{150}$$

where $L_* = \sqrt{A\exp(-\tau^2/C(\beta\vee 1))(n(\beta\vee 1) + p)/n^2}$.

Proof For simplicity of notation, it is convenient to introduce the vector $\mathbf{y}' = \mathbf{y}_{\mathsf{S}}$. Throughout the proof, we will use that $\|\mathbf{y}'\| \le 1$ and $\|\mathbf{y}'\|_\infty \le \sqrt{A/p}$. We compute the gradients as follows:

$$\nabla_{\tilde{\mathbf{z}}_\ell} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle = \begin{cases} \frac{w_\ell}{n} \sum_{i\in\mathsf{Q}^c} y'_i \partial\eta(\langle \tilde{\mathbf{x}}_\ell, \tilde{\mathbf{z}}_i\rangle/n)\tilde{\mathbf{z}}_j & \text{when } \ell\in\mathsf{Q} \\ \frac{y'_\ell}{n} \sum_{i\in\mathsf{Q}} w_i \partial\eta(\langle \tilde{\mathbf{z}}_\ell, \tilde{\mathbf{x}}_i\rangle/n)\tilde{\mathbf{x}}_i & \text{when } \ell\in\mathsf{Q}^c\,. \end{cases} \tag{151}$$

Therefore we have

$$\begin{aligned} \sum_{\ell\in\mathsf{Q}} \|\nabla_{\tilde{\mathbf{z}}_\ell} \langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle\|^2 &\le \sum_{\ell\in\mathsf{Q}} \frac{w_\ell^2}{n^2}\|\mathbf{Z}\|^2 \sum_{i\in\mathsf{Q}^c} \big(y'_i\partial\eta(\tilde{\mathbf{x}}_\ell, \tilde{\mathbf{z}}_\ell)\big)^2 &(152)\\ &\le \frac{A\|\mathbf{Z}\|^2}{pn^2} \max_{\ell\in\mathsf{Q}} \sum_{i\in\mathsf{Q}^c} \partial\eta(\langle \tilde{\mathbf{x}}_\ell, \tilde{\mathbf{z}}_i\rangle/n), &(153)\end{aligned}$$

where we used the fact that $|y'_i| \le \sqrt{A/p}$ and that $\partial\eta(\cdot) \in \{0,1\}$. Similarly, for $\ell \in \mathsf{Q}^c$:

$$\sum_{\ell\in\mathsf{Q}^c} \|\nabla_{\tilde{\mathbf{z}}_\ell}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle\|^2 \le \sum_{\ell\in\mathsf{Q}^c} \frac{(y'_i)^2\|\mathbf{X}_{\mathsf{Q}}\|^2}{n^2} \sum_{i\in\mathsf{Q}} \big(w_i\partial\eta(\langle\tilde{\mathbf{z}}_\ell, \tilde{\mathbf{x}}_i\rangle/n)\big)^2 \tag{154}$$

$$= \sum_{i\in\mathsf{Q}} \frac{w_i^2\|\mathbf{X}_{\mathsf{Q}}\|^2}{n^2} \sum_{\ell\in\mathsf{Q}^c} (y'_\ell)^2 \partial\eta(\langle\tilde{\mathbf{z}}_\ell, \tilde{\mathbf{x}}_i\rangle/n)^2 \tag{155}$$

$$\le \frac{A\|\mathbf{X}_{\mathsf{Q}}\|^2}{pn^2} \max_{\ell\in\mathsf{Q}} \sum_{i\in\mathsf{Q}^c} \partial\eta(\langle\tilde{\mathbf{z}}_i, \tilde{\mathbf{x}}_\ell\rangle/n). \tag{156}$$

Combining the bounds in Eqs.(153), (156), we obtain

$$\|\nabla_{\mathbf{Z}}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle\|_F^2 = \sum_{\ell\in[p]} \|\nabla_{\tilde{\mathbf{z}}_\ell}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle\|^2 \tag{157}$$

$$\le \frac{2A}{pn^2}(\|\mathbf{X}_{\mathsf{Q}}\|^2 + \|\mathbf{Z}\|^2) \max_{i\in\mathsf{Q}} \sum_{j\in\mathsf{Q}^c} \partial\eta(\langle\tilde{\mathbf{x}}_i, \tilde{\mathbf{z}}_j\rangle/n). \tag{158}$$

With $K = C\beta \vee 1$, we define the good set $\mathcal{G}_5$ of pairs $(\mathbf{Z}, \mathbf{Z}')$ satisfying

$$\|\mathbf{Z}\|, \|\mathbf{Z}'\| \le \sqrt{2n} + \sqrt{2p} \tag{159}$$

$$\forall i \in \mathsf{Q}, \quad \frac{1}{p}\sum_{j\in\mathsf{Q}^c} \mathbb{I}(\langle\tilde{\mathbf{x}}_i, \tilde{\mathbf{z}}_j\rangle \ge \tau\sqrt{n}/2) \le 2\exp(-\tau^2/K) \tag{160}$$

$$\forall i \in \mathsf{Q}, \quad \frac{1}{p}\sum_{j\in\mathsf{Q}^c} \mathbb{I}(\langle\tilde{\mathbf{x}}'_i, \tilde{\mathbf{z}}'_j\rangle \ge \tau\sqrt{n}/2) \le 2\exp(-\tau^2/K) \tag{161}$$

$$\forall i \in \mathsf{Q}, \quad \frac{1}{p}\sum_{j\in\mathsf{Q}^c} \mathbb{I}(\langle\tilde{\mathbf{x}}'_i, \tilde{\mathbf{z}}_j\rangle \ge \tau\sqrt{n}/4) \le 2\exp(-\tau^2/K) \tag{162}$$

$$\forall i \in \mathsf{Q}, \quad \frac{1}{p}\sum_{j\in\mathsf{Q}^c} \mathbb{I}(\langle\tilde{\mathbf{x}}_i, \tilde{\mathbf{z}}'_j\rangle \ge \tau\sqrt{n}/4) \le 2\exp(-\tau^2/K). \tag{163}$$

Define $\mathbf{Z}(t) = \sqrt{t}\mathbf{Z} + \sqrt{1-t}\mathbf{Z}'$ with $(\mathbf{Z}, \mathbf{Z}') \in \mathcal{G}_5$. By Eq. (158) the gradient evaluated at $\mathbf{Z}(t)$ is bounded by

$$\|\nabla\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle\|^2 \le \frac{2A}{pn^2}(\|\mathbf{X}_{\mathsf{Q}}(t)\|^2 + \|\mathbf{Z}(t)\|^2) \max_{i\in\mathsf{Q}} \sum_{j\in\mathsf{Q}^c} \partial\eta(\langle\tilde{\mathbf{x}}(t)_i), \tilde{\mathbf{z}}(t)_j\rangle/n) \tag{164}$$

$$\le \frac{CA}{pn^2}((\beta\vee 1)n + p) \max_{i\in\mathsf{Q}} \sum_{j\in\mathsf{Q}^c} \partial\eta(\langle\tilde{\mathbf{x}}(t)_i), \tilde{\mathbf{z}}(t)_j\rangle/n)\,, \tag{165}$$

where we bounded $\|\mathbf{X}_{\mathsf{Q}}(t)\|$ as in Eq. (142), and used $\|\mathbf{Z}(t)\|_{op} \le 2(\sqrt{n}+\sqrt{p})$, which follows from Eq. (159) and triangle inequality. Furthermore, as $\langle\tilde{\mathbf{x}}(t)_i, \tilde{\mathbf{z}}(t)_j\rangle = t\langle\tilde{\mathbf{x}}_i, \tilde{\mathbf{z}}_j\rangle + (1-t)\langle\tilde{\mathbf{x}}'_i, \tilde{\mathbf{z}}'_j\rangle + \sqrt{t(1-t)}(\langle\tilde{\mathbf{x}}_i, \tilde{\mathbf{z}}'_j\rangle + \langle\tilde{\mathbf{x}}'_i, \tilde{\mathbf{z}}_j\rangle)$, we have that:

$$\partial\eta(\langle\tilde{\mathbf{x}}(t)_i), \tilde{\mathbf{z}}(t)_j\rangle/n) = \mathbb{I}(|\langle\tilde{\mathbf{x}}(t)_i, \tilde{\mathbf{z}}(t)_j\rangle| \ge \tau\sqrt{n}) \tag{166}$$

$$\le \mathbb{I}(|\langle\tilde{\mathbf{x}}_i, \tilde{\mathbf{z}}_j\rangle| \ge \tau\sqrt{n}/2) + \mathbb{I}(|\langle\tilde{\mathbf{x}}'_i, \tilde{\mathbf{z}}'_j\rangle| \ge \tau\sqrt{n}/2)$$
$$+ \mathbb{I}(|\langle\tilde{\mathbf{x}}'_i, \tilde{\mathbf{z}}_j\rangle| \ge \tau\sqrt{n}/4) + \mathbb{I}(|\langle\tilde{\mathbf{x}}'_i, \tilde{\mathbf{z}}_j\rangle| \ge \tau\sqrt{n}/4). \tag{167}$$

Hence on the good set $\mathcal{G}_5$, we have:

$$\max_{i\in \mathsf{Q}} \sum_{j\in \mathsf{Q}^c} \partial\eta(\langle \tilde{\mathbf{x}}(t)_i), \tilde{\mathbf{z}}(t)_j\rangle/n) \le 4p\, e^{-\tau^2/K}\,. \tag{168}$$

Therefore the gradient satisfies, on the good set:

$$\|\nabla_{\mathbf{Z}}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle\|^2 \le C\frac{A}{n^2}((\beta\vee 1)n + p)\, e^{-\tau^2/K} = CL_*^2\,. \tag{169}$$

Hence, by Lemma 9, we obtain:

$$\begin{aligned}\widetilde{\mathbb{P}}\Big\{\max_{\mathbf{w}\in T^\varepsilon_{\mathsf{Q}},\mathbf{y}\in T^\varepsilon_p}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle \ge \Delta\Big\} \le & C|T^\varepsilon_{\mathsf{Q}}||T^\varepsilon_p|\exp\Big(-\frac{\Delta^2}{CL_*^2}\Big) \\ & + CL_*^{-2}\widetilde{\mathbb{E}}\{\max\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle^4\}^{1/4}\mathbb{P}\{\mathcal{G}_5^c\}^{1/2}\,.\end{aligned} \tag{170}$$

By Lemma 7, keeping $\varepsilon = 1/4$ we have that the first term is at most $C\exp(Cp+\exp(-\Delta^2/CL_*^2))$. For the second term, we have $|\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle| \le \|\bar{\mathbf{C}}\| \le \|\bar{\mathbf{C}}\|_F \le \|\widehat{\mathbf{\Sigma}}\|_F$. Since $\mathbb{E}\{\|\widehat{\mathbf{\Sigma}}\|_F^4\} \le (np)^C$, we have that $\mathbb{E}\{\max_{\mathbf{w},\mathbf{y}}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}\rangle^4\}^{1/4} \le (np)^C$. Also as $\tau < \sqrt{\log p}$, $L_* \ge (np)^{-C}$, implying that the second term is bounded above by $(np)^C\mathbb{P}\{\mathcal{G}_5^c\}^{1/2}$. Therefore:

$$\widetilde{\mathbb{P}}\Big\{\max_{\mathbf{w}\in T^\varepsilon_{\mathsf{Q}},\mathbf{y}\in T^\varepsilon_p}\langle \mathbf{w}, \bar{\mathbf{C}}\mathbf{y}'\rangle \ge \Delta\Big\} \le C\exp\Big(Cp - \frac{\Delta^2}{CL_*^2}\Big) + (np)^C\mathbb{P}\{\mathcal{G}_5^c\}^{1/2}\,. \tag{171}$$

It remains to control the probability of the bad set $\mathcal{G}_5^c$. For this, we control the probability of violating any one condition among (159), (160), (161), (162) and (163) defining $\mathcal{G}_5$ and then use the union bound. By Lemmas 11, condition (159) hold with probability $1 - C\exp(-cn)$. The argument controlling the probability for conditions (160), (161), (162) and (163) to hold are essentially the same, so we restrict ourselves to condition (160) keeping $i = 1 \in \mathsf{Q}$, without loss of generality. Conditional on $\tilde{\mathbf{x}}_1$, $\langle \tilde{\mathbf{x}}_1, \tilde{\mathbf{z}}_j\rangle$ for $j \in \mathsf{Q}^c$ are independent $\mathsf{N}(0, \|\tilde{\mathbf{x}}_1\|^2)$ variables. Therefore, conditional on $\tilde{\mathbf{x}}_1$, $\mathbb{I}(|\langle \tilde{\mathbf{x}}_1, \tilde{\mathbf{z}}_j\rangle| \ge \tau\sqrt{n}/2)$ are independent Bernoulli random variables with success probability $\Phi\{-\tau\sqrt{n}/2\|\tilde{\mathbf{x}}_1\|\}$. Define h_1 to be the success probability, i.e. $h_1 = \Phi(-\tau\sqrt{n}/(2\|\tilde{\mathbf{x}}_1\|))$.

Since $K = C(\beta \vee 1)$ we can enlarge C to a large absolute constant. Letting $\mathbf{V} \in \mathbb{R}^{n\times r}$ be the matrix with columns $\mathbf{v}_1, \dots, \mathbf{v}_r$, and $\mathbf{B}$ the diagonal matrix with $B_{q,q} = \sqrt{\beta_q}$, we have, with probability at least $1 - \exp(-n/C)$,

$$\|\tilde{\mathbf{x}}_1\| \le \|\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\mathbf{e_1}\| + \|\tilde{\mathbf{z}}_1\| \le \|\mathbf{B}\|\|\mathbf{U}\| + \|\tilde{\mathbf{z}}_1\| \le \sqrt{\frac{Kn}{4}}\,, \tag{172}$$

where the last equality holds since $\mathbf{U} \in \mathcal{U}$ and by tail bounds on chi-squared random variables. Further

$$\begin{aligned}\widetilde{\mathbb{P}}\Big\{\sum_{j\in\mathsf{Q}^c}\mathbb{I}(|\langle \tilde{\mathbf{x}}_1, \tilde{\mathbf{z}}_j\rangle \ge \tau\sqrt{n}/2) \ge |\mathsf{Q}^c|h\Big\} \le & \widetilde{\mathbb{P}}\{\|\tilde{\mathbf{x}}_1\|^2 \ge Kn\} \\ & + \sup_{\|\tilde{\mathbf{x}}_1\|^2\le Kn}\widetilde{\mathbb{P}}\Big\{\sum_{j\in\mathsf{Q}^c}\mathbb{I}(|\langle \tilde{\mathbf{x}}_1, \tilde{\mathbf{z}}_j\rangle \ge \tau\sqrt{n}/2) \ge |\mathsf{Q}^c|h \;\Big|\; \tilde{\mathbf{x}}_1\Big\}.\end{aligned} \tag{173}$$

By the above argument, the first term is at most $\exp(-n/C)$ and we turn to the second term. By the Chernoff bound

$$\widetilde{\mathbb{P}}\Big\{\sum_{j\in\mathsf{Q}^c}\mathbb{I}(|\langle\tilde{\mathbf{x}}_1,\tilde{\mathbf{z}}_j\rangle\geq\tau\sqrt{n}/2)\geq|\mathsf{Q}^c|h\ \Big|\tilde{\mathbf{x}}_1\Big\}\leq\exp\big(-|\mathsf{Q}^c|D(h\|h_1)\big),, \tag{174}$$

with $h_1<\exp(-\tau^2/K)$ when $\|\tilde{\mathbf{x}}_1\|^2\leq Kn/4$. Choosing $h=2\exp(-\tau^2/K)$ implies that $h_1\leq h/2$ when and, thereby, that $D(h\|h_1)\geq h/C$. Further since $\tau<\sqrt{\log p}/2$, $h\geq 1/\sqrt{p}$. This implies that

$$\exp(-|\mathsf{Q}^c|D(h-h_1\|h_1))=\exp(-(p-s_0)h/C)\geq\exp(-\sqrt{p}/C). \tag{175}$$

Combining this with Eq. (173) we have that $\mathbb{P}\{\mathcal{G}^c\}\leq Cp^2\exp(-\min(n,\sqrt{p})/C)$ for some absolute C. Plugging this in Eq. (171) yields the lemma. ■

We are now ready to prove Proposition 15. Indeed, as in Proposition 13, for any unit vector $\mathbf{y}\in\mathbb{R}^{\mathsf{Q}^c}$, let $\mathsf{S}=\{i:|y_i|\geq\sqrt{A/p}\}$ and $\mathbf{y}_{\mathsf{S}},\mathbf{y}_{\mathsf{S}^c}$ denote the projections on the indices in S,S^c respectively.

$$\begin{aligned}\widetilde{\mathbb{P}}\Big\{\|\bar{\mathbf{C}}_1\|\geq\Delta\Big\}&\leq\widetilde{\mathbb{P}}\Big\{\max_{\mathbf{w}\in T^\varepsilon_{\mathsf{Q}},\mathbf{y}\in T^\varepsilon_{\mathsf{Q}^c}}|\langle\mathbf{w},\bar{\mathbf{C}}\mathbf{y}\rangle|\geq\Delta(1-2\varepsilon)\Big\} &(176)\\ &\leq\widetilde{\mathbb{P}}\Big\{\max_{\mathbf{w}\in T^\varepsilon_{\mathsf{Q}},\mathbf{y}\in T^\varepsilon_{\mathsf{Q}^c}}|\langle\mathbf{w},\bar{\mathbf{C}}\mathbf{y}_{\mathsf{S}}\rangle|\geq\Delta(1-2\varepsilon)/2\Big\}\\ &\quad+\mathbb{P}\Big\{\max_{\mathbf{w}\in T^\varepsilon_{\mathsf{Q}},\mathbf{y}\in T^\varepsilon_{\mathsf{Q}^c}}|\langle\mathbf{w},\bar{\mathbf{C}}\mathbf{y}_{\mathsf{S}^c}\rangle|\geq\Delta(1-2\varepsilon)/2\Big\}. &(177)\end{aligned}$$

As before, we will let $\varepsilon=1/4$. The first term is controlled via Lemma 22, while the second is controlled by Lemma 23. We keep $\Delta=\Delta_*$ where

$$\Delta_*=C\Big(L_*\sqrt{p}+L\sqrt{\frac{p\log A}{A}}\Big). \tag{178}$$

so that, via the bounds of Lemmas 22, 23 and that $s_0^2\leq p$:

$$\mathbb{P}\{\|\mathbf{C}_1\|\geq\Delta_*\}\leq C\exp\Big(-c\frac{p\log A}{A}\Big)+L_*^{-2}(np)^C\exp\big(-c\min(\sqrt{p},n)\big). \tag{179}$$

We now set $A=\big((\tau^2/K)\exp(\tau^2/K)\big)^{1/2}$ with $K=C(\beta\vee 1)$ for a suitable constant C and, since $\tau\leq\sqrt{\log p}/2$, we get that $A\leq p^{1/3}$. Furthermore, it is straightforward to see that $L\geq(np)^{-C}$, and this implies that

$$\mathbb{P}\{\|\mathbf{C}_1\|\geq\Delta_*\}\leq(np)^C\exp(-c\min(\sqrt{p},n))=o(1). \tag{180}$$

With this setting of A, we get the form of Δ_* below, as required for the proposition.

$$\begin{aligned}\Delta_*&\leq C\,e^{-c\tau^2/K}\sqrt{\frac{\tau^2\vee 1}{K}\cdot\frac{pn(\beta\vee 1)+p^2}{n^2}} &(181)\\ &\leq C\,(\tau\vee 1)e^{-c\tau^2/K}\sqrt{\frac{p}{n}}\vee\frac{p}{n}. &(182)\end{aligned}$$

6.4 Proof of Proposition 16

Since $\mathbf{D}$ is a diagonal matrix, its spectral norm is bounded by the maximum of its entries. This is easily done as, for every $i \in \mathsf{Q}^c$:

$$|(\mathbf{D})_{ii}| = \left|\eta\Big(\frac{\|\tilde{\mathbf{z}}_i\|^2}{n} - 1; \frac{\tau}{\sqrt{n}}\Big)\right| \tag{183}$$

$$\leq \left|\frac{\|\tilde{\mathbf{z}}_i\|^2 - n}{n}\right|. \tag{184}$$

By the Chernoff bound for χ^2-squared random variables as in Lemma 12 followed by the union bound, with probability $1 - o(1)$:

$$\max_i \left|\frac{\|\tilde{\mathbf{z}}_i\|^2}{n} - 1\right| \leq C\sqrt{\frac{\log p}{n}} \tag{185}$$

for some absolute C. Here we used the fact that $(\log p)/n < 1$.

6.5 Proof of Proposition 17

It suffices to show that with probability $1 - o(1)$

$$\max_{i,j \in \mathsf{F}\cup\mathsf{G}} |\widehat{\Sigma}_{ij}| \leq \frac{\tau}{\sqrt{n}} = C_0(\beta \vee 1)\sqrt{\frac{\log p}{n}}. \tag{186}$$

This is a standard argument (see Bickel and Levina, 2008b, Lemma A.3) where (following the dependence on β) it suffices to take $\tau \geq C_0(\beta \vee 1)\sqrt{\log p}$ for C_0 a sufficiently large absolute constant. We note here that the same can also be proved via the conditioning technique applied in the proofs of Propositions 13 and 15.

7. Proof of Theorems 3

Throughout this section, to lighten notation, we drop the prime from $\widehat{\mathbf{\Sigma}}'$ and $\mathbf{X}'$ while keeping in mind that these are independent from $\widehat{\mathbf{v}}_1, \dots, \widehat{\mathbf{v}}_r$. We further write $\mathbf{X} = \mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}} + \mathbf{Z}$, where $\mathbf{U} \in \mathbb{R}^{n\times r}$ is the matrix with columns $\mathbf{u}_1, \dots, \mathbf{u}_r$, $\mathbf{B}$ is diagonal with $B_{ii} = \sqrt{\beta_i}$ and $\mathbf{V} \in \mathbb{R}^{p\times r}$ has columns $\mathbf{v}_1, \dots, \mathbf{v}_r$.

Define the event

$$\mathcal{U} \equiv \Big\{\mathbf{U} \in \mathbb{R}^{n\times r}: \ \Big\|\frac{1}{n}\mathbf{U}^{\mathsf{T}}\mathbf{U} - \mathrm{I}_{r\times r}\Big\|_{op} \leq 3\sqrt{\frac{r}{n}}\Big\}. \tag{187}$$

By the Bai-Yin law on eigenvalues of Wishart matrices (see Vershynin, 2012), $\lim_{n\to\infty}\mathbb{P}(\mathbf{U} \in \mathcal{U}) = 1$. In the rest of the proof, we will therefore assume $\mathbf{U} \in \mathcal{U}$ fixed, and denote by $\widetilde{\mathbb{P}}(\,\cdot\,) = \mathbb{P}(\,\cdot\,|\mathbf{U})$ the expectation conditional on $\mathbf{U}$. In other words, $\widetilde{\mathbb{P}}(\,\cdot\,)$ denotes expectation with respect to $\mathbf{Z}$.

Note that

$$\widehat{\mathbf{\Sigma}} = \frac{1}{n}\mathbf{V}\mathbf{B}\mathbf{U}^{\mathsf{T}}\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}} + \frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}} + \frac{1}{n}\mathbf{V}\mathbf{B}\mathbf{U}^{\mathsf{T}}\mathbf{Z} + \frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}\,. \tag{188}$$

We then have, for $q \in \{1, \dots, r\}$ and $i \in \{1, \dots, p\}$,

$$\big|(\widehat{\mathbf{\Sigma}}\widehat{\mathbf{v}}_q)_i - \beta_q\langle \mathbf{v}_q, \widehat{\mathbf{v}}_q\rangle\, v_{q,i}\big| \le T^{(1)}_{i,q} + T^{(2)}_{i,q} + T^{(3)}_{i,q}\,, \tag{189}$$

$$T^{(1)}_{i,q} \equiv \Big|\frac{1}{n}\langle \mathbf{e}_i, \mathbf{V}\mathbf{B}\mathbf{U}^{\mathsf{T}}\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\widehat{\mathbf{v}}_q\rangle - \beta_q\langle \mathbf{v}_q, \widehat{\mathbf{v}}_q\rangle v_{q,i}\Big|\,, \tag{190}$$

$$T^{(2)}_{i,q} \equiv \frac{1}{n}\Big|\langle \mathbf{Z}, \big[(\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\mathbf{e}_i)\widehat{\mathbf{v}}_q^{\mathsf{T}} + (\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\widehat{\mathbf{v}}_q)\mathbf{e}_i^{\mathsf{T}}\big]\rangle\Big|\,, \tag{191}$$

$$T^{(3)}_{i,q} \equiv \Big|\langle \mathbf{e}_i, \Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}\Big)\widehat{\mathbf{v}}_q\rangle\Big|\,. \tag{192}$$

We next bound, with high probability, $\max_{i,q} T^{(a)}_{i,q}$ for $a \in \{1, 2, 3\}$. Throughout we let $\varepsilon \equiv \max_{q\in[r]} \|\widehat{\mathbf{v}}_q - \mathbf{v}_q\|$.

Considering the first term, we have

$$T^{(1)}_{i,q} \le \Big|\langle \mathbf{e}_i, \mathbf{V}\mathbf{B}\Big(\frac{1}{n}\mathbf{U}^{\mathsf{T}}\mathbf{U} - \mathrm{I}\Big)\mathbf{B}\mathbf{V}^{\mathsf{T}}\widehat{\mathbf{v}}_q\rangle\Big| + \Big|\langle \mathbf{e}_i, \mathbf{V}\mathbf{B}^2\mathbf{V}^{\mathsf{T}}\widehat{\mathbf{v}}_q\rangle - \beta_q\langle \mathbf{v}_q, \widehat{\mathbf{v}}_q\rangle v_{q,i}\Big| \tag{193}$$

$$\le 2\beta\sqrt{\frac{r}{n}} + \beta\varepsilon\sqrt{r}\max_{q'\in[r]\setminus q} |v_{q',i}|\,, \tag{194}$$

where in the last inequality we used $\sum_{q'\in[r]\setminus q}\langle \mathbf{v}_{q'}, \widehat{\mathbf{v}}_q\rangle^2 \le 1 - \langle \mathbf{v}_q, \widehat{\mathbf{v}}_q\rangle^2 \le \varepsilon^2/2$.

Consider next the second term. Since $Z_{ij} \sim_{iid} \mathsf{N}(0, 1)$, it follows that $T^{(2)}_{i,q} = |W_{i,q}|$, for $W_{i,q} \sim \mathsf{N}(0, \sigma^2_{i,q})$ a Gaussian random variable with variance

$$\sigma^2_{i,q} = \frac{1}{n^2}\big\|(\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\mathbf{e}_i)\widehat{\mathbf{v}}_q^{\mathsf{T}} + (\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\widehat{\mathbf{v}}_q)\mathbf{e}_i^{\mathsf{T}}\big\|_F^2 \tag{195}$$

$$\le \frac{2}{n^2}\Big\{\|\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\mathbf{e}_i\|^2 + \|\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\widehat{\mathbf{v}}_q\|^2\Big\} \tag{196}$$

$$\le \frac{4}{n^2}\|\mathbf{U}\mathbf{B}\mathbf{V}^{\mathsf{T}}\|_{op}^2 \tag{197}$$

$$\le \frac{4}{n^2}\|\mathbf{U}\|_{op}^2\|\mathbf{B}\|_{op}^2 \le \frac{8\beta^2}{n}\,. \tag{198}$$

By union bound over $i \in [p]$, $q \in [r]$ we obtain

$$\max_{i\in[p],q\in[r]} T^{(2)}_{i,q} \le 8\beta\sqrt{\frac{\log p}{n}}\,. \tag{199}$$

Finally, consider the last term. By rotational invariance of $\mathbf{Z}$, the distribution of $T^{(3)}_{i,q}$ only depends on the angle between $\mathbf{e}_i$ and $\widehat{\mathbf{v}}_q$. Calling this angle ϑ, we have

$$T^{(3)}_{i,q} \stackrel{\mathrm{d}}{=} \Big|\langle \mathbf{e}_1, \Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}\Big)\mathbf{e}_1\rangle\cos\vartheta + \langle \mathbf{e}_1, \Big(\frac{1}{n}\mathbf{Z}^{\mathsf{T}}\mathbf{Z} - \mathrm{I}\Big)\mathbf{e}_2\rangle\sin\vartheta\Big| \tag{200}$$

$$\le \Big|\frac{1}{n}\|\tilde{\mathbf{z}}_1\|^2 - 1\Big| + \Big|\frac{1}{n}\langle \tilde{\mathbf{z}}_1, \tilde{\mathbf{z}}_2\rangle\Big|\,. \tag{201}$$

Both of these terms have Bernstein-type tail bonds, whence

$$\widetilde{\mathbb{P}}\Big(T^{(3)}_{i,q} \ge \frac{t}{\sqrt{n}}\Big) \le 2\,\exp\big\{-c\min(t\sqrt{n}, t^2)\big\}\,. \tag{202}$$

Using $t = C_0\sqrt{\log p}$, and recalling that $n \geq C\log p$ for C a large constant, we obtain $\widetilde{\mathbb{P}}\big(T^{(3)}_{i,q} \geq C_0\sqrt{(\log p)/n}\big) \leq 2\,p^{-10}$. Hence by union bound

$$\max_{i\in[p],q\in[r]} T^{(3)}_{i,q} \leq C_0\sqrt{\frac{\log p}{n}}\,. \tag{203}$$

By putting together Eqs. (194), (199), (203), and using assumption $\mathsf{A2}$, we get

$$\big|(\widehat{\mathbf{\Sigma}}\widehat{\mathbf{v}}_q)_i - \beta_q\langle \mathbf{v}_q, \widehat{\mathbf{v}}_q\rangle\, v_{q,i}\big| \leq C\beta\sqrt{\frac{r}{n}} + C(\beta\vee 1)\sqrt{\frac{\log p}{n}} + \beta\varepsilon\gamma\sqrt{r}\,|v_{q,i}|\,\mathbb{I}(i\in\mathsf{Q})\,. \tag{204}$$

Let $\widehat{\mathsf{Q}}_q = \{i \in [p]\,:\; |(\widehat{\mathbf{\Sigma}}'\widehat{\mathbf{v}}_q)_i| \geq \rho\}$. We claim that the above implies that, with high probability, $\mathsf{Q}_q \subseteq \widehat{\mathsf{Q}}_q \subseteq \mathsf{Q}$ for all q.

For $i \notin \mathsf{Q}$, we have

$$\big|(\widehat{\mathbf{\Sigma}}\widehat{\mathbf{v}}_q)_i\big| \leq C\beta\sqrt{\frac{r}{n}} + C(\beta\vee 1)\sqrt{\frac{\log p}{n}} \tag{205}$$

$$< \frac{\beta_{\min}\theta}{2\sqrt{s_0}}\,, \tag{206}$$

where the last inequality follows from Eq. (14).

On the other hand, By Theorem 2 and using the assumption (14), we can guarantee

$$\varepsilon \leq \frac{1}{8}\Big(\frac{\beta_{\min}}{\beta\gamma\sqrt{r}} \wedge 1\Big)\,. \tag{207}$$

Hence for $i \in \mathsf{Q}_q$, and considering –to be definite– $v_{q,i} > 0$, we get

$$(\widehat{\mathbf{\Sigma}}\widehat{\mathbf{v}}_q)_i \geq \beta_q\langle \mathbf{v}_q, \widehat{\mathbf{v}}_q\rangle\, v_{q,i} - C\beta\sqrt{\frac{r}{n}} - C(\beta\vee 1)\sqrt{\frac{\log p}{n}} - \beta\varepsilon\gamma\sqrt{r}\,|v_{q,i}| \tag{208}$$

$$\geq \beta_{\min}\Big(1-\varepsilon-\frac{\beta}{\beta_{\min}}\varepsilon\gamma\sqrt{r}\Big)v_{q,i} - C\beta\sqrt{\frac{r}{n}} - C(\beta\vee 1)\sqrt{\frac{\log p}{n}} \tag{209}$$

$$\geq \frac{3\beta_{\min}\theta}{4\sqrt{s_0}} - C\beta\sqrt{\frac{r}{n}} - C(\beta\vee 1)\sqrt{\frac{\log p}{n}} \tag{210}$$

$$> \frac{\beta_{\min}\theta}{2\sqrt{s_0}}\,. \tag{211}$$

where, in the first inequality, we used $\langle \mathbf{v}_q, \widehat{\mathbf{v}}_q\rangle \geq 1-\varepsilon$.

This concludes the proof. Keeping track of the dependence on θ, γ, β, $\beta_{\min}$, we get that the following conditions are sufficient for the theorem's conclusion to hold (with C a suitable numerical constant):

$$n \geq C \, \frac{(\beta^2 \vee 1)}{\beta_{\min}^2 \theta^2} \, s_0 \log p \,, \tag{212}$$

$$n \geq C \, \frac{\beta^2}{\beta_{\min}^2 \theta^2} \, r s_0 \,, \tag{213}$$

$$n \geq C \, \left\{ \frac{\beta^4 \vee \beta^2}{\beta_{\min}^2} \gamma^2 \right\} \, r \, s_0^2 \log \frac{p}{s_0^2} \,, \tag{214}$$

$$n \geq C \, \frac{(\beta^2 \vee 1)}{\beta_{\min}^2} \, s_0^2 \log \frac{p}{s_0^2} \,. \tag{215}$$

All of these conditions are implied by the assumptions of Theorem 3, namely Eq. (14). In particular, this is shown by using the fact that $s_0 \log p \leq s_0^2 \log(p/s_0^2)$ for $s_0 \leq \sqrt{p}$.

Acknowledgments

We are grateful to David Donoho for his feedback on an early draft of this manuscript. This work was partially supported by the NSF CAREER award CCF-0743978, the NSF grant CCF-1319979, and the grants AFOSR/DARPA FA9550-12-1-0411 and FA9550-13-1-0036.

References

Arash A Amini and Martin J Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *The Annals of Statistics*, 37(5B):2877–2921, 2009.

Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, pages 1643–1697, 2005.

Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227 (1):494–521, 2011.

Quentin Berthet and Philippe Rigollet. Computational lower bounds for sparse pca. *arXiv preprint arXiv:1304.0828*, 2013.

Peter J Bickel and Elizaveta Levina. Covariance regularization by thresholding. *The Annals of Statistics*, pages 2577–2604, 2008a.

Peter J Bickel and Elizaveta Levina. Regularized estimation of large covariance matrices. *The Annals of Statistics*, pages 199–227, 2008b.

T Tony Cai, Cun-Hui Zhang, Harrison H Zhou, et al. Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics*, 38(4):2118–2144, 2010.

T Tony Cai, Harrison H Zhou, et al. Optimal rates of convergence for sparse covariance matrix estimation. *The Annals of Statistics*, 40(5):2389–2420, 2012.

T Tony Cai, Zongming Ma, Yihong Wu, et al. Sparse pca: Optimal rates and adaptive estimation. *The Annals of Statistics*, 41(6):3074–3110, 2013.

Tony Cai and Weidong Liu. Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*, 106(494):672–684, 2011.

Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1–47, 2009.

Xiuyuan Cheng and Amit Singer. The spectrum of random inner-product kernel matrices. *Random Matrices: Theory and Applications*, 2(04):1350010, 2013.

Alexandre d'Aspremont, Laurent El Ghaoui, Michael I Jordan, and Gert RG Lanckriet. A direct formulation for sparse pca using semidefinite programming. *SIAM review*, 49(3): 434–448, 2007.

Alexandre d'Aspremont, Francis Bach, and Laurent El Ghaoui. Optimal solutions for sparse principal component analysis. *The Journal of Machine Learning Research*, 9:1269–1294, 2008.

Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.

Yash Deshpande and Andrea Montanari. Sparse pca via covariance thresholding. In *Advances in Neural Information Processing Systems*, pages 334–342, 2014.

David L. Donoho and Iain M. Johnstone. Minimax risk over l_p balls. *Prob. Th. and Rel. Fields*, 99:277–303, 1994.

Noureddine El Karoui. On information plus noise kernel random matrices. *The Annals of Statistics*, 38(5):3191–3216, 2010a.

Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1):1–50, 2010b.

Zhou Fan and Andrea Montanari. The spectral norm of random inner-product kernel matrices. arXiv:1507.05343, 2015.

Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large wigner matrices. *Communications in mathematical physics*, 272(1):185–228, 2007.

Iain M. Johnstone. Function estimation and gaussian sequence models. *Unpublished manuscript*, 2015.

Iain M Johnstone and Arthur Yu Lu. Sparse principal components analysis. *Unpublished manuscript*, 2004.

Iain M Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486), 2009.

Noureddine El Karoui. Operator norm consistent estimation of large-dimensional sparse covariance matrices. *The Annals of Statistics*, pages 2717–2756, 2008.

Antti Knowles and Jun Yin. The isotropic semicircle law and deformation of wigner matrices. *Communications on Pure and Applied Mathematics*, 2013.

Robert Krauthgamer, Boaz Nadler, Dan Vilenchik, et al. Do semidefinite relaxations solve sparse pca up to the information limit? *The Annals of Statistics*, 43(3):1300–1322, 2015.

Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.

Michel Ledoux. *The concentration of measure phenomenon*. Number 89. American Mathematical Soc., 2005.

Tengyu Ma and Avi Wigderson. Sum-of-squares lower bounds for sparse pca. In *Advances in Neural Information Processing Systems*, pages 1603–1611, 2015.

Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, pages 1436–1462, 2006.

Baback Moghaddam, Yair Weiss, and Shai Avidan. Spectral bounds for sparse pca: Exact and greedy algorithms. In *Advances in neural information processing systems*, pages 915–922, 2005.

Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617, 2007.

R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y.C. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*, pages 210–268. Cambridge University Press, 2012.

Vincent Q Vu and Jing Lei. Minimax rates of estimation for sparse pca in high dimensions. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS) 2012*, 2012.

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

Tengyao Wang, Quentin Berthet, and Richard J Samworth. Statistical and computational trade-offs in estimation of sparse principal components. *arXiv preprint arXiv:1408.5369*, 2014.

Daniela M Witten, Robert Tibshirani, and Trevor Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009.

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