SPARSE PRINCIPAL COMPONENT ANALYSIS AND ITERATIVE THRESHOLDING

By Zongming Ma [*]

University of Pennsylvania

Principal component analysis (PCA) is a classical dimension reduction method which projects data onto the principal subspace spanned by the leading eigenvectors of the covariance matrix. However, it behaves poorly when the number of features p is comparable to, or even much larger than, the sample size n. In this paper, we propose a new iterative thresholding approach for estimating principal subspaces in the setting where the leading eigenvectors are sparse. Under a spiked covariance model, we find that the new approach recovers the principal subspace and leading eigenvectors consistently, and even optimally, in a range of high-dimensional sparse settings. Simulated examples also demonstrate its competitive performance.

1. Introduction. In many contemporary datasets, if we organize the p-dimensional observations $x_1, \ldots, x_n$, into the rows of an $n \times p$ data matrix X, the number of features p is often comparable to, or even much larger than, the sample size n. For example, in biomedical studies, we usually have measurements on the expression levels of tens of thousands of genes, but only for tens or hundreds of individuals. One of the crucial issues in the analysis of such "large p" datasets is dimension reduction of the feature space.

As a classical method, principal component analysis (PCA) [23, 8] reduces dimensionality by projecting the data onto the *principal subspace* spanned by the m leading eigenvectors of the population covariance matrix Σ, which represent the principal modes of variation. In principle, one expects that for some $m < p$, most of the variance in the data is captured by these m modes. Thus, PCA reduces the dimensionality of the feature space while retaining most of the information in data. In addition, projection to a low-dimensional space enables visualization of the data. In practice, Σ is unknown. Classical PCA then estimates the leading population eigenvectors by those of the sample covariance matrix S. It performs well in the traditional data setting where p is small and n is large [2].

1

In high-dimensional settings, a collection of data can be modeled by a low-rank signal plus noise structure, and PCA can be used to recover the low-rank signal. In particular, each observation vector x_i can be viewed as an independent instantiation of the following generative model:

$$(1.1) \qquad\qquad x_i = \mu + A u_i + \sigma z_i.$$

Here, μ is the mean vector, A is a $p \times \bar{m}$ deterministic matrix of factor loadings, u_i is an $\bar{m}$-vector of random factors, $\sigma > 0$ is the noise level, and z_i is a p-vector of white noise. For instance, in chemometrics, x_i can be a vector of the logarithm of the absorbance or reflectance spectra measured with noise, where the columns of A are characteristic spectral responses of different chemical components, and u_i's the concentration levels of these components [31]. The number of observations are relatively few compared with the number of frequencies at which the spectra are measured. In econometrics, x_i can be the returns for a collection of assets, where the u_i's are the unobservable random factors [29]. The assumption of additive white noise is reasonable for asset returns with low frequencies (e.g., monthly returns of stocks). Here, people usually look at tens or hundreds of assets simultaneously, while the number of observations are also at the scale of tens or hundreds. In addition, model (1.1) represents a big class of signal processing problems [32]. Without loss of generality, we assume $\mu = 0$ from now on.

In this paper, our primary interest lies in PCA of high-dimensional data generated as in (1.1). Let the covariance matrix of u_i be Φ which is of full rank. Suppose that A has full column rank and that u_i and z_i are independent. Then the covariance matrix of x_i becomes

$$(1.2) \qquad\qquad \Sigma = A \Phi A' + \sigma^2 I = \sum_{j=1}^{\bar{m}} \lambda_j^2 q_j q_j' + \sigma^2 I.$$

Here, $\lambda_1^2 \geq \cdots \geq \lambda_{\bar{m}}^2 > 0$ are the eigenvalues of $A \Phi A'$, with q_j, $j = 1, \ldots, \bar{m}$, the associated eigenvectors. Therefore, the j-th eigenvalue of Σ is $\lambda_j^2 + \sigma^2$ for $j = 1, \ldots, \bar{m}$, and σ^2 otherwise. Since there are $\bar{m}$ spikes $(\lambda_1^2, \ldots, \lambda_{\bar{m}}^2)$ in the spectrum of Σ, (1.2) has been called the *spiked covariance model* in the literature [10]. Note that we use λ_j^2 to denote the spikes rather than λ_j used previously in the literature [22]. For data with such a covariance structure, it makes sense to project the data onto the low-dimensional subspaces spanned by the first few q_j's. Here and after, $\bar{m}$ denotes the number of spikes in the model, and m is the target dimension of the principal subspace to be estimated, which is no greater than $\bar{m}$.

Classical PCA encounters both practical and theoretical difficulties in high dimensions. On the practical side, the eigenvectors found by classical

PCA involve all the p features, which makes their interpretation challenging. On the theoretical side, the sample eigenvectors are no longer always consistent estimators. Sometimes, they can even be nearly orthogonal to the target direction. When both $n, p \to \infty$ with $n/p \to c \in (0, \infty)$, at different levels of rigor and generality, this phenomenon has been examined by a number of authors [25, 14, 9, 21, 17, 19] under model (1.2). See [13] for similar results when $p \to \infty$ and n is fixed.

In recent years, to facilitate interpretation, researchers have started to develop sparse PCA methodologies, where they seek a set of sparse vectors spanning the low-dimensional subspace that explains most of the variance. See, for example, [12, 36, 3, 27, 30, 34]. These approaches typically start with a certain optimization formulation of PCA and then induce a sparse solution by introducing appropriate penalties or constraints.

On the other hand, when Σ indeed has sparse leading eigenvectors in the current basis (perhaps after transforming the data), it becomes possible to estimate them consistently under high-dimensional settings via new estimation schemes. For example, under normality assumption, when Σ only has a single spike, i.e., when $\bar{m} = 1$ in (1.2), Johnstone and Lu [11] proved consistency of PCA obtained on a subset of features with large sample variances when the leading eigenvalue is fixed and $(\log p)/n \to 0$. Under the same single spike model, if in addition the leading eigenvector has exactly k non-zero loadings, Amini and Wainwright [1] studied conditions for recovering the non-zero locations using the methods in [11] and [3], and Shen, et al [26] established conditions for consistency of a sparse PCA method in [27] when $p \to \infty$ and n is fixed. For the more general multiple component case, Paul and Johnstone [22] proposed an augmented sparse PCA method for estimating each of the leading eigenvectors, and showed that their procedure attains near optimal rate of convergence under a range of high-dimensional sparse settings when the leading eigenvalues are comparable and well separated. Notably, these methods all focus on estimating individual eigenvectors.

In this paper, we focus primarily on finding *principal subspaces* of Σ spanned by sparse leading eigenvectors, as opposed to finding each sparse vector individually. One of the reasons is that individual eigenvectors are not identifiable when some leading eigenvalues are identical or close to each other. Moreover, if we view PCA as a dimension reduction technique, it is the low-dimensional subspace onto which we project data that is of the greatest interest.

We propose a new iterative thresholding algorithm to estimate principal subspaces, which is motivated by the orthogonal iteration method in matrix computation. In addition to the usual orthogonal iteration steps, an

additional thresholding step is added to seek sparse basis vectors for the subspace. When Σ follows the spiked covariance model and the sparsity of the leading eigenvectors are characterized by the weak-ℓ_r condition (3.5), the algorithm leads to a consistent subspace estimator adaptively over a wide range of high-dimensional sparse settings, and the rates of convergence are derived under an appropriate loss function (2.1). Moreover, for any individual leading eigenvector whose eigenvalue is well separated from the rest of the spectrum, our algorithm also yields an eigenvector estimator which adaptively attains optimal rate of convergence derived in [22] up to a multiplicative log factor. In addition, it has appealing model selection property in the sense that the resulting estimator only involves coordinates with large signal-to-noise ratios.

The contribution of the current paper is threefold. First, we propose to estimate principal subspaces. This is natural for the purpose of dimension reduction and visualization, and avoids the identifiability issue for individual eigenvectors. Second, we construct a new algorithm to estimate the subspaces, which is efficient in computation and easy to implement. Last but not least, we derive convergence rates of the resulting estimator under the spiked covariance model when the eigenvectors are sparse.

The rest of the paper is organized as follows. In Section 2, we frame the principal subspace estimation problem and propose the iterative thresholding algorithm. The statistical properties and computational complexity of the algorithm are examined in Sections 3 and 4 under normality assumption. Simulation results in Section 5 demonstrate its competitive performance. Section 6 presents the proof of the main theorems.

Reproducible code: The MATLAB package `SPCALab` implementing the proposed method and producing the tables and figures of the current paper is available at the author's website.

2. Methodology.

2.1. *Notation.* We say x is a p-vector if $x \in \mathbb{R}^p$, and we use $\|x\|_2$ to denote its Euclidean norm. For an $m \times n$ matrix A, its submatrix with rows indexed by I and columns indexed by J is denoted by A_{IJ}. If I or J includes all the indices, we replace it with a dot. For example, $A_{I\cdot}$ is the submatrix of A with rows in I and all columns. The spectral norm of A is $\|A\| = \max_{\|x\|_2=1} \|Ax\|_2$, and the range, i.e., the column subspace, of A is $\mathrm{ran}(A)$. If $m \geq n$, and the columns of A form an orthonormal set in $\mathbb{R}^m$, we say A is orthonormal.

We use C, C_0, C_1, etc. to represent constants, though their values might differ at different occurrences. For real numbers a and b, let $a \vee b = \max(a, b)$

and $a \wedge b = \min(a, b)$. We write $a_n = O(b_n)$, if there is a constant C, such that $|a_n| \leq Cb_n$ for all n, and $a_n = o(b_n)$ if $a_n/b_n \to 0$ as $n \to \infty$. Moreover, we write $a_n \asymp b_n$ if $a_n = O(b_n)$ and $b_n = O(a_n)$. Throughout the paper, we use ν as the generic index for features, i for observations, j for eigenvalues and eigenvectors, and k for iterations in the algorithm to be proposed.

2.2. *Framing the Problem: Principal Subspace Estimation.* When the co-variance matrix Σ follows model (1.2), its j-th largest eigenvalue $\ell_j(\Sigma) = \lambda_j^2 + \sigma^2$ for $j = 1, \ldots, \bar{m}$, and equals σ^2 for all $j > \bar{m}$. Let span$\{\cdot\}$ denote the linear subspace spanned by the vectors in the curly brackets. If for some $m \leq \bar{m}$, $\ell_m(\Sigma) > \ell_{m+1}(\Sigma)$, the *principal subspace*

$$\mathcal{P}_m = \text{span}\{q_1, \ldots, q_m\}$$

is well defined, regardless of the behavior of the other $\ell_j(\Sigma)$'s. Therefore, it is an identifiable object for the purpose of estimation. Note that $\mathcal{P}_{\bar{m}}$ is always identifiable, because $\ell_{\bar{m}}(\Sigma) > \ell_{\bar{m}+1}(\Sigma)$. The primary goal of this paper is to estimate the principal subspace $\mathcal{P}_m$, for some $m \leq \bar{m}$ with $\ell_m(\Sigma) > \ell_{m+1}(\Sigma)$. More precisely, we require the gap $\ell_m(\Sigma) - \ell_{m+1}(\Sigma) = \lambda_m^2 - \lambda_{m+1}^2 \asymp \lambda_1^2$. Note that such an m always exists, e.g., the largest $m \leq \bar{m}$ such that $\lambda_m^2 \asymp \lambda_1^2$. We allow the case of $m < \bar{m}$ partly because under certain circumstances, one might not be interested in $\mathcal{P}_{\bar{m}}$ directly. For example, to visualize the data, one might want to estimate $\mathcal{P}_2$ or $\mathcal{P}_3$ while $\bar{m}$ could be larger than 3. In addition, sometimes $\mathcal{P}_{\bar{m}}$ might not be consistently estimable while some smaller principal subspace $\mathcal{P}_m$ is. In most part of the paper, we assume that an appropriate m is given for convenience. In Section 3.5, we discuss how to choose m and how to estimate $\bar{m}$ under normality assumption.

To measure the accuracy of an estimator $\widehat{\mathcal{S}}$ for a subspace $\mathcal{S}$, note that each linear subspace is associated with a unique projection matrix onto it. Let P and $\widehat{P}$ be the projection matrices associated with $\mathcal{S}$ and $\widehat{\mathcal{S}}$, respectively. The distance between $\mathcal{S}$ and $\widehat{\mathcal{S}}$ is given by the spectral norm of the difference between P and $\widehat{P}$: $\text{dist}(\mathcal{S}, \widehat{\mathcal{S}}) = \|P - \widehat{P}\|$. See [7, Sec. 2.6.3]. Thus, we can define a loss function by the squared distances between $\mathcal{S}$ and $\widehat{\mathcal{S}}$:

$$(2.1) \qquad L(\mathcal{S}, \widehat{\mathcal{S}}) = \text{dist}^2(\mathcal{S}, \widehat{\mathcal{S}}) = \|P - \widehat{P}\|^2.$$

By definition, this loss function measures the maximum possible discrepancy between the projections of any unit vector onto the two subspaces. The loss ranges in $[0, 1]$, and equals zero if and only if $\widehat{\mathcal{S}} = \mathcal{S}$. When $\dim(\widehat{\mathcal{S}}) \neq \dim(\mathcal{S})$, we have $L(\mathcal{S}, \widehat{\mathcal{S}}) = 1$. Geometrically, it equals the squared sine of the largest canonical angle between $\mathcal{S}$ and $\widehat{\mathcal{S}}$ [28, Theorem 5.5]. Throughout the paper, we use the loss function (2.1) for principal subspace estimation.

2.3. *Orthogonal Iteration.*   Given a positive definite matrix A, a standard technique to compute its leading eigenspace is orthogonal iteration [7]. When only the first eigenvector is sought, it is also known as the power method.

To state the orthogonal iteration method, we note that for any $p \times m$ matrix T, when $p \geq m$, we could decompose it into the product of two matrices $T = QR$, where Q is $p \times m$ orthonormal and R is $m \times m$ upper triangular. This decomposition is called QR factorization, and can be computed using Gram-Schmidt orthogonalization and other numerical methods [7]. Suppose A is $p \times p$, and we want to compute its leading eigenspace of dimension m. Starting with a $p \times m$ orthonormal matrix $Q^{(0)}$, orthogonal iteration generates a sequence of $p \times m$ orthonormal matrices $Q^{(k)}$, $k = 1, 2, \ldots$, by alternating the following two steps till convergence:

   (1) Multiplication: $T^{(k)} = AQ^{(k-1)}$;
   (2) QR factorization: $Q^{(k)}R^{(k)} = T^{(k)}$.

Denote the orthonormal matrix at convergence by $Q^{(\infty)}$. Then its columns are the leading eigenvectors of A, and $\mathrm{ran}(Q^{(\infty)})$ gives the eigenspace. In practice, one terminates the iteration once $\mathrm{ran}(Q^{(k)})$ stabilizes.

When we apply orthogonal iteration directly to the sample covariance matrix S, it gives the classical PCA result, which could be problematic in high dimensions. Observe that all the p features are included in orthogonal iteration. When the dimensionality is high, not only the interpretation is hard, but the variance accumulated across all the features becomes so high that it makes consistent estimation impossible.

If the eigenvectors spanning $\mathcal{P}_m$ are sparse in the current basis, one sensible way to reduce estimation error is to focus only on those features at which the leading eigenvectors have large values, and to estimate other features by zeros. Of course, one introduces bias this way, but hopefully it is much smaller compared to the amount of variance thus reduced.

The above heuristics lead to the estimation scheme in the next subsection which incorporates this feature screening idea in orthogonal iteration.

2.4. *Iterative Thresholding Algorithm.*   Let $S = \frac{1}{n}\sum_{i=1}^{n} x_i x_i'$ be the sample covariance matrix. An effective way to incorporate feature screening into orthogonal iteration is to 'kill' small coordinates of the $T^{(k)}$ matrix after each multiplication step, which leads to the estimation scheme summarized in Algorithm 1. Although the later theoretical study is conducted under normality assumption, Algorithm 1 itself is not confined to normal data.

In addition to the two basic orthogonal iteration steps, Algorithm 1 adds a thresholding step in between them, where we threshold each element of

---

Algorithm 1: ITSPCA (Iterative thresholding sparse PCA)

Input:
1. Sample covariance matrix S;
2. Target subspace dimension m;
3. Thresholding function η, and threshold levels γ_{nj}, $j = 1, \ldots, m$;
4. Initial orthonormal matrix $\widehat{Q}^{(0)}$.

Output: Subspace estimator $\widehat{\mathcal{P}}_m = \mathrm{ran}(\widehat{Q}^{(\infty)})$, where $\widehat{Q}^{(\infty)}$ denotes the $\widehat{Q}^{(k)}$ matrix at convergence.

1 repeat

2   Multiplication: $T^{(k)} = (t^{(k)}_{\nu j}) = S\widehat{Q}^{(k-1)}$;

3   Thresholding: $\widehat{T}^{(k)} = (\widehat{t}^{(k)}_{\nu j})$, with $\widehat{t}^{(k)}_{\nu j} = \eta(t^{(k)}_{\nu j}, \gamma_{nj})$;

4   QR factorization: $\widehat{Q}^{(k)}\widehat{R}^{(k)} = \widehat{T}^{(k)}$;

5 until *convergence*;

$T^{(k)}$ with a user-specified thresholding function η which satisfies

$$(2.2) \quad |\eta(t, \gamma) - t| \leq \gamma \quad \text{and} \quad \eta(t, \gamma)1_{(|t| \leq \gamma)} = 0, \quad \text{for all } t \text{ and all } \gamma > 0.$$

Here, $1_{(E)}$ denotes the indicator function of an event E. We note that both hard-thresholding $\eta_H(t, \gamma) = t1_{(|t| > \gamma)}$ and soft-thresholding $\eta_S(t, \gamma) = \mathrm{sgn}(t)(|t| - \gamma)_+$ satisfy (2.2). So does any η sandwiched by them, such as that resulting from a SCAD criterion [6]. In $\eta(t, \gamma)$, the parameter γ is called the threshold level. In Algorithm 1, for each column of $T^{(k)}$, a common threshold level γ_{nj} needs to be specified for all its elements, which remains unchanged across iterations. The subscripts of γ_{nj} indicate that it depends on both the size of the problem n and the index j of the column it is applied to.

REMARK 2.1. The ranges of $\widehat{Q}^{(k)}$ and $\widehat{T}^{(k)}$ are the same, because QR factorization only amounts to a basis change within the same subspace. However, as in orthogonal iteration, the QR step is essential for numerical stability, and should not be omitted. Moreover, although the algorithm is designed for subspace estimation, the column vectors of $\widehat{Q}^{(\infty)}$ can be used as estimators of leading eigenvectors.

Initialization. Algorithm 1 requires an initial orthonormal matrix $\widehat{Q}^{(0)}$. It can be generated from the 'diagonal thresholding' sparse PCA algorithm [11]. Its multiple eigenvector version is summarized in Algorithm 2. Here, for any set I, $\mathrm{card}(I)$ denotes its cardinality. Given the output $\widehat{Q}_B = [\widehat{q}_1, \ldots, \widehat{q}_{\mathrm{card}(B)}]$ of Algorithm 2, we take $\widehat{Q}^{(0)} = [\widehat{q}_1, \ldots, \widehat{q}_m]$. When σ^2 is unknown, we could replace it by an estimator $\widehat{\sigma}^2$ in the definition of B.

Algorithm 2: DTSPCA (Diagonal thresholding sparse PCA)

Input:
>   1. Sample covariance matrix S;
>   2. Diagonal thresholding parameter α_n.

Output: Orthonormal matrix $\widehat{Q}_B$.

1 Variance selection: select the set B of coordinates (which are likely to have "big" signals):
$$B = \{\nu : s_{\nu\nu} \geq \sigma^2(1 + \alpha_n)\};$$

2 Reduced PCA: compute the eigenvectors, $\widehat{q}_1^B, \ldots, \widehat{q}_{\mathrm{card}(B)}^B$, of the submatrix S_{BB};

3 Zero-padding: construct $\widehat{Q}_B = [\widehat{q}_1, \ldots, \widehat{q}_{\mathrm{card}(B)}]$ such that,

$$\widehat{q}_{jB} = \widehat{q}_j^B, \qquad \widehat{q}_{jB^c} = 0, \quad j = 1, \ldots, \mathrm{card}(B).$$

For example, for normal data, [11] suggested

$$(2.3) \qquad\qquad \widehat{\sigma}^2 = \mathrm{median}\big(\tfrac{1}{n}\textstyle\sum_{i=1}^n x_{i\nu}^2\big).$$

When available, subject knowledge could also be incorporated into the construction of $\widehat{Q}^{(0)}$. Algorithm 1 also requires inputs for the γ_{nj}'s and subspace dimension m. Under normality assumption, we give explicit specification for them in (3.3) and (3.15) later. Under the conditions of the later Section 3, B is non-empty with probability tending to 1, and so $\widehat{Q}^{(0)}$ is well-defined.

Convergence. For normal data, to obtain the error rates in later Theorems 3.1 and 3.2, we can terminate Algorithm 1 after K_s iterations with K_s given in (3.4). In practice, one could also stop iterating if the difference between successive iterates becomes sufficiently small, e.g., when $L(\mathrm{ran}(\widehat{Q}^{(k)}), \mathrm{ran}(\widehat{Q}^{(k+1)})) \leq n^{-2}$. We suggest this empirical stopping rule because n^{-2} typically tends to zero faster than the rates we shall obtain, and so intuitively it should not change the statistical performance of the resulting estimator. In simulation studies reported in Section 5, the difference in numerical performance between the outputs based on this empirical stopping rule and those based on the theoretical rule (3.4) is negligible compared to the estimation errors. Whether Algorithm 1 always converges numerically is an interesting question left for possible future research.

Bibliographical note. When $m = 1$, Algorithm 1 is similar to the algorithms proposed in [27], [34] and [35]. When $m > 1$, all these methods propose to iteratively find the first leading eigenvectors of residual covariance matrices, which becomes different from our approach.

3. Statistical Properties. This section is devoted to analyzing the statistical properties of Algorithm 1 under normality assumption. After some preliminaries, we first establish the convergence rates for subspace estimation in a special yet interesting case in Section 3.1. Then we introduce a set of general assumptions in Section 3.2 and a few key quantities in Section 3.3. Section 3.4 states the main results, which include convergence rates for principal subspace estimation under general assumptions and a correct exclusion property. In addition, we derive rates for estimating individual eigenvectors. For conciseness, we first state all the results assuming a suitable target subspace dimension $m \leq \bar{m}$ is given. In Section 3.5, we discuss how to choose m and estimate $\bar{m}$ based on data.

We start with some preliminaries. Under normality assumption, $x_1, \ldots, x_n$ are i.i.d. $N_p(0, \Sigma)$ distributed, with Σ following model (1.2). Further assume σ^2 is known – though this assumption could be removed by estimating σ^2 using, say, $\widehat{\sigma}^2$ in (2.3). Since one can always scale the data first, we assume $\sigma^2 = 1$ from now on. Thus, (1.1) reduces to the orthogonal factor form

$$(3.1) \qquad x_i = \sum_{j=1}^{\bar{m}} \lambda_j v_{ij} q_j + z_i, \quad i = 1, \ldots, n.$$

Here, v_{ij} are i.i.d. standard normal random factors, which are independent of the i.i.d. white noise vectors $z_i \sim N_p(0, I)$, and $\{q_j, 1 \leq j \leq \bar{m}\}$ is a set of leading eigenvectors of Σ. In what follows, we use n to index the size of the problem. So the dimension $p = p(n)$ and the spikes $\lambda_j^2 = \lambda_j^2(n)$ can be regarded as functions of n, while both $\bar{m}$ and m remain fixed as n grows.

Let $p_n = p \vee n$. We obtain the initial matrix $\widehat{Q}^{(0)}$ in Algorithm 1 by applying Algorithm 2 with

$$(3.2) \qquad \alpha_n = \alpha \left[\frac{\log(p_n)}{n} \right]^{1/2}.$$

In Algorithm 1, the threshold levels are set at

$$(3.3) \qquad \gamma_{nj} = \gamma \left[\ell_j^B \frac{\log(p_n)}{n} \right]^{1/2}, \quad j = 1, \ldots, m.$$

Here, α and γ are user specified constants, and $\ell_j^B = \ell_j(S_{BB}) \vee 1$ with $\ell_j(S_{BB})$ the j-th largest eigenvalue of S_{BB}, where the set B is obtained in step 1 of Algorithm 2. For theoretical study, we always stop Algorithm 1 after K_s iterations, where for $h(x) = x^2/(x+1)$,

$$(3.4) \qquad K_s = \frac{1.1 \cdot \ell_1^B}{\ell_m^B - \ell_{m+1}^B} \left[\left(1 + \frac{1}{\log 2} \right) \log n + 0 \vee \log h(\ell_1^B - 1) \right].$$

Under the conditions of Theorem 3.1 or of Theorems 3.2 and 3.3, or when m is defined by (3.15), we have $\ell^B_m \neq \ell^B_{m+1}$ and $K_s < \infty$ with probability 1.

3.1. *A Special Case.* To facilitate understanding, we first state the convergence rates for principal subspace estimation in a special case.

Consider the asymptotic setting where $n \to \infty$ with $p \geq n$ and $(\log p)/n \to 0$, while the spikes $\lambda_1^2 \geq \cdots \geq \lambda_m^2 > \lambda_{m+1}^2 \geq \cdots \lambda_{\bar{m}}^2 > 0$ remain unchanged. Suppose that the q_j's are sparse in the sense that, for some $r \in (0, 2)$, the ℓ_r norm of the eigenvectors are uniformly bounded by s, i.e., $\|q_j\|_r = (\sum_{\nu=1}^p |q_{\nu j}|^r)^{1/r} \leq s$, for $j = 1, \ldots, \bar{m}$, where $s \geq 1$ is an absolute constant.

Recall that $h(x) = x^2/(x+1)$. Under the above setup, we have the following upper bound for subspace estimation error.

THEOREM 3.1. *Under the above setup, for sufficiently large constants $\alpha, \gamma > 2\sqrt{3}$ in (3.2) and (3.3), there exist constants C_0, $C_1 = C_1(\gamma, r, m)$ and C_2, such that for sufficiently large n, uniformly over all Σ with $\|q_j\|_r \leq s$ for $1 \leq j \leq \bar{m}$, with probability at least $1 - C_0 p^{-2}$, we have $K_s \asymp \log n$ and the subspace estimator $\widehat{\mathcal{P}}_m^{(K_s)} = \mathrm{ran}(\widehat{Q}^{(K_s)})$ of Algorithm 1 satisfies*

$$L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s)}) \leq C_1\,\bar{m}s^r \left[\frac{\log p}{nh(\lambda_m^2)}\right]^{1-r/2} + C_2\,g_m(\lambda)\,\frac{\log p}{n},$$

where $g_m(\lambda) = \frac{(\lambda_1^2+1)(\lambda_{m+1}^2+1)}{(\lambda_m^2-\lambda_{m+1}^2)^2}$.

The upper bound in Theorem 3.1 consists of two terms. The first is a "nonparametric" term, which can be decomposed as the product of two components. The first component, $\bar{m}s^r[nh(\lambda_m^2)/\log p]^{r/2}$, up to a multiplicative constant, bounds the number of coordinates used in estimating the subspace, while the second component, $\log p/[nh(\lambda_m^2)]$, gives the average error per coordinate. The second term in the upper bound, $g_m(\lambda)(\log p)/n$, up to a logarithmic factor, has the same form as the cross-variance term in the "fixed p, large n" asymptotic limit for classical PCA (cf., [2, Theorem 1]). We call it a "parametric" error term, because it always arises when we try to separate the first m eigenvectors from the rest, regardless of how sparse they are. Under the current setup, both terms converge to 0 as $n \to \infty$, which establishes the consistency of our estimator.

To better understand the upper bound, we compare it with an existing lower bound. Suppose $\lambda_1^2 > \lambda_2^2$. Consider the simplest case where $m = 1$. Then, estimating $\mathcal{P}_1$ is the same as estimating the first eigenvector q_1. For estimating an individual eigenvector q_j, [22] considered the loss function $l(q_j, \widetilde{q}_j) = \|q_j - \mathrm{sgn}(q_j'\widetilde{q}_j)\widetilde{q}_j\|_2^2$. Here, the λ_j^2's, s and r are fixed and $p \geq n$,

so when n is large, $s^r[nh(\lambda_1^2)/\log p]^{r/2} \leq Cp^{1-c}$ for some $c \in (0,1)$. For this case, Theorem 2 in [22] asserts that for any estimator $\widehat{q}_1$,

$$\sup_{\|q_j\|_r \leq s, \forall j} \mathbb{E}l(q_1, \widehat{q}_1) \geq C_1 s^r \left[\frac{\log p}{nh(\lambda_1^2)}\right]^{1-r/2} + C_2 \frac{g_1(\lambda)}{n}.$$

Let $\widehat{\mathcal{P}}_1 = \text{span}\{\widehat{q}_1\}$. We have $\frac{1}{2}l(q_1, \widehat{q}_1) \leq L(\mathcal{P}_1, \widehat{\mathcal{P}}_1) \leq l(q_1, \widehat{q}_1)$. So the above lower bound also holds for any $\widehat{\mathcal{P}}_1$ and $\mathbb{E}L(\mathcal{P}_1, \widehat{\mathcal{P}}_1)$. Note that in both Theorem 3.1 and the last display, the nonparametric term is dominant, and so both the lower and upper bounds are of order $[(\log p)/n]^{1-r/2}$. Therefore, Theorem 3.1 shows that the estimator from Algorithm 1 is rate optimal.

Since α_n and γ_{nj} and the stopping rule (3.4) do not involve any unknown parameter, the theorem establishes the adaptivity of our estimator: the optimal rate of convergence in Theorem 3.1 is obtained without any knowledge of the power r, the radius s or the spikes λ_j^2. Last but not least, the estimator could be obtained in $O(\log n)$ iterations and holds for all thresholding function η satisfying (2.2).

Later in Section 3.4, Theorem 3.2 establishes analogous convergence rates, but for a much wider range of high-dimensional sparse settings. In particular, the above result will be extended simultaneously along two directions:

1. The spikes $\lambda_1^2, \ldots, \lambda_{\bar{m}}^2$ will be allowed to scale as $n \to \infty$, and $\lambda_{m+1}^2, \ldots \lambda_{\bar{m}}^2$ could even be of smaller order as compared to the first m spikes;
2. Each individual eigenvector q_j will be constrained to a weak-ℓ_r ball of radius s_j (which contains the ℓ_r ball of the same radius), and the radii s_j's will be allowed to diverge as $n \to \infty$.

3.2. *Assumptions.* We now state assumptions for the general theoretical results in Section 3.4.

As outlined above, the first extension of the special case is to allow the spikes $\lambda_j^2 = \lambda_j^2(n) > 0$ to change with n, though the dependence will usually not be shown explicitly. Recall that $p_n = p \vee n$, we impose the following growth rate condition on p and the λ_j^2's.

CONDITION GR. *As $n \to \infty$, we have*

1. *the dimension p satisfies $(\log p)/n = o(1)$;*
2. *the largest spike λ_1^2 satisfies $\lambda_1^2 = O(p_n)$; the smallest spike $\lambda_{\bar{m}}^2$ satisfies $\log(p_n) = o(n\lambda_{\bar{m}}^4)$; and their ratio satisfies $\lambda_1^2/\lambda_{\bar{m}}^2 = O(n[\log(p_n)/n]^{1/2+r/4})$;*
3. *$\lim_{n\to\infty} \lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2) \in [1, \infty]$ exists for $j = 1, \ldots, \bar{m}$, with $\lambda_{\bar{m}+1}^2 = 0$.*

The first part of condition GR requires the dimension to grow at a sub-exponential rate of the sample size. The second part ensures that the spikes grow at most at linear rate with p_n, and are all of larger magnitude than $\sqrt{\log(p_n)/n}$. In addition, the condition on the ratio $\lambda_1^2/\lambda_{\bar{m}}^2$ allows us to deal with the interesting cases where the first several spikes scale at a faster rate with n than the others. This is more flexible than the assumption previously made in [22] that all the spikes grow at the same rate. The third part requires $\lim_{n\to\infty} \lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2)$ to exist for each $1 \le j \le \bar{m}$, but the limit can be infinity.

Turn to the sparsity assumption on the q_j's. We first make a mild extension from ℓ_r ball to weak-ℓ_r ball [5]. To this end, for any p-vector u, order its coordinates by magnitude as $|u|_{(1)} \ge \cdots \ge |u|_{(p)}$. We say that u belongs to the weak-ℓ_r ball of radius s, denoted by $u \in w\ell_r(s)$, if

$$(3.5) \qquad\qquad |u|_{(\nu)} \le s\nu^{-1/r}, \qquad \text{for all } \nu.$$

For $r \in (0,2)$, the above condition implies rapid decay of the ordered coefficients of u, and thus describes its sparsity. For instance, consider $u = (1/\sqrt{k}, \ldots, 1/\sqrt{k}, 0, \ldots, 0)'$ with exactly k non-zero entries all equal to $1/\sqrt{k}$. Then, for fixed $r \in (0,2)$, we have $u \in w\ell_r(k^{1/r-1/2})$. In particular, when $k = 1$, $u \in w\ell_r(1)$. Note that weak-ℓ_r ball extends ℓ_r ball, because $\|u\|_r \le s$, i.e., $u \in \ell_r(s)$, implies $u \in w\ell_r(s)$.

In what follows, we assume that for some fixed $r \in (0,2)$ and all $j \le \bar{m}$, $q_j \in w\ell_r(s_j)$ for some $s_j > 1$. We choose to use the notion of "weak-ℓ_r decay", because it provides a unified framework for several different notions of sparsity, which is convenient for analyzing a statistical estimation problem from a minimax point of view [5]. Hence, at any fixed n, we will consider whether Algorithm 1 performs uniformly well on n i.i.d. observations x_i generated by (3.1) whose covariance matrix Σ belongs to the following uniformity class

$$\mathcal{F}_n = \left\{ \Sigma_{p\times p} = \sum_{j=1}^{\bar{m}} \lambda_j^2 q_j q_j' + I : q_j \in w\ell_r(s_j), \forall j \right\}.$$

For general results, we allow the radii s_j's to depend on or even diverge with n, though we require that they do not grow too rapidly, so the leading eigenvectors are indeed sparse. This leads to the following sparsity condition.

CONDITION SP.   *As $n \to \infty$, the radius s_j of the weak-ℓ_r ball satisfies $s_j \ge 1$ and*

$$s_j^r \left[\frac{\log(p_n)}{n\lambda_j^4} \right]^{1/2-r/4} = o(1 \wedge \lambda_1^4), \quad \text{for } j = 1, \ldots, \bar{m}.$$

This type of condition also appeared in previous study of individual eigen-vector estimation in the multiple component spiked covariance model [22]. The condition is, for example, satisfied if condition GR holds and the largest spike λ_1^2 is bounded away from zero while the radii s_j's are all bounded above by an arbitrarily large constant. That is, if there exists a constant $C > 0$, such that $\lambda_1^2 \geq 1/C$ and $s_j \leq C$ for all $j \leq \bar{m}$ and all n.

It is straightforward to verify that conditions GR and SP are satisfied by the special case in Section 3.1. We conclude this part with an example.

EXAMPLE. When each x_i collects noisy measurements of an underlying random function on a regular grid, model (3.1) becomes discretization of a functional PCA model [24], and the q_j's are discretized eigenfunctions. When the eigenfunctions are smooth or have isolated singularities either in themselves or in their derivatives, their wavelet coefficients belong to some weak ℓ_r ball [5]. So do the discrete wavelet transform of the q_j's. Moreover, the radii of the weak ℓ_r balls are determined by the underlying eigenfunctions and are thus uniformly bounded as the size of the grid p gets larger. In this case, condition SP is satisfied when condition GR holds and λ_1^2 is bounded away from zero. So, for functional data of this type, we could always first transform to the wavelet domain and then apply Algorithm 1.

3.3. *Key Quantities.* We now introduce a few key quantities which appear later in the general theoretical results.

The first quantity gives the rate at which we distinguish high from low signal coordinates. Recall that $h(x) = x^2/(x+1)$. For $j = 1, \ldots, \bar{m}$, define

$$(3.6) \qquad \tau_{nj} = \sqrt{\frac{\log(p_n)}{nh(\lambda_j^2)}}.$$

According to [20], up to a logarithmic factor, τ_{nj}^2 can be interpreted as the average error per coordinate in estimating an eigenvector with eigenvalue $\lambda_j^2 + 1$. Thus, a coordinate can be regarded as of high signal if at least one of the leading eigenvectors is of larger magnitude on this coordinate compared to τ_{nj}. Otherwise, we call it a low signal coordinate. We define $H(\beta)$ to be the set of high signal coordinates:

$$(3.7) \qquad H = H(\beta) = \big\{\nu : |q_{\nu j}| \geq \beta\tau_{nj}, \text{ for some } 1 \leq j \leq \bar{m}\big\}.$$

Here, β is a constant not depending on n, the actual value of which will be specified in Theorem 3.2. If $\bar{m} = 1$ and q_1 has k non-zero entries all equal to $1/\sqrt{k}$, then H contains exactly these k coordinates when $k <$

$nh(\lambda_1^2)/[\beta^2 \log(p_n)]$, which is guaranteed under Condition SP. In addition, let $L = \{1, \ldots, p\} \backslash H$ be the complement of H. Here, H stands for "high", and L for "low" (also recall B in Algorithm 2, where B stands for "big"). The dependence of H, L and B on n is suppressed for notational convenience.

To understand the convergence rate of the subspace estimator stated later in (3.11), it is important to have an upper bound for $\mathrm{card}(H)$, the cardinality of H. To this end, define

$$(3.8) \qquad M_n = p \wedge \sum_{j=1}^{\bar{m}} \frac{s_j^r}{\tau_{nj}^r}.$$

The following lemma shows that a constant multiple of M_n bounds $\mathrm{card}(H)$. The proof of the lemma is given in [15]. Thus, in the general result, M_n plays the same role as the term $\bar{m}s^r[nh(\lambda^2)/\log p]^{r/2}$ has played in Theorem 3.1.

LEMMA 3.1. *For sufficiently large n, the cardinality of $H = H(\beta)$ satisfies $\bar{m} \leq \mathrm{card}(H) \leq CM_n$ for a constant C depending on β and r.*

The last quantity we introduce is related to the "parametric" term in the convergence rate. Let $\lambda_{\bar{m}+1}^2 = 0$. For $j = 1, \ldots, \bar{m}$, define

$$(3.9) \qquad \epsilon_{nj}^2 = \frac{(\lambda_1^2 + 1)(\lambda_{j+1}^2 + 1)}{(\lambda_j^2 - \lambda_{j+1}^2)^2} \frac{\log(p_n)}{n}.$$

So the second term of the upper bound in Theorem 3.1 is $C_2 \epsilon_{nm}^2$. For the interpretation of this quantity, we refer to the discussion after Theorem 3.1.

3.4. *Main Results.* We turn to the statement of main theoretical results.

A key condition for the results is the asymptotic distinguishability (AD) condition introduced below. Recall that all the spikes λ_j^2 (hence all the leading eigenvalues) are allowed to depend on n. The condition AD will guarantee that the largest few eigenvalues are asymptotically well separated from the rest of the spectrum, and so the corresponding principal subspace is distinguishable.

DEFINITION. *We say that* condition $\mathrm{AD}(j, \kappa)$ is satisfied with constant κ, *if there exists a numeric constant $\kappa \geq 1$, such that for sufficiently large n, the gap between the j-th and the $(j+1)$-th eigenvalues satisfies*

$$\lambda_j^2 - \lambda_{j+1}^2 \geq \lambda_1^2/\kappa.$$

We define $\mathrm{AD}(0, \kappa)$ and $\mathrm{AD}(\bar{m}, \kappa)$ by letting $\lambda_0^2 = \infty$, and $\lambda_{\bar{m}+1}^2 = 0$. So $\mathrm{AD}(0, \kappa)$ holds for any $\kappa \geq 1$. Note that there is always some $1 \leq j \leq \bar{m}$ such that condition $\mathrm{AD}(j, \kappa)$ is satisfied. For instance, $\mathrm{AD}(j, \kappa)$ is satisfied with some κ for the largest j such that $\lambda_j^2 \asymp \lambda_1^2$. When the spikes do not change with n, condition $\mathrm{AD}(\bar{m}, \kappa)$ is satisfied with any constant $\kappa \geq \lambda_1^2/\lambda_{\bar{m}}^2$.

Rates of Convergence for Principal Subspace Estimation. Recall the definitions (3.2)–(3.4) and (3.6)–(3.9). The following theorem establishes the rate of convergence of the principal subspace estimator obtained via Algorithm 1 under relaxed assumptions, which generalizes Theorem 3.1.

THEOREM 3.2. *Suppose conditions GR and SP hold, and condition $AD(m, \kappa)$ is satisfied with some constant $\kappa \geq 1$ for the given subspace dimension m. Let the constants $\alpha, \gamma > 2\sqrt{3}$ in (3.2) and (3.3), and for $c = 0.9(\gamma - 2\sqrt{3})$, let $\beta = c/\sqrt{m}$ in H (3.7). Then, there exist constants C_0, $C_1 = C_1(\gamma, r, m, \kappa)$ and C_2, such that for sufficiently large n, uniformly over $\mathcal{F}_n$, with probability at least $1 - C_0 p_n^{-2}$, $K_s \in [K, 2K]$ for*

$$(3.10) \qquad K = \frac{\lambda_1^2 + 1}{\lambda_m^2 - \lambda_{m+1}^2} \left[\left(1 + \frac{1}{\log 2} \right) \log n + 0 \vee \log h(\lambda_1^2) \right],$$

and the subspace estimator $\widehat{\mathcal{P}}_m^{(K_s)} = \mathrm{ran}(\widehat{Q}^{(K_s)})$ satisfies

$$(3.11) \qquad L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s)}) \leq C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2 = o(1).$$

Theorem 3.2 states that for appropriately chosen threshold levels and all thresholding function satisfying (2.2), after enough iterations, Algorithm 1 yields principal subspace estimators whose errors are, with high probability, uniformly bounded over $\mathcal{F}_n$ by a sequence of asymptotically vanishing constants as $n \to \infty$. In addition, the probability that the estimation error is not well controlled vanishes polynomially fast. Therefore, the subspace estimators are uniformly consistent over $\mathcal{F}_n$.

The interpretation of the two terms in the error bound (3.11) is similar to those in Theorem 3.1. Having introduced those quantities in Section 3.3, we could elaborate a little more on the first, i.e., the "nonparametric", term. By Theorem 3.3 below, when estimating $\mathcal{P}_m$, Algorithm 1 focuses only on the coordinates in H, whose cardinality is $\mathrm{card}(H) = O(M_n)$. Though H does not appear explicitly in the rates, the rates depend crucially on its cardinality which is further upper bounded by M_n. Since τ_{nm}^2 can be interpreted as the average error per coordinate, the total estimation error accumulated over all coordinates in H is thus of order $O(M_n \tau_{nm}^2)$. Moreover, as we will

show later, the squared bias induced by focusing only on H is also of order $O(M_n\tau_{nm}^2)$. Thus, this term indeed comes from the bias-variance tradeoff of the nonparametric estimation procedure. The meaning of the second, i.e., the "parametric", term is the same as in Theorem 3.1. Finally, we note that both terms vanish as $n \to \infty$ under conditions GR, SP and AD(m, κ).

The threshold levels α_n and γ_{nj} in (3.2) and (3.3) as well as K_s in (3.4) do not depend on unknown parameters. So the estimation procedure achieves the rates adaptively over a wide range of high-dimensional sparse settings.

In addition, (3.10) implies that Algorithm 1 only needs a relatively small number of iterations to yield the desired estimator. In particular, when the largest spike λ_1^2 is bounded away from zero, (3.10) shows that it suffices to have $K_s \asymp \log n$ iterations. We remark that it is not critical to run precisely K_s iterations. The result holds when we stop anywhere between K and $2K$.

Theorem 3.2 could also be extended to an upper bound for the risk. Note that $p_n^{-2} = o(\tau_{nm}^2 \vee \epsilon_{nm}^2)$, and that the loss function (2.1) is always bounded above by 1. The following result is a direct consequence of Theorem 3.2.

COROLLARY 3.1.   *Under the setup of Theorem 3.2, we have*

$$\sup_{\mathcal{F}_n} \mathbb{E}L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s)}) \leq C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2.$$

Correct Exclusion Property.   We now switch to the model selection property of Algorithm 1. By the discussion in Section 2, an important motivation for the iterative thresholding procedure is to trade bias for variance by keeping low signal coordinates out of the orthogonal iterations. More specifically, it is desirable to restrict our effort to estimating those coordinates in H and simply estimating those coordinates in L with zeros.

By construction, Algorithm 2 yields an initial matrix with a lot of zeros, but Algorithm 1 is at liberty to introduce new non-zero coordinates. The following result shows that with high probability all the non-zero coordinates introduced are in the set H.

THEOREM 3.3.   *Under the setup of Theorem 3.2, uniformly over $\mathcal{F}_n$, with probability at least $1 - C_0 p_n^{-2}$, for all $k = 0, \ldots, K_s$, the orthonormal matrix $\widehat{Q}^{(k)}$ has zeros in all its rows indexed by L, i.e., $\widehat{Q}_{L\cdot}^{(k)} = 0$.*

We call the property in Theorem 3.3 "correct exclusion", because it ensures that all the low signal coordinates in L are correctly excluded from iterations. In addition, Theorem 3.3 shows that the principal subspace estimator is indeed spanned by a set of sparse loading vectors, where all loadings in L are exactly zero.

Note that the initial matrix $\widehat{Q}^{(0)}$ has all its nonzero coordinates in B, which, with high probability, only selects "big" coefficients in the leading eigenvectors, whose magnitudes are no less than $O([\log p_n/(n\lambda_m^4)]^{1/4})$. On the other hand, the set H includes all coordinates with magnitude no less than $O([\log p_n/(nh(\lambda_m^2))]^{1/2})$. Thus, the minimum signal strength for H is of smaller order than that for B. So, with high probability, B is a subset of H consisting only of its coordinates with "big" signals. Thus, though $\widehat{Q}^{(0)}$ excludes all the coordinates in L, it only includes "big" coordinates in H and fails to pick those medium sized ones which are crucial for obtaining the convergence rate (3.11). Algorithm 1 helps to include more coordinates in H along iterations and hence achieves (3.11).

Rates of Convergence for Individual Eigenvector Estimation. The primary focus of this paper is on estimating principal subspaces. However, when an individual eigenvector, say q_j, is identifiable, it is also of interest to see whether Algorithm 1 can estimate it well. The following result shows that for K_s in (3.4), the j-th column of $\widehat{Q}^{(K_s)}$ estimates q_j well, provided that the j-th eigenvalue is well separated from the rest of the spectrum.

COROLLARY 3.2. *Under the setup of Theorem 3.2, suppose for some $j \leq m$, both conditions $AD(j-1,\kappa')$ and $AD(j,\kappa')$ are satisfied for some constant $\kappa' < \lim_{n\to\infty} \lambda_1^2/(\lambda_m^2 - \lambda_{m+1}^2)$. Then uniformly over $\mathcal{F}_n$, with probability at least $1 - C_0 p_n^{-2}$, $\widehat{q}_j^{(K_s)}$, the j-th column of $\widehat{Q}^{(K_s)}$, satisfies*

$$L(\operatorname{span}\{q_j\}, \operatorname{span}\{\widehat{q}_j^{(K_s)}\}) \leq C_1 M_n \tau_{nj}^2 + C_2(\epsilon_{n,j-1}^2 \vee \epsilon_{nj}^2).$$

Moreover, $\sup_{\mathcal{F}_n} \mathbb{E} L(\operatorname{span}\{q_j\}, \operatorname{span}\{\widehat{q}_j^{(K_s)}\})$, the supremum risk over $\mathcal{F}_n$, is also bounded by the right side of the above inequality.

Corollary 3.2 connects closely to the previous investigation [22] on estimating individual sparse leading eigenvectors. Recall their loss function $l(q_j, \widetilde{q}_j) = \|q_j - \operatorname{sgn}(q_j'\widetilde{q}_j)\widetilde{q}_j\|_2^2$. Since $\frac{1}{2}l(q_j, \widetilde{q}_j) \leq L(\operatorname{span}\{q_j\}, \operatorname{span}\{\widetilde{q}_j\}) \leq l(q_j, \widetilde{q}_j)$, l is equivalent to the restriction of the loss function (2.1) to one-dimensional subspaces. Thus, Corollary 3.2 implies that

$$\sup_{\mathcal{F}_n} \mathbb{E}\, l(q_j, \widehat{q}_j^{(K_s)}) \leq C_1 M_n \tau_{nj}^2 + C_2 \frac{(\lambda_1^2 + 1)(\lambda_j^2 + 1)}{[(\lambda_{j-1}^2 - \lambda_j^2) \wedge (\lambda_j^2 - \lambda_{j+1}^2)]^2} \frac{\log(p_n)}{n}.$$

When the radii of the weak-ℓ_r balls grow at the same rate, i.e., $\max_j s_j \asymp \min_j s_j$, the upper bound in the last display matches the lower bound in Theorem 2 of [22] up to a logarithmic factor. Thus, when the j-th eigenvalue is well separated from the rest of the spectrum, Algorithm 1 yields a near optimal estimator of q_j in the adaptive rate minimax sense.

3.5. *Choice of m.*   The main results in this section are stated with the assumption that the subspace dimension m is given. In what follows, we discuss how to choose m and also how to estimate $\bar{m}$ based on data.

Recall ℓ_j^B defined after (3.3) and the set B in step 1 of Algorithm 2. Let

$$(3.12) \qquad \widehat{m} = \max\left\{ j : \ell_j^B > 1 + \delta_{\mathrm{card}(B)} \right\}$$

be an estimator for $\bar{m}$, where for any positive integer k,

$$(3.13) \qquad \delta_k = 2(\sqrt{k/n} + t_k) + (\sqrt{k/n} + t_k)^2, \quad \text{with}$$

$$(3.14) \qquad t_k^2 = \frac{6 \log p_n}{n} + \frac{2k(\log p_n + 1)}{n}.$$

Then in Algorithm 1, for a large constant $\bar{\kappa}$, we define

$$(3.15) \qquad m = \max\left\{ j : 1 \le j \le \widehat{m} \quad \text{and} \quad \frac{\ell_1^B - 1}{\ell_j^B - \ell_{j+1}^B} \le \bar{\kappa} \right\}.$$

Setting $\bar{\kappa} = 15$ works well in simulation. For a given dataset, such a choice of m is intended to lead us to estimate the largest principal subspace such that its eigenvalues maintain a considerable gap from the rest of the spectrum. Note that (3.15) can be readily incorporated into Algorithm 2: we could compute the eigenvalues ℓ_j^B of S_{BB} in step 2, and then obtain $\widehat{m}$ and m.

For $\widehat{m}$ in (3.12), we have the following results.

PROPOSITION 3.1.   *Suppose conditions GR and SP hold. Let $\widehat{m}$ be defined in (3.12) with B obtained by Algorithm 2 with α_n specified by (3.2) for some $\alpha > 2\sqrt{3}$. Then, uniformly over $\mathcal{F}_n$, with probability at least $1 - C_0 p_n^{-2}$,*

(1) $\widehat{m} \le \bar{m}$;
(2) For any m such that the condition $AD(m, \kappa)$ is satisfied with some constant κ, $m \le \widehat{m}$ when n is sufficiently large;
(3) If the condition $AD(\bar{m}, \kappa)$ is satisfied with some constant κ, $\widehat{m} = \bar{m}$ when n is sufficiently large.

By claim (1), any $m \le \widehat{m}$ satisfies $m \le \bar{m}$ with high probability. In addition, claim (2) shows that, for sufficiently large n, any m such that $AD(m, \kappa)$ holds is no greater than $\widehat{m}$. Thus, when restricting to those $m \le \widehat{m}$, we do not miss any m such that Theorems 3.2 and 3.3 hold for estimating $\mathcal{P}_m$. These two claims jointly ensure that we do not need to consider any target dimension beyond $\widehat{m}$. Finally, claim (3) shows that we recover the exact number of spikes with high probability for large samples when $AD(\bar{m}, \kappa)$ is satisfied, i.e., when $\lambda_1^2 \asymp \lambda_{\bar{m}}^2$. Note that this assumption was made in [22].

Turn to the justification of (3.15). We show later in Corollary 6.1 that for $1 \leq j \leq \bar{m}$, $(\ell_1^B - 1)/(\ell_j^B - \ell_{j+1}^B)$ estimates $\lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2)$ consistently under conditions GR and SP. [It is important that the condition $\mathrm{AD}(m, \kappa)$ is not needed for this result!] This implies that for m in (3.15), we have $\lambda_1^2/(\lambda_m^2 - \lambda_{m+1}^2) \leq 1.1\bar{\kappa}$ when n is sufficiently large. Hence, the condition $\mathrm{AD}(m, \kappa)$ is satisfied with the constant $\kappa = 1.1\bar{\kappa}$. Therefore, the main theoretical results in Section 3.4 remain valid when we set m by (3.15) in Algorithm 1.

4. Computational Complexity. We now study the computational complexity of Algorithm 1. Throughout, we assume the same setup as in Section 3, and restrict the calculation to the high probability event on which the conclusions of Theorems 3.2 and 3.3 hold. For any matrix A, we use supp$\{A\}$ to denote the index set of the non-zero rows of A.

Consider a single iteration, say, the k-th. In the multiplication step, the (ν, j)-th element of $T^{(k)}$, $t_{\nu j}^{(k)}$, comes from the inner product of the ν-th row of S and the j-th column of $\widehat{Q}^{(k-1)}$. Though both are p-vectors, Theorem 3.3 asserts that for any column of $\widehat{Q}^{(k-1)}$, at most card(H) of its entries are non-zero. So if we know supp$\{\widehat{Q}^{(k-1)}\}$, then $t_{\nu j}^{(k)}$ can be calculated in $O(\mathrm{card}(H))$ flops, and $T^{(k)}$ in $O(mp\,\mathrm{card}(H))$ flops. Since supp$\{\widehat{Q}^{(k-1)}\}$ can be obtained in $O(mp)$ flops, the multiplication step can be completed in $O(mp\,\mathrm{card}(H))$ flops. Next, the thresholding step performs elementwise operation on $T^{(k)}$, and hence can be completed in $O(mp)$ flops. Turn to the QR step. First, we can obtain supp$\{\widehat{T}^{(k)}\}$ in $O(mp)$ flops. Then QR factorization can be performed on the reduced matrix which only includes the rows in supp$\{\widehat{T}^{(k)}\}$. Since Theorem 3.3 implies supp$\{\widehat{T}^{(k)}\} = \mathrm{supp}\{\widehat{Q}^{(k)}\} \subset H$, the complexity of this step is $O(m^2\mathrm{card}(H))$. Since $m = O(p)$, the complexity of the multiplication step dominates, and so the complexity of each iteration is $O(mp\,\mathrm{card}(H))$. Theorem 3.2 shows that K_s iteration is enough. Therefore, the overall complexity of Algorithm 1 is $O(K_s mp\,\mathrm{card}(H))$.

When the true eigenvectors are sparse, card(H) is of manageable size. In many realistic situations, λ_1^2 is bounded away from 0 and so $K_s \asymp \log n$. For these cases, Algorithm 1 is scalable to very high dimensions.

We conclude the section with a brief discussion on parallel implementation of Algorithm 1. In the k-th iteration, both matrix multiplication and elementwise thresholding can be computed in parallel. For QR factorization, one needs only to communicate the rows of $\widehat{T}^{(k)}$ with non-zero elements, the number of which is no greater than card(H). Thus, the overhead from communication is $O(m\,\mathrm{card}(H))$ for each iteration, and $O(K_s m\,\mathrm{card}(H))$ in total. When the leading eigenvectors are sparse, card(H) is manageable, and parallel computing of Algorithm 1 is feasible.

5. Numerical Experiments.

5.1. *Single Spike Settings.* We first consider the case where each x_i is generated by (3.1) with $\bar{m} = 1$. Motivated by functional data with localized features, four test vectors q_1 are considered, where $q_1 = (f(1/p), \ldots, f(p/p))'$, with f one of the four functions in Fig. 1. For each test vector, the dimension $p = 2048$, the sample size $n = 1024$, and λ_1^2 ranges in $\{100, 25, 10, 5, 2\}$.

Before applying any sparse PCA method, we transform the observed data vectors into the wavelet domain using the Symmlet 8 basis [16], and scale all the observations by $\widehat{\sigma}$ with $\widehat{\sigma}^2$ given in (2.3). The multi-resolution plots of wavelet coefficients of the test vectors are shown in Fig. 2. In the wavelet domain, the four vectors exhibits different levels of sparsity, with `step` the least sparse, and `sing` the most.

Table 1 compares the average loss of subspace estimation over 100 runs for each spike value and each test vector by Algorithm 1 (ITSPCA) with several existing methods: augmented sparse PCA (AUGSPCA) [22], correlation augmented sparse PCA (CORSPCA) [18], and diagonal thresholding sparse PCA (DTSPCA) given in Algorithm 2. For ITSPCA, we computed $\widehat{Q}^{(0)}$ by Algorithm 2. α_n and γ_{n1} are specified by (3.2) and (3.3) with $\alpha = 3$ and $\gamma = 1.5$. These values are smaller than those in theoretical results, but lead to better numerical performance. We stop iterating once $L(\operatorname{ran}(\widehat{Q}^{(k)}), \operatorname{ran}(\widehat{Q}^{(k+1)})) \leq n^{-2}$. Parameters in competing algorithms are all set to the values recommended by their authors.

From Table 1, ITSPCA and CORSPCA outperform the other two methods in all settings. Between the two, CORSPCA only wins by small margins when the spike values are large. Otherwise, ITSPCA wins, sometimes with large margins. For the same algorithm at the same spike value, the sparser the signal, the smaller the estimation error.

Table 1 also presents the average sizes of the sets of selected coordinates. While all methods yield sparse PC loadings, AUGSPCA and DTSPCA seem to select too few coordinates, and thus introduce too much bias. ITSPCA and CORSPCA apparently result in a better bias-variance tradeoff.

5.2. *Multiple Spike Settings.* Next, we simulated data vectors using model (3.1) with $\bar{m} = 4$. The q_j vectors are taken to be the four test vectors used in single spike settings, in the same order as in Fig. 1, up to orthonormalization[1]. We tried four different configurations of the spike values $(\lambda_1^2, \ldots, \lambda_4^2)$, as specified in the first column of Table 2. For each configuration of spike values, the dimension is $p = 2048$, and the sample size is $n = 1024$.

---

[1]The four test vectors are shifted such that the inner product of any pair is close to 0. So the vectors after orthonormalization are visually indistinguishable from those in Fig. 1.
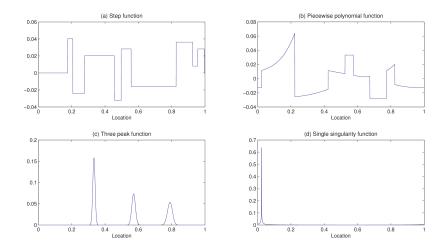
FIG 1. *Four test vectors in the original domain: values at* $p = 2048$ *equispaced points on* $[0,1]$ *of four test functions. (a)* `step`: *step function, (b)* `poly`: *piecewise polynomial function, (c)* `peak`: *three-peak function, and (d)* `sing`: *single singularity function.*
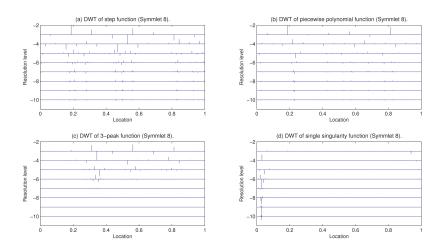

FIG 2. *Discrete wavelet transform of the four test vectors in Fig. 1. In each plot, the length of each stick is proportional to the magnitude of the* `Symmlet 8` *wavelet coefficient at the given location and resolution level.*

For each simulated dataset, we estimate $\mathcal{P}_m$ for $m = 1, 2, 3$, and 4. The last four columns of Table 2 present the losses in estimating subspaces, averaged over 100 runs, using the same sparse PCA methods as in single spike settings. For ITSPCA, we set the thresholds $\{\gamma_{nj}, j = 1, \ldots, 4\}$ as in (3.3) with $\gamma = 1.5$. All other implementation details are the same. Again, we

		ITSPCA		AUGSPCA		CORSPCA		DTSPCA	
Test vector	λ_1^2	Loss	Size	Loss	Size	Loss	Size	Loss	Size
step	100	0.0061	114.2	0.0096	96.5	0.0055	120.1	0.0275	66.6
	25	0.0224	76.3	0.0362	55.4	0.0236	73.9	0.0777	38.3
	10	0.0470	53.4	0.0710	37.4	0.0551	45.9	0.1494	24.1
	5	0.0786	45.5	0.1370	23.7	0.1119	28.7	0.2203	17.1
	2	0.1921	25.4	0.3107	11.4	0.3846	15.2	0.4518	9.7
poly	100	0.0060	83.1	0.0088	66.5	0.0051	92.0	0.0191	49.2
	25	0.0175	52.4	0.0254	41.4	0.0173	53.1	0.0540	28.7
	10	0.0346	38.7	0.0527	27.5	0.0404	34.0	0.0959	20.5
	5	0.0588	30.7	0.0844	20.2	0.0684	24.6	0.1778	14.0
	2	0.1317	20.0	0.2300	10.3	0.2155	16.3	0.3370	8.1
peak	100	0.0019	45.7	0.0032	39.6	0.0016	51.2	0.0075	32.8
	25	0.0071	34.1	0.0099	29.9	0.0069	35.2	0.0226	24.3
	10	0.0158	28.0	0.0222	23.8	0.0165	27.3	0.0592	18.6
	5	0.0283	24.7	0.0449	19.6	0.0320	22.5	0.1161	14.1
	2	0.0927	20.8	0.1887	9.9	0.1176	14.6	0.2702	8.8
sing	100	0.0016	38.0	0.0025	33.2	0.0014	43.6	0.0070	26.3
	25	0.0068	27.1	0.0095	23.1	0.0060	31.8	0.0237	17.5
	10	0.0161	20.3	0.0233	16.6	0.0154	20.9	0.0377	13.6
	5	0.0279	17.3	0.0372	13.2	0.0313	15.2	0.0547	12.7
	2	0.0631	15.2	0.0792	10.9	0.0652	13.0	0.2025	8.8

TABLE 1

Comparison of sparse PCA methods in single spike settings: average loss in estimation and size of selected feature set.

used recommended values for parameters in all other competing methods.

The simulation results reveal two interesting phenomena. First, when the spikes are relatively well separated (the first and the last blocks of Table 2), all methods yield decent estimators of $\mathcal{P}_m$ for all values of m, which implies that the individual eigenvectors are also estimated well. In this case, ITSPCA always outperforms the other three competing methods. Second, when the spikes are not so well separated (the middle two blocks, with $m = 1, 2$, or 3), no method leads to decent subspace estimator. However, all methods give reasonable estimators for $\mathcal{P}_4$, because λ_4^2 in both cases are well above 0. This implies that, under such settings, we fail to recover individual eigenvectors, but can still estimate $\mathcal{P}_4$ well. ITSPCA again gives the smallest average losses. In all configurations, the estimated number of spikes $\widehat{\widehat{m}}$ in (3.12) and the data-based choice of m in (3.15) with $\bar{\kappa} = 15$ consistently picked $m = \widehat{\widehat{m}} = 4$ in all simulated datasets. Therefore, we are always led to estimating the "right" subspace $\mathcal{P}_4$ and ITSPCA performs favorably over the competing methods.

In summary, simulations under multiple spike settings not only demonstrate the competitiveness of Algorithm 1, but also suggest:

1. The quality of principal subspace estimation depends on the gap between successive eigenvalues, in addition to the sparsity of eigenvec-

$(\lambda_1^2, \lambda_2^2, \lambda_3^2, \lambda_4^2)$	m	$L(\mathcal{P}_m, \widehat{\mathcal{P}}_m)$			
		ITSPCA	AUGSPCA	CORSPCA	DTSPCA
$(100, 75, 50, 25)$	1	0.0216	0.0260	0.0240	0.0378
	2	0.0180	0.0213	0.0214	0.0308
	3	0.0094	0.0129	0.0126	0.0234
	4	0.0087	0.0122	0.0181	0.0235
$(60, 55, 50, 45)$	1	0.3100	0.2588	0.2548	0.2831
	2	0.2675	0.2045	0.2095	0.2349
	3	0.1844	0.1878	0.1872	0.1968
	4	0.0157	0.0203	0.0178	0.0333
$(30, 27, 25, 22)$	1	0.3290	0.2464	0.2495	0.2937
	2	0.3147	0.2655	0.2882	0.3218
	3	0.1740	0.1662	0.1708	0.1821
	4	0.0270	0.0342	0.0338	0.0573
$(30, 20, 10, 5)$	1	0.0268	0.0392	0.0380	0.0658
	2	0.0237	0.0353	0.0391	0.0605
	3	0.0223	0.0336	0.0372	0.0599
	4	0.0298	0.0414	0.0717	0.0638

TABLE 2

Comparison of sparse PCA methods in multiple spike settings: average loss in estimation.

tors;

2. Focusing on individual eigenvectors can be misleading for the purpose of finding low-dimensional projections.

6. Proof. This section is devoted to the proofs of Theorems 3.2 and 3.3. We state the main ideas in Section 6.1 and divide the proof into three major steps, which are then completed in sequel in Sections 6.2–6.4. Others results in Section 3.4 are proved in the Supplement [15].

6.1. *Main Ideas and Outline of Proof.* The proof is based on an *oracle sequence approach*, the main ideas of which are as follows. First, assuming oracle knowledge of the set H, we construct a sequence of $p \times m$ orthonormal matrices $\{\widehat{Q}^{(k),\mathrm{o}}, k \geq 0\}$. Then we study how fast the sequence converges, and how well each associated column subspace approximates the principal subspace $\mathcal{P}_m$ of interest. Finally, we show that, with high probability, the first K_s terms of the oracle sequence is exactly the sequence $\{\widehat{Q}^{(k)}, 0 \leq k \leq K_s\}$ obtained by Algorithm 1. The actual estimating sequence thus inherits from the oracle sequence various properties in terms of estimation error and number of steps needed to achieve the desired error rate. The actual sequence mimics the oracle because the thresholding step forces it to only consider the high signal coordinates in H.

In what follows, we first construct the oracle sequence and then lay out a road map of the proof. Here and after, we use an extra superscript "o" to indicate oracle quantities. For example, $\widehat{Q}^{(k),\mathrm{o}}$ denotes the k-th orthonormal matrix in the oracle sequence.

Construction of The Oracle Sequence. First, we construct $\widehat{Q}^{(0),o}$ using an oracle version of Algorithm 2, where the set B is replaced by its oracle version $B^o = B \cap H$. This ensures that $\widehat{Q}_{L\cdot}^{(0),o} = 0$.

To construct the rest of the sequence, suppose that the p features are organized (after reordering) in such a way that those in H always have smaller indices than those in L, and that within H, those in B^o precede those not. Define the oracle sample covariance matrix

$$(6.1) \qquad\qquad S^o = \begin{bmatrix} S_{HH} & 0 \\ 0 & I_{LL} \end{bmatrix}.$$

Here, I_{LL} is the identity matrix of dimension $\mathrm{card}(L)$. Then, the matrices $\{\widehat{Q}^{(k),o}, k \geq 0\}$ are obtained via an oracle version of Algorithm 1, in which the initial matrix is $\widehat{Q}^{(0),o}$, and S is replaced by S^o.

REMARK 6.1. This formal construction does not guarantee that $\widehat{Q}^{(k),o}$ has full column rank or that $\widehat{Q}_{L\cdot}^{(k),o} = 0$ for all k. Later, Lemma 6.3, Proposition 6.1 and Lemma 6.4 show that these statements are true with high probability for all $k \leq K_s$.

Major Steps of The Proof. In the k-th iteration of the oracle Algorithm 1, denote the matrices obtained after multiplication and thresholding by

$$(6.2) \qquad \begin{aligned} T^{(k),o} &= S^o \widehat{Q}^{(k-1),o} = (t_{\nu j}^{(k),o}), \qquad \text{and} \\ \widehat{T}^{(k),o} &= (\widehat{t}_{\nu j}^{(k),o}) \ \text{ with } \widehat{t}_{\nu j}^{(k),o} = \eta(t_{\nu j}^{(k),o}, \gamma_{nj}). \end{aligned}$$

Further denote the QR factorization of $\widehat{T}^{(k),o}$ by $\widehat{T}^{(k),o} = \widehat{Q}^{(k),o} \widehat{R}^{(k),o}$. Last but not least, let $\widehat{\mathcal{P}}_m^{(k),o} = \mathrm{ran}(\widehat{Q}^{(k),o})$.

A joint proof of Theorems 3.2 and 3.3 can then be completed by the following three major steps:

1. Show that the principal subspace of S^o with dimension m, denoted by $\widehat{\mathcal{P}}_m^o$, satisfies the error bound in (3.11) for estimating $\mathcal{P}_m$;
2. Show that for K in (3.10), $K_s \in [K, 2K]$ and that the approximation error of $\widehat{\mathcal{P}}_m^{(k),o}$ to $\widehat{\mathcal{P}}_m^o$ for all $k \geq K$ also satisfies the bound in (3.11);
3. Show that $\widehat{Q}_{L\cdot}^{(k),o} = 0$ for all $k \leq 2K$, and that the oracle and the actual estimating sequences are identical up to $2K$ iterations.

In each step, we only need the result to hold with high probability. By the triangle inequality, steps 1 and 2 imply that the error of $\widehat{\mathcal{P}}_m^{(K_s),o}$ in estimating $\mathcal{P}_m$ satisfies (3.11). Step 3 shows this is also the case for the actual estimator $\widehat{\mathcal{P}}_m^{(K_s)}$. It also implies the correct exclusion property in Theorem 3.3.

In what follows, we complete the three steps in Sections 6.2–6.4.

6.2. *Principal Subspace of S^o.*   To study how well the principal subspace of S^o approximates $\mathcal{P}_m$, we break into a "bias" part and a "variance" part.

Consider the "bias" part first. Define the oracle covariance matrix

$$(6.3) \qquad \Sigma^o = \begin{bmatrix} \Sigma_{HH} & 0 \\ 0 & I_{LL} \end{bmatrix},$$

which is the expected value of S^o. The following lemma gives the error of the principal subspace of Σ^o in approximating $\mathcal{P}_m$, which could be regarded as the "squared bias" induced by feature selection.

LEMMA 6.1.   *Let the eigenvalues of Σ^o be $\ell_1^o \geq \cdots \geq \ell_{\bar{m}}^o \geq \cdots \geq 0$, and $\{q_1^o, \ldots, q_{\bar{m}}^o\}$ be a set of first $\bar{m}$ eigenvectors. Denote $Q^o = [q_1^o, \ldots, q_m^o]$. Then, uniformly over $\mathcal{F}_n$,*

(1) $|\ell_j^o - (\lambda_j^2 + 1)|/\lambda_1^2 \to 0$ as $n \to \infty$, for $j = 1, \ldots, \bar{m} + 1$, with $\lambda_{\bar{m}+1}^2 = 0$;

(2) For sufficiently large n, $Q_{L\cdot}^o = 0$, and for $\mathcal{P}_m^o = \mathrm{ran}(Q^o)$, there exists a constant $C = C(m, r, \kappa)$, s.t. $L(\mathcal{P}_m, \mathcal{P}_m^o) \leq C M_n \tau_{nm}^2$.

A proof is given in the Supplement [15]. Weyl's theorem [28, Corollary 4.4.10] and Davis-Kahan's $\sin \theta$ theorem [4] are the key ingredients in the proof here, and also in the proofs of Lemmas 6.2 and 6.3. Here, claim (1) only requires conditions GR and SP, but not the condition $\mathrm{AD}(m, \kappa)$.

Turn to the "variance" part. We check how well the principal subspace of S^o estimates $\mathcal{P}^o$. Since $\Sigma^o = \mathbb{E}[S^o]$, the error here is analogous to "variance".

LEMMA 6.2.   *Let the eigenvalues of S^o be $\widehat{\ell}_1^o \geq \cdots \geq \widehat{\ell}_{\bar{m}}^o \geq \cdots \geq 0$, and $\{\widehat{q}_1^o, \ldots, \widehat{q}_{\bar{m}}^o\}$ be a set of first $\bar{m}$ eigenvectors. Denote $\widehat{Q}^o = [\widehat{q}_1^o, \ldots, \widehat{q}_m^o]$. Then, uniformly over $\mathcal{F}_n$, with probability at least $1 - C_0 p_n^{-2}$,*

(1) $|\widehat{\ell}_j^o - \ell_j^o|/\lambda_1^2 \to 0$ as $n \to \infty$, for $j = 1, \ldots, \bar{m} + 1$;

(2) For sufficiently large n, $\widehat{Q}_{L\cdot}^o = 0$, and for $\widehat{\mathcal{P}}_m^o = \mathrm{ran}(\widehat{Q}^o)$, there exist constants $C_1 = C_1(m, r, \kappa)$ and C_2, s.t.

$$L(\mathcal{P}_m^o, \widehat{\mathcal{P}}_m^o) \leq C_1 M_n \tau_{nm}^2 / \log(p_n) + C_2 \epsilon_{nm}^2.$$

A proof is given in the Supplement [15]. Again, claim (1) does not require the condition $\mathrm{AD}(m, \kappa)$. By the triangle inequality, the above two lemmas imply the error in estimating $\mathcal{P}_m$ with $\widehat{\mathcal{P}}_m^o$ satisfies the bound in (3.11).

6.3. *Properties of The Oracle Sequence.*   In step 2, we study properties of the oracle sequence. For K in (3.10), the goal is to show that, with high probability, for all $k \geq K$, the error of the oracle subspace estimator $\widehat{\mathcal{P}}_m^{(k),o}$ in approximating $\widehat{\mathcal{P}}_m^o$ satisfies in (3.11). To this end, characterization of the oracle sequence evolution in Proposition 6.1 below plays the key role.

The Initial Point. We start with the initial point. Let

$$(6.4) \qquad \rho = \widehat{\ell}^{\mathrm{o}}_{m+1}/\widehat{\ell}^{\mathrm{o}}_m$$

denote the ratio between the $(m+1)$-th and the m-th largest eigenvalues of S^{o}. The following lemma shows that $\widehat{Q}^{(0),\mathrm{o}}$ is orthonormal, and is a good initial point for (oracle) Algorithm 1.

LEMMA 6.3. *Uniformly over $\mathcal{F}_n$, with probability at least $1 - C_0 p_n^{-2}$,*

(1) $B^{\mathrm{o}} = B$;

(2) $|\ell_j(S_{B^{\mathrm{o}} B^{\mathrm{o}}}) \vee 1 - \widehat{\ell}^{\mathrm{o}}_j|/\lambda_1^2 \to 0$ as $n \to \infty$, for $j = 1, \ldots, \bar{m} + 1$;

(3) For sufficiently large n, $\widehat{Q}^{(0),\mathrm{o}}$ has full column rank, and $L(\widehat{\mathcal{P}}^{\mathrm{o}}_m, \widehat{\mathcal{P}}^{(0),\mathrm{o}}_m) \le (1 - \rho)^2/5$;

(4) For sufficiently large n, $K_s \in [K, 2K]$.

A proof is given in the Supplement [15]. Here, claims (1) and (2) do not require the condition $\mathrm{AD}(m, \kappa)$. In claim (3), the bound $(1 - \rho)^2/5$ is much larger than that in (3.11). For instance, if $\lambda_m^2 + 1 \asymp \lambda_m^2 - \lambda_{m+1}^2$, Lemmas 6.1 and 6.2 imply that $(1 - \rho^2)/5 \asymp 1$ with high probability.

Claims (1) and (2) here, together with claims (1) of Lemmas 6.1 and 6.2, lead to the following result on consistent estimation of $\lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2)$ and λ_j^2, the proof of which is given in the Supplement [15].

COROLLARY 6.1. *Suppose conditions GR and SP hold, and let $\ell_j^B = \ell_j(S_{BB}) \vee 1$. For $1 \le j \le \bar{m}$, if $\lim_{n \to \infty} \lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2) < \infty$, then*

$$\lim_{n \to \infty} \frac{(\ell_1^B - 1)/(\ell_j^B - \ell_{j+1}^B)}{\lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2)} = 1, \quad a.s.$$

Otherwise, $\lim_{n \to \infty} (\ell_1^B - 1)/(\ell_j^B - \ell_{j+1}^B) = \lim_{n \to \infty} \lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2) = \infty$, a.s.

If further the condition $AD(m, \kappa)$ holds for some $m \le \bar{m}$ and $\kappa > 0$, then $\lim_{n \to \infty} (\ell_j^B - 1)/\lambda_j^2 = 1$, a.s., for $1 \le j \le m$.

Evolution of The Oracle Sequence. Next, we study the evolution of the oracle sequence. Let $\theta^{(k)} \in [0, \pi/2]$ be the largest canonical angle between the subspaces $\widehat{\mathcal{P}}^{\mathrm{o}}_m$ and $\widehat{\mathcal{P}}^{(k),\mathrm{o}}_m$. By the discussion after (2.1), we have

$$(6.5) \qquad \sin^2 \theta^{(k)} = L\big(\widehat{\mathcal{P}}^{\mathrm{o}}_m, \widehat{\mathcal{P}}^{(k),\mathrm{o}}_m\big).$$

The following proposition describes the evolution of $\theta^{(k)}$ over iterations.

PROPOSITION 6.1. *Let n be sufficiently large. On the event such that the conclusions of Lemmas 6.1–6.3 hold, uniformly over $\mathcal{F}_n$, for all $k \geq 1$,*

(1) $\widehat{Q}^{(k),\mathrm{o}}$ *is orthonormal, and $\theta^{(k)}$ satisfies*

$$(6.6) \qquad \sin \theta^{(k)} \leq \rho \tan \theta^{(k-1)} + \omega \sec \theta^{(k-1)},$$

where $\omega = (\widehat{\ell}_m^{\mathrm{o}})^{-1}\big[\mathrm{card}(H)\sum_{j=1}^m \gamma_{nj}^2\big]^{1/2}$;
(2) For any $a \in (0, 1/2]$, If

$$(6.7) \qquad \sin^2 \theta^{(k-1)} \leq 1.01(1-a)^{-2}\omega^2(1-\rho)^{-2},$$

then so is $\sin^2 \theta^{(k)}$. Otherwise,

$$(6.8) \qquad \sin^2 \theta^{(k)} / \sin^2 \theta^{(k-1)} \leq [1 - a(1-\rho)]^2.$$

A proof is given in the Supplement [15], the key ingredient of which is Wedin's $\sin\theta$ theorem for singular subspaces [33]. The recursive inequality (6.6) characterizes the evolution of the angles $\theta^{(k)}$, and hence of the oracle subspace $\widehat{\mathcal{P}}_m^{(k),\mathrm{o}}$. It is the foundation of claim (2) in the current proposition and of Proposition 6.2 below.

By (6.5), the inequality (6.8) gives the rate at which the approximation error $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k),\mathrm{o}})$ decreases. For a given $a \in (0, 1/2]$, the rate is maintained until the error becomes smaller than $1.01(1-a)^{-2}\omega^2(1-\rho)^{-2}$. Then the error continues to decrease, but at a slower rate, say, with a replaced by $a/2$ in (6.8), until (6.7) is satisfied with a replaced by $a/2$. The decrease continues at slower and slower rate in this fashion until the approximation error falls into the interval $[0, 1.01\omega^2/(1-\rho)^2]$, and remains inside thereafter.

Together with Lemma 6.3, Proposition 6.1 also justifies the previous claim that elements of the oracle sequence are orthonormal with high probability.

Convergence. Finally, we study how fast the oracle sequence converges to a stable subspace estimator, and how good this estimator is.

To define convergence of the subspace sequence $\{\widehat{\mathcal{P}}_m^{(k),\mathrm{o}}, k \geq 0\}$, we first note that $1.01\omega^2/(1-\rho)^2$ is almost the smallest possible value of $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k),\mathrm{o}})$ that (6.6) could imply. Indeed, when $\sin\theta^{(k)}$ converges and is small, we have $\sin\theta^{(k)} \approx \sin\theta^{(k-1)}$, and $\cos\theta^{(k)} \approx 1$. Consequently, (6.6) reduces to

$$\sin\theta^{(k)} \leq (\rho\sin\theta^{(k)} + \omega)(1 + o(1)).$$

So, $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k),\mathrm{o}}) = \sin^2\theta^{(k)} \leq (1 + o(1))\,\omega^2/(1-\rho)^2$. In addition, Lemma 6.2 suggests that we can stop the iteration as soon as $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k),\mathrm{o}})$ becomes

smaller than a constant multiple of ϵ_{nm}^2, for we always get an error of order $O(\epsilon_{nm}^2)$ for estimating $\mathcal{P}_m$, even if we use $\widehat{\mathcal{P}}_m^{\text{o}}$ directly. In observation of both aspects, we say that $\widehat{\mathcal{P}}_m^{(k),\text{o}}$ has *converged* if

$$(6.9) \qquad L(\widehat{\mathcal{P}}_m^{\text{o}}, \widehat{\mathcal{P}}_m^{(k),\text{o}}) \leq \max\Big\{ \frac{1.01}{(1-n^{-1})^2} \frac{\omega^2}{(1-\rho)^2}, \; \epsilon_{nm}^2 \Big\}.$$

On the event that conclusions of Lemmas 6.1–6.3 hold, we have $\omega^2/(1-\rho)^2 = O(M_n \tau_{nm}^2)$. Under definition (6.9), for K in (3.10), the following proposition shows that it takes K iterations for the oracle sequence to converge, and for all $k \geq K$, the error of approximating $\widehat{\mathcal{P}}_m^{\text{o}}$ by $\widehat{\mathcal{P}}_m^{(k),\text{o}}$ satisfies (3.11).

PROPOSITION 6.2.   *For sufficiently large n, on the event such that the conclusions of Lemmas 6.1–6.3 hold, uniformly over $\mathcal{F}_n$, it takes at most K steps for the oracle sequence to converge. In addition, there exist constants $C_1 = C_1(\gamma, r, m, \kappa)$ and C_2, such that for all $k \geq K$,*

$$(6.10) \qquad \sup_{\mathcal{F}_n} L(\widehat{\mathcal{P}}_m^{\text{o}}, \widehat{\mathcal{P}}_m^{(k),\text{o}}) \leq C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2.$$

A proof is given in the Supplement [15], and this completes step 2.

6.4. *Proof of Main Results.*   We now prove the properties of the actual estimating sequence. The proof relies on the following lemma, which shows the actual and the oracle sequences are identical up to $2K$ iterations.

LEMMA 6.4.   *For sufficiently large n, with probability at least $1 - C_0 p_n^{-2}$, for all $k \leq 2K$, we have $\widehat{Q}_{L\cdot}^{(k),\text{o}} = 0$, $\widehat{Q}^{(k)} = \widehat{Q}^{(k),\text{o}}$, and hence $\widehat{\mathcal{P}}_m^{(k)} = \widehat{\mathcal{P}}_m^{(k),\text{o}}$.*

A proof is given in the Supplement [15], and this completes step 3.

We now prove Theorems 3.2 and 3.3 by showing that the actual sequence inherits the desired properties from the oracle sequence. Since Theorem 3.1 is a special case of Theorem 3.2, we do not give a separate proof.

PROOF OF THEOREM 3.2.   Note that the event on which the conclusions of Lemmas 6.1–6.4 hold has probability at least $1 - C_0 p_n^{-2}$. On this event,

$$
\begin{aligned}
L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s)}) &= L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s),\text{o}}) \\
&\leq \big[L^{1/2}(\mathcal{P}_m, \mathcal{P}_m^{\text{o}}) + L^{1/2}(\mathcal{P}_m^{\text{o}}, \widehat{\mathcal{P}}_m^{\text{o}}) + L^{1/2}(\widehat{\mathcal{P}}_m^{\text{o}}, \widehat{\mathcal{P}}_m^{(K_s),\text{o}}) \big]^2 \\
&\leq C \big[L(\mathcal{P}_m, \mathcal{P}_m^{\text{o}}) + L(\mathcal{P}_m^{\text{o}}, \widehat{\mathcal{P}}_m^{\text{o}}) + L(\widehat{\mathcal{P}}_m^{\text{o}}, \widehat{\mathcal{P}}_m^{(K_s),\text{o}}) \big] \\
&\leq C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2.
\end{aligned}
$$

Here, the first equality comes from Lemma 6.4. The first two inequalities result from the triangle inequality and Jensen's inequality, respectively. Finally, the last inequality is obtained by noting that $K_s \in [K, 2K]$ and by replacing all the error terms by their corresponding bounds in Lemmas 6.1, 6.2 and Proposition 6.2. □

PROOF OF THEOREM 3.3. Again, we consider the event on which the conclusions of Lemmas 6.1–6.4 hold. Then Lemma 6.4 directly leads to the conclusion that $\widehat{Q}_{L\cdot}^{(k)} = \widehat{Q}_{L\cdot}^{(k),\mathrm{o}} = 0$, for all $0 \leq k \leq K_s \leq 2K$. □

SUPPLEMENTARY MATERIAL

Supplement to "Sparse principal component analysis and iterative thresholding"
(). We give in the supplement proofs to Corollaries 3.1 and 3.2, Proposition 3.1 and all the claims in Section 6.

References.

[1] A.A. Amini and M.J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *Ann. Stat.*, 37(5B):2877–2921, 2009.

[2] T.W. Anderson. Asymptotic theory for principal component analysis. *Ann. Math. Stat.*, 34(1):122–148, 1963.

[3] A. d'Aspremont, L. El Ghaoui, M. Jordan, and G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM Rev.*, 49:434–448, 2007.

[4] C. Davis and W.M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.*, 7(1):1–46, 1970.

[5] D. L. Donoho. Unconditional bases are optimal bases for data compression and for statistical estimation. *Appl. Comput. Harmon. Anal.*, 1:100–115, 1993.

[6] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Am. Stat. Assoc.*, 96:1348–1360, 2001.

[7] G.H. Golub and C.F. van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore and London, 3rd edition, 1996.

[8] H. Hotelling. Analysis of a complex of statistical variables into principal components. *J. Educ. Psychol.*, 24:417–441, 498–520, 1933.

[9] D.C. Hoyle and M. Rattray. Principal-component-analysis eigenvalue spectra from data with symmetry-breaking structure. *Phys. Rev. E*, 69(2):026124, 2004.

[10] I.M. Johnstone. On the distribution of the largest eigenvalue in principal component analysis. *Ann. Stat.*, 29:295–327, 2001.

[11] I.M. Johnstone and A.Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *J. Am. Stat. Assoc.*, 104(486):682–693, 2009.

[12] I. Jolliffe, N. Trendafilov, and M. Uddin. A modified principal component technique based on the LASSO. *J. Comput. Graph. Stat.*, 12:531–547, 2003.

[13] S. Jung and J. S. Marron. PCA consistency in high dimension, low sample size context. *Ann. Stat.*, 37(6B):4104–4130, 2009.

[14] A.Y. Lu. *Sparse principal component analysis for functional data*. PhD thesis, Stanford University, Stanford, CA, 2002.

[15] Z. Ma. Supplement to "Sparse principal component analysis and iterative threhsolding". 2011.

[16] S. Mallat. *A Wavelet Tour of Signal Processing: The Sparse Way.* Academic Press, 2009.

[17] B. Nadler. Finite sample approximation results for principal component analysis: A matrix perturbation approach. *Ann. Stat.*, 36(6):2791–2817, 2008.

[18] B. Nadler. Discussion of "On consistency and sparsity for principal components analysis in high dimensions" by Johnstone and Lu. *J. Am. Stat. Assoc.*, 104(486):694–697, 2009.

[19] A. Onatski. Asymptotics of the principal components estimator of large factor models with weak factors. *J. Econometrics*, 168:244–258, 2012.

[20] D. Paul. *Nonparametric estimation of principal components.* PhD thesis, Stanford University, 2005.

[21] D. Paul. Asymptotics of sample eigenstruture for a large dimensional spiked covariance model. *Stat. Sinica*, 17(4):1617–1642, 2007.

[22] D. Paul and I.M. Johnstone. Augmented sparse principal component analysis for high dimensional data. Arxiv preprint: arXiv:1202.1242v1, 2007.

[23] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philos. Mag. Ser. 6*, 2(11):559–572, 1901.

[24] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis.* Springer, New York, NY, 2nd edition, 2005.

[25] P. Reimann, C. Van den Broeck, and G.J. Bex. A Gaussian scenario for unsupervised learning. *J. Phys. A*, 29:3521–3535, 1996.

[26] D. Shen, H. Shen, and J. S. Marron. Consistency of sparse PCA in high dimension, low sample size contexts. 2011.

[27] H. Shen and J. Z. Huang. Sparse principal component analysis via regularized low rank matrix approximation. *J. Multivariate Anal.*, 99:1015–1034, 2008.

[28] G.W. Stewart and J.-G. Sun. *Matrix Perturbation Theory.* Computer science and scientific computing. Academic Press, 1990.

[29] R.S. Tsay. *Analysis of Financial Time Series.* Wiley-Interscience, 2005.

[30] M.O. Ulfarsson and V. Solo. Sparse variable PCA using geodesic steepest descent. *IEEE T. Signal Proces.*, 56(12):5823–5832, 2008.

[31] K. Varmuza and P. Filzmoser. *Introduction to Multivariate Statistical Analysis in Chemometrics.* CRC Press, 2009.

[32] M. Wax and T. Kailath. Detection of signals by information theoretic criteria. *IEEE T. Acoust. Speech*, 33(2):387–392, 1985.

[33] P.-A. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT*, 12:99–111, 1972.

[34] D.M. Witten, R. Tibshirani, and T. Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009.

[35] X.-T. Yuan and T. Zhang. Truncated power method for sparse eigenvalue problems. *arXiv manuscript*, arXiv:1112.2679v1, 2011.

[36] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *J. Comput. Graph. Stat.*, 15:265–286, 2006.

DEPARTMENT OF STATISTICS
THE WHARTON SCHOOL
UNIVERSITY OF PENNSYLVANIA
PHILADELPHIA, PA 19104
E-MAIL: zongming@wharton.upenn.edu

Supplement to "Sparse Principal Component Analysis and Iterative Thresholding"

Zongming Ma
University of Pennsylvania

January 28, 2013

Abstract

We give in the supplement proofs of Corollaries 3.1 and 3.2, Proposition 3.1 and the intermediate results in Section 6 of the paper [7].

A Proofs of Lemma 3.1 and Corollaries 3.1 and 3.2

A.1 Proof of Lemma 3.1

Recall that $H = H(\beta)$ is defined in Eq.(3.7) of [7], and L is the complement of H. To prove Lemma 3.1, the following result is useful.

Lemma A.1. *The quantities $\|q_{jL}\|_2^2$ satisfy $\|q_{jL}\|_2^2 \leq [\frac{2}{2-r} s_j^r (\beta\tau_{nj})^{-r} \wedge p](\beta\tau_{nj})^2$, for $j = 1, \ldots, \bar{m}$.*

Proof. We apply the "ideal risk calculation trick" in [6] to prove the lemma. Let $t_j = t_j(n)$ be the solution of $s_j t^{-1/2} = \beta\tau_{nj}$. We have

$$\|q_{jL}\|_2^2 = \sum_{\nu \in L} q_{\nu j}^2 = \sum_{\nu \in L} \min\{q_{\nu j}^2, \beta^2 \tau_{nj}^2\} = \sum_{\nu \in L} \min\{|q_j|_{(\nu)}^2, \beta^2 \tau_{nj}^2\}$$

$$\leq \int_0^\infty \min\{\beta^2 \tau_{nj}^2, s_j^2 t^{-2/r}\} dt \leq t_j \beta^2 \tau_{nj}^2 + \frac{r}{2-r} s_j^2 t_j^{1-2/r} = \frac{2}{2-r} s_j^r (\beta\tau_{nj})^{2-r}.$$

Moreover, by the definition of L, we always have the bound $\|q_{jL}\|_2^2 \leq \text{card}(L) \cdot \max_{\nu \in L}\{q_{\nu j}^2\} \leq p(\beta\tau_{nj})^2$. This completes the proof. $\square$

We now prove Lemma 3.1. Note that the proof holds for all fixed $\beta > 0$ when n is sufficiently large. First, we show $\text{card}(H) \leq C M_n$. To this end, for $j = 1, \ldots, \bar{m}$, let $H_j = H_j(\beta) = \{\nu : |q_{\nu j}| \geq \beta\tau_{nj}\}$. Note that $H = \bigcup_{j=1}^{\bar{m}} H_j$. So, $\text{card}(H) \leq \sum_{j=1}^{\bar{m}} \text{card}(H_j)$. For each H_j, the weak-ℓ_r constraint implies that $\text{card}(H_j) \leq s_j^r (\beta\tau_{nj})^{-r}$. So, $\text{card}(H) \leq p \wedge \sum_{j=1}^{\bar{m}} s_j^r (\beta\tau_{nj})^{-r} \leq C M_n$, where the constant C depends on β and r.

Next, to show $\text{card}(H) \geq \bar{m}$, it suffices to show that $\bar{Q}_{H\cdot}$ has full column rank. Note that $\bar{Q}$ is orthonormal. So, this is equivalent to showing $\|\bar{Q}_{L\cdot}\| < 1$. For $\|\bar{Q}_{L\cdot}\|$, claim (1) implies that

$$\|\bar{Q}_{L\cdot}\|^2 \leq \|\bar{Q}_{L\cdot}\|_F^2 = \sum_{j=1}^{\bar{m}} \|q_{jL}\|_2^2 \leq \sum_{j=1}^{\bar{m}} \left[p \wedge \frac{2}{2-r} s_j^r (\beta\tau_{nj})^{-r} \right] (\beta\tau_{nj})^2 = o(1) < 1.$$

The last equality comes from condition SP. This completes the proof.

1

A.2 Proof of Corollary 3.1

Under conditions GR, SP and $\mathrm{AD}(m,\kappa)$, we have $p_n^{-2} = o(\tau_{nm}^2 \vee \epsilon_{nm}^2) = o(C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2)$. In addition, the loss function L defined in Eq.(2.1) of [7] is always bounded above by 1. Let E denote the event on which the conclusions of Lemmas 6.1–6.4 in [7] hold. Then we have

$$
\begin{aligned}
\mathbb{E}L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s)}) &= \mathbb{E}L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s)})1_{(\mathrm{E})} + \mathbb{E}L(\mathcal{P}_m, \widehat{\mathcal{P}}_m^{(K_s)})1_{(\mathrm{E}^c)} \\
&\leq (C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2)\mathbb{P}(\mathrm{E}) + \mathbb{P}(\mathrm{E}^c) \leq C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2.
\end{aligned}
$$

Here, the first inequality comes from Proposition 6.2, Lemma 6.4 and the fact that $L \leq 1$, and the second inequality uses $p_n^{-2} = o(C_1 M_n \tau_{nm}^2 + C_2 \epsilon_{nm}^2)$. This completes the proof.

A.3 Proof of Corollary 3.2

Let $\widehat{Q}^{(K_s)} = [\widehat{q}_1^{(K_s)}, \ldots, \widehat{q}_m^{(K_s)}]$. By Algorithm 1 in [7], for any $l < m$, $\mathrm{span}\{\widehat{q}_1^{(K_s)}, \ldots, \widehat{q}_l^{(K_s)}\} = \widehat{\mathcal{P}}_l^{(K_s)}$, where $\widehat{\mathcal{P}}_l^{(K_s)}$ is the subspace obtained after the K_s-th iteration of Algorithm 1 with target rank l. When $\mathrm{AD}(l,\kappa')$ holds for $l = j-1, j$ with some constant $\kappa' < \lim_{n\to\infty} \lambda_1^2/(\lambda_m^2 - \lambda_{m+1}^2)$, the proof of Theorem 3.2 leads to

$$
L(\mathcal{P}_l, \widehat{\mathcal{P}}_l^{(K_s)}) \leq C_1 M_n \tau_{nl}^2 + C_2 \epsilon_{nl}^2, \quad l = j-1, j. \tag{A.1}
$$

For $l = j-1, j$, write the projection matrices associated with $\mathcal{P}_l$ and $\widehat{\mathcal{P}}_l^{(K_s)}$ as P_l and $\widehat{P}_l^{(K_s)}$. Then the projection matrices onto $\mathrm{span}\{q_j\}$ and $\mathrm{span}\{\widehat{q}_j^{(K_s)}\}$ are $P_j - P_{j-1}$ and $\widehat{P}_j^{(K_s)} - \widehat{P}_{j-1}^{(K_s)}$. This, together with the definition of the loss function L, implies

$$
\begin{aligned}
L(\mathrm{span}\{q_j\}, \mathrm{span}\{\widehat{q}_j^{(K_s)}\}) &= \|(P_j - P_{j-1}) - (\widehat{P}_j^{(K_s)} - \widehat{P}_{j-1}^{(K_s)})\|^2 = \|(P_j - \widehat{P}_j^{(K_s)}) - (P_{j-1} - \widehat{P}_{j-1}^{(K_s)})\|^2 \\
&\leq 2\|P_j - \widehat{P}_j^{(K_s)}\|^2 + 2\|P_{j-1} - \widehat{P}_{j-1}^{(K_s)})\|^2 = 2L(\mathcal{P}_j, \widehat{\mathcal{P}}_j^{(K_s)}) + 2L(\mathcal{P}_{j-1}, \widehat{\mathcal{P}}_{j-1}^{(K_s)}) \\
&\leq C_1 M_n \tau_{nj}^2 + C_2(\epsilon_{n,j-1}^2 \vee \epsilon_{nj}^2).
\end{aligned}
$$

The last inequality comes from (A.1) and the fact that $\tau_{nj}^2 > \tau_{n,j-1}^2$.

To further show that the same bound holds for $\sup_{\mathcal{F}_n} \mathbb{E}L(\mathrm{span}\{q_j\}, \mathrm{span}\{\widehat{q}_j^{(K_s)}\})$, we could essentially repeat the proof of Corollary 3.1.

B Proofs of Results in Section 6

This section gives proofs of Lemmas 6.1–6.4, Corollary 6.1 and Propositions 6.1 and 6.2 in [7]. To this end, we note that the data matrix $X_{n\times p}$ admits the representation

$$
X = V\Lambda\bar{Q}' + Z. \tag{B.1}
$$

Here, $V_{n\times\bar{m}} = (v_{ij})$ has i.i.d. $N(0,1)$ entries, $\Lambda_{\bar{m}\times\bar{m}} = \mathrm{diag}(\lambda_1, \ldots, \lambda_{\bar{m}})$, $\bar{Q}_{p\times\bar{m}} = [q_1, \ldots, q_{\bar{m}}]$ is orthonormal, and $Z_{n\times p} = (z_{i\nu})$ has i.i.d. $N(0,1)$ entries.

B.1 Proof of Lemma 6.1

We use Weyl's theorem (Corollary 4.4.10 in [8]) to prove claim (1), and Davis-Kahn's $\sin\theta$ theorem (Theorem D.1) to prove claim (2).

2

Proof of claim (1) Weyl's theorem states that $|\ell_j^o - (\lambda_j^2 + 1)| \leq \|\Sigma^o - \Sigma\|$, for all j. Note that $\Sigma^o - \Sigma = \begin{bmatrix} 0 & -\Sigma_{HL} \\ -\Sigma_{LH} & I_{LL} - \Sigma_{LL} \end{bmatrix}$, where $\Sigma_{HL} = \bar{Q}_{H.}\Lambda^2\bar{Q}'_{L.}$ and $\Sigma_{LL} - I_{LL} = \bar{Q}_{L.}\Lambda^2\bar{Q}'_{L.}$. Thus, $\|\Sigma^o - \Sigma\| \leq 2\|\Sigma_{HL}\| + \|\Sigma_{LL} - I_{LL}\|$. Further note that $\|q_{jH}\|_2, \|q_{jL}\|_2 \leq 1$, and that Lemma A.1, together with condition SP, implies that $\sum_{j=1}^{\bar{m}} \|q_{jL}\|_2 = o(1)$. Thus, $\|\Sigma_{HL}\|$, $\|\Sigma_{LL} - I_{LL}\| \leq \lambda_1^2\sum_{j=1}^{\bar{m}}\|q_{jL}\|_2 = o(\lambda_1^2)$, and so $|\ell_j^o - (\lambda_j^2 + 1)| \leq \|\Sigma^o - \Sigma\| = o(\lambda_1^2)$, for $1 \leq j \leq \bar{m} + 1$. Note that condition AD$(m, \kappa)$ is not needed here.

Proof of claim (2) Claim (1), together with AD(m, κ), implies that $\ell_m^o \geq \lambda_m^2 + 1 - o(\lambda_1^2) > 1$ for large n. So the m leading eigenvectors of Σ^o can be obtained by first taking the m leading eigenvectors of Σ_{HH}, and then augmenting all coordinates in L with zeros. Consequently, $Q_{L.}^o = 0$,
To further bound $L(\mathcal{P}_m, \mathcal{P}_m^o)$, we apply Theorem D.1 (Davis-Kahn's $\sin\theta$ theorem) with $A = \Sigma$, $A + E = \Sigma^o$, $G_0 = Q$, and $F_0 = Q^o$. Since G_0 is orthonormal, the theorem asserts that

$$L(\mathcal{P}_m, \mathcal{P}_m^o) \leq \frac{\|(Q_c^o)'(\Sigma^o - \Sigma)Q\|^2}{[(\lambda_m^2 + 1) - \ell_{m+1}^o]^2}, \tag{B.2}$$

where Q_c^o is a $p \times (p - m)$ orthonormal matrix whose columns are orthogonal to those in Q^o. The subscript 'c' is used to indicate the range of Q_c^o is the orthogonal complement to that of Q^o.
For the right side of (B.2), claim (1) and condition AD(m, κ) imply the denominator

$$|(\lambda_m^2 + 1) - \ell_{m+1}^o| \geq \lambda_m^2 - \lambda_{m+1}^2 - o(\lambda_1^2) = (\lambda_m^2 - \lambda_{m+1}^2)(1 + o(1)). \tag{B.3}$$

Since Q_c^o is orthonormal, we bound the numerator as $\|(Q_c^o)'(\Sigma^o - \Sigma)Q\| \leq \|(\Sigma^o - \Sigma)Q\|$, where

$$(\Sigma^o - \Sigma)Q = \begin{bmatrix} E_{0,H} \\ E_{0,L}^1 + E_{0,L}^2 \end{bmatrix},$$

with $E_{0,H} = -\bar{Q}_{H.}\Lambda^2\bar{Q}'_{L.}Q_{L.}, E_{0,L}^1 = -\bar{Q}_{L.}\Lambda^2\bar{Q}'_{H.}Q_{H.}$ and $E_{0,L}^2 = -\bar{Q}_{L.}\Lambda^2\bar{Q}'_{L.}Q_{L.}$. Thus, the numerator is further bounded by $\|(\Sigma^o - \Sigma)Q\| \leq \|E_{0,H}\| + \|E_{0,L}^1\| + \|E_{0,L}^2\|$. In what follows, we bound each of the three terms on the right side.
Take $E_{0,L}^1$ for example. Let

$$\Lambda_0^2 = \text{diag}(\lambda_1^2, \ldots, \lambda_m^2), \quad \Lambda_1^2 = \text{diag}(\lambda_{m+1}^2, \ldots, \lambda_{\bar{m}}^2), \quad Q_1 = [q_{m+1}, \ldots q_{\bar{m}}]. \tag{B.4}$$

Then $-E_{0,L}^1 = Q_{L.}\Lambda_0^2Q'_{H.}Q_{H.} + Q_{1,L.}\Lambda_1^2Q'_{1,H.}Q_{H.} = Q_{L.}\Lambda_0^2Q'_{H.}Q_{H.} - Q_{1,L.}\Lambda_1^2Q'_{1,L.}Q_{L.}$. Here, the second equality comes from the identity $Q'_{1,H.}Q_{H.} + Q'_{1,L.}Q_{L.} = Q'_1Q = 0$. So,

$$\|E_{0,L}^1\| \leq \|Q_{L.}\|\|\Lambda_0^2\|\|Q_{H.}\|^2 + \|Q_{L.}\|\|\Lambda_1^2\|\|Q_{1,L.}\|^2 \leq \|Q_{L.}\|\|\Lambda_0^2\|(1 + o(1)) \leq \lambda_1^2\|Q_{L.}\|_F(1 + o(1)).$$

Similar arguments show that $\|E_{0,H}\|, \|E_{0,L}^2\| = o(\lambda_1^2\|Q_{L.}\|_F)$. So we finally bound the numerator as

$$\|(Q_c^o)'(\Sigma^o - \Sigma)Q\| \leq \|(\Sigma^o - \Sigma)Q\| \leq \lambda_1^2\|Q_{L.}\|_F(1 + o(1)). \tag{B.5}$$

Together with (B.2) and (B.3), the last display leads to

$$L(\mathcal{P}_m, \mathcal{P}_m^o) \leq \frac{\lambda_1^4\|Q_{L.}\|_F^2}{(\lambda_m^2 - \lambda_{m+1}^2)^2}(1 + o(1)) \leq \left(\frac{\lambda_1^2}{\lambda_m^2 - \lambda_{m+1}^2}\right)^2 \sum_{j=1}^m \|q_{jL}\|_2^2 (1 + o(1)).$$

Since $\lambda_1^2/(\lambda_m^2 - \lambda_{m+1}^2) \leq \kappa$ under condition AD(m, κ), we complete the proof by bounding the $\|q_{jL}\|_2^2$'s using Lemma A.1.

3

B.2 Proof of Lemma 6.2

We prove Lemma 6.2 by a similar approach to that for Lemma 6.1. The major difference here is that we are dealing with a random matrix S^o instead of a deterministic one. Note that the difference $S^\mathrm{o} - \Sigma^\mathrm{o} = \begin{bmatrix} E_{HH} & 0 \\ 0 & 0 \end{bmatrix}$. We decompose E_{HH} as $E_{HH} = \sum_{k=1}^{3} E_k$, where

$$
E_1 = \bar{Q}_{H\cdot}\Lambda \left(\frac{1}{n}V'V - I \right) \Lambda \bar{Q}'_{H\cdot}, \quad E_2 = \frac{1}{n}Z'_{\cdot H}Z_{\cdot H} - I_{HH},
$$

$$
E_3 = \frac{1}{n}\left(\bar{Q}_{H\cdot}\Lambda V'Z_{\cdot H} + Z'_{\cdot H}V\Lambda\bar{Q}'_{H\cdot} \right). \tag{B.6}
$$

Many of the bounds to appear in the proof involve the following two quantities:

$$
\zeta(n,l) = 2\left(\sqrt{\frac{l}{n}} + 2\sqrt{\frac{\log(p_n)}{n}} \right) + \left(\sqrt{\frac{l}{n}} + 2\sqrt{\frac{\log(p_n)}{n}} \right)^2, \tag{B.7}
$$

$$
\xi(n,l,m;b) = \sqrt{1 + \frac{4\log(p_n)}{n}}\left(\sqrt{\frac{l}{n}} + \sqrt{\frac{m}{n}} + b\sqrt{\frac{\log(p_n)}{n}} \right), \tag{B.8}
$$

where $l,m,n \in \mathbb{N}$ and $b > 0$. By Proposition D.1, $\zeta(n,l)$ gives a probabilistic bound on the spectral norm of the difference between a (scaled) Wishart matrix A and its mean I. Moreover, Proposition D.2 shows that $\xi(n,l,m;b)$ is a probabilistic bound on the spectral norm of $Y'Z/n$, where Y and Z are independent $n \times l$ and $n \times m$ matrices with i.i.d. $N(0,1)$ entries. For controlling these quantities, the following lemma will be used multiple times.

Lemma B.1. *There exist constants $C_1 = C_1(r), C'_1 = C'_1(r), C_2$ and C'_2, such that when n is large,*

$$
\zeta(n,\mathrm{card}(H)) \le C_1\frac{\lambda_1^2 M_n^{1/2}\tau_{nm}}{\sqrt{\log(p_n)}} + C_2\sqrt{\frac{\log(p_n)}{n}},
$$

$$
\xi(n,\bar{m},\mathrm{card}(H);2) \le C'_1\frac{\lambda_1^2 M_n^{1/2}\tau_{nm}}{\sqrt{(\lambda_1^2 + 1)\log(p_n)}} + C'_2\sqrt{\frac{\log(p_n)}{n}}.
$$

Proof. By Lemma A.1,

$$
\frac{\mathrm{card}(H)}{n} \le \frac{CM_n}{n} = CM_n\tau_{nm}^2\frac{\lambda_1^4}{(\lambda_1^2 + 1)\log(p_n)}.
$$

Note that condition GR implies that for large n, $\log(p_n)/n \le 1$ and $1/\log p_n < 1$. Therefore,

$$
\zeta(n,\mathrm{card}(H)) \le \frac{2C^{1/2}\lambda_m^2 M_n^{1/2}\tau_{nm}}{\sqrt{\log(p_n)}} + 2CM_n\tau_{nm}^2\frac{\lambda_m^4}{(\lambda_m^2 + 1)\log(p_n)} + C_2\sqrt{\frac{\log(p_n)}{n}}
$$

$$
= \frac{2C^{1/2}\lambda_m^2 M_n^{1/2}\tau_{nm}}{\sqrt{\log(p_n)}}\left(1 + \frac{\lambda_m^2}{(\lambda_m^2 + 1)\sqrt{\log(p_n)}} \right) + C_2\sqrt{\frac{\log(p_n)}{n}}
$$

$$
\le C_1\frac{\lambda_1^2 M_n^{1/2}\tau_{nm}}{\sqrt{\log(p_n)}} + C_2\sqrt{\frac{\log(p_n)}{n}}.
$$

The bound on $\xi(n,\bar{m},\mathrm{card}(H);2)$ follows from similar arguments. $\qquad\square$

Proof of claim (1) As before, Weyl's theorem implies $|\widehat{\ell}_j^{\mathrm{o}} - \ell_j^{\mathrm{o}}| \leq \|S^{\mathrm{o}} - \Sigma^{\mathrm{o}}\| = \|E_{HH}\| \leq \sum_{k=1}^3 \|E_k\|$. In what follows, we derive bounds for each of the $\|E_k\|$'s.

Consider E_1 first. Since $\|\bar{Q}_{H\cdot}\| \leq 1$ and $\|\Lambda\| \leq \lambda_1$, on the event

$$\left\{\|\frac{1}{n}V'V - I\| \leq \zeta(n, \bar{m})\right\}, \tag{B.9}$$

we obtain $\|E_1\| \leq \|\bar{Q}_{H\cdot}\|^2 \|\Lambda\|^2 \|\frac{1}{n}V'V - I\| \leq \lambda_1^2 \zeta(n, \bar{m})$. Next, for E_2, we immediately get the bound $\|E_2\| \leq \zeta(n, \mathrm{card}(H))$ on the event

$$\left\{\|\frac{1}{n}Z'_{\cdot H}Z_{\cdot H} - I_{HH}\| \leq \zeta(n, \mathrm{card}(H))\right\}. \tag{B.10}$$

Finally, for E_3, since $\|\bar{Q}_{H\cdot}\| \leq 1$ and $\|\Lambda\| \leq \lambda_1$, on the event

$$\left\{\|V'Z_{\cdot H}\| \leq n\,\xi(n, \bar{m}, \mathrm{card}(H); 2)\right\}, \tag{B.11}$$

we have $\|E_3\| \leq \frac{2}{n}\|\bar{Q}_{H\cdot}\|\|\Lambda\|\|V'Z_{H\cdot}\| \leq 2\lambda_1\xi(n, \bar{m}, \mathrm{card}(H); 2)$. These bounds and Lemma B.1 lead to $|\widehat{\ell}_j^{\mathrm{o}} - \ell_j^{\mathrm{o}}| \leq \lambda_1^2\left(C_1\frac{M_n^{1/2}\tau_{nm}}{\sqrt{\log(p_n)}} + C_2\sqrt{\frac{\log(p_n)}{n(\lambda_1^2\wedge 1)^2}}\right) = o(\lambda_1^2)$. Here, the last equality comes from conditions GR and SP. Furthermore, observe that Propositions D.1 and D.2 imply the events in (B.9), (B.10) and (B.11) occur with probability at least $1 - C_0 p_n^{-2}$, and so does the intersection of them. This completes the proof of claim (1). Note that condition $\mathrm{AD}(m, \kappa)$ is not needed here.

Proof of claim (2) We use Theorem D.1 again. In the theorem, let $A = \Sigma^{\mathrm{o}}$, $A + E = S^{\mathrm{o}}$, $G_0 = Q^{\mathrm{o}}$, $F_0 = \widehat{Q}^{\mathrm{o}}$. Then $F_1 = \widehat{Q}_c^{\mathrm{o}}$, where $\widehat{Q}_c^{\mathrm{o}}$ is a $p \times (p - m)$ orthonormal matrices whose columns are orthogonal to those in $\widehat{Q}^{\mathrm{o}}$. With this setup, Theorem D.1 asserts that

$$L(\mathcal{P}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{\mathrm{o}}) \leq \frac{\|(\widehat{Q}_c^{\mathrm{o}})'(S^{\mathrm{o}} - \Sigma^{\mathrm{o}})Q^{\mathrm{o}}\|^2}{(\ell_m^{\mathrm{o}} - \widehat{\ell}_{m+1}^{\mathrm{o}})^2}. \tag{B.12}$$

Consider the right side of (B.12). For the denominator, claim (1), Lemma 6.1 and condition $\mathrm{AD}(m, \kappa)$ imply $|\ell_m^{\mathrm{o}} - \widehat{\ell}_{m+1}^{\mathrm{o}}| \geq \lambda_m^2 - \lambda_{m+1}^2 - o(\lambda_1^2) = (\lambda_m^2 - \lambda_{m+1}^2)(1 + o(1))$. For the numerator, (B.6) leads to

$$(\widehat{Q}_c^{\mathrm{o}})'(S^{\mathrm{o}} - \Sigma^{\mathrm{o}})Q^{\mathrm{o}} = (\widehat{Q}_{c,H\cdot}^{\mathrm{o}})'E_{HH}Q_{H\cdot}^{\mathrm{o}} = \sum_{k=1}^3 (\widehat{Q}_{c,H\cdot}^{\mathrm{o}})'E_kQ_{H\cdot}^{\mathrm{o}} = \sum_{k=1}^3 E_k^0, \tag{B.13}$$

where $\widehat{Q}_{c,H\cdot}^{\mathrm{o}}$ denotes the submatrix of $\widehat{Q}_c^{\mathrm{o}}$ that contains all its rows in H. So the triangle inequality would lead to an upper bound on the numerator if we bound each $\|E_k^0\|$, which we carry out below.

By definition, $E_1^0 = (\widehat{Q}_{c,H\cdot}^{\mathrm{o}})'\bar{Q}_{H\cdot}\Lambda\left(\frac{1}{n}V'V - I\right)\Lambda\bar{Q}'_{H\cdot}Q_{H\cdot}^{\mathrm{o}}$. Since $\zeta(n, \bar{m}) \leq C\sqrt{\log(p_n)/n}$, $\|\Lambda\| \leq \lambda_1$ and $\|\bar{Q}_{H\cdot}\|, \|Q_{H\cdot}^{\mathrm{o}}\| \leq 1$, on the event (B.9),

$$\|E_1^0\| \leq \|(\widehat{Q}_{c,H\cdot}^{\mathrm{o}})'\bar{Q}_{H\cdot}\Lambda\|\|\frac{1}{n}V'V - I\|\|\Lambda\|\|\bar{Q}_{H\cdot}\|\|Q_{H\cdot}^{\mathrm{o}}\| \leq C\lambda_1\sqrt{\frac{\log(p_n)}{n}}\|(\widehat{Q}_{c,H\cdot}^{\mathrm{o}})'\bar{Q}_{H\cdot}\Lambda\|.$$

To further control the last term on the right side, we apply an "orthogonal decomposition" trick. Observe that $Q_{H\cdot}^{\mathrm{o}}$ is a $\mathrm{card}(H) \times m$ orthonormal matrix, because Q^{o} is $p \times m$ orthonormal with

5

$Q^o_{L\cdot} = 0$. Define P^o_H as a $\text{card}(H) \times (\text{card}(H) - m)$ orthonormal matrix whose columns are orthogonal to those in $Q^o_{H\cdot}$, then

$$I_{HH} = Q^o_{H\cdot}(Q^o_{H\cdot})' + P^o_H(P^o_H)'. \tag{B.14}$$

Thus $(\widehat{Q}^o_{c,H\cdot})'\bar{Q}_{H\cdot}\Lambda = (\widehat{Q}^o_{c,H\cdot})'Q^o_{H\cdot}(Q^o_{H\cdot})'\bar{Q}_{H\cdot}\Lambda + (\widehat{Q}^o_{c,H\cdot})'P^o_H(P^o_H)'\bar{Q}_{H\cdot}\Lambda$. The first term satisfies $\|(\widehat{Q}^o_{c,H\cdot})'Q^o_{H\cdot}(Q^o_{H\cdot})'\bar{Q}_{H\cdot}\Lambda\| \leq \lambda_1 \|(\widehat{Q}^o_{c,H\cdot})'Q^o_{H\cdot}\| = \lambda_1\|(\widehat{Q}^o_c)'Q^o\|$. The first inequality holds because $\|Q^o_{H\cdot}\|, \|\bar{Q}_{H\cdot}\| \leq 1$ and $\|\Lambda\| \leq \lambda_1$, and the last equality comes from the observation that $Q^o_{L\cdot} = 0$. The second term in the decomposition can be further expanded as $(\widehat{Q}^o_{c,H\cdot})'P^o_H(P^o_H)'Q_{H\cdot}\Lambda_0 + (\widehat{Q}^o_{c,H\cdot})'P^o_H(P^o_H)'Q_{1,H\cdot}\Lambda_1$, where Q_1 is defined in (B.4). Since $\|(P^o_H)'Q_{H\cdot}\| \leq \|(Q^o_c)'Q\|$, $\|\Lambda_0\| \leq \lambda_1$, $\|\Lambda_1\| \leq \lambda_{m+1}$, and the spectrum norms of all the other matrices are uniformly bounded above by 1, we obtain $\|(\widehat{Q}^o_{c,H\cdot})'P^o_H(P^o_H)'\bar{Q}_{H\cdot}\Lambda\| \leq \lambda_1\|(Q^o_c)'Q\| + \lambda_{m+1}$. Combining the two terms leads to

$$\|(\widehat{Q}^o_{c,H\cdot})'\bar{Q}_{H\cdot}\Lambda\| \leq \lambda_{m+1} + \lambda_1\left(\|(\widehat{Q}^o_c)'Q^o\| + \|(Q^o_c)'Q\|\right), \tag{B.15}$$

which in turn implies $\|E^0_1\| \leq C\left[\lambda_1\lambda_{m+1} + \lambda_1^2\left(\|(\widehat{Q}^o_c)'Q^o\| + \|(Q^o_c)'Q\|\right)\right]\sqrt{\frac{\log(p_n)}{n}}$.

Turn to E^0_2. (B.14) leads to $E^0_2 = E^0_{21} + E^0_{22}$, where

$$E^0_{21} = (\widehat{Q}^o_{c,H\cdot})'Q^o_{H\cdot}\left[\frac{1}{n}(Z_{\cdot H}Q^o_{H\cdot})'(Z_{\cdot H}Q^o_{H\cdot}) - I_m\right], \quad E^0_{22} = (\widehat{Q}^o_{c,H\cdot})'P^o_H(P^o_H)'\frac{1}{n}Z'_{\cdot H}Z_{\cdot H}Q^o_{H\cdot}.$$

For E^0_{21}, recall that $Q^o_{L\cdot} = 0$. Thus, $\|(\widehat{Q}^o_{c,H\cdot})'Q^o_{H\cdot}\| = \|(\widehat{Q}^o_c)'Q^o\|$. So, on the event

$$\left\{\|\frac{1}{n}(Z_{\cdot H}Q^o_{H\cdot})'(Z_{\cdot H}Q^o_{H\cdot}) - I_m\| \leq \zeta(n,m)\right\}, \tag{B.16}$$

we obtain $\|E^0_{21}\| \leq \|(\widehat{Q}^o_c)'Q^o\|\zeta(n,m) \leq C\|(\widehat{Q}^o_c)'Q^o\|\sqrt{\log(p_n)/n}$. For E^0_{22}, on the event

$$\left\{\|(Z_{\cdot H}P^o_H)'Z_{\cdot H}Q^o_{H\cdot}\| \leq n\,\xi(n,m,\text{card}(H);2)\right\}, \tag{B.17}$$

we obtain $\|E^0_{22}\| \leq \xi(n,m,\text{card}(H);2) \leq \xi(n,\bar{m},\text{card}(H);2)$. Combine the two parts and apply Lemma B.1. We obtain $\|E^0_2\| \leq C_1\frac{\lambda_1^2 M_n^{1/2}\tau_{nm}}{\sqrt{(\lambda_1^2+1)\log(p_n)}} + C_2\left(1 + \|\widehat{Q}^o_c Q^o\|\right)\sqrt{\frac{\log(p_n)}{n}}$.

For E^0_3, (B.6) and (B.13) lead to $E^0_3 = E^0_{31} + E^0_{32}$, where

$$E^0_{31} = \frac{1}{n}(\widehat{Q}^o_{c,H\cdot})'\bar{Q}_{H\cdot}\Lambda V'Z_{\cdot H}Q^o_{H\cdot}, \quad E^0_{32} = \frac{1}{n}(\widehat{Q}^o_{c,H\cdot})'Z'_{\cdot H}V\Lambda\bar{Q}'_{H\cdot}Q^o_{H\cdot}.$$

For E^0_{31}, on the event

$$\left\{\|V'Z_{\cdot H}Q^o_{H\cdot}\| \leq n\,\xi(n,\bar{m},m;2)\right\}, \tag{B.18}$$

(B.15) and (B.8) imply

$$\|E^0_{31}\| \leq \frac{1}{n}\|(\widehat{Q}^o_{c,H\cdot})'\bar{Q}_{H\cdot}\Lambda\|\|V'Z_{\cdot H}Q^o_{H\cdot}\| \leq \left[\lambda_{m+1} + \lambda_1\left(\|(\widehat{Q}^o_c)'Q^o\| + \|(Q^o_c)'Q\|\right)\right]\xi(n,\bar{m},m;2)$$

$$\leq C\left[\lambda_{m+1} + \lambda_1\left(\|(\widehat{Q}^o_c)'Q^o\| + \|(Q^o_c)'Q\|\right)\right]\sqrt{\frac{\log(p_n)}{n}}.$$

In addition, on the event (B.11), $\|E^0_{32}\| \le \lambda_1 \xi(n, \bar{m}, \mathrm{card}(H); 2)$. Thus the triangle inequality and Lemma B.1 lead to

$$\|E^0_3\| \le C_1 \frac{\lambda_1^3 M_n^{1/2} \tau_{nm}}{\sqrt{(\lambda_1^2 + 1)\log(p_n)}} + \left[C_2(\lambda_1 + \lambda_{m+1}) + C_3 \lambda_1 \left(\|(\widehat{Q}^{\mathrm{o}}_c)' Q^{\mathrm{o}}\| + \|(Q^{\mathrm{o}}_c)' Q\| \right) \right] \sqrt{\frac{\log(p_n)}{n}}.$$

Summing up the above bounds for the $\|E^0_k\|$'s leads to

$$\begin{aligned}
L^{1/2}(\mathcal{P}^{\mathrm{o}}_m, \widehat{\mathcal{P}}^{\mathrm{o}}_m) = \|(\widehat{Q}^{\mathrm{o}}_c)' Q^{\mathrm{o}}\| &\le \frac{\sum_{k=1}^3 \|E^0_k\|}{(\lambda_m^2 - \lambda_{m+1}^2)(1 + o(1))} \\
&\le C_1 \frac{\lambda_1^2}{\lambda_m^2 - \lambda_{m+1}^2} \frac{(\lambda_1 + 1)M_n^{1/2}\tau_{nm}}{\sqrt{(\lambda_1^2 + 1)\log(p_n)}} \\
&\quad + C_2 \frac{\lambda_1\lambda_{m+1} + \lambda_1 + \lambda_{m+1} + 1}{\lambda_m^2 - \lambda_{m+1}^2} \sqrt{\frac{\log(p_n)}{n}} + o\left(\|(\widehat{Q}^{\mathrm{o}}_c)' Q^{\mathrm{o}}\| \right)
\end{aligned}$$

Condition AD(m, κ) implies $\lambda_1^2/(\lambda_m^2 - \lambda_{m+1}^2) \le \kappa$. We collect terms and apply Jensen's inequality to obtain

$$L(\mathcal{P}^{\mathrm{o}}_m, \widehat{\mathcal{P}}^{\mathrm{o}}_m) = \|(\widehat{Q}^{\mathrm{o}}_c)' Q^{\mathrm{o}}\|^2 \le C_1 \frac{M_n \tau_{nm}^2}{\log(p_n)} + C_2 \frac{(\lambda_1^2 + 1)(\lambda_{m+1}^2 + 1)}{(\lambda_m^2 - \lambda_{m+1}^2)^2} \frac{\log(p_n)}{n}.$$

Since $Z_{\cdot H} Q^{\mathrm{o}}_{H\cdot}$ is a $n \times m$ matrix with i.i.d. $N(0, 1)$ entries, independent of the matrices $Z_{\cdot H} P^{\mathrm{o}}_H$ and V which also have i.i.d. $N(0, 1)$ entries, Propositions D.1 and D.2 then ensure that the events in (B.16), (B.17) and (B.18) occur with probability at least $1 - C_0 p_n^{-2}$, and so does the intersection of them. This completes the proof.

B.3 Proof of Lemma 6.3

We first introduce some preliminary results. The first result shows that, with high probability, the set B selects all the "big" coefficients in the leading eigenvectors up to order $O\left(\sqrt[4]{\log(p_n)/(n\lambda_j^4)} \right)$. Indeed, for $0 < a_- < 1 < a_+$, define

$$B_\pm = \left\{ \nu : \sum_{j=1}^{\bar{m}} \lambda_j^2 q_{\nu j}^2 > a_\mp \alpha \sqrt{\frac{\log(p_n)}{n}} \right\}. \tag{B.19}$$

Lemma B.2. *Let $B_\pm$ be defined as above. For appropriately chosen α and $a_\mp$, when n is sufficiently large, with probability at least $1 - C_0 p_n^{-2}$, $B_- \subset B \subset B_+ \subset H$.*

Proof. Note that $S_{\nu\nu} \sim \sigma_\nu^2 \chi_n^2/n$ with $\sigma_\nu^2 = 1 + \sum_{j=1}^{\bar{m}} \lambda_j^2 q_{\nu j}^2$. We then have

$$\mathbb{P}\{B_- \not\subset B\} = \mathbb{P}\left\{ \cup_{\nu \in B_-} \{ S_{\nu\nu} < 1 + \alpha\sqrt{\log(p_n)/n} \} \right\} \le \sum_{\nu \in B_-} \mathbb{P}\{ S_{\nu\nu} < 1 + \alpha\sqrt{\log(p_n)/n} \}$$

$$\le \sum_{\nu \in B_-} \mathbb{P}\left\{ \frac{S_{\nu\nu}}{\sigma_\nu^2} \le \frac{1 + \alpha\sqrt{\log(p_n)/n}}{1 + a_+\alpha\sqrt{\log(p_n)/n}} \right\} \le \sum_{\nu \in B_-} \mathbb{P}\left\{ \frac{\chi_n^2}{n} - 1 \le -\frac{(a_+ - 1)^2\alpha^2\log(p_n)}{4(1 + a_+\alpha\sqrt{\log(p_n)/n})^2} \right\}$$

$$\le p_n^{1 - (a_+ - 1)^2\alpha^2(1 - o(1))/4}.$$

7

Here, the last inequality comes from Lemma D.2. On the other hand, we have

$$\mathbb{P}\{B \not\subseteq B_+\} = \mathbb{P}\left\{\cup_{\nu \in B_+^c}\{S_{\nu\nu} \geq 1 + \alpha\sqrt{\log(p_n)/n}\}\right\} \leq \sum_{\nu \in B_+^c} \mathbb{P}\{S_{\nu\nu} \geq 1 + \alpha\sqrt{\log(p_n)/n}\}$$

$$\leq \sum_{\nu \in B_+^c} \mathbb{P}\left\{\frac{S_{\nu\nu}}{\sigma_\nu^2} \geq \frac{1 + \alpha\sqrt{\log(p_n)/n}}{1 + a_-\alpha\sqrt{\log(p_n)/n}}\right\} \leq \sum_{\nu \in B_+^c} \mathbb{P}\left\{\frac{\chi_n^2}{n} - 1 \geq \frac{(1-a_-)\alpha\sqrt{\log(p_n)/n}}{1 + a_-\alpha\sqrt{\log(p_n)/n}}\right\}$$

$$\leq |B_+^c| \frac{\sqrt{2}(1 + a_-\alpha\sqrt{\log(p_n)/n})}{(1-a_-)\alpha\sqrt{\log(p_n)}} \exp\left(-\frac{(1-a_-)^2\alpha^2\log(p_n)}{4(1 + a_-\alpha\sqrt{\log(p_n)/n})^2}\right)$$

$$\leq p_n^{1-(1-a_-)^2\alpha^2(1-o(1))/4}.$$

Here, the second last inequality comes from Lemma D.2. From the above bounds, if we choose α and $a_\pm$ properly, then $B_- \subset B \subset B_+$ holds with probability at least $1 - C_0 p_n^{-2}$.

To show $B_+ \subset H$, observe that for any $\nu \in B_+$, there exists some $j \in \{1, \ldots, \bar{m}\}$, s.t. $\lambda_j^2 q_{\nu j}^2 > \frac{a_-\alpha}{\bar{m}}\sqrt{\frac{\log(p_n)}{n}}$. By the definition of H, it then suffices to show that for large n, $\frac{a_-\alpha}{\bar{m}\lambda_j^2}\sqrt{\frac{\log(p_n)}{n}} > \beta^2\tau_{nj}^2 = \beta^2\frac{\lambda_j^2+1}{\lambda_j^4}\frac{\log(p_n)}{n}$, for $j = 1, \ldots, \bar{m}$. This is true because under conditions GR and SP, both $\log(p_n)/(n\lambda_{\bar{m}}^4)$ and $\log(p_n)/n$ converge to 0 as $n \to \infty$. □

Lemma B.3. *Let* $e_n^2 = \sum_{j=1}^{\bar{m}} s_j^r[\log(p_n)/(n\lambda_j^4)]^{1/2-r/4}$. *Then* $\|\bar{Q}_{\cdot B_-^c}\|_F^2 = \sum_{j=1}^{\bar{m}} \|q_{jB_-^c}\|_2^2 \leq Ce_n^2$, *and* $\lambda_1^{-i}\sqrt{\mathrm{card}(H)/n}$, $\lambda_1^{-i}\sqrt{\log(p_n)/n} = o(e_n)$ *for* $i = 1, 2$.

Proof. The bound on $\|\bar{Q}_{\cdot B_-^c}\|_F^2$ is obtained by following the lines of the proof of Lemma A.1, but with H replaced by B_-. In addition, conditions GR and SP, together with Lemma A.1, imply the bounds on $\lambda_1^{-i}\sqrt{\mathrm{card}(H)/n}$ and $\lambda_1^{-i}\sqrt{\log(p_n)/n}$. □

We now prove Lemma 6.3. Throughout, we restrict our attention to the high probability event on which the conclusion of Lemma B.2 holds. Let $M = H \backslash B^o$, and define

$$S_0^o = \begin{bmatrix} S_{0,HH} & 0 \\ 0 & I_{LL} \end{bmatrix}, \quad \text{with} \quad S_{0,HH} = \begin{bmatrix} S_{B^o B^o} & 0 \\ 0 & I_{MM} \end{bmatrix}. \tag{B.20}$$

Proof of claim (1) First, claim (1) is a direct consequence of Lemma B.2. Recall that $B^o = B \cap H$. On the event such that the conclusion of Lemma B.2 holds, we get $B \subset H$, and so $B^o = B$. Note that the proof here does not need the condition $\mathrm{AD}(m, \kappa)$.

Proof of claim (2) We first apply Weyl's theorem to obtain

$$|\ell_j(S_0^o) - \widehat{\ell}_j^o| \leq \|S^o - S_0^o\| = \|S_{HH}^o - S_{0,HH}^o\| \leq 2\|S_{MB^o}\| + \|S_{MM} - I_{MM}\|. \tag{B.21}$$

To bound $\|S_{MB^o}\|$, observe that $S_{MB^o} = \sum_{i=1}^4 D_i$, where

$$D_1 = \frac{1}{n}\bar{Q}_{M\cdot}\Lambda V'V\Lambda\bar{Q}_{\cdot B^o\cdot}', \quad D_2 = \frac{1}{n}Z_{\cdot M}'Z_{\cdot B^o}, \quad D_3 = \frac{1}{n}\bar{Q}_{M\cdot}\Lambda V'Z_{\cdot B^o}, \quad D_4 = \frac{1}{n}Z_{\cdot M}'V\Lambda\bar{Q}_{\cdot B^o\cdot}'.$$

Note that $\|\bar{Q}_{B^\circ \cdot}\| \le 1$ and that $M \subset B_-^c$. So, on event (B.9), Lemma B.3 leads to

$$\|D_1\| \le \|\Lambda\|^2 \|\frac{1}{n} V'V\| \|\bar{Q}_{M\cdot}\|_{\mathrm{F}} \le \lambda_1^2 \|\frac{1}{n} V'V\| \|\bar{Q}_{B_-^c \cdot}\|_{\mathrm{F}} \le C\lambda_1^2 e_n.$$

Define $\bar{Z} = [\bar{Z}_{\cdot 1}, \ldots, \bar{Z}_{\cdot p}]$, with $\bar{Z}_{\cdot \nu} = Z_{\cdot \nu} / \|Z_{\cdot \nu}\|_2$. Thus, D_2 satisfies $\|D_2\| \le \frac{1}{n} \max_{\nu \in H} \|Z_{\cdot \nu}\|_2^2 \cdot \|\bar{Z}'_{\cdot M} \bar{Z}_{\cdot B^\circ}\|$. Note that $\|Z_{\cdot \nu}\|_2^2 \sim_{\mathrm{i.i.d.}} \chi_n^2$. Lemma D.2 ensures that with probability at least $1 - C_0 p_n^{-2}$, $\frac{1}{n} \max_{\nu \in H} \|Z_{\cdot \nu}\|_2^2 = 1 + o(1)$. On the other hand, since B° only depends on $\{\|Z_{\cdot \nu}\|_2, \nu \in H\}$, which are independent of $\{\bar{Z}_{\cdot \nu}, \nu \in H\}$, we obtain that $\bar{Z}_{\cdot H}$ is independent of B°. So, $\bar{Z}_{\cdot M}$ is independent of $\bar{Z}_{\cdot B^\circ}$. Let $U = \mathrm{diag}(u_1, \ldots, u_{\mathrm{card}(B^\circ)})$ with u_j^2 i.i.d. χ_n^2 random variables. Then, $\bar{Z}'_{\cdot M} \bar{Z}_{\cdot B^\circ} U$ has i.i.d. $N(0,1)$ entries, and so Lemma [1, Theorem II.13] asserts that with probability at least $1 - C_0 p_n^{-2}$,

$$\|\bar{Z}'_{\cdot M} \bar{Z}_{\cdot B^\circ} U\| \le C\left(\sqrt{\mathrm{card}(M)} + \sqrt{\mathrm{card}(B^\circ)} + \sqrt{\log(p_n)}\right) \le C\left(\sqrt{\mathrm{card}(H)} + \sqrt{\log(p_n)}\right).$$

In addition, Lemma D.2 shows that with probability at least $1 - C_0 p_n^{-2}$, $\min_j u_j = n(1 - o(1))$. Thus, we obtain that, with probability at least $1 - C_0 p_n^{-2}$,

$$\|\bar{Z}'_{\cdot M} \bar{Z}_{\cdot B^\circ}\| \le \frac{\|\bar{Z}'_{\cdot M} \bar{Z}_{\cdot B^\circ} U\|}{\min_j u_j} \le C\left(\sqrt{\frac{\mathrm{card}(H)}{n}} + \sqrt{\frac{\log(p_n)}{n}}\right) = o(\lambda_1^2 e_n).$$

Here, the last inequality comes from Lemma B.3. Thus, we obtain $\|D_2\| = o(e_n)$. For D_3 and D_4, on the event (B.11), Lemmas B.1 and B.3 lead to

$$\|D_3\| \le \frac{1}{n} \|\Lambda\| \|V' Z_{\cdot B^\circ}\| \|\bar{Q}_{M\cdot}\|_{\mathrm{F}} \le \frac{1}{n} \lambda_1 \|V' Z_{\cdot H}\| \|\bar{Q}_{B_-^c \cdot}\|_{\mathrm{F}} \le \lambda_1 \xi(n, \bar{m}, \mathrm{card}(H); 2) \|\bar{Q}_{B_-^c \cdot}\|_{\mathrm{F}} = o(\lambda_1^2 e_n).$$

$$\|D_4\| \le \frac{1}{n} \|\Lambda\| \|Z'_{\cdot M} V\| \|\bar{Q}_{B^\circ \cdot}\| \le \frac{1}{n} \lambda_1 \|Z'_{\cdot H} V\| \le \lambda_1 \xi(n, \bar{m}, \mathrm{card}(H); 2) = o(\lambda_1^2 e_n).$$

Combining the four parts leads to $\|S_{MB^\circ}\| \le C\lambda_1^2 e_n$.

Switch to $S_{MM} - I_{MM}$. We decompose it as $S_{MM} - I_{MM} = \sum_{k=1}^{3} E_{MM,i}$, where

$$E_{MM,1} = \bar{Q}_{M\cdot} \Lambda \left(\frac{1}{n} V'V - I\right) \Lambda \bar{Q}'_{M\cdot}, \quad E_{MM,2} = \frac{1}{n} Z'_{\cdot M} Z_{\cdot M} - I_{MM},$$

$$E_{MM,3} = \frac{1}{n} \left(\bar{Q}_{M\cdot} \Lambda V' Z_{\cdot M} + Z'_{\cdot M} V \Lambda \bar{Q}'_{M\cdot}\right).$$

On the intersection of the events in (B.9), (B.10) and (B.11), Lemmas B.1 and B.3 lead to

$$\|E_{MM,1}\| \le \lambda_1^2 \|\frac{1}{n} V'V\| \|\bar{Q}_{M\cdot}\|_{\mathrm{F}}^2 \le C\lambda_1^2 \|\bar{Q}_{B_-^c \cdot}\|_{\mathrm{F}}^2 = o(\lambda_1^2 e_n),$$

$$\|E_{MM,2}\| \le \|\frac{1}{n} Z'_{\cdot H} Z_{\cdot H} - I_{HH}\| \le \zeta(n, \mathrm{card}(H)) = o(\lambda_1^2 e_n),$$

$$\|E_{MM,3}\| \le \frac{2}{n} \lambda_1 \|V' Z_{\cdot H}\| \|\bar{Q}_{M\cdot}\|_{\mathrm{F}} \le C\lambda_1 \xi(n, \bar{m}, \mathrm{card}(H); 2) \|\bar{Q}_{B_-^c \cdot}\|_{\mathrm{F}} = o(\lambda_1^2 e_n).$$

Thus, we combine the three parts to obtain $\|S_{MM} - I_{MM}\| = o(\lambda_1^2 e_n)$.

Condition SP implies $e_n = o(1)$. So (B.21), together with the bounds on $\|S_{MB^\circ}\|$ and $\|S_{MM} - I_{MM}\|$, leads to $|\ell_j(S_0^\circ) - \widehat{\ell}_j^\circ| = o(\lambda_1^2)$, for $j = 1, \ldots, \bar{m} + 1$. Noting that $\ell_j(S_0^\circ) = \ell_j(S_{B^\circ B^\circ}) \vee 1$, we complete the proof of claim (2). The proof here does not need the condition $\mathrm{AD}(m, \kappa)$.

9

Proof of claim (3) Under the condition $\mathrm{AD}(m, \kappa)$, $\lambda_m^2 \geq \lambda_1^2/\kappa$. This, together with claim (2) and Lemmas 6.1 and 6.2, implies $|\ell_j(S_0^o) - (\lambda_j^2 + 1)| = o(\lambda_j^2)$ for $j = 1, \ldots, m$. Hence, $\ell_j(S_0^o) = \ell_j(S_{B \circ B \circ}) > 1$ for $j = 1, \ldots, m$, and so $\widehat{\mathcal{P}}_m^{(0),o}$ is the leading principal subspace of S_0^o. We then apply Theorem D.1 with $A = S^o$, $A + E = S_0^o$ and $\delta = \ell_m(S_0^o) - \ell_{m+1}(S^o)$. Claim (1) and Lemma 6.2 imply that $\delta = (\lambda_m^2 - \lambda_{m+1}^2)(1 + o(1))$. Therefore, Theorem D.1 leads to

$$L(\widehat{\mathcal{P}}_m^o, \widehat{\mathcal{P}}_m^{(0),o}) \leq \frac{\|S^o - S_0^o\|^2}{(\lambda_m^2 - \lambda_{m+1}^2)^2}(1 + o(1)) \leq \frac{C\lambda_1^4 e_n^2}{(\lambda_m^2 - \lambda_{m+1}^2)^2} \leq Ce_n^2 \leq \frac{1}{5}(1 - \rho)^2.$$

The last inequality comes from conditions SP, $\mathrm{AD}(m, \kappa)$ and the definition of ρ, and holds when n is sufficiently large.

Proof of claim (4) On the event that claims (1) and (2), and conclusions of Lemmas 6.1–6.2 hold, we have $\left|\frac{\ell_1^B}{\lambda_1^2+1} - 1\right| = \frac{|\ell_1^B - (\lambda_1^2 + 1)|}{\lambda_1^2 + 1} = \frac{o(\lambda_1^2)}{\lambda_1^2 + 1} = o(1)$. On the other hand, $\mathrm{AD}(m, \kappa)$ ensures $(\lambda_m^2 - \lambda_{m+1}^2) \geq \lambda_1^2/\kappa$, and so

$$\left|\frac{\ell_m^B - \ell_{m+1}^B}{\lambda_m^2 - \lambda_{m+1}^2} - 1\right| = \frac{|\ell_m^B - (\lambda_m^2 + 1) - [\ell_{m+1}^B - (\lambda_{m+1}^2 + 1)]|}{\lambda_m^2 - \lambda_{m+1}^2}$$

$$\leq \frac{\kappa}{\lambda_1^2}\left(|\ell_m^B - (\lambda_m^2 + 1)| + |\ell_{m+1}^B - (\lambda_{m+1}^2 + 1)|\right) = \frac{\kappa}{\lambda_1^2} \cdot o(\lambda_1^2) = o(1).$$

Compare the definitions of K_s and K. The last two displays imply that as $n \to \infty$, $K_s/K \to 1.1$. Hence, for large enough n, $K_s \in [K, 2K]$.

B.4 Proof of Corollary 6.1

$1°$ We first prove the claims on $(\ell_1^B - 1)/(\ell_j^B - \ell_{j+1}^B)$ only assuming conditions GR and SP. To this end, we note that without the condition $\mathrm{AD}(m, \kappa)$, claims (1) of Lemmas 6.1 and 6.2, and claims (1) and (2) of Lemma 6.3 hold.

Suppose for some $1 \leq j \leq \bar{m}$, $\lim_{n \to \infty} \lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2) < \infty$. Then there exists some constant $C > 0$, such that for large n, $\lambda_j^2 - \lambda_{j+1}^2 > \lambda_1^2/C$. Similar to the proof of claim (4) of Lemma 6.3, on the event such that claims (1) of Lemmas 6.1 and 6.2 and claims (1) and (2) of Lemma 6.3 hold, we have $|\frac{\ell_1^B - 1}{\lambda_1^2} - 1| = \frac{|\ell_1^B - (\lambda_1^2 + 1)|}{\lambda_1^2} = \frac{o(\lambda_1^2)}{\lambda_1^2} = o(1)$, and

$$\left|\frac{\ell_j^B - \ell_{j+1}^B}{\lambda_j^2 - \lambda_{j+1}^2} - 1\right| = \frac{|\ell_j^B - (\lambda_j^2 + 1) - [\ell_{j+1}^B - (\lambda_{j+1}^2 + 1)]|}{\lambda_m^2 - \lambda_{j+1}^2}$$

$$\leq \frac{C}{\lambda_1^2}|\ell_j^B - (\lambda_j^2 + 1) - [\ell_{j+1}^B - (\lambda_{j+1}^2 + 1)]|$$

$$\leq \frac{C}{\lambda_1^2}\left(|\ell_j^B - (\lambda_j^2 + 1)| + |\ell_{j+1}^B - (\lambda_{j+1}^2 + 1)|\right) = \frac{C}{\lambda_1^2} \cdot o(\lambda_1^2) = o(1).$$

The last two displays then jointly imply that $\frac{(\ell_1^B - 1)/(\ell_j^B - \ell_{j+1}^B)}{\lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2)} = \frac{\ell_1^B - 1}{\lambda_1^2} \cdot \frac{\lambda_j^2 - \lambda_{j+1}^2}{\ell_j^B - \ell_{j+1}^B} = 1 + o(1)$. Since the event that the above result holds occurs with probability at least $1 - C_0 p_n^{-2}$ and $p_n^{-2} \leq n^{-2}$ is a summable sequence, Borel-Cantelli implies that the convergence to 1 happens almost surely.

10

Next suppose for some $1 \leq j \leq \bar{m}$, $\lim_{n\to\infty} \lambda_1^2/(\lambda_j^2 - \lambda_{j+1}^2) = \infty$. On the event that claims (1) of Lemmas 6.1 and 6.2 and claims (1) and (2) of Lemma 6.3 hold, $\ell_1^B - 1 = \lambda_1^2(1 + o(1))$, and $\ell_l^B = (\lambda_l^2 + 1) + o(\lambda_1^2)$, for $l = j, j+1$. Therefore, we have $\frac{\ell_j^B - \ell_{j+1}^B}{\ell_1^B - 1} = \frac{\lambda_j^2 - \lambda_{j+1}^2 + o(\lambda_1^2)}{\lambda_1^2(1+o(1))} = o(1)$. Here, the last equality holds as $(\lambda_j^2 - \lambda_{j+1}^2)/\lambda_1^2 \to 0$. Thus, for any constant $K > 0$, $\liminf_{n\to\infty}(\ell_1^B - 1)/(\ell_j^B - \ell_{j+1}^B) > K$. As before, Borel-Cantelli further shows that this happens almost surely. Thus, $\lim_{n\to\infty}(\ell_1^B - 1)/(\ell_j^B - \ell_{j+1}^B) = \infty$, a.s.

$2°$ Turn to the proof for ℓ_j^B. Under condition AD(m, κ), $\lambda_1^2/\lambda_j^2 \leq \kappa$ for all $1 \leq j \leq m$. Then on the event such claims (1) of Lemmas 6.1 and 6.2 and claims (1) and (2) of Lemma 6.3 hold, for any $1 \leq j \leq m$, $|\frac{\ell_j^B - 1}{\lambda_j^2} - 1| \leq \frac{|\ell_j^B - (\lambda_j^2 + 1)|}{\lambda_j^2} \leq \frac{o(\lambda_1^2)}{\lambda_1^2/\kappa} = o(1)$. Again, Borel-Cantelli further shows the convergence occurs almost surely. This completes the proof.

B.5 Proof of Proposition 6.1

Here, we restrict the proof on the event such that the conclusions of Lemmas 6.1–6.3 hold. Thus, all the arguments are deterministic. Recall that $T^{(k),\mathrm{o}}$ denotes the matrix obtained after multiplication in the k-th iteration of oracle Algorithm 1. Denote $\mathrm{ran}(T^{(k),\mathrm{o}})$ by $\mathcal{T}^{(k),\mathrm{o}}$, and let $\phi^{(k)}$ be the largest canonical angle between $\mathcal{T}^{(k),\mathrm{o}}$ and $\widehat{\mathcal{P}}_m^{(k-1),\mathrm{o}}$. So, $\sin^2 \phi^{(k)} = L(\mathcal{T}^{(k),\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k-1),\mathrm{o}})$. In what follows, we first prove both claims for the first iteration, and then extend them to subsequent iterations.

The First Iteration Claim (1) relies on the following lemma, which describes the effects of the multiplication and the thresholding steps separately.

Lemma B.4. *In the first iteration of oracle Algorithm 1,*

(1) *After the multiplication step, $T^{(1),\mathrm{o}}$ has full column rank and $\sin \phi^{(1)} \leq \rho \tan \theta^{(0)}$;*

(2) *After the thresholding step, $\widehat{Q}^{(1),\mathrm{o}}$ has full column rank, and $L(\mathcal{T}^{(1),\mathrm{o}}, \widehat{\mathcal{P}}_m^{(1),\mathrm{o}}) \leq \omega^2 \sec^2 \theta^{(0)}$.*

Proof. Claim (1) is essentially one-step analysis of orthogonal iteration, which is obtained by directly applying Theorem 8.2.2 in [3] to a single iteration.

The proof of claim (2) relies on Wedin's $\sin \theta$ theorem for singular subspaces (Theorem D.2). In the theorem, let $A = T^{(0),\mathrm{o}}$ and $B = \widehat{T}^{(1),\mathrm{o}}$. Then, the theorem gives that

$$L(\mathcal{T}^{(1),\mathrm{o}}, \widehat{\mathcal{P}}_m^{(1),\mathrm{o}}) \leq \frac{\|\widehat{T}^{(1),\mathrm{o}} - T^{(1),\mathrm{o}}\|^2}{\sigma_m^2(T^{(1),\mathrm{o}})}. \tag{B.22}$$

We now bound the numerator and the denominator of the right side separately.

First, we derive a lower bound for $\sigma_m(T^{(1),\mathrm{o}})$. For any unit vector $x \in \mathbb{R}^m$, let $y = \widehat{Q}^{(0),\mathrm{o}}x$, then y is a unit vector in $\mathbb{R}^p$. Decompose y as $y = \tilde{y} + \tilde{y}_c$, with $\tilde{y} \in \mathrm{ran}(\widehat{Q}^{\mathrm{o}})$ and $\tilde{y}_c \in \mathrm{ran}(\widehat{Q}_c^{\mathrm{o}})$. It follows that $\|T^{(1),\mathrm{o}}x\|_2^2 = \|S^{\mathrm{o}}y\|_2^2 = \|S^{\mathrm{o}}\tilde{y}\|_2^2 + \|S^{\mathrm{o}}\tilde{y}_c\|_2^2 \geq \|S^{\mathrm{o}}\tilde{y}\|_2^2 \geq (\widehat{\ell}_m^{\mathrm{o}})^2 \|\tilde{y}\|_2^2$. Since $\|\tilde{y}\|_2^2 \geq \cos^2 \theta^{(0)}$,

$$\sigma_m^2(T^{(1),\mathrm{o}}) \geq \inf_{\|x\|_2 = 1} \|T^{(1),\mathrm{o}}x\|_2^2 \geq (\widehat{\ell}_m^{\mathrm{o}})^2 \cos^2 \theta^{(0)}. \tag{B.23}$$

To bound the numerator, we define matrix $\Delta T \in \mathbb{R}^{p \times m}$, whose (ν, j)-th entry is given by $(\Delta T)_{\nu j} = \gamma_{nj} \mathbf{1}_{(\nu \in H)}$. Then, we obtain

$$\|\widehat{T}^{(1),\mathrm{o}} - T^{(1),\mathrm{o}}\| \leq \|\widehat{T}^{(1),\mathrm{o}} - T^{(1),\mathrm{o}}\|_{\mathrm{F}} \leq \|\Delta T\|_{\mathrm{F}} = \widehat{\ell}_m^{\mathrm{o}} \omega. \tag{B.24}$$

Here, the second inequality comes from Eq.(2.2) and (3.7) of [7].

Plugging the bounds (B.23) and (B.24) into (B.22), we obtain the bound for $L(\mathcal{T}^{(1),\mathrm{o}}, \widehat{\mathcal{P}}_m^{(1),\mathrm{o}})$ in the lemma. Finally, we are to show that $\widehat{T}^{(1),\mathrm{o}}$ has full column rank. By Corollary 8.6.2 of [3] and (B.24), it suffices to show that $\|\widehat{T}^{(1),\mathrm{o}} - T^{(1),\mathrm{o}}\| \leq \widehat{\ell}_m^{\mathrm{o}} \omega < \widehat{\ell}_m^{\mathrm{o}} \cos \theta^{(0)}$, i.e., $\omega < \cos \theta^{(0)}$. To this end, we note that Lemma 6.3 shows $\cos^2 \theta^{(0)} \geq 4/5$, while (B.27) shows that $\omega = o(1)$. Therefore, $\omega < \cos^2 \theta^{(k_0)}$ for large enough n, which completes the proof. $\qquad\square$

By the discussion following the loss function [7, Eq.(2.1)], we have the triangle inequality
$$\sin \theta^{(1)} \leq \sin \phi^{(1)} + L^{1/2}(\mathcal{T}^{(1),\mathrm{o}}, \widehat{\mathcal{P}}_m^{(1),\mathrm{o}}) \leq \rho \tan \theta^{(0)} + \omega \sec \theta^{(0)}.$$
Here, the second inequality comes from Lemma B.4. This completes the proof of claim (1).

Turn to claim (2). We first show that, for any $a \leq 1/2$,
$$\sin^2 \theta^{(0)} > \frac{1.01\omega^2}{(1-a)^2(1-\rho)^2} \tag{B.25}$$
implies $\sin^2 \theta^{(1)} / \sin^2 \theta^{(0)} \leq [1 - a(1-\rho)]^2$. By claim (1), it suffices to show
$$\frac{\rho}{\cos \theta^{(0)}} + \frac{\omega}{\sin \theta^{(0)} \cos \theta^{(0)}} \leq 1 - a(1-\rho).$$
Let $x = \sin \theta^{(0)}$. Multiply both sides of the last display with $x\sqrt{1-x^2}$ and collect terms. Then it is equivalent to $\omega \leq x\{[1 - a(1-\rho)]\sqrt{1-x^2} - \rho\}$. The right side of last inequality is non-negative, because $x^2 \leq (1-\rho)^2/5$ (by Lemma 6.3) and $a \leq 1/2$. So, we can square both sides, and simple algebra shows that the inequality holds if
$$\mathsf{Q}(x^2) = [1 - a(1-\rho)]^2 x^4 - (1-a)^2(1-\rho)^2 x^2 + \omega^2 \leq 0. \tag{B.26}$$
Here, $\mathsf{Q}(x^2)$ is a quadratic form of x^2, with discriminant $\Delta = (1-a)^4(1-\rho)^4 - 4\omega^2[1 - a(1-\rho)]^2$. Note that Lemmas 6.1–6.3 lead to
$$\begin{aligned}
\frac{\omega^2}{(1-\rho)^4} &= \mathrm{card}(H)\gamma^2 \frac{\sum_{j=1}^m (\lambda_j^2 + 1)}{(\lambda_m^2 + 1)^2} \frac{\log(p_n)}{n} \left(1 - \frac{\lambda_{m+1}^2 + 1}{\lambda_m^2 + 1}\right)^{-4} (1 + o(1)) \\
&\leq C(\lambda_m^{-4} \vee 1) M_n \tau_{nm}^2.
\end{aligned} \tag{B.27}$$
Here, $C = C(\gamma, r, m, \kappa)$. Moreover, conditions GR and SP imply that the rightmost side of (B.27) tends to 0 as $n \to \infty$. Therefore, for sufficiently large $n \geq n_0(\gamma, r, m, \kappa)$, we have $\frac{\omega^2}{(1-\rho)^4} \leq \frac{3}{16} \leq \frac{(1-a)^4}{4[1-a(1-\rho)]^2}$, and so $\Delta > 0$. Let the two roots of Q be $x_\pm^2$. We solve Q to obtain that (B.26), and hence the conclusion, holds when
$$x^2 \geq x_-^2 = \frac{\omega^2(1 + o(1))}{(1-a)^2(1-\rho)^2}. \tag{B.28}$$
By (B.27), $x^2 \leq (1-\rho)^2/5 \leq x_+^2$ always holds. So the case is completed as (B.25) implies (B.28).

Next, we show that if $\sin^2 \theta^{(0)} \leq 1.01\omega^2/[(1-a)(1-\rho)]^2$, then so is $\sin^2 \theta^{(1)}$. By claim (1), it suffices to show that $\rho \tan \theta^{(0)} + \omega \sec \theta^{(0)} \leq \sqrt{1.01}\omega/[(1-a)(1-\rho)]$. Multiply both sides by $\cos \theta^{(k_0)} \frac{(1-a)(1-\rho)}{\sqrt{1.01}\omega}$. We obtain that the last inequality holds when $\rho + \frac{(1-a)(1-\rho)}{\sqrt{1.01}} \leq \sqrt{1 - \frac{1.01\omega^2}{(1-a)^2(1-\rho)^2}}$, because the right side is no greater than $\cos \theta^{(0)}$ by the condition on $\sin \theta^{(0)}$. Since $\sqrt{x} \geq x$ for $x \in [0,1]$, it further suffices to have the last inequality hold but with the square root of the right side removed, which is equivalent to have $\frac{1.01\omega^2}{(1-a)^2(1-\rho)^2} \leq (1-\rho)[1 - (1-a)/\sqrt{1.01}]$. By (B.27), the last desired inequality holds when $n \geq n_0(\gamma, r, m, \kappa)$. This completes the case.

Subsequent Iterations Now we have $\widehat{Q}^{(1),\mathrm{o}}$ is orthonormal with $\sin^2\theta^{(1)} \leq (1-\rho)^2/5$. Therefore, we can essentially repeat the whole argument to prove both claims for $k=2$, and so on by induction. This is valid because the minimum sample size needed, $n_0 = n_0(\gamma, r, m, \kappa)$, such that both claims hold does not depend on the iteration index k.

B.6 Proof of Proposition 6.2

We divide the proof into two parts, depending on whether the following condition holds or not:

$$1.01\omega^2/(1-\rho)^2 \geq \epsilon_{nm}^2/4. \tag{B.29}$$

Without loss of generality, we assume that $n = 2^l$ for some $l \geq 1$. So, $l = \log n/\log 2$.

When (B.29) holds, let k_1 be the number of iterations needed to achieve $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k_1),\mathrm{o}}) \leq 1.01\omega^2/[(1-\frac{1}{2})(1-\rho)]^2$. Before the last inequality is satisfied, the decay rate of the approximation error satisfies the bound in claim (2) of Proposition 6.1 with $a = 1/2$. Thus, it suffices to have k_1 satisfy $[1 - (1-\rho)/2]^{2k_1} \leq 1.01\omega^2/[(1-\frac{1}{2})(1-\rho)]^2$, i.e., $2k_1|\log(1 - \frac{1}{2}(1-\rho))| \geq \log[(1-\frac{1}{2})^2(1-\rho)^2/1.01\omega^2]$. Condition (B.29) implies $(1-1/2)^2(1-\rho)^2/1.01\omega^2 \leq \epsilon_{nm}^{-2} \leq nh(\lambda_m^2)$. In addition, $|\log(1-x)| \geq x$, for all $x \in (0,1)$. Thus, it suffices to set

$$k_1 = \frac{\log n + 0 \vee \log h(\lambda_m^2)}{1-\rho} = \frac{\lambda_m^2 + 1}{\lambda_m^2 - \lambda_{m+1}^2}\left[\log n + 0 \vee \log h(\lambda_m^2)\right](1 + o(1)).$$

Now, let $k_2 - k_1$ be the number of additional iterations needed to achieve $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k_2),\mathrm{o}}) \leq 1.01\omega^2/[(1-\frac{1}{4})(1-\rho)]^2$. Then, before the inequality is satisfied, the decay rate satisfies claim (2) of Proposition 6.1 with $a = 1/4$. Thus, it suffices to have $[1 - \frac{1}{4}(1-\rho)]^{2(k_2-k_1)} \leq (1-\frac{1}{2})^2/(1-\frac{1}{4})^2$. This is guaranteed if we set $k_2 - k_1 \geq \frac{4}{1-\rho}\left[\log\left(1-\frac{1}{4}\right) - \log\left(1-\frac{1}{2}\right)\right]$. Recursively, we define k_i for $i = 3, \ldots, l$, such that $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k_i),\mathrm{o}}) \leq 1.01\omega^2/[(1-2^{-i})(1-\rho)]^2$. Repeating the above argument shows it suffices to have $k_i - k_{i-1} \geq \frac{2^i}{1-\rho}\left[\log\left(1-\frac{1}{2^i}\right) - \log\left(1-\frac{1}{2^{i-1}}\right)\right]$, for $i = 3, \ldots, l$. Therefore, if we set

$$k_l - k_1 = \frac{l + 1/2}{1-\rho} = \frac{\lambda_m^2 + 1}{\lambda_m^2 - \lambda_{m+1}^2}\frac{\log n}{\log 2}(1 + o(1)) \geq \sum_{i=2}^{l}\frac{2^i}{1-\rho}\left[\log\left(1-\frac{1}{2^i}\right) - \log\left(1-\frac{1}{2^{i-1}}\right)\right],$$

then $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k),\mathrm{o}}) \leq 1.01\omega^2/[(1-n^{-1})(1-\rho)]^2$ for all $k \geq k_l$. We complete the proof for this case by noting that $K \geq k_l$ for large enough n and that $\frac{\omega^2}{(1-\rho)^2} \leq Cm\gamma^2\mathrm{card}(H)h(\lambda_m^2) \leq CM_n\tau_{nm}^2$.

Turn to the case when (B.29) does not hold. Then, the decay rate in claim (2) of Proposition 6.1 holds with $a = 1/2$ for all k such that $L(\widehat{\mathcal{P}}_m^{\mathrm{o}}, \widehat{\mathcal{P}}_m^{(k),\mathrm{o}}) > \epsilon_{nm}^2$. Since the approximation error is monotone decreasing, it suffices to verify that $[1 - (1-\rho)/2]^{2K} \leq \epsilon_{nm}^2$, i.e.,

$$2K\left|\log\left(1 - \frac{1-\rho}{2}\right)\right| \geq \log\frac{n}{\log(p_n)} + \log\frac{(\lambda_m^2 - \lambda_{m+1}^2)^2}{(\lambda_1^2 + 1)(\lambda_{m+1}^2 + 1)}.$$

Note that the second term on the right side is bounded above by $\log h(\lambda_m^2)$. Since $\log(p_n) = o(n)$ and $|\log(1-x)| \geq x$ for $x \in (0,1)$, it suffices to verify $K \geq \frac{1}{1-\rho}\left[\log n + \log h(\lambda_m^2)\right] = \frac{\lambda_m^2 + 1}{\lambda_m^2 - \lambda_{m+1}^2}\log n\,(1 + o(1))$, which is indeed satisfied by K for large enough n. This completes the proof.

13

B.7 Proof of Lemma 6.4

In the proof, we restrict our attention to the event on which conclusions of Lemmas 6.1–6.3 hold.. Lemma B.2 shows that $B^{\mathrm{o}} = B$, which leads to $\widehat{Q}^{(0)} = \widehat{Q}^{(0),\mathrm{o}}$. In what follows, we first show that $\widehat{Q}^{(1)} = \widehat{Q}^{(1),\mathrm{o}}$, and then extend the argument to subsequent iterations.

The First Iteration Fix $\nu \in L$ and $1 \le j \le m$. Let $\widehat{Q}^{(0),\mathrm{o}} = [\widehat{q}_1^{(0),\mathrm{o}}, \dots, \widehat{q}_m^{(0),\mathrm{o}}]$. Since $\widehat{Q}^{(0)} = \widehat{Q}^{(0),\mathrm{o}}$, the (ν, j)-th element of $T^{(1)}$ is then $t_{\nu j}^{(1)} = S_{\nu \cdot} \widehat{q}_j^{(0),\mathrm{o}} = S_{\nu H} \widehat{q}_{jH}^{(0),\mathrm{o}}$. Here, the second equality holds because $\widehat{q}_{jL}^{(0),\mathrm{o}} = 0$. Our goal is to show that for sufficiently large numeric constant γ, the size of $t_{\nu j}^{(1)}$ is bounded by the threshold γ_{nj} in Theorem 3.2 of [7] with high probability. Following Corollary 6.1, it suffices to show that $t_{\nu j}^{(1)}$ is bounded by a constant multiple of $\sqrt{(\lambda_j^2 + 1) \log(p_n)/n}$.

By (B.14), we decompose $t_{\nu j}^{(1)}$ as

$$t_{\nu j}^{(1)} = S_{\nu H} \left[Q_{H\cdot}^{\mathrm{o}} (Q_{H\cdot}^{\mathrm{o}})' + P_H^{\mathrm{o}} (P_H^{\mathrm{o}})' \right] \widehat{q}_{jH}^{(0),\mathrm{o}} = t_1^{(1)} + t_2^{(1)}. \tag{B.30}$$

Moreover, $S_{\nu H}$ admits the decomposition $S_{\nu H} = \sum_{i=1}^4 S_{\nu H, i}$, where

$$S_{\nu H, 1} = \frac{1}{n} \bar{Q}_{\nu\cdot} \Lambda V' V \Lambda \bar{Q}'_{H\cdot}, \qquad S_{\nu H, 2} = \frac{1}{n} Z'_{\cdot\nu} Z_{\cdot H}$$

$$S_{\nu H, 3} = \frac{1}{n} \bar{Q}_{\nu\cdot} \Lambda V' Z_{\cdot H}, \qquad S_{\nu H, 4} = \frac{1}{n} Z'_{\cdot\nu} V \Lambda \bar{Q}'_{H\cdot}. \tag{B.31}$$

In what follows, we bound $t_i^{(1)}$, $i = 1, 2$, respectively. The argument is lengthy but elementary. The key conclusions are (B.32) and (B.37), which then lead to (B.44).

1^{o}. *Bound for $t_1^{(1)}$.* We show below

$$|t_1^{(1)}| \le (1 + o(1)) \, \|(Q_{H\cdot}^{\mathrm{o}})' \widehat{q}_{jH}^{(0),\mathrm{o}}\|_2 (\beta \sqrt{m} + 2\sqrt{3}\,) \gamma_{nj}/\gamma. \tag{B.32}$$

To this end, using (B.31), we decompose $t_1^{(1)} = \sum_{i=1}^4 S_{\nu H, i} Q_{H\cdot}^{\mathrm{o}} (Q_{H\cdot}^{\mathrm{o}})' \widehat{q}_{jH}^{(0),\mathrm{o}} = \sum_{i=1}^4 t_{1i}^{(1)}$.

For $t_{11}^{(1)}$, we decompose it as $t_{11}^{(1)} = \bar{Q}_{\nu\cdot} \Lambda^2 \bar{Q}'_{H\cdot} Q_{H\cdot}^{\mathrm{o}} (Q_{H\cdot}^{\mathrm{o}})' \widehat{q}_{jH}^{(0),\mathrm{o}} + \bar{Q}_{\nu\cdot} \Lambda \left(\frac{1}{n} V'V - I \right) \Lambda \bar{Q}'_{H\cdot} Q_{H\cdot}^{\mathrm{o}} (Q_{H\cdot}^{\mathrm{o}})' \widehat{q}_{jH}^{(0),\mathrm{o}}$. Recall (B.4). We have $\|\bar{Q}_{\nu\cdot} \Lambda^2 \bar{Q}'_{H\cdot} Q_{H\cdot}^{\mathrm{o}}\| \le \|Q_{\nu\cdot} \Lambda_0^2\|_2 \|Q'_{H\cdot} Q_{H\cdot}^{\mathrm{o}}\| + \|Q_{1,\nu\cdot} \Lambda_1^2\|_2 \|Q'_{1,H\cdot} Q_{H\cdot}^{\mathrm{o}}\|$. Let

$$\rho_{lj} = (\lambda_l^2 + 1)/(\lambda_j^2 + 1), \quad 1 \le l, j \le \bar{m}. \tag{B.33}$$

Since $\nu \in L$, we have $|q_{\nu j}| \le \beta \tau_{nj}$, and so $\|Q_{\nu\cdot} \Lambda_0^2\|_2 \le (\sum_{l=1}^m \beta^2 \tau_{nl}^2 \lambda_l^4)^{1/2} = (\sum_{l=1}^m \rho_{lj})^{1/2} (\beta/\gamma) \gamma_{nj}$, with γ and γ_{nj} specified as in Theorem 3.1 of [7]. Similarly, $\|Q_{1,\nu\cdot} \Lambda_1^2\|_2 \le (\sum_{l=m+1}^{\bar{m}} \rho_{lj})^{1/2} (\beta/\gamma) \gamma_{nj}$. Moreover, we have $\|Q'_{H\cdot} Q_{H\cdot}^{\mathrm{o}}\| \le 1$ and $\|Q'_{1,H\cdot} Q_{H\cdot}^{\mathrm{o}}\| \le \|Q_1 Q^{\mathrm{o}}\| \le L^{1/2}(\mathcal{P}_m, \mathcal{P}_m^{\mathrm{o}}) = o(1)$. Therefore, $\|\bar{Q}_{\nu\cdot} \Lambda^2 \bar{Q}'_{H\cdot} Q_{H\cdot}^{\mathrm{o}}\| \le (1 + o(1)) (\sum_{l=1}^m \rho_{lj})^{1/2} \beta \gamma_{nj}/\gamma$. On the other hand, on the event (B.9), we have $\|\bar{Q}_{\nu\cdot} \Lambda \left(\frac{1}{n} V'V - I \right) \Lambda \bar{Q}'_{H\cdot}\| \le \|\bar{Q}_{\nu\cdot} \Lambda\|_2 \|\frac{1}{n} V'V - I\| \|\Lambda \bar{Q}'_{H\cdot}\| \le \lambda_1 \zeta(n, \bar{m}) \left(\sum_{l=1}^{\bar{m}} \rho_{lj}/\lambda_l^2 \right)^{1/2} \beta \gamma_{nj}/\gamma = o(\gamma_{nj})$. Here, the last equality comes from conditions GR and SP. Finally, the triangle inequality leads to

$$|t_{11}^{(1)}| \le (1 + o(1)) \left(\sum_{l=1}^m \rho_{lj} \right)^{1/2} \|(Q_{H\cdot}^{\mathrm{o}})' \widehat{q}_{jH}^{(0),\mathrm{o}}\|_2 \, \beta \gamma_{nj}/\gamma.$$

14

For $t_{12}^{(1)}$, define event

$$\left\{ \|Z'_{\cdot\nu} Z_{\cdot H} Q^{\mathrm{o}}_{H\cdot}\|_2 \le n\xi(n, 1, m; \sqrt{6}), \forall \nu \in L \right\}. \tag{B.34}$$

Note that $\xi(n, 1, m; \sqrt{6}) = (1 + o(1))[6/(\lambda_j^2 + 1)]^{1/2}(\beta/\gamma)\gamma_{nj}$. On this event, we have $|t_{12}^{(1)}| \le$
$(1 + o(1)) \left(\frac{6}{\lambda_j^2 + 1}\right)^{1/2} \|(Q^{\mathrm{o}}_{H\cdot})' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2 \frac{\gamma_{nj}}{\gamma}$. For $t_{13}^{(1)}$, on the event (B.18), we have

$$
\begin{aligned}
|t_{13}^{(1)}| &\le \frac{1}{n} \|\bar{Q}_{\nu\cdot}\Lambda\|_2 \|V' Z_{\cdot H} Q^{\mathrm{o}}_{H\cdot}\| \|(Q^{\mathrm{o}}_{H\cdot})' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2 \\
&\le (1 + o(1)) \, 2\sqrt{\frac{\log(p_n)}{n}} \left(\sum_{l=1}^{\bar{m}} \frac{\rho_{lj}}{\lambda_l^2}\right)^{1/2} \|(Q^{\mathrm{o}}_{H\cdot})' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2 \frac{\beta\gamma_{nj}}{\gamma} = o(\gamma_{nj}).
\end{aligned}
$$

Switch to $t_{14}^{(1)}$. We note that $\|\Lambda\bar{Q}'_{H\cdot}Q^{\mathrm{o}}_{H\cdot}\| \le \lambda_1(1 + o(1))$. So, on the event

$$\left\{ \|Z'_{\cdot\nu} V\|_2 \le n\xi(n, 1, \bar{m}; \sqrt{6}), \forall \nu \in L \right\}, \tag{B.35}$$

we have $|t_{14}^{(1)}| \le (1 + o(1)) \left(\frac{6\lambda_1^2}{\lambda_j^2 + 1}\right)^{1/2} \|(Q^{\mathrm{o}}_{H\cdot})' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2 \frac{\gamma_{nj}}{\gamma}$. Finally, we apply the triangle inequality
to obtain

$$|t_1^{(1)}| \le \sum_{i=1}^4 |t_{1i}^{(1)}| \le (1 + o(1)) \, \|(Q^{\mathrm{o}}_{H\cdot})' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2 \left[\beta\left(\sum_{l=1}^m \rho_{lj}\right)^{1/2} + 2\sqrt{3}\,\rho_{1j}^{1/2}\right] \frac{\gamma_{nj}}{\gamma}. \tag{B.36}$$

Without loss of generality, we could assume that $\lambda_1/\lambda_m \to 1$ as $n \to \infty$. When this is the case,
(B.36) reduces to (B.32). If this is not the case, let $1 < m_0 \le m$ be the smallest index such that
$\lambda_{m_0}/\lambda_m \to 1$. We further divide the set $\{1, \ldots, m\}$ into $\{1, \ldots, m_0 - 1\} \cup \{m_0, \ldots, m\}$, and similar
but lengthier argument would also lead to (B.32).

2^{o}. *Bound for* $t_2^{(1)}$. Here, we show

$$|t_2^{(1)}| \le (1 + o(1)) \|(P^{\mathrm{o}}_H)' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2 \left[\beta\left(\sum_{l=m+1}^{\bar{m}} \rho_{lj}\right)^{1/2} + \sqrt{7}\,\rho_{1j}^{1/2}\right] \frac{\gamma_{nj}}{\gamma}. \tag{B.37}$$

To this end, observe that (B.31) leads to $t_2^{(1)} = \sum_{i=1}^4 S_{\nu H, i} P^{\mathrm{o}}_H (P^{\mathrm{o}}_H)' \hat{q}_{jH}^{(0),\mathrm{o}} = \sum_{i=1}^4 t_{2i}^{(1)}$.

For $t_{21}^{(1)}$, by argument similar to that for $t_{11}^{(1)}$, we have on event (B.9),

$$|t_{21}^{(1)}| \le (1 + o(1)) \left(\sum_{l=m+1}^{\bar{m}} \rho_{lj}\right)^{1/2} \|(P^{\mathrm{o}}_H)' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2 \beta\gamma_{nj}/\gamma.$$

For $t_{22}^{(1)}$, define

$$y_{\nu j}^{(1)} = nt_{22}^{(1)} / \|Z_{\cdot H} P^{\mathrm{o}}_H (P^{\mathrm{o}}_H)' \hat{q}_{jH}^{(0),\mathrm{o}}\|_2. \tag{B.38}$$

By our construction of the initial value, $\widehat{q}_{jH}^{(0),o}$ is a function of S_{HH} and hence is independent of $Z_{.\nu}$. In addition, $Z_{.H}$ is independent of $Z_{.\nu}$. Thus, we have $y_{\nu j}^{(1)} \sim N(0,1)$. Define events

$$\left\{ |y_{\nu j}^{(1)}| \le \sqrt{7\log(p_n)}, 1 \le j \le m, \forall \nu \in L \right\}, \tag{B.39}$$

$$\left\{ \|Z_{.H}\| \le \sqrt{n} + \sqrt{\mathrm{card}(H)} + 2\sqrt{\log(p_n)} \right\}. \tag{B.40}$$

On the intersection of the above events, we have

$$|t_{22}^{(1)}| \le \frac{1}{n}|y_{\nu j}^{(1)}|\,\|Z_{.H}\|\|(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2 \le (1+o(\lambda_1))\|(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2 \left(\frac{7}{\lambda_j^2+1}\right)^{1/2}\frac{\gamma_{nj}}{\gamma}.$$

On the event (B.11), $|t_{23}^{(1)}| \le \frac{1}{n}\|\bar{Q}_{\nu.}\Lambda\|_2\|V'Z_{.H}\|\|(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2 \le \frac{C}{\lambda_{\bar{m}}}\sqrt{\frac{\mathrm{card}(H)}{n}}\|(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2\frac{\beta\gamma_{nj}}{\gamma}$. Condition (GR) implies

$$\frac{1}{\lambda_{\bar{m}}^2}\frac{\mathrm{card}(H)}{n} \le \frac{\beta^{-r}}{n\lambda_{\bar{m}}^2}\sum_{j=1}^{\bar{m}} s_j^r \tau_{nj}^{-r} = \frac{\beta^{-r}}{n^{1/2-r/4}[\log(p_n)]^{1/2+r/4}}\sum_{j=1}^{\bar{m}} s_j^r\left[\frac{\log(p_n)}{n\lambda_j^4}\right]^{1/2-r/4}\cdot\frac{\lambda_j^2}{\lambda_{\bar{m}}^2}\left(\frac{\lambda_j^2}{\lambda_j^2+1}\right)^{r/2}$$

$$\le \frac{\lambda_1^2}{\lambda_{\bar{m}}^2}\frac{\beta^{-r}}{n^{1/2-r/4}[\log(p_n)]^{1/2+r/4}}\sum_{j=1}^{\bar{m}} s_j^r\left[\frac{\log(p_n)}{n\lambda_j^4}\right]^{1/2-r/4} = o(1).$$

Thus, we obtain $|t_{23}^{(1)}| = o(\gamma_{nj})$.

For $t_{24}^{(1)}$, we could bound it using the same strategy as for $t_{22}^{(1)}$. In particular, let

$$w_{\nu j}^{(1)} = nt_{24}^{(1)}/\|V\Lambda\bar{Q}_{H.}P_H^o(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2 \tag{B.41}$$

For any $\nu \in L$, $Z_{.\nu}$ is independent of both V and $\widehat{q}_{jH}^{(0),o}$, and so $w_{\nu j}^{(1)} \sim N(0,1)$. Define events

$$\left\{ |w_{\nu j}^{(1)}| \le \sqrt{7\log(p_n)}, 1 \le j \le m, \forall \nu \in L \right\}, \tag{B.42}$$

$$\left\{ \|V\| \le \sqrt{n} + \sqrt{\bar{m}} + 2\sqrt{\log(p_n)} \right\}. \tag{B.43}$$

On their intersection, $|t_{24}^{(1)}| \le \frac{1}{n}|w_{\nu j}^{(1)}|\|V\|\|\Lambda\|\|(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2 \le (1+o(1))(\frac{7\lambda_1^2}{\lambda_j^2+1})^{1/2}\|(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2\frac{\gamma_{nj}}{\gamma}$. Finally, the triangle inequality leads to (B.37).

3^o.*Summary.* By Lemmas 6.1–6.3, we have $\|(Q_{H.}^o)'\widehat{q}_{jH}^{(0),o}\|_2 = 1 + o(1)$, $\|(P_H^o)'\widehat{q}_{jH}^{(0),o}\|_2 = o(1)$. This, together with (B.30), (B.32) and (B.37), implies that, if we choose $\gamma > 2\sqrt{3}$ and for $c = 0.9(\gamma - 2\sqrt{3})$, let $\beta = c/\sqrt{m}$, then when n is sufficiently large,

$$|t_{\nu j}^{(1)}| \le (1+o(1))(\beta\sqrt{m} + 2\sqrt{3})\gamma_{nj}/\gamma \le \gamma_{nj}. \tag{B.44}$$

Since the above inequality holds for all $\nu \in L$ and $1 \le j \le m$, we obtain $\widehat{T}_{L.}^{(1)} = 0$, and so, $\widehat{T}^{(1)} = \widehat{T}^{(1),o}$. Therefore, $\widehat{Q}^{(1)} = \widehat{Q}^{(1),o}$. Last but not least, Propositions D.1, D.2, Lemmas D.1 – D.2 and the Davidson–Szarek bound [1, Theorem II.13] ensure that the intersection of all the above defined events has probability at least $1 - C_0 p_n^{-2}$.

16

Subsequent Iterations Fix the choice of γ_{nj} such that (B.44) holds. Here, we extend the above calculation up to the $2K$-th iteration. The goal is to show that with high probability,

$$\widehat{Q}^{(k)} = \widehat{Q}^{(k),\text{o}}, \quad k = 1, \ldots, 2K. \tag{B.45}$$

To this end, we first recall the events defined in (B.39) and (B.42). Note that for any $k > 1$, $y_{\nu j}^{(k)}$ and $w_{\nu j}^{(k)}$ can be defined analogously by replacing $\widehat{q}_{jH}^{(0),\text{o}}$ with $\widehat{q}_{jH}^{(k),\text{o}}$. For any $1 \leq j \leq m$ and $k \geq 0$, $\widehat{q}_{jH}^{(k),\text{o}}$ is a function of V and $Z_{\cdot H}$, which is independent of $Z_{\cdot L}$, and so both $y_{\nu j}^{(k)}$ and $w_{\nu j}^{(k)}$ follows the $N(0,1)$ distribution. Thus, we define events

$$\left\{ |y_{\nu j}^{(k)}| \leq \sqrt{7\log(p_n)}, 1 \leq j \leq m, \forall \nu \in L, 1 \leq k \leq 2K \right\}, \tag{B.46}$$

$$\left\{ |w_{\nu j}^{(k)}| \leq \sqrt{7\log(p_n)}, 1 \leq j \leq m, \forall \nu \in L, 1 \leq k \leq 2K \right\}. \tag{B.47}$$

Under our assumption, for any given ν, j, and k, the tail bound of standard normal distribution gives $\mathbb{P}\left\{ |y_{\nu j}^{(k)}| > \sqrt{7\log(p_n)} \right\} \leq \frac{C}{\sqrt{\log(p_n)}} p_n^{-7/2}$. Since m is fixed, $\text{card}(L) \leq p_n$, and $K = O(\sqrt{n\log(p_n)})$, we obtain that

$$\mathbb{P}\left\{ |y_{\nu j}^{(k)}| \leq \sqrt{7\log(p_n)}, 1 \leq j \leq m, \forall \nu \in L, 1 \leq k \leq 2K \right\}$$

$$\geq 1 - \sum_{j=1}^{m} \sum_{\nu \in L} \sum_{k=1}^{2K} \mathbb{P}\left\{ |y_{\nu j}^{(k)}| > \sqrt{7\log(p_n)} \right\} \geq 1 - C_0 p_n^{-2}.$$

The same argument extends to the event in (B.47). Finally, on the intersection of (B.46), (B.47), and the other events involved in the calculation of the first iteration, we can essentially repeat the arguments for the first iteration $2K - 1$ times to obtain (B.45). This complete the proof.

C Proof of Proposition 3.1

Proof of claim (1) We note that

$$\mathbb{P}\left\{ \widehat{m} > \bar{m} \right\} \leq \mathbb{P}\left\{ \ell_{\bar{m}+1}(S_{BB}) > 1 + \delta_{|B|} \right\} \leq \mathbb{P}\left\{ \max_{|A|=|B|} \ell_{\bar{m}+1}(S_{AA}) > 1 + \delta_{|B|} \right\}$$

$$\leq \sum_{k=\bar{m}+1}^{p} \mathbb{P}\left\{ \max_{|A|=|B|} \ell_{\bar{m}+1}(S_{AA}) > 1 + \delta_{|B|}, |B| = k \right\}$$

$$\leq \sum_{k=\bar{m}+1}^{p} \mathbb{P}\left\{ \max_{|A|=k} \ell_{\bar{m}+1}(S_{AA}) > 1 + \delta_k \right\}.$$

By Proposition 1.2 of [4] and the interlacing property of eigenvalues in general, we have

$$\max_{|A|=k} \ell_{\bar{m}+1}(S_{AA}) \overset{st}{<} \max_{|A|=k-\bar{m}} \ell_1(W_{AA}) \overset{st}{<} \max_{|A|=k} \ell_1(W_{AA}).$$

Here, $nW \sim W_p(n, I_p)$ and $\overset{st}{<}$ means "stochastically smaller". Thus, we obtain

$$\mathbb{P}\left\{ \widehat{m} > \bar{m} \right\} \leq \sum_{k=\bar{m}+1}^{p} \mathbb{P}\left\{ \max_{|A|=k} \ell_1(W_{AA}) > 1 + \delta_k \right\}.$$

17

For each summand on the right side, we have that, for $nW_k \sim W_k(n, I_k)$,

$$\mathbb{P}\left\{\max_{|A|=k} \ell_1(W_{AA}) > 1 + \delta_k\right\} \leq \binom{p}{k} \mathbb{P}\left\{\ell_1(W_k) > 1 + \delta_k\right\}$$

$$\leq \left(\frac{ep}{k}\right)^k e^{-nt_k^2/2} \leq \left(\frac{ep}{k}\right)^k e^{-3\log p_n - k(\log p_n + 1)} \leq p_n^{-3}.$$

Here, the first inequality is union bound, while the second inequality comes from Proposition D.1 and the fact $\binom{p}{k} \leq (ep/k)^k$. Summing over all possible k's, we complete the proof of claim (1).

Proof of claim (2) For here, we need to show that, for sufficiently large n, $\ell_m(S_{BB}) > 1 + \delta_{|B|}$ with high probability. To this end, we first note that Lemmas 6.1–6.3 jointly imply that, for large enough n, with probability at least $1 - C_p p_n^{-2}$,

$$|\ell_m(S_{BB}) - (\lambda_m^2 + 1)| \leq \frac{\lambda_1^2}{3\kappa}. \tag{C.1}$$

On the other hand, following the proof of Lemma 6.3, on this event, we also have $B \subset B_+$, with B_+ the deterministic set defined in (B.19). Similar to the way that we bound $\text{card}(H)$ in Lemma A.1, we have $\text{card}(B_+) \leq \left(\frac{\bar{m}}{\alpha a_-}\right)^{r/2} \sum_{j=1}^{\bar{m}} s_j^r \left(\frac{n\lambda_j^4}{\log p_n}\right)^{r/4} = o\left((\lambda_1^4 \wedge 1)\left(\frac{\log p_n}{n\lambda_1^4}\right)^{-1/2}\right)$. Here, the last equality is guaranteed by condition SP. This leads to the bounds

$$\frac{\text{card}(B_+)\log p_n}{n\lambda_1^i} = o\left((\lambda_1^4 \wedge 1)\lambda_1^{2-i}\sqrt{\frac{\log p_n}{n}}\right) = o(1), \quad \text{for } i = 2, 4.$$

Since $B \subset B_+$, we thus obtain that $\text{card}(B) \leq \text{card}(B_+)$, and the last display leads to $\delta_{\text{card}(B)} = o(\lambda_1^2)$. So, for sufficiently large n, we have $\delta_{\text{card}(B)} \leq \lambda_1^2/(3\kappa)$. Together with (C.1), this leads to

$$\ell_m(S_{BB}) - 1 - \delta_{\text{card}(B)} = \ell_m(S_{BB}) - (\lambda_m^2 + 1) + \lambda_m^2 - \delta_{\text{card}(B)} \geq -\frac{\lambda_1^2}{3\kappa} + \lambda_m^2 - \frac{\lambda_1^2}{3\kappa} \geq \frac{\lambda_1^2}{3\kappa} > 0.$$

Here, the second last inequality comes from the condition $\text{AD}(m, \kappa)$, which implies that $\lambda_m^2 \geq \lambda_1^2/\kappa$. This completes the proof of claim (2).

Proof of claim (3) By claim (2), for sufficiently large n, we have $\bar{m} \leq \widehat{m}$, while claim (1) ensures that $\widehat{m} \leq \bar{m}$. The two inequalities jointly imply $\widehat{m} = \bar{m}$ with probability at least $1 - C_0 p_n^{-2}$.

D   Auxiliary Results

First, we present two $\sin\theta$ theorems. One is for principal subspaces, which comes from [2] and is used extensively in the proofs of Lemmas 6.1–6.3. The other is for singular subspaces, which comes from [9] and is used in the proof of Proposition 6.1.

Theorem D.1 ($\sin\theta$ theorem for principal subspaces). *Let A and $A + E$ be symmetric matrices satisfying $A = \begin{bmatrix} G_0 & G_1 \end{bmatrix} \begin{bmatrix} A_0 & 0 \\ 0 & A_1 \end{bmatrix} \begin{bmatrix} G_0' \\ G_1' \end{bmatrix}$ and $A + E = \begin{bmatrix} F_0 & F_1 \end{bmatrix} \begin{bmatrix} \Lambda_0 & 0 \\ 0 & \Lambda_1 \end{bmatrix} \begin{bmatrix} F_0' \\ F_1' \end{bmatrix}$, with $G_0' G_1 = 0$,*

G_1 orthonormal, and $[F_0 \; F_1]$ orthogonal. If the eigenvalues of $A_0 G_0' G_0$ are contained in an interval (a, b), and the eigenvalues of Λ_1 are excluded from the interval $(a - \delta, b + \delta)$ for some $\delta > 0$, then

$$L(\mathrm{ran}(G_0), \mathrm{ran}(F_0)) = \|F_1' G_0\|^2 \leq \frac{\|F_1' E G_0\|^2}{\delta^2 \sigma_{\min}(G_0)^2}.$$

Theorem D.2 ($\sin \theta$ theorem for singular subspaces). *Suppose that $p \geq k$. Let $p \times k$ matrices A and $B = A + E$ both have full column rank, and let $w_{k+1}, \ldots, w_p$ be orthonormal vectors spanning $\mathrm{ran}(B)^{\perp}$, the orthogonal complement of $\mathrm{ran}(B)$. Take $W = [w_{k+1}, \ldots, w_p]$, and define $R = A'W$. If $\sigma_{\min}(A) \geq \delta > 0$, then $L(\mathrm{ran}(A), \mathrm{ran}(B)) \leq \|R\|^2/\delta^2 \leq \|E\|^2/\delta^2$. The last inequality holds as $R = A'W = -E'W$.*

Next, we present two probabilistic bounds on matrix norms.

Proposition D.1. *Let Y be an $n \times k$ matrix with i.i.d. $N(0, 1)$ entries. For any $t > 0$,*

$$\mathbb{P}\left\{ \|\frac{1}{n} Y'Y - I_k\| \leq 2\left(\sqrt{\frac{k}{n}} + t\right) + \left(\sqrt{\frac{k}{n}} + t\right)^2 \right\} \geq 1 - 2e^{-nt^2/2}.$$

Proof. Let $W_{nk} = \frac{1}{n} Y'Y$, and $\delta_k(t) = 2(\sqrt{\frac{k}{n}} + t) + (\sqrt{\frac{k}{n}} + t)^2$. Note that $\ell_i(W_{nk}) = n^{-1/2}\sigma_i(Y)$ for $i = 1, \ldots, \max(k, n)$. Therefore, we have $\{\ell_1(W_{nk}) > 1 + \delta_k(t)\} \subset \{\sigma_1(Y) > \sqrt{n} + \sqrt{k} + \sqrt{n}t\}$, and $\{\ell_k(W_{nk}) < 1 - \delta_k(t)\} \subset \{\sigma_k(Y) < \sqrt{n} - \sqrt{k} - \sqrt{n}t\}$. Since $\|W_{nk}\| = \max\{\ell_1(W_{nk}) - 1, 1 - \ell_k(W_{nk})\}$, we obtain that

$$\mathbb{P}\{\|W_{nk}\| > \delta_k(t)\} \leq \mathbb{P}\{\ell_1(W_{nk}) > 1 + \delta_k(t)\} + \mathbb{P}\{\ell_k(W_{nk}) < 1 - \delta_k(t)\}$$
$$\leq \mathbb{P}\{\sigma_1(Y) > \sqrt{n} + \sqrt{k} + \sqrt{n}t\} + \mathbb{P}\{\sigma_k(Y) < \sqrt{n} - \sqrt{k} - \sqrt{n}t\} \leq 2e^{-nt^2/2}.$$

Here, the second inequality comes from the inclusion established in the second last display and the last inequality is the Davidson-Szarek bound [1]. This completes the proof. $\square$

Proposition D.2. *Let $Y \in \mathbb{R}^{n \times l}$ and $Z \in \mathbb{R}^{n \times m}$ be two independent matrices with i.i.d. $N(0, 1)$ entries. Then for any $0 < a < \frac{1}{2}\sqrt{n}$ and $b > 0$,*

$$\mathbb{P}\left\{ \|Y'Z\| \geq n\sqrt{1 + \frac{a}{\sqrt{n}}}\left(\sqrt{\frac{l}{n}} + \sqrt{\frac{m}{n}} + \frac{b}{\sqrt{n}}\right) \right\} \leq (l \wedge m)e^{-3a^2/16} + e^{-b^2/2}.$$

Proof. Without loss of generality, suppose $l \leq m$. Define $\bar{Y} = [\bar{y}_1, \ldots, \bar{y}_l]$, where $\bar{y}_i = y_i/\|y_i\|_2$ with y_i the i-th column of Y. Then, we obtain $\|Y'Z\| \leq (\max_{1 \leq i \leq l} \|y_i\|_2)\|\bar{Y}'Z\|$. Note that $\|y_i\|_2^2$ are i.i.d. χ_n^2 random variables. We apply Lemma D.2 to obtain that $\mathbb{P}\{\|y_i\|_2 > \sqrt{n + \sqrt{n}a}, i = 1, \ldots, l\} \leq \sum_{i=1}^l \mathbb{P}\{\|y_i\|_2 > \sqrt{n + \sqrt{n}a}\} \leq le^{-3a^2/16}$. On the other hand, $\bar{Y}'Z$ is an $l \times m$ matrix with i.i.d. $N(0, 1)$ entries. To see this, note that $\bar{y}_i'Z$ has iid $N(0, 1)$ elements since $\bar{y}_i$ has unit norm and Z has iid $N(0, 1)$ entries. In addition, the $\bar{y}_i'Z$'s are mutually independent since the $\bar{y}_i$'s are independent. So, the Davidson–Szarek bound [1, Theorem II.13] leads to $\mathbb{P}\{\|\bar{Y}'Z\| > \sqrt{l} + \sqrt{m} + b\} \leq e^{-b^2/2}$. Note that

$$\mathbb{P}\left\{ \|Y'Z\| \geq n\sqrt{1 + \frac{a}{\sqrt{n}}}\left(\sqrt{\frac{l}{n}} + \sqrt{\frac{m}{n}} + \frac{b}{\sqrt{n}}\right) \right\} \leq \mathbb{P}\{\|y_i\|_2 > \sqrt{n + \sqrt{n}a}, i = 1, \ldots, l\}$$
$$+ \mathbb{P}\{\|\bar{Y}'Z\| > \sqrt{l} + \sqrt{m} + b\}.$$

We complete the proof by plugging in the bounds for the two terms on the right side. $\square$

Finally, we list two probabilistic bounds which are used multiple times.

Lemma D.1. *Let $x \sim N(0,1)$. For any $t \geq 0$, $\mathbb{P}\{|x| > t\} \leq \sqrt{2/\pi}\, t^{-1} e^{-t^2/2}$.*

Lemma D.2 ([5]). *Let χ_n^2 denote a Chi-square random variable with n degrees of freedom. Then*

$$
\begin{aligned}
\mathbb{P}\{\chi_n^2 > n(1-\epsilon)\} &\leq e^{-n\epsilon^2/4}, && \text{when } 0 < \epsilon < 1, \\
\mathbb{P}\{\chi_n^2 > n(1+\epsilon)\} &\leq e^{-3n\epsilon^2/16}, && \text{when } 0 < \epsilon < \tfrac{1}{2}, \\
\mathbb{P}\{\chi_n^2 > n(1+\epsilon)\} &\leq \frac{\sqrt{2}}{\epsilon\sqrt{n}} e^{-n\epsilon^2/4}, && \text{when } 0 < \epsilon < n^{1/16}, n \geq 16.
\end{aligned}
$$

References

[1] K.R. Davidson and S. Szarek. *Handbook on the Geometry of Banach Spaces*, volume 1, chapter Local operator theory, random matrices and Banach spaces, pages 317–366. Elsevier Science, 2001.

[2] C. Davis and W.M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.*, 7(1):1–46, 1970.

[3] G.H. Golub and C.F. van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore and London, 3rd edition, 1996.

[4] I.M. Johnstone. On the distribution of the largest eigenvalue in principal component analysis. *Ann. Stat.*, 29:295–327, 2001.

[5] I.M. Johnstone. Thresholding for weighted chi-squared. *Stat. Sinica*, 11:691–704, 2001.

[6] I.M. Johnstone and A.Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *J. Am. Stat. Assoc.*, 104(486):682–693, 2009.

[7] Z. Ma. Sparse principal component analysis and iterative thresholding.

[8] G.W. Stewart and J.-G. Sun. *Matrix Perturbation Theory*. Computer science and scientific computing. Academic Press, 1990.

[9] P.-A. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT*, 12:99–111, 1972.