

Quantifying the Estimation Error of Principal Component Vectors

Raphael Hauser ^{*}; Jüri Lember [†]; Heinrich Matzinger [‡]; Raul Kangro [§]

April 29, 2019

Abstract

Principal component analysis is an important pattern recognition and dimensionality reduction tool in many applications. Principal components are computed as eigenvectors of a maximum likelihood covariance $\widehat{\Sigma}$ that approximates a population covariance Σ , and these eigenvectors are often used to extract structural information about the variables (or attributes) of the studied population. Since PCA is based on the eigendecomposition of the proxy covariance $\widehat{\Sigma}$ rather than the ground-truth Σ , it is important to understand the approximation error in each individual eigenvector as a function of the number of available samples. The combination of recent results of Koltchinskii & Lounici [8] and Yu, Wang & Samworth [11] yields such bounds. In the present paper we sharpen these bounds and show that eigenvectors can often be reconstructed to a required accuracy from a sample of strictly smaller size order.

^{*}Oxford Mathematical Institute and Alan Turing Institute, London U.K.; hauser@maths.ox.ac.uk

[†]Institute of Mathematics and Statistics, University of Tartu, Tartu, Estonia

[‡]School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332, USA. matzi@math.gatech.edu

[§]Institute of Mathematics and Statistics, University of Tartu, Tartu, Estonia

1 Introduction

Consider a random row vector $\vec{X} = [X_1, X_2, \dots, X_p]$, defined over a probability space (Ω, \mathbb{P}) and representing a data sample of p different items of interest, such as the returns of p different financial assets over a given investment period, the relative frequencies of p different words in a randomly chosen text, or the expression rates of p different genes in a cell line exposed to a randomly chosen chemical compound. In many applications the data are approximately Gaussian with some unknown ground-truth covariance matrix $\Sigma = \text{Cov}(\vec{X})$. The subspaces spanned by the eigenvectors corresponding to the k largest eigenvalues of Σ serve as best representation of the data in a k -dimensional space and represent an important dimensionality reduction technique. Moreover, the eigenvectors themselves are used to reveal underlying structure hidden in the data, such as subsets of genes that tend to be jointly expressed, clusters of texts that belong to a same category, risk factors that drive a financial market and many other quantities of interest in numerous applications. Often, relevant hidden structure is revealed by the eigenvectors corresponding to the second to k -th largest eigenvalues of Σ , but not by the leading eigenvector. For example, the relative size of components in the leading eigenvector of the covariance matrix of monthly returns of S&P500 stocks are approximating the relative weights of the market capitalizations of these companies, hence this vector reproduces the market index. The components of the next few leading eigenvectors reveal nearly market neutral investment portfolios, that is, portfolios that are relatively unaffected by market shocks. Other examples occur in meteorological data, where the second to k -th leading principal components are known to be useful in the automatic detection of nascent storm systems, and in the natural language processing context, where these vectors can be used to cluster texts by topic. In genome expression data, *all* eigenvalues of Σ and associated eigenvectors (respectively, the singular spectrum and associated singular vectors) can be metabolistically relevant, see e.g. Alter et. al. [3, 4] and

Bertagnoli et. al. [5]. For this reason, we are interested in approximating the first few leading eigenvectors of Σ , not just the first.

1.1 Problem Setting

In practice, a population covariance Σ is of course rarely available and in all likelihood only exists in the mind of the human modeller. True data are not exactly Gaussian either, nor can they be expected to have covariances that are invariant over time. In these circumstances it is customary to collect a limited number of sample data vectors $\vec{X}^{(1)}, \dots, \vec{X}^{(n)}$, assumed to be i.i.d. and taken from the same underlying distribution. We will assume that the data vectors have been centralized, which is carried out by subtracting their mean and considering the new data vectors $\vec{Z}^{(i)} = \vec{X}^{(i)} - \frac{1}{n} \sum_i \vec{X}^{(i)}$, so that the expectation of $\vec{Z}^{(i)}$ may be assumed to be negligible for n large enough. For applications in which the signal to noise ratio is low, such as in daily returns of financial assets, this centralisation step is not strictly required. The data vectors may thus be assumed to have zero mean. For notational simplicity we will use the maximum likelihood covariance throughout this paper instead of the unbiased sample covariance estimator, since the two matrices only differ by the factor $n/(n-1)$ and we assume n to be typically of order $O(10^2)$ or larger. A spectral decomposition of the maximum likelihood covariance

$$\hat{\Sigma} := \frac{1}{n} \sum_{i=1}^n [\vec{X}^{(i)}]^\top \vec{X}^{(i)}$$

is thus used as a proxy for Σ , and the leading part of the spectral decomposition of $\hat{\Sigma}$ is computed. This approach is called *principal component analysis* (PCA). In a situation where Σ changes over time, such as in finance, one can only avail a limited number of data points that can meaningfully be considered to have been sampled from the same underlying distribution, and since the PCA is computed from the maximum likelihood covariance $\hat{\Sigma}$, it is important to understand to what accuracy the computed principal components approximate the eigenvectors of the ground-truth

Σ . Thus, we need to study how many sample points suffice to approximate the eigenvectors of Σ by the eigenvectors of $\widehat{\Sigma}$ to a given accuracy. This is an interesting question, since we shall see that not all eigenvectors require the same number of samples, and that the required sample size depends on the distribution of eigenvalues of the population covariance Σ .

1.2 Notation and Simplifying Assumptions

We will henceforth assume that $\vec{X}^{(i)}$, ($i = 1, \dots, n$) are i.i.d. copies of $N(0, \Sigma)$ random row vectors defined on some probability space Ω . Combining these vectors into a data matrix X by row stacking, i.e., $\vec{X}^{(i)}$ is taken as the i -th row of X , the maximum likelihood covariance is given by $\widehat{\Sigma} = \frac{1}{n} X^T X$ and determines the estimation error $E := \widehat{\Sigma} - \Sigma$. Let $\Sigma = Q \text{Diag}(\lambda_1, \dots, \lambda_p) Q^T$ be a spectral decomposition of the population covariance in which the eigenvalues appear in non-increasing order $\lambda_1 \geq \dots \geq \lambda_p \geq 0$. For ease of exposition, we will now change coordinates and express all vectors with respect to the basis given by the columns of the orthogonal factor Q , that is, the eigenvectors of Σ . Instead of working with the data vectors $\vec{X}^{(i)}$, we thus work with $\vec{Y}^{(i)} = \vec{X}^{(i)} Q$ and may assume that $\Sigma = \text{Diag}(\lambda_1, \dots, \lambda_p)$ and

$$\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \left[\vec{Y}^{(i)} \right]^T \vec{Y}^{(i)}.$$

This transformation entails neither a change of the operator norm of the estimation error E , nor of the eigenvalues of Σ or $\widehat{\Sigma}$, nor of any inner products between vectors, as Q is orthogonal, but the new coordinate system renders the analysis notationally simpler, since all off-diagonal coefficients of Σ are zero and its eigenvectors are the canonical unit vectors $\vec{\mu}_i$, ($i = 1, \dots, p$).

Since $\vec{X}^{(i)}$ are random variables, the same is true for $\vec{Y}^{(i)}$, $\widehat{\Sigma}$ and the associated eigenvalues $\widehat{\lambda}_i$ and eigenvectors $\vec{\eta}_i$. Where it is desirable to expressly remind the reader of the stochasticity of these objects, we write them as functions of a generic element ω of the probability space Ω . For example, we will write $\vec{\eta}_i(\omega)$ to indicate that $\vec{\eta}_i$ is a random vector.

1.3 Prior Art

In a recent paper that builds on deep results from the theory of random matrices, Koltchinskii & Lounici proved the existence of a universal constant C_1 , independent of p and n , such that the operator norm of $E = \widehat{\Sigma} - \Sigma$ is probabilistically bounded as follows,

$$\mathbb{P} \left[\|E\| \leq C_1 \|\Sigma\| \max \left(\sqrt{\frac{r(\Sigma)}{n}}, \frac{r(\Sigma)}{n}, \sqrt{\frac{t}{n}}, \frac{t}{n} \right) \right] \geq 1 - e^{-t}, \quad t \geq 1 \quad (1.1)$$

where

$$r(\Sigma) = \frac{\sum_{j=1}^p \lambda_j}{\max_{j=1, \dots, p} \lambda_j} = \frac{\sum_{j=1}^p \lambda_j}{\lambda_1}$$

is the *effective rank* of Σ ; see Corollary 2 in Koltchinskii & Lounici [8]. Note that $r(\Sigma) \leq p$, so that for $n \geq \max(p, t)$ (1.1) becomes

$$\mathbb{P} \left[\|E\| \leq C_1 \|\Sigma\| \max \left(\sqrt{\frac{r(\Sigma)}{n}}, \sqrt{\frac{t}{n}} \right) \right] \geq 1 - e^{-t},$$

and using $\|\Sigma\| = \lambda_1$ and the definition of $r(\Sigma)$, this yields the bound

$$\mathbb{P} \left[\|E\| \leq C_1 \frac{\sqrt{\lambda_1}}{\sqrt{n}} \times \max \left(\sqrt{\sum_{j=1}^p \lambda_j}, \sqrt{\lambda_1 t} \right) \right] \geq 1 - e^{-t}. \quad (1.2)$$

This was an important advance, as previous bounds on $\|E\|$ had only been known for the case where all eigenvalues are of the same order or when p remains bounded as n tends to infinity, the so-called *finite-dimensional case*, while the cited result of Koltchinskii & Lounici [8] applies to the *infinite-dimensional case*, where p tends to infinity as well.

In this paper we adopt the infinite-dimensional setting and study the error that is incurred in approximating individual eigenvectors of the population covariance Σ by computing a spectral decomposition of the maximum likelihood covariance matrix $\widehat{\Sigma}$. Ledoit & P  ch   [9] also studied the convergence of the sample covariance eigenvectors but their results are asymptotic and restricted to the finite-dimensional case, and their motivation is to derive optimal bias-correction terms. The study of the convergence of eigenvalues (as opposed to eigenvectors) is a mature field, see e.g. Bai & Yao [6].

1.4 Contributions and Main Result

Let $\widehat{\lambda}_i$ be the i -th largest eigenvalue of $\widehat{\Sigma}$ and $\vec{\eta}_i$ an associated eigenvector. For convenience, we assume throughout that the corresponding eigenvalue λ_i of the population covariance is non-repeated. Recall that under the notation introduced above $\vec{\mu}_i$ is a unit eigenvector of the population covariance associated with eigenvalue λ_i . In contrast to $\vec{\eta}_i(\omega)$ and $\widehat{\lambda}_i(\omega)$, $\vec{\mu}_i$ and λ_i are deterministic objects. It follows from Corollary 1 of Yu, Wang & Samworth [11] that the (random) principal angle

$$\theta_i(\omega) := \angle(\vec{\mu}_i, \vec{\eta}_i(\omega)) = \cos^{-1} \left(\frac{|\langle \vec{\mu}_i, \vec{\eta}_i(\omega) \rangle|}{\|\vec{\eta}_i(\omega)\|} \right)$$

is bounded by

$$\sin \theta_i(\omega) \leq \frac{2\|E(\omega)\|}{\text{gap}_i} \quad (\omega \in \Omega), \quad (1.3)$$

where $\|E(\omega)\|$ is the operator norm of $E = \Sigma - \widehat{\Sigma}(\omega)$ and $\text{gap}_i := \min_{j \neq i} |\lambda_i - \lambda_j|$ is the spectral gap at eigenvalue λ_i of Σ . Note that in contrast to the classical Davis-Kahan $\sin \theta$ -Theorem [7] the denominator of (1.3) does not depend on the (stochastic) spectrum of $\widehat{\Sigma}$ but only on the (deterministic) spectrum of Σ .

Combining the Corollary 1 of Yu, Wang & Samworth [11] with Corollary 2 of Koltchinskii & Lounici [8], i.e. (1.3) with (1.2), we find that for fixed $q \in (0, 1)$, $t \geq 1$, $p \in \mathbb{N}$ and

$$n \geq \max \left(p, t, \frac{4C_1^2 \lambda_1^2 t}{q^2 \text{gap}_i^2}, \frac{4C_1^2 \lambda_1 \sum_{j=1}^p \lambda_j}{q^2 \text{gap}_i^2} \right), \quad (1.4)$$

the following probabilistic bound applies,

$$\mathbb{P} [\sin \theta_i < q] \geq 1 - e^{-t}.$$

The aim of the present paper is to show that the bound (1.4) on the required sample size can be replaced by

$$n \geq \max \left(p, t, M \times 16C_1^2 \max \left(\left(1 + \frac{\lambda_i}{\text{gap}_i}\right), \frac{2\lambda_i}{q^2 \text{gap}_i} \right) \max \left(\left(1 + \frac{\lambda_i}{\text{gap}_i}\right)t, \frac{3}{2} \sum_{j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_i|} \right) \right), \quad (1.5)$$

where $M > 0$ is a constant that depends on the distribution of the population spectrum (λ_j) . For some distributions of the eigenvalues of Σ , this bound is an improvement of several orders of magnitude over previous bounds, see Examples 1.1 and 1.2 below.

In fact, we will prove a marginally weaker result: For the purposes of our analysis we will work with the index

$$i^*(\omega) = \arg \min_{k \in \{1, \dots, p\}} |\lambda_k - \widehat{\lambda}_i|,$$

which is unique with probability 1 and identifies the population eigenvalue closest to the i -th largest maximum likelihood eigenvalue. Note that this is a random index, and that if λ_i is non-repeated, then $i^*(\omega) = i$ with arbitrarily high probability for large n . We will further comment on this below. By giving up on the requirement that η_i be a unit vector, with probability 1 we are able to choose the unit eigenvector $\vec{\mu}_{i^*}$ such that $\langle \vec{\mu}_{i^*}, \vec{\eta}_i \rangle > 0$, so that $\vec{\eta}_i \in \vec{\mu}_{i^*} + (\vec{\mu}_{i^*})^\perp$, that is, $\Delta \vec{\mu}_{i^*, i} := \vec{\eta}_i - \vec{\mu}_{i^*}$ is orthogonal to $\vec{\mu}_{i^*}$. Then we have

$$\sin \theta_{i^*, i}(\omega) \leq \tan \theta_{i^*, i}(\omega) = \|\Delta \vec{\mu}_{i^*, i}\|_2, \quad (1.6)$$

where $\theta_{i^*, i}$ is defined as

$$\theta_{i^*, i}(\omega) := \angle(\vec{\mu}_{i^*}, \vec{\eta}_i(\omega)) = \cos^{-1} \left(\frac{|\langle \vec{\mu}_{i^*}, \vec{\eta}_i(\omega) \rangle|}{\|\vec{\eta}_i(\omega)\|} \right).$$

Therefore, it suffices to prove that

$$\mathbb{P} \left[\|\Delta \vec{\mu}_{i^*, i}\|_2 < q \mid n \geq \max(p, t, \Psi(q, t, i^*)) \right] \geq 1 - \frac{p e^{-t}}{\mathbb{P}[n \geq \Psi(q, t, i^*)]} \quad (1.7)$$

to establish the following result, which is our main theorem:

Theorem 1.1 *For $q \in (0, 1)$, $t \geq 1$, $p \in \mathbb{N}$, $i \in \{1, \dots, p\}$ and $n \geq \max(t, p, \Psi(q, t, i))$ we have*

$$\mathbb{P} \left[\sin \theta_{i^*, i} \leq q \mid n \geq \Psi(q, t, i^*) \right] \geq 1 - \frac{p e^{-t}}{\mathbb{P}[n \geq \Psi(q, t, i^*)]}, \quad (1.8)$$

where

$$\Psi(q, t, i) := 16C_1^2 \max \left(\left(1 + \frac{\lambda_i}{\text{gap}_i}\right), 2 \frac{\lambda_i}{q^2 \text{gap}_i} \right) \max \left(\left(1 + \frac{\lambda_i}{\text{gap}_i}\right) t, \frac{3}{2} \sum_{j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_i|} \right). \quad (1.9)$$

1.5 Discussion and Practical Significance of Results

The auxiliary parameter t allows users of Theorem 1.1 to control the probability with which the alignment of eigenvectors $\vec{\mu}_{i^*}$ and $\vec{\eta}_i$ exceeds the threshold q : choosing t large renders the probability of this outcome arbitrarily close to 1 except an exponentially small term in t at the expense of requiring a number of samples that scales linearly in t . Note that for nonrepeated λ_i , we have $i^* = i$ for n large enough with high probability, and then $\theta_i = \theta_{i^*,i}$. In the relevant contexts to which Theorem 1.1 applies, one can usually quantify

$$\mathbb{P}[n \geq \Psi(q, t, i^*)] \approx 1$$

as being close to 1 for $n \geq M \times \Psi(q, t, i)$ quite easily, where the constant $M > 0$ depends on the distribution of the spectrum (λ_j) and is independent of p as long as λ_i for fixed i remains bounded. Substituting such a quantitative bound into (1.8), this yields a lower bound on

$$\begin{aligned} \mathbb{P}[\sin \theta_i \leq q] &\geq \mathbb{P}\left[\sin \theta_i \leq q \mid n \geq \Psi(q, t, i^*)\right] \times \mathbb{P}[n \geq \Psi(q, t, i^*)] \\ &\stackrel{(1.8)}{\geq} \mathbb{P}[n \geq \Psi(q, t, i^*)] - p e^{-t} \end{aligned}$$

for all n that satisfy the bound (1.5). Since a comprehensive treatment of all cases is not a reasonable undertaking, we limit the scope of this paper to the derivation of the bound (1.8) and leave the quantification of the constant M to users of Theorem 1.1, but the examples below illustrate the practical relevance even for small values of M .

For eigenvalues λ_i of lower order, the new bound (1.5) is of strictly smaller order than (1.4), because the numerator satisfies $\lambda_i \ll \lambda_1$, and furthermore, the eigenvalues of largest order in the sum $\sum_{j=1}^p \lambda_j$ are replaced by terms of $O(1)$ in the sum $\sum_{j \neq i} \lambda_j / |\lambda_j - \lambda_i|$, so that this sum is generally much smaller when the eigenvalue λ_i is of medium (but not smallest) order. Note that, like the $\sin \theta$ -Theorem of Yu, Wang & Samworth [11], our bound (1.5) only depends on the spectrum of the population covariance, but not on the maximum likelihood covariance.

Example 1.1 *In data science one often deals with data matrices whose scree plots have a short and quickly decreasing initial section, followed by a large slowly decreasing middle section with close to equally spaced eigenvalues, and if the population covariance is a noisy version of a rank $k < p$ matrix, then this is followed by a sharp drop off to background noise after position k . For an example, see e.g. Alter et. al. [10] or [2]*

A reasonable model is therefore to assume that for some fixed $\beta \in (0, 1)$, the population covariance matrix Σ has $O(1)$ eigenvalues of order $O(p)$ that are structurally not important, $O(p^{1-\beta})$ eigenvalues of order $O(p^\beta)$, and $O(p)$ smaller eigenvalues of order $O(1)$, where Σ has been rescaled so that $\sum_{j=1}^p \lambda_j = p$. The latter assumption is reasonable, as PCA is often carried out on the correlation matrix rather than the covariances, to correct for differences in scale between the feature variables.

The relevant structural information is often revealed by eigenvectors corresponding to eigenvalues of the intermediate order $O(p^\beta)$ rather than those of the largest order $O(p)$. For example, when PCA is used to classify texts into different topics, the principal components that correspond to the largest order of eigenvalues typically pick up high frequency words that are common in all texts of the same language, whereas the topic is typically identified by context-specific words that merely appear with higher frequency than in the general text population.

Assuming that the eigenvalues of order $O(p^\beta)$ are locally not very different from a renewal process, or similarly, that they are approximately located on a lattice, their spectral gaps are typically of order $O(p^\beta/p^{1-\beta}) = O(p^{2\beta-1})$, because there are $O(p^{1-\beta})$ of these gaps of equal

length spanning an interval of length $O(p^\beta)$. In the case where λ_i is of order $O(p^\beta)$, this yields

$$\sum_{\lambda_j=O(p)} \frac{\lambda_j}{|\lambda_j - \lambda_i|} = O(1), \quad (1.10)$$

$$\sum_{\lambda_j=O(p^\beta), j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_i|} = O\left(\frac{p^\beta}{p^{2\beta-1}} \times \log p\right) = O\left(\log p \times p^{1-\beta}\right), \quad (1.11)$$

$$\sum_{\lambda_j=O(1)} \frac{\lambda_j}{|\lambda_j - \lambda_i|} = O\left(p \times p^{-\beta}\right) = O\left(p^{1-\beta}\right), \quad (1.12)$$

where (1.10) holds because by assumption there are $O(1)$ summands, all of which are of order $O(1)$, (1.11) holds because the left hand side is a harmonic series with $O(p^{1-\beta})$ terms under the assumption that the eigenvalues of order $O(p^\beta)$ lie on a lattice, i.e.,

$$\sum_{\lambda_j=O(p^\beta), j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_i|} \sim \frac{p^\beta}{p^{2\beta-1}} \sum_{\ell=1}^{p^{1-\beta}} \frac{1}{\ell} \sim p^{1-\beta} \int_1^{p^{1-\beta}} \frac{dx}{x} \sim p^{1-\beta} (1 - \beta) \log p,$$

and (1.12) holds since there are $O(p)$ summands, all of which are of order $p^{-\beta}$. Summing up the bounds (1.10)–(1.12), we obtain

$$\sum_{j=1, j \neq i}^p \frac{\lambda_j}{|\lambda_j - \lambda_i|} = O\left(p^{1-\beta} \times \log p\right).$$

Substituting this bound into (1.5), we find that for $\beta > \frac{1}{2}$ and when λ_i corresponds to one of the eigenvalues of intermediate order $O(p^\beta)$, the number of samples required to guarantee that $\sin \theta_i < q$ with probability larger than $1 - \varepsilon$ is given by

$$n = O\left(\max(-\log \varepsilon + \log p, p, p^{1-\beta} \times p^{1-\beta} \times \log p)\right) = O\left(\max(\log p, p, p^{2(1-\beta)} \log p)\right) = O(p),$$

while the bound (1.4) would imply that the number of samples required is of the order

$$n = O\left(\max\left(-\log p, p, \frac{\log p \times p^2}{p^{4\beta-2}} \times \frac{p^2}{p^{4\beta-2}}\right)\right) = O\left(p^{4(1-\beta)} \log p\right)$$

which is of strictly larger order for $\beta < \frac{3}{4}$. Thus, for $\frac{1}{2} < \beta < \frac{3}{4}$ we can reconstruct the relevant principal components with linearly many samples, whereas the bound (1.4) suggests that the number of samples required would have to be superlinear in the number p of feature variables.

Example 1.2 *Due to the dependence on the random index i^* , our formula 1.8 yields a probabilistic bound on the sample complexity of PCA. We note however that in practical problems the concentration of measure phenomenon causes fluctuations of i^* to be of smaller order than the expectation, that is, $\sqrt{\text{Var}[i^*]}$ is typically much smaller than $E[i^*]$, and that the map $i \mapsto \text{gap}_i$ is very regular, that is, gap_i remains practically constant for small changes of i . Since i^* is the only source of randomness in the bound 1.8, the latter is practically deterministic under these circumstances, because its order of fluctuations is strictly smaller than the order of its expectation. Under these circumstances, the bound can be interpreted as $\text{P}[\sin \theta_i \leq q] \geq 1 - pe^{-t}$ for $n \geq \max(p, t, \Psi(q, t, i))$, which is the same as (1.5).*

To illustrate the practical relevance of this bound, we used 2000 daily returns of 800 stocks to estimate a realistic covariance matrix, and we then used this covariance as population covariance with which to simulate i.i.d. multivariate Gaussian vectors. Taking $q = 0.5$ and $C_1, M = 1$, two bounds on the number of data points required to reduce $\sin \theta_i < q$ with high probability were computed: the bound

$$n_{\text{new}} := \frac{\lambda_i}{\text{gap}_i} \times \sum_{j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_i|}, \quad (1.13)$$

is practical version of (1.5) stripped of constants, and analogously

$$n_{\text{old}} := \frac{\lambda_1}{(\text{gap}_i)^2} \times \sum_{j \neq i} \lambda_j, \quad (1.14)$$

which is a stripped down version of (1.4). Table 1.2 confirms that the new bound (1.13) derived in this paper is of the correct order, while (1.14) is up to 6 orders of magnitude too large. In reading Table 1.2, the reader ought to keep in mind that in high dimensions, any two random vectors are almost orthogonal to one another. For example in dimension $p = 800$, the probability that two uniformly sampled unit vectors \vec{v} and \vec{w} have alignment $|\langle \vec{v}, \vec{w} \rangle| > 0.15$ is smaller than 10^{-4} . The results of the table thus confirms that $n \approx n_{\text{new}}$ is the correct threshold for the number of samples needed to achieve phenomenally good eigenvector alignment, whereas the bound n_{old}

p	i	n_{new}	n_{old}	n	$\cos \theta_i$
800	5	6.51 e 2	1.1 e 6	$n_{\text{new}}/2$	0.79
800	5	6.51 e 2	1.1 e 6	n_{new}	0.9
800	5	6.51 e 2	1.1 e 6	$2n_{\text{new}}$	0.95
800	10	2.458 e 3	2.740 e 9	$n_{\text{new}}/2$	0.56
800	10	2.458 e 3	2.740 e 9	n_{new}	0.78
800	10	2.458 e 3	2.740 e 9	$2n_{\text{new}}$	0.93
200	15	1.31 e 5	1.47 e 8	$n_{\text{new}}/2$	0.5
200	15	1.31 e 5	1.47 e 8	n_{new}	0.6
200	15	1.31 e 5	1.47 e 8	$2n_{\text{new}}$	0.71

Table 1: Alignment between the i -th estimated and ground truth eigenvectors as a function of the sample size n used to estimate the covariance matrix from i.i.d. multivariate Gaussian data with a practically relevant spectrum.

can be many orders of magnitude too large.

2 Main Ideas of the Proof

Before we proceed with a detailed proof of Theorem 1.1, it is helpful to outline of the main ideas. In what follows, we shall assume that λ_i is unrepeated, and we will write $\Delta\lambda_{k,i} = \widehat{\lambda}_i - \lambda_k$, where k is a fixed index that we shall specify later. In analogy to the earlier construction of the vector $\Delta\vec{\mu}_i$, let $\vec{\eta}_i$ be an eigenvector of $\widehat{\Sigma}$ associated with the i -th largest eigenvalue $\widehat{\lambda}_i$, and chosen such that $\langle \vec{\mu}_k, \vec{\eta}_i \rangle \geq 0$ and $\vec{\eta}_i \in \vec{\mu}_k + \vec{\mu}_k^\perp$, and let us write $\Delta\vec{\mu}_{k,i} := \vec{\eta}_i - \vec{\mu}_k$, so that $\Delta\vec{\mu}_{k,i} \in \vec{\mu}_k^\perp$. In the special case where $k = i$ we denote $\Delta\vec{\mu}_i = \Delta\vec{\mu}_{i,i}$ and $\Delta\lambda_i := \Delta\lambda_{i,i}$.

2.1 The Finite-Dimensional Case

For fixed p and $k = i$,

$$\Delta \vec{\mu}_i \approx -\sqrt{\frac{\lambda_i}{n}} Z_i, \quad (2.1)$$

where $Z_i = [Z_{1,i}, \dots, Z_{p,i}]$ is a random vector of size p with coefficients

$$Z_{j,i} = \begin{cases} \frac{\sqrt{\lambda_j}}{\lambda_j - \lambda_i} N_{j,i}, & (j \neq i), \\ 0, & (j = i), \end{cases}$$

and where the random variables $N_{j,i}$ converge (jointly, in distribution) to independent standard Gaussians as $n \rightarrow \infty$ (see, e.g. Anderson [1], Thm. 13.5.1). Thus, to guarantee that $\|\Delta \vec{\mu}_i\| < q$ with high probability, we need $\sqrt{\lambda_i/n} \times \|Z_i\| < q$ with high probability. Assuming n to be large enough for the variables $N_{i,j}$ to be close to i.i.d. standard Gaussians and the approximation (2.1) to hold, one finds

$$\|\Delta \vec{\mu}_i\|^2 \approx \frac{1}{n} \times \sum_{j \neq i} \frac{\lambda_i \lambda_j}{(\lambda_j - \lambda_i)^2} N_{j,i}^2$$

and this approximation can be used to bound $P(\|\Delta \mu_i\| < q)$. In the present case, when p is not fixed, this approach will not hold, but let us gain a quick understanding of how (2.1) arises in the finite-dimensional case in which n is left to tend to infinity for fixed p , and of how the argument has to be amended in the infinite-dimensional case, where p is allowed to tend to ∞ at a linear rate in n : We have

$$\Sigma \vec{\mu}_k = \lambda_k \vec{\mu}_k, \quad (2.2)$$

$$\widehat{\Sigma} [\vec{\mu}_k + \Delta \vec{\mu}_{k,i}] = (\lambda_k + \Delta \lambda_{k,i}) [\vec{\mu}_k + \Delta \vec{\mu}_{k,i}]. \quad (2.3)$$

Subtracting (2.2) from (2.3) and using $\widehat{\Sigma} = \Sigma + E$ yields

$$[\Sigma - (\lambda_k + \Delta \lambda_{k,i}) \mathbf{I}_p] \Delta \vec{\mu}_{k,i} + E \Delta \vec{\mu}_{k,i} = -E \vec{\mu}_k + \Delta \lambda_{k,i} \vec{\mu}_k \quad (2.4)$$

where \mathbf{I}_p is the $p \times p$ identity matrix. In the finite-dimensional case, where p is fixed and n tends to infinity, taking $k = i$ implies that E , $\Delta \vec{\mu}_i$ and $\Delta \lambda_i$ are all of order $1/\sqrt{n}$, and hence the terms

$\Delta\lambda_i\Delta\vec{\mu}_i$ and $E\Delta\vec{\mu}_i$ are of the smaller order $1/n$ and can be neglected in the asymptotics, so as to give the approximation

$$[\Sigma - \lambda_i \mathbf{I}_p] \Delta\vec{\mu}_i \approx -E\vec{\mu}_i + \Delta\lambda_i\vec{\mu}_i \quad (2.5)$$

Using the facts that $\Sigma = \text{Diag}(\lambda_j)$, $\vec{\mu}_i$ is the i -th unit vector and that $\vec{\mu}_i$ and $\Delta\vec{\mu}_i$ are mutually orthogonal by construction, the i -th equation of system (2.5) yields

$$\Delta\lambda_i \approx E_{ii},$$

and dividing the j -th equation of the system (2.5) by $\lambda_j - \lambda_i$ ($j \neq i$) yields

$$\Delta\vec{\mu}_i \approx - \begin{pmatrix} \frac{E_{1,i}}{\lambda_1 - \lambda_i} \\ \frac{E_{2,i}}{\lambda_2 - \lambda_i} \\ \vdots \\ \frac{E_{(i-1),i}}{\lambda_{i-1} - \lambda_i} \\ 0 \\ \frac{E_{(i+1),i}}{\lambda_{i+1} - \lambda_i} \\ \vdots \\ \frac{E_{p,i}}{\lambda_p - \lambda_i} \end{pmatrix} = -\sqrt{\frac{\lambda_i}{n}} \begin{pmatrix} \frac{\sqrt{\lambda_1}}{\lambda_1 - \lambda_i} N_{1,i} \\ \frac{\sqrt{\lambda_2}}{\lambda_2 - \lambda_i} N_{2,i} \\ \vdots \\ \frac{\sqrt{\lambda_{i-1}}}{\lambda_{i-1} - \lambda_i} N_{(i-1),i} \\ 0 \\ \frac{\sqrt{\lambda_{i+1}}}{\lambda_{i+1} - \lambda_i} N_{(i+1),i} \\ \vdots \\ \frac{\sqrt{\lambda_p}}{\lambda_p - \lambda_i} N_{p,i} \end{pmatrix}, \quad (2.6)$$

where

$$N_{s,t} := \sqrt{\frac{n}{\lambda_s \lambda_t}} E_{s,t} \quad (2.7)$$

for all $s, t \in 1, \dots, p$ with $s \neq t$. The random variables N_{ji} ($j \neq i$) converge in joint distribution to i.i.d. standard Gaussians, a fact we only mention for motivational purposes and on which our technical argument below does not rely.

2.2 The Infinite-Dimensional Case

In contrast, in the infinite-dimensional case the terms $\Delta\lambda_{k,i}\Delta\vec{\mu}_{k,i}$ and $E\Delta\vec{\mu}_{k,i}$ can no longer be asymptotically disregarded, as p is also allowed to tend to infinity at up to a linear rate in n .

Let P_k denote the orthogonal projection into the orthogonal complement $\vec{\mu}_k^\perp$ of $\vec{\mu}_k$, and define the operator

$$D_k := -\text{diag}\left(\frac{1}{\lambda_1 - \hat{\lambda}_i}, \dots, \frac{1}{\lambda_{k-1} - \hat{\lambda}_i}, 0, \frac{1}{\lambda_{k+1} - \hat{\lambda}_i}, \dots, \frac{1}{\lambda_p - \hat{\lambda}_i}\right),$$

so that $D_k \vec{\mu}_k = 0$ and D_k is well defined as long as $\hat{\lambda}_i$ is not an eigenvalue (except, possibly, the k -th eigenvalue) of Σ . Using the fact that $\Delta \vec{\mu}_{k,i} \in \vec{\mu}_k^\perp$, multiplying (2.4) by D_k and solving for $\Delta \mu_{k,i}$ yields

$$\begin{aligned} \Delta \vec{\mu}_{k,i} &= -(\mathbf{I}_p - D_k E P_k)^{-1} D_k \begin{pmatrix} E_{1k} & \dots & E_{(k-1)k} & 0 & E_{(k+1)k} & \dots & E_{pk} \end{pmatrix}^T \\ &= -\sqrt{\frac{\lambda_k}{n}} \times (\mathbf{I}_p - D_k E P_k)^{-1} \begin{pmatrix} \frac{\sqrt{\lambda_1}}{\lambda_1 - \hat{\lambda}_i} N_{1,k} \\ \frac{\sqrt{\lambda_2}}{\lambda_2 - \hat{\lambda}_i} N_{2,k} \\ \dots \\ \frac{\sqrt{\lambda_{k-1}}}{\lambda_{k-1} - \hat{\lambda}_i} N_{(k-1),k} \\ 0 \\ \frac{\sqrt{\lambda_{k+1}}}{\lambda_{k+1} - \hat{\lambda}_i} N_{(k+1),k} \\ \dots \\ \frac{\sqrt{\lambda_p}}{\lambda_p - \hat{\lambda}_i} N_{p,k} \end{pmatrix} \quad (2.8) \\ &= -\sqrt{\frac{\lambda_k}{n}} \times (\mathbf{I}_p - D_k E P_k)^{-1} |D_k|^{1/2} J \vec{V}_k, \quad (2.9) \end{aligned}$$

where $N_{s,t}$ is as defined in (2.7), $\vec{V}_k = (v_{j,k})$ is defined by

$$v_{j,k} := \sqrt{\frac{\lambda_j}{|\lambda_j - \hat{\lambda}_i|}} \times N_{j,k}, \quad (j = 1, \dots, k-1, k+1, \dots, p), \quad v_{k,k} := 0, \quad (2.10)$$

$|D_k|$ is the matrix obtained by replacing the coefficients of D_k by their absolute values, and J the diagonal matrix with coefficients $J_{j,j} := \text{sign}(D_k)_{j,j}$.

Example 2.1 Let $\vec{X} = (X_1, X_2, X_3)$ be a 3-dimensional random vector with zero mean and covariance matrix $\Sigma = \text{Diag}(\lambda_1, \lambda_2, \lambda_3)$, where $\lambda_1 > \lambda_2 > \lambda_3$. Using the notation introduced

earlier, and taking n large enough to guarantee that $1^* = 1$ with high probability, the first line of (2.4) for $k = i = 1$ yields with $\vec{\eta}_1 = (1, \eta_{21}, \eta_{31})$

$$\hat{\lambda}_1 - \lambda_1 = \Delta\lambda_1 = E_{11} + E_{12}\eta_{21} + E_{13}\eta_{31} = \frac{\lambda_1}{\sqrt{n}}N_{11} + \frac{\sqrt{\lambda_1}}{\sqrt{n}}(\sqrt{\lambda_2}N_{12}\eta_{21} + \sqrt{\lambda_3}N_{13}\eta_{31}),$$

whereas the second and third lines yield

$$\left[\begin{pmatrix} \lambda_2 - \hat{\lambda}_1 & 0 \\ 0 & \lambda_3 - \hat{\lambda}_1 \end{pmatrix} + \begin{pmatrix} E_{22} & E_{23} \\ E_{32} & E_{33} \end{pmatrix} \right] \begin{pmatrix} \eta_{21} \\ \eta_{31} \end{pmatrix} = - \begin{pmatrix} E_{21} \\ E_{31} \end{pmatrix}.$$

Equivalently

$$\left[\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \begin{pmatrix} \frac{E_{22}}{\lambda_2 - \hat{\lambda}_1} & \frac{E_{23}}{\lambda_2 - \hat{\lambda}_1} \\ \frac{E_{32}}{\lambda_3 - \hat{\lambda}_1} & \frac{E_{33}}{\lambda_3 - \hat{\lambda}_1} \end{pmatrix} \right] \begin{pmatrix} \eta_{21} \\ \eta_{31} \end{pmatrix} = - \begin{pmatrix} \frac{E_{21}}{\lambda_2 - \hat{\lambda}_1} \\ \frac{E_{31}}{\lambda_3 - \hat{\lambda}_1} \end{pmatrix}.$$

Therefore, with

$$D_1 = \text{diag}\left(0, \frac{-1}{\lambda_2 - \hat{\lambda}_1}, \frac{-1}{\lambda_3 - \hat{\lambda}_1}\right), \quad P_1 = \text{diag}(0, 1, 1)$$

we obtain

$$\begin{aligned} \Delta\vec{\mu}_{1,1} &= \begin{pmatrix} 0 \\ \eta_{21} \\ \eta_{31} \end{pmatrix} = \left[\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{E_{22}}{\lambda_2 - \hat{\lambda}_1} & \frac{E_{23}}{\lambda_2 - \hat{\lambda}_1} \\ 0 & \frac{E_{32}}{\lambda_3 - \hat{\lambda}_1} & \frac{E_{33}}{\lambda_3 - \hat{\lambda}_1} \end{pmatrix} \right]^{-1} \begin{pmatrix} 0 \\ \frac{-E_{21}}{\lambda_2 - \hat{\lambda}_1} \\ \frac{-E_{31}}{\lambda_3 - \hat{\lambda}_1} \end{pmatrix} \\ &= [I_3 - D_1 E P_1]^{-1} D_1 \begin{pmatrix} 0 \\ E_{21} \\ E_{31} \end{pmatrix}. \end{aligned}$$

Comparing (2.8) with formula (2.6) from the finite-dimensional case, we note the following differences: Firstly (2.8) is an exact formula, whilst (2.6) is an approximation in the case where $k = i$. Secondly, the appearance of the term $\Delta\lambda_{k,i}$ in the denominators on the r.h.s. of (2.8) is problematic, unless we can bound $|\lambda_j - \hat{\lambda}_i|$ away from zero. We will fix this choice by setting $k = i^*$ and conditioning on i^* , where the random index i^* is defined in Theorem 1.1. As we will show in Section 3, this yields the following bounds:

Lemma 2.1 For all $j \neq i^*$, the following hold true,

$$\frac{1}{|\lambda_j - \hat{\lambda}_i|} \leq \frac{2}{|\lambda_j - \lambda_{i^*}|}, \quad (2.11)$$

$$\frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|} \leq 1 + \frac{\lambda_{i^*}}{\text{gap}_{i^*}}. \quad (2.12)$$

Thirdly, and most significantly, the term $D_{i^*} E P_{i^*}$ appears in the r.h.s. of (2.8). If we managed to prove that $\|D_{i^*} E P_{i^*}\| < 1$, then by the Neuman Series Formula,

$$(I_p - D_{i^*} E P_{i^*})^{-1} = I_p + D_{i^*} E P_{i^*} + (D_{i^*} E P_{i^*})^2 + (D_{i^*} E P_{i^*})^3 + \dots, \quad (2.13)$$

and then we could argue along the lines of the finite-dimensional case. However, instead of bounding $\|D_{i^*} E P_{i^*}\|$, we will bound $\|\Lambda_{i^*}\|$, defined as

$$\Lambda_{i^*} := |D_{i^*}|^{1/2} E |D_{i^*}|^{1/2}, \quad (2.14)$$

where $|D_{i^*}|$ denotes the matrix obtained by replacing the coefficients of D_{i^*} by their absolute values. By using the properties

$$D_{i^*} = |D_{i^*}|^{1/2} J |D_{i^*}|^{1/2}, \quad P_{i^*} D_{i^*} = D_{i^*}, \quad P_{i^*} = |D_{i^*}|^{1/2} |D_{i^*}|^{-1/2},$$

where $|D_{i^*}|^{-1/2}$ is the diagonal matrix containing values $\sqrt{|\lambda_j - \hat{\lambda}_i|}$ at positions $j \neq i^*$, and 0 at the position i^* , we easily get

$$(D_{i^*} E P_{i^*})^k = |D_{i^*}|^{1/2} (J \Lambda_{i^*})^k |D_{i^*}|^{-1/2}, \quad k \geq 1.$$

The Neumann series (2.13) may thus be rewritten as

$$(I_p - D_{i^*} E P_{i^*})^{-1} = I_p + |D_{i^*}|^{1/2} \left(\sum_{k=1}^{\infty} (J \Lambda_{i^*})^k \right) |D_{i^*}|^{-1/2} \quad (2.15)$$

implying that

$$(I_p - D_{i^*} E P_{i^*})^{-1} |D_{i^*}|^{1/2} = |D_{i^*}|^{1/2} \left(I_p + \sum_{k=1}^{\infty} (J \Lambda_{i^*})^k \right).$$

If $\|\Lambda_{i^*}\| \in (0, 1)$, then this series converges, and using (2.11) in

$$\| |D_{i^*}|^{1/2} \| = \max_{j \neq i^*} \frac{1}{\sqrt{|\lambda_j - \hat{\lambda}_i|}} \leq \sqrt{\frac{2}{\text{gap}_{i^*}}},$$

the taking of norms on both sides of (2.9) yields

$$\|\Delta \vec{\mu}_{i^*, i}\| \leq \sqrt{\frac{2}{\text{gap}_{i^*}}} \times \frac{1}{1 - \|\Lambda_{i^*}\|} \times \sqrt{\frac{\lambda_{i^*}}{n}} \|\vec{V}_{i^*}\|. \quad (2.16)$$

A crucial observation is that for i^* fixed, the norm of Λ_{i^*} can be bounded by the norm of a matrix, which has an interpretation as covariance matrix estimation error for a multivariate Gaussian random vector with zero mean, independent coefficients, and whose j -th coefficient has variance

$$\frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|}, \quad (j \neq i^*). \quad (2.17)$$

Its operator norm can therefore be bounded using the technique of Koltchinskii & Lounici that allowed for the derivation of (1.2). To exploit this mechanism, we will have to condition on the value of i^* , but for intuitive purposes, the reader may keep in mind that for n that satisfy (1.5), we have $i^* = i$ with high probability.

3 Further Technical Details

We will now fill in the missing details of the proof we outlined so far.

3.1 Preliminary Results

We start with two technical lemmas.

Proof. (Lemma 2.1) By definition of i^* , for all $j \neq i^*$ we have

$$|\lambda_j - \lambda_{i^*}| \leq |\lambda_j - \hat{\lambda}_i| + |\hat{\lambda}_i - \lambda_{i^*}| \leq 2|\lambda_j - \hat{\lambda}_i|$$

which shows (2.11). As

$$\frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|} = \frac{\lambda_j - \lambda_{i^*}}{|\lambda_j - \lambda_{i^*}|} + \frac{\lambda_{i^*}}{|\lambda_j - \lambda_{i^*}|} \leq 1 + \frac{\lambda_{i^*}}{|\lambda_j - \lambda_{i^*}|},$$

we have also proved the inequality (2.12) ■

Lemma 3.1 *Let $\text{Diag}(\nu)$ be a diagonal matrix of size p with diagonal coefficients $\nu_j \geq 0$, ($j = 1, \dots, p$), and let $E(\nu) = \text{Diag}(\nu)E\text{Diag}(\nu)$, where $E = \widehat{\Sigma} - \Sigma$ is the covariance matrix error used throughout this paper. Then for all $t \geq 1$ and $n \in \mathbb{N}$, $n \geq \max(p, t)$ it is true that*

$$\mathbb{P} [\|E(\nu)\|^2 \leq \Xi(n, t, \nu)] \geq 1 - e^{-t},$$

where

$$\Xi(n, t, \nu) = \frac{C_1^2 \|\nu^2 \lambda\|_\infty}{n} \times \max \left(\|\nu^2 \lambda\|_\infty \times t, \sum_{j=1}^p \nu_j^2 \lambda_j \right)$$

with

$$\|\nu^2 \lambda\|_\infty = \max_{j=1,2,\dots,p} \nu_j^2 \lambda_j.$$

Proof. Recall that

$$E = \frac{1}{n} \sum_{i=1}^n \left[\vec{Y}^{(i)} \right]^T \vec{Y}^{(i)} - \Sigma,$$

where $\vec{Y}^{(i)}$ are i.i.d. multivariate normal row vectors with zero mean and covariance matrix $\Sigma = \text{Diag}(\lambda)$. Therefore,

$$E(\nu) = \frac{1}{n} \sum_{i=1}^n \left[\vec{Z}^{(i)} \right]^T \vec{Z}^{(i)} - \Sigma(\nu),$$

where $\vec{Z}^{(i)} = \vec{Y}^{(i)} \text{Diag}(\nu)$ are i.i.d. multivariate normal random vectors with covariance matrix $\Sigma(\nu) := \text{Diag}(\nu) \Sigma \text{Diag}(\nu) = \text{Diag}(\nu_1^2 \lambda_1, \dots, \nu_p^2 \lambda_p)$. Thus, $E(\nu)$ is the error matrix of the maximum likelihood estimator of $\Sigma(\nu)$, so that the claim follows from (1.2). ■

3.2 Core Part of the Analysis

Lemma 3.1 suggests an approach to proving Theorem 1.1 based on “recycling” the results of Koltchinskii & Lounici [8] as follows: Inequality (2.16) shows that our aim to prove (1.7) can be achieved by bounding $\sqrt{\lambda_{i^*}/n}\|\vec{V}_{i^*}\|$ and $\|\Lambda_{i^*}\|$ with high probability. To do this, we need to construct an appropriate scaling vector ν and apply Lemma 3.1. Recall that $\vec{V}_{i^*} = (v_{j,i^*})$ satisfies

$$\sqrt{\frac{\lambda_{i^*}}{n}}v_{j,i^*} = \sqrt{\frac{\lambda_j\lambda_{i^*}}{n|\lambda_j - \hat{\lambda}_i|}} \cdot N_{j,i^*} = \sqrt{\frac{1}{|\lambda_j - \hat{\lambda}_i|}} \cdot E_{j,i^*}, \quad (j \neq i^*).$$

Let $\vec{W} = (w_j)$ be defined by $w_j := \tilde{\nu}_j \tilde{\nu}_{i^*} E_{j,i^*}$, where $\tilde{\nu}_j := |\lambda_j - \hat{\lambda}_i|^{-1/2}$ for $(j \neq i^*)$, and

$$\tilde{\nu}_{i^*} := \max_{j \neq i^*} \frac{\sqrt{\lambda_j}}{\sqrt{|\lambda_j - \lambda_{i^*}|} \sqrt{\lambda_{i^*}}}.$$

Under the notation of Lemma 3.1 the vector \vec{W} is then the i^* -th column of $E(\tilde{\nu})$. By construction, $\tilde{\nu}_{i^*} \cdot \sqrt{\lambda_{i^*}/n} \cdot v_{j,i^*} = w_j$ for $(j \neq i^*)$ and $v_{i^*,i^*} = 0$, so that

$$\frac{\tilde{\nu}_{i^*}^2}{\lambda_{i^*}} n \|\vec{V}_{i^*}\|^2 \leq \|\vec{W}\|^2 \leq \|E(\tilde{\nu})\|^2, \quad (3.1)$$

Next, recall that

$$\Lambda_{i^*} = |D_{i^*}|^{1/2} E |D_{i^*}|^{1/2} = P_{i^*} \text{Diag}(\tilde{\nu}) E \text{Diag}(\tilde{\nu}) P_{i^*}.$$

Therefore,

$$\|\Lambda_{i^*}\|^2 \leq \|E(\tilde{\nu})\|^2. \quad (3.2)$$

Now let $\nu_j := \sqrt{2}|\lambda_j - \lambda_{i^*}|^{-1/2}$ for $(j \neq i^*)$, and $\nu_{i^*} := \tilde{\nu}_{i^*}$, so that by virtue of (2.11), we have $\nu_j/\tilde{\nu}_j \geq 1$ for all j , and $E(\nu) = \text{Diag}((\nu_j/\tilde{\nu}_j)) E(\tilde{\nu}) \text{Diag}((\nu_j/\tilde{\nu}_j))$. This implies

$$\|E(\tilde{\nu})\|^2 \leq \|E(\nu)\|^2. \quad (3.3)$$

In order to prove our main result, we would like to show that for n large enough we have $\Delta\mu_{i^*,i} \leq q$ with a high probability. Taking into account (3.1), (3.2), (3.3) and (2.16) this is true,

if the conditions

$$\|E(\nu)\| \leq \frac{1}{2} \quad (3.4)$$

and

$$\|E(\nu)\| \leq q \cdot \frac{\sqrt{\text{gap}_{i^*}} \cdot \tilde{\nu}_{i^*}}{2\sqrt{2}} \quad (3.5)$$

hold with a high probability. Lemma 3.1 gives us hope that both conditions are satisfied with a high probability, if for a sufficiently large t and for $n \geq \max(p, t)$ we have

$$\Xi(n, t, \nu) \leq \min\left(\frac{1}{4}, \frac{q^2 \text{gap}_{i^*} \tilde{\nu}_{i^*}^2}{8}\right).$$

By definition of ν , $\|\nu^2 \lambda\|_\infty = 2\lambda_{i^*} \tilde{\nu}_{i^*}^2$ and by the definition of $\tilde{\nu}_{i^*}$,

$$\lambda_{i^*} \tilde{\nu}_{i^*}^2 = \max_{j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|} \leq \sum_{j \neq i^*} \frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|}, \quad \lambda_{i^*} \tilde{\nu}_{i^*}^2 \leq \left(1 + \frac{\lambda_{i^*}}{\text{gap}_{i^*}}\right),$$

where the last inequality follows from (2.12). Hence

$$\begin{aligned} \Xi(n, t, \nu) &= \frac{2C_1^2 \lambda_{i^*} \tilde{\nu}_{i^*}^2}{n} \times \max\left(2\lambda_{i^*} \tilde{\nu}_{i^*}^2 t, 2 \sum_{j \neq i^*} \frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|} + \lambda_{i^*} \tilde{\nu}_{i^*}^2\right) \\ &\leq \frac{4C_1^2 \lambda_{i^*} \tilde{\nu}_{i^*}^2}{n} \times \max\left(\lambda_{i^*} \tilde{\nu}_{i^*}^2 t, \frac{3}{2} \sum_{j \neq i^*} \frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|}\right) \\ &\leq \frac{4C_1^2}{n} \times (\lambda_{i^*} \tilde{\nu}_{i^*}^2) \times \max\left(\left(1 + \frac{\lambda_{i^*}}{\text{gap}_{i^*}}\right)t, \frac{3}{2} \sum_{j \neq i^*} \frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|}\right) =: \Phi(n, t, i^*). \end{aligned}$$

Therefore, we get that both (3.4) and (3.5) are satisfied with high probability, if in addition to $n \geq \max(p, t)$ we have

$$n \geq \Psi(q, t, i^*),$$

where

$$\Psi(q, t, i) = 16C_1^2 \max\left(\left(1 + \frac{\lambda_i}{\text{gap}_i}\right), 2 \frac{\lambda_i}{q^2 \text{gap}_i}\right) \max\left(\left(1 + \frac{\lambda_i}{\text{gap}_i}\right)t, \frac{3}{2} \sum_{j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_i|}\right).$$

Remark: Observe that $\Xi(n, t, \nu)$ could be bounded above also by

$$\frac{4C_1^2}{n} \times (\lambda_{i^*} \tilde{\nu}_{i^*}^2) \times t \times \left(\frac{5}{2} \sum_{j \neq i^*} \frac{\lambda_j}{|\lambda_j - \lambda_{i^*}|}\right) \quad (3.6)$$

Also note that when $\frac{\lambda_i}{\text{gap}_i} \geq 1$, then

$$\max \left(\left(1 + \frac{\lambda_i}{\text{gap}_i}\right), 2 \frac{\lambda_i}{q^2 \text{gap}_i} \right) = 2 \frac{\lambda_i}{q^2 \text{gap}_i}$$

and so with the bound (3.6), $\Psi(q, t, i)$ reduces to

$$80C_1^2 t \frac{\lambda_i}{q^2 \text{gap}_i} \left(\sum_{j \neq i} \frac{\lambda_j}{|\lambda_j - \lambda_i|} \right).$$

Note that we can not apply Lemma 3.1 directly, since ν is a random vector. So to apply the lemma, we have to get rid of randomness of ν which we achieve by conditioning i^* . To simplify notations, let us define the vector $\nu(k)$ by $\nu_j(k) = \sqrt{2}|\lambda_j - \lambda_k|^{-1/2}$, $j \neq k$ and

$$\nu_k(k) = \max_{j \neq k} \frac{\sqrt{\lambda_j}}{\sqrt{\lambda_j - \lambda_k} \sqrt{\lambda_k}}.$$

Thus, conditional on $i^* = k$ we have $\nu = \nu(k)$ and $\|E(\nu)\| = \|E(\nu(k))\|$, as well as

$$\Xi(n, t, \nu) = \Xi(n, t, \nu(k)) \leq \Phi(n, t, k).$$

Thus, for fixed $p \in \mathbb{N}$, $t \geq 1$, $n \geq \max(p, t)$ and $k \in \{1, \dots, p\}$ we have

$$\begin{aligned} \mathbb{P} \left[\|E(\nu)\|^2 \leq \Phi(n, t, i^*) \mid i^* = k \right] &\geq \mathbb{P} \left[\|E(\nu(k))\|^2 \leq \Xi(n, t, \nu(k)) \mid i^* = k \right] \\ &= \frac{\mathbb{P} \left[\{i^* = k\} \setminus \{\|E(\nu(k))\|^2 > \Xi(n, t, \nu(k))\} \right]}{\mathbb{P}[i^* = k]} \\ &\geq \frac{\mathbb{P}[i^* = k] - \mathbb{P} \left[\|E(\nu(k))\|^2 > \Xi(n, t, \nu(k)) \right]}{\mathbb{P}[i^* = k]} \\ &\geq 1 - \frac{\min(e^{-t}, \mathbb{P}[i^* = k])}{\mathbb{P}[i^* = k]}. \end{aligned} \tag{3.7}$$

3.3 Proof of Theorem 1.1

We are ready to prove Theorem 1.1, the main theorem of this paper: For all n, k, t and q we have

$$n \geq \Psi(q, t, k) \Rightarrow \Phi(n, t, k) \leq \frac{1}{4} \tag{3.8}$$

and

$$n \geq \Psi(q, t, k) \Rightarrow \Phi(n, t, k) \leq \frac{q^2 \text{gap}_k \nu_k^2(k)}{8} \quad (3.9)$$

Therefore, we have

$$\begin{aligned} \|E(\nu)\|^2 &\leq \Phi(n, t, i^*), \quad n \geq \Psi(q, t, i^*) \\ &\stackrel{(3.1),(3.2),(3.3)}{(3.8), (3.9)} \Rightarrow \|\Lambda_{i^*}\| \leq \frac{1}{2}, \quad \sqrt{\frac{\lambda_{i^*}}{n}} \|\vec{V}_{i^*}\| \leq q \sqrt{\frac{\text{gap}_{i^*}}{8}} \\ &\stackrel{(2.16)}{\Rightarrow} \|\Delta \vec{\mu}_{i^*, i}\| \leq \sqrt{\frac{2}{\text{gap}_{i^*}}} \times \frac{1}{1 - \frac{1}{2}} \times q \sqrt{\frac{\text{gap}_{i^*}}{8}} = q \\ &\stackrel{(1.6)}{\Rightarrow} \sin \theta_i \leq q. \end{aligned} \quad (3.10)$$

Moreover, since $\Psi(q, t, i^*)$ is stochastic only through its dependence on i^* , we have for k which satisfy $n \geq \Psi(q, t, k)$ the equality

$$\mathbb{P} \left[\|E(\nu)\|^2 \leq \Phi(n, t, i^*) \mid i^* = k, n \geq \Psi(q, t, i^*) \right] = \mathbb{P} \left[\|E(\nu)\|^2 \leq \Phi(n, t, i^*) \mid i^* = k \right]. \quad (3.11)$$

This yields (recall that by assumption $n \geq \Psi(q, t, i)$)

$$\begin{aligned} &\mathbb{P} \left[\sin \theta_i \leq q \mid n \geq \Psi(q, t, i^*) \right] \\ &\stackrel{(3.10)}{\geq} \mathbb{P} \left[\|E(\nu)\|^2 \leq \Phi(n, t, i^*) \mid n \geq \Psi(q, t, i^*) \right] \\ &\stackrel{(3.11)}{\geq} \sum_{k: n \geq \Psi(q, t, k)} \mathbb{P} \left[\|E(\nu)\|^2 \leq \Phi(n, t, i^*) \mid i^* = k \right] \times \mathbb{P} \left[i^* = k \mid n \geq \Psi(q, t, i^*) \right] \\ &\stackrel{(3.7)}{\geq} \sum_{k: n \geq \Psi(q, t, k)} \left(1 - \frac{\min(e^{-t}, \mathbb{P}[i^* = k])}{\mathbb{P}[i^* = k]} \right) \times \mathbb{P} \left[i^* = k \mid n \geq \Psi(q, t, i^*) \right] \\ &= 1 - \sum_{k: n \geq \Psi(q, t, k)} \frac{\min(e^{-t}, \mathbb{P}[i^* = k])}{\mathbb{P}[n \geq \Psi(q, t, i^*)]} \times \frac{\mathbb{P}[i^* = k, n \geq \Psi(q, t, i^*)]}{\mathbb{P}[i^* = k]} \\ &\geq 1 - \frac{e^{-t}}{\mathbb{P}[n \geq \Psi(q, t, i^*)]} \times \left(\sum_{k: n \geq \Psi(q, t, k)} \mathbb{P} \left[n \geq \Psi(q, t, i^*) \mid i^* = k \right] \right). \end{aligned}$$

4 Outlook and Open Problems

Example 1.2 shows that the bound (1.8) is of the correct order when the random index i^* is replaced by the deterministic index i . For conceptual reasons it would be desirable to be able to bound the number of samples required for strong eigenvector alignment in terms of the index i alone, instead of the weaker theory based on i^* developed in this paper. It would also be interesting to have theoretical lower bounds that show that the improved bound (1.5) is tight up to a multiplicative constant. Finally, any practically applicable version of (1.5) has to rely on replacing the unknown population eigenvalues λ_j by de-biased estimates, and one could work out how many samples are needed for the replaced values to yield a bound that comes sufficiently close to the bound that corresponds to the population eigenvalues. We plan to address these tasks in future research.

References

- [1] T.W. Anderson, An introduction to multivariate statistical analysis. Wiley, Third Edition, 2003.
- [2] Dimension reduction with PCA. <http://www.polarmicrobes.org/dimension-reduction-with-pca/>. Accessed: 2017-07-12.
- [3] ALTER, O. & BROWN, P.O. & BOTSTEIN, D. (2000) Singular value decomposition for genome-wide expression data processing and modeling. PNAS, 97(18): pp 10101–10106.
- [4] ALTER, O. & GOLUB, G. (2006) Singular value decomposition of genome-scale MRNA lengths distribution reveals asymmetry in RNA gel electrophoresis band broadening. PNAS, 103(32): pp 11828–11833.

- [5] BERTAGNOLI, N.M. & DRAKE, J.A. & TENNENSEN, J.M. & ALTER, O. (2013) SVD identifies transcript length distribution functions from DNA microarray data and reveals evolutionary forces globally affecting GBM metabolism. *PLOS One*, 8(e78913).
- [6] BAI, Z. & YAO, J.-F. (2008) Central limit theorems for eigenvalues in a spiked population model. *Ann. Inst. Henri Poincaré Probab. Stat.*, 44(3): pp 447–474, 2008. DOI: 10.1214/07-AIHP118.
- [7] DAVIS, C. & KAHAN, W. (1970) The rotation of eigenvectors by a perturbation. *SIAM J. Numer. Anal.*, (7): pp 1–46.
- [8] KOLTSCHINKSKII, V. & LOUNICI, K. (2017) Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23(1): pp 110–133.
- [9] LEDOIT, O. & PECHE, S. (2011) Eigenvectors of some large sample covariance matrix ensembles. *Probab. Theory Relat. Fields*, 151(1-2): pp 233–264, <https://doi.org/10.1007/s00440-010-0298-3>.
- [10] OMBERG, L. & MEYERSON, J.R. & KOBAYASHI, K. & DRURY, L.S. & DIF-FLEY, J.F.X. & ALTER, O. (2009) Global effects of DNA replication and DNA replication origin activity on eukaryotic gene expression. *Molecular Systems Biology*, 312(5). DOI 10.1038/msb.2009.70.
- [11] YU, Y. & WANG, T. & SAMWORTH, R. (2015) A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika*, 102(2), pp 315–323, <https://doi.org/10.1093/biomet/asv008>.

Acknowledgment: The authors wish to thank the anonymous reviewers for their valuable comments that led to several substantial improvements. This work was supported by EPSRC

[EP/N510129/1 to R.H. through the Alan Turing Institute]; and Estonian Institutional Research
Funding [IUT34-5 to J.L. and R.K.]