On the Quality of a Semidefinite Programming Bound for Sparse Principal Component Analysis

Laurent El Ghaoui
EECS Department, UC Berkeley
elghaoui@eecs.berkeley.edu

February 1, 2006

Abstract

We examine the problem of approximating a positive, semidefinite matrix Σ by a dyad xx^T, with a penalty on the cardinality of the vector x. This problem arises in sparse principal component analysis, where a decomposition of Σ involving sparse factors is sought. We express this hard, combinatorial problem as a maximum eigenvalue problem, in which we seek to maximize, over a box, the largest eigenvalue of a symmetric matrix that is linear in the variables. This representation allows to use the techniques of robust optimization, to derive a bound based on semidefinite programming. The quality of the bound is investigated using a technique inspired by Nemirovski and Ben-Tal (2002).

Contents

Notation

The notation $\mathbf{1}$ denotes the vector of ones (with size inferred from context), while $\mathbf{Card}(x)$ denotes the cardinality of a vector x (number of non-zero elements), and $D(x)$ the diagonal matrix with the elements of x on its diagonal. We denote by e_i the unit vectors of $\mathbb{R}^n$. For a $n \times n$ matrix X, $X \succeq 0$ means X is symmetric and positive semi-definite. The notation B_+, for a symmetric matrix B, denotes the matrix obtained from B by replacing negative eigenvalues by 0. The notation has precedence over the trace operator, so that $\mathbf{Tr}\, B_+$ denotes the sum of positive eigenvalues of B if any, and 0 otherwise. Throughout, the symbol $\mathbf{E}$ refers to expectations taken with respect to the normal Gaussian distribution of dimension inferred from context. Finally, the support of a vector x is defined to be the set of indices corresponding to its non-zero elements.

1 Introduction

Given a non-zero $n \times n$ positive semi-definite symmetric matrix Σ and a scalar $\rho > 0$, we consider the *cardinality-penalized variational problem*

$$\phi(\rho) := \max_x \; x^T \Sigma x - \rho\, \mathbf{Card}(x) \;:\; \|x\|_2 = 1. \tag{1}$$

This problem is equivalent to solving the sparse rank-one approximation problem

$$\min_z \; \|\Sigma - zz^T\|_F^2 - \rho\, \mathbf{Card}(z),$$

which arises in the *sparse PCA* problem [4, 2], where a "decomposition" of Σ into sparse factors is sought. We refer to [2] for a motivation of the sparse PCA problem, and an overview of its many applications.

In the paper [2], the authors have developed the "direct sparse PCA" approach, which leads to the following convex relaxation for the problem (1):

$$\max_X \; \mathbf{Tr}\, X\Sigma - \rho \|X\|_1 \;:\; X \succeq 0, \;\; \mathbf{Tr}\, X = 1.$$

The above problem is amenable to both general-purpose semidefinite programming (SDP) interior-point codes, and more recent first-order algorithms such as Nesterov's smooth minimization technique [3]. Unfortunately, the quality of the relaxation seems to be hard to analyze at present.

In this paper, we introduce two new representations of the problem, and a new SDP bound, based on robust optimization ideas [1]. Our main goal is to use the new representations of the problem to analyze the quality of the corresponding bound.

The paper is organized as follows. Section 2 develops some preliminary results allowing to restrict our attention to the case when $\rho < \max_i \Sigma_{ii}$. Section 3 then proposes two new representations for $\phi(\rho)$, one based on largest eigenvalue maximization, and the other on a thresholded version of the Rayleigh quotient. In section 4, we derive an SDP-based upper bound on $\phi(\rho)$, and in section 5, we analyze its quality: as a function of the penalty parameter ρ first, then in terms of structural conditions on matrix Σ.

It will be helpful to describe Σ in terms of the Cholesky factorization $\Sigma = A^T A$, where $A = [a_1 \dots a_n]$, with $a_i \in \mathbf{R}^m$, $i = 1, \dots, n$, where $m = \mathbf{Rank}(\Sigma)$. Further, we will assume, without loss of generality, that the diagonal of Σ is ordered, and none of the diagonal elements is zero, so that $\Sigma_{11} \geq \dots \geq \Sigma_{nn} > 0$. Finally, we define the set $\mathcal{I}(\rho) := \{i \ : \ \Sigma_{ii} > \rho\}$, and let $n(\rho) := \mathbf{Card}\,\mathcal{I}(\rho)$.

2 Equality vs. Inequality Models

In the sequel we will develop SDP bounds for the related quantity

$$\overline{\phi}(\rho) := \max_x \ x^T \Sigma x - \rho\,\mathbf{Card}(x) \ : \ \|x\|_2 \leq 1. \tag{2}$$

The following theorem says that when $\rho < \Sigma_{11}$, the two quantities $\phi(\rho)$, $\overline{\phi}(\rho)$ are positive and equal; otherwise, both $\phi(\rho)$ and $\overline{\phi}(\rho)$ have trivial solutions.

Theorem 1 *If $\rho < \Sigma_{11}$, we have $\phi(\rho) = \overline{\phi}(\rho) > 0$, and the optimal sets of problems (1) and (2) are the same. Conversely, if $\rho \geq \Sigma_{11}$, we have $\overline{\phi}(\rho) = 0 \geq \phi(\rho) = \Sigma_{11} - \rho$, and a corresponding optimal vector for $\phi(\rho)$ (resp. $\overline{\phi}(\rho)$) is $x = e_1$, the first basis vector in $\mathbb{R}^n$ (resp. $x = 0$).*

Proof: If $\rho < \Sigma_{11}$, then the choice $x = e_1$ in (2) implies $\overline{\phi}(\rho) > 0$, which in turn implies that an optimal solution x^* for (2) is not zero. Since the **Card** function is scale-invariant, it is easy to show that without loss of generality, we can assume that x^* has l_2-norm equal to one, which then results in $\phi(\rho) = \overline{\phi}(\rho) > 0$.

Let us now turn to the case when $\rho \geq \Sigma_{11}$. We develop an expression for $\phi(\rho)$ as follows. First observe that, since $\Sigma \succeq 0$,

$$\max_{\|x\|_1 = 1} \ x^T \Sigma x = \Sigma_{11},$$

which implies that, for every x,

$$\Sigma_{11} \|x\|_1^2 \geq x^T \Sigma x. \tag{3}$$

Now let $t \geq 0$. The condition $\phi(\rho) \leq -t$ holds if and only if

$$\forall\, x, \ \|x\|_2 = 1 \ : \ \rho\,\mathbf{Card}(x) \geq t + x^T \Sigma x.$$

Specializing the above condition to $x = e_1$, we obtain that $\phi(\rho) \leq -t$ implies $\rho \geq \Sigma_{11} + t$. Conversely, assume that $\rho \geq \Sigma_{11} + t$. Using (3), we have for every x, $\|x\|_2 = 1$:

$$\rho\,\mathbf{Card}(x) \geq \rho \|x\|_1^2 \geq (\Sigma_{11} + t)\|x\|_1^2 \geq x^T \Sigma x + t,$$

where we have used the fact that $\|x\|_1 \geq 1$ whenever $\|x\|_2 = 1$. Thus we have obtained that $\phi(\rho) \leq -t$ with $t \geq 0$ if and only if $\rho \geq \Sigma_{11} + t$, which means that $\phi(\rho) = \Sigma_{11} - \rho$ whenever $\rho \geq \Sigma_{11}$.

Finally, let us prove that $\overline{\phi}(\rho) = 0$ when $\rho \geq \Sigma_{11}$. For every $x \neq 0$ such that $\|x\|_2 \leq 1$, we have

$$\rho\,\mathbf{Card}(x) \geq \frac{\rho}{\|x\|_2^2} \|x\|_1^2 \geq \rho \|x\|_1^2 \geq x^T \Sigma x,$$

which shows that $\overline{\phi}(\rho) \leq 0$, and concludes our proof. ■

In the sequel, we will make the following assumption.

Assumption 1 *We assume that $\rho < \Sigma_{11}$, that is, the set $\mathcal{I}(\rho) := \{i : \Sigma_{ii} > \rho\}$ is not empty.*

3 New Representations

3.1 Largest eigenvalue maximization

The following theorem shows that the problem of computing $\phi(\rho)$ can be expressed as a eigenvalue maximization problem, where the sparsity pattern is the decision variable.

Theorem 2 *For $\rho \in [0, \Sigma_{11}[$, $\phi(\rho)$ can be expressed as the maximum eigenvalue problem*

$$\phi(\rho) = \max_{u \in [0,1]^n} \lambda_{\max} \left(\sum_{i=1}^n u_i B_i \right), \tag{4}$$

where $B_i := a_i a_i^T - \rho \cdot I_m$, $i = 1, \ldots, n$.

An optimal solution to the original problem (1) is obtained from a sparsity pattern vector u that is optimal for (4), by finding an eigenvector y corresponding to the largest eigenvalue of $D(u)\Sigma D(u)$, and setting $x = D(u)y/\|D(u)y\|_2$, where $D(u) := \mathbf{diag}(u)$.

Proof. Since $\rho < \Sigma_{11}$, the result of Theorem 1 implies that $\phi(\rho)$ is equal to $\overline{\phi}(\rho)$ defined in in (2). Let us now prove that $\phi(\rho) = \tilde{\phi}(\rho)$, where

$$\tilde{\phi}(\rho) := \max_{u \in \{0,1\}^n} \max_{y^T y \leq 1} y^T D(u) \Sigma D(u) y - \rho \cdot \mathbf{1}^T u. \tag{5}$$

To prove this intermediate result, first note that if x is optimal for $\phi(\rho)$, that is, for (2), then we can set $u_i = 1$ if $x_i \neq 0$, $u_i = 0$ otherwise, so that $\mathbf{Card}(x) = \mathbf{1}^T u$; then, we set $y = x$ and obtain that the pair (u, y) is feasible for $\tilde{\phi}(\rho)$, and achieves the objective value $\phi(\rho)$, hence $\phi(\rho) \leq \tilde{\phi}(\rho)$. Conversely, if (u, y) is optimal for $\tilde{\phi}(\rho)$, then $x = D(u)y$ is feasible for $\phi(\rho)$ (as expressed in (2)), and satisfies $\mathbf{Card}(x) \leq \mathbf{Card}(u) = \mathbf{1}^T u$, thus

$$\tilde{\phi}(\rho) = y^T D(u) \Sigma D(u) y - \rho \mathbf{1}^T u \leq x^T \Sigma x - \rho \, \mathbf{Card}(x) \leq \phi(\rho),$$

This concludes the proof that $\phi(\rho) = \tilde{\phi}(\rho)$.

We proceed by eliminating y from (5), as follows:

$$\begin{aligned} \phi(\rho) &= \max_{u \in \{0,1\}^n} \lambda_{\max}(D(u)\Sigma D(u)) - \rho \cdot \mathbf{1}^T u \\ &= \max_{u \in \{0,1\}^n} \lambda_{\max}(D(u) A^T A D(u)) - \rho \cdot \mathbf{1}^T u \\ &= \max_{u \in \{0,1\}^n} \lambda_{\max}(A D(u) A^T) - \rho \cdot \mathbf{1}^T u \\ &= \max_{u \in \{0,1\}^n} \lambda_{\max}(\sum_{i=1}^n u_i a_i a_i^T) - \rho \cdot \mathbf{1}^T u, \end{aligned}$$

in virtue of $\Sigma = A^T A$, and $D(u)^2 = D(u)$ for every feasible u. Invoking the convexity of the largest eigenvalue function, we can replace the set $\{0,1\}^n$ by $[0,1]^n$ in the above expression, and obtain (4). ■

3.2 Thresholded Rayleigh quotient

The following theorem shows that $\phi(\rho)$ can be expressed as a maximal "thresholded Rayleigh quotient", which for $\rho = 0$ reduces to the ordinary Rayleigh quotient.

Theorem 3 *For $\rho \in [0, \Sigma_{11}[$, we have*

$$\begin{aligned} \phi(\rho) &= \max_{\xi^T\xi=1} \sum_{i=1}^{n} ((a_i^T\xi)^2 - \rho)_+, & (6)\\ &= \max_{X} \sum_{i=1}^{n} (a_i^T X a_i - \rho)_+ \;:\; X \succeq 0, \;\; \mathbf{Tr}\, X = 1. & (7) \end{aligned}$$

An optimal solution x for (1) is obtained from an optimal solution ξ to problem (6) by setting $u_i = 1$ if $(a_i^T\xi)^2 > \rho$, $u_i = 0$ otherwise; then, finding an eigenvector y corresponding to the largest eigenvalue of $D(u)\Sigma D(u)$, and setting $x = D(u)y/\|D(u)y\|_2$.

Proof: From the expression (4), we derive

$$\begin{aligned} \phi(\rho) &= \max_{u\in[0,1]^n} \max_{\xi^T\xi\le 1} \xi^T \left(\sum_{i=1}^{n} u_i a_i a_i^T\right) \xi - \rho \cdot \mathbf{1}^T u \\ &= \max_{\xi^T\xi\le 1} \sum_{i=1}^{n} ((a_i^T\xi)^2 - \rho\xi^T\xi)_+ \\ &= \max_{\xi^T\xi=1} \sum_{i=1}^{n} ((a_i^T\xi)^2 - \rho)_+, & (8) \end{aligned}$$

where the last equality derives from the fact that $\phi(\rho) > 0$ (which is in turn the consequence of our assumption that $\Sigma_{11} = \max_i a_i^T a_i > \rho$). Finally, the equivalence between (6) and (7) stems from convexity of the objective function in problem (7), which implies that without loss of generality, we can impose X to be of rank one in (7). ■

The following corollary shows that we can safely remove columns and rows in Σ that have variance below the threshold ρ.

Corollary 1 *Without loss of generality, we can assume that every optimal solution to the original problem (1) has a support included in the set $\mathcal{I}(\rho) := \{i : \Sigma_{ii} > \rho\}$. Thus, if $\Sigma_{ii} \le \rho$, the corresponding column and row can be safely removed from Σ.*

Proof: This is a direct implication of the fact that for every i, if $\rho \ge a_i^T a_i$, then we have $(a_i^T\xi)^2 \le \rho$ for every ξ such that $\xi^T\xi = 1$. Hence, the corresponding term does not appear in the sum in (8). ■

3.3 Exact Solutions in Some Special cases

Theorems 2 and 3 allows to solve exactly the problem in some special cases.

First, Theorem 2 can be invoked when Σ is diagonal, in which case the optimal vector x turns out to be simply the first unit vector, e_1.

Next, consider the case when the matrix Σ has rank one, that is, $m = 1$. Then, the a_i's are scalars, and the representation given in Theorem 3 yields

$$\phi(\rho) = \max_{\xi^2=1} \sum_{i=1}^{n} ((a_i \xi)^2 - \rho)_+ = \sum_{i=1}^{n} (a_i^2 - \rho)_+.$$

A corresponding optimal solution for $\phi(\rho)$ is obtained by setting $u_i = 1$ if $\rho < a_i^2$, $u_i = 0$ otherwise, and then setting $x = \tilde{a}/\|\tilde{a}\|_2$, with $\tilde{a}$ obtained from a by thresholding a with absolute level $\sqrt{\rho}$. In the sequel, we assume that $m > 1$.

A similar result holds when Σ has the form $\Sigma = I + aa^T$, when a is a given n-vector, since then the problem trivially reduces to the rank-one case.

4 SDP relaxation

A relaxation inspired by [1] is given by the following theorem.

Theorem 4 *For every $\rho \in [0, \Sigma_{11}[$, we have $\phi(\rho) \leq \psi(\rho)$, where $\psi(\rho)$ is the solution to a semidefinite program:*

$$\psi(\rho) := \min_{(Y_i)_{i=1}^n} \lambda_{\max}\left(\sum_{i=1}^{n} Y_i\right) \;:\; Y_i \succeq B_i, \;\; Y_i \succeq 0, \;\; i = 1, \ldots, n. \tag{9}$$

The problem can be represented in dual form, as the convex problem

$$\psi(\rho) = \max_{X} \sum_{i=1}^{n} \mathbf{Tr}\left(X^{1/2} a_i a_i^T X^{1/2} - \rho X\right)_+ \;:\; X \succeq 0, \;\; \mathbf{Tr}\, X = 1. \tag{10}$$

Proof: If $(Y_i)_{i=1}^n$ is feasible for the above SDP, then for every $\xi \in \mathbf{R}^m$, $\xi^T \xi \leq 1$, and $u \in [0,1]^n$, we have

$$\xi^T \left(\sum_{i=1}^{n} u_i B_i\right) \xi \leq \sum_{i=1}^{n} (\xi^T B_i \xi)_+ \leq \xi^T \left(\sum_{i=1}^{n} Y_i\right) \xi \leq \psi(\rho),$$

which proves $\phi(\rho) \leq \psi(\rho)$. The dual of the SDP (9) is given by

$$\psi(\rho) = \max_{X, (P_i)_{i=1}^n} \sum_{i=1}^{n} \langle P_i, B_i \rangle \;:\; X \succeq P_i \succeq 0, \;\; i = 1, \ldots, n, \;\; \mathbf{Tr}\, X = 1. \tag{11}$$

Using the fact that, for any symmetric matrix B, and positive semi-definite matrix X,

$$\max_{P} \{\langle P, B \rangle \;:\; X \succeq P \succeq 0\} = \mathbf{Tr}\left(X^{1/2} B X^{1/2}\right)_+,$$

allows to represent the dual problem in the form (10). Note that the convexity of the representation (10) is not immediately obvious. ■

A few comments are in order.

The fact that $\phi(\rho) \leq \psi(\rho)$ can also be inferred directly from the dual expression (10): we have, by convexity, and using the representation (7) for $\phi(\rho)$,

$$\psi(\rho) \geq \max_X \left\{ \sum_{i=1}^n \left(a_i^T X a_i - \rho\right)_+ \ : \ X \succeq 0, \ \ \mathbf{Tr}\, X = 1 \right\} = \phi(\rho).$$

From the representation (10) and this, we obtain that if the rank k of X at the optimum of the dual problem (10) is one, then our relaxation is exact: $\phi(\rho) = \psi(\rho)$.

In fact, problem (10) can be obtained as a rank relaxation of the following exact representation of ϕ:

$$\phi = \max_X \sum_{i=1}^n \mathbf{Tr} \left(X^{1/2} a_i a_i^T X^{1/2} - \rho X\right)_+ \ : \ X \succeq 0, \ \ \mathbf{Tr}\, X = 1, \ \ \mathbf{Rank}(X) = 1.$$

In contrast, applying a direct rank relaxation to problem (6) (that is, writing the problem in terms of letting $X = \xi\xi^T$ and dropping the rank constraint on X) would be useless: it would yield (7), which is $\phi(\rho)$ itself.

Finally, note that our relaxation shares the property of the exact formulation (6) observed in Corollary 1, that indices i such that $\rho \geq \Sigma_{ii}$ can be simply ignored, since then $B_i \preceq 0$.

5 Quality of the SDP relaxation

In this section, we seek to estimate a lower bound on the *quality* of the SDP relaxation, which we define to be a scalar $\theta \in [0,1]$ such that

$$\theta\psi(\rho) \leq \phi(\rho) \leq \psi(\rho). \tag{12}$$

Thus, $(1-\theta)/\theta$ is a upper bound on the relative approximation error, $(\psi(\rho) - \phi(\rho))/\phi(\rho)$.

5.1 Quality estimate as a function of the penalty parameter

Our first result gives a bound on the relaxation quality conditional on a bound on ρ. We begin by making the following assumption:

Assumption 2 *We assume that* $0 < \rho < \min_{1\leq i\leq n} \Sigma_{ii} = \Sigma_{nn}$*, and* $m = \mathbf{Rank}(\Sigma) > 1$.

From the result of Corollary 1, we can always reduce the problem so that the above assumption holds, by removing appropriate columns and rows of Σ if necessary.

Theorem 5 *With assumption 2 in force, for every value of the penalty parameter* $\rho \in [0, \Sigma_{nn}[$*, and for every* $\gamma \geq 0$ *such that*

$$\rho \leq \frac{\gamma}{n+\gamma}\Sigma_{11}, \tag{13}$$

the bound (12) holds with θ set to $\theta_m(\gamma)$, where for $m > 1$ and $\gamma \geq 0$, we define

$$\theta_m(\gamma) \quad := \quad \mathbf{E}\left(\xi_1^2 - \frac{\gamma}{m-1} \cdot \sum_{j=2}^m \xi_j^2\right)_+ , \tag{14}$$

which can be computed by the formula

$$\theta_m(\gamma) \quad = \quad \frac{\displaystyle\int_0^{\pi/2} \left(\cos^2(t) - \frac{\gamma}{m-1}\sin^2(t)\right)_+ \sin^{m-2}(t)dt}{\displaystyle\int_0^{\pi/2} \cos^2(t)\sin^{m-2}(t)dt}. \tag{15}$$

For every $\gamma \geq 0$, the value $\theta_m(\gamma)$ decreases with m, and admits the bound

$$\theta_m(\gamma) \geq \frac{1}{2}\left(1 - \gamma + \frac{2}{\pi}\sqrt{1 + \frac{\gamma^2}{m-1}}\right)_+ . \tag{16}$$

In particular, if ρ satisfies (13) with $\gamma = 1$, that is, $\rho \leq \Sigma_{11}/(n+1)$, then bound (12) holds with $\theta \geq 1/\pi$.

Before we prove the theorem, let us make a few comments.

First, as will be apparent from the proof, the value of m can be safely replaced by the rank k of an optimal solution to the SDP (10). This can only improve the quality estimate, as $k \leq m$ and $\theta_m(\gamma)$ is a decreasing function of m for every $\gamma \geq 0$.

Second, the smaller m is, and the larger γ is, the smaller the corresponding quality estimate. However, a small value for γ does not allow for a large range of ρ values via (13), and this effect is becomes more pronounced as n grows. The theorem presents the result in such a way that the respective contributions of m, n to the deterioration of the quality estimate are separated. A plot of the function θ_m for various values of m is shown in Figure 1.

Third, the theorem allows to plot the predicted quality estimate θ as a function of the penalty parameter, in the interval $[0, \Sigma_{nn}[$. Leveraging these results to the entire range $[0, \Sigma_{11}[$ will be straightforward, but will require us to be careful about the sizes n and m, as they change as ρ crosses the values $\Sigma_{n-1,n-1}, \ldots, \Sigma_{11}$, in view of Corollary 1. We formalize the argument in Corollary 2.

Finally, the theorem allows to derive conditions on the structure of Σ that guarantee a prescribed value of the quality. We describe such a condition in Corollary 3.

Proof of theorem 5: The approach we use in our proof is inspired by that of Theorem 2.1 in [1]. Let $X \succeq 0$, $\mathbf{Tr}\, X = 1$, be optimal for the upper bound $\psi(\rho)$ in dual form (10), so that

$$\psi(\rho) = \sum_{i=1}^n \mathbf{Tr}(B_i(X)_+),$$

where $B_i(X) := X^{1/2} B_i X^{1/2}$. Let $k = \mathbf{Rank}(X)$. We have seen that if $k = 1$, then our relaxation is exact: $\phi(\rho) = \psi(\rho)$. If the rest of the proof, we will assume that $k > 1$. We thus have $1 < k \leq m = \mathbf{Rank}(\Sigma) \leq n$.

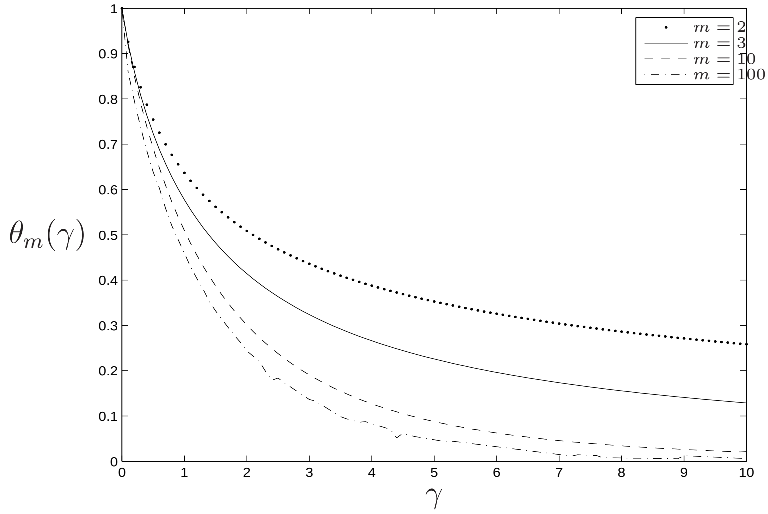

Figure 1: Plot of function $\theta_m(\gamma)$, as defined in (15), for various values of m.

Assume that we find a scalar $\theta \in [0, 1]$ such that:

$$\mathbf{E} \sum_{i=1}^{n} (\xi^T B_i(X)\xi)_+ > (\theta\psi(\rho)) \cdot \mathbf{E}(\xi^T X \xi), \tag{17}$$

where ξ follows the normal distribution in $\mathbb{R}^m$. The bound above implies that there exist a non-zero $\xi \in \mathbf{R}^m$ such that

$$\sum_{i=1}^{n} (\xi^T B_i(X)\xi)_+ > (\theta\psi(\rho)) \cdot (\xi^T X \xi).$$

Thus, with $u_i = 1$ if $\xi^T B_i(X)\xi > 0$, $u_i = 0$ otherwise, we obtain that there exist a non-zero $\xi \in \mathbf{R}^m$ and $u \in [0, 1]^n$ such that

$$\xi^T \left(\sum_{i=1}^{n} u_i B_i(X) \right) \xi > (\theta\psi(\rho)) \cdot (\xi^T X \xi).$$

With $z = X^{1/2}\xi$:

$$z^T \left(\sum_{i=1}^{n} u_i B_i \right) z > (\theta\psi(\rho)) \cdot (z^T z).$$

The above implies that $z \neq 0$, so we conclude that there exist $u \in [0, 1]^n$ such that

$$\lambda_{\max} \left(\sum_{i=1}^{n} u_i B_i \right) > \theta\psi(\rho),$$

from which we obtain the quality estimate $\theta\psi(\rho) \leq \phi(\rho) \leq \psi(\rho)$. By a continuity argument, this result still holds if (17) is satisfied, but not strictly. The rest of the proof is dedicated to finding a scalar θ such that the bound (17) holds.

Fix $i \in \{1,\dots,n\}$. It is easy to show that $B_i(X)$ has exactly one positive eigenvalue α_i, since assumption 2 holds. Thus $\alpha_i = \mathbf{Tr}\, B_i(X)_+$. Further, $B_i(X)$ has exactly rank $k = \mathbf{Rank}(X)$. Denote by $(-\beta_j^i)_{j=1}^{k-1}$ the negative eigenvalues of $B_i(X)$, ordered such that $\beta_1^i \geq \ldots \geq \beta_{k-1}^i$. Likewise, denote by $\{\lambda_j\}_{j=1}^k$ the non-zero eigenvalues of X, ordered such that $\lambda_1 \geq \ldots \geq \lambda_k$. Using an interlacing property of eigenvalues, we can show that

$$\beta_j^i \leq \rho\lambda_j, \quad j = 1,\dots,k-1.$$

Thus,

$$\sum_{j=1}^{k-1} \beta_j^i \leq \rho \sum_{j=1}^{k-1} \lambda_j = \rho(1-\lambda_k) \leq \rho.$$

Let $\xi \sim \mathcal{N}(0, I_m)$. By rotational invariance of the normal distribution, we have:

$$\mathbf{E}(\xi^T B_i(X)\xi)_+ \;=\; \mathbf{E}\left(\alpha_i\xi_1^2 - \sum_{j=1}^{k-1}\beta_j^i\xi_{j+1}^2\right)_+ .$$

Thus,

$$\begin{aligned}
\mathbf{E}(\xi^T B_i(X)\xi)_+ &\geq \min_{\beta\in\mathbb{R}^k}\left\{\mathbf{E}\left(\alpha_i\xi_1^2 - \sum_{j=1}^{k-1}\beta_j\xi_{j+1}^2\right)_+ \;:\; \beta \geq 0,\; \sum_{j=1}^{k-1}\beta_j \leq \rho\right\} && (18)\\
&\geq \min_{\beta\in\mathbb{R}^m}\left\{\mathbf{E}\left(\alpha_i\xi_1^2 - \sum_{j=1}^{m-1}\beta_j\xi_{j+1}^2\right)_+ \;:\; \beta \geq 0,\; \sum_{j=1}^{m-1}\beta_j \leq \rho\right\} && (19)\\
&= \mathbf{E}\left(\alpha_i\xi_1^2 - \frac{\rho}{m-1}\sum_{j=1}^{m-1}\xi_{j+1}^2\right)_+ , && (20)
\end{aligned}$$

where we have exploited the convexity and symmetry in problem (19). (As claimed in the first remark made after Theorem 5, we could safely keep k instead of m in the remaining of the proof.)

Summing over i, and in view of $\psi(\rho) = \sum_{i=1}^n \alpha_i$, we get:

$$\begin{aligned}
\mathbf{E}\sum_{i=1}^n(\xi^T B_i(X)\xi)_+ &\geq \sum_{i=1}^n \mathbf{E}\left(\alpha_i\xi_1^2 - \frac{\rho}{m-1}\sum_{j=1}^{m-1}\xi_{j+1}^2\right)_+ && \text{(by the bound (20))}\\
&\geq \mathbf{E}\left(\psi(\rho)\xi_1^2 - \frac{n\rho}{m-1}\sum_{j=1}^{m-1}\xi_{j+1}^2\right)_+ && \text{(by homogeneity and convexity)}\\
&\geq \theta_m(\gamma)\cdot\psi(\rho), && (21)
\end{aligned}$$

provided $\gamma \geq n\rho/(m-1)\psi(\rho)$. Using the fact that $\psi(\rho) \geq \phi(\rho) \geq \Sigma_{11} - \rho$, we obtain that the bound (12) holds with $\theta = \theta_m(\gamma)$ whenever (13) does, as claimed in the theorem. The expression (15) of the function θ_m is proved in Appendix A, while the bound (16) is proved in Appendix B. ■

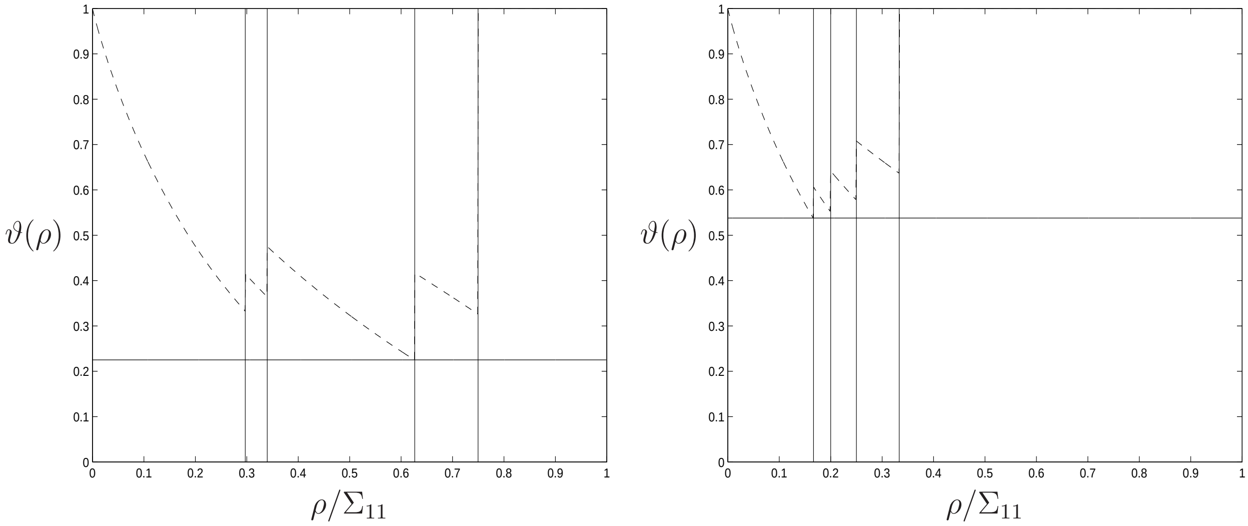

Figure 2: Plot of the function $\vartheta(\rho)$ defined in Corollary 2, for a specific 5×5 covariance matrix Σ. The left pane corresponds to a random matrix, and the right pane, to a random matrix that satisfies the conditions of Corollary 3.

The following corollary allows to plot the quality estimate, as derived from Theorem 5, as a function of ρ across the entire range $[0, \Sigma_{11}[$. We do not make the assumption 2 anymore, but do keep assumption 1.

Corollary 2 *Let* $\rho \in [0, \Sigma_{11}[$, *and define* $n(\rho) = \mathbf{Card}\{i : \Sigma_{ii} > \rho\} > 0$ *and* $m(\rho) = \mathbf{Rank}(\Sigma(\rho))$, *where* $\Sigma(\rho)$ *is the* $n(\rho) \times n(\rho)$ *matrix obtained by removing the last* $n - n(\rho)$ *rows and columns in* Σ. *The bound (12) holds for* $\theta = \vartheta(\rho)$, *where*

$$\vartheta(\rho) = \begin{cases} \theta_{m(\rho)}(\gamma(\rho)), \quad \gamma(\rho) = \dfrac{n(\rho)}{m(\rho)-1} \cdot \dfrac{\rho}{\Sigma_{11} - \rho} & \text{if } m(\rho) > 1, \\ 1 & \text{otherwise.} \end{cases}$$

An example of the resulting plot is shown in Figure 2.

5.2 Quality estimate based on the structure of Σ

The next result illustrates how to obtain a quality estimate based on structural assumptions on Σ, requiring that its ordered diagonal decreases fast enough.

Corollary 3 *Assume* $\Sigma_{11} > \ldots > \Sigma_{nn}$. *If* $\Sigma_{22} \leq \rho < \Sigma_{11}$, *then the bounds (12) hold with* $\theta = 1$, *that is,* $\phi(\rho) = \psi(\rho)$. *If in addition, we have, for every* $h \in \{2, \ldots, n\}$

$$\Sigma_{hh} \leq \frac{1}{h+1} \Sigma_{11}, \tag{22}$$

then, whenever $0 < \rho < \Sigma_{22}$, *the bounds (12) hold with* $\theta \geq 1/\pi$.

Proof: In the case $\rho \in [\Sigma_{22}, \Sigma_{11}[$, $n(\rho) = 1$, so that $m(\rho) = 1$, and the bound (12) holds with $\theta = 1$. Now let ρ be such that $0 < \rho < \Sigma_{22}$. Then there exist $h \in \{2, \ldots, n\}$ such that $\Sigma_{h+1,h+1} \leq \rho < \Sigma_{hh}$, with the convention $\Sigma_{n+1,n+1} = 0$. In this case, $n(\rho) = \mathbf{Card}\{i : \Sigma_{ii} > \rho\} = h$, so that the sufficient condition (13) with $\gamma = 1$ writes

$$\rho \leq \frac{1}{(h+1)}\Sigma_{11},$$

which, in view of $\Sigma_{h+1,h+1} \leq \rho < \Sigma_{hh}$, holds when (22) holds, independent of ρ. Applying the bound (16) ends the proof. ■

An example corresponding to the situation of Corollary 3 is shown in Figure 2 (left pane).

Acknowledgements

References

[1] A. Ben-Tal and A. Nemirovski. On tractable approximations of uncertain linear matrix inequalities affected by interval uncertainty. *SIAM J. on Optimization*, 12(3):811–833, 2002.

[2] Alexandre d'Aspremont, Laurent El Ghaoui, Michael I. Jordan, and Gert R.G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. In *Advances in Neural Information Processing Systems 17*, pages 41–48. MIT Press, Cambridge, MA, 2005.

[3] Yu. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, May 2005.

[4] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Technical report, statistics department, Stanford University*, 2004.

A A Formula for θ_m

Let $\gamma > 0$, $m > 1$. Let us prove that the function θ_m defined in (14) can be represented as in (15). Using the hyperspherical change of variables

$$\begin{aligned}
\xi_1 &= r\cos(\phi_1) \\
\xi_2 &= r\sin(\phi_1)\cos(\phi_2) \\
&\vdots \\
\xi_{m-1} &= r\sin(\phi_1)\ldots\sin(\phi_{m-2})\cos(\phi_{m-1}) \\
\xi_m &= r\sin(\phi_1)\ldots\sin(\phi_{m-2})\sin(\phi_{m-1}),
\end{aligned}$$

with $\phi_j \in [0, \pi]$, $j = 1, \ldots, m-2$, $\phi_{m-1} \in [0, 2\pi]$, and with the corresponding change of measure

$$d\xi = r^{m-1}\sin^{m-2}(\phi_1)\ldots\sin(\phi_{m-2})d\phi_1\ldots d\phi_{m-1},$$

we obtain

$$\begin{aligned}\theta_m(\gamma) &= (2\pi)^{-m/2}\int_{\mathbb{R}^m}\left(\xi_1^2 - \frac{\gamma}{m-1}\sum_{j=2}^m \xi_j^2\right)_+ e^{-\|\xi\|_2^2/2}d\xi \\ &= I_m \cdot J_m(\gamma),\end{aligned}$$

where I_m is some constant, independent of γ, and

$$J_m(\gamma) := \int_0^\pi \left(\cos^2(\phi_1) - \frac{\gamma}{m-1}\sin^2(\phi_1)\right)_+ \sin^{m-2}(\phi_1)d\phi_1.$$

Since $\theta_m(0) = 1$, we have $I_m = 1/J_m(0)$. Exploiting symmetry to reduce the integration interval from $[0,\pi]$ to $[0,\pi/2]$, proves the formula (15).

B A bound on θ_m

The bound stems from the identity $a_+ = (a+|a|)/2$, valid for every $a \in \mathbb{R}$, and the following result, found in the proof of Theorem 2.1 of [1]:

$$\forall\, y \in \mathbb{R}^m \ : \ \mathbf{E}\left|\sum_{i=1}^m y_i \xi_i^2\right| \geq \frac{2}{\pi}\|y\|_2.$$