



Sparse PSD approximation of the PSD cone

Grigoriy Blekherman¹ · Santanu S. Dey¹ · Marco Molinaro² · Shengding Sun¹ 

Received: 7 February 2020 / Accepted: 7 October 2020 / Published online: 16 October 2020
© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2020

Abstract

While semidefinite programming (SDP) problems are polynomially solvable in theory, it is often difficult to solve large SDP instances in practice. One technique to address this issue is to relax the global positive-semidefiniteness (PSD) constraint and only enforce PSD-ness on smaller $k \times k$ principal submatrices—we call this the *sparse SDP relaxation*. Surprisingly, it has been observed empirically that in some cases this approach appears to produce bounds that are close to the optimal objective function value of the original SDP. In this paper, we formally attempt to compare the strength of the sparse SDP relaxation vis-à-vis the original SDP from a theoretical perspective. In order to simplify the question, we arrive at a data independent version of it, where we compare the sizes of SDP cone and the k -PSD closure, which is the cone of matrices where PSD-ness is enforced on all $k \times k$ principal submatrices. In particular, we investigate the question of how far a matrix of unit Frobenius norm in the k -PSD closure can be from the SDP cone. We provide two incomparable upper bounds on this farthest distance as a function of k and n . We also provide matching lower bounds, which show that the upper bounds are tight within a constant in different regimes of k and n . Other than linear algebra techniques, we extensively use probabilistic methods to arrive at these bounds. One of the lower bounds is obtained by observing a connection between matrices in the k -PSD closure and matrices satisfying the restricted isometry property.

Mathematics Subject Classification 15A42 · 15B48 · 15B57 · 90C22

✉ Shengding Sun
ssun313@gatech.edu

¹ Georgia Institute of Technology, Atlanta, USA

² PUC-Rio, Rio de Janeiro, Brazil

1 Introduction

1.1 Motivation

Semidefinite programming (SDP) relaxations are an important tool to provide dual bounds for many discrete and continuous non-convex optimization problems [35]. These SDP relaxations have the form

$$\begin{aligned} \min \quad & \langle C, X \rangle \\ \text{s.t.} \quad & \langle A^i, X \rangle \leq b_i \quad \forall i \in \{1, \dots, m\} \\ & X \in \mathcal{S}_+^n, \end{aligned} \tag{1}$$

where C and the A^i 's are $n \times n$ matrices, $\langle M, N \rangle := \sum_{i,j} M_{ij}N_{ij}$, and \mathcal{S}_+^n denotes the cone of $n \times n$ symmetric positive semidefinite (PSD) matrices:

$$\mathcal{S}_+^n = \{X \in \mathbb{R}^{n \times n} \mid X = X^T, x^T X x \geq 0, \forall x \in \mathbb{R}^n\}.$$

In practice, it is often computationally challenging to solve large-scale instances of SDPs due to the global PSD constraint $X \in \mathcal{S}_+^n$. One technique to address this issue is to consider a further relaxation that replaces the PSD cone by a larger one $\mathcal{S} \supseteq \mathcal{S}_+^n$. In particular, one can enforce PSD-ness on (some or all) smaller $k \times k$ principal submatrices of X , i.e., we consider the problem

$$\begin{aligned} \min \quad & \langle C, X \rangle \\ \text{s.t.} \quad & \langle A^i, X \rangle \leq b_i \quad \forall i \in \{1, \dots, m\} \\ & \text{selected } k \times k \text{ principal submatrices of } X \in \mathcal{S}_+^k. \end{aligned} \tag{2}$$

We call such a relaxation the *sparse SDP relaxation*.

One reason why these relaxations may be solved more efficiently in practice is that we can enforce PSD constraints by iteratively separating linear constraints. Enforcing PSD-ness on smaller $k \times k$ principal submatrices leads to linear constraints that are sparser, an important property leveraged by linear programming solvers that greatly improves their efficiency [3,7,17,29,32]. This is an important motivation for using sparse SDP relaxations [4,18,28]. (This is also the motivation for studying approximations of polytopes [20], convex hulls of integer linear programs [19,21,32], and integer programming formulations [22] by sparse linear inequalities.) This is our reason for calling the relaxation obtained by enforcing the SDP constraints on smaller $k \times k$ principal submatrices of X as the sparse SDP relaxation.

It has been observed that sparse SDP relaxations not only can be solved much more efficiently in practice, but in some cases they produce bounds that are close to the optimal value of the original SDP. See [4,18,28] for successful applications of this technique for solving box quadratic programming instances, and [26,30] for solving the optimal power flow problem in power systems.

Despite their computational success, theoretical understanding of sparse SDP relaxations remains quite limited. In this paper, we initiate such theoretical investigation. Ideally we would like to compare the objective function values of (1) and (2), but this appears to be a very challenging problem. Therefore, we consider a simpler data-independent question, where we ignore the data of the SDP and the particular selected principal submatrices, to arrive at the following:

How close to the PSD cone \mathcal{S}_+^n do we get when we only enforce PSD-ness on $k \times k$ principal submatrices?

To formalize this question, we begin by defining the *k-PSD closure*, namely matrices that satisfy all $k \times k$ principal submatrices PSD constraints.

Definition 1 (*k-PSD closure*) Given positive integers n and k where $2 \leq k \leq n$, the k -PSD closure $\mathcal{S}^{n,k}$ is the set of all $n \times n$ symmetric real matrices where all $k \times k$ principal submatrices are PSD.

It is clear that the k -PSD closure is a relaxation of the PSD cone (i.e., $\mathcal{S}^{n,k} \supseteq \mathcal{S}_+^n$ for all $2 \leq k \leq n$) and is an increasingly better approximation as the parameter k increases, i.e., we enforce that larger chunks of the matrix are PSD (in particular $\mathcal{S}^{n,n} = \mathcal{S}_+^n$). The SOCP relaxation formulated in [30] is equivalent to using the k -PSD closure with $k = 2$ to approximate the PSD cone. Our definition is a generalization of this construction.

It is worth noting that the dual cone of $\mathcal{S}^{n,k}$ is the set of symmetric matrices with factor width k , defined and studied in [8,27]. In particular, the set of symmetric matrices with factor width 2 is the set of scaled diagonally dominant matrices [8,33], i.e., symmetric matrices A such that DAD is diagonally dominant for some positive diagonal matrix D . Note that [2] uses scaled diagonally dominant for constructing inner approximation of the SDP cones for use in solving polynomial optimization problems.

1.2 Problem setup

We are interested in understanding how well the k -PSD closure approximates the PSD cone for the different values of k and n . To measure this approximation we would like to consider the matrix in the k -PSD closure that is farthest from the PSD cone. We need to make two choices here: the norm to measure this distance and a normalization method (since otherwise there is no upper bound on the distance between matrices in the PSD cone and the k -PSD closure).

We will use the Frobenius norm $\|\cdot\|_F$ for both purposes. That is, the distance between a matrix M and the PSD cone is measured as $\text{dist}_F(M, \mathcal{S}_+^n) = \inf_{N \in \mathcal{S}_+^n} \|M - N\|_F$, and we restrict our attention to matrices in k -PSD closure with Frobenius norm equal to 1. Thus we arrive at the (*normalized*) *Frobenius distance* between the k -PSD closure and the PSD cone, namely the largest distance between a unit-norm matrix M in $\mathcal{S}^{n,k}$ and the cone \mathcal{S}_+^n :

$$\begin{aligned} \overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n) &= \sup_{M \in \mathcal{S}^{n,k}, \|M\|_F=1} \text{dist}_F(M, \mathcal{S}_+^n) \\ &= \sup_{M \in \mathcal{S}^{n,k}, \|M\|_F=1} \inf_{N \in \mathcal{S}_+^n} \|M - N\|_F. \end{aligned}$$

Note that since the origin belongs to \mathcal{S}_+^n this distance is at most 1.

The rest of the paper is organized as follows: Sect. 2 presents all our results and Sect. 3 concludes with some open questions. Then Sect. 4 presents additional notation and background results needed for proving the main results. The remaining sections present the proofs of the main results.

2 Our results

In order to understand how well the k -PSD closure approximates the PSD cone we present:

- Matching upper and lower bounds on $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n)$ for different regimes of k .
- Show that a polynomial number of $k \times k$ PSD constraints are sufficient to provide a good approximation (in Frobenius distance) to the full k -PSD closure (which has $\binom{n}{k} \approx (\frac{en}{k})^k$ such constraints).

We present these result in more details in the following subsections.

2.1 Upper bounds

First we show that the distance between the k -PSD closure and the SDP cone is at most roughly $\approx \frac{n-k}{n}$. In particular, this bound approximately goes from 1 to 0 as the parameter k goes from 2 to n , as expected.

Theorem 1 *For all $2 \leq k < n$ we have*

$$\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n) \leq \frac{n - k}{n + k - 2}. \tag{3}$$

The idea for obtaining this upper bound is the following: given any matrix M in the k -PSD closure $\mathcal{S}^{n,k}$, we construct a PSD matrix \tilde{M} by taking the average of the (PSD) matrices obtained by zeroing out all entries of M but those in a $k \times k$ principal submatrix; the distance between M and \tilde{M} provides an upper bound on $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n)$. The proof of Theorem 1 is provided in Sect. 5.

It appears that for k close to n this upper bound is not tight. In particular, our next upper bound is of the form $(\frac{n-k}{n})^{3/2}$, showing that the gap between the k -PSD closure and the PSD cone goes to 0 as $n - k \rightarrow n$ at a faster rate than that prescribed by the previous theorem. In particular, for $k = n - c$ for a constant c , Theorem 1 gives an upper bound of $O(\frac{1}{n})$ whereas the next Theorem gives an improved upper bound of $O(\frac{1}{n^{3/2}})$.

Theorem 2 Assume $n \geq 97$ and $k \geq \frac{3n}{4}$. Then

$$\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n) \leq 96 \left(\frac{n-k}{n}\right)^{3/2}. \tag{4}$$

It is easy to verify that for sufficiently large r if $k > rn$, then the upper bound given by Theorem 2 dominates the upper bound given by Theorem 1.

The proof of Theorem 2 is more involved than that of Theorem 1. The high-level idea is the following: Using Cauchy’s Interlace Theorem for eigenvalues of hermitian matrices, we first verify that every matrix in $\mathcal{S}^{n,k}$ has at most $n - k$ negative eigenvalues. Since the PSD cone consists of symmetric matrices with non-negative eigenvalues, it is now straightforward to see that the distance from a unit-norm matrix $M \in \mathcal{S}^{n,k}$ to \mathcal{S}_+^n is upper bounded by the absolute value of the most negative eigenvalue of M times $\sqrt{n - k}$. To bound a negative eigenvalue $-\lambda$ of M (where $\lambda \geq 0$), we consider an associated eigenvector $v \in \mathbb{R}^n$ and randomly sparsify it to obtain a random vector V that has at most k non-zero entries. By construction we ensure that $V \approx v$, and that V remains almost orthogonal to all other eigenvectors of M . This guarantees that $V^T M V \approx -\lambda +$ “small error”. On the other hand, since only k entries of V are non-zero, it guarantees that $V^T M V$ only depends on a $k \times k$ submatrix of M , which is PSD by the definition of the k -PSD closure; thus, we have $V^T M V \geq 0$. Combining these observations we get that $\lambda \leq$ “small error”. This eigenvalue bound is used to upper bound the distance from M to the PSD cone. A proof of Theorem 2 is provided in Sect. 6.

We briefly comment on extensions of the results of Theorems 1 and 2 when the distance is measured in other norms. More generally, consider:

$$\overline{\text{dist}}_{\mathcal{N}_1, \mathcal{N}_2}(\mathcal{S}^{n,k}, \mathcal{S}_+^n) := \sup_{M \in \mathcal{S}^{n,k}, \|M\|_{\mathcal{N}_1} = 1} \inf_{N \in \mathcal{S}_+^n} \|M - N\|_{\mathcal{N}_2},$$

where $\|\cdot\|_{\mathcal{N}_1}$ and $\|\cdot\|_{\mathcal{N}_2}$ are two given matrix norms. We call \mathcal{N}_1 the normalizing norm and \mathcal{N}_2 the distance-measuring norm.

The proof of Theorem 1 uses two properties of the norms: (i) the normalizing norm and the distance-measuring norm are the same and (ii) this norm does not change if all off-diagonal terms of the matrix are negated. A class of matrix norms where (ii) holds is element-wise norms, i.e., a norm of the form $\|A\| = \left(\sum_{i,j} |A_{ij}|^p\right)^{\frac{1}{p}}$. Thus, Theorem 1 may be more generally stated as:

$$\overline{\text{dist}}_{\mathcal{N}, \mathcal{N}}(\mathcal{S}^{n,k}, \mathcal{S}_+^n) \leq \frac{n-k}{n+k-2},$$

where \mathcal{N} is an element-wise norm.

The proof of Theorem 2, as discussed above, uses two key arguments: (i) the number of negative eigenvalues can be bounded by $n - k$ for matrices in $\mathcal{S}^{n,k}$ and (ii) the most negative eigenvalue of matrices in $\mathcal{S}^{n,k}$ with Frobenius norm 1 is at most $\frac{96(n-k)}{n^{3/2}}$ in absolute value. Therefore, as pointed out by an anonymous reviewer, Theorem 2 easily generalizes to the case where the distance-measuring norm is any other unitarily invariant matrix norm (thus, a norm depending entirely on the eigenvalues). As an

example, consider the Schatten p -norm $\|A\|_{S_p} := (\sum_{i=1}^n \sigma_i(A)^p)^{\frac{1}{p}}$ where $\sigma_i(A)$ is the i^{th} singular value of A . Then, we obtain the following result:

$$\overline{\text{dist}}_{F,S_p}(\mathcal{S}^{n,k}, \mathcal{S}_+^n) \leq \frac{96(n-k)(n-k)^{1/p}}{n^{3/2}}, \quad \text{for all } k \geq \frac{3}{4}n, n \geq 97.$$

2.2 Lower bounds

We next provide lower bounds on $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n)$ that show that the upper bounds presented in Sect. 2.1 are tight for various regimes of k . The first lower bound, presented in the next theorem, is obtained by a simple construction of an explicit matrix in the k -PSD closure that is far from being PSD. Its proof is provided in Sect. 7.

Theorem 3 *For all $2 \leq k < n$, we have*

$$\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n) \geq \frac{n-k}{\sqrt{(k-1)^2 n + n(n-1)}}. \tag{5}$$

Notice that for small values of k the above lower bound is approximately $\approx \frac{n-k}{n}$ which matches the upper bound from Theorem 1. For very large values of k , i.e. $k = n - c$ for a constant c , the above lower bound is approximately $\approx \frac{c}{n^{3/2}}$ which matches the upper bound by Theorem 2.

Now consider the regime where k is a constant fraction of n . While our upper bounds give $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n) = O(1)$, Theorem 3 only shows that this distance is at least $\Omega(\frac{1}{\sqrt{n}})$, leaving open the possibility that the k -PSD closure provides a sublinear approximation of the PSD cone in this regime. Unfortunately, our next lower bound shows that this is not that case: the upper bounds are tight (up to a constant) in this regime.

Theorem 4 *Fix a constant $r < \frac{1}{93}$ and let $k = rn$. Then for all $k \geq 2$,*

$$\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n) > \frac{\sqrt{r-93r^2}}{\sqrt{162r+3}},$$

which is independent of n .

For this construction we establish a connection with the *Restricted Isometry Property* (RIP) [12,13], a very important notion in signal processing and recovery [14,15]. Roughly speaking, these are matrices that approximately preserve the ℓ_2 norm of sparse vectors. The details of this connection and the proof of Theorem 4 are provided in Sect. 8.

2.3 Achieving the strength of $\mathcal{S}^{n,k}$ by a polynomial number of PSD constraints

In practice one is unlikely to use the full k -PSD closure, since it involves enforcing the PSD-ness for all $\binom{n}{k} \approx \left(\frac{en}{k}\right)^k$ principal submatrices. Is it possible to achieve the upper bounds mentioned above while enforcing PSD-ness on fewer principal submatrices?

2.3.1 Deterministic approach

For a subset $I \subseteq [n]$ of size k and a symmetric matrix M , let M^I be the matrix where we zero out all the rows and columns of M except the ones belonging to I . As we discuss after the proof of Theorem 1, the bound in (3) is obtained by estimating the distance between a given matrix M in $\mathcal{S}^{n,k}$ and the PSD matrix \bar{M} , which is the average of $\binom{n}{k}$ (PSD) matrices M^I corresponding to every k -element subset I of $[n]$. Now suppose that we have a collection of $k \times k$ principal submatrices (and the corresponding collection I of k -subsets of $[n]$) with the property that *every off-diagonal entry appears in exactly the same number of principal submatrices in the collection*. It is then straightforward to show that the average of the matrices M^I , where I is in the above collection, is the same as the average \bar{M} taken over all $\binom{n}{k}$ principal submatrices. In other words, the upper bound (3) can be achieved by any collection of $k \times k$ principal submatrices with the above property.

This property is formally captured by the notion of 2-designs: Recall that a collection \mathcal{D} of k -subsets of $[n]$ (called blocks) is called a 2-design (also called a balanced incomplete block design or BIBD) if every pair of elements in $[n]$ belongs to the same number of blocks, denoted λ . It follows that every element of $[n]$ belongs to the same number of blocks, denoted r . Let b be the total number of blocks. From the discussion above, b principal submatrices corresponding to a 2-design is sufficient to give a bound of (3). For background on block designs we refer to [31], (Chapters 1 and 2).

It is known from the work of Wilson [34],(Corollary A and B) that, a 2-design with $b = n(n - 1)$ exists for all sufficiently large values of n , although to the best of our knowledge no explicit construction is known. (Wilson’s theorem gives a much more general statement for existence of 2-designs). Therefore, for almost all n we can achieve the strength of bound (3) while only using $n(n - 1)$ submatrices. Fisher’s inequality states that $b \geq n$, so we need to enforce PSD-ness of at least n minors if we use a 2-design. A 2-design is called *symmetric* if $b = n$. Bruck-Ryser-Chowla Theorem [10,16] gives necessary conditions on b, k and λ , for which a symmetric 2-designs exist, and this is certainly a limited set of parameters. Nevertheless, symmetric 2-designs may be of use in practice, as they give us the full strength of (3) while enforcing PSD-ness of only $n k \times k$ submatrices. Some important examples of symmetric 2-designs are finite projective planes (symmetric 2-designs with $\lambda = 1$), biplanes ($\lambda = 2$) and Hadamard 2-designs.

2.3.2 Randomized approach

Another way to achieve the bound in (3) is to randomly select $k \times k$ principal submatrices. We show that the upper bound given by (3) can also be achieved within factor

$1 + \epsilon$ and probability at least $1 - \delta$ by randomly sampling $O\left(\frac{n^2}{\epsilon^2} \ln \frac{n}{\delta}\right)$ of the $k \times k$ principal submatrices.

Theorem 5 *Let \mathcal{S}^n denote the set of real symmetric matrices. Let $2 \leq k \leq n - 1$. Consider $\epsilon, \delta > 0$ and let*

$$m := \frac{12n(n - 1)^2}{\epsilon^2(n - k)^2k} \ln \frac{2n^2}{\delta} \in O\left(\frac{n^2}{\epsilon^2} \ln \frac{n}{\delta}\right).$$

Let $\mathcal{I} = (I_1, \dots, I_m)$ be a sequence of random k -sets independently uniformly sampled from all subsets of $\{1, \dots, n\}$ of size k , and define $\mathcal{S}_{\mathcal{I}}$ as the set of symmetric matrices satisfying the PSD constraints for the principal submatrices indexed by the I_i 's, namely

$$\mathcal{S}_{\mathcal{I}} := \{M \in \mathcal{S}^n : M_{I_i} \text{ is PSD, } \forall i \in [m]\}.$$

Then with probability at least $1 - \delta$ we have

$$\overline{\text{dist}}_F(\mathcal{S}_{\mathcal{I}}, \mathcal{S}_+^n) \leq (1 + \epsilon) \frac{n - k}{n + k - 2}.$$

Remark 1 Since the zero matrix is PSD, by definition we always have $\overline{\text{dist}}_F(\mathcal{S}_{\mathcal{I}}, \mathcal{S}_+^n) \leq 1$. So in order for the bound given by Theorem 5 to be of interest, we need $(1 + \epsilon) \frac{n - k}{n + k - 2} \leq 1$, which means $\epsilon \leq \frac{2k - 2}{n - k}$. Plugging this into m , we see that we need at least $\frac{3n(n - 1)^2}{k(k - 1)^2} \ln \frac{2n^2}{\delta} = \tilde{O}\left(\frac{n^3}{k^3}\right)$ samples to obtain a nontrivial upper bound on the distance.

3 Conclusion and open questions

In this paper, we have been able to provide various upper and lower bounds on $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n)$. In two regimes our bounds on $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n)$ are quite tight. These are: (i) k is small, i.e., $2 \leq k \leq \sqrt{n}$ and (ii) k is quite large, i.e., $k = n - c$ where c is a constant. These are shown in the first two rows of Table 1. When k/n is a constant, we have also established upper and lower bounds on $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n)$ that are independent of n . However, our upper and lower bounds are not quite close when viewed as a function of the ratio k/n . Improving these bounds as a function of this ratio is an important open question.

We also showed that instead of selecting all minors, only a polynomial number of randomly selected minors realizes upper bound (3) within factor $1 + \epsilon$ with high probability. An important question in this direction is to deterministically and strategically determine principal submatrices to impose PSD-ness, so as to obtain the best possible bound for (2) As discussed earlier, such questions are related to exploring 2-designs and perhaps further generalizations of results presented in [25].

Table 1 Bounds on $\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n)$ for some regimes

Regime	Upper bound	Lower bound
(small k) $2 \leq k \leq \sqrt{n}$	$\frac{n-k}{n}$ (Simplified from Thm 1)	$\frac{1}{\sqrt{2}} \frac{n-k}{n}$ (Simplified from Thm 3)
(large k) $k \geq n - c$ ($n \geq 97, k \geq 0.75n$)	$96 \left(\frac{c}{n}\right)^{3/2}$ (Simplified from Thm 2)	$\frac{1}{\sqrt{2}} \frac{c}{n^{3/2}}$ (Simplified from Thm 3)
(k/n is a constant) $k = rn$	Constant, independent of n	Constant, independent of n
($r < \frac{1}{93}$)	$1 - r$ (Simplified from Thm 1)	$\sqrt{\frac{r-93r^2}{5}}$ (Simplified from Thm 4)

4 Notation and preliminaries

The *support* of a vector is the set of its non-zero coordinates, and we call a vector *k-sparse* if its support has size at most k . We will use $[n]$ to denote the set $\{1, \dots, n\}$. A *k-set* of a set A is a subset $B \subset A$ with $|B| = k$. We also use $\binom{[n]}{k}$ to denote the set of all k -sets of $[n]$. Given any vector $x \in \mathbb{R}^n$ and a k -set $J \subset [n]$ we define $x_J \in \mathbb{R}^k$ as the vector where we remove the coordinates whose indices are not in J . Similarly, for a matrix $M \in \mathbb{R}^{n \times n}$ and a k -set $J \subset [n]$, we denote the principal submatrix of M corresponding to the rows and columns in J by M_J .

4.1 Linear algebra

Given any $n \times n$ matrix $A = [a_{ij}]$ its trace (the sum of its diagonal entries) is denoted as $\text{Tr}(A)$. Recall that $\text{Tr}(A)$ is also equal to the sum of all eigenvalues of A , counting multiplicities. Given a symmetric matrix A , we use $\lambda_1(A) \geq \lambda_2(A) \geq \dots$ to denote its eigenvalues in non-increasing order.

We remind the reader that, a real symmetric $n \times n$ matrix M is said to be PSD if $x^T M x \geq 0$ for all $x \in \mathbb{R}^n$, or equivalently all of its eigenvalues are non-negative. We also use the notation that $A \succeq B$ if $A - B$ is PSD.

We next present the famous Cauchy’s Interlace Theorem which will be important for obtaining an upper bound on the number of negative eigenvalues of matrices in $\mathcal{S}^{n,k}$. A proof can be found in [23].

Theorem 6 (Cauchy’s Interlace Theorem) *Consider an $n \times n$ symmetric matrix A and let A_J be any of its $k \times k$ principal submatrix. Then for all $1 \leq i \leq k$,*

$$\lambda_{n-k+i}(A) \leq \lambda_i(A_J) \leq \lambda_i(A).$$

4.2 Probability

These following concentration inequalities will be used throughout, and can be found in [9].

Theorem 7 (Markov’s inequality) *Let X be a non-negative random variable. Then for all $a \geq 1$,*

$$\Pr(X \geq a\mathbb{E}(X)) \leq \frac{1}{a}.$$

Theorem 8 (Chebyshev’s inequality) *Let X be a random variable with finite mean and variance. Then for all $a > 0$,*

$$\Pr(|X - \mathbb{E}(X)| \geq a) \leq \frac{\text{Var}(X)}{a^2}.$$

Theorem 9 (Chernoff bound) *Let X_1, \dots, X_n be i.i.d. Bernoulli random variables, with $\Pr(X_i = 1) = \mathbb{E}(X_i) = p$ for all i . Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}(X) = np$. Then for any $0 < \delta < 1$,*

$$\Pr(|X - \mu| > \delta\mu) \leq 2 \exp\left(-\frac{\mu\delta^2}{3}\right).$$

5 Proof of Theorem 1: averaging operator

Consider a matrix M in the k -PSD closure $\mathfrak{S}^{n,k}$ with $\|M\|_F = 1$. To upper bound its distance to the PSD cone we transform M into a “close by” PSD matrix \tilde{M} .

The idea is clear: since all $k \times k$ principal submatrices of M are PSD, we define \tilde{M} as the average of these minors. More precisely, for a set $I \subseteq [n]$ of k indices, let M^I be the matrix where we zero out all the rows and columns of M except those indexed by indices in I ; then \tilde{M} is the average of all such matrices:

$$\tilde{M} := \frac{1}{\binom{n}{k}} \sum_{I \in \binom{[n]}{k}} M^I.$$

Notice that indeed since the principal submatrix M_I is PSD, M^I is PSD as well: for all vectors $x \in \mathbb{R}^n$, $x^\top M^I x = x_I^\top M_I x_I \geq 0$. Since the average of PSD matrices is also PSD, we have that \tilde{M} is PSD, as desired.

Moreover, notice that the entries of \tilde{M} are just scalings of the entries of M , depending on how many terms of the average are not zeroed out:

1. **Diagonal terms:** These are scaled by the factor

$$\frac{\binom{n}{k} - \binom{n-1}{k}}{\binom{n}{k}} = \frac{k}{n},$$

that is, $\tilde{M}_{ii} = \frac{k}{n} M_{ii}$ for all $i \in [n]$.

2. **Off-diagonal terms:** These are scaled by the factor

$$\frac{\binom{n}{k} - (2\binom{n-1}{k} - \binom{n-2}{k})}{\binom{n}{k}} = \frac{k(k-1)}{n(n-1)},$$

that is, $\tilde{M}_{ij} = \frac{k(k-1)}{n(n-1)}M_{ij}$ for all $i \neq j$.

To even out these factors, we define the scaling $\alpha := \frac{2n(n-1)}{k(n+k-2)}$ and consider $\alpha\tilde{M}$. Now we have that the difference between M and $\alpha\tilde{M}$ is a uniform scaling (up to sign) of M itself: $(M - \alpha\tilde{M})_{ii} = (1 - \alpha\frac{k}{n})M_{ii} = -\frac{n-k}{n+k-2}M_{ii}$, and $(M - \alpha\tilde{M})_{ij} = (1 - \alpha\frac{k(k-1)}{n(n-1)})M_{ij} = \frac{n-k}{n+k-2}M_{ij}$ for $i \neq j$. Therefore, we have

$$\text{dist}_F(M, S_+^n) \leq \|M - \alpha\tilde{M}\|_F = \frac{n-k}{n+k-2} \|M\|_F = \frac{n-k}{n+k-2}.$$

Since this holds for all unit-norm matrix $M \in \mathfrak{S}^{n,k}$, this upper bound also holds for $\overline{\text{dist}}_F(\mathfrak{S}^{n,k}, S_+^n)$. This concludes the proof of Theorem 1.

6 Proof of Theorem 2: randomized sparsification

Let $M \in \mathfrak{S}^{n,k}$ be a matrix in the k -PSD closure with $\|M\|_F = 1$. To prove Theorem 2, we show that the Frobenius distance from M to the PSD cone is at most $O((\frac{n-k}{n})^{3/2})$. We assume that M is not PSD, otherwise we are done, and hence it has a negative eigenvalue. We write M in terms of its eigendecomposition: Let $-\lambda_1 \leq -\lambda_2 \leq \dots \leq -\lambda_\ell$ and $\mu_1, \dots, \mu_{n-\ell}$ be the negative and non-negative eigenvalues of M , and let $v^1, \dots, v^\ell \in \mathbb{R}^n$ and $w^1, \dots, w^{n-\ell} \in \mathbb{R}^n$ be orthonormal eigenvectors relative to these eigenvalues. Thus

$$M = -\sum_{i \leq \ell} \lambda_i v^i (v^i)^\top + \sum_{i \leq n-\ell} \mu_i w^i (w^i)^\top. \tag{6}$$

Notice that since $\|M\|_F = 1$ we have

$$\sum_{i \leq \ell} \lambda_i^2 + \sum_{i \leq n-\ell} \mu_i^2 = 1. \tag{7}$$

We first relate the distance from M to the PSD cone to its negative eigenvalues.

6.1 Distance to PSD cone and negative eigenvalues

We start with the following general observation.

Proposition 1 Suppose M is a symmetric $n \times n$ matrix with $\ell \leq n$ negative eigenvalues. Let $-\lambda_1 \leq -\lambda_2 \leq \dots \leq -\lambda_\ell < 0$ and $\mu_1, \dots, \mu_{n-\ell} \geq 0$ be the negative and non-negative eigenvalues of M . Then

$$\text{dist}_F(M, \mathcal{S}_+^n) = \sqrt{\sum_{i=1}^{\ell} \lambda_i^2}.$$

Proof Let V be the orthonormal matrix that diagonalizes M , i.e.,

$$V^T M V = D := \text{diag}(-\lambda_1, \dots, -\lambda_\ell, \mu_1, \dots, \mu_{n-\ell}).$$

It is well-known that the Frobenius norm is invariant under orthonormal transformation. Therefore, for any $N \in \mathcal{S}_+^n$ we have

$$\text{dist}_F(M, N) = \|M - N\|_F = \|V^T(M - N)V\|_F = \text{dist}_F(D, V^T N V).$$

Since $N \in \mathcal{S}_+^n$ iff $V^T N V \in \mathcal{S}_+^n$, we see that $\text{dist}_F(M, \mathcal{S}_+^n) = \text{dist}_F(D, \mathcal{S}_+^n)$. So we only need to show that the latter is $\sqrt{\sum_{i=1}^{\ell} \lambda_i^2}$.

Let $D_+ = \text{diag}(0, \dots, 0, \mu_1, \dots, \mu_{n-\ell})$ be obtained from D by making all negative eigenvalues zero. Then

$$\|D - D_+\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n (D - D_+)_{ij}^2} = \sqrt{\sum_{i=1}^{\ell} \lambda_i^2}.$$

It then suffices to show that D_+ is the PSD matrix closest to D . For that, let N be any PSD matrix. Then $N_{ii} = e_i^T N e_i \geq 0$ for all i , where e_i is the standard unit vector on i^{th} coordinate. Thus we have

$$\|D - N\|_F = \sqrt{\sum_{i=1}^{\ell} (N_{ii} + \lambda_i)^2 + \sum_{i=\ell+1}^n (\mu_{i-\ell} - N_{ii})^2 + \sum_{i=1}^n \sum_{j \neq i} N_{ij}^2} \geq \sqrt{\sum_{i=1}^{\ell} \lambda_i^2}.$$

This concludes the proof. □

In addition, Cauchy’s Interlace Theorem gives an upper bound on the number of negative eigenvalues of matrices in $\mathcal{S}^{n,k}$.

Proposition 2 Any $A \in \mathcal{S}^{n,k}$ has at most $n - k$ negative eigenvalues.

Proof Let J be any k -subset of $[n]$. Since $A \in \mathcal{S}^{n,k}$ we have that A_J is PSD, so in particular $\lambda_k(A_J) \geq 0$. Thus, by Theorem 6 the original matrix A also has $\lambda_k(A) \geq 0$, and so the first k eigenvalues of A are nonnegative. □

Using Propositions 1 and 2, given any symmetric matrix $M \in \mathcal{S}^{n,k}$ we can get an upper bound on $\text{dist}_F(M, \mathcal{S}_+^n)$ using its smallest eigenvalue.

Proposition 3 Consider a matrix $M \in \mathcal{S}^{n,k}$ with smallest eigenvalue $-\lambda_1 < 0$. Then

$$\text{dist}_F(M, \mathcal{S}_+^n) \leq \sqrt{n-k} \cdot \lambda_1.$$

Proof Letting $-\lambda_1, \dots, -\lambda_\ell$ be the negative eigenvalues of M , we have from Proposition 1 that $\text{dist}_F(M, \mathcal{S}_+^n) = \sqrt{\sum_{i=1}^\ell \lambda_i^2} \leq \sqrt{\ell} \cdot \lambda_1$, since $-\lambda_1$ is the smallest eigenvalue. Since $\ell \leq n - k$, because of Proposition 2, we obtain the result. \square

6.2 Upper bounding λ_1

Given the previous proposition, fix throughout this section a (non PSD) matrix $M \in \mathcal{S}^{n,k}$ with smallest eigenvalue $-\lambda_1 < 0$. Our goal is to upper bound λ_1 .

The first observation is the following: Consider a symmetric matrix A and a set of coordinates $I \subseteq [n]$, and notice that for every vector $x \in \mathbb{R}^n$ supported in I we have $x^\top Ax = x_I^\top A_I x_I$. Thus, the principal submatrix A_I is PSD iff for all vectors $x \in \mathbb{R}^n$ supported in I we have $x^\top Ax \geq 0$. Applying this to all principal submatrices gives a characterization of the k -PSD closure via k -sparse test vectors.

Observation 1 A symmetric real matrix A belongs to $\mathcal{S}^{n,k}$ iff for all k -sparse vectors $x \in \mathbb{R}^n$ we have $x^\top Ax \geq 0$.

Using this characterization, and the fact that $M \in \mathcal{S}^{n,k}$, the idea to upper bound λ_1 is to find a vector \bar{v} with the following properties (informally):

1. \bar{v} is k -sparse
2. \bar{v} is similar to the eigenvector v^1 relative to λ_1
3. \bar{v} is almost orthogonal to the eigenvectors of M relative to its non-negative eigenvalues.

Such vector gives a bound on λ_1 because using the eigendecomposition (6)

$$0 \stackrel{\text{Obs(1)}}{\leq} \bar{v}^\top M \bar{v} = - \sum_{i \leq \ell} \lambda_i \langle v^i, \bar{v} \rangle^2 + \sum_{i \leq n-k} \mu_i \langle w^i, \bar{v} \rangle^2 \lesssim -\lambda_1 + \text{“small error”},$$

and hence $\lambda_1 \lesssim \text{“small error”}$.

We show the existence of such k -sparse vector \bar{v} via the probabilistic method by considering a random sparsification of v^1 . More precisely, define the random vector $V \in \mathbb{R}^n$ as follows: in hindsight set $p := 1 - \frac{2(n-k)}{n}$, which is always at least $\frac{1}{2}$ due to our assumption in Theorem 2 that $k \geq \frac{3n}{4}$, and let V have independent entries satisfying

$$V_i = \begin{cases} v_i^1 & \text{if } (v_i^1)^2 > 2/n, \\ \frac{v_i^1}{p} \text{ with probability } p & \text{if } (v_i^1)^2 \leq \frac{2}{n}, \\ 0 \text{ with probability } 1 - p & \text{if } (v_i^1)^2 \leq \frac{2}{n}. \end{cases}$$

The choice of p guarantees that V is k -sparse with good probability.

Lemma 1 V is k -sparse with probability at least $\frac{1}{2}$.

Proof Let m be the number of entries in v^1 with $(v_i^1)^2 \leq \frac{2}{n}$. Since $\|v^1\|_2 = 1$ we have $m \geq \frac{n}{2}$. By the randomized construction, the number of coordinates of value 0 in V is lower bounded by a binomial random variable B with m trials and success probability $1 - p$. Using the definition of p we have the expectation

$$\mathbb{E}B = m(1 - p) \geq \frac{n}{2} \cdot \frac{2(n - k)}{n} = n - k;$$

since $n - k$ is integer we have $\lfloor \mathbb{E}B \rfloor \geq n - k$. Moreover, it is known that the median of a binomial distribution is at least the expectation rounded down to the nearest integer [24], hence $\Pr(B \geq \lfloor \mathbb{E}B \rfloor) \geq \frac{1}{2}$. Chaining these observations we have

$$\Pr(\text{\# of coordinates of value 0 in } V \geq n - k) \geq \Pr(B \geq n - k) \geq \Pr(B \geq \lfloor \mathbb{E}B \rfloor) \geq \frac{1}{2}.$$

In other words, our randomized vector V is k -sparse with probability at least $\frac{1}{2}$. □

Next, we show that with good probability V and v_1 are in a “similar direction”.

Lemma 2 With probability $> 1 - \frac{1}{6}$ we have $\langle V, v^1 \rangle \geq \frac{1}{2}$.

Proof To simplify the notation we use v to denote v^1 . By definition of V , for each coordinate we have $\mathbb{E}[V_i v_i] = v_i^2$, and hence $\mathbb{E}\langle V, v \rangle = \|v\|_2^2 = 1$.

In addition, let I be the set of coordinates i where $v_i^2 \leq \frac{2}{n}$. Then for $i \notin I$ we have $\text{Var}(V_i v_i) = 0$, and for $i \in I$ we have $\text{Var}(V_i v_i) = v_i^2 \text{Var}(V_i) \leq \frac{2}{n} \text{Var}(V_i)$. Moreover, since $p \geq \frac{1}{2}$ (implied by the assumption $k \geq \frac{3n}{4}$) we have by construction $V_i \leq \frac{v_i}{p} \leq 2v_i$, and hence

$$\text{Var}(V_i) \leq \mathbb{E}V_i^2 \leq 2v_i \mathbb{E}V_i = 2v_i^2.$$

So using the independence of the coordinates of V we have

$$\text{Var}\langle V, v \rangle = \sum_{i \in I} \text{Var}(V_i v_i) \leq \frac{4}{n} \sum_i v_i^2 = \frac{4}{n}.$$

Then by Chebyshev’s inequality we obtain that

$$\Pr\left(\langle V, v \rangle \leq \frac{1}{2}\right) \leq \Pr\left(|\langle V, v \rangle - 1| \geq \frac{1}{2}\right) \leq \frac{16}{n}.$$

Since $n \geq 97$, this proves the lemma. □

Finally, we show that V is almost orthogonal to the eigenvectors of M relative to non-negative eigenvalues.

Lemma 3 *With probability $\geq 1 - \frac{1}{3}$ we have $\sum_{i \leq n-\ell} \mu_i \langle V, w^i \rangle^2 \leq \frac{24(n-k)}{n^{3/2}}$.*

Proof Again we use v to denote v^1 . Define the matrix $\bar{M} := \sum_{i \leq n-\ell} \mu_i w_i w_i^\top$, so we want to upper bound $V^\top \bar{M} V$. Moreover, let $\Delta = V - v$; since v and the w_i 's are orthogonal we have $\bar{M} v = 0$ and hence

$$V^\top \bar{M} V = v^\top \bar{M} v + 2\Delta^\top \bar{M} v + \Delta^\top \bar{M} \Delta = \Delta^\top \bar{M} \Delta, \tag{8}$$

so it suffices to upper bound the right-hand side.

For that, notice that Δ has independent entries with the form

$$\Delta_i = \begin{cases} 0 & \text{if } v_i^2 > \frac{2}{n}, \\ \frac{v_i(1-p)}{p} \text{ with probability } p & \text{if } v_i^2 \leq \frac{2}{n}, \\ -v_i \text{ with probability } 1-p & \text{if } v_i^2 \leq \frac{2}{n}. \end{cases}$$

So $\mathbb{E}[\Delta_i \Delta_j] = \mathbb{E}\Delta_i \mathbb{E}\Delta_j = 0$ for all $i \neq j$. In addition $\mathbb{E}\Delta_i^2 = 0$ for indices where $v_i^2 > \frac{2}{n}$, and for other indices

$$\mathbb{E}\Delta_i^2 \leq \frac{v_i^2(1-p)^2}{p} + v_i^2(1-p) = v_i^2 \frac{1-p}{p} \leq \frac{2(1-p)}{np}.$$

Using these we can expand $\mathbb{E}[\Delta^\top \bar{M} \Delta]$ as

$$\begin{aligned} \mathbb{E}[\Delta^\top \bar{M} \Delta] &= \mathbb{E}\left[\sum_{i,j} \bar{M}_{ij} \Delta_i \Delta_j\right] = \sum_{i,j} \bar{M}_{ij} \mathbb{E}[\Delta_i \Delta_j] = \sum_{i=1}^n \bar{M}_{ii} \mathbb{E}\Delta_i^2 \\ &\leq \frac{2(1-p)}{np} \text{Tr}(\bar{M}) \\ &= \frac{4(n-k)}{n^2 p} \text{Tr}(\bar{M}), \end{aligned} \tag{9}$$

where the last equation uses the definition of p .

Since the μ_i 's are the eigenvalues of \bar{M} , we can therefore bound the trace as

$$\begin{aligned} \text{Tr}(\bar{M}) &= \sum_{i \leq n-\ell} \mu_i \leq \sqrt{n-\ell} \cdot \sqrt{\sum_{i \leq n-\ell} \mu_i^2} \\ &\leq \sqrt{n-\ell} \leq \sqrt{n}, \end{aligned}$$

where the first inequality follows from the well-known inequality that $\|u\|_1 \leq \sqrt{n}\|u\|_2$ for all $u \in \mathbb{R}^n$ and the second inequality comes from (7). Further using the assumption that $p \geq \frac{1}{2}$, we get from (9) that

$$\mathbb{E}[\Delta^\top \bar{M} \Delta] \leq \frac{8(n-k)}{n^{3/2}}.$$

Finally, since all the eigenvalues μ_i of \overline{M} are non-negative, this matrix is PSD and hence the random variable $\Delta^\top \overline{M} \Delta$ is non-negative. Markov’s inequality then gives that

$$\Pr \left(\Delta^\top \overline{M} \Delta \geq \frac{24(n-k)}{n^{3/2}} \right) \leq \Pr \left(\Delta^\top \overline{M} \Delta \geq 3 \mathbb{E}[\Delta^\top \overline{M} \Delta] \right) \leq \frac{1}{3}.$$

This concludes the proof of the lemma. □

With these properties of V we can finally upper bound the absolute value of the most negative eigenvalue of M .

Lemma 4 $\lambda_1 \leq \frac{96(n-k)}{n^{3/2}}$.

Proof We take the union bound over Lemmas 1–3. In other words, the probability that V fails at least one of the properties in above three lemmas is strictly less than $\frac{1}{2} + \frac{1}{6} + \frac{1}{3} = 1$. Therefore, with strictly positive probability V satisfies all these properties. That is, there is a vector $\bar{v} \in \mathbb{R}^n$ that is k -sparse, has $\langle \bar{v}, v^1 \rangle \geq \frac{1}{2}$ and $\sum_{i \leq n-\ell} \mu_i \langle \bar{v}, w^i \rangle^2 \leq \frac{24(n-k)}{n^{3/2}}$. Then using Observation 1 and the eigendecomposition (6)

$$0 \stackrel{\text{Obs(1)}}{\leq} \bar{v}^\top M \bar{v} = - \sum_{i \leq \ell} \lambda_i \langle \bar{v}, v^i \rangle^2 + \sum_{i \leq n-\ell} \mu_i \langle \bar{v}, w^i \rangle^2 \leq -\frac{\lambda_1}{4} + \frac{24(n-k)}{n^{3/2}}.$$

Reorganizing the terms proves the lemma. □

6.3 Concluding the proof of Theorem 2

Plugging the upper bound on λ_1 from Lemma 4 into Proposition 3 we obtain that

$$\text{dist}_F(M, \mathcal{S}_+^n) \leq 96 \left(\frac{n-k}{n} \right)^{3/2}.$$

Since this holds for all unit-norm $M \in \mathfrak{S}^{n,k}$, we have that $\overline{\text{dist}}_F(\mathfrak{S}^{n,k}, \mathcal{S}_+^n)$ also satisfies the same upper bound. This concludes the proof.

7 Proof of Theorem 3: a specific family of matrices in $\mathfrak{S}^{n,k}$

To prove the lower bounds on $\overline{\text{dist}}_F(\mathfrak{S}^{n,k}, \mathcal{S}_+^n)$ we construct specific families of matrices in $\mathfrak{S}^{n,k}$ with Frobenius norm 1, and then lower bound their distance to the PSD cone.

For the first lower bound in Theorem 3, we consider the construction where all diagonal entries are the same, and all off-diagonal ones are also the same. More precisely, given scalars $a, b \geq 0$ we define the matrix

$$G(a, b, n) := (a + b)I_n - a\mathbf{1}\mathbf{1}^\top, \tag{10}$$

where I_n is the $n \times n$ identity matrix, and $\mathbf{1}$ is the column vector with all entries equal to 1. In other words, all diagonal entries of $G(a, b, n)$ are b , and all off-diagonal ones are $-a$.

The parameter a will control how far this matrix is from PSD: for $a = 0$ it is PSD, and if a is much bigger than b it should be ‘‘far’’ from the PSD cone. We then directly compute its eigenvalues, as well as its Frobenius distance to the PSD cone.

Proposition 4 *The eigenvalues of $G(a, b, n)$ are $b - (n - 1)a$ with multiplicity 1, and $b + a$ with multiplicity $n - 1$.*

Proof Let $\{v^1, \dots, v^n\}$ be an orthonormal basis of \mathbb{R}^n such that $\sqrt{n}v^1 = \mathbf{1}$. Then we can rewrite $G(a, b, n)$ as

$$\begin{aligned} G(a, b, n) &= (a + b) \sum_{i=1}^n v^i(v^i)^\top - nav^1(v^1)^\top \\ &= (b - (n - 1)a)v^1(v^1)^\top + (a + b) \sum_{i=2}^n v^i(v^i)^\top. \end{aligned}$$

This gives a spectral decomposition of $G(a, b, n)$, so it has the aforementioned set of eigenvalues. □

The next two corollaries immediately follow from Proposition 4.

Corollary 1 *If $a, b \geq 0$, then $G(a, b, n) \in \mathcal{S}^{n,k}$ iff $b \geq (k - 1)a$. In particular, since $\mathcal{S}^{n,n} = \mathcal{S}_+^n$, $G(a, b, n) \in \mathcal{S}_+^n$ iff $b \geq (n - 1)a$.*

Proof Note that every $k \times k$ principal submatrix of $G(a, b, n)$ is just the matrix $G(a, b, k)$, which belongs to \mathcal{S}_+^k iff $b - (k - 1)a \geq 0$, since $a, b \geq 0$. □

Corollary 2 *If $a, b \geq 0$, then $\text{dist}_F(G(a, b, n), \mathcal{S}_+^n) = \max\{(n - 1)a - b, 0\}$.*

Proof If $b \geq (n - 1)a$, then $G(a, b, n) \in \mathcal{S}_+^n$ from first corollary, so $\text{dist}_F(G(a, b, n), \mathcal{S}_+^n) = 0$ by definition.

If $b < (n - 1)a$, then $G(a, b, n)$ has only one negative eigenvalue $b - (n - 1)a$. Thus using Proposition 1 we get $\text{dist}_F(G(a, b, n), \mathcal{S}_+^n) = (n - 1)a - b$. □

To conclude the proof of Theorem 3, let $\bar{a} = \frac{1}{\sqrt{(k-1)^2n+n(n-1)}}$ and $\bar{b} = (k - 1)\bar{a}$. From Corollary 1 we know that $G(\bar{a}, \bar{b}, n)$ belongs to the k -PSD closure $\mathcal{S}^{n,k}$, and it is easy to check that it has Frobenius norm 1. Then using Corollary 2 we get

$$\overline{\text{dist}}_F(\mathcal{S}^{n,k}, \mathcal{S}_+^n) \geq \text{dist}_F(G(\bar{a}, \bar{b}, n), \mathcal{S}_+^n) = (n - k)\bar{a} = \frac{n - k}{\sqrt{(k - 1)^2n + n(n - 1)}}.$$

This concludes the proof.

8 Proof of Theorem 4: RIP construction when $k = O(n)$

Again, to prove the lower bound $\overline{\text{dist}}_F(\mathfrak{S}^{n,k}, \mathcal{S}_+^n) \geq cst$ for a constant cst we will construct (randomly) a unit-norm matrix M in the k -PSD closure $\mathfrak{S}^{n,k}$ that has distance at least cst from the PSD cone \mathcal{S}_+^n ; we will use its negative eigenvalues to assess this distance, via Proposition 1. Motivation for connection with RIP property. Before presenting the actual construction, we give the high-level idea of how the RIP property (Definition 2 below) fits into the picture. For simplicity, assume $k = n/2$. (The actual proof will not have this value of k). The idea is to construct a matrix M where about half of its eigenvalues take the negative value $-\frac{1}{\sqrt{n}}$, with orthonormal eigenvectors $v^1, v^2, \dots, v^{n/2}$, and rest take a positive value $\frac{1}{\sqrt{n}}$, with orthonormal eigenvectors $w^1, w^2, \dots, w^{n/2}$. This normalization makes $\|M\|_F = \Theta(1)$, so the reader can just think of M being unit-norm, as desired. In addition, from Proposition 1 this matrix is far from the PSD cone: $\text{dist}_F(M, \mathcal{S}_+^n) \gtrsim \sqrt{\left(\frac{1}{\sqrt{n}}\right)^2 \cdot \frac{n}{2}} = cst$. So we only need to guarantee that M belongs to the k -PSD closure; for that we need to carefully choose its positive eigenspace, namely the eigenvectors $w^1, w^2, \dots, w^{n/2}$.

Recall that from Observation 1, M belongs to the k -PSD closure iff $x^\top Mx$ for all k -sparse vectors $x \in \mathbb{R}^n$. Letting V be the matrix with rows v^1, v^2, \dots , and W the matrix with rows w^1, w^2, \dots , the quadratic form $x^\top Mx$ is

$$x^\top Mx = -\frac{1}{\sqrt{n}} \sum_i \langle v^i, x \rangle^2 + \frac{1}{\sqrt{n}} \sum_i \langle w^i, x \rangle^2 = -\frac{1}{\sqrt{n}} \|Vx\|_2^2 + \frac{1}{\sqrt{n}} \|Wx\|_2^2.$$

Since the rows of V are orthonormal we have $\|Vx\|_2^2 \leq \|x\|_2^2$. Therefore, if we could construct the matrix W so that for all k -sparse vectors $x \in \mathbb{R}^n$ we had $\|Wx\|_2^2 \approx \|x\|_2^2$, we would be in good shape, since we would have

$$x^\top Mx \gtrsim -\frac{1}{\sqrt{n}} \|x\|_2^2 + \frac{1}{\sqrt{n}} \|x\|_2^2 \gtrsim 0 \quad \text{for all } k\text{-sparse vectors } x, \quad (11)$$

thus M would be (approximately) in the k -PSD closure. This approximate preservation of norms of sparse vectors is precisely the notion of the *Restricted Isometry Property* (RIP) [12,13].

Definition 2 (RIP) Given $k < m < n$, an $m \times n$ matrix A is said to be (k, δ) -RIP if for all k -sparse vectors $x \in \mathbb{R}^n$, we have

$$(1 - \delta)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta)\|x\|_2^2.$$

This definition is very important in signal processing and recovery [12–15], and there has been much effort trying to construct deterministic [5,11] or randomized [6] matrices satisfying given RIP guarantees.

The following theorem in [6] provides a probabilistic guarantee for a random Bernoulli matrix to have the RIP.

Theorem 10 ((4.3) and (5.1) in [6], see also [1]) *Let A be an $m \times n$ matrix where each entry is independently $\pm 1/\sqrt{m}$ with probability $1/2$. Then A is (k, δ) -RIP with probability at least*

$$1 - 2 \left(\frac{12}{\delta} \right)^k e^{-(\delta^2/16 - \delta^3/48)m}. \tag{12}$$

Proof of Theorem 4 After we have observed the above connection between matrices in $\mathcal{S}^{n,k}$ and RIP matrices, in the actual proof we adopt a strategy that does not “flow” exactly as described above but is easier to analyze. We will: 1) select W , a RIP matrix by selecting parameters m and δ and applying Theorem 10; 2) use it to construct a matrix $M \in \mathcal{S}^{n,k}$; 3) rescale the resulting matrix so that its Frobenius norm is 1, and; 4) finally compute its distance from \mathcal{S}_+^n and show that this is a constant independent of n .

Actual construction of M . Set $m = 93k$ and $\delta = 0.9$. Then we can numerically verify that whenever $k \geq 2$, the probability (12) is at least $0.51 > \frac{1}{2}$. Then let W be a random $m \times n$ matrix as in Theorem 10, and define the matrix

$$M := -(1 - \delta)I + W^\top W.$$

First observe that M has a large relative distance to the PSD cone and with good probability belongs to the k -PSD closure.

Lemma 5 *The matrix M satisfies the following:*

1. *With probability at least 0.51, $M \in \mathcal{S}^{n,k}$*
2. *$\text{dist}_F(M, \mathcal{S}_+^n) \geq \sqrt{n - m} (1 - \delta)$.*

Proof Whenever W is (k, δ) -RIP, by definition, for all k -sparse x we have $x^\top W^\top W x = \|Wx\|^2 \geq (1 - \delta)x^\top x$. Therefore $x^\top Mx \geq 0$ for all k -sparse x , and hence $M \in \mathcal{S}^{n,k}$ by Observation 1. This gives the first item of the lemma.

For the second item, notice that all vectors in the kernel of W , which has dimension $n - m$, are eigenvectors of M with eigenvalue $-(1 - \delta)$. So the negative eigenvalues of M include at least $n - m$ copies of $-(1 - \delta)$, and the result follows from Proposition 1. □

Now we need to normalize M , and for that we need to control its Frobenius norm.

Lemma 6 *With probability at least $\frac{1}{2}$, $\|M\|_F^2 \leq 2n\delta^2 + \frac{2n(n-1)}{m}$.*

Proof Notice that the diagonal entries of $W^\top W$ equal 1, so

$$\|M\|_F^2 = \sum_{i=1}^n M_{ii}^2 + \sum_{i,j \in [n], i \neq j} M_{ij}^2 = n\delta^2 + \sum_{i,j \in [n], i \neq j} (W^\top W)_{ij}^2.$$

We upper bound the last sum. Let the columns of W be C^1, \dots, C^n , and denote by $X_{ij} = \langle C^i, C^j \rangle$ the ij -th entry of $W^\top W$. Notice that when $i \neq j$, X_{ij} is the sum of m independent random variables $C_\ell^i C_\ell^j$ that take values $\{-\frac{1}{m}, \frac{1}{m}\}$ with equal probability, where ℓ ranges from 1 to m . Therefore,

$$\mathbb{E}X_{ij}^2 = \text{Var}(X_{ij}) = \sum_{\ell=1}^m \text{Var}(C_\ell^i C_\ell^j) = m \frac{1}{m^2} = \frac{1}{m}.$$

This gives that

$$\mathbb{E} \|M\|_F^2 = n\delta^2 + \frac{n(n-1)}{m}.$$

Since $\|M\|_F^2$ is non-negative, from Markov’s inequality $\|M\|_F^2 \leq 2\mathbb{E} \|M\|_F^2$ with probability at least $1/2$. This gives the desired bound, concluding the proof.

Taking a union bound over Lemmas 5 and 6, with strictly positive probability the normalized matrix $\frac{M}{\|M\|_F}$ belongs to $\mathcal{S}^{n,k}$ and has

$$\text{dist}_F \left(\frac{M}{\|M\|_F}, \mathcal{S}_+^n \right) \geq \frac{\sqrt{n-m} (1-\delta)}{\sqrt{2n(n-1)/m + 2n\delta^2}} \geq \frac{\sqrt{n-m} (1-\delta)}{\sqrt{2n^2/m + 2n\delta^2}}.$$

The first inequality comes from the fact that \mathcal{S}_+^n is a pointed convex cone, and for any constant $a > 0$ we have $\text{dist}_F(aM, \mathcal{S}_+^n) = a \text{dist}_F(M, \mathcal{S}_+^n)$. Thus, there is a matrix with such properties.

Now plugging in $k = rn, m = 93k, \delta = 0.9$, the right hand side is at least $\frac{\sqrt{r-93r^2}}{\sqrt{162r+3}}$. This concludes the proof of Theorem 4.

9 Proof of Theorem 5

The idea of the proof is similar to that of Theorem 1 (in Sect. 5), with the following difference: Given a unit-norm matrix $M \in \mathcal{S}_{\mathcal{I}}$, we construct a matrix \tilde{M} by averaging over the principal submatrices indexed by *only the k -sets in \mathcal{I}* instead of considering all k -sets, and upper bound the distance from M to the PSD cone by $\text{dist}_F(M, \alpha \tilde{M})$. Then we need to provide a uniform upper bound on $\text{dist}_F(M, \tilde{M})$ that holds for all M ’s *simultaneously with good probability* (with respect to the samples \mathcal{I}). This will then give an upper bound on $\overline{\text{dist}_F}(\mathcal{S}_{\mathcal{I}}, \mathcal{S}_+^n)$.

Recall that $\mathcal{I} = (I_1, \dots, I_m)$ is a sequence of independent uniform samples from the k -sets of $[n]$. As defined in Sect. 5, let M^I be the matrix where we zero out all the rows and columns of M except those indexed by indices in I . Let $T_{\mathcal{I}}$ be the (random) partial averaging operator, namely for every matrix $M \in \mathbb{R}^{n \times n}$

$$T_{\mathcal{I}}(M) := \frac{1}{|\mathcal{I}|} \sum_{I \in \mathcal{I}} M^I.$$

As we showed in Sect. 5 for the full average $\tilde{M} := T_{\binom{[n]}{k}}(M)$, the first observation is that if $M \in \mathcal{S}_{\mathcal{I}}$, that is, all principal submatrices $\{M_I\}_{I \in \mathcal{I}}$ are PSD, then the partial average $T_{\mathcal{I}}(M)$ is also PSD.

Lemma 7 *If $M \in \mathcal{S}_{\mathcal{I}}$, then $T_{\mathcal{I}}(M)$ is PSD.*

Proof This is straightforward, since each M^I is PSD.

Consider a unit-norm matrix M . Now we need to upper bound $\text{dist}_F(M, \alpha T_{\mathcal{I}}(M))$, for a scaling α , in a way that is “independent” of M . In order to achieve this goal, notice that $(T_{\mathcal{I}}(M))_{ij} = f_{ij}M_{ij}$, where f_{ij} is the fraction of sets in \mathcal{I} that contain $\{i, j\}$. Then it is not difficult to see that the Frobenius distance between M and $T_{\mathcal{I}}(M)$ can be controlled using only these fractions $\{f_{ij}\}$, since the Frobenius norm of M is fixed to be 1.

The next lemma makes this observation formal. Since the fractions $\{f_{ij}\}$ are random (they depend on \mathcal{I}), the lemma focuses on the typical scenarios where they are close to their expectations.

Notice that the probability that a fixed index i belongs to I_{ℓ} is $\frac{k}{n}$, so the fraction f_{ii} is $\frac{k}{n}$ in expectation. Similarly, the expected value of f_{ij} is $\frac{k(k-1)}{n(n-1)}$ when $i \neq j$. In other words, the expectation of $T_{\mathcal{I}}(M)$ is \tilde{M} .

Lemma 8 *Consider $\varepsilon \in [0, 1)$ and let $\gamma := \frac{k(n-k)}{2n(n-1)}$. Consider a scenario where \mathcal{I} satisfies the following for some $\varepsilon \in [0, 1)$:*

1. *For every $i \in [n]$, the fraction of the sets in \mathcal{I} containing i is in the interval $[\frac{k}{n} - \varepsilon\gamma, \frac{k}{n} + \varepsilon\gamma]$.*
2. *For every pair of distinct indices $i, j \in [n]$, the fraction of the sets in \mathcal{I} containing both i and j is in the interval $[\frac{k(k-1)}{n(n-1)} - \varepsilon\gamma, \frac{k(k-1)}{n(n-1)} + \varepsilon\gamma]$.*

Then there is a scaling $\alpha > 0$ such that for all matrices $M \in \mathbb{R}^{n \times n}$ we have

$$\text{dist}_F(M, \alpha T_{\mathcal{I}}(M)) \leq (1 + \varepsilon) \frac{n - k}{n + k - 2} \|M\|_F.$$

Proof As in Sect. 5, let $\tilde{M} = T_{\binom{[n]}{k}}(M)$ be the full average matrix. Recall that $\tilde{M}_{ii} = \tilde{f}_{ii}M_{ii}$ for $\tilde{f}_{ii} = \frac{k}{n}$, and $\tilde{M}_{ij} = \tilde{f}_{ij}M_{ij}$ for $\tilde{f}_{ij} = \frac{k(k-1)}{n(n-1)}$ when $i \neq j$. Also let $\alpha := \frac{2n(n-1)}{k(n+k-2)}$. Finally, define $\Delta := \tilde{M} - T_{\mathcal{I}}(M)$ as the error between the full and partial averages.

From triangle inequality we have

$$\|M - \alpha T_{\mathcal{I}}(M)\|_F \leq \|M - \alpha \tilde{M}\|_F + \alpha \|\Delta\|_F.$$

Moreover, in Sect. 5 we proved the full average bound $\|M - \alpha \tilde{M}\|_F \leq \frac{n-k}{n+k-2} \|M\|_F$. Moreover, from our assumptions we have $f_{ij} \in [\tilde{f}_{ij} - \varepsilon\gamma, \tilde{f}_{ij} + \varepsilon\gamma]$ for all i, j , and hence $|\Delta_{ij}| \leq \varepsilon\gamma |M_{ij}|$; this implies the norm bound $\|\Delta\|_F \leq \varepsilon\gamma \|M\|_F$. Putting these bounds together in the previous displayed inequality gives

$$\|M - \alpha T_{\mathcal{I}}(M)\|_F \leq \left(\frac{n - k}{n + k - 2} + \varepsilon\alpha\gamma \right) \|M\|_F = (1 + \varepsilon) \frac{n - k}{n + k - 2} \|M\|_F.$$

This concludes the proof. □

Finally, we use concentration inequalities to show that the “typical” scenario assumed in the previous lemma holds with good probability.

Lemma 9 *With probability at least $1 - \delta$ and the parameter m given in Theorem 5, the sequence \mathcal{I} is in a scenario satisfying the assumptions of Lemma 8.*

Proof As stated in Lemma 8, we only need that for all entries i, j the fraction f_{ij} deviates from its expectation by at most $+\varepsilon\gamma$, with failure probability at most δ . From union bound, this can be achieved if for each entry, the probability that the deviation of its fraction f_{ij} fails to be within $[-\varepsilon\gamma, \varepsilon\gamma]$ is at most $\frac{\delta}{n^2}$. Now we consider both diagonal and off-diagonal terms:

1. Diagonal terms f_{ii} : For each k -set sample I , let X_I be the indicator variable that is 0 if $i \notin I$, and 1 if $i \in I$. Notice that they are independent, with expectation $\frac{k}{n}$. Let $X = \sum_{I \in \mathcal{I}} X_I$ be the sum of these variables.

From definition of f_{ii} we have $X = f_{ii}m$, where m is the total number of samples. From Chernoff bound, have that

$$\begin{aligned} \Pr \left(\left| f_{ii} - \frac{k}{n} \right| > \varepsilon \frac{(n - k)k}{2n(n - 1)} \right) &= \Pr \left(\left| X - \frac{mk}{n} \right| > \varepsilon m \frac{(n - k)k}{2n(n - 1)} \right) \\ &\leq 2 \exp \left(- \frac{\varepsilon^2 (n - k)^2 km}{12n(n - 1)^2} \right) \\ &\leq \frac{\delta}{n^2} \end{aligned}$$

as long as

$$m \geq \frac{12n(n - 1)^2}{\varepsilon^2(n - k)^2k} \ln \frac{2n^2}{\delta}.$$

2. Off-diagonal terms f_{ij} : Similar to first case, now for each k -set sample I , let X_I be the indicator variable that is 1 if $\{i, j\} \subset I$, and 0 otherwise. Now the expectation of each X_I becomes $\frac{k(k-1)}{n(n-1)}$. Again let $X = \sum_{I \in \mathcal{I}} X_I$.

Using same argument as above, $X = f_{ij}m$. From Chernoff bound we get

$$\begin{aligned} \Pr \left(\left| f_{ij} - \frac{k(k - 1)}{n(n - 1)} \right| > \varepsilon \frac{(n - k)k}{2n(n - 1)} \right) &= \Pr \left(\left| X - \frac{mk(k - 1)}{n(n - 1)} \right| > \varepsilon m \frac{(n - k)k}{2n(n - 1)} \right) \\ &\leq 2 \exp \left(- \frac{\varepsilon^2 (n - k)^2 km}{12n(n - 1)(k - 1)} \right) \\ &\leq \frac{\delta}{n^2} \end{aligned}$$

as long as

$$m \geq \frac{12n(n-1)(k-1)}{\varepsilon^2(n-k)^2k} \ln \frac{2n^2}{\delta}.$$

Since we chose m large enough so it satisfies both of these cases, taking a union bound over all i, j 's we get that the probability that any of the f_{ij} 's is $+\varepsilon\gamma$ more than their expectations is at most δ . This concludes the proof. \square

Combining this with Lemma 8, we conclude the proof of Theorem 5.

Acknowledgements Grigoriy Blekherman was partially supported by NSF grant DMS-1901950. Marco Molinaro was supported in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES, Brasil) - Finance Code 001, by CNPq grant Universal #431480/2016-8, and by the CAPES-Print program. This research was partially supported by ONR under grant No. N4458-NV-ONR.

References

1. Achlioptas, D.: Database-friendly random projections. In: Proceedings of the Twentieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS'01, New York, NY, USA, Association for Computing Machinery, pp. 274–281 (2001)
2. Ahmadi, A.A., Majumdar, A.: Dsos and sdsos optimization: more tractable alternatives to sum of squares and semidefinite optimization. *SIAM J. Appl. Algebra Geom.* **3**, 193–230 (2019)
3. Amaldi, E., Coniglio, S., Gualandi, S.: Coordinated cutting plane generation via multi-objective separation. *Math. Program.* **143**, 87–110 (2014)
4. Baltean-Lugojan, R., Bonami, P., Misener, R., Tramontani, A.: Selecting cutting planes for quadratic semidefinite outer-approximation via trained neural networks (2018)
5. Bandeira, A.S., Fickus, M., Mixon, D.G., Wong, P.: The road to deterministic matrices with the restricted isometry property. *J. Fourier Anal. Appl.* **19**, 1123–1149 (2013)
6. Baraniuk, R., Davenport, M., DeVore, R., Wakin, M.: A simple proof of the restricted isometry property for random matrices. *Construct. Approx.* **28**, 253–263 (2008)
7. Bixby, R.E.: Solving real-world linear programs: a decade and more of progress. *Oper. Res.* **50**, 3–15 (2002)
8. Boman, E.G., Chen, D., Parekh, O., Toledo, S.: On factor width and symmetric h-matrices. *Linear Algebra Appl.* **405**, 239–248 (2005)
9. Boucheron, S., Lugosi, G., Massart, P.: Concentration Inequalities: A Nonasymptotic Theory of Independence. OUP, Oxford (2013)
10. Bruck, R.H., Ryser, H.J.: The nonexistence of certain finite projective planes. *Can. J. Math.* **1**, 88–93 (1949)
11. Calderbank, R., Howard, S., Jafarpour, S.: Construction of a large class of deterministic sensing matrices that satisfy a statistical isometry property. *IEEE J. Sel. Top. Signal Process.* **4**, 358–374 (2010)
12. Candes, E.J., Romberg, J.K., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. *Commun. Pure Appl. Math. J. Iss. Courant Inst. Math. Sci.* **59**, 1207–1223 (2006)
13. Candes, E.J., Tao, T.: Decoding by linear programming. *IEEE Trans. Inf. Theory* **51**, 4203–4215 (2005)
14. Candès, E.J.: The restricted isometry property and its implications for compressed sensing. *C. R. Math.* **346**, 589–592 (2008)
15. Chartrand, R., Staneva, V.: Restricted isometry properties and nonconvex compressive sensing. *Inverse Prob.* **24**, 035020 (2008)
16. Chowla, S., Ryser, H.J.: Combinatorial problems. *Can. J. Math.* **2**, 93–99 (1950)
17. Coleman, T.F., Li, Y., et al.: Large-Scale Numerical Optimization, vol. 46, Siam (1990)
18. Dey, S.S., Kazachkov, A., Lodi, A., Munoz, G.: Sparse cutting planes for quadratically-constrained quadratic programs (2019)
19. Dey, S.S., Molinaro, M.: Theoretical challenges towards cutting-plane selection. *Math. Program.* **170**, 237–266 (2018)

20. Dey, S.S., Molinaro, M., Wang, Q.: Approximating polyhedra with sparse inequalities. *Math. Program.* **154**, 329–352 (2015)
21. Dey, S.S., Molinaro, M., Wang, Q.: Analysis of sparse cutting planes for sparse milps with applications to stochastic milps. *Math. Oper. Res.* **43**, 304–332 (2017)
22. Hojny, C., Lüthen, H., Pfetsch, M. E.: On the size of integer programs with bounded coefficients or sparse constraints. Technical Report, Technical report, Optimization Online (2017)
23. Horn, R., Johnson, C.: *Matrix Analysis*. Cambridge University Press, Cambridge (1985)
24. Kaas, R., Buhrman, J.M.: Mean, median and mode in binomial distributions. *Stat. Neerl.* **34**, 13–18 (1980)
25. Kim, S., Kojima, M., Mevissen, M., Yamashita, M.: Exploiting sparsity in linear and nonlinear matrix inequalities via positive semidefinite matrix completion. *Math. Program.* **129**, 33–68 (2011)
26. Kocuk, B., Dey, S.S., Sun, X.A.: Strong socp relaxations for the optimal power flow problem. *Oper. Res.* **64**, 1177–1196 (2016)
27. Permenter, F., Parrilo, P.: Partial facial reduction: simplified, equivalent sdps via approximations of the psd cone. *Math. Program.* **171**, 1–54 (2018)
28. Qualizza, A., Belotti, P., Margot, F.: Linear programming relaxations of quadratically constrained quadratic programs. In: *Mixed Integer Nonlinear Programming*, Springer, pp. 407–426 (2012)
29. Reid, J.K.: A sparsity-exploiting variant of the bartels-golub decomposition for linear programming bases. *Math. Program.* **24**, 55–69 (1982)
30. Sojoudi, S., Lavaei, J.: Exactness of semidefinite relaxations for nonlinear optimization problems with underlying graph structure. *SIAM J. Optim.* **24**, 1746–1778 (2014)
31. Stinson, D.R.: *Combinatorial Designs*. Springer, New York (2004)
32. Walter, M.: Sparsity of lift-and-project cutting planes. In: *Operations Research Proceedings*, vol. 2014, Springer, pp. 9–14 (2012)
33. Wang, Y., Tanaka, A., Yoshise, A.: Polyhedral approximations of the semidefinite cone and their applications. arXiv preprint [arXiv:1905.00166](https://arxiv.org/abs/1905.00166) (2019)
34. Wilson, R.M.: An existence theory for pairwise balanced designs: III—proof of the existence conjectures. *J. Combin. Theory Ser. A* **18**, 71–79 (1975)
35. Wolkowicz, H., Saigal, R., Vandenberghe, L.: *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*, vol. 27, Springer, Berlin (2012)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.