JMLR: Workshop and Conference Proceedings vol 49:1–24, 2016

On the Approximability of Sparse PCA

Siu On Chan SOCHAN@GMAIL.COM
Chinese University of Hong Kong. Most of this work was done while the author was a postdoc at Microsoft Research New England.

Dimitris Papailliopoulos DIMITRISP@BERKELEY.EDU
UC Berkeley. Supported by NSF awards CCF-1217058 and CCF-1116404, and MURI AFOSR grant 556016.

Aviad Rubinstein AVIAD@EECS.BERKELEY.EDU
UC Berkeley. Parts of this work were done at the Simons Institute for the Theory of Computing and Microsoft Research New England. Supported by NSF grants CCF0964033 and CCF1408635, and by Templeton Foundation grant 3966.

Abstract

It is well known that Sparse PCA (Sparse Principal Component Analysis) is NP-hard to solve exactly on worst-case instances. What is the complexity of solving Sparse PCA approximately? Our contributions include:

1. a simple and efficient algorithm that achieves an $n^{-1/3}$-approximation;
2. NP-hardness of approximation to within $(1-\varepsilon)$, for some small constant $\varepsilon > 0$;
3. SSE-hardness of approximation to within *any* constant factor; and
4. an $\exp\exp\left(\Omega\left(\sqrt{\log\log n}\right)\right)$ ("quasi-quasi-polynomial") gap for the standard semidefinite program.

Keywords: Sparse PCA; hardness of approximation.

1. Introduction

Principal component analysis (PCA) is one of the most popular tools for data analytics. PCA operates on data point vectors supported on features, and outputs orthogonal directions (*i.e.,* principal components) that maximize the *explained variance*. A limitation of PCA is that —in many cases of interest — the extracted principal components (PCs) are dense. However, in applications such as text analysis, or gene expression analytics, having only a few non-zero features per extracted PC, offers significantly higher interpretabilty. For example, in text analysis where PCs are supported on words, if they consist of only a few of them, then these words can be used to detect frequently occurring topics.

Sparse PCA addresses the issue of interpretability directly by enforcing a sparsity constraint on the extracted PCs. Given a matrix of centered data samples $\mathbf{S} \in \mathbb{R}^{n\times p}$, let us denote by $\mathbf{A} = \frac{1}{n}\mathbf{S}\mathbf{S}^T$ the sample covariance matrix of the data set. The leading sparse principal component is the solution to the following sparsity constrained, quadratic form maximization

$$\max_{\|\mathbf{x}\|_2=1, \|\mathbf{x}\|_0\le k} \mathbf{x}^{\mathsf{T}}\mathbf{A}\mathbf{x} \qquad \text{(SparsePCA)}$$

where $\|x\|_2$ is the ℓ_2-norm, $\|\mathbf{x}\|_0$ denotes the number of nonzero entries in $\mathbf{x}$, and $\mathbf{A} \succcurlyeq 0$.

The objective in the above optimization is usually referred to as the *explained variance*. This metric has an operational meaning: if a linear combination of k features has high explained variance, then it captures a representative behavior of the samples. Typically, this means that these features "interact" significantly with each other. As an example consider the case where $\mathbf{A}$ is a covariance matrix of a gene expression data set. Then, the (i, j) entry of $\mathbf{A}$ is a proxy for the positive or negative interaction between the ith and jth gene. In this case, if a subset of k genes "explains" a lot of variance, then these genes have strong pairwise interactions.

There has been a large volume of work on sparse PCA: from heuristic algorithms, to statistical guarantees, and conditional approximation ratios. Yet, there are remarkably few worst-case approximability bounds, and many questions remain open. Does sparse PCA admit a nontrivial worst-case approximation ratio? Are there significant computational barriers? How does it relate to other problems? In this work we take a modest first step towards a better understanding of these questions. Our contributions are summarized below.

1. We show that a simple spectral technique that is popular in practice, combined with a column selection procedure, achieves an $n^{-1/3}$-approximation ratio for SparsePCA.
2. We establish that, assuming P $\neq$ NP, SparsePCA does not admit a PTAS.
3. We further prove that, assuming Small Set Expansion (SSE) Hypothesis (Raghavendra and Steurer, 2010a), SparsePCA is hard to approximate to within *any* constant factor.
4. We construct an $e^{e^{\Omega(\sqrt{\log \log n})}}$ (i.e., a "quasi-quasi-polynomial") gap instance for the following commonly used SDP relaxation of (d'Aspremont et al., 2007a)

$$\begin{aligned} \max \quad & \operatorname{tr}(\mathbf{AX}) \\ \text{such that} \quad & \operatorname{tr}(\mathbf{X}) = 1,\ \mathbf{1}^{\mathsf{T}}|\mathbf{X}|\mathbf{1} \leqslant k,\ \mathbf{X} \succcurlyeq 0 \end{aligned}.$$

   Here $|\mathbf{X}|$ denotes the matrix obtained from $\mathbf{X}$ by taking entry-wise absolute value.

1.1. Discussion of techniques and connections to other sparsity problems

A recurring theme in our technical discussion is the comparison of SparsePCA to (variants of) the Densest k-Subgraph (DkS) problem: given a graph G, find the k-vertex subgraph that contains the highest number of edges. Notice that DkS can be stated as a quadratic form maximization, similar to SparsePCA:

$$\max_{\mathbf{x} \in \{0,1\}^n, \|\mathbf{x}\|_0 \leq k} \mathbf{x}^T \mathbf{A} \mathbf{x} \qquad \text{(DkS)}$$

(where we use $\mathbf{A}$ to denote the adjacency matrix instead of the covariance matrix).

The connection between the two problems has been observed before. For example, it has been noted by many authors that the k-Clique problem, a decision variant of DkS[1], reduces to solving SparsePCA exactly, thus the latter is NP-hard. Then, the Planted Clique, an average-case variant

1. Notice that k-Clique is an exact variant of both Max-Clique and DkS. By now, the inapproximability of Max-Clique is relatively well understood (e.g. (Håstad, 1999; Khot, 2001; Zuckerman, 2007)), but these results do not translate to inapproximability of DkS (or SparsePCA).

of DkS, was recently used to establish statistical recovery hardness results in the sparse spiked-covariance model (Berthet and Rigollet, 2013b,a; Wang et al., 2014; Gao et al., 2014).

Then, why are algorithmic and inapproximability DkS results not directly applicable to SparsePCA? From a computational standpoint, the most important difference between the two problems is the restriction on the input matrix $\mathbf{A}$: In SparsePCA, $\mathbf{A}$ is required to be positive semi-definite, whereas in DkS, $\mathbf{A}$ is required to be entry-wise nonnegative.

With the above comparison to DkS in mind, we are now ready to discuss our results and techniques.

$n^{-1/3}$-approximation algorithm Our $n^{-1/3}$-approximation scheme outputs the best solution among the following three procedures: *i)* pick the best standard basis vector; *ii)* pick the largest k entries in any column vector of $\mathbf{A}$; and *iii)* pick the largest k entries of the leading eigenvector of $\mathbf{A}$.

Our algorithm is inspired by (but is substantially different from) a combinatorial $\Omega\left(n^{-1/3}\right)$-approximation algorithm for DkS, due to Feige et al. (Feige et al., 2001). The aforementioned ratio for DkS was further improved in the same paper to $\Omega\left(n^{-1/3+\varepsilon}\right)$, and later to $\Omega\left(n^{-1/4+\varepsilon}\right)$ (Bhaskara et al., 2010). It is an open question whether similar ideas can improve the approximation guarantees for SparsePCA.

NP-hardness Our NP-hardness of approximation reduction begins from MAX-E2SAT-d, the problem of maximizing the number of satisfied clauses in a CNF formula, where every clause contains exactly two distinct variables, and every variable appears in exactly $d = O(1)$ clauses. We set $A_{i,j}$ to be higher if literals i and j satisfy some clause, and a consistent assignment is ensured by having large negative values whenever indices i and j correspond to a literal and its negation. A PSD matrix is obtained by adding a large multiple of the identity. As we discuss below, this last step seems to be the main obstacle to obtaining a stronger inapproximability factor.

Interestingly, this result highlights an important difference between SparsePCA and DkS: for the latter, even proving NP-hardness of approximation to within a $(1-\varepsilon)$ factor (for a small constant $\varepsilon > 0$) remains a major open problem.

The PSD challenge The biggest challenge to obtaining inapproximability results for SparsePCA, when reducing from say DkS, is achieving $\mathbf{A} \succcurlyeq 0$. One naive approach is to add a large multiple of the identity matrix and force diagonal dominance (as we do for our NP-hardness result). Unfortunately, this ruins our inapproximability factor: the large entries on the diagonal outweigh the interactions between different features. In particular, *every* vector achieves a reasonably high explained variance.

A second approach to obtain a PSD matrix is by squaring the adjacency matrix. When we start from Planted Clique, or other known hard DkS instances (e.g. (Bhaskara et al., 2012; Alon et al., 2011; Khot, 2006; Braverman et al., 2015a)), squaring the adjacency matrix gives weak inapproximability results, as in the case of (Berthet and Rigollet, 2013b,a; Krauthgamer et al., 2015; Wang et al., 2014; Gao et al., 2014) (see also discussion of impossibility results for the sparse spiked covariance model below). To understand why, it is helpful to consider random walks on regular graphs. The density of a subgraph is proportional to the probability that a length-1 random walk remains in the subgraph. (Thus the densest k-subgraph is also the least *expanding* k-subgraph.) Similarly, when we restrict $\mathbf{A}^2$ to the same k-tuple of vertices, the density corresponds to the probability of remaining in the subgraph after a random walk of length 2. Intuitively, squaring the adjacency matrices of the instances mentioned at the beginning of this paragraph is ineffective, because even their dense subgraphs are very expanding: most length-2 walks that start and end inside the densest k-subgraph, take their middle step outside the subgraph. Thus the density of the subgraph has only a small effect on the density with respect to $\mathbf{A}^2$. To overcome this difficulty, we want the "good" subgraph to have very small expansion.

SSE-hardness and SDP gap The Small Set Expansion Hypothesis (Raghavendra and Steurer, 2010a) postulates that it is hard to find a linear-size subgraph with a very small expansion. Intuitively, if the expansion of a particular k-subgraph is sufficiently small, then, even after taking two steps, the random walk should remain inside the subgraph; thus the corresponding k-sparse vector should continue to give an exceptionally high value for DkS/SparsePCA with $\mathbf{A}^2$. To formalize this intuition, we apply a recent result of Raghavendra and Schramm (Raghavendra and Schramm, 2014) on the expansion of random walk graphs.

Finally, the gap for the standard semidefinite program for SparsePCA builds on known integrality gap instances for SSE, in particular the Short Code graph (Barak et al., 2012b). Notice that the "quasi-quasi-polynomial" factor ($e^{e^{\Omega(\sqrt{\log\log n})}}$) is slightly smaller than polynomial and "quasi-polynomial" ($e^{\Omega(\sqrt{\log n})}$) factors, but much larger than polylogarithmic.

Additive PTAS To complete the picture of our current understanding of worst-case approximability of SparsePCA, let us also mention a recent additive PTAS due to (Asteris et al., 2015). By additive PTAS, we mean that if all the entries of $\mathbf{A}$ are bounded in [-1, 1], the optimum explained variance can be approximated in polynomial time to within an additive error of εk, for any constant $\varepsilon > 0$ (compare to an optimum of at most k in the case of an all-ones $k \times k$ submatrix). In contrast, note that a corresponding additive PTAS for DkS is unlikely (Braverman et al., 2015a).

1.2. Worst-case, average-case, and best-case approximability of SparsePCA

The performance of many algorithms has been analyzed under the *sparse spiked covariance* and related models. For example, under this model Amini et al. (Amini and Wainwright, 2008) develop the first theoretical guarantees for simple thresholding and the SDP of (d'Aspremont et al., 2007a). Several statistical analyses were carried for more general settings, while using a variety of different algorithms, under various metrics of interest (Ma, 2013; d'Aspremont et al., 2014; Cai et al., 2013a,b; Deshpande and Montanari, 2013; Krauthgamer et al., 2015; Ma and Wigderson, 2015). A recent celebrated line of works (Berthet and Rigollet, 2013b,a; Wang et al., 2014; Gao et al., 2014), initiated by Berthet and Rigollet, also establish a gap between the threshold of samples where detection is information theoretically possible, and that where it is computationally feasible, assuming hardness of the Planted Clique problem.

This excellent body of work is often described as "average-case analysis". We find this term somewhat confusing, as *average* may simultaneously refer to two or more of the following:

Average-case hardness assumptions: As we mentioned above, (Berthet and Rigollet, 2013b,a; Wang et al., 2014; Gao et al., 2014) prove hardness results for SparsePCA assuming hardness of the Planted Clique problem. Average-case hardness assumptions like Planted Clique are more likely to be false than worst-case assumptions like P $\neq$ NP. In this sense replacing average-case hardness assumptions with worst-case hardness is desirable whenever possible (Braverman et al., 2015b). (Note, however, that there are much more substantial differences between this line of "average case" papers and our work, and replacing average-case hardness assumption is far from being our central goal.)

Average over random noise: The main goal of all the works we mentioned above is to characterize the resilience of different algorithms to random noise, where the noise distribution may be slightly different in each paper. Understanding the tradeoffs between resilience to noise (or number of samples) and computational complexity is an exciting research direction. In our case however, our hardness results hold even in the presence of no noise! (This corresponds to making no assumptions about the availability of samples.) In this sense, our work may be better described as *best-case analysis*.

Average over instances: In the sparse spiked covariance model, the input is a collection of samples from a distribution with a covariance matrix that is equal to the identity plus a sparse rank-1 matrix (the spike). The goal is to identify (or detect) the rank-1 sparse "spike" from the samples. (If we could observe the true covariance matrix, the algorithmic task would be trivial. However, when the input to this problem is a finite number of samples, then there exist sharp information theoretic, and computational barriers on the identifiability of the spike.) The works mentioned above draw their instances from this specific distribution, whereas our instances are "worst-case", aka they are drawn from less canonical distributions.

A classic response in computer science to impossibility results as in the "average case" works is *approximation*. We may not be able to find the optimum solution, but can we find a solution which is almost as good?

1.3. Additional Related work

The algorithmic tapestry for sparse PCA is rich and diverse. Early heuristics used rotation and thresholding of eigenvectors (Kaiser, 1958; Jolliffe, 1995; Cadima and Jolliffe, 1995) and LASSO heuristics (Ando et al., 2009; Jolliffe et al., 2003). Then, in (Zou et al., 2006), a nonconvex ℓ_1 penalized approximation, re-generated a lot of interest in the problem. A great variety of greedy, spectral, and nonconvex heuristics were presented in the past decade (Sriperumbudur et al., 2007; Moghaddam et al., 2006, 2007; Shen and Huang, 2008; Journée et al., 2010; Yuan and Zhang, 2013; Kuleshov, 2013). There has also been a significant body of work on semidefinite programming (SDP) approaches (d'Aspremont et al., 2007a; Zhang et al., 2012; d'Aspremont et al., 2008, 2014). Some recent works established conditional approximation guarantees for sparse PCA using spectral ϵ-net search algorithms, under the assumption of a decaying matrix spectrum (Asteris et al., 2014, 2015).

We should also mention some recent inapproximability results in the general case where $\mathbf{A}$ is not necessarily positive semi-definite (PSD) (Magdon-Ismail, 2015). (Recall that in typical applications $\mathbf{A}$ is a covariance matrix and thus necessarily PSD.) We note that in this general matrix setting, it is even hopeless to determine, in polynomial time, the sign of the optimal value, unless $\mathsf{P} = \mathsf{NP}$.

1.4. Organization

Our approximation algorithm is described in Section 2. In Section 3 we prove our NP-hardness result, and our SSE-hardness result appears in Section 4. Finally, in Section 5 we prove the quasi-quasi-polynomial gap for the standard SDP. For completeness, we also briefly describe in Section 6 the additive PTAS due to (Asteris et al., 2015), and shortly discuss in Section 7 the case where the input matrix is not PSD.

2. $n^{-1/3}$-approximation algorithm

Theorem 1 *SparsePCA can be approximated to within $n^{-1/3}$ in deterministic polynomial time.*

The rest of this section is devoted to the proof of Theorem 1. Our algorithm takes the best of two options: a truncation of one of $\mathbf{A}$'s columns in the standard basis, and a truncation of one of $\mathbf{A}$'s eigenvectors. We present and analyze the guarantees for each algorithm, and then show that together they give the bound on the approximation ratio.

Let $\mathbf{y}_*$ denote an optimum solution to the SPARSE PCA instance, and let $OPT = \mathbf{y}_*^\intercal \mathbf{A} \mathbf{y}_*$ denote its value.

2.1. Truncation in the standard basis

Algorithm 1 For each $i \in [n]$, let $\mathbf{A}_{\cdot,i}$ be the i-th column of $\mathbf{A}$, and let $\mathbf{x}_i$ be the unit-norm, k-sparse truncation of $\mathbf{A}_{\cdot,i}$. That is, let

$$[\hat{\mathbf{x}}_i]_j = \begin{cases} \mathbf{A}_{i,j} & \text{if } |\mathbf{A}_{i,j}| \text{ is one of the } k \text{ largest (in absolute value) entries of } \mathbf{A}_{\cdot,i} \\ 0 & \text{otherwise} \end{cases}$$

and $\mathbf{x}_i = \hat{\mathbf{x}}_i / \|\hat{\mathbf{x}}_i\|_2$.

Return the best out of all $\mathbf{x}_i$'s and $\mathbf{e}_i$'s, where $\mathbf{e}_i$ is the i-th standard basis vector.

Lemma 2 *Algorithm 1 returns a solution with value $V_1(\mathbf{A}, k) \geqslant \frac{OPT}{\sqrt{k}}$*

Proof First, we claim that for each i, $\mathbf{x}_i$ maximizes $\mathbf{e}_i^\intercal \mathbf{A} \mathbf{x}_i$ among all feasible (k-sparse and unit-norm) vectors. By Cauchy-Schwartz inequality, for any choice of support S of size k, the unit-norm vector that maximizes the inner product with $\mathbf{A}_{\cdot,i}$ is the restriction of $\mathbf{A}_{\cdot,i}$ to S, normalized. The inner product is thus $\sqrt{\sum_{j \in S} \mathbf{A}_{j,i}^2}$; this is indeed maximized when S is the set of entries with largest absolute value.

Now, rewrite $\mathbf{y}_* = \sum y_i \mathbf{e}_i$ as a linear combination of (at most k) standard basis vectors. By Cauchy-Schwartz inequality, we have

$$OPT = \sum y_i \left(\mathbf{e}_i^\intercal \mathbf{A} \mathbf{y}_*\right) \leqslant \sqrt{\sum y_i^2} \sqrt{\sum \left(\mathbf{e}_i^\intercal \mathbf{A} \mathbf{y}_*\right)^2}.$$

Plugging in $\sqrt{\sum y_i^2} = \|\mathbf{y}_*\|_2 = 1$, we get

$$OPT \leqslant \sqrt{\sum \left(\mathbf{e}_i^\intercal \mathbf{A} \mathbf{y}_*\right)^2} \leqslant \sqrt{k} \max_i \mathbf{e}_i^\intercal \mathbf{A} \mathbf{y}_*.$$

In particular, this means that for some i, then $\mathbf{e}_i^\intercal \mathbf{A} \mathbf{x}_i \geqslant OPT/\sqrt{k}$, where $\mathbf{x}_i$, is as defined above.

Finally, since $\mathbf{A} \succcurlyeq 0$, we have

$$0 \leqslant (\mathbf{e}_i - \mathbf{x}_i)^\intercal \mathbf{A} (\mathbf{e}_i - \mathbf{x}_i) = \mathbf{e}_i^\intercal \mathbf{A} \mathbf{e}_i + \mathbf{x}_i^\intercal \mathbf{A} \mathbf{x}_i - 2\mathbf{e}_i^\intercal \mathbf{A} \mathbf{x}_i.$$

Rearranging, we get

$$\max \{\mathbf{e}_i^\intercal \mathbf{A} \mathbf{e}_i, \mathbf{x}_i^\intercal \mathbf{A} \mathbf{x}_i\} \geqslant OPT/\sqrt{k}.$$

∎

2.2. Truncation in the eigenspace basis

Algorithm 2 Let $(\mathbf{v}_1, \lambda_1)$ be the top eigenvector and eigenvalue of $\mathbf{A}$. Return the unit-norm, k-sparse truncation of $\mathbf{v}_1$. That is, let

$$[\hat{\mathbf{x}}]_j = \begin{cases} [\mathbf{v}_1]_j & \text{if } [\mathbf{v}_1]_j \text{ is one of the } k \text{ largest (in absolute value) entries of } [\mathbf{v}_1]_j \\ 0 & \text{otherwise} \end{cases}$$

and $\mathbf{x} = \hat{\mathbf{x}} / \|\hat{\mathbf{x}}\|_2$. Return $\mathbf{x}$.

Lemma 3 *Algorithm 2 returns a solution with value $V_2(\mathbf{A}, k) \geqslant \frac{k}{n} OPT$.*

Proof First, notice that

$$\mathbf{x}^{\mathsf{T}} \mathbf{A} \mathbf{v}_1 = \lambda_1 \mathbf{x}^{\mathsf{T}} \mathbf{v}_1 = \lambda_1 \mathbf{x}^{\top} \hat{\mathbf{x}} = \lambda_1 \cdot \|\hat{\mathbf{x}}\|_2 \geqslant \lambda_1 \sqrt{k/n},$$

where the last inequality follows by the greedy construction of $\hat{\mathbf{x}}$. Since $\mathbf{A} \succcurlyeq 0$, it induces an inner product over $\mathbb{R}^n$. Thus we can apply the Cauchy Schwartz inequality to get:

$$\lambda_1 \sqrt{k/n} \leqslant \mathbf{x}^{\top} \mathbf{A} \mathbf{v}_1 \leqslant \sqrt{\mathbf{x}^{\top} \mathbf{A} \mathbf{x}} \cdot \sqrt{\mathbf{v}_1^{\top} \mathbf{A} \mathbf{v}_1} = \sqrt{\mathbf{x}^{\top} \mathbf{A} \mathbf{x}} \cdot \sqrt{\lambda_1}.$$

Rearranging, we have

$$\mathbf{x}^{\top} \mathbf{A} \mathbf{x} \geqslant \frac{k}{n} \lambda_1.$$

Finally, to complete the proof recall that $\lambda_1 = \mathbf{v}_1^{\top} \mathbf{A} \mathbf{v}_1 \geqslant OPT$ since $\mathbf{v}_1$ maximizes the objective function among all (not necessarily k-sparse) unit-norm vectors. ■

2.3. Putting it altogether

Our final algorithm simply takes the best out of the outputs of Algorithms 1 and 2. We now have

$$\begin{aligned} V(\mathbf{A}, k) &= \max\{V_1(\mathbf{A}, k), V_2(\mathbf{A}, k)\} \\ &\geqslant (V_1(\mathbf{A}, k))^{2/3} \cdot (V_2(\mathbf{A}, k))^{1/3} \\ &\geqslant \frac{OPT^{2/3}}{k^{1/3}} \cdot \frac{k^{1/3}}{n^{1/3}} OPT^{1/3} = OPT / n^{1/3} \end{aligned}$$

3. NP-hardness

Theorem 4 *There exists a constant $\varepsilon > 0$ such that SparsePCA is NP-hard to approximate to within $(1 - \varepsilon)$.*

Proof

We reduce from MAX-E2SAT-d: given a 2CNF over n variables where every variable appears in exactly d distinct clauses, maximize the number of satisfied clauses.

Lemma 5 *There exist constants $0 < s < c < 1$ and d such that given a MAX-E2SAT-d instance over n clauses, it is NP-hard to decide whether at least cn clauses can be satisfied ("yes" case), or at most sn ("no" case).*

Theorem 5 follows from standard techniques. We briefly sketch the proof below for completeness.

Proof [Proof sketch of Theorem 5] By, e.g. (Feige, 1998), MAX-3SAT-5 is NP-hard to approximate to within some constant factor. We can convert each 3SAT clause $C = (x \vee y \vee z)$ into 10 2SAT clauses (introducing one additional variable h_C),

$$(x) \wedge (y) \wedge (z) \wedge (h_C) \wedge (\neg x \vee \neg y) \wedge (\neg x \vee \neg z) \wedge (\neg y \vee \neg z) \wedge (x \vee \neg h_C) \wedge (y \vee \neg h_C) \wedge (z \vee \neg h_C)$$

with the following guarantee: the optimal assignment to the 2SAT instance satisfies at most 7 out of 10 clauses for every satisfied 3SAT clause, and at most 6 out of 10 clauses for every unsatisfied 3SAT clause (e.g. (Safra)).

This establishes the result for MAX-2SAT with bounded degree. Add a linear number of variables and trivially satisfied clauses to get a MAX-E2SAT-d instance. ■

Given a 2CNF ψ, we construct a symmetric $2n \times 2n$ matrix $\mathbf{A}^{(0)} = \mathbf{A}^{(0)}(\psi)$ as follows: every row/column corresponds to a literal of ψ; if row i and column j correspond to an assignment that satisfies some clause, then $\mathbf{A}^{(0)}_{i,j} = 1$, and $\mathbf{A}^{(0)}_{i,j} = 0$ otherwise. Let $\mathcal{Y}$ denote the set of vectors that correspond to legal assignments to ψ, i.e.

$$\mathcal{Y} = \left\{ \mathbf{y} : \substack{|\mathbf{y}|_2 = 1;\, \mathbf{y} \in \left\{0, 1/\sqrt{n}\right\}^{2n};\\ \forall \text{i} \;\; \mathbf{y}_{x_i} = 0 \iff \mathbf{y}_{\neg x_i} = 1/\sqrt{n}} \right\}$$

By Theorem 5 it is NP-hard to distinguish between

$$\text{“yes”:} \max_{\mathbf{y} \in \mathcal{Y}} \mathbf{y}^\intercal \mathbf{A}^{(0)} \mathbf{y} \geqslant c \qquad\qquad \text{“no”:} \max_{\mathbf{y} \in \mathcal{Y}} \mathbf{y}^\intercal \mathbf{A}^{(0)} \mathbf{y} \leqslant s.$$

The proof continues by adding the following matrices to $\mathbf{A}^{(0)}$: a matrix $\mathbf{C}$ with large negative entries that enforces a consistent assignment; a larger scalar times the identity matrix that ensures our input is PSD; and an even larger (yet still constant) scalar times the all-ones matrix that guarantees the optimal solution uses a large support. While adding these matrices preserves the qualitative properties of the instance, they significantly weaken our inapproximability factor.

ENFORCING A CONSISTENT ASSIGNMENT

Our first step is to enforce consistency using the objective function instead of restricting the input to be from $\mathcal{Y}$. Let $\mathbf{C}_{i,j} = -2d$ if i and j correspond to a literal and its negation, and $\mathbf{C}_{i,j} = 0$, otherwise. We claim that among all unit-norm vectors $\mathbf{z} \in \{0, 1/\sqrt{n}\}^{2n}$, the objective $\mathbf{z}\left(\mathbf{A}^{(0)} + \mathbf{C}\right)\mathbf{z}$ is maximized by some legal assignment $\mathbf{z}^* \in \mathcal{Y}$. Assume by contradiction that the objective is maximized by some $\mathbf{z}$ which assigns $1/\sqrt{n}$ to some variable x_i and its negation; since $\mathbf{z}$ is exactly n-sparse, it must also assign 0 to another variable x_j and its negation. However, the objective value can be increased by considering $\mathbf{z}'$ which assigns $1/\sqrt{n}$ to x_i and x_j, 0 to their negations, and is equal to $\mathbf{z}$ everywhere else. Therefore, for $\mathbf{A}^{(1)} \triangleq \mathbf{A}^{(0)} + \mathbf{C}$, we have

$$\text{“yes”:} \max_{\substack{|\mathbf{z}|_2 = 1\\ \mathbf{z} \in \left\{0, 1/\sqrt{n}\right\}^{2n}}} \mathbf{z}^\intercal \mathbf{A}^{(1)} \mathbf{z} \geqslant c \qquad\qquad \text{“no”:} \max_{\substack{|\mathbf{z}|_2 = 1\\ \mathbf{z} \in \left\{0, 1/\sqrt{n}\right\}^{2n}}} \mathbf{z}^\intercal \mathbf{A}^{(1)} \mathbf{z} \leqslant s.$$

PSD INPUT

$\mathbf{A}^{(1)}$ is not a legitimate input to SparsePCA because it is not be positive semi-definite. Fortunately, $\mathbf{A}^{(2)} \triangleq 3d\mathbf{I} + \mathbf{A}^{(1)}$ is positive semi-definite because it is symmetric and diagonally-dominant. The identity matrix adds exactly 1 to the objective function for any input. Therefore we also have

$$\text{“yes”:} \quad \max_{\substack{|\mathbf{z}|_2=1 \\ \mathbf{z}\in\left\{0,1/\sqrt{n}\right\}^{2n}}} \mathbf{z}^\intercal \mathbf{A}^{(2)} \mathbf{z} \geqslant 3d + c \qquad \text{“no”:} \quad \max_{\substack{|\mathbf{z}|_2=1 \\ \mathbf{z}\in\left\{0,1/\sqrt{n}\right\}^{2n}}} \mathbf{z}^\intercal \mathbf{A}^{(2)} \mathbf{z} \leqslant 3d + s \tag{1}$$

ENFORCING A (NEARLY) n-UNIFORM OPTIMUM

Now, we would of course like to replace $\left\{0, 1/\sqrt{n}\right\}^{2n}$ with the set of all n-sparse vectors, while maintaining (approximately) the same optima. Consider the positive semi-definite matrix $\mathbf{J} = \mathbf{1}\mathbf{1}^\intercal$; the objective $\mathbf{x}^\intercal \mathbf{J} \mathbf{x} = |\mathbf{x}|_1^2$ is maximized by an n-uniform vector in $\left\{0, 1/\sqrt{n}\right\}^{2n}$.

We define our final hard instance input matrix to be $\mathbf{A}^{(3)} \triangleq \frac{\alpha}{n}\mathbf{J} + \mathbf{A}^{(2)}$, for a sufficiently large (but constant) α. As we show below, the objective is now maximized by a vector $\mathbf{x}$ that is approximately n-uniform.

Formally, observe that $\mathbf{A}^{(2)}$ induces an inner product over $\mathbb{R}^n$; thus for any $|\mathbf{x}|_2^2 = |\mathbf{z}|_2^2 = 1$ we can use the Cauchy-Schwartz inequality to get:

$$\begin{aligned} \mathbf{x}^\intercal \mathbf{A}^{(2)} \mathbf{x} - \mathbf{z}^\intercal \mathbf{A}^{(2)} \mathbf{z} &= (\mathbf{x} - \mathbf{z})^\intercal \mathbf{A}^{(2)} (\mathbf{x} + \mathbf{z}) \\ &\leqslant \sqrt{(\mathbf{x} - \mathbf{z})^\intercal \mathbf{A}^{(2)} (\mathbf{x} - \mathbf{z})} \cdot \sqrt{(\mathbf{x} + \mathbf{z})^\intercal \mathbf{A}^{(2)} (\mathbf{x} + \mathbf{z})} \\ &\leqslant \left\|\mathbf{A}^{(2)}\right\|_2^2 \cdot |\mathbf{x} + \mathbf{z}|_2 \, |\mathbf{x} - \mathbf{z}|_2 \, , \end{aligned}$$

where $\left\|\mathbf{A}^{(2)}\right\|_2$ is the l^2 operator norm of $\mathbf{A}^{(2)}$, and is bounded by:

$$\left\|\mathbf{A}^{(2)}\right\|_2^2 = \max_{|\mathbf{x}|_2^2=1} \mathbf{x}^\intercal \mathbf{A}^{(2)} \mathbf{x} = 3d + \max_{|\mathbf{x}|_2^2=1} \mathbf{x}^\intercal \mathbf{A}^{(1)} \mathbf{x} < 5d + \max_{|\mathbf{x}|_2^2=1} \mathbf{x}^\intercal \mathbf{A}^{(0)} \mathbf{x} \leqslant 6d.$$

By triangle inequality, $|\mathbf{x} + \mathbf{z}|_2 \leqslant 2$, and therefore

$$\mathbf{x}^\intercal \mathbf{A}^{(2)} \mathbf{x} - \mathbf{z}^\intercal \mathbf{A}^{(2)} \mathbf{z} \leqslant 12d \, |\mathbf{x} - \mathbf{z}|_2 \, . \tag{2}$$

Suppose further that $\mathbf{z}$ is a rounding of $\mathbf{x}$ to $\left\{0, 1/\sqrt{n}\right\}^{2n}$. In particular, $\operatorname{supp}(\mathbf{x}) \subseteq \operatorname{supp}(\mathbf{z})$ (we have equality if $\mathbf{x}$ is exactly n-sparse) and $\mathbf{x}^\intercal \mathbf{z} = \lambda_\mathbf{z} \geqslant 0$. Let us decompose $\mathbf{x} = \lambda_\mathbf{z} \mathbf{z} + \lambda_\mathbf{w} \mathbf{w}$ where $\mathbf{w}$ is a unit-norm vector orthogonal to $\mathbf{z}$ (i.e. $|\mathbf{w}|_2 = 1$ and $\mathbf{w}^\intercal \mathbf{z} = 0$). Since all the vectors have unit norm, $\lambda_\mathbf{z}^2 + \lambda_\mathbf{w}^2 = \lambda_\mathbf{z}^2 |\mathbf{z}|_2^2 + \lambda_\mathbf{w}^2 |\mathbf{w}|_2^2 = |\mathbf{x}|_2^2 = 1$. We can now write the difference between $\mathbf{x}$ and $\mathbf{z}$ as,

$$\begin{aligned} |\mathbf{x} - \mathbf{z}|_2^2 &= \left|(1 - \lambda_\mathbf{z}) \, \mathbf{z} + \sqrt{1 - \lambda_\mathbf{z}^2} \cdot \mathbf{w}\right|_2^2 \\ &= (1 - \lambda_\mathbf{z})^2 + \left(1 - \lambda_\mathbf{z}^2\right) \\ &\leqslant 2\left(1 - \lambda_\mathbf{z}^2\right) . \end{aligned} \tag{3}$$

Since $\text{supp}(\mathbf{x}) \subseteq \text{supp}(\mathbf{z})$, we also have that $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{z})$. Thus $\mathbf{w}^\intercal \mathbf{z} = 0$ is equivalent to $\mathbf{w}^\intercal \mathbf{1} = 0$. We therefore have:

$$\begin{aligned}
\mathbf{z}^\intercal \mathbf{J} \mathbf{z} - \mathbf{x}^\intercal J \mathbf{x} &= \mathbf{z}^\intercal \mathbf{J} \mathbf{z} - (\lambda_\mathbf{z} \mathbf{z} + \lambda_\mathbf{w} \mathbf{w})^\intercal \mathbf{J} (\lambda_\mathbf{z} \mathbf{z} + \lambda_\mathbf{w} \mathbf{w}) \\
&\underbrace{=}_{\mathbf{w}^\intercal \mathbf{1} = 0} \mathbf{z}^\intercal \mathbf{J} \mathbf{z} - (\lambda_\mathbf{z} \mathbf{z})^\intercal \mathbf{J} (\lambda_\mathbf{z} \mathbf{z}) \\
&= \left(1 - \lambda_\mathbf{z}^2\right) \cdot \mathbf{z}^\intercal \mathbf{J} \mathbf{z} \\
&\underbrace{\geqslant}_{Eq. (3)} \frac{|\mathbf{x} - \mathbf{z}|_2^2}{2} \cdot \|\mathbf{J}\|_2 = \frac{n}{2} |\mathbf{x} - \mathbf{z}|_2^2 .
\end{aligned} \tag{4}$$

Recall that $\mathbf{A}^{(3)} = \frac{\alpha}{n}\mathbf{J} + \mathbf{A}^{(2)}$. Combining Eq. (2) and Eq. (4), we have that for every n-sparse, unit-norm $\mathbf{x}$,

$$\max_{\substack{|\mathbf{z}|_2^2 = 1 \\ \mathbf{z} \in \left\{0, 1/\sqrt{n}\right\}^{2n}}} \mathbf{z}^\intercal \mathbf{A}^{(3)} \mathbf{z}. - \mathbf{x}^\intercal \mathbf{A}^{(3)} \mathbf{x} \geqslant \frac{\alpha}{2} \cdot |\mathbf{x} - \mathbf{z}|_2^2 - 12d \cdot |\mathbf{x} - \mathbf{z}|_2 \tag{5}$$

Let $\alpha \triangleq 144d^2 / (c - s)$. Then,

$$\begin{aligned}
\frac{\alpha}{2} \cdot |\mathbf{x} - \mathbf{z}|_2^2 - 12d \cdot |\mathbf{x} - \mathbf{z}|_2 &= \left(\frac{72}{c - s}\right) d^2 \cdot |\mathbf{x} - \mathbf{z}|_2^2 - 12d \cdot |\mathbf{x} - \mathbf{z}|_2 \\
&= \left(\frac{2}{c - s}\right) \left(36d^2 \cdot |\mathbf{x} - \mathbf{z}|_2^2 - 6d \cdot |\mathbf{x} - \mathbf{z}|_2 (c - s) + \left(\frac{c - s}{2}\right)^2\right) - \frac{c - s}{2} \\
&= \left(\frac{2}{c - s}\right) \left(6d \cdot |\mathbf{x} - \mathbf{z}|_2 - \frac{c - s}{2}\right)^2 - \frac{c - s}{2} \\
&\geqslant -\frac{c - s}{2}.
\end{aligned}$$

Plugging into Eq. (5), we have

$$\max_{\substack{|\mathbf{z}|_2^2 = 1 \\ \mathbf{z} \in \left\{0, 1/\sqrt{n}\right\}^{2n}}} \mathbf{z}^\intercal \mathbf{A}^{(3)} \mathbf{z} \leqslant \max_{\substack{|\mathbf{x}|_2^2 = 1 \\ |\mathbf{x}|_0 \leqslant n}} \mathbf{x}^\intercal \mathbf{A}^{(3)} \mathbf{x} \leqslant \frac{c - s}{2} + \max_{\substack{|\mathbf{z}|_2^2 = 1 \\ \mathbf{z} \in \left\{0, 1/\sqrt{n}\right\}^{2n}}} \mathbf{z}^\intercal \mathbf{A}^{(3)} \mathbf{z}.$$

Finally, by Eq. (1), it is NP-hard to distinguish between:

$$\text{“yes”: } \max_{\substack{|\mathbf{x}|_2^2 = 1 \\ |\mathbf{x}|_0 \leqslant n}} \mathbf{x}^\intercal \mathbf{A}^{(3)} \mathbf{x} \geqslant \alpha + 3d + c \qquad \text{“no”: } \max_{\substack{|\mathbf{x}|_2^2 = 1 \\ |\mathbf{x}|_0 \leqslant n}} \mathbf{x}^\intercal \mathbf{A}^{(3)} \mathbf{x} \leqslant \alpha + 3d + \frac{c + s}{2}.$$

■

3.0.1. A REMARK ON HARDNESS FOR LOWER SPARSITY

In the reduction above, we set the sparsity parameter to $k = n/2$, but in typical applications we're interested in k which is much smaller than n. We note that for any polynomial sparsity parameter $k = n^\delta$ (for any constant $\delta > 0$), SparsePCA is still NP-hard to approximate by a straightforward padding argument. Furthermore, for any $k = \omega(\log n)$, one can use the same padding argument (with more padding) to show that SparsePCA doew not admit a PTAS assuming the exponential time hypothesis. Looking at smaller values of k, it is an interesting open problem whether SparsePCA can be approximated in fixed parameter time.

4. Small-Set Expansion hardness

Throughout this section, we will consider edge-weighted 1-regular graphs $G = (V, E)$, whose adjacency matrix/probability transition matrix $\mathbf{G}$ has every row sum equal to 1.

Recall that for a 1-regular graph $G = (V, E)$ on n vertices, the expansion of $S \subseteq V$ is

$$\Phi_G(S) \triangleq \frac{|E(S, V \setminus S)|}{|S|},$$

where $|E(S,T)| \triangleq \sum_{i\in S, j\in T} \mathbf{G}_{ij}$ denotes the total weight of edges with one end point in S and one end point in T. The expansion profile of G is

$$\Phi_G(\delta) \triangleq \min_{S:|S|\leqslant \delta n} \Phi_G(S).$$

Recall the Small-Set Expansion Hypothesis (Raghavendra and Steurer, 2010b)[2]:

Problem 6 (SSE(η, δ)) *Given a regular graph $G = (V, E)$, distinguish between the following two cases:*

1. *Yes: Some subset $S \subseteq V$ with $|S| = \delta n$ has $\Phi_G(S) \leqslant \eta$*
2. *No: Any set $S \subseteq V$ with $|S| \leqslant 2\delta n$ has $\Phi_G(S) \geqslant 1 - \eta$*

Conjecture 7 (Small-Set Expansion Hypothesis (Raghavendra and Steurer, 2010b)) *For any $\eta > 0$, there is $\delta > 0$ such that SSE(η, δ) is NP-hard.*

There is little consensus among researchers whether this conjecture is true. At any rate, if the conjecture turns out to be false, significantly new algorithmic or analytic ideas will be needed. See e.g. (Arora et al., 2010; Barak et al., 2012a) on efforts to refute the conjecture and pointers to the literature.

It is more convenient to work with the following version of Small-Set Expansion, where in the No case the subset size can be an arbitrarily large constant multiple of the subset size in the Yes case.

2. This formulation comes from the full version of the paper on Prasad Raghavendra's homepage. This formulation has a different soundness condition than the one in the conference version of (Raghavendra and Steurer, 2010b). Furthermore, (Raghavendra et al., 2012) shows that the two formulations are equivalent.

Problem 8 ($\mathsf{SSE}(\eta, \delta, M)$) *Given a regular graph $G = (V, E)$, distinguish between the following two cases:*

1. *Yes: Some subset $S \subseteq V$ with $|S| = \delta n$ has $\Phi_G(S) \leqslant \eta$*
2. *No: Any set $S \subseteq V$ with $|S| \leqslant M\delta n$ has $\Phi_G(S) \geqslant 1 - \eta$*

The following reduction in (Raghavendra et al., 2012, Proposition 5.8) shows that the two versions of Small-Set Expansion are equivalent.

Claim 9 *For all $\eta, \delta > 0$, $M \geqslant 1$, there is a polynomial time reduction from $\mathsf{SSE}(\eta/M, \delta)$ to $\mathsf{SSE}(\eta, \delta, M)$.*

We note that our statement is slightly different from (Raghavendra et al., 2012, Proposition 5.8), due to our different version of Small-Set Expansion Hypothesis, but the proof of the above claim is the same.

We will use the following lemma from (Raghavendra and Schramm, 2014). Here the lazy random walk $\mathbf{G}_{\text{lazy}}$ corresponds to staying at the current vertex with probability $1/2$, otherwise moving according to the probability transition matrix $\mathbf{G}$. Therefore the probability transition matrix is given by $\mathbf{G}_{\text{lazy}} \triangleq (\mathbf{I} + \mathbf{G})/2$. For any $t \in \mathbb{N}$, define the t-step lazy random walk as $\mathbf{G}^t_{\text{lazy}} \triangleq (\mathbf{G}_{\text{lazy}})^t$, and let G^t_{lazy} denote the corresponding graph.

Lemma 10 *(Raghavendra and Schramm, 2014, Lemma 13) For all $t \in \mathbb{N}$ and $\eta, \delta \in (0, 1]$,*

$$\Phi_{G^t_{\text{lazy}}}(\delta) \geqslant \min\left(1 - \left(1 - \frac{\Phi^2_G(4\delta/\eta)}{32}\right)^t, 1 - \eta\right).$$

We define $\mathsf{PSD\text{-}SSE}(\eta, \delta)$ as the special case of Theorem 6 where the adjacency matrix of the graph is positive semidefinite. We now show that this special case is again equivalent to the general case.

Theorem 11 *For any $\eta > 0$, there is $\eta' > 0$ such that for any $\delta > 0$, $\mathsf{SSE}(\eta', \delta)$ is polynomial-time reducible to $\mathsf{PSD\text{-}SSE}(\eta, \delta)$.*

Proof Fix $\eta > 0$. Thanks to Theorem 9, it suffices to reduce from Theorem 8. That is, we will show that there are $\eta' > 0$ and $M \geqslant 1$ such that $\mathsf{SSE}(\eta', \delta, M)$ is polynomial-time reducible to $\mathsf{PSD\text{-}SSE}(\eta, \delta)$.

We will assume $\eta \leqslant 1/2$ (if $\mathsf{PSD\text{-}SSE}(\eta, \delta)$ is hard then so is the same problem with larger η). Let $t \triangleq 128 \log(1/\eta)$, $\eta' \triangleq \min(\eta, 2\eta/t)$, $M \triangleq 4/\eta$.

The reduction takes an instance G of $\mathsf{SSE}(\eta', \delta, M)$ and outputs G^t_{lazy}. The lazy random walk matrix $\mathbf{G}_{\text{lazy}}$ is positive semidefinite, and hence so is $\mathbf{G}^t_{\text{lazy}}$. As a result, the output is an instance of $\mathsf{PSD\text{-}SSE}$.

Yes case: By (Raghavendra and Schramm, 2014, Lemma 12), for every subset S, $\Phi_{G^t_{\text{lazy}}}(S) \leqslant t\Phi_G(S)/2$. In particular, if G is a Yes case of $\mathsf{SSE}(\eta', \delta, M)$, then for some subset S of size δn, has $\Phi_G(S) \leqslant \eta'$, and thus also $\Phi_{G^t_{\text{lazy}}}(S) \leqslant \eta$.

No case: This follows from Theorem 10. Indeed, $\Phi_G(4\delta/\eta) \geqslant 1-\eta' \geqslant 1-\eta \geqslant 1/2$ by assumptions. Thus

$$\left(1-\frac{\Phi_G^2(4\delta/\eta)}{32}\right)^t \leqslant \left(1-\frac{1}{128}\right)^t \leqslant \exp(-t/128) \leqslant \eta.$$

By Theorem 10, $\Phi_{G^t_{\text{lazy}}}(\delta) \geqslant 1-\eta$, and G^t_{lazy} is a No instance of PSD-SSE(η, δ). ■

Let us mention that a variant of the previous lemma follows from the techniques of (Chan et al., 2015; Kwok and Lau, 2014), and in fact without making the graph lazy at all.

Given a PSD matrix $\mathbf{A}$ of size n, let us define the sparse PCA objective $\textsc{Val}_{\mathbf{A}}(\delta) \triangleq \max_{\|\mathbf{x}\|_2=1,\|\mathbf{x}\|_0\leqslant\delta n} \mathbf{x}^\intercal \mathbf{A}\mathbf{x}$.

We also need the local version of Cheeger–Alon–Milman inequality (Natarajan and Wu, 2014, Theorem 1.7).

Lemma 12 *Let $\mathbf{L} = \mathbf{I} - \mathbf{G}$ be the normalized Laplacian matrix of a regular graph G on n vertices. For any $\delta \leqslant 1/2$, let $\lambda_\delta = \min \mathbf{x}^\intercal\mathbf{L}\mathbf{x}/\mathbf{x}^\intercal\mathbf{x} \mid \|\mathbf{x}\|, \leqslant \delta\backslash$. Then*

$$\Phi_G(\delta) \leqslant \sqrt{(2-\lambda_\delta)\lambda_\delta}.$$

Theorem 13 *If G is a Yes instance of* PSD-SSE(η,δ)*, then* $\textsc{Val}_{\mathbf{G}}(\delta) \geqslant 1-\eta$*. If G is a No instance of* PSD-SSE(η,δ)*, then* $\textsc{Val}_{\mathbf{G}}(\delta) \leqslant \sqrt{1-(1-\eta)^2}$.

Proof **Yes case:** Let S be a subset with $|S| \leqslant \delta n$ and $\Phi_G(S) \leqslant \eta$. Consider the normalized indicator function $\mathbf{1}_S : V \to \mathbb{R}$ for S. $\mathbf{1}_S$ has at most δn non-zero entries, and by normalization, $\|\mathbf{1}_S\|_2 = 1$. Furthermore,

$$\mathbf{1}_S^\intercal\mathbf{G}\mathbf{1}_S = \frac{\sum_{i,j\in S}\mathbf{G}_{ij}}{|S|} = \frac{\sum_{i\in S}(1-\sum_{j\notin S}\mathbf{G}_{ij})}{|S|} = 1-\frac{\sum_{i\in S,j\notin S}\mathbf{G}_{ij}}{|S|} = 1-\Phi_G(S).$$

Therefore $\textsc{Val}_{\mathbf{G}}(\delta) \geqslant 1-\eta$.

No case: Let $\mathbf{x}$ be any δn-sparse vector. Then

$$\frac{\mathbf{x}^\intercal\mathbf{G}\mathbf{x}}{\mathbf{x}^\intercal\mathbf{x}} = 1-\frac{\mathbf{x}^\intercal\mathbf{L}\mathbf{x}}{\mathbf{x}^\intercal\mathbf{x}} \leqslant 1-\lambda_\delta,$$

where λ_δ is as defined in Theorem 12 and satisfies

$$\sqrt{(2-\lambda_\delta)\lambda_\delta} \geqslant \Phi_G(\delta) \geqslant 1-\eta.$$

Letting $\rho \triangleq 1-\lambda_\delta$, the previous inequality becomes $1-\rho^2 = (1+\rho)(1-\rho) \geqslant (1-\eta)^2$, and hence $\mathbf{x}^\intercal\mathbf{G}\mathbf{x}/\mathbf{x}^\intercal\mathbf{x} \leqslant \rho \leqslant \sqrt{1-(1-\eta)^2}$. ■

Theorem 11 implies SparsePCA is hard to solve within any constant factor C. Indeed, let $\eta \triangleq \min(1-\sqrt{1-1/4C^2}, 1/2)$. Theorems 7, 11 and 13 imply that given the matrix $\mathbf{G}$ in the output of Theorem 13, it is NP-hard to tell whether $\textsc{Val}_{\mathbf{G}}(\delta) \geqslant 1-\eta \geqslant 1/2$, or $\textsc{Val}_{\mathbf{G}}(\delta) \leqslant \sqrt{1-(1-\eta)^2} = 1/2C$.

5. SDP gap

Recall the SDP for sparse PCA proposed by (d'Aspremont et al., 2007b):

$$\begin{aligned} \max \quad & \mathrm{tr}(\mathbf{A}\mathbf{X}) \\ \text{such that} \quad & \mathrm{tr}(\mathbf{X}) = 1 \\ & \mathbf{1}^\intercal |\mathbf{X}| \mathbf{1} \leqslant k \\ & X \succcurlyeq 0 \end{aligned} \tag{6}$$

In this section, we will show that the SDP has a factor $\exp\exp(\Omega(\sqrt{\log\log n}))$ gap.

If $\mathbf{A}$ is the adjacency matrix of a graph, then the SDP is essentially identical to the SDP for small-set expansion in (Raghavendra et al., 2010). Gap instances for the latter problem therefore imply strong rank gap for sparse PCA, provided the adjacency matrix is PSD. A typical gap instance for small-set expansion SDP is the noisy hypercube of dimension $\log n$ with n vertices. It is not hard to see that its adjacency matrix leads to $(\log n)^{\Omega(1)}$ gap for sparse PCA SDP. Below we use a more sophisticated graph G that can be considered as a small induced subgraph of the noisy hypercube (even though formally G is not such a subgraph). This will lead to $\exp\exp(\Omega(\sqrt{\log\log n}))$ gap for sparse PCA SDP, where n is the number of vertices in this graph. This gap factor is super-polylogarithmic but sub-polynomial.

CONSTRUCTION

The gap instance $\mathbf{A}$ for the SDP is derived from the short code graph G from (Barak et al., 2012b), also known as the low-degree long code. Its vertex set is the Reed–Muller code RM(m, d) (evaluations of polynomials of (total) degree $\leqslant d$ over $\mathbb{F}_2$ in m variables $x_1, \ldots, x_m$). Two vertices are connected if their corresponding polynomials differ by a product of exactly d linearly independent affine forms. Call T the collection of all such affine forms. Therefore G is the Cayley graph on RM(m, d) with generating set T.

The matrix $\mathbf{A}$ will be the adjacency matrix for continuous-time random walk on G. That is, $\mathbf{A} = e^{-t(\mathbf{I}-\mathbf{G})}$ for some $t \geqslant 0$. Here we denote by $\mathbf{G}$ the probability transition matrix for the graph G. Therefore $\mathbf{G}$ is a matrix where every row and every column sum to 1. As in (Barak et al., 2012b), taking a continuous-time random walk significantly reduces the value of the quadratic form for sparse vectors. For our application, continuous-time random walk has the additional benefit that $\mathbf{A}$ is guaranteed to be PSD because $\mathbf{A}$ is the exponentiation of a real symmetric matrix.

It will be more convenient to transform Eq. (6) into the following SDP:

$$\begin{aligned} \max \quad & \mathsf{E}_f \langle \mathbf{w}_f, (\mathbf{A}\mathbf{w})_f \rangle \\ \text{such that} \quad & \mathsf{E}_f \langle \mathbf{w}_f, \mathbf{w}_f \rangle = 1 \\ & \mathsf{E}_{f,g} |\langle \mathbf{w}_f, \mathbf{w}_g \rangle| \leqslant \delta = k/n \end{aligned} \tag{7}$$

The SDPs in Eqs. (6) and (7) are indeed equivalent, because any SDP solution $\mathbf{X}$ to Eq. (6) is the (scaled) Gram matrix

$$\mathbf{X}_{f,g} = \langle \mathbf{w}_f, \mathbf{w}_g \rangle / n, \tag{8}$$

of some vectors $\mathbf{w}_f \in \mathbb{R}^n$, and vice versa.

Choice of parameters: m is a free parameter that all other parameters depend on. Let $\delta \triangleq 1/2^{m/2}$ be the fractional sparsity parameter. Let $\eta \triangleq \delta^{1/(4\log 3)}$ be the eigenvalue threshold. Let

$\varepsilon_2 = \min\{\varepsilon_1, 1/20\}$, where ε_1 is the constant from (Bhattacharyya et al., 2010, Theorem 1). Let $d \triangleq \log\log(1/\eta) + \log(1/\varepsilon_2) - 1$ be the degree of the Reed Muller code, and let $t \triangleq 2^{d-1}$ be the time parameter for the continuous random walk. Let $n \triangleq |\mathrm{RM}(m,d)| = 2^{\binom{m}{\leqslant d}}$ be the size of $\mathbf{A}$. Here $\binom{m}{\leqslant d} \triangleq \sum_{r \leqslant d} \binom{m}{r}$ denotes the number of ways to choose a subset of size $\leqslant d$ out of m elements. Let $k \triangleq n/2^{m/2}$.

Proposition 14 *The SDP in Eq. (7) has a solution of value* $1/e = \Omega(1)$.

Proof Let $\mathbf{w}_f$ by the standard embedding of $f \in \mathrm{RM}(m,d)$. That is, $\mathbf{w}_f : \mathbb{F}_2^m \to \mathbb{R}$ is the vector/function such that its x-coordinate is $\mathbf{w}_f(x) = (-1)^{f(x)} \in \pm\infty$ for $x \in \mathbb{F}_2^m$. This defines a solution to Eq. (8). In Eq. (8) and below, the inner product $\langle \cdot, \cdot \rangle$ on $\mathbb{F}_2^m \to \mathbb{R}$ is defined as $\langle \mathbf{w}, \mathbf{w}' \rangle \triangleq \mathsf{E}_{x \in \mathbb{F}_2^m} \mathbf{w}(x)\mathbf{w}'(x)$.

We now verify that $\mathbf{X}$ is a feasible solution to the SDP. As a Gram matrix, $\mathbf{X}$ is clearly PSD. Also

$$\mathsf{E}_{f \in \mathrm{RM}(m,d)} \langle \mathbf{w}_f, \mathbf{w}_f \rangle = \mathsf{E}_{f \in \mathrm{RM}(m,d)} \mathsf{E}_{x \in \mathbb{F}_2^m} [((-1)^{f(x)})^2] = 1,$$

and

$$\mathsf{E}_{f,g \in \mathrm{RM}(m,d)} |\langle \mathbf{w}_f, \mathbf{w}_g \rangle| = \mathsf{E}_{f,g \in \mathrm{RM}(m,d)} \left| \mathsf{E}_{x \in \mathbb{F}_2^m} (-1)^{f(x)-g(x)} \right| = \mathsf{E}_{h \in \mathrm{RM}(m,d)} \left| \mathsf{E}_{x \in \mathbb{F}_2^m} (-1)^{h(x)} \right|, \tag{9}$$

where in the last equality we let $h = f - g$. Using Cauchy–Schwarz, the right-hand-side is at most

$$\sqrt{\mathsf{E}_h \left(\mathsf{E}_{x \in F_2^m} (-1)^{h(x)} \right)^2} = \sqrt{\mathsf{E}_{x,y \in F_2^m} \mathsf{E}_h (-1)^{h(x)-h(y)}}. \tag{10}$$

We now analyze the term inside the square root. When $x \neq y$,

$$\mathsf{E}_h (-1)^{h(x)-h(y)} = 0,$$

thanks to pairwise independence of $\mathrm{RM}(m,d)$. When $x = y$ (which happens with probability $1/2^m$), the same expectation is 1. Therefore Eq. (10) is at most $1/2^{m/2}$, and so is Eq. (9). Then $\mathbf{X}$ satisfies the sparsity constraint with $k/n = 1/2^{m/2}$.

We now bound the SDP value. Let $\varphi_x(f) \triangleq (-1)^{f(x)}$. Then

$$\mathsf{E}_f \langle \mathbf{w}_f, (\mathbf{A}\mathbf{w})_f \rangle = \mathsf{E}_f \mathsf{E}_{x \in \mathbb{F}_2^m} (\varphi_x)(f)(\mathbf{A}\varphi_x)(f).$$

We claim that φ_x is an eigenfunction of $\mathbf{A}$ with eigenvalue $1/e$. Assuming this claim, the right-hand side becomes

$$(1/e) \cdot \mathsf{E}_f \mathsf{E}_{x \in \mathbb{F}_2^m} [(\varphi_x(f))^2] = 1/e,$$

giving an SDP solution of value $1/e$.

We now verify the claim. For every $x \in \mathbb{F}_2^m$, the function $\varphi_x(f) = (-1)^{f(x)}$ is an eigenvector of $\mathbf{G}$ because

$$(\mathbf{G}\varphi_x)(f) = \mathsf{E}_{g \in T} (-1)^{f(x)-g(x)} = \varphi_x(f) \cdot \mathsf{E}_{g \in T} (-1)^{g(x)}.$$

It has eigenvalue

$$\lambda_x \triangleq \mathsf{E}_{g \in T} (-1)^{g(x)} = 1 - 2 \mathop{\mathbb{P}}_{g \in T} [g(x) = 1] = 1 - 2^{1-d}.$$

Since $\mathbf{G}$ and $\mathbf{A}$ have the same eigenvectors, φ_x is also an eigenvector of $\mathbf{A}$ with eigenvalue

$$e^{-t(1-\lambda_x)} = e^{-t2^{1-d}} = 1/e.$$

■

Proposition 15 *Any k-sparse rank-1 solution $\mathbf{w} : \mathrm{RM}(m,d) \to \mathbb{R}$ to Eq. (7) has value $\leqslant \eta + (1/\eta)^{\log 3}\sqrt{k/n}$.*

Since the proof is quite technical, let us recall main ideas in (Barak et al., 2012b). Intuitively, the sparse PCA instance $\mathbf{A}$ has low value for rank-1 sparse vector for the following reason. The inner product space $V(G) \to \mathbb{R}$ can be decomposed into a sum of the subspace V_ℓ and its orthogonal complement $V_\ell^\perp$. One can show that V_ℓ does not contain any sparse vector (more precisely, has bounded 2-to-4 norm). Therefore any sparse vector must be essentially contained in (i.e. has large projection to) $V_\ell^\perp$. $V_\ell^\perp$ will be the span of eigenvectors of $\mathbf{A}$ whose eigenvalues are small, say at most a small positive number η. This ensures all sparse vectors have small objective value under the quadratic form, as desired.

Proof This is essentially Theorem 4.14 in (Barak et al., 2012b). Even though their statement only concerns $\prime, \infty$-valued sparse vectors, their proof also works for real-valued sparse vectors , as we now show.

SETTING UP THE FOURIER EXPANSION

Let $M \triangleq 2^m$ and $C = \mathrm{RM}(m,d)$. We first think of the elements of C as functions $\mathbb{F}_2^m \to \mathbb{F}_2$; later it will be more convenient to think of them as vectors in F_2^M. For $c_1, c_2 : \mathbb{F}_2^m \to \mathbb{F}_2$ denote the inner product $(c_1, c_2)_2 \triangleq \sum_{x\in\mathbb{F}_2^m} c_1(x)c_2(x) \pmod 2$

Denote by $C^\perp \triangleq \dashv \in \mathbb{F}_\in^{\mathcal{M}} \mid (\dashv, \rfloor)_\in = \prime$ for all $c \in C$ the orthogonal subspace of C.

Any function $\mathbf{w} : C \to \mathbb{R}$ has a Fourier expansion, as follows. For every coset $\alpha + C^\perp \in \mathbb{F}_2^M/C^\perp$, we choose an arbitrary representative α in $\alpha + C^\perp$, and let $\chi_\alpha(f) = (-1)^{(\alpha,f)_2}$ be its character. Its degree is $\deg_{\mathbb{R}}(\chi_\alpha) \triangleq \min_{c^\perp \in C^\perp} |\alpha + c^\perp|$, where $|\alpha|$ denotes the Hamming weight (i.e. number of non-zero coordinates) of α. (Do not confuse this degree with the degrees of polynomials in the Reed Muller code!) Any function $\mathbf{w} : C \to \mathbb{R}$ is a unique linear combination of characters $\chi_{\alpha_{\alpha\in\mathbb{F}_2^M/C^\perp}}$,

$$\mathbf{w}(f) = \sum_{\alpha\in\mathbb{F}_2^M/C^\perp} \hat{\mathbf{w}}(\alpha)\chi_\alpha(f),$$

where $\hat{\mathbf{w}}(\alpha) \triangleq \langle \chi_\alpha, \mathbf{w}\rangle$ is the Fourier transform of $\mathbf{w}$ over the abelian group C.

Set the character degree bound $\ell \triangleq \varepsilon_2 2^{d+1}$. Consider the subspace $V_\ell \triangleq \operatorname{span} \chi_\alpha \mid \deg_{\mathbb{R}}(\chi_\alpha) \leqslant \ell$ of functions of degree at most ℓ. Note that V_ℓ and $V_\ell^\perp$ are both invariant subspaces of $\mathbf{A}$.

Given any vector $\mathbf{w}$, we expand it as $\mathbf{w} = \mathbf{w}^\parallel + \mathbf{w}^\perp$ where $\mathbf{w}^\parallel \in V_\ell$ and $\mathbf{w}^\perp \in V_\ell^\perp$. Then

$$\langle \mathbf{w}, \mathbf{A}\mathbf{w}\rangle = \langle \mathbf{w}^\parallel, \mathbf{A}\mathbf{w}^\parallel\rangle + \langle \mathbf{w}^\perp, \mathbf{A}\mathbf{w}^\perp\rangle. \tag{11}$$

Below, we separately bound the contribution of $\langle \mathbf{w}^\parallel, \mathbf{A}\mathbf{w}^\parallel\rangle$ and $\langle \mathbf{w}^\perp, \mathbf{A}\mathbf{w}^\perp\rangle$.

THE LOW-DEGREE SUBSPACE V_ℓ

Consider the projection operator P_ℓ to the subspace V_ℓ. The p-to-q norm of P_ℓ is defined as

$$\|P_\ell\|_{p\to q} \triangleq \max_{\mathbf{w}:C\to\mathbb{R}} \frac{\|P_\ell \mathbf{w}\|_q}{\|\mathbf{w}\|_p},$$

where in the case of a function $\mathbf{w} : C \to \mathbb{R}$, we define $\|\mathbf{w}\|_p \triangleq \mathsf{E}_{x\in C}[|\mathbf{w}(x)|^p]^{1/p}$.

We use the following bound on the 2-to-4 norm of P_ℓ, from (Barak et al., 2012b, Lemma 4.9): For any $\ell < (2^{d-1}-1)/4$,

$$\|P_\ell\|_{2\to 4} \leqslant 3^{\ell/2}. \tag{12}$$

For any k-sparse vector $\mathbf{w} : C \to \mathbb{R}$, let $S = \S \in \mathcal{C} \mid \mathbf{w}(\S) \neq \prime$ be the set of nonzero entries. By Hölder's inequality,

$$\|\mathbf{w}\|_{4/3} = \|\mathbf{1}_S \cdot \mathbf{w}\|_{4/3} \leqslant \|\mathbf{1}_S\|_4 \|\mathbf{w}\|_2 = (k/n)^{1/4}\|\mathbf{w}\|_2. \tag{13}$$

Recall that $\mathbf{A} = e^{-t(\mathbf{I}-\mathbf{G})}$. Since $(\mathbf{I}-\mathbf{G})$ is PSD, we have that all of $\mathbf{A}$'s eigenvalues are at most 1, i.e. $\mathbf{I} \succcurlyeq \mathbf{A}$. Therefore,

$$\langle \mathbf{w}^{\|}, \mathbf{A}\mathbf{w}^{\|}\rangle \leqslant \|\mathbf{w}^{\|}\|_2^2 \leqslant \|P_\ell\|_{4/3\to 2}^2 \|\mathbf{w}\|_{4/3}^2,$$

Together with $\|P_\ell\|_{4/3\to 2} \leqslant \|P_\ell\|_{2\to 4}$ (Barak et al., 2012b, Lemma 4.2) and Eqs. (12) and (13), we get

$$\langle \mathbf{w}^{\|}, \mathbf{A}\mathbf{w}^{\|}\rangle \leqslant 3^\ell \sqrt{k/n}\|\mathbf{w}\|_2^2 = (1/\eta)^{\log 3}\sqrt{k/n}\|\mathbf{w}\|_2^2, \tag{14}$$

where the last equation follows from $3^\ell = 3^{\varepsilon_2 2^{d+1}} = 3^{\log(1/\eta)}$.

THE HIGH-DEGREE SUBSPACE $V_\ell^\perp$

We now bound the second term $\langle \mathbf{w}^\perp, \mathbf{A}\mathbf{w}^\perp\rangle$. $\mathbf{w}^\perp$ is a linear combination of characters of degree $> \ell$. Recall that T, the generating set of G, is the set of products of exactly d linearly independent affine forms. Any character χ_α is an eigenvector of $\mathbf{G}$ because

$$(\mathbf{G}\chi_\alpha)(f) = \mathsf{E}_{g\in T}[\chi_\alpha(f+g)] = \mathsf{E}_{g\in T}[(-1)^{(\alpha,f+g)_2}] = (-1)^{(\alpha,f)_2}\mathsf{E}_{g\in T}[(-1)^{(\alpha,g)_2}] = \chi_\alpha(f)\mathsf{E}_{g\in T}\chi_\alpha(g),$$

and its eigenvalue is

$$\lambda_\alpha \triangleq \mathsf{E}_{g\in T}\chi_\alpha(g) = \mathsf{E}_{g\in T}[(-1)^{(\alpha,g)_2}] = 1 - 2\mathop{\mathbb{P}}_{g\in T}[(\alpha,g)_2 = 1],$$

We now use a theorem about Reed Muller code testers to bound $\mathbb{P}_{g\in T}[(\alpha,g)_2 = 1]$. An important problem in the intersection of coding theory and property testing is as follows: given a code $C^\perp$ and a word α, query a small number of α's bits to decide whether α belongs to the code, or is far from the code. By "far" from the code, it is meant that it has a large Hamming distance from any $c^\perp \in C^\perp$. When $C^\perp$ is a Reed-Muller code, in particular RM$(m, m-d-1)$, this is equivalent to testing whether α is a low $(m-d-1)$ degree polynomial, or far from every low degree polynomial. A canonical test for this problem is as follows: pick a random $(m-d)$-dimensional affine subspace S_g, and test whether α restricted to this subset is a degree-$(m-d-1)$ polynomial.

It turns out that having degree $\leqslant m-d-1$ over S_g corresponds exactly to having $\sum_{x\in S_g} \alpha(x) = 0 \pmod 2$ (Bhattacharyya et al., 2010). (Proof sketch: any monomial of degree $\leqslant m-d-1$ does not contain at least one of the $m-d$ variables, and thus zeros out when we sum modulo 2 over that variable; in the other direction, there is only one homogenous full-degree monomial, and it is nonzero only on the all-ones input.)

Furthermore, picking a random $(m-d)$-dimensional affine subspace S_g corresponds precisely to picking a random $g \in T$ and letting $S_g \triangleq \{x : g(x) = 1\}$. (This is related to "dual codes"; see also (Alon et al., 2005).) In other words, the test is the same as verifying that $(\alpha, g)_2 = 0$.

Bhattacharyya et al. (Bhattacharyya et al., 2010) analyze the probability that the above test rejects polynomials that are far from the code, i.e. precisely the quantity $\mathbb{P}_{g\in T}[(\alpha, g)_2 = 1]$. Recall that the degree of χ_α was defined as the Hamming distance of α from $C^\perp$. By our assumption that $\chi_\alpha \in V_\ell^\perp$, we have that $\deg_{\mathbb{R}}(\chi_\alpha) \geqslant \ell = \varepsilon_2 2^{d+1}$; that is α disagrees with every $c^\perp \in C^\perp$ on at least $\varepsilon_2 2^{d+1}/2^m = \varepsilon_2 2^{-(m-d-1)}$-fraction of the entries. Therefore, by (Bhattacharyya et al., 2010, Theorem 1), we have that $\mathbb{P}_{g\in T}[(\alpha, g)_2 = 1] \geqslant \varepsilon_2$.

As a result, any χ_α with $\deg_{\mathbb{R}}(\chi_\alpha) \geqslant \varepsilon_2 2^{d+1}$ is also an eigenvector of $\mathbf{A}$ with eigenvalue

$$\mu_\alpha \triangleq e^{-t(1-\lambda_\alpha)} \leqslant e^{-\varepsilon_2 2^{d+1}} \leqslant \eta.$$

Therefore

$$\langle \mathbf{w}^\perp, \mathbf{A}\mathbf{w}^\perp \rangle = \sum_{\substack{\alpha\in\mathbb{F}_2^M/C^\perp \\ \deg_{\mathbb{R}}(\chi_\alpha)\geqslant \ell}} \mu_\alpha \hat{\mathbf{w}}(\alpha)^2 \leqslant \eta \sum_{\substack{\alpha\in\mathbb{F}_2^M/C^\perp \\ \deg_{\mathbb{R}}(\chi_\alpha)\geqslant \ell}} \hat{\mathbf{w}}(\alpha)^2 = \eta \|\mathbf{w}^\perp\|_2^2 \leqslant \eta \|\mathbf{w}\|_2^2. \tag{15}$$

Finally, Theorem 15 follows from Eqs. (11), (14) and (15) and the constraint $\|\mathbf{w}\|_2^2 \leqslant 1$. ∎

We remark that an alternative proof of the previous proposition (with a slightly different bound) can be obtained by combining Theorem 4.14 in (Barak et al., 2012b) and local Cheeger–Alon–Milman inequality (Natarajan and Wu, 2014, Theorem 1.7).

Theorem 16 *Let $\mathbf{A}$ be the matrix defined above. The SDP in Eq. (7) has an SDP solution of value $\Omega(1)$, but any rank-1 solution has value $1/\exp\exp(\Omega(\sqrt{\log\log n}))$.*

Proof The SDP solution is given in Theorem 14. On the other hand, Theorem 15 shows that any rank-1 solution has value $(k/n)^{\Omega(1)}$. Since $\log n = \binom{m}{\leqslant d}$, we have $\log\log n = (\log m)^2(1+o_m(1))$ and $(k/n)^{\Omega(1)} = 1/\exp(\Omega(m)) = 1/\exp\exp(\Omega(\sqrt{\log\log n}))$. ∎

6. Additive PTAS

To complete the approximability picture for SparsePCA, we briefly sketch the proof of the additive PTAS due to Asteris et al. (2015). The algorithm first approximates $\mathbf{A}$ with a low-rank sketch, and then finds approximate solutions via an ϵ-net search of the low dimensional space. (We note that a similar approach was previously presented in Alon et al. (2013), for the closely related problem of DkS on a PSD adjacency matrix.)

The existence of a low-rank sketch, due to Alon et al., is via an application of the Johnson-Lindenstrauss Lemma:

Lemma 17 (Alon et al. (2013)) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be PSD matrix with entries in $[-1, 1]$. Then, we can construct in polynomial time a PSD matrix $\mathbf{A}_\epsilon$ with rank $O(\frac{\log n}{\epsilon^2})$ such that*

$$|[\mathbf{A}]_{i,j} - [\mathbf{A}_\epsilon]_{i,j}| \leq \epsilon$$

for all i, j with high probability.

The above low-rank approximation to $\mathbf{A}$ preserves all k-sparse quadratic forms to within an additive error term:

$$|\mathbf{x}^\intercal \mathbf{A}\mathbf{x} - \mathbf{x}^\intercal \mathbf{A}_\epsilon \mathbf{x}| = \left| \sum_{i,j} \mathbf{x}_i \mathbf{x}_j ([\mathbf{A}]_{ij} - [\mathbf{A}_\epsilon]_{ij}) \right| \leq \epsilon \left| \sum_{i=1}^n |\mathbf{x}_i| \sum_{j=1}^n |\mathbf{x}_j| \right| = \epsilon \|\mathbf{x}\|_1^2 \leq \epsilon k. \quad (16)$$

Since $\mathbf{A}$ is PSD, one can rewrite $\mathbf{A} = \mathbf{B}^\intercal \mathbf{B}$, where $\mathbf{A}$'s low-rank property translates to $\mathbf{B}$ having few columns. Enumerating over an ϵ-net on the low dimension of $\mathbf{B}$ now gives, results in the following:

Lemma 18 (Asteris et al. (2015)) *Let $\mathbf{A}_d \in \mathbb{R}^{n \times n}$ be PSD matrix of rank d. Then, we can construct a vector $\mathbf{x}_d$, in time $O(\epsilon^{-d} \cdot n \log n)$, such that*

$$\mathbf{x}_d^\intercal \mathbf{A}_d \mathbf{x}_d > (1 - \epsilon) \cdot OPT_d.$$

Finally, combining the above two results gives the additive PTAS.

Theorem 19 (Asteris et al. (2015)) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be PSD matrix with entries in $[-1, 1]$. Then, we can compute in $n^{O(\mathsf{poly}(1/\epsilon))}$ time a k-sparse unit norm vector $\mathbf{x}_\epsilon$ such that*

$$\mathbf{x}_\epsilon^\intercal \mathbf{A} \mathbf{x}_\epsilon \geq OPT - \epsilon \cdot k$$

with high probability.

7. When the input matrix is not PSD

In this section, we briefly remark that although the SparsePCA optimization problem can be defined when $\mathbf{A}$ is not required to be PSD, no meaningful multiplicative approximation guarantees are possible (in polynomial time, assuming P $\neq$ NP).

Theorem 20 *When $\mathbf{A}$ is not positive semi-definite, it is NP-hard to decide whether the SparsePCA objective is positive or negative.*

Proof Let $\text{VAL}_{\mathbf{A}}(k) \triangleq \max_{\|\mathbf{x}\|_2 = 1, \|\mathbf{x}\|_0 \leqslant k} \mathbf{x}^T \mathbf{A} \mathbf{x}$. It is well known that solving the SparsePCA exactly is NP-hard even in the PSD case; i.e. it is NP-hard to distinguish between $\text{VAL}_{\mathbf{A}}(k) \geqslant c$ and $\text{VAL}_{\mathbf{A}}(k) \leqslant s$ for some (potentially very close) $c < s$. Consider the modified matrix $\mathbf{A}' = \mathbf{A} - \left(\frac{c+s}{2}\right) \cdot \mathbf{I}$. Conclude that it is NP-hard to distinguish $\text{VAL}_{\mathbf{A}'}(k) \geqslant \frac{c-s}{2}$ and $\text{VAL}_{\mathbf{A}'}(k) \leqslant \frac{s-c}{2}$. ■

Acknowledgments

We would like to thank Robert Krauthgamer, Tsz Chiu Kwok, Lap Chi Lau, Prasad Raghavendra, Aaron Schild, and anonymous reviewers for numerous suggestions and helpful discussions.

References

Noga Alon, Tali Kaufman, Michael Krivelevich, Simon Litsyn, and Dana Ron. Testing Reed-Muller codes. *IEEE Transactions on Information Theory*, 51(11):4032–4039, 2005. doi: 10.1109/TIT.2005.856958. URL http://dx.doi.org/10.1109/TIT.2005.856958.

Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri Weinstein. Inapproximability of densest κ-subgraph from average case hardness. *Unpublished manuscript*, 2011.

Noga Alon, Troy Lee, Adi Shraibman, and Santosh Vempala. The approximate rank of a matrix and its algorithmic applications: approximate rank. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 675–684. ACM, 2013.

Arash Amini and Martin Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. In *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, pages 2454–2458. IEEE, 2008.

Ei Ando, Toshio Nakata, and Masafumi Yamashita. Approximating the longest path length of a stochastic DAG by a normal distribution in linear time. *Journal of Discrete Algorithms*, 7(4):420 – 438, 2009. ISSN 1570-8667. doi: http://dx.doi.org/10.1016/j.jda.2009.01.001. URL http://www.sciencedirect.com/science/article/pii/S1570866709000033.

Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. In *Symposium on Foundations of Computer Science*, FOCS, pages 563–572, Washington, DC, USA, 2010. IEEE. ISBN 978-0-7695-4244-7.

Megasthenis Asteris, Dimitris Papailiopoulos, and Alexandros Dimakis. Nonnegative sparse PCA with provable guarantees. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1728–1736, 2014.

Megasthenis Asteris, Dimitris Papailiopoulos, Anastasios Kyrillidis, and Alexandros G. Dimakis. Sparse pca via bipartite matchings. *preprint*, 2015.

Boaz Barak, Fernando Guadalupe dos Santos Lins Brandão, Aram Wettroth Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Symposium on Theory of Computing*, STOC, pages 307–326, New York, NY, USA, 2012a. ACM. ISBN 978-1-4503-1245-5.

Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David Steurer. Making the long code shorter. In *FOCS*, pages 370–379, 2012b. URL http://dx.doi.org/10.1109/FOCS.2012.83. Lemma numbers refer to the full version available at arXiv:1111.0405 [cs.CC].

Q. Berthet and P. Rigollet. Complexity theoretic lower bounds for sparse principal component detection. *Journal of Machine Learning Research (JMLR)*, 30:1046–1066 (electronic), 2013a.

Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *Ann. Statist.*, 41(1):1780–1815, 2013b.

Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. Detecting high log-densities: an $O(n^{1/4})$ approximation for densest k-subgraph. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, (STOC) 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 201–210, 2010. doi: 10.1145/1806689.1806718. URL http://doi.acm.org/10.1145/1806689.1806718.

Aditya Bhaskara, Moses Charikar, Aravindan Vijayaraghavan, Venkatesan Guruswami, and Yuan Zhou. Polynomial integrality gaps for strong SDP relaxations of densest k-subgraph. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 388–405, 2012. URL http://portal.acm.org/citation.cfm?id=2095150&CFID=63838676&CFTOKEN=79617016.

Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of Reed-Muller codes. In *FOCS*, pages 488–497, 2010. URL http://dx.doi.org/10.1109/FOCS.2010.54.

Mark Braverman, Young Kun-Ko, Aviad Rubinstein, and Omri Weinstein. ETH Hardness for Densest-k-Subgraph with Perfect Completeness. volume abs/1504.08352, 2015a. URL http://arxiv.org/abs/1504.08352.

Mark Braverman, Young Kun-Ko, and Omri Weinstein. Approximating the best nash equilibrium in $n^{o(\log n)}$-time breaks the exponential time hypothesis. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 970–982, 2015b. doi: 10.1137/1.9781611973730.66. URL http://dx.doi.org/10.1137/1.9781611973730.66.

J. Cadima and I.T. Jolliffe. Loading and correlations in the interpretation of principle compenents. *Journal of Applied Statistics*, 22(2):203–214, 1995.

T. Tony Cai, Zongming Ma, and Yihong Wu. Sparse PCA: Optimal rates and adaptive estimation. *The Annals of Statistics*, 41(6):3074–3110, 12 2013a. doi: 10.1214/13-AOS1178. URL http://dx.doi.org/10.1214/13-AOS1178.

Tony Cai, Zongming Ma, and Yihong Wu. Optimal estimation and rank detection for sparse spiked covariance matrices. *arXiv preprint arXiv:1305.3235*, 2013b.

Siu On Chan, Tsz Chiu Kwok, and Lap Chi Lau. Random walks and evolving sets: Faster convergences and limitations. manuscript, 2015.

A. d'Aspremont, L. El Ghaoui, M.I. Jordan, and G.R.G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM review*, 49(3):434–448, 2007a.

Alexandre d'Aspremont, Laurent El Ghaoui, Michael I. Jordan, and Gert R. G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM Rev.*, 49(3):434–448, July 2007b. ISSN 0036-1445. URL http://dx.doi.org/10.1137/050645506.

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

Alexandre d'Aspremont, Francis R. Bach, and Laurent El Ghaoui. Approximation Bounds for Sparse Principal Component Analysis. *Math. Program.*, 148(1-2):89–110, 2014. doi: 10.1007/s10107-014-0751-7. URL http://dx.doi.org/10.1007/s10107-014-0751-7.

Yash Deshpande and Andrea Montanari. Sparse PCA via covariance thresholding. *arXiv preprint arXiv:1311.5179*, 2013.

Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.

Uriel Feige, Guy Kortsarz, and David Peleg. The dense k-subgraph problem. *Algorithmica*, 29 (3):410–421, 2001. doi: 10.1007/s004530010050. URL http://dx.doi.org/10.1007/s004530010050.

Chao Gao, Zongming Ma, and Harrison H. Zhou. Sparse CCA: Adaptive Estimation and Computational Barriers. *arXiv preprint arXiv:1409.8565*, 2014.

Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105–142, 1999.

I.T. Jolliffe. Rotation of principal components: choice of normalization constraints. *Journal of Applied Statistics*, 22(1):29–35, 1995.

I.T. Jolliffe, N.T. Trendafilov, and M. Uddin. A modified principal component technique based on the lasso. *Journal of Computational and Graphical Statistics*, 12(3):531–547, 2003.

M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *The Journal of Machine Learning Research*, 11:517–553, 2010.

H.F. Kaiser. The varimax criterion for analytic rotation in factor analysis. *Psychometrika*, 23(3): 187–200, 1958.

Subhash Khot. Improved inaproximability results for maxclique, chromatic number and approximate graph coloring. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 600–609, 2001. doi: 10.1109/SFCS.2001.959936. URL http://dx.doi.org/10.1109/SFCS.2001.959936.

Subhash Khot. Ruling out PTAS for graph min-bisection, dense k-subgraph, and bipartite clique. *SIAM Journal on Computing*, 36(4):1025–1071, 2006.

Robert Krauthgamer, Boaz Nadler, and Dan Vilenchik. Do semidefinite relaxations solve sparse PCA up to the information limit? *Annals of Probability*, 43:1300–1322, 2015.

Volodymyr Kuleshov. Fast algorithms for sparse principal component analysis based on rayleigh quotient iteration. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1418–1425, 2013.

Tsz Chiu Kwok and Lap Chi Lau. Lower bounds on expansions of graph powers. In *APPROX*, pages 313–324, 2014. URL http://dx.doi.org/10.4230/LIPIcs.APPROX-RANDOM.2014.313.

Tengyu Ma and Avi Wigderson. Sum-of-Squares Lower Bounds for Sparse PCA. In *NIPS*, 2015. URL http://arxiv.org/abs/1507.06370. To appear.

Zongming Ma. Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*, 41(2):772–801, 2013.

Malik Magdon-Ismail. NP-hardness and inapproximability of sparse PCA. *CoRR*, abs/1502.05675, 2015. URL http://arxiv.org/abs/1502.05675.

B. Moghaddam, Y. Weiss, and S. Avidan. Spectral bounds for sparse pca: Exact and greedy algorithms. *NIPS*, 18:915, 2006.

B. Moghaddam, Y. Weiss, and S. Avidan. Fast pixel/part selection with sparse eigenvectors. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007.

Abhiram Natarajan and Yi Wu. Computational complexity of certifying restricted isometry property. In *APPROX*, pages 371–380, 2014. URL http://dx.doi.org/10.4230/LIPIcs.APPROX-RANDOM.2014.371.

Prasad Raghavendra and Tselil Schramm. Gap amplification for small-set expansion via random walks. In *APPROX*, pages 381–391, 2014. URL http://dx.doi.org/10.4230/LIPIcs.APPROX-RANDOM.2014.381.

Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, (STOC) 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 755–764, 2010a. doi: 10.1145/1806689.1806788. URL http://dblp.uni-trier.de/rec/bib/conf/stoc/RaghavendraS10.

Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *STOC*, pages 755–764, New York, NY, USA, 2010b. ACM. ISBN 978-1-4503-0050-6. URL http://doi.acm.org/10.1145/1806689.1806792.

Prasad Raghavendra, David Steurer, and Prasad Tetali. Approximations for the isoperimetric and spectral profile of graphs and related parameters. In *STOC*, pages 631–640, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0050-6. URL http://doi.acm.org/10.1145/1806689.1806776.

Prasad Raghavendra, David Steurer, and Madhur Tulsiani. Reductions between expansion problems. In *CCC*, pages 64–73, Washington, DC, USA, 2012. IEEE Computer Society. ISBN 978-0-7695-4708-4. URL http://dx.doi.org/10.1109/CCC.2012.43.

Muli Safra. Lecture notes in computational complexity theory. URL http://www.tau.ac.il/~safra/Complexity/pdf/2SAT_Handouts.pdf.

H. Shen and J.Z. Huang. Sparse principal component analysis via regularized low rank matrix approximation. *Journal of multivariate analysis*, 99(6):1015–1034, 2008.

B.K. Sriperumbudur, D.A. Torres, and G.R.G. Lanckriet. Sparse eigen methods by D.C. programming. In *Proceedings of the 24th international conference on Machine learning*, pages 831–838. ACM, 2007.

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

Xiao-Tong Yuan and Tong Zhang. Truncated power method for sparse eigenvalue problems. *The Journal of Machine Learning Research*, 14(1):899–925, 2013.

Y. Zhang, A. d'Aspremont, and L.E. Ghaoui. Sparse PCA: Convex relaxations, algorithms and applications. *Handbook on Semidefinite, Conic and Polynomial Optimization*, pages 915–940, 2012.

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

David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007. doi: 10.4086/toc.2007.v003a006. URL http://dx.doi.org/10.4086/toc.2007.v003a006.