

Sum-of-Squares Lower Bounds for Sparse PCA

Tengyu Ma^{*1} and Avi Wigderson^{†2}

¹Department of Computer Science, Princeton University

²School of Mathematics, Institute for Advanced Study

October 17, 2015

Abstract

This paper establishes a statistical versus computational trade-off for solving a basic high-dimensional machine learning problem via a basic convex relaxation method. Specifically, we consider the *Sparse Principal Component Analysis* (Sparse PCA) problem, and the family of *Sum-of-Squares* (SoS, aka Lasserre/Parillo) convex relaxations. It was well known that in large dimension p , a planted k -sparse unit vector can be *in principle* detected using only $n \approx k \log p$ (Gaussian or Bernoulli) samples, but all *efficient* (polynomial time) algorithms known require $n \approx k^2$ samples. It was also known that this quadratic gap cannot be improved by the most basic *semi-definite* (SDP, aka spectral) relaxation, equivalent to a degree-2 SoS algorithms. Here we prove that also degree-4 SoS algorithms cannot improve this quadratic gap. This average-case lower bound adds to the small collection of hardness results in machine learning for this powerful family of convex relaxation algorithms. Moreover, our design of moments (or “pseudo-expectations”) for this lower bound is quite different than previous lower bounds. Establishing lower bounds for higher degree SoS algorithms for remains a challenging problem.

1 Introduction

We start with a general discussion of the tension between sample size and computational efficiency in statistical and learning problems. We then describe the concrete model and problem at hand: Sum-of-Squares algorithms and the Sparse-PCA problem. All are broad topics studied from different viewpoints, and the given references provide more information.

1.1 Statistical vs. computational sample-size

Modern machine learning and statistical inference problems are often high dimensional, and it is highly desirable to solve them using far less samples than the ambient dimension. Luckily, we often know, or assume, some underlying structure of the objects sought, which allows such savings *in principle*. Typical such assumption is that the number of *real* degrees of freedom is far smaller than the dimension; examples include sparsity constraints for vectors, and low rank for matrices and tensors. The main difficulty that occurs in nearly all these problems is that while information

^{*}Supported in part by Simons Award for Graduate Students in Theoretical Computer Science

[†]Supported in part by NSF grant CCF-1412958

theoretically the sought answer is present (with high probability) in a small number of samples, actually computing (or even approximating) it from these many samples is a computationally hard problem. It is often expressed as a non-convex optimization program which is NP-hard in the worst case, and seemingly hard even on random instances.

Given this state of affairs, *relaxed* formulations of such non-convex programs were proposed, which can be solved efficiently, but sometimes to achieve accurate results seem to require far more samples than existential bounds provide. This phenomenon has been coined the “statistical versus computational trade-off” by Chandrasekaran and Jordan [CJ13], who motivate and formalize one framework to study it in which efficient algorithms come from the Sum-of-Squares family of convex relaxations (which we shall presently discuss). They further give a detailed study of this trade-off for the basic *de-noising problem* [Joh02, Don95, DJ98] in various settings (some exhibiting the trade-off and others that do not). This trade-off was observed in other practical machine learning problems, in particular for the Sparse PCA problem that will be our focus, by Berthet and Rigollet [BR13a].

As it turns out, the study of the same phenomenon was proposed even earlier in computational complexity, primarily from theoretical motivations. Decatur, Goldreich and Ron [DGR97] initiate the study of “computational sample complexity” to study statistical versus computation trade-offs in sample-size. In their framework efficient algorithms are arbitrary polynomial time ones, not restricted to any particular structure like convex relaxations. They point out for example that in the distribution-free PAC-learning framework of Vapnik-Chervonenkis and Valiant, there is often no such trade-off. The reason is that the number of samples is essentially determined (up to logarithmic factors, which we will mostly ignore here) by the VC-dimension of the given concept class learned, and moreover, an “Occam algorithm” (computing *any* consistent hypothesis) suffices for classification from these many samples. So, in the many cases where efficiently finding a hypothesis consistent with the data is possible, enough samples to learn are enough to do so efficiently! This paper also provide examples where this is not the case in PAC learning, and then turns to an extensive study of possible trade-offs for learning various concept classes under the uniform distribution. This direction was further developed by Servedio [Ser00].

The fast growth of Big Data research, the variety of problems successfully attacked by various heuristics and the attempts to find efficient algorithms with provable guarantees is a growing area of interaction between statisticians and machine learning researchers on the one hand, and optimization and computer scientists on the other. The trade-offs between sample size and computational complexity, which seems to be present for many such problems, reflects a curious “conflict” between these fields, as in the first more data is good news, as it allows more accurate inference and prediction, whereas in the second it is bad news, as a larger input size is a source of increased complexity and inefficiency. More importantly, understanding this phenomenon can serve as a guide to the design of better algorithms from both a statistical and computational viewpoints, especially for problems in which data acquisition itself is costly, and not just computation. A basic question is thus for which problems is such trade-off inherent, and to establish the limits of what is achievable by efficient methods.

Establishing a trade-off has two parts. One has to prove an existential, information theoretic upper bound on the number of samples needed when efficiency is not an issue, and then prove a computational lower bound on the number of samples for the class of efficient algorithms at hand. Needless to say, it is desirable that the lower bounds hold for as wide a class of algorithms as possible, and that it will match the best known upper bound achieved by algorithms from this class. The most general one, the computational complexity framework of [DGR97, Ser00] allows all polynomial-time

algorithms. Here one cannot hope for unconditional lower bounds, and so existing lower bounds rely on computational assumptions, e.g. “cryptographic assumptions”, e.g. that factoring integers has no polynomial time algorithm, or other average case assumptions. For example, hardness of refuting random 3CNF was used for establishing the sample-computational tradeoff for learning halfspaces [DLS13], and hardness of finding planted clique in random graphs was used for tradeoff in sparse PCA [BR13a, GMZ14]. On the other hand, in frameworks such as [CJ13], where the class of efficient algorithms is more restricted (e.g. a family of convex relaxations), one can hope to prove unconditional lower bounds, which are called “integrality gaps” in the optimization and algorithms literature. Our main result is of this nature, adding to the small number of such lower bounds for machine learning problems.

We now turn to describe and motivate SoS convex relaxations algorithms, and then the Sparse PCA problem.

1.2 Sum-of-Squares convex relaxations

Sum-of-Squares algorithms (sometimes called the Lasserre hierarchy) encompasses perhaps the strongest known algorithmic technique for a diverse set of optimization problems. It is a family of convex relaxations introduced independently around the year 2000 by Lasserre [Las01], Parrillo [Par00], and in the (equivalent) context of proof systems by Grigoriev [Gri01b]. These papers followed better and better understanding in real algebraic geometry [Art27, Kri64, Ste74, Sho87, Sch91, Put93, Nes00] of David Hilbert’s famous 17th problem on certifying the non-negativity of a polynomial by writing it as a *sum of squares* (which explains the name of this method). We only briefly describe this important class of algorithms; far more can be found in the book [Las15] and the excellent extensive survey [Lau09].

The SoS method provides a principled way of adding constraints to a linear or convex program in a way that obtains tighter and tighter convex sets containing all solutions of the original problem. This family of algorithms is parametrized by their *degree d* (sometimes called the *number of rounds*); as d gets larger, the approximation becomes better, but the running time becomes slower, specifically $n^{O(d)}$. Thus in practice one hopes that small degree (ideally constant) would provide sufficiently good approximation, so that the algorithm would run in polynomial time. This method extends the standard semi-definite relaxation (SDP, sometimes called spectral), that is captured already by degree-2 SoS algorithms. Moreover, it is more powerful than two earlier families of relaxations: the Sherali-Adams [SA90] and Lovász-Scrijver [LS91] hierarchies.

The introduction of these algorithms has made a huge splash in the optimization community, and numerous applications of it to problems in diverse fields were found that greatly improve solution quality and time performance over all past methods. For large classes of problems they are considered the strongest algorithmic technique known. Relevant to us is the very recent growing set of applications of constant-degree SoS algorithms to machine learning problems, such as [BKS15, BKS14, BM15]. The survey [BS14] contains some of these exciting developments. Section 2.3 contains some self-contained material about the general framework SoS algorithms as well.

Given their power, it was natural to consider proving lower bounds on what SoS algorithms can do. There has been an impressive progress on SoS degree lower bounds (via beautiful techniques) for a variety of combinatorial optimization problems [Gri01a, Gri01b, Sch08, MPW15]. However, for machine learning problems relatively few such lower bounds (above SDP level) are known [BM15, WGL15] and follow via reductions to the above bounds. So it is interesting to enrich the set of techniques for proving such limits on the power of SoS for ML. The lower bound we prove indeed

seem to follow a different route than previous such proofs.

1.3 Sparse PCA

Sparse principal component analysis, the version of the classical PCA problem which assumes that the direction of variance of the data has a sparse structure, is by now a central problem of high-dimensional statistical analysis. In this paper we focus on the single-spiked covariance model introduced by Johnstone [Joh01]. One observes n samples from p -dimensional Gaussian distribution with covariance

$$\Sigma = \lambda vv^T + I \tag{1.1}$$

where (the *planted* vector) v is assumed to be a unit-norm *sparse* vector with at most k non-zero entries, and $\lambda > 0$ represents the strength of the signal. The task is to find (or estimate) the sparse vector v . More general versions of the problem allow several sparse directions/components and general covariance matrix [Ma13, VL13]. Sparse PCA and its variants have a wide variety of applications ranging from signal processing to biology: see, e.g., [ABN⁺99, JL09, Che11, JOB10].

The hardness of Sparse PCA, at least in the worst case, can be seen through its connection to the (NP-hard) Clique problem in graphs. Note that if Σ is a $\{0, 1\}$ adjacency matrix of a graph (with 1's on the diagonal), then it has a k -sparse eigenvector v with eigenvalue k if and only if the graph has a k -clique. This connection between these two problems is actually deeper, and will appear again below, for our real, average case version above.

From a theoretical point of view, Sparse PCA is one of the simplest examples where we observe a gap between the number of samples needed information theoretically and the number of samples needed for a polynomial time estimator: It has been well understood [VL12, PJ12, BR13b] that information theoretically, given $n = O(k \log p)$ samples¹, one can estimate v up to constant error (in euclidean norm), using a non-convex (therefore not polynomial time) optimization algorithm. On the other hand, all the existing provable polynomial time algorithms [JL09, AW09, VL13, DM14], which use either diagonal thresholding (for the single spiked model) or semidefinite programming (for general covariance), first introduced for this problem in [dGJL07], need at least quadratically many samples to solve the problem, namely $n = O(k^2)$. Moreover, Krauthgamer, Nadler and Vilenchik [KNV15] and Berthet and Rigollet [BR13b] have shown that for semi-definite programs (SDP) this bound is tight. Specifically, the natural SDP cannot even solve the *detection problem*: to distinguish the data in equation 1.1 above from the null hypothesis in which no sparse vector is planted, namely the n samples are drawn from the Gaussian distribution with covariance matrix I .

Recall that the natural SDP for this problem (and many others) is just the first level of the SoS hierarchy, namely degree-2. Given the importance of the Sparse PCA, it is an intriguing question whether one can solve it efficiently with far fewer samples by allowing degree- d SoS algorithms with larger d . A very interesting *conditional* negative answer was suggested by Berthet and Rigollet [BR13b]. They gave an efficient reduction from *Planted Clique*² problem to Sparse PCA, which shows in particular that degree- d SoS algorithms for Sparse PCA will imply similar ones for Planted Clique. Gao, Ma and Zhou [GMZ14] strengthen the result by establishing the hardness of the Gaussian single-spiked covariance model, which is an interesting subset of models considered by [BR13a]. These are useful as nontrivial constant-degree SoS lower bounds for

¹We treat λ as a constant so that we omit the dependence on it for simplicity throughout the introduction section

²An average case version of the Clique problem in which the input is a random graph in which a much larger than expected clique is planted.

Planted Clique were recently proved by [MPW15, DM15] (see there for the precise description, history and motivation for Planted Clique). As [BR13b, GMZ14] argue, strong yet *believed* bounds, if true, would imply that the quadratic gap is tight for any constant d . Before the submission of this paper, the known lower bounds above for planted clique were not strong enough yet to yield any lower bound for Sparse PCA beyond the minimax sample complexity. We also note that the recent progress [RS15, HKP15] that show the tight lower bounds for planted clique, together with the reductions of [BR13a, GMZ14], also imply the tight lower bounds for Sparse PCA, as shown in this paper.

1.4 Our contribution

We give a direct, unconditional lower bound proof for computing Sparse PCA using degree-4 SoS algorithms, showing that they too require $n = \tilde{\Omega}(k^2)$ samples to solve the detection problem (Theorem 3.1), which is tight up to polylogarithmic factors when the strength of the signal λ is a constant. Indeed the theorem gives a lower bound for every strength λ , which becomes weaker as λ gets larger. Our proof proceeds by constructing the necessary pseudo-moments for the SoS program that achieve too high an objective value (in the jargon of optimization, we prove an “integrality gap” for these programs). As usual in such proofs, there is tension between having the pseudo-moments satisfy the constraints of the program and keeping them positive semidefinite (PSD). Differing from past lower bound proofs, we construct *two* different PSD moments, each approximately satisfying one sets of constraints in the program and is negligible on the rest. Thus, their sum give PSD moments which *approximately* satisfy all constraints. We then perturb these moments to satisfy constraints *exactly*, and show that with high probability over the random data, this perturbation leaves the moments PSD.

We note several features of our lower bound proof which makes the result particularly strong and general. First, it applies not only for the Gaussian distribution, but also for Bernoulli and other distributions. Indeed, we give a set of natural (pseudorandomness) conditions on the sampled data vectors under which the SoS algorithm is “fooled”, and show that these conditions are satisfied with high probability under many similar distributions (possessing strong concentration of measure). Next, our lower bound holds even if the hidden sparse vector is discrete, namely its entries come from the set $\{0, \pm \frac{1}{\sqrt{k}}\}$. We also extend the lower bound for the detection problem to apply also to the estimation problem, in the regime when the ambient dimension is linear in the number of samples, namely $n \leq p \leq Bn$ for constant B (see Theorem 3.2).

Organization: Section 2 provides more backgrounds of sparse PCA and SoS algorithms. Then we state our main results in Section 3. In Section 4, we design the pseudo-moments and state their properties and then in Section 5 we prove our main theorems using these moments. Section 6 and 7 contain the analysis of the moments. Section 8 lists the tools that we heavily used for proving concentration inequalities in the analysis. Finally we conclude with a discussion of further directions of study in Section 9.

2 Formal description of the model and problem

Notation: We use $\|\cdot\|$ to denote the euclidean norm of a vector and spectral norm of a matrix, $\|\cdot\|_q$ to denote the q -norm of a vector, and $|\cdot|_0$ is the number of nonzero entries of a vector. We use $[m]$ to denote the set of integers $\{1, \dots, m\}$.

We write $M \succeq 0$ if M is a positive semidefinite matrix.

$\mathbb{R}_n[x]_d$ is used to denote the set of real polynomials with n variables and degree at most d . We will drop the subscript n when it is clear from context. We will assume that n, k, p are all sufficiently large³, and that $n \leq p$.

Throughout this paper, by “with high probability some event happens”, we mean the failure probability is bounded by p^{-c} for every constant c , as p tends to infinity.

We use the asymptotic notation $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ to hide the logarithmic dependency (in p). That is, $m \leq \tilde{O}(f(n, p, k))$ means that there exists universal constant $r \geq 0$ (which is less than 3 typically in this paper) and C such that $m \leq C f(n, p, k) \log^r p$, and $m \geq \tilde{\Omega}(f(n, p, k))$ means that there exist constants r and c such that $m \geq c f(n, p, k) / \log^r p$.

2.1 Sparse PCA estimation and detection problems

We will consider the simplest setting of sparse PCA, which is called single-spiked covariance model in literature [Joh01] (note that restricting to a special case makes our lower bound hold in all generalizations of this simple model). In this model, the task is to recover a single sparse vector from noisy samples as follows. The “hidden data” is an unknown k -sparse vector $v \in \mathbb{R}^p$ with $|v|_0 = k$ and $\|v\| = 1$. To make the task easier (and so the lower bound stronger), we even assume that v has discrete entries, namely that $v_i \in \{0, \pm \frac{1}{\sqrt{k}}\}$ for all $i \in [p]$. We observe n noisy samples $X^1, \dots, X^n \in \mathbb{R}^p$ that are generated as follows. Each is independently drawn as

$$X^j = \sqrt{\lambda} g^j v + \xi^j \quad (2.1)$$

from a distribution which generalizes both Gaussian and Bernoulli noise to v . Namely, the g^j 's are i.i.d real random variable with mean 0 and variance 1, and ξ^j 's are i.i.d random vectors which have independent entries with mean zero and variance 1. Therefore under this model, the covariance of X^i is equal to $\lambda v v^T + I$. Moreover, we assume that g^j and entries of ξ^j are sub-gaussian⁴ with variance proxy $O(1)$. Given these samples, the *estimation* problem is to approximate the unknown sparse vector v .

It is also interesting to also consider the sparse component *detection* problem [BR13b, BR13a], which is the decision problem of distinguishing from random samples the following two distributions

H_0 : data $X^j = \xi^j$ is purely random

H_v : data $X^j = \xi^j + \sqrt{\lambda} g^j v$ contains a hidden sparse signal with strength λ .

Rigollet [MR14] observed that a polynomial time algorithm for estimation version of sparse PCA with constant error implies that an algorithm for the detection problem with twice number of the samples. Thus, for polynomial time lower bounds, it suffices to consider the detection problem.

We will use X as a shorthand for the $p \times n$ matrix $[X^1, \dots, X^n]$. We denote the rows of X as X_1^T, \dots, X_p^T , therefore X_i 's are n -dimensional column vectors. The empirical covariance matrix is defined as $\hat{\Sigma} = \frac{1}{n} X X^T$.

³Or we assume that they go to infinity as typically done in statistics.

⁴A real random variable X is subgaussian with variance proxy σ^2 if it has similar tail behavior as gaussian distribution with variance σ^2 . More formally, if for any $t \in \mathbb{R}$, $\mathbb{E}[\exp(tX)] \leq \exp(t^2 \sigma^2 / 2)$

2.2 Statistically optimal estimator/detector

It is well known that the following non-convex program achieves optimal statistical minimax rate for the estimation problem and the optimal sample complexity for the detection problem. Note that we scale the variables x up by a factor of \sqrt{k} for simplicity (the hidden vector now has entries from $\{0, \pm 1\}$).

$$\lambda_{\max}^k(\hat{\Sigma}) = \frac{1}{k} \cdot \max \quad \langle \hat{\Sigma}, xx^T \rangle \quad (2.2)$$

$$\text{subject to} \quad \|x\|_2^2 = k \quad (2.3)$$

$$\|x\|_0 = k \quad (2.4)$$

Proposition 2.1 ([AW09], [BR13b], [VL12] informally stated). The non-convex program (2.2) statistically optimally solves the sparse PCA problem when $n \geq Ck/\lambda^2 \log p$ for some sufficiently large C . Namely, the following hold with high probability. If X is generated from H_v , then optimal solution x_{opt} of program (2.2) satisfies $\|\frac{1}{k} \cdot x_{\text{opt}} x_{\text{opt}}^T - vv^T\| \leq \frac{1}{3}$, and the objective value $\lambda_{\max}^k(\hat{\Sigma})$ is at least $1 + \frac{2\lambda}{3}$. On the other hand, if X is generated from null hypothesis H_0 , then $\lambda_{\max}^k(\hat{\Sigma})$ is at most $1 + \frac{\lambda}{3}$.

Therefore, for the detection problem, one can simply use the test $\lambda_{\max}^k(\hat{\Sigma}) > 1 + \frac{\lambda}{2}$ to distinguish the case of H_0 and H_v , with $n = \tilde{\Omega}(k/\lambda^2)$ samples. However, this test is highly inefficient, as the best known ways for computing $\lambda_{\max}^k(\hat{\Sigma})$ take exponential time! We now turn to consider efficient ways of solving this problem.

2.3 Sum of Squares (Lasserre) Relaxations

Here we will only briefly introduce the basic ideas of Sum-of-Squares (Lasserre) relaxation that will be used for this paper. We refer readers to the extensive [Las15, Lau09, BS14] for detailed discussions of sum of squares algorithms and proofs and their applications to algorithm design.

Let $\mathbb{R}[x]_d$ denote the set of all real polynomials of degree at most d with n variables x_1, \dots, x_n . We start by defining the notion of *pseudo-moment* (sometimes called *pseudo-expectation*). The intuition is that these pseudo-moments behave like the actual first d moments of a real probability distribution.

Definition 2.2 (pseudo-moment). A degree- d pseudo-moments M is a linear operator that maps $\mathbb{R}[x]_d$ to \mathbb{R} and satisfies $M(1) = 1$ and $M(p^2(x)) \geq 0$ for all real polynomials $p(x)$ of degree at most $d/2$.

For a multi-set $S \subset [n]$, we use x^S to denote the monomial $\prod_{i \in S} x_i$. Since M is a linear operator, it can be clearly described by all the values of M on the monomial of degree d , that is, all the values of $M(x^S)$ for multi-set S of size at most d uniquely determines M . Moreover, the nonnegativity constraint $M(p(x)^2) \geq 0$ is equivalent to the positive semidefiniteness of the matrix-form (as defined below), and therefore the set of all pseudo-moments is convex.

Definition 2.3 (matrix-form). For an even integer d and any degree- d pseudo-moments M , we define the matrix-form of M as the trivial way of viewing all the values of M on monomials as a matrix: we use $\text{mat}(M)$ to denote the matrix that is indexed by multi-subset S of $[n]$ with size at most $d/2$, and $\text{mat}(M)_{S,T} = M(x^S x^T)$.

Given polynomials $p(x)$ and $q_1(x), \dots, q_m(x)$ of degree at most d , and a polynomial program,

$$\begin{aligned} & \text{Maximize} && p(x) \\ & \text{Subject to} && q_i(x) = 0, \forall i \in [m] \end{aligned} \tag{2.5}$$

We can write a sum of squares based relaxation in the following way: Instead of searching over $x \in \mathbb{R}^n$, we search over all the possible “pseudo-moments” M of a hypothetical distribution over solutions x , that satisfy the constraints above. The key of the relaxation is to consider only moments up to degree d . Concretely, we have the following semidefinite program in roughly n^d variables.

$$\begin{aligned} & \text{Variables} && M(x^S) && \forall S : |S| \leq d \\ & \text{Maximize} && M(p(x)) \\ & \text{Subject to} && M(q_i(x)x^K) = 0 && \forall i, K : |K| + \deg(q_i) \leq d \\ & && \text{mat}(M) \succeq 0 \end{aligned} \tag{2.6}$$

Note that (2.6) is a valid relaxation because for any solution x_* of (2.5), if we define $M(x^S)$ to be $M(x^S) = x_*^S$, then M satisfies all the constraints and the objective value is $p(x_*)$. Therefore it is guaranteed that the optimal value of (2.6) is always larger than that of (2.5).

Finally, the key point is that this program can be solved efficiently, in polynomial time in its size, namely in time $n^{O(d)}$. As d grows, the constraints added make the “pseudo-distribution” defined by the moments closer and closer to an actual distribution, thus providing a tighter relaxation, at the cost of a larger running time to solve it.

In the next section we apply this relaxation to the Sparse PCA problem and state our results.

3 Main Results

To exploit the sum of squares relaxation framework as described in Section 2.3], we first convert the statistically optimal estimator/detector (2.2) into the “polynomial” program version below.

$$\text{Maximize} \quad \langle \hat{\Sigma}, xx^T \rangle \tag{3.1}$$

$$\text{subject to} \quad \|x\|_2^2 = k \tag{3.2}$$

$$x_i^3 = x_i, \forall i \in [p] \tag{3.3}$$

$$|x|_1 \leq k \tag{3.4}$$

Note that the non-convex sparsity constraint (2.4) is replaced by the polynomial constraint 3.3, which ensures that any solution vector x has entries in $\{0, \pm 1\}$, and so together with the constraint (3.2) guarantees that it has precisely k non-zero entries, each of absolute value 1. Note that constraint (3.3) implies other natural constraints that one may add to the program in order to make it stronger: for example, the upper bound on each entry x_i , the lower bound on the non-zero entries of x_i , and the constraint $\|x\|^4 \geq k$ which has been used as a surrogate for k -sparse vectors in [BKS14, BKS15]. Note that we have also added an ℓ_1 sparsity constraint (3.4) (which can be easily made into a polynomial constraint) as is often used in practice and makes our lower bound even stronger. Of course, it is formally implied by the other constraints, but not in low-degree SoS.

Now we are ready to apply the sum-of-squares relaxation scheme described in Section 2.3) to the polynomial program above as . For degree-4 relaxation we obtain the following semidefinite

program $\text{SoS}_4(\hat{\Sigma})$, which we view as an algorithm for both detection and estimation problems. Note that the same objective function, with only the three constraints (C1), (C2), (C6) gives the degree-2 relaxation, which is precisely the standard SDP relaxation of Sparse PCA studied in [AW09, BR13b, KNV15]. So clearly $\text{SoS}_4(\hat{\Sigma})$ subsumes the SDP relaxation.

Algorithm 1 $\text{SoS}_4(\hat{\Sigma})$: Degree-4 Sum of Squares Relaxation

Input: $\hat{\Sigma} = \frac{1}{n}XX^T$ where $X = [X^1, \dots, X^n] \in \mathbb{R}^{p \times n}$

Solve the following semidefinite programming and obtain optimal objective value $\text{SoS}_4(\hat{\Sigma})$ and maximizer M^* .

Variables: $M(S)$, for all mutli-sets S of size at most 4.

$$\text{SoS}_4(\hat{\Sigma}) = \max \sum_{i,j} M(x_i x_j) \hat{\Sigma}_{ij} \quad (\text{Obj})$$

$$\text{subject to} \quad \sum_{i \in [p]} M(x_i^2) = k \quad (\text{C1})$$

$$\sum_{i,j \in [p]} |M(x_i x_j)| \leq k^2 \quad (\text{C2})$$

$$M(x_i^3 x_j) = M(x_i x_j), \quad \forall i, j \in [p] \quad (\text{C3})$$

$$\sum_{i \in [p]} M(x_i^2 x_s x_t) = k \cdot M(x_s x_t), \quad \forall s, t \in [p] \quad (\text{C4})$$

$$\sum_{i,j,s,t \in [p]} |M(x_i x_j x_s x_t)| \leq k^4 \quad (\text{C5})$$

$$M \succeq 0 \quad (\text{C6})$$

Output: 1. For detection problem : output H_v if $\text{SoS}_4(\hat{\Sigma}) > (1 + \frac{1}{2}\lambda)k$, H_0 otherwise
 2. For estimation problem: output $M_2^* = (M^*(x_i x_j))_{i,j \in [p]}$

Before stating the lower bounds for both detection and estimation in the next two subsections, we comment on the choices made for the outputs of the algorithm in both, as clearly other choices can be made that would be interesting to investigate. For *detection*, we pick the natural threshold $(1 + \frac{1}{2}\lambda)k$ from the statistically optimal detection algorithm of Section 2.2. Our lower bound of the objective under H_0 is actually a large constant multiple of λk , so we could have taken a higher threshold. To analyze even higher ones would require analyzing the behavior of SoS_4 under the (planted) alternative distribution H_v . For *estimation* we output the maximizer M_2^* of the objective function, and prove that it is not too correlated with the rank-1 matrix vv^T in the planted distribution H_v . This suggest, but does not prove, that the leading eigenvector of M_2^* (which is a natural estimator for v) is not too correlated with v . We finally note that Rigollet's efficient reduction from detection to estimation is not in the SoS framework, and so our detection lower bound does not automatically imply the one for estimation.

3.1 Lower bounds for detection problem

For the detection problem, we prove that $\text{SoS}_4(\hat{\Sigma})$ gives a large objective value on null hypothesis H_0 .

Theorem 3.1. There exists absolute constant C and r such that for $1 \leq \lambda < \min\{k^{1/4}, \sqrt{n}\}$ and any $p \geq C\lambda n$, $k \geq C\lambda^{7/6}\sqrt{n}\log^r p$, the following holds. When the data X is drawn from the null hypothesis H_0 , then with high probability $(1 - p^{-10})$, the objective value of degree-4 sum of squares relaxation $\text{SoS}_4(\hat{\Sigma})$ is at least $10\lambda k$. Consequently, Algorithm 1 can't solve the detection problem.

To parse the theorem and to understand its consequence, consider first the case when λ is a constant (which is also arguably the most interesting regime). Then the theorem says that when we have only $n \ll k^2$ samples, degree-4 SoS relaxation SoS_4 still overfits heavily to the randomness of the data X under the null hypothesis H_0 . Therefore, using $\text{SoS}_4(\hat{\Sigma}) > (1 + \frac{\lambda}{2})k$ (or even $10\lambda k$) as a threshold will fail with high probability to distinguish H_0 and H_v .

We note that for constant λ our result is essentially tight in terms of the dependencies between n, k, p . The condition $p = \tilde{\Omega}(n)$ is necessary since otherwise when $p = o(n)$, even without the sum of squares relaxation, the objective value is controlled by $(1 + o(1))k$ since $\hat{\Sigma}$ has maximum eigenvalue $1 + o(1)$ in this regime. Furthermore, as mentioned in the introduction, $k \geq \tilde{\Omega}(\sqrt{n})$ is also necessary (up to poly-logarithmic factors), since when $n \gg k^2$, a simple diagonal thresholding algorithm works for this simple single-spike model.

When λ is not considered as a constant, the dependence of the lower bound on λ is not optimal, but close. Ideally one could expect that as long as $k \gg \lambda\sqrt{n}$, and $p \geq \lambda n$, the objective value on the null hypothesis is at least $\Omega(\lambda k)$. Tightening the $\lambda^{1/6}$ slack, and possibly extending the range of λ are left to future study.

3.2 Lower bounds for the estimation problem

For estimation problem, we prove that M_2^* output by Algorithm 1 is not too correlated with the desired rank-1 matrix vv^T .

Theorem 3.2. For any constant B there exists absolute constants C and r such that for $\lambda \leq B/2$, $Bn \geq p \geq 2\lambda n$ and $o(p) \geq k \geq C\sqrt{n}\log^r p$, suppose the data X is drawn hypothesis H_v (model (2.1)), then with high probability $(1 - p^{-10})$ over the randomness of the data, Algorithm 1 will output M_2^* such that $\|\frac{1}{k} \cdot M_2^* - vv^T\| \geq 1/5$.

We observe that the result is of the same nature (and arguably near-optimal for estimation problem) as [KNV15] achieve the for degree-2 SoS relaxation. The proof follows simply from combining our detection lower bound Theorem 3.1 and arguments similar to [KNV15]. Finally we address a threshold-like behavior of the estimation error. Note that while our Theorem proves that $n = \tilde{\Omega}(k^2)$ samples is necessary for efficient algorithms to get even constant estimation error, it is known [YZ13, Ma13, WLL14] that slightly more samples, $n = \tilde{O}(k^2)$, can already achieve in polynomial time a much smaller (and optimal) estimation error, namely $O(\sqrt{(k \log p)/n})$.

4 Design of Pseudo-moments

We start with a sketch of our approach to the design of the moments M at a very high level, highlighting aspects of their design which are different than in previous lower bounds. First, there

are some natural choices to make. We define the degree-2 moments \widetilde{M} from the input as the empirical covariance matrix, as was done in the proof of the SDP lower bound. This already gives a large objective value (see Lemma 4.2). We also define taking odd moments (degree 1 and 3) to be 0. The difficult part is designing the degree-4 moments consistently with the constraints and \widetilde{M} . We do this in stages, first approximating the constraints (indeed even \widetilde{M} only approximately satisfies, in a way we will specify in Section 4.1, constraints (C1) and (C2)) and later in Section 4.2 correcting the moments to satisfy the constraints precisely. Moreover, we separately use different 4-moments for different constraints and then combine them, as follows. We define two different degree-4 PSD moments P and Q such that (with high probability) P *almost* satisfies constraints (C3), (C5) and (C6), and *negligible* for constraint (C4) (see Lemma 4.4), whereas Q *almost* satisfies constraints (C5), (C4) and (C6), and *negligible* for (C3) (Lemma 4.5). Therefore taking the sum $P + Q$ will *almost* satisfy all the constraints (Lemma 4.6), which completes the design of the approximate moments. Finally we “locally” adjust $P + Q$ so that the resulting moments M exactly satisfy all the constraints (Theorem 4.7), and remain PSD with high probability.

All moments will be defined from the data matrix X , to which we first apply a simple pre-processing step: we scale all its rows to have square norm n (around which they are concentrated). We abuse notation and call the scaled matrix X as well. Note that when the noise model in the null hypothesis H_0 is Bernoulli, namely the entries of X are chosen as unbiased independent ± 1 variables, the rows are automatically scaled, which motivates our abuse of notation. We suggest that the reader thinks of this distribution, even though the proof works for a much wider class of distributions.

The properties above of our moments will be proved under the assumption that the scaled matrix X satisfies the “pseudo-randomness” condition below. This set-up allows us to encapsulate what we really need the data to satisfy, and thus prove our lower bound not only for Gaussian or Bernoulli noise, but actually for a larger family containing both. Namely, we later prove in Section 7, via a series of concentration inequalities, that when data is drawn from null hypothesis H_0 , its scaling X satisfies the pseudorandomness condition with very high probability under all these noise models. Note that this condition is actually a sequence of statements about deviation from the mean of various polynomials in the data - these will become natural once we define our moments.

Condition 4.1 (Pseudorandomness Condition). Our constructions of the moments will only require

the following pseudorandom conditions about the (scaled) data matrix X ,

$$\|X_i\|^2 = n \quad \forall i \in [p] \quad (\text{P1})$$

$$|\langle X_i, X_j \rangle| \leq \tilde{O}(\sqrt{n}) \quad \forall i \neq j \quad (\text{P2})$$

$$\left| \sum_{\ell \in [p] \setminus i \cup j} \langle X_i, X_\ell \rangle^3 \langle X_j, X_\ell \rangle \right| \leq \tilde{O}(n^{1.5}p), \quad \forall i \neq j \quad (\text{P3})$$

$$\left| \sum_{\ell \in [p]} \langle X_i, X_\ell \rangle \langle X_j, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \right| \leq \tilde{O}(n^2 p^5) \quad \forall \text{ distinct } i, j, s, t \quad (\text{P4})$$

$$\left| \sum_{i \in [p]} \langle X_i, X_s \rangle \langle X_i, X_t \rangle \right| \leq \tilde{O}(n^5 p) \quad \forall \text{ distinct } s, t \quad (\text{P5})$$

$$\sum_{i, \ell \in [p]} \langle X_i, X_\ell \rangle^2 \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \leq \tilde{O}(n^{1.5} p^2), \quad \forall \text{ distinct } s, t \quad (\text{P6})$$

$$\|X X^T\|_F^2 \geq (1 - o(1)) n p^2 \quad (\text{P7})$$

4.1 Approximate Pseudo-moments

In this section, we design a pseudo-moments \widetilde{M} that approximately satisfies all the constraints. Then in the next subsection we will locally adjust it to obtain one that exactly satisfies all of the constraints.

We begin by designing a (partial) degree-2 moments that gives large objective value, which will be later used for the degree-4 moments. The design is essentially the same as [KNV15] though we only work with null hypothesis for now. For the purpose of this section, we suggest the reader to think of X as having uniform $\{\pm 1\}$ entries for simplicity, though as we will see later, we assume that X satisfies certain pseudorandomness condition which holds if X is chosen from a variety of natural stochastic models (with row normalization). We define $\widetilde{M} : \mathbb{R}[x]_2 \rightarrow \mathbb{R}$ as follows:

$$\begin{aligned} \widetilde{M}(x_i x_j) &\triangleq \frac{\gamma k n}{p^2} \hat{\Sigma}_{ij} = \frac{\gamma k}{p^2} \langle X_i, X_j \rangle \quad \forall i, j \in [p] \\ \widetilde{M}(x_i) &\triangleq 0 \quad \forall i \in [p] \\ \widetilde{M}(1) &\triangleq 1 \end{aligned} \quad (4.1)$$

where γ is a constant that to be tuned later according to the signal strength λ . Note that by design $\text{mat}(\widetilde{M})$ is a PSD matrix. We can check straightforwardly that \widetilde{M} satisfies constraint (C2) and gives a large objective value (Obj).

Lemma 4.2. There exists constant C such that for $p \geq \gamma n$ and $k \geq C \gamma \sqrt{n} \log p$, suppose X satisfies Condition 4.1, then \widetilde{M} is a valid degree-2 pseudo-moments and satisfies the sparsity constraint (C2).

$$\sum_{i, j \in [p]} |\widetilde{M}(x_i x_j)| \leq k^2 / 2 \quad (4.2)$$

and has objective value

$$\sum_{i, j \in [p]} \widetilde{M}(x_i x_j) \hat{\Sigma}_{ij} \geq (1 - o(1)) \gamma k$$

Moreover, we also have $\widetilde{M}(x_i^2) = \frac{\gamma kn}{p^2} \leq \frac{k}{p}$, and $\widetilde{M}(x_i x_j) \leq \widetilde{O}(\frac{\gamma k \sqrt{n}}{p^2})$.

Proof. The proof follows simple calculation and concentration inequality. Since $\|X_i\|^2 = n$ for all i and with high probability over the randomness of X , for all $i \neq j$, $|\langle X_i, X_j \rangle| \leq \widetilde{O}(\sqrt{n})$, we obtain that $\widetilde{M}(x_i^2) = \frac{\gamma kn}{p^2} \leq \frac{k}{p}$, and $\widetilde{M}(x_i x_j) \leq \widetilde{O}(\frac{\gamma k \sqrt{n}}{p^2})$. Then to verify equation (4.2), we have

$$\sum_{i,j \in [p]} |\widetilde{M}(x_i x_j)| \leq \sum_i |M(x_i^2)| + \sum_{i \neq j} |M(x_i x_j)| \leq k + \widetilde{O}(\gamma k \sqrt{n}) \leq k^2/2$$

when $k \gg \gamma \sqrt{n}$. Finally, we can verify the objective value is large

$$\sum_{i,j \in [p]} \widetilde{M}(x_i x_j) \hat{\Sigma}_{ij} = \frac{\gamma k}{p^2 n} \sum_{i,j} \langle X_i, X_j \rangle^2 = \frac{\gamma k}{p^2 n} \|X X^T\|_F^2 \geq (1 - o(1)) \gamma k$$

where we use the fact that $\|X X^T\|_F^2 \geq (1 - o(1)) p^2 n$ (see property (P7) in Condition 4.1). \square

Note that \widetilde{M} doesn't satisfies constraint (C1) exactly. However, we could simply fix this by defining $M'(x_i x_j) = \widetilde{M}(x_i x_j)$ for all $i \neq j$ and $M'(x_i^2) = k/p$. However, note that we will use a perturbation of M' in our final design in Section 4.2 so that it is consistent with the degree-4 moments.

Corollary 4.3. There exists absolute constant C such that for $p \geq \gamma n$ and $k \geq C \gamma \sqrt{n} \log p$, there exists a degree-2 pseudo-moments M' that satisfies constraints (C1), (C2) and give objective value at least $(1 - o(1)) \gamma k$.

Now we define a degree-4 pseudo-moment that approximately satisfies all the constraints in $\text{SoS}_4(\widetilde{\Sigma})$ and give a large objective value. We keep the current (approximate) design \widetilde{M} for degree-2 moments, since the degree-2 moments defined in previous section seems to be nearly optimal and enjoys many good properties. Then we define $\widetilde{M}(S) = 0$ for any multi-set S of size 3, because apparently degree-3 moments don't play any role the semidefinite relaxation.

The main difficulty is to define $\widetilde{M}(S)$ for S of size 4. Here we have three constraints (C3), (C4), and (C5), and the PSDness constraint that implicitly compete with each other. We took the following approach. We let \widetilde{M} be a sum of two matrices matrix P and Q . We ensure that P “almost” (as will be specified later) satisfies (C3) and (C5), and is negligible for constraints C4. In turn Q is negligible for constraints (C3) and “almost” satisfies constraint (C4) and (C5). Therefore $P + Q$ will “almost” satisfies constraints (C3)) and (C4)), and satisfy the sparsity constraint (C5). Moreover, P and Q will be PSD by definition. Concretely, we define

$$\widetilde{M}(i, j, s, t) \triangleq P(i, j, s, t) + Q(i, j, s, t)$$

where P and Q are defined as

$$P(i, j, s, t) = \frac{\gamma k}{p^2 n^3} \sum_{\ell \in [p]} \langle X_i, X_\ell \rangle \langle X_j, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle$$

$$Q(i, j, s, t) = \frac{\gamma k^2}{p^3 n} (\langle X_s, X_t \rangle \langle X_i, X_j \rangle + \langle X_i, X_t \rangle \langle X_s, X_j \rangle + \langle X_j, X_t \rangle \langle X_i, X_s \rangle)$$

We note that P and Q are well defined pseudo-moments because they are invariant to the permutation of indices and naturally PSD. To see the PSDness, note that P is a sum of p rank-1 PSD matrices. Moreover, Q is also PSD: the part that corresponds to $\langle X_s, X_t \rangle \langle X_i, X_j \rangle$ is simply a rank-1 PSD matrix; $\langle X_i, X_t \rangle \langle X_s, X_j \rangle$ can be written as $\langle X_t \otimes X_s, X_i \otimes X_j \rangle$ and therefore it also contributes a PSD matrix to Q . Similarly, $\langle X_j, X_t \rangle \langle X_s, X_i \rangle$ can be written as $\langle X_s \otimes X_t, X_i \otimes X_j \rangle$, and it also contributes a PSD matrix.

In the next two lemmas (one for P and one for Q), we formalize the intuition above by showing that, the deviation from P and Q exactly satisfying the constraints is captured by error matrices $\mathcal{E}, \mathcal{F}, \mathcal{G}$. We bound the magnitude of these error matrices and establish the PSDness of some of them so that later we can fix them for the exact satisfaction of the constraints.

Lemma 4.4. There exists some absolute constant C and r such that for $1 \leq \gamma \leq \min\{k^{1/4}, \sqrt{n}\}$, $1 \leq \gamma \leq n$, $p = 1.1\gamma n$, and $k \geq C \cdot \gamma^{7/6} \sqrt{n} \log^r p$, suppose X satisfies pseudorandomness condition 4.1, then P almost satisfies constraint (C3) and (C5), naturally satisfies PSD constraint (C6), and is negligible for constraint (C4) in the sense that

$$P(x_i^3 x_j) = \widetilde{M}(x_i x_j) + \mathcal{F}_{ij}, \quad \forall i, j \in [p] \quad (4.3)$$

$$\sum_i P(x_i^2 x_s x_t) = \mathcal{E}_{st} \quad \forall s, t \in [p] \quad (4.4)$$

$$\sum_{i,j,s,t} |P(x_i x_j x_s x_t)| \leq k^4/3 \quad (4.5)$$

where \mathcal{F} and \mathcal{E} are $p \times p$ matrices that satisfy

1. $0 \leq \mathcal{F}_{ii} \leq \widetilde{O}\left(\frac{\gamma k}{pn}\right)$, $|\mathcal{F}_{ij}| \leq \widetilde{O}\left(\frac{\gamma k}{pn^{1.5}}\right)$ for any i and $j \neq i$.
2. \mathcal{E} is PSD with $|\mathcal{E}_{ss}| \leq \widetilde{O}\left(\frac{\gamma k}{n}\right)$, and $|\mathcal{E}_{st}| \leq \widetilde{O}\left(\frac{\gamma k}{n\sqrt{n}}\right)$ for any $s \neq t$.

Lemma 4.5. There exists some absolute constant C and r such that for $1 \leq \gamma \leq n$, $p = 1.1\gamma n$ and $k \geq C \cdot \gamma \sqrt{n} \log^r p$, suppose X satisfies pseudorandomness condition 4.1, then Q is negligible for constraint (C3) and almost satisfies constraint (C4) and (C5) in the sense that,

$$Q(x_i^3 x_j) = \frac{3k}{p} \widetilde{M}(x_i x_j) \quad \forall i, j \quad (4.6)$$

$$\sum_i Q(x_i^2 x_s x_t) = k \widetilde{M}(x_s x_t) + \mathcal{G}_{st}, \quad \forall s, t \quad (4.7)$$

$$\sum_{i,j,s,t} |Q(x_i x_j x_s x_t)| \leq k^4/3 \quad (4.8)$$

where \mathcal{G} is a $p \times p$ PSD matrix $|\mathcal{G}_{ss}| \leq \widetilde{O}\left(\frac{\gamma k^2}{p^2}\right)$ and $|\mathcal{G}_{st}| \leq \widetilde{O}\left(\frac{\gamma k^2}{p^2 \sqrt{n}}\right)$ for any $s \neq t$.

Now we are ready to prove that $\widetilde{M} = P + Q$ almost satisfies all other constraints (C1)- (C6) approximately.

Lemma 4.6. Define $\widetilde{M}(x_i x_j x_s x_t) = P(x_i x_j x_s x_t) + Q(x_i x_j x_s x_t)$ for all $i, j, s, t \in [p]$, then we have under the condition of Lemma 4.4,

$$\widetilde{M}(x_i^3 x_j) = \widetilde{M}(x_i x_j) + \mathcal{F}'_{ij} \quad (4.9)$$

$$\sum_{i \in [p]} \widetilde{M}(x_i^2 x_s x_t) = k \widetilde{M}(x_s x_t) + \mathcal{E}'_{st} \quad \forall s, t \quad (4.10)$$

$$\sum_{i, j, s, t} |\widetilde{M}(x_i x_j x_s x_t)| \leq 2k^4/3 \quad (4.11)$$

where \mathcal{F}' and \mathcal{E}' are $p \times p$ matrices that satisfy

1. $|\mathcal{F}'_{ii}| \leq \widetilde{O}\left(\frac{\gamma k^2 n}{p^3}\right)$ and $|\mathcal{F}'_{ij}| \leq \widetilde{O}\left(\frac{\gamma k^2 \sqrt{n}}{p^3}\right)$ for all $i \neq j$.
2. \mathcal{E}' is a PSD matrix with $\mathcal{E}'_{ss} \leq \widetilde{O}\left(\frac{\gamma k}{n}\right)$ and $|\mathcal{E}'_{st}| \leq \widetilde{O}\left(\frac{\gamma k}{n\sqrt{n}}\right)$ for $s \neq t$.

Proof of Lemma 4.6 using Lemma 4.4 and Lemma 4.5. Note that by definition of \widetilde{M} and Lemma 4.4 and Lemma 4.5, we have $\mathcal{F}'_{ij} = \mathcal{F}_{ij} + \frac{3k}{p} \widetilde{M}(x_i x_j)$ and $\mathcal{E}' = \mathcal{E} + \mathcal{G}$. The bound for \mathcal{F}' follows the bound for \mathcal{F} and the facts that $\widetilde{M}(x_i^2) = \frac{\gamma k n}{p^2}$ and $|\widetilde{M}(x_i x_j)| \leq \widetilde{O}\left(\frac{\gamma k \sqrt{n}}{p^2}\right)$. The PSDness of \mathcal{E}' and the bounds for it follows straightforwardly from those of \mathcal{E} and \mathcal{G} . Equation (4.11) follows equation (4.4) and equation (4.8). \square

4.2 Exact Pseudo-moments

Note that \widetilde{M} only satisfies the constraints approximately up to some additive errors (which are carefully bounded for the purpose of the next theorem). We fix this issue by defining the actual pseudo-moments M based (on a carefully chosen) local adjustment of \widetilde{M} . Concretely, we define $M(1) = 1$ and for all add degree monomial x^α , $M(x^\alpha) = 0$. For distinct i, j, s, t , we define $M(x_i x_j x_s x_t) \triangleq \widetilde{M}(x_i x_j x_s x_t)$ and $M(x_i^2 x_s x_t) \triangleq \widetilde{M}(x_i^2 x_s x_t)$. For distinct s, t , we define

$$M(x_s^3 x_t) = M(x_s x_t) \triangleq \widetilde{M}(x_s x_t) + \frac{1}{k-2} (\mathcal{E}'_{st} - 2\mathcal{F}'_{st}) \quad (4.12)$$

and $M(x_s^2 x_t^2) \triangleq \widetilde{M}(x_s^2 x_t^2) + \delta$ where δ a constant (will be proved to be nonnegative) such that

$$\sum_{i \neq j} M(x_s^2 x_t^2) = \sum_{s \neq t} (\widetilde{M}(x_s^2 x_t^2) + \delta) = k^2 - k \quad (4.13)$$

Then we define

$$M(x_i^4) = M(x_i^2) \triangleq \frac{1}{k-1} \sum_{j: j \neq i} M(x_i^2 x_j^2) \quad (4.14)$$

Therefore we can see by construction, it is almost obvious that M satisfies all the linear constraints (C1), (C3), (C4) exactly. Moreover, since \mathcal{E}' and \mathcal{F}' are small error matrices, most of the entries $M(x_i x_j x_s x_t)$ are equal or close to $\widetilde{M}(x_i x_j x_s x_t)$. Note that \widetilde{M} satisfies the rest of constraints (C2), (C5) and (C6) (even with some slackness). We will prove that the difference between M and \widetilde{M} is small enough so that these constraints are still satisfied by M .

Theorem 4.7. Under the condition of Lemma 4.4, suppose X satisfies pseudorandomness condition 4.1, then the pseudo-moments M defined above satisfies all the constraint (C1)-(C6) of the semidefinite programming and has objective value larger than $(1 - o(1))\gamma k$.

Proof. We prove that M satisfies all the constraints in an order that is most convenient for the proof, and check the objective value at the end.

- Constraint (C3): This is satisfied by the definition of M .
- Constraint (C4): By the definition, we can see that $M(x_s^3 x_t)$ is also a perturbation of $\widetilde{M}(x_s^3 x_t)$:

$$M(x_s^3 x_t) = \widetilde{M}(x_s x_t) + \frac{1}{k-2} (\mathcal{E}'_{st} - 2\mathcal{F}'_{st}) = \widetilde{M}(x_s^3 x_t) + \frac{\mathcal{E}'_{st}}{k-2} - \frac{k}{k-2} \mathcal{F}'_{st} \quad (4.15)$$

It follows that for $s \neq t$,

$$\begin{aligned} \sum_{i \in [p]} M(x_i^2 x_s x_t) &= 2M(x_s x_t) + \sum_{i \in [p] \setminus \{s, t\}} \widetilde{M}(x_i^2 x_s x_t) \\ &= 2\widetilde{M}(x_s x_t) + \frac{2}{k-2} (\mathcal{E}'_{st} - 2\mathcal{F}'_{st}) + \sum_{i \in [p] \setminus \{s, t\}} \widetilde{M}(x_i^2 x_s x_t) \\ &= \widetilde{M}(x_s^3 x_t) + \widetilde{M}(x_s x_t^3) - 2\mathcal{F}'_{st} + \frac{2}{k-2} (\mathcal{E}'_{st} - 2\mathcal{F}'_{st}) + \sum_{i \in [p] \setminus \{s, t\}} \widetilde{M}(x_i^2 x_s x_t) \\ &= k\widetilde{M}(x_s x_t) + \mathcal{E}'_{st} - 2\mathcal{F}'_{st} + \frac{2}{k-2} (\mathcal{E}'_{st} - 2\mathcal{F}'_{st}) \\ &= kM(x_s x_t) \end{aligned}$$

where the second equality uses definition (4.12) and the third uses equation (4.9), and the fourth uses (4.10) and the last equality uses the definition (4.12) again.

Moreover, for the case when $s = t$, we have that

$$\begin{aligned} \sum_{i \in [p]} M(x_i^2 x_s^2) &= M(x_s^4) + \sum_{i \in [p] \setminus \{s\}} M(x_i^2 x_s^2) \\ &= M(x_s^2) + (k-1)M(x_s^2) = kM(x_s^2) \end{aligned}$$

where we used the definition (4.14) of $M(x_s^4)$ and $M(x_s^2)$. Therefore we verified that M satisfies constraint (C4).

- Constraint (C1): Using equation (4.13) and (4.14), we have

$$\sum_{i \in [p]} M(x_i^2) = \frac{1}{k-1} \sum_{i \neq j} M(x_i^2 x_j^2) = k$$

- Constraint (C6):

Next we check the PSDness of matrix $\text{mat}(M)$. Note that $\text{mat}(M)$ is indexed by all the mutli subset of $[p]$ of size at most 2, and it consists of 3 blocks $\text{mat}(M) = \text{blkdiag}(M_4, M_2, M_0)$, where

$$M_4 = (\text{mat}(M)_{S,T})_{|S|=2, |T|=2}$$

$$M_2 = (\text{mat}(M)_{S,T})_{|S|=1, |T|=1}$$

$$M_0 = 1$$

Therefore it suffices to check that M_0 , M_2 and M_4 are all PSD. M_0 is trivially PSD. We can write M_2 in the following form

$$M_2 = (M(x_s x_t))_{s,t \in [p]} = \left(\widetilde{M}(x_s x_t) \right)_{s,t \in [p]} + \Delta$$

where $\Delta = M_2 - \left(\widetilde{M}(x_s x_t) \right)_{s,t \in [p]}$. By equation (4.12), we have that for $s \neq t$, $\Delta_{st} = \frac{1}{k-2} (\mathcal{E}'_{st} - 2\mathcal{F}'_{st})$ for all $s \neq t$. Moreover, by definition of $M(x_s^2)$ and $M(x_s^2 x_t^2)$, we have that

$$\begin{aligned} M(x_s^2) &= \frac{1}{k-1} \sum_{s:t \neq s} M(x_s^2 x_t^2) = \frac{1}{k-1} \sum_{s:t \neq s} \left(\widetilde{M}(x_s^2 x_t^2) + \delta \right) \\ &= \frac{1}{k-1} \left(k \widetilde{M}(x_s^2) + \mathcal{E}'_{ss} - \widetilde{M}(x_s^4) \right) + \frac{p-1}{k-1} \cdot \delta \\ &= \frac{1}{k-1} \left(k \widetilde{M}(x_s^2) + \mathcal{E}'_{ss} - \widetilde{M}(x_s^2) - \mathcal{F}'_{ss} \right) + \frac{p-1}{k-1} \cdot \delta \\ &= \widetilde{M}(x_s^2) + \frac{1}{k-1} (\mathcal{E}'_{ss} - \mathcal{F}'_{ss}) + \frac{p-1}{k-1} \cdot \delta \end{aligned} \quad (4.16)$$

where second line uses equation (4.10) and the third line uses (4.9), and therefore $\Delta_{ss} = \frac{p-1}{k-1} \cdot \delta + \frac{1}{k-1} (\mathcal{E}'_{ss} - \mathcal{F}'_{ss})$. We extract the PSD matrix $\frac{1}{k-2} \cdot \mathcal{E}'$ form Δ and obtain $\Delta' = \Delta - \frac{1}{k-2} \cdot \mathcal{E}'$. Then by this definition, $\Delta'_{ss} = \frac{p-1}{k-1} \cdot \delta + \frac{1}{k-1} (\mathcal{E}'_{ss} - \mathcal{F}'_{ss}) - \frac{1}{k-2} \mathcal{E}'_{ss}$, and $\Delta'_{st} = -\frac{2}{k-2} \mathcal{F}'_{st}$. We use Gershgorin Circle Theorem to establish the PSDness of Δ' . By Lemma 4.6, we have $|\mathcal{F}_{ij}| \leq \widetilde{O}\left(\frac{\gamma k^2 \sqrt{n}}{p^3}\right)$ Therefore

$$\sum_{j:j \neq i} |\Delta'_{ij}| \leq p \cdot \frac{4}{k-2} \widetilde{O}\left(\frac{\gamma k^2 \sqrt{n}}{p^3}\right) \leq o\left(\frac{k}{p}\right)$$

where we used the fact that $\sqrt{n}/p = o(1)$ which follows from $p = 1.1\gamma n$ and $\gamma \leq \sqrt{n}$.

Using equation (4.16) and constrain (C1) we have that

$$k = \sum_s M(x_s^2) = \sum_s \widetilde{M}(x_s^2) + \sum_s \frac{1}{k-1} (\mathcal{E}'_{ss} - \mathcal{F}'_{ss}) + \frac{p(p-1)}{k-1} \cdot \delta \quad (4.17)$$

$$\leq \frac{\gamma k n}{p} + \widetilde{O}(1) + \frac{p(p-1)}{k-1} \cdot \delta \quad (4.18)$$

It follows that $\delta \geq (1-o(1)) \cdot \frac{k(k-1)}{12p(p-1)}$. Therefore we obtain that $\Delta'_{ii} = \frac{p-1}{k-1} \cdot \delta + \frac{1}{k-1} (\mathcal{E}'_{ss} - \mathcal{F}'_{ss}) - \frac{1}{k-2} \mathcal{E}'_{ss} \geq \frac{1}{12}(1-o(1)) \frac{k}{p} - \tilde{O}(\frac{\gamma}{n}) - \tilde{O}\left(\frac{\gamma kn}{p^3}\right) = \frac{1}{12}(1-o(1)) \frac{k}{p}$. Therefore we obtain $\Delta'_{ii} \geq \sum_{j:j \neq i} |\Delta'_{ij}|$ and by Gershgorin Circle Theorem Δ' is PSD.

Now we examine M_4 . we write M_4 as

$$M_4 = \text{mat}(P) + \text{mat}(Q) + \Gamma$$

where $\Gamma = M_4 - (\text{mat}(P) + \text{mat}(Q))$. One can observe that Γ has only non-zero entries of the form

$$\begin{aligned} \Gamma_{ii,ii} &= M(x_i^4) - P(x_i^4) - Q(x_i^4) = M(x_i^4) - \widetilde{M}(x_i^4) = M(x_i^2) - \widetilde{M}(x_i^2) - \mathcal{F}'_{ii} \\ &= \frac{(p-1)}{k-1} \cdot \delta + \frac{1}{k-1} \mathcal{E}'_{ii} - \frac{k}{k-1} \mathcal{F}_{ii} \end{aligned} \quad (4.19)$$

and

$$\begin{aligned} \forall i \neq j, \Gamma_{ii,jj} = \Gamma_{ij,ij} = \Gamma_{ij,ji} &= M(x_i^2 x_j^2) - P(x_i^2 x_j^2) - Q(x_i^2 x_j^2) \\ &= M(x_i^2 x_j^2) - \widetilde{M}(x_i^2 x_j^2) = \delta \end{aligned} \quad (4.20)$$

and

$$\begin{aligned} \forall i \neq j, \Gamma_{ii,ij} = \Gamma_{ii,ji} &= M(x_i^3 x_j) - P(x_i^3 x_j) - Q(x_i^3 x_j) \\ &= M(x_i^3 x_j) - \widetilde{M}(x_i^3 x_j) = \frac{\mathcal{E}'_{st}}{k-2} - \frac{k}{k-2} \mathcal{F}_{st} \end{aligned} \quad (4.21)$$

where the last equality uses equation (4.15).

Now we are ready to prove PSDness of Γ . We further decompose Γ as $\Gamma = \Gamma' + \text{blkdiag}(\Lambda', 0)$ where Λ' is the $p \times p$ matrix with $\Lambda' = \delta \mathbf{1}\mathbf{1}^T$ ⁵. Note that Λ' is a PSD matrix and therefore it suffices to prove that $\Gamma' = \Gamma - \text{blkdiag}(\Lambda', 0)$ is a PSD matrix.

Note that Γ' has ij -th column the same as ji -th column, and therefore it's only of rank at most $p + p(p-1)/2$. We define Γ'' be the $p + p(p-1)/2$ by $p + p(p-1)/2$ submatrix of Γ' , that is indexed by subsets (i, i) for $i \in [p]$ and (i, j) for $i < j$. Therefore it suffices to prove that Γ'' is PSD. We prove it using Gershgorin Circle Theorem.

Note that by equation (4.19), we have that $\Gamma''_{ii,ii} = \Gamma'_{ii,ii} - \Lambda'_{ii,ii} = \frac{(p-k)}{k-1} \cdot \delta + \frac{1}{k-1} \mathcal{E}'_{ii} - \frac{k}{k-1} \mathcal{F}_{ii}$. Therefore by the lower bound for δ and Lemma 4.6, we obtain, $\Gamma''_{ii,ii} \geq (1-o(1)) \frac{\epsilon k}{p}$. Moreover, $\Gamma''_{ii,ij} = \Gamma_{ii,ij} = \frac{\mathcal{E}'_{st}}{k-2} - \frac{k}{k-2} \mathcal{F}_{st}$ and therefore $|\Gamma''_{ii,ij}| \leq |\frac{\mathcal{E}'_{st}}{k-2}| + |\frac{k}{k-2} \mathcal{F}_{st}| \leq \tilde{O}(\frac{\gamma}{n^{1.5}}) + \tilde{O}\left(\frac{\gamma k^2 \sqrt{n}}{p^3}\right) \leq \tilde{O}(\frac{\gamma}{n^{1.5}})$. Furthermore, for $i < j$, $\Gamma''_{ij,ij} = \Gamma_{ij,ij} = \delta \geq (1-o(1)) \frac{\epsilon k^2}{p^2}$. Finally observe that $\Gamma''_{ii,jj} = 0$ by definition and all other entries of Γ'' are trivially 0 because the corresponding entries of Γ and Λ' vanish. Therefore we are ready to use a variant of Gershgorin Circle Theorem (Lemma 8.8) to prove the PSDness of Γ' . Taking $\alpha = 1/\gamma^2$, we have for any i ,

⁵ Λ' is index by ii , $i = 1, \dots, p$

$$\begin{aligned}
\alpha \sum_{s,t:(s,t) \neq (i,i), s < t} |\Gamma''_{ii,st}| &= \sum_{j \in [p]} |\Gamma''_{ii,jj}| + \sum_{j:j > i} |\Gamma''_{ii,ij}| + \sum_{j:j < i} |\Gamma''_{ii,ji}| \\
&\leq \alpha p \cdot \tilde{O}\left(\frac{\gamma}{n^{1.5}}\right) = o\left(\frac{k}{p}\right) \leq \Gamma''_{ii,ii}
\end{aligned}$$

where we used the fact that $k \gg \gamma\sqrt{n}$ and ϵ is a constant.

Moreover, for any $i < j$, we have that

$$\begin{aligned}
\frac{1}{\alpha} \sum_{(s,t):(s,t) \neq (i,j), s < t} |\Gamma''_{ij,st}| &\leq |\Gamma''_{ij,ii}| + |\Gamma''_{ij,jj}| \\
&\leq \tilde{O}\left(\frac{\gamma^3}{n^{1.5}}\right) = o\left(\frac{k^2}{p^2}\right) \leq \Gamma''_{ij,ij}
\end{aligned}$$

where we used $k \geq \gamma^4$ and $k \gg \gamma\sqrt{n}$. Therefore by Lemma 8.8, we obtain that Γ'' is PSD.

- Constraint (C2): Using Lemma 4.2 and equation (4.12), we have that

$$\begin{aligned}
\sum_{s,t} |M(x_s x_t)| &\leq \sum_{ss} M(x_s^2) + \sum_{s \neq t} |M(x_s x_t) - \widetilde{M}(x_s x_t)| + \sum_{s \neq t} |\widetilde{M}(x_s x_t)| \\
&\leq k + p^2 \tilde{O}\left(\frac{\gamma}{n^{1.5}}\right) + k^2/2 \leq k^2
\end{aligned}$$

- Constraint (C5): Finally, we check that M satisfies the sparsity constraint (C5).

$$\begin{aligned}
\sum_{i,j,s,t} |M(x_i x_j x_s x_t)| &\leq \sum_{i,j,s,t} |\Gamma_{ij,st}| + \sum_{i,j,s,t} |\widetilde{M}(x_i x_j x_s x_t)| \\
&\leq k^4
\end{aligned}$$

where we used (4.11) and the (trivial) facts that $\Gamma_{ij,st} \leq O(k/p)$ for any i, j, s, t and there are only at most $O(p^2)$ nonzero entries in Γ .

- Objective value (Obj): Note that by constraint (C1) and Lemma 4.2, we have that $\sum_i M(x_i^2) \hat{\Sigma}_{ii} = k \geq \sum_i \widetilde{M}(x_i^2) \hat{\Sigma}_{ii}$, then

$$\begin{aligned}
\sum_{i,j} M(x_i x_j) \hat{\Sigma}_{i,j} &\geq \sum_{i,j} \widetilde{M}(x_i x_j) \hat{\Sigma}_{i,j} - \sum_{i \neq j} |M(x_i x_j) - \widetilde{M}(x_i x_j)| |\hat{\Sigma}_{ij}| \\
&\geq (1 - o(1)) \gamma k - p^2 \cdot \tilde{O}\left(\frac{\gamma}{n\sqrt{n}}\right) \cdot \tilde{O}\left(\frac{1}{\sqrt{n}}\right) \\
&\geq (1 - o(1)) \gamma k
\end{aligned}$$

where in the second inequality we used Lemma 4.2 and the facts that $\hat{\Sigma}_{ij} = \frac{1}{n} \langle X_i, X_j \rangle \leq \tilde{O}(1/\sqrt{n})$ and $|\mathcal{E}'_{ij}| + |\mathcal{F}'_{ij}| \leq \tilde{O}\left(\frac{\gamma}{n^{1.5}}\right)$, and the last line uses the fact that $\gamma^4 \leq k$.

□

5 Proof of Theorem 3.1 and Theorem 3.2

In this section, we prove our main Theorems using the technical results of the previous sections. Before getting in the proof, we start with the observation that in order to get a lower bound of objective value $10\lambda k$, it suffices to consider the special case when $p = 10\lambda n$. The reason is that the objective value of SoS_4 is increasing in p while fixing all other parameters: Suppose $p' \leq p$ and $\hat{\Sigma}' \in \mathbb{R}^{p' \times p'}$ is a submatrix of $\hat{\Sigma}$, and $M^{*'} : \mathbb{R}_{p'}[x]_4$ is the maximizer of $\text{SoS}_4(\hat{\Sigma}')$. Then we can extend M' to $M : \mathbb{R}_p[x]_4 \rightarrow \mathbb{R}$ by simply defining that $M(x^S) = M^{*'}(x^S)$ if $S \subset [p']$ and 0 otherwise. This preserves all the constraint and objective value. Thus we proved that the objective value for Σ is at least the one for Σ' . Formally, we have

Proposition 5.1. Fixing λ, k, n , given a data matrix $X \in \mathbb{R}^{p \times n}$, and any submatrix matrix $Y \in \mathbb{R}^{p' \times n}$ of X with $p' \leq p$, let $\hat{\Sigma}_X$ and $\hat{\Sigma}_Y$ be the covariance matrices of X and Y , then we have that $\text{SoS}_4(\hat{\Sigma}_X) \geq \text{SoS}_4(\hat{\Sigma}_Y)$.

Now we are ready to prove our main Theorem 3.1. The idea is very simple: we normalize the data matrix X so that the resulting matrix \bar{X} satisfies the pseudorandomness condition 4.1. Then we apply Theorem 4.7 and obtain a moment matrix which give large objective value with respect to \bar{X} . Then we argue that the difference between \bar{X} from X is negligible so that the same moment matrix has also large objective value with respect to X .

Proof of Theorem 3.1. Using the observation above, we take $p = 1.1\gamma n$ with $\gamma = 11\lambda$, and we define \bar{X} to matrix obtained by normalizing rows of X to euclidean norm \sqrt{n} . Then by Theorem 7.1 it satisfies the pseudorandomness condition 4.1. Let $\hat{\Sigma}' = \frac{1}{n} \bar{X} \bar{X}^T$ be the covariance matrix defined by \bar{X} . By Theorem 4.7 we have that $\text{SoS}_4(\hat{\Sigma}') \geq (1 - o(1))\gamma k \geq 11\lambda k$. Moreover, let M be the moment defined in Theorem 4.7, and M_2 its restriction to degree-2 moments, that is, $M_2 = (\text{mat}(M)_{S,T})_{|S|=|T|=1}$. We are going to show that the entry-wise difference between $\hat{\Sigma}$ and $\hat{\Sigma}'$ are small enough so that $\langle M_2, \hat{\Sigma} \rangle$ is close to $\langle M_2, \hat{\Sigma}' \rangle$.

Note that since $\|X_i\|^2 = n \pm \tilde{O}(\sqrt{n})$, therefore for any $i \neq j$, $\hat{\Sigma}'_{ij} = \frac{\hat{\Sigma}_{ij}}{\|X_i\| \|X_j\|} = \hat{\Sigma}_{ij} \pm \tilde{O}(\frac{1}{\sqrt{n}}) |\hat{\Sigma}|_{ij} = \hat{\Sigma}_{ij} \pm \tilde{O}(\frac{1}{n})$. For $i = j$, we have similarly that $\hat{\Sigma}'_{ii} = \hat{\Sigma}_{ii} \pm \tilde{O}(\frac{1}{\sqrt{n}})$. We bound the difference between $\langle M_2, \hat{\Sigma}' \rangle$ and $\langle M_2, \hat{\Sigma} \rangle$ by the sum of the entry-wise differences:

$$\begin{aligned} |\langle M_2, \hat{\Sigma}' - \hat{\Sigma} \rangle| &\leq \sum_i M(x_i^2) |\hat{\Sigma}_{ii} - \hat{\Sigma}'_{ii}| + \sum_{i \neq j} M(x_i x_j) |\hat{\Sigma}_{ij} - \hat{\Sigma}'_{ij}| \\ &\leq p \cdot O(k/p) \cdot \tilde{O}(\frac{1}{\sqrt{n}}) + p^2 \cdot \tilde{O}(\frac{\gamma k \sqrt{n}}{p^2}) \cdot \tilde{O}(\frac{1}{n}) = o(k) \end{aligned}$$

Therefore $\langle M_2, \hat{\Sigma} \rangle \geq (1 - o(1))\gamma k - o(k) = (1 - o(1))\gamma k$. Therefore the moment M gives objective value $(1 - o(1))\gamma k$ for data $\hat{\Sigma}$, and therefore $\text{SoS}_4(\hat{\Sigma}) \geq (1 - o(1))\gamma k \geq 10\lambda k$. \square

Then we prove that Theorem 3.1 together with the arguments in [KNV15] implies Theorem 3.2. The intuition behind is the following: Suppose M_2^* is very close to vv^T , then it is close to rank-1 and its leading eigenvector is close to \hat{v} . However, since we prove that the objective value is large (which is true also in the planted vector case), M_2^* needs to be highly correlated with $\hat{\Sigma}$, which implies its leading eigenvector \hat{v} needs to be correlated with $\hat{\Sigma}$, which in turns implies that v is correlated with $\hat{\Sigma}$. However, it turns out that v is not correlated enough with $\hat{\Sigma}$, which leads to a contradiction.

Proof of Theorem 3.2. We first prove that the optimal value of $\text{SoS}_4(\hat{\Sigma})$ for hypothesis H_v is also at least $0.99kp/n$. Suppose v has support S of size k . We consider the restriction of linear operator M to the subset $T = [p] \setminus S$, denoted by M_T . That is, we have that $M_T(x^\alpha) = 0$ for any monomial x^α that contains a factor x_i with $i \in S$, and otherwise $M_T(x^\alpha) = M(x^\alpha)$. We also consider the data matrix X_T obtained by restricting to the rows indexed by T . Note that since X_T doesn't contain the signal, and $k \gg \sqrt{n}$, using Theorem 4.7 with $\gamma = (p - k)/(1.01n)$, we have that there exists pseudo-moment M_T^* which gives objective value $\geq (1 - o(1))\gamma k \geq 0.99pk/n$ with respect to covariance matrix $\hat{\Sigma}_T = \frac{1}{n}X_TX_T^T$. Note that by Proposition 5.1, $\text{SoS}_4(\hat{\Sigma}) \geq \text{SoS}_4(\hat{\Sigma}_T)$ and therefore we obtain that under hypothesis H_v , with high probability, $\text{SoS}_4(\hat{\Sigma}) \geq 0.99kp/n$.

Now suppose M^* is the maximizer of $\text{SoS}_4(\hat{\Sigma})$, and $M_2^* = (M^*(x_i x_j))_{i,j \in [p]}$. For the sake of contradiction, we assume that $\|\frac{1}{k}M_2^* - vv^T\| \leq 1/5$. We first show that this implies that M_2^* has an eigenvector \hat{v} that is close to v and its eigenvalue is large. Indeed we have $\|\frac{1}{k}M_2^*\| \geq \|vv^T\| - 1/5 = 4/5$. Therefore the top eigenvector of $\frac{1}{k}M_2^*$ has eigenvalue larger than $4/5$. Then we can decompose the difference between $\frac{1}{k} \cdot M_2^*$ and vv^T into $\frac{1}{k} \cdot M_2^* - vv^T = \frac{4}{5} \cdot (\hat{v}\hat{v}^T - vv^T) + ((\frac{1}{k}M_2^* - \frac{4}{5}\hat{v}\hat{v}^T) - \frac{1}{5}vv^T)$. Note that since $(\frac{1}{k}M_2^* - \frac{4}{5}\hat{v}\hat{v}^T)$ is a PSD matrix with eigenvalue at most $1/5$, we have $\|((\frac{1}{k}M_2^* - \frac{4}{5}\hat{v}\hat{v}^T) - \frac{1}{5}vv^T)\| \leq 2/5$ by triangle inequality. Then by triangle inequality and our assumption again we obtain that

$$\frac{1}{5} \geq \|\frac{1}{k} \cdot M_2^* - vv^T\| \geq \|\frac{4}{5}(\hat{v}\hat{v}^T - vv^T)\| - \left\| \left(\left(\frac{1}{k}M_2^* - \frac{4}{5}\hat{v}\hat{v}^T \right) - \frac{1}{5}vv^T \right) \right\| \geq \|\frac{4}{5}(\hat{v}\hat{v}^T - vv^T)\| - \frac{2}{5}$$

Therefore we obtain that $\|vv^T - \hat{v}\hat{v}^T\| \leq 3/5$ and therefore $\|vv^T - \hat{v}\hat{v}^T\|_F^2 \leq 2\|vv^T - \hat{v}\hat{v}^T\|^2 \leq 1$. It follows that $|\langle v, \hat{v} \rangle|^2 = 1 - \frac{1}{2}\|vv^T - \hat{v}\hat{v}^T\|_F^2 \geq 1/2$.

Next we are going to show that it is impossible for M_2^* to have an eigenvector that is close to v with a large eigenvalue and therefore we will get a contradiction. Let $\hat{v} = \alpha v + \beta s$ where s is orthogonal to v and $\alpha^2 + \beta^2 = 1$ and $\alpha \geq \sqrt{1/2}$, and $\beta \leq \sqrt{1/2}$. Then using triangle inequality we have that $\|\hat{v}\|_{\hat{\Sigma}} \leq \|\alpha v\|_{\hat{\Sigma}} + \|\beta s\|_{\hat{\Sigma}} \leq \sqrt{O(\lambda)\alpha} + \sqrt{\|\hat{\Sigma}\|\beta}$. Proposition 5.3 of [KNV15] implies that for sufficiently large C and $\lambda \geq 1$, when $p/n \geq C\lambda$, $\|\hat{\Sigma}\| \leq 1.01p/n$. Therefore under our assumption we have that $\|\hat{\Sigma}\| \leq 1.01p/n$. It follows $\beta \leq \sqrt{1/2}$ that $\|\hat{v}\|_{\hat{\Sigma}} \leq \sqrt{O(\lambda)\alpha} + \sqrt{\beta} \cdot \sqrt{p/n} \leq \sqrt{O(\lambda)} + \sqrt{p/2n}$. Therefore, we have that

$$\begin{aligned} 0.99p/n &\leq \frac{1}{k} \cdot \text{SoS}_4(\hat{\Sigma}) = \frac{1}{k} \cdot \langle M_2, \hat{\Sigma} \rangle = \frac{1}{k} \cdot \langle M_2^* - \frac{4k}{5} \cdot \hat{v}\hat{v}^T, \hat{\Sigma} \rangle + \frac{4}{5} \langle \hat{v}\hat{v}^T, \hat{\Sigma} \rangle \\ &\leq \frac{1}{k} \text{tr}(M_2^* - \frac{4}{5} \cdot \hat{v}\hat{v}^T) \|\hat{\Sigma}\|_2 + \frac{4}{5} \|\hat{v}\|_{\hat{\Sigma}}^2 \\ &\leq \frac{1}{5} \|\hat{\Sigma}\| + O(\alpha^2 \lambda) + O(\sqrt{\lambda p/n}) + \frac{2}{5} \cdot p/n \\ &\leq \frac{4}{5} \cdot \frac{p}{n} + O(\sqrt{\lambda p/n}) \end{aligned}$$

where in the third line we used the fact that $\|\hat{v}\|_{\hat{\Sigma}} \leq \sqrt{O(\lambda)} + \sqrt{p/2n}$, and the last line we used $\|\hat{\Sigma}\| \leq 1.01p/n$. Note that this is a contradiction since we assumed that $p/n \geq C\lambda$ for sufficiently large C . \square

6 Analysis of matrices P and Q

In this section we prove Lemma 4.4 and 4.5. They basically follow direct calculation and the pseudorandomness properties of data matrix X listed in Condition 4.1.

Proof of Lemma 4.4. Note that since $p = 1.1\gamma n$ and $1 \leq \gamma \leq n$, we have that $O(n^2) \geq p \geq n$. We verify equations (4.3), (4.4) and (4.5) and the bounds for \mathcal{F} and \mathcal{E} one by one.

- Equation (4.3):

For the case when $i = j$, we verify $P(x_i^4)$ using property (P1) and (P2),

$$\begin{aligned} P(x_i^4) &= \frac{\gamma k}{p^2 n^3} \left(\langle X_i, X_i \rangle^4 + \sum_{\ell \in [p] \setminus i} \langle X_i, X_\ell \rangle^4 \right) \\ &\leq \frac{\gamma k}{p^2 n^3} \left(n^3 \langle X_i, X_i \rangle + \tilde{O}(pn^2) \right) \\ &= \widetilde{M}(x_i^2) + \tilde{O}\left(\frac{\gamma k}{pn}\right) \end{aligned}$$

For distinct i, j , we have that

$$\begin{aligned} P(x_i^3 x_j) &= \frac{\gamma k}{p^2 n^3} \left(\langle X_i, X_i \rangle^3 \langle X_i, X_j \rangle + \langle X_i, X_j \rangle^3 \langle X_i, X_i \rangle + \sum_{\ell \in [p] \setminus i \cup j} \langle X_i, X_\ell \rangle^3 \langle X_j, X_\ell \rangle \right) \\ &= \frac{\gamma k}{p^2 n^3} \left(n^3 \langle X_i, X_j \rangle \pm \tilde{O}(n^{2.5}) \pm \tilde{O}(pn^{1.5}) \right) \\ &= \widetilde{M}(x_i x_j) \pm \tilde{O}\left(\frac{\gamma k}{p^2 n^{1.5}}\right) \pm \tilde{O}\left(\frac{\gamma k}{pn^{1.5}}\right) \\ &= \widetilde{M}(x_i x_j) \pm \tilde{O}\left(\frac{\gamma k}{pn^{1.5}}\right) \end{aligned}$$

where in the second equality we use equation (P3), and $p \geq n$.

- Equation (4.5):

Note that for distinct i, j, s, t , by equation (P4), we have

$$|P(x_i x_j x_s x_t)| = \frac{\gamma k}{p^2 n^3} \left| \sum_{\ell \in [p]} \langle X_i, X_\ell \rangle \langle X_j, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \right| \leq \tilde{O}\left(\frac{\gamma k}{p^{1.5} n}\right)$$

Therefore taking the sum over all distinct i, j, s, t we have

$$\sum_{i, j, s, t \text{ distinct}} |P(x_i x_j x_s x_t)| \leq \tilde{O}\left(\frac{\gamma k}{p^{1.5} n}\right) \cdot p^4 = \tilde{O}\left(\frac{\gamma p^{2.5} k}{n}\right) \leq k^4/4 \quad (6.1)$$

where we used $k \gg \gamma^{7/6} \sqrt{n}$, which implies that $k^3 \gg \gamma p^{2.5}/n$.

By equation (4.2) and equation (4.3), we have that

$$\sum_{i,j} |P(x_i^3 x_j)| \leq \sum_{i,j} |\widetilde{M}(x_i x_j)| + p^2 \cdot \widetilde{O}\left(\frac{\gamma k}{pn^{1.5}}\right) + p \cdot \widetilde{O}\left(\frac{\gamma k}{pn}\right) \leq k^2/2 + \widetilde{O}(\gamma k \sqrt{n}) \leq k^2 \quad (6.2)$$

where we used the fact that $p \leq n^2$ and $k \gg \gamma \sqrt{n}$.

For distinct i, s, t , we have that

$$\begin{aligned} \left| \sum_{\ell \in [p]} \langle X_i, X_\ell \rangle^2 \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \right| &= \left| \sum_{\ell \neq i, s, t} \langle X_i, X_\ell \rangle^2 \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \right| + |\langle X_i, X_i \rangle^2 \langle X_s, X_i \rangle \langle X_t, X_i \rangle| \\ &\quad + |\langle X_i, X_s \rangle^2 \langle X_s, X_s \rangle \langle X_t, X_s \rangle| + |\langle X_i, X_t \rangle^2 \langle X_s, X_t \rangle \langle X_t, X_t \rangle| \\ &= \widetilde{O}(pn^2) + \widetilde{O}(n^3) + \widetilde{O}(n^{2.5}) = \widetilde{O}(pn^2) \end{aligned}$$

It follows that

$$|P(x_i^2 x_s x_t)| = \frac{\gamma k}{p^2 n^3} \left| \sum_{\ell \in [p]} \langle X_i, X_\ell \rangle^2 \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \right| \leq \widetilde{O}\left(\frac{\gamma k}{pn}\right)$$

and therefore,

$$\sum_{i, s, t \text{ distinct}} |P(x_i^2 x_s x_t)| \leq p^3 \cdot \widetilde{O}\left(\frac{\gamma k}{pn}\right) = \widetilde{O}\left(\frac{\gamma k p^2}{n}\right) \quad (6.3)$$

Therefore, combining equation (6.1), (6.2), (6.3), we obtain that

$$\sum_{i, j, s, t} |P(x_i x_j x_s x_t)| \leq k^2 + \widetilde{O}\left(\frac{\gamma k p^2}{n}\right) + k^4/4 \leq k^4/3$$

- Equation (4.4):

Finally it remains to bound \mathcal{E} . Note that \mathcal{E} is a sum of submatrices of P and therefore it is PSD. Moreover,

$$\begin{aligned} \mathcal{E}_{ss} = \sum_{i \in [p]} P(x_i^2 x_s^2) &= \frac{\gamma k}{p^2 n^3} \sum_i \sum_\ell \langle X_i, X_\ell \rangle^2 \langle X_s, X_\ell \rangle^2 \\ &= \frac{\gamma k}{p^2 n^3} \sum_i \langle X_i, X_i \rangle^2 \sum_\ell \langle X_s, X_\ell \rangle^2 \\ &\leq \frac{\gamma k}{p^2 n^3} \cdot \widetilde{O}(p^2 n^2) = \widetilde{O}\left(\frac{\gamma k}{n}\right) \end{aligned}$$

where the last inequality uses equation (P2). Finally we bound \mathcal{E}_{st} using equation (P6)

$$\begin{aligned} \sum_{i \in [p]} P(x_i^2 x_s x_t) &= \frac{\gamma k}{p^2 n^3} \sum_i \sum_\ell \langle X_i, X_\ell \rangle \langle X_i, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \\ &\leq \frac{\gamma k}{p^2 n^3} \tilde{O}(p^2 n^{1.5}) = \tilde{O}\left(\frac{\gamma k}{n^{1.5}}\right) \end{aligned}$$

□

Proof of Lemma 4.5. Again we verify equation (4.6), (4.7) and (4.8) in order.

- Equation (4.6): By definition we have that for any i, j ,

$$Q(x_i^3 x_j) = \frac{\gamma k^2}{p^3 n} \cdot 3n \langle X_i, X_j \rangle = \frac{3\gamma k^2}{p^3} \langle X_i, X_j \rangle = \frac{3k}{p} \widetilde{M}(x_i x_j)$$

- Equation (4.8): For the sparsity constraint, we note first that for distinct i, j, s, t , using property (P2), we have

$$|Q(x_i x_j x_s x_t)| \leq \frac{\gamma k^2}{p^3 n} \cdot \tilde{O}(n) = \tilde{O}\left(\frac{\gamma k^2}{p^3}\right)$$

and therefore taking sum, we have

$$\sum_{i, j, s, t \text{ distinct}} |Q(x_i x_j x_s x_t)| \leq \tilde{O}(\gamma k^2 p) \leq k^4/6 \quad (6.4)$$

where we used the fact that $k^2 \gg \gg c^2 n$. We bound other terms as follows:

For any $i, j, s, t \in [p]$, we have that

$$Q(x_i x_j x_s x_t) \leq \frac{\gamma k^2}{p^3 n} \cdot 3n^2 = \frac{3\gamma k^2 n}{p^3}$$

There are only at most $O(p^3)$ different choices of i, j, s, t such that i, j, s, t are not distinct, therefore we have

$$\sum_{i, j \text{ not distinct}} |Q(x_i^3 x_j)| \leq \frac{3\gamma k^2 n}{p^3} \cdot O(p^3) \leq k^4/6 \quad (6.5)$$

where we used the fact that $k \gg \gamma \sqrt{n}$ and $\gamma \geq 1$.

Combining equation (6.4) and (6.5), we obtain that

$$\sum_{i, j, s, t} |Q(x_i x_j x_s x_t)| \leq k^4/3$$

- Equation (4.7): For any s, t , we have

$$\begin{aligned}
\sum_i Q(x_i^2 x_s x_t) &= \frac{\gamma k^2}{p^3 n} \sum_{i \in [p]} (n \langle X_s, X_t \rangle + 2 \langle X_i, X_s \rangle \langle X_i, X_t \rangle) \\
&= \frac{\gamma k^2}{p^2} \langle X_s, X_t \rangle + \frac{2\gamma k^2}{p^3 n} \sum_{i \in [p]} \langle X_i, X_s \rangle \langle X_i, X_t \rangle
\end{aligned}$$

Therefore $\mathcal{G}_{st} = \frac{2\gamma k^2}{p^3 n} \sum_{i \in [p]} \langle X_i, X_s \rangle \langle X_i, X_t \rangle$ forms a PSD matrix. Moreover, when $s \neq t$, using equation (P5), we have that

$$\begin{aligned}
\sum_i Q(x_i^2 x_s x_t) &= k \widetilde{M}(x_s x_t) \pm \frac{2\gamma k^2}{p^3 n} \cdot \widetilde{O}(p\sqrt{n}) \\
&= k \widetilde{M}(x_s x_t) \pm \widetilde{O}\left(\frac{\gamma k^2}{p^2 \sqrt{n}}\right)
\end{aligned}$$

When $s = t$, we have that

$$\begin{aligned}
\sum_i Q(x_i^2 x_s^2) &= k \widetilde{M}(x_s x_t) \pm \frac{2\gamma k^2}{p^3 n} \cdot \widetilde{O}(pn) \\
&= k \widetilde{M}(x_s x_t) \pm \widetilde{O}\left(\frac{\gamma k^2}{p^2}\right)
\end{aligned}$$

□

7 Pseudo-randomness of X

In this section, we prove that basically as long as the noise model is subgaussian and has variance 1 (which generalizes the standard Bernoulli and Gaussian distributions), after normalizing the rows of the data matrix $X \sim H_0$, it satisfies the pseudorandomness condition 4.1.

Theorem 7.1. Suppose independent random variables $X_1, \dots, X_p \in \mathbb{R}^n$ satisfy that for any i , X_i has a i.i.d entries with mean zero, variance 1, and subgaussian variance proxy⁶ $O(1)$, then the matrix \bar{X} with $\frac{X_i^T}{\|X_i\|}$ as rows satisfies the pseudorandomness condition 4.1.

The proof of the Theorem relies on the following Proposition and Theorem 7.4. The proposition says that $\frac{X_i^T}{\|X_i\|}$ still satisfies *good* properties like symmetry and that each entries has a subgaussian tail, even though its entries are no longer independent due to normalization. These properties will be encapsulated in the definition of a *good* random variable following the proposition. Then we prove in Theorem 7.4 that these properties suffice for establishing the pseudorandomness Condition 4.1 with high probability. We will heavily use the ψ_α -Orlicz norm (denoted $\|\cdot\|_{\psi_\alpha}$) of a random variable, defined in Definition 8.1, and its properties, summarized in the next (toolbox) section. Intuitively, $\|\cdot\|_{\psi_2}$ norm is a succinct and convenient way to capture the “subgaussianity” of a random variable.

⁶A real random variable X is subgaussian with variance proxy σ^2 if it has similar tail behavior as gaussian distribution with variance σ^2 , and formally if for any $t \in \mathbb{R}$, $\mathbb{E}[\exp(tX)] \leq \exp(t^2 \sigma^2 / 2)$

Proposition 7.2. Suppose $y \in \mathbb{R}^n$ has i.i.d entries with variance 1 and mean zero, and gaussian variance proxy $O(1)$, then random variable $x = \frac{y}{\|y\|}$ satisfies the following properties:

1. $\|x\|^2 = n$, almost surely.
2. for any vector $a \in \mathbb{R}^n$ with $\|a\|^2 \leq 2n$, $\|\langle x, a \rangle\|_{\psi_2}^2 \leq O(n)$.
3. $\|x\|_\infty \leq \tilde{O}(1)$
4. $\mathbb{E}[x_i^2] = 1$, $\mathbb{E}[x_i^4] = C_4$, and $\mathbb{E}[x_i^2 x_j^2] = C_{2,2}$ for all i and $j \neq i$, where $C_4, C_{2,2} = O(1)$ are constants that don't depend on i, j
5. For any monomial x^α with an odd degree, $\mathbb{E}[x^\alpha] = 0$.

For simplicity, we call a random variable *good* if it satisfies the five properties listed in the proposition above. Goodness will be invoked in most statements below.

Definition 7.3 (goodness). A random variable $x \in \mathbb{R}^n$ is called *good*, if it satisfies the conclusion of Proposition 7.2.

We will show a random matrix X with *good* rows satisfies the pseudo-randomness Condition 4.1 with high probability.

Theorem 7.4. Suppose independent n -dimensional random vectors X_1, \dots, X_p with $p \geq n$ are all *good*, then X_1, \dots, X_p satisfies the pseudorandomness condition 4.1 with high probability.

The general approach to prove the theorem is just to use the concentration of measure. The only caveat here is that in most of cases, the random variables that we are dealing with are not bounded a.s. so we can't use Chernoff bound or Bernstein inequality directly. However, though these random variables are not bounded a.s., they typically have a light tail, that is, their ψ_α norms can be bounded. Then we are going to apply Theorem 8.4 of Ledoux and Talagrand's, a extended version of Bernstein inequality with only ψ_α norm boundedness required. We will also use other known technical results listed in the toolbox Section 8.

Proof of Theorem 7.4. Equation (P1) and (P2) follows the assumptions on X_i 's and union bound. Equation (P3) is proved in Lemma 7.5 by taking $u = X_s$ and $v = X_t$ and view the rest of X_i 's as Z_j 's in the statement of Lemma 7.5. Equation (P4) is proved in Lemma 7.6, (P5) in Lemma 7.8, (P6) in Lemma 7.10, and equation P7 is proved in Lemma 7.15. \square

Lemma 7.5. For any *good* random variable x , we have that for fixed u, v with $\|u\|^2 = \|v\|^2 = n$, $|u|_\infty \leq \tilde{O}(1)$, $|v|_\infty \leq \tilde{O}(1)$, and $\langle u, v \rangle \leq \tilde{O}(\sqrt{n})$,

$$|\mathbb{E}[\langle x, u \rangle^3 \langle x, v \rangle]| \leq \tilde{O}(n^{1.5})$$

and moreover, for $p \geq n$ and a sequence of *good* independent random variables Z_1, \dots, Z_p , we have that with high probability,

$$\left| \sum_{i=1}^p \langle Z_i, u \rangle^3 \langle Z_i, v \rangle \right| \leq \tilde{O}(n^{1.5} p)$$

Proof. We calculate the expectation as follows

$$\begin{aligned}
\mathbb{E} [\langle x, u \rangle^3 \langle x, v \rangle] &= \mathbb{E} \left[\left(\sum_i u_i^2 x_i^2 + 2 \sum_{i < j} u_i u_j x_i x_j \right) \left(\sum_i v_i u_i x_i^2 + \sum_{i \neq j} u_i v_j x_i x_j \right) \right] \\
&= \mathbb{E} \left[\sum_i u_i^3 v_i x_i^4 \right] + \mathbb{E} \left[\sum_{i \neq j} u_i^2 u_j v_j x_i^2 x_j^2 \right] + \mathbb{E} \left[\sum_{i \neq j} u_i u_j (u_i v_j + v_i u_j) x_i^2 x_j^2 \right] \\
&= (C_4 - C_{2,2}) \sum_i u_i^3 v_i + C_{2,2} n \sum_i u_i v_i + C_{2,2} \sum_{i \neq j} (u_i^2 u_j v_j + u_j^2 u_i v_i) \\
&= (C_4 - 3C_{2,2}) \sum_i u_i^3 v_i + 3C_{2,2} n \sum_i u_i v_i
\end{aligned}$$

Therefore by our assumption on u and v we obtain that

$$|\mathbb{E} [\langle x, u \rangle^3 \langle x, v \rangle]| \leq \tilde{O}(n) + O(n)|\langle u, v \rangle| \leq \tilde{O}(n^{1.5})$$

Now we prove the second statement. Since $\|\langle Z_i, u \rangle\|_{\psi_2} \leq O(\sqrt{n})$, by Lemma 8.5 we have that $\|\langle Z_i, u \rangle^3 \langle Z_i, v \rangle\|_{\psi_{1/2}} \leq O(n^2)$, and it follows Lemma 8.6 that $\|\langle Z_i, u \rangle^3 \langle Z_i, v \rangle - \mathbb{E} [\langle Z_i, u \rangle^3 \langle Z_i, v \rangle]\|_{\psi_{1/2}} \leq O(n^2)$. Then by Lemma 8.4 we obtain that with high probability,

$$\sum_{i=1}^p \langle Z_i, u \rangle^3 \langle Z_i, v \rangle - \mathbb{E} \left[\sum_{i=1}^p \langle Z_i, u \rangle^3 \langle Z_i, v \rangle \right] \leq \tilde{O}(n^2 \sqrt{p})$$

Note that we have proved that $|\mathbb{E} [\sum_{i=1}^p \langle Z_i, u \rangle^3 \langle Z_i, v \rangle]| = \tilde{O}(n^{1.5})$, therefore we obtain the desired result. \square

Lemma 7.6. Suppose $p \geq n$ and X_1, \dots, X_p are *good* independent random variables, then with high probability, for any distinct i, j, s, t ,

$$\left| \sum_{\ell \in [p]} \langle X_i, X_\ell \rangle \langle X_j, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \right| \leq \tilde{O}(n^2 \sqrt{p})$$

Proof. Fixing i, j, s, t , we can write

$$\begin{aligned}
\sum_{\ell \in [p]} \langle X_i, X_\ell \rangle \langle X_j, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle &= \sum_{\ell \in [p] \setminus \{i, j, s, t\}} \langle X_i, X_\ell \rangle \langle X_j, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \\
&\quad + n \langle X_j, X_i \rangle \langle X_s, X_i \rangle \langle X_t, X_i \rangle + n \langle X_i, X_j \rangle \langle X_s, X_j \rangle \langle X_t, X_j \rangle \\
&\quad + n \langle X_i, X_s \rangle \langle X_j, X_s \rangle \langle X_t, X_s \rangle + n \langle X_i, X_t \rangle \langle X_j, X_t \rangle \langle X_s, X_t \rangle
\end{aligned}$$

Using Lemma 7.7, the first term on RHS is bounded by $\tilde{O}(n^2 \sqrt{p})$ with high probability over the randomness of $X_\ell, \ell \in [p] \setminus \{i, j, s, t\}$. The rest of the four terms are bounded by $\tilde{O}(n^{2.5})$. Therefore putting together $\|\sum_{\ell \in [p]} \langle X_i, X_\ell \rangle \langle X_j, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle\| \leq \tilde{O}(n^2 \sqrt{p})$ for any fixed i, j, s, t with high probability and taking union bound we get the result. \square

Lemma 7.7. For any good random variable x , and for fixed a, b, c, d such that $\max\{|a|_\infty, |b|_\infty, |c|_\infty, |d|_\infty\} = \tilde{O}(1)$, and all the pair-wise inner products between a, b, c, d have magnitude at most $\tilde{O}(\sqrt{n})$, we have that

$$|\mathbb{E}[\langle x, a \rangle \langle x, b \rangle \langle x, c \rangle \langle x, d \rangle]| = \tilde{O}(n)$$

and moreover, for $p \geq n$ and a sequence independent random variable Z_1, \dots, Z_p such that each Z_i satisfies the conclusion of proposition 7.2, we have that with high probability,

$$\left| \sum_{i=1}^p \langle Z_i, a \rangle \langle Z_i, b \rangle \langle Z_i, c \rangle \langle Z_i, d \rangle \right| \leq \tilde{O}(n^2 \sqrt{p})$$

Proof. We calculate the mean

$$\mathbb{E}[\langle x, a \rangle \langle x, b \rangle \langle x, c \rangle \langle x, d \rangle] = \mathbb{E} \left[\sum_{i \in [p]} a_i b_i c_i d_i x_i^4 + \left\{ \sum_{i \neq j} a_i b_i c_j d_j x_i^2 x_j^2 \right\} \right]$$

where we use $\left\{ \sum_{i \neq j} a_i b_i c_j d_j x_i^2 x_j^2 \right\}$ to denote the sum of $a_i b_i c_j d_j x_i^2 x_j^2$ and all its permutations with respect to a, b, c, d .

Note that

$$\left| \mathbb{E} \left[\sum_{i \neq j} a_i b_i c_j d_j x_i^2 x_j^2 \right] \right| = C_{2,2} \left| \langle a, b \rangle \langle c, d \rangle - \sum_{i \in [p]} a_i b_i c_i d_i \right| \leq \tilde{O}(n)$$

and

$$\left| \mathbb{E} \left[\sum_{i \in [p]} a_i b_i c_i d_i x_i^4 \right] \right| = \left| C_4 \sum_{i \in [p]} a_i b_i c_i d_i \right| \leq \tilde{O}(n)$$

and therefore we have $|\mathbb{E}[\langle x, a \rangle \langle x, b \rangle \langle x, c \rangle \langle x, d \rangle]| \leq \tilde{O}(n)$.

Since $\langle x, a \rangle$ has ψ_2 norm \sqrt{n} and similar for the other three terms, we have that by Lemma 8.5 that $\|\langle x, a \rangle \langle x, b \rangle \langle x, c \rangle \langle x, d \rangle\|_{\psi_{1/2}} \leq O(n^2)$. Therefore using Theorem 8.4 we have that

$$\left\| \sum_{i=1}^p \langle Z_i, a \rangle \langle Z_i, b \rangle \langle Z_i, c \rangle \langle Z_i, d \rangle - \mathbb{E} \left[\sum_{i=1}^p \langle Z_i, a \rangle \langle Z_i, b \rangle \langle Z_i, c \rangle \langle Z_i, d \rangle \right] \right\|_{\psi_{1/2}} \leq \tilde{O}(n^2 \sqrt{p})$$

□

Lemma 7.8. Suppose $p \geq n$ and X_1, \dots, X_p are *good* independent random variables, then with high probability, for any distinct s, t ,

$$\left| \sum_{i \in [p]} \langle X_i, X_s \rangle \langle X_i, X_t \rangle \right| \leq \tilde{O}(p\sqrt{n})$$

Proof. With high probability over the randomness of $X_i, i \in [p] \setminus \{s, t\}$,

$$\sum_{i \in [p]} \langle X_i, X_s \rangle \langle X_i, X_t \rangle = \sum_{i \in [p] \setminus \{s, t\}} \langle X_i, X_s \rangle \langle X_i, X_t \rangle + 2\langle X_s, X_t \rangle \leq \tilde{O}(p\sqrt{n}) + \tilde{O}(\sqrt{n})$$

where the last inequality is by Lemma 7.9. Taking union bound we complete the proof. □

Lemma 7.9. For $p \geq n$ and a sequence of *good* independent random variable Z_1, \dots, Z_p , and any two fixed vectors u, v with $|u|_\infty \leq \tilde{O}(1)$ and $|v|_\infty \leq \tilde{O}(1)$, and $\langle u, v \rangle \leq \tilde{O}(\sqrt{n})$, we have that with high probability,

$$\left| \sum_{i \in [p]} \langle Z_i, u \rangle \langle Z_i, v \rangle \right| \leq \tilde{O}(p\sqrt{n})$$

Proof. $\mathbb{E}[\langle Z_i, u \rangle \langle Z_i, v \rangle] = \langle u, v \rangle \leq \tilde{O}(\sqrt{n})$, and therefore $\left| \mathbb{E} \left[\sum_{i \in [p]} \langle Z_i, u \rangle \langle Z_i, v \rangle \right] \right| \leq \tilde{O}(p\sqrt{n})$. Note that $\|\langle Z_i, u \rangle\|_{\psi_2} \leq O(\sqrt{n})$ and therefore $\|\langle Z_i, u \rangle \langle Z_i, v \rangle\|_{\psi_1} \leq O(n)$. By Theorem 8.4, we have the desired result. \square

Lemma 7.10. Suppose $p \geq n$ and X_1, \dots, X_p are *good* independent random variables, then with high probability, for any distinct s, t ,

$$\sum_{i, \ell \in [p]} \langle X_i, X_\ell \rangle \langle X_i, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \leq \tilde{O}(p^2 n^{1.5})$$

Proof. We expand the target as follows:

$$\begin{aligned} \sum_{i, \ell \in [p]} \langle X_i, X_\ell \rangle \langle X_i, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle &= \sum_{i \in [p], \ell \in [p] \setminus s \cup t} \langle X_i, X_\ell \rangle \langle X_i, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \\ &\quad + \sum_i \langle X_i, X_s \rangle^2 \langle X_s, X_s \rangle \langle X_t, X_s \rangle + \sum_i \langle X_i, X_t \rangle^2 \langle X_t, X_t \rangle \langle X_s, X_t \rangle \\ &= \sum_{i \in [p] \setminus s \cup t, \ell \in [p] \setminus s \cup t} \langle X_i, X_\ell \rangle \langle X_i, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \\ &\quad + \sum_i \langle X_i, X_s \rangle^2 \langle X_s, X_s \rangle \langle X_t, X_s \rangle + \sum_i \langle X_i, X_t \rangle^2 \langle X_t, X_t \rangle \langle X_s, X_t \rangle \\ &\quad + \sum_{\ell \in [p] \setminus s \cup t} \langle X_s, X_\ell \rangle^3 \langle X_\ell, X_t \rangle + \sum_{\ell \in [p] \setminus s \cup t} \langle X_t, X_\ell \rangle^3 \langle X_\ell, X_s \rangle \end{aligned}$$

By equation (P3), we have that

$$\sum_{\ell \in [p] \setminus s \cup t} \langle X_s, X_\ell \rangle^3 \langle X_\ell, X_t \rangle \leq \tilde{O}(pn^{1.5})$$

Since $\langle X_s, X_t \rangle \leq \tilde{O}(\sqrt{n})$ and $\sum_{i \in [p]} \langle X_i, X_s \rangle^2 = n^2 + \sum_{i \neq s} \langle X_i, X_s \rangle^2 \leq \tilde{O}(np)$, we have that

$$\sum_i \langle X_i, X_s \rangle^2 \langle X_s, X_s \rangle \langle X_t, X_s \rangle \leq \tilde{O}(pn^{2.5})$$

Invoking Lemma 7.11 with $u = X_s$ and $v = X_t$ fixed and view $X_\ell, \ell \in [p] \setminus s \cup t$ as random variables Z_i 's, we have that with high probability,

$$\sum_{i \in [p] \setminus s \cup t, \ell \in [p] \setminus s \cup t} \langle X_i, X_\ell \rangle \langle X_i, X_\ell \rangle \langle X_s, X_\ell \rangle \langle X_t, X_\ell \rangle \leq \tilde{O}(p^2 n^{1.5})$$

Hence combining the three equations above, taking union bound over all choices of s, t , we obtain the desired result. \square

Lemma 7.11. For $p \geq n$ and a sequence of *good* independent random variables Z_1, \dots, Z_p , and any two fixed vectors u, v with $|u|_\infty \leq \tilde{O}(1)$ and $|v|_\infty \leq \tilde{O}(1)$, and $\langle u, v \rangle \leq \tilde{O}(\sqrt{n})$, we have that with high probability,

$$\sum_{i \in [p]} \sum_{j \in [p]} \langle Z_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle \leq \tilde{O}(p^2 n^{1.5})$$

Proof. We first extract the consider those cases with $i = j$ separately by expanding

$$\begin{aligned} \sum_{i \in [p]} \sum_{j \in [p]} \langle Z_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle &= \sum_{i \neq j} \langle Z_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle + \sum_i \langle Z_i, Z_i \rangle^2 \langle Z_i, u \rangle \langle Z_i, v \rangle \\ &= \sum_{i \neq j} \langle Z_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle + \tilde{O}(pn^{2.5}) \end{aligned} \quad (7.1)$$

where the the last line uses Lemma 7.9. Let Y_1, \dots, Y_p are independent random variables that have the same distribution as Z_1, \dots, Z_p , respectively, then by Theorem 8.7, we can decouple the sum of functions of Z_i, jZ_j into a sum that of functions of Z_i and Y_j ,

$$\Pr \left[\sum_{i \neq j} \langle Z_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle \geq t \right] \leq C \Pr \left[\sum_{i \neq j} \langle Y_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle \geq t/C \right]$$

Now we can invoke Lemma 7.12 which deals with RHS of the equation above, and obtain that with high probability

$$\sum_{i \neq j} \langle Y_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle \leq \tilde{O}(p^2 n^{1.5})$$

Therefore, with high probability,

$$\sum_{i \neq j} \langle Z_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle \leq \tilde{O}(p^2 n^{1.5})$$

Then combine with equation (7.1) we obtain the desired result. \square

Lemma 7.12. For $p \geq n$ and a sequence of *good* independent random variables Z_1, \dots, Z_p , let Y_1, \dots, Y_p be independent random variables which have the same distribution as Z_1, \dots, Z_p , respectively, then for any two fixed vectors u, v with $|u|_\infty \leq \tilde{O}(1)$ and $|v|_\infty \leq \tilde{O}(1)$, and $\langle u, v \rangle \leq \tilde{O}(\sqrt{n})$, with high probability,

$$\sum_{i \in [p]} \sum_{j \in [p]} \langle Y_i, Z_j \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle \leq \tilde{O}(p^2 n^{1.5})$$

Proof. Let $B = \sum_{i \in [p]} Y_i Y_i^T$. Therefore by Lemma 7.13, we have that with high probability over the randomness of Y , $\|B\|_2 \leq \tilde{O}(p)$, $\text{tr}(B) = pn$. Moreover, by Lemma 7.9, we have that with high probability, $|u^T B v| \leq \tilde{O}(p\sqrt{n})$. Note that these bounds only depend on the randomness of Y , and conditioning on all these bounds are true, we can still use the randomness of Z_i 's for concentration. We invoke Lemma 7.14 and obtain that

$$\left| \sum_{i, j \in [p]} \langle Z_j, Y_i \rangle^2 \langle Z_j, u \rangle \langle Z_j, v \rangle \right| = \left| \sum_{j=1}^p Z_j^T B Z_j \langle Z_j, u \rangle \langle Z_j, v \rangle \right| \leq \tilde{O}(p^2 n^{1.5})$$

\square

Lemma 7.13. For $p \geq n$ and a sequence of *good* independent random variables Z_1, \dots, Z_p , we have that with high probability,

$$\left\| \sum_{i \in [p]} Z_i Z_i^T \right\| \leq \tilde{O}(p)$$

Proof. We use matrix Bernstein inequality. First of all, we have that $\mathbb{E}[Z_i Z_i^T] = I_{n \times n}$, and therefore $\mathbb{E} \left[\sum_{i \in [p]} Z_i Z_i^T \right] = p I_{n \times n}$. Moreover, we check the variance of the $Z_i Z_i^T$:

$$\mathbb{E}[Z_i Z_i^T Z_i Z_i^T] = n \mathbb{E}[Z_i Z_i^T] = n I_{n \times n}$$

Finally we observe that $\|Z_i Z_i^T\| \leq n$. Thus applying matrix Bernstein inequality we obtain that with high probability,

$$\left\| \sum_{i \in [p]} Z_i Z_i^T - p I_{n \times n} \right\| \leq \tilde{O}(\sqrt{np} + n) = \tilde{O}(\sqrt{np})$$

□

Lemma 7.14. For $p \geq n$ and a sequence of *good* independent random variables Z_1, \dots, Z_p , and for any fixed symmetric PSD matrix $B \in \mathbb{R}^{n \times n}$ with $\|B\| \leq \tilde{O}(p)$, $\text{tr}(B) \leq 2pn$, and any two fixed vectors u, v with $|u|_\infty \leq \tilde{O}(1)$ and $|v|_\infty \leq \tilde{O}(1)$, and $\langle u, v \rangle \leq \tilde{O}(\sqrt{n})$, we have that with high probability over the randomness of Z_i 's,

$$\left| \sum_{i=1}^p Z_i^T B Z_i \langle Z_i, u \rangle \langle Z_i, v \rangle \right| \leq \tilde{O}(p^2 n^{1.5})$$

Proof. Let $W = x^T B x \langle x, u \rangle \langle x, v \rangle$, where x is a random variable that satisfies the conclusion of Proposition 7.2. We first calculate the expectation of W ,

$$\begin{aligned} \mathbb{E}[W] &= \mathbb{E} \left[\left(\sum_i B_{ii} x_i^2 + \sum_{i \neq j} x_i x_j B_{ij} \right) \left(\sum_i u_i v_i x_i^2 + \sum_{i \neq j} x_i x_j u_i v_j \right) \right] \\ &= (C_4 - C_{2,2}) \sum_i B_{ii} u_i v_i + C_{2,2} \text{tr}(B) \langle u, v \rangle + \mathbb{E} \left[\sum_{i \neq j} B_{ij} (u_i v_j + u_j v_i) x_i^2 x_j^2 \right] \\ &= (C_4 - 3C_{2,2}) \sum_i B_{ii} u_i v_i + \text{tr}(B) \langle u, v \rangle \end{aligned}$$

Therefore by the fact that $|u|_\infty \leq \tilde{O}(1)$ and $\text{tr}(B) \leq 2pn$, we obtain that $|\mathbb{E}[W]| \leq \tilde{O}(pn^{1.5})$.

Observe that $Z_j^T B Z_j \leq \tilde{O}(pn)$ a.s. (with respect to the randomness of Z_j), and $\|\langle Z_i, u \rangle \langle Z_i, v \rangle\|_{\psi_1} \leq O(n)$, therefore we have that $\|Z_j^T B Z_j \langle Z_i, u \rangle \langle Z_i, v \rangle\|_{\psi_1} \leq O(pn^2)$. Using Theorem 8.4, we obtain that with high probability,

$$\left| \sum_{i=1}^p Z_i^T B Z_i \langle Z_i, u \rangle \langle Z_i, v \rangle - \mathbb{E} \left[\sum_{i=1}^p Z_i^T B Z_i \langle Z_i, u \rangle \langle Z_i, v \rangle \right] \right| \leq \tilde{O}(n^2 p^{1.5})$$

Using the fact that $\mathbb{E} [Z_i^T B Z_i \langle Z_i, u \rangle \langle Z_i, v \rangle] \leq \tilde{O}(pn^{1.5})$ we obtain that with high probability

$$\left| \sum_{i=1}^p Z_i^T B Z_i \langle Z_i, u \rangle \langle Z_i, v \rangle \right| \leq \tilde{O}(p^2 n^{1.5})$$

□

Lemma 7.15. Suppose $p \geq n$ and X_1, \dots, X_p are *good* independent random variables, then with high probability,

$$\|X X^T\|_F^2 \geq (1 - o(1))p^2 n$$

Proof. We first i and examine $\sum_{j \neq i} \langle X_j, X_i \rangle^2$ first. We have that $\mathbb{E}[\sum_{j \neq i} \langle X_j, X_i \rangle^2] = (p-1)\|X_i\|^2 = (p-1)n$. Moreover, $\|\langle X_j, X_i \rangle^2\|_{\psi_1} \leq O(n)$ (where X_j is viewed as random and X_i is viewed as fixed). Therefore by Theorem 8.4, we obtain that with high probability over the randomness of X_j 's, ($j \neq i$), $\sum_{j \neq i} \langle X_j, X_i \rangle^2 = (p-1)n \pm \tilde{O}(n\sqrt{p}) = (1 \pm o(1))pn$. Therefore taking union bound over all i , and taking the sum we obtain that

$$\|X X^T\|_F^2 \geq \sum_i \sum_{j \neq i} \langle X_j, X_i \rangle^2 \geq (1 - o(1))p^2 n$$

□

8 Toolbox

This section contains a collection of known technical results which are useful in proving the concentration bounds of Section 7. We note that when the data matrix X takes uniformly $\{\pm 1\}$ entries, then X satisfies Proposition 7.2 without any normalization and actually due to the independence of the entries, it's much easier to prove that it satisfies Condition 4.1.

Definition 8.1 (Orlicz norm $\|\cdot\|_{\psi_\alpha}$). For $1 \leq \alpha < \infty$, let $\psi_\alpha(x) = \exp(x^\alpha) - 1$. For $0 < \alpha < 1$, let $\psi_\alpha(x) = x^\alpha - 1$ for large enough $x \geq x_\alpha$, and ψ_α is linear in $[0, x_\alpha]$. The Orlicz norm ψ_α of a random variable X is defined as

$$\|X\|_{\psi_\alpha} \triangleq \inf\{c \in (0, \infty) \mid \mathbb{E}[\psi_\alpha(|X|/c)] \leq 1\} \quad (8.1)$$

Note that by definition ψ_α is convex and increasing. The following Theorem of Ledoux and Talagrand's is our main tool for proving concentration inequalities in Section 7.

Theorem 8.2 (Theorem 6.21 of [LT13]). There exists a constant K_α depending on α such that for a sequence of independent mean zero random variables X_1, \dots, X_n in L_{ψ_α} , if $0 < \alpha \leq 1$,

$$\left\| \sum_i X_i \right\|_{\psi_\alpha} \leq K_\alpha \left(\left\| \sum_i X_i \right\|_1 + \left\| \max_i \|X_i\| \right\|_{\psi_\alpha} \right) \quad (8.2)$$

and if $1 < \alpha \leq 2$,

$$\left\| \sum_i X_i \right\|_{\psi_\alpha} \leq K_\alpha \left(\left\| \sum_i X_i \right\|_1 + \left(\sum_i \|X_i\|_{\psi_\alpha}^\beta \right)^{1/\beta} \right) \quad (8.3)$$

where $1/\alpha + 1/\beta = 1$.

The following convenient Lemma allows us to control the second part of RHS of (8.2) easily.

Lemma 8.3 ([vdVW00]). There exists absolute constant c , such that for any real valued random variables X_1, \dots, X_n , we have that

$$\left\| \max_{1 \leq i \leq n} |X_i| \right\|_{\psi_\alpha} \leq c\psi_\alpha^{-1}(n) \max_{1 \leq i \leq n} \|X_i\|_{\psi_\alpha}$$

Using Lemma 8.3 and Theorem 8.2, we obtain straightforwardly the following theorem that will be used many times for proving concentration bounds in this paper.

Theorem 8.4. For any $0 < \alpha \leq 1$, there exists a constant K_α such that for a sequence of independent random variables X_1, \dots, X_n ,

$$\left\| \sum_i X_i - \mathbb{E}[\sum_i X_i] \right\|_{\psi_\alpha} \leq K_\alpha \sqrt{n} \log n \cdot \max_i \|X_i\|_{\psi_\alpha} \quad (8.4)$$

which implies that with high probability over the randomness of X_i 's,

$$\left| \sum_i X_i - \mathbb{E}[\sum_i X_i] \right| \leq \tilde{O}(K_\alpha \sqrt{n} \cdot \max_i \|X_i\|_{\psi_\alpha})$$

The following two lemmas are used to bound the Orlicz norms of random variables.

Lemma 8.5. There exists constant D_α depending on α such that, if two (possibly correlated) random variables X, Y have ψ_α Orlicz norm bounded by $\|X\|_{\psi_\alpha} \leq a$ and $\|Y\|_{\psi_\alpha} \leq b$ then $\|XY\|_{\psi_{\alpha/2}} \leq D_\alpha ab$

Proof. For any $x, y, a, b, \alpha > 0$,

$$\begin{aligned} \exp(|xy|)^{\alpha/2} - 1 &\leq \exp\left(\frac{1}{2}|x|^\alpha + \frac{1}{2}|y|^\alpha\right) - 1 \\ &\leq \frac{1}{2}((\exp|x|^\alpha - 1) + (\exp|y|^\alpha - 1)) \end{aligned}$$

Moreover, note that by definition of ψ_α , there exists constant C_α and C'_α such that for $x \geq 0$, $C'_\alpha(\exp(x^\alpha) - 1) \geq \psi_\alpha(x) \geq C_\alpha(\exp(x^\alpha) - 1)$. Therefore we have that there exists a constant E_α such that $\psi_{\alpha/2}(|xy|) \leq \frac{E_\alpha}{2}(\psi_\alpha(|x|) + \psi_\alpha(|y|))$. Also note that for any constant c , there exists constant c' such that $\psi_\alpha(x/c') \leq \psi_\alpha(x)/c$. Therefore, choosing D_α such that $\psi_{\alpha/2}(x/D_\alpha) \leq \psi_\alpha(x)/E_\alpha$ for all $x \geq 0$ we obtain that

$$\mathbb{E} \left[\psi_{\alpha/2} \left(\frac{|XY|}{abD_\alpha} \right) \right] \leq \mathbb{E} \left[\psi_{\alpha/2} \left(\frac{|XY|}{ab} \right) \right] / E_\alpha \leq \frac{1}{2} (\mathbb{E}[\psi_\alpha(|X|/a)] + \mathbb{E}[\psi_\alpha(|Y|/b)]) \leq 1$$

□

Lemma 8.6. Suppose random variable X has ψ_α -Orlicz norm a , then $X - \mathbb{E}[X]$ has ψ_α Orlicz norm at most $2a$.

Proof. First of all, since ψ_α is convex and increasing on $[0, \infty)$, we have that $\mathbb{E}[\psi_\alpha(|X|/a)] \geq \psi_\alpha(\mathbb{E}[|X|]/a) \geq \psi_\alpha(|\mathbb{E}[X]|/a)$. Then we have that

$$\mathbb{E} \left[\psi_\alpha \left(\frac{|X - \mathbb{E}[X]|}{2a} \right) \right] \leq \mathbb{E} \left[\psi_\alpha \left(\frac{|X|}{2a} + \frac{|\mathbb{E}[X]|}{2a} \right) \right] \leq \mathbb{E} \left[\frac{1}{2} \psi_\alpha(|X|/a) + \frac{1}{2} \psi_\alpha(|\mathbb{E}[X]|/a) \right] \leq \mathbb{E}[\psi_\alpha(|X|/a)] \leq 1$$

where we used the convexity of ψ_α and the fact that $\mathbb{E}[\psi_\alpha(|X|/a)] \geq \psi_\alpha(|\mathbb{E}[X]|/a)$ \square

The following Theorem of [PMS95] is useful to decouple the randomness of a sum of correlated random variables into a form that is easier to control.

Theorem 8.7 (Special case of Theorem 1 of [PMS95]). Let $X_1, \dots, X_n, Y_1, \dots, Y_n$ are independent random variables on a measurable space over S , where X_i and Y_i has the same distribution for $i = 1, \dots, n$. Let $f_{ij}(\cdot, \cdot)$ be a family of functions taking $S \times S$ to a Banach space $(B, \|\cdot\|)$. Then there exists absolute constant C , such that for all $n \geq 2, t > 0$,

$$\Pr \left[\left\| \sum_{i \neq j} f_{ij}(X_i, X_j) \right\| \geq t \right] \leq C \Pr \left[\left\| \sum_{i \neq j} f_{ij}(X_i, Y_j) \right\| \geq t/C \right]$$

The following lemma provides a simple way to prove the PSDness of a matrix that has large value on the diagonal and small off-diagonal values.

Lemma 8.8 (Consequence of Gershgorin Circle Theorem). Suppose a matrix Γ is of the form $\Gamma = \begin{bmatrix} A & B \\ C & D \end{bmatrix}$ where A, D are square diagonal matrices, and C is of dimension $n \times m$. Then Γ is PSD if there exists $\alpha > 0$ such that the following holds: $A_{ii} \geq \frac{1}{\alpha} \sum_{j \in [n]} |C_{ij}|, \forall i \in [p]$ and $D_{jj} \geq \alpha \sum_{i \in [m]} |C_{ij}|, \forall j \in [p]$.

Proof. Let vector $u = (\alpha \mathbf{1}_m, \alpha^{-1} \mathbf{1}_n)$ and $v = (\alpha^{-1} \mathbf{1}_m, \alpha \mathbf{1}_n)$, where $\mathbf{1}_n$ is n -dimensional all 1's vector. Then Γ can be written as $\Gamma = vv^T \odot (uu^T \odot \Gamma)$, where \odot denotes the entries-wise product of two matrices (That is, $A \odot B$ is a matrix with entry $A_{ij}B_{ij}$). Using the Gershgorin Circle Theorem and the conditions of the Lemma we obtain that $uu^T \odot \Gamma$ is PSD and therefore Γ is PSD. \square

9 Conclusions and future directions

In this paper we prove a lower bounds on the number of samples required to solve the Sparse PCA problem by degree-4 SoS algorithms. This extends the (spectral) degree-2 SoS lower bound for the problem, establishing the quadratic gap from the number of samples required by the (inefficient) information theoretic bound. It remains an interesting problem to extend our lower bounds to higher degree SoS algorithms (or even better, show that with some constant degree, one can solve the problem with fewer samples). One specific difficulty we encountered in trying to extend the lower bound to higher degree was the polynomial constraint $x_i^3 = x_i$, capturing the discreteness of the hidden sparse vector. The SoS formulation of the problem without this condition is interesting as well, and lower bound for it may be easier.

As mentioned, it is possible that the best way to prove strong SoS lower bounds for Sparse PCA is via the reduction of Berthet and Rigollet's [BR13a], namely by improving existing lower bounds for the Planted Clique problem. However, we note that this approach is limited as well, as it seems

that sparse PCA is significantly harder. Specifically, Planted Clique has a simple $O(\log n)$ -degree SoS algorithm (and thus a quasi-polynomial time) *optimal* solution, whereas for Sparse PCA we know of no better sample-optimal algorithm than one running in exponential $p^{O(k)}$ time. It is thus conceivable that one can even prove $\Omega(k)$ -degree SoS lower bounds for this problem.

More generally, we believe that statistical and machine learning problems provide a new and challenging setting for testing the power and limits of SoS algorithms. While we have fairly strong techniques for proving optimal SoS lower bounds for combinatorial optimization problems, we lack similar ones for ML problems. In particular, many other problems besides Sparse PCA seem to exhibit the apparent trade-off between the number of samples required information theoretically versus via computationally efficient techniques, offering fertile ground for attempting SoS lower bounds establishing such trade-offs.

Finally it would be nice to see more reductions between problems of statistical and ML nature, as the one by [BR13a]. Efficient reductions have proved extremely powerful in computational complexity theory and optimization, enabling the framework of complexity classes and complete problems. Creating such a framework within machine learning will hopefully expose structure on the relative difficulty of problems in this vast area, highlighting some problems as more central to attack, and enabling both new algorithms and new lower bounds.

Acknowledgments: We would like to thank Sanjeev Arora, Boaz Barak, Philippe Rigollet and David Steurer for helpful discussions throughout various stages of this work.

References

- [ABN⁺99] U. Alon, N. Barkai, D. A. Notterman, K. Gish, S. Ybarra, D. Mack, and A. J. Levine. Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. *Proceedings of the National Academy of Sciences*, 96(12):6745–6750, 1999.
- [Art27] Emil Artin. Über die zerlegung definiter funktionen in quadrate. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 5, pages 100–115. Springer, 1927.
- [AW09] Arash A. Amini and Martin J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *Ann. Statist.*, 37(5B):2877–2921, 10 2009.
- [BKS14] Boaz Barak, Jonathan A. Kelner, and David Steurer. Rounding sum-of-squares relaxations. In *STOC*, pages 31–40, 2014.
- [BKS15] Boaz Barak, Jonathan A. Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, STOC ’15, 2015.
- [BM15] Boaz Barak and Ankur Moitra. Tensor prediction, rademacher complexity and random 3-xor. *CoRR*, abs/1501.06521, 2015.
- [BR13a] Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *COLT 2013 - The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA*, pages 1046–1066, 2013.

- [BR13b] Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*, 41(4):1780–1815, 2013.
- [BS14] Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. In *Proceedings of International Congress of Mathematicians (ICM)*, 2014. To appear.
- [Che11] Xi Chen. Adaptive elastic-net sparse principal component analysis for pathway association testing. *Statistical Applications in Genetics and Molecular Biology*, 10, 2011.
- [CJ13] Venkat Chandrasekaran and Michael I. Jordan. Computational and statistical tradeoffs via convex relaxation. *Proceedings of the National Academy of Sciences*, 110(13):E1181–E1190, 2013.
- [dGJL07] Alexandre d’Aspremont, Laurent El Ghaoui, Michael I. Jordan, and Gert R. G. Lanckriet. A direct formulation for sparse pca using semidefinite programming. *SIAM Review*, 49(3):434–448, 2007.
- [DGR97] Scott Decatur, Oded Goldreich, and Dana Ron. Computational sample complexity. In *Proceedings of the Tenth Annual Conference on Computational Learning Theory, COLT ’97*, pages 130–142, New York, NY, USA, 1997. ACM.
- [DJ98] David L. Donoho and Iain M. Johnstone. Minimax estimation via wavelet shrinkage. *Ann. Statist.*, 26(3):879–921, 06 1998.
- [DLS13] Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. More data speeds up training time in learning halfspaces over sparse vectors. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 145–153, 2013.
- [DM14] Yash Deshpande and Andrea Montanari. Sparse PCA via covariance thresholding. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 334–342, 2014.
- [DM15] Y. Deshpande and A. Montanari. Improved Sum-of-Squares Lower Bounds for Hidden Clique and Hidden Submatrix Problems. *ArXiv e-prints*, February 2015.
- [Don95] D. L. Donoho. De-noising by soft-thresholding. *IEEE Trans. Inf. Theor.*, 41(3):613–627, May 1995.
- [GMZ14] C. Gao, Z. Ma, and H. H. Zhou. Sparse CCA: Adaptive Estimation and Computational Barriers. *ArXiv e-prints*, September 2014.
- [Gri01a] D. Grigoriev. Complexity of positivstellensatz proofs for the knapsack. *computational complexity*, 10(2):139–154, 2001.
- [Gri01b] Dima Grigoriev. Linear lower bound on degrees of positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1):613–622, 2001.

- [HKP15] Samuel B. Hopkins, Pravesh K. Kothari, and Aaron Potechin. Sos and planted clique: Tight analysis of MPW moments at all degrees and an optimal lower bound at degree four. *CoRR*, abs/1507.05230, 2015.
- [JL09] Iain M. Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):pp. 682–703, 2009.
- [JOB10] Rodolphe Jenatton, Guillaume Obozinski, and Francis R. Bach. Structured sparse principal component analysis. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, pages 366–373, 2010.
- [Joh01] Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.*, 29(2):295–327, 04 2001.
- [Joh02] IM Johnstone. Function estimation and gaussian sequence models. *Unpublished manuscript*, 2002.
- [KNV15] Robert Krauthgamer, Boaz Nadler, and Dan Vilenchik. Do semidefinite relaxations solve sparse pca up to the information limit? *The Annals of Statistics*, 43(3):1300–1322, 2015.
- [Kri64] Jean-Louis Krivine. Anneaux préordonnés. *Journal d’analyse mathématique*, 12(1):307–326, 1964.
- [Las01] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- [Las15] Jean Bernard Lasserre. *An introduction to polynomial and semi-algebraic optimization*. Cambridge Texts in Applied Mathematics. Cambridge: Cambridge University Press. , 2015.
- [Lau09] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In Mihai Putinar and Seth Sullivan, editors, *Emerging Applications of Algebraic Geometry*, volume 149 of *The IMA Volumes in Mathematics and its Applications*, pages 157–270. Springer New York, 2009.
- [LS91] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 01 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.
- [LT13] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer Science & Business Media, 2013.
- [Ma13] Zongming Ma. Sparse principal component analysis and iterative thresholding. *Ann. Statist.*, 41(2):772–801, 04 2013.
- [MPW15] Raghu Meka, Aaron Potechin, and Avi Wigderson. Sum-of-squares lower bounds for planted clique. *CoRR*, abs/1503.06447, 2015.
- [MR14] Tengyu Ma and Philippe Rigollet. personal communication, 2014.

- [Nes00] Yurii Nesterov. Squared functional systems and optimization problems. In Hans Frenk, Kees Roos, Tams Terlaky, and Shuzhong Zhang, editors, *High Performance Optimization*, volume 33 of *Applied Optimization*, pages 405–440. Springer US, 2000.
- [Par00] Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, 2000.
- [PJ12] Debashis Paul and Iain M Johnstone. Augmented sparse principal component analysis for high dimensional data. *arXiv preprint arXiv:1202.1242*, 2012.
- [PMS95] Victor H. de la Pena and S. J. Montgomery-Smith. Decoupling inequalities for the tail probabilities of multivariate u-statistics. *The Annals of Probability*, 23(2):pp. 806–816, 1995.
- [Put93] Mihai Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3):969–984, 1993.
- [RS15] Prasad Raghavendra and Tselil Schramm. Tight lower bounds for planted clique in the degree-4 SOS program. *CoRR*, abs/1507.05136, 2015.
- [SA90] Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990.
- [Sch91] Konrad Schmüdgen. Thek-moment problem for compact semi-algebraic sets. *Mathematische Annalen*, 289(1):203–206, 1991.
- [Sch08] Grant Schoenebeck. Linear level lasserre lower bounds for certain k-csps. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS '08*, pages 593–602, Washington, DC, USA, 2008. IEEE Computer Society.
- [Ser00] Rocco A. Servedio. Computational sample complexity and attribute-efficient learning. *Journal of Computer and System Sciences*, 60(1):161 – 178, 2000.
- [Sho87] N.Z. Shor. An approach to obtaining global extremums in polynomial mathematical programming problems. *Cybernetics*, 23(5):695–700, 1987.
- [Ste74] Gilbert Stengle. A nullstellensatz and a positivstellensatz in semialgebraic geometry. *Mathematische Annalen*, 207(2):87–97, 1974.
- [vdVW00] Aad van der Vaart and Jon Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics (Springer Series in Statistics)*. Springer, corrected edition, November 2000.
- [VL12] Vincent Q. Vu and Jing Lei. Minimax rates of estimation for sparse PCA in high dimensions. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2012, La Palma, Canary Islands, April 21-23, 2012*, pages 1278–1286, 2012.

- [VL13] Vincent Q. Vu and Jing Lei. Minimax sparse principal subspace estimation in high dimensions. *Ann. Statist.*, 41(6):2905–2947, 12 2013.
- [WGL15] Z. Wang, Q. Gu, and H. Liu. Statistical Limits of Convex Relaxations. *ArXiv e-prints*, March 2015.
- [WLL14] Zhaoran Wang, Huanran Lu, and Han Liu. Tighten after relax: Minimax-optimal sparse PCA in polynomial time. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3383–3391, 2014.
- [YZ13] Xiao-Tong Yuan and Tong Zhang. Truncated power method for sparse eigenvalue problems. *J. Mach. Learn. Res.*, 14(1):899–925, April 2013.