

Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors

Samuel B. Hopkins* Tselil Schramm[†] Jonathan Shi[‡] David Steurer[§]

May 14, 2016

Abstract

We consider two problems that arise in machine learning applications: the problem of recovering a planted sparse vector in a random linear subspace and the problem of decomposing a random low-rank overcomplete 3-tensor. For both problems, the best known guarantees are based on the sum-of-squares method. We develop new algorithms inspired by analyses of the sum-of-squares method. Our algorithms achieve the same or similar guarantees as sum-of-squares for these problems but the running time is significantly faster.

For the planted sparse vector problem, we give an algorithm with running time nearly linear in the input size that approximately recovers a planted sparse vector with up to constant relative sparsity in a random subspace of \mathbb{R}^n of dimension up to $\tilde{\Omega}(\sqrt{n})$. These recovery guarantees match the best known ones of Barak, Kelner, and Steurer (STOC 2014) up to logarithmic factors.

For tensor decomposition, we give an algorithm with running time close to linear in the input size (with exponent ≈ 1.125) that approximately recovers a component of a random 3-tensor over \mathbb{R}^n of rank up to $\tilde{\Omega}(n^{4/3})$. The best previous algorithm for this problem due to Ge and Ma (RANDOM 2015) works up to rank $\tilde{\Omega}(n^{3/2})$ but requires quasipolynomial time.

*Cornell University, samhop@cs.cornell.edu. Supported by an NSF Graduate Research Fellowship (NSF award no. 1144153) and by David Steurer's NSF CAREER award.

[†]UC Berkeley, tschramm@cs.berkeley.edu. Supported by an NSF Graduate Research Fellowship (NSF award no. 1106400).

[‡]Cornell University, jshi@cs.cornell.edu. Supported by David Steurer's NSF CAREER award.

[§]Cornell University, dsteurer@cs.cornell.edu. Supported by a Microsoft Research Fellowship, a Alfred P. Sloan Fellowship, an NSF CAREER award, and the Simons Collaboration for Algorithms and Geometry.

Contents

1	Introduction	1
1.1	Planted sparse vector in random linear subspace	2
1.2	Overcomplete tensor decomposition	3
1.3	Tensor principal component analysis	4
1.4	Related work	5
2	Techniques	6
2.1	Planted sparse vector in random linear subspace	8
2.2	Overcomplete tensor decomposition	10
2.3	Tensor principal component analysis	12
3	Preliminaries	13
4	Planted sparse vector in random linear subspace	14
4.1	Algorithm succeeds on good basis	16
4.2	Closeness of input basis and good basis	19
5	Overcomplete tensor decomposition	20
5.1	Proof of Theorem 5.3	22
5.2	Spectral gap for diagonal terms: proof of Proposition 5.4	24
5.3	Bound for cross terms: proof of Proposition 5.5	28
5.4	Full algorithm and proof of Theorem 5.2	32
6	Tensor principal component analysis	38
6.1	Spiked tensor model	38
6.2	Linear-time algorithm	38
	References	40
A	Additional preliminaries	43
A.1	Linear algebra	43
A.2	Concentration tools	46
B	Concentration bounds for planted sparse vector in random linear subspace	52
C	Concentration bounds for overcomplete tensor decomposition	55
D	Concentration bounds for tensor principal component analysis	62

1 Introduction

The sum-of-squares (SoS) method (also known as the Lasserre hierarchy) [Sho87, Par00, Nes00, Las01] is a powerful, semidefinite-programming based meta-algorithm that applies to a wide-range of optimization problems. The method has been studied extensively for moderate-size polynomial optimization problems that arise for example in control theory and in the context of approximation algorithms for combinatorial optimization problems, especially constraint satisfaction and graph partitioning (see e.g. the survey [BS14]). For the latter, the SoS method captures and generalizes the best known approximation algorithms based on linear programming (LP), semidefinite programming (SDP), or spectral methods, and it is in many cases the most promising approach to obtain algorithms with better guarantees—especially in the context of Khot’s Unique Games Conjecture [BBH⁺12].

A sequence of recent works applies the sum-of-squares method to basic problems that arise in unsupervised machine learning: in particular, recovering sparse vectors in linear subspaces and decomposing tensors in a robust way [BKS14, BKS15, HSS15, BM15, GM15]. For a wide range of parameters of these problems, SoS achieves significantly stronger guarantees than other methods, in polynomial or quasi-polynomial time.

Like other LP and SDP hierarchies, the sum-of-squares method comes with a degree parameter $d \in \mathbb{N}$ that allows for trading off running time and solution quality. This trade-off is appealing because for applications the additional utility of better solutions could vastly outweigh additional computational costs. Unfortunately, the computational cost grows rather steeply in terms of the parameter d : the running time is $n^{O(d)}$ where n is the number of variables (usually comparable to the instance size). Further, even when the SDP has size polynomial in the input (when $d = O(1)$), solving the underlying semidefinite programs is prohibitively slow for large instances.

In this work, we introduce spectral algorithms for planted sparse vector, tensor decomposition, and tensor principal components analysis (PCA) that exploit the same high-degree information as the corresponding sum-of-squares algorithms without relying on semidefinite programming, and achieve the same (or close to the same) guarantees. The resulting algorithms are quite simple (a couple of lines of MATLAB code) and have considerably faster running times—quasi-linear or close to linear in the input size.

A surprising implication of our work is that for some problems, spectral algorithms can exploit information from larger values of the parameter d without spending time $n^{O(d)}$. For example, our algorithm for the planted sparse vector problem runs in nearly-linear time in the input size, even though it uses properties that the sum-of-squares method can only use for degree parameter $d \geq 4$. (In particular, the guarantees that the algorithm achieves are strictly stronger than the guarantees that SoS achieves for values of $d < 4$.)

The initial successes of SoS in the machine learning setting gave hope that techniques developed in the theory of approximation algorithms, specifically the techniques of hierarchies of convex relaxations and rounding convex relaxations, could broadly impact the practice of machine learning. This hope was dampened by the fact that in general, algorithms that rely on solving large semidefinite programs are too slow to be practical for the large-scale problems that arise in machine learning. Our work brings this hope back into focus by demonstrating for the first time that with some care SoS algorithms can be made practical for large-scale problems.

In the following subsections we describe each of the problems that we consider, the prior best-known guarantee via the SoS hierarchy, and our results.

1.1 Planted sparse vector in random linear subspace

The problem of finding a sparse vector planted in a random linear subspace was introduced by Spielman, Wang, and Wright as a way of learning sparse dictionaries [SWW12]. Subsequent works have found further applications and begun studying the problem in its own right [DH14, BKS14, QSW14]. In this problem, we are given a basis for a d -dimensional linear subspace of \mathbb{R}^n that is random except for one planted sparse direction, and the goal is to recover this sparse direction. The computational challenge is to solve this problem even when the planted vector is only mildly sparse (a constant fraction of non-zero coordinates) and the subspace dimension is large compared to the ambient dimension ($d \geq n^{\Omega(1)}$).

Several kinds of algorithms have been proposed for this problem based on linear programming (LP), basic semidefinite programming (SDP), sum-of-squares, and non-convex gradient descent (alternating directions method).

An inherent limitation of simpler convex methods (LP and basic SDP) [SWW12, dGJL04] is that they require the relative sparsity of the planted vector to be polynomial in the subspace dimension (less than n/\sqrt{d} non-zero coordinates).

Sum-of-squares and non-convex methods do not share this limitation. They can recover planted vectors with constant relative sparsity even if the subspace has polynomial dimension (up to dimension $O(n^{1/2})$ for sum-of-squares [BKS14] and up to $O(n^{1/4})$ for non-convex methods [QSW14]).

We state the problem formally:

Problem 1.1 (Planted sparse vector problem with ambient dimension $n \in \mathbb{N}$, subspace dimension $d \leq n$, sparsity $\varepsilon > 0$, and accuracy $\eta > 0$). Given an arbitrary orthogonal basis of a subspace spanned by vectors $v_0, v_1, \dots, v_{d-1} \in \mathbb{R}^n$, where v_0 is a vector with at most εn non-zero entries and v_1, \dots, v_{d-1} are vectors sampled independently at random from the standard Gaussian distribution on \mathbb{R}^n , output a unit vector $v \in \mathbb{R}^n$ that has correlation $\langle v, v_0 \rangle^2 \geq 1 - \eta$ with the sparse vector v_0 .

Our Results. Our algorithm runs in nearly linear time in the input size, and matches the best-known guarantees up to a polylogarithmic factor in the subspace dimension [BKS14].

Theorem 1.2 (Planted sparse vector in nearly-linear time). *There exists an algorithm that, for every sparsity $\varepsilon > 0$, ambient dimension n , and subspace dimension d with $d \leq \sqrt{n}/(\log n)^{O(1)}$, solves the planted sparse vector problem with high probability for some accuracy $\eta \leq O(\varepsilon^{1/4}) + o_{n \rightarrow \infty}(1)$. The running time of the algorithm is $\tilde{O}(nd)$.*

We give a technical overview of the proof in [Section 2](#), and a full proof in [Section 4](#).

Previous work also showed how to recover the planted sparse vector exactly. The task of going from an approximate solution to an exact one is a special case of standard compressed sensing (see e.g. [BKS14]).

Table 1: Comparison of algorithms for the planted sparse vector problem with ambient dimension n , subspace dimension d , and relative sparsity ε .

Reference	Technique	Runtime	Largest d	Largest ε
Demanet, Hand [DH14]	linear programming	poly	any	$\Omega(1/\sqrt{d})$
Barak, Kelner, Steurer [BKS14]	SoS, general SDP	poly	$\Omega(\sqrt{n})$	$\Omega(1)$
Qu, Sun, Wright [QSW14]	alternating minimization	$\tilde{O}(n^2 d^5)$	$\Omega(n^{1/4})$	$\Omega(1)$
this work	SoS, partial traces	$\tilde{O}(nd)$	$\tilde{\Omega}(\sqrt{n})$	$\Omega(1)$

1.2 Overcomplete tensor decomposition

Tensors naturally represent multilinear relationships in data. Algorithms for tensor decompositions have long been studied as a tool for data analysis across a wide-range of disciplines (see the early work of Harshman [Har70] and the survey [KB09]). While the problem is NP-hard in the worst-case [Hås90, HL13], algorithms for special cases of tensor decomposition have recently led to new provable algorithmic results for several unsupervised learning problems [AGH⁺14, BCMV14, GVX14, AGHK14] including independent component analysis, learning mixtures of Gaussians [GHK15], Latent Dirichlet topic modeling [AFH⁺15] and dictionary learning [BKS15]. Some previous learning algorithms can also be reinterpreted in terms of tensor decomposition [Cha96, MR06, NR09].

A key algorithmic challenge for tensor decompositions is *overcompleteness*, when the number of components is larger than their dimension (i.e., the components are linearly dependent). Most algorithms that work in this regime require tensors of order 4 or higher [LCC07, BCMV14]. For example, the FOABI algorithm of [LCC07] can recover up to $\Omega(d^2)$ components given an order-4 tensor in dimension d under mild algebraic independence assumptions for the components—satisfied with high probability by random components. For overcomplete 3-tensors, which arise in many applications of tensor decompositions, such a result remains elusive.

Researchers have therefore turned to investigate average-case versions of the problem, when the components of the overcomplete 3-tensor are random: Given a 3-tensor $T \in \mathbb{R}^{d^3}$ of the form

$$T = \sum_{i=1}^n a_i \otimes a_i \otimes a_i,$$

where a_1, \dots, a_n are random unit or Gaussian vectors, the goal is to approximately recover the components a_1, \dots, a_n .

Algorithms based on *tensor power iteration*—a gradient-descent approach for tensor decomposition—solve this problem in polynomial time when $n \leq C \cdot d$ for any constant $C \geq 1$ (the running time is exponential in C) [AGJ15]. Tensor power iteration also admits local convergence analyses for up to $n \leq \tilde{\Omega}(d^{1.5})$ components [AGJ15, AGJ14]. Unfortunately, these analyses do not give polynomial-time algorithms because it is not known how to efficiently obtain the kind of initializations assumed by the analyses.

Recently, Ge and Ma [GM15] were able to show that a tensor-decomposition algorithm [BKS15] based on sum-of-squares solves the above problem for $n \leq \tilde{\Omega}(d^{1.5})$ in quasi-polynomial time $n^{O(\log n)}$. The key ingredient of their elegant analysis is a subtle spectral concentration bound for a particular degree-4 matrix-valued polynomial associated with the decomposition problem of random

Table 2: Comparison of decomposition algorithms for overcomplete 3-tensors with n components in dimension d .

Reference	Technique	Runtime	Largest n	Components
Anandkumar et al. [AGJ15] ^a	tensor power iteration	poly	$C \cdot d$	incoherent
Ge, Ma [GM15]	SoS, general SDP	$n^{O(\log n)}$	$\tilde{\Omega}(d^{3/2})$	$\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$
this work ^b	SoS, partial traces	$\tilde{O}(nd^{1+\omega})$	$\tilde{\Omega}(d^{4/3})$	$\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$

^a The analysis shows that for every constant $C \geq 1$, the running time is polynomial for $n \leq C \cdot d$ components, assuming that the components also satisfy other random-like properties besides incoherence.

^b Here, $\omega \leq 2.3729$ is the constant so that $d \times d$ matrices can be multiplied in $O(d^\omega)$ arithmetic operations.

overcomplete 3-tensors.

We state the problem formally:

Problem 1.3 (Random tensor decomposition with dimension d , rank n , and accuracy η). Let $a_1, \dots, a_n \in \mathbb{R}^d$ be independently sampled vectors from the Gaussian distribution $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$, and let $\mathbf{T} \in (\mathbb{R}^d)^{\otimes 3}$ be the 3-tensor $\mathbf{T} = \sum_{i=1}^n a_i^{\otimes 3}$.

Single component: Given \mathbf{T} sampled as above, find a unit vector b that has correlation $\max_i \langle a_i, b \rangle \geq 1 - \eta$ with one of the vectors a_i .

All components: Given \mathbf{T} sampled as above, find a set of unit vectors $\{b_1, \dots, b_n\}$ such that $\langle a_i, b_i \rangle \geq 1 - \eta$ for every $i \in [n]$.

Our Results. We give the first polynomial-time algorithm for decomposing random overcomplete 3-tensors with up to $\omega(d)$ components. Our algorithm works as long as the number of components satisfies $n \leq \tilde{\Omega}(d^{4/3})$, which comes close to the bound $\tilde{\Omega}(d^{1.5})$ achieved by the aforementioned quasi-polynomial algorithm of Ge and Ma. For the single-component version of the problem, our algorithm runs in time close to linear in the input size.

Theorem 1.4 (Fast random tensor decomposition). *There exist randomized algorithms that, for every dimension d and rank n with $d \leq n \leq d^{4/3}/(\log n)^{O(1)}$, solve the random tensor decomposition problem with probability $1 - o(1)$ for some accuracy $\eta \leq \tilde{O}(n^3/d^4)^{1/2}$. The running time for the single-component version of the problem is $\tilde{O}(d^{1+\omega})$, where $d^\omega \leq d^{2.3279}$ is the time to multiply two d -by- d matrices. The running time for the all-components version of the problem is $\tilde{O}(n \cdot d^{1+\omega})$.*

We give a technical overview of the proof in [Section 2](#), and a full proof in [Section 5](#).

We remark that the above algorithm only requires access to the input tensor with some fixed inverse polynomial accuracy because each of its four steps amplifies errors by at most a polynomial factor (see [Algorithm 5.17](#)). In this sense, the algorithm is robust.

1.3 Tensor principal component analysis

The problem of tensor principal component analysis is similar to the tensor decomposition problem. However, here the focus is not on the number of components in the tensor, but about recovery in the presence of a large amount of random noise. We are given as input a tensor $\tau \cdot v^{\otimes 3} + \mathbf{A}$, where $v \in \mathbb{R}^n$ is a unit vector and the entries of \mathbf{A} are chosen iid from $\mathcal{N}(0, 1)$. This *spiked tensor* model was introduced by Montanari and Richard [RM14], who also obtained the first algorithms to

Table 3: Comparison of algorithms for principal component analysis of 3-tensors in dimension d and with signal-to-noise ratio τ .

Reference	Technique	Runtime	Smallest τ
Montanari, Richard [RM14]	spectral	$\tilde{O}(d^3)$	n
Hopkins, Shi, Steurer [HSS15]	SoS, spectral	$\tilde{O}(d^3)$	$O(n^{3/4})$
this work	SoS, partial traces	$O(d^3)$	$\tilde{O}(n^{3/4})$

solve the model with provable statistical guarantees. The spiked tensor model was subsequently addressed by a subset of the present authors [HSS15], who applied the SoS approach to improve the signal-to-noise ratio required for recovery from odd-order tensors.

We state the problem formally:

Problem 1.5 (Tensor principal components analysis with signal-to-noise ratio τ and accuracy η). Let $\mathbf{T} \in (\mathbb{R}^d)^{\otimes 3}$ be a tensor so that $\mathbf{T} = \tau \cdot v^{\otimes 3} + \mathbf{A}$, where \mathbf{A} is a tensor with independent standard gaussian entries and $v \in \mathbb{R}^d$ is a unit vector. Given \mathbf{T} , recover a unit vector $v' \in \mathbb{R}^d$ such that $\langle v', v \rangle \geq 1 - \eta$.

Our results. For this problem, our improvements over the previous results are more modest—we achieve signal-to-noise guarantees matching [HSS15], but with an algorithm that runs in linear time rather than near-linear time (time $O(d^3)$ rather than $O(d^3 \text{ polylog } d)$, for an input of size d^3).

Theorem 1.6 (Tensor principal component analysis in linear time). *There is an algorithm which solves the tensor principal component analysis problem with accuracy $\eta > 0$ whenever the signal-to-noise ratio satisfies $\tau \geq O(n^{3/4} \cdot \eta^{-1} \cdot \log^{1/2} n)$. Furthermore, the algorithm runs in time $O(d^3)$.*

Though for tensor PCA our improvement over previous work is modest, we include the results here as this problem is a pedagogically poignant illustration of our techniques. We give a technical overview of the proof in [Section 2](#), and a full proof in [Section 6](#).

1.4 Related work

Foremost, this work builds upon the SoS algorithms of [BKS14, BKS15, GM15, HSS15]. In each of these previous works, a machine learning decision problem is solved using an SDP relaxation for SoS. In these works, the SDP value is large in the YES case and small in the NO case, and the SDP value can be bounded using the spectrum of a specific matrix. This was implicit in [BKS14, BKS15], and in [HSS15] it was used to obtain a fast algorithm as well. In our work, we design spectral algorithms which use smaller matrices, inspired by the SoS certificates in previous works, to solve these machine-learning problems much faster, with almost matching guarantees.

A key idea in our work is that given a large matrix with information encoded in the matrix’s spectral gap, one can often efficiently “compress” the matrix to a much smaller one without losing that information. This is particularly true for problems with planted solutions. In this way, we are able to improve running time by replacing an $n^{O(d)}$ -sized SDP with an eigenvector computation for an $n^k \times n^k$ matrix, for some $k < d$.

The idea of speeding up LP and SDP hierarchies for specific problems has been investigated in a series of previous works [dIVK07, BRS11, GS12], which shows that with respect to local analyses of the sum-of-squares algorithm it is sometimes possible to improve the running time from $n^{O(d)}$ to $2^{O(d)} \cdot n^{O(1)}$. However, the scopes and strategies of these works are completely different from ours. First, the notion of local analysis from these works does not apply to the problems considered here. Second, these works employ the ellipsoid method with a separation oracle inspired by rounding algorithms, whereas we reduce the problem to an ordinary eigenvector computation.

It would also be interesting to see if our methods can be used to speed up some of the other recent successful applications of SoS to machine-learning type problems, such as [BM15], or the application of [BKS14] to tensor decomposition with components that are well-separated (rather than random). Finally, we would be remiss not to mention that SoS lower bounds exist for several of these problems, specifically for tensor principal components analysis, tensor prediction, and sparse PCA [HSS15, BM15, MW15]. The lower bounds in the SoS framework are a good indication that we cannot expect spectral algorithms achieving better guarantees.

2 Techniques

Sum-of-squares method (for polynomial optimization over the sphere). The problems we consider are connected to optimization problems of the following form: Given a homogeneous n -variate real polynomial f of constant degree, find a unit vector $x \in \mathbb{R}^n$ so as to maximize $f(x)$. The sum-of-squares method allows us to efficiently compute upper bounds on the maximum value of such a polynomial f over the unit sphere.

For the case that $k = \deg(f)$ is even, the most basic upper bound of this kind is the largest eigenvalue of a *matrix representation* of f . A *matrix representation* of a polynomial f is a symmetric matrix M with rows and columns indexed by monomials of degree $k/2$ so that $f(x)$ can be written as the quadratic form $f(x) = \langle x^{\otimes k/2}, Mx^{\otimes k/2} \rangle$, where $x^{\otimes k/2}$ is the $k/2$ -fold tensor power of x . The largest eigenvalue of a matrix representation M is an upper bound on the value of $f(x)$ over all unit vectors $x \in \mathbb{R}^n$ because

$$f(x) = \langle x^{\otimes k/2}, Mx^{\otimes k/2} \rangle \leq \lambda_{\max}(M) \cdot \|x^{\otimes k/2}\|_2^2 = \lambda_{\max}(M).$$

The sum-of-squares method improves on this basic spectral bound systematically by associating a large family of polynomials (potentially of degree higher than $\deg(f)$) with the input polynomial f and computing the best possible spectral bound within this family of polynomials. Concretely, the sum-of-squares method with degree parameter d applied to a polynomial f with $\deg(f) \leq d$ considers the affine subspace of polynomials $\{f + (1 - \|x\|_2^2) \cdot g \mid \deg(g) \leq d - 2\} \subseteq \mathbb{R}[x]$ and minimizes $\lambda_{\max}(M)$ among all matrix representations ¹ M of polynomials in this space.² The

¹ Earlier we defined matrix representations only for homogeneous polynomials of even degree. In general, a matrix representation of a polynomial g is a symmetric matrix M with rows and columns indexed by monomials of degree at most $\ell = \deg(g)/2$ such that $g(x) = \langle x^{\otimes \leq \ell}, Mx^{\otimes \leq \ell} \rangle$ (as a polynomial identity), where $x^{\otimes \leq \ell} = (x^{\otimes 0}, x^{\otimes 1}, \dots, x^{\otimes \ell})/\sqrt{\ell + 1}$ is the vector of all monomials of degree at most ℓ . Note that $\|x^{\otimes \leq \ell}\| = 1$ for all x with $\|x\| = 1$.

² The name of the method stems from the fact that this last step is equivalent to finding the minimum number λ such that the space contains a polynomial of the form $\lambda - (g_1^2 + \dots + g_t^2)$, where g_1, \dots, g_t are polynomials of degree at most $d/2$.

problem of searching through this affine linear space of polynomials and their matrix representations and finding the one of smallest maximum eigenvalue can be solved using semidefinite programming.

Our approach for faster algorithms based on SoS algorithms is to construct specific matrices (polynomials) in this affine linear space, then compute their top eigenvectors. By designing our matrices carefully, we ensure that our algorithms have access to the same higher degree information that the sum-of-squares algorithm can access, and this information affords an advantage over the basic spectral methods for these problems. At the same time, our algorithms avoid searching for the best polynomial and matrix representation, which gives us faster running times since we avoid semidefinite programming. This approach is well suited to average-case problems where we avoid the problem of adversarial choice of input; in particular it is applicable to machine learning problems where noise and inputs are assumed to be random.

Compressing matrices with partial traces. A serious limitation of the above approach is that the representation of a degree- d , n -variate polynomial requires size roughly n^d . Hence, even avoiding the use of semidefinite programming, improving upon running time $O(n^d)$ requires additional ideas.

In each of the problems that we consider, we have a large matrix (suggested by a SoS algorithm) with a “signal” planted in some amount of “noise”. We show that in some situations, this large matrix can be compressed significantly without loss in the signal by applying *partial trace* operations. In these situations, the partial trace yields a smaller matrix with the same signal-to-noise ratio as the large matrix suggested by the SoS algorithm, even in situations when lower degree sum-of-squares approaches are known to fail (as for the planted sparse vector and tensor PCA problems).³

The partial trace $\text{Tr}_{\mathbb{R}^d} : \mathbb{R}^{d^2 \times d^2} \rightarrow \mathbb{R}^{d \times d}$ is the linear operator that satisfies $\text{Tr}_{\mathbb{R}^d} A \otimes B = (\text{Tr} A) \cdot B$ for all $A, B \in \mathbb{R}^{d \times d}$. To see how the partial trace can be used to compress large matrices to smaller ones with little loss, consider the following problem: Given a matrix $M \in \mathbb{R}^{d^2 \times d^2}$ of the form $M = \tau \cdot (v \otimes v)(v \otimes v)^\top + A \otimes B$ for some unit vector $v \in \mathbb{R}^d$ and matrices $A, B \in \mathbb{R}^{d \times d}$, we wish to recover the vector v . (This is a simplified version of the planted problems that we consider in this paper, where $\tau \cdot (v \otimes v)(v \otimes v)^\top$ is the signal and $A \otimes B$ plays the role of noise.)

It is straightforward to see that the matrix $A \otimes B$ has spectral norm $\|A \otimes B\| = \|A\| \cdot \|B\|$, and so when $\tau \gg \|A\| \|B\|$, the matrix M has a noticeable spectral gap, and the top eigenvector of M will be close to $v \otimes v$. If $|\text{Tr} A| \approx \|A\|$, the matrix $\text{Tr}_{\mathbb{R}^d} M = \tau \cdot v v^\top + \text{Tr}(A) \cdot B$ has a matching spectral gap, and we can still recover v , but now we only need to compute the top eigenvector of a $d \times d$ (as opposed to $d^2 \times d^2$) matrix.⁴

If A is a Wigner matrix (e.g. a symmetric matrix with iid ± 1 entries), then both $\text{Tr}(A)$, $\|A\| \approx \sqrt{n}$, and the above condition is indeed met. In our average case/machine learning settings the “noise” component is not as simple as $A \otimes B$ with A a Wigner matrix. Nonetheless, we are able to ensure that the noise displays a similar behavior under partial trace operations. In some cases, this requires additional algorithmic steps, such as random projection in the case of tensor decomposition, or

³ For both problems we use matrices with dimensions corresponding to degree-2 SoS programs. An argument of Spielman et al. ([SWW12], Theorem 9) shows that degree-2 sum-of-squares can only find sparse vectors with sparsity $k \leq \tilde{O}(\sqrt{n})$, whereas we achieve sparsity as large as $k = \Theta(n)$. For tensor PCA, the degree-2 SoS program cannot even express the objective function.

⁴In some of our applications, the matrix M is only represented implicitly and has size super-linear in the size of the input, but nevertheless we can compute the top eigenvector of the partial trace $\text{Tr}_{\mathbb{R}^d} M$ in nearly-linear time.

centering the matrix eigenvalue distribution in the case of the planted sparse vector.

It is an interesting question if there are general theorems describing the behavior of spectral norms under partial trace operations. In the current work, we compute the partial traces explicitly and estimate their norms directly. Indeed, our analyses boil down to concentration bounds for special matrix polynomials. A general theory for the concentration of matrix polynomials is a notorious open problem (see [MW13]).

Partial trace operations have previously been applied for rounding SoS relaxations. Specifically, the operation of *reweighing* and *conditioning*, used in rounding algorithms for sum-of-squares such as [BRS11, RT12, BKS14, BKS15, LR15], corresponds to applying a partial trace operation to the moments matrix returned by the sum-of-squares relaxation.

We now give a technical overview of our algorithmic approach for each problem, and some broad strokes of the analysis for each case. Our most substantial improvements in runtime are for the planted sparse vector and overcomplete tensor decomposition problems (Section 2.1 and Section 2.2 respectively). Our algorithm for tensor PCA is the simplest application of our techniques, and it may be instructive to skip ahead and read about tensor PCA first (Section 2.3).

2.1 Planted sparse vector in random linear subspace

Recall that in this problem we are given a linear subspace U (represented by some basis) that is spanned by a k -sparse unit vector $v_0 \in \mathbb{R}^d$ and random unit vectors $v_1, \dots, v_{d-1} \in \mathbb{R}^d$. The goal is to recover the vector v_0 approximately.

Background and SoS analysis. Let $A \in \mathbb{R}^{n \times d}$ be a matrix whose columns form an orthonormal basis for U . Our starting point is the polynomial $f(x) = \|Ax\|_4^4 = \sum_{i=1}^n (Ax)_i^4$. Previous work showed that for $d \ll \sqrt{n}$ the maximizer of this polynomial over the sphere corresponds to a vector close to v_0 and that degree-4 sum-of-squares is able to capture this fact [BBH⁺12, BKS14]. Indeed, typical random vectors v in \mathbb{R}^n satisfy $\|v\|_4^4 \approx 1/n$ whereas our planted vector satisfies $\|v_0\|_4^4 \geq 1/k \gg 1/n$, and this degree-4 information is leveraged by the SoS algorithms.

The polynomial f has a convenient matrix representation $M = \sum_{i=1}^n (a_i a_i^T)^{\otimes 2}$, where a_1, \dots, a_n are the rows of the generator matrix A . It turns out that the eigenvalues of this matrix indeed give information about the planted sparse vector v_0 . In particular, the vector $x_0 \in \mathbb{R}^d$ with $Ax_0 = v_0$ witnesses that M has an eigenvalue of at least $1/k$ because M 's quadratic form with the vector $x_0^{\otimes 2}$ satisfies $\langle x_0^{\otimes 2}, Mx_0^{\otimes 2} \rangle = \|v_0\|_4^4 \geq 1/k$. If we let M' be the corresponding matrix for the subspace U without the planted sparse vector, M' turns out to have only eigenvalues of at most $O(1/n)$ up to a single spurious eigenvalue with eigenvector far from any vector of the form $x \otimes x$ [BBH⁺12].

It follows that in order to distinguish between a random subspace with a planted sparse vector (YES case) and a completely random subspace (NO case), it is enough to compute the second-largest eigenvalue of a d^2 -by- d^2 matrix (representing the 4-norm polynomial over the subspace as above). This decision version of the problem, while strictly speaking easier than the search version above, is at the heart of the matter: one can show that the large eigenvalue for the YES case corresponds to an eigenvector which encodes the coefficients of the sparse planted vector in the basis.

Improvements. The best running time we can hope for with this basic approach is $O(d^4)$ (the size of the matrix). Since we are interested in $d \leq O(\sqrt{n})$, the resulting running time $O(nd^2)$ would be subquadratic but still super-linear in the input size $n \cdot d$ (for representing a d -dimensional subspace of \mathbb{R}^n). To speed things up, we use the partial trace approach outlined above. We will begin by applying the partial trace approach naively, obtaining reasonable bounds, and then show that a small modification to the matrix before the partial trace operation allows us to achieve even smaller signal-to-noise ratios.

In the planted case, we may approximate $M \approx \frac{1}{k}(x_0 x_0^\top)^{\otimes 2} + Z$, where x_0 is the vector of coefficients of v_0 in the basis representation given by A (so that $Ax_0 = v_0$), and Z is the noise matrix. Since $\|x_0\| = 1$, the partial trace operation preserves the projector $(x_0 x_0^\top)^{\otimes 2}$ in the sense that $\text{Tr}_{\mathbb{R}^d}(x_0 x_0^\top)^{\otimes 2} = x_0 x_0^\top$. Hence, with our heuristic approximation for M above, we could show that the top eigenvector of $\text{Tr}_{\mathbb{R}^d} M$ is close to x_0 by showing that the spectral norm bound $\|\text{Tr}_{\mathbb{R}^d} Z\| \leq o(1/k)$.

The partial trace of our matrix $M = \sum_{i=1}^n (a_i a_i^\top) \otimes (a_i a_i^\top)$ is easy to compute directly,

$$N = \text{Tr}_{\mathbb{R}^d} M = \sum_{i=1}^n \|a_i\|_2^2 \cdot a_i a_i^\top.$$

In the YES case (random subspace with planted sparse vector), a direct computation shows that

$$\lambda_{\text{YES}} \geq \langle x_0, N x_0 \rangle \approx \frac{d}{n} \cdot \left(1 + \frac{n}{d} \|v_0\|_4^4\right) \geq \frac{d}{n} \left(1 + \frac{n}{dk}\right).$$

Hence, a natural approach to distinguish between the YES case and NO case (completely random subspace) is to upper bound the spectral norm of N in the NO case.

In order to simplify the bound on the spectral norm of N in the NO case, suppose that the columns of A are iid samples from the Gaussian distribution $\mathcal{N}(0, \frac{1}{d} \text{Id})$ (rather than an orthogonal basis for the random subspace)—[Lemma 4.6](#) establishes that this simplification is legitimate. In this simplified setup, the matrix N in the NO case is the sum of n iid matrices $\{\|a_i\|^2 \cdot a_i a_i^\top\}$, and we can upper bound its spectral norm λ_{NO} by $d/n \cdot (1 + O(\sqrt{d/n}))$ using standard matrix concentration bounds. Hence, using the spectral norm of N , we will be able to distinguish between the YES case and the NO case as long as

$$\sqrt{d/n} \ll n/(dk) \implies \lambda_{\text{NO}} \ll \lambda_{\text{YES}}.$$

For linear sparsity $k = \varepsilon \cdot n$, this inequality is true so long as $d \ll (n/\varepsilon^2)^{1/3}$, which is somewhat worse than the bound \sqrt{n} bound on the dimension that we are aiming for.

Recall that $\text{Tr} B = \sum_i \lambda_i(B)$ for a symmetric matrix B . As discussed above, the partial trace approach works best when the noise behaves as the tensor of two Wigner matrices, in that there are cancellations when the eigenvalues of the noise are summed. In our case, the noise terms $(a_i a_i^\top) \otimes (a_i a_i^\top)$ do not have this property, as in fact $\text{Tr} a_i a_i^\top = \|a_i\|^2 \approx d/n$. Thus, in order to improve the dimension bound, we will center the eigenvalue distribution of the noise part of the matrix. This will cause it to behave more like a Wigner matrix, in that the spectral norm of the noise will not increase after a partial trace. Consider the partial trace of a matrix of the form

$$M - \alpha \cdot \text{Id} \otimes \sum_i a_i a_i^\top,$$

for some constant $\alpha > 0$. The partial trace of this matrix is

$$N' = \sum_{i=1}^n (\|a_i\|_2^2 - \alpha) \cdot a_i a_i^\top.$$

We choose the constant $\alpha \approx d/n$ such that our matrix N' has expectation 0 in the NO case, when the subspace is completely random. In the YES case, the Rayleigh quotient of N' at x_0 simply shifts as compared to N , and we have $\lambda_{\text{YES}} \geq \langle x_0, N'x_0 \rangle \approx \|v_0\|_4^4 \geq 1/k$ (see [Lemma 4.5](#) and sublemmas). On the other hand, in the NO case, this centering operation causes significant cancellations in the eigenvalues of the partial trace matrix (instead of just shifting the eigenvalues). In the NO case, N' has spectral norm $\lambda_{\text{NO}} \leq O(d/n^{3/2})$ for $d \ll \sqrt{n}$ (using standard matrix concentration bounds; again see [Lemma 4.5](#) and sublemmas). Therefore, the spectral norm of the matrix N' allows us to distinguish between the YES and NO case as long as $d/n^{3/2} \ll 1/k$, which is satisfied as long as $k \ll n$ and $d \ll \sqrt{n}$. We give the full formal argument in [Section 4](#).

2.2 Overcomplete tensor decomposition

Recall that in this problem we are given a 3-tensor \mathbf{T} of the form $\mathbf{T} = \sum_{i=1}^n a_i^{\otimes 3} \in \mathbb{R}^{d^3}$, where $a_1, \dots, a_n \in \mathbb{R}^d$ are independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id})$. The goal is to find a unit vector $a \in \mathbb{R}^d$ that is highly correlated with one⁵ of the vectors a_1, \dots, a_n .

Background. The starting point of our algorithm is the polynomial $f(x) = \sum_{i=1}^n \langle a_i, x \rangle^3$. It turns out that for $n \ll d^{1.5}$ the (approximate) maximizers of this polynomial are close to the components a_1, \dots, a_n , in the sense that $f(x) \approx 1$ if and only if $\max_{i \in [n]} \langle a_i, x \rangle^2 \approx 1$. Indeed, Ge and Ma [\[GM15\]](#) show that the sum-of-squares method already captures this fact at degree 12, which implies a quasipolynomial time algorithm for this tensor decomposition problem via a general rounding result of Barak, Kelner, and Steurer [\[BKS15\]](#).

The simplest approach to this problem is to consider the tensor representation of the polynomial $\mathbf{T} = \sum_{i \in [n]} a_i^{\otimes 3}$, and flatten it, hoping the singular vectors of the flattening are correlated with the a_i . However, this approach is doomed to failure for two reasons: firstly, the simple flattenings of \mathbf{T} are $d^2 \times d$ matrices, and since $n \gg d$ the $a_i^{\otimes 2}$ collide in the column space, so that it is impossible to determine $\text{Span}\{a_i^{\otimes 2}\}$. Secondly, even for $n \leq d$, because the a_i are random vectors, their norms concentrate very closely about 1. This makes it difficult to distinguish any one particular a_i even when the span is computable.

Improvements. We will try to circumvent both of these issues by going to higher dimensions. Suppose, for example, that we had access to $\sum_{i \in [n]} a_i^{\otimes 4}$.⁶ The eigenvectors of the flattenings of this matrix are all within $\text{Span}_{i \in [n]} \{a_i^{\otimes 2}\}$, addressing our first issue, leaving us only with the trouble of extracting individual $a_i^{\otimes 2}$ from their span. If furthermore we had access to $\sum_{i \in [n]} a_i^{\otimes 6}$, we could perform a partial random projection $(\Phi \otimes \text{Id} \otimes \text{Id}) \sum_{i \in [n]} a_i^{\otimes 6}$ where $\Phi \in \mathbb{R}^{d \times d}$ is a matrix with independent Gaussian entries, and then taking a partial trace, we end up with

$$\text{Tr}_{\mathbb{R}^d} \left((\Phi \otimes \text{Id} \otimes \text{Id}) \sum_{i \in [n]} a_i^{\otimes 6} \right) = \sum_{i \in [n]} \langle \Phi, a_i^{\otimes 2} \rangle a_i^{\otimes 4}.$$

⁵ We can then approximately recover all the components a_1, \dots, a_n by running independent trials of our randomized algorithm repeatedly on the same input.

⁶ As the problem is defined, we assume that we do not have access to this input, and in many machine learning applications this is a valid assumption, as gathering the data necessary to generate the 4th order input tensor requires a prohibitively large number of samples.

With reasonable probability (for exposition’s sake, say with probability $1/n^{10}$), Φ is closer to some $a_i^{\otimes 2}$ than to all of the others so that $\langle \Phi, a_i^{\otimes 2} \rangle \geq 100 \langle \Phi, a_j^{\otimes 2} \rangle$ for all $j \in [n]$, and then $a_i^{\otimes 2}$ is distinguishable from the other vectors in the span of our matrix, taking care of the second issue. As we show, a much smaller gap is sufficient to distinguish the top a_i from the other a_j , and so the higher-probability event that Φ is only slightly closer to a_i suffices (allowing us to recover all vectors at an additional runtime cost of a factor of $\tilde{O}(n)$). This discussion ignores the presence of a single spurious large eigenvector, which we address in the technical sections.

Of course, we do not have access to the higher-order tensor $\sum_{i \in [n]} a_i^{\otimes 6}$. Instead, we can obtain a noisy version of this tensor. Our approach considers the following matrix representation of the polynomial f^2 ,

$$M = \sum_{i,j} a_i a_j^\top \otimes (a_i a_i^\top) \otimes (a_j a_j^\top) \in \mathbb{R}^{d^3 \times d^3}.$$

Alternatively, we can view this matrix as a particular flattening of the Kronecker-squared tensor $\mathbf{T}^{\otimes 2}$. It is instructive to decompose $M = M_{\text{diag}} + M_{\text{cross}}$ into its diagonal terms $M_{\text{diag}} = \sum_i (a_i a_i^\top)^{\otimes 3}$ and its cross terms $M_{\text{cross}} = \sum_{i \neq j} a_i a_j^\top \otimes (a_i a_i^\top) \otimes (a_j a_j^\top)$. The algorithm described above is already successful for M_{diag} ; we need only control the eigenvalues of the partial trace of the “noise” component, M_{cross} . The main technical work will be to show that $\|\text{Tr}_{\mathbb{R}^d} M_{\text{diag}}\|$ is small. In fact, we will have to choose Φ from a somewhat different distribution—observing that $\text{Tr}_{\mathbb{R}^d}(\Phi \otimes \text{Id} \otimes \text{Id}) = \sum_{i,j} \langle a_i, \Phi a_j \rangle \cdot (a_i \otimes a_j)(a_i \otimes a_j)^\top$, we will sample Φ so that $\langle a_i, \Phi a_i \rangle \gg \langle a_i, \Phi a_j \rangle$. We give a more detailed overview of this algorithm in the beginning of [Section 5](#), explaining in more detail our choice of Φ and justifying heuristically the boundedness of the spectral norm of the noise.

Connection to SoS analysis. To explain how the above algorithm is a speedup of SoS, we give an overview of the SoS algorithm of [\[GM15, BKS15\]](#). There, the degree- t SoS SDP program is used to obtain an order- t tensor χ_t (or a *pseudodistribution*). Informally speaking, we can understand χ_t as a proxy for $\sum_{i \in [n]} a_i^{\otimes t}$, so that $\chi_t = \sum_{i \in [n]} a_i^{\otimes t} + N$, where N is a noise tensor. While the precise form of N is unclear, we know that N must obey a set of constraints imposed by the SoS hierarchy at degree t . For a formal discussion of pseudodistributions, see [\[BKS15\]](#).

To extract a single component a_i from the tensor $\sum_{i \in [n]} a_i^{\otimes t}$, there are many algorithms which would work (for example, the algorithm we described for M_{diag} above). However, any algorithm extracting an a_i from χ_t must be robust to the noise tensor N . For this it turns out the following algorithm will do: suppose we have the tensor $\sum_{i \in [n]} a_i^{\otimes t}$, taking $t = O(\log n)$. Sample $g_1, \dots, g_{\log(n)-2}$ random unit vectors, and compute the matrix $M = \sum_i (\prod_{1 \leq j \leq \log(n)-2} \langle g_j, a_i \rangle) \cdot a_i a_i^\top$. If we are lucky enough, there is some a_i so that every g_j is a bit closer to a_i than any other $a_{i'}$, and $M = a_i a_i^\top + E$ for some $\|E\| \ll 1$. The proof that $\|E\|$ is small can be made so simple that it applies also to the SDP-produced proxy tensor $\chi_{\log n}$, and so this algorithm is robust to the noise N . This last step is very general and can handle tensors whose components a_i are less well-behaved than the random vectors we consider, and also more overcomplete, handling tensors of rank up to $n = \tilde{Q}(d^{1.5})$.⁷

Our subquadratic-time algorithm can be viewed as a low-degree, spectral analogue of the [\[BKS15\]](#) SoS algorithm. However, rather than relying on an SDP to produce an object close to $\sum_{i \in [n]} a_i^{\otimes t}$, we manufacture one ourselves by taking the Kronecker square of our input tensor. We

⁷ It is an interesting open question whether taking $t = O(\log n)$ is really necessary, or whether this heavy computational requirement is simply an artifact of the SoS proof.

explicitly know the form of the deviation of $\mathbf{T}^{\otimes 2}$ from $\sum_{i \in [n]} a_i^{\otimes 6}$, unlike in [BKS15], where the deviation of the SDP certificate χ_t from $\sum_{i \in [n]} a_i^{\otimes t}$ is poorly understood. We are thus able to control this deviation (or “noise”) in a less computationally intensive way, by cleverly designing a partial trace operation which decreases the spectral norm of the deviation. Since the tensor handled by the algorithm is much smaller—order 6 rather than order $\log n$ —this provides the desired speedup.

2.3 Tensor principal component analysis

Recall that in this problem we are given a tensor $\mathbf{T} = \tau \cdot v^{\otimes 3} + \mathbf{A}$, where $v \in \mathbb{R}^d$ is a unit vector, \mathbf{A} has iid entries from $\mathcal{N}(0, 1)$, and $\tau > 0$ is the signal-to-noise ratio. The aim is to recover v approximately.

Background and SoS analysis. A previous application of SoS techniques to this problem discussed several SoS or spectral algorithms, including one that runs in quasi-linear time [HSS15]. Here we apply the partial trace method to a subquadratic spectral SoS algorithm discussed in [HSS15] to achieve nearly the same signal-to-noise guarantee in only linear time.

Our starting point is the polynomial $\mathbf{T}(x) = \tau \cdot \langle v, x \rangle^3 + \langle x^{\otimes 3}, \mathbf{A} \rangle$. The maximizer of $\mathbf{T}(x)$ over the sphere is close to the vector v so long as $\tau \gg \sqrt{n}$ [RM14]. In [HSS15], it was shown that degree-4 SoS maximizing this polynomial can recover v with a signal-to-noise ratio of at least $\tilde{\Omega}(n^{3/4})$, since there exists a suitable SoS bound on the noise term $\langle x^{\otimes 3}, \mathbf{A} \rangle$.

Specifically, let A_i be the i th slice of \mathbf{A} , so that $\langle x, A_i x \rangle$ is the quadratic form $\sum_{j,k} \mathbf{A}_{ijk} x_j x_k$. Then there is a SoS proof that $\mathbf{T}(x)$ is bounded by $|\mathbf{T}(x) - \tau \cdot \langle v, x \rangle^3| \leq f(x)^{1/2} \cdot \|x\|$, where $f(x)$ is the degree-4 polynomial $f(x) = \sum_i \langle x, A_i x \rangle^2$. The polynomial f has a convenient matrix representation: $f(x) = \langle x^{\otimes 2}, (\sum_i A_i \otimes A_i) x^{\otimes 2} \rangle$: since this matrix is a sum of iid random matrices $A_i \otimes A_i$, a matrix Chernoff bound shows that this matrix spectrally concentrates to its expectation. So with high probability one can show that the eigenvalues of $\sum_i A_i \otimes A_i$ are at most $\approx d^{3/2} \log(d)^{1/2}$ (except for a single spurious eigenvector), and it follows that degree-4 SoS solves tensor PCA so long as $\tau \gg d^{3/4} \log(d)^{1/4}$.

This leads the authors to consider a slight modification of $f(x)$, given by $g(x) = \sum_i \langle x, T_i x \rangle^2$, where T_i is the i th slice of \mathbf{T} . Like \mathbf{T} , the function g also contains information about v , and the SoS bound on the noise term in \mathbf{T} carries over as an analogous bound on the noise in g . In particular, expanding $T_i \otimes T_i$ and ignoring some negligible cross-terms yields

$$\sum_i T_i \otimes T_i \approx \tau^2 \cdot (v \otimes v)(v \otimes v)^\top + \sum_i A_i \otimes A_i.$$

Using $v \otimes v$ as a test vector, the quadratic form of the latter matrix can be made at least $\tau^2 - O(d^{3/2} \log(d)^{1/2})$. Together with the boundedness of the eigenvalues of $\sum_i A_i \otimes A_i$ this shows that when $\tau \gg d^{3/4} \log(d)^{1/4}$ there is a spectral algorithm to recover v . Since the matrix $\sum_i T_i \otimes T_i$ is $d^2 \times d^2$, computing the top eigenvector requires $\tilde{O}(d^4 \log n)$ time, and by comparison to the input size d^3 the algorithm runs in subquadratic time.

Improvements. In this work we speed this up to a linear time algorithm via the partial trace approach. As we have seen, the heart of the matter is to show that taking the partial trace of

$\tau^2 \cdot (v \otimes v)(v \otimes v)^\top + \sum_i A_i \otimes A_i$ does not increase the spectral noise. That is, we require that

$$\left\| \text{Tr}_{\mathbb{R}^d} \sum_i A_i \otimes A_i \right\| = \left\| \sum_i \text{Tr}(A_i) \cdot A_i \right\| \leq O(d^{3/2} \log(d)^{1/2}).$$

The A_i have iid Gaussian entries, and so as in the case of Wigner matrices, it is roughly true that $|\text{Tr}(A_i)| \approx \|A_i\|$. Thus the situation is very similar to our toy example of the application of partial traces in [Section 2](#).

Heuristically, because $\sum_{i \in [n]} A_i \otimes A_i$ and $\sum_{i \in [n]} \text{Tr}(A_i) \cdot A_i$ are random matrices, we expect that their eigenvalues are all of roughly the same magnitude. This means that their spectral norm should be close to their Frobenius norm divided by the square root of the dimension, since for a matrix M with eigenvalues $\lambda_1, \dots, \lambda_n$, $\|M\|_F = \sqrt{\sum_{i \in [n]} \lambda_i^2}$. By estimating the sum of the squared entries, we expect that the Frobenius norm of $\sum_i \text{Tr}(A_i) \cdot A_i$ is less than that of $\sum_i A_i \otimes A_i$ by a factor of \sqrt{d} after the partial trace, while the dimension decreases by a factor of d , and so assuming that the eigenvalues are all of the same order, a typical eigenvalue should remain unchanged. We formalize these heuristic calculations using standard matrix concentration arguments in [Section 6](#).

3 Preliminaries

Linear algebra. We will work in the real vector spaces given by \mathbb{R}^n . A vector of indeterminates may be denoted $x = (x_1, \dots, x_n)$, although we may sometimes switch to parenthetical notation for indexing, i.e. $x = (x(1), \dots, x(n))$ when subscripts are already in use. We denote by $[n]$ the set of all valid indices for a vector in \mathbb{R}^n . Let e_i be the i th canonical basis vector so that $e_i(i) = 1$ and $e_i(j) = 0$ for $j \neq i$.

For a vectors space V , we may denote by $\mathcal{L}(V)$ the space of linear operators from V to V . The space orthogonal to a vector v is denoted v^\perp .

For a matrix M , we use M^{-1} to denote its inverse or its Moore-Penrose pseudoinverse; which one it is will be clear from context. For M PSD, we write $M^{-1/2}$ for the unique PSD matrix with $(M^{-1/2})^2 = M^{-1}$.

Norms and inner products. We denote the usual entrywise inner product by $\langle \cdot, \cdot \rangle$, so that $\langle u, v \rangle = \sum_{i \in [n]} u_i v_i$ for $u, v \in \mathbb{R}^n$. The ℓ_p -norm of a vector $v \in \mathbb{R}^n$ is given by $\|v\|_p = (\sum_{i \in [n]} v_i^p)^{1/p}$, with $\|v\|$ denoting the ℓ_2 -norm by default. The matrix norm used throughout the paper will be the operator / spectral norm, denoted by $\|M\| = \|M\|_{op} := \max_{x \neq 0} \|Mx\| / \|x\|$.

Tensor manipulation. Boldface variables will reserved for tensors $\mathbf{T} \in \mathbb{R}^{n \times n \times n}$, of which we consider only order-3 tensors. We denote by $\mathbf{T}(x, y, z)$ the multilinear function in $x, y, z \in \mathbb{R}^n$ such that $\mathbf{T}(x, y, z) = \sum_{i, j, k \in [n]} T_{i, j, k} x_i y_j z_k$, applying x, y , and z to the first, second, and third modes of the tensor \mathbf{T} respectively. If the arguments are matrices P, Q , and R instead, this lifts $\mathbf{T}(P, Q, R)$ to the unique multilinear tensor-valued function such that $[\mathbf{T}(P, Q, R)](x, y, z) = \mathbf{T}(Px, Qy, Rz)$ for all vectors x, y, z .

Tensors may be flattened to matrices in the multilinear way such that for every $u \in \mathbb{R}^{n \times n}$ and $v \in \mathbb{R}^n$, the tensor $u \otimes v$ flattens to the matrix $uv^\top \in \mathbb{R}^{n^2 \times n}$ with u considered as a vector.

There are 3 different ways to flatten a 3-tensor \mathbf{T} , corresponding to the 3 modes of \mathbf{T} . Flattening may be understood as reinterpreting the indices of a tensor when the tensor is expressed as an 3-dimensional array of numbers. The expression $v^{\otimes 3}$ refers to $v \otimes v \otimes v$ for a vector v .

Probability and asymptotic bounds. We will often refer to collections of independent and identically distributed (or *iid*) random variables. The Gaussian distribution with mean μ and variance σ^2 is denoted $\mathcal{N}(\mu, \sigma^2)$. Sometimes we state that an event happens with overwhelming probability. This means that its probability is at least $1 - n^{-\omega(1)}$. A function is $\tilde{O}(g(n))$ if it is $O(g(n))$ up to polylogarithmic factors.

4 Planted sparse vector in random linear subspace

In this section we give a nearly-linear-time algorithm to recover a sparse vector planted in a random subspace.

Problem 4.1. Let $v_0 \in \mathbb{R}^n$ be a unit vector such that $\|v_0\|_4^4 \geq \frac{1}{\varepsilon n}$. Let $v_1, \dots, v_{d-1} \in \mathbb{R}^n$ be iid from $\mathcal{N}(0, \frac{1}{n} \text{Id}_n)$. Let w_0, \dots, w_{d-1} be an orthogonal basis for $\text{Span}\{v_0, \dots, v_{d-1}\}$. **Given:** w_0, \dots, w_{d-1}
Find: a vector v such that $\langle v, v_0 \rangle^2 \geq 1 - o(1)$.

Sparse Vector Recovery in Nearly-Linear Time

Algorithm 4.2. Input: w_0, \dots, w_{d-1} as in [Problem 4.1](#). Goal: Find v with $\langle \hat{v}, v_0 \rangle^2 \geq 1 - o(1)$.

- Compute leverage scores $\|a_1\|^2, \dots, \|a_n\|^2$, where a_i is the i th row of the $n \times d$ matrix $S := \begin{pmatrix} w_0 & \cdots & w_{d-1} \end{pmatrix}$.
- Compute the top eigenvector u of the matrix

$$A \stackrel{\text{def}}{=} \sum_{i \in [n]} \left(\|a_i\|_2^2 - \frac{d}{n} \right) \cdot a_i a_i^\top.$$

- Output Su .

Remark 4.3 (Implementation of [Algorithm 4.2](#) in nearly-linear time). The leverage scores $\|a_1\|^2, \dots, \|a_n\|^2$ are clearly computable in time $O(nd)$. In the course of proving correctness of the algorithm we will show that A has constant spectral gap, so by a standard analysis $O(\log d)$ matrix-vector multiplies suffice to recover its top eigenvector. A single matrix-vector multiply Ax requires computing $c_i := \left(\|a_i\|^2 - \frac{d}{n} \right) \langle a_i, x \rangle$ for each i (in time $O(nd)$) and summing $\sum_{i \in [n]} c_i x_i$ (in time $O(nd)$). Finally, computing Su requires summing d vectors of dimension n , again taking time $O(nd)$.

The following theorem expresses correctness of the algorithm.

Theorem 4.4. Let $v_0 \in \mathbb{R}^n$ be a unit vector with $\|v_0\|_4^4 \geq \frac{1}{\varepsilon n}$. Let $v_1, \dots, v_{d-1} \in \mathbb{R}^n$ be iid from $\mathcal{N}(0, \frac{1}{n} \text{Id}_n)$. Let w_0, \dots, w_{d-1} be an orthogonal basis for $\text{Span}\{v_0, \dots, v_{d-1}\}$. Let a_i be the i -th row of the $n \times d$ matrix $S := \begin{pmatrix} w_0 & \cdots & w_{d-1} \end{pmatrix}$.

When $d \leq n^{1/2} / \text{polylog}(n)$, for any sparsity $\varepsilon > 0$, w.ov.p. the top eigenvector u of $\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top$ has $\langle Su, v_0 \rangle^2 \geq 1 - O(\varepsilon^{1/4}) - o(1)$.

We have little control over the basis vectors the algorithm is given. However, there is a particularly nice (albeit non-orthogonal) basis for the subspace which exposes the underlying randomness. Suppose that we are given the basis vectors v_0, \dots, v_d , where v_0 is the sparse vector normalized so that $\|v_0\| = 1$, and v_1, \dots, v_{d-1} are iid samples from $\mathcal{N}(0, \frac{1}{n} \text{Id}_n)$. The following lemma shows that if the algorithm had been handed this good representation of the basis rather than an arbitrary orthogonal one, its output would be the correlated to the vector of coefficients giving of the planted sparse vector (in this case the standard basis vector e_1).

Lemma 4.5. *Let $v_0 \in \mathbb{R}^n$ be a unit vector. Let $v_1, \dots, v_{d-1} \in \mathbb{R}^n$ be iid from $\mathcal{N}(0, \frac{1}{n} \text{Id})$. Let a_i be the i th row of the $n \times d$ matrix $S := \begin{pmatrix} v_0 & \cdots & v_{d-1} \end{pmatrix}$. Then there is a universal constant $\varepsilon^* > 0$ so that for any $\varepsilon \leq \varepsilon^*$, so long as $d \leq n^{1/2} / \text{polylog}(n)$, w.ov.p.*

$$\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top = \|v_0\|_4^4 \cdot e_1 e_1^\top + M,$$

where e_1 is the first standard basis vector and $\|M\| \leq O(\|v_0\|_4^3 \cdot n^{-1/4} + \|v_0\|_4^2 \cdot n^{-1/2} + \|v_0\|_4 \cdot n^{-3/4} + n^{-1})$.

The second ingredient we need is that the algorithm is robust to exchanging this good basis for an arbitrary orthogonal basis.

Lemma 4.6. *Let $v_0 \in \mathbb{R}^n$ have $\|v_0\|_4^4 \geq \frac{1}{\varepsilon n}$. Let $v_1, \dots, v_{d-1} \in \mathbb{R}^n$ be iid from $\mathcal{N}(0, \frac{1}{n} \text{Id}_n)$. Let w_0, \dots, w_{d-1} be an orthogonal basis for $\text{Span}\{v_0, \dots, v_{d-1}\}$. Let a_i be the i th row of the $n \times d$ matrix $S := \begin{pmatrix} v_0 & \cdots & v_{d-1} \end{pmatrix}$. Let a'_i be the i th row of the $n \times d$ matrix $S' := \begin{pmatrix} w_0 & \cdots & w_{d-1} \end{pmatrix}$. Let $A := \sum_i a_i a_i^\top$. Let $Q \in \mathbb{R}^{d \times d}$ be the orthogonal matrix so that $SA^{-1/2} = S'Q$, which exists since $SA^{-1/2}$ is orthogonal, and which has the effect that $a'_i = QA^{-1/2}a_i$. Then when $d \leq n^{1/2} / \text{polylog}(n)$, w.ov.p.*

$$\left\| \sum_{i=1}^n (\|a'_i\|^2 - \frac{d}{n}) \cdot a'_i a_i'^\top - Q \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) Q^\top \right\| \leq O\left(\frac{1}{n}\right) + o(\|v\|_4^4)$$

Last, we will need the following fact, which follows from standard concentration. The proof is in [Section B](#).

Lemma 4.7. *Let $v \in \mathbb{R}^n$ be a unit vector. Let $b_1, \dots, b_n \in \mathbb{R}^{d-1}$ be iid from $\mathcal{N}(0, \frac{1}{n} \text{Id}_{d-1})$. Let $a_i \in \mathbb{R}^d$ be given by $a_i := (v(i) \ b_i)$. Then w.ov.p. $\|\sum_{i=1}^n a_i a_i^\top - \text{Id}_d\| \leq \tilde{O}(d/n)^{1/2}$. In particular, when $d = o(n)$, this implies that w.ov.p. $\|(\sum_{i=1}^n a_i a_i^\top)^{-1} - \text{Id}_d\| \leq \tilde{O}(d/n)^{1/2}$ and $\|(\sum_{i=1}^n a_i a_i^\top)^{-1/2} - \text{Id}_d\| \leq \tilde{O}(d/n)^{1/2}$.*

We are ready to prove [Theorem 4.4](#).

Proof of [Theorem 4.4](#). Let b_1, \dots, b_n be the rows of the matrix $S' := \begin{pmatrix} v_0 & \cdots & v_{d-1} \end{pmatrix}$. Let $B = \sum_i b_i b_i^\top$. Note that $S'B^{-1/2}$ has columns which are an orthogonal basis for $\text{Span}\{w_0, \dots, w_{d-1}\}$. Let $Q \in \mathbb{R}^{d \times d}$ be the rotation so that $S'B^{-1/2} = SQ$.

By [Lemma 4.5](#) and [Lemma 4.6](#), we can write the matrix $A = \sum_{i=1}^n (\|a_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top$ as

$$A = \|v_0\|_4^4 \cdot Q e_1 e_1^\top Q^\top + M,$$

where w.ov.p.

$$\|M\| \leq O(\|v_0\|_4^3 \cdot n^{-1/4} + \|v_0\|_4^2 \cdot n^{-1/2} + \|v_0\|_4 \cdot n^{-3/4} + n^{-1}) + o(\|v\|_4^4).$$

We have assumed that $\|v_0\|_4^4 \geq (\varepsilon n)^{-1}$, and so since A is an almost-rank-one matrix ([Lemma A.3](#)), the top eigenvector u of A has $\langle u, Qe_1 \rangle^2 \geq 1 - O(\varepsilon^{1/4})$, so that $\langle Su, SQe_1 \rangle^2 \geq 1 - O(\varepsilon^{1/4})$ by column-orthogonality of S .

At the same time, $SQe_1 = S'B^{-1/2}e_1$, and by [Lemma 4.7](#), $\|B^{-1/2} - \text{Id}\| \leq \tilde{O}(d/n)^{1/2}$ w.ov.p., so that $\langle Su, S'e_1 \rangle^2 \geq \langle Su, SQe_1 \rangle^2 - o(1)$. Finally, $S'e_1 = v_0$ by definition, so $\langle Su, v_0 \rangle^2 \geq 1 - O(\varepsilon^{1/4}) - o(1)$. \square

4.1 Algorithm succeeds on good basis

We now prove [Lemma 4.5](#). We decompose the matrix in question into a contribution from $\|v_0\|_4^4$ and the rest: explicitly, the decomposition is $\sum(\|a_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top = \sum v(i)^2 \cdot a_i a_i^\top + \sum(\|b_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top$. This first lemma handles the contribution from $\|v_0\|_4^4$.

Lemma 4.8. *Let $v \in \mathbb{R}^n$ be a unit vector. Let $b_1, \dots, b_n \in \mathbb{R}^{d-1}$ be random vectors iid from $\mathcal{N}(0, \frac{1}{n} \cdot \text{Id}_{d-1})$. Let $a_i = (v(i) \ b_i) \in \mathbb{R}^d$. Suppose $d \leq n^{1/2} / \text{polylog}(n)$. Then*

$$\sum_{i=1}^n v(i)^2 \cdot a_i a_i^\top = \|v\|_4^4 \cdot e_1 e_1^\top + M',$$

where $\|M'\| \leq O(\|v\|_4^3 n^{-1/4} + \|v\|_4^2 n^{-1/2})$ w.ov.p..

Proof of Lemma 4.8. We first show an operator-norm bound on the principal submatrix $\sum_{i=1}^n v(i)^2 \cdot b_i b_i^\top$ using the truncated matrix Bernstein inequality [Proposition A.7](#). First, the expected operator norm of each summand is bounded:

$$\mathbb{E} v(i)^2 \|b_i\|_2^2 \leq (\max_j v(j)^2) \cdot O\left(\frac{d}{n}\right) \leq \|v\|_4^2 \cdot O\left(\frac{d}{n}\right).$$

The operator norms are bounded by constant-degree polynomials in Gaussian variables, so [Lemma A.8](#) applies to truncate their tails in preparation for application of a Bernstein bound. We just have to calculate the variance of the sum, which is at most

$$\left\| \mathbb{E} \sum_{i=1}^n v(i)^4 \|b_i\|_2^2 \cdot b_i b_i^\top \right\| = \|v\|_4^4 \cdot O\left(\frac{d}{n^2}\right).$$

The expectation $\mathbb{E} \sum_{i=1}^n v(i)^2 \cdot b_i b_i^\top$ is $\frac{\|v\|_4^2}{n} \cdot \text{Id}$. Applying a matrix Bernstein bound ([Proposition A.7](#)) to the deviation from expectation, we get that w.ov.p.,

$$\left\| \left(\sum_{i=1}^n v(i)^2 \cdot b_i b_i^\top \right) - \frac{1}{n} \cdot \text{Id} \right\| \leq \|v\|_4^2 \cdot \tilde{O}\left(\frac{d}{n}\right) \leq O(\|v\|_4^2 n^{-1/2})$$

for appropriate choice of $d \leq n^{-1/2} / \text{polylog}(n)$. Hence, by triangle inequality, $\|\sum_{i=1}^n v(i)^2 \cdot b_i b_i^\top\| \leq \|v\|_4^2 n^{-1/2}$ w.ov.p..

Using a Cauchy-Schwarz-style inequality ([Lemma A.1](#)) we now show that the bound on this principal submatrix is essentially enough to obtain the lemma. Let $p_i, q_i \in \mathbb{R}^d$ be given by

$$p_i \stackrel{\text{def}}{=} v_0(i) \cdot \begin{pmatrix} v_0(i) \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad q_i \stackrel{\text{def}}{=} v_0(i) \cdot \begin{pmatrix} 0 \\ b_i \end{pmatrix}.$$

Then

$$\sum_{i=1}^n v(i)^2 \cdot b_i b_i^\top = \|v\|_4^4 + \sum_{i=1}^n p_i q_i^\top + q_i p_i^\top + q_i q_i^\top.$$

We have already bounded $\sum_{i=1}^n q_i q_i^\top = \sum_{i=1}^n v(i)^2 \cdot b_i b_i^\top$. At the same time, $\|\sum_{i=1}^n p_i p_i^\top\| = \|v\|_4^4$. By [Lemma A.1](#), then,

$$\left\| \sum_{i=1}^n p_i q_i^\top + q_i p_i^\top \right\| \leq O(\|v\|_4^3 n^{-1/4})$$

w.ov.p.. A final application of triangle inequality gives the lemma. \square

Our second lemma controls the contribution from the random part of the leverage scores.

Lemma 4.9. *Let $v \in \mathbb{R}^n$ be a unit vector. Let $b_1, \dots, b_n \in \mathbb{R}^{d-1}$ be random vectors iid from $\mathcal{N}(0, \frac{1}{n} \cdot \text{Id}_{d-1})$. Let $a_i = (v(i) \ b_i) \in \mathbb{R}^d$. Suppose $d \leq n^{1/2} / \text{polylog}(n)$. Then w.ov.p.*

$$\left\| \sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top \right\| \leq \|v\|_4^2 \cdot O(n^{-3/4}) + \|v\|_4 \cdot O(n^{-1}) + O(n^{-1}).$$

Proof. Like in the proof of [Lemma 4.8](#), $\sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top$ decomposes into a convenient block structure; we will bound each block separately.

$$\sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top = \sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) \cdot \begin{pmatrix} v(i)^2 & v(i) \cdot b_i^\top \\ v(i) \cdot b_i & b_i b_i^\top \end{pmatrix}. \quad (4.1)$$

In each block we can apply a (truncated) Bernstein inequality. For the large block $\sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) b_i b_i^\top$, the choice $\frac{d}{n}$ ensures that $\mathbb{E}(\|b_i\|_2^2 - \frac{d}{n}) b_i b_i^\top = O(\frac{1}{n^2}) \cdot \text{Id}$. The expected operator norm of each summand is small:

$$\begin{aligned} \mathbb{E} \|(\|b_i\|_2^2 - \frac{d}{n}) b_i b_i^\top\| &= \mathbb{E} |(\|b_i\|_2^2 - \frac{d}{n})| \|b_i\|_2^2 \\ &\leq (\mathbb{E}(\|b_i\|_2^2 - \frac{d}{n})^2)^{1/2} (\mathbb{E} \|b_i\|_2^4)^{1/2} \quad \text{by Cauchy-Schwarz} \\ &\leq O\left(\frac{d^{1/2}}{n}\right) \cdot O\left(\frac{d}{n}\right) \quad \text{variance of } \chi^2 \text{ with } k \text{ degrees of freedom is } O(k) \\ &= O\left(\frac{d^{3/2}}{n^2}\right). \end{aligned}$$

The termwise operator norms are bounded by constant-degree polynomials in Gaussian variables, so [Lemma A.8](#) applies to truncate the tails of the summands in preparation for a Bernstein bound.

We just have to compute the variance of the sum, which is small because we have centered the coefficients:

$$\left\| \sum_i \mathbb{E}(\|b_i\|_2^2 - \frac{d}{n})^2 \|b_i\|_2^2 \cdot b_i b_i^\top \right\| \leq O\left(\frac{d^2}{n^3}\right)$$

by direct computation of $\mathbb{E}(\|b_i\|_2^2 - \frac{d}{n})^2 \|b_i\|_2^2 b_i b_i^\top$ using [Fact A.6](#). These facts together are enough to apply the matrix Bernstein inequality ([Proposition A.7](#)) and conclude that w.ov.p.

$$\left\| \sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) \cdot b_i b_i^\top \right\| \leq \tilde{O}\left(\frac{d}{n^{3/2}}\right) \leq O\left(\frac{1}{n}\right)$$

for appropriate choice of $d \leq n / \text{polylog}(n)$.

We turn to the other blocks from [\(4.1\)](#). The upper-left block contains just the scalar $\sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) v(i)^2$. By standard concentration each term is bounded: w.ov.p.,

$$(\|b_i\|_2^2 - \frac{d}{n}) v(i)^2 \leq (\max_i v(i)^2) \cdot \tilde{O}\left(\frac{d^{1/2}}{n}\right) \leq \|v\|_4^2 \cdot \tilde{O}\left(\frac{d^{1/2}}{n}\right).$$

The sum has variance at most $\sum_{i=1}^n v(i)^4 \mathbb{E}(\|b_i\|_2^2 - \frac{d}{n})^2 \leq \|v\|_4^4 \cdot O(d/n^2)$. Again using [Lemma A.8](#) and [Proposition A.7](#), we get that w.ov.p.

$$\left| \sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) v(i)^2 \right| \leq \|v\|_4^2 \cdot \tilde{O}\left(\frac{d^{1/2}}{n}\right).$$

It remains just to address the block $\sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) v(i) \cdot b_i$. Each term in the sum has expected operator norm at most

$$(\max_i v(i)^2)^{1/2} \cdot O\left(\frac{d}{n^{3/2}}\right) \leq \|v\|_4 \cdot O\left(\frac{d}{n^{3/2}}\right),$$

and once again since the summands' operator norms are bounded by constant-degree polynomials of Gaussian variables [Lemma A.8](#) applies to truncate their tails in preparation to apply a Bernstein bound. The variance of the sum is at most $\|v\|_2^2 \cdot O(d^2/n^3)$, again by [Fact A.6](#). Finally, [Lemma A.8](#) and [Proposition A.7](#) apply to give that w.ov.p.

$$\left\| \sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) v(i) \cdot b_i \right\| \leq \|v\|_4 \cdot \tilde{O}\left(\frac{d}{n^{3/2}}\right) + \tilde{O}\left(\frac{d}{n^{3/2}}\right) = \|v\|_4 \cdot n^{-1} + n^{-1}$$

for appropriate choice of $d \leq n^{1/2} / \text{polylog}(n)$. Putting it all together gives the lemma. \square

We are now ready to prove [Lemma 4.5](#)

Proof of [Lemma 4.5](#). We decompose $\|a_i\|_2^2 = v_0(i)^2 + \|b_i\|_2^2$ and use [Lemma 4.8](#) and [Lemma 4.9](#).

$$\begin{aligned} \sum_{i=1}^n (\|a_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top &= \left(\sum_{i=1}^n v_0(i)^2 \cdot a_i a_i^\top \right) + \left(\sum_{i=1}^n (\|b_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) \\ &= \|v_0\|_4^4 \cdot e_1 e_1^\top + M, \end{aligned}$$

where

$$\|M\| \leq O(\|v_0\|_4^3 \cdot n^{-1/4} + \|v_0\|_4^2 \cdot n^{-1/2}) + O(\|v_0\|_4 \cdot n^{-1} + n^{-1}).$$

Since $\|v_0\|_4^4 \geq (\varepsilon n)^{-1}$, we get $\|v_0\|_4^4 / \|M\| \geq \frac{1}{\varepsilon^{1/4}}$, completing the proof. \square

4.2 Closeness of input basis and good basis

We turn now to the proof of [Lemma 4.6](#). We recall the setting. We have two matrices: M , which the algorithm computes, and M' , which is induced by a basis for the subspace which reveals the underlying randomness and which we prefer for the analysis. M' differs from M by a rotation and a basis orthogonalization step (the good basis is only almost orthogonal). The rotation is easily handled. The following lemma gives the critical fact about the orthogonalization step: orthogonalizing does not change the leverage scores too much. ⁸

Lemma 4.10 (Restatement of [Lemma B.4](#)). *Let $v \in \mathbb{R}^n$ be a unit vector and let $b_1, \dots, b_n \in \mathbb{R}^{d-1}$ be iid from $\mathcal{N}(0, \frac{1}{n} \text{Id}_{d-1})$. Let $a_i \in \mathbb{R}^d$ be given by $a_i := (v(i) b_i)$. Let $A := \sum_i a_i a_i^\top$. Let $c \in \mathbb{R}^{d-1}$ be given by $c := \sum_i v(i) b_i$. Then for every index $i \in [n]$, w.o.v.p.,*

$$\left| \|A^{-1/2} a_i\|^2 - \|a_i\|^2 \right| \leq \tilde{O} \left(\frac{d + \sqrt{n}}{n} \right) \cdot \|a_i\|^2.$$

The proof again uses standard concentration and matrix inversion formulas, and can be found in [Section B](#). We are ready to prove [Lemma 4.6](#).

Proof of [Lemma 4.6](#). The statement we want to show is

$$\left\| \sum_{i=1}^n (\|a'_i\|^2 - \frac{d}{n}) \cdot a'_i a'_i{}^\top - Q \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i{}^\top \right) Q^\top \right\| \leq O \left(\frac{1}{n} \right) + o(\|v\|_4^4).$$

Conjugating by Q and multiplying by -1 does not change the operator norm, so that this is equivalent to

$$\left\| \sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i{}^\top - Q^\top \left(\sum_{i=1}^n (\|a'_i\|^2 - \frac{d}{n}) \cdot a'_i a'_i{}^\top \right) Q \right\| \leq O \left(\frac{1}{n} \right) + o(\|v\|_4^4).$$

Finally, substituting $a'_i = QA^{-1/2} a_i$, and using the fact that Q is a rotation, it will be enough to show

$$\left\| \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i{}^\top \right) - A^{-1/2} \left(\sum_{i=1}^n (\|A^{-1/2} a_i\|^2 - \frac{d}{n}) \cdot a_i a_i{}^\top \right) A^{-1/2} \right\| \leq O \left(\frac{1}{n} \right) + o(\|v\|_4^4). \quad (4.2)$$

We write the right-hand matrix as

$$\begin{aligned} & A^{-1/2} \left(\sum_{i=1}^n (\|A^{-1/2} a_i\|^2 - \frac{d}{n}) \cdot a_i a_i{}^\top \right) A^{-1/2} \\ &= A^{-1/2} \left(\sum_{i=1}^n (\|A^{-1/2} a_i\|^2 - \|a_i\|^2) \cdot a_i a_i{}^\top \right) A^{-1/2} + A^{-1/2} \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i{}^\top \right) A^{-1/2}. \end{aligned}$$

The first of these we observe has bounded operator norm w.o.v.p.:

$$\left\| A^{-1/2} \left(\sum_{i=1}^n (\|A^{-1/2} a_i\|^2 - \|a_i\|^2) \cdot a_i a_i{}^\top \right) A^{-1/2} \right\| \leq \left\| A^{-1/2} \left(\sum_{i=1}^n |\|A^{-1/2} a_i\|^2 - \|a_i\|^2| \cdot a_i a_i{}^\top \right) A^{-1/2} \right\|$$

⁸Strictly speaking the good basis does not have leverage scores since it is not orthogonal, but we can still talk about the norms of the rows of the matrix whose columns are the basis vectors.

$$\leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \left\| \sum_{i=1}^n \|a_i\|^2 \cdot a_i a_i^\top \right\|$$

where we have used [Lemma 4.7](#) to find that $A^{1/2}$ is close to identity, and [Lemma 4.10](#) to simplify the summands

$$\begin{aligned} &= \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \left(\left\| \sum_{i=1}^n v_0(i)^2 \cdot a_i a_i^\top \right\| + \left\| \sum_{i=1}^n \|b_i\|_2^2 \cdot a_i a_i^\top \right\| \right) \\ &\leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \left(O(\|v\|_4^4) + \tilde{O}\left(\frac{d}{n}\right) \right), \end{aligned}$$

using in the last step [Lemma 4.8](#) and standard concentration to bound $\sum_{i=1}^n \|b_i\|_2^2 \cdot a_i a_i^\top$ ([Lemma 4.7](#)). Thus, by triangle inequality applied to [\(4.2\)](#), we get

$$\begin{aligned} &\left\| \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) - A^{-1/2} \left(\sum_{i=1}^n (\|A^{-1/2} a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) A^{-1/2} \right\| \\ &\leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \left(O(\|v\|_4^4) + \tilde{O}\left(\frac{d}{n}\right) \right) + \left\| \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) - A^{-1/2} \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) A^{-1/2} \right\|. \end{aligned}$$

Finally, since w.ov.p. $\|A^{-1/2} - \text{Id}\| = \tilde{O}(d/n)^{1/2}$, we get

$$\begin{aligned} &\left\| \left(\sum_{i=1}^n (\|a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) - A^{-1/2} \left(\sum_{i=1}^n (\|A^{-1/2} a_i\|^2 - \frac{d}{n}) \cdot a_i a_i^\top \right) A^{-1/2} \right\| \\ &\leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \left(O(\|v\|_4^4) + \tilde{O}\left(\frac{d}{n}\right) \right) + \tilde{O}\left(\frac{d}{n}\right)^{1/2} \cdot \left\| \sum_{i=1}^n (\|a_i\|_2^2 - \frac{d}{n}) \cdot a_i a_i^\top \right\| \\ &\leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \left(O(\|v\|_4^4) + \tilde{O}\left(\frac{d}{n}\right) \right) + \tilde{O}\left(\frac{d}{n}\right)^{1/2} \cdot O(\|v\|_4^4). \end{aligned}$$

using [Lemma 4.5](#) in the last step. For appropriate choice of $d \leq n^{-1/2}/\text{polylog}(n)$, this is at most $O(n^{-1}) + o(\|v\|_4^4)$. \square

5 Overcomplete tensor decomposition

In this section, we give a polynomial-time algorithm for the following problem when $n \leq d^{4/3}/(\text{polylog } d)$:

Problem 5.1. Given an order-3 tensor $\mathbf{T} = \sum_{i=1}^n a_i \otimes a_i \otimes a_i$, where $a_1, \dots, a_n \in \mathbb{R}^d$ are iid vectors sampled from $\mathcal{N}(0, \frac{1}{d} \text{Id})$, find vectors $b_1, \dots, b_n \in \mathbb{R}^n$ such that for all $i \in [n]$,

$$\langle a_i, b_i \rangle \geq 1 - o(1).$$

We give an algorithm that solves this problem, so long as the overcompleteness of the input tensor is bounded such that $n \ll d^{4/3}/\text{polylog } d$.

Theorem 5.2. *Given as input the tensor $\mathbf{T} = \sum_{i=1}^n a_i \otimes a_i \otimes a_i$ where $a_i \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ with $d \leq n \leq d^{4/3} / \text{polylog } d$,⁹ there is an algorithm which may run in time $\tilde{O}(nd^{1+\omega})$ or $\tilde{O}(nd^{3.257})$, where d^ω is the time to multiply two $d \times d$ matrices, which with probability $1 - o(1)$ over the input \mathbf{T} and the randomness of the algorithm finds unit vectors $b_1, \dots, b_n \in \mathbb{R}^d$ such that for all $i \in [n]$,*

$$\langle a_i, b_i \rangle \geq 1 - \tilde{O}\left(\frac{n^{3/2}}{d^2}\right).$$

We remark that this accuracy can be improved from $1 - \tilde{O}(n^{3/2}/d^2)$ to an arbitrarily good precision using existing local search methods with local convergence guarantees—see [Corollary 5.23](#).

As discussed in [Section 2](#), to decompose the tensor $\sum_i a_i^{\otimes 6}$ (note we do not actually have access to this input!) there is a very simple tensor decomposition algorithm: sample a random $g \in \mathbb{R}^d$ and compute the matrix $\sum_i \langle g, a_i^{\otimes 2} \rangle (a_i a_i^\top)^{\otimes 2}$. With probability roughly $n^{-O(\varepsilon)}$ this matrix has (up to scaling) the form $(a_i a_i^\top)^{\otimes 2} + E$ for some $\|E\| \leq 1 - \varepsilon$, and this is enough to recover a_i .

However, instead of $\sum_i a_i^{\otimes 6}$, we have only $\sum_{i,j} (a_i \otimes a_j)^{\otimes 3}$. Unfortunately, running the same algorithm on the latter input will not succeed. To see why, consider the extra terms $E' := \sum_{i \neq j} \langle g, a_i \otimes a_j \rangle (a_i \otimes a_j)^{\otimes 2}$. Since $|\langle g, a_i \otimes a_j \rangle| \approx 1$, it is straightforward to see that $\|E'\|_F \approx n$. Since the rank of E' is clearly d^2 , even if we are lucky and all the eigenvalues have similar magnitudes, still a typical eigenvalue will be $\approx n/d \gg 1$, swallowing the $\sum_i a_i^{\otimes 6}$ term.

A convenient feature separating the signal terms $\sum_i (a_i \otimes a_i)^{\otimes 3}$ from the crossterms $\sum_{i \neq j} (a_i \otimes a_j)^{\otimes 3}$ is that the crossterms are not within the span of the $a_i \otimes a_i$. Although we cannot algorithmically access $\text{Span}\{a_i \otimes a_i\}$, we have access to something almost as good: the unfolded input tensor, $T = \sum_{i \in [n]} a_i (a_i \otimes a_i)^\top$. The rows of this matrix lie in $\text{Span}\{a_i \otimes a_i\}$, and so for $i \neq j$, $\|T(a_i \otimes a_i)\| \gg \|T(a_i \otimes a_j)\|$. In fact, careful computation reveals that $\|T(a_i \otimes a_i)\| \geq \tilde{\Omega}(\sqrt{n}/d) \|T(a_i \otimes a_j)\|$.

The idea now is to replace $\sum_{i,j} \langle g, a_i \otimes a_j \rangle (a_i \otimes a_j)^{\otimes 2}$ with $\sum_{i,j} \langle g, T(a_i \otimes a_j) \rangle (a_i \otimes a_j)^{\otimes 2}$, now with $g \sim \mathcal{N}(0, \text{Id}_d)$. As before, we are hoping that there is i_0 so that $\langle g, T(a_{i_0} \otimes a_{i_0}) \rangle \gg \max_{j \neq i_0} \langle g, T(a_j \otimes a_j) \rangle$. But now we also require $\|\sum_{i \neq j} \langle g, T(a_i \otimes a_j) \rangle (a_i \otimes a_j) (a_i \otimes a_j)^\top\| \ll \langle g, T(a_{i_0} \otimes a_{i_0}) \rangle \approx \|T(a_{i_0} \otimes a_{i_0})\|$. If we are lucky and all the eigenvalues of this cross-term matrix have roughly the same magnitude (indeed, we will be lucky in this way), then we can estimate heuristically that

$$\begin{aligned} \left\| \sum_{i \neq j} \langle g, T(a_i \otimes a_j) \rangle (a_i \otimes a_j) (a_i \otimes a_j)^\top \right\| &\approx \frac{1}{d} \left\| \sum_{i \neq j} \langle g, T(a_i \otimes a_j) \rangle (a_i \otimes a_j) (a_i \otimes a_j)^\top \right\|_F \\ &\leq \frac{1}{d} \cdot \frac{\sqrt{n}}{d} |\langle g, T(a_{i_0} \otimes a_{i_0}) \rangle| \left\| \sum_{i \neq j} (a_i \otimes a_j) (a_i \otimes a_j)^\top \right\|_F \\ &\leq \frac{n^{3/2}}{d^2} |\langle g, T(a_{i_0} \otimes a_{i_0}) \rangle|, \end{aligned}$$

suggesting our algorithm will succeed when $n^{3/2} \ll d^2$, which is to say $n \ll d^{4/3}$.

The following theorem, which formalizes the intuition above, is at the heart of our tensor decomposition algorithm.

⁹The lower bound $d \leq n$ on n , is a matter of technical convenience, avoiding separate concentration analyses and arithmetic in the undercomplete ($n < d$) and overcomplete ($n \geq d$) settings. Indeed, our algorithm still works in the undercomplete setting (tensor decomposition is easier in the undercomplete setting than the overcomplete one), but here other algorithms based on local search also work [\[AGJ15\]](#).

Theorem 5.3. Let a_1, \dots, a_n be independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ with $d \leq n \leq d^{4/3}/(\text{polylog } d)$ and let g be a random vector from $\mathcal{N}(0, \text{Id}_d)$. Let $\Sigma := \mathbb{E}_{x \sim \mathcal{N}(0, \text{Id}_d)}(xx^\top)^{\otimes 2}$ and let $R := \sqrt{2} \cdot (\Sigma^+)^{1/2}$. Let $T = \sum_{i \in [n]} a_i(a_i \otimes a_i)^\top$. Define the matrix $M \in \mathbb{R}^{d^2 \times d^2}$,

$$M = \sum_{i, j \in [n]} \langle g, T(a_i \otimes a_j) \rangle \cdot (a_i \otimes a_j)(a_i \otimes a_j)^\top.$$

With probability $1 - o(1)$ over the choice of a_1, \dots, a_n , for every $\text{polylog } d / \sqrt{d} < \varepsilon < 1$, the spectral gap of RMR is at least $\lambda_2/\lambda_1 \leq 1 - O(\varepsilon)$ and the top eigenvector $u \in \mathbb{R}^{d^2}$ of RMR satisfies, with probability $\tilde{\Omega}(1/n^{O(\varepsilon)})$ over the choice of g ,

$$\max_{i \in [n]} \langle Ru, a_i \otimes a_i \rangle^2 / (\|u\|^2 \cdot \|a_i\|^4) \geq 1 - \tilde{O}\left(\frac{n^{3/2}}{\varepsilon d^2}\right).$$

Moreover, with probability $1 - o(1)$ over the choice of a_1, \dots, a_n , for every $\text{polylog } d / \sqrt{d} < \varepsilon < 1$ there are events E_1, \dots, E_n so that $\mathbb{P}_g E_i \geq \tilde{\Omega}(1/n^{1+O(\varepsilon)})$ for all $i \in [n]$ and when E_i occurs, $\langle Ru, a_i \otimes a_i \rangle^2 / \|u\|^2 \cdot \|a_i\|^4 \geq 1 - \tilde{O}\left(\frac{n^{3/2}}{\varepsilon d^2}\right)$.

We will eventually set $\varepsilon = 1/\log n$, which gives us a spectral algorithm for recovering a vector $(1 - \tilde{O}(n/d^{3/2}))$ -correlated to some $a_i^{\otimes 2}$. Once we have a vector correlated with each $a_i^{\otimes 2}$, obtaining vectors close to the a_i is straightforward. We will begin by proving this theorem, and defer the algorithmic details to section [Section 5.4](#).

The rest of this section is organized as follows. In [Section 5.1](#) we prove [Theorem 5.3](#) using two core facts: the Gaussian vector g is closer to some a_i than to any other with good probability, and the noise term $\sum_{i \neq j} \langle g, T(a_i \otimes a_j) \rangle \cdot (a_i \otimes a_j)(a_i \otimes a_j)^\top$ is bounded in spectral norm. In [Section 5.2](#) we prove the first of these two facts, and in [Section 5.3](#) we prove the second. In [Section 5.4](#), we give the full details of our tensor decomposition algorithm, then prove [Theorem 5.2](#) using [Theorem 5.3](#). Finally, [Section C](#) contains proofs of elementary or long-winded lemmas we use along the way.

5.1 Proof of [Theorem 5.3](#)

The strategy to prove [Theorem 5.3](#) is to decompose the matrix M into two parts $M = M_{\text{diag}} + M_{\text{cross}}$, one formed by diagonal terms $M_{\text{diag}} = \sum_{i \in [n]} \langle g, T(a_i \otimes a_i) \rangle \cdot (a_i \otimes a_i)(a_i \otimes a_i)^\top$ and one formed by cross terms $M_{\text{cross}} = \sum_{i \neq j} \langle g, T(a_i \otimes a_j) \rangle \cdot (a_i \otimes a_j)(a_i \otimes a_j)^\top$. We will use the fact that the top eigenvector M_{diag} is likely to be correlated with one of the vectors $a_j^{\otimes 2}$, and also the fact that the spectral gap of M_{diag} is noticeable.

The following two propositions capture the relevant facts about the spectra of M_{diag} and M_{cross} , and will be proven in [Section 5.2](#) and [Section 5.3](#).

Proposition 5.4 (Spectral gap of diagonal terms). Let $R = \sqrt{2} \cdot ((\mathbb{E}(xx^\top)^{\otimes 2})^+)^{1/2}$ for $x \sim \mathcal{N}(0, \text{Id}_d)$. Let a_1, \dots, a_n be independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ with $d \leq n \leq d^{2-\Omega(1)}$ and let $g \sim \mathcal{N}(0, \text{Id}_d)$ be independent of all the others. Let $T := \sum_{i \in [n]} a_i(a_i \otimes a_i)^\top$. Suppose $M_{\text{diag}} = \sum_{i \in [n]} \langle g, T a_i^{\otimes 2} \rangle \cdot (a_i a_i^\top)^{\otimes 2}$. Let also v_j be such that $v_j v_j^\top = \langle g, T a_j^{\otimes 2} \rangle \cdot (a_j a_j^\top)^{\otimes 2}$. Then, with probability $1 - o(1)$ over a_1, \dots, a_n , for each $\varepsilon > \text{polylog } d / \sqrt{d}$ and each $j \in [n]$, the event

$$E_{j, \varepsilon} \stackrel{\text{def}}{=} \left\{ \left\| RM_{\text{diag}} R - \varepsilon \cdot R v_j v_j^\top R \right\| \leq \left\| RM_{\text{diag}} R \right\| - (\varepsilon - \tilde{O}(\sqrt{n}/d)) \cdot \left\| R v_j v_j^\top R \right\| \right\}$$

has probability at least $\tilde{\Omega}(1/n^{1+O(\varepsilon)})$ over the choice of g .

Second, we show that when $n \ll d^{4/3}$ the spectral norm of M_{cross} is negligible compared to this spectral gap.

Proposition 5.5 (Bound on crossterms). *Let a_1, \dots, a_n be independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$, and let g be a random vector from $\mathcal{N}(0, \text{Id}_d)$. Let $T := \sum_{i \in [n]} a_i (a_i \otimes a_i)^\top$. Let $M_{\text{cross}} := \sum_{i \neq j \in [n]} \langle g, T(a_i \otimes a_j) \rangle a_i a_i^\top \otimes a_j a_j^\top$. Suppose $n \geq d$. Then with w.ov.p.,*

$$\|M_{\text{cross}}\| \leq \tilde{O}\left(\frac{n^3}{d^4}\right)^{1/2}.$$

Using these two propositions we will conclude that the top eigenvector of RMR is likely to be correlated with one of the vectors $a_j^{\otimes 2}$. We also need two simple concentration bounds; we defer the proof to the appendix.

Lemma 5.6. *Let a_1, \dots, a_n be independently sampled vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$, and let g be sampled from $\mathcal{N}(0, \text{Id}_d)$. Let $T = \sum_i a_i (a_i \otimes a_i)^\top$. Then with overwhelming probability, for every $j \in [n]$,*

$$|\langle g, T(a_j \otimes a_j) \rangle - \langle g, a_j \rangle \|a_j\|^4| \leq \tilde{O}\left(\frac{\sqrt{n}}{d}\right).$$

Fact 5.7 (Simple version of [Fact C.1](#)). *Let $x, y \sim \mathcal{N}(0, \frac{1}{d} \text{Id})$. With overwhelming probability, $|1 - \|x\|^2| \leq \tilde{O}(1/\sqrt{d})$ and $\langle x, y \rangle^2 = \tilde{O}(1/d)$.*

As a last technical tool we will need a simple claim about the fourth moment matrix of the multivariate Gaussian:

Fact 5.8 (simple version of [Fact C.4](#)). *Let $\Sigma = \mathbb{E}_{x \sim \mathcal{N}(0, \text{Id}_d)} (xx^\top)^{\otimes 2}$ and let $R = \sqrt{2}(\Sigma^+)^{1/2}$. Then $\|R\| = 1$, and for any $v \in \mathbb{R}^d$,*

$$\|R(v \otimes v)\|_2^2 = \left(1 - \frac{1}{d+2}\right) \cdot \|v\|^4.$$

We are prepared prove [Theorem 5.3](#).

Proof of [Theorem 5.3](#). Let $d \leq n \leq d^{4/3}/(\text{polylog } d)$ for some polylog d to be chosen later. Let a_1, \dots, a_n be independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ and let $g \sim \mathcal{N}(0, \text{Id}_d)$ be independent of the others. Let

$$M_{\text{diag}} := \sum_{i \in [n]} \langle g, T(a_i \otimes a_i) \rangle \cdot (a_i a_i^\top)^{\otimes 2} \quad \text{and} \quad M_{\text{cross}} := \sum_{i \neq j \in [n]} \langle g, T(a_i \otimes a_j) \rangle \cdot a_i a_i^\top \otimes a_j a_j^\top.$$

Note that $M := M_{\text{diag}} + M_{\text{cross}}$.

[Proposition 5.5](#) implies that

$$\mathbb{P}\left\{\|M_{\text{cross}}\| \leq \tilde{O}(n^{3/2}/d^2)\right\} \geq 1 - d^{-\omega(1)}. \quad (5.1)$$

Recall that $\Sigma = \mathbb{E}_{x \sim \mathcal{N}(0, \text{Id}_d)} (xx^\top)^{\otimes 2}$ and $R = \sqrt{2} \cdot (\Sigma^+)^{1/2}$. By [Proposition 5.4](#), with probability $1 - o(1)$ over the choice of a_1, \dots, a_n , each of the following events $E_{j,\varepsilon}$ for $j \in [n]$ and $\varepsilon > \text{polylog}(d)/\sqrt{d}$ has probability at least $\tilde{\Omega}(1/n^{1+O(\varepsilon)})$ over the choice of g :

$$E_{j,\varepsilon}^0: \quad \left\|R \left(M_{\text{diag}} - \varepsilon \langle g, T a_j^{\otimes 2} \rangle (a_j a_j^\top)^{\otimes 2}\right) R\right\|$$

$$\leq \|RM_{\text{diag}}R\| - (\varepsilon - \tilde{O}(\frac{n^{1/2}}{d})) \cdot |\langle g, Ta_j^{\otimes 2} \rangle| \cdot \|Ra_j^{\otimes 2}\|^2.$$

Together with (5.1), with probability $1 - o(1)$ over the choice of a_1, \dots, a_n , each of the following events $E_{j,\varepsilon}^*$ has probability at least $\tilde{\Omega}(1/n^{1+O(\varepsilon)}) - d^{-\omega(1)} \geq \tilde{\Omega}(1/n^{1+O(\varepsilon)})$ over the choice of g ,

$$\begin{aligned} E_{j,\varepsilon}^* : \quad & \left\| R \left(M - \varepsilon \langle g, Ta_j^{\otimes 2} \rangle (a_j a_j^\top)^{\otimes 2} \right) R \right\| \\ & \leq \|R \cdot M \cdot R\| - (\varepsilon - \tilde{O}(\frac{n^{1/2}}{d})) \cdot |\langle g, Ta_j^{\otimes 2} \rangle| \cdot \|Ra_j^{\otimes 2}\|^2 + \tilde{O}(n^{3/2}/d^2) \end{aligned}$$

Here, we used that $M = M_{\text{diag}} + M_{\text{cross}}$ and that $\|R \cdot M_{\text{cross}} \cdot R\| \leq \|M_{\text{cross}}\|$ as $\|R\| \leq 1$ (Fact 5.8).

By standard reasoning about the top eigenvector of a matrix with a spectral gap (recorded in Lemma A.3), the event $E_{j,\varepsilon}^*$ implies that the top eigenvector $u \in \mathbb{R}^{d^2}$ of $R \cdot M \cdot R$ satisfies

$$\left\langle u, \frac{Ra_j^{\otimes 2}}{\|Ra_j^{\otimes 2}\|} \right\rangle^2 \geq 1 - \frac{\tilde{O}(\sqrt{n}/d)}{\varepsilon \|Ra_j^{\otimes 2}\|^2} - \frac{\tilde{O}(n^{3/2}/d^2)}{\varepsilon \|Ra_j^{\otimes 2}\|^2 |\langle g, Ta_j^{\otimes 2} \rangle|}.$$

Since $\|Ra_j^{\otimes 2}\|^2 \geq \Omega(\|a_j\|^4)$ (by Fact 5.8), and since $\|a_j\| \geq 1 - \tilde{O}(1/\sqrt{d})$ (by Fact 5.7),

$$\geq 1 - \tilde{O}\left(\frac{\sqrt{n}}{\varepsilon d}\right) - \frac{\tilde{O}(n^{3/2}/d^2)}{\varepsilon \cdot |\langle g, Ta_j^{\otimes 2} \rangle|}$$

Now, by Lemma 5.6 we have that for all $j \in [n]$, $|\langle g, Ta_j^{\otimes 2} \rangle - \langle g, a_j \rangle \|a_j\|^4| \leq \tilde{O}(\sqrt{n}/d)$ with probability $1 - n^{-\omega(1)}$. By standard concentration (see Fact C.1 for a proof) $|\langle g, a_j \rangle \|a_j\|^4 - 1| \leq \tilde{O}(1/\sqrt{d})$ for all $j \in [n]$ with probability $1 - n^{-\omega(1)}$. Therefore with overwhelming probability, the final term is bounded by $\tilde{O}(n^{3/2}/\varepsilon d^2)$. A union bound now gives the desired conclusion.

Finally, we give a bound on the spectral gap. We note that the second eigenvector w has $\langle u, w \rangle = 0$, and therefore

$$\left\langle w, \frac{Ra_j^{\otimes 2}}{\|Ra_j^{\otimes 2}\|} \right\rangle = \left\langle w, \frac{Ra_j^{\otimes 2}}{\|Ra_j^{\otimes 2}\|} - u \right\rangle \leq \left\| \frac{Ra_j^{\otimes 2}}{\|Ra_j^{\otimes 2}\|} - u \right\| \leq \tilde{O}(n^{3/2}/\varepsilon d^2).$$

Thus, using our above bound on $\|R(M - \varepsilon \langle g, Ta_j^{\otimes 2} \rangle (a_j a_j^\top)^{\otimes 2})R\|$ and the concentration bounds we have already applied for $\|a_j\|$, $\langle g, Ta_j^{\otimes 2} \rangle$, and $\|Ra_j^{\otimes 2}\|$, we have that

$$\begin{aligned} \lambda_2(RMR) &= w^\top RMRw \\ &= w^\top R \left(M - \varepsilon \langle g, Ta_j^{\otimes 2} \rangle \cdot (a_j a_j^\top)^{\otimes 2} \right) R w + \varepsilon \langle g, Ta_j^{\otimes 2} \rangle \langle w, Ra_j^{\otimes 2} \rangle^2 \\ &\leq \left\| R \left(M - \varepsilon \langle g, Ta_j^{\otimes 2} \rangle \cdot (a_j a_j^\top)^{\otimes 2} \right) R \right\| + \tilde{O}(n^{3/2}/\varepsilon d^2) \\ &\leq 1 - \tilde{O}(\varepsilon) + \tilde{O}(n^{3/2}/\varepsilon d^2). \end{aligned}$$

We conclude that the above events also imply that $\lambda_2(RMR)/\lambda_1(RMR) \leq 1 - O(\varepsilon)$. \square

5.2 Spectral gap for diagonal terms: proof of Proposition 5.4

We now prove that the signal matrix, when preconditioned by R , has a noticeable spectral gap:

Proposition (Restatement of [Proposition 5.4](#)). Let $R = \sqrt{2} \cdot ((\mathbb{E}(xx^\top)^{\otimes 2})^+)^{1/2}$ for $x \sim \mathcal{N}(0, \text{Id}_d)$. Let a_1, \dots, a_n be independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ with $d \leq n \leq d^{2-\Omega(1)}$ and let $g \sim \mathcal{N}(0, \text{Id}_d)$ be independent of all the others. Let $T := \sum_{i \in [n]} a_i (a_i \otimes a_i)^\top$. Suppose $M_{\text{diag}} = \sum_{i \in [n]} \langle g, T a_i^{\otimes 2} \rangle \cdot (a_i a_i^\top)^{\otimes 2}$. Let also v_j be such that $v_j v_j^\top = \langle g, T a_j^{\otimes 2} \rangle \cdot (a_j a_j^\top)^{\otimes 2}$. Then, with probability $1 - o(1)$ over a_1, \dots, a_n , for each $\varepsilon > \text{polylog } d / \sqrt{d}$ and each $j \in [n]$, the event

$$E_{j,\varepsilon} \stackrel{\text{def}}{=} \left\{ \left\| R M_{\text{diag}} R - \varepsilon \cdot R v_j v_j^\top R \right\| \leq \left\| R M_{\text{diag}} R \right\| - (\varepsilon - \tilde{O}(\sqrt{n}/d)) \cdot \left\| R v_j v_j^\top R \right\| \right\}$$

has probability at least $\tilde{\Omega}(1/n^{1+O(\varepsilon)})$ over the choice of g .

The proof has two parts. First we show that for $a_1, \dots, a_n \sim \mathcal{N}(0, \text{Id}_d)$ the matrix $P := \sum_{i \in [n]} (a_i a_i^\top)^{\otimes 2}$ has tightly bounded spectral norm when preconditioned with R : more precisely, that $\|RPR\| \leq 1 + \tilde{O}(n/d^{3/2})$.

Lemma 5.9. Let $a_1, \dots, a_n \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ be independent random vectors with $d \leq n$. Let $R := \sqrt{2} \cdot ((\mathbb{E}(aa^\top)^{\otimes 2})^+)^{1/2}$ for $a \sim \mathcal{N}(0, \text{Id}_d)$. For $S \subseteq [n]$, let $P_S = \sum_{i \in S} (a_i a_i^\top)^{\otimes 2}$ and let Π_S be the projector into the subspace spanned by $\{R a_i^{\otimes 2} \mid i \in S\}$. Then, with probability $1 - o(1)$ over the choice of a_1, \dots, a_n ,

$$\forall S \subseteq [n]. \quad (1 - \tilde{O}(n/d^{3/2})) \cdot \Pi_S \leq R P_S R \leq (1 + \tilde{O}(n/d^{3/2})) \cdot \Pi_S.$$

Remark 5.10. In [\[GM15, Lemma 5\]](#) a similar lemma to this one is proved in the context of the SoS proof system. However, since Ge and Ma leverage the full power of the SoS algorithm their proof goes via a spectral bound on a different (but related) matrix. Since our algorithm avoids solving an SDP we need a bound on this matrix in particular.

The proof of [Lemma 5.9](#) proceeds by standard spectral concentration for tall matrices with independent columns (here the columns are $R a_i^{\otimes 2}$). The arc of the proof is straightforward but it involves some bookkeeping; we have deferred it to [Section C.0.2](#).

We also need the following lemma on the concentration of some scalar random variables involving R ; the proof is straightforward by finding the eigenbasis of R and applying standard concentration, and it is deferred to the appendix.

Lemma 5.11. Let $a_1, \dots, a_n \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$. Let Σ, R be as in [Fact 5.8](#). Let $u_i = a_i \otimes a_i$. With overwhelming probability, every $j \in [n]$ satisfies $\sum_{i \neq j} \langle u_j, R^2 u_i \rangle^2 = \tilde{O}(n/d^2)$ and $|1 - \|R u_j\|^2| \leq \tilde{O}(1/\sqrt{d})$.

The next lemma is the linchpin of the proof of [Proposition 5.4](#): one of the inner products $\langle g, T a_j^{\otimes 2} \rangle$, is likely to be a $\approx (1 + 1/\log(n))$ -factor larger than the maximum of the inner products $\langle g, T a_i^{\otimes 2} \rangle$ over $i \neq j$. Together with standard linear algebra these imply that the matrix $M_{\text{diag}} = \sum_{i \in [n]} \langle g, T a_i^{\otimes 2} \rangle (a_i a_i^\top)^{\otimes 2}$ has top eigenvector highly correlated or anticorrelated with some a_i .

Lemma 5.12. Let $a_1, \dots, a_n \in \mathbb{R}^d$ be independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$, and let g be a random vector from $\mathcal{N}(0, \text{Id}_d)$. Let $T = \sum_{i \in [n]} a_i (a_i \otimes a_i)^\top$. Let $\varepsilon > 0$ and $j \in [n]$. Then with overwhelming probability over a_1, \dots, a_n , the following event $\hat{E}_{j,\varepsilon}$ has probability $1/n^{1+O(\varepsilon)+\tilde{O}(1/\sqrt{d})}$ over the choice of g ,

$$\hat{E}_{j,\varepsilon} = \left\{ \langle g, T a_j^{\otimes 2} \rangle \geq (1 + \varepsilon)(1 - \tilde{O}(1/\sqrt{d})) \cdot \max_{i \neq j} |\langle g, T a_i^{\otimes 2} \rangle| \right\}.$$

Now we can prove [Proposition 5.4](#).

Proof of Proposition 5.4. Let $u_i := a_i^{\otimes 2}$. Fix $j \in [n]$. We begin by showing a lower bound on the spectral norm $\|RM_{\text{diag}}R\|$.

$$\begin{aligned} \|RM_{\text{diag}}R\| &= \max_{\|v\|=1} |\langle v, RM_{\text{diag}}Rv \rangle| \\ &\geq \frac{\langle Ru_j, (RM_{\text{diag}}R)Ru_j \rangle}{\|Ru_j\|^2} \\ &= \frac{1}{\|Ru_j\|^2} \left(\langle g, Tu_j \rangle \|Ru_j\|^4 + \langle Ru_j, \sum_{i \neq j} \langle g, T_i \rangle Ru_i u_i^\top R \cdot Ru_j \rangle \right) \end{aligned}$$

From Lemma 5.12, the random vector g is closer to Tu_j than to all Tu_i for $i \neq j, i \in [n]$ with reasonable probability. More concretely there is some polylog d so that as long as $\varepsilon > \text{polylog } d / \sqrt{d}$ there is some $\alpha = \Theta(\varepsilon)$ with $1 - \varepsilon = 1/[(1 + \alpha)(1 - \tilde{O}(d^{-1/2}))]$ so that with w.ov.p. over a_1, \dots, a_n the following event (a direct consequence of $\hat{E}_{j,\varepsilon}$) has probability $\tilde{\Omega}(1/n^{1+O(\alpha)+\tilde{O}(d^{-1/2})}) = \tilde{\Omega}(1/n^{1+O(\varepsilon)})$ over g :

$$-(1 - \varepsilon) |\langle g, Tu_j \rangle| \cdot \left(\sum_{i \neq j} Ru_i u_i^\top R \right) \leq \sum_{i \neq j} \langle g, Tu_i \rangle \cdot Ru_i u_i^\top R \leq (1 - \varepsilon) |\langle g, Tu_j \rangle| \cdot \left(\sum_{i \neq j} Ru_i u_i^\top R \right). \quad (5.2)$$

When (5.2) occurs,

$$\begin{aligned} \|RM_{\text{diag}}R\| &\geq \frac{1}{\|Ru_j\|^2} \left(|\langle g, Tu_j \rangle| \|Ru_j\|^4 - (1 - \varepsilon) |\langle g, Tu_j \rangle| \left\langle Ru_j, \sum_{i \neq j} Ru_i u_i^\top R \cdot Ru_j \right\rangle \right) \\ &= \frac{|\langle g, Tu_j \rangle|}{\|Ru_j\|^2} \left(\|Ru_j\|^4 - (1 - \varepsilon) \sum_{i \neq j} \langle u_j, R^2 u_i \rangle^2 \right) \\ &\geq \frac{|\langle g, Tu_j \rangle| \left(1 - \tilde{O}(1/\sqrt{d}) - (1 - \varepsilon) \tilde{O}(n/d^2) \right)}{1 + \tilde{O}(1/\sqrt{d})} \quad \text{w.ov.p. over } a_1, \dots, a_n \text{ (Lemma 5.11)} \\ &\geq |\langle g, Tu_j \rangle| \cdot (1 - \eta_{\text{norm}}), \end{aligned} \quad (5.3)$$

where we have chosen some $0 \leq \eta_{\text{norm}} \leq \tilde{O}(1/\sqrt{d}) + \tilde{O}(n/d^2)$ (since for any $x \in \mathbb{R}$, $(1+x)(1-x) \leq 1$).

Next we exhibit an upper bound on $\|RM_{\text{diag}}R - \varepsilon \langle g, Tu_j \rangle Ru_j u_j^\top R\|$. Again when (5.2) occurs,

$$\begin{aligned} &\|RM_{\text{diag}}R - \varepsilon \langle g, Tu_j \rangle Ru_j u_j^\top R\| \quad (5.4) \\ &= \left\| (1 - \varepsilon) \langle g, Tu_j \rangle Ru_j u_j^\top R + \sum_{i \neq j} \langle g, T_i \rangle Ru_i u_i^\top R \right\| \\ &\leq (1 - \varepsilon) |\langle g, Tu_j \rangle| \left\| \sum_{i \in [n]} Ru_i u_i^\top R \right\| \quad \text{when (5.2) occurs} \\ &\leq (1 - \varepsilon) |\langle g, Tu_j \rangle| (1 + \tilde{O}(n/d^{1.5})) \quad \text{w.p. } 1 - o(1) \text{ over } a_1, \dots, a_n \text{ by Lemma 5.9} \\ &\leq (1 - \varepsilon) |\langle g, Tu_j \rangle| (1 + \eta_{\text{gap}}) \end{aligned} \quad (5.5)$$

where we have chosen some $0 \leq \eta_{\text{gap}} \leq \tilde{O}(n/d^{1.5})$.

Putting together (5.3) and (5.5) with our bounds on η_{norm} and η_{gap} and recalling the conditions on (5.2), we have shown that

$$\begin{aligned} \mathbb{P}_{a_1, \dots, a_n} \left\{ \mathbb{P}_g \left\{ \|RM_{\text{diag}} R - \varepsilon \langle g, Tu_j \rangle Ru_j u_j^\top R\| \leq \|RM_{\text{diag}} R\| - (\varepsilon - \tilde{O}(\sqrt{n}/d)) \cdot |\langle g, Tu_j \rangle| \cdot \|Ru_j\|^2 \right\} \right. \\ \left. \geq \tilde{\Omega}(1/n^{1+O(\varepsilon)}) \right\} \geq 1 - o(1). \end{aligned}$$

This concludes the argument. \square

We now turn to proving that with reasonable probability, g is closer to some $Ta_j^{\otimes 2}$ than all others.

Proof of Lemma 5.12. To avoid proliferation of indices, without loss of generality fix $j = 1$. We begin by expanding the expression $\langle g, Ta_i^{\otimes 2} \rangle$,

$$\langle g, Ta_i^{\otimes 2} \rangle = \sum_{\ell \in [n]} \langle g, a_\ell \rangle \langle a_\ell, a_i \rangle^2 = \|a_i\|^4 \langle g, a_i \rangle + \sum_{\ell \neq i} \langle g, a_\ell \rangle \langle a_\ell, a_i \rangle^2.$$

The latter sum is bounded by

$$\left| \sum_{\ell \neq i} \langle g, a_\ell \rangle \langle a_\ell, a_i \rangle^2 \right| \leq \tilde{O} \left(\frac{\sqrt{n}}{d} \right),$$

with overwhelming probability for all i and choices of g ; this follows from a Bernstein bound, given in Lemma 5.6.

For ease of notation, let $\hat{a}_i \stackrel{\text{def}}{=} a_i / \|a_i\|_2$. We conclude from Fact 5.7 that with overwhelming probability, $1 - \tilde{O}(1/\sqrt{d}) \leq \|a_i\|_2 \leq 1 + \tilde{O}(1/\sqrt{d})$ for all $i \in [n]$. Thus $\|a_i\|_2$ is roughly equal for all i , and we may direct our attention to $\langle g, \hat{a}_i \rangle$.

Let \mathcal{G}_1 be the event that $\sqrt{2\alpha} \log^{1/2} n \leq |\langle g, \hat{a}_1 \rangle| \leq d^{1/4}$ for some $\alpha \leq d^{1/2 - \Omega(1)}$ to be chosen later. We note that $\langle g, \hat{a}_1 \rangle$ is distributed as a standard gaussian, and that g is independent of a_1, \dots, a_n . Thus, we can use standard tail estimates on univariate Gaussians (Lemma A.4) to conclude that

$$\mathbb{P} \left(|\langle g, \hat{a}_1 \rangle| \geq \sqrt{2\alpha} \log^{1/2} n \right) = \tilde{\Theta}(n^{-\alpha}) \quad \text{and} \quad \mathbb{P} \left(|\langle g, \hat{a}_1 \rangle| \geq d^{1/4} \right) = \Theta \left(\frac{\exp(-\sqrt{d}/2)}{d^{1/4}} \right).$$

So by a union bound, $\mathbb{P}(\mathcal{G}_1) \geq \tilde{\Omega}(n^{-\alpha}) - O(e^{-d^{1/2}/3}) = \tilde{\Omega}(n^{-\alpha})$.

Now, we must obtain an estimate for the probability that all other inner products with g are small. Let $\mathcal{G}_{i>1}$ be the event that $|\langle g, \hat{a}_i \rangle| \leq \sqrt{(2+\rho)} \log^{1/2} n$ for all $i \in [n], i > 1$ and for some ρ to be chosen later. We will show that conditioned on \mathcal{G}_1 , $\mathcal{G}_{i>1}$ occurs with probability $1 - O(n^{1-(2+\rho)/2})$. Define $g_1 := \langle g, \hat{a}_1 \rangle \hat{a}_1$ to be the component of g parallel to a_1 , let $g_\perp := g - g_1$ be the component of g orthogonal to \hat{a}_1 , and similarly let $\hat{a}_2^\perp, \dots, \hat{a}_n^\perp$ be the components of $\hat{a}_2, \dots, \hat{a}_n$ orthogonal to a_1 . Because g_\perp is independent of g_1 , even conditioned on \mathcal{G}_1 we may apply the standard tail bound for univariate Gaussians (Lemma A.4), concluding that for all $i > 1$,

$$\mathbb{P} \left(|\langle g_\perp, \hat{a}_i \rangle| \geq \sqrt{(2+\rho)} \log^{1/2} n \mid \mathcal{G}_1 \right) = \tilde{\Theta}(n^{-(2+\rho)/2}).$$

Thus, a union bound over $i \neq 1$ allows us to conclude that conditioned on \mathcal{G}_1 , with probability $1 - \tilde{O}(n^{-\rho/2})$ every $i \in [n]$ with $i > 1$ has $|\langle g_\perp, \hat{a}_i^\perp \rangle| \leq \sqrt{(2+\rho)} \log^{1/2} n$.

On the other hand, let $\hat{a}_2^\parallel, \dots, \hat{a}_n^\parallel$ be the components of the \hat{a}_i parallel to \hat{a}_1 . We compute the projection of \hat{a}_i onto \hat{a}_1 . With overwhelming probability,

$$\begin{aligned} \langle \hat{a}_1, \hat{a}_i \rangle &= \frac{\langle a_1, a_i \rangle}{\|a_1\|_2 \cdot \|a_i\|_2} \\ &= (1 \pm \tilde{O}(1/\sqrt{d})) \cdot \langle a_1, a_i \rangle \quad \text{w.ov.p. by } \|a_i\|, \|a_1\| = 1 \pm \tilde{O}(1/\sqrt{d}) \text{ (Fact 5.7)} \\ &= (1 \pm \tilde{O}(1/\sqrt{d})) \cdot \tilde{O}(1/\sqrt{d}) \quad \text{w.ov.p. by } \langle a_1, a_i \rangle = \tilde{O}(1/\sqrt{d}) \text{ (Fact 5.7),} \end{aligned}$$

Thus w.ov.p.,

$$\langle g_1, \hat{a}_i^\parallel \rangle = \langle g, \hat{a}_1 \rangle \cdot \langle \hat{a}_1, \hat{a}_i \rangle \leq \langle g, \hat{a}_1 \rangle \cdot \tilde{O}(1/\sqrt{d}),$$

for all $i \in [n]$. Now we can analyze $\mathcal{G}_{i>1}$. Taking a union bound over the overwhelmingly probable events (including $\|a_i\| \leq 1 + \tilde{O}(1/\sqrt{d})$) and the event that $\langle g_\perp, a_i \rangle$ is small for all i , we have that with probability $1 - \tilde{O}(n^{-\rho/2})$, for every $i \in [n]$ with $i > 1$,

$$\begin{aligned} |\langle g, \hat{a}_i \rangle| &\leq |\langle g_\perp, \hat{a}_i \rangle| + |\langle g_1, \hat{a}_i \rangle| \\ &\leq \sqrt{(2 + \rho) \log^{1/2} n} + \tilde{O}(1/\sqrt{d}) \cdot \langle g, \hat{a}_1 \rangle \\ &\leq \sqrt{(2 + \rho) \log^{1/2} n} + \tilde{O}(1/d^{1/4}). \end{aligned}$$

We conclude that

$$\begin{aligned} \mathbb{P}(\mathcal{G}_1, \mathcal{G}_{i>1}) &= \mathbb{P}(\mathcal{G}_{i>1} | \mathcal{G}_1) \cdot \mathbb{P}(\mathcal{G}_1) \\ &\geq (1 - O(n^{-\rho/2})) \cdot \tilde{\Omega}(n^{-\alpha}) \end{aligned}$$

Setting $\rho = 2 \frac{\log \log n}{\log n}$ and $\alpha = (1 + \varepsilon)^2 (1 + \log \log n / \log n + \tilde{O}(1/\sqrt{d}))$, the conclusion follows. \square

5.3 Bound for cross terms: proof of Proposition 5.5

We proceed to the bound on the cross terms M_{cross} .

Proposition (Restatement of Proposition 5.5). *Let a_1, \dots, a_n be independent random vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$, and let g be a random vector from $\mathcal{N}(0, \text{Id}_d)$. Let $T := \sum_{i \in [n]} a_i (a_i \otimes a_i)^\top$. Let $M_{\text{cross}} := \sum_{i \neq j \in [n]} \langle g, T(a_i \otimes a_j) \rangle a_i a_i^\top \otimes a_j a_j^\top$. Suppose $n \geq d$. Then with w.ov.p.,*

$$\|M_{\text{cross}}\| \leq \tilde{O} \left(\frac{n^3}{d^4} \right)^{1/2}.$$

The proof will use two iterations of Matrix Rademacher bounds. The first step will be to employ a classical decoupling inequality that has previously been used in a tensor decomposition context [GM15].

Theorem 5.13 (Special Case of Theorem 1 in [dlPMS95]). *Let $\{s_i\}, \{t_i\}$ be independent iid sequences of random signs. Let $\{M_{ij}\}$ be a family of matrices. There is a universal constant C so that for every $t > 0$,*

$$\mathbb{P} \left(\left\| \sum_{i \neq j} s_i s_j M_{ij} \right\|_{\text{op}} > t \right) \leq C \cdot \mathbb{P} \left(C \left\| \sum_{i \neq j} s_i t_j M_{ij} \right\|_{\text{op}} > t \right).$$

Once the simplified cross terms are decoupled, we can use a matrix Rademacher bound on one set of signs.

Theorem 5.14 (Adapted from Theorem 4.1.1 in [Tro12]¹⁰). Consider a finite sequence $\{M_i\}$ of fixed $m \times m$ Hermitian matrices. Let s_i be a sequence of independent sign variables. Let $\sigma^2 := \|\sum_i M_i^2\|$. Then for every $t \geq 0$,

$$\mathbf{P}\left(\left\|\sum_i s_i M_i\right\|_{op} \geq t\right) \leq 2m \cdot e^{-t^2/2\sigma^2}.$$

Also,

$$\mathbb{E}\left\|\sum_i s_i M_i\right\| \leq \sqrt{8\sigma^2 \log d}.$$

Corollary 5.15. Let s_1, \dots, s_n be independent signs in $\{-1, 1\}$. Let A_1, \dots, A_n and B_1, \dots, B_n be Hermitian matrices. Then w.o.v.p.,

$$\left\|\sum_i s_i \cdot A_i \otimes B_i\right\| \leq \tilde{O}\left(\max_i \|B_i\| \cdot \left\|\sum_i A_i^2\right\|^{1/2}\right).$$

Proof. We use a matrix Rademacher bound and standard manipulations:

$$\begin{aligned} \left\|\sum_i s_i \cdot A_i \otimes B_i\right\| &\stackrel{\text{w.o.v.p.}}{\leq} \tilde{O}\left(\left\|\sum_i A_i^2 \otimes B_i^2\right\|^{1/2}\right) \\ &\leq \tilde{O}\left(\left\|\sum_i \|B_i\|^2 \cdot (A_i^2 \otimes \text{Id})\right\|^{1/2}\right) \quad \text{since } A_i^2 \text{ is PSD for all } i \\ &\leq \tilde{O}\left(\max_i \|B_i\|^2 \cdot \left\|\sum_i A_i^2\right\|^{1/2}\right) \quad \text{since } A_i^2 \otimes \text{Id} \text{ is PSD for all } i. \quad \square \end{aligned}$$

We also need a few further concentration bounds on matrices which will come up as parts of M_{cross} . These can be proved by standard inequalities for sums of independent matrices.

Lemma 5.16 (Restatement of [Fact C.2](#) and [Lemma C.3](#)). Let a_1, \dots, a_n be independent from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ with $n \geq d$ polylog(d). With overwhelming probability, $\tilde{\Omega}(n/d) \cdot \text{Id} \leq \sum_{i \in [n]} a_i a_i^\top \leq \tilde{O}(n/d) \cdot \text{Id}$. Additionally, if $g \sim \mathcal{N}(0, \text{Id}_d)$ is independent of the rest, for every $j \in [n]$ w.o.v.p.

$$\left\|\sum_{\substack{i \in [n] \\ i \neq j}} \langle g, a_i \rangle \|a_i\|^2 \langle a_i, a_j \rangle \cdot a_i a_i^\top\right\| \leq \tilde{O}(n/d^2)^{1/2}.$$

Proof of [Proposition 5.5](#). We expand M_{cross} :

$$M_{\text{cross}} = \sum_{i \neq j} \langle g, T(a_i \otimes a_j) \rangle \cdot a_i a_i^\top \otimes a_j a_j^\top$$

¹⁰We remark that Tropp's bound is phrased in terms of $\lambda_{\max} \sum_i s_i M_i$. Since $\lambda_{\max} \sum_i s_i M_i = \lambda_{\min} \sum_i -s_i M_i$, and the distribution of $s_i M_i$ is negation-invariant, the result we state here follows from an easy union bound.

$$= \sum_{i \neq j} \left(\sum_{\ell \in [n]} \langle a_\ell, a_i \rangle \langle a_\ell, a_j \rangle \langle g, a_\ell \rangle \right) \cdot a_i a_i^\top \otimes a_j a_j^\top.$$

Since the joint distribution of (a_1, \dots, a_n) is identical to that of $(s_1 a_1, \dots, s_n a_n)$, this is distributed identically to

$$M'_{\text{cross}} = \sum_{\ell \in [n]} \sum_{i \neq j} s_i s_j s_\ell \langle g, a_\ell \rangle \langle a_\ell, a_i \rangle \langle a_\ell, a_j \rangle \cdot a_i a_i^\top \otimes a_j a_j^\top,$$

(where we have also swapped the sums over ℓ and $i \neq j$). We split M'_{cross} into M_{diff} , for which $i \neq \ell$ and $j \neq \ell$, and M_{same} , for which $\ell = i$ or $\ell = j$, and bound the norm of each of these sums separately. We begin with M_{same} .

$$M_{\text{same}} \stackrel{\text{def}}{=} \sum_{i \neq j} s_i^2 s_j \langle g, a_i \rangle \langle a_i, a_i \rangle \langle a_i, a_j \rangle \cdot a_i a_i^\top \otimes a_j a_j^\top + \sum_{i \neq j} s_j^2 s_i \langle g, a_j \rangle \langle a_j, a_j \rangle \langle a_i, a_j \rangle \cdot a_i a_i^\top \otimes a_j a_j^\top.$$

By a union bound and an application of the triangle inequality it will be enough to show that just one of these two sums is $\tilde{O}(n^3/d^4)^{1/2}$ w.ov.p.. We rewrite the left-hand one:

$$\sum_{i \neq j} s_i^2 s_j \langle g, a_i \rangle \langle a_i, a_i \rangle \langle a_i, a_j \rangle \cdot a_i a_i^\top \otimes a_j a_j^\top = \sum_{j \in [n]} s_j a_j a_j^\top \otimes \left(\sum_{i \neq j} \langle g, a_i \rangle \|a_i\|^2 \langle a_i, a_j \rangle \cdot a_i a_i^\top \right).$$

Define

$$M_j \stackrel{\text{def}}{=} \sum_{i \neq j} \langle g, a_i \rangle \|a_i\|^2 \langle a_i, a_j \rangle \cdot a_i a_i^\top$$

so that now we need to bound $\sum_{j \in [n]} s_j a_j a_j^\top \otimes M_j$. By [Corollary 5.15](#),

$$\begin{aligned} \left\| \sum_{j \in [n]} s_j a_j a_j^\top \otimes M_j \right\| &\stackrel{\text{w.ov.p.}}{\leq} \tilde{O}(\max_j \|M_j\|) \cdot \tilde{O} \left(\left\| \sum_{j \in [n]} \|a_j\|^2 a_j a_j^\top \right\|^{1/2} \right) \\ &\leq \tilde{O}(\max_j \|M_j\|) \cdot \max_j \|a_j\| \cdot \tilde{O} \left(\left\| \sum_{j \in [n]} a_j a_j^\top \right\|^{1/2} \right) \end{aligned}$$

In [Lemma 5.16](#), we bound $\max_j \|M_j\| \leq \tilde{O}(n/d^2)^{1/2}$ w.ov.p. using a matrix Bernstein inequality. Combining this bound with the concentration of $\|a_j\|$ around 1 ([Fact 5.7](#)), we obtain

$$\begin{aligned} &\stackrel{\text{w.ov.p.}}{\leq} \tilde{O}(n/d^2)^{1/2} \cdot \tilde{O}(n/d)^{1/2} \\ &= \tilde{O}(n/d^{1.5}). \end{aligned}$$

Having finished with M_{same} , we turn to M_{diff} .

$$\|M_{\text{diff}}\| = \left\| \sum_{\ell \neq i \neq j} s_i s_j s_\ell \langle g, a_\ell \rangle \langle a_\ell, a_i \rangle \langle a_\ell, a_j \rangle \cdot a_i a_i^\top \otimes a_j a_j^\top \right\|$$

$$= \left\| \sum_{\ell} s_{\ell} \langle g, a_{\ell} \rangle \left(\sum_{i \neq \ell} s_i \langle a_{\ell}, a_i \rangle a_i a_i^{\top} \otimes \left(\sum_{j \neq \ell, i} s_j \langle a_{\ell}, a_j \rangle a_j a_j^{\top} \right) \right) \right\|.$$

Letting t_1, \dots, t_n and r_1, \dots, r_n be independent uniformly random signs, by [Theorem 5.13](#), it will be enough to bound the spectral norm after replacing the second and third occurrences of s_i for t_i and r_i . To this end, we define

$$M'_{\text{diff}} \stackrel{\text{def}}{=} \sum_{\ell} s_{\ell} \langle g, a_{\ell} \rangle \left(\sum_{i \neq \ell} t_i \langle a_{\ell}, a_i \rangle a_i a_i^{\top} \otimes \left(\sum_{j \neq \ell, i} r_j \langle a_{\ell}, a_j \rangle a_j a_j^{\top} \right) \right).$$

Let

$$N_{\ell} \stackrel{\text{def}}{=} \sum_{i \neq \ell} t_i \langle a_{\ell}, a_i \rangle a_i a_i^{\top} \otimes \left(\sum_{j \neq \ell, i} r_j \langle a_{\ell}, a_j \rangle a_j a_j^{\top} \right)$$

so that we are to bound $\left\| \sum_{\ell \in [n]} s_{\ell} \langle g, a_{\ell} \rangle \cdot N_{\ell} \right\|$. By a matrix Rademacher bound and elementary manipulations,

$$\begin{aligned} \left\| \sum_{\ell \in [n]} s_{\ell} \langle g, a_{\ell} \rangle \cdot N_{\ell} \right\| &\stackrel{\text{w.ov.p.}}{\leq} \tilde{O} \left(\left\| \sum_{\ell \in [n]} \langle g, a_{\ell} \rangle^2 \cdot N_{\ell}^2 \right\| \right)^{1/2} \\ &\leq \tilde{O}(\sqrt{n}) \cdot \max_{\ell \in [n]} |\langle g, a_{\ell} \rangle| \cdot \max_{\ell \in [n]} \|N_{\ell}\| \\ &\stackrel{\text{w.ov.p.}}{\leq} \tilde{O}(\sqrt{n}) \cdot \max_{\ell \in [n]} \|N_{\ell}\| \quad \text{since } |\langle g, a_i \rangle| \leq \tilde{O}(1) \text{ (Fact 5.7)}. \end{aligned}$$

The rest of the proof is devoted to bounding $\|N_{\ell}\|$.

We start with [Corollary 5.15](#) to get

$$\|N_{\ell}\| \stackrel{\text{w.ov.p.}}{\leq} \tilde{O} \left(\left(\max_i \left\| \sum_{j \neq \ell, i} r_j \langle a_{\ell}, a_j \rangle \cdot a_j a_j^{\top} \right\| \right) \cdot \left\| \sum_{i \neq \ell} \langle a_{\ell}, a_i \rangle^2 \|a_i\|^2 \cdot a_i a_i^{\top} \right\|^{1/2} \right)$$

We use a matrix Rademacher bound for the left-hand matrix,

$$\begin{aligned} \left\| \sum_{j \neq \ell, i} r_j \langle a_{\ell}, a_j \rangle \cdot a_j a_j^{\top} \right\| &\stackrel{\text{w.ov.p.}}{\leq} \tilde{O} \left(\left\| \sum_{j \neq \ell, i} \langle a_{\ell}, a_j \rangle^2 \|a_j\|^2 \cdot a_j a_j^{\top} \right\| \right)^{1/2} \\ &\leq \tilde{O} \left(\max_{j \neq \ell} \langle a_{\ell}, a_j \rangle^2 \|a_j\|^2 \left\| \sum_j a_j a_j^{\top} \right\| \right)^{1/2} \\ &\stackrel{\text{w.ov.p.}}{\leq} \tilde{O} \left(\frac{\sqrt{n}}{d} \right), \end{aligned}$$

where we have used that $\langle a_{\ell}, a_i \rangle^2$ concentrates around $\frac{1}{d}$ ([Fact 5.7](#)), that $\|a_i\|^2$ concentrates around 1 ([Fact 5.7](#)), and that $\left\| \sum_i a_i a_i^{\top} \right\|$ concentrates around $\frac{n}{d}$ ([Lemma 5.16](#)) within logarithmic factors all with overwhelming probability.

For the right-hand matrix, we use the fact that the summands are PSD to conclude that

$$\left\| \sum_{i \neq \ell} \langle a_\ell, a_i \rangle^2 \|a_i\|^2 \cdot a_i a_i^\top \right\| \leq \max_{i \neq \ell} \langle a_\ell, a_i \rangle^2 \|a_i\|^2 \cdot \left\| \sum_{i \neq \ell} a_i a_i^\top \right\|$$

$$\stackrel{\text{w.ov.p.}}{\leq} \tilde{O}(1/d) \cdot \tilde{O}(n/d),$$

using the same concentration facts as earlier.

Putting these together, w.ov.p.

$$\|N_\ell\| \leq \tilde{O}(\sqrt{n}/d) \cdot \tilde{O}(\sqrt{n}/d) = \tilde{O}(n/d^2).$$

Now we are ready to make the final bound on M'_{diff} . With overwhelming probability,

$$\|M'_{\text{diff}}\| \leq \tilde{O}(\sqrt{n}) \cdot \max_{\ell \in [n]} \|N_\ell\| \leq \tilde{O}(n^3/d^4)^{1/2}$$

and hence by [Theorem 5.13](#), $\|M_{\text{diff}}\| \leq \tilde{O}(\sqrt{n}) \cdot \max_{\ell \in [n]} \|N_\ell\| \leq \tilde{O}(n^3/d^4)^{1/2}$ w.ov.p..

Finally, by triangle inequality and all our bounds thus far, w.ov.p.

$$\|M_{\text{cross}}\| \leq \|M_{\text{same}}\| + \|M_{\text{diff}}\| \leq \tilde{O}(n/d^{1.5}) + \tilde{O}(n^3/d^4)^{1/2} \leq \tilde{O}(n^3/d^4)^{1/2}. \quad \square$$

5.4 Full algorithm and proof of [Theorem 5.2](#)

In this subsection we give the full details of our tensor decomposition algorithm. As discussed above, the algorithm proceeds by constructing a random matrix from the input tensor, then computing and post-processing its top eigenvector.

Spectral Tensor Decomposition (One Attempt)

This is the main subroutine of our algorithm—we will run it $\tilde{O}(n)$ times to recover all of the components a_1, \dots, a_n .

Algorithm 5.17. Input: $\mathbf{T} = \sum_{i=1}^n a_i \otimes a_i \otimes a_i$. Goal: Recover a_i for some $i \in [n]$.

- Compute the matrix unfolding $T \in \mathbb{R}^{d^2 \times d}$ of \mathbf{T} . Then compute a 3-tensor $\mathbf{S} \in \mathbb{R}^{d^2 \times d^2 \times d^2}$ by starting with the 6-tensor $\mathbf{T} \otimes \mathbf{T}$, permuting indices, and flattening to a 3-tensor. Apply T in one mode of \mathbf{S} to obtain $\mathbf{M} \in \mathbb{R}^{d \otimes d^2 \otimes d^2}$, so that:

$$T = \sum_{i \in [n]} a_i (a_i \otimes a_i)^\top, \quad \mathbf{S} = \mathbf{T}^{\otimes 2} = \sum_{i,j=1}^n (a_i \otimes a_j)^{\otimes 3},$$

$$\mathbf{M} = \mathbf{S}(T, \text{Id}_{d^2}, \text{Id}_{d^2}) = \sum_{i,j \in [n]} T(a_i \otimes a_j) \otimes (a_i \otimes a_j) \otimes (a_i \otimes a_j).$$

- Sample a vector $g \in \mathbb{R}^d$ with iid standard gaussian entries. Evaluate \mathbf{M} in its first mode in the direction of g to obtain $M \in \mathbb{R}^{d^2 \times d^2}$:

$$M := \mathbf{M}(g, \text{Id}_{d^2}, \text{Id}_{d^2}) = \sum_{i,j \in [n]} \langle g, T(a_i \otimes a_j) \rangle \cdot (a_i \otimes a_j)(a_i \otimes a_j)^\top.$$

- Let $\Sigma \stackrel{\text{def}}{=} \mathbb{E}[(aa^\top)^{\otimes 2}]$ for $a \sim \mathcal{N}(0, \text{Id}_d)$. Let $R \stackrel{\text{def}}{=} \sqrt{2} \cdot (\Sigma^+)^{1/2}$. Compute the top eigenvector $u \in \mathbb{R}^{d^2}$ of RMR , and reshape Ru to a matrix $U \in \mathbb{R}^{d \times d}$.
- For each of the signings of the top 2 unit left (or right) singular vectors $\pm u_1, \pm u_2$ of U , check if $\sum_{i \in [n]} \langle a_i, \pm u_j \rangle^3 \geq 1 - c(n, d)$ where $c(n, d) = \Theta(n/d^{3/2})$ is an appropriate threshold. If so, output $\pm u_j$. Otherwise output nothing.

[Theorem 5.3](#) gets us most of the way to the correctness of [Algorithm 5.17](#), proving that the top eigenvector of the matrix RMR is correlated with some $a_i^{\otimes 2}$ with reasonable probability. We need a few more ingredients to prove [Theorem 5.2](#). First, we need to show a bound on the runtime of [Algorithm 5.17](#).

Lemma 5.18. *Algorithm 5.17 can be implemented in time $\tilde{O}(d^{1+\omega}) \leq \tilde{O}(d^{3.3729})$, where d^ω is the runtime for multiplying two $d \times d$ matrices.*

Proof. To run the algorithm, we only require access to power iteration using the matrix RMR . We first give a fast implementation for power iteration with the matrix M , and handle the multiplications with R separately.

Consider a vector $v \in \mathbb{R}^{d^2}$, and a random vector $g \sim \mathcal{N}(0, \text{Id}_d)$, and let $V, G \in \mathbb{R}^{d \times d}$ be the reshaping of v and gT respectively into matrices. Call $\mathbf{T}_v = \mathbf{T}(\text{Id}_d, V, G)$, where we have applied V and G in the second and third modes of \mathbf{T} , and call T_v the reshaping of \mathbf{T}_v into a $d \times d^2$ matrix. We have that

$$T_v = \sum_{i \in [n]} a_i (V a_i \otimes G a_i)^\top.$$

We show that the matrix-vector multiply Mv can be computed as a flattening of the following product:

$$\begin{aligned} T_v T^\top &= \left(\sum_{i \in [n]} a_i (V a_i \otimes G a_i)^\top \right) \left(\sum_{j \in [n]} (a_j \otimes a_j) a_j^\top \right) \\ &= \sum_{i, j \in [n]} \langle a_j, V a_i \rangle \cdot \langle a_j, G a_i \rangle \cdot a_i a_j^\top \\ &= \sum_{i, j \in [n]} \langle a_i \otimes a_j, v \rangle \cdot \langle gT, a_i \otimes a_j \rangle \cdot a_i a_j^\top. \end{aligned}$$

Flattening $T_v T^\top$ from a $d \times d$ matrix to a vector $v_{TT} \in \mathbb{R}^{d^2}$, we have that

$$v_{TT} = \sum_{i, j \in [n]} \langle gT, a_i \otimes a_j \rangle \cdot \langle a_i \otimes a_j, v \rangle \cdot a_i \otimes a_j = Mv.$$

So we have that Mv is a flattening of the product $T_v T^\top$, which we will compute as a proxy for computing Mv via direct multiplication.

Computing $T_v = \mathbf{T}(\text{Id}, V, G)$ can be done with two matrix multiplication operations, both times multiplying a $d^2 \times d$ matrix with a $d \times d$ matrix. Computing $T_v T^\top$ is a multiplication of a $d \times d^2$ matrix by a $d^2 \times d$ matrix. Both these steps may be done in time $O(d^{1+\omega})$, by regarding the $d \times d^2$ matrices as block matrices with blocks of size $d \times d$. The asymptotically fastest known algorithm for matrix multiplication gives a time of $O(d^{3.3729})$ [Gal14].

Now, to compute the matrix-vector multiply $RMRu$ for any vector $u \in \mathbb{R}^{d^2}$, we may first compute $v = Ru$, perform the operation Mv in time $O(d^{1+\omega})$ as described above, and then again multiply by R . The matrix R is sparse: it has $O(d)$ entries per row (see [Fact C.4](#)), so the multiplication Ru requires time $O(d^3)$.

Performing the update $RMRv$ a total of $O(\log^2 n)$ times is sufficient for convergence, as we have that with reasonable probability, the spectral gap $\lambda_2(RMR)/\lambda_1(RMR) \leq 1 - O(\frac{1}{\log n})$, as a result of applying [Theorem 5.3](#) with the choice of $\varepsilon = O(\frac{1}{\log n})$.

Finally, checking the value of $\sum_i \langle a_i, x \rangle^3$ requires $O(d^3)$ operations, and we do so a constant number of times, once for each of the signings of the top 2 left (or right) singular vectors of U . \square

Next, we need to show that given u with $\langle Ru, a_i \otimes a_i \rangle^2 \geq (1 - \tilde{O}(n^{3/2}/\varepsilon d^2)) \cdot \|u\|^2 \cdot \|a_i\|^4$ we can actually recover the tensor component a_i . Here [Algorithm 5.17](#) reshapes Ru to a $d \times d$ matrix and checks the top two left- or right-singular vectors; the next lemma shows one of these singular vectors must be highly correlated with a_i . (The proof is deferred to [Section A.1](#).)

Lemma 5.19. *Let $M \in \mathbb{R}^{d^2 \times d^2}$ be a symmetric matrix with $\|M\| \leq 1$, and let $v \in \mathbb{R}^d$ and $u \in \mathbb{R}^{d^2}$ be vectors. Furthermore, let U be the reshaping of the vector $Mu \in \mathbb{R}^{d^2}$ to a matrix in $\mathbb{R}^{d \times d}$. Fix $c > 0$, and suppose that $\langle Mu, v \otimes v \rangle^2 \geq c^2 \cdot \|u\|^2 \cdot \|v\|^4$. Then U has some left singular vector a and some right singular vector b such that*

$$|\langle a, v \rangle|, |\langle b, v \rangle| \geq c \cdot \|v\|.$$

Furthermore, for any $0 < \alpha < 1$, there are a', b' among the top $\lfloor \frac{1}{\alpha c^2} \rfloor$ singular vectors of U with

$$|\langle a', v \rangle|, |\langle b', v \rangle| \geq \sqrt{1 - \alpha} \cdot c \cdot \|v\|.$$

If $c \geq \sqrt{\frac{1}{2}(1+\eta)}$ for some $\eta > 0$, then a, b are amongst the top $\lfloor \frac{(1+\eta)}{\eta c^2} \rfloor$ singular vectors.

Since here $c^2 = 1 - o(1)$, we can choose $\eta = 1 - o(1)$ and check only the top 2 singular vectors.

Next, we must show how to choose the threshold $c(n, d)$ so that a big enough value $\sum_{i \in [n]} \langle a_i, u_j \rangle^3$ ensures that u_j is close to a tensor component. The proof is at the end of this section. (A very similar fact appears in [GM15]. We need a somewhat different parameterization here, but we reuse many of their results in the proof.)

Lemma 5.20. Let $T = \sum_{i \in [n]} a_i \otimes a_i \otimes a_i$ for normally distributed vectors $a_i \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$. For all $0 < \gamma, \gamma' < 1$,

1. With overwhelming probability, for every $v \in \mathbb{R}^d$ such that $\sum_{i \in [n]} \langle a_i, v \rangle^3 \geq 1 - \gamma$,

$$\max_{i \in [n]} |\langle a_i, v \rangle| \geq 1 - O(\gamma) - \tilde{O}(n/d^{3/2}).$$

2. With overwhelming probability over a_1, \dots, a_n if $v \in \mathbb{R}^d$ with $\|v\| = 1$ satisfies $\langle v, a_j \rangle \geq 1 - \gamma'$ for some j then $\sum_i \langle a_i, v \rangle^3 \geq 1 - O(\gamma') - \tilde{O}(n/d^{3/2})$.

We are now ready to prove [Theorem 5.2](#).

Proof of [Theorem 5.2](#). By [Theorem 5.3](#), with probability $1 - o(1)$ over a_1, \dots, a_n there are events E_1, \dots, E_n so that $\mathbb{P}_g(E_i) \geq \tilde{O}(1/n^{1+O(\varepsilon)})$ such that when event E_i occurs the top eigenvector u of RMR satisfies

$$\frac{\langle Ru, a_i \otimes a_i \rangle^2}{\|u\|^2 \cdot \|a_i\|^4} \geq 1 - \tilde{O}\left(\frac{n^{3/2}}{\varepsilon d^2}\right).$$

For a particular sample $g \sim \mathcal{N}(0, \text{Id}_d)$, let u_g be this eigenvector.

The algorithm is as follows. Sample $g_1, \dots, g_r \sim \mathcal{N}(0, \text{Id}_d)$ independently for some r to be chosen later. Compute $Ru_{g_1}, \dots, Ru_{g_r}$, reshape each to a $d \times d$ matrix, and compute its singular value decomposition. This gives a family of (right) singular vectors v_1, \dots, v_{dr} . For each, evaluate $\sum_i \langle a_i, v_j \rangle^3$. Let $c(n, d)$ be a threshold to be chosen later. Initialize $S \subset \mathbb{R}^d$ to the empty set. Examining each $1 \leq j \leq dr$ in turn, add v_j to S if $\sum_i \langle a_i, v_j \rangle^3 \geq 1 - c(n, d)$ and for every v already in S , $\langle v, v_j \rangle^2 \leq 1/2$. Output the set S .

Choose $\varepsilon = 1/\log n$. By [Lemma 5.19](#), when E_i occurs for g_j one of $v \in \{\pm v_{jr}, \dots, \pm v_{(j+1)r}\}$ has $\langle v, a_i \rangle \geq (1 - \tilde{O}(n^{3/2}/d^2))(\|u_j\|^2 \cdot \|a_j\|^4)$. Then by [Lemma 5.20](#), when E_i occurs for g_j , this v we will have $\sum_i \langle a_i, \pm v \rangle^3 \geq 1 - \tilde{O}(n/d^{3/2})$. Choose $c(n, d) = \tilde{\Theta}(n^{3/2}/d^2)$ so that when E_i occurs for g_j , so long as it has not previously occurred for some $j' < j$, the algorithm adds $\pm v$ to S .

The events $E_i^{(t)}$ and $E_i^{(t')}$ are independent for any two executions of the algorithm t and t' and have probability $\tilde{\Omega}(1/n)$. Thus, after $r = \tilde{O}(n)$ executions of the algorithm, with high probability for every $i \in [n]$ there is $j \in [r]$ so that E_i occurs for g_j . Finally, by [Lemma 5.20](#), the algorithm can never add to S a vector which is not $(1 - \tilde{O}(n/d^{3/2}))$ -close to some a_i . \square

It just remains to prove [Lemma 5.20](#).

Proof of [Lemma 5.20](#). We start with the first claim. By [[GM15](#), Lemma 2, (proof of) Lemma 8, Theorem 4.2], the following inequalities all hold w.ov.p..

$$\sum_{i \in [n]} \langle a_i, x \rangle^4 \leq 1 + \tilde{O}(n/d^{3/2}) \quad \text{for all } \|x\| = 1, \tag{5.6}$$

$$\sum_{i \in [n]} \langle a_i, x \rangle^6 \geq 1 - O\left(\sum_{i \in [n]} \langle a_i, x \rangle^3 - 1\right) - \tilde{O}(n/d^{3/2}) \quad \text{for all } \|x\| = 1, \quad (5.7)$$

$$\left| \sum_{i \in [n]} \langle a_i, x \rangle^3 \right| \leq 1 + \tilde{O}(n/d^{3/2}) \quad \text{for all } \|x\| = 1. \quad (5.8)$$

To begin,

$$\sum_{i \in [n]} \langle a_i, v \rangle^6 \leq \left(\max_{i \in [n]} \langle a_i, v \rangle^2 \right) \cdot \left(\sum_{i \in [n]} \langle a_i, v \rangle^4 \right).$$

By (5.6), this implies

$$\max_{i \in [n]} \langle a_i, v \rangle^2 \geq (1 - \tilde{O}(n/d^{3/2})) \cdot \sum_{i \in [n]} \langle v, a_i \rangle^6. \quad (5.9)$$

Now combining (5.7) with (5.9) we have

$$\max_{i \in [n]} \langle a_i, v \rangle^2 \geq (1 - \tilde{O}(n/d^{3/2})) \cdot (1 - O(1 - \sum_i \langle a_i, v \rangle^3) - \tilde{O}(n/d^{3/2})).$$

Together with (5.8) this concludes the of the first claim.

For the second claim, we note that by (5.8), and homogeneity, $|\sum_{i \neq j} \langle a_i, x \rangle^3| \leq \|x\|^3(1 + \tilde{O}(n/d^{3/2}))$ w.ov.p.. We write $v = \langle a_j, x \rangle a_j + x^\perp$, where $\langle x^\perp, a_j \rangle = 0$. Now we expand

$$\begin{aligned} \sum_i \langle a_i, v \rangle^3 &\geq (1 - \gamma')^3 + \sum_{i \neq j} \langle \langle a_j, x \rangle a_j + x^\perp, a_i \rangle^3 \\ &= (1 - \gamma')^3 + \sum_{i \neq j} \langle a_j, x \rangle^3 \langle a_j, a_i \rangle^3 + 3 \langle a_j, x \rangle^2 \langle a_j, a_i \rangle^2 \langle x^\perp, a_i \rangle \\ &\quad + 3 \langle a_j, x \rangle \langle a_j, a_i \rangle \langle x^\perp, a_i \rangle^2 + \langle x^\perp, a_i \rangle^3. \end{aligned}$$

We estimate each term in the expansion:

$$\left| \sum_{i \neq j} \langle a_j, x \rangle^3 \langle a_j, a_i \rangle^3 \right| \leq |\langle a_j, x \rangle^3| \sum_{i \neq j} |\langle a_j, a_i \rangle|^3 \leq \tilde{O}\left(\frac{n}{d^{3/2}}\right)$$

w.ov.p. by Cauchy-Schwarz and standard concentration.

$$\begin{aligned} \left| \sum_{i \neq j} \langle a_j, x \rangle^2 \langle a_j, a_i \rangle^2 \langle x^\perp, a_i \rangle \right| &\leq \left(\sum_{i \neq j} \langle a_j, x \rangle^4 \langle a_j, a_i \rangle^4 \right)^{1/2} \left(\sum_{i \neq j} \langle x^\perp, a_i \rangle^2 \right)^{1/2} \quad \text{by Cauchy-Schwarz} \\ &\leq O(\sqrt{n}) \cdot \max_{i \neq j} \langle a_j, a_i \rangle^2 \cdot \tilde{O}\left(\frac{n}{d}\right)^{1/2} \quad \text{w.ov.p. by standard concentration.} \\ &\leq \tilde{O}\left(\frac{n}{d^{3/2}}\right) \quad \text{w.ov.p. by standard concentration} \end{aligned}$$

$$\begin{aligned} \left| \sum_{i \neq j} \langle a_j, x \rangle \langle a_j, a_i \rangle \langle x^\perp, a_i \rangle^2 \right| &\leq O(1) \cdot \max_{i \neq j} |\langle a_j, a_i \rangle| \cdot \sum_{i \neq j} \langle x^\perp, a_i \rangle^2 \quad \text{w.ov.p. by standard concentration} \\ &\leq \tilde{O}\left(\frac{1}{\sqrt{d}}\right) \cdot \tilde{O}\left(\frac{n}{d}\right) \quad \text{w.ov.p. by standard concentration} \end{aligned}$$

$$\begin{aligned} &\leq \tilde{O}\left(\frac{n}{d^{3/2}}\right) \\ \left| \sum_{i \neq j} \langle x^\perp, a_i \rangle^3 \right| &\leq \gamma' + \tilde{O}\left(\frac{n}{d^{3/2}}\right) \quad \text{w.ov.p. by (5.8) and homogeneity.} \end{aligned}$$

Now we estimate

$$\sum_i \langle a_i, v \rangle^3 \geq (1 - \gamma')^3 + \sum_{i \neq j} \langle a_i, x \rangle^3 \geq (1 - \gamma)^3 - \gamma' - \tilde{O}(n/d^{3/2}) \geq 1 - O(\gamma') - \tilde{O}(n/d^{3/2}).$$

since $\gamma' < 1$. □

5.4.1 Boosting Accuracy with Local Search

We remark that [Algorithm 5.17](#) may be used in conjunction with a local search algorithm to obtain much greater guarantees on the accuracy of the recovered vectors. Previous progress on the tensor decomposition problem has produced iterative algorithms that provide local convergence guarantees given a good enough initialization, but which leave the question of how to initialize the procedure up to future work, or up to the specifics of an implementation. In this context, our contribution can be seen as a general method of obtaining good initializations for these local iterative procedures.

In particular, Anandkumar et al. [\[AGJ15\]](#) give an algorithm that combines tensor power iteration and a form of coordinate descent, which when initialized with the output of [Algorithm 5.17](#), achieves a linear convergence rate to the true decomposition within polynomial time.

Theorem 5.21 (Adapted from Theorem 1 in [\[AGJ15\]](#)). *Given a rank- n tensor $\mathbf{T} = \sum_i a_i \otimes a_i \otimes a_i$ with random Gaussian components $a_i \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$. There is a constant $c > 0$ so that if a set of unit vectors $\{x_i \in \mathbb{R}^d\}_i$ satisfies*

$$\langle x_i, a_i \rangle \geq 1 - c, \quad \forall i \in [n],$$

then there exists a procedure which with overwhelming probability over \mathbf{T} and for any $\varepsilon > 0$, recovers a set of vectors $\{\hat{a}_i\}$ such that

$$\langle \hat{a}_i, a_i \rangle \geq 1 - \varepsilon, \quad \forall i \in [n],$$

in time $O(\text{poly}(d) + nd^3 \log \varepsilon)$.

Remark 5.22. Theorem 1 of Anandkumar et al. is stated for random asymmetric tensors, but the adaptation to symmetric tensors is stated in equations (14) and (27) in the same paper.

The theorem of Anandkumar et al. allows for a perturbation tensor Φ , which is just the zero tensor in our setting. Additionally, the weight ratios specifying the weight of each rank-one component in the input tensor are $w_{\max} = w_{\min} = 1$. Lastly, the initialization conditions are given in terms of the distance between the initialization vectors and the true vectors $|x_i - a_i|$, which is related to our measure of closeness $\langle x_i, a_i \rangle$ by the equation $|x_i - a_i|^2 = |x_i|^2 + |a_i|^2 - 2\langle x_i, a_i \rangle$.

The linear convergence guarantee is stated in Lemma 12 of Anandkumar et al.

Corollary 5.23 (Corollary of [Theorem 5.2](#)). *Given as input the tensor $\mathbf{T} = \sum_{i=1}^n a_i \otimes a_i \otimes a_i$ where $a_i \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ with $d \leq n \leq d^{4/3} / \text{polylog } d$, there is a polynomial-time algorithm which with probability*

$1 - o(1)$ over the input \mathbf{T} and the algorithm randomness finds unit vectors $\hat{a}_1, \dots, \hat{a}_n \in \mathbb{R}^d$ such that for all $i \in [n]$,

$$\langle \hat{a}_i, a_i \rangle \geq 1 - O(2^{-n}).$$

Proof. We repeatedly invoke [Algorithm 5.17](#) until we obtain a full set of n vectors as characterized by [Theorem 5.2](#). Apply [Theorem 5.21](#) to the recovered set of vectors until the desired accuracy is obtained. \square

6 Tensor principal component analysis

The Tensor PCA problem in the spiked tensor model is similar to the setting of tensor decomposition, but here the goal is to recover a single large component with all smaller components of the tensor regarded as random noise.

Problem 6.1 (Tensor PCA in the Order-3 Spiked Tensor Model). Given an input tensor $\mathbf{T} = \tau \cdot v^{\otimes 3} + \mathbf{A}$, where $v \in \mathbb{R}^n$ is an arbitrary unit vector, $\tau \geq 0$ is the signal-to-noise ratio, and \mathbf{A} is a random noise tensor with iid standard Gaussian entries, recover the signal v approximately.

Using the partial trace method, we give the first linear-time algorithm for this problem that recovers v for signal-to-noise ratio $\tau = O(n^{3/4} / \text{poly log } n)$. In addition, the algorithm requires only $O(n^2)$ auxiliary space (compared to the input size of n^3) and uses only one non-adaptive pass over the input.

6.1 Spiked tensor model

This spiked tensor model (for general order- k tensors) was introduced by Montanari and Richard [\[RM14\]](#), who also obtained the first algorithms to solve the model with provable statistical guarantees. Subsequently, the SoS approach was applied to the model to improve the signal-to-noise ratio required for odd-order tensors [\[HSS15\]](#); for 3-tensors reducing the requirement from $\tau = \Omega(n)$ to $\tau = \Omega(n^{3/4} \log(n)^{1/4})$.

Using the linear-algebraic objects involved in the analysis of the SoS relaxation, the previous work has also described algorithms with guarantees similar to those of the SoS SDP relaxation, while requiring only nearly subquadratic or linear time [\[HSS15\]](#).

The algorithm here improves on the previous results by use of the partial trace method, simplifying the analysis and improving the runtime by a factor of $\log n$.

6.2 Linear-time algorithm

Linear-Time Algorithm for Tensor PCA

Algorithm 6.2. Input: $\mathbf{T} = \tau \cdot v^{\otimes 3} + \mathbf{A}$. Goal: Recover v' with $\langle v, v' \rangle \geq 1 - o(1)$.

- Compute the partial trace $M := \text{Tr}_{\mathbb{R}^n} \sum_i T_i \otimes T_i \in \mathbb{R}^{n \times n}$, where T_i are the first-mode slices of \mathbf{T} .
- Output the top eigenvector v' of M .

Theorem 6.3. When \mathbf{A} has iid standard Gaussian entries and $\tau \geq Cn^{3/4} \log(n)^{1/2}/\varepsilon$ for some constant C , [Algorithm 6.2](#) recovers v' with $\langle v, v' \rangle \geq 1 - O(\varepsilon)$ with high probability over \mathbf{A} .

Theorem 6.4. [Algorithm 6.2](#) can be implemented in linear time and sublinear space.

These theorems are proved by routine matrix concentration results, showing that in the partial trace matrix, the signal dominates the noise.

To implement the algorithm in linear time it is enough to show that this (sublinear-sized) matrix has constant spectral gap; then a standard application of the matrix power method computes the top eigenvector.

Lemma 6.5. For any v , with high probability over \mathbf{A} , the following occur:

$$\begin{aligned} \left\| \sum_i \text{Tr}(A_i) \cdot A_i \right\| &\leq O(n^{3/2} \log^2 n) \\ \left\| \sum_i v(i) \cdot A_i \right\| &\leq O(\sqrt{n} \log n) \\ \left\| \sum_i \text{Tr}(A_i) v(i) \cdot vv^\top \right\| &\leq O(\sqrt{n} \log n). \end{aligned}$$

The proof may be found in [Appendix D](#).

Proof of [Theorem 6.3](#). We expand the partial trace $\text{Tr}_{\mathbb{R}^n} \sum_i T_i \otimes T_i$.

$$\begin{aligned} \text{Tr}_{\mathbb{R}^n} \sum_i T_i \otimes T_i &= \sum_i \text{Tr}(T_i) \cdot T_i \\ &= \sum_i \text{Tr}(\tau \cdot v(i)vv^\top + A_i) \cdot (\tau \cdot v(i)vv^\top + A_i) \\ &= \sum_i (\tau v(i)\|v\|^2 + \text{Tr}(A_i)) \cdot (\tau \cdot v(i)vv^\top + A_i) \\ &= \tau^2 vv^\top + \tau \left(\sum_i v(i) \cdot A_i + \sum_i \text{Tr}(A_i) v(i)vv^\top \right) + \sum_i \text{Tr}(A_i) \cdot A_i. \end{aligned}$$

Applying [Lemma 6.5](#) and the triangle inequality, we see that

$$\left\| \tau \left(\sum_i v(i) \cdot A_i + \sum_i \text{Tr}(A_i) v(i)vv^\top \right) + \sum_i \text{Tr}(A_i) \cdot A_i \right\| \leq O(n^{3/2} \log n)$$

with high probability. Thus, for appropriate choice of $\tau = \Omega(n^{3/4} \sqrt{(\log n)/\varepsilon})$, the matrix $\text{Tr}_{\mathbb{R}^n} \sum_i T_i \otimes T_i$ is close to rank one, and the result follows by standard manipulations. \square

Proof of [Theorem 6.4](#). Carrying over the expansion of the partial trace from above and setting $\tau = O(n^{3/4} \sqrt{(\log n)/\varepsilon})$, the matrix $\text{Tr}_{\mathbb{R}^n} \sum_i T_i \otimes T_i$ has a spectral gap ratio equal to $\Omega(1/\varepsilon)$ and so the matrix power method finds the top eigenvector in $O(\log(n/\varepsilon))$ iterations. This matrix has

dimension $n \times n$, so a single iteration takes $O(n^2)$ time, which is sublinear in the input size n^3 . Finally, to construct $\text{Tr}_{\mathbb{R}^n} \sum_i T_i \otimes T_i$ we use

$$\text{Tr}_{\mathbb{R}^n} \sum_i T_i \otimes T_i = \sum_i \text{Tr}(T_i) \cdot T_i$$

and note that to construct the right-hand side it is enough to examine each entry of \mathbf{T} just $O(1)$ times and perform $O(n^3)$ additions. At no point do we need to store more than $O(n^2)$ matrix entries at the same time. \square

Acknowledgements

We would like to thank Rong Ge for very helpful discussions. We also thank Jonah Brown Cohen, Pasin Manurangsi and Aviad Rubinfeld for helpful comments in the preparation of this manuscript.

References

- [AFH⁺15] Anima Anandkumar, Dean P. Foster, Daniel Hsu, Sham M. Kakade, and Yi-Kai Liu, *A spectral algorithm for latent dirichlet allocation*, *Algorithmica* **72** (2015), no. 1, 193–214. [3](#)
- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky, *Tensor decompositions for learning latent variable models*, *Journal of Machine Learning Research* **15** (2014), no. 1, 2773–2832. [3](#)
- [AGHK14] Animashree Anandkumar, Rong Ge, Daniel Hsu, and Sham M. Kakade, *A tensor approach to learning mixed membership community models*, *Journal of Machine Learning Research* **15** (2014), no. 1, 2239–2312. [3](#)
- [AGJ14] Anima Anandkumar, Rong Ge, and Majid Janzamin, *Analyzing tensor power method dynamics: Applications to learning overcomplete latent variable models*, *CoRR* **abs/1411.1488** (2014). [3](#)
- [AGJ15] Animashree Anandkumar, Rong Ge, and Majid Janzamin, *Learning overcomplete latent variable models through tensor methods*, *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pp. 36–112. [3](#), [4](#), [21](#), [37](#)
- [AW02] Rudolf Ahlswede and Andreas J. Winter, *Strong converse for identification via quantum channels*, *IEEE Transactions on Information Theory* **48** (2002), no. 3, 569–579. [48](#)
- [BBH⁺12] Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou, *Hypercontractivity, sum-of-squares proofs, and their applications*, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pp. 307–326. [1](#), [8](#)
- [BCMV14] Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan, *Smoothed analysis of tensor decompositions*, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pp. 594–603. [3](#)

- [BKS14] Boaz Barak, Jonathan A. Kelner, and David Steurer, *Rounding sum-of-squares relaxations*, Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014, pp. 31–40. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#), [43](#)
- [BKS15] ———, *Dictionary learning and tensor decomposition via the sum-of-squares method*, Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015, pp. 143–151. [1](#), [3](#), [5](#), [8](#), [10](#), [11](#), [12](#)
- [BM15] Boaz Barak and Ankur Moitra, *Tensor prediction, rademacher complexity and random 3-xor*, CoRR [abs/1501.06521](#) (2015). [1](#), [6](#)
- [BRS11] Boaz Barak, Prasad Raghavendra, and David Steurer, *Rounding semidefinite programming hierarchies via global correlation*, IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011, pp. 472–481. [6](#), [8](#)
- [BS14] Boaz Barak and David Steurer, *Sum-of-squares proofs and the quest toward optimal algorithms*, CoRR [abs/1404.5236](#) (2014). [1](#)
- [Cha96] J. T. Chang, *Full reconstruction of markov models on evolutionary trees: Identifiability and consistency*, Math Biosci. **137** (1996), 51–73. [3](#)
- [dGJL04] Alexandre d’Aspremont, Laurent El Ghaoui, Michael I. Jordan, and Gert R. G. Lanckriet, *A direct formulation for sparse PCA using semidefinite programming*, Advances in Neural Information Processing Systems 17 [Neural Information Processing Systems, NIPS 2004, December 13-18, Vancouver, British Columbia, Canada], 2004, pp. 41–48. [2](#)
- [DH14] Laurent Demanet and Paul Hand, *Scaling law for recovering the sparsest element in a subspace*, Information and Inference **3** (2014), no. 4, 295–309. [2](#), [3](#)
- [dIPMS95] Victor H de la Peña and Stephen J Montgomery-Smith, *Decoupling inequalities for the tail probabilities of multivariate u -statistics*, The Annals of Probability (1995), 806–816. [28](#)
- [dIVK07] Wenceslas Fernandez de la Vega and Claire Kenyon-Mathieu, *Linear programming relaxations of maxcut*, Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007, pp. 53–61. [6](#)
- [Gal14] François Le Gall, *Powers of tensors and fast matrix multiplication*, International Symposium on Symbolic and Algebraic Computation, ISSAC ’14, Kobe, Japan, July 23-25, 2014, pp. 296–303. [34](#)
- [GHK15] Rong Ge, Qingqing Huang, and Sham M. Kakade, *Learning mixtures of gaussians in high dimensions*, Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015, pp. 761–770. [3](#)
- [GM15] Rong Ge and Tengyu Ma, *Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms*, Approximation, Randomization, and Combinatorial Optimization.

- Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, Princeton, NJ, USA, 2015, pp. 829–849. [1](#), [3](#), [4](#), [5](#), [10](#), [11](#), [25](#), [28](#), [35](#)
- [GS12] Venkatesan Guruswami and Ali Kemal Sinop, *Faster SDP hierarchy solvers for local rounding algorithms*, 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012, pp. 197–206. [6](#)
- [GVX14] Navin Goyal, Santosh Vempala, and Ying Xiao, *Fourier PCA and robust tensor decomposition*, Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014, pp. 584–593. [3](#)
- [Har70] Richard A Harshman, *Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis*, UCLA Working Papers in Phonetics **16** (1970), 1–84. [3](#)
- [Hås90] Johan Håstad, *Tensor rank is np-complete*, J. Algorithms **11** (1990), no. 4, 644–654. [3](#)
- [HL13] Christopher J. Hillar and Lek-Heng Lim, *Most tensor problems are np-hard*, J. ACM **60** (2013), no. 6, 45. [3](#)
- [HSS15] Samuel B. Hopkins, Jonathan Shi, and David Steurer, *Tensor principal component analysis via sum-of-square proofs*, Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015, pp. 956–1006. [1](#), [5](#), [6](#), [12](#), [38](#)
- [Jan97] Svante Janson, *Gaussian hilbert spaces*, vol. 129, Cambridge university press, 1997. [47](#)
- [KB09] Tamara G. Kolda and Brett W. Bader, *Tensor decompositions and applications*, SIAM Review **51** (2009), no. 3, 455–500. [3](#)
- [Las01] Jean B. Lasserre, *Global optimization with polynomials and the problem of moments*, SIAM J. Optim. **11** (2000/01), no. 3, 796–817. MR 1814045 (2002b:90054) [1](#)
- [LCC07] Lieven De Lathauwer, Joséphine Castaing, and Jean-François Cardoso, *Fourth-order cumulant-based blind identification of underdetermined mixtures*, IEEE Transactions on Signal Processing **55** (2007), no. 6-2, 2965–2973. [3](#)
- [LM00] B. Laurent and P. Massart, *Adaptive estimation of a quadratic functional by model selection*, Ann. Statist. **28** (2000), no. 5, 1302–1338. [52](#)
- [LR15] Elaine Levey and Thomas Rothvoss, *A lasserre-based $(1 + \epsilon)$ -approximation for $P_m \mid p_j = 1, prec \mid C_{\max}$* , CoRR [abs/1509.07808](#) (2015). [8](#)
- [MR06] Elchanan Mossel and Sébastien Roch, *Learning nonsingular phylogenies and hidden markov models*, Ann. Appl. Probab. **16** (2006), no. 2, 583–614. [3](#)
- [MW13] Raghu Meka and Avi Wigderson, *Association schemes, non-commutative polynomial concentration, and sum-of-squares lower bounds for planted clique*, CoRR [abs/1307.7615](#) (2013). [8](#)

- [MW15] Tengyu Ma and Avi Wigderson, *Sum-of-squares lower bounds for sparse PCA*, CoRR [abs/1507.06370](#) (2015). [6](#)
- [Nes00] Yurii Nesterov, *Squared functional systems and optimization problems*, High performance optimization, Appl. Optim., vol. 33, Kluwer Acad. Publ., Dordrecht, 2000, pp. 405–440. MR 1748764 (2001b:90063) [1](#)
- [NR09] Phong Q. Nguyen and Oded Regev, *Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures*, J. Cryptology **22** (2009), no. 2, 139–160. [3](#)
- [Par00] Pablo A Parrilo, *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*, Ph.D. thesis, Citeseer, 2000. [1](#)
- [QSW14] Qing Qu, Ju Sun, and John Wright, *Finding a sparse vector in a subspace: Linear sparsity using alternating directions*, Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, 2014, pp. 3401–3409. [2](#), [3](#)
- [RM14] Emile Richard and Andrea Montanari, *A statistical model for tensor PCA*, Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, 2014, pp. 2897–2905. [4](#), [5](#), [12](#), [38](#)
- [RT12] Prasad Raghavendra and Ning Tan, *Approximating csps with global cardinality constraints using SDP hierarchies*, Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012, pp. 373–387. [8](#)
- [Sho87] N. Z. Shor, *An approach to obtaining global extrema in polynomial problems of mathematical programming*, Kibernetika (Kiev) (1987), no. 5, 102–106, 136. MR 931698 (89d:90202) [1](#)
- [SWW12] Daniel A. Spielman, Huan Wang, and John Wright, *Exact recovery of sparsely-used dictionaries*, COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, Edinburgh, Scotland, 2012, pp. 37.1–37.18. [2](#), [7](#)
- [Tro12] Joel A. Tropp, *User-friendly tail bounds for sums of random matrices*, Foundations of Computational Mathematics **12** (2012), no. 4, 389–434. [29](#), [48](#)
- [Ver10] Roman Vershynin, *Introduction to the non-asymptotic analysis of random matrices*, CoRR [abs/1011.3027](#) (2010). [52](#), [54](#), [61](#)

A Additional preliminaries

A.1 Linear algebra

Here we provide some lemmas in linear algebra.

This first lemma is closely related to the sos Cauchy-Schwarz from [BKS14], and the proof is essentially the same.

Lemma A.1 (PSD Cauchy-Schwarz). Let $M \in \mathbb{R}^{d \times d}$, $M \geq 0$ and symmetric. Let $p_1, \dots, p_n, q_1, \dots, q_n \in \mathbb{R}^d$. Then

$$\langle M, \sum_{i=1}^n p_i q_i^\top \rangle \leq \langle M, \sum_{i=1}^n p_i p_i^\top \rangle^{1/2} \langle M, \sum_{i=1}^n q_i q_i^\top \rangle^{1/2}.$$

In applications, we will have $\sum_i p_i q_i$ as a single block of a larger block matrix containing also the blocks $\sum_i p_i p_i^\top$ and $\sum_i q_i q_i^\top$.

Proof. We first claim that

$$\langle M, \sum_{i=1}^n p_i q_i^\top \rangle \leq \frac{1}{2} \langle M, \sum_{i=1}^n p_i p_i^\top \rangle + \frac{1}{2} \langle M, \sum_{i=1}^n q_i q_i^\top \rangle.$$

To see this, just note that the right-hand side minus the left is exactly

$$\langle M, \sum_{i=1}^n (p_i - q_i)(p_i - q_i)^\top \rangle = \sum_i (p_i - q_i)^\top M (p_i - q_i) \geq 0.$$

The lemma follows now by applying this inequality to

$$p'_i = \frac{p_i}{\langle M, \sum_{i=1}^n p_i p_i^\top \rangle^{1/2}} \quad q'_i = \frac{q_i}{\langle M, \sum_{i=1}^n q_i q_i^\top \rangle^{1/2}}. \quad \square$$

Lemma A.2 (Operator Norm Cauchy-Schwarz for Sums). Let $A_1, \dots, A_m, B_1, \dots, B_m$ be real random matrices. Then

$$\left\| \sum_i \mathbb{E} A_i B_i \right\| \leq \left\| \sum_i \mathbb{E} A_i^\top A_i \right\|^{1/2} \left\| \sum_i \mathbb{E} B_i^\top B_i \right\|^{1/2}.$$

Proof. We have for any unit x, y ,

$$\begin{aligned} x^\top \sum_i \mathbb{E} A_i B_i x &= \sum_i \mathbb{E} \langle A_i x, B_i y \rangle \\ &\leq \sum_i \mathbb{E} \|A_i x\| \|B_i y\| \\ &\leq \sum_i (\mathbb{E} \|A_i x\|^2)^{1/2} (\mathbb{E} \|B_i y\|^2)^{1/2} \\ &\leq \sqrt{\sum_i \mathbb{E} \|A_i x\|^2} \sqrt{\sum_i \mathbb{E} \|B_i y\|^2} \\ &= \sqrt{\mathbb{E} x^\top \sum_i A_i^\top A_i x} \sqrt{\mathbb{E} y^\top \sum_i B_i^\top B_i y} \\ &\leq \left\| \sum_i \mathbb{E} A_i^\top A_i \right\|^{1/2} \left\| \sum_i \mathbb{E} B_i^\top B_i \right\|^{1/2}. \end{aligned}$$

where the nontrivial inequalities follow from Cauchy-Schwarz for expectations, vectors and scalars, respectively. \square

The following lemma allows to argue about the top eigenvector of matrices with spectral gap.

Lemma A.3 (Top eigenvector of gapped matrices). *Let M be a symmetric r -by- r matrix and let u, v be a vectors in \mathbb{R}^r with $\|u\| = 1$. Suppose u is a top singular vector of M so that $|\langle u, Mu \rangle| = \|M\|$ and v satisfies for some $\varepsilon > 0$,*

$$\|M - vv^\top\| \leq \|M\| - \varepsilon \cdot \|v\|^2$$

Then, $\langle u, v \rangle^2 \geq \varepsilon \cdot \|v\|^2$.

Proof. We lower bound the quadratic form of $M - vv^\top$ evaluated at u by

$$|\langle u, (M - vv^\top)u \rangle| \geq |\langle u, Mu \rangle| - \langle u, v \rangle^2 = \|M\| - \langle u, v \rangle^2.$$

At the same time, this quadratic form evaluated at u is upper bounded by $\|M\| - \varepsilon \cdot \|v\|^2$. It follows that $\langle u, v \rangle^2 \geq \varepsilon \cdot \|v\|^2$ as desired. \square

The following lemma states that a vector in \mathbb{R}^{d^2} which is close to a symmetric vector $v^{\otimes 2}$, if flattened to a matrix, has top eigenvector correlated with the symmetric vector.

Lemma (Restatement of [Lemma 5.19](#)). *Let $M \in \mathbb{R}^{d^2 \times d^2}$ be a symmetric matrix with $\|M\| \leq 1$, and let $v \in \mathbb{R}^d$ and $u \in \mathbb{R}^{d^2}$ be vectors. Furthermore, let U be the reshaping of the vector $Mu \in \mathbb{R}^{d^2}$ to a matrix in $\mathbb{R}^{d \times d}$. Fix $c > 0$, and suppose that $\langle Mu, v \otimes v \rangle^2 \geq c^2 \cdot \|u\|^2 \cdot \|v\|^4$. Then U has some left singular vector a and some right singular vector b such that*

$$|\langle a, v \rangle|, |\langle b, v \rangle| \geq c \cdot \|v\|.$$

Furthermore, for any $0 < \alpha < 1$, there are a', b' among the top $\lfloor \frac{1}{\alpha c^2} \rfloor$ singular vectors of U with

$$|\langle a', v \rangle|, |\langle b', v \rangle| \geq \sqrt{1 - \alpha} \cdot c \cdot \|v\|.$$

If $c \geq \sqrt{\frac{1}{2}(1 + \eta)}$ for some $\eta > 0$, then a, b are amongst the top $\lfloor \frac{(1+\eta)}{\eta c^2} \rfloor$ singular vectors.

Proof. Let $\hat{v} = v/\|v\|$. Let (σ_i, a_i, b_i) be the i th singular value, left and right (unit) singular vectors of U respectively.

Our assumptions imply that

$$|\hat{v}^\top U \hat{v}| = |\langle Mu, \hat{v} \otimes \hat{v} \rangle| \geq c \cdot \|u\|.$$

Furthermore, we observe that $\|U\|_F = \|Mu\| \leq \|M\| \cdot \|u\|$, and that therefore $\|U\|_F \leq \|u\|$. We thus have that,

$$c \cdot \|u\| \leq |\hat{v}^\top U \hat{v}| = \left| \sum_{i \in [d]} \sigma_i \cdot \langle \hat{v}, a_i \rangle \langle \hat{v}, b_i \rangle \right| \leq \|u\| \cdot \sqrt{\sum_{i \in [d]} \langle \hat{v}, a_i \rangle^2 \langle \hat{v}, b_i \rangle^2},$$

where to obtain the last inequality we have used Cauchy-Schwarz and our bound on $\|U\|_F$. We may thus conclude that

$$c^2 \leq \sum_{i \in [d]} \langle \hat{v}, a_i \rangle^2 \langle \hat{v}, b_i \rangle^2 \leq \max_{i \in [d]} \langle a_i, \hat{v} \rangle^2 \cdot \sum_{i \in [d]} \langle b_i, \hat{v} \rangle^2 = \max_{i \in [d]} \langle a_i, \hat{v} \rangle^2, \quad (\text{A.1})$$

where we have used the fact that the left singular values of U are orthonormal. The argument is symmetric in the b_i .

Furthermore, we have that

$$c^2 \cdot \|u\|^2 \leq |\hat{v}^\top U \hat{v}|^2 = \left| \sum_{i \in [d]} \sigma_i \cdot \langle \hat{v}, a_i \rangle \langle \hat{v}, b_i \rangle \right|^2 \leq \left(\sum_{i \in [d]} \sigma_i^2 \langle \hat{v}, a_i \rangle^2 \right) \cdot \left(\sum_{i \in [d]} \langle \hat{v}, b_i \rangle^2 \right) = \sum_{i \in [d]} \sigma_i^2 \langle \hat{v}, a_i \rangle^2,$$

where we have applied Cauchy-Schwarz and the orthonormality of the b_i . In particular,

$$\sum_{i \in [d]} \sigma_i^2 \langle \hat{v}, a_i \rangle^2 \geq c^2 \|u\|^2 \geq c^2 \|U\|_F^2.$$

On the other hand, let S be the set of $i \in [d]$ for which $\sigma_i^2 \leq \alpha c^2 \|U\|_F^2$. By substitution,

$$\sum_{i \in S} \sigma_i^2 \langle \hat{v}, a_i \rangle^2 \leq \alpha c^2 \|U\|_F^2 \sum_{i \in S} \langle \hat{v}, a_i \rangle^2 \leq \alpha c^2 \|U\|_F^2,$$

where we have used the fact that the right singular vectors are orthonormal. The last two inequalities imply that $S \neq [d]$. Letting $T = [d] \setminus S$, it follows from subtraction that

$$(1 - \alpha) c^2 \|U\|_F^2 \leq \sum_{i \in T} \sigma_i^2 \langle \hat{v}, a_i \rangle^2 \leq \max_{i \in T} \langle \hat{v}, a_i \rangle^2 \sum_{i \in T} \sigma_i^2 = \max_{i \in T} \langle \hat{v}, a_i \rangle^2 \|U\|_F^2,$$

so that $\max_{i \in T} \langle \hat{v}, a_i \rangle^2 \geq (1 - \alpha) c^2$. Finally,

$$|T| \cdot \alpha c^2 \|U\|_F^2 \leq |T| \cdot \min_{i \in T} \sigma_i^2 \leq \sum_{i \in [d]} \sigma_i^2 = \|U\|_F^2,$$

so that $|T| \leq \lfloor \frac{1}{\alpha c^2} \rfloor$. Thus, one of the top $\lfloor \frac{1}{\alpha c^2} \rfloor$ right singular vectors a has correlation $|\langle \hat{v}, a \rangle| \geq \sqrt{(1 - \alpha)c}$. The same proof holds for the b .

Furthermore, if $c^2 > \frac{1}{2}(1 + \eta)$ for some $\eta > 0$, and $(1 - \alpha)c^2 > \frac{1}{2}$, then by (A.1) it must be that $\max_{i \in T} \langle \hat{v}, a_i \rangle^2 = \max_{i \in [d]} \langle \hat{v}, a_i \rangle^2$, as \hat{v} cannot have square correlation larger than $\frac{1}{2}$ with more than one left singular vector. Taking $\alpha = \frac{\eta}{1 + \eta}$ guarantees this. The conclusion follows. \square

A.2 Concentration tools

We require a number of tools from the literature on concentration of measure.

A.2.1 For scalar-valued polynomials of Gaussians

We need the some concentration bounds for certain polynomials of Gaussian random variables.

The following lemma gives standard bounds on the tails of a standard gaussian variable—somewhat more precisely than other bounds in this paper. Though there are ample sources, we repeat the proof here for reference.

Lemma A.4. *Let $X \sim \mathcal{N}(0, 1)$. Then for $t > 0$,*

$$\mathbb{P}(X > t) \leq \frac{e^{-t^2/2}}{t\sqrt{2\pi}},$$

and

$$\mathbb{P}(X > t) \geq \frac{e^{-t^2/2}}{\sqrt{2\pi}} \cdot \left(\frac{1}{t} - \frac{1}{t^3} \right).$$

Proof. To show the first statement, we apply an integration trick,

$$\begin{aligned}\mathbb{P}(X > t) &= \frac{1}{\sqrt{2\pi}} \int_t^\infty e^{-x^2/2} dx \\ &\leq \frac{1}{\sqrt{2\pi}} \int_t^\infty \frac{x}{t} e^{-x^2/2} dx \\ &= \frac{e^{-t^2/2}}{t\sqrt{2\pi}},\end{aligned}$$

where in the third step we have used the fact that $\frac{x}{t} \leq x$ for $t \geq x$. For the second statement, we integrate by parts and repeat the trick,

$$\begin{aligned}\mathbb{P}(X > t) &= \frac{1}{\sqrt{2\pi}} \int_t^\infty e^{-x^2/2} dx \\ &= \frac{1}{\sqrt{2\pi}} \int_t^\infty \frac{1}{x} \cdot x e^{-x^2/2} dx \\ &= \frac{1}{\sqrt{2\pi}} \left[-\frac{1}{x} e^{-x^2/2} \right]_t^\infty - \frac{1}{\sqrt{2\pi}} \int_t^\infty \frac{1}{x^2} \cdot e^{-x^2/2} dx \\ &\geq \frac{1}{\sqrt{2\pi}} \left[-\frac{1}{x} e^{-x^2/2} \right]_t^\infty - \frac{1}{\sqrt{2\pi}} \int_t^\infty \frac{x}{t^3} \cdot e^{-x^2/2} dx \\ &= \frac{1}{\sqrt{2\pi}} \left(\frac{1}{t} - \frac{1}{t^3} \right) e^{-t^2/2}.\end{aligned}$$

This concludes the proof. \square

The following is a small modification of Theorem 6.7 from [Jan97] which follows from Remark 6.8 in the same.

Lemma A.5. *For each $\ell \geq 1$ there is a universal constant $c_\ell > 0$ such that for every f a degree- ℓ polynomial of standard Gaussian random variables X_1, \dots, X_m and $t \geq 2$,*

$$\mathbb{P}(|f(X)| > t \mathbb{E}|f(X)|) \leq e^{-c_\ell t^{2/\ell}}.$$

The same holds (with a different constant c_ℓ) if $\mathbb{E}|f(x)|$ is replaced by $(\mathbb{E}f(x)^2)^{1/2}$.

In our concentration results, we will need to calculate the expectations of multivariate Gaussian polynomials, many of which share a common form. Below we give an expression for these expectations.

Fact A.6. *Let x be a d -dimensional vector with independent identically distributed gaussian entries with variance σ^2 . Let u be a fixed unit vector. Then setting $X = (\|x\|^2 - c)^p \|x\|^{2m} x x^T$, and setting $U = (\|x\|^2 - c)^p \|x\|^{2m} u u^T$, we have*

$$\mathbb{E}[X] = \left(\sum_{0 \leq k \leq p} \binom{p}{k} (-1)^k c^k (d+2) \cdots (d+2p+2m-2k) \sigma^{2(p+m-k+1)} \right) \cdot \text{Id},$$

and

$$\mathbb{E}[U] = \left(\sum_{0 \leq k \leq p} \binom{p}{k} (-1)^k c^k d(d+2) \cdots (d+2p+2m-2k-2) \sigma^{2(p+m-k)} \right) \cdot u u^T$$

Proof.

$$\begin{aligned}\mathbb{E}[X] &= \mathbb{E}[(\|x\|^2 - c)^p \|x\|^{2m} x_1^2] \cdot \text{Id} \\ &= \text{Id} \cdot \sum_{0 \leq k \leq p} \binom{p}{k} (-1)^k c^k \mathbb{E} \left[\left(\sum_{\ell \in [d]} x_\ell^2 \right)^{p+m-k} x_1^2 \right]\end{aligned}$$

Since $(\sum_{i \in [d]} x_i^2)^{p+m-k}$ is symmetric in x_1, \dots, x_d , we have

$$= \text{Id} \cdot \frac{1}{d} \sum_{0 \leq k \leq p} \binom{p}{k} (-1)^k c^k \mathbb{E} \left[\left(\sum_{i \in [d]} x_i^2 \right)^{p+m-k+1} \right]$$

We have reduced the computation to a question of the moments of a Chi-squared variable with d degrees of freedom. Using these moments,

$$\begin{aligned}&= \text{Id} \cdot \frac{1}{d} \sum_{0 \leq k \leq p} \binom{p}{k} (-1)^k c^k d(d+2) \cdots (d+2p+2m-2k) \sigma^{2(p+m-k+1)} \\ &= \text{Id} \cdot \left(\sum_{0 \leq k \leq p} \binom{p}{k} (-1)^k c^k (d+2) \cdots (d+2p+2m-2k) \sigma^{2(p+m-k+1)} \right).\end{aligned}$$

A similar computation yields the result about $\mathbb{E}[U]$. □

A.2.2 For matrix-valued random variables

On several occasions we will need to apply a Matrix-Bernstein-like theorem to a sum of matrices with an unfortunate tail. To this end, we prove a “truncated Matrix Bernstein Inequality.” Our proof uses an standard matrix Bernstein inequality as a black box. The study of inequalities of this variety—on tails of sums of independent matrix-valued random variables—was initiated by Ahlswede and Winter [AW02]. The excellent survey of Tropp [Tro12] provides many results of this kind.

In applications of the following the operator norms of the summands X_1, \dots, X_n have well-behaved tails and so the truncation is a routine formality. Two corollaries following the proposition and its proof capture truncation for all the matrices we encounter in the present work.

Proposition A.7 (Truncated Matrix Bernstein). *Let $X_1, \dots, X_n \in \mathbb{R}^{d_1 \times d_2}$ be independent random matrices, and suppose that*

$$\mathbb{P} \left[\|X_i - \mathbb{E}[X_i]\|_{op} \geq \beta \right] \leq p \text{ for all } i \in [n].$$

Furthermore, suppose that for each X_i ,

$$\|\mathbb{E}[X_i] - \mathbb{E}[X_i \mathbb{I}[\|X_i\|_{op} < \beta]]\| \leq q.$$

Denote

$$\sigma^2 = \max \left\{ \left\| \sum_{i \in [n]} \mathbb{E}[X_i X_i^T] - \mathbb{E}[X_i] \mathbb{E}[X_i^T] \right\|_{op}, \left\| \sum_{i \in [n]} \mathbb{E}[X_i^T X_i] - \mathbb{E}[X_i]^T \mathbb{E}[X_i] \right\|_{op} \right\}.$$

Then for $X = \sum_{i \in [n]} X_i$, we have

$$\mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t] \leq n \cdot p + (d_1 + d_2) \cdot \exp\left(\frac{-(t - nq)^2}{2(\sigma^2 + \beta(t - nq)/3)}\right).$$

Proof. For simplicity we start by centering the variables X_i . Let $\tilde{X}_i = X_i - \mathbb{E} X_i$ and $\tilde{X} = \sum_{i \in [n]} \tilde{X}_i$. The proof proceeds by a straightforward application of the noncommutative Bernstein's Inequality. We define variables Y_1, \dots, Y_n , which are the truncated counterparts of the \tilde{X}_i s in the following sense:

$$Y_i = \begin{cases} \tilde{X}_i & \|\tilde{X}_i\|_{op} < \beta, \\ 0 & \text{otherwise.} \end{cases}$$

Define $Y = \sum_{i \in [n]} Y_i$. We claim that

$$\left\| \sum_i \mathbb{E} Y_i Y_i^T - \mathbb{E}[Y_i] \mathbb{E}[Y_i]^T \right\|_{op} \leq \left\| \sum_i \mathbb{E} \tilde{X}_i \tilde{X}_i^T \right\|_{op} \leq \sigma^2 \text{ and} \quad (\text{A.2})$$

$$\left\| \sum_i \mathbb{E} Y_i^T Y_i - \mathbb{E}[Y_i]^T \mathbb{E}[Y_i] \right\|_{op} \leq \left\| \sum_i \mathbb{E} \tilde{X}_i^T \tilde{X}_i \right\|_{op} \leq \sigma^2, \quad (\text{A.3})$$

which, together with the fact that $\|Y_i\| \leq \beta$ almost surely, will allow us to apply the non-commutative Bernstein's inequality to Y . To see (A.2) ((A.3) is similar), we expand $\mathbb{E} Y_i Y_i^T$ as

$$\mathbb{E} Y_i Y_i^T = \mathbb{P}[\|\tilde{X}_i\|_{op} < \beta] \mathbb{E}[\tilde{X}_i \tilde{X}_i^T \mid \|\tilde{X}_i\|_{op} < \beta].$$

Additionally expanding $\mathbb{E}[\tilde{X}_i \tilde{X}_i^T]$ as

$$\mathbb{E}[\tilde{X}_i \tilde{X}_i^T] = \mathbb{P}[\|\tilde{X}_i\|_{op} < \beta] \mathbb{E}[\tilde{X}_i \tilde{X}_i^T \mid \|\tilde{X}_i\|_{op} < \beta] + \mathbb{P}[\|\tilde{X}_i\|_{op} \geq \beta] \mathbb{E}[\tilde{X}_i \tilde{X}_i^T \mid \|\tilde{X}_i\|_{op} \geq \beta],$$

we note that $\mathbb{E}[\tilde{X}_i \tilde{X}_i^T \mid \|\tilde{X}_i\|_{op} \geq \beta]$ is PSD. Thus, $\mathbb{E}[Y_i Y_i^T] \geq \mathbb{E}[X_i X_i^T]$. But by definition $\mathbb{E}[Y_i Y_i^T]$ is still PSD (and hence $\|\sum_i \mathbb{E}[Y_i Y_i^T]\|_{op}$ is given by the maximum eigenvalue of $\mathbb{E}[Y_i Y_i^T]$), so

$$\left\| \sum_i \mathbb{E} Y_i Y_i^T \right\|_{op} \leq \left\| \sum_i \mathbb{E} \tilde{X}_i \tilde{X}_i^T \right\|_{op}.$$

Also PSD are $\mathbb{E}[Y_i] \mathbb{E}[Y_i]^T$ and $\mathbb{E}[(Y_i - \mathbb{E}[Y_i])(Y_i - \mathbb{E}[Y_i])^T] = \mathbb{E}[Y_i Y_i^T] - \mathbb{E}[Y_i] \mathbb{E}[Y_i]^T$. By the same reasoning again, then, we get $\|\sum_i \mathbb{E} Y_i Y_i^T - \mathbb{E}[Y_i] \mathbb{E}[Y_i]^T\|_{op} \leq \|\sum_i \mathbb{E}[Y_i Y_i^T]\|_{op}$. Putting this all together gives (A.2).

Now we are ready to apply the non-commutative Bernstein's inequality to Y . We have

$$\mathbb{P}[\|Y - \mathbb{E}[Y]\|_{op} \geq \alpha] \leq (d_1 + d_2) \cdot \exp\left(\frac{-\alpha^2/2}{\sigma^2 + \beta \cdot \alpha/3}\right).$$

Now, we have

$$\mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t] = \mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t \mid X = Y] \cdot \mathbb{P}[X = Y]$$

$$\begin{aligned}
& + \mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t \mid X \neq Y] \cdot \mathbb{P}[X \neq Y], \\
& \leq \mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t \mid X = Y] + n \cdot p
\end{aligned}$$

by a union bound over the events $\{X_i \neq Y_i\}$. It remains to bound the conditional probability $\mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t \mid X = Y]$. By assumption, $\|\mathbb{E}[X] - \mathbb{E}[Y]\|_{op} \leq nq$, and so by the triangle inequality,

$$\|X - \mathbb{E}[X]\|_{op} \leq \|X - \mathbb{E}[Y]\|_{op} + \|\mathbb{E}[Y] - \mathbb{E}[X]\|_{op} \leq \|X - \mathbb{E}[Y]\|_{op} + nq.$$

Thus,

$$\begin{aligned}
\mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t \mid X = Y] & \leq \mathbb{P}[\|X - \mathbb{E}[Y]\|_{op} + nq \geq t \mid X = Y] \\
& = \mathbb{P}[\|Y - \mathbb{E}[Y]\|_{op} \geq t - nq \mid X = Y].
\end{aligned}$$

Putting everything together and setting $\alpha = t - nq$,

$$\mathbb{P}[\|X - \mathbb{E}[X]\|_{op} \geq t] \leq n \cdot p + (d_1 + d_2) \cdot \exp\left(\frac{-(t - nq)^2/2}{\sigma^2 + \beta(t - nq)/3}\right),$$

as desired. \square

The following lemma helps achieve the assumptions of [Proposition A.7](#) easily for a useful class of thin-tailed random matrices.

Lemma A.8. *Suppose that X is a matrix whose entries are polynomials of constant degree ℓ in unknowns x , which we evaluate at independent Gaussians. Let $f(x) := \|X\|_{op}$ and $g(x) := \|XX^T\|_{op}$, and either f is itself a polynomial in x of degree at most 2ℓ or g is a polynomial in x of degree at most 4ℓ . Then if $\beta = R \cdot \alpha$ for $\alpha \geq \min\{\mathbb{E}[|f(x)|], \sqrt{\mathbb{E}[g(x)]}\}$ and $R = \text{polylog}(n)$,*

$$\mathbb{P}(\|X\|_{op} \geq \beta) \leq n^{-\log n}, \tag{A.4}$$

and

$$\mathbb{E}[\|X \cdot \mathbb{I}\{\|X\|_{op} \geq \beta\}\|_{op}] \leq (\beta + \alpha)n^{-\log n}. \tag{A.5}$$

Proof. We begin with [\(A.4\)](#). Either $f(x)$ is a polynomial of degree at most 2ℓ , or $g(x)$ is a polynomial of degree at most 4ℓ in gaussian variables. We can thus use [Lemma A.5](#) to obtain the following bound,

$$\mathbb{P}(|f(x)| \geq t\alpha) \leq \exp(-ct^{1/(2\ell)}), \tag{A.6}$$

where c is a universal constant. Taking $t = R = \text{polylog}(n)$ gives us [\(A.4\)](#).

We now address [\(A.5\)](#). To this end, let $p(t)$ and $P(t)$ be the probability density function and cumulative density function of $\|X\|_{op}$, respectively. We apply Jensen's inequality and instead bound

$$\|\mathbb{E}[X \mathbb{I}\{\|X\|_{op} \geq \beta\}]\| \leq \mathbb{E}[\|X\|_{op} \mathbb{I}\{\|X\|_{op} \geq \beta\}] = \int_0^\infty t \cdot \mathbb{I}\{t \geq \beta\} p(t) dt$$

since the indicator is 0 for $t \leq \beta$,

$$= \int_{\beta}^{\infty} (-t)(-p(t))dt$$

integrating by parts,

$$= -t \cdot (1 - P(t)) \Big|_{\beta}^{\infty} + \int_{\beta}^{\infty} (1 - P(t))dt$$

and using the equality of $1 - P(t)$ with $\mathbb{P}(\|X\|_{op} \geq t)$ along with (A.4),

$$\leq \beta n^{-\log n} + \int_{\beta}^{\infty} \mathbb{P}(\|X\|_{op} \geq t)dt$$

Applying the change of variables $t = \alpha s$ so as to apply (A.6),

$$\begin{aligned} &= \beta n^{-\log n} + \alpha \int_R^{\infty} \mathbb{P}(\|X\|_{op} \geq \alpha s)ds \\ &\leq \beta n^{-\log n} + \alpha \int_R^{\infty} \exp(-cs^{1/(2\ell)})ds \end{aligned}$$

Now applying a change of variables so $s = (\frac{u \log n}{c})^{2\ell}$,

$$\begin{aligned} &= \beta n^{-\log n} + \alpha \int_{\frac{cR^{1/(2\ell)}}{\log n}}^{\infty} n^{-u} \cdot 2\ell \left(\frac{\log n}{c}\right)^{2\ell} u^{2\ell-1} du \\ &\leq \beta n^{-\log n} + \alpha \int_{\frac{cR^{1/(2\ell)}}{\log n}}^{\infty} n^{-u/2} du, \end{aligned}$$

where we have used the assumption that ℓ is constant. We can approximate this by a geometric sum,

$$\begin{aligned} &\leq \beta n^{-\log n} + \alpha \sum_{u=\frac{cR^{1/(2\ell)}}{\log n}}^{\infty} n^{-u/2} \\ &\leq \beta n^{-\log n} + \alpha \cdot n^{-cR^{1/(2\ell)}/(2\log n)} \end{aligned}$$

Evaluating at $R = \text{polylog } n$ for a sufficiently large polynomial in the log gives us

$$\mathbb{E} [\|X \cdot \mathbb{I}\{\|X\|_{op} \geq \beta\}\|_{op}] \leq (\beta + \alpha)n^{-\log n},$$

as desired. □

B Concentration bounds for planted sparse vector in random linear subspace

Proof of Lemma 4.7. Let $c := \sum_{i=1}^n v(i)b_i$. The matrix in question has a nice block structure:

$$\sum_{i=1}^n a_i a_i^\top = \begin{pmatrix} \|v\|_2^2 & c^\top \\ c & \sum_{i=1}^n b_i b_i^\top \end{pmatrix}.$$

The vector c is distributed as $\mathcal{N}(0, \frac{1}{n} \text{Id}_{d-1})$ so by standard concentration has $\|c\| \leq \tilde{O}(d/n)^{1/2}$ w.ov.p.. By assumption, $\|v\|_2^2 = 1$. Thus by triangle inequality w.ov.p.

$$\left\| \sum_{i=1}^n a_i a_i^\top - \text{Id}_d \right\| \leq \tilde{O}\left(\frac{d}{n}\right)^{1/2} + \left\| \sum_{i=1}^n b_i b_i^\top - \text{Id}_{d-1} \right\|.$$

By [Ver10, Corollary 5.50] applied to the subgaussian vectors nb_i , w.ov.p.

$$\left\| \sum_{i=1}^n b_i b_i^\top - \text{Id}_{d-1} \right\| \leq O\left(\frac{d}{n}\right)^{1/2}$$

and hence $\left\| \sum_{i=1}^n a_i a_i^\top - \text{Id}_d \right\| \leq \tilde{O}(d/n)^{1/2}$ w.ov.p.. This implies $\left\| (\sum_{i=1}^n a_i a_i^\top)^{-1} - \text{Id}_d \right\| \leq \tilde{O}(d/n)^{1/2}$ and $\left\| (\sum_{i=1}^n a_i a_i^\top)^{-1/2} - \text{Id}_d \right\| \leq \tilde{O}(d/n)^{1/2}$ when $d = o(n)$ by the following facts applied to the eigenvalues of $\sum_{i=1}^n a_i a_i^\top$. For $0 \leq \varepsilon < 1$,

$$\begin{aligned} (1 + \varepsilon)^{-1} &= 1 - O(\varepsilon) & \text{and} & & (1 - \varepsilon)^{-1} &= 1 + O(\varepsilon), \\ (1 + \varepsilon)^{-1/2} &= 1 - O(\varepsilon) & \text{and} & & (1 - \varepsilon)^{-1/2} &= 1 + O(\varepsilon). \end{aligned}$$

These are proved easily via the identity $(1 + \varepsilon)^{-1} = \sum_{k=1}^{\infty} \varepsilon^k$ and similar. \square

Orthogonal subspace basis

Lemma B.1. *Let $a_1, \dots, a_n \in \mathbb{R}^d$ be independent random vectors from $\mathcal{N}(0, \frac{1}{n} \text{Id})$ with $d \leq n$ and let $A = \sum_{i=1}^n a_i a_i^\top$. Then for every unit vector $x \in \mathbb{R}^d$, with overwhelming probability $1 - d^{-\omega(1)}$,*

$$\left| \langle x, A^{-1}x \rangle - \|x\|^2 \right| \leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \|x\|^2.$$

Proof. Let $x \in \mathbb{R}^d$. By scale invariance, we may assume $\|x\| = 1$.

By standard matrix concentration bounds, the matrix $B = \text{Id} - A$ has spectral norm $\|B\| \leq \tilde{O}(d/n)^{1/2}$ w.ov.p. [Ver10, Corollary 5.50]. Since $A^{-1} = (\text{Id} - B)^{-1} = \sum_{k=0}^{\infty} B^k$, the spectral norm of $A^{-1} - \text{Id} - B$ is at most $\sum_{k=2}^{\infty} \|B\|^k$ (whenever the series converges). Hence, $\|A^{-1} - \text{Id} - B\| \leq \tilde{O}(d/n)$ w.ov.p..

It follows that it is enough to show that $|\langle x, Bx \rangle| \leq \tilde{O}(1/n)^{1/2}$ w.ov.p.. The random variable $n - n\langle x, Bx \rangle = \sum_{i=1}^n \langle \sqrt{n} \cdot a_i, x \rangle^2$ is χ^2 -distributed with n degrees of freedom. Thus, by standard concentration bounds, $n|\langle x, Bx \rangle| \leq \tilde{O}(\sqrt{n})$ w.ov.p. [LM00].

We conclude that with overwhelming probability $1 - d^{-\omega(1)}$,

$$\left| \langle x, A^{-1}x \rangle - \|x\|^2 \right| \leq |\langle x, Bx \rangle| + \tilde{O}(d/n) \leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right).$$

\square

Lemma B.2. Let $a_1, \dots, a_n \in \mathbb{R}^d$ be independent random vectors from $\mathcal{N}(0, \frac{1}{n} \text{Id})$ with $d \leq n$ and let $A = \sum_{i=1}^n a_i a_i^\top$. Then for every index $i \in [n]$, with overwhelming probability $1 - d^{\omega(1)}$,

$$\left| \langle a_j, A^{-1} a_j \rangle - \|a_j\|^2 \right| \leq \tilde{O} \left(\frac{d + \sqrt{n}}{n} \right) \cdot \|a_j\|^2.$$

Proof. Let $A_{-j} = \sum_{i \neq j} a_i a_i^\top$. By Sherman–Morrison,

$$A^{-1} = (A_{-j} + a_j a_j^\top)^{-1} = A_{-j}^{-1} - \frac{1}{1 + a_j^\top A_{-j}^{-1} a_j} A_{-j}^{-1} a_j a_j^\top A_{-j}^{-1}$$

Thus, $\langle a_j, A^{-1} a_j \rangle = \langle a_j, A_{-j}^{-1} a_j \rangle - \langle a_j, A_{-j}^{-1} a_j \rangle^2 / (1 + \langle a_j, A_{-j}^{-1} a_j \rangle)$. Since $\|\frac{n}{n-1} A_{-j} - \text{Id}\| = \tilde{O}(d/n)^{1/2}$ w.ov.p., we also have $\|A_{-j}^{-1}\| \leq 2$ with overwhelming probability. Therefore, w.ov.p.,

$$\left| \langle a_j, A^{-1} a_j \rangle - \langle a_j, A_{-j}^{-1} a_j \rangle \right| \leq \langle a_j, A_{-j}^{-1} a_j \rangle^2 \leq 4 \|a_j\|^4 \leq \tilde{O}(d/n) \cdot \|a_j\|^2.$$

At the same time, by [Lemma B.1](#), w.ov.p.,

$$\left| \langle a_j, \frac{n}{n-1} A_{-j}^{-1} a_j \rangle - \|a_j\|^2 \right| \leq \tilde{O} \left(\frac{d + \sqrt{n}}{n} \right) \cdot \|a_j\|^2.$$

We conclude that, w.ov.p.,

$$\begin{aligned} \left| \langle a_j, A^{-1} a_j \rangle - \|a_j\|^2 \right| &\leq \left| \langle a_j, A^{-1} a_j \rangle - \langle a_j, A_{-j}^{-1} a_j \rangle \right| + \left| \langle a_j, A_{-j}^{-1} a_j \rangle - \frac{n-1}{n} \|a_j\|^2 \right| + \frac{1}{n} \|a_j\|^2 \\ &\leq \tilde{O} \left(\frac{d + \sqrt{n}}{n} \right). \end{aligned}$$

□

Lemma B.3. Let A be a block matrix where one of the diagonal blocks is the 1×1 identity; that is,

$$A = \begin{pmatrix} \|v\|^2 & c^\top \\ c & B \end{pmatrix} = \begin{pmatrix} 1 & c^\top \\ c & B \end{pmatrix}.$$

for some matrix B and vector c . Let x be a vector which decomposes as $x = (x(1) \ x')$ where $x(1) = \langle x, e_1 \rangle$ for e_1 the first standard basis vector.

Then

$$\langle x, A^{-1} x \rangle = \langle x', \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) x' \rangle + 2x(1) \langle \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) c, x' \rangle + (1 - c^\top B^{-1} c)^{-1} x(1)^2.$$

Proof. By the formula for block matrix inverses,

$$A^{-1} = \begin{pmatrix} (1 - c^\top B^{-1} c)^{-1} & c^\top (B - c c^\top)^{-1} \\ (B - c c^\top)^{-1} c & (B - c c^\top)^{-1} \end{pmatrix}.$$

The result follows by Sherman–Morrison applied to $(B - c c^\top)^{-1}$ and the definition of x . □

Lemma B.4. Let $v \in \mathbb{R}^n$ be a unit vector and let $b_1, \dots, b_n \in \mathbb{R}^{d-1}$ have iid entries from $\mathcal{N}(0, 1/n)$. Let $a_i \in \mathbb{R}^d$ be given by $a_i := (v(i) b_i)$. Let $A := \sum_i a_i a_i^\top$. Let $c \in \mathbb{R}^{d-1}$ be given by $c := \sum_i v(i) b_i$. Then for every index $i \in [n]$, w.ov.p.,

$$|\langle a_i, A^{-1} a_i \rangle - \|a_i\|^2| \leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \|a_i\|^2.$$

Proof. Let $B := \sum_i b_i b_i^\top$. By standard concentration, $\|B^{-1} - \text{Id}\| \leq \tilde{O}(d/n)^{1/2}$ w.ov.p. [Ver10, Corollary 5.50]. At the same time, since v has unit norm, the entries of c are iid samples from $\mathcal{N}(0, 1/n)$, and hence $n\|c\|^2$ is χ^2 -distributed with d degrees of freedom. Thus w.ov.p. $\|c\|^2 \leq \frac{d}{n} + \tilde{O}(dn)^{-1/2}$. Together these imply the following useful estimates, all of which hold w.ov.p.:

$$\begin{aligned} |c^\top B^{-1} c| &\leq \|c\|^2 \|B^{-1}\|_{op} \leq \frac{d}{n} + \tilde{O}\left(\frac{d}{n}\right)^{3/2} \\ \|B^{-1} c c^\top B^{-1}\|_{op} &\leq \|c\|^2 \|B^{-1}\|_{op}^2 \leq \frac{d}{n} + \tilde{O}\left(\frac{d}{n}\right)^{3/2} \\ \left\| \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right\|_{op} &\leq \frac{d}{n} + \tilde{O}\left(\frac{d}{n}\right)^{3/2}, \end{aligned}$$

where the first two use Cauchy-Schwarz and the last follows from the first two.

We turn now to the expansion of $\langle a_i, A^{-1} a_i \rangle$ offered by Lemma B.3,

$$\langle a_i, A^{-1} a_i \rangle = \langle b_i, \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) b_i \rangle \quad (\text{B.1})$$

$$+ 2v(i) \left\langle \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) c, b_i \right\rangle \quad (\text{B.2})$$

$$+ (1 - c^\top B^{-1} c)^{-1} v(i)^2. \quad (\text{B.3})$$

Addressing (B.1) first, by the above estimates and Lemma B.2 applied to $\langle b_i, B^{-1} b_i \rangle$,

$$\left| \langle b_i, \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) b_i \rangle - \|b_i\|^2 \right| \leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \|b_i\|^2$$

w.ov.p.. For (B.2), we pull out the important factor of $\|c\|$ and separate $v(i)$ from b_i : w.ov.p.,

$$\begin{aligned} \left| 2v(i) \left\langle \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) c, b_i \right\rangle \right| &= \left| 2\|c\| v(i) \left\langle \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) \frac{c}{\|c\|}, b_i \right\rangle \right| \\ &\leq \left| \|c\|^2 \left(v(i)^2 + \left\langle \left(B^{-1} + \frac{B^{-1} c c^\top B^{-1}}{1 - c^\top B^{-1} c} \right) \frac{c}{\|c\|}, b_i \right\rangle^2 \right) \right| \\ &\leq \tilde{O}\left(\frac{d}{n}\right) (v(i)^2 + \|b_i\|^2) \\ &= \tilde{O}\left(\frac{d}{n}\right) \|a_i\|^2, \end{aligned}$$

where the last inequality follows from our estimates above and Cauchy-Schwarz.

Finally, for (B.3), since $(1 - c^\top B^{-1}c) \geq 1 - \tilde{O}(d/n)$ w.ov.p., we have that

$$|(1 - c^\top B^{-1}c)^{-1}v(i)^2 - v(i)^2| \leq \tilde{O}\left(\frac{d}{n}\right)v(i)^2.$$

Putting it all together,

$$\begin{aligned} |\langle a_i, A^{-1}a_i \rangle - \|a_i\|^2| &\leq \left| \langle b_i, \left(B^{-1} + \frac{B^{-1}cc^\top B^{-1}}{1 - c^\top B^{-1}c} \right) b_i \rangle - \|b_i\|^2 \right| \\ &\quad + \left| 2v(i) \langle \left(B^{-1} + \frac{B^{-1}cc^\top B^{-1}}{1 - c^\top B^{-1}c} \right) c, b_i \rangle \right| \\ &\quad + |(1 - c^\top B^{-1}c)^{-1}v(i)^2 - v(i)^2| \\ &\leq \tilde{O}\left(\frac{d + \sqrt{n}}{n}\right) \cdot \|a_i\|^2. \quad \square \end{aligned}$$

C Concentration bounds for overcomplete tensor decomposition

We require some facts about the concentration of certain scalar-and matrix-valued random variables, which generally follow from standard concentration arguments. We present proofs here for completeness.

The first lemma captures standard facts about random Gaussians.

Fact C.1. Let $a_1, \dots, a_n \in \mathbb{R}^d$ be sampled $a_i \sim \mathcal{N}(0, \frac{1}{d} \text{Id})$.

1. Inner products $|\langle a_i, a_j \rangle|$ are all $\approx 1/\sqrt{d}$:

$$\mathbb{P} \left\{ \langle a_i, a_j \rangle^2 \leq \tilde{O}\left(\frac{1}{d}\right) \mid \forall i, j \in [n], i \neq j \right\} \geq 1 - n^{-\omega(1)}.$$

2. Norms are all about $\|a_i\| \approx 1 \pm \tilde{O}(1/\sqrt{d})$:

$$\mathbb{P} \left\{ 1 - \tilde{O}(1/\sqrt{d}) \leq \|a_i\|_2^2 \leq 1 + \tilde{O}(1/\sqrt{d}) \mid \forall i \in [n] \right\} \geq 1 - n^{-\omega(1)}.$$

3. Fix a vector $v \in \mathbb{R}^d$. Suppose $g \in \mathbb{R}^d$ is a vector with entries identically distributed $g_i \sim \mathcal{N}(0, \sigma)$. Then $\langle g, v \rangle^2 \approx \sigma^2 \cdot \|v\|_2^2$:

$$\mathbb{P} \left\{ \left| \langle g, v \rangle^2 - \sigma^2 \cdot \|v\|_2^4 \right| \leq \tilde{O}(\sigma^2 \cdot \|v\|_4^2) \right\} \geq 1 - n^{-\omega(1)}.$$

Proof of Fact C.1. We start with [Item 1](#). Consider the quantity $\langle a_i, a_j \rangle^2$. We calculate the expectation,

$$\mathbb{E}[\langle a_i, a_j \rangle^2] = \sum_{k, \ell \in [d]} \mathbb{E}[a_i(k)a_i(\ell)a_j(k)a_j(\ell)] = \sum_{k \in [d]} \mathbb{E}[a_i(k)^2] \cdot \mathbb{E}[a_j(k)^2] = d \cdot \frac{1}{d^2} = \frac{1}{d}.$$

Since this is a degree-4 square polynomial in the entries of a_i and a_j , we may apply [Lemma A.5](#) to conclude that

$$\mathbb{P} \left(\langle a_i, a_j \rangle^2 \geq t \cdot \frac{1}{d} \right) \leq \exp(-O(t^{1/2})).$$

Applying this fact with $t = \text{polylog}(n)$ and taking a union bound over pairs $i, j \in [n]$ gives us the desired result.

Next is [Item 2](#). Consider the quantity $\|a_i\|_2^2$. We will apply [Lemma A.5](#) in order to obtain a tail bound for the value of the polynomial $(\|a_i\|_2^2 - 1)^2$. We have

$$\mathbb{E}[(\|a_i\|_2^2 - 1)^2] = O\left(\frac{1}{d}\right),$$

and now applying [Lemma A.5](#) with the square root of this expectation, we have

$$\mathbb{P}\left(\left|\|a_i\|_2^2 - 1\right| \geq \tilde{O}\left(\frac{1}{\sqrt{d}}\right)\right) \leq n^{-\log n}.$$

This gives both bounds for a single a_i . The result now follows from taking a union bound over all i .

Moving on to [Item 3](#), we view the expression $f(g) := (\langle g, v \rangle^2 - \sigma^2 \|v\|^2)^2$ as a polynomial in the gaussian entries of g . The degree of $f(g)$ is 4, and $\mathbb{E}[|f(g)|] = 3\sigma^4 \cdot \|v\|_4^4$, and so we may apply [Lemma A.5](#) to conclude that

$$\mathbb{P}(|f(g)| \geq t \cdot 3\sigma^4 \cdot \|v\|_4^4) \leq \exp(-c_4 t^{1/2}),$$

and taking $t = \text{polylog}(n)$ the conclusion follows. \square

We also use the fact that the covariance matrix of a sum of sufficiently many gaussian outer products concentrates about its expectation.

Fact C.2. Let $a_1, \dots, a_n \in \mathbb{R}^d$ be vectors with iid gaussian entries such that $\mathbb{E}[\|a_i\|_2^2] = 1$, and $n = \Omega(d)$. Let \mathcal{E} be the event that the sum $\sum_{i \in [n]} a_i a_i^\top$ is close to $\frac{n}{d} \cdot \text{Id}$, that is

$$\mathbb{P}\left\{\tilde{\Omega}(n/d) \cdot \text{Id} \leq \sum_{i \in [n]} a_i a_i^\top \leq \tilde{O}(n/d) \cdot \text{Id}\right\} \geq 1 - n^{-\omega(1)}.$$

Proof of [Fact C.2](#). We apply a truncated matrix bernstein inequality. For convenience, $A := \sum_{i \in [n]} a_i a_i^\top$ and let $A_i := a_i a_i^\top$ be a single summand. To begin, we calculate the first and second moments of the summands,

$$\begin{aligned} \mathbb{E}[A_i] &= \frac{1}{d} \cdot \text{Id} \\ \mathbb{E}[A_i A_i^\top] &= O\left(\frac{1}{d}\right) \cdot \text{Id}. \end{aligned}$$

So we have $\mathbb{E}[A] = \frac{n}{d} \cdot \text{Id}$ and $\sigma^2(A) = O\left(\frac{n}{d}\right)$.

We now show that each summand is well-approximated by a truncated variable. To calculate the expected norm $\|A_i\|_{op}$, we observe that A_i is rank-1 and thus $\mathbb{E}[\|A_i\|_{op}] = \mathbb{E}[\|a_i\|_2^2] = 1$. Applying [Lemma A.8](#), we have

$$\mathbb{P}(\|A_i\|_{op} \geq \tilde{O}(1)) \leq n^{-\log n},$$

and also

$$\mathbb{E}[\|A_i\|_{op} \cdot \mathbb{I}\{\|A_i\|_{op} \geq \tilde{O}(1)\}] \leq n^{-\log n}.$$

Thus, applying the truncated matrix Bernstein inequality from [Proposition A.7](#) with $\sigma^2 = O(\frac{n}{d})$, $\beta = \tilde{O}(1)$, $p = n^{-\log n}$, $q = n^{-\log n}$, and $t = \tilde{O}\left(\frac{n^{1/2}}{d^{1/2}}\right)$, we have that with overwhelming probability,

$$\left\|A - \frac{n}{d} \cdot \text{Id}\right\|_{op} \leq \tilde{O}\left(\frac{n^{1/2}}{d^{1/2}}\right).$$

□

We now show that among the terms of the polynomial $\langle g, T a_i^{\otimes 2} \rangle$, those that depend on a_j with $j \neq i$ have small magnitude. This polynomial appears in the proof that M_{diag} has a noticeable spectral gap.

Lemma (Restatement of [Lemma 5.6](#)). *Let a_1, \dots, a_n be independently sampled vectors from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$, and let g be sampled from $\mathcal{N}(0, \text{Id}_d)$. Let $T = \sum_i a_i (a_i \otimes a_i)^\top$. Then with overwhelming probability, for every $j \in [n]$,*

$$|\langle g, T(a_j \otimes a_j) \rangle - \langle g, a_j \rangle \|a_j\|^4| \leq \tilde{O}\left(\frac{\sqrt{n}}{d}\right).$$

Proof. Fixing a_i and g , the terms in the summation are independent, and we may apply a Bernstein inequality. A straightforward calculation shows that the expectation of the sum is 0 and the variance is $\tilde{O}(\frac{n}{d^2}) \cdot \|g\|^2 \|a_i\|^4$. Additionally, each summand is a polynomial in Gaussian variables, the square of which has expectation $\tilde{O}(\frac{1}{d^2}) \cdot \|g\|^2 \|a_i\|^4$. Thus [Lemma A.5](#) allows us to truncate each summand appropriately so as to employ [Proposition A.7](#). An appropriate choice of logarithmic factors and the concentration of $\|g\|^2$ and $\|a_i\|^2$ due to [Fact C.1](#) gives the result for each $i \in [n]$. A union bound over each choice of i gives the final result. □

Finally, we prove that a matrix which appears in the expression for M_{same} has bounded norm w.ov.p.

Lemma C.3. *Let a_1, \dots, a_n be independent from $\mathcal{N}(0, \frac{1}{d} \text{Id}_d)$. Let $g \sim \mathcal{N}(0, \text{Id}_d)$. Fix $j \in [n]$. Then w.ov.p.*

$$\left\| \sum_{\substack{i \in [n] \\ i \neq j}} \langle g, a_i \rangle \|a_i\|^2 \langle a_i, a_j \rangle \cdot a_i a_i^\top \right\| \leq \tilde{O}(n/d^2)^{1/2}.$$

Proof. The proof proceeds by truncated matrix Bernstein, since the summands are independent for fixed g, a_j . For this we need to compute the variance:

$$\sigma^2 = \left\| \sum_{\substack{i \in [n] \\ i \neq j}} \mathbb{E} \langle g, a_i \rangle^2 \|a_i\|^6 \langle a_i, a_j \rangle^2 \cdot a_i a_i^\top \right\| \leq O(1/d) \cdot \left\| \sum_{\substack{i \in [n] \\ i \neq j}} \mathbb{E} a_i a_i^\top \right\| \leq O(1/d) \cdot n/d \leq O(n/d^2).$$

The norm of each term in the sum is bounded by a constant-degree polynomial of Gaussians. Straightforward calculations show that in expectation each term is $O(\frac{1}{d} \langle g, a_i \rangle)$ in norm; w.ov.p. this

is $O(\sigma)$. So [Lemma A.5](#) applies to establish the hypothesis of truncated Bernstein [Proposition A.7](#). In turn, [Proposition A.7](#) yields that w.ov.p.

$$\left\| \sum_{\substack{i \in [n] \\ i \neq j}} \langle g, a_i \rangle \|a_i\|^2 \langle a_i, a_j \rangle \cdot a_i a_i^\top \right\| \leq \tilde{O}(\sigma) = \tilde{O}(n/d^2)^{1/2}.$$

□

C.0.1 Proof of [Fact C.4](#)

Here we prove the following fact.

Fact C.4. Let $\Sigma = \mathbb{E}_{x \sim \mathcal{N}(0, \text{Id}_d)}(xx^\top)^{\otimes 2}$ and let $\tilde{\Sigma} = \mathbb{E}_{x \sim \mathcal{N}(0, \text{Id}_d)}(xx^\top)^{\otimes 2} / \|x\|^4$. Let $\Phi = \sum_i e_i^{\otimes 2} \in \mathbb{R}^{d^2}$ and let Π_{sym} be the projector to the symmetric subspace of \mathbb{R}^{d^2} (the span of vectors of the form $x^{\otimes 2}$ for $x \in \mathbb{R}^d$). Then

$$\begin{aligned} \Sigma &= 2 \Pi_{\text{sym}} + \Phi \Phi^\top, & \tilde{\Sigma} &= \frac{2}{d^2+2d} \Pi_{\text{sym}} + \frac{1}{d^2+2d} \Phi \Phi^\top, \\ \Sigma^+ &= \frac{1}{2} \Pi_{\text{sym}} - \frac{1}{2(d+2)} \Phi \Phi^\top, & \tilde{\Sigma}^+ &= \frac{d^2+2d}{2} \Pi_{\text{sym}} - \frac{d}{2} \Phi \Phi^\top. \end{aligned}$$

In particular,

$$R = \sqrt{2} (\Sigma^+)^{1/2} = \Pi_{\text{sym}} - \frac{1}{d} \left(1 - \sqrt{\frac{2}{d+2}}\right) \Phi \Phi^\top \quad \text{has} \quad \|R\| = 1$$

and for any $v \in \mathbb{R}^d$,

$$\|R(v \otimes v)\|_2^2 = \left(1 - \frac{1}{d+2}\right) \cdot \|v\|^4.$$

We will derive [Fact C.4](#) as a corollary of a more general claim about rotationally symmetric distributions.

Lemma C.5. Let \mathcal{D} be a distribution over \mathbb{R}^d which is rotationally symmetric; that is, for any rotation R , $x \sim \mathcal{D}$ is distributed identically to Rx . Let $\Sigma = \mathbb{E}_{x \sim \mathcal{D}}(xx^\top)^{\otimes 2}$, let $\Phi = \sum_i e_i^{\otimes 2} \in \mathbb{R}^{d^2}$ and let Π_{sym} be the projector to the symmetric subspace of \mathbb{R}^{d^2} (the span of vectors of the form $x^{\otimes 2}$ for $x \in \mathbb{R}^d$). Then there is a constant r so that

$$\Sigma = 2r \Pi_{\text{sym}} + r \Phi \Phi^\top.$$

Furthermore, r is given by

$$r = \mathbb{E}\langle x, a \rangle^2 \langle x, b \rangle^2 = \frac{1}{3} \mathbb{E}\langle x, a \rangle^4$$

where a, b are orthogonal unit vectors.

Proof. First, Σ is symmetric and operates nontrivially only on the symmetric subspace (in other words $\ker \Pi_{\text{sym}} \subseteq \ker \Sigma$). This follows from Σ being an expectation over symmetric matrices whose kernels always contain the complement of the symmetric subspace.

Let $\hat{a}, \hat{b}, \hat{c}, \hat{d} \in \mathbb{R}^d$ be any four orthogonal unit vectors. Let R be any rotation of \mathbb{R}^d that takes \hat{a} to $-\hat{a}$, but fixes \hat{b}, \hat{c} , and \hat{d} (this rotation exists for $d \geq 5$, but a different argument holds for $d \leq 4$). By rotational symmetry about R , all of these quantities are 0:

$$\mathbb{E}\langle \hat{a}, x \rangle \langle \hat{b}, x \rangle \langle \hat{c}, x \rangle \langle \hat{d}, x \rangle = 0,$$

$$\mathbb{E}\langle \hat{a}, x \rangle \langle \hat{b}, x \rangle \langle \hat{c}, x \rangle^2 = 0, \quad \mathbb{E}\langle \hat{a}, x \rangle \langle \hat{b}, x \rangle^3 = 0.$$

Furthermore, let Q be a rotation of \mathbb{R}^d that takes \hat{a} to $(\hat{a} + \hat{b})/\sqrt{2}$. Then by rotational symmetry about Q ,

$$\mathbb{E}\langle \hat{a}, x \rangle^4 = \mathbb{E}\langle \hat{a}, Qx \rangle^4 = \mathbb{E}\langle \frac{1}{\sqrt{2}}(\hat{a} + \hat{b}), x \rangle^4 = \mathbb{E}\frac{1}{4}[\langle \hat{a}, x \rangle^4 + \langle \hat{b}, x \rangle^4 + 6\langle \hat{a}, x \rangle^2 \langle \hat{b}, x \rangle^2]$$

Thus, since $\mathbb{E}\langle \hat{a}, x \rangle^4 = \mathbb{E}\langle \hat{b}, x \rangle^4$ by rotational symmetry, we have

$$\mathbb{E}\langle \hat{a}, x \rangle^4 = 3\mathbb{E}\langle \hat{a}, x \rangle^2 \langle \hat{b}, x \rangle^2.$$

So let $r := \mathbb{E}\langle \hat{a}, x \rangle^2 \langle \hat{b}, x \rangle^2 = \frac{1}{3}\mathbb{E}\langle \hat{a}, x \rangle^4$. By rotational symmetry, r is constant over choice of orthogonal unit vectors \hat{a} and \hat{b} .

Since Σ operates only on the symmetric subspace, let $u \in \mathbb{R}^{d^2}$ be any unit vector in the symmetric subspace. Such a u unfolds to a symmetric matrix in $\mathbb{R}^{d \times d}$, so that it has an eigendecomposition $u = \sum_{i=1}^d \lambda_i u_i \otimes u_i$. Evaluating $\langle u, \Sigma u \rangle$,

$$\begin{aligned} \langle u, \Sigma u \rangle &= \sum_{i,j=1}^d \mathbb{E} \lambda_i \lambda_j \langle x, u_i \rangle^2 \langle x, u_j \rangle^2 \quad \text{other terms are 0 by above} \\ &= 3r \sum_{i=1}^d \lambda_i^2 + r \sum_{i \neq j} \lambda_i \lambda_j \\ &= 2r \sum_{i=1}^d \lambda_i^2 + r \left(\sum_{i=1}^d \lambda_i \right)^2 \\ &= 2r \|u\|^2 + r \left(\sum_{i=1}^d \lambda_i \right)^2 \quad \text{Frobenious norm is sum of squared eigenvalues} \\ &= 2r \|u\|^2 + r \left(\sum_i u_{i,i} \right)^2 \quad \text{trace is sum of eigenvalues} \\ &= 2r \langle u, \Pi_{\text{sym}} u \rangle + r \langle u, \Phi \Phi^\top u \rangle, \end{aligned}$$

so therefore $\Sigma = 2r \Pi_{\text{sym}} + r \Phi \Phi^\top$. □

Proof of Fact C.4. When $x \sim \mathcal{N}(0, \text{Id}_d)$, the expectation $\mathbb{E}\langle x, a \rangle^2 \langle x, b \rangle^2 = 1$ is just a product of independent standard Gaussian second moments. Therefore by Lemma C.5, $\Sigma = 2\Pi_{\text{sym}} + \Phi \Phi^\top$.

To find $\tilde{\Sigma}$ where x is uniformly distributed on the unit sphere, we compute

$$1 = \mathbb{E} \|x\|^4 = \sum_{i,j} \mathbb{E} x_i^2 x_j^2 = d \mathbb{E} x_1^4 + (d^2 - d) \mathbb{E} x_1^2 x_2^2$$

and use the fact that $\mathbb{E} x_1^4 = 3 \mathbb{E} x_1^2$ (by Lemma C.5) to find that $\mathbb{E} x_1^2 x_2^2 = \frac{1}{d^2+2d}$, and therefore by Lemma C.5, $\tilde{\Sigma} = \frac{2}{d^2+2d} \Pi_{\text{sym}} + \frac{1}{d^2+2d} \Phi \Phi^\top$.

To verify the pseudoinverses, it is enough to check that $MM^+ = \Pi_{\text{sym}}$ for each matrix M and its claimed pseudoinverse M^+ .

To show that

$$\|R(v \otimes v)\|_2^2 = \left(1 - \frac{1}{d+2}\right) \cdot \|v\|^4,$$

for any $v \in \mathbb{R}^d$, we write $\|R(v \otimes v)\|_2^2 = (v \otimes v)^\top R^2(v \otimes v)$ and use the substitution $R^2 = 2\Sigma^+$, along with the facts that $\Pi_{\text{sym}}(v \otimes v) = v \otimes v$ and $\langle \Phi, v \otimes v \rangle = \|v\|^2$.

□

Now we can prove some concentration claims we deferred:

Lemma (Restatement of [Lemma 5.11](#)). *Let $a_1, \dots, a_n \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$. Let Σ, R be as in [Fact 5.8](#). Let $u_i = a_i \otimes a_i$. With overwhelming probability, every $j \in [n]$ satisfies $\sum_{i \neq j} \langle u_j, R^2 u_i \rangle^2 = \tilde{O}(n/d^2)$ and $|1 - \|Ru_j\|^2| \leq \tilde{O}(1/\sqrt{d})$.*

Proof of [Lemma 5.11](#). We prove the first item:

$$\begin{aligned} \sum_{i \neq j} \langle u_j, R^2 u_i \rangle^2 &= \sum_{i \neq j} \langle u_j, 2\Sigma^+ u_i \rangle^2 \\ &= \sum_{i \neq j} \langle u_j, (\Pi_{\text{sym}} - \frac{1}{d+2} \Phi \Phi^\top) u_i \rangle^2 \quad \text{by [Fact C.4](#)} \\ &= \sum_{i \neq j} (\langle a_j, a_i \rangle^2 - \frac{1}{d+2} \|u_j\|^2 \|u_i\|^2)^2 \\ &= \sum_{i \neq j} \tilde{O}(1/d)^2 \quad \text{w.ov.p. by [Fact C.1](#)} \\ &= \tilde{O}(n/d^2). \end{aligned}$$

And one direction of the second item, using [Fact C.4](#) and [Fact C.1](#) (the other direction is similar):

$$\|Ru_j\|^2 = \langle u_j, R^2 u_j \rangle = \langle u_j, (\Pi_{\text{sym}} + \frac{1}{d+2} \Phi \Phi^\top) u_j \rangle = (1 - \Theta(1/d)) \|a_j\|^4 = 1 - \tilde{O}(1/\sqrt{d})$$

where the last equality holds w.ov.p..

□

C.0.2 Proof of [Lemma 5.9](#)

To prove [Lemma 5.9](#) we will begin by reducing to the case $S = [n]$ via the following.

Lemma C.6. *Let $v_1, \dots, v_n \in \mathbb{R}^d$. Let A_S have columns $\{v_i\}_{i \in S}$. Let Π_S be the projector to $\text{Span}\{v_i\}_{i \in S}$. Suppose there is $c \geq 0$ so that $\|A_{[n]}^\top A_{[n]} - \text{Id}_n\| \leq c$. Then for every $S \subseteq [n]$, $\|A_S A_S^\top - \Pi_S\| \leq c$*

Proof. If the hypothesized bound $\|A_{[n]}^\top A_{[n]} - \text{Id}_n\| \leq c$ holds then for every $S \subseteq [n]$ we get $\|A_S^\top A_S - \text{Id}_{|S|}\| \leq c$ since $A_S^\top A_S$ is a principal submatrix of $A_{[n]}^\top A_{[n]}$. If $\|A_S^\top A_S - \text{Id}_{|S|}\| \leq c$, then because $A_S A_S^\top$ has the same nonzero eigenvalues as $A_S^\top A_S$, we must have also $\|A_S A_S^\top - \Pi_S\| \leq c$. □

It will be convenient to reduce concentration for matrices involving $a_i \otimes a_i$ to analogous matrices where the vectors $a_i \otimes a_i$ are replaced by isotropic vectors of constant norm. The following lemma shows how to do this.

Lemma C.7. *Let $a \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$. Let $\tilde{\Sigma} := \mathbb{E}_{x \sim \mathcal{N}(0, \text{Id}_d)} (xx^\top)^{\otimes 2} / \|x\|^4$. Then $u := (\tilde{\Sigma}^+)^{1/2} a \otimes a / \|a\|^2$ is an isotropic random vector in the symmetric subspace $\text{Span}\{y \otimes y \mid y \in \mathbb{R}^d\}$ with $\|u\| = \sqrt{\dim \text{Span}\{y \otimes y \mid y \in \mathbb{R}^d\}}$.*

Proof. The vector u is isotropic by definition so we prove the norm claim. Let $\tilde{\Phi} = \Phi/\|\Phi\|$. By [Fact C.4](#),

$$\tilde{\Sigma}^+ = \frac{d^2+2d}{2}\Pi_{\text{sym}} - \frac{d}{2}\Phi\Phi^\top$$

Thus,

$$\|u\|^2 = \left\langle \frac{a \otimes a}{\|a\|^2}, \tilde{\Sigma}^+ \frac{a \otimes a}{\|a\|^2} \right\rangle = \frac{d^2+2d}{2} - \frac{d}{2} = \frac{d^2+d}{2} = \dim \text{Span}\{y \otimes y \mid y \in \mathbb{R}^d\}. \quad \square$$

The last ingredient to finish the spectral bound is a bound on the incoherence of independent samples from $(\tilde{\Sigma}^+)^{1/2}$.

Lemma C.8. *Let $\tilde{\Sigma} = \mathbb{E}_{a \sim \mathcal{N}(0, \text{Id}_d)}(aa^\top \otimes aa^\top)/\|a\|^4$. Let $a_1, \dots, a_n \sim \mathcal{N}(0, \text{Id}_d)$ be independent, and let $u_i = (\tilde{\Sigma}^+)^{1/2}(a_i \otimes a_i)/\|a_i\|^2$. Let $d' = \dim \text{Span}\{y \otimes y \mid y \in \mathbb{R}^d\} = \frac{1}{2}(d^2 + d)$. Then*

$$\frac{1}{d'} \mathbb{E} \max_i \sum_{j \neq i} \langle u_i, u_j \rangle^2 \leq \tilde{O}(n).$$

Proof. Expanding $\langle u_i, u_j \rangle^2$ and using $\tilde{\Sigma}^+ = \frac{d^2+2d}{2}\Pi_{\text{sym}} - \frac{d}{2}\Phi\Phi^\top$, we get

$$\langle u_i, u_j \rangle^2 = \left(\frac{d^2+2d}{2} \left\langle \frac{a_i \otimes a_i}{\|a_i\|^2}, \frac{a_j \otimes a_j}{\|a_j\|^2} \right\rangle - \frac{d}{2} \right)^2 = \left(\frac{d^2+2d}{2} \cdot \frac{\langle a_i, a_j \rangle^2}{\|a_i\|^2 \|a_j\|^2} - \frac{d}{2} \right)^2$$

From elementary concentration, $\mathbb{E} \max_{i \neq j} \langle a_i, a_j \rangle^2 / \|a_i\|^2 \|a_j\|^2 \leq \tilde{O}(1/d)$, so the lemma follows by elementary manipulations. \square

We need the following bound on the deviation from expectation of a tall matrix with independent columns.

Theorem C.9 (Theorem 5.62 in [\[Ver10\]](#)). *Let A be an $N \times n$ matrix ($N \geq n$) whose columns A_j are independent isotropic random vectors in \mathbb{R}^N with $\|A_j\|_2 = \sqrt{N}$ almost surely. Consider the incoherence parameter*

$$m \stackrel{\text{def}}{=} \frac{1}{N} \mathbb{E} \max_{i \in [n]} \sum_{j \neq i} \langle A_i, A_j \rangle^2.$$

Then $\mathbb{E} \left\| \frac{1}{N} A^T A - \text{Id} \right\| \leq C_0 \sqrt{\frac{m \log n}{N}}$.

We are now prepared to handle the case of $S = [n]$ via spectral concentration for matrices with independent columns, [Theorem C.9](#).

Lemma (Restatement of [Lemma 5.9](#)). *Let $a_1, \dots, a_n \sim \mathcal{N}(0, \frac{1}{d} \text{Id}_d)$ be independent random vectors with $d \leq n$. Let $R := \sqrt{2} \cdot ((\mathbb{E}(aa^\top)^{\otimes 2})^+)^{1/2}$ for $a \sim \mathcal{N}(0, \text{Id}_d)$. For $S \subseteq [n]$, let $P_S = \sum_{i \in S} (a_i a_i^\top)^{\otimes 2}$ and let Π_S be the projector into the subspace spanned by $\{R a_i^{\otimes 2} \mid i \in S\}$. Then, with probability $1 - o(1)$ over the choice of a_1, \dots, a_n ,*

$$\forall S \subseteq [n]. \quad (1 - \tilde{O}(n/d^{3/2})) \cdot \Pi_S \leq R P_S R \leq (1 + \tilde{O}(n/d^{3/2})) \cdot \Pi_S.$$

Proof of Lemma 5.9. By [Lemma C.6](#) it is enough to prove the lemma in the case of $S = [n]$. For this we will use [Theorem C.9](#). Let A be the matrix whose columns are given by $a_i \otimes a_i$, so that $P_{[n]} = P = AA^\top$. Because $RAA^\top R$ and $A^\top RRA$ have the same nonzero eigenvalues, it will be

enough to show that $\|A^\top R^2 A - \text{Id}\| \leq \tilde{O}(\sqrt{n}/d) + \tilde{O}(n/d^{3/2})$ with probability $1 - o(1)$. (Since $n \leq d$ we have $\sqrt{n}/d = \tilde{O}(n/d^{3/2})$ so this gives the theorem.)

The columns of RA are independent, given by $R(a_i \otimes a_i)$. However, they do not quite satisfy the normalization conditions needed for [Theorem C.9](#). Let D be the diagonal matrix whose i -th diagonal entry is $\|a_i\|^2$. Let $\tilde{\Sigma} = \mathbb{E}_{x \sim \mathcal{N}(0, \text{Id})}(xx^\top)^{\otimes 2} / \|x\|^4$. Then by [Lemma C.7](#) the matrix $(\tilde{\Sigma}^+)^{1/2} D^{-1} A$ has independent columns from an isotropic distribution with a fixed norm d' . Together with [Lemma C.8](#) this is enough to apply [Theorem C.9](#) to conclude that $\mathbb{E} \|\frac{1}{(d')^2} A^\top D^{-1} \tilde{\Sigma}^+ D^{-1} A - \text{Id}\| \leq \tilde{O}(\sqrt{n}/d)$. By Markov's inequality, $\|\frac{1}{(d')^2} A^\top D^{-1} \tilde{\Sigma}^+ D^{-1} A - \text{Id}\| \leq \tilde{O}(\sqrt{n}/d)$ with probability $1 - o(1)$.

We will show next that $\|A^\top R^2 A - \frac{1}{(d')^2} A^\top D^{-1} \tilde{\Sigma}^+ D^{-1} A\| \leq \tilde{O}(n/d^{3/2})$ with probability $1 - o(1)$; the lemma then follows by triangle inequality. The expression inside the norm expands as

$$A^\top (R^2 - \frac{1}{(d')^2} D^{-1} \tilde{\Sigma}^+ D^{-1}) A.$$

and so

$$\|A^\top R^2 A - \frac{1}{(d')^2} A^\top D^{-1} \tilde{\Sigma}^+ D^{-1} A\| \leq \|A\|^2 \|R^2 - \frac{1}{(d')^2} D^{-1} \tilde{\Sigma}^+ D^{-1}\|$$

By [Fact C.1](#), with overwhelming probability $\|D - \text{Id}\| \leq \tilde{O}(1/\sqrt{d})$. So $\|(1/d')^2 D^{-1} \tilde{\Sigma}^+ D^{-1} - (1/d')^2 \tilde{\Sigma}^+\| \leq \tilde{O}(1/\sqrt{d})$ w.ov.p.. We recall from [Fact C.4](#), given that $R = \sqrt{2} \cdot (\Sigma^+)^{1/2}$, that

$$R^2 = \Pi_{\text{sym}} - \frac{1}{d+2} \Phi \Phi^\top \quad \text{and} \quad \frac{1}{(d')^2} \tilde{\Sigma}^+ = \frac{d+2}{d+1} \Pi_{\text{sym}} - \frac{1}{d+1} \Phi \Phi^\top.$$

This implies that $\|R^2 - (1/d')^2 \tilde{\Sigma}^+\| \leq O(1/d)$. Finally, by an easy application of [Proposition A.7](#), $\|A\|^2 = \|\sum_i (a_i a_i^\top)^{\otimes 2}\| \leq \tilde{O}(n/d)$ w.ov.p.. All together, $\|A^\top R^2 A - \frac{1}{(d')^2} A^\top D^{-1} \tilde{\Sigma}^+ D^{-1} A\| \leq \tilde{O}(n/d^{3/2})$. \square

D Concentration bounds for tensor principal component analysis

For convenience, we restate [Lemma 6.5](#) here.

Lemma D.1 (Restatement of [Lemma 6.5](#)). *For any v , with high probability over \mathbf{A} , the following occur:*

$$\begin{aligned} \left\| \sum_i \text{Tr}(A_i) \cdot A_i \right\| &\leq O(n^{3/2} \log^2 n) \\ \left\| \sum_i v(i) \cdot A_i \right\| &\leq O(\sqrt{n} \log n) \\ \left\| \sum_i \text{Tr}(A_i) v(i) \cdot v v^\top \right\| &\leq O(\sqrt{n} \log n). \end{aligned}$$

Proof of Lemma 6.5. We begin with the term $\sum_i \text{Tr}(A_i) \cdot A_i$. It is a sum of iid matrices $\text{Tr}(A_i) \cdot A_i$. A routine computation gives $\mathbb{E} \text{Tr}(A_i) \cdot A_i = \text{Id}$. We will use the truncated matrix Bernstein's inequality ([Proposition A.7](#)) to bound $\|\sum_i \text{Tr}(A_i) A_i\|$.

For notational convenience, let A be distributed like a generic A_i . By a union bound, we have both of the following:

$$\mathbb{P}(\|\text{Tr}(A) \cdot A\| \geq tn) \leq \mathbb{P}(|\text{Tr}(A)| \geq \sqrt{tn}) + \mathbb{P}(\|A\| \geq \sqrt{tn})$$

$$\mathbb{P}(\|\text{Tr}(A) \cdot A - \text{Id}\| \geq (t+1)n) \leq \mathbb{P}(|\text{Tr}(A)| \geq \sqrt{tn}) + \mathbb{P}(\|A\| \geq \sqrt{tn}).$$

Since $\text{Tr}(A)$ the sum of iid Gaussians, $\mathbb{P}(|\text{Tr}(A)| \geq \sqrt{tn}) \leq e^{-c_1 t}$ for some constant c_1 . Similarly, since the maximum eigenvalue of a matrix with iid entries has a subgaussian tail, $\mathbb{P}(\|A\| \geq \sqrt{tn}) \leq e^{-c_2 t}$ for some c_2 . All together, for some c_3 , we get $\mathbb{P}(\|\text{Tr}(A) \cdot A\| \geq tn) \leq e^{-c_3 t}$ and $\mathbb{P}(\|\text{Tr}(A) \cdot A - \text{Id}\| \geq (t+1)n) \leq e^{-c_3 t}$.

For a positive parameter β , let \mathbb{I}_β be the indicator variable for the event $\|\text{Tr}(A) \cdot A\| \leq \beta$. Then

$$\begin{aligned} \mathbb{E} \|\text{Tr}(A) \cdot A\| - \mathbb{E} \|\text{Tr}(A) \cdot A\| \mathbb{I}_\beta &= \int_0^\infty [\mathbb{P}(\|\text{Tr} A \cdot A\| > s) - \mathbb{P}(\|\text{Tr} A \cdot A\| \mathbb{I}_\beta > s)] ds \\ &= \beta \mathbb{P}(\|\text{Tr} A \cdot A\| > \beta) + \int_\beta^\infty \mathbb{P}(\|\text{Tr} A \cdot A\| > s) ds \\ &\leq \beta e^{-c_3 \beta/n} + \int_\beta^\infty \mathbb{P}(\|\text{Tr} A \cdot A\| > s) ds \\ &= \beta e^{-c_3 \beta/n} + \int_{\beta/n}^\infty \mathbb{P}(\|\text{Tr} A \cdot A\| \geq tn) n dt \\ &\leq \beta e^{-c_3 \beta/n} + \int_{\beta/n}^\infty n e^{-c_3 t} dt \\ &= \beta e^{-c_3 \beta/n} + \frac{n}{c_3} e^{-c_3 \beta/n}. \end{aligned}$$

Thus, for some $\beta = O(n \log n)$ we may take the parameters p, q of [Proposition A.7](#) to be $O(n^{-150})$. The only thing that remains is to bound the parameter σ^2 . Since $(\mathbb{E} \text{Tr}(A) \cdot A)^2 = \text{Id}$, it is enough just to bound $\|\mathbb{E} \text{Tr}(A)^2 A A^T\|$. We use again a union bound:

$$\mathbb{P}(\|\text{Tr}(A)^2 A A^T\| > tn^2) \leq \mathbb{P}(|\text{Tr}(A)| > t^{1/4} \sqrt{n}) + \mathbb{P}(\|A\| > t^{1/4} \sqrt{n}).$$

By a similar argument as before, using the Gaussian tails of $\text{Tr} A$ and $\|A\|$, we get $\mathbb{P}(\|\text{Tr}(A)^2 A A^T\| > tn^2) \leq e^{-c_4 \sqrt{t}}$. Then starting out with the triangle inequality,

$$\begin{aligned} \sigma^2 &= \|n \cdot \mathbb{E} \text{Tr}(A)^2 A A^T\| \\ &\leq n \cdot \mathbb{E} \|\text{Tr}(A)^2 A A^T\| \\ &= n \cdot \int_0^\infty \mathbb{P}(\text{Tr}(A)^2 A A^T > s) ds \\ &= n \cdot \int_0^\infty \mathbb{P}(\text{Tr}(A)^2 A A^T > tn^2) n^2 dt \\ &\leq n \cdot \int_0^\infty e^{-c_4 \sqrt{t}} n^2 dt \\ &= n \cdot \left[-\frac{2n^2(c_4 \sqrt{t} + 1)}{c_4^2} e^{-c_4 \sqrt{t}} \right]_{t=0}^{t=\infty} \\ &\leq O(n^3). \end{aligned}$$

This gives that with high probability,

$$\left\| \sum_i \text{Tr}(A_i) \cdot A_i \right\| \leq O(n^{3/2} \log^2 n).$$

The other matrices are easier. First of all, we note that the matrix $\sum_i v(i) \cdot A_i$ has independent standard Gaussian entries, so it is standard that with high probability $\|\sum_i v(i) \cdot A_i\| \leq O(\sqrt{n} \log n)$. Second, we have

$$\sum_i v(i) \operatorname{Tr}(A_i) v v^T = v v^T \sum_i v(i) \operatorname{Tr}(A_i).$$

The random variable $\operatorname{Tr}(A_i)$ is a centered Gaussian with variance n , and since v is a unit vector, $\sum_i v(i) \operatorname{Tr}(A_i)$ is also a centered Gaussian with variance n . So with high probability we get

$$\left\| v v^T \sum_i v(i) \operatorname{Tr}(A_i) \right\| = \left| \sum_i v(i) \operatorname{Tr}(A_i) \right| \leq O(\sqrt{n} \log n)$$

by standard estimates. This completes the proof. □