THE UNIVERSITY OF CHICAGO

NONLINEAR MATRIX CONCENTRATION AND ITS APPLICATIONS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL AND APPLIED MATHEMATICS

BY

JUNE QIONG WU

CHICAGO, ILLINOIS

JUNE 2025

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I would like to express my deepest gratitude to Madhur Tulsiani and Shmuel Weinberger for their time and guidance. Their wisdom, professionalism, and integrity are truly inspiring. I thank the Committee on Computational and Applied Mathematics for hosting me as a graduate student. I also extend my sincere thanks to the National Defense Science and Engineering Graduate program, Sandia National Lab and Los Alamos National Lab for their generous funding and stimulating summer programs.

ABSTRACT

The main purpose of this dissertation is to introduce a new method for obtaining norm bounds for random matrices, where each entry is a low-degree polynomial in an underlying set of independent real-valued random variables. Such matrices arise in a variety of settings in the analysis of spectral and optimization algorithms. Using ideas of decoupling and linearization, we show a simple way of expressing norm bounds for such matrices, in terms of their higher-order derivatives. Some of the highlighted applications include graph matrices, tensor networks and smoothed analysis. In addition, we present a concise analysis of the ellipsoid fitting problem using the concentration bound of random matrices with covariance structure.

CHAPTER 1

INTRODUCTION

Matrix concentration inequalities try to answer the following question: Given a vector of random variables $\mathbf{x} = (x_1, \ldots, x_n)$ taking values in some measure space Ω^n, and a measurable map $\mathbf{F} : \Omega^n \to \mathbb{R}^{d_1 \times d_2}$, what is an explicit bound on

$$\mathbb{P}(\|\mathbf{F}(\mathbf{x}) - \mathbb{E}\mathbf{F}(\mathbf{x})\| \geq t) \leq \ ? \tag{1.1}$$

where $\|\cdot\|$ denotes the spectral norm of a matrix.

Matrix concentration inequalities are used extensively in theoretical computer science, applied mathematics, and a wide range of other fields. In contrast to the classical random matrix theory, the random matrix $\mathbf{F}(\mathbf{x})$ may not have i.i.d entries and many applications seek the tail probability bound (1.1) in the non-asymptotic regime. Usually a sub-Gaussian bound on the tail probability is desirable.

This chapter will review some of the major scalar and matrix concentration inequalities and discuss tools for proving these inequalities. We delay the discussion of applications to Chapter 2 and Chapter 3.

1.1 Scalar concentration inequalities

Before delving into matrix concentration inequalities, we briefly mention some important scalar concentration inequalities. Some of the material in this section is taken from the wonderful book by Boucheron, Lugosi and Massart [BLM13]. We kept the notations and language intact at some places.

For sums of independent real-valued random variables, one of the classical inequalities for tail probabilities is Hoeffding's inequality.

Theorem 1.1.1 (Hoeffding's Ineqality [Hoe63]). *Let $x_1, \ldots, x_n$ be independent random variables such that x_i takes its value in $[a_i, b_i]$ almost surely for all $i \leq n$. Let*

$$S = \sum_{i=1}^{n} (x_i - \mathbb{E}x_i).$$

Then for every $t > 0$,

$$\mathbb{P}(S \ \geq \ t) \ \leq \ \exp\left(-\frac{2t^2}{\sum_{i=1}^{n}(b_i - a_i)^2}\right).$$

Hoeffding's inequality implies a sub-Gaussian tail inequality. However, the variance of S may be much smaller than $\sum_{i=1}^{n}(b_i - a_i)^2$. In such cases, we ask if a sharper bound is obtainable? Bennett's inequality provides such an improvement. Instead of boundness, Bennett's inequality only requires an appropriate control of moments.

Theorem 1.1.2 (Bennett's Inequality [Ben62]). *Let $x_1, \ldots, x_n$ be independent random variables with finite variance such that $x_i \leq C$ for some $C > 0$ almost surely for all $i < n$. Let*

$$S = \sum_{i=1}^{n} (x_i - \mathbb{E}x_i) \quad and \quad \sigma^2 = \sum_{i=1}^{n} \mathbb{E}x_i^2.$$

Then for every $t \geq 0$,

$$\mathbb{P}(S \ \geq \ t) \ \leq \ \exp\left(-\frac{\sigma^2}{C^2}\, h\left(\frac{Ct}{\sigma^2}\right)\right),$$

where $h(u) = (1+u)\log(1+u) - u$ for $u > 0$.

Bernstein's inequality is very similar to Bennett's inequality. Both inequalities provide a sub-Gaussian type inequality. Now instead of simply the sums of independent real-valued random variables, what about more general functions of independent real-valued random variables?

Formally, given a vector of random variables $\mathbf{x} = (x_1, \ldots, x_n)$ taking values in some measure space Ω^n, and a measurable map $f : \Omega^n \to \mathbb{R}$, we seek an explicit bound on

$$\mathbb{P}(|f(\mathbf{x}) - \mathbb{E}f(\mathbf{x})| \ \geq \ t) \ \leq \ ?$$

We emphasize that the x_i's may have different distributions, the only essential assumption is independence. Based on martingale methods, the following inequality was first discovered by Efron and Stein. We present Steele's version of the celebrated Efron-Stein inequality.

Theorem 1.1.3 (Efron-Stein Inequality [ES81, Ste86]). *Let $\mathbf{x} = (x_1, \ldots, x_n)$ be a vector of independent random variables and let $f(\mathbf{x})$ be a square-integrable function of $\mathbf{x}$. Then*

$$\mathrm{Var}(f(\mathbf{x})) \ \leq \ \sum_{i=1}^{n} \mathbb{E}\left[\left(f(\mathbf{x}) - \mathbb{E}^{(i)} f(\mathbf{x})\right)^2\right] := \nu,$$

where

$$\mathbb{E}^{(i)} f(\mathbf{x}) \ = \ \int_\Omega f(x_1, \ldots, x_{i-1}, X_i, x_{i+1}, \ldots, x_n) \ \ \mathrm{d}\mu_i(X_i).$$

ν is often called the variance proxy which captures the idea of summing up "local" variations of the function f. Efron-Stein inequality provides a bound in terms of the variance proxy. Once the variance of f is controlled, Chebyshev's inequality can be used to derive upper bounds for the tail probability. However, Efron-Stein inequality does not capture the phenomenon that the tail probability decreases at an exponential rate. The exponential Efron-Stein was proved by Boucheron, Lugosi and Massart [BLM03] using the entropy method.

Efron-Stein inequality is one of many results in which concentration properties may be controlled by studying the local behavior of the function at hand. For smooth functions of independent normal random variables, the Gaussian Poincaré inequality shares the same spirit.

Theorem 1.1.4 (Gaussian Poincaré Inequality). *Let $\mathbf{x} = (x_1, \ldots, x_n)$ be a vector of i.i.d standard Gaussian random variables. Let $f : \mathbb{R}^n \to \mathbb{R}$ be any continuously differentiable function. Then*

$$\mathrm{Var}(f(\mathbf{x})) \leq \mathbb{E}\left[\|\nabla f(\mathbf{x})\|^2\right].$$

Similar to Efron-Stein inequality, the Gaussian Poincaré inequality does not capture the right tail behavior either. If we restrict our attention to Lipschitz functions, the Gaussian concentration inequality provides a subguassian tail bound.

Theorem 1.1.5 (Gaussian Concentration Inequality [Bor75]). *Let $\mathbf{x} = (x_1, \ldots, x_n)$ be a vector of i.i.d standard Gaussian random variables. Let $f : \mathbb{R}^n \to \mathbb{R}$ be a L-Lipschitz function. Then, for all $t > 0$,*

$$\mathbb{P}(|f(\mathbf{x}) - \mathbb{E}f(\mathbf{x})| \geq t) \leq 2\exp\left(-\frac{t^2}{2L^2}\right).$$

However, Lipschitz functions are too limited when it comes to applications. Many combinatorial problems, for example, the subgraph counting, need to consider polynomials in independent random variables. Methods for obtaining concentration of (scalar) polynomial functions have been the subject of a large body of work, including results by Kim and Vu [KV00], Latała [Lat06], Schudy and Sviridenko [SS11, SS12], Adamczak and Wolff [AW15], and Bobkov, Götze and Sambale [BGS19]. Concentration inequalities for non-Lipschitz functions are fairly intricate. The work by Adamczak and Wolff [AW15] provides tail bounds for non-Lipschitz functions with bounded derivatives of higher order. Before we present a simplified version of their result, we introduce some notations.

Consider a homogeneous polynomial f of degree d. Let $A = (a_{\mathbf{i}})_{\mathbf{i} \in [n]^d}$ be a d-indexed matrix that stores all the coefficents of f. Let P be the set of all possible partitions of $\{1, \ldots, d\}$ into nonempty, pairwise disjoint sets. Let $\mathcal{J} = \{J_1, \ldots, J_k\}$ be a partition in P. Let $|\mathcal{J}|$ denote the cardinality of $\mathcal{J}$. Define the injective tensor product norm of A with respect to $\mathcal{J}$ as

$$\|A\|_{\mathcal{J}} = \sup \left\{ \sum_{\mathbf{i} \in [n]^d} a_{\mathbf{i}} \prod_{l=1}^{k} x^{(l)}_{\mathbf{i}_{J_l}} \ : \ \|(x^{(l)}_{\mathbf{i}_{J_l}})\|_2 \leq 1, 1 \leq l \leq k \right\}.$$

Theorem 1.1.6 (Polynomial Concentration Inequality [AW15]). *Assume that a random vector* $\mathbf{x}$ *in* $\mathbb{R}^n$ *satisfies the logarithmic Sobolev inequality with constant* L*. Let* $f : \mathbb{R}^n \to \mathbb{R}$ *be a homogeneous polynomial of degree* d*, then*

$$\mathbb{P}(|f(\mathbf{x}) - \mathbb{E}f(\mathbf{x})| \ \geq \ t) \ \leq \ 2\exp\left(-\frac{\eta_f(t)}{C_d}\right),$$

where

$$\eta_f(t) \ = \ \min_{\mathcal{J} \in P} \left(\frac{t}{(L/2)^{(d/2)} \|\mathbf{D}^d f(\mathbf{x})\|_{\mathcal{J}}} \right)^{\frac{2}{|\mathcal{J}|}}.$$

If the underlying measure satisfies the logarithmic Sobolev inequality (or the Poincaré inequality), the polynomial concentration inequality of [AW15] can be expressed in terms of the tensor product norms of the derivatives of f.

We omit other interesting concentration inequalities as the aforementioned inequalities are most relevant to this thesis.

1.2 Linear matrix concentration inequalities

As we move on to matrix concentration inequalities, one natural question to ask is if scalar concentration inequalities can be lifted directly to matrices. For sums of independent random matrices, the answer is yes!

The matrix analogues of the classical inequalities due to Hoeffding, Berstein, Khintchine and Rosenthal were proved by Mackey *et al.* [MJC+12] using the generalized Stein's method of exchangeable pairs. See the monograph [Tro15] for a complete bibliography.

Theorem 1.2.1 (Matrix Berstein [Tro15]). *Let $\mathbf{X}_1, \ldots, \mathbf{X}_n$ be independent, centered random matrices with common dimension $d_1 \times d_2$, and assume that each one is uniformly bounded*

$$\mathbb{E}\mathbf{X}_k = \mathbf{0} \quad \textit{and} \quad \|\mathbf{X}_k\| \leq L \quad \textit{for each } k = 1, \ldots, n.$$

introduce the sum

$$\mathbf{Z} = \sum_{k=1}^{n} \mathbf{X}_k,$$

and let $\nu(\mathbf{Z})$ denote the matrix variance statistics of the sum:

$$\begin{aligned} \nu(\mathbf{Z}) &= \max\{\|\mathbb{E}(\mathbf{Z}\mathbf{Z}^*)\|, \ \|\mathbb{E}(\mathbf{Z}^*\mathbf{Z})\|\} \\ &= \max\left\{\left\| \sum_{k=1}^{n} \mathbb{E}(\mathbf{X}_k\mathbf{X}_k^*)\right\|, \ \left\| \sum_{k=1}^{n} \mathbb{E}(\mathbf{X}_k^*\mathbf{X}_k)\right\|\right\}. \end{aligned}$$

Then,

$$\mathbb{P}(\|\mathbf{Z}\| \geq t) \leq (d_1 + d_2) \cdot \exp\left(\frac{-t^2/2}{\nu(\mathbf{Z}) + Lt/3}\right) \quad \textit{for all } t \geq 0.$$

Notice that the Matrix Berstein inequality is expressed in terms of the ambient dimension d_1 and d_2 of the random matrix $\mathbf{Z}$. The dependence on the ambient dimension is also present in the noncommutative Khintchine inequality.

Theorem 1.2.2 (Noncommutative Khintchine Inequality [Pis03]). *Let $\mathbf{A}_1, \ldots \mathbf{A}_n$ be self-adjoint $d \times d$ matrices and $g_1, \ldots, g_n$ be standard Gaussian random variables,. Then*

$$C_1 \left\| \left(\sum_{i=1}^{n} \mathbf{A}_i^2\right)^{\frac{1}{2}} \right\| \leq \mathbb{E} \left\| \sum_{i=1}^{n} g_i \mathbf{A}_i \right\| \leq C_2 \sqrt{\log d} \left\| \left(\sum_{i=1}^{n} \mathbf{A}_i^2\right)^{\frac{1}{2}} \right\|,$$

where C_1 and C_2 are some universal constants.

The noncommutative Khintchine inequality yields two-sided estimates on the spectral norm of sums of Gaussian random matrices. The dependence on the ambient dimension is only sharp when $\mathbf{A}_i$'s are commutative, which is not optimal in many cases. Recently, [BBvH23] provides a stronger noncommutative Khintchine inequality which takes into account the noncommutativity of $\mathbf{A}_i$'s.

1.3 Nonlinear matrix concentration inequalities

While there are abundant concentration inequalities for general scalar functions of random variables and sophisticated techniques for proving them, the matrix analogues are still very scarce.

Using the generalized Stein's method of exchangeable pairs again, Paulin, Mackey and Tropp [PMT16] proved the matrix analogues of exponential (and polynomial) Efron-Stein inequality.

Let $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{X}$ be a vector of mutually independent random variables in a Polish space. Let $\mathbf{H} : \mathcal{X} \to \mathbb{H}^d$ be a measurable function that takes values in the space of Hermitian matrices, and construct the centered random matrix

$$\mathbf{Z} \ := \ \mathbf{H}(\mathbf{x}) - \mathbb{E}\mathbf{H}(\mathbf{x}).$$

Similar to the scalar Efron-Stein inequality, we quantify the "local" variation of $\mathbf{H}$ by replacing each coordinate with an independent copy. Construct the random vector

$$\mathbf{x}^{(j)} \ := (x_1, \ldots, \, x_{j-1}, \, x_j', \, x_{j+1}, \ldots, \, x_n) \ \in \ \mathcal{X},$$

where x_j' is an independent copy of x_j. Clearly $\mathbf{x}$ and $\mathbf{x}^{(j)}$ have the same distribution, hence they are "exchangeable" in some sense. Form the centered random matrix

$$\mathbf{Z}^{(j)} \ := \ \mathbf{H}(\mathbf{x}^{(j)}) - \mathbb{E}\mathbf{H}(\mathbf{x}^{(j)}).$$

The total "local" variation is captured by the variance proxy

$$\mathbf{V} \ := \ \frac{1}{2}\sum_{j=1}^{n}\mathbb{E}\left[\left(\mathbf{Z}-\mathbf{Z}^{(j)}\right)^2 \,\middle|\, \mathbf{x}\right].$$

Theorem 1.3.1 (Matrix Efron-Stein [PMT16])**.** *Instate the aforementioned notation and assume that* $\mathbb{E}\|\mathbf{Z}\|^2 < \infty$*. For each natural number* $p \geq 1$*,*

$$\left(\mathbb{E}\|\mathbf{Z}\|_{2p}^{2p}\right)^{\frac{1}{2p}} \ \leq \ \sqrt{2(2p-1)}\left(\|\mathbf{V}\|_{2p}^{p}\right)^{\frac{1}{2p}},$$

where $\|\cdot\|_{2p}$ *denotes the Schatten-*$2p$ *norm.*

Recently, [HT20, ABY20] proved matrix generalizations of the Poincaré inequality using the Markov semi-group method. Although the inequality holds for any measure μ that satisfies the Poincaré inequality, we shall present the Gaussian version of it for ease of presentation.

Theorem 1.3.2 (Matrix Poincaré Inequality [HT20])**.** *Let* $\mathbf{x}$ *be a vector of standard Gaussain variables, and let* $\mathbf{F} : \Omega^n \to \mathbb{H}_d$ *be a matrix-valued function. Then for all* $t > 0$*,*

$$\mathbb{P}(\|\mathbf{F}(\mathbf{x})) - \mathbb{E}\mathbf{F}(\mathbf{x})\| \ \geq \ \sqrt{\nu}\cdot t) \ \leq \ 6d \cdot e^{-t},$$

where the variance proxy

$$\nu := \operatorname*{ess\,sup}_{\mathbf{x}\in\Omega^n}\left\|\sum_{i=1}^{n}(\partial_i\mathbf{F}(\mathbf{x}))^2\right\|.$$

In the language of the Markov semi-group, the variance proxy ν is associated with the carré du champ operator of the process whose stationary measure is μ. Notice that the matrix Poincaré inequality only implies subexponential concentration. In a later paper [HT21], Huang and Tropp improved the bound to subguassian with local Poincaré inequality.

So far, the matrix Poincaré inequality is the only functional inequality that led to desirable matrix concentration results. Chen and Tropp [TC13] developed matrix generalization of (modified) log-Sobolev inequality. Unfortunately, the concentration bounds they derive require artificial assumptions. It remains open whether functional inequalities could be used to the same extent as they were in the scalar setting.

1.4 Organization of the thesis

In Chapter 2, we prove Theorem 2.4.5 for obtaining norm bounds for random matrices, where each entry is a low-degree polynomial in an underlying set of independent real-valued random variables. Using ideas of decoupling and linearization from analysis, we show a simple way of expressing norm bounds for such matrices, in terms of matrices of their higher-order derivatives. In Section 2.6, we use Theorem 2.4.5 to obtain norm bounds for sparse graph matrices which have applications in complexity theory. In Section 2.7, we use Theorem 2.5.3 to obtain norm bound for a simple tensor network which has applications in the analysis of spectral algorithms. In Section 2.8, we study the smoothed analysis and obtain the least singular value bound for smoothed Khatri-Rao products using Theorem 2.4.5. Most of the content in this chapter is a joint work [TW25b] with Madhur Tulsiani that has appeared at 16th Innovations in Theoretical Computer Science Conference (ITCS 2025). Section 2.8 is a joint work with Madhur Tulsiani and Jeff Xu.

In Chapter 3, we discuss the ellipsoid fitting conjecture of Saunderson, Chandrasekaran, Parrilo and Willsky. The ellipsoid fitting conjecture considers the maximum number n random Gaussian points in $\mathbb{R}^d$, such that with high probability, there exists an origin-symmetric ellipsoid passing through all the points. We give a simple proof (Theorem 3.3.1) based on concentration of sample covariance matrices, that with probability $1 - o_d(1)$, it is possible to fit an ellipsoid through d^2/C random Gaussian points. In Section 3.4, we mention some directions for future work. Most of the content in this chapter is a joint work [TW25a] with Madhur Tulsiani that has appeared at 2025 Symposium on Simplicity in Algorithms (SOSA).

CHAPTER 2

POLYNOMIAL RANDOM MATRICES

In this chapter, we present a new method for obtaining norm bounds for random matrices, where each entry is a low-degree polynomial in an underlying set of independent real-valued random variables. Such matrices arise in a variety of settings in the analysis of spectral and optimization algorithms, which require understanding the spectrum of a random matrix depending on data obtained as independent samples.

Using ideas of decoupling and linearization from analysis, we show a simple way of expressing norm bounds for such matrices, in terms of matrices of lower-degree polynomials corresponding to derivatives. Iterating this method gives a simple bound with an elementary proof, which can recover many bounds previously required more involved techniques.

2.1 Introduction

For fixed matrices $\mathbf{C_1}, \ldots, \mathbf{C_n}$ and independent scalar random variables $x_1, \ldots, x_n$, consider the problem of analyzing the random matrix

$$\mathbf{M} \;=\; \mathbf{C_1} \cdot x_1 + \cdots + \mathbf{C_n} \cdot x_n \,.$$

Note that the entries of the random matrix $\mathbf{M}$ are not necessarily independent, but are (possibly correlated) linear functions of the independent random variables $x_1, \ldots, x_n$. Such matrices, which arise in a variety of applications in algorithms, statistics, and numerical linear algebra, can often be shown to be concentrated around the mean, using a rich selection of matrix deviation inequalities [Tro15]. Moreover, when the random variables $x_1, \ldots, x_n$ are Gaussian, recent breakthrough results have obtained even sharper concentration guarantees depending on the structure of the matrices $\mathbf{C_1}, \ldots, \mathbf{C_n}$ [BBvH23, BvH22] already leading to applications in discrepancy theory [BJM23] and quantum information theory [LY23].

In contrast to the above random matrices which are linear functions of independent random variables, several recent applications in spectral algorithms and lower bounds for statistical problems require understanding the (expected) norms of random matrices of the form

$$\mathbf{F} = \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} \mathbf{C}_S \cdot \prod_{i \in S} x_i \,.$$

The above random matrix, which denotes as a matrix-valued function $\mathbf{F}(\mathbf{x})$, is a (multilinear, in the example above) *low-degree polynomial* in the vector of random variables $\mathbf{x} = (x_1, \ldots, x_n)$. Such matrix-valued polynomial functions arise naturally in a variety of applications, including algorithms for tensor decomposition and completion [HSSS16, HSS19, KP20, DdL$^+$22], orbit recovery [MW19], power-sum decompositions and learning of Gaussian mixtures [BHKX22], and lower bounds for Sum-of-Squares hierarchies [AMP21, BHK$^+$19, JPR$^+$21, Raj22]. In general, such matrices are often an important technical component in the analysis of spectral algorithms, which require understanding the spectrum of some random matrix depending on data obtained as independent samples. Given the expressive power of polynomials, one can often write the matrix entries as low-degree polynomials in the data.

In several applications above, one is interested in obtaining bounds on the spectral norm of the deviation matrix $\mathbf{F}(\mathbf{x}) - \mathbb{E}\,\mathbf{F}(\mathbf{x})$, which hold with high probability over the choice of $\mathbf{x}$. Although matrix-deviation inequalities for nonlinear random matrices have been a subject of active research in recent work [PMT16, MJC$^+$12, ABY20, HT20, HT21], the analyses in the applications above have often needed to rely on estimating matrix norms via direct computations of trace moments. While these do yield sharp bounds required for applications, they require somewhat intricate computations and ingenious combinatorial arguments tailored to the applications at hand.

Analyzing concentration via moment estimates. We consider concentration bounds on a *scalar* polynomial $f(\mathbf{x})$ with mean zero. Using Markov's inequality, this can be reduced to computing moment estimates, since

$$\mathbb{P}\left[|f(\mathbf{x})| \geq \lambda\right] \;=\; \mathbb{P}\left[(f(\mathbf{x}))^{2t} \geq \lambda^{2t}\right] \;\leq\; \lambda^{-2t} \cdot \mathbb{E}\left[(f(\mathbf{x}))^{2t}\right]$$

Note that while in some cases $\mathbb{E}[(f(\mathbf{x}))^{2t}]$ can be computed by direct expansion, it often involves an intricate analysis of the structure of terms with degrees growing with t, and therefore indirect methods may be more convenient. Methods for obtaining concentration of (scalar) polynomial functions have been the subject of a large body of work, including results by Kim and Vu [KV00], Latała [Lat06], Schudy and Sviridenko [SS11, SS12], Adamczak and Wolff [AW15], and Bobkov, Götze and Sambale [BGS19]. A useful comparison to direct computation, is the method of Adamczak and Wolff, which is applicable to functions $f(\mathbf{x})$ of a broad range of random vectors $\mathbf{x}$ (with not necessarily independent coordinates) from a distribution obeying the Poincaré inequality. They obtain such moment bounds by recursive applications of the Poincaré inequality, reducing the computation of the moments of f to that of its derivatives. For a degree-d polynomial function f, one can consider its coefficients as forming a (symmetric) order-d tensor, and their results obtain bounds on the moments of f in terms of the norms of various lower-order tensors "flattening" of this coefficient tensor. Thus, the problem of estimating moments of the random variable $f(\mathbf{x})$, can be reduced in a black-box way, to understanding the norms of a small number (depending on d) of *deterministic* tensors. Any dependence on the problem structure can then be limited to understanding these norms, where often crude estimates can suffice.

The matrix analog of the Markov argument involves the Schatten-$2t$ norm $\|.\|_{2t}$, which is defined for a matrix $\mathbf{M}$ with non-zero singular values $\sigma_1, \ldots, \sigma_r$ as $\|\mathbf{M}\|_{2t}^{2t} := \sum_{j\in[r]} \sigma_j^{2t}$. For a function $\mathbf{F}$ with $\mathbb{E}[\mathbf{F}(Z)] = \mathbf{0}$, we have the following bound using Schatten norms.

$$\mathbb{P}\left[\sigma_1(\mathbf{F}) \geq \lambda\right] \;\leq\; \lambda^{-2t} \cdot \mathbb{E}\left\|\mathbf{F}\right\|_{2t}^{2t} \;=\; \lambda^{-2t} \cdot \mathbb{E}\,\mathrm{tr}\left[\left(\mathbf{F}(\mathbf{x})\mathbf{F}(\mathbf{x})^{\mathsf{T}}\right)^t\right]$$

Several applications requiring norm bounds for matrix-valued polynomial functions, start with the above inequality, and often rely on direct expansion of the trace for the matrix power $\left(\mathbf{F}(\mathbf{x})\mathbf{F}(\mathbf{x})^{\mathsf{T}}\right)^t$. Recall that when $\mathbf{A}$ is the adjacency matrix of a graph, $\mathrm{tr}(\mathbf{A}^{2t})$ can be interpreted as the number of cycles of length $2t$. Similarly, when working with matrices $\mathbf{F}(\mathbf{x})$ where each entry can be interpreted as arising from a combinatorial pattern, $\left(\mathbf{F}(\mathbf{x})\mathbf{F}(\mathbf{x})^{\mathsf{T}}\right)^t$ can be viewed as the number of copies of such patterns "chained" together in a certain way. Estimating the trace then amounts to estimating the (expected) number of such chainings [AMP21].

Obtaining bounds on the $\mathbb{E}\left\|\mathbf{F}\right\|_{2t}^{2t}$ via this method requires first using problem structure to decompose $\mathbf{F}$ into matrices with such combinatorial patterns, and then using (ingenious) combinatorial arguments to understand the expected number of occurrences of such chains of patterns. The focus of this work is understanding alternative general methods for estimating such norm bounds, when the underlying problem may not necessarily have a combinatorial flavor, or when such decompositions may be hard to analyze.

Efron-Stein inequalities. Efron-Stein inequalities bound the global variance of a function of independent random variables, in terms of local variance estimates obtained by changing one variable at a time. For $i \in [n]$ and tuple $\mathbf{x} = (x_1, \ldots, x_n)$, let $\mathbf{x}^{(i)}$ denote the tuple $(x_1, \ldots, x_{i-1}, \widetilde{x}_i, x_{i+1}, \ldots, x_n)$, where $\widetilde{x}_i$ is an independent copy of x_i. For a scalar function

$f(\mathbf{x})$, the Efron-Stein inequality states that

$$\mathrm{Var}\left[f(\mathbf{x})\right] \;=\; \mathbb{E}\left[(f(\mathbf{x})-\mathbb{E}f)^2\right] \;\leq\; \frac{1}{2}\cdot\sum_{i\in[n]}\mathbb{E}\left[\left(f(\mathbf{x})-f\left(\mathbf{x}^{(i)}\right)\right)^2\right] \;=\; \mathbb{E}\left[V(\mathbf{x})\right],$$

where $V(\mathbf{x}) \;:=\; \sum_{i\in[n]}\mathbb{E}\left[(f(\mathbf{x})-f(\mathbf{x}^{(i)}))^2 \;\middle|\; \mathbf{x}\right]$. Moment versions of the Efron-Stein inequality were developed by Boucheron et al. [BBLM05], who proved[1]

$$\mathbb{E}\left[(f(\mathbf{x})-\mathbb{E}f)^{2t}\right] \;\leq\; (C_0\cdot t)^t\cdot\mathbb{E}\left[(V(\mathbf{x}))^t\right].$$

A beautiful matrix generalization of the above inequality was obtained by Paulin, Mackey and Tropp [PMT16], via the method of exchangeable pairs. This was also extended in later works [ABY20, HT20] to settings where the random input vector $\mathbf{x}$ does not necessarily come from a product distribution, but from a distribution satisfying a Poincaré inequality.

Rajendran and Tulsiani [RT23] used the matrix Efron-Stein inequality of Paulin, Mackey and Tropp [PMT16] to obtain norm bounds for matrix-valued polynomial functions $\mathbf{F}(\mathbf{x})$ of independent variables $x_1,\ldots,x_n$. For (say) a degree-d homogeneous function $\mathbf{F}(\mathbf{x})$ their method relied on iteratively applying the matrix Efron-Stein inequality, and at each step interpreting the variance proxy $\mathbf{V}(\mathbf{x})$ as the square of a degree-$(d-1)$ polynomial matrix. While this method indeed yields simple bounds when $x_1,\ldots,x_n$ are independent Rademacher random variables, this simple recursion based on the black-box use of the inequality from [PMT16] unfortunately stalls for other product distributions. As a result, the proof and also the form of the result in [RT23] was significantly more involved for other product distributions (including Gaussians). For the general case, the proof required modifying the exchangeable pairs argument of [PMT16], and the bounds reduced the question of estimating moments of matrix-valued functions to those of scalar polynomials

1. In fact, the estimates of [BBLM05] are in terms of more refined quantities $V_+(\mathbf{x})$ and $V_-(\mathbf{x})$, which are both upper bounded by $V(\mathbf{x})$ but can be much smaller.

(which were then analyzed using known methods).

Recent and independent work. Shortly after our work, an independent work by Bandeira, Lucca, Nizić-Nikolac and van Handel [ASBvH25] provided general matrix concentration inequalities for matrix chaoses in a similar vein.

Decoupling. Decoupling inequalities were developed in the study of U-statistics [PG99], multiple stochastic integration [MT84], and polynomial chaos [Kwa87], and have found important applications in applied mathematics, theoretical computer science, applied probability theory and statistics. Some examples include the study of compressed sensing [Rau10], query complexity [OZ16], the proof of Hanson-Wright inequality [Ver18], learning mixture of Gaussians [GHK15] and so on. Moreover, the inequalities are applicable in both scalar and matrix settings, and even more broadly in Banach spaces.

For a degree-d homogeneous *multilinear* polynomial $f(\mathbf{x})$ in n independent random variables, standard decoupling inequalities (see Section 2.3) can be used to obtain

$$\mathop{\mathbb{E}}_{\mathbf{x}}\left[(f(\mathbf{x}))^{2t}\right] \;\le\; C_d^{2t}\cdot \mathop{\mathbb{E}}_{\mathbf{x}^{(1)},\ldots,\mathbf{x}^{(d)}}\left[\left(f(\mathbf{x}^{(1)},\ldots,\mathbf{x}^{(d)})\right)^{2t}\right],$$

where $f(\mathbf{x}^{(1)},\ldots,\mathbf{x}^{(d)})$ denotes the polynomial in $d\cdot n$ random variables obtained by replacing the d variables in each (ordered) monomial by the corresponding d coordinates coming from the d independent random vectors $\mathbf{x}^{(1)},\ldots,\mathbf{x}^{(d)}$, and C_d is a constant depending on the degree d. One can now fix the vectors $\mathbf{x}^{(1)},\ldots,\mathbf{x}^{(d-1)}$, and consider the expectation

$$\mathop{\mathbb{E}}_{\mathbf{x}^{(d)}}\left[\left(f(\mathbf{x}^{(1)},\ldots,\mathbf{x}^{(d)})\right)^{2t}\right] \;=\; \mathop{\mathbb{E}}_{\mathbf{x}^{(d)}}\left[\left(\sum_{i\in[n]} f_i(\mathbf{x}^{(1)},\ldots,\mathbf{x}^{(d-1)})\cdot x_i^{(d)}\right)^{2t}\right],$$

where we write $f(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)})$ as a linear form in $\mathbf{x}^{(d)}$ with coefficients $f_i(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d-1)})$. The moments of such a linear form can be understood (for example) using the Khintchine inequality, which bounds it in terms of a polynomial depending only on $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d-1)}$.

Similarly, in the matrix case, we use decoupling inequalities to reduce the problem of estimating moments for degree-d homogeneous (and multilinear) matrix-valued polynomials, to that of estimating moments for *linear* matrix-valued functions. At this point, one can apply standard and well-known linear matrix concentration inequalities (we use the matrix Rosenthal inequality) which yield a bound in terms of a degree-$(d-1)$ matrix-valued polynomial function of the vectors $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d-1)}$. Iterating this process gives a simple method for obtaining norm bounds, applicable for *multilinear* polynomial functions on n independent random variables (with any reasonable distribution).

While this method does not immediately apply for non-multilinear polynomials, it suffices for many of the applications mentioned above. Moreover, for arbitrary polynomials in Gaussian random variables, we can also obtain a simple recursion as an immediate consequence of the more recent Poincaré inequalities of Huang and Tropp [HT21]. Together, these suffice for most applications of interest.

2.1.1 Methods and results: a technical overview

Let $x_1, \ldots, x_n$ be i.i.d real-valued random variables with $\mathbb{E}[x_i] = 0, \mathbb{E}[x_i^2] = 1$ and $|x_i| \leq L$ for each $i \in [n]$. Let $\mathcal{T}_n^d \subseteq [n]^d$ denote the set of ordered d-tuples $\mathbf{i} = (i_1, \ldots, i_d)$ with $i_1, \ldots, i_d$ all distinct. For a fixed sequence of (deterministic) matrices $\{\mathbf{A_i}\}_{\mathbf{i} \in \mathcal{T}_n^d} \subseteq \mathbb{R}^{d_1 \times d_2}$, we consider a matrix-valued multilinear polynomial function defined as

$$\mathbf{F}(\mathbf{x}) \;=\; \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{A_i} \cdot \prod_{j \in [d]} x_{i_j} \,.$$

We write the polynomial in terms of ordered tuples for decoupling, but since the variables $x_1, \ldots, x_n$ commute, we can assume that for any permutation $\sigma : [d] \to [d]$ permuting the d coordinates of each tuple, $\mathbf{A}_{\mathbf{i}} = \mathbf{A}_{\sigma(\mathbf{i})}$ (this is referred to as *permutation symmetry* of $\mathbf{F}$). We can use the decoupling inequalities from Section 2.3 to show that

$$\|\mathbf{F}(\mathbf{x})\|_{4t} \;\leq\; C_d \cdot \left\|\mathbf{F}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)})\right\|_{4t} \;=\; C_d \cdot \left\| \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{A}_{\mathbf{i}} \cdot \prod_{j \in [d]} x_{i_j}^{(j)} \right\|_{4t} ,$$

where $\|\cdot\|_{4t}$ denotes the (expected) Schatten norm $\left(\mathbb{E}[\mathrm{tr}(\mathbf{F}(\mathbf{x})^{\mathsf{T}}\mathbf{F}(\mathbf{x}))^{2t}]\right)^{1/4t}$.

We remark that while the constant C_d is of the form d^d in general, this can be improved when the underlying set of indices $[n]$ has more structure. For example, when $[n]$ corresponds to the set of *pairs* in a base set $[r]$ and each of the monomials involves at most k elements of $[r]$ (has *index degree* k), it is possible to improve the constant to k^k. This is essentially the "vertex partitioning" argument used in works on graph matrices, and is discussed in Section 2.3.1 in the language of decoupling.

As mentioned earlier, we can now "linearize" the function $\mathbf{F}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)})$ by fixing $\mathbf{x}^{(d)}, \ldots, \mathbf{x}^{(d-1)}$ and treating it only as a function of $\mathbf{x}^{(d)}$ i.e., we consider

$$\begin{aligned}
\mathbb{E}\left\|\mathbf{F}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)})\right\|_{4t}^{4t} &= \mathbb{E}\left\| \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{A}_{\mathbf{i}} \cdot \prod_{j \in [d]} x_{i_j}^{(j)} \right\|_{4t}^{4t} \\
&= \mathop{\mathbb{E}}_{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d-1)}} \; \mathop{\mathbb{E}}_{\mathbf{x}^{(d)}} \left\| \sum_{k \in [n]} \left(\sum_{\substack{\mathbf{i} \in \mathcal{T}_n^d \\ i_d = k}} \mathbf{A}_{\mathbf{i}} \cdot \prod_{j \in [d-1]} x_{i_j}^{(j)} \right) \cdot x_k^{(d)} \right\|_{4t}^{4t} \\
&= \mathop{\mathbb{E}}_{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d-1)}} \; \mathop{\mathbb{E}}_{\mathbf{x}^{(d)}} \left\| \sum_{k \in [n]} \left(\partial_{x_k^{(d)}} \mathbf{F} \right) \cdot x_k^{(d)} \right\|_{4t}^{4t} .
\end{aligned}$$

To bound the inner expectation, we can now use the matrix Rosenthal inequality (Lemma 2.4.1) which says that for a collection of centered, independent random matrices $\{\mathbf{Y}_k\}$ with finite moments, we have

$$\mathbb{E}\left\|\sum_k \mathbf{Y}_k\right\|_{4t}^{4t} \leq (16t)^{3t} \cdot \left\{\left\|\left(\sum_k \mathbb{E}\mathbf{Y}_k\mathbf{Y}_k^\top\right)^{1/2}\right\|_{4t}^{4t} + \left\|\left(\sum_k \mathbb{E}\mathbf{Y}_k^\top\mathbf{Y}_k\right)^{1/2}\right\|_{4t}^{4t}\right\} + (8t)^{4t} \cdot \left(\sum_k \mathbb{E}\left\|\mathbf{Y}_k\right\|_{4t}^{4t}\right).$$

Taking $\mathbf{Y}_k = \partial_{x_k^{(d)}}\mathbf{F} \cdot x_k^{(d)}$ we can now compute the expectations (over $\mathbf{x}^{(d)}$) in the RHS as

$$\begin{aligned}
\sum_k \mathbb{E}\mathbf{Y}_k\mathbf{Y}_k^\top &= \begin{bmatrix}\partial_{x_1^{(d)}}\mathbf{F} & \dots & \partial_{x_n^{(d)}}\mathbf{F}\end{bmatrix}\begin{bmatrix}\partial_{x_1^{(d)}}\mathbf{F} & \dots & \partial_{x_n^{(d)}}\mathbf{F}\end{bmatrix}^\top \\
\sum_k \mathbb{E}\mathbf{Y}_k^\top\mathbf{Y}_k &= \begin{bmatrix}\partial_{x_1^{(d)}}\mathbf{F}^\top & \dots & \partial_{x_n^{(d)}}\mathbf{F}^\top\end{bmatrix}\begin{bmatrix}\partial_{x_1^{(d)}}\mathbf{F}^\top & \dots & \partial_{x_n^{(d)}}\mathbf{F}^\top\end{bmatrix}^\top \\
\sum_k \mathbb{E}\left\|\mathbf{Y}_k\right\|_{4t}^{4t} &= \mathbb{E}\left\|\begin{bmatrix} x_1^{(d)} \cdot \partial_{x_1^{(d)}}\mathbf{F} & & \\ & \ddots & \\ & & x_n^{(d)} \cdot \partial_{x_n^{(d)}}\mathbf{F}\end{bmatrix}\right\|_{4t}^{4t} \\
&\leq L^{4t} \cdot \left\|\begin{bmatrix}\partial_{x_1^{(d)}}\mathbf{F} & & \\ & \ddots & \\ & & \partial_{x_n^{(d)}}\mathbf{F}\end{bmatrix}\right\|_{4t}^{4t}
\end{aligned}$$

Iterating the above argument, we can prove a bound in terms of a collection of matrices

$\mathbf{F}_{a,b,c}$ constructed inductively as

$$\begin{aligned}
\mathbf{F}_{a+1,b,c} &= \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b,c} & \dots & \partial_{x_n}\mathbf{F}_{a,b,c} \end{bmatrix}^{\top} \\
\mathbf{F}_{a,b+1,c} &= \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b,c} & \dots & \partial_{x_n}\mathbf{F}_{a,b,c} \end{bmatrix} \\
\mathbf{F}_{a,b,c+1} &= \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b,c} & & \\ & \ddots & \\ & & \partial_{x_n}\mathbf{F}_{a,b,c} \end{bmatrix}.
\end{aligned}$$

Note that in the above definition, we do not specify the order of differentiation, which can be ignored due to the permutation symmetry of $\mathbf{F}$ and the fact that changing the order does not change the Schatten norms (see Section 2.2 and Section 2.4 for details). This argument yields the following bound for all homogeneous multilinear polynomial functions.

Theorem 2.1.1 (Restatement of Theorem 2.4.5). *Let $\mathbf{x} = \{x_i\}_{i=1}^n$ be a sequence of i.i.d random variables with $\mathbb{E}x_i = 0$, $\mathbb{E}x_i^2 = 1$ and $|x_i| \le L$ for all $1 \le i \le n$. Let $\{\mathbf{A_i}\}_{\mathbf{i}\in\mathcal{T}_n^d}$ be a multi-indexed sequence of deterministic matrices of the same dimension. Define a permutation symmetric, homogeneous multilinear polynomial random matrix of degree d as*

$$\mathbf{F}(\mathbf{x}) = \sum_{\mathbf{i}\in\mathcal{T}_n^d} \left(\mathbf{A_i} \cdot \prod_{j\in\{i_1,\dots,i_d\}} x_j \right).$$

Let $a, b, c \in \mathbb{Z}_{\ge 0}$ and $d = a + b + c$. Then for $2 \le t \le \infty$,

$$\mathbb{E}\, \|\mathbf{F}(\mathbf{x})\|_{4t}^{4t} \le \sum_{\substack{a,b,c: \\ a+b+c=d}} (48dt)^{4dt} \cdot L^{4ct} \left\|\mathbf{F}_{a,b,c}\right\|_{4t}^{4t}.$$

We remark that the condition that $|x_i| \le L$ can be replaced by a condition on the growth of moments of x_i, which is true for subgaussian variables. The above bound is in terms of the norms of a constant number of *deterministic* matrices $\mathbf{F}_{a,b,c}$, as is also the case for general moment bounds on scalar polynomials [AW15]. We will see later that these deterministic matrices can easily be interpreted in several cases of interest, to recover the bounds obtained via combinatorial methods.

Gaussian polynomial matrices. While the above bounds are only for multilinear polynomials, they can also be extended to *arbitrary* polynomials of independent Gaussian random variables, using standard techniques to approximate them by multilinear polynomials. However, for the case of Gaussian polynomials, the recent Poincaré inequalities of Huang and Tropp [HT20] directly yield a simple bound, which is easier to apply. They show that for a Hermitian matrix-valued function $\mathbf{H}$ of Gaussian valued random variables

$$\mathbb{E}\,\|\mathbf{H}(\mathbf{x}) - \mathbb{E}\mathbf{H}(\mathbf{x})\|_{2t}^{2t} \;\leq\; (\sqrt{2}t)^{2t} \cdot \mathbb{E}\,\mathrm{tr}\left(\sum_{i=1}^{n} (\partial_i \mathbf{H}(\mathbf{x}))^2\right)^{t}.$$

Given a (not necessarily Hermitian) matrix-valued degree-d polynomial function $\mathbf{F}(\mathbf{x}) \in \mathbb{R}^{d_1 \times d_2}$, we can apply the above bound to the "Hermitian dilation" $\mathbf{H} = \begin{bmatrix} \mathbf{0} & \mathbf{F} \\ \mathbf{F}^\top & \mathbf{0} \end{bmatrix}$, to get

$$2\,\|\mathbf{F} - \mathbb{E}\mathbf{F}\|_{2t}^{2t} \;\leq\; (\sqrt{2}t)^{2t} \cdot \left\|\left(\sum_{i=1}^{n} \partial_i \mathbf{F}\, \partial_i \mathbf{F}^\top\right)^{1/2}\right\|_{2t}^{2t} + (\sqrt{2}t)^{2t} \cdot \left\|\left(\sum_{i=1}^{n} \partial_i \mathbf{F}^\top\, \partial_i \mathbf{F}\right)^{1/2}\right\|_{2t}^{2t}$$

As before, one can apply this argument inductively to obtain a bound in terms of the matrices $\mathbf{F}_{a,b}$ which are defined similarly as the matrices $\mathbf{F}_{a,b,c}$ before, with $c = 0$. Since we no longer have $\mathbb{E}\mathbf{F}_{a,b} = \mathbf{0}$ for $a + b < d$, the bound is in terms of the expected matrices $\mathbb{E}\mathbf{F}_{a,b}$ for all a, b with $a + b \leq d$.

Theorem 2.1.2 (Restatement of Theorem 2.5.3)**.** *Let* $\mathbf{x} \sim \mathcal{N}(0, \mathbb{I}_n)$ *and* $\{\mathbf{A}_{\mathbf{i}}\}_{\mathbf{i} \in [n]^d}$ *be a sequence of deterministic matrices of the same dimension. Define a degree-d homogeneous Gaussian* *polynomial random matrix as*

$$\mathbf{F}(\mathbf{x}) = \sum_{\mathbf{i}\in[n]^d} \left(\mathbf{A}_{\mathbf{i}} \cdot \prod_{j\in\{i_1,\dots,i_d\}} x_j \right).$$

Let $a, b \in \mathbb{Z}_{\geq 0}$. Then for $2 \leq t \leq \infty$,

$$\mathbb{E}\,\|\mathbf{F}(\mathbf{x}) - \mathbb{E}\mathbf{F}(\mathbf{x})\|_{2t}^{2t} \;\leq\; (2^d\sqrt{2}t)^{2t} \left(\sum_{1\leq a+b\leq d} \left\|\mathbb{E}\mathbf{F}_{a,b}\right\|_{2t}^{2t} \right).$$

While the above theorem is stated for homogeneous polynomials, the proof is identical also for the non-homogeneous ones, with the sum in the definition being over all tuple sizes.

2.1.2 *Applying the framework*

We now discuss how to interpret the matrices $\mathbf{F}_{a,b,c}$ arising in the bound in Theorem 2.1.1 to recover the norm bounds for graph matrices derived via combinatorial methods. This is discussed with more formal details in Section 2.6, but we present an intuitive (at least for the authors) version of the argument here.

Graph matrices are defined using constant-sized template "shapes" and provide a convenient basis for expressing (large) random matrices where the the entries are low-degree polynomials in the (normalized) indicators of edges in a $G_{n,p}$ random graph. Such matrices and their norm bounds have been used for a large number of applications [AMP21].

Let $N = \binom{n}{2}$ and fix a canonical bijection between the space $[N]$ and $\binom{[n]}{2}$, and let $[N]$ index the space of all possible edges in a random graph on n vertices. For each $e = \{i, j\} \in \binom{[n]}{2}$, for $\Delta_p = \sqrt{(1-p)/p}$ let x_e be independently $1/\Delta_p$ with probability $(1-p)$ and $-\Delta_p$ with probability p, so that $\mathbb{E}x_e = 0$, $\mathbb{E}x_e^2 = 1$, and $|x_e| \leq \Delta_p$.

Let $(V(\tau), E(\tau))$ be a graph on k vertices with d edges, and let U_τ, V_τ be *ordered* subsets of $V(\tau)$ of size k_1 and k_2 respectively. The tuple $\tau = (V(\tau), E(\tau), U_\tau, V_\tau)$ is referred to as a "shape" and is used to define the following graph matrix $\mathbf{M}_\tau$ with rows and columns indexed by $\mathcal{T}_n^{k_1}$ and $\mathcal{T}_n^{k_2}$ respectively

$$\mathbf{M}_\tau[\mathbf{i}, \mathbf{j}] := \sum_{\psi \in \mathcal{T}_n^k} \mathbb{1}_{\{\psi(U_\tau) = \mathbf{i}\}} \cdot \mathbb{1}_{\{\psi(V_\tau) = \mathbf{j}\}} \cdot \prod_{e_0 \in E(\tau)} x_{\psi(e_0)},$$

where we interpret a tuple $\psi \in \mathcal{T}_n^k$ as an injective function $\psi : V(\tau) \to [n]$ in the canonical way. If the function ψ maps U_τ to $\mathbf{i}$ (the row-index) and V_τ to $\mathbf{j}$ (the column index), we add the monomial to $\mathbf{M}_\tau[\mathbf{i}, \mathbf{j}]$ corresponding to the product of the images of all edges in the "pattern" $E(\tau)$ under ψ. Note that each nonzero entry of the matrix-valued function $\mathbf{M}$ is a multilinear homogeneous polynomial in the variables $\{x_e\}_{e \in \binom{[n]}{2}}$ with degree $d = |E(\tau)|$. We denote $\mathbf{M}_\tau$ by $\mathbf{F}$ for ease of notation in the discussion below.

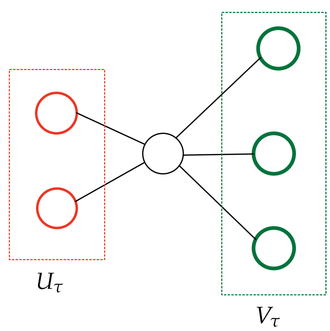

Figure 2.1: A shape τ

As will be apparent from the analysis below, the norm bounds for such matrices behave different characterizations in the "dense" regime where $p = \Omega(1)$ and the "sparse" one where $p = o(1)$. This is because the subgaussian norm $L = \Delta_p \approx p^{-1/2}$ is bounded in the first case, and growing in the second case. The bounds for the first case are stated in terms of the size the minimum vertex separator separating U_τ from V_τ in the shape τ.

For the second case, one needs to augment the defintion of minimum vertex separator, to use a cost based on the subgaussian norm Δ_p. We will show how to recover both these characterizations, using the matrices $\mathbf{F}_{a,b,c}$ given by Theorem 2.1.1.

Dense graph matrices. We start by considering $p = 1/2$ so that the random variables x_e are just independent Rademacher variables. Taking d to be constant, the bound in Theorem 2.1.1 involves only a constant number of matrices, and will depend on the dominant term. To understand each of the terms, we first consider the matrices derived after one step of the recursion.

$$\mathbf{F}_{1,0,0} = \begin{bmatrix} \partial_{x_1}\mathbf{F} & \dots & \partial_{x_N}\mathbf{F} \end{bmatrix}^{\top}, \ \mathbf{F}_{0,1,0} = \begin{bmatrix} \partial_{x_1}\mathbf{F} & \dots & \partial_{x_N}\mathbf{F} \end{bmatrix}$$

and

$$\mathbf{F}_{0,0,1} = \begin{bmatrix} \partial_{x_1}\mathbf{F} & & \\ & \ddots & \\ & & \partial_{x_N}\mathbf{F} \end{bmatrix}.$$

Since $\mathrm{tr}\left(\sum_i \mathbf{A}_i\mathbf{A}_i^{\top}\right)^{2t} + \mathrm{tr}\left(\sum_i \mathbf{A}_i^{\top}\mathbf{A}_i\right)^{2t} \ \geq \ \sum_i \mathrm{tr}\left(\mathbf{A}_i\mathbf{A}_i^{\top}\right)^{2t}$ and $L = 1$, we will have that $\left\|\mathbf{F}_{1,0,0}\right\|_{4t}^{4t} + \left\|\mathbf{F}_{0,1,0}\right\|_{4t}^{4t} \geq L^{4t} \cdot \left\|\mathbf{F}_{0,0,1}\right\|_{4t}^{4t}$. This will also be true (up to constant factors) for any constant L, since we interested in the $4t$-th root of the above trace. Using a similar argument, we can ignore terms with $c > 0$ for now (we will come back to these in the sparse case) and focus on matrices $\mathbf{F}_{a,b,0}$.

We now interpret the matrices $\mathbf{F}_{1,0,0}$ and $\mathbf{F}_{0,1,0}$. The row-space of $\mathbf{F}_{1,0,0}$ is now indexed by $(\mathbf{i}, e)$ for $e \in \binom{[n]}{2}$ (respectively $(\mathbf{j}, e)$ for the column space of $\mathbf{F}_{0,1,0}$). We now include terms corresponding to maps ψ for which $\psi(U_\tau) = \mathbf{i}$, $\psi(V_\tau) = \mathbf{j}$ and $\psi(e_0) = e$ for some $e_0 \in E(\tau)$ so that $\partial x_e \mathbf{F}[\mathbf{i}, \mathbf{j}] \neq 0$.

Decomposing $\mathbf{F}_{1,0,0}$ into $|E(\tau)|$ matrices (one for each choice of e_0), we can consider the $\mathbf{F}_{1,0,0,e_0}$ as simply a new graph matrix we have *deleted* the edge e_0 and included both its end points in U_τ (respectively in V_τ for $\mathbf{F}_{0,1,0}$) to obtain a new shape τ'. Similarly, each $\mathbf{F}_{a,b,0}$ can be decomposed into $|E(\tau)|^{a+b}$ graph matrices, where we delete $a+b$ edges in total, and include the endpoints in U_τ for a of them and V_τ for b of them. This is illustrated in Fig. 2.2 where we color the edges red or green depending on their inclusion in U_τ or V_τ.

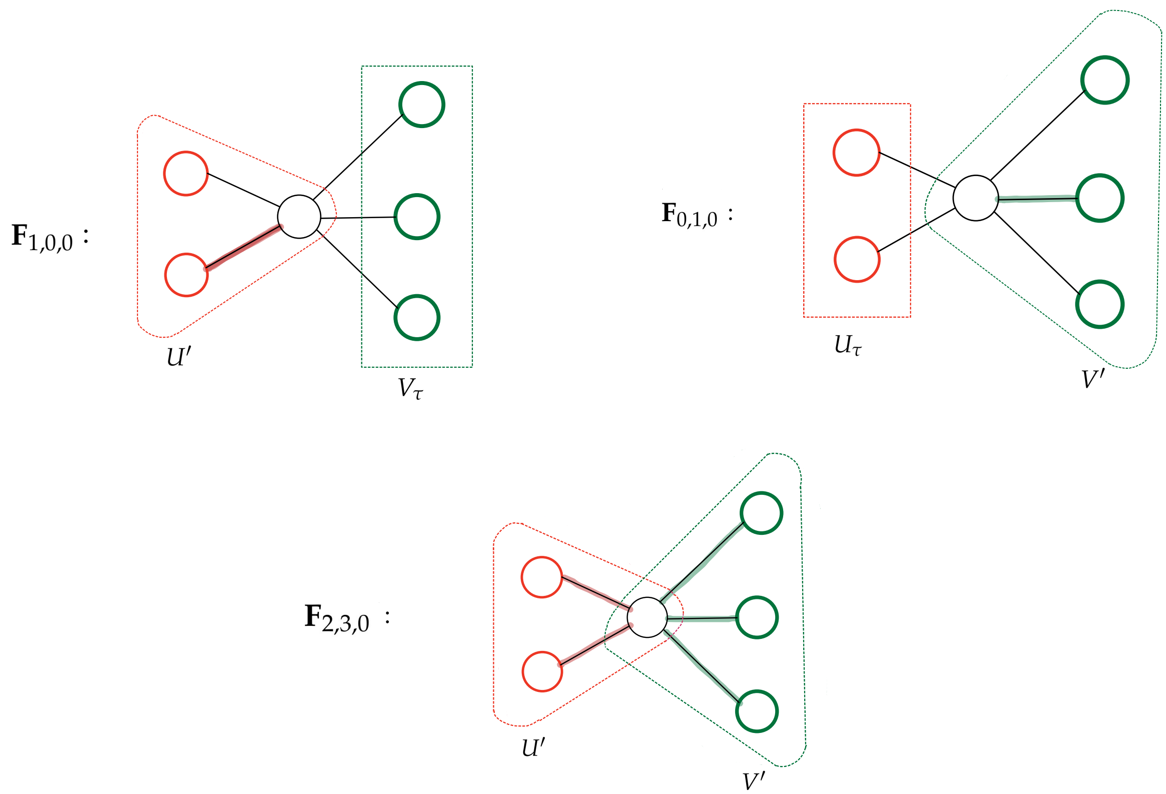

Figure 2.2: $\mathbf{F}_{1,0,0}$, $\mathbf{F}_{0,1,0}$ and $\mathbf{F}_{2,3,0}$

The bound in Theorem 2.1.1 is then in terms of these new graph matrices where we delete all d edges, and split their endpoints between U_τ or V_τ to obtain $U', V' \subseteq V(\tau) = [k]$. For an entry $\mathbf{M}[\mathbf{i}', \mathbf{j}']$ of any such matrix, there is at most one $\psi : V(\tau) \to [n]$ such that $\psi(U') = \mathbf{i}'$ and $\psi(V') = \mathbf{j}'$ (assuming τ has no isolated vertices), and $\mathbf{M}[\mathbf{i}', \mathbf{j}'] = \mathbb{1}_{\{\psi(U')=\mathbf{i}'\}} \cdot \mathbb{1}_{\{\psi(V')=\mathbf{j}'\}}$. Taking $S = U' \cap V'$, this matrix can be simply written as a block-diagonal matrix with $n^{|S|}$ of size $n^{|U' \setminus S|} \times n^{|V' \setminus S|}$. It is easy to check that the spectral norm of such a matrix is $n^{(|U' \setminus S| + |V' \setminus S|)/2} = n^{(k-|S|)/2}$ since $|U' \cup V'| = k$ and $S = U' \cap V'$.

Thus, each of the terms in the bound will correspond to graph matrices with different U' and V' (obtained by coloring edges red or green according to the above process) and the dominant term will be the one with the smallest value of $|U' \cap V'|$. We now claim that $U' \cap V'$ must be a vertex separator separating U_τ and V_τ in the shape τ i.e., any path between them must pass through $U' \cap V'$. If the path contains both red and green edges, then it contains one vertex with both red and green edges incident on it, which is then in $U' \cap V'$ (since U' is obtain by adding endpoint of red edges to U_τ and V' similarly for green edges). If path is entirely red, then we have an endpoint of a red edge in V_τ which is a vertex in $U' \cap V'$, and similarly for green paths.

Thus, we have $n^{(k-|S|)/2} \leq n^{(k-r)/2}$, where r is the size of the minimum vertex separator between U_τ and V_τ, which recovers the known bound for dense graph matrices [AMP21].

Sparse graph matrices. The argument for sparse graph matrices is almost the same as above except now we also need to consider terms of the form $\mathbf{F}_{a,b,c}$ for $c > 0$. Just as we interpreted $\mathbf{F}_{a,b,0}$ as coloring a edges from $E(\tau)$ as red and including in $U' \supseteq U_\tau$ and b green edges in $V' \supseteq V_\tau$, we can now interpret $\mathbf{F}_{a,b,c}$ as having c edges whose endpoints are included in *both* U' and V' (say these edges are colored yellow). This simply reflects the fact that the derivatives in the definition of $\mathbf{F}_{a,b,c+1}$ are placed as diagonal blocks.

As we saw above, increasing the intersection of U' and V' decreases the norm, and so the matrices with $c > 0$ should have a smaller Schatten norm. However, they are now included in the bound with a multiplicative factor of L^c. Thus, we simply look for a vertex separator S maximizing $L^{e(S)} \cdot n^{(k-|S|)/2}$ where $e(S)$ counts the (yellow) edges contained in $S = U' \cap V'$. This is precisely the "sparse vertex separator" which determines the norm bound for such sparse graph matrices [JPR+21, RT23].

$\mathbf{F}_{0,0,1}$:

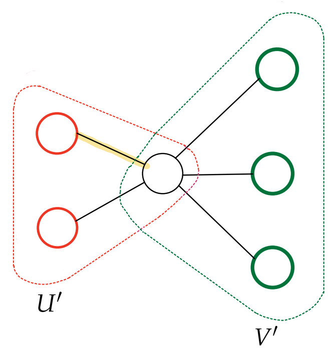

Figure 2.3: $\mathbf{F}_{0,0,1}$

2.2 Preliminaries and Notation

Vectors are denoted by bold face lower case letters $\mathbf{x}, \mathbf{y}$ and $\mathbf{z}$. Deterministic matrices are denoted by bold face upper case letters $\mathbf{A}$ and $\mathbf{B}$. Matrix-valued functions are denoted by bold face upper case letters $\mathbf{X}$, $\mathbf{Y}$, $\mathbf{F}$, $\mathbf{M}$, $\mathbf{H}$ and $\mathbf{P}$.

Sets and indices. For a set S, let $|S|$ denote the number of distinct elements in S. We denote $[n] = \{1, \ldots, n\}$ for any positive integer n. For $\mathbf{i} := (i_1, \ldots, i_m) \in [n]^m$, define

$$\mathcal{T}_n^m := \{(i_1, \ldots, i_m) : \forall j, k \in [m],\ j \neq k \Rightarrow i_j \neq i_k\}.$$

Matrix norms. Let $\mathbb{R}^{d_1 \times d_2}$ be the space of all $d_1 \times d_2$ real matrices. Let $\mathbb{H}^d$ be the subspace of $\mathbb{R}^{d \times d}$ containing all Hermitian matrices. We write $\|\cdot\|_2$ for the ℓ_2 operator norm, $\|\cdot\|_F$ for the Frobenius norm, and $\mathrm{tr}(\cdot)$ for the trace.

Definition 2.2.1 (Schatten norm). *For $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$ and $t \geq 1$, the Schatten $2t$-norm is defined as*

$$\|\mathbf{A}\|_{2t} := (\mathrm{tr}(\mathbf{A}^\top \mathbf{A})^t)^{1/2t}.$$

Polynomial random matrices. Let $(\Omega, \mathcal{F}, \mu)$ be a probability space. Introduce the random vector $\mathbf{x} := (x_1, \ldots, x_n) \in \Omega$ where $x_1, \ldots, x_n$ are independent random variables. A polynomial random matrix $\mathbf{F}(\mathbf{x})$ is a matrix-valued polynomial

$$\mathbf{F} : \Omega \to \mathbb{R}^{d_1 \times d_2}.$$

We can construct a polynomial random matrix $\mathbf{F}(\mathbf{x})$ by drawing a copy of $\mathbf{x} \sim \mu$. Note that despite $\mathbf{x}$ having independent entries, the entries of $\mathbf{F}(\mathbf{x})$ can be dependent (or correlated) with each other.

Example 2.2.2 (Rademacher chaos of order 2)**.** *We give an example of a degree-2 polynomial random matrices in 3 variables. Let x_1, x_2, x_3 be i.i.d. Rademacher random variables and $\{\mathbf{B}_\mathbf{i}\}_{\mathbf{i} \in \mathcal{T}_3^2}$ be a double sequence of deterministic matrices (explicitly chosen here) of the same dimension.*

$$\sum_{\mathbf{i} \in \mathcal{T}_3^2} \mathbf{B}_\mathbf{i} \cdot x_i x_j = x_1 x_2 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + x_1 x_3 \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} + x_2 x_3 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} x_1 x_2 & x_2 x_3 \\ x_2 x_3 & x_1 x_3 \end{bmatrix}. \tag{2.1}$$

Remark 2.2.3. *Unlike Wigner matrices whose entries are i.i.d., the polynomial random matrices can have dependencies among its entries.*

Definition 2.2.4 (Permutation Symmetric Property)**.** *Suppose $\mathbf{F}(\mathbf{x})$ is a degree-d homogeneous multilinear polynomial random matrix, i.e.,*

$$\mathbf{F}(\mathbf{x}) = \sum_{\mathbf{i} \in \mathcal{T}_n^d} \left(\mathbf{A}_\mathbf{i} \cdot \prod_{j \in \{i_1, \ldots, i_d\}} x_j \right),$$

where $\{\mathbf{A}_\mathbf{i}\}_{\mathbf{i} \in \mathcal{T}_n^d}$ is a multi-indexed sequence of deterministic matrices of the same dimension. $\mathbf{F}(\mathbf{x})$ is permutation symmetric if $\mathbf{A}_{i_1, \ldots, i_d} = \mathbf{A}_{i_{\sigma(1)}, \ldots, i_{\sigma(d)}}$ for any permutation $\sigma \in \mathfrak{S}_d$ and any $(i_1, \ldots, i_d) \in \mathcal{T}_n^d$.

Remark 2.2.5. *The polynomial random matrix (2.1) in Example 2.2.2 is permutation symmetric if we formally rewrite it as*

$$
\begin{aligned}
\sum_{\mathbf{i}\in\mathcal{T}_3^2} \mathbf{B}_{\mathbf{i}} \cdot x_i x_j &= x_1x_2 \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & 0 \end{bmatrix} + x_2x_1 \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & 0 \end{bmatrix} + x_1x_3 \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{2} \end{bmatrix} + x_3x_1 \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{2} \end{bmatrix} \\
&+ x_2x_3 \begin{bmatrix} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{bmatrix} + x_3x_2 \begin{bmatrix} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{bmatrix}.
\end{aligned}
$$

The purpose of this technical detail will become clear in Theorem 2.4.2 and Theorem 2.4.5.

Partial derivatives. Let $\mathbf{F}(\mathbf{x})$ be an n-variate polynomial random matrix of degree D. For $a, b, c \in \mathbb{Z}_{\geq 0}$ and $d = a + b + c$, we consider the block matrices $\mathbf{F}_{a,b,c}$ containing all d-th order partial derivatives of $\mathbf{F}(\mathbf{x})$. We define $\mathbf{F}_{a,b,c}$ recursively as $\mathbf{F}_{0,0,0} = \mathbf{F}(\mathbf{x})$ and

$$\mathbf{F}_{a+1,b,c} = \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b,c} & \dots & \partial_{x_n}\mathbf{F}_{a,b,c} \end{bmatrix}^{\mathsf{T}},$$

$$\mathbf{F}_{a,b+1,c} = \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b,c} & \dots & \partial_{x_n}\mathbf{F}_{a,b,c} \end{bmatrix},$$

and

$$\mathbf{F}_{a,b,c+1} = \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b,c} & & \\ & \ddots & \\ & & \partial_{x_n}\mathbf{F}_{a,b,c} \end{bmatrix}.$$

It is evident from the definition that the block matrices $\mathbf{F}_{a+1,b,c}$, $\mathbf{F}_{a,b+1,c}$, $\mathbf{F}_{a,b,c+1}$ are assembled from sub-blocks $\{\partial_{x_i}\mathbf{F}_{a,b,c}\}_{i=1}^{n}$ arranged vertically, horizontally and diagonally respectively. The order of increment between a and b doesn't affect the resulting matrix, i.e.,

$$\mathbf{F}_{a+1,b+1,c} = \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b+1,c} \\ \vdots \\ \partial_{x_n}\mathbf{F}_{a,b+1,c} \end{bmatrix} = \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a+1,b,c} & \dots & \partial_{x_n}\mathbf{F}_{a+1,b,c} \end{bmatrix}$$
$$= \begin{bmatrix} \partial_{x_1}\partial_{x_1}\mathbf{F}_{a,b,c} & \dots & \partial_{x_1}\partial_{x_n}\mathbf{F}_{a,b,c} \\ & \vdots & \\ \partial_{x_n}\partial_{x_1}\mathbf{F}_{a,b,c} & \dots & \partial_{x_n}\partial_{x_n}\mathbf{F}_{a,b,c} \end{bmatrix}.$$

However, the order of increment involving c affects the resulting matrix, i.e.,

$$\begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b,c+1} & \dots & \partial_{x_n}\mathbf{F}_{a,b,c+1} \end{bmatrix} \neq \begin{bmatrix} \partial_{x_1}\mathbf{F}_{a,b+1,c} & & \\ & \ddots & \\ & & \partial_{x_n}\mathbf{F}_{a,b+1,c} \end{bmatrix}.$$

Nevertheless, the order of increment doesn't affect the Schatten norm. Since we will only be interested in the Schatten norms of these matrices, we will simply label them as $\mathbf{F}_{a,b+1,c+1}$ without specifying the order. It is easy to see that $\mathbf{F}_{a,b,c}$ is a deterministic matrix if $a+b+c=D$ and $\mathbf{F}_{a,b,c}=\mathbf{0}$ if $a+b+c>D$.

Non-Hermitian matrices. Many results for Hermitian matrices can be easily extended to non-Hermitian matrices by the following argument.

Definition 2.2.6 (Hermitian Dilation, see [Tro15] Sec. 2.1.16)**.** *For each $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$, the Hermitian dilation $\mathcal{H} : \mathbb{R}^{d_1 \times d_2} \to \mathbb{H}^{(d_1+d_2)\times(d_1+d_2)}$ is defined as*

$$\mathcal{H}(\mathbf{A}) = \begin{bmatrix} \mathbf{0} & \mathbf{A} \\ \mathbf{A}^\top & \mathbf{0} \end{bmatrix}.$$

Since the square of the Hermitian dilation satisfies $\mathcal{H}(\mathbf{A})^2 = \begin{bmatrix} \mathbf{A}\mathbf{A}^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{A}^\top\mathbf{A} \end{bmatrix}$, it follows that $\|\mathcal{H}(\mathbf{A})\| = \|\mathbf{A}\|$.

Constants. We denote C for some universal constants and C_a for constants depending only on some parameter a. The values of the constants may differ from one instance to another.

2.3 Decoupling Inequalities

In this work, we focus on decoupling inequalities for moments. The following decoupling inequality holds in surprising generality.

Lemma 2.3.1 (Decoupling Inequality, see [PG99] Theorem 3.1.1)**.** *For natural numbers $n \geq m$, let $\mathbf{x} = \{x_i\}_{i=1}^n$ be a sequence of n independent random variables in a measurable space $(S, \mathcal{S})$, and let $\{x_i^{(k)}\}_{i=1}^n$, $k = 1, \dots, m$, be m independent copies of this sequence. Let B be a separable Banach space, and let $h_\mathbf{i} : S^m \to B$ be a measurable function such that $\mathbb{E}(\left\|h_\mathbf{i}(x_{i_1}, \dots, x_{i_m})\right\|) \leq \infty$ for each $\mathbf{i} \in \mathcal{T}_n^m$. Let $\Phi : [0, \infty) \to [0, \infty)$ be a convex nondecreasing function such that $\mathbb{E}\Phi(\left\|h_\mathbf{i}(x_{i_1}, \dots, x_{i_m})\right\|) \leq \infty$ for all $\mathbf{i} \in \mathcal{T}_n^m$. Then,*

$$\mathbb{E}\Phi\left(\left\|\sum_{\mathbf{i}\in\mathcal{T}_n^m} h_\mathbf{i}(x_{i_1}, \dots, x_{i_m})\right\|\right) \leq \mathbb{E}\Phi\left(C_m \left\|\sum_{\mathbf{i}\in\mathcal{T}_n^m} h_\mathbf{i}(x_{i_1}^{(1)}, \dots, x_{i_m}^{(m)})\right\|\right),$$

where $C_m = 2^m \cdot (m^m - 1) \cdot ((m-1)^{(m-1)} - 1) \cdot \dots \cdot 3$.

In general, the decoupling constant C_m is super-exponential in m, which might seem quite large. But nevertheless C_m is independent of the dimension n. Furthermore, C_m can be improved significantly for some special classes of functions h and certain underlying distributions of $\mathbf{x}$. For instance, [Kwa87] showed that if h is a homogeneous multilinear form of degree m, then $C_m = m^m$. Moreover, if we assume $\mathbf{x}$ is a Gaussian random vector, C_m further improves to $m^{m/2}$.

While the proof of decoupling inequality for tail probability is slightly involved [PMS95], the proof of decoupling inequalities for moments are fairly elementary and can be interpreted combinatorially as partitioning the random vector $\mathbf{x}$. We present the following lemma that generalizes Lemma 6.21 [Rau10] to arbitrary degree d.

Lemma 2.3.2. *Let $\mathbf{x} = \{x_j\}_{j=1}^n$ be a sequence of independent random variables with $\mathbb{E}x_j = 0$ for all $j = 1, \ldots, n$. Let $\{\mathbf{B_i}\}_{\mathbf{i} \in \mathcal{T}_n^d}$ be a multi-indexed sequence of deterministic matrices of the same dimension. Then for $1 \leq t \leq \infty$*

$$\mathbb{E} \left\| \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{B_i} \cdot x_{i_1} \ldots x_{i_d} \right\|_{2t} \leq d^d \cdot \mathbb{E} \left\| \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{B_i} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)} \right\|_{2t},$$

where $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)}$ denote d independent copies of $\mathbf{x}$.

Proof. Consider a random partition of $\mathbf{x}$ into d parts and let $\mathbf{r}$ be the *partitioner*. Formally, $\mathbf{r} = (r_1, \ldots, r_n)$ is a sequence of independent random variables with each r_k uniformly distributed on $\{1, \ldots, d\}$, i.e., for $1 \leq k \leq n$,

$$\mathbb{P}(r_k = 1) = \mathbb{P}(r_k = 2) = \cdots = \mathbb{P}(r_k = d) = 1/d.$$

Define an event $E := \{r_{i_1} = 1, \cdots, r_{i_d} = d\}$ for each $\mathbf{i} \in \mathcal{T}_n^d$. Conditioned on $\mathbf{x}$,

$$\mathbb{E}_{\mathbf{r}} \left[\mathbb{1}_E(r_{i_1}, \ldots, r_{i_d}) \right] = 1/d^d.$$

It follows that

$$
\begin{aligned}
F := \; & \mathbb{E}_{\mathbf{x}} \left\| \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{B}_{\mathbf{i}} \cdot x_{i_1} \ldots x_{i_d} \right\|_{2t} \\
= \; & d^d \cdot \mathbb{E}_{\mathbf{x}} \left\| \mathbb{E}_{\mathbf{r}} \left[\sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbb{1}_E(r_{i_1}, \ldots, r_{i_d}) \cdot \mathbf{B}_{\mathbf{i}} \cdot x_{i_1} \ldots x_{i_d} \right] \right\|_{2t} \\
\leq \; & d^d \cdot \mathbb{E}_{\mathbf{r}} \mathbb{E}_{\mathbf{x}} \left\| \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbb{1}_E(r_{i_1}, \ldots, r_{i_d}) \cdot \mathbf{B}_{\mathbf{i}} \cdot x_{i_1} \ldots x_{i_d} \right\|_{2t},
\end{aligned}
$$

where the last step is due to Jensen's inequality and Fubini's theorem. Since the expectation over $\mathbf{r}$ satisfies the inequality, this implies the existence of an $\mathbf{r}^*$ satisfying the inequality.

Fix $\mathbf{r}^*$ and define the *partitioning* with respect to $\mathbf{r}^*$ as

$$
P_1 = \{k \in [n] : r_k^* = 1\}, \; \ldots \;, P_d = \{k \in [n] : r_k^* = d\}.
$$

Since these partitions have no intersections and x_j's are independent random variables, replacing $x_{i_1} \cdots x_{i_d}$ with $x_{i_1}^{(1)} \cdots x_{i_d}^{(d)}$ for any $\mathbf{i} \in P_1 \times \cdots \times P_d$ will not change the distribution. Hence

$$
F \; \leq \; d^d \cdot \mathbb{E} \left\| \sum_{\mathbf{i} \in P_1 \times \cdots \times P_d} \mathbf{B}_{\mathbf{i}} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)} \right\|_{2t}. \tag{2.2}
$$

Denote all variables in $P_1 \times \cdots \times P_d$ by

$$
\mathcal{X} := \{x_{i_1}^{(1)}, \ldots, x_{i_d}^{(d)} \, : \, \forall \, i_1 \in P_1, \; \ldots \;, i_d \in P_d\},
$$

and denote the rest of the variables by $\mathcal{X}^c$. It remains to show that the sum on the right-hand side of (2.2) is over all $\mathbf{i} \in \mathcal{T}_n^d$. Observe that

$$\mathbb{E}_{\mathcal{X}^c}\left[\sum_{\mathbf{i} \notin P_1 \times \cdots \times P_d} \mathbf{B_i} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)} \;\middle|\; \mathcal{X}\right] = \mathbf{0}. \tag{2.3}$$

Since the sum in (3.1) has conditional expectation zero, we can add it to (2.2) to get

$$F \leq d^d \cdot \mathbb{E}_{\mathcal{X}}\left\|\sum_{\mathbf{i} \in P_1 \times \cdots \times P_d} \mathbf{B_i} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)} + \mathbb{E}_{\mathcal{X}^c}\left[\sum_{\mathbf{i} \notin P_1 \times \cdots \times P_d} \mathbf{B_i} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)} \;\middle|\; \mathcal{X}\right]\right\|_{2t}.$$

Since the two sums on the right-hand side are independent conditioned on $\mathcal{X}$,

$$F \leq d^d \cdot \mathbb{E}_{\mathcal{X}}\left\|\mathbb{E}_{\mathcal{X}^c}\left[\sum_{\mathbf{i} \in P_1 \times \cdots \times P_d} \mathbf{B_i} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)} + \sum_{\mathbf{i} \notin P_1 \times \cdots \times P_d} \mathbf{B_i} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)} \;\middle|\; \mathcal{X}\right]\right\|_{2t}.$$

A conditional application of Jensen's inequality yields the desired inequality,

$$F \leq d^d \cdot \mathbb{E}_{\mathbf{x}^{(1)} \ldots \mathbf{x}^{(d)}}\left\|\sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{B_i} \cdot x_{i_1}^{(1)} \ldots x_{i_d}^{(d)}\right\|_{2t}. \qquad \blacksquare$$

2.3.1 An improved decoupling inequality

For a generic degree d multilinear polynomial random matrices, an application of Lemma 2.3.2 yields a decoupling constant of d^d, which might not be optimal for polynomial random matrices with additional structures. We will show that the additional structures enable us to give a tighter decoupling inequality through a careful partitioning.

Building on Lemma 2.3.2, we prove a decoupling inequality for polynomial random matrices, in which the indices of random variables have a graph structure. Let $G = (V, E)$ be a fixed simple graph with $|V| = k$, and $|E| = d$. We denote the vertices by $V(G) = \{v_1, \ldots, v_k\}$ and the ordered edge set by $E(G) = \{(i_1, j_1), \ldots, (i_d, j_d)\}$. For any $\varphi \in \mathcal{T}_n^k$, we write $\varphi(i)$ for the i-th component of φ and $\varphi(E)$ for $\{(\varphi(i_1), \varphi(j_1)), \dots, (\varphi(i_d), \varphi(j_d))\}$.

Lemma 2.3.3. *Let $\mathbf{z} = \{Z_{\varphi(i),\varphi(j)}\}_{\varphi \in \mathcal{T}_n^k, (i,j) \in E}$ be a double sequence of independent random variables with $\mathbb{E} Z_{\varphi(i),\varphi(j)} = 0$ for all $(i,j) \in E$ and $\varphi \in \mathcal{T}_n^k$. Let $\{\mathbf{B}_{\varphi(E)}\}_{\varphi \in \mathcal{T}_n^k}$ be a multi-indexed sequence of deterministic matrices of the same dimension. Then for $d \le k(k-1)$ and $1 \le t \le \infty$,*

$$\mathbb{E} \left\| \sum_{\varphi \in \mathcal{T}_n^k} \mathbf{B}_{\varphi(E)} \cdot \prod_{(i,j) \in E} Z_{\varphi(i),\varphi(j)} \right\|_{2t} \le k^k \cdot \mathbb{E} \left\| \sum_{\varphi \in \mathcal{T}_n^k} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \right\|_{2t},$$

where $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(d)}$ denote d independent copies of $\mathbf{z}$.

Remark 2.3.4. *It is crucial for the indices of all monomials to share the same structure so that d independent copies of $\mathbf{z}$ suffices. Otherwise, $k(k-1)$ independent copies are needed in general.*

Proof. Consider a random partition of the ground set $[n]$ into k parts and let $\mathbf{r}$ be the *vertex partitioner*. Formally, let $\mathbf{r} = (r_1, \dots, r_n)$ be a sequence of independent random variables with each r_p uniformly distributed on $\{1, \dots, k\}$, i.e., for $1 \le p \le n$,

$$\mathbb{P}(r_p = 1) = \mathbb{P}(r_p = 2) = \cdots = \mathbb{P}(r_p = k) = 1/k.$$

Define an event $A := \{r_{\varphi(v_1)} = 1, \cdots, r_{\varphi(v_k)} = k\}$ for each $\varphi \in \mathcal{T}_n^k$. Conditioned on $\mathbf{z}$,

$$\mathbb{E}_{\mathbf{r}} \left[\mathbb{1}_A (r_{\varphi(v_1)}, \dots, r_{\varphi(v_k)}) \right] = 1/k^k.$$

It follows that

$$
\begin{aligned}
F :&= \mathbb{E}_{\mathbf{z}} \left\| \sum_{\varphi \in \mathcal{T}_n^k} \mathbf{B}_{\varphi(E)} \cdot \prod_{(i,j) \in E} Z_{\varphi(i),\varphi(j)} \right\|_{2t} \\
&= k^k \cdot \mathbb{E}_{\mathbf{z}} \left\| \mathbb{E}_{\mathbf{r}} \left[\sum_{\varphi \in \mathcal{T}_n^k} \mathbb{1}_A(r_{\varphi(v_1)}, \ldots, r_{\varphi(v_k)}) \cdot \mathbf{B}_{\varphi(E)} \cdot \prod_{(i,j) \in E} Z_{\varphi(i),\varphi(j)} \right] \right\|_{2t} \\
&\le k^k \cdot \mathbb{E}_{\mathbf{r}} \mathbb{E}_{\mathbf{z}} \left\| \sum_{\varphi \in \mathcal{T}_n^k} \mathbb{1}_A(r_{\varphi(v_1)}, \ldots, r_{\varphi(v_k)}) \cdot \mathbf{B}_{\varphi(E)} \cdot \prod_{(i,j) \in E} Z_{\varphi(i),\varphi(j)} \right\|_{2t},
\end{aligned}
$$

where the last step is due to Jensen's inequality and Fubini's theorem. Since the expectation over $\mathbf{r}$ satisfies the inequality, this implies the existence of an $\mathbf{r}^*$ satisfying the inequality.

Fix $\mathbf{r}^*$ and define the *vertex partitioning* with respect to $\mathbf{r}^*$ as

$$
P_1 = \{p \in [n] : r_p^* = 1\}, \ \ldots \ , P_k = \{p \in [n] : r_p^* = k\},
$$

which in turn induces an *edge partitioning*,

$$
P_e \ := \ \{(P_i, P_j)\}_{(i,j) \in \mathcal{T}_k^2}.
$$

Notice that there are $k(k-1)$ elements in P_e, but we only have d edges to partition. So we select d elements out of P_e in the following way. Fix an arbitrary $\varphi^* \in \mathcal{T}_n^k$, choose

$$
P_1^* = (P_{\varphi^*(i_1)}, P_{\varphi^*(j_1)}), \ \ldots \ , P_d^* = (P_{\varphi^*(i_d)}, P_{\varphi^*(j_d)}).
$$

Let's call $\{P_1^*, \ldots P_d^*\}$ a *d-edge partitioning*. In other words, we may select any d elements out of P_e to form a *d-edge partitioning* so long as the pattern of the indices of (P_i, P_j)'s matches that of $E(G)$. Since the elements in P_e have no intersections, the partitions in $\{P_1^*, \ldots, P_d^*\}$ also have no intersections. Since $Z_{\varphi(i),\varphi(j)}$'s are independent random variables, replacing

$Z_{\varphi(i_1),\varphi(j_1)} \cdots Z_{\varphi(i_d),\varphi(j_d)}$ with $Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)}$ for any $\varphi(E) \in P_1^* \times \cdots \times P_d^*$ will not change the distribution. Hence

$$F \leq k^k \cdot \mathbb{E} \left\| \sum_{\varphi(E) \in P_1^* \times \cdots \times P_d^*} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \right\|_{2t}. \tag{2.4}$$

Denote all variables in $P_1^* \times \cdots \times P_d^*$ by

$$\mathcal{Z} := \left\{ Z^{(1)}_{\varphi(i_1),\varphi(j_1)}, \ \ldots, \ Z^{(d)}_{\varphi(i_d),\varphi(j_d)} : \forall \ (\varphi(i_1), \varphi(j_1)) \in P_1^*, \ \ldots \ , (\varphi(i_d), \varphi(j_d)) \in P_d^* \right\},$$

and denote the rest of the variables by $\mathcal{Z}^c$. It remains to show that the sum on the right-hand side of (2.4) is over all $\varphi \in \mathcal{T}_n^k$. Observe that

$$\mathbb{E}_{\mathcal{Z}^c} \left[\sum_{\varphi(E) \notin P_1^* \times \cdots \times P_d^*} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \ \middle| \ \mathcal{Z} \right] = \mathbf{0}. \tag{2.5}$$

Denote $P_1^* \times \cdots \times P_d^*$ by $\mathcal{P}$. Since the sum in (2.5) has conditional expectation zero, we can add it to (2.4) to get

$$\begin{aligned} F \ &\leq k^k \cdot \mathbb{E}_{\mathcal{Z}} \left\| \sum_{\varphi(E) \in \mathcal{P}} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \right. \\ &\qquad \left. + \mathbb{E}_{\mathcal{Z}^c} \left[\sum_{\varphi(E) \notin \mathcal{P}} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \ \middle| \ \mathcal{Z} \right] \right\|_{2t}. \end{aligned}$$

Since the two sums on the right-hand side are independent conditioned on $\mathcal{Z}$,

$$\begin{aligned} F \ &\leq k^k \cdot \mathbb{E}_{\mathcal{Z}} \left\| \mathbb{E}_{\mathcal{Z}^c} \left[\sum_{\varphi(E) \in \mathcal{P}} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \right. \right. \\ &\qquad \left. \left. + \sum_{\varphi(E) \notin \mathcal{P}} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \ \middle| \ \mathcal{Z} \right] \right\|_{2t}. \end{aligned}$$

A conditional application of Jensen's inequality yields the desired inequality,

$$F \leq k^k \cdot \mathbb{E}_{\mathbf{z}^{(1)},\ldots,\mathbf{z}^{(d)}} \left\| \sum_{\varphi \in \mathcal{T}_n^k} \mathbf{B}_{\varphi(E)} \cdot Z^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots Z^{(d)}_{\varphi(i_d),\varphi(j_d)} \right\|_{2t}. \qquad \blacksquare$$

2.4 Multilinear Polynomial Random Matrices

In this section, we prove a moment bound for multilinear polynomial random matrices with bounded, normalized random variables. We follow our recursion framework by decoupling the polynomial random matrices first and then apply the matrix Rosenthal inequality recursively to sums of (conditionally) independent random matrices.

To this end, we derive the non-Hermitian matrix Rosenthal inequality from the Hermitian matrix Rosenthal inequality by Hermitian dilation.

Lemma 2.4.1 (Non-Hermitian Matrix Rosenthal Inequality)**.** *Suppose that $t = 1$ or $t \geq 1.5$. Consider a finite sequence $\{\mathbf{Y}_k\}_{k\geq 1}$ of centered, independent, random matrices, and assume that $\mathbb{E}\,\|\mathbf{Y}_k\|_{4t}^{4t} < \infty$. Then*

$$\mathbb{E}\left\|\sum_k \mathbf{Y}_k\right\|_{4t}^{4t} \leq (16t)^{3t} \cdot \left\{ \left\| \left(\sum_k \mathbb{E}\mathbf{Y}_k\mathbf{Y}_k^\top\right)^{1/2} \right\|_{4t}^{4t} + \left\| \left(\sum_k \mathbb{E}\mathbf{Y}_k^\top\mathbf{Y}_k\right)^{1/2} \right\|_{4t}^{4t} \right\}$$
$$+ (8t)^{4t} \cdot \left(\sum_k \mathbb{E}\,\|\mathbf{Y}_k\|_{4t}^{4t}\right).$$

Proof. For a sequence of Hermitian matrices $\{\mathbf{X}_k\}_{k\geq 1}$ satisfying all the assumptions above, we have (Hermitian) matrix Rosenthal inequality [MJC+12] Corollary 7.4,

$$\mathbb{E}\left\|\sum_k \mathbf{X}_k\right\|_{4t}^{4t} \leq (16t)^{3t} \cdot \left\| \left(\sum_k \mathbb{E}\mathbf{X}_k^2\right)^{1/2} \right\|_{4t}^{4t} + (8t)^{4t} \cdot \sum_k \mathbb{E}\,\|\mathbf{X}_k\|_{4t}^{4t}.$$

We use Hermitian dilation to extend the inequality to non-Hermitian matrices by setting

$\mathbf{X}_k = \mathcal{H}(\mathbf{Y}_k)$ and notice that

$$\mathbb{E}\left\|\sum_k \begin{bmatrix} \mathbf{0} & \mathbf{Y}_k \\ \mathbf{Y}_k^\top & \mathbf{0} \end{bmatrix}\right\|_{4t}^{4t} = \mathbb{E}\,\mathrm{tr}\left(\begin{bmatrix} ((\sum_k \mathbf{Y}_k)(\sum_k \mathbf{Y}_k^\top))^{2t} & \mathbf{0} \\ \mathbf{0} & ((\sum_k \mathbf{Y}_k^\top)(\sum_k \mathbf{Y}_k))^{2t} \end{bmatrix}\right)$$

$$= 2\mathbb{E}\left\|\sum_k \mathbf{Y}_k\right\|_{4t}^{4t},$$

$$\left\|\left(\sum_k \mathbb{E}\begin{bmatrix} \mathbf{Y}_k\mathbf{Y}_k^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{Y}_k^\top\mathbf{Y}_k \end{bmatrix}\right)^{1/2}\right\|_{4t}^{4t} = \mathrm{tr}\left(\begin{bmatrix} \sum_k \mathbb{E}\mathbf{Y}_k\mathbf{Y}_k^\top & \mathbf{0} \\ \mathbf{0} & \sum_k \mathbb{E}\mathbf{Y}_k^\top\mathbf{Y}_k \end{bmatrix}\right)^{2t}$$

$$= \left\|\left(\sum_k \mathbb{E}\mathbf{Y}_k\mathbf{Y}_k^\top\right)^{1/2}\right\|_{4t}^{4t} + \left\|\left(\sum_k \mathbb{E}\mathbf{Y}_k^\top\mathbf{Y}_k\right)^{1/2}\right\|_{4t}^{4t},$$

$$\sum_k \mathbb{E}\left\|\begin{bmatrix} \mathbf{0} & \mathbf{Y}_k \\ \mathbf{Y}_k^\top & \mathbf{0} \end{bmatrix}\right\|_{4t}^{4t} = 2\sum_k \mathbb{E}\left\|\mathbf{Y}_k\right\|_{4t}^{4t}.$$

Hence, the result follows. ■

For the ease of notation, we demonstrate our recursion framework in a special case of quadratic form first. The multilinear form follows exactly the same reasoning.

Theorem 2.4.2 (Quadratic Recursion)**.** *Let $\{x_i\}_{i=1}^n$ be a sequence of i.i.d random variables with $\mathbb{E}x_i = 0$, $\mathbb{E}x_i^2 = 1$ and $|x_i| \le L$ for all $1 \le i \le n$. Let $\{\mathbf{A}_{i,j}\}_{(i,j)\in\mathcal{T}_n^2}$ be a double sequence of deterministic matrices of the same dimension. Define*

$$\mathbf{F}(\mathbf{x}) = \sum_{(i,j)\in\mathcal{T}_n^2} \mathbf{A}_{i,j}x_ix_j.$$

Suppose that $\mathbf{F}(\mathbf{x})$ is also permutation symmetric, i.e., $\mathbf{A}_{i,j} = \mathbf{A}_{j,i}$, $\forall (i,j) \in \mathcal{T}_n^2$. Then for

$a, b, c \in \mathbb{Z}_{\geq 0}$ and $2 \leq t < \infty$,

$$\mathbb{E}\,\|\mathbf{F}(\mathbf{x}) - \mathbb{E}\mathbf{F}(\mathbf{x})\|_{4t} \;\leq\; 2(32t)^2 \cdot \left(\sum_{\substack{a,b,c: \\ a+b+c=2}} L^c \cdot \left\|\mathbf{F}_{a,b,c}\right\|_{4t} \right).$$

Remark 2.4.3. *For $a, b, c \in \mathbb{Z}_{\geq 0}$ such that $a + b + c = 2$, the partial derivative matrices $\{\mathbf{F}_{a,b,c}\}$ are block matrices whose blocks consist of $\{\mathbf{A}_{i,j}\}_{(i,j)\in\mathcal{T}_n^2}$ and zero matrices. For instance,*

$$\mathbf{F}_{1,1,0} = \begin{bmatrix} \mathbf{0} & \mathbf{A}_{1,2} & \dots & \mathbf{A}_{1,n} \\ \mathbf{A}_{2,1} & \mathbf{0} & \dots & \mathbf{A}_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{n,1} & \mathbf{A}_{n,2} & \dots & \mathbf{0} \end{bmatrix},$$

$$\mathbf{F}_{0,2,0} = \begin{bmatrix} \mathbf{0} & \mathbf{A}_{1,2} & \dots & \mathbf{A}_{1,n} | & \dots & |\mathbf{A}_{n,1} & \dots & \mathbf{A}_{n,n-1} & \mathbf{0} \end{bmatrix},$$

and

$$\mathbf{F}_{0,1,1} = \begin{bmatrix} \mathbf{0} \;\; \mathbf{A}_{1,2} \;\; \dots \;\; \mathbf{A}_{1,n} & & & \\ & \mathbf{A}_{2,1} \;\; \mathbf{0} \;\; \dots \;\; \mathbf{A}_{2,n} & & \\ & & \ddots & \\ & & & \mathbf{A}_{n,1} \;\; \mathbf{A}_{n,2} \;\; \dots \;\; \mathbf{0} \end{bmatrix}.$$

Proof. Let $\mathbf{x}' = \{x_j'\}_{j=1}^n$ be an independent copy of $\mathbf{x} = \{x_i\}_{i=1}^n$. We decouple $\mathbf{F}(\mathbf{x})$ by Lemma 2.3.2,

$$E := \mathbb{E}\left\|\sum_{(i,j)\in\mathcal{T}_n^2} \mathbf{A}_{i,j} x_i x_j\right\|_{4t}^{4t} \leq 4^{4t} \cdot \mathbb{E}\left\|\sum_{(i,j)\in\mathcal{T}_n^2} \mathbf{A}_{i,j} x_i x_j'\right\|_{4t}^{4t}.$$

Since the decoupled $\mathbf{F}(\mathbf{x})$ on the right hand side is a function of $\mathbf{x}$ and $\mathbf{x}'$, the decoupled

$\mathbf{F}(\mathbf{x})$ can be viewed as taking two arguments instead. Thus we will refer to the decoupled $\mathbf{F}(\mathbf{x})$ as $\mathbf{F}(\mathbf{x}, \mathbf{x}')$.

By the law of total expectation,

$$E \le 4^{4t} \cdot \mathbb{E}\left(\mathbb{E}_{\mathbf{x}}\left(\left\|\sum_{i=1}^{n} x_i \left(\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right)\right\|_{4t}^{4t} \,\middle|\, \mathbf{x}'\right)\right).$$

Since $\mathbf{x}$ and $\mathbf{x}'$ are independent, if we fix $\mathbf{x}'$, the summation becomes a sum of centered, (conditionally) independent random matrices in $\mathbf{x}$. So we can apply matrix Rosenthal inequality (Lemma 2.4.1) and use the fact that $\mathbb{E}x_i^2 = 1$, and x_i's are independent to get

$$E \le 4^{4t}(16t)^{3t} \cdot \mathbb{E}\left\|\left(\sum_{i=1}^{n}\left(\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right)\left(\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right)^{\mathsf{T}}\right)^{1/2}\right\|_{4t}^{4t} \tag{2.6}$$

$$+ 4^{4t}(16t)^{3t} \cdot \mathbb{E}\left\|\left(\sum_{i=1}^{n}\left(\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right)^{\mathsf{T}}\left(\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right)\right)^{1/2}\right\|_{4t}^{4t} \tag{2.7}$$

$$+ 4^{4t}(8t)^{4t} L^{4t} \cdot \sum_{i=1}^{n} \mathbb{E}\left\|\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right\|_{4t}^{4t}. \tag{2.8}$$

Recall that $\mathbf{F}(\mathbf{x}, \mathbf{x}')$ refers to the decoupled $\mathbf{F}(\mathbf{x})$. The expression in (2.6) satisfies

$$\sum_{i=1}^{n}\left(\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right)\left(\sum_{j=1}^{n} \mathbf{A}_{i,j} x_j'\right)^{\mathsf{T}} = \begin{bmatrix}\partial_{x_1}\mathbf{F}(\mathbf{x}, \mathbf{x}') & \dots & \partial_{x_n}\mathbf{F}(\mathbf{x}, \mathbf{x}')\end{bmatrix}\begin{bmatrix}\partial_{x_1}\mathbf{F}^{\mathsf{T}}(\mathbf{x}, \mathbf{x}') \\ \vdots \\ \partial_{x_n}\mathbf{F}^{\mathsf{T}}(\mathbf{x}, \mathbf{x}')\end{bmatrix}. \tag{2.9}$$

Notice that the two sides are merely equal to each other in a formal way. We didn't actually take the partial derivatives since x_i's could be discrete random variables. But conceptually, it is very convenient to think of this step as taking partial derivatives in $\mathbf{x}$.

Following (2.9) and using our notation for partial derivative block matrices introduced in Section 2.2, we can write (2.6) compactly as

$$
\begin{aligned}
&\mathbb{E}\left\|\left(\sum_{i=1}^{n}\left(\sum_{j=1}^{n}\mathbf{A}_{i,j}x'_j\right)\left(\sum_{j=1}^{n}\mathbf{A}_{i,j}x'_j\right)^{\top}\right)^{1/2}\right\|_{4t}^{4t}\\
=&\mathbb{E}\left\|\left(\begin{bmatrix}\partial_{x_1}\mathbf{F}(\mathbf{x},\mathbf{x}') & \dots & \partial_{x_n}\mathbf{F}(\mathbf{x},\mathbf{x}')\end{bmatrix}\begin{bmatrix}\partial_{x_1}\mathbf{F}^{\top}(\mathbf{x},\mathbf{x}')\\ \vdots \\ \partial_{x_n}\mathbf{F}^{\top}(\mathbf{x},\mathbf{x}')\end{bmatrix}\right)^{1/2}\right\|_{4t}^{4t} &(2.10)\\
=&\mathbb{E}\left\|\mathbf{F}_{0,1,0}(\mathbf{x}')\right\|_{4t}^{4t}. &(2.11)
\end{aligned}
$$

Similarly, (2.7) and (2.8) can be written compactly as

$$
\begin{aligned}
&\mathbb{E}\left\|\left(\sum_{i=1}^{n}\left(\sum_{j=1}^{n}\mathbf{A}_{i,j}x'_j\right)^{\top}\left(\sum_{j=1}^{n}\mathbf{A}_{i,j}x'_j\right)\right)^{1/2}\right\|_{4t}^{4t}\\
=&\mathbb{E}\left\|\left(\begin{bmatrix}\partial_{x_1}\mathbf{F}^{\top}(\mathbf{x},\mathbf{x}') & \dots & \partial_{x_n}\mathbf{F}^{\top}(\mathbf{x},\mathbf{x}')\end{bmatrix}\begin{bmatrix}\partial_{x_1}\mathbf{F}(\mathbf{x},\mathbf{x}')\\ \vdots \\ \partial_{x_n}\mathbf{F}(\mathbf{x},\mathbf{x}')\end{bmatrix}\right)^{1/2}\right\|_{4t}^{4t} &(2.12)\\
=&\mathbb{E}\left\|\mathbf{F}_{1,0,0}(\mathbf{x}')\right\|_{4t}^{4t}, &(2.13)
\end{aligned}
$$

and

$$
\begin{aligned}
&\sum_{i=1}^{n} \mathbb{E} \left\| \sum_{j=1}^{n} \mathbf{A}_{i,j} x'_j \right\|_{4t}^{4t} \\
=&\mathbb{E} \operatorname{tr} \begin{bmatrix} \partial_{x_1} \mathbf{F}(\mathbf{x}, \mathbf{x}') \cdot \partial_{x_1} \mathbf{F}^{\top}(\mathbf{x}, \mathbf{x}') & & \\ & \ddots & \\ & & \partial_{x_n} \mathbf{F}(\mathbf{x}, \mathbf{x}') \cdot \partial_{x_n} \mathbf{F}^{\top}(\mathbf{x}, \mathbf{x}') \end{bmatrix}^{2t} &(2.14)\\
=&\mathbb{E} \left\| \mathbf{F}_{0,0,1}(\mathbf{x}') \right\|_{4t}^{4t}. &(2.15)
\end{aligned}
$$

Hence the inequality (2.6-2.8) can be written succinctly as

$$
E \leq 4^{4t}(16t)^{3t} \cdot \left(\mathbb{E} \left\| \mathbf{F}_{1,0,0} \right\|_{4t}^{4t} + \mathbb{E} \left\| \mathbf{F}_{0,1,0} \right\|_{4t}^{4t} \right) + 4^{4t}(8t)^{4t} L^{4t} \cdot \mathbb{E} \left\| \mathbf{F}_{0,0,1} \right\|_{4t}^{4t}.
$$

Since $\mathbf{F}(\mathbf{x}, \mathbf{x}')$ is multilinear and we have already taken the partial derivatives in $\mathbf{x}$, $\mathbf{F}_{1,0,0}$, $\mathbf{F}_{0,1,0}$, and $\mathbf{F}_{0,0,1}$ are sums of centered, independent random matrices in $\mathbf{x}'$; i.e.,

$$
\mathbf{F}_{1,0,0}(\mathbf{x}') = \sum_{j=1}^{n} x'_j \begin{bmatrix} \mathbf{A}_{1,j}^{\top} & \dots & \mathbf{A}_{n,j}^{\top} \end{bmatrix}^{\top},
$$

$$
\mathbf{F}_{0,1,0}(\mathbf{x}') = \sum_{j=1}^{n} x'_j \begin{bmatrix} \mathbf{A}_{1,j} & \dots & \mathbf{A}_{n,j} \end{bmatrix},
$$

and

$$
\mathbf{F}_{0,0,1}(\mathbf{x}') = \sum_{j=1}^{n} x'_j \begin{bmatrix} \mathbf{A}_{1,j} & & \\ & \ddots & \\ & & \mathbf{A}_{n,j} \end{bmatrix}.
$$

Therefore, we can apply matrix Rosenthal inequality (Lemma 2.4.1) again to bound

$\mathbb{E}\left\|\mathbf{F}_{1,0,0}\right\|_{4t}^{4t}$, $\mathbb{E}\left\|\mathbf{F}_{0,1,0}\right\|_{4t}^{4t}$, and $\mathbb{E}\left\|\mathbf{F}_{0,0,1}\right\|_{4t}^{4t}$. Observe that

$$\begin{bmatrix} \partial_{x_1'} \mathbf{F}_{0,1,0} \\ \vdots \\ \partial_{x_n'} \mathbf{F}_{0,1,0} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{A}_{2,1} & \dots & \mathbf{A}_{n,1} \\ \mathbf{A}_{1,2} & \mathbf{0} & \dots & \mathbf{A}_{n,2} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{1,n} & \mathbf{A}_{2,n} & \dots & \mathbf{0} \end{bmatrix} = \mathbf{F}_{1,1,0}$$

and

$$\begin{bmatrix} \partial_{x_1'} \mathbf{F}_{1,0,0}^\top \\ \vdots \\ \partial_{x_n'} \mathbf{F}_{1,0,0}^\top \end{bmatrix}^\top = \begin{bmatrix} \mathbf{0} & \mathbf{A}_{1,2} & \dots & \mathbf{A}_{1,n} \\ \mathbf{A}_{2,1} & \mathbf{0} & \dots & \mathbf{A}_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{n,1} & \mathbf{A}_{n,2} & \dots & \mathbf{0} \end{bmatrix} = \mathbf{F}_{1,1,0},$$

where the equality is due to the permutation symmetric property of $\mathbf{F}(\mathbf{x})$, i.e., $\mathbf{A}_{i,j} = \mathbf{A}_{j,i}$, $\forall (i,j) \in \mathcal{T}_n^2$. After applying two rounds of matrix Rosenthal inequality (Lemma 2.4.1), we have

$$\begin{aligned} E < 2 \cdot 4^{4t}(16t)^{8t} \cdot \Big(&\left\|\mathbf{F}_{2,0,0}\right\|_{4t}^{4t} + \left\|\mathbf{F}_{0,2,0}\right\|_{4t}^{4t} + \left\|\mathbf{F}_{1,1,0}\right\|_{4t}^{4t} + L^{4t}\left\|\mathbf{F}_{1,0,1}\right\|_{4t}^{4t} \\ &+ L^{4t}\left\|\mathbf{F}_{0,1,1}\right\|_{4t}^{4t} + L^{8t}\left\|\mathbf{F}_{0,0,2}\right\|_{4t}^{4t} \Big). \end{aligned}$$

We complete the proof by observing that $\mathbb{E}\mathbf{F}(\mathbf{x}) = \mathbf{0}$ due to multilinearity of $\mathbf{F}(\mathbf{x})$. ■

In general, differentiating the decoupled $\mathbf{F}(\mathbf{x}, \mathbf{x}')$ with respect to $\mathbf{x}$ and $\mathbf{x}'$ can be quite different from taking second order derivatives of $\mathbf{F}(\mathbf{x})$ with respect to $\mathbf{x}$. For example, it is clear that,

$$\partial_{x_1}\partial_{x_2} \begin{bmatrix} 0 & x_1x_2 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} = \partial_{x_2}\partial_{x_1} \begin{bmatrix} 0 & x_1x_2 \\ 0 & 0 \end{bmatrix},$$

since ∂x_1 and ∂x_2 commute. But taking partial derivatives after decoupling gives,

$$\partial_{x_1}\partial_{x_2'}\begin{bmatrix}0 & x_1x_2' \\ 0 & 0\end{bmatrix} = \begin{bmatrix}0 & 1 \\ 0 & 0\end{bmatrix} \neq \begin{bmatrix}0 & 0 \\ 0 & 0\end{bmatrix} = \partial_{x_2}\partial_{x_1'}\begin{bmatrix}0 & x_1x_2' \\ 0 & 0\end{bmatrix},$$

where x_1, x_2, x_1', x_2' are actually four distinct variables. So one can't expect the result of taking partial derivatives under two scenarios are the same. We could directly work with this by recording $\partial_{x_1}\partial_{x_2'}\mathbf{F}(\mathbf{x})$ and $\partial_{x_2}\partial_{x_1'}\mathbf{F}(\mathbf{x})$ separately, but the notation gets cumbersome very quickly. We would prefer to state our results in regular partial derivatives, which is more intuitive. So instead, we write $\mathbf{F}(\mathbf{x})$ in the permutation symmetric way. Continuing with our example, we write

$$\begin{bmatrix}0 & x_1x_2 \\ 0 & 0\end{bmatrix} = \frac{1}{2}\begin{bmatrix}0 & x_1x_2 \\ 0 & 0\end{bmatrix} + \frac{1}{2}\begin{bmatrix}0 & x_2x_1 \\ 0 & 0\end{bmatrix}.$$

This doesn't affect the second order derivatives, but notice now what we get taking partial derivatives after decoupling,

$$\partial_{x_1}\partial_{x_2'}\left(\frac{1}{2}\begin{bmatrix}0 & x_1x_2' \\ 0 & 0\end{bmatrix} + \frac{1}{2}\begin{bmatrix}0 & x_2x_1' \\ 0 & 0\end{bmatrix}\right) = \frac{1}{2}\begin{bmatrix}0 & 1 \\ 0 & 0\end{bmatrix}.$$

and

$$\partial_{x_2}\partial_{x_1'}\left(\frac{1}{2}\begin{bmatrix}0 & x_1x_2' \\ 0 & 0\end{bmatrix} + \frac{1}{2}\begin{bmatrix}0 & x_2x_1' \\ 0 & 0\end{bmatrix}\right) = \frac{1}{2}\begin{bmatrix}0 & 1 \\ 0 & 0\end{bmatrix}$$

Therefore, the permutation symmetric property of $\mathbf{F}(\mathbf{x})$ makes the resulting partial derivative matrices $\{\mathbf{F}_{a,b,c}\}$ under two scenarios differ only by a constant factor. This also applies to $\mathbf{F}(\mathbf{x})$ of higher degree.

Claim 2.4.4. *Let $\mathbf{F}(\mathbf{x})$ be a homogeneous multilinear polynomial random matrix of degree D. Let*

$\mathbf{F}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(D)})$ *be the decoupled* $\mathbf{F}(\mathbf{x})$*, i.e.,*

$$\mathbf{F}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(D)}) = \sum_{\mathbf{j} \in \mathcal{T}_n^D} \mathbf{A}_{\mathbf{j}} \cdot x_{j_1}^{(1)} \cdots x_{j_D}^{(D)},$$

where $\{\mathbf{A}_{\mathbf{k}}\}_{\mathbf{k} \in \mathcal{T}_n^D}$ *is a multi-indexed sequence of deterministic matrices of the same dimension. For some fixed* $a, b, c \in \mathbb{Z}_{\geq 0}$ *and* $k = a + b + c < D$*, let* $\mathbf{F}_{a,b,c}$ *be the block matrix of* k*-th order partial derivatives of* $\mathbf{F}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(D)})$*. Then* $\mathbf{F}_{a,b,c}$ *is a homogeneous multilinear polynomial random matrix of degree* $d = D - k$*. Furthermore,*

$$\mathbb{E} \left\|\mathbf{F}_{a,b,c}\right\|_{4t}^{4t} \leq (16t)^{3t} \cdot \left(\mathbb{E} \left\|\mathbf{F}_{a+1,b,c}\right\|_{4t}^{4t} + \mathbb{E} \left\|\mathbf{F}_{a,b+1,c}\right\|_{4t}^{4t}\right) + (8t)^{4t} \cdot \mathbb{E} \left\|\mathbf{F}_{a,b,c+1}\right\|_{4t}^{4t}.$$

Proof. Without loss of generality, we differentiate $\mathbf{F}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(D)})$ with respect to $\mathbf{x}^{(D)}$ first, then $\mathbf{x}^{(D-1)}$, and so on. It is straightforward to see that after $k = a + b + c$ rounds of differentiation, $\mathbf{F}_{a,b,c}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)})$ is a homogeneous polynomial random matrix of degree $d = D - k$. And

$$\mathbf{F}_{a,b,c}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)}) = \sum_{\mathbf{i} \in \mathcal{T}_n^d} \mathbf{B}_{\mathbf{i}} \cdot x_{i_1}^{(1)} \cdots x_{i_d}^{(d)},$$

where $\{\mathbf{B}_{\mathbf{i}}\}_{\mathbf{i} \in \mathcal{T}_n^d}$ is a multi-indexed sequence of deterministic matrices of the same dimension. More specifically, each $\mathbf{B}_{\mathbf{i}}$ is a block matrix whose blocks consist of $\{\mathbf{A}_{\mathbf{j}}\}_{\mathbf{j} \in \mathcal{T}_n^D}$ such that $j_1 = i_1, \ldots, j_d = i_d$.

Conditioned on $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d-1)}$, $\mathbf{F}_{a,b,c}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(d)})$ is a sum of centered, (conditionally) independent random matrices in $\mathbf{x}^{(d)}$. It follows that

$$
\begin{aligned}
&\mathbb{E}\left(\left\|\mathbf{F}_{a,b,c}(\mathbf{x}^{(1)},\dots,\mathbf{x}^{(d)})\right\|_{4t}^{4t} \,\middle|\, \mathbf{x}^{(1)},\dots \mathbf{x}^{(d-1)}\right) \\
&= \mathbb{E}\left(\mathbb{E}_{\mathbf{x}^{(d)}}\left(\left\|\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_d}^{(d)}\right\|_{4t}^{4t} \,\middle|\, \mathbf{x}^{(1)},\dots \mathbf{x}^{(d-1)}\right)\right) \\
&\le (16t)^{3t}\cdot\mathbb{E}\left\|\left(\sum_{i_d=1}^{n}\left(\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_{d-1}}^{(d-1)}\right)\left(\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_{d-1}}^{(d-1)}\right)^{\mathsf{T}}\right)^{1/2}\right\|_{4t}^{4t} \\
&\quad + (16t)^{3t}\cdot\mathbb{E}\left\|\left(\sum_{i_d=1}^{n}\left(\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_{d-1}}^{(d-1)}\right)^{\mathsf{T}}\left(\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_{d-1}}^{(d-1)}\right)\right)^{1/2}\right\|_{4t}^{4t} \\
&\quad + (8t)^{4t}L^{4t}\cdot\sum_{i_d=1}^{n}\mathbb{E}\left\|\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_{d-1}}^{(d-1)}\right\|_{4t}^{4t},
\end{aligned}
$$

where the inequality is due to matrix Rosenthal inequality (Lemma 2.4.1). Similar to observations (2.11), (2.13), and (2.15) in the proof of Theorem 2.4.2, we have

$$
\begin{aligned}
&\mathbb{E}\left\|\left(\sum_{i_d=1}^{n}\left(\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_{d-1}}^{(d-1)}\right)\left(\sum_{\mathbf{i}\in\mathcal{T}_n^d}\mathbf{B}_{\mathbf{i}}\cdot x_{i_1}^{(1)}\cdots x_{i_{d-1}}^{(d-1)}\right)^{\mathsf{T}}\right)^{1/2}\right\|_{4t}^{4t} \\
&= \mathbb{E}\left\|\left(\begin{bmatrix}\partial_{x_1^{(d)}}\mathbf{F}_{a,b,c} & \dots & \partial_{x_n^{(d)}}\mathbf{F}_{a,b,c}\end{bmatrix}\begin{bmatrix}\partial_{x_1^{(d)}}\mathbf{F}_{a,b,c}^{\mathsf{T}} \\ \vdots \\ \partial_{x_n^{(d)}}\mathbf{F}_{a,b,c}^{\mathsf{T}}\end{bmatrix}\right)^{1/2}\right\|_{4t}^{4t} = \mathbb{E}\left\|\mathbf{F}_{a,b+1,c}\right\|_{4t}^{4t},
\end{aligned}
$$

$$
\mathbb{E}\left\| \left(\sum_{i_d=1}^{n} \left(\sum_{\mathbf{i}\in\mathcal{T}_n^d} \mathbf{B}_{\mathbf{i}} \cdot x_{i_1}^{(1)} \cdots x_{i_{d-1}}^{(d-1)} \right)^{\top} \left(\sum_{\mathbf{i}\in\mathcal{T}_n^d} \mathbf{B}_{\mathbf{i}} \cdot x_{i_1}^{(1)} \cdots x_{i_{d-1}}^{(d-1)} \right) \right)^{1/2} \right\|_{4t}^{4t}
$$

$$
= \mathbb{E}\left\| \left(\begin{bmatrix} \partial_{x_1^{(d)}} \mathbf{F}_{a,b,c}^{\top} & \cdots & \partial_{x_n^{(d)}} \mathbf{F}_{a,b,c}^{\top} \end{bmatrix} \begin{bmatrix} \partial_{x_1^{(d)}} \mathbf{F}_{a,b,c} \\ \vdots \\ \partial_{x_n^{(d)}} \mathbf{F}_{a,b,c} \end{bmatrix} \right)^{1/2} \right\|_{4t}^{4t} = \mathbb{E}\left\| \mathbf{F}_{a+1,b,c} \right\|_{4t}^{4t},
$$

and

$$
\sum_{i_d=1}^{n} \mathbb{E}\left\| \sum_{\mathbf{i}\in\mathcal{T}_n^d} \mathbf{B}_{\mathbf{i}} \cdot x_{i_1}^{(1)} \cdots x_{i_{d-1}}^{(d-1)} \right\|_{4t}^{4t}
$$

$$
= \mathbb{E}\,\mathrm{tr} \begin{bmatrix} \partial_{x_1^{(d)}} \mathbf{F}_{a,b,c} \cdot \partial_{x_1^{(d)}} \mathbf{F}_{a,b,c}^{\top} & & \\ & \ddots & \\ & & \partial_{x_n^{(d)}} \mathbf{F}_{a,b,c} \cdot \partial_{x_n^{(d)}} \mathbf{F}_{a,b,c}^{\top} \end{bmatrix}^{2t} = \mathbb{E}\left\| \mathbf{F}_{a,b,c+1} \right\|_{4t}^{4t}.
$$

Hence, the result follows. ■

We present a moment bound for permutation symmetric, homogeneous multilinear polynomial random matrices of degree d.

Theorem 2.4.5 (Homogeneous Multilinear Recursion)**.** *Let $\mathbf{x} = \{x_i\}_{i=1}^n$ be a sequences of i.i.d random variables with $\mathbb{E}x_i = 0$, $\mathbb{E}x_i^2 = 1$ and $|x_i| \leq L$ for all $1 \leq i \leq n$. Let $\{\mathbf{A}_{\mathbf{i}}\}_{\mathbf{i}\in\mathcal{T}_n^d}$ be a multi-indexed sequence of deterministic matrices of the same dimension. Define a permutation symmetric, homogeneous multilinear polynomial random matrix of degree d as*

$$
\mathbf{F}(\mathbf{x}) = \sum_{\mathbf{i}\in\mathcal{T}_n^d} \left(\mathbf{A}_{\mathbf{i}} \cdot \prod_{j\in\{i_1,\ldots,i_d\}} x_j \right).
$$

Let $a, b, c \in \mathbb{Z}_{\geq 0}$ and $d = a + b + c$. Then for $2 \leq t \leq \infty$,

$$\mathbb{E} \left\|\mathbf{F}(\mathbf{x}) - \mathbb{E}\mathbf{F}(\mathbf{x})\right\|_{4t}^{4t} \leq \sum_{\substack{a,b,c: \\ a+b+c=d}} (48dt)^{4dt} \cdot L^{4ct} \left\|\mathbf{F}_{a,b,c}\right\|_{4t}^{4t}.$$

Proof. Let $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(d)}$ be d independent copies of $\mathbf{x}$. We decouple $\mathbf{F}(\mathbf{x})$ by Lemma 2.3.2,

$$E := \mathbb{E} \left\|\mathbf{F}((\mathbf{x})\right\|_{4t}^{4t} \leq d^{4dt} \cdot \mathbb{E} \left\|\mathbf{F}(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(d)})\right\|_{4t}^{4t}.$$

We take the partial derivative with respect to $\mathbf{x}^{(d)}$ by applying Lemma 2.4.1 to get,

$$E \leq d^{4dt}(16t)^{3t} \cdot \left(\mathbb{E} \left\|\mathbf{F}_{1,0,0}\right\|_{4t}^{4t} + \mathbb{E} \left\|\mathbf{F}_{0,1,0}\right\|_{4t}^{4t}\right) + d^{4dt}(8t)^{4t} L^{4t} \cdot \mathbb{E} \left\|\mathbf{F}_{0,0,1}\right\|_{4t}^{4t}.$$

Note that $\mathbf{F}_{1,0,0}$, $\mathbf{F}_{0,1,0}$, and $\mathbf{F}_{0,0,1}$ are functions in variables $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(d-1)}$. We take the partial derivative with respect to the rest of the variables until $\mathbf{F}_{a,b,c}$'s become deterministic matrices. Apply Claim 2.4.4 recursively and use the permutation symmetric property of $\mathbf{F}(\mathbf{x})$, we have

$$E \leq \sum_{\substack{a,b,c: \\ a+b+c=d}} (16dt)^{4dt} \cdot L^{4ct} \frac{d!}{a!b!c!} \left\|\mathbf{F}_{a,b,c}\right\|_{4t}^{4t}.$$

Since $\frac{d!}{a!b!c!} \leq \frac{d!}{(d/3)!(d/3)!(d/3)!}$ and $\frac{d!}{(d/3)!(d/3)!(d/3)!} \sim \frac{3\sqrt{3}}{2\pi d} \cdot 3^d$,

$$E \leq \sum_{\substack{a,b,c: \\ a+b+c=d}} (48dt)^{4dt} \cdot L^{4ct} \left\|\mathbf{F}_{a,b,c}\right\|_{4t}^{4t}.$$

∎

As a corollary, we derive a moment bound for general multilinear polynomial random matrices of degree D. The polynomials are split into homogeneous parts and each parts are bounded separately. Let $\mathbf{F}^{=d}(\mathbf{x})$ denote the degree-d homogeneous part of $\mathbf{F}(\mathbf{x})$.

We will need the following inequality.

Lemma 2.4.6 (Trace Inequality, see [SA13] Theorem 3.1)**.** *Let* $\mathbf{A}_i \in \mathbb{R}^{n\times n}$ *for* $i = 1, 2, \ldots, m$ *and* $r \geq 1$. *Then*

$$\operatorname{tr}\left|\sum_{i=1}^{m} \mathbf{A}_i\right|^r \;\leq\; m^{r-1} \operatorname{tr}\left(\sum_{i=1}^{m} |\mathbf{A}_i|^r\right),$$

where $|\mathbf{A}_i| = (\mathbf{A}_i{}^\top \mathbf{A}_i)^{1/2}$.

Corollary 2.4.7 (Multilinear Recursion)**.** *Let* $\mathbf{x} = \{x_i\}_{i=1}^n$ *be a sequences of i.i.d random variables with* $\mathbb{E}x_i = 0$, $\mathbb{E}x_i^2 = 1$ *and* $|x_i| \leq L$ *for all* $1 \leq i \leq n$. *Let* $\{\mathbf{A}_\mathbf{i}\}_{\mathbf{i}\in\mathcal{T}_n^d}$ *be multi-indexed sequences of deterministic matrices of the same dimension. Define a degree-D multilinear polynomial random matrix as*

$$\mathbf{F}(\mathbf{x}) = \sum_{d=1}^{D} \sum_{\mathbf{i}\in\mathcal{T}_n^d} \left(\mathbf{A}_\mathbf{i} \cdot \prod_{j\in\{i_1,\ldots,i_d\}} x_j\right).$$

Suppose $\mathbf{F}^{=d}(\mathbf{x})$ *is permutation symmetric for* $1 \leq d \leq D$. *Let* $a, b, c \in \mathbb{Z}_{\geq 0}$ *and* $d = a + b + c$. *Then for* $2 \leq t \leq \infty$,

$$\mathbb{E}\,\|\mathbf{F}(\mathbf{x}) - \mathbb{E}\mathbf{F}(\mathbf{x})\|_{4t}^{4t} \;\leq\; \sum_{d=1}^{D} D^{4t}(48dt)^{4dt} \left(\sum_{\substack{a,b,c:\\ a+b+c=d}} L^{4ct} \left\|\mathbf{F}_{a,b,c}^{=d}\right\|_{4t}^{4t}\right).$$

Proof. We rewrite $\mathbf{F}(\mathbf{x})$ into a formal sum of its homogeneous parts $\mathbf{F}(\mathbf{x}) = \sum_{d=1}^{D} \mathbf{F}^{=d}(\mathbf{x})$. By the trace inequality (Lemma 2.4.6), we have

$$\mathbb{E}\,\|\mathbf{F}(\mathbf{x}) - \mathbb{E}\mathbf{F}(\mathbf{x})\|_{4t}^{4t} \leq D^{4t} \sum_{d=1}^{D} \mathbb{E}\left\|\mathbf{F}^{=d}(\mathbf{x}) - \mathbb{E}\mathbf{F}^{=d}(\mathbf{x})\right\|_{4t}^{4t}.$$

By Theorem 2.4.5,

$$\mathbb{E}\,\|\mathbf{F}(\mathbf{x})-\mathbb{E}\mathbf{F}(\mathbf{x})\|_{4t}^{4t} \le D^{4t}\sum_{d=1}^{D}(48dt)^{4dt}\left(\sum_{\substack{a,b,c:\\ a+b+c=d}} L^{4ct}\left\|\mathbf{F}_{a,b,c}^{=d}\right\|_{4t}^{4t}\right).$$

■

2.5 Gaussian Polynomial Random Matrices

In this section, we prove a moment bound for polynomial random matrices with Gaussian random variables. While the decoupling technique could work for Gaussian polynomial random matrices as well, we base our recursion framework on the following bound for its simplicity.

Lemma 2.5.1 (Polynomial Moments, see [HT20] Theorem 7.1)**.** *Let $\mathbf{H}(\mathbf{x}) : \Omega \to \mathbb{H}^d$ be a function with $\mathbf{x} \sim \mathcal{N}(0, \mathbb{I}_n)$. For $t = 1$ and $t \ge 1.5$,*

$$\mathbb{E}\,\|\mathbf{H}(\mathbf{x})-\mathbb{E}\mathbf{H}(\mathbf{x})\|_{2t}^{2t} \;\le\; (\sqrt{2}t)^{2t}\cdot \mathbb{E}\,\mathrm{tr}\left(\sum_{i=1}^{n}(\partial_i\mathbf{H}(\mathbf{x}))^2\right)^t. \tag{2.16}$$

We derive the non-Hermitian version of Lemma 2.5.1.

Lemma 2.5.2 (Non-Hermitian Polynomial Moments)**.** *Let $\mathbf{F}(\mathbf{x}) : \Omega \to \mathbb{R}^{d_1\times d_2}$ be a function with $\mathbf{x} \sim \mathcal{N}(0, \mathbb{I}_n)$. For $t = 1$ and $t \ge 1.5$,*

$$\begin{aligned}\mathbb{E}\,\|\mathbf{F}(\mathbf{x})-\mathbb{E}\mathbf{F}(\mathbf{x})\|_{2t}^{2t} \;\le\; (\sqrt{2}t)^{2t}\cdot\Bigg(&\mathbb{E}\left\|\left(\sum_{i=1}^{n}\partial_i\mathbf{F}(\mathbf{x})\,\partial_i\mathbf{F}^{\mathsf{T}}(\mathbf{x})\right)^{1/2}\right\|_{2t}^{2t}\\ &+\mathbb{E}\left\|\left(\sum_{i=1}^{n}\partial_i\mathbf{F}^{\mathsf{T}}(\mathbf{x})\,\partial_i\mathbf{F}(\mathbf{x})\right)^{1/2}\right\|_{2t}^{2t}\Bigg).\end{aligned}$$

Proof. Apply Hermitian dilation. The left hand side of (2.16) becomes

$$\begin{aligned}\mathbb{E}\left\|\begin{bmatrix}\mathbf{0} & \mathbf{F}-\mathbb{E}\mathbf{F}\\ (\mathbf{F}-\mathbb{E}\mathbf{F})^\top & \mathbf{0}\end{bmatrix}\right\|_{2t}^{2t} &= \mathbb{E}\,\mathrm{tr}\begin{bmatrix}(\mathbf{F}-\mathbb{E}\mathbf{F})(\mathbf{F}-\mathbb{E}\mathbf{F})^\top & \mathbf{0}\\ \mathbf{0} & (\mathbf{F}-\mathbb{E}\mathbf{F})^\top(\mathbf{F}-\mathbb{E}\mathbf{F})\end{bmatrix}^t\\ &= 2\mathbb{E}\left\|\mathbf{F}-\mathbb{E}\mathbf{F}\right\|_{2t}^{2t},\end{aligned}$$

and the right hand side of (2.16) becomes

$$\begin{aligned}\mathbb{E}\,\mathrm{tr}\left(\sum_{i=1}^n\begin{bmatrix}\mathbf{0} & \partial_i\mathbf{F}\\ \partial_i\mathbf{F}^\top & \mathbf{0}\end{bmatrix}^2\right)^t &= \mathbb{E}\,\mathrm{tr}\begin{bmatrix}\sum_{i=1}^n\partial_i\mathbf{F}\,\partial_i\mathbf{F}^\top & \mathbf{0}\\ \mathbf{0} & \sum_{i=1}^n\partial_i\mathbf{F}^\top\,\partial_i\mathbf{F}\end{bmatrix}^t\\ &= \mathbb{E}\left\|\left(\sum_{i=1}^n\partial_i\mathbf{F}\,\partial_i\mathbf{F}^\top\right)^{1/2}\right\|_{2t}^{2t} + \mathbb{E}\left\|\left(\sum_{i=1}^n\partial_i\mathbf{F}^\top\,\partial_i\mathbf{F}\right)^{1/2}\right\|_{2t}^{2t}.\end{aligned}$$

Combine both sides and the result follows. ■

Theorem 2.5.3 (Gaussian Recursion). *Let $\mathbf{x}\sim\mathcal{N}(0,\mathbb{I}_n)$ and $\{\mathbf{A_i}\}_{\mathbf{i}\in[n]^d}$ be multi-indexed sequences of deterministic matrices of the same dimension. Define a degree-d homogeneous Gaussian polynomial random matrix as*

$$\mathbf{P}(\mathbf{x}) = \sum_{\mathbf{i}\in[n]^d}\left(\mathbf{A_i}\cdot\prod_{j\in\{i_1,\dots,i_d\}}x_j\right).$$

Let $a,b\in\mathbb{Z}_{\geq 0}$. Then for $2\leq t\leq\infty$,

$$\mathbb{E}\left\|\mathbf{P}(\mathbf{x})-\mathbb{E}\mathbf{P}(\mathbf{x})\right\|_{2t}^{2t} \;\leq\; (2^d\sqrt{2}t)^{2t}\left(\sum_{1\leq a+b\leq d}\left\|\mathbb{E}\mathbf{P}_{a,b}\right\|_{2t}^{2t}\right).$$

Proof. Start by applying Claim 2.5.4 with $k=0$. Apply Claim 2.5.4 recursively until $k=d-1$ to get the desired bound. ■

Claim 2.5.4. *Let* $\mathbf{x} \sim \mathcal{N}(0, \mathbb{I}_n)$ *and* $\{\mathbf{A}_{\mathbf{i}}\}_{\mathbf{i}\in[n]^d}$ *be multi-indexed sequences of deterministic matrices of the same dimension. Define a degree-d homogeneous Gaussian polynomial random matrix as*

$$\mathbf{P}(\mathbf{x}) = \sum_{\mathbf{i}\in[n]^d} \left(\mathbf{A}_{\mathbf{i}} \cdot \prod_{j\in\{i_1,\ldots,i_d\}} x_j \right).$$

Let $a, b \in \mathbb{Z}_{\geq 0}$ *and* $a + b = k < d$. *Let* $\mathbf{P}_{a,b}(\mathbf{x})$ *be the k-th partial derivative block matrix associated to* $\mathbf{P}(\mathbf{x})$. *Then for* $2 \leq t \leq \infty$,

$$\begin{aligned} \mathbb{E} \left\|\mathbf{P}_{a,b} - \mathbb{E}\mathbf{P}_{a,b}\right\|_{2t}^{2t} \;\leq\; & (2\sqrt{2}t)^{2t} \cdot \left(\mathbb{E} \left\|\mathbf{P}_{a+1,b} - \mathbb{E}\mathbf{P}_{a+1,b}\right\|_{2t}^{2t} + \left\|\mathbb{E}\mathbf{P}_{a+1,b}\right\|_{2t}^{2t}\right) \\ & + (2\sqrt{2}t)^{2t} \cdot \left(\mathbb{E} \left\|\mathbf{P}_{a,b+1} - \mathbb{E}\mathbf{P}_{a,b+1}\right\|_{2t}^{2t} + \left\|\mathbb{E}\mathbf{P}_{a,b+1}\right\|_{2t}^{2t}\right). \end{aligned}$$

Proof. By the non-hermitian polynomial moment (Lemma 2.5.2), we have

$$\begin{aligned} E :&= \mathbb{E} \left\|\mathbf{P}_{a,b} - \mathbb{E}\mathbf{P}_{a,b}\right\|_{2t}^{2t} \\ &\leq (\sqrt{2}t)^{2t} \cdot \left(\mathbb{E} \left\| \left(\sum_{i=1}^{n} \partial_i \mathbf{P}_{a,b}\, \partial_i \mathbf{P}_{a,b}^{\mathsf{T}} \right)^{1/2} \right\|_{2t}^{2t} + \mathbb{E} \left\| \left(\sum_{i=1}^{n} \partial_i \mathbf{P}_{a,b}^{\mathsf{T}}\, \partial_i \mathbf{P}_{a,b} \right)^{1/2} \right\|_{2t}^{2t} \right) \\ &= (\sqrt{2}t)^{2t} \cdot \left(\mathbb{E} \left\| \begin{bmatrix} \partial_1 \mathbf{P}_{a,b} & \ldots & \partial_n \mathbf{P}_{a,b} \end{bmatrix} \right\|_{2t}^{2t} + \mathbb{E} \left\| \begin{bmatrix} \partial_1 \mathbf{P}_{a,b} \\ \vdots \\ \partial_n \mathbf{P}_{a,b} \end{bmatrix} \right\|_{2t}^{2t} \right) \\ &= (\sqrt{2}t)^{2t} \cdot \left(\mathbb{E} \left\|\mathbf{P}_{a+1,b}\right\|_{2t}^{2t} + \mathbb{E} \left\|\mathbf{P}_{a,b+1}\right\|_{2t}^{2t} \right), \end{aligned}$$

where the last equality uses our notation for partial derivative block matrices introduced in Section 2.2. Notice that $\mathbf{P}_{a+1,b}$ and $\mathbf{P}_{a,b+1}$ are homogeneous Gaussian polynomial random matrix of degree $k-1$. Since $\mathbf{P}_{a+1,b}$ and $\mathbf{P}_{a,b+1}$ are not necessarily centered,

$$\begin{aligned} E &\leq (\sqrt{2}t)^{2t} \cdot \Big(\mathbb{E} \left\| \mathbf{P}_{a+1,b} - \mathbb{E}\mathbf{P}_{a+1,b} + \mathbb{E}\mathbf{P}_{a+1,b} \right\|_{2t}^{2t} \\ &\qquad\qquad + \mathbb{E} \left\| \mathbf{P}_{a,b+1} - \mathbb{E}\mathbf{P}_{a,b+1} + \mathbb{E}\mathbf{P}_{a,b+1} \right\|_{2t}^{2t} \Big) \\ &\leq (2\sqrt{2}t)^{2t} \cdot \Big(\mathbb{E} \left\| \mathbf{P}_{a+1,b} - \mathbb{E}\mathbf{P}_{a+1,b} \right\|_{2t}^{2t} + \left\| \mathbb{E}\mathbf{P}_{a+1,b} \right\|_{2t}^{2t} \\ &\qquad\qquad + \mathbb{E} \left\| \mathbf{P}_{a,b+1} - \mathbb{E}\mathbf{P}_{a,b+1} \right\|_{2t}^{2t} + \left\| \mathbb{E}\mathbf{P}_{a,b+1} \right\|_{2t}^{2t} \Big), \end{aligned}$$

where the last inequality is due to the trace inequality (Lemma 2.4.6). ■

2.6 Application I: Graph Matrices

2.6.1 Background

Denote the Erdős-Rényi random graph on n vertices as $\mathcal{G}_{n,p}$, where each edge is present with probability p independent of all other edges. When $\mathcal{G}_{n,p}$ is viewed as a probability space, it is equal to $(\Omega, \mathcal{F}, \mu)$ where Ω is the sample space of all possible graphs on n vertices. Any $G \in \Omega$ is a random vector representing all edges. Each coordinate in G, denoted by G_{ij}, is an independent random variable representing a single edge. We can construct a random graph by drawing a copy of $G \sim \mu$. To adopt the convention in p-biased Fourier analysis, we normalize G_{ij} such that $\mathbb{E}G_{ij} = 0$ and $\mathbb{E}G_{ij}^2 = 1$, which leads to the sample space $\Omega = \left\{ -\sqrt{\frac{1-p}{p}}, \sqrt{\frac{p}{1-p}} \right\}^{\binom{n}{2}}$.

Definition 2.6.1 (Shape, see e.g. [RT23] Definition 4.2)**.** *A shape is a tuple $\tau = (V(\tau), E(\tau), U_\tau, V_\tau)$ where $(V(\tau), E(\tau))$ is a graph and U_τ, V_τ are ordered subsets of the vertices.*

Definition 2.6.2 (Graph matrix, see e.g. [RT23] Definition 4.4)**.** *Given a shape τ, the associated* *graph matrix* $\mathbf{M} : \Omega \to \mathbb{R}^{d_1 \times d_2}$ *is a matrix-valued function such that*

$$\mathbf{M}[I, J] = \sum_{\substack{\varphi \in \mathcal{T}_n^k: \\ \varphi(U_\tau) = I, \varphi(V_\tau) = J}} \prod_{(u,v) \in E(\tau)} G_{\varphi(u), \varphi(v)}.$$

In other words, $\mathbf{M}$ *maps an input graph* $G \in \Omega$ *to a* $\frac{n!}{(n-|I|)!} \times \frac{n!}{(n-|J|)!}$ *matrix whose rows and columns are indexed by* I *and* J *respectively.*

Definition 2.6.3 (Vertex Separator, see e.g. [RT23] Definition 4.8)**.** *For any shape* τ*, a vertex separator is a subset of vertices* $S \subseteq V(\tau)$ *such that there is no path from* U_τ *to* V_τ *in* $\tau \setminus S$*, which is the shape obtained by deleting* S *and all edges adjacent to* S*. We write* S_τ *for a vertex separator of the minimum size.*

2.6.2 Sparse graph matrices

Claim 2.6.4. *For any shape* τ*, let* $\{E_1, E_2\}$ *be an arbitrary cover of* $E(\tau)$*. Let* $S = E_1 \cap E_2$ *be the vertex set of the overlapping edges between* E_1 *and* E_2*. If* S *contains* $E_1 \cap V_\tau$ *and* $E_2 \cap U_\tau$*, then* S *is a vertex separator.*

Proof. Since $S = E_1 \cap E_2$, there is no path from $E_1 \backslash S$ to $E_2 \backslash S$. Additionally, if S contains $E_1 \cap V_\tau$, then $E_1 \backslash S$ is not adjacent to V_τ and $E_1 \backslash S$ contains $U_\tau \backslash S$. If S also contains $E_2 \cap U_\tau$, then $E_2 \backslash S$ is not adjacent to U_τ and $E_2 \backslash S$ contains $V_\tau \backslash S$. Since there is no path from $U_\tau \backslash S$ to $V_\tau \backslash S$, S is a vertex separator. ∎

Theorem 2.6.5. *For any shape* τ *and* $2 \leq t \leq \infty$*,*

$$\mathbb{E} \|\mathbf{M}\|_{4t}^{4t} \leq \left((48t|V(\tau)|)^{4t|V(\tau)|} (C|E(\tau)|)^{|E(\tau)|} n^{|V(\tau)|} \right) \left(\sqrt{\frac{1-p}{p}} \right)^{4t|E(S_\tau)|} n^{2t(|V(\tau)| - |S_\tau|)}$$

where C *is an absolute constant and* $E(S_\tau)$ *are all edges adjacent to* S_τ*.*

Proof. First, let $|V(\tau)| = k$ and we write $\mathbf{M}$ as a polynomial whose coefficients are matrices,

$$\mathbf{M} = \sum_{\varphi \in \mathcal{T}_n^k} \mathbf{B}_{\varphi(E)} \prod_{(i,j)\in E(\tau)} G_{\varphi(i),\varphi(j)},$$

where the $[I, J]$-entry of $\mathbf{B}_{\varphi(E)}$ is

$$\mathbf{B}_{\varphi(E)}[I, J] = \begin{cases} 1 & \text{if } \varphi(U_\tau) = I \ \text{ and } \ \varphi(V_\tau) = J \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\mathbf{B}_{\varphi(E)}$'s have the same dimension as $\mathbf{M}$ and the rows and columns of $\mathbf{B}_{\varphi(E)}$'s are indexed in the same way as $\mathbf{M}$. Since U_τ and V_τ are ordered and each φ assigns one set of value to vertices in U_τ and V_τ, there is only one nonzero entry in each $\mathbf{B}_{\varphi(E)}$. But there might be multiple φ's for which $\mathbf{B}_{\varphi(E)}$'s are identical due to free vertices.

Next, let $|E(\tau)| = d$ and we rewrite $\mathbf{M}$ in a permutation symmetric way,

$$\begin{aligned} \mathbf{M} &= \sum_{\varphi \in \mathcal{T}_n^k} \left(\frac{1}{d!} \sum_{\sigma \in \mathfrak{S}_d} \mathbf{B}_{\sigma(\varphi(E))} G_{\varphi(i_{\sigma(1)}),\varphi(j_{\sigma(1)})} \ \cdots \ G_{\varphi(i_{\sigma(d)}),\varphi(j_{\sigma(d)})} \right) \\ &= \frac{1}{d!} \sum_{\varphi \in \mathcal{T}_n^k} \mathbf{A}_{\varphi(E)} \cdot G_{\varphi(i_1),\varphi(j_1)} \ \cdots \ G_{\varphi(i_d),\varphi(j_d)}, \end{aligned}$$

where $\mathbf{A}_{\varphi(E)} = \sum_{\sigma \in \mathfrak{S}_d} \mathbf{B}_{\sigma(\varphi(E))}$. Notice that the indices of the random variables of $\mathbf{M}$ have the same structure as in Lemma 2.3.3, so we can decouple $\mathbf{M}$ using Lemma 2.3.3,

$$F := \mathbb{E}\,\|\mathbf{M}\|_{4t}^{4t} \ \le \ |V(\tau)|^{4t|V(\tau)|} \left(\frac{1}{d!}\right)^{4t} \cdot \mathbb{E} \left\| \sum_{\varphi \in \mathcal{T}_n^k} \mathbf{A}_{\varphi(E)} \cdot G^{(1)}_{\varphi(i_1),\varphi(j_1)} \cdots G^{(d)}_{\varphi(i_d),\varphi(j_d)} \right\|_{4t}^{4t},$$

where $\{G^{(1)}, \ldots, G^{(d)}\}$ denote d independent copies of G. Since $\mathbf{M}$ is a multilinear polynomial random matrix, $\mathbb{E}\mathbf{M} = \mathbf{0}$. An application of Theorem 2.4.5 yields,

$$F \leq \sum_{\substack{a,b,c: \\ a+b+c=|E(\tau)|}} (48t|V(\tau)|)^{4t|V(\tau)|} \left(\frac{1}{|E(\tau)|!}\right)^{4t} \left(\sqrt{\frac{1-p}{p}}\right)^{4ct} \|\mathbf{M}_{a,b,c}\|_{4t}^{4t},$$

where $\{\mathbf{M}_{a,b,c}\}$ are partial derivative block matrices associated to $\mathbf{M}$.

Fix a set of a, b, c, $\mathbf{M}_{a,b,c}$ is a block matrix whose blocks are comprised of $\{\mathbf{A}_{\varphi(E)}\}_{\varphi \in \mathcal{T}_n^k}$. More specifically, the $[I, J]$-th block of $\mathbf{M}_{a,b,c}$ is

$$\mathbf{M}_{a,b,c}[I,J] = \partial_{G_{I_1}} \cdots \partial_{G_{I_{a+c}}} \partial_{G_{J_1}} \cdots \partial_{G_{J_b}} \mathbf{M} = \mathbf{A}_{I \cup J} = \sum_{\sigma \in \mathfrak{S}_d} \mathbf{B}_{\sigma(I \cup J)}. \tag{2.17}$$

Algebraically, the last equality is due to the commutativity of partial derivative operators. Combinatorially, if we fix an arbitrary ordering on the edge set $E(\tau)$, the last expression is summing up all permutations on d edges in $E(\tau)$. Now let's delve into the combinatorics. Let $E(\tau) = E_1 \cup E_2$, where E_1 contains $a + c$ edges and E_2 contains $b + c$ edges with c number of overlapping edges. Let $S = E_1 \cap E_2$ be the set of vertices shared by E_1 and E_2. By permuting the edges in $E(\tau)$, $\{\sigma(E_1), \sigma(E_2)\}$ are all possible covers of $E(\tau)$ by two overlapping sets of ordered edges. The row and column blocks of $\mathbf{M}_{a,b,c}$ are indexed by $\varphi(E_1)$ and $\varphi(E_2)$ respectively.

Let's write

$$\mathbf{M}_{a,b,c} = \sum_{\sigma \in \mathfrak{S}_d} \mathbf{M}_{a,b,c,\sigma},$$

where $\mathbf{M}_{a,b,c,\sigma}[I,J] = \mathbf{B}_{\sigma(I \cup J)}$. By the trace inequality (Lemma 2.4.6),

$$\|\mathbf{M}_{a,b,c}\|_{4t}^{4t} \leq (d!)^{4t} \sum_{\sigma \in \mathfrak{S}_d} \|\mathbf{M}_{a,b,c,\sigma}\|_{4t}^{4t}.$$

It follows that

$$F \le \sum_{\substack{a,b,c:\\ a+b+c=|E(\tau)|}} \sum_{\sigma\in\mathfrak{S}_d} (48t|V(\tau)|)^{4t|V(\tau)|} \left(\sqrt{\frac{1-p}{p}}\right)^{4ct} \left\|\mathbf{M}_{a,b,c,\sigma}\right\|_{4t}^{4t}. \tag{2.18}$$

For any two distinct σ_1 and $\sigma_2 \in \mathfrak{S}_d$, $\left\|\mathbf{M}_{a,b,c,\sigma_1}\right\|_{4t}^{4t} = \left\|\mathbf{M}_{a,b,c,\sigma_2}\right\|_{4t}^{4t}$ since $\mathbf{M}_{a,b,c,\sigma_1}$ and $\mathbf{M}_{a,b,c,\sigma_2}$ are identical after permuting rows and columns. So we will focus on bounding $\left\|\mathbf{M}_{a,b,c,\sigma_0}\right\|_{4t}^{4t}$ where σ_0 is the identity permutation. For identity permutation, we have $\mathbf{M}_{a,b,c,\sigma_0}[I,J] = \mathbf{B}_{I\cup J}$ by (2.17). Denote $\mathbf{M}_{a,b,c,\sigma_0}$ by $\mathbf{F}$,

$$\begin{aligned}
E &:= \left\|\mathbf{M}_{a,b,c,\sigma_0}\right\|_{4t}^{4t} = \|\mathbf{F}\|_{4t}^{4t} = \operatorname{tr}\left(\mathbf{F}^\top\mathbf{F}\right)^{2t} \\
&= \operatorname{tr} \sum_{I_1,\ldots,I_{2t}\in[n]^{2(a+c)}} \sum_{J_1,\ldots,J_{2t}\in[n]^{2(b+c)}} \mathbf{F}^\top[I_1,J_1]\,\mathbf{F}[I_1,J_2]\,\mathbf{F}^\top[I_2,J_2]\,\mathbf{F}[I_2,J_3]\cdots\mathbf{F}^\top[I_{2t},J_{2t}]\,\mathbf{F}[I_{2t},J_1] \\
&= \operatorname{tr} \sum_{I_1,\ldots,I_{2t}\in[n]^{2(a+c)}} \sum_{J_1,\ldots,J_{2t}\in[n]^{2(b+c)}} \mathbf{B}^\top_{I_1\cup J_1}\,\mathbf{B}_{I_1\cup J_2}\,\mathbf{B}^\top_{I_2\cup J_2}\,\mathbf{B}_{I_2\cup J_3}\cdots\mathbf{B}^\top_{I_{2t}\cup J_{2t}}\,\mathbf{B}_{I_{2t}\cup J_1}
\end{aligned} \tag{2.19}$$

We take a look at for which $I_1,\ldots I_{2t}$ and $J_1,\ldots,J_{2t}$ that the summands in (2.19) are nonzero and the structure of these nonzero summands. First of all, $\mathbf{B}_{I_i\cup J_i} \neq \mathbf{0}$ implies that $I_i\cup J_i = \varphi(E_1\cup E_2)$, for some $\varphi\in\mathcal{T}_n^k$ and $I_i(S) = J_i(S)$ for $1\le i\le 2t$. Similarly, $\mathbf{B}_{I_i\cup J_{i+1}} \neq \mathbf{0}$ implies that $I_i\cup J_{i+1} = \varphi'(E_1\cup E_2)$, for some $\varphi'\in\mathcal{T}_n^k$ and $I_i(S) = J_{i+1}(S)$ for $1\le i\le 2t$ (the additions in the subscripts are in mod $2t$, *i.e.*, $J_{2t+1} = J_1$). Secondly, for each $\mathbf{B}_{I_i\cup J_i}\neq\mathbf{0}$, there is only one nonzero entry, namely $\mathbf{B}_{I_i\cup J_i}[I_i\cup J_i(U_\tau), I_i\cup J_i(V_\tau)] = 1$. So $\mathbf{B}^\top_{I_i\cup J_i}\mathbf{B}_{I_i\cup J_{i+1}}\neq\mathbf{0}$ if $I_i\cup J_i(U_\tau) = I_i\cup J_{i+1}(U_\tau)$ and the only nonzero entry is

$$\mathbf{B}_{I_i\cup J_i}{}^\top\mathbf{B}_{I_i\cup J_{i+1}}[I_i\cup J_i(V_\tau), I_i\cup J_{i+1}(V_\tau)] = 1.$$

It follows that $\mathbf{B}^\top_{I_i\cup J_i}\mathbf{B}_{I_i\cup J_{i+1}} \neq \mathbf{0}$ for $1\le i\le 2t$ if $J_1(E_2\cap U_\tau) = J_2(E_2\cap U_\tau) =$

$\cdots = J_{2t}(E_2 \cap U_\tau)$. Lastly, $\mathbf{B}_{I_i \cup J_i}^\top \mathbf{B}_{I_i \cup J_{i+1}} \mathbf{B}_{I_{i+1} \cup J_{i+1}}^\top \mathbf{B}_{I_{i+1} \cup J_{i+2}} \neq \mathbf{0}$ if $I_i \cup J_{i+1}(V_\tau) = I_{i+1} \cup J_{i+1}(V_\tau)$ and the only nonzero entry is

$$\mathbf{B}_{I_i \cup J_i}^\top \mathbf{B}_{I_i \cup J_{i+1}} \mathbf{B}_{I_{i+1} \cup J_{i+1}}^\top \mathbf{B}_{I_{i+1} \cup J_{i+2}} [I_i \cup J_i(V_\tau), I_{i+1} \cup J_{i+2}(V_\tau)] = 1.$$

It follows that $\mathbf{B}_{I_1 \cup J_1}^\top \mathbf{B}_{I_1 \cup J_2} \mathbf{B}_{I_2 \cup J_2}^\top \mathbf{B}_{I_2 \cup J_3} \cdots \mathbf{B}_{I_{2t} \cup J_{2t}}^\top \mathbf{B}_{I_{2t} \cup J_1} \neq \mathbf{0}$ if $I_1(E_1 \cap V_\tau) = I_2(E_1 \cap V_\tau) = \cdots = I_{2t}(E_1 \cap V_\tau)$ and the only nonzero entry is

$$\mathbf{B}_{I_1 \cup J_1}^\top \mathbf{B}_{I_1 \cup J_2} \mathbf{B}_{I_2 \cup J_2}^\top \mathbf{B}_{I_2 \cup J_3} \cdots \mathbf{B}_{I_{2t} \cup J_{2t}}^\top \mathbf{B}_{I_{2t} \cup J_1} [I_1 \cup J_1(V_\tau), I_{2t} \cup J_1(V_\tau)] = 1.$$

Since $I_1(E_1 \cap V_\tau) = I_{2t}(E_1 \cap V_\tau)$, we have $I_1 \cup J_1(V_\tau) = I_{2t} \cup J_1(V_\tau)$. So if the summand is nonzero, it has a 1 on its diagonal.

To summarize, each summand in (2.19) is nonzero if and only if $I_i \cup J_i = \varphi_i(E_1 \cup E_2)$, $I_i \cup J_{i+1} = \varphi_i'(E_1 \cup E_2)$ for some $\varphi_i, \varphi_i' \in \mathcal{T}_n^k$, I_i's and J_i's agree on S, I_i's agree on $E_1 \cap V_\tau$ and J_i's agree on $E_2 \cap U_\tau$ for $1 \leq i \leq 2t$. Since each nonzero summand contributes a 1 on the diagonal, we can simply count the number of nonzero summands to compute the trace in (2.19). Notice that the number of nonzero summands is equal to the number of $\varphi_1, \ldots, \varphi_{2t}, \varphi_1', \ldots, \varphi_{2t}'$ that satisfy all the constraints. Hence

$$E \leq n^{|S|} \, n^{|E_1 \cap V_\tau|} \, n^{|E_2 \cap U_\tau|} \, n^{2t|E_1 \setminus V_\tau \setminus S|} \, n^{2t|E_2 \setminus U_\tau \setminus S|}. \tag{2.20}$$

For a fixed set of a, b, c, (2.20) provides an upper bound for $\|\mathbf{M}_{a,b,c,\sigma}\|_{4t}^{4t}$ for any permutation $\sigma \in \mathfrak{S}_d$. But what is an upper bound for $\|\mathbf{M}_{a,b,c,\sigma}\|_{4t}^{4t}$ among all possible sets of a, b, c? By varying a, b, c, we can always make $S = E_1 \cap E_2$ large enough to contain vertices in $E_1 \cap V_\tau$ and $E_2 \cap U_\tau$. Thus S is a vertex separator by Claim 2.6.4. It follows that

$E_1 \backslash V_\tau \backslash S = E_1 \backslash S$ and $E_2 \backslash U_\tau \backslash S = E_2 \backslash S$. Hence

$$E \;\le\; n^{|V(\tau)|}\, n^{2t|E_1 \backslash S|}\, n^{2t|E_2 \backslash S|} \;\le\; n^{|V(\tau)|}\, n^{2t(|V(\tau)|-|S_\tau|)}, \tag{2.21}$$

where S_τ is a vertex separator of minimum size. Substituting (2.21) into (2.18) yields the desired bound. ■

Corollary 2.6.6. *For any given shape τ, any $\varepsilon > 0$, with probability $1-\varepsilon$,*

$$\|\mathbf{M}\| \;\le\; |V(\tau)|^{|V(\tau)|} \left(C \log\left(|E(\tau)|^{|E(\tau)|} n^{|V(\tau)|}/\varepsilon\right)\right)^{|V(\tau)|} \left(\sqrt{\frac{1-p}{p}}\right)^{|E(S_\tau)|} \sqrt{n}^{|V(\tau)|-|S_\tau|},$$

where $C > 0$ is an absolute constant.

Proof. By Markov's inequality and Theorem 2.6.5, we have

$$\begin{aligned} &\mathbb{P}(\|\mathbf{M}\| \ge \theta) \;\le\; \mathbb{P}(\|\mathbf{M}\|_{4t}^{4t} \ge \theta^{4t}) \;\le\; \theta^{-4t}\mathbb{E}\,\|\mathbf{M}\|_{4t}^{4t} \\ &= \theta^{-4t}\left((48t|V(\tau)|)^{4t|V(\tau)|}(C|E(\tau)|)^{|E(\tau)|} n^{|V(\tau)|}\right)\left(\sqrt{\frac{1-p}{p}}\right)^{4t|E(S_\tau)|} n^{2t(|V(\tau)|-|S_\tau|)}. \end{aligned} \tag{2.22}$$

Set the right hand side of (2.22) to ε by taking

$$\theta = \left(\varepsilon^{-\frac{1}{4t}}(48t|V(\tau)|)^{|V(\tau)|}(C|E(\tau)|)^{\frac{|E(\tau)|}{4t}} n^{\frac{|V(\tau)|}{4t}}\right)\left(\sqrt{\frac{1-p}{p}}\right)^{|E(S_\tau)|} \sqrt{n}^{|V(\tau)|-|S_\tau|}.$$

Take

$$t = \frac{1}{4}\log\left(|E(\tau)|^{|E(\tau)|} n^{|V(\tau)|}/\varepsilon\right)$$

to complete the proof. ■

2.7 Application II: Tensor Networks

2.7.1 Background

A tensor network is a diagram of a collection of tensors connected together by contractions, which results in another tensor, matrix, vector or scalar. For example, Figure 2.4 represents a tensor $\mathbf{T}$ of order three.

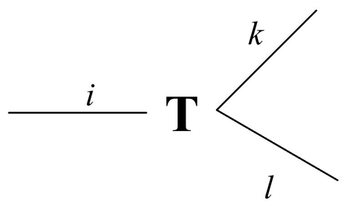

Figure 2.4: Tensor of order three

Now we have two copies of $\mathbf{T}$ and we would like to contract them together according to Figure 2.5.

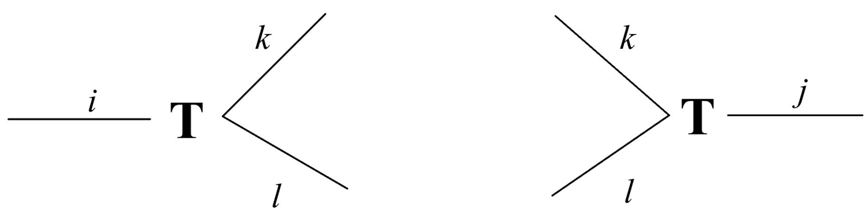

Figure 2.5: Contraction of two copies of $\mathbf{T}$

We simply follow the Einstein summation convention by summing up repeated indices, in this case, k and l. The nonrepeated indices are called free indices, in this case, i and j. So the resulting matrix (two free indices) is

$$\mathbf{A} = \sum_{k,l} \mathbf{T}_{ikl}\, \mathbf{T}_{klj},$$

and the tensor network representation of $\mathbf{A}$ is called the melon graph (Figure 2.6).

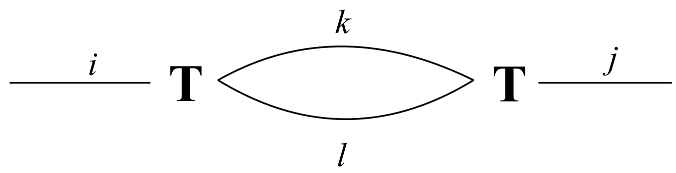

Figure 2.6: The Melon Graph

Aside from being an interesting notation, tensor networks are important for the desgin and analysis of spectral algorithms. Many existing algorithms for tensor decomposition and tensor PCA analyze spectral properties of matrices arising from tensor networks [RM14, HSS15, HSSS16, MW19, DdL$^+$22, OTR22]. The spectral properties are directly related to the signal-to-noise ratio of the algorithms.

2.7.2 *A simple tensor network: the melon graph*

The melon graph (Figure 2.6) is similar to the tensor unfolding method used in [RM14]. The spectral norm of the matrices associated to the melon graph is studied in [OTR22] using combinatorial tools developed in Random Tensor Theory [Gur16]. In this section, we will demonstrate our recursion framework for Gaussian polynomial random matrices (Theorem 2.5.3) in analyzing the spectral norm of the pure noise matrix associated to the melon graph.

Theorem 2.7.1. *(See also [Oue22] Theorem 7) Let $\mathbf{T} \in (\mathbb{R}^n)^{\otimes 3}$ be a random noise tensor with i.i.d. standard Gaussian entries. For $i, j \in [n]$, let T_i and T_j be the first-mode slices of $\mathbf{T}$. The pure noise matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ associated to the melon graph has entries $\mathbf{M}[i,j] = \langle T_i, T_j \rangle$. Then*

$$\mathbb{E}\,\|\mathbf{M} - \mathbb{E}\mathbf{M}\|_{2t}^{2t} \;\leq\; \mathcal{O}(n^{3t}).$$

Proof. Since $\mathbf{M}$ is a degree 2 homogeneous Gaussian random matrix, $\mathbb{E}\mathbf{M}_{1,0} = \mathbb{E}\mathbf{M}_{0,1} = \mathbf{0}$.

By Theorem 2.5.3,

$$\mathbb{E}\,\|\mathbf{M}-\mathbb{E}\mathbf{M}\|_{2t}^{2t} \;\le\; (8t)^{2t}\left(\|\mathbf{M}_{2,0}\|_{2t}^{2t}+\|\mathbf{M}_{0,2}\|_{2t}^{2t}+\|\mathbf{M}_{1,1}\|_{2t}^{2t}\right).$$

Let $\{\mathbf{B}_{i,j,k,l}\}$ be the second order partial derivatives of $\mathbf{M}$, i.e.,

$$\mathbf{B}_{i,j,k,l} = \partial_{\mathbf{T}_{i,k,l}}\,\partial_{\mathbf{T}_{j,k,l}}\mathbf{M}.$$

Note that $\mathbf{M}_{2,0}$, $\mathbf{M}_{0,2}$ and $\mathbf{M}_{1,1}$ are made up with $\{\mathbf{B}_{i,j,k,l}\}$ stacked in different ways. Let $E_{i,j}$ be the matrix basis with 1 in its $[i,j]$-th entry and 0 otherwise. For $i=j$,

$$\mathbf{B}_{i,i,k,l} \;=\; \partial^2_{\mathbf{T}_{i,k,l}}\mathbf{M} \;=\; 2\cdot E_{i,i},$$

For $i\neq j$,

$$\mathbf{B}_{i,j,k,l} \;=\; \partial_{\mathbf{T}_{i,k,l}}\,\partial_{\mathbf{T}_{j,k,l}}\mathbf{M} \;=\; E_{i,j}+E_{j,i}.$$

Hence,

$$\begin{aligned}
\|\mathbf{M}_{0,2}\|_{2t}^{2t} &= \mathrm{tr}\left(\mathbf{M}_{0,2}^{\top}\,\mathbf{M}_{0,2}\right)^{t}\\
&= \mathrm{tr}\left(\sum_{i,j,k,l\in[n]}\mathbf{B}_{i,j,k,l}^{\top}\mathbf{B}_{i,j,k,l}\right)^{t}\\
&= \mathrm{tr}\left(\sum_{i=j,k,l\in[n]}\mathbf{B}_{i,j,k,l}^{\top}\mathbf{B}_{i,j,k,l}+\sum_{i\neq j,k,l\in[n]}\mathbf{B}_{i,j,k,l}^{\top}\mathbf{B}_{i,j,k,l}\right)^{t}\\
&= \mathrm{tr}\left(\sum_{i,k,l\in[n]}2\cdot E_{i,i}+\sum_{i\neq j,k,l\in[n]}\left(E_{i,i}+E_{j,j}\right)\right)^{t}\\
&\le \mathrm{tr}\left(2n^2\cdot\mathbb{I}_n+n^3\cdot\mathbb{I}_n\right)^{t}\\
&= \mathcal{O}(n^{3t}).
\end{aligned}$$

Since $\mathbf{M}$ is hermitian, $\mathbf{B}^{\top}_{i,j,k,l} = \mathbf{B}_{i,j,k,l}$, so $\|\mathbf{M}_{0,2}\|_{2t}^{2t} = \|\mathbf{M}_{2,0}\|_{2t}^{2t}$. We proceed to bound $\|\mathbf{M}_{1,1}\|_{2t}^{2t}$. Since permutation of rows and columns doesn't affect Schatten norm, we permute the rows and columns of $\mathbf{M}_{1,1}$ into a block diagonal matrix with its diagonal block $\mathbf{D}_{k,l}$ indexed by k and l.

$$E := \|\mathbf{M}_{1,1}\|_{2t}^{2t} = \operatorname{tr}\left(\mathbf{M}_{1,1}^{\top}\mathbf{M}_{1,1}\right)^{t} = \sum_{k,l\in[n]} \operatorname{tr}\left(\mathbf{D}_{k,l}^{\top}\mathbf{D}_{k,l}\right)^{t} \leq n^2 \cdot \|\mathbf{D}_{k,l}\|_{2t}^{2t} \leq n^2 \cdot \left(\|\mathbf{D}_{k,l}\|_2^2\right)^{t},$$

where we used the Frobenius norm to upper bound the Schatten norm in the last inequality. $\mathbf{D}_{k,l}$ is a block matrix with $\mathbf{D}_{k,l}[i,i] = 2 \cdot E_{i,i}$ and $\mathbf{D}_{k,l}[i,j] = E_{i,j} + E_{j,i}$. So $\|\mathbf{D}_{k,l}\|_2^2 \leq 4n + 2n^2$. Thus

$$E \;\leq\; n^2(4n + 2n^2)^t \;\leq\; \mathcal{O}(n^{2t}).$$

Hence $\|\mathbf{M}_{0,2}\|_{2t}^{2t} \leq \mathcal{O}(n^{3t})$ dominates and the result follows. ∎

2.8 Application III: Smoothed Analysis

2.8.1 Background

It is not uncommon to observe gaps between theory and practice. The simplex algorithm invented by Dantzig [Dan51] is a great example. It is very efficient in practice, while it has exponential worst-case complexity [KM70]. In [ST04], Spielman and Teng introduced the smoothed analysis as a new measure of complexity to capture the practical usefulness of algorithms.

Formally, let A be an algorithm with input space X. For $\mathbf{x} \in X$, let $C(A, \mathbf{x})$ denote the run time of A on $\mathbf{x}$. The worst-case complexity of A is

$$\max_{\mathbf{x}\in X} \; C(A, \mathbf{x}).$$

The worst-case complexity is somewhat pessimistic. The issue is that the worst case is often quite contrived and does not arise naturally in practice.

Now let us take a look at the average-case complexity of A on a given distribution μ over X,

$$\mathbb{E}_{\mathbf{x}\sim\mu} \;\; C(A,\mathbf{x}).$$

Notice that the algorithm minimizing $\mathbb{E}_{\mathbf{x}\sim\mu}C(A,\mathbf{x})$ will typically be different for different distributions μ. The issue with average-case complexity is that we usually do not know what the appropriate distributions are.

The smoothed complexity is a hybrid model. Let D be the set of all possible distributions over X. Conceptually, the smoothed complexity can be thought of as

$$\max_{\mu\in D} \mathbb{E}_{\mathbf{x}\in\mu} \;\; C(A,\mathbf{x}). \tag{2.23}$$

If D contains only a single distribution μ, (2.23) reduces to $\mathbb{E}_{\mathbf{x}\in\mu}C(A,\mathbf{x})$, which is just the average-case complexity. If D contains all possible distributions over μ, (2.23) reduces to $\max_{\mathbf{x}\in X} C(A,\mathbf{x})$, which is just the worst-case complexity. Hence, the smoothed complexity "interpolates" between the worst-case and the average-case complexity.

More specifically, the input to the smoothed analysis is a randomly perturbed $\mathbf{x}$. The idea is that if the worst cases are sufficiently "isolated", then most randomly perturbed inputs are "well-behaved". The smoothed complexity of A is defined as

$$\max_{\mathbf{x}\in X} \mathbb{E}_{\mathbf{g}} \;\; C(A,\mathbf{x}+\mathbf{g}), \tag{2.24}$$

where each coordinate of $\mathbf{g}$ is an independent Gaussian random variable with mean 0 and variance σ^2. Let n be the input size. We say A has polynomial smoothed complexity, if (2.24) is polynomial in n and $1/\sigma$. See Part Four of the book [Rou21] edited by Roughgarden for further discussion on smoothed analysis and applications.

2.8.2 Smoothed Khatri–Rao product

The Khatri-Rao product is defined as the column-wise Kronecker product.

Definition 2.8.1 (Khatri-Rao product)**.** *Given* $\mathbf{A} \in \mathbb{R}^{p\times n}$ *and* $\mathbf{B} \in \mathbb{R}^{q\times n}$*. Let* $\{a_1, \ldots, a_n\}$ *be the columns of* $\mathbf{A}$ *and* $\{b_1, \ldots, b_n\}$ *be the columns of* $\mathbf{B}$*. The Khatri-Rao product is defined as*

$$\mathbf{A} \odot \mathbf{B} \;=\; \left[a_1 \otimes b_1 \middle| \;\; \ldots \;\; \ldots \;\; \middle| a_n \otimes b_n \right] \in \mathbb{R}^{pq\times n}.$$

Given a set of arbitrary vectors $v_1, \ldots, v_m \in \mathbb{R}^n$. We "smooth" each v_i by perturbing it with a small Gaussian noise, i.e., $v_i + g_i$. The set of corresponding smoothed vectors is $\{v_1 + g_1, \ldots, v_m + g_m\}$.

Definition 2.8.2 (Smoothed Khatri-Rao product)**.** *Given a set of arbitrary vectors* $v_1, \ldots v_m \in \mathbb{R}^n$*, the kth order smoothed Khatri-Rao product is defined as*

$$\mathbf{A}^{\odot k} \;=\; \left[(v_1 + g_1)^{\otimes k} \middle| \;\; \ldots \;\; \ldots \;\; \middle| (v_m + g_m)^{\otimes k} \right] \in \mathbb{R}^{n^k\times m},$$

where $g_1, \ldots, g_m$ *are independent centered normal random vectors.*

The goal is to study how large m can be so that the least singular value of $\mathbf{A}^{\odot k}$ remains bounded away from 0 with high probability. In other words, we want to know for what m a set of smoothed vectors $\{(v_i + g_i)^{\otimes k}\}_{i=1}^m$ is robustly linearly independent.

The Weyl's perturbation theorem shall help us in bounding $\sigma_{min}(\mathbf{A}^{\odot k})$. When $\mathbf{B}$ is the perturbation matrix with small spectral norm, we expect the least eigenvalue of $\mathbf{A} + \mathbf{B}$ is close to that of $\mathbf{A}$.

Theorem 2.8.3 (Weyl's Inequality)**.** *For $n \times n$ Hermitian matrices $\mathbf{A}$ and $\mathbf{B}$ with spectrum ordered in descending order $\lambda_1 \geq \cdots \geq \lambda_n$, we have*

$$|\lambda_i(\mathbf{A}+\mathbf{B}) - \lambda_i(\mathbf{A})| \;\leq\; \|\mathbf{B}\| \qquad\qquad \forall i \in [n].$$

The smoothed Khatri-Rao product of order one. This is a simple case in which we study $\sigma_{min}(\mathbf{A})$ where

$$\mathbf{A} \;=\; \left[v_1 + g_1 \middle| \;\; \ldots \;\; \ldots \;\; \middle| v_m + g_m \right] \in \mathbb{R}^{n \times m},$$

and g_i's are independent standard normal random vectors.

Let $S = span\{v_1, \ldots, v_m\}$ and Π be a projection matrix mapping $\mathbb{R}^n$ to $S^\perp$. Notice that we have no control over v_i's (they are simply given). Hence, we study the projection of $\{v_i + g_i\}_{i \in [m]}$ onto $S^\perp$ instead. Now we want to show that the projected vectors are linearly independent. Formally, we want to show that $\sigma_{min}(\mathbf{G})$ is bounded away from 0, where

$$\begin{aligned} \mathbf{G} \;&=\; \left[\Pi(v_1 + g_1) \middle| \;\; \ldots \;\; \ldots \;\; \middle| \Pi(v_m + g_m) \right] \\ &=\; \left[\Pi g_1 \middle| \;\; \ldots \;\; \ldots \;\; \middle| \Pi g_m \right]. \end{aligned}$$

We decompose $\mathbf{G}^\top\mathbf{G}$ into the diagonal and off-diagonal part and by Weyl's inequality,

$$\lambda_{min}(\mathbf{G}^\top\mathbf{G}) \;=\; \lambda_{min}\Big(diag(\mathbf{G}^\top\mathbf{G}) + offdiag(\mathbf{G}^\top\mathbf{G}) \Big) \tag{2.25}$$

$$\geq\; \lambda_{min}\Big(diag(\mathbf{G}^\top\mathbf{G}) \Big) - \lambda_{max}\Big(offdiag(\mathbf{G}^\top\mathbf{G}) \Big) \tag{2.26}$$

Notice that the diagonal entries $\mathbf{G}^\top\mathbf{G}[i,i] = \langle \Pi g_i, \Pi g_i \rangle = n \pm \mathcal{O}(\sqrt{n})$. Thus,

$$\lambda_{min}\Big(diag(\mathbf{G}^\top\mathbf{G}) \Big) \;=\; n \pm \mathcal{O}(\sqrt{n}).$$

If we want $\lambda_{min}(\mathbf{G}^\intercal \mathbf{G})$ to be bounded below by a polynomial in n, it is clear from (2.26) that we need

$$\left\| offdiag(\mathbf{G}^\intercal \mathbf{G}) \right\| \;\leq\; \left\| offdiag(\mathbf{G}^\intercal \mathbf{G}) \right\|_{2t} \;\leq\; \mathcal{O}(n).$$

Now we use Theorem 2.4.5 to bound $\|offdiag(\mathbf{G}^\intercal \mathbf{G})\|$. Let $\mathbf{F} = offdiag(\mathbf{G}^\intercal \mathbf{G}) \in \mathbb{R}^{m\times m}$. After decouping, we get

$$\mathbf{F}[i,j] \;=\; \langle \Pi g_i, \Pi g'_j \rangle \qquad \forall i \;\neq\; j \in [m].$$

$\mathbf{F}_{1,1} \in \mathbb{R}^{nm^2 \times nm^2}$ is a block matrix with

$$\mathbf{F}_{1,1}[(i,i),(j,j)] \;=\; \frac{\partial \mathbf{F}}{\partial g_i (\partial g_j)^\intercal} \;=\; \Pi \in \mathbb{R}^{n\times n}.$$

After an appropriate change of basis,

$$\mathbf{F}_{1,1}[(i,i),(j,j)] \;=\; \frac{\partial \mathbf{F}}{\partial g_i (\partial g_j)^\intercal} \;=\; I_n.$$

It follows that

$$\mathbf{F}_{1,1}^\intercal \mathbf{F}_{1,1}[(j,j),(j,j)] \;=\; \sum_i I_n \;=\; m \cdot I_n$$

Hence, for an appropriate t,

$$\|\mathbf{F}_{1,1}\|_{2t} \;\leq\; \mathcal{O}(m).$$

Now let us look at $\mathbf{F}_{2,0} \in \mathbb{R}^{n^2 m^3 \times m}$.

$$\mathbf{F}_{2,0}[i,j,l,k] \;=\; \frac{\partial \mathbf{F}}{\partial g_i[l] \partial g_j[k]} = \Pi_{lk} \cdot E_{ij} \in \mathbb{R}^{m\times m}$$

$$\begin{aligned}
\mathbf{F}_{2,0}^{\mathsf{T}}\mathbf{F}_{2,0} &= \sum_{i,j,l,k} \mathbf{F}_{2,0}^{\mathsf{T}}[i,j,l,k]\mathbf{F}_{2,0}[i,j,l,k] \\
&= \sum_{i,j}\sum_{l,k} \Pi_{l,k}^2 \cdot E_{j,i}E_{i,j} \\
&= \sum_{j} \|\Pi\|_F^2 \cdot E_{j,j} \\
&= n \cdot m \cdot I_m
\end{aligned}$$

Thus $\|\mathbf{F}_{2,0}\|_{2t} \leq \mathcal{O}(\sqrt{nm})$. It follows that

$$\left\| offdiag(\mathbf{G}^{\mathsf{T}}\mathbf{G}) \right\|_{2t} \ \leq \ \mathcal{O}(m + \sqrt{nm}).$$

If we choose $m = \mathcal{O}(n)$, we shall get the desired result.

CHAPTER 3

ELLIPSOID FITTING

The ellipsoid fitting conjecture of Saunderson, Chandrasekaran, Parrilo and Willsky considers the maximum number n random Gaussian points in $\mathbb{R}^d$, such that with high probability, there exists an origin-symmetric ellipsoid passing through all the points. They conjectured a threshold of $n = (1 - o_d(1)) \cdot d^2/4$, while until recently, known lower bounds on the maximum possible n were of the form $d^2/(\log d)^{O(1)}$.

In this chapter, we give a simple proof based on concentration of sample covariance matrices, that with probability $1 - o_d(1)$, it is possible to fit an ellipsoid through d^2/C random Gaussian points. Similar results were also obtained in two recent independent works by Hsieh, Kothari, Potechin and Xu [HKPX] and by Bandeira, Maillard, Mendelson, and Paquette [BMMP24].

3.1 Introduction

In this work, we consider the problem of fitting an ellipsoid to random Gaussian points in d dimensions. An origin-symmetric ellipsoid in $\mathbb{R}^d$ is specified by the equation $\langle x, Qx \rangle = 1$, for a positive semidefinite matrix $Q \succeq 0$. The following problem was considered by Saunderson, Chandrasekaran, Parrilo, and Willsky [SCPW12, SPW13] in the context of recovering a subspace from noisy measurements, using semidefinite programs: Given $z_1, \ldots, z_n \in \mathbb{R}^d$, does there exist $Q \succeq 0$ such that $\langle z_i, Qz_i \rangle = 1 \ \forall i \in [n]$?

They proved that for a d-dimensional subspace, fitting the projections of n canonical basis vectors V to an ellipsoid, is equivalent to recovering $V^\perp$ using a canonical semidefinite program. Thus, one is interested in determining how large can n be taken relative to d, to ensure small co-dimension for the subspace $V^\perp$ (see [Sau11] for a detailed discussion).

Saunderson et al. considered both average-case and worst-case versions of the above problem, and here we discuss average-case version where the points $z_1, \ldots, z_n$ can be taken to be independent samples from a Gaussian distribution. Since the problem is invariant under linear transformations (as Q can be modified to $LQL^\top$ for any linear transformation L), it is convenient to consider $z_1, \ldots, z_n \sim \mathcal{N}(0, \frac{1}{d} \cdot I_d)$, so that in expectation, the points lie on the unit sphere S^{d-1}. One then considers the problem of determining the largest n as a function of d, such that with probability close to 1 (say $1 - o_d(1)$), there exists an ellipsoid passing through the random points $z_1, \ldots, z_n$. It is easy to see that we must have $n < \binom{d+1}{2}$ via dimension arguments. Saunderson et al. [SCPW12, Sau11] made the following conjecture based on experimental results.

Conjecture 3.1.1 ([SCPW12, Sau11]). *For arbitrarily small $\varepsilon > 0$, if $n < (1/4 - \varepsilon) \cdot d^2$, then with high probability over $z_1, \ldots, z_n \sim \mathcal{N}(0, \frac{1}{d} \cdot I_d)$, there exists an ellipsoid passing through all the points.*

In addition to being an intriguing conjecture in its own right, this problem has also attracted recent attention because of its connections to techniques used in the study of lower bounds for semidefinite programs. In such lower bounds, the key technical challenge is often to come up with a candidate matrix, *which may depend on random data*, and then to prove the positivity of such a matrix (with high probability). The task of finding the matrix $Q \succeq 0$ also requires exhibiting a solution to a semidefinite program, and in fact some of the known results for the ellipsoid fitting problem have a significant overlap with semidefinite programming lower bounds, in terms of the techniques.

Known results. Saunderson et al. [SPW13] proved the ellipsoid fitting property (with high probability) for n random points when $n \leq d^{6/5-\eta}$ for arbitrarily small $\eta > 0$. This bound was improved to $n \leq d^{3/2-\eta}$ by a result of Ghosh et al. [GJJ+20] in the context of lower bounds for the Sum-of-Squares semidefinite programming hierarchy, for the closely related “Planted Affine Planes" problem. These bounds were further improved to $n \leq d^2/(\log d)^{O(1)}$ by two independent works of Potechin et al. [PTVW22], and of Kane and Diakonikolas [KD23]. The work of Potechin et al. relied on techniques used in proofs of Sum-of-Squares lower bounds, such as decompositions of random matrices into so called “graph matrices" (and a careful analysis of the coefficients in such decompositions, as well as trade-off between various terms [AMP21, PR23]). On the other hand, Kane and Diakonikolas gave a more direct proof based on the spectral analysis of a single random matrix arising from the random data.

Recent and independent work. Very recently, two independent works by Hsieh, Kothari, Potechin and Xu [HKPX], and by Bandeira, Maillard, Mendelson, and Paquette et al. [BMMP24] improved the bound on the number of points to $n \leq d^2/C$ for some (possibly large) constant $C > 0$, in both cases by obtaining a sharper analysis of the construction in [KD23]. The result of Hsieh et al. relies on a more careful application of the graph matrix decompositions arising in [PTVW22], while the work of Bandeira et al. uses a different approach based on concentration bounds for Gram matrices of independently sampled vectors from distributions with good tail behavior.

In this work, we present independent results, also proving that with high probability, there exists an ellipsoid passing through n random points, for $n \leq d^2/C$ (see Theorem 3.3.1 for a formal statement of our result). The approach in this work also gives a sharper analysis of the construction of Kane and Diakonikolas, obtaining spectral norm bounds for random matrices arising in their construction, by using results on the concentration of empirical covariance matrices. In addition to sharper bounds, this approach also yields a significantly simpler analysis, avoiding the need for norm bounds via the trace power method, as well as the careful accounting for trade-offs between terms in graph matrix decompositions. Similar ideas were also used independently in the work of Bandeira et al. [BMMP24], although the approaches differ in the specifics of the concentration results used to obtain the relevant norm bounds.

3.2 Preliminaries and notation

The Euclidean norm $\|\cdot\|_2$ and scalar product $\langle\cdot,\cdot\rangle$ in $\mathbb{R}^d$ are defined here using the counting measure on the coordinates. For a matrix $A \in \mathbb{R}^{m\times n}$, $\|A\|$ denotes its spectral norm, and $\|A\|_F$ denotes the Frobenius norm. For a real symmetric matrix B, $diag(B)$ denotes the diagonal matrix obtained from B replacing all non-diagonal entries by 0. We denote various positive absolute constants by $C, C', C_0, C_1, \ldots$ when the constants are greater than 1, and $c, c', c_0, c_1, \ldots$ when they are smaller than 1. The values of the constants may change from one instance to another.

Orlicz norms, sub-gaussian and sub-exponential random variables. Let X be a real-valued random variable with mean zero. We define the Orlicz α-norm ψ_α for $\alpha \geq 1$ as

$$\|X\|_{\psi_\alpha} = \inf\{t > 0 : \mathbb{E}\exp(|X|^\alpha/t^\alpha) \leq 2\}.$$

Let Z be a random vector in $\mathbb{R}^n$. ψ_α norm of a random vector is characterized by its one-dimensional projections,

$$\|Z\|_{\psi_\alpha} = \sup_{x\in S^{n-1}} \|\langle Z, x\rangle\|_{\psi_\alpha}.$$

A random variable X with bounded $\|X\|_{\psi_1}$ is known as a sub-exponential random variable, while one with bounded $\|X\|_{\psi_2}$ is known as a sub-gaussian random variable. The above definitions are also known to be equivalent to decay of the the tail probabilities $\mathbb{P}[|X| \geq t]$ as $\exp(-t/a_1)$ for sub-exponential random variables, and $\exp\left(-t^2/a_2^2\right)$ for sub-gaussian random variables, where a_1 and a_2 are within constant factors of the ψ_1 and ψ_2 norms respectively (see e.g., [Ver18]).

3.3 Fitting an ellipsoid to d^2/C random points

We will prove that with high probability over the choice of n independent Gaussian vectors in $\mathbb{R}^d$, there exists an origin-symmetric ellipsoid passing through all the points, when $n \leq d^2/C$ for some fixed constant C.

Theorem 3.3.1 (Main result)**.** *There exists a universal constant $C > 4$ such that for $n \leq d^2/C$, and with probability $1 - o_d(1)$ over the choice of i.i.d. samples $G_1, \ldots, G_n$ from $\mathcal{N}(0, \frac{1}{d} \cdot I_d)$, there exists a positive semidefinite matrix $Q \succeq 0$ satisfying $\langle G_i, QG_i \rangle \;=\; 1 \;\; \forall i \in [n]$.*

We use the identity-perturbation construction of Kane and Diakonikolas [KD23], which can be derived by considering the dual of the semidefinite program for finding the matrix Q as described above. Since the program has a constraint $\langle G_i, QG_i \rangle \;=\; 1$ for each $i \in [n]$, the construction is defined by dual variables $\delta_i \in \mathbb{R}$ for each $i \in [n]$. In particular, let $G_i \;=\; \|G_i\|_2 \cdot X_i$, where $X_i \in S^{d-1}$ are random unit vectors, independent of the norms $\|G_i\|_2$. Kane and Diakonikolas consider the matrix

$$Q \;=\; I_d + \sum_{i=1}^{n} \delta_i \cdot X_i {X_i}^{\mathsf{T}} \,,$$

and show that with high probability the (unique) choice of the variables δ_i satisfying $\langle G_i, QG_i \rangle = 1$ also satisfies $Q \succeq 0$. Considering the linear constraints $\langle G_i, QG_i \rangle = 1$ on the matrix Q above yields the linear system $M\delta = \varepsilon$, for the matrix M and the vector ε defined as

$$M_{ij} \;=\; \left\langle X_i, X_j \right\rangle^2 \;=\; \left\langle X_i^{\otimes 2}, X_j^{\otimes 2} \right\rangle \qquad \text{and} \qquad \varepsilon_i \;=\; \frac{1}{\|G_i\|_2^2} - 1 \,.$$

The proof then has two key components. The first is to show that the matrix M is invertible, and in fact the random matrix M^{-1} has bounded spectral norm with high probability. The second is to obtain bounds on the random vector ε (independent of M) of deviations in norms of G_i, to bound the norm of the perturbation $\sum_i \delta_i \cdot X_i X_i^{\mathsf{T}}$.

It is in the first part of the proof that our approach differs significantly from that of Diakonikolas and Kane [KD23]. Their proof relied on using the trace power method (which is the most technically involved part of the proof) and yielded bounds for $n \leq d^2/(\log d)^{O(1)}$. On the other hand, we rely on the fact that the vectors $X_i^{\otimes 2}$ have a sub-exponential distribution, and use existing results on the concentration of empirical covariance matrices for random sub-exponential vectors. The second part of our proof is similar to that of Kane and Diakonikolas, but needs to develop improved tail estimates for the vector ε, to avoid the loss of $(\log d)^{O(1)}$ factors.

3.3.1 Spectral norm bound for the Gram matrix

We will bound the norm of M^{-1} for the Gram matrix M defined as $M_{ij} = \left\langle X_i^{\otimes 2}, X_j^{\otimes 2} \right\rangle$ by using the results of Adamczak et al. [ALPTJ10a, ALPTJ10b] on the concentration of empirical covariance matrices, for independent sub-exponential random vectors. For independent (mean-zero) random vectors $Y_1, \ldots, Y_n$, forming columns of a matrix Y, the empirical covariance matrix is given (up to scaling) by YY^{T}. Since YY^{T} has the same non-zero eigenvalues as the the Gram matrix $Y^{\mathsf{T}}Y$, their proofs in fact proceed by understanding the spectrum of the Gram matrix, which implies the results needed for the matrix M above. We will need the following inequality.

Theorem 3.3.2 (Hanson-Wright Inequality, see e.g. [Ver18] Theorem 6.2.1)**.** *Let* $Z = (z_1, \ldots, z_d) \in \mathbb{R}^d$ *be a random vector with independent mean-zero sub-gaussian coordinates.*

Let A be a $d \times d$ matrix. Then, for every $t \geq 0$, we have

$$\mathbb{P}\left[\left|Z^{\mathsf{T}}AZ - \mathbb{E}\left[Z^{\mathsf{T}}AZ\right]\right| \geq t\right] \;\leq\; 2\exp\left(-c\min\left(\frac{t^2}{K^4\|A\|_F^2}, \frac{t}{K^2\|A\|}\right)\right),$$

where $K = \max_i \|z_i\|_{\psi_2}$.

We now consider the centered versions of the vectors $X^{\otimes 2}$, defined as $Y = X^{\otimes 2} - \mathbb{E}\,X^{\otimes 2}$, and show for any unit vector $u \in \mathbb{R}^{d^2}$, the random variable $\langle u, Y\rangle$ has sub-exponential tail. Since a vector $u \in \mathbb{R}^{d^2}$ can be thought of as a matrix $A \in \mathbb{R}^{d\times d}$, we can obtain the required tail behavior using Theorem 3.3.2.

Lemma 3.3.3 (Sub-exponential behavior)**.** *Let $X = (x_1, \ldots, x_d)$ be a random vector uniformly distributed on the unit sphere: $X \sim Unif(S^{d-1})$. Then $Y = X^{\otimes 2} - \mathbb{E}X^{\otimes 2}$ is a sub-exponential random vector satisfying $\left\|X^{\otimes 2} - \mathbb{E}X^{\otimes 2}\right\|_{\psi_1} \leq C/d$.*

Proof. Consider a standard normal random vector $G \sim \mathcal{N}(0, I_d)$. The direction $X = G/\|G\|_2$ is uniformly distribution on S^{d-1}. We need to show that $\langle X^{\otimes 2} - \mathbb{E}X^{\otimes 2}, u\rangle$ is sub-exponential for all unit vectors $u \in \mathbb{R}^{d^2}$. Reshape u into a $d \times d$ matrix A in which $a_{ij} = \langle u, e_i \otimes e_j\rangle$. Hence

$$\begin{aligned}\langle X^{\otimes 2} - \mathbb{E}X^{\otimes 2}, u\rangle \;&=\; \sum_{i,j=1}^{d}(x_i x_j - \mathbb{E}x_i x_j)\langle u, e_i \otimes e_j\rangle \\ &=\; X^{\mathsf{T}}AX - \mathbb{E}X^{\mathsf{T}}AX.\end{aligned}$$

We want to bound the tail probability,

$$\begin{aligned}\mathbb{P}\left[|\langle X^{\otimes 2} - \mathbb{E}X^{\otimes 2}, u\rangle| \geq \frac{t}{d}\right] \;&=\; \mathbb{P}\left[|X^{\mathsf{T}}AX - \mathbb{E}X^{\mathsf{T}}AX| \geq \frac{t}{d}\right] \\ &=\; \mathbb{P}\left[\left|\frac{G^{\mathsf{T}}AG}{\|G\|_2^2} - \mathbb{E}\frac{G^{\mathsf{T}}AG}{\|G\|_2^2}\right| \geq \frac{t}{d}\right]\end{aligned}$$

Notice that $\|G\|_2$ is well concentrated around $\sqrt{d}$. Thus the event $E := \{\|G\|_2 \geq \sqrt{d}/2\}$ is very likely and its complement $\mathbb{P}(E^c) \leq 2\exp(-c_1 \cdot d)$. Then

$$\begin{aligned}
\mathbb{P}\left[\left|\frac{G^\top AG}{\|G\|_2^2} - \mathbb{E}\frac{G^\top AG}{\|G\|_2^2}\right| \geq \frac{t}{d}\right] \;&\leq\; \mathbb{P}\left[\left\{\left|\frac{G^\top AG}{\|G\|_2^2} - \mathbb{E}\frac{G^\top AG}{\|G\|_2^2}\right| \geq \frac{t}{d}\right\} \wedge E\right] + \mathbb{P}(E^c) \\
&\leq\; \mathbb{P}\left[\left|G^\top AG - \mathbb{E}G^\top AG\right| \geq \frac{t}{4}\right] + 2 \cdot \exp(-c_1 \cdot d) \\
&\leq\; 2\exp\left(-c_0 \cdot \min(t^2, t)\right) + 2\exp(-c_1 \cdot d)\,,
\end{aligned}$$

where the last inequality is due to Theorem 3.3.2 with $\|A\| \leq \|A\|_F = 1$ and $K = \max_i \|G_i\|_{\psi_2} = 1$.

For the first term in the last inequality, take c_2 such that $2\exp(-c_2) \geq 1$ and let $c' = \min(c_0, c_2)$. The first term is bounded by $2\exp(-c' \cdot t)$. For the second term in the last inequality, notice that $|\langle X^{\otimes 2}, u\rangle| \leq \|X^{\otimes 2}\|_2 \|u\|_2 = 1$. This implies that $t \leq 2d$, otherwise the tail probability $\mathbb{P}\left(|\langle X^{\otimes 2} - \mathbb{E}X^{\otimes 2}, u\rangle| > 2\right) = 0$. If $t \leq 2d$, then $2\exp(-c_1 \cdot d) \leq 2\exp(-(c_1/2) \cdot t)$. Now let $c = \min(c_1/2, c')$ and substitute $\theta = t/d \geq 0$. It follows that

$$\mathbb{P}\left(|\langle X^{\otimes 2} - \mathbb{E}X^{\otimes 2}, u\rangle| \geq \theta\right) \leq 4\exp(-c \cdot \theta d).$$

It is a standard argument (see e.g. [Ver18] Proposition 2.7.1) that if the tail probability is bounded by $2\exp(-c\theta d)$ for all $\theta \geq 0$, then $\|X^{\otimes 2} - \mathbb{E}X^{\otimes 2}\|_{\psi_1} \leq C/d$. ∎

With the above lemma, we can now apply a result of Adamczak et al. [ALPTJ10a] which yields bounds on the singular values of a random matrix Y with columns of Y being independent sub-exponential random variables. We present a simplified version of Theorem 3.13 in [ALPTJ10a] below.

Theorem 3.3.4 ([ALPTJ10a], Simplified version of Theorem 3.13). *Let $Y_1, \ldots, Y_N$ be independent random vectors in $\mathbb{R}^D$ such that* $\max_{1\leq i\leq N} \|Y_i\|_{\psi_1} \leq \psi$ *and* $\max_{1\leq i\leq N} \|Y_i\|_2 \leq 1$. *Let Y*

be a random $D \times N$ matrix whose columns are Y_i's. Then

$$\mathbb{P}\left[\sigma_1(Y) \le C\psi \cdot \sqrt{N} + 6\right] \;\ge\; 1 - \exp(-c\sqrt{N})\,.$$

For our application, we will actually need a bound on the singular values which decreases with ψ, and it is in fact easy to check that the additive term of 6 in their bound arises from diagonal entries of the matrix $Y^{\mathsf{T}}Y$, which can be $\Omega(1)$ if the columns are unit vectors. We will need the following version of their result, bounding the contribution of the off-diagonal terms.

Lemma 3.3.5 (Off-diagonal estimate)**.** *Let $Y_1, \dots, Y_N$ be independent random vectors in $\mathbb{R}^D$ such that $\max_{1\le i\le N} \|Y_i\|_{\psi_1} \le \psi$ and $\max_{1\le i\le N} \|Y_i\|_2 \le 1$. Let Y be a random $n \times N$ matrix whose columns are Y_i's. Then*

$$\mathbb{P}\left[\left\|Y^{\mathsf{T}}Y - diag(Y^{\mathsf{T}}Y)\right\| \le C\max\{\psi^2 \cdot N, \psi \cdot \sqrt{N}\}\right] \;\ge\; 1 - \exp(-c \cdot \sqrt{N})\,.$$

Proof. We will use the tail bound for $\sigma_1(Y)$ from Theorem 3.3.4 to obtain a tail bound for $\|Y^{\mathsf{T}}Y - diag(Y^{\mathsf{T}}Y)\|$. In the proof of Theorem 3.13 [ALPTJ10a], $\left\|Y^{\mathsf{T}}Y - diag(Y^{\mathsf{T}}Y)\right\|$ and $\left\|diag(Y^{\mathsf{T}}Y)\right\|$ are estimated separately in terms of $\sigma_1(Y)$. They are combined at the last step to give the bound on $\sigma_1(Y)$ by the triangle inequality $\|Y^{\mathsf{T}}Y\| \le \|Y^{\mathsf{T}}Y - diag(Y^{\mathsf{T}}Y)\| + \|diag(Y^{\mathsf{T}}Y)\|$ together with an approximation argument. To simplify the bound presented in the proof of Theorem 3.13 [ALPTJ10a], we set the parameters exactly the same as in Theorem 3.3.4:

$$\mathbb{P}\left[\left\|Y^{\mathsf{T}}Y - diag(Y^{\mathsf{T}}Y)\right\| \le C_0\sigma_1(Y) \cdot \psi\sqrt{N}\right] \;\ge\; 1 - \exp(-c_0 \cdot \sqrt{N}). \tag{3.1}$$

Let $E_1 := \{\|Y^\top Y - diag(Y^\top Y)\| \le C_0\sigma_1(A) \cdot \psi\sqrt{N}\}$ and $E_2 := \{\sigma_1(Y) \le C\psi \cdot \sqrt{N} + 6\}$.

$$\mathbb{P}(E_1 \wedge E_2) \;=\; \mathbb{P}(E_1) - \mathbb{P}(E_1 \wedge E_2^c) \;\ge\; \mathbb{P}(E_1) - \mathbb{P}(E_2^c)$$

Substituting the probability bounds from Theorem 3.3.4 and Equation 3.1, we have

$$\left\|Y^\top Y - diag(Y^\top Y)\right\| \le C_0 \cdot \psi\sqrt{N} \cdot (C\psi\sqrt{N} + 6) \;\le\; C_1 \cdot \max\{\psi^2 N, \psi\sqrt{N}\}$$

with high probability. ∎

We can now apply the bound from the above lemma to control the spectral norm of M^{-1} for the Gram matrix M. We denote the event that $\left\|M^{-1}\right\| \le 3$ by E_1.

Lemma 3.3.6 (Operator norm bound: Event E_1). *Let M be the matrix defined as $M_{ij} = \left\langle X_i^{\otimes 2}, X_j^{\otimes 2}\right\rangle$, for $X_1, \ldots, X_n \sim Unif(S^{d-1})$. Then, for $n = d^2/C$ for a sufficient large constant C, the matrix M satisfies $M \succeq \frac{1}{3} \cdot I_n + \frac{1}{d} \cdot J_n$, and thus $\|M^{-1}\| \le 3$, with probability at least $1 - \exp(-c \cdot d)$.*

Proof. Let $Y_i = X_i^{\otimes 2} - \mathbb{E}\, X_i^{\otimes 2}$, where $\mathbb{E}\, X_i^{\otimes 2} = \frac{1}{d}\sum_{r=1}^{d} e_r^{\otimes 2}$. Note that

$$Y^\top Y \;=\; M + \frac{1}{d} \cdot J_n \qquad \text{and} \qquad \mathbb{E}\, M \;=\; \left(1 - \frac{1}{d}\right) \cdot I_n + \frac{1}{d} \cdot J_n\,,$$

where J_n denotes the all-ones matrix in $\mathbb{R}^{n \times n}$. Since the diagonal entries of M are 1, we also get that $diag(Y^\top Y) = \left(1 - \frac{1}{d}\right) \cdot I_n$, and thus $Y^\top Y - diag(Y^\top Y) = M - \mathbb{E}\, M$.

By Lemma 3.3.3, $\{Y_i\}_{i \in [n]}$ are independent sub-exponential random vectors with $\|Y_i\|_{\psi_1} \le C/d$. Thus Lemma 3.3.5 implies that

$$\mathbb{P}\left[\|Y^\top Y - diag(Y^\top Y)\| \le C \max\left\{\frac{n}{d^2}, \frac{\sqrt{n}}{d}\right\}\right] \;\ge\; 1 - \exp(-c\sqrt{n}).$$

Choose $n = d^2/C_1$ for large enough C_1 such that $C\sqrt{n}/d \le 1/2$. We have that with probability at least $1 - \exp(-c \cdot d)$,

$$\|Y^\mathsf{T} Y - diag(Y^\mathsf{T} Y)\| \;=\; \|M - \mathbb{E}M\| \;\leq\; 1/2\,.$$

This implies

$$M \;\succeq\; \left(\frac{1}{2} - \frac{1}{d}\right) \cdot I_n + \frac{1}{d} \cdot J_n \;\succeq\; \frac{1}{3} \cdot I_n$$

which gives $\left\|M^{-1}\right\| \leq 3$. ∎

3.3.2 Estimates for deviations in norms

We next develop some bounds for the random variables $\varepsilon_i = \frac{1}{\|G_i\|_2^2} - 1$, where $G_i \sim \mathcal{N}(0, 1/d \cdot I_d)$. Together with the bounds for $\left\|M^{-1}\right\|$ developed above, these will help us analyze the perturbation depending on the vector $\delta = M^{-1}\varepsilon$. We also note that the M and ε are independent random variables.

Lemma 3.3.7 (Sub-gaussian behavior of ε)**.** *The independent random variables $\varepsilon_1, \ldots, \varepsilon_n$ all satisfy*

1. $\mathbb{E}[\varepsilon_i] = \frac{2}{d-2}$ *and* $\operatorname{Var}[\varepsilon_i] = O(\frac{1}{d})$.
2. $\mathbb{P}[|\varepsilon_i - \mathbb{E}\,\varepsilon_i| \geq t] \;\leq\; 2 \cdot \exp\left(-t^2 \cdot d/32\right)$ *for all* $t \leq 1/2$.

Proof. We fix an $i \in [n]$ and write $\varepsilon_i = \frac{d}{\|G\|_2^2} - 1$ for $G \sim \mathcal{N}(0, I_d)$. The random variable $1/\left\|G\right\|_2^2$ has the inverse-χ^2 distribution with d degrees of freedom, for which the moments are known to be $\mathbb{E}[(1/\left\|G\right\|_2^2)^k] = 2^{-k} \cdot \frac{\Gamma((d/2)-k)}{\Gamma(d/2)}$ when $k \leq d/2$ [Wika, Wikb]. A direct computation yields

$$\mathbb{E}\left[\varepsilon_i\right] \;=\; \frac{2}{d-2} \qquad \text{and} \qquad \operatorname{Var}\left[\varepsilon_i\right] \;=\; \frac{2 \cdot d^2}{(d-2)^2 \cdot (d-4)} \;=\; O(1/d)\,.$$

Let $\varepsilon_0 = \mathbb{E}\varepsilon_i$ and $Z_i = \varepsilon_i - \mathbb{E}\varepsilon_i = \varepsilon_i - \varepsilon_0$. We want to bound the tail probability of Z_i. We have

$$\mathbb{P}\left[Z_i \geq t\right] \;=\; \mathbb{P}\left[\frac{d}{\|G\|_2^2} - 1 \geq t + \varepsilon_0\right] \;=\; \mathbb{P}\left[\|G\|_2^2 - d \leq -\frac{d(t+\varepsilon_0)}{1+t+\varepsilon_0}\right],$$

and

$$\mathbb{P}\left[Z_i \leq -t\right] \;=\; \mathbb{P}\left[\frac{d}{\|G\|_2^2} - 1 \leq -t + \varepsilon_0\right] \;=\; \mathbb{P}\left[\|G\|_2^2 - d \geq \frac{d(t-\varepsilon_0)}{1-t+\varepsilon_0}\right].$$

The tail bound for χ^2 distribution (see [Ver18] Theorem 3.1.1) implies that

$$\mathbb{P}\left[|Z_i| \geq t\right] \;\leq\; 2 \cdot \exp\left(-\frac{d}{8}\min\left\{\frac{t-\varepsilon_0}{1-t+\varepsilon_0}, \frac{t+\varepsilon_0}{1+t+\varepsilon_0}\right\}^2\right).$$

Using the fact that

$$\min\left\{\frac{t-\varepsilon_0}{1-t+\varepsilon_0}, \frac{t+\varepsilon_0}{1+t+\varepsilon_0}\right\}^2 \;\geq\; \left(\frac{t}{2}\right)^2, \quad \forall t \in [2\varepsilon_0, 1/2],$$

we can simplify the tail bound to

$$\mathbb{P}\left(|Z_i| \geq t\right) \;\leq\; 2\exp\left(-\frac{dt^2}{32}\right), \quad \forall t \leq \frac{1}{2},$$

since the probability is trivial for $t \leq 2\varepsilon_0$. ■

The above lemma also shows that with high probability $\|\varepsilon\|_\infty = O(\sqrt{(\log d)/d})$, which we denote as the event E_2. Actually, we will use E_2 as a shorthand for $\bigwedge E_{2,i}$, where $E_{2,i}$ is defined by a bound on $|\varepsilon_i|$ (so that the collection of random variables $\{\varepsilon_i | E_{2,i}\}_{i\in[n]}$ remain independent).

Corollary 3.3.8 (ℓ_∞ bound on ε: Event E_2). *With probability* $1 - d^{-\Omega(1)}$ *we have* $\|\varepsilon\|_\infty \leq$

$C\sqrt{\log(d)/d}$. *Also, conditioned on* E_2, *we have* $\|\varepsilon_i - \mathbb{E}\,\varepsilon_i\|_{\psi_2} \le C/\sqrt{d}$.

Proof. Using the tail bound from Lemma 3.3.7, and $n \le d^2$ we have that

$$\mathbb{P}\left[\exists i \in [n].\ \ |\varepsilon_i| \ge C_0 \cdot \frac{\sqrt{\log d}}{\sqrt{d}}\right] \ \le\ d^2 \cdot \exp\left(-C_1 \cdot \log d\right) \ \le\ d^{-C}.$$

Also, the tail bound for $Z_i = \varepsilon_i - \mathbb{E}\,\varepsilon_i$, conditioned on E_2 is:

$$\mathbb{P}\left[|Z_i| \ge t \mid E_2\right] \ =\ \frac{\mathbb{P}[|Z_i| \ge t \wedge E_2]}{\mathbb{P}[E_2]} \ \le\ \frac{2\exp(-dt^2/32)}{1 - o_d(1)} \ \le\ 2\exp(-dt^2/64), \quad \forall t.$$

The ψ_2 norm follows directly from the tail bound. ∎

Lemma 3.3.9 (ℓ^∞ bound on δ: Event E_3)**.** *With probability* $1 - o_d(1)$, *the vector* $\delta = M^{-1}\varepsilon$ *satisfies* $\|\delta\|_\infty \le C \cdot (\log d)/\sqrt{d}$.

Proof. We condition on events E_1 and E_2 both of which happen with probability $1 - o_d(1)$. We fix the matrix $A = M^{-1}$ satisfying $\|A\| \le 3$, and consider ε (which is independent of M) with independent entries $\varepsilon_1, \ldots, \varepsilon_n$ satisfying $\|\varepsilon_i - \mathbb{E}\,\varepsilon_i\|_{\psi_2} \le C/\sqrt{d}$ by Corollary 3.3.8. Then $\delta_i = \langle A_i, \varepsilon\rangle$, where A_i denotes the i-th row of A with $\|A_i\|_2 \le \|A\| \le 3$. By Hoeffding's inequality for independent sub-gaussian variables

$$\begin{aligned}
\mathbb{P}\left[|\delta_i - \mathbb{E}\,\delta_i| \ge t\right] \ &=\ \mathbb{P}\left[|\langle A_i, \varepsilon - \mathbb{E}\,\varepsilon\rangle| \ge t\right] \\
&\le\ 2\exp\left(-\frac{c \cdot t^2}{\|A_i\|_2^2 \cdot \max_i \|\varepsilon_i - \mathbb{E}\,\varepsilon_i\|_{\psi_2}^2}\right) \\
&\le\ 2\exp\left(-c' \cdot t^2 \cdot d\right).
\end{aligned}$$

To bound $\|\mathbb{E}\,\delta\|_\infty = \left\|M^{-1}\,\mathbb{E}\,\varepsilon\right\|_\infty$, we note that by Lemma 3.3.7, $\mathbb{E}\,\varepsilon = (2/d-2) \cdot \mathbf{1}$, and that conditioned on E_1 we actually have $M \succeq (1/3) \cdot I_n + (1/d) \cdot J_n$. Thus,

$$\|\mathbb{E}\,\delta\|_\infty \;=\; \left\|M^{-1}\,\mathbb{E}\,\varepsilon\right\|_\infty \;\leq\; \frac{2}{d-2} \cdot \left\|M^{-1}\mathbf{1}\right\|_2 \;\leq\; \frac{2}{d-2} \cdot \frac{d}{n} \cdot \|\mathbf{1}\|_2 \;\leq\; \frac{C_1}{\sqrt{n}}\,.$$

Choosing $t = C'(\log d)/\sqrt{d}$ in the Hoeffding estimate then proves the claim, since $n = d^2/C$. ■

3.3.3 Spectral bounds on the perturbation

To prove Theorem 3.3.1, we need to show that $Q = I_d + \sum_{i=1}^n \delta_i \cdot X_i {X_i}^\mathsf{T}$ is positive semidefinite. It suffices to show that the perturbation matrix $P = \sum_{i=1}^n \delta_i \cdot X_i {X_i}^\mathsf{T}$ satisfies $\|P\| \leq 1$. To obtain a bound on the spectral norm of the above matrix P, we will use the standard argument that it suffices to bound $\|Pu\|$ for all u in a κ-net of the sphere S^{d-1} i.e., a set of points N such that for all $x \in S^{d-1}$, there exists $y \in N$ with $\|x - y\|_2 \leq \kappa$. In particular, we will use the following result.

Lemma 3.3.10 ([Ver18]). *Let $A \in \mathbb{R}^{d\times d}$ be a symmetric matrix and let N be a κ-net of S^{d-1}. Then,*

$$\|A\| \;=\; \sup_{x\in S^{d-1}} |\langle x, Ax\rangle| \;\leq\; \frac{1}{1-2\kappa} \cdot \sup_{u\in N} |\langle u, Au\rangle|$$

Proof. Let $x \in S^{d-1}$ be the maximizer of the first quadratic form with $|\langle x, Ax\rangle| = \|A\|$, and let $u \in N$ be the closest point. Then

$$\begin{aligned}\|A\| - |\langle u, Au\rangle| \;&\leq\; |\langle (x-u), Ax\rangle| + |\langle x, A(x-u)\rangle| \\ &\leq\; 2 \cdot \|Ax\| \cdot \|x-u\| \\ &\leq\; 2\kappa \cdot \|A\|\;.\end{aligned}$$

■

Combined with known estimates on sizes of κ-nets, stating for example that there exists a 1/4-net size at most 9^d (see e.g., [Ver18], Corollary 4.2.13), it suffices to prove that for a *fixed* u in a 1/4-net N, we have $|\langle u, Pu\rangle| \leq 1/2$ with probability at least (say) $1 - \exp(-4d)$ (to allow for a union bound).

We start by considering $\langle u, Pu\rangle = \sum_{i=1}^{n} \delta_i \cdot \langle u, X_i\rangle^2$. We will consider the vector β with $\beta_i = \langle u, X_i\rangle^2$, and split it in two parts according to a threshold t_0 to be chosen later.

$$\beta_{heavy} = \sum_{i\in[n]} \beta_i \cdot e_i \cdot \mathbb{1}_{\{\beta_i > t_0\}} \qquad \text{and} \qquad \beta_{light} = \beta - \beta_{heavy}\,.$$

We first prove that $\mathrm{Supp}(\beta_{heavy}) = \{i | \beta_i > t_0\}$ is bounded in size.

Claim 3.3.11. *There exist $C_1, C_2 > 0$ such that over the choice of $X_1, \ldots, X_n$, we have for $t_0 \geq C_1/d$,*

$$\mathop{\mathbb{P}}_{X_1,\ldots,X_n}\left[\left|\mathrm{Supp}(\beta_{heavy})\right| \geq C_2/t_0\right] \;\leq\; \exp(-100\cdot d)\,.$$

Proof. Since $X_i \sim Unif(S^{d-1})$, $\|X_i\|_{\psi_2} \leq C/\sqrt{d}$, which implies $\mathbb{P}[\langle u, X_i\rangle^2 \geq t] \leq \exp(-c \cdot t \cdot d)$. By independence, the probability of having k heavy coordinates is at most $n^k \cdot \exp(-c \cdot t \cdot d \cdot k)$, which can be made smaller than $\exp(-100 \cdot d)$ for $k = O(1/t_0)$. ∎

We next prove that $\|\beta_{light}\|_2$ is bounded, which allows for concentration bounds on $\left\langle \delta, \beta_{light}\right\rangle$.

Lemma 3.3.12 (ℓ_2 bound on β_{light})**.** *There exist constants $c_0, C_0 > 0$ such that for $t_0 < c_0$, we have*

$$\mathop{\mathbb{P}}_{X_1,\ldots,X_n}\left[\left\|\beta_{light}\right\|_2^2 \geq C_0 \cdot \frac{n}{d^2}\right] \;\leq\; \exp(-100\cdot d)\,.$$

Proof. We will use again that sub-gaussian behavior of X_i implies $\mathbb{P}[\beta_i \geq t] \leq \exp(-c \cdot td)$.

For $\lambda \in (0, cd/(2t_0))$, we have

$$\begin{aligned}
\mathbb{E}\left[\exp\left(\lambda \cdot \min\{\beta_i^2, t_0^2\}\right)\right] &= \int_0^{exp(\lambda t_0^2)} \mathbb{P}\left(\exp(\lambda\beta_i^2) \geq s\right) \mathrm{d}s \\
&\leq 1 + \int_1^{exp(\lambda t_0^2)} \mathbb{P}\left(\exp(\lambda\beta_i^2) \geq s\right) \mathrm{d}s \\
&= 1 + \int_0^{t_0} \mathbb{P}\left(\beta_i^2 \geq t\right) 2\lambda t \exp(\lambda t^2)\, \mathrm{d}t \\
&\leq 1 + \int_0^{t_0} \exp(\lambda t^2 - cdt)\lambda t\, \mathrm{d}t \\
&\leq 1 + \int_0^{t_0} \exp(-cdt/2)\lambda t\, \mathrm{d}t \;=\; 1 + \frac{8\lambda}{c^2 d^2}
\end{aligned}$$

where the last inequality used $\lambda < cd/(2t_0)$. We thus have

$$\begin{aligned}
\mathbb{E}[\exp(\lambda \cdot \|\beta_{light}\|_2^2)] &\leq \prod_{i=1}^{n} \mathbb{E}\left[\exp\left(\lambda \cdot \min\{\beta_i^2, t_0^2\}\right)\right] \\
&\leq (1 + 8\lambda/c^2 d^2)^n \\
&\leq \exp(8\lambda n/c^2 d^2)\,.
\end{aligned}$$

Taking $c_0 < c/2$ and choosing $\lambda = d$ gives by Markov's inequality,

$$\mathbb{P}\left[\left\|\beta_{light}\right\|_2^2 \geq C_0 \cdot \frac{n}{d^2}\right] \;\leq\; \exp\left(\frac{8d \cdot n}{c^2 d^2} - \frac{d \cdot C_0 \cdot n}{d^2}\right) \;\leq\; \exp(-100 \cdot d)\,,$$

since the probability can be made small by an appropriate choice of C_0. ■

We now have all the ingredients to prove Theorem 3.3.1.

Proof of Theorem 3.3.1: Let $E = E_1 \wedge E_2 \wedge E_3$ and note that $\mathbb{P}(E) = 1 - o_d(1)$. Conditioned on E, we will show that for any $u \in N$, $|\langle u, Pu\rangle| = |\langle \delta, \beta\rangle| \leq 1/2$ with exponentially high probability, and then take a union bound over N. From Lemma 3.3.9 we know that $\|\delta\|_\infty = O(\log(d)/\sqrt{d})$ and Claim 3.3.11 implies that $\mathrm{Supp}(\beta_{heavy})$ is bounded by

$O(1/t_0)$ with probability $1 - \exp(-100 \cdot d)$. Thus,

$$|\langle \delta, \beta_{heavy} \rangle| \;\le\; \|\delta\|_\infty \cdot \|\beta_{heavy}\|_1 \;\le\; O\left(\log(d)/(\sqrt{d} \cdot t_0)\right) \;\le\; \frac{1}{6} \quad \text{for } t_0 = O(d^{-1/4})$$

Now we turn to the contribution from β_{light}, which can be written as

$$|\langle \delta, \beta_{light} \rangle| \;=\; |\langle M^{-1}\varepsilon, \beta_{light} \rangle| \;=\; |\langle M^{-1}\varepsilon, \beta_{light} \rangle| \;=\; |\langle \varepsilon, M^{-1}\beta_{light} \rangle|$$

By Lemma 3.3.6 and Lemma 3.3.12, we have $\|M^{-1}\beta_{light}\|_2^2 \;\le\; \|M^{-1}\|_2^2 \|\|\beta_{light}\|_2^2 \;=\; O(n/d^2)$. Since ε and $M^{-1}\beta_{light}$ are independent, and ε conditioned on E_2 has sub-gaussian norm $O(1/\sqrt{d})$ by Corollary 3.3.8, Hoeffding's inequality implies that,

$$\begin{aligned}
\mathbb{P}\left[\left|\langle \varepsilon - \mathbb{E}\,\varepsilon, M^{-1}\beta_{light} \rangle\right| > \frac{1}{6}\right] \;&\le\; 2 \cdot \exp\left(-\frac{c/36}{\|M^{-1}\beta_{light}\|_2^2 \cdot \|\varepsilon - \mathbb{E}\,\varepsilon\|_{\psi_2}^2}\right) \\
&=\; 2 \cdot \exp(-c' \cdot d^3/n)\,,
\end{aligned}$$

which is at most $\exp(-100 \cdot d)$ when $n < d^2/C$ for an appropriate C. Finally, note that using Lemma 3.3.7 and Lemma 3.3.12, with probability $1 - \exp(-100 \cdot d)$,

$$\begin{aligned}
\left|\left\langle \mathbb{E}\,\varepsilon, M^{-1}\beta_{light} \right\rangle\right| \;&=\; \frac{2}{d-2} \cdot \left|\left\langle \mathbf{1}, M^{-1}\beta_{light} \right\rangle\right| \\
&\le\; \frac{2}{d-2} \cdot \sqrt{n} \cdot \left\|M^{-1}\beta_{light}\right\| \\
&=\; O\left(\frac{1}{d} \cdot \sqrt{n} \cdot \frac{\sqrt{n}}{d}\right)\,.
\end{aligned}$$

Thus, we have that conditioned on E, $\mathbb{P}[|\langle \delta, \beta \rangle| \ge 1/2] \le \exp\left(-10 \cdot d\right)$ when $n <$

d^2/C for a large enough C. We can now take the union bound on the net to show

$$\begin{aligned}\mathbb{P}\left(\sup_{u\in N}\left|\sum_{i=1}^{n}\delta_i\left\langle u,X_i\right\rangle^2\right|\geq\frac{1}{2}\right) &\leq \mathbb{P}\left(\sup_{u\in N}\left|\sum_{i=1}^{n}\delta_i\left\langle u,X_i\right\rangle^2\right|\geq\frac{1}{2}\middle|E\right)\cdot\mathbb{P}(E)+\mathbb{P}(E^c)\\ &\leq 9^d\cdot\left[\mathbb{P}\Big(|\langle\delta,\beta\rangle|\geq 1/2\Big|E\Big)\cdot\mathbb{P}(E)\right]+\mathbb{P}(E^c)\\ &\leq 9^d\cdot\exp(-10\cdot d)+o_d(1) \;=\; o_d(1)\,.\end{aligned}$$

Combined with Lemma 3.3.10, this implies $\|P\|\leq 1$ and hence $Q\succeq 0$ with probability $1-o_d(1)$. ∎

3.4 Other explicit constructions

The ellipsoid fitting problem can be viewed as a semidefinite program with random linear constraints and a postivitiy constraint, i.e., find Q such that

$$\begin{cases} Q \;=\; Q^{\mathsf{T}} \\ \langle G_i, QG_i\rangle \;=\; 1 \qquad \forall i\in[n] \\ Q \;\succeq\; 0. \end{cases} \tag{3.2}$$

See [BMMP24] Corollary 1.3 for the dual program of (3.2).

The common approach in [KD23, PTVW22, BMMP24, TW25a] is to construct Q explicitly and prove that Q is positive semidefinite with high probability. This approach relaxes the positivty constraint in (3.2) and seeks to minimize some "potential" functions of Q instead. There are two types of "potential" functions discussed in [MK24].

1. $V_1(Q) \;=\; \|Q-I_d\|_t$ for some $t\geq 1$.

2. $V_2(Q) \;=\; \|Q\|_t$ for some $t\geq 1$.

$\| \cdot \|_t$ denotes Schatten-t norm.

We briefly show the derivation of the identity perturbation construction from $V_1(Q)$. Now we set to solve

$$\begin{cases} \min & \frac{1}{2}\|Q - I_d\|_2^2 \\ s.t. & \langle G_i, QG_i \rangle \ = \ 1 \qquad \forall i \in [n]. \end{cases} \tag{3.3}$$

The Lagrangian of (3.3) is

$$\mathcal{L}(Q, \lambda) \ = \ \frac{1}{2} \operatorname{tr}((Q - I_d)^\mathsf{T}(Q - I_d)) - \sum_{i=1}^{n} \lambda_i \langle G_i, QG_i \rangle + \sum_{i=1}^{n} \lambda_i.$$

By taking the derivative of $\mathcal{L}(Q, \lambda)$ with respect to Q and setting the derivatives to 0, We get

$$\frac{\partial \mathcal{L}(Q, \lambda)}{\partial Q} = Q - I_d - \sum_{i=1}^{n} \lambda_i G_i G_i^\mathsf{T} = 0.$$

It follows that

$$Q_{IP} \ = \ I_d + \sum_{i=1}^{n} \lambda_i G_i G_i^\mathsf{T}.$$

Since these explicit constructions are obtained by relaxing (3.2), each construction has different phase transition threshold. See [MK24, BM23] for more detailed discussion on various thresholds. From the numerical experiments in [MK24, PTVW22], the identity perturbation construction of (3.3) undergo phase transition around $d^2/10$. Beyond $d^2/10$, Q_{IP} ceases to be positive semidefinite with high probability. A surprising result from [MK24, BM23] is that the phase transition threshold for the minimal nuclear norm construction seems to be $d^2/4$ from numerical experiments, which coincides with the original ellipsoid fitting problem.

The minimal nuclear norn construction from $V_2(Q)$ is

$$\begin{cases} \min & \|Q\|_1 \\ s.t. & \langle G_i, QG_i \rangle \ = \ 1 \qquad \forall i \in [n]. \end{cases} \tag{3.4}$$

The Lagrangian of (3.4) is

$$\mathcal{L}(Q, \lambda) \ = \ \mathrm{tr}(Q^\intercal Q)^{\frac{1}{2}} - \sum_{i=1}^{n} \lambda_i \langle G_i, QG_i \rangle + \sum_{i=1}^{n} \lambda_i.$$

By taking the derivative of $\mathcal{L}(Q, \lambda)$ with respect to Q and setting the derivatives to 0, We get

$$\frac{\partial \mathcal{L}(Q, \lambda)}{\partial Q} = Q(Q^\intercal Q)^{-\frac{1}{2}} - \sum_{i=1}^{n} \lambda_i G_i G_i^\intercal = 0.$$

It follows that

$$Q(Q^\intercal Q)^{-\frac{1}{2}} \ = \ \sum_{i=1}^{n} \lambda_i G_i G_i^\intercal,$$

and

$$Q \ = \ \sum_{i=1}^{n} \lambda_i G_i G_i^\intercal (Q^\intercal Q)^{\frac{1}{2}}.$$

Substitute Q back into the random linear constraint, we get that

$$\begin{cases} \left\langle G_1, \quad \sum_{i=1}^{n} \lambda_i G_i G_i^\intercal (Q^\intercal Q)^{\frac{1}{2}} G_1 \right\rangle = 1 \\ \qquad \vdots \\ \left\langle G_n, \quad \sum_{i=1}^{n} \lambda_i G_i G_i^\intercal (Q^\intercal Q)^{\frac{1}{2}} G_n \right\rangle = 1 \end{cases} \tag{3.5}$$

If there is an intermediate solution for $(Q^\intercal Q)^{\frac{1}{2}}$, we can solve the system of equations above to get λ_i. Then we can compute Q by substituting in λ_i and the intermediate solution for$(Q^\intercal Q)^{\frac{1}{2}}$. This motivates an iterative algorithm for solving Q. Of course, it remains to show that Q is positive semi-definite with high probability. Although it looks mathematically daunting, we think it is an interesting direction for future work.

REFERENCES

[ABY20] Richard Aoun, Marwa Banna, and Pierre Youssef. Matrix poincaré inequalities and concentration. *Advances in Mathematics*, 371:107251, 2020. 8, 12, 15

[Ada15] Radosław Adamczak. A note on the hanson-wright inequality for random vectors with dependencies. *Electronic Communications in Probability*, 20:1–13, 2015.

[ALPTJ10a] Radosław Adamczak, Alexander E. Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of The American Mathematical Society*, 23(2):535–561, 2010. 75, 77, 78

[ALPTJ10b] Radosław Adamczak, Alexander E. Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Sharp bounds on the rate of convergence of the empirical covariance matrix. *Comptes Rendus Mathematique*, 349, 2010. 75

[AMP21] Kwangjun Ahn, Dhruv Medarametla, and Aaron Potechin. Graph matrices: Norm bounds and applications. *arXiv preprint https://arxiv.org/abs/1604.03423*, 2021. 12, 14, 22, 26, 72

[ASBvH25] Petar Nizić-Nikolac Afonso S. Bandeira, Kevin Lucca and Ramon van Handel. Matrix chaos inequalities and chaos of combinatorial type. In *Proceedings of the Fifty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '25, 2025. 16

[AW15] Radosław Adamczak and Paweł Wolff. Concentration inequalities for non-lipschitz functions with bounded derivatives of higher order. *Probability Theory and Related Fields*, 162(3):531–586, 2015. 4, 5, 13, 21

[BBLM05] Stéphane Boucheron, Olivier Bousquet, Gábor Lugosi, and Pascal Massart. Moment inequalities for functions of independent random variables. *The Annals of Probability*, 33(2):514 – 560, 2005. 15

[BBvH23] Afonso S Bandeira, March T Boedihardjo, and Ramon van Handel. Matrix concentration inequalities and free probability. *Inventiones Mathematicae*, pages 1–69, 2023. 7, 12

[Ben62] George Bennett. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57(297):33–45, 1962. 2

[BGS19] Sergey G Bobkov, Friedrich Götze, and Holger Sambale. Higher order concentration of measure. *Communications in Contemporary Mathematics*, 21(03):1850043, 2019. 4, 13

[BHK+19] Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019. 12

[BHKX22] Mitali Bafna, Jun-Ting Hsieh, Pravesh K. Kothari, and Jeff Xu. Polynomial-time power-sum decomposition of polynomials. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 956–967, 2022. 12

[BJM23] Nikhil Bansal, Haotian Jiang, and Raghu Meka. Resolving matrix Spencer conjecture up to poly-logarithmic rank. pages 1814–1819, 2023. 12

[BLM03] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. Concentration inequalities using the entropy method. *The Annals of Probability*, 31(3):1583 – 1614, 2003. 3

[BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 02 2013. 1

[BM23] Afonso S Bandeira and Antoine Maillard. Exact threshold for approximate ellipsoid fitting of random points. *arXiv preprint arXiv:2310.05787*, 2023. 88

[BMMP24] Afonso S Bandeira, Antoine Maillard, Shahar Mendelson, and Elliot Paquette. Fitting an ellipsoid to a quadratic number of random points. *Latin American Journal of Probability and Mathematical Statistics*, 21(2):1835, 2024. 70, 72, 73, 87

[Bor75] Christer Borell. The brunn-minkowski inequality in gauss space. *Inventiones mathematicae*, 30:207–216, 1975. 4

[BvH22] Tatiana Brailovskaya and Ramon van Handel. Universality and sharp matrix concentration inequalities. *arXiv preprint arXiv:2201.05142*, 2022. 12

[Dan51] G.B. Dantzig. Maximization of a linear function of variables subject to linear inequalities. *Activity Analysis of Production and Allocation*, 1951. 64

[DdL$^+$22] Jingqiu Ding, Tommaso d'Orsi, Chih-Hung Liu, David Steurer, and Stefan Tiegel. Fast algorithm for overcomplete order-3 tensor decomposition. In *Conference on Learning Theory*, pages 3741–3799. PMLR, 2022. 12, 62

[ES81] B. Efron and C. Stein. The Jackknife Estimate of Variance. *The Annals of Statistics*, 9(3):586 – 596, 1981. 3

[GHK15] Rong Ge, Qingqing Huang, and Sham M. Kakade. Learning mixtures of gaussians in high dimensions. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, page 761–770, New York, NY, USA, 2015. Association for Computing Machinery. 16

[GJJ+20] Mrinalkanti Ghosh, Fernando Granha Jeronimo, Chris Jones, Aaron Potechin, and Goutham Rajendran. Sum-of-squares lower bounds for Sherrington-Kirkpatrick via planted affine planes. In *61st Annual Symposium on Foundations of Computer Science (FOCS)*, page 954–965. IEEE, 2020. 71

[Gur16] Răzvan Gheorghe Gurău. *Random Tensors*. Oxford University Press, 10 2016. 62

[HKPX] Jun-Ting Hsieh, Pravesh K. Kothari, Aaron Potechin, and Jeff Xu. Ellipsoid Fitting up to a Constant. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. 70, 72

[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. 2

[HSS15] Samuel B Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-square proofs. In *Conference on Learning Theory*, pages 956–1006. PMLR, 2015. 62

[HSS19] Samuel B Hopkins, Tselil Schramm, and Jonathan Shi. A robust spectral algorithm for overcomplete tensor decomposition. In *Conference on Learning Theory*, pages 1683–1722. PMLR, 2019. 12

[HSSS16] Samuel B Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. pages 178–191, 2016. 12, 62

[HT20] De Huang and Joel A. Tropp. From poincaré inequalities to nonlinear matrix concentration. *Bernoulli*, 2020. 8, 12, 15, 21, 51

[HT21] De Huang and Joel A. Tropp. Nonlinear matrix concentration via semigroup methods. *Electronic Journal of Probability*, 26:Art. No. 8, Jan 2021. 8, 12, 17

[JPR+21] Chris Jones, Aaron Potechin, Goutham Rajendran, Madhur Tulsiani, and Jeff Xu. Sum-of-squares lower bounds for sparse independent set. 2021. 12, 26

[KD23] Daniel Kane and Ilias Diakonikolas. A nearly tight bound for fitting an ellipsoid to gaussian random points. In *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 3014–3028. PMLR, 12–15 Jul 2023. 72, 74, 75, 87

[KM70] Victor Klee and George J. Minty. How good is the simplex algorithm. In *Inequalities*, 1970. 64

[KP20] Bohdan Kivva and Aaron Potechin. Exact nuclear norm, completion and decomposition for random overcomplete tensors via degree-4 sos. *arXiv preprint arXiv:2011.09416*, 2020. 12

[KV00] Jeong Han Kim and Van H Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000. 4, 13

[Kwa87] Stanislaw Kwapien. Decoupling Inequalities for Polynomial Chaos. *The Annals of Probability*, 15(3):1062 –1071, 1987. 16, 31

[Lat99] Rafał Latała. Tail and moment estimates for some types of chaos. *Studia Mathematica*, 135(1):39–53, 1999.

[Lat06] Rafał Latała. Estimates of moments and tails of gaussian chaoses. *The Annals of Probability*, 34(6):2315–2331, 2006. 4, 13

[LM00] Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics*, 28(5):1302–1338, 2000.

[LY23] Cécilia Lancien and Pierre Youssef. A note on quantum expanders, 2023. 12

[MJC+12] Lester Mackey, Michael Jordan, Richard Chen, Brendan Farrell, and Joel Tropp. Matrix concentration inequalities via the method of exchangeable pairs. *The Annals of Probability*, 42, 2012. 5, 12, 38

[MK24] Antoine Maillard and Dmitriy Kunisky. Fitting an ellipsoid to random points: predictions using the replica method. *IEEE Transactions on Information Theory*, 2024. 87, 88

[MT84] Terry R. McConnell and Murad S. Taqqu. Double integration with respect to symmetric stable processes. Technical report, Technical Report 618, Cornell Univ., 1984. 16

[MW19] Ankur Moitra and Alexander S Wein. Spectral methods from tensor networks. pages 926–937, 2019. 12, 62

[OTR22] Mohamed Ouerfelli, Mohamed Tamaazousti, and V. Rivasseau. Random tensor theory for tensor decomposition. In *AAAI Conference on Artificial Intelligence*, 2022. 62

[Oue22] Mohamed Ouerfelli. *New perspectives and tools for Tensor Principal Component Analysis and beyond*. PhD thesis, Université Paris-Saclay, 2022. 62

[OZ16] Ryan O'Donnell and Yu Zhao. Polynomial bounds for decoupling, with applications. In *Proceedings of the 31st Conference on Computational Complexity*, CCC '16, Dagstuhl, DEU, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. 16

[PG99] Víctor H. Peña and Evarist Giné. *Decoupling: From dependence to independence*. Springer-Verlag, 1999. 16, 31

[Pis03] Gilles Pisier. *Introduction to Operator Space Theory*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2003. 6

[PMS95] Victor H. Pena and S. J. Montgomery-Smith. Decoupling Inequalities for the Tail Probabilities of Multivariate U-Statistics. *The Annals of Probability*, 23(2):806 – 816, 1995. 32

[PMT16] Daniel Paulin, Lester Mackey, and Joel A. Tropp. Efron–Stein inequalities for random matrices. *The Annals of Probability*, 44(5):3431 – 3473, 2016. 7, 8, 12, 15

[PR23] Aaron Potechin and Goutham Rajendran. Machinery for proving sum-of-squares lower bounds on certification problems. 2023. 72

[PTVW22] Aaron Potechin, Paxton Turner, Prayaag Venkat, and Alexander S Wein. Near-optimal fitting of ellipsoids to random points. *arXiv preprint arXiv:2208.09493*, 2022. 72, 87, 88

[Raj22] Goutham Rajendran. *Nonlinear Random Matrices and Applications to the Sum of Squares Hierarchy*. PhD thesis, University of Chicago, 2022. 12

[Rau10] Holger Rauhut. Compressive sensing and structured random matrices. In *Theoretical Foundations and Numerical Methods for Sparse Recovery*, pages 1–92. De Gruyter, Berlin, New York, 2010. 16, 32

[RM14] Emile Richard and Andrea Montanari. A statistical model for tensor pca. In *Neural Information Processing Systems*, 2014. 62

[Rou21] Tim Roughgarden, editor. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2021. 65

[RT23] Goutham Rajendran and Madhur Tulsiani. Concentration of polynomial random matrices via efron-stein inequalities. *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2023. 15, 26, 54, 55

[SA13] Khalid Shebrawi and Hussien Albadawi. Trace inequalities for matrices. *Bulletin of the Australian Mathematical Society*, 87(1):139–148, 2013. 50

[Sau11] James Francis Saunderson. *Subspace identification via convex optimization*. PhD thesis, Massachusetts Institute of Technology, 2011. 70, 71

[SCPW12] James Saunderson, Venkat Chandrasekaran, Pablo A. Parrilo, and Alan S. Willsky. Diagonal and low-rank matrix decompositions, correlation matrices, and ellipsoid fitting. *SIAM Journal on Matrix Analysis and Applications*, 33(4):1395–1416, 2012. 70, 71

[SPW13] James Saunderson, Pablo A. Parrilo, and Alan S. Willsky. Diagonal and low-rank decompositions and fitting ellipsoids to random point. In *52nd IEEE Conference on Decision and Control*, page 6031–6036. IEEE, 2013. 70, 71

[SS11] Warren Schudy and Maxim Sviridenko. Bernstein-like concentration and moment inequalities for polynomials of independent random variables: multilinear case. *arXiv preprint arXiv:1109.5193*, 2011. 4, 13

[SS12] Warren Schudy and Maxim Sviridenko. Concentration and moment inequalities for polynomials of independent random variables. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 437–446. SIAM, 2012. 4, 13

[ST04] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM*, 51(3):385–463, 2004. 64

[Ste86] J. Michael Steele. An Efron-Stein Inequality for Nonsymmetric Statistics. *The Annals of Statistics*, 14(2):753 – 758, 1986. 3

[TC13] Joel Tropp and Richard Chen. Subadditivity of matrix phi-entropy and concentration of random matrices. *Electronic Journal of Probability [electronic only]*, 19, 08 2013. 9

[Tro15] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015. 5, 6, 11, 30

[TW25a] Madhur Tulsiani and June Wu. Ellipsoid fitting up to constant via empirical covariance estimation. In *2025 Symposium on Simplicity in Algorithms (SOSA)*, pages 134–143, 2025. 10, 87

[TW25b] Madhur Tulsiani and June Wu. Simple norm bounds for polynomial random matrices via decoupling. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, Leibniz International Proceedings in Informatics (LIPIcs), 2025. 9

[Ver12] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing: Theory and Applications*, page 210–268. Cambridge University Press, 2012.

[Ver18] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2018. 16, 74, 75, 77, 81, 83, 84

[Wika] Wikipedia. Inverse-chi-squared distribution. 80

[Wikb] Wikipedia. Inverse-gamma distribution. 80