

MATRIX CHAOS INEQUALITIES AND CHAOS OF COMBINATORIAL TYPE

AFONSO S. BANDEIRA, KEVIN LUCCA, PETAR NIZIĆ-NIKOLAC, AND RAMON VAN HANDEL

ABSTRACT. Matrix concentration inequalities and their recently discovered sharp counterparts provide powerful tools to bound the spectrum of random matrices whose entries are linear functions of independent random variables. However, in many applications in theoretical computer science and in other areas one encounters more general random matrix models, called matrix chaoses, whose entries are polynomials of independent random variables. Such models have often been studied on a case-by-case basis using ad-hoc methods that can yield suboptimal dimensional factors.

In this paper we provide general matrix concentration inequalities for matrix chaoses, which enable the treatment of such models in a systematic manner. These inequalities are expressed in terms of flattenings of the coefficients of the matrix chaos. We further identify a special family of matrix chaoses of combinatorial type for which the flattening parameters can be computed mechanically by a simple rule. This allows us to provide a unified treatment of and improved bounds for matrix chaoses that arise in a variety of applications, including graph matrices, Khatri-Rao matrices, and matrices that arise in average case analysis of the sum-of-squares hierarchy.

CONTENTS

1. Introduction	1
2. Matrix chaos inequalities	3
2.1. Matrix chaos and decoupling	4
2.2. Flattenings	4
2.3. The main inequalities	5
2.4. Iteration approach	7
3. Chaos of combinatorial type	9
3.1. Definition and guiding example	9
3.2. How to compute norms of flattenings	10
3.3. Chaos of nearly combinatorial type	12
4. Applications	12
4.1. Khatri-Rao matrices	12
4.2. The sum-of-squares algorithm for tensor PCA	13
4.3. Graph matrices	14
4.4. Sharper bounds on graph matrices and ellipsoid fitting	19
Acknowledgements	20
References	20
Appendix A. Proofs of main results and supporting lemmas	22

1. INTRODUCTION

Classical random matrix theory is largely concerned with special models, such as matrices with i.i.d. entries, whose spectral properties are understood asymptotically with stunning precision. However, random matrices that appear in applications in theoretical computer science and in other

fields often fall outside the scope of the classical models; moreover, these applications typically require an understanding of such models in the nonasymptotic regime.

One of the key advances from this perspective has been the development of a large family of *matrix concentration inequalities* that are widely used in applications. These inequalities can be applied to random matrices whose entries are very general *linear* functions of independent random variables. A prototypical example of such a model is any random matrix with centered jointly gaussian entries (with arbitrary covariance), which can always be represented as

$$X = \sum_{i \in [m]} g_i A_i$$

where g_1, \dots, g_m are i.i.d. standard gaussians and A_i are deterministic matrix coefficients. In this setting, the noncommutative Khintchine (NCK) inequality of Lust-Piquard and Pisier [Pis03, §9.8] provides explicitly computable upper and lower bounds on the spectral norm $\|X\|$ that differ only by a logarithmic dimensional factor. This inequality has been extended to non-gaussian models that can be expressed as sums of independent random matrices [Tro15]. These results are used in numerous applications, including average case analysis of spectral methods, algorithms in the sum-of-squares hierarchy, and randomized linear algebra.

Matrix concentration inequalities are usually not sharp and often introduce mild but spurious dimensional factors in the analysis. In recent years, new kinds of inequalities have been developed that are applicable to the same models, but eliminate these dimensional factors and give rise to sharp bounds in many applications [BBvH23, BvH24, BCSv24]. This is achieved by introducing additional parameters that quantify the degree to which random matrices behave like idealized models from free probability theory. We will refer to these inequalities as *strong matrix concentration inequalities*, to distinguish them from their classical counterparts.

While matrix concentration inequalities are extremely versatile, there are large classes of models that cannot be readily understood with these tools. One such class, which we call *matrix chaos*, are matrices whose entries are *polynomials* of independent random variables. For example, in the gaussian case, we will consider random matrices X whose entries are homogeneous square-free polynomials of independent gaussian variables g_1, \dots, g_m :

$$X = \sum_{\substack{i_1, \dots, i_q \in [m] \\ i_1, \dots, i_q \text{ distinct}}} g_{i_1} \cdots g_{i_q} A_{i_1, \dots, i_q}.$$

Here q is the order of the chaos and A_{i_1, \dots, i_q} are deterministic matrix coefficients. Such models and their non-gaussian counterparts appear in many applications; a prominent example are the graph matrices of Potechin et al. [MPW15, AMP16].

The aim of this paper is to develop general *matrix chaos inequalities* that enable the treatment of matrix chaos models in a systematic manner, that are easily applicable in concrete situations, and that give rise to sharp bounds in a variety of applications.

Contributions and prior work. It was understood long ago in operator theory that when linear inequalities of NCK type are available, these can be iterated by means of a systematic procedure to obtain chaos inequalities; see [HP93] and [Pis03, Remark 9.8.9]. However, the resulting inequalities were not fully spelled out, and their significance to applications does not appear to have been realized. Consequently, many (special cases of) inequalities of this kind were repeatedly re-discovered in applied mathematics; see, for example, [Rau09, Theorem 4.3], [MSS16, Theorem 6.8], [MW19], [DNY20, §4.4], [RT23, Theorem 6.7], [FM24], and [TW24].¹ Furthermore, special matrix chaos models, such as graph matrices [MPW15, AMP16], have often been investigated using ad-hoc methods without the benefit of generally applicable tools.

¹Among these references, the closest in spirit to the approach developed here is the recent work [TW24], which appeared on arxiv after the present paper was submitted for publication.

In this paper, we revisit the original operator-theoretic approach for deriving matrix chaos inequalities from their linear counterparts. Besides drawing attention to this simple and natural method, it enables us to achieve a significantly improved toolbox for the study of matrix chaoses that arise in applications. The main contributions of this paper are twofold:

(i) We will show in section 2 that the operator-theoretic approach can be adapted to apply not only to NCK-type inequalities, but also to the recent theory of strong matrix concentration inequalities. This gives rise to *strong matrix chaos inequalities* that yield bounds without spurious dimensional factors in various settings. By using the linear inequalities as a black box, these inequalities leverage sophisticated tools of random matrix theory and free probability to obtain general bounds that would be difficult to achieve using ad-hoc methods.

(ii) The basic construction that underpins the operator-theoretic approach will naturally lead us in section 3 to the consideration of a special class of models that we call *chaos of combinatorial type*, for which all the parameters that appear in our bounds can be computed explicitly by a simple rule. Many matrix chaoses that we have encountered in theoretical computer science applications turn out to be special cases of this class. When that is the case, our methods reduce the study of such models to a nearly trivial computation that often yields improved bounds.

It is worth emphasizing that our inequalities provide both upper and lower bounds on the spectral norm of matrix chaoses, which suffice to show in most cases that our bounds are optimal either up to a universal constant or a logarithmic dimensional factor.

In section 4, we will illustrate our main results in the context of two notable examples of chaos of combinatorial type: *graph matrices* [MP16, AMP16], which are ubiquitous in the average case analysis of sum-of-squares algorithms [MPW15, BHK⁺16, PR20]; and *Khatri-Rao matrices* [KR68], which have been used in the context of differential privacy [KRSU10, De12]. Beside providing a unified analysis of the spectrum of these matrices, our techniques yield, in several applications, inequalities without spurious dimensional factors. We will illustrate the latter in the context of the ellipsoid fitting problem (recovering, with a simplified argument, the sharper analysis of graph matrices in [HKPX23]); and in a matrix chaos arising in the analysis of a sum-of-squares algorithm for tensor PCA (resulting in a correct-up-to-constants algorithmic guarantee).

In this paper, we focus for simplicity on bounding the spectral norm of homogeneous and square-free matrix chaoses. Our results can be extended to treat more general models, as well as more general spectral statistics such as the smallest singular value. We defer such extensions and other applications of our techniques to a longer companion manuscript [BLNv].

Notation. The following notations will be used throughout this paper.

We write $x \lesssim_q y$ if $x \leq C_q y$ for a universal constant C_q that depends only on q . When $x \lesssim_q y$ and $y \lesssim_q x$, we write $x \asymp_q y$. We use $x \lesssim y$ and $x \asymp y$ when the constants are universal. We denote by $a : b = \{a, a + 1, \dots, b\}$, by $[n] = 1 : n$, and by $|I|$ the cardinality of a finite set I .

We will always work with real matrices for simplicity (the results of this paper extend readily to complex matrices). The entries of a matrix M will be denoted $M[i, j]$ or $M_{i,j}$, its adjoint is denoted M^\top , and its operator norm is denoted $\|M\|$. For a scalar random variable h , we denote by $\|h\|_{\psi_2}$ its subgaussian constant and by $\|h\|_{L^p} = (\mathbb{E}|h|^p)^{1/p}$ its L^p -norm.

2. MATRIX CHAOS INEQUALITIES

The aim of this section is to formulate the main inequalities of this paper. In section 2.1, we first introduce the general matrix chaos model and its decoupled version. In section 2.2, we introduce the basic notation for tensor flattenings that will be used throughout this paper. The main inequalities are stated in section 2.3. Finally, we will outline the basic approach to the proofs in section 2.4. Most of the proof details will be deferred to Appendix A.

2.1. Matrix chaos and decoupling. The basic model of this paper is a *matrix chaos*

$$X = \sum_{\substack{i_1, \dots, i_q \in [m] \\ i_1, \dots, i_q \text{ distinct}}} h_{i_1} \cdots h_{i_q} A_{i_1, \dots, i_q} \quad (1)$$

of order q . Here h_1, \dots, h_m are i.i.d. copies of a random variable h with zero mean, and A_{i_1, \dots, i_q} are deterministic $d_1 \times d_2$ matrix coefficients (we will write $d = d_1 \vee d_2$).

We will often consider a decoupled variant of the above model. To this end, let $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(q)}$ denote i.i.d. copies of $\mathbf{h} := (h_1, \dots, h_m)$. We define the *decoupled matrix chaos* as

$$Y = \sum_{i_1, \dots, i_q \in [m]} h_{i_1}^{(1)} \cdots h_{i_q}^{(q)} A_{i_1, \dots, i_q}. \quad (2)$$

Note that in the decoupled case, the coordinates i_1, \dots, i_q need not be distinct.

The connection between coupled and decoupled chaoses is captured by classical decoupling inequalities. In the present setting, [dlPG12, Theorem 3.1.1] yields the following.

Theorem 2.1 (Decoupling inequalities). *Let X be any matrix chaos as in (1), and let Y be the decoupled matrix chaos as in (2) defined by the same random variables h_1, \dots, h_m and matrix coefficients A_{i_1, \dots, i_q} (where we set $A_{i_1, \dots, i_q} = 0$ when i_1, \dots, i_q are not distinct). Then we have*

$$\mathbb{E} \|X\| \lesssim_q \mathbb{E} \|Y\|.$$

Moreover, this inequality can be reversed

$$\mathbb{E} \|Y\| \lesssim_q \mathbb{E} \|X\|$$

provided that the matrix coefficients are assumed to be symmetric in the sense that $A_{i_1, \dots, i_q} = A_{i_{\pi(1)}, \dots, i_{\pi(q)}}$ for every permutation π of $[q]$.

The iteration argument that forms the basis for our proofs will rely crucially on the independence structure of the decoupled model. As decoupled chaoses arise in applications in their own right, we will formulate our main inequalities for decoupled chaoses (2) and take for granted in the sequel that these inequalities can also be applied for coupled chaoses (1) by virtue of Theorem 2.1.

Remark 2.2 (Lower bounding $\mathbb{E} \|X\|$). *The lower bound in Theorem 2.1 has an additional assumption that the matrix coefficients are symmetric. This assumption is necessary: consider, for example, the chaos $Y = g_1^{(1)} g_2^{(2)} - g_2^{(1)} g_1^{(2)}$ whose coupled version vanishes $X = 0$. On the other hand, as the matrix coefficients in (1) may clearly be chosen to be symmetric without loss of generality, this does not present any fundamental restriction to obtaining lower bounds on $\mathbb{E} \|X\|$.*

Remark 2.3 (More general chaoses). *In this paper, we work only with homogeneous square-free chaoses (1). However, non-homogeneous or non-square-free matrix chaoses can often be treated using similar methods. Such models are addressed in the longer companion manuscript [BLNv].*

2.2. Flattenings. Fix a decoupled matrix chaos as in (2). It will be convenient to view the matrix coefficients A_{i_1, \dots, i_q} of the chaos as defining a tensor \mathcal{A} of order $q + 2$ by

$$\mathcal{A}_{i_1, \dots, i_q, i_{q+1}, i_{q+2}} := (A_{i_1, \dots, i_q})_{i_{q+1}, i_{q+2}}.$$

Here the first q coordinates (which we call *chaos coordinates*) range from 1 to m and the last two (which we call *matrix coordinates*) range from 1 to d_1 and 1 to d_2 , respectively.

The main inequalities of this paper will be defined in terms of the norms of *flattenings* of the tensor \mathcal{A} that are defined as follows. Denote by e_i the i th element of the standard coordinate basis, viewed as a column vector. Then for any subsets $R, C \subseteq [q + 2]$, we define the matrix

$$\mathcal{A}_{[R|C]} := \sum_{\substack{i_1, \dots, i_q \in [m] \\ i_{q+1} \in [d_1], i_{q+2} \in [d_2]}} \left(\bigotimes_{t \in R} e_{i_t} \right) \otimes \left(\bigotimes_{t \in C} e_{i_t}^\top \right) \mathcal{A}_{i_1, \dots, i_{q+2}}. \quad (3)$$

This definition is easiest to interpret when $R = [q+2] \setminus C$: in this case, $\mathcal{A}_{[R|C]}$ is the matrix whose rows are indexed by the coordinates in the row set R , whose columns are indexed by the coordinates in the column set C , and whose entries are the corresponding entries of \mathcal{A} . For example, if $q = 2$ and $R = \{1, 3\}$, $C = \{2, 4\}$, then the associated flattening $\mathcal{A}_{[R|C]}$ is the $md_1 \times md_2$ matrix with entries $(\mathcal{A}_{[R|C]})_{(i_1, i_3), (i_2, i_4)} = \mathcal{A}_{i_1, i_2, i_3, i_4}$. However, we will also encounter flattenings where the same coordinate may appear simultaneously in R and C , which corresponds to diagonalization. For example, if $q = 1$ and $R = \{1, 2\}$, $C = \{1, 3\}$, then $(\mathcal{A}_{[R|C]})_{(i_1, i_2), (i'_1, i_3)} = \mathbb{1}_{i_1=i'_1} \mathcal{A}_{i_1, i_2, i_3}$.

2.3. The main inequalities. We now formulate the main inequalities of this paper, which bound the spectral norm of a matrix chaos in terms of the spectral norms of flattenings of the coefficient tensor \mathcal{A} . All these inequalities will be derived by an iteration argument from an underlying matrix concentration inequality for linear random matrices. The basic idea behind the iteration method will be explained in section 2.4. We postpone the detailed proofs to Appendix A.

Our main inequalities will be stated for decoupled chaoses as in (2). Their extension to coupled chaoses as in (1) is immediate by Theorem 2.1. We focus for simplicity on expectation bounds; tail bounds can then be deduced using concentration tools (e.g., [ALM21] or as in [Pis14, Lemma 7.6]).

2.3.1. Iterated NCK. The simplest matrix concentration inequality for linear random matrices is the noncommutative Khintchine inequality (NCK) [Pis03, §9.8], see Theorem A.3 in the appendix. We begin by formulating the corresponding matrix chaos inequality.

To this end, we now introduce the basic parameter that controls the leading order behavior of matrix chaoses. A flattening is said to be a **σ -flattening** if $R = [q+2] \setminus C$ and if the original matrix coordinates are kept as row and column coordinates, that is, $q+1 \in R$ and $q+2 \in C$. In this case, $\mathcal{A}_{[R|C]}$ is an $m^{|R|-1}d_1 \times m^{|C|-1}d_2$ matrix. We now define

$$\sigma(\mathcal{A}) := \max_{\substack{R=[q+2] \setminus C \\ q+1 \in R, q+2 \in C}} \left\| \mathcal{A}_{[R|C]} \right\|, \quad (4)$$

that is, $\sigma(\mathcal{A})$ is the largest spectral norm of all σ -flattenings. We can now formulate the iterated NCK inequality, whose proof will be given in Appendix A.2.

Theorem 2.4 (Iterated NCK). *Let Y be a decoupled chaos as in (2). Then*

$$\|h\|_{L^1}^q \sigma(\mathcal{A}) \lesssim_q \mathbb{E} \|Y\| \lesssim_q \|h\|_{\psi_2}^q \log(d+m)^{\frac{q}{2}} \sigma(\mathcal{A}).$$

Alternatively, the upper bound remains valid if $\|h\|_{\psi_2}$ is replaced by $\|h\|_{L^{\log m}}$.

Note, for example, that $\|h\|_{L^1}, \|h\|_{\psi_2} \asymp 1$ if h is a standard Gaussian or Rademacher variable. When this is the case, Theorem 2.4 states that the parameter $\sigma(\mathcal{A})$ captures the spectral norm of any matrix chaos up to a logarithmic dimensional factor.

2.3.2. Iterated strong NCK. The drawback of Theorem 2.4 is that the dimensional factor in the upper bound often proves to be suboptimal. For subgaussian random matrices, sharp bounds can often be achieved by using instead the strong NCK inequality of [BBvH23], see Theorem A.5 in the appendix. We now formulate a corresponding matrix chaos inequality.

A flattening is said to be a **v -flattening** if $R = [q+2] \setminus C$ is nonempty and if the original matrix coordinates are both assigned to be column coordinates, that is, $q+1 \in C$ and $q+2 \in C$. Define

$$v(\mathcal{A}) := \max_{\substack{R=[q+2] \setminus C \\ q+1, q+2 \in C \\ R \neq \emptyset}} \left\| \mathcal{A}_{[R|C]} \right\|, \quad (5)$$

that is, $v(\mathcal{A})$ is the largest spectral norm of all v -flattenings. We can now formulate the iterated strong NCK inequality, whose proof will be given in Appendix A.3.

Theorem 2.5 (Iterated strong NCK). *Let Y be a decoupled chaos as in (2). Then*

$$\mathbb{E} \|Y\| \lesssim_q \|h\|_{\psi_2}^q \left(\sigma(\mathcal{A}) + \log(d+m)^{\frac{q+2}{2}} v(\mathcal{A}) \right).$$

The corresponding lower bound on $\mathbb{E} \|Y\|$ follows already from Theorem 2.4. The significance of Theorem 2.5 is that when $v(\mathcal{A}) \ll \sigma(\mathcal{A})$, the logarithmic factor in Theorem 2.4 is eliminated.

2.3.3. Iterated matrix Rosenthal. The above results yield matching upper and lower bounds for matrix chaoses that are based on regularly behaved random variables h_1, \dots, h_m , such as Gaussians or Rademachers. However, they may result in poor bounds in situations where $\|h\|_{\psi_2}$ is very large or $\|h\|_{L^1}$ is very small. A typical situation of this kind that arises frequently in practice is in the study of sparse models, where h is a standardized (i.e., normalized to have zero mean and unit variance) Bern(p) random variable. In this case, it is readily verified that $\|h\|_{\psi_2} \rightarrow \infty$ and $\|h\|_{L^1} \rightarrow 0$ in the sparse regime $p \rightarrow 0$, which causes the previous bounds to diverge.

For linear random matrices, this issue can be surmounted by using inequalities of Rosenthal type [JZ13, MJC⁺14, BvH24], see Theorem A.6 in the appendix. We now formulate a corresponding matrix chaos inequality. The strong form will be given in the next section.

A flattening is said to be an **r -flattening** if the original matrix coordinates are kept as row and column coordinates, that is, $q+1 \in R$ and $q+2 \in C$, but there is at least one of the q chaos coordinates that appears both in R and C . We now define

$$r(\mathcal{A}) := \max_{\substack{R \cup C = [q+2] \\ q+1 \in R, q+2 \in C \\ \emptyset \neq R \cap C \subseteq [q]}} \left\| \mathcal{A}_{[R|C]} \right\|, \quad (6)$$

that is, $r(\mathcal{A})$ as the largest spectral norm of all r -flattening. We can now formulate an iterated matrix Rosenthal inequality, whose proof is given in Appendix A.4.

Theorem 2.6 (Iterated matrix Rosenthal). *Let Y be a decoupled chaos as in (2). Assume that h has unit variance, and define the parameter $\alpha(h) = \|h\|_{L^{\log(d+m)}}$. Then we have*

$$\sigma(\mathcal{A}) - C_q \alpha(h)^q \log(d+m)^{\frac{q}{2}} r(\mathcal{A}) \lesssim_q \mathbb{E} \|Y\| \lesssim_q \log(d+m)^{\frac{q}{2}} \sigma(\mathcal{A}) + \alpha(h)^q \log(d+m)^{\frac{q+1}{2}} r(\mathcal{A}),$$

where C_q is a constant that depends only on q .

This result may be viewed as an analogue of the iterated NCK inequality (Theorem 2.4) where the distributional parameter $\alpha(h)$ only appears in the second-order term that is controlled by $r(\mathcal{A})$. Therefore, when $r(\mathcal{A}) \ll \sigma(\mathcal{A})$, the parameter $\sigma(\mathcal{A})$ captures the spectral norm up to a logarithmic factor even in (e.g., sparse) situations where $\alpha(h)$ may diverge.

2.3.4. Iterated strong matrix Rosenthal. Just as the strong NCK inequality eliminates the dimensional factor in the NCK inequality in many situations, there is an analogous strong form of the matrix Rosenthal inequality [BvH24], see Theorem A.7. We now formulate a corresponding matrix chaos inequality, whose proof will be given in Appendix A.5.

Theorem 2.7 (Iterated strong Matrix Rosenthal). *Let Y be a decoupled chaos as in (2). Assume that h has unit variance, and define the parameter $\alpha(h) = \|h\|_{L^{\log(d+m)}}$. Then*

$$\mathbb{E} \|Y\| \lesssim_q \sigma(\mathcal{A}) + \alpha(h)^q \log(d+m)^{\frac{q+3}{2}} v(\mathcal{A}).$$

Let us note that the inequality $r(\mathcal{A}) \leq v(\mathcal{A})$ always holds (Lemma A.2), so that $r(\mathcal{A})$ need not be computed when applying an inequality in which $v(\mathcal{A})$ already appears. In particular, the lower bound corresponding to Theorem 2.7 follows from Theorem 2.6. The significance of Theorem 2.7 is that when $\alpha(h)^q v(\mathcal{A}) \ll \sigma(\mathcal{A})$, we obtain $\mathbb{E} \|Y\| \asymp_q \sigma(\mathcal{A})$ without a logarithmic factor.

2.4. Iteration approach. We now outline the basic iteration approach to the proofs of our matrix chaos inequalities. For NCK this approach dates back at least to [HP93], and we will explain how it can be adapted to capture the strong inequalities. We defer detailed proofs to Appendix A.

2.4.1. The linear case. The key observation behind the proofs is that the linear matrix concentration inequalities can be reinterpreted in terms of flattenings. Once they have been reformulated in this manner, the chaos inequalities will follow seamlessly by induction.

Let us begin by illustrating the case of NCK. Consider a matrix chaos Y as in (2) of order $q = 1$, that is, $Y = \sum_{i \in [m]} h_i A_i$. The NCK inequality (Theorem A.3) states that

$$\mathbb{E}\|Y\| \lesssim \|h\|_{\psi_2} \log(d)^{\frac{1}{2}} (\sigma_R(Y) \vee \sigma_C(Y)) \quad (7)$$

with

$$\sigma_R(Y) := \left\| \sum_{i \in [m]} A_i^\top A_i \right\|^{\frac{1}{2}}, \quad \sigma_C(Y) := \left\| \sum_{i \in [m]} A_i A_i^\top \right\|^{\frac{1}{2}}.$$

To reformulate this inequality in terms of flattenings, note that

$$\mathcal{A}_{[1,2|3]} = \sum_{i \in [m]} e_i \otimes A_i = \begin{pmatrix} A_1 \\ \vdots \\ A_m \end{pmatrix}, \quad \mathcal{A}_{[2|1,3]} = \sum_{i \in [m]} e_i^\top \otimes A_i = (A_1 \ \cdots \ A_m).$$

As

$$\mathcal{A}_{[1,2|3]}^\top \mathcal{A}_{[1,2|3]} = \sum_{i \in [m]} A_i^\top A_i, \quad \mathcal{A}_{[2|1,3]} \mathcal{A}_{[2|1,3]}^\top = \sum_{i \in [m]} A_i A_i^\top,$$

we have clearly shown that $\sigma_R(Y) = \|\mathcal{A}_{[1,2|3]}\|$ and $\sigma_C(Y) = \|\mathcal{A}_{[2|1,3]}\|$. Note that in this notation, the NCK inequality (7) is essentially recovered as the $q = 1$ case of Theorem 2.4.

The strong NCK, Rosenthal, and strong Rosenthal inequalities (Theorems A.5, A.6, and A.7) involve two additional matrix parameters

$$v(Y) := \|\text{Cov}(Y)\|^{\frac{1}{2}}, \quad r(Y) := \max_{i \in [m]} \|A_i\|,$$

where $\text{Cov}(Y)$ denotes the covariance matrix of the entries of Y . We now observe that these parameters can also be reformulated in terms of flattenings. To this end, note that

$$\begin{aligned} \mathcal{A}_{[1|2,3]} &= \sum_{i \in [m]} e_i \otimes \text{vec}(A_i)^\top = \begin{pmatrix} \text{vec}(A_1)^\top \\ \vdots \\ \text{vec}(A_m)^\top \end{pmatrix}, \\ \mathcal{A}_{[1,2|1,3]} &= \sum_{i \in [m]} e_i \otimes e_i^\top \otimes A_i = \begin{pmatrix} A_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & A_m \end{pmatrix}, \end{aligned}$$

where $\text{vec}(\cdot)$ denotes the operation that arranges all the entries of a matrix in a column vector. As $\mathcal{A}_{[1|2,3]}^\top \mathcal{A}_{[1|2,3]} = \text{Cov}(Y)$, it follows directly that $v(Y) = \|\mathcal{A}_{[1|2,3]}\|$. The operator norm of a block-diagonal matrix equals the maximum operator norm of its blocks, hence $r(Y) = \|\mathcal{A}_{[1,2|1,3]}\|$. In this notation, the strong NCK and (strong) Rosenthal inequalities are again essentially recovered as the $q = 1$ case of the corresponding matrix chaos inequalities.




FIGURE 1. Intermediate flattenings that arise from each matrix parameter. Here \mathcal{B} is the (random) tensor of order 3 associated to the linear chaos $\sum_{i_1} h_{i_1} B_{i_1}$ in (8).

2.4.2. *Iteration.* Now let Y be a decoupled chaos as in (2) of order $q \geq 2$. If we condition on the random variables associated with the first $q - 1$ chaos coordinates (the vectors $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(q-1)}$), then Y can be written as a linear chaos with random coefficients:

$$Y = \sum_{i_q} h_{i_q}^{(q)} \left(\sum_{i_1, \dots, i_{q-1}} h_{i_1}^{(1)} \dots h_{i_{q-1}}^{(q-1)} A_{i_1, \dots, i_q} \right) = \sum_{i_q} h_{i_q}^{(q)} B_{i_q}. \quad (8)$$

Applying the linear inequalities to the random matrix $Y = \sum_{i_q} h_{i_q}^{(q)} B_{i_q}$ yields upper bounds in terms of four possible flattenings, each of which can itself be interpreted as a matrix chaos of order $q - 1$. For example, the $\sigma_R(Y)$ parameter of this random matrix is the norm of

$$\begin{aligned} \sum_{i_q} e_{i_q} \otimes B_{i_q} &= \sum_{i_q} e_{i_q} \otimes \left(\sum_{i_1, \dots, i_{q-1}} h_{i_1}^{(1)} \dots h_{i_{q-1}}^{(q-1)} A_{i_1, \dots, i_q} \right) \\ &= \sum_{i_q} e_{i_q} \otimes \left(\sum_{i_1, \dots, i_{q-1}} h_{i_1}^{(1)} \dots h_{i_{q-1}}^{(q-1)} \left(\sum_{i_{q+1}, i_{q+2}} e_{i_{q+1}} \otimes e_{i_{q+2}}^\top \mathcal{A}_{i_1, \dots, i_{q+2}} \right) \right) \\ &= \sum_{i_1, \dots, i_{q-1}} h_{i_1}^{(1)} \dots h_{i_{q-1}}^{(q-1)} \left(\sum_{i_q, i_{q+1}, i_{q+2}} (e_{i_q} \otimes e_{i_{q+1}}) \otimes e_{i_{q+2}}^\top \mathcal{A}_{i_1, \dots, i_{q+2}} \right), \end{aligned} \quad (9)$$

which is a decoupled matrix chaos of order $q - 1$ with matrix coefficients of dimension $md_1 \times d_2$. Analogous expressions hold for the remaining matrix parameters. In this manner, the expected norm of a matrix chaos of order q is bounded by the expected norms of matrix chaoses of order $q - 1$, and the proofs can proceed by induction on q .

To formalize the above procedure, we introduce the following notation. Given $Z, R, C \subseteq [q + 2]$ with $Z \subseteq [q]$, we define the **intermediate flattening**

$$Y_{[Z|R|C]} := \sum_{i_1, \dots, i_{q+2}} \left(\bigotimes_{t \in R} e_{i_t} \right) \otimes \left(\bigotimes_{t \in C} e_{i_t}^\top \right) \left(\prod_{t \in Z} h_{i_t}^{(t)} \right) \mathcal{A}_{i_1, \dots, i_{q+2}}, \quad (10)$$

Note that $Y_{[Z|R|C]}$ (with $Z \neq \emptyset$) is a decoupled matrix chaos of order $|Z|$. We will denote by $\mathcal{A}_{[Z|R|C]}$ the tensor of order $|Z| + 2$ associated with the chaos $Y_{[Z|R|C]}$. The intermediate flattening in (9) corresponds precisely to $Y_{[1:q-1|q, q+1|q+2]}$.

Using this notation, applying the linear matrix concentration inequalities to the original chaos $Y = Y_{[1:q-1, q | q+1 | q+2]}$ of order q yields bounds in terms of four intermediate chaoses of order $q - 1$ as described in Figure 1. To prove our main inequalities, we can iterate this procedure until the order of the (intermediate) flattenings has been reduced to zero. The resulting final flattenings $Y_{[\emptyset|R|C]}$ are deterministic and are equal to both $\mathcal{A}_{[\emptyset|R|C]}$ and $\mathcal{A}_{[R|C]}$. In practice, this procedure is easily implemented by induction on q . The details are deferred to Appendix A.

3. CHAOS OF COMBINATORIAL TYPE

While the matrix chaos inequalities of the previous section can capture a large class of models, their application may appear daunting due to the large number of flattenings that must be controlled. However, the construction of these flattenings in section 2.2 by means of tensor products of canonical basis vectors e_i and their transposes e_i^\top suggests that the norms of the flattenings should be especially easy to control if the matrix coefficients A_{i_1, \dots, i_q} can themselves be expressed as tensor products of e_i and e_i^\top , resulting in $\{0, 1\}$ -matrices with many symmetries.

This observation naturally leads us in this section to define a special class of *matrix chaoses of combinatorial type*, for which the parameters in all our matrix chaos inequalities can be computed mechanically by a simple rule. Remarkably, it turns out that many matrix chaoses that arise in theoretical computer science applications are of this special form: two important examples are graph matrices [MPW15, AMP16] and Khatri-Rao matrices [KR68, KRSU10, De12, Rud11]. Whenever this structure is present, our methods will reduce the study of such models to a nearly trivial computation. As will be illustrated in section 4, this enables us to achieve the best known results in the literature for several applications using a unified and remarkably simple analysis.

3.1. Definition and guiding example. In order to motivate the general definition of chaos of combinatorial type, we begin by introducing the guiding example of Khatri-Rao matrices. These are random matrices with dependent entries whose study dates back to at least the 1960s [KR68], and have more recently been used in the context of differential privacy [KRSU10, De12].

Example 3.1 (Khatri-Rao matrices). *We begin by stating the definition.*

Definition 3.2. Let q, n, d be positive integers, let h be a scalar random variable with zero mean and unit variance, and let $W^{(1)}, \dots, W^{(q)}$ be $d \times n$ random matrices whose entries are i.i.d. copies of h . The *Khatri-Rao matrix* Y is the $d^q \times n$ random matrix obtained by taking the column-wise Kronecker product of $W^{(1)}, \dots, W^{(q)}$, that is, the matrix whose entries are defined by

$$Y[(j_1, \dots, j_q), k] = \prod_{t=1}^q W^{(t)}[j_t, k], \quad (11)$$

for any $j_1, \dots, j_q \in [d]$ and $k \in [n]$.

The Khatri-Rao matrix (11) can be equivalently expressed as a decoupled matrix chaos

$$Y = \sum_{j_1, \dots, j_q \in [d], k \in [n]} \left(\prod_{t=1}^q W^{(t)}[j_t, k] \right) e_{j_1} \otimes \dots \otimes e_{j_q} \otimes e_k^\top \quad (12)$$

with the special property that every matrix coefficient $A_{(j_1, k), \dots, (j_q, k)} = e_{j_1} \otimes \dots \otimes e_{j_q} \otimes e_k^\top$ is a tensor product of coordinate basis vectors and their transposes.²

A characteristic feature of the above example is that even though each matrix coefficient is a tensor product of coordinate basis vectors and their adjoints, the indices of these coordinate vectors may simultaneously appear in the coordinates corresponding to distinct random vectors $\mathbf{h}^{(t)}$ in the definition (2) of a decoupled matrix chaos. In this example, the coordinate basis vectors that define the matrix coefficients are indexed by j_1, \dots, j_q, k , which we call **summation indices**. Each random vector $\mathbf{h}^{(t)}$ is indexed by an ordered subset (j_t, k) of the summation indices, which we call **chaos coordinates**. Finally, the entries of the matrix coefficients are also indexed by ordered subsets (j_1, \dots, j_q) and k of the summation indices, which we call **matrix coordinates**.

The above structure is generalized by the notion of matrix chaos of combinatorial type.

²Formally speaking, the definition (2) of a matrix chaos requires us to assign independent indices to each coordinate of $A_{(j_1, k_1), \dots, (j_q, k_q)}$. In order to capture (11), we would then set $A_{(j_1, k_1), \dots, (j_q, k_q)} = 0$ except when $k_1 = \dots = k_q = k$. To lighten the notation, however, we will generally drop these zero coefficients from the summation as in (12).

Definition 3.3 (Matrix Chaos of Combinatorial type). Let h be a scalar random variable with zero mean, let q, p, S_1, \dots, S_p be positive integers, and let I_1, \dots, I_{q+2} be ordered subsets of $[p]$. A *matrix chaos of combinatorial type* is defined by

$$Y = \sum_{\mathbf{s} \in [S_1] \times \dots \times [S_p]} h_{I_1(\mathbf{s})}^{(1)} \cdots h_{I_q(\mathbf{s})}^{(q)} e_{I_{q+1}(\mathbf{s})} \otimes e_{I_{q+2}(\mathbf{s})}^\top, \quad (13)$$

where for $I = (i_1, \dots, i_k)$ and $\mathbf{s} = (s_1, \dots, s_p)$ we define $I(\mathbf{s}) := (s_{i_1}, \dots, s_{i_k})$,

$$e_{(s_{i_1}, \dots, s_{i_k})} := e_{s_{i_1}} \otimes e_{s_{i_2}} \otimes \cdots \otimes e_{s_{i_k}},$$

and $h_{(s_{i_1}, \dots, s_{i_k})}^{(t)}$ are independent copies of h . Here s_1, \dots, s_p are *summation indices*; $I_1(\mathbf{s}), \dots, I_q(\mathbf{s})$ are *chaos coordinates*; and $I_{q+1}(\mathbf{s}), I_{q+2}(\mathbf{s})$ are *matrix coordinates*.

Further examples of chaos of combinatorial type will be treated in section 4.

3.2. How to compute norms of flattenings. The aim of this section is to develop a user-friendly procedure to compute the norms of flattenings of chaoses of combinatorial type (Algorithm 3.5).

3.2.1. The Khatri-Rao example as a warm-up. We again use the guiding example of Khatri-Rao matrices to illustrate the procedure. We focus on the case $q = 2$ for simplicity.

Let Y be a Khatri-Rao matrix as in (12) with $q = 2$. In the notation of Definition 3.3 we have $q = 2$ and $p = 3$; the summation indices are $j_1, j_2 \in [d]$ and $k \in [n]$; the chaos coordinates are given by $I_1(j_1, j_2, k) = (j_1, k)$ and $I_2(j_1, j_2, k) = (j_2, k)$, and the matrix coordinates are given by $I_3(j_1, j_2, k) = (j_1, j_2)$ and $I_4(j_1, j_2, k) = k$; and $h_{(j_t, k)}^{(t)} = W^{(t)}[j_t, k]$. We can therefore write

$$Y = \sum_{j_1, j_2 \in [d], k \in [n]} h_{(j_1, k)}^{(1)} h_{(j_2, k)}^{(2)} e_{(j_1, j_2)} \otimes e_k^\top.$$

For the sake of exposition, let us focus on the flattening $\mathcal{A}_{[1,2,3|4]}$. This is the σ -flattening where both chaos coordinates are in the row set R (see (3)). It is given by

$$\begin{aligned} \mathcal{A}_{[1,2,3|4]} &= \sum_{j_1, j_2 \in [d], k \in [n]} e_{(j_1, k)} \otimes e_{(j_2, k)} \otimes e_{(j_1, j_2)} \otimes e_k^\top \\ &= \sum_{j_1, j_2 \in [d], k \in [n]} e_{j_1} \otimes e_k \otimes e_{j_2} \otimes e_k \otimes e_{j_1} \otimes e_{j_2} \otimes e_k^\top, \end{aligned}$$

where in the last line we used $e_{(j_1, k)} = e_{j_1} \otimes e_k$ (and similarly for other coordinates).

By permuting the order of tensor products (which corresponds to reordering rows and columns, and so preserves all singular values of the matrix), we obtain

$$\mathcal{A}_{[1,2,3|4]} \simeq \sum_{j_1, j_2 \in [d], k \in [n]} e_{j_1} \otimes e_{j_1} \otimes e_{j_2} \otimes e_{j_2} \otimes e_k \otimes e_k \otimes e_k^\top \quad (14)$$

$$= \left(\sum_{j_1 \in [d]} e_{j_1} \otimes e_{j_1} \right) \otimes \left(\sum_{j_2 \in [d]} e_{j_2} \otimes e_{j_2} \right) \otimes \left(\sum_{k \in [n]} e_k \otimes e_k \otimes e_k^\top \right), \quad (15)$$

where $A \simeq B$ means that the two matrices are related by a unitary change of basis. We thus have

$$\|\mathcal{A}_{[1,2,3|4]}\| = \underbrace{\left\| \sum_{j_1 \in [d]} e_{j_1} \otimes e_{j_1} \right\|}_{=\sqrt{d}} \underbrace{\left\| \sum_{j_2 \in [d]} e_{j_2} \otimes e_{j_2} \right\|}_{=\sqrt{d}} \underbrace{\left\| \sum_{k \in [n]} e_k \otimes e_k \otimes e_k^\top \right\|}_{=1} = d,$$

where we used that $\{e_k \otimes e_k\}$ are orthonormal vectors and therefore, by a unitary change of basis and restriction to a subspace, $\sum_{j \in [d]} e_j \otimes e_j$ and $\sum_{k \in [n]} e_k \otimes e_k \otimes e_k^\top$ may be viewed as a d -dimensional vector of ones and an n -dimensional identity matrix, respectively.

Repeating this procedure for the other 11 flattenings (see Table 1 below), we readily obtain

$$\sigma(\mathcal{A}) = \max \left\{ d, n^{\frac{1}{2}} \right\}, \quad v(\mathcal{A}) = r(\mathcal{A}) = d^{\frac{1}{2}}. \quad (16)$$

Since $\sigma(\mathcal{A})$ dominates both $v(\mathcal{A})$ and $r(\mathcal{A})$, Theorems 2.6 and 2.7 imply that

$$\mathbb{E} \|Y\| \asymp_q \max \left\{ d, n^{\frac{1}{2}} \right\} \quad (17)$$

for $d, n \rightarrow \infty$ at any relative speed, provided that $\alpha(h)$ is sub-polynomial in d, n .

3.2.2. The general case. Analogously to (14), any flattening of a general chaos of combinatorial type can be written (after reordering the tensor products) as

$$\mathcal{A}_{[R|C]} \simeq \bigotimes_{u=1}^p \sum_{s_u \in [S_u]} \left(\underbrace{e_{s_u} \otimes \cdots \otimes e_{s_u}}_{\mu_u \text{ factors}} \otimes \underbrace{e_{s_u}^\top \otimes \cdots \otimes e_{s_u}^\top}_{\nu_u \text{ factors}} \right),$$

where μ_u and ν_u are non-negative integers. By convention, if $\mu_u = \nu_u = 0$, the tensor product inside the brackets is to be interpreted as the scalar 1.

The calculation now proceeds by noting, as in (15), that

$$\left\| \sum_{s_u \in [S_u]} \left(\underbrace{e_{s_u} \otimes \cdots \otimes e_{s_u}}_{\mu_u \text{ factors}} \otimes \underbrace{e_{s_u}^\top \otimes \cdots \otimes e_{s_u}^\top}_{\nu_u \text{ factors}} \right) \right\| = \begin{cases} 1 & \text{if } \mu_u > 0 \text{ and } \nu_u > 0 \\ \sqrt{S_u} & \text{if } \mu_u > 0 \text{ xor } \nu_u > 0 \\ S_u & \text{if } \mu_u = 0 \text{ and } \nu_u = 0. \end{cases} \quad (18)$$

This can be conveniently summarized by defining by \mathcal{R} and \mathcal{C} the sets of summation indices that appear in R and C , respectively. This yields the following result, which we prove in section A.6.

Proposition 3.4. *Let Y be a chaos of combinatorial type as in (13) of order q with p summation indices. Let $R, C \subseteq [q+2]$ with $R \cup C = [q+2]$, and let $\mathcal{R} = \cup_{t \in R} I_t$ and $\mathcal{C} = \cup_{t \in C} I_t$. Then*

$$\left\| \mathcal{A}_{[R|C]} \right\|^2 = \left(\prod_{u \in \mathcal{R}^c} S_u \right) \left(\prod_{u \in \mathcal{C}^c} S_u \right). \quad (19)$$

This proposition yields a straightforward algorithm to compute the norms of flattenings of chaoses of combinatorial type: given a set of choices of whether each particular chaos or matrix coordinate is in R or C , the sets $\mathcal{R} = \cup_{t \in R} I_t$ and $\mathcal{C} = \cup_{t \in C} I_t$ determine which summation indices belong to row and/or column matrix coordinates, and the norm of the flattening is given by (19).

Algorithm 3.5. *Construct a table with the following data:*

- *The flattening type: σ , v or r ;*
- *For each flattening type, list all possible assignments of the chaos (I_1, \dots, I_q) and matrix (I_{q+1}, I_{q+2}) coordinates to R , C , or $R \cap C$.*
- *Next, for each summation index, list whether it appears in \mathcal{R} , \mathcal{C} , or $\mathcal{R} \cap \mathcal{C}$.*
- *Finally, $\left\| \mathcal{A}_{[R|C]} \right\|$ can be computed directly using the formula (19).³*

In Table 1, we illustrate the application of this algorithm to the $q = 2$ case of the Khatri-Rao matrix, recovering the manual computation of (16).

In practice, it is generally not necessary in applications to list every possible flattening, as one can directly analyze using (19) which flattenings will dominate in the matrix chaos inequalities. This will be illustrated in section 4.1, where we will analyze the Khatri-Rao model for all $q \geq 2$.

³In applications, it is often the case that every summation index appears in at least one of the (chaos or matrix) coordinates, so that $\mathcal{R}^c \cap \mathcal{C}^c = \emptyset$. In this case, the right-hand side of (19) is simply the product of the dimensions of all the summation indices that appear in coordinates assigned only to R or only to C .

type	coordinates				summation indices			norm ²
	j_1k	j_2k	j_1j_2	k	j_1	j_2	k	
σ	R	R	R	C	R	R	RC	d^2
	R	C	R	C	R	RC	RC	d
	C	R	R	C	RC	R	RC	d
	C	C	R	C	RC	RC	C	n
v	R	R	C	C	RC	RC	RC	1
	R	C	C	C	RC	C	RC	d
	C	R	C	C	C	RC	RC	d
r	R	RC	R	C	R	RC	RC	d
	C	RC	R	C	RC	RC	RC	1
	RC	R	R	C	RC	R	RC	d
	RC	C	R	C	RC	RC	RC	1
	RC	RC	R	C	RC	RC	RC	1

TABLE 1. Flattenings of Khatri-Rao matrices (Example 3.1) with $q = 2$ as produced by Algorithm 3.5. The σ , v , and r parameters are the maxima of the norms of the respective flattenings. The two dominant flattenings are shaded.

3.3. Chaos of nearly combinatorial type. It will be useful (see Sections 4.2, 4.3, and 4.4) to consider a slightly more general class of chaoses that include a weight function.

Definition 3.6 (Matrix Chaos of nearly Combinatorial type). Let $h, q, p, S_1, \dots, S_p, I_1, \dots, I_{q+2}$ be as in Definition 3.3, and $f: [S_1] \times \dots \times [S_p] \rightarrow \mathbb{R}$ be a weight function. A matrix chaos of nearly combinatorial type with weight function f is a chaos of the form:

$$Y^f = \sum_{\mathbf{s} \in [S_1] \times \dots \times [S_p]} f(\mathbf{s}) h_{I_1(\mathbf{s})}^{(1)} \cdots h_{I_q(\mathbf{s})}^{(q)} e_{I_{q+1}(\mathbf{s})} \otimes e_{I_{q+2}(\mathbf{s})}^\top. \quad (20)$$

By pointwise bounding $|f|$ by its maximum $\|f\|_\infty$, we can generalize Proposition 3.4 to the following bound, whose proof we defer to section A.6.

Proposition 3.7 (Flattenings of chaoses of nearly combinatorial type). *Let Y^f be a chaos of nearly combinatorial type as in (20) of order q, p summation indices, and weight function f . Let $R, C \subseteq [q+2]$ with $R \cup C = [q+2]$, and let $\mathcal{R} = \cup_{t \in R} I_t$ and $\mathcal{C} = \cup_{t \in C} I_t$. Then*

$$\left\| \mathcal{A}_{[R|C]}^f \right\|^2 \leq \|f\|_\infty^2 \left(\prod_{u \in \mathcal{R}^c} S_u \right) \left(\prod_{u \in \mathcal{C}^c} S_u \right). \quad (21)$$

While the norms of flattenings could be computed exactly, Proposition 3.7 provides a user-friendly upper bound that enables one to directly apply Algorithm 3.5 to chaoses of nearly combinatorial type. In sections 4.2, 4.3, and 4.4, we will apply Proposition 3.7 to chaoses whose weight function is almost always equal to its maximum value, for which this procedure is nearly optimal.

4. APPLICATIONS

In this section we focus on four illustrative applications of our techniques. Further applications and extensions are deferred to a longer companion manuscript [BLNv].

4.1. Khatri-Rao matrices. Algorithm 3.5 provides a simple recipe for computing the norms of flattenings of chaoses of combinatorial type, which can be applied manually to chaoses of small order q . This recipe can however also be used to reason about chaoses of arbitrary order without having to explicitly write a table for each q . In particular, as only the largest norm in each class of

flattening must be computed to bound $\sigma(\mathcal{A}), v(\mathcal{A}), r(\mathcal{A})$, it suffices to analyze which choices of R and C minimize the number of summation indices that end up in $\mathcal{R} \cap \mathcal{C}$.

To illustrate this procedure, we will generalize the Khatri-Rao bound (17) for $q = 2$ to arbitrary $q \geq 2$. An analogous bound was originally derived by Rudelson [Rud11, Theorem 1.3] under more restrictive assumptions. The present bound is considerably stronger; for example, unlike the bound of [Rud11], it remains valid for a large class of sparse entry distributions.

Theorem 4.1. *Let Y be a Khatri-Rao random matrix as defined in Definition 3.2. Then*

$$\mathbb{E} \|Y\| \asymp_q \max \left\{ d^{\frac{q}{2}}, n^{\frac{1}{2}} \right\}$$

provided that $\|h\|_{L^q \log(d+n)}^q \log(d+n)^{\frac{q+3}{2}} d^{\frac{q-1}{2}} = o(\max\{d^{\frac{q}{2}}, n^{\frac{1}{2}}\})$.

Proof. For this chaos of combinatorial type, the summation indices are j_1, \dots, j_q and k . We claim that the following two final flattenings are dominant:⁴

(1) the σ -flattening $\mathcal{A}_{[1:q, q+1 | q+2]}$ with all chaos coordinates $I_t = (j_t, k)$ being in R , has

$$\mathcal{R} = \{j_1, \dots, j_q, k\}, \mathcal{C} = \{k\} \implies \|\mathcal{A}_{[1:q, q+1 | q+2]}\|^2 = d^q;$$

(2) the σ -flattening $\mathcal{A}_{[q+1 | 1:q, q+2]}$ with all chaos coordinates $I_t = (j_t, k)$ being in C , has

$$\mathcal{R} = \{j_1, \dots, j_q\}, \mathcal{C} = \{j_1, \dots, j_q, k\} \implies \|\mathcal{A}_{[q+1 | 1:q, q+2]}\|^2 = n.$$

Indeed, for any other σ - or r -flattening $\mathcal{A}_{[R|C]}$, there are $t, t' \in [q]$ (possibly equal) such that $t \in R$ and $t' \in C$. Hence both summation indices k and $j_{t'}$ appear in $\mathcal{R} \cap \mathcal{C}$, and thus $\|\mathcal{A}_{[R|C]}\|^2 \leq d^{q-1}$. Similarly, given an arbitrary v -flattening $\mathcal{A}_{[R|C]}$, there must be some $t \in [q] \cap R$ (as $R \neq \emptyset$), so both k and j_t appear in $\mathcal{R} \cap \mathcal{C}$, and thus $\|\mathcal{A}_{[R|C]}\|^2 \leq d^{q-1}$.

We have therefore shown that $\sigma(\mathcal{A}) = \max\{d^{\frac{q}{2}}, n^{\frac{1}{2}}\}$ and that $v(\mathcal{A}), r(\mathcal{A}) \leq d^{\frac{q-1}{2}}$. The conclusion now follows readily from Theorems 2.6 and 2.7. \square

Remark 4.2. *One of the main contributions of [Rud11] is to show that the smallest singular value $s_n(Y)$ is lower bounded up to an absolute constant by $d^{\frac{q}{2}}$ whenever $n \lesssim_{q,s} \frac{d^q}{\log_{(s)}(d)}$, where $\log_{(s)}(\cdot)$ is the iterated logarithm function. As will be shown in the companion paper [BLNv], a variant of our main results for the smallest singular value makes it possible to remove the $\log_{(s)}(\cdot)$ factor.*

4.2. The sum-of-squares algorithm for tensor PCA. Another important example of a matrix chaos arises in the analysis of a sum-of-squares algorithm for tensor PCA [HSS15, Hop18]. While graph matrices (see Section 4.3) are often used to provide algorithmic lower bounds, the chaos in this section is used to prove upper bounds (i.e., algorithmic guarantees).

Hopkins and collaborators [Hop18] (see also [HSS15, Section 6]) prove upper bounds on the performance of the sum-of-squares hierarchy for tensor PCA via an upper bound on the norm of $X := \sum_{i \in [n]} (W_i \otimes W_i - \mathbb{E}[W_i \otimes W_i])$, where W_1, \dots, W_n are i.i.d. $d \times d$ matrices with i.i.d. standard gaussian entries ([HSS15, Theorem B.5] and [Hop18, Theorem 6.7.1 and Lemma 6.3.4]). Their bounds are optimal up to a logarithmic factor. Using the methods of this paper, we can easily remove the spurious logarithmic factor in their bound.

Theorem 4.3. *Let W_1, \dots, W_n be i.i.d. $d \times d$ random matrices with i.i.d. $N(0, 1)$ entries. Then*

$$\mathbb{E} \left\| \sum_{i \in [n]} (W_i \otimes W_i - \mathbb{E}[W_i \otimes W_i]) \right\| \lesssim d\sqrt{n},$$

provided that $n, d \gtrsim \log(d+n)^4$.

⁴For clarity of exposition, we indicate informally for I_t and \mathcal{R}, \mathcal{C} which summation indices appear in them, rather than specifying the label of the summation index as in the formal Definition 3.3.

type	coordinates			summation			norm ²
	chaos	matrix		indices			
	ijk	j	k	i	j	k	
σ	R	R	C	R	R	RC	nd
	C	R	C	C	RC	C	nd
r	RC	R	C	RC	RC	RC	1

TABLE 2. Flattenings of \mathcal{A}'_1 : $\sigma(\mathcal{A}'_1) = \sqrt{nd}$, $r(\mathcal{A}'_1) = 1$, used in (22).

Using Theorem 4.3, it is straightforward to remove the logarithmic factor in the sum-of-squares algorithmic guarantee of [HSS15, Hop18]. Let us note that the regime of interest in this application is $d^{\tau_-} \leq n \leq d^{\tau_+}$ for fixed $0 < \tau_- < \tau_+$, so that the assumption on n, d is automatically satisfied.

Proof of Theorem 4.3. Let $g_{i,j,k} = W_i[j, k]$. Note that X naturally decomposes as

$$X = \sum_{\substack{i \in [n] \\ j, k \in [d]}} (g_{i,j,k}^2 - 1) e_{(j,j)} \otimes e_{(k,k)}^\top + \sum_{\substack{i \in [n] \\ j_1, k_1, j_2, k_2 \in [d]}} \mathbf{1}_{(j_1, k_1) \neq (j_2, k_2)} g_{i, j_1, k_1} g_{i, j_2, k_2} e_{(j_1, j_2)} \otimes e_{(k_1, k_2)}^\top,$$

and denote by X_1 and X_2 the two terms on the right-hand side.

- (1) After reordering rows and columns, and using $e_j \otimes e_j \simeq e_j$, we can express X_1 as

$$Y'_1 := \sum_{\substack{i \in [n] \\ j, k \in [d]}} (g_{i,j,k}^2 - 1) e_j \otimes e_k^\top.$$

This is a chaos of combinatorial type with $h_{i,j,k} = g_{i,j,k}^2 - 1$, for which Algorithm 3.5 outputs Table 2. As $\alpha(h) \lesssim \log(d + nd^2)$ for $p = \log(d + nd^2)$, Theorem 2.6 yields

$$\mathbb{E} \|Y'_1\| \lesssim \log(d + nd^2)^{\frac{1}{2}} \sqrt{dn} + \log(d + nd^2)^2. \quad (22)$$

- (2) After decoupling, X_2 corresponds to the chaos

$$Y_2 := \sum_{\substack{i \in [n] \\ j_1, k_1, j_2, k_2 \in [d]}} \underbrace{\mathbf{1}_{(j_1, k_1) \neq (j_2, k_2)}}_{\text{weight function } f} g_{i, j_1, k_1}^{(1)} g_{i, j_2, k_2}^{(2)} e_{(j_1, j_2)} \otimes e_{(k_1, k_2)}^\top.$$

This is a chaos of nearly combinatorial type. We can therefore use Proposition 3.7 to to upper bound the parameters by the output of Algorithm 3.5, which is given in Table 3. The iterated strong NCK inequality (Theorem 2.5) yields

$$\mathbb{E} \|Y_2\| \lesssim d\sqrt{n} + \log(d^2 + nd^2)^2 (d \vee \sqrt{n}). \quad (23)$$

Combining (22) and (23) gives the desired bound. \square

We note that the same random matrix, and an analogous bound, also appears in work on quantum expanders [LY23, Theorem 1]. Our aim here is to illustrate that we readily recover the correct bound by a mechanical application of matrix chaos inequalities.

4.3. Graph matrices. The standard framework [PR20] for obtaining algorithmic lower bounds in the sum-of-squares hierarchy is to construct a candidate pseudo-expectation, and to show that its moment matrix is positive semidefinite. When providing lower bounds for average case instances, a now standard way to construct candidate pseudo-expectation matrices is through matrix chaoses. A major challenge in this area has been that most classical random matrix inequalities were not able to analyze the spectrum of these chaoses.

This bottleneck was resolved [MPW15, BHK⁺16] by the development of a theory of the so-called *graph matrices* [MP16, AMP16]. One can think of these as a natural basis in which moment matrices

type	coordinates				summation indices					norm ²
	chaos		matrix		i	j_1	j_2	k_1	k_2	
	ij_1k_1	ij_2k_2	j_1j_2	k_1k_2						
σ	R	R	R	C	R	R	R	RC	RC	nd^2
	R	C	R	C	RC	R	RC	RC	C	d^2
	C	R	R	C	RC	RC	R	C	RC	d^2
	C	C	R	C	C	RC	RC	C	C	nd^2
v	R	R	C	C	R	RC	RC	RC	RC	n
	R	C	C	C	RC	RC	C	RC	C	d^2
	C	R	C	C	RC	C	RC	C	RC	d^2

 TABLE 3. Flattenings of \mathcal{A}_2 : $\sigma(\mathcal{A}_2) = d\sqrt{n}$, $v(\mathcal{A}_2) = d \vee \sqrt{n}$, used in (23).

(at least those that possess “enough symmetry”) can be expressed. For any graph matrix, norm bounds are known [AMP16] (see section 4.3.2 below) which in certain cases translate to bounds for moment matrices. This approach is used in showing several of the state of the art lower bounds for average case complexity in the sum-of-squares hierarchy, see [PR20, AMP16].

4.3.1. *Definition.* Graph matrices are random matrices that depend on an *input distribution* of $\binom{n}{2}$ i.i.d. Rademacher random variables (each corresponding to an edge of a complete graph on n nodes), and a small sized graph α called a *shape*, with identified subsets $U_\alpha, V_\alpha \subseteq V(\alpha)$ of, respectively, left and right vertices. The shape will be fixed, while n is best thought of as arbitrarily large (in other words, we will not aim to optimize the dependency of our bounds on constants depending on α).

Definition 4.4 (Shape). A *shape* is a graph, that has a subset $U_\alpha \subseteq V(\alpha)$ of *left vertices*, and another subset $V_\alpha \subseteq V(\alpha)$ of *right vertices*.⁵

Definition 4.5 (Graph matrices). Let α be a shape and n be a large integer.

- (1) The set of *middle vertices* is given by $W_\alpha = V(\alpha) \setminus (U_\alpha \cup V_\alpha)$.
- (2) The *ground set* is the set of indices $[n] = \{1, \dots, n\}$, which we also interpret as the vertices of the complete graph K_n .
- (3) The *input distribution* $\varepsilon = (\varepsilon_e)_{e \in E(K_n)}$ is a collection of i.i.d. Rademachers indexed by the edges of K_n (i.e. unordered pairs of distinct numbers).
- (4) A *realization* is any injective map $\varphi: V(\alpha) \rightarrow [n]$ from the shape vertices to the ground set.
- (5) The *graph matrix* M_α is the $n^{|U_\alpha|} \times n^{|V_\alpha|}$ random matrix, whose rows and columns are indexed by ordered subsets of $[n]$ with cardinality $|U_\alpha|$ and $|V_\alpha|$, respectively, given by

$$M_\alpha := \sum_{\text{realization } \varphi} \left(\prod_{(i,j) \in E(\alpha)} \varepsilon_{\varphi(i), \varphi(j)} \right) e_{\varphi(U_\alpha)} \otimes e_{\varphi(V_\alpha)}^\top. \quad (24)$$

Remark 4.6 (Identification in notation). *In the discussion that follows, we will treat graph matrices within the framework developed in Section 3. Observe that summing over realizations in (24) corresponds to different choices for summation indices in (13). Thus we will, in a slight abuse of notation, use the same symbols to denote vertices from $V(\alpha)$ and summation indices.*

Example 4.7 (Examples of graph matrices). *These examples are also represented in Figure 2.*

- (1) (*Wigner without a diagonal*) If $U_\beta = \{i\}, V_\beta = \{j\}, W_\beta = \emptyset, E(\beta) = \{(i, j)\}$, then

$$M_\beta = \sum_{i \neq j} \varepsilon_{i,j} e_i \otimes e_j^\top \quad (25)$$

is an $n \times n$ Wigner matrix with zeros on the diagonal.

⁵The sets U_α and V_α can intersect, their union is not necessarily $V(\alpha)$, and their sizes are not necessarily equal.




FIGURE 2. Using Theorem 4.8 on graph matrices from Example 4.7 yields the following bounds (logarithmic factors omitted): $\|M_\beta\| \approx \sqrt{n}$, $\|M_\gamma\| \approx n$, $\|M_\delta\| \approx n^3$.

(2) (*Z-shaped graph matrix*) If $U_\gamma = \{i, j\}$, $V_\gamma = \{k, l\}$, $W_\gamma = \{\emptyset\}$, $E(\gamma) = \{(i, k), (j, k), (j, l)\}$,

$$M_\gamma = \sum_{i,j,k,l \text{ distinct}} \varepsilon_{i,k} \varepsilon_{j,k} \varepsilon_{j,l} e_{(i,j)} \otimes e_{(k,l)}^\top$$

is an $n^2 \times n^2$ asymmetric matrix that was studied in the context of free probability [CP22].

(3) (*Example of a graph matrix with middle vertices*) If $U_\delta = \{i, j\}$, $V_\delta = \{k, l\}$, $W_\delta = \{m, o\}$, $E(\delta) = \{(i, m), (j, m), (k, m), (l, m)\}$, then

$$M_\delta = \sum_{i,j,k,l,m,o \text{ distinct}} \varepsilon_{i,m} \varepsilon_{j,m} \varepsilon_{k,m} \varepsilon_{l,m} e_{(i,j)} \otimes e_{(k,l)}^\top$$

is an $n^2 \times n^2$ symmetric matrix. Note that $|W_\delta|$ has no effect on the dimension of M_δ .

4.3.2. *Norm bounds.* We are ready to state a general bound on the norm of graph matrices. Recall that a set of vertices S is a $U-V$ vertex separator if all paths from U to V pass through S .

Theorem 4.8 (Graph matrix norm bounds). *Given a shape α , let M_α be the associated graph matrix as in (24). Then we have*

$$n^{\frac{1}{2}(|V(\alpha)| - |S_{\min}| + |W_{\text{iso}}|)} \lesssim_\alpha \mathbb{E} \|M_\alpha\| \lesssim_\alpha n^{\frac{1}{2}(|V(\alpha)| - |S_{\min}| + |W_{\text{iso}}|)} \cdot (\log n)^{\frac{1}{2}f(\alpha)}, \quad (26)$$

where $f(\alpha) = |S_{\min}| - |U_\alpha \cap V_\alpha| + |W_\alpha| - |W_{\text{iso}}|$. Here $|S_{\min}|$ is the size of the minimal $U_\alpha - V_\alpha$ vertex separator and W_{iso} is the set of all isolated middle vertices.

The upper bound in Theorem 4.8 first appeared in [AMP16, MP16] where an intricate moment method argument is used. The minimal vertex separator S_{\min} appears indirectly as a consequence of duality between max-flow and min-cut. A lower bound was shown there for most shapes. More recently, a similar upper bound was obtained in [RT23] by analyzing matrices of partial derivatives that arise by iterating Efron-Stein inequalities. In the case of Rademachers, these matrices are deterministic (and coincide with the flattenings discussed here), and S_{\min} naturally arises in computations of Frobenius norms. This provides a more direct proof of the upper bound, but with a larger power of the logarithm: an edge quantity $|E(\alpha)|$ replaces the vertex quantity $f(\alpha)$.

Our tools allow a direct proof of Theorem 4.8 with the $f(\alpha)$ logarithmic power and provide a lower bound for all shapes. In this manuscript we provide a proof of the upper bound, whereas a proof of the lower bound is deferred to [BLNv] (see Remark 4.9).

It should be emphasized, however, that the proof of Theorem 4.8 only uses the simplest iterated NCK inequality to achieve universal bounds. The main benefit of our framework is that it provides an effortless way to remove logarithmic factors in instances where v -flattenings are negligible, by using instead the iterated strong inequalities. The latter provides a systematic method for achieving improved bounds for graph matrices, as will be illustrated in Section 4.4 below.

4.3.3. *Flattenings of graph matrices.* One small obstacle to directly proving Theorem 4.8 using Proposition 3.7 is the fact that the input distribution is indexed by edges, not by ordered pairs—in other words $\varepsilon_{\varphi(i),\varphi(j)} \equiv \varepsilon_{\varphi(j),\varphi(i)}$. This obstacle is already present when trying to upper bound the spectral norm in the example of M_β defined in (25). However, if we decompose

$$M_\beta = \sum_{i,j} \varepsilon_{i,j} e_i \otimes e_j^\top = \sum_{i,j} \mathbf{1}_{i<j} \varepsilon_{i,j} e_i \otimes e_j^\top + \sum_{i,j} \mathbf{1}_{i>j} \varepsilon_{i,j} e_i \otimes e_j^\top,$$

then each summand is a chaos of nearly combinatorial type, and its parameters can be analyzed with Proposition 3.7. We will apply a similar idea in the general setting.

Proof of Theorem 4.8: upper bound with $(\log n)^{\frac{1}{2}|E(\alpha)}$ factor. Given a graph matrix M_α , we have

$$\begin{aligned} M_\alpha &= \sum_{\text{realization } \varphi} \left(\prod_{(i,j) \in E(\alpha)} \varepsilon_{\varphi(i),\varphi(j)} \right) e_{\varphi(U_\alpha)} \otimes e_{\varphi(V_\alpha)}^\top \\ &= \sum_{E \subseteq E(\alpha)} \sum_{\varphi} \left(\prod_{(i,j) \in E} \mathbf{1}_{\varphi(i) > \varphi(j)} \prod_{(i,j) \in E(\alpha) \setminus E} \mathbf{1}_{\varphi(i) < \varphi(j)} \right) \left(\prod_{(i,j) \in E(\alpha)} \varepsilon_{\varphi(i),\varphi(j)} \right) e_{\varphi(U_\alpha)} \otimes e_{\varphi(V_\alpha)}^\top. \end{aligned}$$

The summand $M_{\alpha,E}$ associated to each (possibly empty) subset of edges $E \subseteq E(\alpha)$ is, after decoupling, is a chaos of nearly combinatorial type (Definition 3.6). Each chaos has

- $p = |V(\alpha)|$ summation indices $(s_v)_{v \in V(\alpha)}$ (which correspond to $s_v := \varphi(v)$, see Remark 4.6);
- $q = |E(\alpha)|$ chaos coordinates, which correspond to shape edges:

$$I_e(\mathbf{s}) = (s_u, s_v) \text{ for } e = (u, v) \in E(\alpha);$$

- matrix coordinates given by

$$I_{q+1}(\mathbf{s}) = (s_u)_{u \in U_\alpha}, \quad I_{q+2}(\mathbf{s}) = (s_u)_{u \in V_\alpha};$$

- a weight function whose ℓ_∞ norm is 1.

Consider any final σ -flattening $\mathcal{A}_{[R|C]}$ of $M_{\alpha,E}$. Then the formula (21) yields

$$\left\| \mathcal{A}_{[R|C]} \right\| \leq n^{\frac{1}{2}|\mathcal{R}^c|} n^{\frac{1}{2}|\mathcal{C}^c|} = n^{\frac{1}{2}(|V(\alpha)| - |\mathcal{R} \cap \mathcal{C}| + |\mathcal{R}^c \cap \mathcal{C}^c|)}.$$

The following two key inequalities explain the polynomial power in (26):

- (1) $|\mathcal{R} \cap \mathcal{C}| \geq |S_{\min}|$ holds as vertices in $\mathcal{R} \cap \mathcal{C}$ form a vertex separator between U_α and V_α : indeed, any path in α that starts in $U_\alpha \subseteq \mathcal{R}$ and ends in $V_\alpha \subseteq \mathcal{C}$ has a vertex in $\mathcal{R} \cap \mathcal{C}$. The equality $|\mathcal{R} \cap \mathcal{C}| = |S_{\min}|$ is attained whenever R consists precisely of all edges that are accessible from U_α without passing through S_{\min} .
- (2) $|\mathcal{R}^c \cap \mathcal{C}^c| \leq |W_{\text{iso}}|$ holds as summation indices that do not appear in \mathcal{R} nor \mathcal{C} must correspond to isolated middle vertices, as they do not have an incident edge (in I_e for some $e \in E(\alpha)$) and do not appear on the left or right sides of the shape (in I_{q+1} or I_{q+2}).

Thus

$$\sigma(\mathcal{A}) \leq n^{\frac{1}{2}(|V(\alpha)| - |S_{\min}| + |W_{\text{iso}}|)}, \quad (27)$$

and an upper bound as in (26) with the multiplicative factor $\log(n)^{\frac{1}{2}|E(\alpha)|}$ follows by using the iterated NCK inequality (Theorem 2.4) and the triangle inequality over all 2^q choices of E . \square

4.3.4. *Intermediate flattenings of graph matrices.* We now focus our attention on improving the logarithmic factor. Recall from Section 2.4.2 that iterating the NCK inequality yields a bound on the norm of a matrix chaos in terms of its intermediate flattenings. More precisely, after performing $k \leq q$ iterations of the NCK inequality, one obtains a *partially iterated NCK inequality*:

$$\mathbb{E} \|Y\| \lesssim_q \log(d+m)^{\frac{k}{2}} \max_{R' \sqcup C' = \{q-k+1, \dots, q\}} \mathbb{E} \left\| Y_{[1:q-k | R' \cup \{q+1\} | C' \cup \{q+2\}]} \right\|. \quad (28)$$

When $k = q$, this reduces to the iterated NCK inequality of Theorem 2.4.

In the present setting, however, it will be useful to apply this bound with $k < q$. The reason is that when the random variables $\mathbf{h}^{(t)}$ in the matrix chaos are uniformly bounded (as is the case for the Rademacher variables that appear here), we can upper bound $Y_{[Z | R | C]}$ entrywise to recover a regular flattening whose norm can be computed using the formula (21) (see Remark A.10). We will show that the chaos variables of graph matrices can always be ordered so that $k \leq f(\alpha)$ iterations suffice to achieve the same upper bound on the partial flattenings as was obtained in the previous section for the final flattenings, resulting in an improved power of the logarithm.

Proof of Theorem 4.8: upper bound with $(\log n)^{\frac{1}{2}f(\alpha)}$ factor. We begin by choosing a special ordering of the edges $E(\alpha)$ of the given shape α , as follows.

- (1) By Menger's theorem (Theorem A.11), there is a family of $|S_{\min}|$ vertex-disjoint paths from U_α to V_α , each of which contains exactly one point from U_α and one point from V_α . We place the union of all k_1 edges in these paths last in our ordering of $E(\alpha)$.
- (2) Next, we choose the smallest number k_2 of additional edges, so that every non-isolated middle vertex that is not contained in one of the above paths is incident to one of the additional edges. We place the additional edges in the middle of our ordering of $E(\alpha)$.
- (3) All remaining edges are placed at the beginning of our ordering of $E(\alpha)$.

We claim that $k = k_1 + k_2 \leq f(\alpha)$. Indeed, by construction, the set of paths constructed in the first step contains exactly $|S_{\min}| - |U_\alpha \cap V_\alpha|$ paths of lengths $\ell_i \geq 2$, each of which contains exactly $\ell_i - 1$ edges and $\ell_i - 2$ middle vertices. The union of these paths therefore contain exactly

$$\sum_{i=1}^{|S_{\min}| - |U_\alpha \cap V_\alpha|} (\ell_i - 2) = k_1 - |S_{\min}| + |U_\alpha \cap V_\alpha|$$

(necessarily non-isolated) middle vertices. As the total number of non-isolated middle vertices is $|W_\alpha| - |W_{\text{iso}}|$, we must therefore choose at most

$$k_2 \leq |W_\alpha| - |W_{\text{iso}}| - (k_1 - |S_{\min}| + |U_\alpha \cap V_\alpha|) = f(\alpha) - k_1$$

additional edges in the second step. This establishes the claim.

Now let $M_{\alpha, E}$ be as in the proof of Theorem 4.8, and let Y be its decoupled version which is a chaos of nearly combinatorial type. Then any intermediate flattening $\mathcal{A}_{[Z | R | C]}$ that appears in (28) has the last k shape edges (chaos coordinates) assigned to either R or C . Therefore:

- (1) Each path constructed in the first step above contains at least one vertex (summation index) in $\mathcal{R} \cap \mathcal{C}$, so $|\mathcal{R} \cap \mathcal{C}| \geq |S_{\min}|$;
- (2) every non-isolated middle vertex is in $\mathcal{R} \cup \mathcal{C}$, so $|\mathcal{R}^c \cap \mathcal{C}^c| \leq |W_{\text{iso}}|$.

By upper bounding $Y_{[Z | R | C]}$ entrywise and applying (21) (see Remark A.10), we obtain

$$\max_{R' \sqcup C' = \{q-k+1, \dots, q\}} \mathbb{E} \left\| Y_{[1:q-k | R' \cup \{q+1\} | C' \cup \{q+2\}]} \right\| \leq n^{\frac{1}{2}(|V(\alpha)| - |S_{\min}| + |W_{\text{iso}}|)}$$

precisely as in (27). The conclusion now follows from the partially iterated NCK inequality (28). \square

Remark 4.9. *The lower bound in Theorem 4.8 can be proved by considering a chaos of combinatorial type that is obtained from M_α by considering only a subset of the summands (by restricting the input distribution to the edge set of a $V(\alpha)$ -partite graph on n nodes). We defer the details of this argument to [BLNv]; a similar idea is used in [AMP16].*

4.4. Sharper bounds on graph matrices and ellipsoid fitting. An important example where a sharper bound on the spectral norm of a graph matrix was derived is in the context of the ellipsoid fitting problem [HKPX23]. The ellipsoid fitting conjecture is a question in stochastic geometry that has received considerable attention recently (see [TW23, HKPX23, BMMP24] and references therein). In order to obtain a lower bound of the correct asymptotic order,⁶ the authors of [HKPX23] developed techniques to remove spurious logarithmic factors from the bound on the spectrum of certain graph matrices. These arguments involve sophisticated refinements of moment method calculations. In this section we show how Theorem 2.7 and Algorithm 3.5 can be used to effortlessly recover these improvements as a mechanical application of our general theory.

The two random matrices that need to be analyzed in this procedure (we refer the reader to [HKPX23], in particular Proposition 2.3 in this reference, for the derivation of how these matrices arise) are the $m \times m$ random matrices M_ϕ and M_ψ given by

$$M_\phi = \sum_{i \neq j \in [m]} \sum_{a \neq b \in [d]} (g_{i,a} g_{i,b} g_{j,a} g_{j,b}) e_i \otimes e_j^\top, \quad (29)$$

$$M_\psi = \sum_{i \neq j \in [m]} \sum_{a \in [d]} (g_{i,a}^2 - 1) (g_{j,a}^2 - 1) e_i \otimes e_j^\top, \quad (30)$$

where $(g_{i,a})_{i \in [m], a \in [d]}$ are md i.i.d. standard gaussian variables. The motivating example has $m \asymp d^2$, see [HKPX23], so that the assumption of Theorem 4.11 below is automatically satisfied.

Remark 4.10. *While M_ϕ and M_ψ are not precisely graph matrices in the sense of Definition 4.5, they may be viewed as generalized graph matrices in the sense of [AMP16]. Here we gloss over the distinction and simply view these matrices as special instances of chaoses of combinatorial type.*

Using our tools we provide an alternative proof of Lemma 2.7 from [HKPX23] (note that there is an additional scaling by d^2 to obtain random variables of unit variance).

Theorem 4.11. *Let M_ϕ and M_ψ be the random matrices in (29) and (30). We have*

$$\mathbb{E} \|M_\phi\| \lesssim d\sqrt{m} \vee m, \quad \mathbb{E} \|M_\psi\| \lesssim m \vee \sqrt{md},$$

provided that $d, m \gtrsim \log(d + m)^9$.

Proof. The proof is similar to that of Theorem 4.3. Note that M_ϕ and M_ψ are square-free matrix chaoses, whose decoupled versions are chaoses of combinatorial type.

- (1) The decoupled version of M_ϕ is a gaussian matrix chaos of order 4. Algorithm 3.5 outputs Table 4, and the iterated strong NCK inequality (Theorem 2.5) and Theorem 2.1 yield

$$\mathbb{E} \|M_\phi\| \lesssim (d\sqrt{m} \vee m) + \log(md + m)^3 (\sqrt{md} \vee d). \quad (31)$$

- (2) The decoupled version of M_ψ is a matrix chaos of order 2 whose random variables are given by $h_{i,a} = g_{i,a}^2 - 1$. Algorithm 3.5 outputs Table 5, and the iterated strong matrix Rosenthal inequality (Theorem 2.7) and Theorem 2.1 yield

$$\mathbb{E} \|M_\psi\| \lesssim (m \vee \sqrt{md}) + \log(md + m)^{\frac{9}{2}} (\sqrt{m} \vee \sqrt{d}). \quad (32)$$

The first term in (31) and in (32) dominates under the assumption on d, m , concluding the proof. \square

⁶A lower bound with the correct asymptotic order was concurrently obtained in [TW23, HKPX23, BMMP24].

tp.	coordinates				matrix		summation indices				norm ²
	chaos	chaos	chaos	chaos	i	j	i	j	a	b	
σ	R	R	R	R	R	C	R	RC	R	R	md^2
	R	R	R	C	R	C	R	RC	R	RC	md
	R	R	C	R	R	C	R	RC	RC	R	md
	R	R	C	C	R	C	R	C	RC	RC	m^2
	R	C	R	R	R	C	RC	RC	R	RC	d
	R	C	R	C	R	C	RC	RC	R	C	d^2
	R	C	C	R	R	C	RC	RC	RC	RC	1
	R	C	C	C	R	C	RC	C	RC	C	md
	C	R	R	R	R	C	RC	RC	RC	R	d
	C	R	R	C	R	C	RC	RC	RC	RC	1
	C	R	C	R	R	C	RC	RC	C	R	d^2
	C	R	C	C	R	C	RC	C	C	RC	md
	C	C	R	R	R	C	RC	RC	RC	RC	1
	C	C	R	C	R	C	RC	RC	RC	C	d
	C	C	C	R	R	C	RC	RC	C	RC	d
	C	C	C	C	R	C	RC	C	C	C	md^2
v	R	R	R	R	C	C	RC	RC	R	R	d^2
	R	R	R	C	C	C	RC	RC	R	RC	d
	R	R	C	R	C	C	RC	RC	RC	R	d
	R	R	C	C	C	C	RC	C	RC	RC	m
	R	C	R	R	C	C	RC	RC	R	RC	d
	R	C	R	C	C	C	RC	RC	R	C	d^2
	R	C	C	R	C	C	RC	RC	RC	RC	1
	R	C	C	C	C	C	RC	C	RC	C	md
	C	R	R	R	C	C	RC	RC	RC	R	d
	C	R	R	C	C	C	RC	RC	RC	RC	1
	C	R	C	R	C	C	RC	RC	C	R	d^2
	C	R	C	C	C	C	RC	C	C	RC	md
	C	C	R	R	C	C	C	C	RC	RC	m
	C	C	R	C	C	C	C	RC	RC	C	md
	C	C	C	R	C	C	C	RC	C	RC	md

TABLE 4. Flattenings of \mathcal{A}_ϕ : $\sigma(\mathcal{A}_\phi) = d\sqrt{m} \vee m$, $v(\mathcal{A}_\phi) = \sqrt{md} \vee d$, used in (31).

type	coordinates		matrix		summation indices			norm ²
	chaos	chaos	i	j	i	j	a	
σ	R	R	R	C	R	RC	R	md
	R	C	R	C	R	C	RC	m^2
	C	R	R	C	RC	RC	RC	1
	C	C	R	C	RC	C	C	md
v	R	R	C	C	RC	RC	R	d
	R	C	C	C	RC	C	RC	m
	C	R	C	C	C	RC	RC	m

TABLE 5. Flattenings of \mathcal{A}_ψ : $\sigma(\mathcal{A}_\psi) = m \vee \sqrt{md}$, $v(\mathcal{A}_\psi) = \sqrt{m} \vee \sqrt{d}$, used in (32).

Acknowledgements. ASB would like to thank Sam Hopkins, Ankur Moitra, and Holger Rauhut for asking insightful questions in conversations over the past few years that helped motivate the methods of this paper. RvH was supported in part by NSF grant DMS-2347954.

REFERENCES

- [ALM21] Radosław Adamczak, Rafał Łatała, and Rafał Meller. Moments of Gaussian chaoses in Banach spaces. *Electron. J. Probab.*, 26:Paper No. 11, 36, 2021.
- [AMP16] Kwangjun Ahn, Dhruv Medarametla, and Aaron Potechin. Graph matrices: Norm bounds and applications. *arXiv preprint arXiv:1604.03423*, 2016.
- [BBvH23] Afonso S Bandeira, March T Boedihardjo, and Ramon van Handel. Matrix concentration inequalities and free probability. *Inventiones mathematicae*, 234(1):419–487, 2023.
- [BCSv24] A. S. Bandeira, G. Cipolloni, D. Schröder, and R. van Handel. Matrix concentration inequalities and free probability II. Two-sided bounds and applications, 2024. Preprint arxiv:2406.11453.
- [BHK⁺16] Boaz Barak, Samuel B. Hopkins, Jonathan A. Kelner, Pravesh Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 428–437, 2016.
- [BLNv] Afonso S. Bandeira, Kevin Lucca, Petar Nizić-Nikolac, and Ramon van Handel. Matrix chaos inequalities. *Forthcoming*.
- [BMMP24] Afonso S Bandeira, Antoine Maillard, Shahar Mendelson, and Elliot Paquette. Fitting an ellipsoid to a quadratic number of random points. *To appear in Latin American Journal of Probability and Mathematical Statistics.*, 2024.
- [BvH24] Tatiana Brailovskaya and Ramon van Handel. Universality and Sharp Matrix Concentration Inequalities. *Geom. Funct. Anal.*, 34(6):1734–1838, 2024.

- [CP22] Wenjun Cai and Aaron Potechin. On mixing distributions via random orthogonal matrices and the spectrum of the singular values of multi-z shaped graph matrices. *arXiv preprint arXiv:2206.02224*, 2022.
- [De12] Anindya De. Lower bounds in differential privacy. In *Theory of Cryptography: 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings 9*, pages 321–338. Springer, 2012.
- [dlPG12] V. de la Peña and E. Giné. *Decoupling: From Dependence to Independence*. Probability and Its Applications. Springer New York, 2012.
- [DNY20] Yu Deng, Andrea R. Nahmod, and Haitian Yue. Random tensors, propagation of randomness, and nonlinear dispersive equations. *Inventiones mathematicae*, 228:539 – 686, 2020.
- [FM24] Zhou Fan and Renyuan Ma. Kronecker-product random matrices and a matrix least squares problem. *arXiv preprint arXiv:2406.00961*, 2024.
- [Gö2] Frank Göring. A proof of Menger’s theorem by contraction. *Discuss. Math. Graph Theory*, 22(1):111–112, 2002. Conference on Graph Theory (Elgersburg, 2000).
- [HKPX23] Jun-Ting Hsieh, Pravesh K. Kothari, Aaron Potechin, and Jeff Xu. Ellipsoid Fitting up to a Constant. In Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 78:1–78:20, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [Hop18] Samuel Hopkins. Statistical inference and the sum of squares method. *Dissertation, Cornell University*, 2018.
- [HP93] Uffe Haagerup and Gilles Pisier. Bounded linear operators between C^* -algebras. *Duke Mathematical Journal*, 71(3):889 – 925, 1993.
- [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.
- [JZ13] Marius Junge and Qiang Zeng. Noncommutative Bennett and Rosenthal inequalities. *Ann. Probab.*, 41(6):4287–4316, 2013.
- [KR68] C. G. Khatri and C. Radhakrishna Rao. Solutions to some functional equations and their applications to characterization of probability distributions. *Sankhya: The Indian Journal of Statistics, Series A (1961-2002)*, 30(2):167–180, 1968.
- [KRSU10] Shiva Prasad Kasiviswanathan, Mark Rudelson, Adam Smith, and Jonathan Ullman. The price of privately releasing contingency tables and the spectra of random matrices with correlated rows. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing, STOC ’10*, page 775–784, New York, NY, USA, 2010. Association for Computing Machinery.
- [LO94] Rafał Łatała and Krzysztof Oleszkiewicz. On the best constant in the khinchin-kahane inequality. *Studia Mathematica*, 109(1):101–104, 1994.
- [LY23] Cécilia Lancien and Pierre Youssef. A note on quantum expanders. *arXiv preprint arXiv:2302.07772*, 2023.
- [MJC⁺14] Lester Mackey, Michael I. Jordan, Richard Y. Chen, Brendan Farrell, and Joel A. Tropp. Matrix concentration inequalities via the method of exchangeable pairs. *Ann. Probab.*, 42(3):906–945, 2014.
- [MP16] Dhruv Medarametla and Aaron Potechin. Bounds on the norms of uniform low degree graph matrices. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*, volume 60 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 40, 26. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016.
- [MPW15] Raghu Meka, Aaron Potechin, and Avi Wigderson. Sum-of-squares lower bounds for planted clique. In *STOC’15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 87–96. ACM, New York, 2015.
- [MSS16] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 438–446, 2016.
- [MW19] Stanislav Minsker and Xiaohan Wei. Moment inequalities for matrix-valued U-statistics of order 2. *Electron. J. Probab.*, 24:Paper No. 133, 32, 2019.
- [Pis03] Gilles Pisier. *Introduction to Operator Space Theory*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2003.
- [Pis14] Gilles Pisier. Random matrices and subexponential operator spaces. *Israel J. Math.*, 203(1):223–273, 2014.
- [PR20] Aaron Potechin and Goutham Rajendran. Machinery for proving sum-of-squares lower bounds on certification problems. *ArXiv*, abs/2011.04253, 2020.

- [Rau09] Holger Rauhut. Circulant and Toeplitz Matrices in Compressed Sensing. In Rémi Gribonval, editor, *SPARS'09 - Signal Processing with Adaptive Sparse Structured Representations*, Saint Malo, France, April 2009. Inria Rennes - Bretagne Atlantique.
- [RT23] Goutham Rajendran and Madhur Tulsiani. Concentration of polynomial random matrices via efron-stein inequalities. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3614–3653. SIAM, 2023.
- [Rud11] Mark Rudelson. Row products of random matrices. *Advances in Mathematics*, 231:3199–3231, 2011.
- [Tro15] Joel A Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- [TW23] Madhur Tulsiani and June Wu. Ellipsoid fitting up to constant via empirical covariance estimation. *arXiv preprint arXiv:2307.10941*, 2023.
- [TW24] Madhur Tulsiani and June Wu. Simple norm bounds for polynomial random matrices via decoupling. *arXiv preprint arXiv:2412.07936*, 2024.
- [Ver18] Roman Vershynin. *High-dimensional probability*. Cambridge University Press, Cambridge, 2018.

APPENDIX A. PROOFS OF MAIN RESULTS AND SUPPORTING LEMMAS

A.1. The iteration scheme. The basic approach to all our main results was outlined in Section 2.4. For each of the iterated inequalities, we start with an inequality for linear random matrices (i.e., for chaos of order $q = 1$). The linear inequalities involve four parameters σ_R, σ_C, v, r defined in Section 2.4.1. Applying these bounds conditionally on all but one of the chaos coordinates gives rise to four intermediate flattenings as shown in Figure 1. As the intermediate flattenings are themselves matrix chaoses of smaller order, the proofs proceed by induction.

The following lemma formalizes the fact, used in the induction step, that the final flattenings of intermediate flattenings coincide with the final flattenings of the original chaos.

Lemma A.1 (σ, v and r of intermediate flattenings). *Let Y be a decoupled chaos as in (2). Given an intermediate flattening $Y_{[Z|R|C]}$, which is a chaos of order $|Z|$, we have*

$$\begin{aligned} \sigma\left(\mathcal{A}_{[Z|R|C]}\right) &= \max_{R' \sqcup C' = Z} \left\| \mathcal{A}_{[R \cup R' | C \cup C']}\right\|, \\ v\left(\mathcal{A}_{[Z|R|C]}\right) &= \max_{\substack{R' \sqcup C' = Z \\ R' \neq \emptyset}} \left\| \mathcal{A}_{[R' | R \cup C \cup C']}\right\|, \\ r\left(\mathcal{A}_{[Z|R|C]}\right) &= \max_{\substack{R' \cup C' = Z \\ R' \cap C' \neq \emptyset}} \left\| \mathcal{A}_{[R \cup R' | C \cup C']}\right\|. \end{aligned}$$

Proof. By its definition (10), the intermediate flattening $Y_{[Z|R|C]}$ is a matrix chaos with the same coefficients $\mathcal{A}_{i_1, \dots, i_{q+2}}$ as Y , but where the $|Z|$ chaos coordinates are indexed by i_t for $t \in Z$ and the two matrix coordinates are indexed by $(i_t : t \in R)$ and $(i_t : t \in C)$, respectively. The conclusion now follows readily from the definitions (4), (5) and (6) of the chaos parameters. \square

For completeness, we record here two basic relations between the chaos parameters.

Lemma A.2. *For any chaos Y as in (2), we have*

$$r(\mathcal{A}) \leq \sigma(\mathcal{A}) \quad \text{and} \quad r(\mathcal{A}) \leq v(\mathcal{A}).$$

Proof. In the case $q = 1$, we may readily read off from the expressions in section 2.4.1 that

$$r(Y) = \left\| \mathcal{A}_{[1,2|1,3]}\right\| \leq \left\| \mathcal{A}_{[1,2|3]}\right\| \leq \sigma(Y), \quad r(Y) = \left\| \mathcal{A}_{[1,2|1,3]}\right\| \leq \left\| \mathcal{A}_{[1|2,3]}\right\| = v(Y),$$

where we used that $\|A_i\| \leq \|A_i\|_F = \|\text{vec}(A_i)\|$ in the second inequality.

Now let $q \geq 2$, and let $\mathcal{A}_{[R|C]}$ be any r -flattening. Consider a 3-tensor \mathcal{B} whose entries are given by $\mathcal{B}_{i_{R \cap C}, i_{R \setminus C}, i_{C \setminus R}} = \mathcal{A}_i$, where $\mathbf{i} = (i_1, \dots, i_{q+2})$ and $\mathbf{i}_T = (i_t : t \in T)$. Then

$$\mathcal{B}_{[1,2|1,3]} = \mathcal{A}_{[R|C]}, \quad \mathcal{B}_{[1,2|3]} = \mathcal{A}_{[R|C \setminus R]}, \quad \mathcal{B}_{[1|2,3]} = \mathcal{A}_{[R \cap C | R^c \cup C^c]}.$$

As $\mathcal{A}_{[R|C\setminus R]}$ is a σ -flattening of \mathcal{A} and $\mathcal{A}_{[R\cap C|R^c\cup C^c]}$ is a v -flattening of \mathcal{A} , we obtain

$$\begin{aligned}\|\mathcal{A}_{[R|C]}\| &= \|\mathcal{B}_{[1,2|1,3]}\| \leq \|\mathcal{B}_{[1,2|3]}\| = \|\mathcal{A}_{[R|C\setminus R]}\| \leq \sigma(\mathcal{A}) \\ \|\mathcal{A}_{[R|C]}\| &= \|\mathcal{B}_{[1,2|1,3]}\| \leq \|\mathcal{B}_{[1|2,3]}\| = \|\mathcal{A}_{[R\cap C|R^c\cup C^c]}\| \leq v(\mathcal{A})\end{aligned}$$

by applying the inequalities for the case $q = 1$ to the tensor \mathcal{B} . \square

A.2. Proof of Iterated NCK. We start by stating the linear theorem. The following result is classical, but we spell it out in a slightly more general setting than is customary.

Theorem A.3 (Noncommutative Khintchine (NCK) inequality). *Let $X = \sum_{i \in [m]} h_i A_i$, where h_1, \dots, h_m are i.i.d. copies of a centered random variable h and A_1, \dots, A_m are $d_1 \times d_2$ matrix coefficients (we define $d := d_1 \vee d_2$). Then we have*

$$\|h\|_{L^1} (\sigma_R(X) + \sigma_C(X)) \lesssim \mathbb{E} \|X\| \lesssim \|h\|_{\psi_2} \log(d)^{\frac{1}{2}} (\sigma_R(X) + \sigma_C(X)).$$

Alternatively, the upper bound remains valid if $\|h\|_{\psi_2}$ is replaced by $\|h\|_{L^{\log m}}$.

Proof. We begin with the lower bound. Let ε_i be i.i.d. Rademachers variables independent of h_i , and define $\tilde{X} = \sum_{i \in [m]} \varepsilon_i h_i A_i$. Then $\mathbb{E} \|\tilde{X}\| \leq 2 \mathbb{E} \|X\|$ by a standard symmetrization argument [Ver18, Lemma 6.3.2]. Taking the expectation only with respect to ε , we can estimate

$$\mathbb{E}_\varepsilon \|\tilde{X}\| \gtrsim 2 \left(\mathbb{E}_\varepsilon \|\tilde{X}\|^2 \right)^{\frac{1}{2}} = \left(\mathbb{E}_\varepsilon \|\tilde{X}^\top \tilde{X}\| \right)^{\frac{1}{2}} + \left(\mathbb{E}_\varepsilon \|\tilde{X} \tilde{X}^\top\| \right)^{\frac{1}{2}} \geq \left\| \sum_i h_i^2 A_i^\top A_i \right\|^{\frac{1}{2}} + \left\| \sum_i h_i^2 A_i A_i^\top \right\|^{\frac{1}{2}},$$

where we used the Khintchine-Kahane inequality [LO94] in the first step and Jensen's inequality in the last step. As the right-hand side is a convex function of $(|h_1|, \dots, |h_m|)$, taking the expectation and applying Jensen's inequality yields $\mathbb{E} \|X\| \gtrsim \|h\|_{L^1} (\sigma_R(X) + \sigma_C(X))$.

We now turn to the upper bound. The classical form of the noncommutative Khintchine inequality [Pis03, §9.8] (and $\text{Tr}[|M|^p]^{1/p} \lesssim \|M\|$ for any $d \times d$ matrix M and $p \gtrsim \log d$) yields the upper bound in the special case that h_i are Rademacher or standard Gaussian variables. The subgaussian upper bound then follows as $\mathbb{E} \|X\| \lesssim \|h\|_{\psi_2} \mathbb{E} \|X_G\|$, where $X_G = \sum_{i \in [m]} g_i A_i$ with g_1, \dots, g_m i.i.d. standard Gaussians, by the subgaussian comparison theorem [Ver18, Corollary 8.6.3].

Alternatively, by symmetrizing as in the lower bound, we can estimate

$$\begin{aligned}\mathbb{E} \|X\| &\leq 2 \mathbb{E} \|\tilde{X}\| \lesssim \log(d)^{\frac{1}{2}} \left(\mathbb{E} \left\| \sum_i h_i^2 A_i^\top A_i \right\|^{\frac{1}{2}} + \mathbb{E} \left\| \sum_i h_i^2 A_i A_i^\top \right\|^{\frac{1}{2}} \right) \\ &\leq \mathbb{E} \left(\max_{i \in [m]} |h_i| \right) \log(d)^{\frac{1}{2}} (\sigma_R(X) + \sigma_C(X))\end{aligned}$$

by applying the Rademacher form of NCK conditionally on \mathbf{h} . It remains to note that we can estimate $\mathbb{E} \max_{i \in [m]} |h_i| \leq (\mathbb{E} \max_{i \in [m]} |h_i|^p)^{1/p} \leq (m \mathbb{E} |h|^p)^{1/p} \lesssim \|h\|_{L^p}$ for $p = \log m$. \square

We can now complete the proof of Theorem 2.4.

Proof of Theorem 2.4. The proof is by induction on q . The base case $q = 1$ is given by Theorem A.3. For the induction step, let $q \geq 2$. We start with the lower bound.

If we condition on $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(q-1)}$, and treat Y as a linear chaos, then applying the lower bound in NCK with respect to the randomness of $\mathbf{h}^{(q)}$ yields (see Figure 1)

$$\mathbb{E}_{\mathbf{h}^{(q)}} \|Y\| = \mathbb{E}_{\mathbf{h}^{(q)}} \left\| Y_{[1:q|q+1|q+2]} \right\| \gtrsim \|h\|_{L^1} \left(\left\| Y_{[1:q-1|q,q+1|q+2]} \right\| + \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| \right).$$

Taking expectations and using the induction hypothesis yields

$$\mathbb{E} \|Y\| \gtrsim_q \|h\|_{L^1}^q \left(\sigma \left(\mathcal{A}_{[1:q-1|q,q+1|q+2]} \right) \vee \sigma \left(\mathcal{A}_{[1:q-1|q+1|q,q+2]} \right) \right),$$

and the conclusion follows from Lemma A.1.

The proof of the upper bound follows similarly. The NCK upper bound yields

$$\mathbb{E}_{\mathbf{h}^{(q)}} \|Y\| \lesssim \|h\|_{\psi_2} \log(d)^{\frac{1}{2}} \left(\left\| Y_{[1:q-1|q,q+1|q+2]} \right\| + \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| \right).$$

Taking expectations and using the induction hypothesis, we obtain

$$\mathbb{E} \|Y\| \lesssim_q \|h\|_{\psi_2}^q \log(d)^{\frac{1}{2}} \log(dm+m)^{\frac{q-1}{2}} \left(\sigma \left(\mathcal{A}_{[1:q-1|q,q+1|q+2]} \right) \vee \sigma \left(\mathcal{A}_{[1:q-1|q+1|q,q+2]} \right) \right),$$

where we used that the largest dimension of the intermediate flattenings is at most dm . The conclusion follows from Lemma A.1 and using $\log(d)^{\frac{1}{2}} \log(dm+m)^{\frac{q-1}{2}} \lesssim_q \log(d+m)^{\frac{q}{2}}$. The identical proof yields the variant of the upper bound where $\|h\|_{\psi_2}$ is replaced by $\|h\|_{L^{\log m}}$. \square

Remark A.4. *It is readily verified in the proof that the iterated NCK inequality also remains valid if $\|h\|_{\psi_2}$ is replaced by $C_c \|h\|_{L^{c \log m}}$ for any $c > 0$, where C_c is a constant that depends on c only. This variant will be used below in the proofs of the iterated Rosenthal inequalities.*

A.3. Proof of iterated strong NCK. We start by stating the linear theorem.

Theorem A.5 (Strong Noncommutative Khintchine inequality). *Let $X = \sum_{i \in [m]} h_i A_i$, where h_1, \dots, h_m are i.i.d. copies of a centered random variable h and A_1, \dots, A_m are $d_1 \times d_2$ matrix coefficients (we define $d := d_1 \vee d_2$). Then we have*

$$\mathbb{E} \|X\| \lesssim \|h\|_{\psi_2} \left(\sigma_R(X) + \sigma_C(X) + \log(d)^{\frac{3}{2}} v(X) \right).$$

Proof. We may estimate $\mathbb{E} \|X\| \lesssim \|h\|_{\psi_2} \mathbb{E} \|X_G\|$ as in the proof of Theorem A.3. For the Gaussian random matrix X_G , applying [BBvH23, Corollary 2.2 and Lemma 2.5] yields

$$\mathbb{E} \|X_G\| \lesssim \sigma_R(X) + \sigma_C(X) + \log(d)^{\frac{3}{4}} (\sigma_R(X) \vee \sigma_C(X))^{\frac{1}{2}} v(X)^{\frac{1}{2}},$$

and the conclusion follows by applying Young's inequality to the last term. \square

We can now complete the proof of Theorem 2.5.

Proof of Theorem 2.5. The proof is by induction on q . The base case $q = 1$ is given by Theorem A.5. For the induction step, let $q \geq 2$. If we condition on $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(q-1)}$, and treat Y as a linear chaos, then applying Theorem A.5 with respect to $\mathbf{h}^{(q)}$ yields (see Figure 1)

$$\mathbb{E}_{\mathbf{h}^{(q)}} \|Y\| \lesssim \|h\|_{\psi_2} \left(\left\| Y_{[1:q-1|q,q+1|q+2]} \right\| + \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| + \log(d)^{\frac{3}{2}} \left\| Y_{[1:q-1|q|q+1,q+2]} \right\| \right).$$

We now take the expectation and bound the norm of each intermediate flattening. For the first two terms, we use the induction hypothesis to estimate

$$\begin{aligned} \mathbb{E} \left\| Y_{[1:q-1|q,q+1|q+2]} \right\| &\lesssim_q \|h\|_{\psi_2}^{q-1} \left(\sigma \left(\mathcal{A}_{[1:q-1|q,q+1|q+2]} \right) + \log(dm+m)^{\frac{q+1}{2}} v \left(\mathcal{A}_{[1:q-1|q,q+1|q+2]} \right) \right), \\ \mathbb{E} \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| &\lesssim_q \|h\|_{\psi_2}^{q-1} \left(\sigma \left(\mathcal{A}_{[1:q-1|q+1|q,q+2]} \right) + \log(dm+m)^{\frac{q+1}{2}} v \left(\mathcal{A}_{[1:q-1|q+1|q,q+2]} \right) \right). \end{aligned}$$

For the last term, we use the iterated NCK inequality (Theorem 2.4) to estimate

$$\mathbb{E} \left\| Y_{[1:q-1|q|q+1,q+2]} \right\| \lesssim_q \|h\|_{\psi_2}^{q-1} \log(d^2 \vee m+m)^{\frac{q-1}{2}} \sigma \left(\mathcal{A}_{[1:q-1|q|q+1,q+2]} \right).$$

Here we used that the largest dimension of the first two intermediate flattenings is at most dm and of the last intermediate flattening is at most $d^2 \vee m$. To conclude, it remains to apply Lemma A.1, and to note that all flattenings that appear in the leading terms are of σ -type and that all flattenings that appear in the terms with logarithmic factors are of v -type. \square

A.4. Proof of iterated Rosenthal inequality. We begin by stating the linear theorem. The upper bound follows from the matrix Rosenthal inequality that may be found in [JZ13, MJC⁺14] (see also [BvH24, Example 2.15]). We were unable to locate a reference for the lower bound.

Theorem A.6 (Matrix Rosenthal inequality). *Let $X = \sum_{i \in [m]} h_i A_i$ where h_1, \dots, h_m are i.i.d. copies of a centered unit-variance random variable h , and A_i are $d_1 \times d_2$ matrix coefficients (set $d := d_1 \vee d_2$). Let $\alpha_c(h) := \|h\|_{L^c \log(d+m)}$ for $c > 0$. Then we have*

$$\mathbb{E} \|X\| \lesssim_c \log(d+m)^{\frac{1}{2}} (\sigma_R(X) + \sigma_C(X)) + \alpha_c(h) \log(d+m) r(X)$$

and

$$\mathbb{E} \|X\| \gtrsim \sigma_R(X) + \sigma_C(X) - C_c \alpha_c(h) \log(d+m)^{\frac{1}{2}} r(X),$$

where C_c is a constant that depends only on c .

Proof. The upper bound follows by applying [BvH24, Example 2.15 and Remark 2.1] with $2p = \lfloor c \log(d+m) \rfloor$ (and $\text{Tr}[|M|^p]^{1/p} \lesssim \|M\|$ for any $d \times d$ matrix M and $p \gtrsim \log d$).

For the lower bound, we begin by estimating

$$\begin{aligned} \mathbb{E} \|X\| &\gtrsim \mathbb{E} \left\| \sum_i h_i^2 A_i^\top A_i \right\|^{\frac{1}{2}} + \mathbb{E} \left\| \sum_i h_i^2 A_i A_i^\top \right\|^{\frac{1}{2}} \\ &\geq \sigma_R(X) + \sigma_C(X) - \mathbb{E} \left\| \sum_i (h_i^2 - 1) A_i^\top A_i \right\|^{\frac{1}{2}} - \mathbb{E} \left\| \sum_i (h_i^2 - 1) A_i A_i^\top \right\|^{\frac{1}{2}}, \end{aligned}$$

where the first line follows from the proof of Theorem A.3 and the second line uses the triangle inequality. We can now apply the matrix Rosenthal upper bound to estimate

$$\mathbb{E} \left\| \sum_i (h_i^2 - 1) A_i^\top A_i \right\| \lesssim_c \log(d+m)^{\frac{1}{2}} \sigma_R(X) r(X) + \alpha_c(h)^2 \log(d+m) r(X)^2,$$

where we used $\left\| \sum_i (A_i^\top A_i)^2 \right\| \leq \sigma_R(X)^2 r(X)^2$. Estimating the remaining term similarly, we obtain

$$\mathbb{E} \|X\| \gtrsim \sigma_R(X) + \sigma_C(X) - C_c \log(d+m)^{\frac{1}{4}} (\sigma_R(X) + \sigma_C(X))^{\frac{1}{2}} r(X)^{\frac{1}{2}} - C_c \alpha_c(h) \log(d+m)^{\frac{1}{2}} r(X)$$

and the conclusion follows by Young's inequality. \square

We can now complete the proof of Theorem 2.6.

Proof of Theorem 2.6. We first prove the upper bound. We aim to prove by induction on q that

$$\mathbb{E} \|Y\| \lesssim_{q,c} \log(d+m)^{\frac{q}{2}} \sigma(\mathcal{A}) + \alpha_c(h)^q \log(d+m)^{\frac{q+1}{2}} r(\mathcal{A}),$$

for every $c > 0$, from which the conclusion follows by choosing $c = 1$.

The base case $q = 1$ is given by Theorem A.6. For the induction step, let $q \geq 2$. If we condition on $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(q-1)}$, then applying Theorem A.6 with respect to $\mathbf{h}^{(q)}$ yields (see Figure 1)

$$\begin{aligned} \mathbb{E}_{\mathbf{h}^{(q)}} \|Y\| &\lesssim_c \log(d+m)^{\frac{1}{2}} \left(\left\| Y_{[1:q-1|q,q+1|q+2]} \right\| + \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| \right) \\ &\quad + \alpha_c(h) \log(d+m) \left\| Y_{[1:q-1|q,q+1|q,q+2]} \right\|. \end{aligned}$$

Now note that the largest dimension of the intermediate flattenings that appear above is md . As $\frac{c}{2} \log(md+m) \leq c \log(d+m)$, the induction hypothesis with $c \leftarrow \frac{c}{2}$ and Lemma A.1 yield

$$\begin{aligned} \mathbb{E} \left\| Y_{[1:q-1|q,q+1|q+2]} \right\| &\lesssim_{q,c} \log(md+m)^{\frac{q-1}{2}} \sigma(\mathcal{A}) + \alpha_c(h)^{q-1} \log(md+m)^{\frac{q}{2}} r(\mathcal{A}), \\ \mathbb{E} \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| &\lesssim_{q,c} \log(md+m)^{\frac{q-1}{2}} \sigma(\mathcal{A}) + \alpha_c(h)^{q-1} \log(md+m)^{\frac{q}{2}} r(\mathcal{A}). \end{aligned}$$

On the other hand, the iterated NCK inequality (Theorem 2.4 and Remark A.4) yields

$$\mathbb{E} \left\| Y_{[1:q-1|q,q+1|q,q+2]} \right\| \lesssim_{q,c} \alpha_c(h)^{q-1} \log(md+m)^{\frac{q-1}{2}} \sigma(\mathcal{A}_{[1:q-1|q,q+1|q,q+2]}).$$

Using again Lemma A.1 yields $\sigma(\mathcal{A}_{[1:q-1|q,q+1|q,q+2]}) \leq r(\mathcal{A})$. The proof of the upper bound is readily concluded by combining the above estimates.

The proof of the lower bound is very similar. We first estimate

$$\mathbb{E}_{\mathbf{h}^{(q)}} \|Y\| \gtrsim \left\| Y_{[1:q-1|q,q+1|q+2]} \right\| + \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| - C_c \alpha_c(h) \log(d+m)^{\frac{1}{2}} \left\| Y_{[1:q-1|q,q+1|q,q+2]} \right\|$$

using Theorem A.6. The proof is concluded by lower bounding the expectation of the first two terms using the induction hypothesis and Lemma A.1, and bounding the expectation of the last term by the iterated NCK inequality as in the upper bound. \square

A.5. Proof of iterated strong matrix Rosenthal inequality. We first state the linear theorem.

Theorem A.7 (Strong matrix Rosenthal inequality). *Let $X = \sum_{i \in [m]} h_i A_i$ where h_1, \dots, h_m are i.i.d. copies of a centered unit-variance random variable h , and A_i are $d_1 \times d_2$ matrix coefficients (set $d := d_1 \vee d_2$). Let $\alpha_c(h) := \|h\|_{L^c \log(d+m)}$ for $c > 0$. Then we have*

$$\mathbb{E} \|X\| \lesssim_c \sigma_R(X) + \sigma_C(X) + \alpha_c(h) \log(d+m)^2 v(X).$$

Proof. Combining [BvH24, Theorem 2.9 and Remark 2.1] and [BBvH23, Theorem 2.7 and Lemma 2.5] with $2p = \lfloor c \log(d+m) \rfloor$ (and $\text{Tr}[|M|^p]^{1/p} \lesssim \|M\|$ for any $d \times d$ matrix M and $p \gtrsim \log d$) yields

$$\mathbb{E} \|X\| \lesssim_c \sigma_R(X) + \sigma_C(X) + \log(d+m)^{\frac{3}{4}} (\sigma_R(X) \vee \sigma_C(X))^{\frac{1}{2}} v(X)^{\frac{1}{2}} + \alpha_c(h) \log(d+m)^2 r(X).$$

The conclusion follows using Young's inequality and as $r(X) \leq v(X)$ (Lemma A.2). \square

We can now complete the proof of Theorem 2.7.

Proof of Theorem 2.7. We aim to prove by induction on q that

$$\mathbb{E} \|Y\| \lesssim_{q,c} \sigma(\mathcal{A}) + \alpha_c(h)^q \log(d+m)^{\frac{q+3}{2}} v(\mathcal{A}),$$

for every $c > 0$, from which the conclusion follows by choosing $c = 1$.

The base case $q = 1$ is given by Theorem A.7. For the induction step, let $q \geq 2$. If we condition on $\mathbf{h}^{(1)}, \dots, \mathbf{h}^{(q-1)}$, then applying Theorem A.7 with respect to $\mathbf{h}^{(q)}$ yields (see Figure 1)

$$\mathbb{E}_{\mathbf{h}^{(q)}} \|Y\| \lesssim_c \left\| Y_{[1:q-1|q,q+1|q+2]} \right\| + \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| + \alpha_c(h) \log(d+m)^2 \left\| Y_{[1:q-1|q|q+1,q+2]} \right\|.$$

As in the proof of Theorem 2.6, the induction hypothesis with $c \leftarrow \frac{c}{2}$ and Lemma A.1 yield

$$\begin{aligned} \mathbb{E} \left\| Y_{[1:q-1|q,q+1|q+2]} \right\| &\lesssim_{q,c} \sigma(\mathcal{A}) + \alpha_c(h)^{q-1} \log(md+m)^{\frac{q+2}{2}} v(\mathcal{A}), \\ \mathbb{E} \left\| Y_{[1:q-1|q+1|q,q+2]} \right\| &\lesssim_{q,c} \sigma(\mathcal{A}) + \alpha_c(h)^{q-1} \log(md+m)^{\frac{q+2}{2}} v(\mathcal{A}). \end{aligned}$$

On the other hand, the iterated NCK inequality (Theorem 2.4 and Remark A.4) yields

$$\mathbb{E} \left\| Y_{[1:q-1|q|q+1,q+2]} \right\| \lesssim_{q,c} \alpha_c(h)^{q-1} \log(md+m)^{\frac{q-1}{2}} \sigma \left(\mathcal{A}_{[1:q-1|q|q+1,q+2]} \right).$$

Using again Lemma A.1 yields $\sigma(Y_{[1:q-1|q|q+1,q+2]}) \leq v(\mathcal{A})$, concluding the proof. \square

Remark A.8. *In the strong matrix Rosenthal inequality that appears in the proof of Theorem A.7, the distributional parameter $\alpha_c(h)$ appears only in the term that is controlled by $r(X)$. We simplified this inequality by estimating $r(X) \leq v(X)$. This simplification can be lossy when $r(X) \ll v(X)$, particularly in sparse situations when the parameter $\alpha_c(h)$ may be very large.*

One may hope that exploiting the sharper form of the strong matrix Rosenthal inequality could give rise to an improved form of Theorem 2.7 where the distributional parameter appears only in a term controlled by $r(\mathcal{A})$. It is not possible to iterate the sharper inequality, however, as it is not true in general that $r(\mathcal{A}_{[1:q-1|q|q+1,q+2]})$ can be controlled by $r(\mathcal{A})$.

It is possible to obtain improved chaos inequalities by introducing additional chaos parameters that control such terms, but we do not at present know of a compelling application of such inequalities.

A.6. Norms of flattenings. We first consider chaoses of combinatorial type.

Proof of Proposition 3.4. Given a chaos of combinatorial type (13), its flattenings (3) are

$$\mathcal{A}_{[R|C]} = \sum_{\mathbf{s} \in [S_1] \times \cdots \times [S_p]} \left(\bigotimes_{t \in R} e_{I_t(\mathbf{s})} \right) \otimes \left(\bigotimes_{t \in C} e_{I_t(\mathbf{s})}^\top \right).$$

Using the natural identification $e_{I_t(\mathbf{s})} \simeq e_{I_t(\mathbf{s})_1} \otimes \cdots \otimes e_{I_t(\mathbf{s})_{|I_t|}}$ and permuting the order of tensor products (which corresponds to reordering rows and columns), we obtain

$$\mathcal{A}_{[R|C]} \simeq \sum_{\mathbf{s} \in [S_1] \times \cdots \times [S_p]} \bigotimes_{u \in [p]} \left(e_{s_u}^{\otimes \mu_u} \otimes (e_{s_u}^\top)^{\otimes \nu_u} \right), \quad (33)$$

where μ_u and ν_u denote the number of times the summation index s_u appears as a row or column index respectively in the tensor product. Distributivity yields

$$\mathcal{A}_{[R|C]} \simeq \bigotimes_{u \in [p]} B_u \quad \text{with} \quad B_u = \sum_{s \in [S_u]} \left(e_s^{\otimes \mu_u} \otimes (e_s^\top)^{\otimes \nu_u} \right).$$

In particular, we have

$$\|\mathcal{A}_{[R|C]}\| = \prod_{u \in [p]} \|B_u\|.$$

Now note that, by definition, $\mu_u = 0$ if and only if $u \in \mathcal{R}^c$, and $\nu_u = 0$ if and only if $u \in \mathcal{C}^c$. We can therefore compute using (18) that $\|B_u\| = (\sqrt{S_u})^{1_{u \in \mathcal{R}^c} + 1_{u \in \mathcal{C}^c}}$, and the conclusion follows. \square

To proceed, we need the following.

Lemma A.9. *If M, M' are (real) matrices so that $|M_{i,j}| \leq M'_{i,j}$ for all i, j , then $\|M\| \leq \|M'\|$.*

Proof. Note that $\|M\| = \sup_{\|x\|=\|y\|=1} \sum_{i,j} x_i M_{i,j} x_j \leq \sup_{\|x\|=\|y\|=1} \sum_{i,j} |x_i| |M_{i,j}| |x_j| \leq \|M'\|$. \square

We can now extend the bound to chaoses of nearly combinatorial type.

Proof of Proposition 3.7. Given a chaos of nearly combinatorial type, following the same steps as in the proof of Proposition 3.4, we obtain an analogue of (33):

$$\mathcal{A}_{[R|C]}^f \simeq \sum_{\mathbf{s}} f(\mathbf{s}) \bigotimes_{u \in [p]} \left(e_{s_u}^{\otimes \mu_u} \otimes (e_{s_u}^\top)^{\otimes \nu_u} \right).$$

Define the associated flattening of combinatorial type by replacing $f \leftarrow 1$:

$$\mathcal{A}_{[R|C]} \simeq \sum_{\mathbf{s}} \bigotimes_{u \in [p]} \left(e_{s_u}^{\otimes \mu_u} \otimes (e_{s_u}^\top)^{\otimes \nu_u} \right).$$

Lemma A.9 yields $\|\mathcal{A}_{[R|C]}^f\| \leq \|f\|_\infty \|\mathcal{A}_{[R|C]}\|$, and we conclude using Proposition 3.4. \square

Remark A.10 (Analogue of Proposition 3.7 for intermediate flattenings). *Consider a chaos Y of nearly combinatorial type, and let $Y_{[Z|R|C]}$ be an intermediate flattening as in (10). Then*

$$Y_{[Z|R|C]} \simeq \sum_{\mathbf{s}} f(\mathbf{s}) \left(\prod_{t \in Z} h_{I_t(\mathbf{s})}^{(t)} \right) \bigotimes_{u \in [p]} \left(e_{s_u}^{\otimes \mu_u} \otimes (e_{s_u}^\top)^{\otimes \nu_u} \right).$$

If h is uniformly bounded, then we can argue as in the proof of Proposition 3.7 that

$$\|Y_{[Z|R|C]}\| \leq \|f\|_\infty \|h\|_\infty^{|Z|} \|\mathcal{A}_{[R|C]}\| \leq \|f\|_\infty \|h\|_\infty^{|Z|} \left(\prod_{u \in \mathcal{R}^c} \sqrt{S_u} \right) \left(\prod_{u \in \mathcal{C}^c} \sqrt{S_u} \right),$$

where we used Proposition 3.4 in the second inequality.

Note that while the definitions of the parameters $\sigma(\mathcal{A})$, $v(\mathcal{A})$, $r(\mathcal{A})$ only involved flattenings $\mathcal{A}_{[R|C]}$ with $R \cup C = [q + 2]$, this need not be the case when considering intermediate flattenings. This is not a problem, as neither the definitions nor the arguments in the proof rely on this assumption.

A.7. Menger’s theorem. The following classical result is used in Section 4.3.

Theorem A.11 (Menger’s theorem, [Gö2]). *Let G be a finite graph and $U, V \subseteq V(G)$ be two subsets of vertices. We say S is a $U - V$ vertex separator if all paths from U to V pass through S . Then the minimal size of a $U - V$ vertex separator equals the maximal number of vertex-disjoint paths from U to V that contain exactly one point in U and one point in V .*

It should be noted that U, V need not be disjoint in Theorem A.11. In this case, any vertex in $U \cap V$ defines a path from U to V of length one. Such a point/path must therefore be contained in any vertex separator, and in any maximal collection of disjoint paths as in the theorem statement.

DEPARTMENT OF MATHEMATICS, ETH ZÜRICH, SWITZERLAND
 Email address: `bandeira@math.ethz.ch`

DEPARTMENT OF MATHEMATICS, ETH ZÜRICH, SWITZERLAND
 Email address: `kevin.lucca@ifor.math.ethz.ch`

DEPARTMENT OF MATHEMATICS, ETH ZÜRICH, SWITZERLAND
 Email address: `petar.nizic-nikolac@ifor.math.ethz.ch`

PACM, PRINCETON UNIVERSITY, PRINCETON, NJ 08544, USA
 Email address: `rvan@math.princeton.edu`