

CONCENTRATION OF BOUNDED SPARSE CHAOS AND SPARSE KHATRI–RAO EMBEDDINGS

GUOZHENG DAI, YIYUN HE, KE WANG, AND YIZHE ZHU

ABSTRACT. We establish moment and mixed-tail inequalities for fixed-order decoupled homogeneous chaoses generated by independent, centered, sparse bounded random variables. Our bounds apply to arbitrary real rectangular coefficient tensors and describe the fluctuation scales through weighted slice and partition norms, with a Bennett-type logarithmic improvement in the largest-entry term. As an application, we derive guarantees for sparse Khatri–Rao embeddings that explicitly account for sparsity and input geometry.

1. INTRODUCTION AND MAIN RESULTS

Concentration inequalities quantify how far a random quantity can deviate from its mean or another typical value. They provide tools for controlling estimation error, studying random matrices, and choosing sample sizes in randomized algorithms [35, 10, 43]. For nonlinear statistics, the useful bounds must reflect both the input distributions and the structure of the function being studied. Polynomials of independent variables are an important example: they include subgraph counts in random graphs and many estimators built from random vectors. Concentration and large deviations for random graph polynomials have been studied, among others, by Kim and Vu [29], Chatterjee [12, 13], and Lubetzky and Zhao [36].

Much of the theory of polynomial concentration proceeds through L_r -moment estimates, from which tail bounds follow by Markov’s inequality. For sums of independent variables, Montgomery-Smith [37] obtained tail estimates in the Rademacher case. Gluskin and Kwapien [21] and Hitczenko, Montgomery-Smith, and Oleszkiewicz [26] obtained sharp estimates for symmetric summands with log-concave and log-convex tails, respectively. More generally, Latała [32] established sharp moment estimates for sums of independent nonnegative or symmetric random variables. For quadratic forms in independent centered sub-Gaussian variables, the Hanson–Wright inequality [23, 40] expresses concentration in terms of the Frobenius and operator norms of the coefficient matrix. Further results on the concentration of quadratic forms can be found in [33, 45, 1, 30, 15].

Beyond quadratic forms, Latała [34] obtained sharp moment estimates for Gaussian chaoses of arbitrary order, using tensor partition norms. For symmetric inputs with log-concave tails, Adamczak and Latała [2] obtained sharp bounds for chaoses of order at most three and for exponential chaoses of arbitrary order. Kolesko and Latała [31] developed the corresponding moment theory for log-convex tails. For general polynomials, Adamczak and Wolff [4] developed bounds in terms of derivative tensors, and Götze, Sambale, and Sinulis [22] obtained multilevel concentration inequalities for independent α -sub-exponential inputs. Related moment inequalities for U -statistics and general functions of independent variables can be found in [20, 9, 41]. Alongside these results, concentration inequalities for simple random tensors were established

2020 *Mathematics Subject Classification.* Primary 60E15; Secondary 60G50, 65F55.

Key words and phrases. Polynomial chaos, Bennett inequality, sparse random variable, Khatri–Rao embedding, tensor sketching.

by Vershynin [44] and Bamberger, Krahmer, and Ward [6]. Adamczak, Latała, and Meller [3] obtained moment and tail estimates for Banach-valued Gaussian chaoses.

Sparsity introduces information that bounds based only on uniform or global Orlicz norms can miss. For inputs $X_i = \delta_i \xi_i$, where $\delta_i \sim \text{Bernoulli}(p_i)$, the probabilities p_i should enter the coefficient scales. Dai and Wang [17] incorporated these probabilities into partition norms of expected derivative tensors and obtained moment and tail bounds for general sparse polynomials with α -sub-exponential amplitudes, $0 < \alpha \leq 1$. Their estimates provide the input to the present work. Here we use boundedness to obtain an additional Bennett-type logarithmic correction.

The benefit of adapting to sparsity is already visible for a sum. Let $\delta_1, \dots, \delta_n$ be independent Bernoulli(p) variables, with $0 < p < 1$, and put $S = \sum_i (\delta_i - p)$ and $v = np(1 - p)$. Hoeffding's inequality [27] and Bennett's inequality [7] give, respectively,

$$\mathbb{P}(S \geq t) \leq \exp(-2t^2/n), \quad \mathbb{P}(S \geq t) \leq \exp(-vh(t/v)), \quad t > 0,$$

where $h(u) = (1+u) \log(1+u) - u$. When $p \rightarrow 0$, $np \rightarrow \infty$, and $t \asymp np$, the two decay exponents have orders np^2 and np . Thus Bennett's bound improves the exponent by a factor of order p^{-1} in this sparse regime. This comparison illustrates the importance of the variance scale. Bennett's rate also captures a further improvement at larger deviations: since $h(u) \sim u \log u$ as $u \rightarrow \infty$, its exponent has order $t \log(t/v)$ when $t \gg v$, whereas Bernstein's bound gives an exponent of order t .

For quadratic forms, Zhou [46] initiated the study of the Hanson–Wright inequality for sparse sub-Gaussian random vectors. He, Wang, and Zhu [25] established sparse Hanson–Wright inequalities and obtained a Bennett-type logarithmic gain through a combinatorial moment argument. Further results on sparse quadratic and bilinear forms and related missing-data models can be found in [47, 38, 16]. Our aim is to obtain such a gain for sparse chaoses of arbitrary fixed order. We prove explicit moment and tail bounds for decoupled homogeneous chaoses with arbitrary real coefficient tensors and independent centered bounded inputs satisfying coordinate-dependent second-moment bounds. The new improvement concerns the largest-entry term; the other terms match the weighted slice and partition scales supplied by Dai and Wang [17].

Application. For simplicity, consider order- d tensors with equal mode dimensions $n_s = n$. A dense random projection acts on vectors of length n^d , which can make both storage and multiplication expensive. A Khatri–Rao map represents each row as a tensor product of d short vectors. Sparse factors make this representation cheaper to store and apply. For a rank-one input supplied through its factors, the expected computational cost drops from $O(Mdn)$ with dense factors to $O(Md[np + 1])$, where M is the number of sketch rows and p is the probability that a factor entry is nonzero. The key question is whether this saving per row survives the additional rows needed to preserve norms with the same accuracy and success probability.

Sparse Khatri–Rao sketches have been studied previously. Sun, Guo, Tropp, and Udell [42] introduce tensor random projections and discuss sparse sign factors. Chen and Jin [14] analyze normalized products of Bernoulli and sub-Gaussian variables, including the two-factor sparse-sign construction studied here. They prove finite-set, subspace, and constrained least-squares guarantees. Rakhshan and Rabusseau [39] study Gaussian projections in rank-one, sums-of-rank-one, and tensor-train forms. Ahle and Knudsen [5] establish strong Johnson–Lindenstrauss moment guarantees for tensor-product sketches. For subspaces, Bujanović, Grubišić, Kressner, and Lam [11] treat two Gaussian factors, while Beretta and Musco [8] obtain nearly linear dependence on the subspace dimension for arbitrary fixed tensor order, using independent isotropic factors with uniformly bounded sub-Gaussian norm.

Our contribution is a sharper, input-dependent analysis of these sparse maps. We identify tensors whose energy is sufficiently spread across entries, slices, and singular directions that

sparse factors achieve a dense-scale measurement bound. For two-factor rank-one inputs with no dominant coordinate in either factor, we also show when taking $p = o(1)$ reduces the sufficient computational cost, at the same accuracy and success probability, compared with the bounds of Chen and Jin [14]. Section 2 gives explicit input conditions and flop comparisons.

Beyond randomized numerical linear algebra, we expect our inequalities to be useful for studying subgraph counts in sparse random graphs and hypergraphs, a classical setting for polynomial concentration and upper-tail problems [28, 19, 24]. We leave these developments to future work.

Notation. For an integer $m \geq 1$, write $[m] = \{1, \dots, m\}$, and adopt the convention that empty products equal one. For real numbers a, b , write $a \vee b = \max\{a, b\}$. Unless otherwise specified, the letters $C, c > 0$ denote constants whose values may change from line to line; subscripts indicate their permitted parameter dependence. We write $f \lesssim_\alpha g$ if $f \leq C_\alpha g$, $f \asymp_\alpha g$ if both comparison inequalities hold, and use $O_\alpha(\cdot)$ with the same convention; omitted subscripts indicate universal constants. The notation $f \ll g$ means that f/g is sufficiently small in the stated regime. For a finite array $T = (t_\alpha)$, set $\|T\|_F = \left(\sum_\alpha |t_\alpha|^2\right)^{1/2}$ and $\|T\|_\infty = \max_\alpha |t_\alpha|$. For a matrix $B = (b_{ij})$, let $\|B\|_{\text{op}}$ denote its Euclidean operator norm and set $\|B\|_{2,\infty} = \max_i \left(\sum_j |b_{ij}|^2\right)^{1/2}$. For a random variable Y and $r \geq 1$, write $\|Y\|_{L_r} = (\mathbb{E}|Y|^r)^{1/r}$. Throughout the paper, $\ell(x) = 1 + \log x$ for $x \geq 1$.

1.1. Bennett-type inequalities for chaoses. Fix an order $d \geq 1$ and dimensions $n_1, \dots, n_d$, and write

$$\mathbf{i} = (i_1, \dots, i_d), \quad a_{\mathbf{i}} = a_{i_1, \dots, i_d},$$

and consider the *decoupled homogeneous polynomial chaos of order d* :

$$Z_A(X) := \sum_{\mathbf{i} \in [n_1] \times \dots \times [n_d]} a_{\mathbf{i}} \prod_{s=1}^d X_{i_s}^{(s)}.$$

Here $A = (a_{\mathbf{i}})$ is a deterministic rectangular d -tensor, and $X = (X_i^{(s)})_{s \in [d], i \in [n_s]}$ is a mutually independent family satisfying

$$(1) \quad \mathbb{E}X_i^{(s)} = 0, \quad \left|X_i^{(s)}\right| \leq 1 \text{ a.s.}, \quad \mathbb{E}(X_i^{(s)})^2 \leq p_i^{(s)}, \quad 0 < p_i^{(s)} \leq 1.$$

Here *homogeneous* means that every monomial has degree d , and *decoupled* means that its d factors come from independent input vectors. For example, $\sum_{i,j} a_{ij} X_i^{(1)} X_j^{(2)}$ is a decoupled quadratic form. The numbers $p_i^{(s)}$ are upper bounds on second moments. A basic example is $X_i^{(s)} = \delta_i^{(s)} \varepsilon_i^{(s)}$, where $\delta_i^{(s)} \sim \text{Bernoulli}(p_i^{(s)})$ are Bernoulli random variables and $\varepsilon_i^{(s)}$ are Rademacher random variables.

Before stating the main theorem, we introduce notation of tensor norms consistent with [17, 22]. For $I \subseteq [d]$, write $I^c = [d] \setminus I$ and $\mathbf{i}_I = (i_s)_{s \in I}$. For a given d -tensor A , we use

$$A_{\mathbf{i}_{I^c}} := (a_{\mathbf{i}})_{\mathbf{i}_I}$$

for the slice obtained by fixing the coordinates in I and allowing those in I^c to vary. Thus the subscript records the free coordinates; the fixed coordinates are understood from context. When $I = \emptyset$, the slice is the full tensor A , and the maximum over fixed indices is omitted. For any finite set S , let $\Pi(S)$ be the set of partitions of S , with the convention that $\Pi(\emptyset) = \{\emptyset\}$.

For $U \subseteq [d]$ and $\mathcal{J} = \{J_1, \dots, J_m\} \in \Pi(U)$, write $|\mathcal{J}| = m$ for its number of blocks. If $T = (t_{i_U})$ is a U -tensor, define the *partition norm associated with $\mathcal{J}$* by

$$\|T\|_{\mathcal{J}} := \sup \left\{ \sum_{i_U} t_{i_U} \prod_{q=1}^m x_{i_{J_q}}^{(q)} : \|x^{(q)}\|_2 \leq 1, q \in [m] \right\}.$$

For $U = \emptyset$, this norm is the absolute value of the scalar. The one-block and singleton partitions give the Frobenius and injective multilinear norms, respectively. For a matrix, these are precisely

$$\|A\|_{\{\{1,2\}\}} = \|A\|_F, \quad \|A\|_{\{\{1\},\{2\}\}} = \|A\|_{\text{op}}.$$

More generally, a two-block partition is an operator norm after reshaping the tensor as a matrix. Partition norms thus extend familiar matrix norms by specifying which indices are grouped together.

For $I \subseteq [d]$ and $\mathcal{J} \in \Pi(I^c)$, define the weighted slice norm $M_{I,\mathcal{J}}(A)$ and, for $r \geq 2$, the associated moment scale $\Delta_I(A; r)$ by

$$(2) \quad \begin{aligned} M_{I,\mathcal{J}}(A) &:= \max_{i_I} \left\| \left(a_i \prod_{s \in I^c} \sqrt{p_{i_s}^{(s)}} \right)_{i_{I^c}} \right\|_{\mathcal{J}}, \\ \Delta_I(A; r) &:= \max_{\mathcal{J} \in \Pi(I^c)} r^{|I|+|\mathcal{J}|/2} M_{I,\mathcal{J}}(A). \end{aligned}$$

We omit the dependence on the variance proxies $p_i^{(s)}$ from the notation. In particular, $M_{[d],\emptyset}(A) = \|A\|_{\infty}$, and $\Delta_{[d]}(A; r) = r^d \|A\|_{\infty}$ has the highest order in r .

The definition has a useful interpretation: fix the coordinates in I , multiply each remaining coordinate by its standard-deviation bound $\sqrt{p_i^{(s)}}$, and take a familiar Euclidean or operator norm of the resulting slice. Each fixed coordinate contributes a factor r to the moment estimate; each block of free coordinates contributes $\sqrt{r}$. For a matrix and a common variance bound p for all input coordinates, the four resulting scales are

$$\sqrt{r} p \|A\|_F, \quad r p \|A\|_{\text{op}}, \quad r^{3/2} \sqrt{p} \max \left\{ \|A\|_{2,\infty}, \|A^{\top}\|_{2,\infty} \right\}, \quad r^2 \|A\|_{\infty}.$$

Only the last scale receives the logarithmic correction in the theorem.

Write $m_A = \|A\|_{\infty}$. For $1 \leq k \leq d$, collect the scales with k free coordinates:

$$(3) \quad B_k(A; r) = \max_{\substack{I \subseteq [d] \\ |I^c|=k}} \Delta_I(A; r), \quad \lambda_A(r) = 1 \vee \min_{1 \leq k \leq d} \left(\frac{r^d m_A}{B_k(A; r)} \right)^{2/k}.$$

Also set $B_0(A; r) = \Delta_{[d]}(A; r) = r^d m_A$. The quantity $\lambda_A(r)$ then measures how far the largest-entry contribution $r^d m_A$ exceeds the other associated moment scales.

Theorem 1.1 (Bennett-type moment bound). *For $d \geq 1$, let A be a nonzero real d -tensor, and let the mutually independent family $X = (X_i^{(s)})_{s \in [d], i \in [n_s]}$ satisfy (1). For every $r \geq 2$,*

$$(4) \quad \|Z_A(X)\|_{L_r} \lesssim_d \max_{1 \leq k \leq d} B_k(A; r) + \frac{r^d m_A}{\ell(\lambda_A(r))^d}.$$

The first term of (4) collects the usual weighted slice and partition scales. When $\lambda_A(r) > 1$, the largest-entry contribution dominates, and the bound simplifies to

$$(5) \quad \|Z_A(X)\|_{L_r} \lesssim_d \frac{r^d m_A}{\ell(\lambda_A(r))^d}.$$

Indeed, $\lambda_A(r) > 1$ implies $B_k(A; r) \leq r^d m_A \lambda_A(r)^{-k/2}$, and these terms are absorbed into the right-hand side up to a constant depending only on d . The gain over $r^d m_A$ grows with $\lambda_A(r)$. Next, we state the tail inequality that follows directly from the moment bounds in Theorem 1.1.

Corollary 1.2 (Bennett-type tail bound). *Under the assumptions of Theorem 1.1, write $\kappa_{I,\mathcal{J}} = |I| + |\mathcal{J}|/2$ and, for $t > 0$, set*

$$(6) \quad \eta_A(t) = 1 \vee \min_{\substack{I \subseteq [d] \\ \mathcal{J} \in \Pi(I^c)}} \left[\frac{m_A}{M_{I,\mathcal{J}}(A)} \left(\frac{t}{m_A} \right)^{1-\kappa_{I,\mathcal{J}}/d} \right]^{2/|I^c|}.$$

Then

$$(7) \quad \mathbb{P}(|Z_A(X)| \geq t) \leq 2 \exp \left[-c_d \min \left\{ \min_{\substack{I \subseteq [d] \\ \mathcal{J} \in \Pi(I^c)}} \left(\frac{t}{M_{I,\mathcal{J}}(A)} \right)^{1/\kappa_{I,\mathcal{J}}}, \left(\frac{t}{m_A} \right)^{1/d} \ell(\eta_A(t)) \right\} \right].$$

For Bernoulli random variables times independent bounded amplitudes, the terms $B_k(A; r)$, $1 \leq k \leq d$, match the corresponding scales in the moment estimates of Dai and Wang [17], up to constants depending only on d .

Both Theorem 1.1 and Corollary 1.2 concern decoupled chaoses. Our proof uses the sparse-exponential moment estimate of Dai and Wang in this setting, recorded in Lemma 3.3. The new feature of Theorem 1.1 is the logarithmic improvement of the largest-entry contribution, replacing $r^d m_A$ by $r^d m_A / \ell(\lambda_A(r))^d$, without changing the other scales. Correspondingly, the logarithm in (7) improves the largest-entry rate $(t/m_A)^{1/d}$ when $\eta_A(t)$ is large.

Extensions. Two extensions are recorded in Appendix A. First, standard decoupling [18] transfers the bounded-input estimates, including the logarithmic improvement, to diagonal-free non-decoupled chaoses. This includes the quadratic form below. Second, we extend the Bennett-type moment bounds to sparse α -sub-exponential inputs with $1 < \alpha \leq \infty$, in both decoupled and non-decoupled form. The logarithmic power is determined by $\beta = 1 - 1/\alpha$, with $\alpha = \infty$ recovering the bounded case.

Below, we unpack the main result in the cases $d = 1, 2, 3$. After recalling the familiar degree-one bound, we focus on the quadratic and cubic moment bounds in the regimes where the logarithmic correction is effective, together with explicit tail bounds valid for every $t > 0$. The tail bounds use the thresholds $\gamma_2(A)$ and $\gamma_3(A)$ defined below to express the logarithmic correction without an additional parameter.

The linear case ($d = 1$). For $S = \sum_i a_i X_i$, consider $V^2 := \sum_i a_i^2 p_i > 0$ and $a_{\max} := \max_i |a_i|$. The theorem recovers, up to universal constants, the moment bound implied by Bennett's inequality [7]:

$$(8) \quad \|S\|_{L_r} \lesssim \sqrt{r} V + \frac{r a_{\max}}{\ell(1 \vee r a_{\max}^2 / V^2)}, \quad r \geq 2.$$

The quadratic case ($d = 2$). The following Bennett-type moment and tail bounds give a sparse analogue of the classical Hanson–Wright inequality [40] for bounded inputs. For nonnegative coefficient matrices, the comparison below shows how these bounds refine the corresponding estimates of He, Wang, and Zhu [25].

Corollary 1.3 (Quadratic Bennett bounds). *Let $A \in \mathbb{R}^{n \times n}$ be a nonzero symmetric matrix with zero diagonal entries. Let $X_1, \dots, X_n$ be independent and centered random variables with*

$|X_i| \leq 1$ and $\mathbb{E}X_i^2 \leq p_i \in (0, 1]$. Set $m_A = \|A\|_\infty$, $D_p = \text{diag}(\sqrt{p_1}, \dots, \sqrt{p_n})$, and

$$\gamma_2(A) = \max \left\{ \frac{\|AD_p\|_{2,\infty}^2}{m_A^2}, \left(\frac{\|D_pAD_p\|_F}{m_A} \right)^{2/3}, \frac{\|D_pAD_p\|_{\text{op}}}{m_A} \right\}.$$

For every $r \geq \max\{2, e\gamma_2(A)\}$,

$$(9) \quad \|X^\top AX\|_{L_r} \lesssim \frac{r^2 m_A}{[1 + \log(r/\gamma_2(A))]^2}.$$

Moreover, for every $t > 0$,

$$(10) \quad \mathbb{P} \left(|X^\top AX| \geq t \right) \leq 2 \exp \left[-c \min \left\{ \frac{t^2}{\|D_pAD_p\|_F^2}, \frac{t}{\|D_pAD_p\|_{\text{op}}}, \left(\frac{t}{\|AD_p\|_{2,\infty}} \right)^{2/3}, \sqrt{\frac{t}{m_A}} \ell \left(1 \vee \frac{\sqrt{t/m_A}}{\gamma_2(A)} \right) \right\} \right].$$

For nonnegative A , the two intermediate scales for inputs of the form $X_i = \delta_i \xi_i$, with independent Bernoulli random variables δ_i and centered amplitudes ξ_i bounded by one, improve on the weighted row-sum control in He, Wang, and Zhu [25, Lemma 6.3]. Indeed, with $\gamma = \max_i \sum_{j \neq i} a_{ij} p_j$,

$$\|AD_p\|_{2,\infty}^2 \leq \|A\|_\infty \gamma, \quad \|D_pAD_p\|_{\text{op}} \leq \gamma.$$

The first inequality uses $a_{ij}^2 \leq \|A\|_\infty a_{ij}$; the second follows because D_pAD_p is similar to the nonnegative matrix AD_p^2 , whose spectral radius is at most its largest row sum. Thus both the row-norm and operator-norm terms in $\gamma_2(A)$ are bounded by $\gamma/\|A\|_\infty$, the corresponding parameters in [25]. The Frobenius contribution is unchanged, while the row and operator scales are no larger and can be strictly smaller.

The cubic case ($d = 3$). For simplicity, take $n_1 = n_2 = n_3 = n$ and let $A \in \mathbb{R}^{n \times n \times n}$. Write

$$\|A\|_{\text{inj}} = \sup_{\|x\|_2, \|y\|_2, \|z\|_2 \leq 1} \left| \sum_{i,j,k} a_{ijk} x_i y_j z_k \right|$$

for its injective norm. The random variable of interest is the cubic chaos

$$Z_A(X) = \sum_{i,j,k=1}^n a_{ijk} X_i^{(1)} X_j^{(2)} X_k^{(3)},$$

where all coordinates of the three input vectors $X^{(s)} = (X_i^{(s)})_{i=1}^n$, $s = 1, 2, 3$, are mutually independent, centered, and bounded in absolute value by one. For simplicity, assume all coordinates share a single variance upper bound $p \in (0, 1]$:

$$\text{Var}(X_i^{(s)}) = \mathbb{E}[(X_i^{(s)})^2] \leq p, \quad i = 1, \dots, n, \quad s = 1, 2, 3.$$

Corollary 1.4 (Cubic Bennett bounds). *Suppose that the assumptions of Theorem 1.1 hold with $d = 3$ and $n_1 = n_2 = n_3 = n$, and that $p_i^{(s)} = p \in (0, 1]$ for all $s \in [3]$ and $i \in [n]$. Set $m_A = \|A\|_\infty$ and define*

$$\begin{aligned} V_1 &= \sqrt{p} \max_{s \in [3]} \max_{i_{\{s\}^c}} \|A_{i_{\{s\}}}\|_2, & V_2 &= p \max_{s < u} \max_{i_{\{s,u\}^c}} \|A_{i_{\{s,u\}}}\|_F, \\ W_2 &= p \max_{s < u} \max_{i_{\{s,u\}^c}} \|A_{i_{\{s,u\}}}\|_{\text{op}}, & V_3 &= p^{3/2} \|A\|_F, \\ W_3 &= p^{3/2} \max_{s \in [3]} \|A\|_{\{\{s\}, [3] \setminus \{s\}\}}, & I_3 &= p^{3/2} \|A\|_{\text{inj}}. \end{aligned}$$

Define

$$(11) \quad \gamma_3(A) = \max \left\{ \left(\frac{V_1}{m_A} \right)^2, \left(\frac{V_2}{m_A} \right)^{2/3}, \frac{W_2}{m_A}, \left(\frac{V_3}{m_A} \right)^{2/5}, \left(\frac{W_3}{m_A} \right)^{1/2}, \left(\frac{I_3}{m_A} \right)^{2/3} \right\}.$$

For every $r \geq \max\{2, e\gamma_3(A)\}$,

$$(12) \quad \|Z_A(X)\|_{L_r} \lesssim \frac{r^3 m_A}{[1 + \log(r/\gamma_3(A))]^3}.$$

Moreover, for every $t > 0$,

$$(13) \quad \mathbb{P}(|Z_A(X)| \geq t) \leq 2 \exp \left[-c \min \left\{ \left(\frac{t}{V_1} \right)^{2/5}, \left(\frac{t}{V_2} \right)^{2/3}, \left(\frac{t}{W_2} \right)^{1/2}, \right. \right. \\ \left. \left. \frac{t^2}{V_3^2}, \frac{t}{W_3}, \left(\frac{t}{I_3} \right)^{2/3}, \left(\frac{t}{m_A} \right)^{1/3} \ell \left(1 \vee \frac{(t/m_A)^{1/3}}{\gamma_3(A)} \right) \right\} \right].$$

Here $A_{i_{\{s\}}}$ is a vector obtained by fixing the other two coordinates, while $A_{i_{\{s,u\}}}$ is a matrix obtained by fixing the remaining coordinate. The two-block norm $\|A\|_{\{\{s\}, [3] \setminus \{s\}\}}$ is the operator norm of the matrix with mode s indexing rows and the other two modes indexing columns. Here V_1 measures fibers, V_2, W_2 measure matrix slices, and V_3, W_3, I_3 measure the full tensor.

Proof overview. The proof of Theorem 1.1 compares the original chaos with a larger one whose moments can be estimated using known results. We build this comparison chaos from sums of independent sparse exponential random variables with random signs. Expanding the sums repeats each tensor coefficient without changing its value; we call this construction *replication*.

For a simple example, start with the linear form $aX_1 + bX_2$. Suppose X_1, X_2 are independent and centered, with $|X_j| \leq 1$ and $\mathbb{E}X_j^2 \leq p_j \in (0, 1]$. Define the reference variables by

$$Y_i = \delta_i \varepsilon_i E_i, \quad i = 1, \dots, 4,$$

where $\delta_1, \delta_2 \sim \text{Bernoulli}(p_1)$ and $\delta_3, \delta_4 \sim \text{Bernoulli}(p_2)$ are Bernoulli random variables, the ε_i take the values ± 1 with equal probability, and $E_i \sim \text{Exp}(1)$ have mean one. All these factors are mutually independent. Thus each Y_i is zero or a randomly signed exponential random variable. We compare $aX_1 + bX_2$ with

$$a(Y_1 + Y_2) + b(Y_3 + Y_4).$$

The coefficient vector changes from (a, b) to (a, a, b, b) . Its largest absolute entry is unchanged, whereas its Euclidean norm grows by a factor of $\sqrt{2}$.

Why can this help? More generally, for an integer $N \geq 1$, take N copies of Y_1 for the first input and N copies of Y_3 for the second, with all copies mutually independent. Write S_1, S_2 for these two sums. The estimate follows in two steps. First, for an even integer $r \geq 2$, the moment comparison in Lemma 3.2 gives

$$\|aX_1 + bX_2\|_{L_r} \leq \frac{2\sqrt{e}}{1 + \log N} \|aS_1 + bS_2\|_{L_r}.$$

Second, $aS_1 + bS_2$ is a sum of independent random variables with variance $2N(a^2 p_1 + b^2 p_2)$. The standard Bernstein moment estimate, also the degree-one case of Lemma 3.3, gives

$$\|aS_1 + bS_2\|_{L_r} \leq C \left[\sqrt{Nr(a^2 p_1 + b^2 p_2)} + r \max\{|a|, |b|\} \right].$$

Combining the two bounds and rounding r up to the next even integer gives, for every $r \geq 2$,

$$\|aX_1 + bX_2\|_{L_r} \leq C \frac{\sqrt{Nr(a^2p_1 + b^2p_2)} + r \max\{|a|, |b|\}}{1 + \log N}.$$

For $(a, b) \neq (0, 0)$, the choice $N = \left\lceil \frac{r \max\{a^2, b^2\}}{a^2p_1 + b^2p_2} \right\rceil$ balances the two numerator terms up to a universal constant and gives

$$\|aX_1 + bX_2\|_{L_r} \leq \frac{Cr \max\{|a|, |b|\}}{1 + \log \left(\frac{r \max\{a^2, b^2\}}{a^2p_1 + b^2p_2} \right)}, \quad r \geq 2.$$

Since the argument of the logarithm is at least one, this makes the logarithmic improvement of the largest-entry term explicit. For an order- d chaos, using N copies in each mode gives the divisor $(1 + \log N)^d$, with the cost of repetition measured by the weighted slice and partition norms; the details are given in Section 3.

2. SPARSE KHATRI–RAO EMBEDDINGS

We apply our concentration inequalities to sparse Khatri–Rao sketches. Sparsity reduces the cost per measurement for rank-one inputs, but may require more measurements. Our input-dependent bounds identify when this tradeoff lowers the total computational cost at fixed accuracy and success probability.

2.1. Construction and computational cost. Throughout this section, take $n_s = n$ in every mode and write $N = n^d$. We identify $\mathbb{R}^N$ with $\mathbb{R}^{n \times \dots \times n}$, so vector Euclidean norms are tensor Frobenius norms. Our goal is a random map $\Phi \in \mathbb{R}^{M \times N}$ that approximately preserves the norm using $M \ll N$ measurements. For a fixed pair of tensors x, y , linearity gives $\Phi x - \Phi y = \Phi(x - y)$. Applying a relative-error guarantee for squared norms to $x - y$ therefore preserves their squared distance, with the same success probability:

$$(1 - \varepsilon) \|x - y\|_{\mathbb{F}}^2 \leq \|\Phi x - \Phi y\|_2^2 \leq (1 + \varepsilon) \|x - y\|_{\mathbb{F}}^2.$$

A *Khatri–Rao sketch* represents each row as a tensor product of short vectors. This representation avoids forming a length- n^d row and, for a rank-one input, reduces each measurement to a product of d short inner products. The same idea is used in the tensor random projections of Sun, Guo, Tropp, and Udell [42] and the tensor-structured sketches of Chen and Jin [14]. We make the factors sparse to further reduce storage and arithmetic work. Fix a common sampling probability $p \in (0, 1]$ and a number of sketch rows $M \geq 1$. For $s \in [d]$, $m \in [M]$, and $i \in [n]$, let the variables $\xi_{mi}^{(s)}$ be mutually independent, with

$$\mathbb{P}\{\xi_{mi}^{(s)} = 1\} = \mathbb{P}\{\xi_{mi}^{(s)} = -1\} = \frac{p}{2}, \quad \mathbb{P}\{\xi_{mi}^{(s)} = 0\} = 1 - p.$$

Write $\xi_m^{(s)} = (\xi_{mi}^{(s)})_{i=1}^n \in \mathbb{R}^n$. The sketching matrix $\Phi \in \mathbb{R}^{M \times N}$ has rows

$$\Phi_{m,:} = \frac{1}{\sqrt{Mp^d}} (\xi_m^{(1)} \otimes \dots \otimes \xi_m^{(d)})^\top, \quad m \in [M].$$

On an order- d tensor $x = (x_i)$, this map acts by

$$(14) \quad (\Phi x)_m = \frac{1}{\sqrt{Mp^d}} \sum_{i \in [n]^d} x_i \prod_{s=1}^d \xi_{m i_s}^{(s)}, \quad m \in [M].$$

We represent Φ through its Md many length- n factor vectors, without forming the full sketching matrix. These vectors have Mdn coordinates in total, of which $Mdnp$ are nonzero on average.

For a rank-one input $x = u^{(1)} \otimes \cdots \otimes u^{(d)}$, we can apply the map without forming its full rows, since

$$(\Phi x)_m = \frac{1}{\sqrt{Mp^d}} \prod_{s=1}^d \left(\sum_{i=1}^n \xi_{mi}^{(s)} u_i^{(s)} \right).$$

Each inner product uses only sampled coordinates, giving the expected computational cost

$$(15) \quad O(Md(np+1)) \quad \text{flops.}$$

For $np \geq 1$, sparsity saves a factor of order p per row; the next result controls how many rows are needed.

2.2. Input geometry and embedding guarantees. The normalization in (14) gives

$$(16) \quad \mathbb{E} \|\Phi x\|_2^2 = \|x\|_F^2.$$

The correct mean alone does not ensure concentration: a large entry of x may be sampled rarely and then make a large contribution. Our bound accounts for this through the *input geometry*: the fractions of the squared Frobenius norm carried by one entry, one slice, or one singular direction of a matrix flattening of a slice. These are properties of the fixed input, unchanged by rescaling; they contain more information than its number of nonzero entries.

We use the largest-entry, slice Frobenius, and flattening operator norms to measure these features. Recall that x_{i_U} is the slice with the modes in U free and the other coordinates fixed. For a nonempty proper subset $V \subset U$, its two-block norm $\|x_{i_U}\|_{\{V, U \setminus V\}}$ is the operator norm of the matrix with modes in V indexing rows and those in $U \setminus V$ indexing columns.

For fixed tensor $x \neq 0$ and failure probability $0 < \delta < 1$, set

$$q := 2 \log(e^2/\delta).$$

Then $q \asymp \log(2/\delta)$, and $\|\cdot\|_{L_q}$ will be the highest moment norm for each row of the sketching matrix that we need to control. Define the entry scale a_0 and the k -mode slice scales a_k , $1 \leq k \leq d$, by

$$a_0 = \frac{q^{2d-1} \|x\|_\infty^2}{p^d},$$

$$a_k = \left(\frac{q^2}{p} \right)^{d-k} \max_{\substack{U \subseteq [d] \\ |U|=k}} \left\{ \max_{i_{U^c}} \|x_{i_U}\|_F^2 + q^{k-1} \max_{\substack{i_{U^c} \\ \emptyset \neq V \subsetneq U}} \|x_{i_U}\|_{\{V, U \setminus V\}}^2 \right\}.$$

The second term is zero for $k = 1$; no coordinates are fixed when $k = d$. Suppressing the dependence of a_k on x, δ, p , set

$$(17) \quad \rho_x(\delta) = 1 \vee \min_{1 \leq k \leq d} \left(\frac{a_0}{a_k} \right)^{1/k},$$

$$G_x(\delta) = \frac{1}{\|x\|_F^2} \max \left\{ a_1, \dots, a_d, \frac{a_0}{[1 + \log \rho_x(\delta)]^{2d}} \right\}.$$

The dependence of $G_x(\delta)$ and $\rho_x(\delta)$ on the sampling probability p is suppressed in the notation. The logarithmic divisor reduces the largest-entry contribution; the remaining terms measure concentration of energy in slices and singular directions.

The central quantity is $G_x(\delta)$: it summarizes the input geometry at sampling probability p and failure probability δ . Smaller $G_x(\delta)$ means fewer measurements suffice.

For $d = 2$, write the input as a matrix X . The scales then involve only familiar matrix norms:

$$\begin{aligned} a_0 &= \frac{q^3 \|X\|_\infty^2}{p^2}, \\ a_1 &= \frac{q^2}{p} \max \left\{ \|X\|_{2,\infty}^2, \|X^\top\|_{2,\infty}^2 \right\}, \\ a_2 &= \|X\|_F^2 + q \|X\|_{\text{op}}^2. \end{aligned}$$

Here the entry and row/column norms measure how much energy can lie in one entry, row, or column. The ratio $\|X\|_{\text{op}}^2 / \|X\|_F^2$ is the fraction of the total energy carried by the largest singular value.

The following proposition gives the guarantee for one tensor and for all tensors in a fixed subspace simultaneously.

Proposition 2.1 (Sparse Khatri–Rao embedding). *Let $0 < \delta < 1$. For every nonzero tensor x ,*

$$(18) \quad \mathbb{P} \left(\left| \frac{\|\Phi x\|_2^2}{\|x\|_F^2} - 1 \right| > C_d G_x(\delta) \left\{ \sqrt{\frac{\log(2/\delta)}{M}} + \frac{\log(2/\delta)}{M} \right\} \right) \leq \delta.$$

Consequently, for $0 < \varepsilon \leq 1$,

$$(19) \quad M \geq C_d G_x(\delta)^2 \varepsilon^{-2} \log \frac{2}{\delta}$$

implies

$$\mathbb{P} \left(\left| \frac{\|\Phi x\|_2^2}{\|x\|_F^2} - 1 \right| > \varepsilon \right) \leq \delta.$$

Let W be a fixed k -dimensional linear subspace of $\mathbb{R}^{n \times \dots \times n}$, and suppose that

$$\sup_{x \in W \setminus \{0\}} G_x \left(\frac{\delta}{9^k} \right) \leq G.$$

Then, for $0 < \varepsilon \leq 1$, $M \geq C_d G^2 \varepsilon^{-2} (k + \log \frac{2}{\delta})$ implies, with probability at least $1 - \delta$,

$$(1 - \varepsilon) \|x\|_F^2 \leq \|\Phi x\|_2^2 \leq (1 + \varepsilon) \|x\|_F^2, \quad \forall x \in W.$$

For the subspace guarantee, the argument $\delta/9^k$ accounts for a union bound over a $1/4$ -net of the unit sphere of W , with at most 9^k points.

2.3. Inputs that tolerate sparsification. Proposition 2.1 gives a smaller sufficient row count when $G_x(\delta)$ is smaller. Recall that $q = 2 \log(e^2/\delta)$ is the highest order moment we consider. The first two examples keep $G_x(\delta)$ bounded by spreading energy across coordinates and singular directions. The final, order-three example shows when the logarithmic correction improves the sufficient row count.

Spread-out matrices. For a nonzero matrix X , it is enough to have

$$\frac{\|X\|_\infty^2}{\|X\|_F^2} \leq \frac{p^2}{q^3}, \quad \frac{\max\{\|X\|_{2,\infty}^2, \|X^\top\|_{2,\infty}^2\}}{\|X\|_F^2} \leq \frac{p}{q^2}, \quad \frac{\|X\|_{\text{op}}^2}{\|X\|_F^2} \leq \frac{1}{q}.$$

These conditions limit the energy in each entry, row or column, and singular direction. They give $a_0, a_1 \leq \|X\|_F^2$ and $a_2 \leq 2 \|X\|_F^2$, so

$$G_X(\delta) \leq 2, \quad M \gtrsim \varepsilon^{-2} \log(2/\delta) \quad \text{is sufficient.}$$

The entrywise condition is sufficient but not necessary: the logarithmic correction can accommodate larger entries.

A concrete example is $X = H/n$, where n is a power of two and $H \in \{-1, 1\}^{n \times n}$ is a Sylvester Hadamard matrix, so $HH^\top = nI$. Then

$$\|X\|_F^2 = 1, \quad \|X\|_\infty^2 = \frac{1}{n^2}, \quad \|X\|_{2,\infty}^2 = \|X^\top\|_{2,\infty}^2 = \|X\|_{\text{op}}^2 = \frac{1}{n}.$$

All three conditions hold whenever $np \geq q^2$. At fixed δ , we may therefore take $p = q^2/n \rightarrow 0$ without increasing the sufficient row-count scale.

Diagonal tensors and subspaces. For an example valid for every $d \geq 2$, let $x_{j,\dots,j} = h^{-1/2}$, $1 \leq j \leq h$, and all other entries be zero, where $h \leq n$. Then $\|x\|_F = 1$, every proper slice has at most one nonzero entry, and every nontrivial matrix flattening has operator norm $h^{-1/2}$. Substitution gives

$$hp^d \geq q^{2d-1} \implies \rho_x(\delta) = 1, \quad G_x(\delta) = 1 + \frac{q^{d-1}}{h} < 2.$$

Here hp^d is the expected number of sampled diagonal entries per row. Under this condition, the sufficient row count remains $M = O_d(\varepsilon^{-2} \log(2/\delta))$, even if $p \rightarrow 0$.

For $n \geq kh$, let e_i denote the i -th standard basis vector in $\mathbb{R}^n$, and define

$$x^{(\ell)} = \frac{1}{\sqrt{h}} \sum_{j=1}^h e_{(\ell-1)h+j}^{\otimes d}, \quad \ell \in [k].$$

The tensor $x^{(\ell)}$ uses only the coordinates $(\ell-1)h+1, \dots, \ell h$ in every mode. These tensors are orthonormal, so $W = \text{span}\{x^{(1)}, \dots, x^{(k)}\}$ has dimension k . Suppose that

$$hp^d \geq (q + 2k \log 9)^{2d-1}.$$

Every unit tensor $x = \sum_{\ell=1}^k c_\ell x^{(\ell)} \in W$ satisfies $\sum_{\ell=1}^k c_\ell^2 = 1$, hence $|c_\ell| \leq 1$. Its entries, proper slices, and nontrivial matrix flattenings therefore satisfy the same norm bounds as above, giving $G_x(\delta/9^k) < 2$. Proposition 2.1 gives a subspace embedding with probability at least $1 - \delta$ and

$$M = O_d(\varepsilon^{-2} [k + \log(2/\delta)]).$$

Overlapping-window tensors: a logarithmic gain. The preceding examples have $\rho_x(\delta) = 1$. For a genuinely coupled order-three example, take $K = \lceil q^4 \rceil$ and any integer $b \geq 2$. Set $n = Kb$, $p = 1/b$, and let $I_i = \{i, \dots, i+b-1\}$, with indices interpreted cyclically modulo n . Define

$$x_{ijk} = \frac{\mathbf{1}_{\{j,k \in I_i\}}}{b\sqrt{n}}, \quad \|x\|_F = 1.$$

Each fixed first coordinate supports a $b \times b$ window; overlapping windows couple the three modes. The maximum squared fiber norm is $1/(bn)$, and the maximum squared slice Frobenius and operator norms are both $1/n$. To compute the full-flattening norms, first put the first mode in the rows. The Gram matrix has entries $|I_i \cap I_{i'}|^2 / (b^2 n)$ and constant row sum

$$\frac{b^2 + 2 \sum_{t=1}^{b-1} t^2}{b^2 n} = \frac{2b^2 + 1}{3bn}.$$

Putting either of the other modes in the rows gives a Gram matrix with constant row sum b/n . These Gram matrices are entrywise nonnegative, so their operator norms equal their row sums. The maximum squared operator norm of a full flattening is therefore $b/n = 1/K$. Hence

$$(a_0, a_1, a_2, a_3) = \frac{1}{K} (q^5, q^4, q^2(1+q), K+q^2).$$

In particular, no full flattening concentrates more than a fraction $1/K$ of the energy in one singular direction. The scales above give

$$\rho_x(\delta) = \left(\frac{q^5}{K + q^2} \right)^{1/3} \asymp q^{1/3}, \quad G_x(\delta) \asymp \max\{1, q/\ell(q)^6\}.$$

Proposition 2.1 therefore guarantees relative squared-norm error at most ε with probability at least $1 - \delta$ whenever

$$(20) \quad M \geq C\varepsilon^{-2} \max\{q, q^3/\ell(q)^{12}\}.$$

Without the logarithmic correction, the entry scale would give $G_x(\delta) \asymp q$ and the sufficient row bound $M \geq C\varepsilon^{-2}q^3$. Thus the correction improves this row bound by a factor of order $\ell(q)^{12}$ as $q \rightarrow \infty$, or equivalently as $\delta \downarrow 0$. At fixed δ , increasing b makes $p = 1/b$ smaller without changing $\rho_x(\delta)$ or $G_x(\delta)$; the logarithmic gain grows with the confidence requirement, not with this decrease in p .

2.4. Comparison of computational costs. Fix an order $d \geq 2$, a mode dimension $n \geq 2$, a deterministic nonzero rank-one input $x = u^{(1)} \otimes \cdots \otimes u^{(d)}$, and $0 < \varepsilon \leq 1$, $0 < \delta < 1$. Throughout this comparison, assume $np \geq 1$. The common target is

$$(21) \quad \mathbb{P} \left(\left| \frac{\|\Phi x\|_2^2}{\|x\|_F^2} - 1 \right| \leq \varepsilon \right) \geq 1 - \delta.$$

We compare sufficient computational cost bounds for this fixed input. In the cost model above, each sparse row uses $O(dnp)$ expected flops. Combining this with (19) gives our sufficient cost $O_d(npG_x(\delta)^2\varepsilon^{-2}\log(2/\delta))$.

Chen–Jin [14]: the same sparse construction. For this comparison, take $d = 2$ and write $x = u \otimes v$. The proof of Chen–Jin’s Theorem 2.1 [14], applied to the singleton $\{x\}$, gives (21) when

$$(22) \quad M \gtrsim \max \left\{ \frac{\log(2/\delta)}{\varepsilon^2 p^4}, \frac{\log^2(2/\delta)}{\varepsilon p^2} \right\}.$$

Multiplying the two row-count bounds by the per-row cost gives, up to universal constants,

$$(23) \quad \begin{aligned} \text{Our bound at } p : & \quad \frac{npG_x(\delta)^2}{\varepsilon^2} \log \frac{2}{\delta}, \\ \text{Chen–Jin at the same } p : & \quad n \max \left\{ \frac{\log(2/\delta)}{\varepsilon^2 p^3}, \frac{\log^2(2/\delta)}{\varepsilon p} \right\}. \end{aligned}$$

Both terms in their sparse budget increase as p decreases, so that bound does not certify a benefit from sparsification. In contrast, at fixed ε and δ , our sufficient flop bound is asymptotically smaller than the dense benchmark whenever $pG_x(\delta)^2 = o(1)$, for example when $G_x(\delta) = O(1)$ and $p = o(1)$.

Dense fixed-input benchmark. For the fixed-input problem, we use the bound from earlier work recalled in [8, equation (1)]; see also [5, Theorem 2] for dense Rademacher factors. Multiplying its row count by the $O(dn)$ cost per dense row gives the sufficient computational cost

$$O_d \left(n \left[\varepsilon^{-2} \log \frac{2}{\delta} + \varepsilon^{-1} \log^d \frac{2}{\delta} \right] \right).$$

For $d = 2$, this agrees, up to constants, with the dense specialization $p = 1$ in [14]. Dividing our sufficient cost scale by the dense benchmark gives, up to constants depending only on d ,

$$\frac{pG_x(\delta)^2}{1 + \varepsilon \log^{d-1}(2/\delta)}.$$

Consequently, as $n \rightarrow \infty$, the condition

$$(24) \quad pG_x(\delta)^2 = o\left(1 + \varepsilon \log^{d-1}(2/\delta)\right)$$

certifies an asymptotically smaller budget than the dense benchmark above.

3. PROOFS OF THE MAIN RESULTS

3.1. The replication estimate. Recall the weighted norms $M_{I,\mathcal{J}}(A)$, the moment scales $\Delta_I(A; r)$ from (2), and $\ell(N) = 1 + \log N$.

Proposition 3.1 (Replication bound). *Let A be a real d -tensor and let the mutually independent inputs satisfy (1). For every $r \geq 2$ and $N \geq 1$,*

$$(25) \quad \|Z_A(X)\|_{L_r} \lesssim_d \frac{\max_{I \subseteq [d]} N^{|I^c|/2} \Delta_I(A; r)}{\ell(N)^d}.$$

Consequently, for every $t > 0$,

$$(26) \quad \mathbb{P}(|Z_A(X)| \geq t) \leq 2 \exp \left[-c_d \min_{\substack{I \subseteq [d] \\ \mathcal{J} \in \Pi(I^c)}} \left(\frac{t \ell(N)^d}{N^{|I^c|/2} M_{I,\mathcal{J}}(A)} \right)^{1/\kappa_{I,\mathcal{J}}} \right], \quad \kappa_{I,\mathcal{J}} = |I| + |\mathcal{J}|/2.$$

The parameter N balances two effects. Replication enlarges a norm with k free coordinates by $N^{k/2}$, but the moment comparison gains the divisor $\ell(N)^d$. In particular, the largest-entry term, for which $k = 0$, is not enlarged. We first prove the estimate for integer N , then extend it by rounding; Section 3.3 chooses N to obtain the main results.

3.2. Proof of the replication estimate. The argument combines a coordinatewise moment comparison, a known bound for sparse-exponential chaoses, and an exact calculation of how replication changes the weighted norms.

3.2.1. A scalar comparison and its multilinear consequence. Let ε be a Rademacher sign and $E \sim \text{Exp}(1)$, independent of each other. The symmetric exponential variable $\zeta = \varepsilon E$ satisfies

$$\mathbb{E} \zeta^m = 0 \quad (m \text{ odd}), \quad \mathbb{E} \zeta^m = m! \quad (m \text{ even}), \quad \mathbb{E} \exp(|\zeta|/2) = 2.$$

For an integer $N \geq 1$, replace each mode- s coordinate i_s by $(i_s, u_s) \in [n_s] \times [N]$, and define

$$a_{(i_1, u_1), \dots, (i_d, u_d)}^{[N]} = a_{i_1, \dots, i_d}, \quad Y_{i,u}^{(s)} = \delta_{i,u}^{(s)} \varepsilon_{i,u}^{(s)} E_{i,u}^{(s)}, \quad \delta_{i,u}^{(s)} \sim \text{Bernoulli}(p_i^{(s)}).$$

All Bernoulli variables, signs, and exponential variables are mutually independent. The replicated tensor $A^{[N]}$ has dimensions $n_1 N \times \dots \times n_d N$; each replica inherits the original probability $p_i^{(s)}$.

Lemma 3.2 (Replicated moment comparison). *Under the assumptions of Proposition 3.1, for every even integer $r \geq 2$ and integer $N \geq 1$,*

$$\|Z_A(X)\|_{L_r} \leq \frac{(2\sqrt{e})^d}{\ell(N)^d} \|Z_{A^{[N]}}(Y)\|_{L_r}.$$

Proof. Step 1: symmetrize the bounded inputs. Let X' be an independent copy of the entire input family and put $V_i^{(s)} = (X_i^{(s)} - X_i'^{(s)})/2$. These coordinates are independent and symmetric, and

$$|V_i^{(s)}| \leq 1, \quad \mathbb{E}(V_i^{(s)})^2 = \frac{1}{2} \mathbb{E}(X_i^{(s)})^2 \leq p_i^{(s)}.$$

For each monomial, independence and centering give

$$\mathbb{E}_{X'} \prod_{s=1}^d (X_{i_s}^{(s)} - X_{i_s}'^{(s)}) = \prod_{s=1}^d X_{i_s}^{(s)}.$$

Thus $Z_A(X) = \mathbb{E}_{X'} Z_A(X - X')$. Conditional Jensen's inequality, followed by degree- d homogeneity, yields

$$\mathbb{E}|Z_A(X)|^r \leq \mathbb{E}|Z_A(X - X')|^r = 2^{dr} \mathbb{E}|Z_A(V)|^r.$$

Step 2: compare the scalar moments. Put $S_i^{(s)} = \sum_{u=1}^N Y_{i,u}^{(s)}$. For every even integer $m \geq 2$, independence gives the multinomial expansion

$$\mathbb{E}(S_i^{(s)})^m = \sum_{\substack{m_1 + \dots + m_N = m \\ m_u \geq 0}} \frac{m!}{m_1! \dots m_N!} \prod_{u=1}^N \mathbb{E}(Y_{i,u}^{(s)})^{m_u}.$$

Terms with an odd m_u vanish, and all remaining terms are nonnegative. Keeping the N terms with one exponent equal to m and all others zero gives

$$\mathbb{E}(S_i^{(s)})^m \geq N p_i^{(s)} m! \geq \frac{p_i^{(s)} \ell(N)^m}{e}.$$

The second inequality follows from the exponential series: $\exp(\ell(N)) = eN \geq \ell(N)^m/m!$. On the other hand, boundedness gives $\mathbb{E}|V_i^{(s)}|^m \leq \mathbb{E}(V_i^{(s)})^2 \leq p_i^{(s)}$. Since $e^{m/2-1} \geq 1$, we obtain

$$(27) \quad \mathbb{E}(V_i^{(s)})^m \leq \mathbb{E} \left(\frac{\sqrt{e} S_i^{(s)}}{\ell(N)} \right)^m \quad (m \geq 2 \text{ even}).$$

Both sides have vanishing odd moments.

Step 3: replace the coordinates one at a time. Take the comparison variables independent of V . After conditioning on every coordinate except one, the chaos is $a + bv$, where a, b are fixed real numbers under the conditioning. For symmetric v and even r ,

$$\mathbb{E}_v |a + bv|^r = \sum_{\substack{0 \leq m \leq r \\ m \text{ even}}} \binom{r}{m} a^{r-m} b^m \mathbb{E} v^m.$$

Both m and $r-m$ are even, so every coefficient is nonnegative. Consequently, (27) allows $v = V_i^{(s)}$ to be replaced by $\sqrt{e} S_i^{(s)}/\ell(N)$ without decreasing the conditional r -th moment. Integrate over the other coordinates and repeat for the entire finite input family. Every monomial contains exactly d coordinates, hence

$$\|Z_A(V)\|_{L_r} \leq \frac{e^{d/2}}{\ell(N)^d} \|Z_A(S)\|_{L_r}.$$

Step 4: identify the replicated chaos. Distributing the sums gives the exact identity

$$Z_A(S) = \sum_{\mathbf{i}} a_{\mathbf{i}} \prod_{s=1}^d \left(\sum_{u_s=1}^N Y_{i_s, u_s}^{(s)} \right) = \sum_{\mathbf{i}} \sum_{\mathbf{u} \in [N]^d} a_{\mathbf{i}} \prod_{s=1}^d Y_{i_s, u_s}^{(s)} = Z_{A^{[N]}}(Y).$$

Combining this with Steps 1 and 3 proves the lemma. $\square$

3.2.2. *The reference bound and the cost of replication.* The following specialization of Dai and Wang [17, Theorem 1] estimates the comparison chaos.

Lemma 3.3 (Sparse-exponential chaos). *Let $W_i^{(s)} = \delta_i^{(s)} \xi_i^{(s)}$, where all Bernoulli variables and amplitudes are mutually independent, $\delta_i^{(s)} \sim \text{Bernoulli}(p_i^{(s)})$, $\mathbb{E} \xi_i^{(s)} = 0$, and $\mathbb{E} \exp(|\xi_i^{(s)}|) / K_{\text{exp}} \leq 2$. For every real d -tensor A and every $r \geq 2$,*

$$(28) \quad \|Z_A(W)\|_{L_r} \leq C_d K_{\text{exp}}^d \max_{I \subseteq [d]} \Delta_I(A; r).$$

Lemma 3.4 (Exact norm scaling). *For integer $N \geq 1$, evaluate the norms of $A^{[N]}$ using the inherited probabilities $p_{(i,u)}^{[N],(s)} = p_i^{(s)}$. Then*

$$M_{I,\mathcal{J}}(A^{[N]}) = N^{|I^c|/2} M_{I,\mathcal{J}}(A), \quad \Delta_I(A^{[N]}; r) = N^{|I^c|/2} \Delta_I(A; r).$$

Proof. Fix I , the values of $\mathbf{i}_I$, and $\mathcal{J} \in \Pi(I^c)$. Write $U = I^c$ and consider the weighted slice

$$b_{\mathbf{i}_U} = a_{\mathbf{i}} \prod_{s \in U} \sqrt{p_{i_s}^{(s)}}.$$

The values of the fixed replica indices $\mathbf{u}_I$ do not affect this slice. If $U = \emptyset$, its norm is the absolute value of one entry, which is unchanged. Assume $U \neq \emptyset$.

For each block $J \in \mathcal{J}$, take a test array $x_{\mathbf{i}_J, \mathbf{u}_J}^{(J)}$ with Euclidean norm at most one, and set

$$y_{\mathbf{i}_J}^{(J)} = N^{-|J|/2} \sum_{\mathbf{u}_J \in [N]^J} x_{\mathbf{i}_J, \mathbf{u}_J}^{(J)}.$$

Cauchy–Schwarz, applied to the $N^{|J|}$ replica indices, gives

$$\sum_{\mathbf{i}_J} |y_{\mathbf{i}_J}^{(J)}|^2 \leq \sum_{\mathbf{i}_J, \mathbf{u}_J} |x_{\mathbf{i}_J, \mathbf{u}_J}^{(J)}|^2 \leq 1.$$

Because the blocks partition U , direct expansion yields

$$\sum_{\mathbf{i}_U, \mathbf{u}_U} b_{\mathbf{i}_U} \prod_{J \in \mathcal{J}} x_{\mathbf{i}_J, \mathbf{u}_J}^{(J)} = N^{|U|/2} \sum_{\mathbf{i}_U} b_{\mathbf{i}_U} \prod_{J \in \mathcal{J}} y_{\mathbf{i}_J}^{(J)}.$$

Taking the supremum over admissible $x^{(J)}$ proves that the replicated partition norm is at most $N^{|U|/2} \|b\|_{\mathcal{J}}$. Conversely, for any admissible $y^{(J)}$, take

$$x_{\mathbf{i}_J, \mathbf{u}_J}^{(J)} = N^{-|J|/2} y_{\mathbf{i}_J}^{(J)}.$$

These arrays have the same Euclidean norms as the corresponding $y^{(J)}$, and the preceding identity gives the reverse inequality. Hence the norm scales exactly by $N^{|U|/2}$. Maximizing over the fixed original and replica indices proves the assertion for $M_{I,\mathcal{J}}$; maximizing over partitions after multiplying by $r^{\kappa_{I,\mathcal{J}}}$ proves the assertion for Δ_I . $\square$

3.2.3. *Completion of the replication estimate.*

Proof of Proposition 3.1. Step 1: even moments and integer replication. Apply the preceding three lemmas in order, using $K_{\text{exp}} = 2$:

$$\|Z_A(X)\|_{L_r} \lesssim_d \frac{\|Z_{A^{[N]}}(Y)\|_{L_r}}{\ell(N)^d} \lesssim_d \frac{\max_{I \subseteq [d]} \Delta_I(A^{[N]}; r)}{\ell(N)^d} = \frac{\max_{I \subseteq [d]} N^{|I^c|/2} \Delta_I(A; r)}{\ell(N)^d}.$$

The first two constants depend only on d .

Step 2: real r and N . Set $\bar{r} = 2\lceil r/2 \rceil$ and $\bar{N} = \lceil N \rceil$. Since $\bar{r} \leq 2r$, $\bar{N} \leq 2N$, and $\kappa_{I,\mathcal{J}} \leq d$, rounding increases the numerator by at most $2^{3d/2}$ and cannot decrease $\ell(N)^d$. Monotonicity of L_r -norms therefore proves (25) for real r, N .

Step 3: convert moments into tails. For $A = 0$, the tail bound is immediate. Otherwise, since $Z_A(X)$ is centered, the standard moment-to-tail implication [22, Proposition 3.3] applied to (25) gives (26). $\square$

3.3. Choosing the replication level.

Proof of Theorem 1.1. In Proposition 3.1, group the terms by the number $k = |I^c|$ of free coordinates. The case $k = 0$ is the largest-entry term $\Delta_{[d]}(A; r) = r^d m_A$, and the maximum over the terms with $k \geq 1$ is $B_k(A; r)$. Thus

$$(29) \quad \|Z_A(X)\|_{L_r} \lesssim_d \frac{\max\{r^d m_A, \max_{1 \leq k \leq d} N^{k/2} B_k(A; r)\}}{\ell(N)^d}.$$

To keep the numerator equal to $r^d m_A$, it suffices that

$$N^{k/2} B_k(A; r) \leq r^d m_A \quad \text{for every } k, \quad \text{or equivalently} \quad N \leq \min_{1 \leq k \leq d} \left(\frac{r^d m_A}{B_k(A; r)} \right)^{2/k}.$$

If $\lambda_A(r) > 1$, its definition makes $N = \lambda_A(r)$ the largest admissible common replication level. Substitution into (29) proves (5) and hence also (4). If $\lambda_A(r) = 1$, use $N = 1$ instead. Since $\ell(1) = 1$,

$$\max\{r^d m_A, \max_k B_k(A; r)\} \leq r^d m_A + \max_k B_k(A; r),$$

which is precisely the right-hand side of (4) in this case. $\square$

Proof of Corollary 1.2. Put $q_0 = (t/m_A)^{1/d}$. If $\eta_A(t) = 1$, take $N = 1$ in (26). The entry term gives the rate q_0 ; every other term gives $(t/M_{I,\mathcal{J}}(A))^{1/\kappa_{I,\mathcal{J}}}$. Since $\ell(\eta_A(t)) = 1$, this is exactly the claimed minimum in (7).

Now suppose $\eta_A(t) > 1$, and take $N = \eta_A(t)$, $L = \ell(N)$. For $I \subsetneq [d]$, write $k = |I^c| \geq 1$ and $\kappa = \kappa_{I,\mathcal{J}}$. The definition of $\eta_A(t)$ implies

$$N \leq \left(\frac{m_A q_0^{d-\kappa}}{M_{I,\mathcal{J}}(A)} \right)^{2/k}, \quad \text{hence} \quad N^{k/2} M_{I,\mathcal{J}}(A) q_0^\kappa \leq t.$$

Since $L \geq 1$ and $\kappa \leq d$,

$$\frac{N^{k/2} M_{I,\mathcal{J}}(A) (q_0 L)^\kappa}{\ell(N)^d} \leq t L^{\kappa-d} \leq t.$$

The entry term satisfies equality: $m_A (q_0 L)^d / \ell(N)^d = t$. Consequently, every rate in (26) is at least $q_0 L$. This proves the stronger estimate $\mathbb{P}(|Z_A(X)| \geq t) \leq 2e^{-c_d q_0 L}$, which implies (7) because its displayed minimum is at most $q_0 L$. $\square$

For the quadratic corollary we also need forward decoupling for diagonal-free polynomials [18, Theorem 1]. If A is diagonal-free and $Q_A(X) = \sum_i a_i X_{i_1} \cdots X_{i_d}$, then

$$(30) \quad \begin{aligned} \mathbb{P}(|Q_A(X)| \geq t) &\leq C_d \mathbb{P}\left(\left|Z_A(X^{(1)}, \dots, X^{(d)})\right| \geq t/C_d\right), \\ \|Q_A(X)\|_{L_r} &\lesssim_d \left\|Z_A(X^{(1)}, \dots, X^{(d)})\right\|_{L_r}, \quad r \geq 2, \end{aligned}$$

where the input vectors on the right are independent copies. Only the forward inequality is used; it does not require symmetry of A . The moment comparison follows by integrating the cited tail inequality, with a constant depending only on d .

Proof of Corollaries 1.3 and 1.4. Step 1: reduce the quadratic form to the decoupled bound. The zero diagonal allows us to apply the moment comparison in (30) to $Q_A(X) = X^\top A X$. Combining it with (25) gives the same common- N moment estimate for the quadratic form, with a different universal constant. The same moment-to-tail implication and choices of N therefore give the same moment and tail conclusions for Q_A . In particular, decoupling is applied before the moment-to-tail conversion, so its constants do not change the arguments of λ_A or η_A . The cubic chaos is already decoupled and needs no such step.

Step 2: identify all non-entry norms. For $d = 2$, fixing one coordinate leaves a weighted row or column. Symmetry of A makes their maximal Euclidean norms equal to $\|AD_p\|_{2,\infty}$. With neither coordinate fixed, the one-block partition gives the Frobenius norm of $D_p A D_p$, and the two-block partition gives its operator norm. Since $\kappa = 2 - k + |\mathcal{J}|/2$, this yields

$$B_1(A; r) = r^{3/2} \|AD_p\|_{2,\infty}, \quad B_2(A; r) = \max\{\sqrt{r} \|D_p A D_p\|_F, r \|D_p A D_p\|_{\text{op}}\}.$$

For $d = 3$, a slice with k free coordinates has weight $p^{k/2}$, and $\kappa = 3 - k + |\mathcal{J}|/2$. With one free coordinate, the slices are vectors, so the only partition gives V_1 with exponent $\kappa = 5/2$. With two free coordinates, the slices are matrices: the one-block partition gives their weighted Frobenius norm V_2 , with $\kappa = 3/2$, while the two-block partition gives their weighted operator norm W_2 , with $\kappa = 2$.

With all three coordinates free, one block gives the weighted Frobenius norm V_3 , with $\kappa = 1/2$; two blocks give the weighted matrix-flattening norms, whose maximum is W_3 , with $\kappa = 1$; and three blocks give the weighted injective norm I_3 , with $\kappa = 3/2$. These scales already include the maxima over the applicable fixed coordinates and choices of modes. Therefore

$$\begin{aligned} B_1(A; r) &= r^{5/2} V_1, \\ B_2(A; r) &= \max\{r^{3/2} V_2, r^2 W_2\}, \\ B_3(A; r) &= \max\{\sqrt{r} V_3, r W_3, r^{3/2} I_3\}. \end{aligned}$$

Step 3: bound the replication parameter. Each non-entry term above has the form $r^\kappa M$, where $k \geq 1$ coordinates remain free. The corresponding term in $\gamma_d(A)$ is exactly $(M/m_A)^{1/(d-\kappa)}$. Thus $M \leq m_A \gamma_d(A)^{d-\kappa}$. For $u = r/\gamma_d(A) \geq 1$, the associated candidate in the minimum defining $\lambda_A(r)$ satisfies

$$\left(\frac{r^d m_A}{r^\kappa M} \right)^{2/k} \geq u^{2(d-\kappa)/k} \geq u.$$

The last inequality uses $d - \kappa = k - |\mathcal{J}|/2 \geq k/2$, because $|\mathcal{J}| \leq k$. Taking the minimum over all terms proves $\lambda_A(r) \geq r/\gamma_d(A)$. If $r \geq \max\{2, e\gamma_d(A)\}$, this lower bound is strictly greater than one, and

$$\ell(\lambda_A(r)) \geq 1 + \log(r/\gamma_d(A)).$$

Substitution into (5) proves both stated moment bounds.

Step 4: expand the tail rates. Put $q_0 = (t/m_A)^{1/d}$. For a non-entry scale M , its candidate in the definition of $\eta_A(t)$ is

$$\left(\frac{m_A}{M} \left(\frac{t}{m_A} \right)^{1-\kappa/d} \right)^{2/k} = \left(\frac{m_A q_0^{d-\kappa}}{M} \right)^{2/k}.$$

If $q_0/\gamma_d(A) \geq 1$, the calculation in Step 3, with r replaced by q_0 , bounds every candidate below by $q_0/\gamma_d(A)$. Otherwise the cutoff in $\eta_A(t)$ gives the bound one. Hence, for every $t > 0$,

$$\eta_A(t) \geq 1 \vee \frac{(t/m_A)^{1/d}}{\gamma_d(A)}.$$

This is an algebraic comparison and does not require $q_0 \geq 2$. Finally, the non-entry rates in (7) are $(t/M)^{1/\kappa}$. For $d = 2$, their exponents are $2/3, 2, 1$, respectively; for $d = 3$, the corresponding exponents are $2/5, 2/3, 1/2, 2, 1, 2/3$. Inserting these rates and the lower bound on $\eta_A(t)$ proves (10) and (13). $\square$

4. PROOFS FOR SPARSE KHATRI–RAO EMBEDDINGS

The proof has two parts. First, the replication estimate controls the moments of one sketch row through the geometry factor $G_x(\delta)$. Second, independence across rows turns this finite range of moment bounds into concentration of the average squared norm. A net argument then gives the subspace guarantee.

Proof of Proposition 2.1. Fix $x \neq 0$ and $0 < \delta < 1$. Write $R = \log(e^2/\delta) > 2$, so the parameter in the definition of $G_x(\delta)$ is $q = 2R$.

Step 1: identify the mean and the independent summands. For each row put

$$Z_m(x) = \sqrt{M}(\Phi x)_m = p^{-d/2} \sum_{\mathbf{i} \in [n]^d} x_{\mathbf{i}} \prod_{s=1}^d \xi_{m i_s}^{(s)}.$$

The variables $Z_1(x), \dots, Z_M(x)$ are independent and identically distributed because they use disjoint, identically distributed families of sketch entries. For fixed s, m , $\mathbb{E}[\xi_{m i_s}^{(s)} \xi_{m j_s}^{(s)}] = p \mathbf{1}_{\{i=j\}}$. Expanding the square and using independence across modes gives

$$\mathbb{E} Z_m(x)^2 = p^{-d} \sum_{\mathbf{i}, \mathbf{j}} x_{\mathbf{i}} x_{\mathbf{j}} \prod_{s=1}^d \mathbb{E}[\xi_{m i_s}^{(s)} \xi_{m j_s}^{(s)}] = \sum_{\mathbf{i}} x_{\mathbf{i}}^2 = \|x\|_{\mathbb{F}}^2.$$

Consequently,

$$\|\Phi x\|_2^2 = \frac{1}{M} \sum_{m=1}^M Z_m(x)^2, \quad \mathbb{E} \|\Phi x\|_2^2 = \|x\|_{\mathbb{F}}^2.$$

Step 2: control one row up to moment order q . Fix $2 \leq r \leq q$, and set $\rho = \rho_x(\delta)$. The sketch entries are centered, bounded by one, and have variance p . Apply Proposition 3.1 with coefficient tensor x and replication level $N = \rho$. Writing $U = I^c$, the normalization $p^{-d/2}$ in $Z_m(x)$ and the weight $p^{|U|/2}$ in $M_{I, \mathcal{J}}(x)$ give

$$\|Z_m(x)\|_{L_r} \leq \frac{C_d}{\ell(\rho)^d} \max_{\substack{U \subseteq [d] \\ \mathcal{J} \in \Pi(U)}} \left[\rho^{|U|/2} r^{d-|U|+|\mathcal{J}|/2} p^{-(d-|U|)/2} \max_{\mathbf{i}_{U^c}} \|x_{\mathbf{i}_U}\|_{\mathcal{J}} \right].$$

We square this estimate and divide by r to compare its terms with the scales $a_0, \dots, a_d$.

If $U = \emptyset$, the term before the common divisor $\ell(\rho)^{2d}$ is

$$r^{2d-1} p^{-d} \|x\|_{\infty}^2 \leq a_0.$$

If $k = |U| \geq 1$ and $\mathcal{J}$ has one block, its norm is the slice Frobenius norm. Its contribution is at most

$$\rho^k r^{2(d-k)} p^{-(d-k)} \max_{\mathbf{i}_{U^c}} \|x_{\mathbf{i}_U}\|_{\mathbb{F}}^2 \leq \rho^k \left(\frac{q^2}{p} \right)^{d-k} \max_{\mathbf{i}_{U^c}} \|x_{\mathbf{i}_U}\|_{\mathbb{F}}^2.$$

Finally, suppose $j = |\mathcal{J}| \geq 2$. Partition norms increase when blocks are merged [17, Section 2], so

$$\|x_{i_U}\|_{\mathcal{J}} \leq \max_{\emptyset \neq V \subsetneq U} \|x_{i_U}\|_{\{V, U \setminus V\}}.$$

Since $j \leq k$ and $2 \leq r \leq q$,

$$r^{2(d-k)+j-1} \leq q^{2(d-k)+k-1}.$$

The resulting contribution is therefore bounded by

$$\rho^k \left(\frac{q^2}{p}\right)^{d-k} q^{k-1} \max_{\substack{i_U^c \\ \emptyset \neq V \subsetneq U}} \|x_{i_U}\|_{\{V, U \setminus V\}}^2.$$

These last two bounds are precisely the two parts of $\rho^k a_k$. It follows that

$$\frac{\|Z_m(x)\|_{L_r}^2}{r} \leq C_d \frac{\max\{a_0, \rho a_1, \dots, \rho^d a_d\}}{\ell(\rho)^{2d}} \leq C_d G_x(\delta) \|x\|_F^2.$$

For the last inequality, if $\rho > 1$, its definition gives $\rho^k a_k \leq a_0$ for every $k \geq 1$. If $\rho = 1$, the quotient is simply $\max_{0 \leq k \leq d} a_k$, which appears in the definition of $G_x(\delta)$. We have proved

$$(31) \quad \|Z_m(x)\|_{L_r} \lesssim_d \sqrt{r G_x(\delta)} \|x\|_F, \quad 2 \leq r \leq q = 2R.$$

Step 3: center and average the squared rows. Put $Y_m = Z_m(x)^2 - \mathbb{E} Z_m(x)^2$. These variables are independent, identically distributed, and centered. Also $G_x(\delta) \geq 1$, because $a_d \geq \|x\|_F^2$. For $2 \leq s \leq R$, Step 2 applies at order $2s \leq q$, so

$$(32) \quad \begin{aligned} \|Y_m\|_{L_s} &\leq \|Z_m(x)\|_{L_{2s}}^2 + \mathbb{E} Z_m(x)^2 \\ &\leq C_d G_x(\delta) s \|x\|_F^2. \end{aligned}$$

By Latała's moment estimate and its centered extension by symmetrization [32, Corollary 2 and Remark 2], if $\|Y_1\|_{L_s} \leq Ks$ for $2 \leq s \leq R$, then

$$(33) \quad \begin{aligned} \left\| \frac{1}{M} \sum_{m=1}^M Y_m \right\|_{L_R} &\lesssim \frac{1}{M} \sup_{\max\{2, R/M\} \leq s \leq R} \frac{R}{s} (M/R)^{1/s} \|Y_1\|_{L_s} \\ &\lesssim K \left\{ \sqrt{\frac{R}{M}} + \frac{R}{M} \right\}. \end{aligned}$$

Here $(M/R)^{1/s} \leq \max\{1, \sqrt{M/R}\}$ for $s \geq 2$, so only moments through order R are needed. Applying (33) with $K = C_d G_x(\delta) \|x\|_F^2$ bounds the L_R -norm of $\|\Phi x\|_2^2 - \|x\|_F^2$.

Step 4: obtain the failure probability and row count. Markov's inequality at e times this L_R -bound gives

$$\mathbb{P} \left(\left| \|\Phi x\|_2^2 - \|x\|_F^2 \right| > C_d G_x(\delta) \|x\|_F^2 \left\{ \sqrt{\frac{R}{M}} + \frac{R}{M} \right\} \right) \leq e^{-R} = \delta/e^2 \leq \delta.$$

Since $R \asymp \log(2/\delta)$, division by $\|x\|_F^2$ proves (18). For the relative-error assertion, suppose $M \geq C_0 G_x(\delta)^2 \varepsilon^{-2} R$. Then

$$G_x(\delta) \sqrt{\frac{R}{M}} \leq \frac{\varepsilon}{\sqrt{C_0}}, \quad G_x(\delta) \frac{R}{M} \leq \frac{\varepsilon^2}{C_0 G_x(\delta)} \leq \frac{\varepsilon}{C_0},$$

where $G_x(\delta) \geq 1$ and $\varepsilon \leq 1$ were used. Choosing C_0 sufficiently large in terms of d proves (19).

Step 5: pass from one tensor to a subspace. The case $W = \{0\}$ is immediate. Otherwise, choose a $1/4$ -net $\mathcal{N}$ of its unit sphere with $|\mathcal{N}| \leq 9^k$ [43, Corollary 4.2.13]. Apply the fixed-vector result on $\mathcal{N}$ with error $\varepsilon/2$ and failure probability $\delta/9^k$. The assumed bound on the geometry factors and $\log(2 \cdot 9^k/\delta) = k \log 9 + \log(2/\delta)$ show that the stated row count suffices. A union bound therefore gives, with probability at least $1 - \delta$,

$$\sup_{v \in \mathcal{N}} |\|\Phi v\|_2^2 - 1| \leq \varepsilon/2.$$

For the self-adjoint operator $T = (\Phi|_W)^*(\Phi|_W) - I_W$, the net bound [43, Exercise 4.4.3(b)] yields

$$\|T\|_{\text{op}} \leq 2 \sup_{v \in \mathcal{N}} |\langle Tv, v \rangle| \leq \varepsilon.$$

This is the asserted norm preservation on W . $\square$

APPENDIX A. MOMENT BOUNDS BEYOND THE BOUNDED DECOUPLED SETTING

A.1. Bounded non-decoupled chaoses. For a diagonal-free polynomial, decoupling replaces the repeated input vector X by independent copies $X^{(1)}, \dots, X^{(d)}$, at a cost depending only on the degree. Every copy of X_i inherits the same variance proxy p_i . Throughout this subsection, evaluate $M_{I, \mathcal{J}}(A)$, $B_k(A; r)$, and $\lambda_A(r)$ with $p_i^{(s)} = p_i$.

Proposition A.1 (Bounded non-decoupled moment bound). *Let $A = (a_{i_1, \dots, i_d})$ be a nonzero real diagonal-free d -tensor: $a_{i_1, \dots, i_d} = 0$ whenever two indices coincide. Let $X_1, \dots, X_n$ be independent, with*

$$\mathbb{E}X_i = 0, \quad |X_i| \leq 1 \text{ a.s.}, \quad \mathbb{E}X_i^2 \leq p_i, \quad 0 < p_i \leq 1,$$

and define

$$Q_A(X) = \sum_{i_1, \dots, i_d=1}^n a_{i_1, \dots, i_d} X_{i_1} \cdots X_{i_d}.$$

Writing $m_A = \|A\|_\infty$, for every $r \geq 2$ we have

$$(34) \quad \|Q_A(X)\|_{L_r} \lesssim_d \max_{1 \leq k \leq d} B_k(A; r) + \frac{r^d m_A}{\ell(\lambda_A(r))^d}.$$

When $\lambda_A(r) > 1$, this improves to

$$(35) \quad \|Q_A(X)\|_{L_r} \lesssim_d \frac{r^d m_A}{\ell(\lambda_A(r))^d}.$$

Proof. Let $X^{(1)}, \dots, X^{(d)}$ be independent copies of X . By the forward decoupling inequality (30),

$$\|Q_A(X)\|_{L_r} \lesssim_d \left\| Z_A(X^{(1)}, \dots, X^{(d)}) \right\|_{L_r}.$$

The copied family satisfies (1) with $p_i^{(s)} = p_i$, so both conclusions follow directly from Theorem 1.1 and its effective form (5). In particular, the individual variance bounds, and hence all weighted norms and $\lambda_A(r)$, are unchanged. No symmetry of A is needed for forward decoupling. $\square$

For $d = 2$, this is the quadratic form $X^\top A X$ with zero diagonal, as used in Corollary 1.3.

A.2. Sparse ψ_α inputs. Let $1 < \alpha \leq \infty$, with $1/\infty = 0$, and put $\beta = 1 - 1/\alpha \in (0, 1]$. For finite α , use the standard Orlicz norm

$$\|\xi\|_{\psi_\alpha} = \inf \{K > 0 : \mathbb{E} \exp((|\xi|/K)^\alpha) \leq 2\}; \quad \|\xi\|_{\psi_\infty} = \|\xi\|_{L_\infty}.$$

The cases $\alpha = 2$ and $\alpha = \infty$ correspond to sub-Gaussian and bounded amplitudes. Relative to the exponential reference law, their smaller moments yield a logarithmic gain of power β per mode.

Proposition A.2 (Bennett-type bound for sparse sub-exponential chaoses). *Let $K > 0$ and $X_i^{(s)} = \delta_i^{(s)} \xi_i^{(s)}$, where all Bernoulli random variables and amplitudes are mutually independent and*

$$\delta_i^{(s)} \sim \text{Bernoulli}(p_i^{(s)}), \quad 0 < p_i^{(s)} \leq 1, \quad \mathbb{E} \xi_i^{(s)} = 0, \quad \|\xi_i^{(s)}\|_{\psi_\alpha} \leq K.$$

For a nonzero real rectangular tensor A , evaluate $M_{I,\mathcal{J}}(A)$, $B_k(A; r)$, and $\lambda_A(r)$ using the Bernoulli probabilities, and write $m_A = \|A\|_\infty$. Then, for every $r \geq 2$,

$$(36) \quad \|Z_A(X)\|_{L_r} \lesssim_d K^d \left[\max_{1 \leq k \leq d} B_k(A; r) + \frac{r^d m_A}{\ell(\lambda_A(r))^{d\beta}} \right].$$

When $\lambda_A(r) > 1$,

$$(37) \quad \|Z_A(X)\|_{L_r} \lesssim_d \frac{K^d r^d m_A}{\ell(\lambda_A(r))^{d\beta}}.$$

The constants are uniform over $1 < \alpha \leq \infty$. The same bounds hold for $Q_A(X)$ when A is diagonal-free, the Bernoulli random variables and centered amplitudes satisfy the same assumptions, and $p_i^{(s)} = p_i$.

These norms use Bernoulli probabilities. For unbounded amplitudes, $p_i^{(s)}$ need not itself bound $\mathbb{E}(X_i^{(s)})^2$; the amplitude scale is carried by K^d . The reference law remains exponential, so the coefficient scales and the choice $\lambda_A(r)$ stay the same. Only the power of the logarithm changes.

Proof. By homogeneity, first take $K = 1$. For finite $\alpha > 1$ and integers $m \geq 2$, the Orlicz assumption gives $\mathbb{E}|\xi_i^{(s)}|^{\alpha m} \leq 2m!$. Thus

$$(38) \quad \mathbb{E}|X_i^{(s)}|^m = p_i^{(s)} \mathbb{E}|\xi_i^{(s)}|^m \leq p_i^{(s)} (2m!)^{1/\alpha} \leq 2p_i^{(s)} (m!)^{1-\beta}.$$

For $\alpha = \infty$, boundedness gives $\mathbb{E}|X_i^{(s)}|^m \leq p_i^{(s)}$, so the same bound holds.

As in the proof of Lemma 3.2, put $V = (X - X')/2$, where X' is an independent copy of X . Symmetrization gives $\|Z_A(X)\|_{L_r} \leq 2^d \|Z_A(V)\|_{L_r}$, and convexity gives $\mathbb{E}|V_i^{(s)}|^m \leq \mathbb{E}|X_i^{(s)}|^m$. For integer $N \geq 1$, let $S_i^{(s)} = \sum_{u=1}^N Y_{i,u}^{(s)}$ be the sparse-exponential replica sums from Section 3.2.1, independent of V , with inherited probabilities $p_i^{(s)}$. The same scalar calculation gives $\mathbb{E}(S_i^{(s)})^m \geq N p_i^{(s)} m!$ for even $m \geq 2$. Consequently,

$$\frac{\mathbb{E}|V_i^{(s)}|^m}{\mathbb{E}(S_i^{(s)})^m} \leq \frac{2}{N(m!)^\beta} \leq \frac{2e^\beta}{\ell(N)^{m\beta}} \leq \left(\frac{\sqrt{2e}}{\ell(N)^\beta} \right)^m,$$

where $N(m!)^\beta \geq (Nm!)^\beta \geq e^{-\beta} \ell(N)^{m\beta}$. This replaces the scalar comparison in the proof of Lemma 3.2; its coordinate-replacement argument therefore yields, for even r ,

$$\|Z_A(X)\|_{L_r} \leq \frac{(2\sqrt{2e})^d}{\ell(N)^{d\beta}} \|Z_{A^{[N]}}(Y)\|_{L_r}.$$

Applying Lemmas 3.3 and 3.4 and restoring K gives

$$(39) \quad \|Z_A(X)\|_{L_r} \lesssim_d K^d \frac{\max_{I \subseteq [d]} N^{|I^c|/2} \Delta_I(A; r)}{\ell(N)^{d\beta}}.$$

The rounding argument in Section 3.2.3 extends this estimate to all real $r \geq 2$, $N \geq 1$. Now choose $N = \lambda_A(r)$ when $\lambda_A(r) > 1$, and $N = 1$ otherwise, exactly as in the proof of Theorem 1.1. The numerator is unchanged, so this proves both moment bounds. All constants are uniform in α : the scalar comparison uses the fixed constant $\sqrt{2e}$, and the remaining steps use only $0 < \beta \leq 1$.

For the non-decoupled assertion, apply the moment comparison in (30) and then the preceding bounds. Each independent copy retains the selector probability p_i and amplitude bound K . $\square$

product with a Bernoulli random variable is not essential. For example, independent The proof only used scalar moment bounds, so a representation as a centered inputs satisfying

$$|X_i^{(s)}| \leq K \quad \text{a.s.}, \quad \mathbb{E}(X_i^{(s)})^2 \leq K^2 p_i^{(s)}$$

obey $\mathbb{E}|X_i^{(s)}|^m \leq K^{m-2} \mathbb{E}(X_i^{(s)})^2 \leq K^m p_i^{(s)}$. The proof therefore applies with $\beta = 1$, recovering Theorem 1.1 when $K = 1$. At the other endpoint, $\beta = 0$ gives the sparse-exponential estimate with no logarithmic gain.

Tail bound in the logarithmic regime. Let η_A be defined by (6), using the Bernoulli probabilities. Under Proposition A.2, whenever $\eta_A(t/K^d) \geq e$,

$$(40) \quad \mathbb{P}(|Z_A(X)| \geq t) \leq 2 \exp \left[-c_d \left(\frac{t}{K^d m_A} \right)^{1/d} \ell(\eta_A(t/K^d))^\beta \right].$$

The same bound holds for $Q_A(X)$ under the non-decoupled assumptions. Thus the logarithm enters the tail rate with power $1/2$ for sub-Gaussian amplitudes and power one for bounded amplitudes.

Proof. We adapt the proof of Corollary 1.2, putting

$$w = \left(\frac{t}{K^d m_A} \right)^{1/d}, \quad N = \eta_A(t/K^d), \quad L = \ell(N).$$

Since $N > 1$, its definition gives, for every $I, \mathcal{J}$,

$$K^d N^{|I^c|/2} M_{I, \mathcal{J}}(A) w^{\kappa_{I, \mathcal{J}}} \leq t,$$

with equality for the entry term. Choose $r = cwL^\beta$, where $0 < c \leq 1$ depends only on d . If $r \geq 2$, (39) implies

$$\|Z_A(X)\|_{L_r} \leq C_d t \max_{I, \mathcal{J}} c^{\kappa_{I, \mathcal{J}}} L^{\beta(\kappa_{I, \mathcal{J}} - d)} \leq C_d \sqrt{c} t,$$

because $1/2 \leq \kappa_{I, \mathcal{J}} \leq d$. Taking c sufficiently small and applying Markov's inequality gives the claimed rate wL^β . If $r < 2$, the bound is trivial after decreasing c_d , as in Section 3.2.3. For $Q_A(X)$, apply the moment comparison (30) before Markov's inequality and absorb its constant into c ; the argument t/K^d of η_A remains unchanged. $\square$

Acknowledgements. GPT-5.6 assisted in developing the proof. The authors independently verified the argument and take full responsibility for the final proof. K.W. is supported by Hong Kong RGC grant GRF 16305526. Y.Z. was partially supported by the Simons Grant MPS-TSM-00013944 and NSF DMS-2606337. This work was carried out while Y.Z. was visiting the Simons Institute for the Theory of Computing during the Spectral Theory Beyond Graphs program in Fall 2026.

REFERENCES

- [1] R. Adamczak. A note on the Hanson-Wright inequality for random vectors with dependencies. *Electronic Communications in Probability*, 72:1–13, 2015.
- [2] R. Adamczak and R. Łatała. Tail and moment estimates for chaoses generated by symmetric random variables with logarithmically concave tails. *Annales de l’IHP Probabilités et statistiques*, 48(4):1103–1136, 2012.
- [3] R. Adamczak, R. Łatała, and R. Meller. Moments of Gaussian chaoses in Banach spaces. *Electronic Journal of Probability*, 26(11):1–36, 2021.
- [4] R. Adamczak and P. Wolff. Concentration inequalities for non-Lipschitz functions with bounded derivatives of higher order. *Probability Theory and Related Fields*, 162:531–586, 2015.
- [5] T. D. Ahle and J. B. T. Knudsen. Almost optimal tensor sketch. *arXiv preprint arXiv:1909.01821*, 2019.
- [6] S. Bamberger, F. Krahmer, and R. Ward. The Hanson–Wright inequality for random tensors. *Sampling Theory, Signal Processing, and Data Analysis*, 20(2):14, 2022.
- [7] G. Bennett. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57(297):33–45, 1962.
- [8] L. Beretta and C. Musco. A tight analysis of Khatri–Rao oblivious subspace embeddings. *arXiv preprint arXiv:2608.28094v2*, 2026.
- [9] S. Boucheron, O. Bousquet, G. Lugosi, and P. Massart. Moment inequalities for functions of independent random variables. *The Annals of Probability*, 33(2):514–560, 2005.
- [10] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, Oxford, 2013.
- [11] Z. Bujanović, L. Grubišić, D. Kressner, and H. Y. Lam. Subspace embedding with random Khatri–Rao products and its application to eigensolvers. *IMA Journal of Numerical Analysis*, 46(3):1328–1355, 2026.
- [12] S. Chatterjee. The missing log in large deviations for triangle counts. *Random Structures & Algorithms*, 40(4):437–451, 2012.
- [13] S. Chatterjee. An introduction to large deviations for random graphs. *Bulletin of the American Mathematical Society*, 53(4):617–642, 2016.
- [14] K. Chen and R. Jin. Tensor-structured sketching for constrained least squares. *SIAM J. Matrix Anal. Appl.*, 42(4):1703–1731, 2021.
- [15] X. Chen and Y. Yang. Hanson–Wright inequality in Hilbert spaces with application to K-means clustering for non-Euclidean data. *Bernoulli*, 27(1):586–614, 2021.
- [16] G. Dai, Y. He, K. Wang, and Y. Zhu. A note on the improved sparse Hanson–Wright inequalities. *arXiv preprint arXiv:2505.20799*, 2025.
- [17] G. Dai and K. Wang. Sparsity-adaptive concentration inequalities for random polynomials. *arXiv preprint arXiv:2606.24090v1*, 2026.
- [18] V. H. de la Peña and S. J. Montgomery-Smith. Decoupling inequalities for the tail probabilities of multivariate U-statistics. *The Annals of Probability*, 23(2):806–816, 1995.
- [19] R. DeMarco and J. Kahn. Tight upper tail bounds for cliques. *Random Structures & Algorithms*, 41(4):469–487, 2012.
- [20] E. Giné, R. Łatała, and J. Zinn. Exponential and moment inequalities for U-statistics. In *High Dimensional Probability II*, pages 13–38. Springer, 2000.
- [21] E. D. Gluskin and S. Kwapień. Tail and moment estimates for sums of independent random variables with logarithmically concave tails. *Studia Mathematica*, 114(3):303–309, 1995.
- [22] F. Götze, H. Sambale, and A. Sinulis. Concentration inequalities for polynomials in α -sub-exponential random variables. *Electronic Journal of Probability*, 26:1–22, 2021.
- [23] D. L. Hanson and F. T. Wright. A bound on tail probabilities for quadratic forms in independent random variables. *The Annals of Mathematical Statistics*, 42(3):1079–1083, 1971.
- [24] M. Harel, F. Mousset, and W. Samotij. Upper tails via high moments and entropic stability. *Duke Mathematical Journal*, 171(10):2089–2192, 2022.
- [25] Y. He, K. Wang, and Y. Zhu. Sparse Hanson–Wright inequalities with applications. *Electron. J. Probab.*, 31(39):1–49, 2026.
- [26] P. Hitczenko, S. J. Montgomery-Smith, and K. Oleszkiewicz. Moment inequalities for sums of certain independent symmetric random variables. *Studia Math*, 123(1):15–42, 1997.
- [27] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58(301):13–30, 1963.
- [28] S. Janson and A. Ruciński. The infamous upper tail. *Random Structures & Algorithms*, 20(3):317–342, 2002.
- [29] J. H. Kim and V. H. Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000.

- [30] Y. Klochkov and N. Zhivotovskiy. Uniform Hanson-Wright type concentration inequalities for unbounded entries via the entropy method. *Electronic Journal of Probability*, 25(22):1–30, 2020.
- [31] K. Kolesko and R. Latała. Moment estimates for chaoses generated by symmetric random variables with logarithmically convex tails. *Statistics & Probability Letters*, 107:210–214, 2015.
- [32] R. Latała. Estimation of moments of sums of independent real random variables. *The Annals of Probability*, 25(3):1502–1513, 1997.
- [33] R. Latała. Tail and moment estimates for some types of chaos. *Studia mathematica*, 135:39–53, 1999.
- [34] R. Latała. Estimates of moments and tails of Gaussian chaoses. *Ann. Probab.*, 34(6):2315–2331, 2006.
- [35] M. Ledoux. *The concentration of measure phenomenon*, volume 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2001.
- [36] E. Lubetzky and Y. Zhao. On the variational problem for upper tails in sparse random graphs. *Random Structures & Algorithms*, 50(3):420–436, 2017.
- [37] S. J. Montgomery-Smith. The distribution of Rademacher sums. *Proceedings of the American Mathematical Society*, 109(2):517–522, 1990.
- [38] S. Park, X. Wang, and J. Lim. Sparse Hanson–Wright inequality for a bilinear form of sub-Gaussian variables. *Stat*, 12(1):e539, 2023.
- [39] B. Rakhshan and G. Rabusseau. Tensorized random projections. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3306–3316. PMLR, 2020.
- [40] M. Rudelson and R. Vershynin. Hanson–Wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18(82):1–9, 2013.
- [41] W. Schudy and M. Sviridenko. Concentration and moment inequalities for polynomials of independent random variables. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 437–446. SIAM, 2012.
- [42] Y. Sun, Y. Guo, J. A. Tropp, and M. Udell. Tensor random projection for low memory dimension reduction. *arXiv preprint arXiv:2105.00105*, 2021.
- [43] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [44] R. Vershynin. Concentration inequalities for random tensors. *Bernoulli*, 26(4):3139–3162, 2020.
- [45] V. Vu and K. Wang. Random weighted projections, random quadratic forms and random eigenvectors. *Random Structures & Algorithms*, 47(4):792–821, 2015.
- [46] S. Zhou. Sparse Hanson–Wright inequalities for subgaussian quadratic forms. *Bernoulli*, 25(3):1603–1639, 2019.
- [47] S. Zhou. Concentration of measure bounds for matrix-variate data with missing values. *Bernoulli*, 30(1):198–226, 2024.

DEPARTMENT OF MATHEMATICS, HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY, CLEAR WATER BAY, KOWLOON, HONG KONG

Email address: guozhengdai@ust.hk

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, SAN DIEGO, LA JOLLA, CA 92093 USA

Email address: yih130@ucsd.edu

DEPARTMENT OF MATHEMATICS, HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY, CLEAR WATER BAY, KOWLOON, HONG KONG

Email address: kewang@ust.hk

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF SOUTHERN CALIFORNIA, 3620 VERMONT AVENUE, LOS ANGELES, CA 90089, USA

Email address: yizhezhu@usc.edu