

Smoothed Analysis of the Komlós Conjecture

Nikhil Bansal ✉

University of Michigan, Ann Arbor, MI, USA

Haotian Jiang ✉

University of Washington, Seattle, WA, USA

Raghu Meka ✉

University of California, Los Angeles, CA, USA

Sahil Singla ✉

Georgia Institute of Technology, Atlanta, GA, USA

Makrand Sinha ✉

Simons Institute and University of California, Berkeley, CA, USA

Abstract

The well-known Komlós conjecture states that given n vectors in \mathbb{R}^d with Euclidean norm at most one, there always exists a ± 1 coloring such that the ℓ_∞ norm of the signed-sum vector is a constant independent of n and d . We prove this conjecture in a smoothed analysis setting where the vectors are perturbed by adding a small Gaussian noise and when the number of vectors $n = \omega(d \log d)$. The dependence of n on d is the best possible even in a completely random setting.

Our proof relies on a weighted second moment method, where instead of considering uniformly random colorings we apply the second moment method on an implicit distribution on colorings obtained by applying the Gram-Schmidt walk algorithm to a suitable set of vectors. The main technical idea is to use various properties of these colorings, including subgaussianity, to control the second moment.

2012 ACM Subject Classification Theory of computation \rightarrow Randomness, geometry and discrete structures

Keywords and phrases Komlós conjecture, smoothed analysis, weighted second moment method, subgaussian coloring

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.14

Category Track A: Algorithms, Complexity and Games

Funding *Nikhil Bansal*: Supported in part by the NWO VICI grant 639.023.812.

Haotian Jiang: Supported by NSF awards CCF-1749609, DMS-1839116, DMS-2023166, CCF-2105772.

Makrand Sinha: Supported by a Simons-Berkeley postdoctoral fellowship.

1 Introduction

A central question in discrepancy theory is the following Komlós problem: given vectors $v_1, \dots, v_n \in \mathbb{R}^d$ with Euclidean length at most 1, i.e., $\|v_i\|_2 \leq 1$ for all $i \in [n]$, find signs $x_i \in \{-1, 1\}$ for $i \in [n]$ to minimize the discrepancy $\|\sum_{i=1}^n x_i v_i\|_\infty$. The long-standing Komlós conjecture says that the discrepancy of any collection of such vectors is $O(1)$, independent of n and d . An important special case (up to scaling by $t^{1/2}$) is the Beck-Fiala problem, where the vectors $v_1, \dots, v_n \in \{0, 1\}^d$ and each v_i has at most t ones, so $\|v_i\|_2 \leq t^{1/2}$. Here, the Komlós conjecture reduces to the Beck-Fiala conjecture [8], which says that the discrepancy is $O(t^{1/2})$. The question of either proving or disproving these conjectures has received a lot of attention, and after a long line of work, the current best bounds for the Komlós and the Beck-Fiala problem are $O((\log n)^{1/2})$ and $O((t \log n)^{1/2})$ respectively, due to Banaszczyk [4].



© Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha; licensed under Creative Commons License CC-BY 4.0

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).

Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;

Article No. 14; pp. 14:1–14:12



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Motivated by the lack of progress for general worst-case instances, there has been a lot of recent work on these problems for random instances, with several interesting results and techniques, see e.g., [9, 7, 13, 10, 16, 18, 3, 14]. In this work, we consider the Komlós problem in the more general setting of smoothed analysis, where the input is generated by taking an arbitrary worst case Komlós instance and perturbing it randomly. The smoothed analysis model was first introduced by Spielman and Teng [17], and it interpolates nicely between worst case and average case analysis, and has been used extensively since then to study various problems. Recently smoothed analysis models have also been considered in discrepancy theory in a few other works [6, 11] – however the setting and focus of these results is quite different, and in particular they are not directly related to the Komlós or Beck-Fiala conjectures.

Random instances. To put our results in the proper context, we first describe the results on random instances. In general, these results depend on the different regimes of the parameters d, n and t , and we focus here on the more interesting case of $n \gg d$.

A natural model for random Beck-Fiala instances is where each entry is 1 with probability $p = t/d$, so that each column has t ones in expectation. In a surprising result, Hoberg and Rothvoss [13] showed that $\text{disc}(A) \leq 1$ w.h.p.¹ if $n = \Omega(d^2 \log d)$. Independently, Franks and Saks [10] showed that $\text{disc}(A) \leq 2$ w.h.p. if $n = \Omega(d^3 \log^2 d)$, for a more general class of instances. Both these results use interesting Fourier analysis based techniques.

It is not hard to see² that $n = \Omega(d \log d)$ is necessary for $O(1)$ discrepancy (provided p is not too small). An important step towards obtaining this optimal dependence was made by Potukuchi [15], who showed that $\text{disc}(A) \leq 1$ if $n = \Omega(d \log d)$ for the dense case of $p = 1/2$, using the second moment method. However, the sparse setting with $p \ll 1$ turns out to be more subtle, and was only recently resolved by Altschuler and Weed [3] using a more sophisticated conditional second moment method together with Stein’s method of exchangeable pairs. They show that $\text{disc}(A) \leq 1$ w.h.p. for $n \geq \Omega(d \log d)$, for every p .

The case of Gaussian matrices with i.i.d. $\mathcal{N}(0, 1)$ entries has also been considered, where Turner, Meka and Rigollet [18] give almost tight bounds for the entire regime, and in particular show that for $n = \Omega(d \log d)$ a discrepancy bound of $1/\text{poly}(d)$ holds.

The smoothed Komlós model. We now define our model formally. The input matrix is of the form $A = M + R$, where $M \in \mathbb{R}^{d \times n}$ is some worst-case matrix with columns of ℓ_2 -norm at most 1 and $R \in \mathbb{R}^{d \times n}$ is a random matrix with i.i.d. Gaussian entries distributed as $\mathcal{N}(0, \sigma^2/d)$, where $\sigma \leq 1$. The σ^2/d variance ensures that each column of R has ℓ_2 -norm roughly σ (and hence much less than that of M). Our goal is to understand the discrepancy of A . We will only be interested in showing the existence of a low discrepancy coloring for A , and not in algorithmically finding it (this seems far beyond the current techniques).

1.1 Results and Techniques

Our main result is the following.

► **Theorem 1** (Smoothed Komlós). *Let $\sigma > 0$ and $n = \frac{\omega(d \log d)}{\sigma^{4/3}}$. Then with probability $1 - o_d(1)$, the discrepancy of $M + R$ is at most $1/\text{poly}(d)$, where $M \in \mathbb{R}^{d \times n}$ is an arbitrary Komlós instance and $R \in \mathbb{R}^{d \times n}$ has i.i.d. $\mathcal{N}(0, \sigma^2/d)$ entries.*

¹ This is much better than the $O(t^{1/2})$ bound in the Beck-Fiala conjecture.

² If we fix any coloring x and consider a random instance, a fixed row has discrepancy $O(1)$ with probability $\approx (pn)^{-1/2}$, so the probability that each row has discrepancy $O(1)$ is $(pn)^{-\Omega(d)}$. As there are (only) 2^n possible colorings, a first moment argument already requires that $2^n (pn)^{-d} = \Omega(1)$.

An interpretation of Theorem 1 is that any counter-example to the Komlós conjecture (if it exists) will be rigid, or have $n \approx d$. Also, notice that dependence of n on d in Theorem 1 is essentially the best possible, as already evident in the very special case of $M = \mathbf{0}$ and $\sigma = 1$, i.e., a random matrix with i.i.d. $\mathcal{N}(0, 1)$ entries, where $n = \Omega(d \log d)$ is necessary to achieve $1/\text{poly}(d)$ discrepancy, as discussed earlier.

► **Remark 2.** Our proof techniques also give a high probability bound when $n = \Omega_\sigma(d^{1+\epsilon})$ for any constant $\epsilon > 0$. However, we do not explore this direction here. It would be interesting to know if the result also holds with high probability when $n = \omega(d \log d)$, as in the (fully) random setting.

► **Remark 3.** The dependence on the noise parameter σ in $n = \Omega_d(1/\sigma)$ in Theorem 1 is necessary, otherwise this would imply an $O(1)$ bound for the worst case Komlós problem. In particular, each row of the random part R must at least have enough ℓ_1 norm to offset the discrepancy from the worst case part M (which can be $O((\log n)^{1/2})$ given the currently known results). As each entry of R has magnitude about $\sigma d^{-1/2}$, we thus require $n = \Omega_d(1/\sigma)$ for each row of $M + R$ to have discrepancy $O(1)$.

The proof of Theorem 1 is based on the classical second moment method, however, it requires several additional ideas beyond those used for random instances, to handle the effect of the worst case part and its interplay with the random part. We describe these briefly next, and discuss them in more detail in Section 1.2.

- **Weighted second moment method.** Instead of applying the second moment method to the uniform distribution on the 2^n colorings, we consider a distribution on low-discrepancy colorings for M . This is necessary as for a random coloring $x \in \{-1, 1\}^n$, a typical entry of Mx will scale as $\sqrt{n/d}$, which is very unlikely to be *cancelled* by the discrepancy of the random part Rx , which typically scales as $\sigma\sqrt{n/d}$ (note that we want to show the existence of some x such that $(Rx)_i \approx -(Mx)_i$ for each coordinate $i \in [d]$).
- **Subgaussianity of colorings.** To ensure that $\|Mx\|_\infty$ is typically small, we consider the (implicit) distribution on colorings produced by the Gram-Schmidt (GS) algorithm [5] applied to M , which ensures that Mx is a 1-subgaussian vector [12] (details in Section 1.2). However, apriori the GS algorithm only guarantees that Mx is subgaussian, and says nothing about the distribution on the colorings x . For instance, it could be that any two colorings in the support have the first $9n/10$ coordinates identical, and thus look very non-random. This makes the second moment bounds much worse and harder to control. To handle this, we use a simple but useful trick to ensure that the distribution on the colorings x produced by the GS algorithm is also $O(1)$ -subgaussian. Roughly, this allows us to pretend that colorings x in the GS distribution behave randomly.
- **Exploiting subgaussianity to get cancellations across rows.** Most importantly, due to the worst case part M , doing a row by row analysis as is typically done in second moment computations for random instances, only works when $n = \Omega(d^2/\sigma^2)$ (details in Section 1.2). Roughly, the problem is that considering each coordinate of Mx separately completely ignores the global properties across the different coordinates that subgaussianity of Mx implies.

To get the optimal dependence of n on d , a key conceptual idea is to analyze *all* the rows together and use the subgaussianity of Mx and x carefully to get various cancellations across the different rows in the second moment computation. Exploiting subgaussianity also leads to various technical difficulties, as subgaussian vectors can differ from fully random Gaussian vectors in various non-trivial ways.

Notation. Throughout this paper, \log denotes the natural logarithm. We use the asymptotic notation $\omega(\cdot)$ or $o(\cdot)$ where the growth is always with respect to d – sometimes to emphasize this dependence we will also write $\omega_d(\cdot)$ or $o_d(\cdot)$. We write $\mathbb{E}_{x \sim \mathcal{G}}[f(x)]$ to denote the expectation of a function f where x is sampled from the distribution \mathcal{G} and we abbreviate this to $\mathbb{E}[f(x)]$ when the distribution is clear from the context. For reals $a, b \in \mathbb{R}$, the notation $[a \pm b]$ is used as a shorthand to denote the interval $[a - b, a + b]$. For a set $S \in \mathbb{R}^d$, we write $\delta S = \{\delta x \mid x \in S\}$ to denote the δ scaling of S .

1.2 Overview and Preliminaries

We now give a more detailed overview of the proof and the ideas. We also briefly describe the second moment method and some concepts we need such as subgaussianity and properties of the Gram-Schmidt algorithm.

Second moment method. The second moment method (e.g. [2]) is based on the following Paley-Zygmund inequality. For any non-negative random variable Z , we have that

$$\mathbb{P}[Z > 0] \geq \frac{(\mathbb{E}[Z])^2}{\mathbb{E}[Z^2]}.$$

So, if $\mathbb{E}[Z^2] = (1 + o(1))(\mathbb{E}[Z])^2$, then this implies that $\mathbb{P}[Z > 0] \geq 1 - o(1)$.

For constraint satisfaction problems, a standard way to use this to show that most random instances are feasible is by defining $S = S(R)$ as the number of solutions to an instance R , and showing that

$$\mathbb{E}_R[S^2] = (1 + o(1))(\mathbb{E}_R[S])^2, \tag{1}$$

which gives that $\mathbb{P}_R[S(R) > 0] = \mathbb{P}_R[S(R) \geq 1] \geq 1 - o(1)$.

Let us consider (1) more closely and define $S(R, x) = 1$ if x is a valid solution for instance R , and 0 otherwise. Then $S(R) = \sum_x S(R, x)$ and (1) can be written as

$$\mathbb{E}_R [\mathbb{P}_{x, y \sim \mathcal{U}} [S(R, x) = 1, S(R, y) = 1]] = (1 + o(1)) (\mathbb{E}_R [\mathbb{P}_{x \sim \mathcal{U}} [S(R, x) = 1]])^2, \tag{2}$$

where \mathcal{U} is the uniform distribution over all inputs.

Second moment method for smoothed Komlós. Let Δ denote the desired discrepancy bound. In our setting, denote by $S(R, x) = 1$ that $x \in \{\pm 1\}^n$ is a feasible coloring for the smoothed Komlós instance $M + R$, that is, if $\|(M + R)x\|_\infty \leq \Delta$. Roughly, this condition means that $Rx = -Mx$ and hence the discrepancy of the random part R cancels that of the worst case part M .

However, if x is chosen uniformly from $\{\pm 1\}^n$, it is not hard to see that this cannot work. The entries $(Mx)_i$ will be distributed roughly as $\mathcal{N}(0, m_i^2)$ where $m_i = (\sum_j M_{ij}^2)^{1/2}$ is the ℓ_2 -norm of row i of M , and in general will be much larger (around $1/\sigma \gg 1$ times) than the entries $(Rx)_i$.

Weighted second moment. To allow a reasonable probability of Rx cancelling Mx , a natural idea is to consider a distribution that is mostly supported over colorings x with low discrepancy on M . So, we will show (2) where x, y are sampled from some another suitable distribution \mathcal{G} instead of the uniform distribution \mathcal{U} . Similar ideas have also been used in other contexts such as [1]. Notice that this does not affect Rx , as for any fixed x , the contribution of the random part $(Rx)_i$ is still distributed as $N(0, n\sigma^2/d)$ (over the randomness of R).

A natural candidate is the distribution on colorings produced by the Gram-Schmidt (GS) walk algorithm [5]. In particular, we use the following result.

► **Theorem 4** ([12]). *Given vectors $v_1, \dots, v_n \in \mathbb{R}^m$ with $\|v_j\|_2 \leq 1$, the Gram-Schmidt walk algorithm outputs a random coloring $x \in \{-1, 1\}^n$ such that $\sum_{j=1}^n x_j v_j$ is 1-subgaussian.*

Recall that a random vector $Y \in \mathbb{R}^m$ is α -subgaussian if for all test vectors $\theta \in \mathbb{R}^m$,

$$\mathbb{E}[\exp(\langle \theta, Y \rangle)] \leq \exp\left(\frac{\alpha^2 \|\theta\|_2^2}{2}\right).$$

Roughly, this means that Y looks like a Gaussian with variance at most α^2 in every direction.

Let \mathcal{G} denote the (implicit) distribution over the colorings output by the GS walk algorithm. For a coloring x , let us denote $P_x := \mathbb{P}_R[S(R, x) = 1]$ and for two colorings x and y , let

$$P_{x,y} := \mathbb{P}_R[S(R, x) = 1, S(R, y) = 1].$$

Then changing the order of expectation in (2) and substituting, our goal is to show that

$$\mathbb{E}_{x,y \sim \mathcal{G}}[P_{xy}] = (1 + o(1)) \cdot \mathbb{E}_{x,y \sim \mathcal{G}}[P_x P_y]. \tag{3}$$

However, the set of low discrepancy colorings for M and the distribution \mathcal{G} can be quite complicated and hard to work with. Later, we will ensure that \mathcal{G} is also $O(1)$ -subgaussian, which will suffice for our purposes. Let us first consider (3) more closely.

The key computation. As R has i.i.d. Gaussian entries, the quantities P_x and P_{xy} can be written in a very clean way. In particular, as $(Rx)_i \sim \mathcal{N}(0, \sigma^2 n/d)$ for any coloring x , and the $(Rx)_i$ are independent for $i \in [d]$, we can write

$$P_x = \mathbb{P}_R \left[\bigcap_{i=1}^d ((Rx)_i \in (Mx)_i \pm \Delta) \right] = \prod_{i=1}^d \mathbb{P}_R [g_i \in (Mx)_i \pm \Delta],$$

where $g_i \sim \mathcal{N}(0, \sigma^2 n/d)$ and g_i 's are independent.

Similarly, for any fixed colorings x and y , writing $g_i = (Rx)_i$ and $g'_i = (Ry)_i$ we have

$$P_{xy} = \prod_{i=1}^d \mathbb{P}_R [g_i \in (Mx)_i \pm \Delta, g'_i \in (My)_i \pm \Delta],$$

where g_i and g'_i are correlated with $\mathbb{E}[g_i g'_i] = \langle x, y \rangle \cdot \sigma^2/d$.

A standard computation of 2-dimensional gaussian probabilities over rectangles (and ignoring some less crucial terms for the discussion here) gives

$$\frac{\mathbb{P}[g_i \in (Mx)_i \pm \Delta, g'_i \in (My)_i \pm \Delta]}{\mathbb{P}[g_i \in (Mx)_i \pm \Delta] \cdot \mathbb{P}[g'_i \in (My)_i \pm \Delta]} \approx \exp\left(\frac{d \langle x, y \rangle (Mx)_i (My)_i}{\sigma^2 n^2}\right). \tag{4}$$

So to prove (3), we could try to show that for each $i \in [d]$,

$$\mathbb{E}_{x,y \sim \mathcal{G}} \left[\frac{d \langle x, y \rangle (Mx)_i (My)_i}{\sigma^2 n^2} \right] = o\left(\frac{1}{d}\right). \tag{5}$$

Indeed, as $|\langle x, y \rangle| \leq n$ and $(Mx)_i, (My)_i$ are typically $O(1)$ (as Mx and My are subgaussian), setting $n = \omega(d^2/\sigma^2)$ would suffice to complete the second moment proof. However, this does not give us the optimal $d \log d$ dependence.

Next, we sketch the two ideas to obtain the optimal dependence.

14:6 Smoothed Analysis of the Komlós Conjecture

Subgaussianity of the distribution \mathcal{G} . If x and y were random colorings, we would typically expect that $|\langle x, y \rangle| \approx \sqrt{n}$ instead of n above. To achieve this, we apply the GS walk algorithm to the $(d+n) \times n$ matrix with M in top d rows and I_n in the bottom n rows. (Note that each column still has $O(1)$ length.) This ensures that the resulting distribution \mathcal{G} on the colorings x is $O(1)$ -subgaussian, while ensuring that Mx is also $O(1)$ -subgaussian.

Handling the rows together. Next, to exploit the subgaussianity of Mx and My , we look at all the rows together in (5) and consider

$$\sum_i \mathbb{E}_{x,y \sim \mathcal{G}} \left[\frac{d \langle x, y \rangle (Mx)_i (My)_i}{\sigma^2 n^2} \right] = \mathbb{E}_{x,y \sim \mathcal{G}} \left[\frac{d \langle x, y \rangle \langle Mx, My \rangle}{\sigma^2 n^2} \right]. \quad (6)$$

By the subgaussianity of the colorings x, y and discrepancy vectors Mx, My , we expect that $\mathbb{E}_{x,y \sim \mathcal{G}} |\langle x, y \rangle| \approx \sqrt{n}$ and $\mathbb{E}_{x,y} |\langle Mx, My \rangle| \approx \sqrt{d}$. Roughly speaking, this implies that the right side of (6) is typically $d^{3/2}/(\sigma^2 n^{3/2})$, and hence $n \gg d/\sigma^{4/3}$ suffices.

The formal argument needs some more care as $\langle x, y \rangle$ and $\langle Mx, My \rangle$ are correlated, and as we need to bound the exponential moment of $d \langle x, y \rangle \langle Mx, My \rangle / (\sigma^2 n^2)$ in (4), instead of the expectation, which gives the additional (necessary) logarithmic factor of $\log d$.

2 Proof of the Smoothed Komlós Conjecture

We use a weighted version of the second moment method as mentioned in the proof overview. Let \mathcal{G} be a distribution over coloring that will be specified later. We define the following random variable S which depends only on the randomness of R ,

$$S = S(R) := \mathbb{E}_{x \sim \mathcal{G}} [\mathbf{1}\{\|(M+R)x\|_\infty \leq \Delta\}],$$

for some parameter $\Delta = 1/\text{poly}(d)$ to be chosen later. The purpose of this variable is that the event $\{S > 0\}$ implies there exists a coloring $x \in \text{supp}(\mathcal{G})$ with discrepancy at most Δ . Our goal is to show that $\mathbb{P}(S > 0) = 1 - o(1)$. As explained in the proof overview, this would follow from the Paley-Zygmund inequality if we can establish that the first moment $\mathbb{E}_R[S]$ is always positive, and the second moment satisfies $\mathbb{E}_R[S^2] = (1 + o(1)) \cdot (\mathbb{E}_R[S])^2$. We next compute the moments.

First moment computation. We can compute

$$\mathbb{E}_R[S] = \mathbb{E}_{x \sim \mathcal{G}} \mathbb{E}_R[\mathbf{1}\{\|(M+R)x\|_\infty \leq \Delta\}] > 0,$$

where the strict inequality follows because fixing any outcome $x \sim \mathcal{G}$, the event $\{\|(M+R)x\|_\infty \leq \Delta\}$ happens with positive probability (recall that R is a Gaussian random matrix with each entry $\mathcal{N}(0, \sigma^2/d)$).

Second moment computation. For any $i \in [d]$, denote by m_i and r_i the i^{th} row of the matrices M and R respectively. The second moment is given by

$$\begin{aligned} \mathbb{E}_R[S^2] &= \mathbb{E}_R[\mathbb{E}_x[\mathbf{1}\{\|(M+R)x\|_\infty \leq \Delta\}] \cdot \mathbb{E}_y[\mathbf{1}\{\|(M+R)y\|_\infty \leq \Delta\}]] \\ &= \mathbb{E}_R \mathbb{E}_{x,y}[\mathbf{1}\{\|(M+R)x\|_\infty \leq \Delta, \|(M+R)y\|_\infty \leq \Delta\}] \\ &= \mathbb{E}_{x,y}[\mathbb{P}_R(\|(M+R)x\|_\infty \leq \Delta, \|(M+R)y\|_\infty \leq \Delta)] = \mathbb{E}_{x,y}[P_{x,y}], \end{aligned}$$

where we define

$$P_{x,y} := \mathbb{P}_R(\|(M+R)x\|_\infty \leq \Delta, \|(M+R)y\|_\infty \leq \Delta).$$

Similarly, denoting

$$P_x := \mathbb{P}_R(\|(M + R)x\|_\infty \leq \Delta),$$

we also have

$$\begin{aligned} (\mathbb{E}_R[S])^2 &= (\mathbb{E}_x \mathbb{P}_R(\|(M + R)x\|_\infty \leq \Delta)) \cdot (\mathbb{E}_y \mathbb{P}_R(\|(M + R)y\|_\infty \leq \Delta)) \\ &= \mathbb{E}_{x,y} [\mathbb{P}_R(\|(M + R)x\|_\infty \leq \Delta) \cdot \mathbb{P}_R(\|(M + R)y\|_\infty \leq \Delta)] \\ &= \mathbb{E}_{xy}[P_x \cdot P_y]. \end{aligned}$$

To compare the quantities $\mathbb{E}_R[S^2]$ and $(\mathbb{E}_R[S])^2$, we first consider a distribution over colorings. A natural distribution to consider is the distribution on colorings derived from the Gram-Schmidt walk which ensures that the discrepancy vector Mx is 1-subgaussian if x is sampled from this distribution. However, we shall also need that the colorings x themselves have a subgaussian tail as well as some additional nice properties that will be useful to compute the second moment. In particular, we prove the following lemma in Section 2.1.

► **Lemma 5 (Truncated Gram-Schmidt Distribution).** *Let $M \in \mathbb{R}^{d \times n}$ be a worst-case Komlós instance. Then, for any constant $C' > 1$ there exists a distribution \mathcal{G} over colorings $x \in \{\pm 1\}^n$ satisfying the following properties:*

- *Almost Constant Euclidean Norm for the discrepancy vectors: for every $x \in \text{supp}(\mathcal{G})$, we have $\|Mx\|_2 \in [r \pm \Delta]$ where $r = O(d^{1/2})$ and $\Delta = d^{-C'}$.*
- *Almost subgaussian tails for the colorings and discrepancy vectors: there exists a constant C depending on C' , such that for every $u \in \mathbb{S}^{n-1}$,*

$$\mathbb{P}_{x \sim \mathcal{G}}[|\langle x, u \rangle| \geq t] \leq 2d^C \cdot e^{-t^2/8} \text{ and } \mathbb{P}_{x \sim \mathcal{G}}[|\langle Mx, u \rangle| \geq t] \leq 2d^C \cdot e^{-t^2/8}.$$

Since the colorings sampled from the above distribution are subgaussian, $|\langle x, y \rangle| \leq n/2$ holds with high probability. To compute the second moment to a good precision, we need a careful comparison of the ratio $P_{xy}/(P_x \cdot P_y)$ for any two colorings x and y where this event occurs. We show the following bound in this case (proof in Section 2.2).

▷ **Claim 6 (Strong bound).** For any two colorings $x, y \in \text{supp}(\mathcal{G})$, denote $\epsilon = \epsilon(x, y) = \langle x, y \rangle/n$. If $|\epsilon| \leq 1/2$, then we have

$$P_{x,y} \leq P_x P_y \cdot \beta(x, y) \text{ where } \beta(x, y) = \exp(\delta_1 + d\epsilon^2 + d\delta^2\epsilon^2 + \delta^2\epsilon \cdot \langle Mx, My \rangle),$$

where the scaling factor $\delta := \frac{\sqrt{d}}{\sigma\sqrt{n}}$ and the error parameter $\delta_1 \leq 1/\text{poly}(d)$.

When the low probability event $|\epsilon| \geq 1/2$ occurs, we use the weak bound $P_{xy} \leq \min\{P_x, P_y\}$.

As x is sampled from the truncated Gram-Schmidt distribution, the probabilities P_x turn out to be almost constant for all colorings $x \in \text{supp}(\mathcal{G})$ as the following claim shows.

▷ **Claim 7.** For any coloring $x \in \text{supp}(\mathcal{G})$,

$$P_x = \exp(\delta_x) \cdot p \text{ where } p := \left(\frac{\delta\Delta}{\sqrt{2\pi}}\right)^d \exp\left(-\frac{\delta^2 r^2}{2}\right), \quad (7)$$

with the scaling factor $\delta := \frac{\sqrt{d}}{\sigma\sqrt{n}}$ and the error parameter δ_x satisfying $|\delta_x| \leq \delta_1 \leq 1/\text{poly}(d)$.

The proof of this claim is in Section 2.2.

14:8 Smoothed Analysis of the Komlós Conjecture

We now focus on the case when $|\epsilon| \leq 1/2$. When we take $x, y \sim \mathcal{G}$, as P_x and P_y are essentially constant, by Claim 6, applying the second moment method reduces to bounding $\beta(x, y)$ as defined in Claim 6. To do this, we will use the properties, as described in Lemma 5, of the underlying random variables x and Mx . The following technical lemma gives a bound on the exponential moment for such random variables.

► **Lemma 8.** *Let X be a non-negative random variable X that satisfies*

$$\mathbb{P}(X \geq t) \leq d^{C_1} \cdot e^{-t^2/8} \text{ for any } t > 0,$$

for some fixed constant $C_1 > 0$. Then for any $\lambda = c_2 \sqrt{\log d}$ with $c_2 \geq \sqrt{32C_1}$,

$$\mathbb{E}[\exp(X^2/\lambda^2)] \leq 1 + 32C_1/c_2^2 + o_d(1).$$

We shall prove this lemma in Section 2.1.

We can now complete the proof of Theorem 1 by comparing $\mathbb{E}_R[S^2]$ and $(\mathbb{E}_R[S])^2$. We show that

► **Lemma 9.** *For $n = \omega(d \log d)\sigma^{-4/3}$, we have*

$$(\mathbb{E}_R[S])^2 = p^2(1 - o_d(1)) \text{ and } \mathbb{E}_R[S^2] = p^2(1 + o_d(1)).$$

The above implies that $\mathbb{E}_R[S^2] = (1 + o(1))(\mathbb{E}_R[S])^2$, and thus the Paley-Zygmund inequality implies Theorem 1 as discussed in the proof overview.

Proof of Lemma 9. For the first moment, Claim 7 implies that

$$(\mathbb{E}_R[S])^2 = \mathbb{E}_{x, y \sim \mathcal{G}}[P_x P_y] = p^2 \mathbb{E}[\exp(\delta_x + \delta_y)] \geq p^2 \exp(-2\delta_1).$$

Since $0 < \delta_1 \leq 1/\text{poly}(d)$, the bound follows.

To compute the second moment, $\mathbb{E}_R[S^2] = \mathbb{E}_{x, y \sim \mathcal{G}}[P_{xy}]$, we define \mathcal{E} to be the event that the colorings $x, y \sim \mathcal{G}$ satisfy $|\langle x, y \rangle| > n/2$ and compute the contribution to the expectation under \mathcal{E} and its complement separately. In particular, using Claim 7 and Claim 6, we have

$$\mathbb{E}_R[S^2] = \mathbb{E}_{x, y \sim \mathcal{G}}[P_{x, y}] \leq \mathbb{P}_{x, y \sim \mathcal{G}}[\mathcal{E}] \cdot p + \mathbb{E}_{x, y \sim \mathcal{G}}[P_x P_y \beta(x, y) \cdot \mathbf{1}[\bar{\mathcal{E}}]] \quad (8)$$

For the first term in (8), since $n \geq d$, Lemma 5 implies that

$$\mathbb{P}_{x, y \sim \mathcal{G}}[\mathcal{E}] \leq \text{poly}(d) \cdot e^{-n/4} \leq e^{-n/8}.$$

Thus, using the exact bound for p from Claim 7 and that $\delta\Delta \leq \text{poly}(\sigma/d)$, the first term

$$\begin{aligned} \mathbb{P}_{x, y \sim \mathcal{G}}[\mathcal{E}] \cdot p &= p^2 \cdot \mathbb{P}[\mathcal{E}] \cdot p^{-1} \leq p^2 \cdot e^{-n/8} \cdot \left(\frac{\sqrt{2\pi}}{\delta\Delta}\right)^d \exp\left(\frac{\delta^2 r}{2}\right) \\ &\leq p^2 \cdot e^{-n/8} \cdot \exp(O(d \log(dn/\sigma) + d^2/(\sigma^2 n))) \\ &= p^2 \cdot o_d(1), \end{aligned} \quad (9)$$

when $n = \omega(d \log d)\sigma^{-4/3}$. In particular, as $\sigma \leq 1$, we have $n/8 \gg d \log(dn/\sigma) + d^2/(\sigma^2 n)$.

For the second term in (8), using Claim 7, we have that $P_x = p \cdot \exp(\delta_x)$ where $|\delta_x| \leq |\delta_1| \leq 1/\text{poly}(d)$. Thus,

$$\begin{aligned} \mathbb{E}_{x, y \sim \mathcal{G}}[P_x P_y \beta(x, y) \cdot \mathbf{1}[\bar{\mathcal{E}}]] &\leq p^2 \cdot \exp(2|\delta_1|) \cdot \mathbb{E}[\beta(x, y) \cdot \mathbf{1}[\bar{\mathcal{E}}]] \\ &\leq \exp(2|\delta_1|) \cdot \mathbb{E}[\beta(x, y)], \end{aligned} \quad (10)$$

since $\beta(x, y)$ is a non-negative random variable. Recall from Claim 6 that $\beta(x, y) \leq \exp(\delta_1) \cdot \exp(Z)$ where $\delta_1 \leq 1/\text{poly}(d)$ and

$$Z = d\epsilon^2 + 2\delta^2\epsilon^2r^2 + 2\delta^2|\epsilon\langle Mx, My \rangle|.$$

Renormalizing $\bar{\epsilon} = \langle x, n^{-1/2}y \rangle$ and $\bar{\theta} = \langle Mx, r^{-1}My \rangle$ and using that $\delta = \frac{\sqrt{d}}{\sigma\sqrt{n}}$ and $r \leq \sqrt{d}$, we have

$$Z \leq \left(\frac{d}{n} + \frac{2d^2}{\sigma^2n^2} \right) \cdot \bar{\epsilon}^2 + \frac{2d\sqrt{d}}{\sigma^2n\sqrt{n}} \cdot |\bar{\epsilon} \cdot \bar{\theta}| \leq (|\bar{\epsilon}| + |\bar{\theta}|)^2 / \lambda_{\min}^2,$$

where we denote

$$\lambda_{\min} = \frac{1}{3} \sqrt{\min \left\{ \frac{n}{d}, \frac{\sigma^2n^2}{2d^2}, \frac{\sigma^2n^{1.5}}{2d^{1.5}} \right\}}.$$

Note that $\lambda_{\min} = \omega_d(1) \cdot \sqrt{\log d}$ when $n = \omega(d \log d) \sigma^{-4/3}$.

We now bound the tails of $\bar{\epsilon}$ and $\bar{\theta}$, which will allow us to bound $\mathbb{E}[\exp(Z)]$. Conditioned on any outcome of $y \sim \mathcal{G}$, and as $\|y\|n^{-1/2} = 1$, the second property in Lemma 5 gives that $\mathbb{P}_{x \sim \mathcal{G}}[|\langle x, n^{-1/2}y \rangle| \geq t] \leq 2d^C \exp(-t^2/8)$. Averaging over y thus gives that

$$\mathbb{P}[|\bar{\epsilon}| \geq t] \leq 2d^C \cdot e^{-t^2/8}.$$

Similarly, as $\|My\|r^{-1} \leq 1$ for any y in the support of \mathcal{G} , we have that $\mathbb{P}_{x \sim \mathcal{G}}[|\langle Mx, r^{-1}My \rangle| \geq t] \leq 2d^C \exp(-t^2/8)$, and averaging over y gives that

$$\mathbb{P}[|\bar{\theta}| \geq t] \leq 2d^C \cdot e^{-t^2/8}.$$

By a union bound, it follows that the random variable $X := |\bar{\epsilon}| + |\bar{\theta}|$ satisfies the tail condition of Lemma 8 with constant $C_1 = 2C$. So when $\lambda_{\min} = \omega_d(1) \cdot \sqrt{\log d}$, the parameter $c_2 := \lambda_{\min}/\sqrt{\log d}$ in Lemma 8 satisfies $32C_1/c_2^2 = o_d(1)$. Therefore, Lemma 8 implies that

$$\mathbb{E}[\beta(x, y)] \leq \exp(\delta_1) \cdot \mathbb{E}[\exp(Z)] \leq (1 + o_d(1))(1 + o_d(1)) = 1 + o_d(1).$$

Plugging the above in (10), it follows that the second term

$$\mathbb{E}_{x, y \sim \mathcal{G}}[P_x P_y \beta(x, y) \cdot \mathbf{1}[\bar{\mathcal{E}}]] \leq p^2(1 + o_d(1)).$$

Combining this with (8) and (9), we get that $\mathbb{E}_R[S^2] \leq p^2(1 + o_d(1))$. \blacktriangleleft

We now prove the lemmas and claims used in the proof of Theorem 1 above.

2.1 Truncated Gram-Schmidt Distribution and Exponential Moments

Proof of Lemma 5. Consider running the Gram-Schmidt walk algorithm on the matrix M stacked with the identity matrix, i.e. $\begin{pmatrix} M \\ I_n \end{pmatrix}$, and let \mathcal{G}_0 be the distribution over colorings obtained as an output of the algorithm.

Since each column of the stacked matrix has Euclidean norm at most 2, the properties of the Gram-Schmidt walk (Theorem 4) guarantees that $(x, Mx) \in \mathbb{R}^{n+d}$ where $x \in \{\pm 1\}^n$ and $Mx \in \mathbb{R}^d$ is 2-subgaussian. It follows that both x and Mx are 2-subgaussian as well when $(x, Mx) \sim \mathcal{G}_0$.

14:10 Smoothed Analysis of the Komlós Conjecture

To obtain a distribution \mathcal{G} where $\|Mx\|_2$ is almost constant for each coloring $x \in \text{supp}(\mathcal{G})$, we will truncate the distribution \mathcal{G}_0 in such a way that the tails are also preserved up to poly(d) factors. Towards this end, we first note that with probability $1 - e^{-cd}$, we have that $\|Mx\|_2 \leq c'\sqrt{d}$ for constants c and c' . This is because for any σ -subgaussian mean-zero random vector X , the Euclidean norm of $\|X\|$ has a subgaussian tail (e.g. Exercise 6.3.5 in [19]). In particular, $\mathbb{P}[\|X\|_2 \geq c_1\sigma\sqrt{d} + t] \leq e^{-c_2t^2/\sigma^2}$ for some universal constants $c_1, c_2 > 0$. Now, by a pigeonhole argument, for a large enough constant C' there exist an annulus W with width $\Delta = d^{-C'}$ and inner radius $r \leq c'\sqrt{d}$ such that $\mathbb{P}_{x \sim \mathcal{G}_0}(x \in W) \geq d^{-C}$ for a constant C depending on C' .

We take the distribution \mathcal{G} to be the probability measure of \mathcal{G}_0 conditioned on the event that $Mx \in W$. It then follows that for any coloring $x \sim \mathcal{G}$, we have $|\|Mx\|_2 - r| \leq \Delta$. Moreover, since x and Mx were 2-subgaussian prior to conditioning, and the probability mass of the annulus is at least d^{-C} , conditioning can only increase the probability of any event by a factor of d^C . Thus, the tail bounds as stated in the statement of the lemma also follow. \blacktriangleleft

Proof of Lemma 8. The assumption on X implies that for any $t \geq 4 \cdot \sqrt{C_1 \log d}$, we have

$$\mathbb{P}(X \geq t) \leq \exp(-t^2/16). \quad (11)$$

We express the expectation as an integration

$$\begin{aligned} \mathbb{E}[\exp(X^2/\lambda^2)] &= \int_0^\infty \mathbb{P}[\exp(X^2/\lambda^2) > s] ds = \int_0^\infty \mathbb{P}(X \geq \lambda\sqrt{\log s}) ds \\ &\leq 1 + c_3 + \int_{1+c_3}^\infty \mathbb{P}(X \geq \lambda\sqrt{\log s}) ds. \end{aligned}$$

Let us set $c_3 = 32C_1/c_2^2$, so that $c_3 \leq 1$ as $c_2 \geq \sqrt{32C_1}$. For $s \geq 1 + c_3$, we have

$$\lambda\sqrt{\log s} \geq c_2\sqrt{\log d} \cdot \sqrt{c_3/2} \geq 4 \cdot \sqrt{C_1 \log d}$$

using that $\sqrt{\log(1+x)} \geq \sqrt{x/2}$ for $x \in [0, 1]$ and as $c_3 \leq 1$. So the condition $t \geq 4 \cdot \sqrt{C_1 \log d}$ for (11) is satisfied whenever $s \geq 1 + c_3$, and applying (11) to the above integration gives

$$\begin{aligned} \int_{1+c_3}^\infty \mathbb{P}(X \geq \lambda\sqrt{\log s}) ds &\leq \int_{1+c_3}^\infty \exp(-\lambda^2 \log s/16) ds \\ &= \left(\frac{\lambda^2}{16} - 1\right)^{-1} \cdot (1+c_3)^{-\lambda^2/16+1} \leq \exp(-c_2^2 c_3 \log d/16). \end{aligned}$$

By our choice of c_3 , the above is at most d^{-2C_1} . Thus, it follows that $\mathbb{E}[\exp(X^2/\lambda^2)] \leq 1 + 32C_1/c_2^2 + d^{-2C_1}$. This proves the lemma. \blacktriangleleft

2.2 Proof of Claims from Section 2

Proof of Claim 6. Since the rows of R are independent, to compute the above ratio, it suffices to compute the ratio for a single row of $M + R$. Fix $i \in [d]$, and let $m = m_i$ and $r = r_i$ denote the i^{th} row and define $a = a_i(x) := -m^\top x$ and $b = b_i(y) := -m^\top y$. We want to compare the ratio of $\mathbb{P}(r^\top x \in [a \pm \Delta], r^\top y \in [b \pm \Delta])$ to $\mathbb{P}(r^\top x \in [a \pm \Delta]) \cdot \mathbb{P}(r^\top y \in [b \pm \Delta])$.

Notice that $r^\top x$ and $r^\top y$ are Gaussian random variables with mean 0, variance $1/\delta^2$, and covariance $\mathbb{E}_r[r^\top x r^\top y] = \mathbb{E}_r[x^\top r r^\top y] = \epsilon/\delta^2$. Denoting the square $K := [a \pm \Delta] \times [b \pm \Delta]$, we have that

$$\mathbb{P}(r^\top x \in [a \pm \Delta], r^\top y \in [b \pm \Delta]) = \mu_\epsilon(\delta K),$$

where μ_ϵ is the 2-dimensional centered Gaussian measure with covariance matrix $\begin{pmatrix} 1 & \epsilon \\ \epsilon & 1 \end{pmatrix}$ and δK denotes the δ scaling of K . Similarly, we can write $\mathbb{P}(r^\top x \in [a \pm \Delta]) \cdot \mathbb{P}(r^\top y \in [b \pm \Delta]) = \mu(\delta K)$, where μ is the *standard* 2-dimensional Gaussian measure.

We will compare the ratio $\mu_\epsilon(\delta K)/\mu(\delta K)$ by approximating the Gaussian measure over δK with the density at the center and show the following bound

$$\mu_\epsilon(\delta K)/\mu(\delta K) \leq \exp(3\alpha + \epsilon^2 + \delta^2\epsilon^2(a^2 + b^2) + 2\delta^2\epsilon ab). \quad (12)$$

Since $(a_1(x), \dots, a_d(x)) = Mx$ and $(b_1(x), \dots, b_d(x)) = My$, using the above bound for all the rows $i \in [d]$, we have

$$\frac{P_{x,y}}{P_x P_y} \leq \exp(3d\alpha + d\epsilon^2 + \delta^2\epsilon^2 \cdot (\|Mx\|_2^2 + \|My\|_2^2) + \delta^2\epsilon \cdot \langle Mx, My \rangle).$$

Since $\|Mx\|_2^2 \leq d + \text{poly}(1/d)$ for every $x \in \text{supp}(\mathcal{G})$, taking $\delta_1 = 4d\alpha$ gives the statement of the claim. To finish the proof we prove (12) now.

Abusing notation and denoting by $\mu(s, t)$ and $\mu_\epsilon(s, t)$ the corresponding densities at $(s, t) \in \mathbb{R}^2$, we have the following explicit formula for the density μ_ϵ :

$$\mu_\epsilon(s, t) = \frac{1}{2\pi\sqrt{1-\epsilon^2}} \cdot \exp\left(-\frac{s^2 + t^2 - 2\epsilon st}{2(1-\epsilon^2)}\right).$$

Since the edge length of the square δK is $2\delta\Delta$, whenever $|\epsilon| \leq 1/2$, a direct calculation with the densities shows that

$$\frac{\sup_{(s,t) \in \delta K} \mu(s, t)}{\inf_{(s,t) \in \delta K} \mu(s, t)} = \exp(2\delta^2\Delta(|a| + |b|)) \leq \exp(2\delta^2\Delta(|a| + |b| + 2\Delta)) \leq \exp(\delta_1),$$

and that

$$\frac{\sup_{(s,t) \in \delta K} \mu_\epsilon(s, t)}{\inf_{(s,t) \in \delta K} \mu_\epsilon(s, t)} \leq \exp(4\delta^2\Delta(|a| + |b| + 2\Delta)) \leq \exp(2\delta_1),$$

where δ_1 is as defined in the claim. It follows that whenever $|\epsilon| \leq 1/2$, we can use the density at the center of K to obtain

$$\begin{aligned} \frac{\mu_\epsilon(K)}{\mu(K)} &\leq \exp(3\delta_1) \cdot \frac{\mu_\epsilon(\delta a, \delta b)}{\mu(\delta a, \delta b)} = \frac{1}{\sqrt{1-\epsilon^2}} \cdot \exp\left(3\delta_1 + \frac{\delta^2\epsilon^2(a^2 + b^2)}{2(1-\epsilon^2)} + \frac{\delta^2\epsilon ab}{1-\epsilon^2}\right) \\ &\leq \exp(3\delta_1 + \epsilon^2 + \delta^2\epsilon^2(a^2 + b^2) + 2\delta^2\epsilon ab), \end{aligned}$$

thus proving (12). ◁

Proof of Claim 7. We have $P_x = \prod_{i \in [d]} \mathbb{P}[r_i^\top x \in [a_i \pm \Delta]]$. For any fixed $i \in [d]$, $r_i^\top x$ is distributed as $\mathcal{N}(0, 1/\delta^2)$, so after scaling the quantity $\mathbb{P}[r_i^\top x \in [a_i \pm \Delta]] = \mu(\delta \cdot I)$ where $I = [a_i \pm \Delta]$ and μ is the standard Gaussian measure in \mathbb{R} . Analogous to the proof of Claim 6, one can approximate the Gaussian density at any at the point in I by the center point a , and compute similarly to the proof of Claim 6 that

$$P_x = \prod_{i \in [d]} \mathbb{P}[r_i^\top x \in [a_i \pm \Delta]] = \left(\frac{\delta\Delta}{\sqrt{2\pi}}\right)^d \exp\left(\alpha_x - \frac{\delta^2\|Mx\|_2^2}{2}\right),$$

for some small error $|\alpha_x| \leq 2\delta^2\Delta(\|Mx\|_1 + d\Delta)$. As $\|Mx\|_2 \in [r \pm \Delta]$ and $r = O(\sqrt{d})$, we have that $\|Mx\|_1 = O(d)$ and the statement of the claim follows for some $\delta_x \leq |\alpha_x| + 1/\text{poly}(d) \leq 1/\text{poly}(d)$. ◁

3 Conclusion

For the Komlós problem, as studied in this paper, Gaussian noise is a natural way to model a smoothed analysis setting since the input vectors have Euclidean norm at most one. One can wonder whether similar results can be obtained with more general noise models, for instance, Bernoulli or other discrete noise models. Such noise models are also more conducive for smoothed analysis in other discrepancy settings, such as for the Beck-Fiala problem. The weighted second moment approach used here can also handle Bernoulli noise when the number of vectors $n \gg d^2$ but the second moment becomes difficult to control when n is smaller. It remains an interesting open problem to see if Bernoulli or other discrete noise models can be handled for the regime $n \gg d \log d$.

References

- 1 Dimitris Achlioptas and Yuval Peres. The threshold for random k -sat is $2^k \log 2 - o(k)$. *J. Amer. Math. Soc.*, 4:947–973, 2004.
- 2 Noga Alon and Joel H. Spencer. *The Probabilistic Method*. John Wiley & Sons, 2016.
- 3 Dylan J. Altschuler and Jonathan Niles-Weed. The discrepancy of random rectangular matrices. *CoRR*, abs/2101.04036, 2021. [arXiv:2101.04036](https://arxiv.org/abs/2101.04036).
- 4 Wojciech Banaszczyk. Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Struct. Algorithms*, 12(4):351–360, 1998.
- 5 Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt walk: A cure for the Banaszczyk blues. *Theory Comput.*, 15:1–27, 2019.
- 6 Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha. Prefix discrepancy, smoothed analysis, and combinatorial vector balancing. In *Proceedings of ITCS*, pages 13:1–13:22, 2022.
- 7 Nikhil Bansal and Raghu Meka. On the discrepancy of random low degree set systems. In *Proceedings of SODA 2019*, pages 2557–2564, 2019.
- 8 József Beck and Tibor Fiala. “Integer-making” theorems. *Discrete Appl. Math.*, 3(1):1–8, 1981.
- 9 Esther Ezra and Shachar Lovett. On the Beck-Fiala conjecture for random set systems. *Random Struct. Algorithms*, 54(4):665–675, 2019.
- 10 Cole Franks and Michael Saks. On the discrepancy of random matrices with many columns. *Random Struct. Algorithms*, 57(1):64–96, 2020.
- 11 Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. In *Proceedings of FOCS*, 2021.
- 12 Christopher Harshaw, Fredrik Sävje, Daniel Spielman, and Peng Zhang. Balancing covariates in randomized experiments using the gram-schmidt walk. *arXiv e-prints*, 2019. [arXiv:1911.03071](https://arxiv.org/abs/1911.03071).
- 13 Rebecca Hoberg and Thomas Rothvoss. A Fourier-Analytic Approach for the Discrepancy of Random Set Systems. In *Proceedings of SODA*, pages 2547–2556, 2019.
- 14 Calum MacRury, Tomáš Masařík, Leilani Pai, and Xavier Pérez-Giménez. The phase transition of discrepancy in random hypergraphs, 2021. [arXiv:2102.07342](https://arxiv.org/abs/2102.07342).
- 15 Aditya Potukuchi. Discrepancy in random hypergraph models. *CoRR*, abs/1811.01491, 2018.
- 16 Aditya Potukuchi. A spectral bound on hypergraph discrepancy. In *Proceedings of ICALP*, pages 93:1–93:14, 2020.
- 17 Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- 18 Paxton Turner, Raghu Meka, and Philippe Rigollet. Balancing gaussian vectors in high dimension. In *Proceedings of COLT*, pages 3455–3486, 2020.
- 19 Roman Vershynin. *High-Dimensional Probability*, volume 1. Cambridge University Press, 2018.